
*HiCommand® Storage Services
Manager
User Guide*

**Version 5.1
MK-93AI002-07**

© Copyright 2002-2006 Hewlett-Packard Development Company, L.P.

All rights reserved.

Content in this document is subject to change without notice. Hewlett-Packard Development Company assumes no responsibility or liability for any errors or inaccuracies that may appear in this document. Information in this document is furnished under license and must be used in accordance with such license. No part of this document shall be reproduced, stored in a retrieval system, or transmitted by any means, electronic, mechanical, photocopying, recording, other than for the purchaser's personal use without the written permission of Hewlett-Packard Development Company. Microsoft and Windows are registered trademarks of Microsoft Corporation. Oracle is a registered trademark of Oracle Corporation. Sun, Solaris, Sun StorEdge, and Java are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. AIX and IBM are registered trademarks of International Business Machines Corporation in the United States, other countries or both. SGI and IRIX are registered trademarks of Silicon Graphics, Inc. Netscape is a registered trademark of Netscape Communications Corporation in the United States and other countries. HDS and HiCommand are registered trademarks of Hitachi Data Systems. HP, HP-UX, and OpenVMS, Tru64 UNIX are registered trademarks of Hewlett-Packard Development Company. QLogic is a trademark of QLogic Corporation. Emulex is a registered trademark of Emulex Corporation. HBAnyware is a trademark of Emulex Corporation. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Other product and company names mentioned herein may be the trademarks of their respective owners.

First Edition

Table of Contents

Preface	<i>About This Guide</i>	<i>li</i>
	About the User Guide	li
	Suggested Topics for First-Time Users	lii
	Before Using the Management Server	liii
	Intended Audience	liii
	Where to Find Additional Documentation	liv
 Chapter 1	 <i>Overview</i>	 <i>1</i>
	About This Product	1
	Suggested Topics for First-Time Users	3
	Product Components	3
	<i>Management Server Components</i>	4
	About the User Interface	6
	<i>About the Top Pane</i>	6

<i>About the Left Pane</i>	7
Opening and Closing the Left Pane	9
About the Home Page	11
Accessing the Management Server	14
Installing the Java Plug-in	16
Installing the Software Security Certificate	17
<i>Installing the Certificate by Using Microsoft</i> <i>Explorer 6.0</i>	17
<i>Installing the Certificate by Using Netscape</i> <i>Navigator 7</i>	18
<i>Changing the Security Certificate to Match</i> <i>the Name of the Server</i>	18
Stopping the Service for the Management Server	19
Signing Out of the Management Server	20

Chapter 2	<i>Discovering NAS Devices, Tape Libraries, Switches and Storage Systems</i>	21
	Discovery Steps	22
	<i>Overall Discovery Tasks</i>	22
	Overview of Discovery Features	25
	<i>Setting Default User Names and Passwords</i>	26
	<i>Adding an IP Range for Scanning</i>	28
	<i>Adding a Single IP Address or DNS Name for Discovery</i>	29
	<i>Modifying a Single IP Address Entry for Discovery</i>	31
	<i>Removing Elements from the Addresses to Discover List</i>	32
	<i>Importing Discovery Settings from a File</i>	32
	<i>Saving Discovery Settings to a File</i>	33

Step 1 - Discover Switches	34
<i>SMI-S Switches Must Be Removed and Rediscovered After Upgrading</i>	37
<i>Discovering Brocade Switches</i>	37
Brocade SMI-S Provider Installation Requirements (New Installations of the Management Server)	38
Verifying Brocade Rapid Program Is Set to 1	39
About Brocade Discovery (Specifying Brocade Discovery Using Fabric Access API or SMI-S)	39
<i>Discovering CNT Switches</i>	42
<i>Discovering Cisco Switches</i>	43
<i>Discovering Sun StorEdge and QLogic Switches</i>	45
<i>Changing the SNMP Trap Listener Port for Sun StorEdge Switches</i>	46
<i>Discovering McDATA and EMC Connectrix Switches</i>	47
SWAPI Setting Through a Proxy	49
SNMP Setting Through a Proxy	53
Contacting a McDATA or Connectrix Switch Directly	56
Changing the Discovery Settings	58
<i>Excluding McDATA and EMC Connectrix Switches from Discovery</i>	59
<i>Viewing Log Messages</i>	60
<i>Viewing the Status of System Tasks</i>	61
Duplicate Logs for Brocade Switches in Same Fabric	61
Step 2 - Discover Storage Systems, NAS Devices and Tape Libraries.	62
<i>Discovering 3PAR Storage Systems</i>	65
<i>Discovering EMC Solutions Enabler 5.1</i>	66
<i>Excluding EMC Symmetrix Storage Systems from Discovery</i>	68
<i>Discovering EMC CLARiiON Storage Systems</i>	70
<i>Discovering Engenio Storage Systems</i>	70

<i>Discovering HDS Storage Systems</i>	72
<i>Excluding HDS Storage Systems from Discovery</i>	73
<i>Discovering HP StorageWorks EVA or MSA Arrays</i>	74
<i>Discovering HP StorageWorks XP Arrays</i>	75
Discovering HP XP Arrays by Using Command View XP	76
Discovering HP XP Arrays by Using Command View XP Advanced Edition	77
Discovering HP XP Arrays by using the XP Provider	78
<i>Discovering IBM Storage Systems</i>	78
<i>Discovering Sun StorEdge 3510 Storage Systems</i>	80
<i>Discovering Sun StorEdge 6920 and 6940 Storage Systems</i>	82
<i>Discovering Sun StorEdge 6130 Storage Systems</i>	82
<i>Discovering Xiotech Storage Systems</i>	83
<i>Discovering HP NAS Devices on Windows</i>	84
<i>Discovering HP NAS Devices on Linux</i>	85
<i>Discovering NetApp NAS Devices</i>	86
<i>Discovering Sun NAS Devices</i>	87
<i>Discovering HP and IBM Tape Libraries</i>	88
Step 3 - Build the Topology	89
<i>Building the Topology View</i>	89
<i>Modifying the Properties of a Discovered Address</i>	91
<i>Deleting Elements from the Product</i>	92
Deleting an Element Using System Explorer or Chargeback	92
Deleting Elements Using Discovery Step 2 (Topology)	93
Step 4 - Get Details	94
<i>Get Details</i>	94
<i>Stopping the Gathering of Details</i>	96

	<i>Excluding EMC Symmetrix Storage Systems from Force Device Manager Refresh</i>	<i>97</i>
	<i>Excluding HDS Storage Systems from Force Device Manager Refresh</i>	<i>98</i>
	Troubleshooting Mode	100
	Managing McDATA and EMC Connectrix Switches	100
	<i>About Managing McDATA and EMC Connectrix Switches</i>	<i>101</i>
	<i>Adding McDATA and EMC Connectrix Switches</i>	<i>101</i>
	<i>Removing McDATA and EMC Connectrix Switches</i>	<i>102</i>
	<i>Swapping McDATA and EMC Connectrix Switches</i>	<i>102</i>
	Assigning a File Extension in Netscape 7	103
	Filtering Discovery Groups	104
	Moving Elements to Another Discovery Group	104
	Placing an Element in Quarantine	105
	Removing an Element from Quarantine	106
	Updating the Database with Element Changes	106
	Notifying the Software of a New Element	108
Chapter 3	<i>Installing and Discovering the Windows Proxy</i>	<i>111</i>
	Installing the Windows Proxy	112
	Discovering the Windows Proxy	113
	Configuring Windows Proxy Authentication	114
	Decreasing the Maximum Java Heap Size	116
	Removing the Windows Proxy	116
Chapter 4	<i>Discovering Applications and Hosts</i>	<i>117</i>
	Step 1 - Discovering Your Hosts	117
	<i>Step A - Set Up Discovery for Hosts</i>	<i>119</i>
	<i>Step B - Build the Topology</i>	<i>123</i>
	<i>(Optional) Step C - View the Topology</i>	<i>123</i>

<i>Step D - Obtain Details</i>	124
Step 2 - Setting Up Discovery for Applications.	125
<i>Monitoring Oracle</i>	126
Step A - Create the APPIQ_USER Account for Oracle	127
Removing the APPIQ_USER Account for Oracle	128
Step B - Provide the TNS Listener Port.	130
Step C - Set up Discovery for Oracle 10g	130
Discovering Oracle Clusters	131
<i>Monitoring Microsoft SQL Server</i>	132
Switching to Mixed Mode Authentication.	133
Step A - Create the APPIQ_USER for the SQL Server	133
Step B - Provide the Microsoft SQL Server Name and Port Number .	135
Removing the APPIQ_USER Account for SQL Server	137
Deleting SQL Server Information.	137
<i>Monitoring Sybase Adaptive Server Enterprise</i>	138
Step A - Create the APPIQ_USER account for Sybase	138
Removing the APPIQ_USER Account for Sybase	140
Step B - Provide the Sybase Server Name and Port Number.	141
Deleting Sybase Information	141
<i>Monitoring Microsoft Exchange</i>	142
Adding Microsoft Exchange Domain Controller Access	142
Deleting a Microsoft Exchange Domain Controller	143
Step 3 - Discovering Applications	143
<i>Step A - Detect Your Applications</i>	144
<i>Step B - Obtain the Topology</i>	145
<i>Step C - Obtain Get Details</i>	145
Changing the Oracle TNS Listener Port.	147
Adding/Modifying Microsoft Exchange Domain Controller Access . . .	148

Changing the Password for the Managed Database Account	149
Obtaining Disk Drive Statistics from Engenio Storage Systems	150

Chapter 5	<i>Managing Security</i>	153
	About the Security for the Management Server	154
	<i>About Roles</i>	154
	<i>About Organizations</i>	157
	<i>Planning Your Hierarchy</i>	160
	<i>Naming Organizations</i>	161
	Managing User Accounts	161
	<i>Adding Users</i>	162
	<i>Editing a User Account</i>	163
	<i>Changing the Password for a User Account</i>	164
	<i>Changing Your Password</i>	165
	<i>Deleting Users</i>	165
	<i>Modifying Your User Profile</i>	166
	<i>Modifying Your User Preferences</i>	167
	System Explorer and Element Topology Preferences	167
	Event Manager Preferences	167
	Warnings for Slow Systems Operations	168
	<i>Viewing the Properties of a Role</i>	169
	<i>Viewing the Properties of an Organization</i>	169
	Managing Roles	170
	<i>Adding Roles</i>	170
	<i>Editing Roles</i>	171
	<i>Deleting Roles</i>	172
	Managing Organizations	173

<i>Adding an Organization</i>	173
Adding Storage Volumes to an Organization	174
<i>Viewing Organizations</i>	175
<i>Editing Organizations</i>	176
<i>Deleting an Organization</i>	177
<i>Removing Members from an Organization</i>	177
<i>Filtering Organizations</i>	178
Changing the Password of System Accounts	181
Using Active Directory/LDAP for Authentication	182
<i>Step 1 - Configure the Management Server to Use Active Directory or LDAP</i>	183
Active Directory	184
LDAP	188
<i>Step 2 - Restart the AppStorManager Service and Login as the Designated Admin Account</i>	192
<i>Step 3 - Add Users to the Management Server</i>	193
<i>Step 4 - Provide Login Information to Your Users</i>	193
 Chapter 6 <i>Managing Licenses</i>	195
About the License	195
Importing a License File	200
Viewing Cumulative Licenses	200
Viewing a Specific License	201
Deleting a License	202
Import License File Window Appears Too Small	202
 Chapter 7 <i>Configuring the Management Server</i>	205
Trap Generation	206

<i>Changing the Default to SNMPv2</i>	206
Setting Up E-mail Notification	207
Managing Getting Discovery Details	208
<i>Adding a Discovery Schedule</i>	209
<i>Disabling a Schedule</i>	210
<i>Editing a Schedule</i>	211
<i>Removing a Schedule</i>	211
Managing Product Health	211
<i>Enabling Disk Space Monitoring</i>	212
<i>Viewing the Results of Disk Space Monitoring</i>	213
<i>Advanced Settings</i>	214
Modifying Java Memory Settings	214
Customizing Properties	214
Managing Logging	215
<i>Accessing the Log Files</i>	215
<i>Downloading Logs to a File</i>	216
<i>Downloading the User Audit Log</i>	217
<i>Downloading the Discovery Summary Log</i>	217
<i>Displaying a Log File in a Command Prompt Window</i>	218
<i>Changing the Provider Log Level</i>	218
<i>Enabling the Scanning of Critical Events</i> <i>of the Management Server Database</i>	220
<i>Viewing the Results of Logging</i>	221
Managing the Display of Events	221
<i>Controlling the Display of Cleared and Deleted Events</i>	222
<i>Modifying the Clearing and Deletion Frequency</i>	222
<i>Configuring the Clearing of Events</i>	224

<i>Configuring the Deletion of Events</i>	225
Managing File Server SRM	226
<i>Managing Collectors for File Server SRM</i>	226
<i>Adding File Servers</i>	227
<i>Setting Data Collection Properties</i>	228
<i>File Server SRM Scan Status</i>	230
<i>Performing a Test File Server SRM Scan</i>	231
Managing Reports	232
<i>Architectural Overview of Report Views and Report Cache Refresh</i>	232
<i>Managing Collectors for Reports</i>	236
Starting Collectors	238
Stopping Collectors	238
<i>Setting the Date and Time for Report Collectors</i>	239
<i>Viewing Scheduled E-mail Deliveries for Reports</i>	240
<i>Editing E-mail Schedules for Reports</i>	242
<i>Viewing Data Aging Statistics for Reports</i>	243
<i>Scheduling Report Cleanup</i>	245
<i>Refreshing the Report Cache</i>	245
Refreshing the Report Cache Immediately	246
Scheduling a Report Cache Refresh	246
<i>Setting Up Global Reporter</i>	248
<i>Editing Remote Site Information</i>	253
<i>Remote Sites Are Not Removed from the tnsnames.ora File</i>	254
Managing Performance Collection	255
<i>Managing Performance Collectors</i>	256
<i>Starting Performance Collectors</i>	258
<i>Stopping Performance Collectors</i>	259

	<i>Setting the Date and Time for Performance Collectors</i>	259
	<i>Viewing Data Aging Statistics for Performance</i>	261
	Editing the Locale and Currency Settings	263
	Process Names	264
Chapter 8	<i>Database Maintenance and Management</i>	267
	Database Maintenance Window	267
	Overview of Backups	269
	Database Mode	271
	Architectural Overview of RMAN Backups	272
	Performing an RMAN Hot Backup	275
	Scheduling RMAN Hot Backups of the Database	275
	Viewing Results from RMAN Backup	277
	About the Database Admin Utility	277
	<i>Accessing the Database Admin Utility</i>	278
	<i>Refreshing the Database Admin Utility</i>	279
	<i>Checking the Database and Listener Status</i>	279
	<i>Changing the Password of System Accounts</i>	279
	<i>Exporting the Database</i>	280
	<i>Importing the Database</i>	281
	<i>Re-initializing the Database</i>	282
	<i>Resetting the Temp and Undo Tablespace</i>	283
	<i>Restarting the Database</i>	283
	<i>Restoring a Cold Backup</i>	284
	<i>Setting the Archive Mode</i>	284
	<i>RMAN Restore</i>	285
	<i>Running a Cold Backup</i>	286
	<i>Changing the Archive Destination</i>	287

<i>Viewing Logs from the Database Admin Utility</i>	288
<i>Changing the Oracle Listener Password</i>	288
<i>Resetting the Admin Password for the Management Server</i>	289
<i>Resetting the Log</i>	289
<i>Warning Messages During Reinitializing the Database</i>	290
Database Password	290
Troubleshooting Listener and Database Connection Problems	291

Chapter 9 *Viewing Element Topology and Properties* **295**

About System Explorer	296
<i>Grey Screen When Attempting to Access System Explorer</i>	298
Accessing System Explorer	298
About Cisco Switches and VSANs in System Explorer	299
About the User Interface	299
<i>About the User Interface for System Explorer</i>	300
<i>About the Toolbar in System Explorer</i>	300
<i>Description of the Icons Displayed in the Topology</i>	304
<i>About the List Tab</i>	304
Viewing Elements by Type	306
<i>About the Access Tab</i>	307
Obtaining Information About Zone Entries	307
Obtaining Information About Host Bindings	310
Obtaining Information About Storage System LUN Masking	313
<i>About the Path Tab</i>	315
<i>About the Right-Click Menu Options</i>	317
Viewing Storage Elements	324
<i>Adding a Virtual Application</i>	325

<i>Adding Information for “Discovered” Hosts</i>	326
<i>Arranging Elements in the Topology</i>	327
<i>Closing Topology Windows</i>	329
<i>Using the Global View</i>	330
<i>Printing Elements in System Explorer</i>	330
<i>Exporting the Topology to Visio</i>	332
Viewing the Topology in Microsoft Visio	333
Installing Storage Planner	333
Configuring Visio to View Exported Topology	334
<i>Updating Element Data</i>	335
<i>Viewing Ports</i>	335
<i>Showing the Impact of an Element</i>	336
<i>Assigning a Business Cost to an Application</i>	338
<i>Expanding the Topology Pane</i>	340
<i>Filtering Fabrics</i>	340
<i>Viewing Event Status in the Topology</i>	341
<i>Custom Name for Switch Truncated in Topology</i>	344
<i>Managing Groups</i>	344
About Groups	345
Grouping "Discovered" Hosts	345
Ungrouping "Discovered" Hosts	346
Grouping "Discovered" Storage Systems	347
Ungrouping "Discovered" Storage Systems	348
<i>Managing Fabrics</i>	349
Changing the Fabric Name	349
Deleting Fabrics	349
<i>Hiding and Showing Generic Hosts</i>	350

About Hiding Generic Hosts	350
Hiding Generic Hosts for a Switch	351
Expanding Generic Hosts for a Switch	352
Hiding Generic Hosts for All Switches	352
Expanding Generic Hosts for All Switches	353
Setting Up Custom Commands	353
<i>About Custom Commands</i>	354
<i>Adding a Custom Command.</i>	355
<i>Editing a Custom Command</i>	358
<i>Deleting a Custom Command</i>	358
<i>Software Environment Variables for Scripting.</i>	358
<i>Using the Remote Console</i>	362
About the Remote Console	362
Keeping the Remote Console Active	363
About the Buttons on the Remote Console	364
About the Drop-Down Menu Options.	365
Copying Text from the Remote Console	366
Using External Tools	366
<i>About the External Tools Feature.</i>	366
<i>Setting up External Tools</i>	367
About the Navigation Tab.	367
<i>Finding the Status of a Port on a Switch</i>	372
<i>Accessing the Navigation Tab</i>	373
Viewing Element Properties.	374
<i>About the Properties Tab</i>	374
<i>Accessing the Properties Tab</i>	375
<i>Viewing Fabric Properties</i>	376

<i>Assigning a Custom Name</i>	377
Viewing Element Topology	377
<i>About the Topology Tab</i>	378
Accessing the Topology	385
<i>About the New Window Option</i>	386
<i>Printing the Topology</i>	387
Creating a Virtual Application	390
About the Provisioning Tab	391
About the Events Tab	391
Asset Attributes of an Element	392
About the Collectors Tab	394
About the Monitoring Tab	395
About the Policies Tab	396
Determining If a Host Belongs to a File System	396
About the Data from CXFS File Systems	397
 Chapter 10 <i>Provisioning</i>	 399
About Provisioning	399
<i>About Provisioning Brocade Switches</i>	
<i>After Upgrading</i>	400
Managing Zones	401
<i>SAN Zoning Overview</i>	401
<i>Accessing Information About Zone Aliases</i>	408
<i>Creating Zone Aliases</i>	409
<i>Modifying a Zone Alias</i>	410
<i>Deleting Zone Aliases</i>	411
<i>Accessing Information About Zoning</i>	411
<i>Creating a Zone in a Fabric</i>	413
<i>Adding and Removing Zone Members</i>	415

<i>Deleting Zones</i>	415
<i>Accessing Information About Zone Sets</i>	416
<i>Creating a Zone Set</i>	417
<i>Modifying a Zone Set</i>	419
<i>Deleting Zone Sets</i>	419
<i>Copying a Zone Set</i>	420
<i>Activating a Zone Set</i>	421
<i>Zones and Zone Sets Are Sometimes Listed Twice</i>	423
<i>Changing the Amount of Information Collected from the Inactive Zone Database (Cisco Switches)</i>	424
<i>About the Messages Displayed in the Brocade Console</i>	426
Managing Storage	426
<i>Setting Up Storage Partitioning</i>	427
<i>Modifying the Cache Settings (Engenio and Sun 6130)</i>	433
<i>Changing the Owner of a Volume (Engenio, CLARiiON and Sun 6130)</i> ...	434
<i>Managing Storage Pools</i>	434
Creating a Storage Pool (Engenio, CLARiiON, Sun 6130 and Sun 35xx)	434
Accessing Information About Storage Pools	436
Deleting a Storage Pool (Engenio and CLARiiON Only)	437
<i>Managing Volumes</i>	438
Accessing Information About Volumes	438
Filtering Volumes	440
Creating a Storage Volume	441
Deleting a Storage Volume	445
Changing the Cache Block Size for a Storage System	446
Modifying the Cache Settings (Engenio and Sun 6130 Only)	447

	<i>Rules for Creating Host Security Groups</i>	448
	<i>Managing Host Security Groups</i>	452
	Accessing Information About Host Security Groups	452
	Creating Host Security Groups	454
	Editing Host Security Groups	457
	Deleting Host Security Groups	460
	Setting the Host Mode for IBM Storage Systems	460
	<i>General Provisioning Issues</i>	461
	Provisioning Can Make a Device Inaccessible	461
	Provisioning Does Not Make an Operating System Aware of a Device	461
	<i>Provisioning Issues by Vendor</i>	462
	Issues Specific to CLARiiON Storage Systems	462
	Issues Specific to EMC Symmetrix Storage Systems	463
	Issues Specific to HDS Storage Systems	465
	Issues Specific to HP Storage Systems	469
	Issues Specific to Engenio Storage Systems	470
Chapter 11	<i>Path Provisioning</i>	473
	About Path Provisioning	474
	How Path Provisioning Works	476
	How to Use Path Provisioning	477
	Default System Action Templates	479
	<i>Volume Creation, LUN Security, and Zone Operation</i>	479
	Step 1 - Select Storage System	480
	Step 2 - Select a Host	481
	Step 3 - Select a Volume	482
	Step 4 - Select a Host Security Group	483

Step 5 - Select a Zone	484
<i>Meta Volume Creation</i>	<i>486</i>
<i>LUN Security</i>	<i>488</i>
Step 1 - Select Storage System	489
Step 2 - Select a Host	490
Step 3 - Select a Volume	492
Step 4 - Select a Host Security Group	493
<i>Zone Operation</i>	<i>493</i>
Step 1 - Select Storage System	494
Step 2 - Select a Host	495
Step 3 - Select a Zone	497
<i>Volume Creation and LUN Security</i>	<i>499</i>
Step 1 - Select Storage System	500
Step 2 - Select a Volume	501
Step 3 - Select a Host Security Group	504
<i>LUN Security and Zone Operation</i>	<i>504</i>
Step 1 - Select Storage System	505
Step 2 - Select a Host	506
Step 3 - Select a Host Security Group	508
Step 4 - Select a Zone	509
<i>Volume Assignment</i>	<i>511</i>
Step 1 - Select Storage System	512
Step 2 - Select a Volume	513
Step 3 - Select a Host Security Group	516
<i>Providing a LUN Number</i>	<i>516</i>
Creating a System Action Template	517
Modifying a System Action Template	518

Adding a Host	519
Creating a Host Security Group	520
Scheduling Provisioning Jobs	521
Executing Provisioning Jobs	523
Monitoring Provisioning Jobs	524
Deleting Multiple Jobs	525
Naming Conventions for Zones	525
Using Multipathing with Path Provisioning	526
Customizing Path Provisioning	527
<i>Storage System Customize Dialog</i>	528
<i>Host Customize Dialog</i>	528
<i>Volume Customize Dialog</i>	529
<i>LUN Customize Dialog</i>	529
<i>Zone Customize Dialog</i>	530
Automatically Configure Zoning	531
Manually Configure Zoning	531
Assigning Templates to a Role	533
 Chapter 12 <i>Event Management</i>	 535
About Event Manager	536
Viewing Event Details	541
Clearing Events	542
Unclearing Events	543
Configuring the Clearing of Events	543
Configuring the Deletion of Events	544
Deleting Events	545
Sorting Events	545
Adding Journal Entries	546
Changing the CLARiiON Event Polling Interval	546
Brocade Events	548

<i>Brocade Switch Events</i>	548
<i>Supported Brocade Events</i>	549
Filtering Events	550
<i>About Filtering Events</i>	550
<i>Selecting a Severity Level</i>	551
<i>Customizing the Severity Level Filter</i>	552
<i>Selecting an Element Type for Events</i>	553
<i>Customizing the Element Type Filter for Events</i>	554
<i>Filtering Events by Summary Text or Element Name</i>	555
<i>Hiding Filters in Event Manager</i>	556
Enabling the Rank Column in Event Manager	556

Chapter 13	<i>Running Reports</i>	559
	About Reporter	560
	Accessing Reporter	563
	Viewing Reports.	563
	<i>Report Parameters</i>	564
	Viewing Report Collectors for an Element	565
	Refreshing a Report.	565
	Changing the Formatting of a Report.	566
	Opening a Report in a New Window	566
	Maximizing the Screen Space for a Report	567
	Filtering Data in Global Reports.	569
	Sending a Report by E-mail.	570
	Managing E-mail Schedules for Reports	571
	<i>Adding an E-mail Schedule for a Report</i>	571
	<i>Editing an E-mail Schedule for a Report</i>	574
	<i>Deleting E-mail Schedules for a Report</i>	577
	<i>Viewing E-mail Schedules for a Report</i>	578

Creating Custom Reports	579
<i>About Creating Custom Reports</i>	579
<i>Configuring Report Designer to Work with the Management Server</i>	581
<i>Designing Custom Reports</i>	583
Creating Standard Reports	583
<i>Integrating Custom Reports</i>	592
<i>Detailed Schema Information</i>	595
<i>Views from Previous Releases</i>	626
<i>Implementing Custom Reports on Sun Solaris</i>	628
 Chapter 14 <i>Viewing Performance Data</i>	 629
About Performance Explorer	630
Accessing Performance Explorer	631
Creating Performance Charts	632
About the Toolbars in Performance Explorer	632
Comparing the Performance of Different Elements	637
Viewing Summary Charts	638
Viewing Trending Information for Performance	639
Removing Performance Data from a Graph	640
Setting a Custom Period	640
About the Monitoring Options	642
Task Already Running on Element	653
Monitoring with Direct Attached Storage	654
Supported Host Configurations for Monitoring	654
Sudden Dips Displayed in Certain Charts in Performance Explorer	657
Values Continue to Increase in Charts for Aggregated Drives and Aggregate Volumes	658

Chapter 15	<i>Finding an Element's Storage Capacity</i>	659
	About Capacity Explorer	660
	Accessing Capacity Explorer	662
	About the Toolbars in Capacity Explorer	663
	Finding the Capacity of an Element	667
	<i>Capacity Information for Applications</i>	668
	<i>Capacity Information for Hosts</i>	669
	<i>Capacity Information for NetApp NAS Devices</i>	670
	<i>Capacity Information for Storage Systems</i>	671
	Obtaining Utilization Reports	673
	Printing Elements in Capacity Explorer	674
	Viewing Capacity Charts	676
	Viewing Trending Information for Storage Capacity	678
	Different Results for the df -k Command and Capacity Explorer	679
Chapter 16	<i>Managing Policies</i>	681
	About Policy Manager	682
	Accessing Policy Manager	682
	Creating Policies	683
	<i>Creating a Utilization Policy</i>	683
	<i>About the Policy Templates</i>	686
	<i>Creating Policies for Discovery</i>	690
	<i>Creating Policies for Provisioning</i>	692
	<i>Creating Policies for Events</i>	694
	<i>Testing a Utilization Policy</i>	697
	Modifying Policies	697
	<i>Modifying Utilization Policies</i>	698
	<i>Modifying Discovery Policies</i>	700

	<i>Modifying Provisioning Policies</i>	702
	<i>Modifying Policies for Events</i>	703
	Viewing Policies	706
	Deactivating a Policy	707
	Deleting Policies	707
	Providing E-mail Notification for a Policy	708
	Providing Event Generation for a Policy	708
	Providing a Custom Command for a Policy	710
Chapter 17	<i>Chargeback</i>	713
	About Chargeback	714
	Setting Up Chargeback	716
	Accessing Chargeback	716
	Creating an Asset Record	717
	Changing the Status of an Element	718
	Saving Chargeback Information	719
	Viewing Assets	720
	Defining Storage Tiers	721
	<i>Creating a New Storage Tier</i>	722
	<i>Adding Elements to a Storage Tier</i>	722
	<i>Removing Elements from a Storage Tier</i>	723
	<i>Editing a Storage Tier</i>	723
	<i>Deleting a Storage Tier</i>	724
	Adding Asset Information	724
	<i>Adding Asset Information</i>	724
	<i>Adding General Information</i>	727
	<i>Adding Staff Information</i>	728
	<i>Adding Geographic Information</i>	729
	<i>Adding Licensing and Warranty Information</i>	729

<i>Adding Custom Information</i>	729
Managing Departments	730
<i>Adding Departments</i>	730
<i>Editing a Department</i>	731
<i>Removing a Department from Chargeback</i>	731
Setting the Infrastructure Cost	731
<i>About Asset-based Infrastructure Cost</i>	732
<i>About Storage-based Infrastructure Cost</i>	732
Setting Up Asset and Storage Based Chargeback	733
<i>Setting Up Asset-Based Chargeback</i>	733
Step 1 - Specify Financial information	735
Step 2 - Assign Departmental Ownership Percentage	736
Step 3 - Review Asset-based Chargeback Result	737
<i>Setting Up Storage-Based Chargeback</i>	739
Step 1 - Assign Departmental Ownership Percentage	741
Step 2 - Review Storage Tier Cost	742
Step 3 - Review Storage Dependency and Cost	742
Step 4 - Review Storage-Based Chargeback Result	743
<i>Editing Percentage of Ownership</i>	743
<i>Removing Department Ownership of an Element</i>	744
<i>Capacity in Chargeback and Capacity Explorer Differ</i>	744
<i>How a Depreciation Method Is Calculated</i>	745
Calculating Straight Line Depreciation	746
Calculating Fixed Declining Balance	747
Calculating Double Declining Balance	749
Viewing Chargeback	752
<i>Viewing Chargeback by Element</i>	752

<i>Viewing Chargeback by Department</i>	753
<i>Viewing Chargeback by Owner</i>	754
Chargeback Reports	756
<i>Viewing Chargeback Reports</i>	756
<i>E-mailing a Chargeback Report</i>	758
<i>Managing E-mail Schedules for Chargeback Reports</i>	759
Adding an E-mail Schedule for a Chargeback Report	759
Editing an E-mail Schedule for a Chargeback Report	761
Deleting E-mail Schedules for a Chargeback Report	763
Viewing E-mail Schedules for a Chargeback Report	763
Viewing the History of an E-mail Chargeback Schedule	765
Filtering Assets	765
<i>About Filtering Assets</i>	766
<i>Selecting an Element Type for Chargeback</i>	766
<i>Customizing the Element Type Filter for Chargeback</i>	767
<i>Filtering Assets by Status</i>	767
<i>Customizing the Asset Status Filter for Chargeback</i>	768
<i>Hiding Filters in Chargeback</i>	768
Chapter 18 Business Tools	771
About the Business Tools	772
Using the HBA Replacement Automator	773
<i>Installing New HBA with Old HBA</i>	773
<i>Installing New HBA by Itself</i>	774
Using the Brocade Firmware Update Utility	774
<i>Updating the Firmware Without Modifying the Configuration File</i>	775
<i>Updating the Firmware With Modifying the Configuration File</i>	776

	Setting up Risk Analysis	777
	Global Change Management	778
Chapter 19	<i>Storage Terminology Table</i>	781
	Storage Terminology Table.	781
Chapter 20	<i>Troubleshooting.</i>	785
	“Data is late or an error occurred” Message	786
	appiq.log Filled with Connection Exceptions	786
	Receiving “HTTP ERROR: 503” When Accessing the Management Server	787
	<i>Errors in the Logs.</i>	788
	Permanently Changing the Port a CIM Extension Uses (UNIX Only)	789
	Configuring UNIX CIM Extensions to Run Behind Firewalls.	791
	Volume Names from Ambiguous Automounts Are Not Displayed	798
	Solaris Management Server Suddenly Restarts	798
	Installing the Software Security Certificate.	799
	<i>Installing the Certificate by Using Microsoft Explorer 6.0.</i>	799
	<i>Installing the Certificate by Using Netscape Navigator 7.</i>	800
	<i>Changing the Security Certificate to Match the Name of the Server.</i>	800
	Troubleshooting After Upgrading	801
	<i>SMI-S Switches Must Be Removed and Rediscovered After Upgrading</i>	802
	Troubleshooting Discovery and Get Details.	802

<i>Names are Changed After Running Get Details for Cisco SMI-S Switches</i>	803
<i>Configuring E-mail Notification for Get Details</i>	804
<i>Increasing the Time-out Period and Number of Retries for Switches in Progress</i>	805
<i>"Connection to the Database Server Failed" Error</i>	807
<i>Using the Test Button to Troubleshoot Discovery</i>	807
<i>DCOM Unable to Communicate with Computer</i>	809
<i>Duplicate Listings for Brocade Switches in Same Fabric</i>	810
<i>Element Logs Authentication Errors During Discovery</i>	810
<i>EMC Device Masking Database Does Not Appear in Topology (ALX Only)</i>	810
<i>Management Server Does Not Discover Another Management Server's Database</i>	811
<i>Microsoft Exchange Drive Shown as a Local Drive</i>	811
<i>Unable to Discover Microsoft Exchange Servers</i>	811
<i>Nonexistent Oracle Instance Is Displayed</i>	811
<i>Requirements for Discovering Oracle</i>	811
<i>Do Not Run Overlapping Discovery Schedules</i>	812
<i>"This storage system uses unsupported firmware. ManagementClassName: ???" Message</i>	812
Troubleshooting Topology Issues	813
<i>About the Topology</i>	813
<i>Undiscovered Hosts Display as Storage Systems</i>	817
<i>Solaris Machines Appear to Have Extra QLogic HBAs</i>	818
<i>No Stitching for Brocade Switches with Firmware 3.2.0</i>	818

<i>Link Between a Brocade Switch and a Host Disappears from the Topology</i>	818
<i>Incorrect Topology Sometimes Displayed for CNT Switches</i>	819
<i>Unable to Find Elements on the Network</i>	819
<i>Unable to See Path Information</i>	819
<i>Device Locking Mechanism for Brocade Element Manager Query/Reconfiguration</i>	819
<i>A Discovered Sun StorEdge A5000 JBOD Does Not Display Its WWN Properly</i>	820
<i>Unable to Monitor McDATA Switches</i>	820
<i>Unable to Detect a Host Bus Adapter</i>	820
<i>Navigation Tab Displays Removed Drives as Disk Drives</i>	821
<i>Unable to Obtain Information from a CLARiiON Storage System</i>	821
<i>Discovery Fails Too Slowly for a Nonexistent IP Address</i>	821
<i>“CIM_ERR_FAILED” Message</i>	823
<i>Re-establishing Communication with EFCM</i>	824
<i>Communicating with HiCommand Device Manager Over SSL</i>	825
<i>Unable to Discover a UNIX Host Because of DNS or Routing Issues</i>	826
Troubleshooting the Java Plug-in	828
<i>Java Applet Has Data from a Different Version of Management Server Software</i>	828
<i>OutOfMemoryException Messages</i>	828
<i>Unable to View System Explorer After Upgrade</i>	828
<i>Improving Reload Performance in System Explorer</i>	829
Troubleshooting Provisioning	829
<i>Cannot Access a Resource Owned by Another Controller</i>	829
<i>Error -56</i>	830
<i>“Can't delete this zone” Message</i>	830

<i>Changes in EFC Manager Requiring Get Details</i>	830
Troubleshooting Hardware	830
<i>About Swapping Host Bus Adapters</i>	831
<i>"Fork Function Failed" Message on AIX Hosts</i>	831
<i>Known Driver Issues</i>	831
<i>Known Device Issues</i>	831
<i>"mailbox command 17 failure status FFF7" Message</i>	835
<i>"Process Has an Exclusive Lock" Message</i>	836
 <i>Glossary</i>	 837

List of Figures

Figure 1-1:	<i>Status Light.</i>	7
Figure 1-2:	<i>Closing the Left Pane.</i>	9
Figure 1-3:	<i>Opening the Left Pane</i>	11
Figure 2-1:	<i>Setting Default User Names and Passwords</i>	27
Figure 2-2:	<i>Adding an IP Range for Scanning</i>	29
Figure 2-3:	<i>Deleting Elements from the Management Server.</i>	93
Figure 5-1:	<i>Parent-Child Hierarchy for Organizations</i>	158
Figure 5-2:	<i>Children in Multiple Organizations.</i>	159
Figure 5-3:	<i>Changing Your User Profile.</i>	166
Figure 5-4:	<i>Accessing the User Preferences Tab</i>	167
Figure 5-5:	<i>Clicking the Organization Link</i>	179
Figure 5-6:	<i>Filtering Organizations</i>	180
Figure 5-7:	<i>Active Organization</i>	181

Figure 6-1:	<i>An Example of Direct Attached Storage</i>	199
Figure 7-1:	<i>Selecting the Date and Time from the Calendar.....</i>	229
Figure 7-2:	<i>Report Views and Report Cache Refresh.....</i>	233
Figure 7-3:	<i>Selecting the Date and Time from the Calendar.....</i>	240
Figure 7-4:	<i>An Example of Global Reporting</i>	249
Figure 7-5:	<i>Selecting the Date and Time from the Calendar.....</i>	261
Figure 9-1:	<i>Expanding the Fabric Node</i>	305
Figure 9-2:	<i>Fabric Members Are Highlighted.....</i>	306
Figure 9-3:	<i>Highlighting the Applications in the Topology.....</i>	307
Figure 9-4:	<i>Members of a Zone Set</i>	308
Figure 9-5:	<i>Displaying Zone Members</i>	309
Figure 9-6:	<i>Zone Member to Switch.....</i>	310
Figure 9-7:	<i>Obtaining Information About a Zone Member's Adapter....</i>	310
Figure 9-8:	<i>Highlighting Elements with Host Bindings.....</i>	311

Figure 9-9:	<i>Displaying Host Bindings</i>	312
Figure 9-10:	<i>HBA Port Properties</i>	313
Figure 9-11:	<i>WWN Properties</i>	314
Figure 9-12:	<i>WWN Properties</i>	315
Figure 9-13:	<i>Obtaining Path Information</i>	316
Figure 9-14:	<i>Path Information Visible in the Tree</i>	317
Figure 9-15:	<i>Enclosing the Elements</i>	328
Figure 9-16:	<i>Dragging Multiple Elements to Their New Location</i>	329
Figure 9-17:	<i>Showing the Impact of an Application</i>	337
Figure 9-18:	<i>Determining Business Cost</i>	339
Figure 9-19:	<i>Expanding the Topology Pane</i>	340
Figure 9-20:	<i>Filtering Fabrics</i>	341
Figure 9-21:	<i>Remote Console</i>	363
Figure 9-22:	<i>Obtaining Information About a Host</i>	369

Figure 9-23:	<i>Details of a Host Connected to a Switch</i>	370
Figure 9-24:	<i>Finding the Status of a Port</i>	372
Figure 9-25:	<i>Topology of a Server</i>	379
Figure 9-26:	<i>Topology of a Server (Continued)</i>	380
Figure 9-27:	<i>Multipathing Displayed in the Topology</i>	382
Figure 9-28:	<i>Multipathing Displayed in the Topology (Continued)</i>	383
Figure 9-29:	<i>Direct Attached Storage in the Topology</i>	384
Figure 9-30:	<i>New Window Option</i>	387
Figure 9-31:	<i>Viewing Asset Records</i>	393
Figure 9-32:	<i>CXFS File System</i>	398
Figure 10-1:	<i>Resources in Two Zones</i>	403
Figure 10-2:	<i>Overview of Zoning Capabilities</i>	404
Figure 10-3:	<i>Selecting a Setting and Size for a Storage Pool</i>	435
Figure 10-4:	<i>Deleting Several Volumes at Once</i>	446

Figure 10-5:	<i>Selecting All Volumes</i>	446
Figure 11-1:	<i>Selecting a Storage System</i>	481
Figure 11-2:	<i>Selecting a Fabric</i>	485
Figure 11-3:	<i>Selecting a Storage System</i>	490
Figure 11-4:	<i>Selecting a Storage System</i>	495
Figure 11-5:	<i>Selecting a Fabric</i>	498
Figure 11-6:	<i>Selecting a Storage System</i>	501
Figure 11-7:	<i>Selecting a Volume</i>	503
Figure 11-8:	<i>Selecting a Storage System</i>	506
Figure 11-9:	<i>Selecting a Fabric</i>	510
Figure 11-10:	<i>Selecting a Storage System</i>	513
Figure 11-11:	<i>Selecting a Volume</i>	515
Figure 11-12:	<i>Specifying a LUN Number</i>	517
Figure 11-13:	<i>Selecting a Port</i>	520

Figure 11-14:	<i>Deleting Multiple Provisioning Jobs</i>	525
Figure 11-15:	<i>Selecting the Second Path</i>	527
Figure 11-16:	<i>Selecting a Fabric</i>	532
Figure 12-1:	<i>Event Manager Summary</i>	536
Figure 12-2:	<i>Right-clicking Event Details</i>	541
Figure 12-3:	<i>Accessing Event Details</i>	541
Figure 12-4:	<i>Selecting Event Filters</i>	553
Figure 13-1:	<i>Hiding the Middle Pane</i>	567
Figure 13-2:	<i>Closing the Left Pane in Reporter</i>	568
Figure 13-3:	<i>Selecting Organizations Used in This Report</i>	574
Figure 13-4:	<i>Selecting Organizations Used in This Report</i>	577
Figure 13-5:	<i>Report Architecture</i>	580
Figure 13-6:	<i>Choosing a Standard Report</i>	584
Figure 13-7:	<i>Adding Tables for a Standard Report</i>	585

Figure 13-8:	<i>Linking Common Data in Tables for a Standard Report</i>	586
Figure 13-9:	<i>Creating Search Criteria for Standard Reports</i>	587
Figure 13-10:	<i>Deciding Which Data Should Appear in the Report</i>	588
Figure 13-11:	<i>Sorting Information in the Report</i>	589
Figure 13-12:	<i>Selecting the Layout of the Report</i>	590
Figure 13-13:	<i>Report Template Displayed</i>	591
Figure 13-14:	<i>Clicking the View Tab</i>	592
Figure 15-1:	<i>Legend for Capacity Explorer</i>	661
Figure 15-2:	<i>Capacity of an Element</i>	662
Figure 15-3:	<i>Viewing the Capacity of Elements in a Fabric</i>	668
Figure 15-4:	<i>Viewing a Utilization Report</i>	674
Figure 16-1:	<i>Testing a Newly Created Utilization Policy</i>	686
Figure 16-2:	<i>Testing a Utilization Policy</i>	697
Figure 16-3:	<i>Testing a Modified Utilization Policy</i>	700

Figure 17-1:	<i>Selecting an Element</i>	719
Figure 17-2:	<i>Accessing an Element's Asset Information</i>	725
Figure 17-3:	<i>Asset Record Node</i>	726
Figure 17-4:	<i>Accessing an Element's Asset Information</i>	734
Figure 17-5:	<i>Clicking the Asset-based Node</i>	734
Figure 17-6:	<i>Ownership Cost</i>	738
Figure 17-7:	<i>Accessing an Element's Asset Information</i>	740
Figure 17-8:	<i>Clicking Storage-Based Node</i>	740
Figure 17-9:	<i>Chargeback Capacity</i>	745
Figure 17-10:	<i>Capacity in Capacity Explorer</i>	745
Figure 17-11:	<i>Accessing Storage-Based Chargeback</i>	753

List of Tables

Table 1:	<i>Additional Documentation</i>	<i>liv</i>
Table 1-1:	<i>Status Light Settings</i>	<i>7</i>
Table 1-2:	<i>Buttons Displayed When the Left Pane Is Closed</i>	<i>9</i>
Table 1-3:	<i>Icons on the Home Page</i>	<i>11</i>
Table 2-1:	<i>Discovery Steps for Switches, NAS Devices, and Storage Systems</i>	<i>24</i>
Table 2-2:	<i>Discovery Requirements for Switches</i>	<i>35</i>
Table 2-3:	<i>Brocade Discovery Methods</i>	<i>37</i>
Table 2-4:	<i>Required Switch Models and InVsn Versions for Discovery</i>	<i>42</i>
Table 2-5:	<i>Discovery Settings for McDATA and Connectrix Switches</i>	<i>48</i>
Table 2-6:	<i>Task Status Descriptions</i>	<i>61</i>
Table 2-7:	<i>Discovery Requirements for Storage Systems, Tape Libraries & NAS Devices</i>	<i>63</i>
Table 4-1:	<i>Making the Management Server Aware of Hosts</i>	<i>119</i>

Table 5-1:	<i>Default Role Privileges.....</i>	155
Table 5-2:	<i>Default Role Privileges with Elements.....</i>	156
Table 5-3:	<i>Changing User Preferences for Event Manager</i>	168
Table 6-1:	<i>License Requirements.....</i>	196
Table 6-2:	<i>Determining Managed Access Points.....</i>	198
Table 7-1:	<i>Logging Levels</i>	219
Table 7-2:	<i>Default Settings for Clearing Events</i>	224
Table 7-3:	<i>About the Collectors for File Server SRM</i>	227
Table 7-4:	<i>About File Server SRM Scan Status</i>	231
Table 7-5:	<i>About Collectors for Reports</i>	237
Table 7-6:	<i>Viewing System Wide E-mail Schedules.....</i>	241
Table 7-7:	<i>Description of “An Example of Global Reporting”</i>	250
Table 7-8:	<i>About Performance Collectors.....</i>	257
Table 7-9:	<i>Currency Settings</i>	263

Table 7-10:	<i>Process Names on Windows</i>	264
Table 7-11:	<i>Process Names on Solaris</i>	265
Table 8-1:	<i>Description of Backups</i>	270
Table 8-2:	<i>Backup Directories in %ORA_HOME%\rman</i>	273
Table 8-3:	<i>Backup Example</i>	273
Table 8-4:	<i>Default Passwords for Database Accounts</i>	290
Table 9-1:	<i>Description of Toolbar in System Explorer</i>	301
Table 9-2:	<i>Elements Displayed in the Topology</i>	304
Table 9-3:	<i>Drop-Down Menu Options Accessible from the Topology*</i> ...	318
Table 9-4:	<i>Drop-Down Menu Options on the Access Tab</i>	323
Table 9-5:	<i>Show Impact Results</i>	338
Table 9-6:	<i>Legend for Severity Level Icons</i>	341
Table 9-7:	<i>Descriptions of the Severity Levels Presented in the Severity Menu</i>	343

Table 9-8:	<i>Variables for All Elements*</i>	359
Table 9-9:	<i>Variables for Only Storage Systems, Switches, and Hosts*</i> ...	359
Table 9-10:	<i>Variables for Only Switches</i>	360
Table 9-11:	<i>Variables for Only Hosts*</i>	360
Table 9-12:	<i>Variables for Only Applications</i>	361
Table 9-13:	<i>Buttons on the Remote Console</i>	365
Table 9-14:	<i>Drop-Down Menu Options</i>	365
Table 9-15:	<i>Information Available from the Navigation Page</i>	371
Table 9-16:	<i>Port Status Definitions</i>	373
Table 9-17:	<i>About the Toolbar in the Topology Tab</i>	385
Table 10-1:	<i>Setting Up Zoning</i>	405
Table 10-2:	<i>Zoning Support</i>	406
Table 10-3:	<i>Provisioning and Pool Support</i>	427
Table 10-4:	<i>Volume and HSG Support</i>	429

Table 10-5:	<i>Setting Up Storage Partitioning</i>	431
Table 10-6:	<i>Allowed Initiators in Host Security Groups</i>	448
Table 10-7:	<i>Volume Usage</i>	471
Table 11-1:	<i>Overview for Path Provisioning</i>	478
Table 11-2:	<i>Zone Icons</i>	484
Table 11-3:	<i>Zone Icons</i>	497
Table 11-4:	<i>Zone Icons</i>	509
Table 11-5:	<i>Zone Icons</i>	532
Table 12-1:	<i>About the Icons in Event Manager</i>	537
Table 12-2:	<i>Supported Hardware</i>	537
Table 12-3:	<i>Default Settings for Clearing Events</i>	543
Table 12-4:	<i>Brocade Switch Events</i>	548
Table 12-5:	<i>Severity Levels</i>	551
Table 12-6:	<i>Element Types</i>	554

Table 13-1:	<i>Buttons Displayed When the Left Pane Is Closed.....</i>	568
Table 13-2:	<i>Viewing E-mail Schedules for a Report</i>	578
Table 13-3:	<i>Description of the Report Views.....</i>	595
Table 13-4:	<i>MVC_HOSTSUMMARYVW.....</i>	600
Table 13-5:	<i>MVC_CARDSUMMARYVW.....</i>	601
Table 13-6:	<i>MVC_HOSTVOLUMESUMMARYVW (logical volumes)....</i>	601
Table 13-7:	<i>MVC_HOSTDISKDRIVEVW.....</i>	602
Table 13-8:	<i>MVC_STORAGESYSTEMSUMMARYVW</i>	602
Table 13-9:	<i>MVC_STORGAEPOOLSUMMARYVW.....</i>	603
Table 13-10:	<i>MVC_STORAGEVOLUMESUMMARYVW.....</i>	604
Table 13-11:	<i>MVC_SWITCHSUMMARYVW.....</i>	605
Table 13-12:	<i>MVC_PORTSUMMARYVW</i>	606
Table 13-13:	<i>MVC_ZONESUMMARYVW.....</i>	607
Table 13-14:	<i>MVC_ZONEVW.....</i>	608

Table 13-15:	<i>MVC_PATHVW</i>	608
Table 13-16:	<i>MVC_SUBPATHVW</i>	609
Table 13-17:	<i>MVC_MULTIPATHVW</i>	610
Table 13-18:	<i>MVC_EVENTSVW</i>	610
Table 13-19:	<i>MVC_ORGANIZATIONVW</i>	611
Table 13-20:	<i>MVC_ORGRELATIONVW</i>	611
Table 13-21:	<i>MVC_HOSTCAPACITYVW</i>	611
Table 13-22:	<i>MVC_STORAGESYSTEMCONFIGVW</i>	612
Table 13-23:	<i>MVC_STORAGEPOOLCONFIGVW</i>	612
Table 13-24:	<i>MVC_SWITCHCONFIGVW</i>	613
Table 13-25:	<i>MVC_OPTIONALTABLEVW</i>	613
Table 13-26:	<i>MVC_DISKEXTENTSUMMARYVW</i>	613
Table 13-27:	<i>MVC_STORAGEVOLUMEPORTS</i>	614
Table 13-28:	<i>MVC_VOLUMEDISKDRIVEVW</i>	615

Table 13-29:	<i>MVC_STORAGEPROCESSORSUMMARYVW.....</i>	615
Table 13-30:	<i>MVC_DISKDRIVESUMMARYVW.....</i>	616
Table 13-31:	<i>MVC_DISK_EXTENTVW.....</i>	617
Table 13-32:	<i>MVC_ASSETSUMMARY.....</i>	617
Table 13-33:	<i>MVC_APPLICATIONSUMMARYVW.....</i>	619
Table 13-34:	<i>MVC_UNITACCESSVW.....</i>	620
Table 13-35:	<i>MVCA_DBAPPCAPACITYVW.....</i>	621
Table 13-36:	<i>MVCA_EXCHAPPCAPACITYVW.....</i>	621
Table 13-37:	<i>MVCA_VIRTUALAPPCAPACITYVW.....</i>	622
Table 13-38:	<i>MVCA_FSRM_VOLUMESUMMARYVW.....</i>	622
Table 13-39:	<i>MVCA_FSRM_AGESUMMARYVW.....</i>	623
Table 13-40:	<i>MVCA_FSRM_EXTDETAILSUMMARYVW.....</i>	623
Table 13-41:	<i>MVCA_FSRM_DIRDETAILSUMMARYVW.....</i>	623
Table 13-42:	<i>MVCA_FSRM_USERSUMMARYVW.....</i>	624

Table 13-43:	<i>MVCA_FSRM_GROUPSUMMARYVW.....</i>	624
Table 13-44:	<i>MVCA_FSRM_TOPNFILES</i>	625
Table 13-45:	<i>MVCA_FSRM_AGEDFILEDETAILS</i>	625
Table 13-46:	<i>MVCA_FSRM_LARGEDIRINFO.....</i>	625
Table 13-47:	<i>MVIEWCORE_STATUS.....</i>	626
Table 13-48:	<i>MVIEW_STATUS.....</i>	626
Table 13-49:	<i>Views from Previous Releases</i>	627
Table 14-1:	<i>Toolbar in Lower Pane of Performance Explorer.....</i>	633
Table 14-2:	<i>About the Monitoring Options</i>	642
Table 14-3:	<i>Host Monitoring Support.....</i>	655
Table 15-1:	<i>Toolbar in Lower Pane of Capacity Explorer</i>	663
Table 15-2:	<i>Explanation of the Properties of the Capacity Levels for HDS Array Groups</i>	672
Table 16-1:	<i>Policy Templates</i>	687

Table 16-2:	<i>Policy Table Description</i>	706
Table 17-1:	<i>Setting Up Chargeback</i>	716
Table 17-2:	<i>Element Type Icons</i>	721
Table 17-3:	<i>Viewing E-mail Schedules for a Chargeback Report</i>	764
Table 17-4:	<i>Element Types</i>	767
Table 19-1:	<i>Storage Terminology Table</i>	782
Table 20-1:	<i>Troubleshooting Firewalls</i>	792
Table 20-2:	<i>Time-out Properties</i>	806
Table 20-3:	<i>Retry Properties</i>	806
Table 20-4:	<i>Troubleshooting Discovery and Get Details</i>	814
Table 20-5:	<i>Known Device Issues</i>	832

This preface describes the following:

- About the User Guide on page li
- Suggested Topics for First-Time Users on page lii
- Before Using the Management Server on page liii
- Intended Audience on page liii
- Where to Find Additional Documentation on page liv

About the User Guide

Welcome. This User Guide explains how to do the following:

- Discover and obtain information from elements
 - Manage applications
 - Monitor SAN elements
 - View elements in the SAN (Storage Area Network)
 - Obtain information about the capacity of an element
 - Manage events
-

- Provision storage
- Manage assets and obtain cost reports
- Run storage reports

Suggested Topics for First-Time Users

First become familiar with the discovery and get details features in the management server. You must configure the management server to discover the devices on the network so that the management server becomes aware of them. Then, you must get details so that the management server is aware of the various types of elements on the network.

To learn more about discovery and getting all elements details, see the following topics:

- Chapter 2, “Discovering NAS Devices, Tape Libraries, Switches and Storage Systems” on page 21
- Chapter 4, “Discovering Applications and Hosts” on page 117

Once you have discovered and obtained details about the devices in the network, begin adding users and adding them to roles and organizations. See “About the Security for the Management Server” on page 154 for more information.

Once you are done adding users and roles, use the following list as a guideline for the topics you should learn about:

- About Event Manager on page 536
 - About Provisioning on page 399
 - About Reporter on page 560
 - About System Explorer on page 296
 - About Policy Manager on page 682
 - About Chargeback on page 714
 - About Capacity Explorer on page 660
-

Before Using the Management Server

Before you use the software for the first time, you must configure the software and install the CIM Extensions, as described in the Installation Guide.

Intended Audience

This guide assumes you are familiar with the following:

- IP networks
- Storage area networks (SANs), including your own.
- The operating systems used by the hosts in your network, for example, Solaris and Microsoft Windows. You should know how to access Microsoft Windows Explorer and the Microsoft Services window.
- Database applications, such as Oracle and Microsoft Exchange.

Where to Find Additional Documentation

The management server also ships with the following documentation:

Table 1: Additional Documentation

Document	Description
Release Notes	Provides late-breaking issues that might impact the usage of the product.
Installation Guide	Provides information on how to configure the management server and install the CIM Extensions
Online help	An online help system containing information from the PDFs in HTML format.
File Servers Guide	Describes how to use the file server storage resource management (SRM) functionality in the product.
CLI Guide	Provides information about the CLI commands.
Guide for monitoring applications	Describes how to use the management server to monitor applications.

The documentation listed above, except for the help system, can be found in two locations:

- On the CD-ROM used to install the CIM Extensions. The help system is not accessible from the CD-ROM.
 - From the online help system, which is accessible by clicking the **Help** menu in the upper-right corner of the main pane.
-

This chapter describes the following:

- About This Product on page 1
- Suggested Topics for First-Time Users on page 3
- Product Components on page 3
- About the User Interface on page 6
- Accessing the Management Server on page 13
- Installing the Java Plug-in on page 14
- Installing the Software Security Certificate on page 15
- Stopping the Service for the Management Server on page 18
- Signing Out of the Management Server on page 19

About This Product

This product can simplify your complex environment and lower your cost of management with CIM-based integrated storage management. The management software integrates the management of

applications, servers, storage networks and storage systems in a single, easy to implement and intuitive solution.

The management server integrates the various components in the storage area network infrastructure into a CIM/WBEM/SMI standards based database so you can eliminate any vendor dependencies and view and manage your infrastructure as a whole. A SAN is a network configuration that is dedicated to transporting storage data among network devices, such as storage systems, servers, tape libraries, and switches. Since the SAN is dedicated to transporting storage data, it frees up the data network for regular TCP/IP traffic.

By giving your administrators a single, integrated console to manage tactical activities such as provisioning storage, managing real time events, installing new applications, and migrating servers and storage, as well as strategic activities such as forecasting, planning and cost analysis. The management server's integrated storage management lowers your cost of acquiring and managing a heterogeneous storage environment.

Storage Management Terms

- **CIM** - A common data model of an implementation-neutral schema for describing overall management information in a network/enterprise environment.
- **Web-Based Enterprise Management (WBEM)** - An initiative based on a set of management and Internet standard technologies developed to unify the management of enterprise computing environments.
- **Storage Management Initiative (SMI)** - A SNIA standard for implementing data storage management using CIM.

See the glossary in this user guide for additional definitions.

Key Benefits

- More efficient use of existing assets
 - Increased application availability and performance
 - Quicker deployment of storage infrastructure and business applications
 - Protection of customer flexibility and investments with a standards-based interface
-

Suggested Topics for First-Time Users

First become familiar with the discovery and Get Details in the management server. You must configure the management server to discover the devices on the network so that the management server becomes aware of them. Then, you must Get Details so that the management server is aware of the various types of elements on the network.

To learn more about discovery and Get Details, see the following topics:

- Chapter 2, “Discovering NAS Devices, Tape Libraries, Switches and Storage Systems” on page 21
- Chapter 4, “Discovering Applications and Hosts” on page 117

Once you have discovered and obtained details about the devices in the network, begin adding users and adding them to roles and organizations. See Chapter 5, “Managing Security” on page 153 for more information.

Once you are done adding users and roles, use the following list as a guideline for the topics you should learn about:

- Chapter 10, “Provisioning” on page 399
- Chapter 12, “Event Management” on page 535
- Chapter 13, “Running Reports” on page 559
- Chapter 14, “Viewing Performance Data” on page 629

Product Components

This product ships with the following software:

- **Management server** - The management server provides various tools to let you monitor and manage your SAN devices. See “Management Server Components” on page 4 for more information about these tools.
 - **CIM Extensions** - A CIM Extension gathers information from the operating system and host bus adapters. It then makes the information available to the management server. For more information, see the following in the glossary. Refer to the *Installation Guide* for information on how to install the CIM Extensions.
-

- **(Optional) Module for managing Microsoft Exchange Server** - Management software for Microsoft Exchange lets the administrator actively manage the data storage requirements for Microsoft Exchange.
- **(Optional) Module for managing Oracle Database** - Management software for Oracle reduces database storage cost and improves performance, availability and reliability by assisting the Oracle database administrator in administering Oracle instances, particularly storage.
- **(Optional) Module for managing Microsoft SQL Server** - Management software for Microsoft SQL Server lets the administrator manage and monitor Microsoft SQL Server.
- **(Optional) Module for managing Sybase Adaptive Server Enterprise** - Management software for Sybase lets the administrator actively manage the data requirements for Sybase.

Management Server Components

You may not have access to all of the features described in this section depending on the following:

- The type of license you have. Depending on your license, all features may not be available. See the “List of Features” to determine if you have access to all features. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center**).
- The role assigned to your user account. For example, users assigned to the Help Desk role by default have access to Application Explorer and Event Manager, but not to System Explorer, Provisioning, Policy Manager, and Reporter.

The following is the full feature set available:

- **Application Explorer** - Application Explorer lets you monitor and display data from applications. You can access Application Explorer by clicking **Application Explorer**.
 - **Business Tools** - The management server provides advisors and automators. Advisors provides detailed information for you to make informed decisions about your network, such as non-compliant HBA firmware. Automators let you automate particular tasks. The advisors and automators available to you depend upon your release. See Chapter 18, “Business Tools” on page 771 for more information.
 - **Capacity Explorer** - Capacity Explorer provides a graphical representation of an element's storage capacity in the storage network. See Chapter 15, “Finding an Element's Storage Capacity” on page 659 for more information.
 - **Chargeback** - Chargeback helps you manage departmental ownership, track cost and assemble business reports making inquiries, such as audits and inventory reviews, easier. See Chapter 17, “Chargeback” on page 713 for more information.
-

- **Command Line Interface** - The command line interface (CLI) provides an alternate way for you to manage elements the management server monitors. You can use the CLI commands in scripts to manage your storage.
- **File Server SRM** - File Server SRM does a recursive lookup on the file system and stores the information in an embedded database. File Server SRM can scan files very quickly because of its structure in the database and because it uses a multi-threaded process. More than one process can be used at a time to scan the files. Refer to the File Servers Guide for more information.
- **Event Manager** - Event Manager lets you view, clear, sort and filter events from managed elements. An event can be anything that occurs on the element, such as a device connected to a Brocade switch has gone off-line. See Chapter 12, “Event Management” on page 535 for more information.
- **Path Provisioning** - Path Provisioning lets you schedule provisioning task, such as creating zones, to run at a later time. See Chapter 11, “Path Provisioning” on page 473 for more information.
- **Performance Explorer** - Performance Explorer provides a graphical representation of the performance history of an element, such as bytes transmitted per second for a switch. See Chapter 14, “Viewing Performance Data” on page 629 for more information.
- **Policy Manager** - Policy Manager can automatically send an e-mail, generate an event, or run a remote script when an element is being overused or when one of the following occurs:
 - A new element is discovered
 - Successful provisioning occurred
 - An event occurred on one or more specified elementsSee Chapter 16, “Managing Policies” on page 681 for more information.
- **Provisioning** - Provisioning assists you in creating zones, zone sets, and zone aliases, in addition to storage pools, volumes, and host security groups. For more information, see Chapter 10, “Provisioning” on page 399.
- **Reporter** - Reporter provides a variety of detailed reports, such as dependency, event, and utilization reports for discovered elements. See Chapter 13, “Running Reports” on page 559 for more information.
- **System Explorer** - System Explorer is the gateway to many features that let you view details about the discovered elements. System Explorer provides a topology that lets you view how the devices in your network are connected. For more information, see Chapter 9, “Viewing Element Topology and Properties” on page 295.

About the User Interface

The Home page provides a gateway into the functionality for the product. The user interface for the Home page is split into several panes:

- **Top pane** - to discovery and configuration features, in addition The management server can be configured to display the icons for the management server's utilities.
- **Left pane** - Displays the status light and the buttons for the management server's utilities.
- **Right pane** - Displays the output of a feature, such as the topology in System Explorer.

About the Top Pane

The menus and button in the upper-right corner of the page provide the following functionality.

Important: The **Configuration**, **Security**, and **Discovery** buttons only appear if you belong to a role that has System Configuration selected on the Edit Role page (such as the Domain Administrator role). See Chapter 5, “Managing Security” on page 153 for more information.

- **Configuration** - This menu provides the tools for you to manage the management server, such as saving the database. See Chapter 7, “Configuring the Management Server” on page 205.
- **Security** - This menu lets you manage users, organizations, roles and licenses. See Chapter 5, “Managing Security” on page 153.
- **Discovery**- This menu provides the tools for the management server to discover and obtain information from elements in your network. See Chapter 2, “Discovering NAS Devices, Tape Libraries, Switches and Storage Systems” on page 21 and Chapter 4, “Discovering Applications and Hosts” on page 117.
- **Sign out** - Use this button to log out of the management server. See “Signing Out of the Management Server” on page 19.
- **Help** - This menu provides access to the online help and the copyright information.

The links in the upper-left corner let you modify your profile and sign out of the management server.

The status light in the upper-left corner indicates the status of the management server. It usually displays green under normal conditions, as shown in the following figure:



Figure 1-1: Status Light

The status light changes color during the following:

Table 1-1: Status Light Settings

Scenario	Status Light
Normal	
Discovery	
Getting Topology	
Include infrastructure details	

Note: When the status light is orange or red, you may want to click the text to the left of the light to access discovery logs quickly (**Discovery > View Logs**).

About the Left Pane

The left pane displays the buttons used to access the management server's main utilities.

Important: You may not see all of the following utilities depending on the role assigned to your user account. For example, users assigned to the Help Desk role by default have access to Application Explorer and Event Manager, but not to System Explorer, Provisioning, Policy Manager, and Reporter.

- **Application Explorer** - Application Explorer lets you monitor and display information obtained from discovered databases.
 - **System Explorer** - System Explorer lets you access systems, and view assets by fabric and logical path. See “About System Explorer” on page 296.
 - **Capacity Explorer** - Capacity Explorer provides a graphical representation of an element's storage capacity in the storage network. See “About Capacity Explorer” on page 660.
 - **PerfExplorer** - Performance Explorer provides a graphical representation of the results obtained from monitoring your elements. See “Creating Performance Charts” on page 632.
 - **Event Manager** - Event Manager lets you view, clear, sort and filter API-generated events. See “About Event Manager” on page 536.
 - **Provisioning** - Provisioning assists you in creating zones, zone sets, and zone aliases, in addition to storage pools, volumes, and host security groups. See “About Provisioning” on page 399.
 - **Policy Manager** - Policy Manager lets you set up rules so that an automated response occurs when a particular event happens or a value triggers the system. See “About Policy Manager” on page 682.
 - **Business Tools** - The management server provides advisors and automators. Advisors provides detailed information for you to make informed decisions about your network, such as non-compliant HBA firmware. Automators let you automate particular tasks. The advisors and automators available to you depend upon your release. See “About the Business Tools” on page 772.
 - **Chargeback** - Chargeback helps you manage ownership by department, track costs and assemble business reports. You can view data gathered by Chargeback by element, department and enterprise-wide. You can show the data Chargeback gathered to upper-management through reports, which can be e-mailed on a regular schedule. See “About Chargeback” on page 714 for more information.
 - **Reporter** - Reporter provides detailed reporting on the infrastructure, such as statistics, and usage trends. See “About Reporter” on page 560.
-

Opening and Closing the Left Pane

The management server lets you open and close the left pane. If you want to obtain additional screen space for the main pane, close the left pane by clicking the section between the arrows at the upper-right border for the left pane, as shown in the following figure:



Figure 1-2: Closing the Left Pane

The buttons in the left pane are moved to the top of the page. See the following table for an explanation of the buttons.

Table 1-2: Buttons Displayed When the Left Pane Is Closed

Button	Provides Access to...
	The Home page
	Application Explorer
	System Explorer
	Capacity Explorer
	Performance Explorer
	Event Manager
	Provisioning
	Policy Manager

Table 1-2: Buttons Displayed When the Left Pane Is Closed (Continued)

Button	Provides Access to...
	Business Tools
	Chargeback
	Reporter

To open the left pane, click the section between the arrows in the upper-left border, as shown in the following figure. In some builds, the left pane may be intentionally disabled.




Figure 1-3: Opening the Left Pane

About the Home Page

After you log into the management server, you are shown the Home page. The Home page provides an overview of the main features for the management server. To access a feature, click its icon

Table 1-3: Icons on the Home Page

Feature	Description	Where to Find More Information
 Policy Manager	Policy Manager can e-mail you, generate an event or run a custom script when policy conditions are met. You can define policies to monitor provisioning status, element creation, storage utilization and incoming events.	Chapter 16, “Managing Policies” on page 681.
 Application Explorer	Application Explorer lists the applications and their instances running in the storage area network (SAN) and any events associated with them.	“Accessing Information About Applications” in the Application Guide.

Table 1-3: Icons on the Home Page (Continued)

Feature	Description	Where to Find More Information
 System Explorer	System Explorer shows you the topology of your SAN and gives you the ability to explore details about each element.	Chapter 9, “Viewing Element Topology and Properties” on page 295
 Performance Explorer	Performance Explorer provides detailed performance management capabilities enabling you to visualize what you have and how it is performing.	Chapter 14, “Viewing Performance Data” on page 629
 Capacity Explorer	Capacity Explorer provides a graphical representation of an element's storage capacity and utilization in the storage network.	Chapter 15, “Finding an Element's Storage Capacity” on page 659
 Event Manager	Events Manager keeps you informed of what is happening with your managed elements. Its filter and report format allows you to view, clear, and sort the events you are interested in easily.	Chapter 12, “Event Management” on page 535
 Provisioning	Provisioning make your SAN-zoning tasks and storage management tasks a snap. The format minimizes errors by giving you easy-to-follow instructions and step-by-step screens.	Chapter 10, “Provisioning” on page 399

Table 1-3: Icons on the Home Page (Continued)

Feature	Description	Where to Find More Information
 Chargeback	Chargeback helps you manage departmental ownership, track costs and assemble business reports making inquiries, such as audits and inventory reviews, easier.	Chapter 17, "Chargeback" on page 713
 Business Tools	Business Tools helps you manage the business aspect of your network. You can use Advisors to retrieve important network information, and Automators to automate tasks.	Chapter 18, "Business Tools" on page 771
 Reporter	Reporter provides reports using data collected by the management server. Some of these reports give you enterprise views of your hosts, switches, storage systems or applications while others give you an at-a-glance analysis based on assets, ownerships, chargeback or performance information.	Chapter 13, "Running Reports" on page 559

You may not see all of these features, depending on the following:

- The type of license you have. Depending on your license, all features may not be available. See the "List of Features" to determine if you have access to all features. The "List of Features" is accessible from the Documentation Center (**Help > Documentation Center**).
- The role assigned to your user account. For example, users assigned to the Help Desk role by default have access to Application Explorer and Event Manager, but not to System Explorer, Provisioning, Policy Manager, and Reporter. See "About the Security for the Management Server" on page 154 for more information.

For first-time users, see the following topics:

- Chapter 2, “Discovering NAS Devices, Tape Libraries, Switches and Storage Systems” on page 21
- Chapter 4, “Discovering Applications and Hosts” on page 117
- Chapter 5, “Managing Security” on page 153
- “About System Explorer” on page 296
- “About Event Manager” on page 536
- “About Provisioning” on page 399
- “About Reporter” on page 560

Accessing the Management Server

Keep in mind the following:

- If you do not have a license installed, you are asked to install the license. If you do not have a valid license, contact customer support, as mentioned in the Documentation Center (**Help > Documentation Center**). To install the license, click the **Import License File** button on the Licenses tab (**Configuration > Licenses**).
- Make sure you do not have pop-up blocking software enabled. If your Web browser has an option for blocking pop-ups, disable it. The management server uses pop-ups for dialog boxes.
- If you are using a Web browser on Sun Solaris, you must manually install the Java Plug-in to access several components on the management server. See “Installing the Java Plug-in” on page 14 for more information.
- It may take more time to log into the management server during a topology recalculation.

To access the management server:

1. Type one of the following in a Web browser:

- ☐ For secure connections:

`https://machinename`

where `machinename` is the name of the management server.

To stop receiving a Security Alert message each time you use the HTTPS logon, install the security certificate as described in “Installing the Software Security

Certificate” on page 15. Install the security certificate after you have completed the steps in this chapter.

- For nonsecure connections:

Enter the DNS name of the computer in the URL instead of localhost, even if you are running a Web browser directly on the management server. If you use `https://localhost` to access the management server, you will receive a “Hostname Mismatch” error when you attempt to use System Explorer or Performance Explorer in the management server.

`http://machinename`

where `machinename` is the name of the management server.

2. If you receive an error message when you attempt to connect to the management server, the AppStorManager service might be still starting. Wait for it to complete its start script.

Note: If you see a message resembling the following, see “Receiving “HTTP ERROR: 503” When Accessing the Management Server” on page 787: Receiving HTTP ERROR: 503 javax.ejb.EJBException: null; CausedByException is: Unexpected Error; nested exception is: java.lang.NoClassDefFoundError

3. In the management server login page, type `admin` in the **Name** field and `password` in the **Password** field. Then, click **Login**.
4. If you are shown the software license agreement and you agree with its terms, click the **Accept** button.

Note: To prevent the license agreement from being displayed each time you log on to the management server, select **Do not show me this again**.

5. If the management server does not detect a license, you are asked to import the license. Click the **Import License File** button to install the license. The license file can be obtained from customer support.

Installing the Java Plug-in

Java 2 Runtime Environment is required to access several features in the management server, such as System Explorer. If your Web browser is running on Sun Solaris, you must manually install the Java plug-in.

To install the Java plug-in:

1. Go to the following URL and download the installation file for the Sun JRE when asked:
`http://<management_server>/appiq/j2re-1_4_2_04-solaris-sparc.sh`
where `<management_server>` is the hostname of the management server.
2. In a terminal window, execute the downloaded file in a directory where you want the JRE installed.
This executable installs the Sun JRE on your computer.

The Java plug-in for your Web browser is available in the following file:
`$JRE_HOME/plugin/sparc/ns610/libjava_oji.so`
where `$JRE_HOME` is the directory containing the JRE installation.
3. In a terminal window, go to the `$HOME/.mozilla/plugins` directory. Create a `plugins` directory if it does not exist.
4. Remove any existing links to the Java plug-in in this directory.
5. Create a symbolic link to the Java plug-in by using the following command:
`ln -s $JRE_HOME/plugin/sparc/ns610/libjava_oji.so`

Note: Remember the dot at the end of the command.

6. If you are a root user on the server and you want to make the plug-in available to all users, create a symbolic link in the `plugins` directory under the browser's installation directory, typically `/opt/SUNWns/plugins`.

Note: Any existing plug-ins in a user's home directory take precedence over this system-wide plug-in.

7. Restart your Web browser.
-

Installing the Software Security Certificate

To stop receiving a Security Alert message each time you use the HTTPS logon, install the software security certificate, as described in this section.

- Installing the Certificate by Using Microsoft Explorer 6.0 on page 16
- Installing the Certificate by Using Netscape Navigator 7 on page 17
- Changing the Security Certificate to Match the Name of the Server on page 17

Keep in mind the following:

- Enter the DNS name of the computer in the URL instead of localhost. If you use `https://localhost` to access the management server, you are shown a "Hostname Mismatch" error.
- For customers to receive a trusted certificate, they need to purchase a certificate from a trusted entity (most browsers have trust relationships set up for Verisign, Entrust, and Thawte, among others.) with the Common Name (CN) set to the name of their management server. Also note that the Common Name in the certificate must match the name in the URL.

Installing the Certificate by Using Microsoft Explorer 6.0

1. Access the management server by typing the following:
`https://machinename`
where `machinename` is the name of the management server.
2. When the security alert message appears, click **OK**.
If you do not want the Web browser to warn you about a secure connection at any Web site, select the **In the future, do not show this warning** option.
3. When you are told there is a problem with the site's security certificate, click the **View Certificate** button.
4. When you are shown the certificate information, click the **Install Certificate** button at the bottom of the screen.
5. When you are shown the Certificate Import Wizard, click **Next** to continue the installation process.
6. Select one of the following:

- ☐ **Automatically select the certificate store based on the type of certificate** - This option places the certificate automatically in the appropriate location.
 - ☐ **Place all certificates in the following store** - This option lets you pick the store where the certificate will be stored.
7. Click **Finish**.
 8. When you are asked if you want to install the certificate, click **Yes**.
You are shown the following message when the certificate is installed.

Installing the Certificate by Using Netscape Navigator 7

1. Access the management server by typing the following:
`https://machinename`
where `machinename` is the name of the management server.
2. When the security alert message appears, click the **Always** button.
3. When you are told you are requesting an encrypted page, click **OK**.
4. Click the **Always** button when you are asked if you want to accept the certificate.
5. When asked if you wanted to trust the signed applet, click the **Always** button.

Changing the Security Certificate to Match the Name of the Server

If your users are shown a Security Alert window with the following message, you might want to modify the security certificate so users feel more comfortable with installing the certificate:

"The name of the security certificate is invalid or does not match the name of the site."

You can change the security certificate so that users receive the following message instead:

"The security certificate has a valid name matching the name of the page you are trying to view."

When you change the certificate, you must use the `generateAppiqKeystore.bat` program to delete the original certificate. Then, use the `generateAppiqKeystore.bat` program to create a new certificate and to copy the new certificate to the management server.

To change the certificate:

1. Go to the %MGR_DIST%\Tools directory.
2. To delete the original certificate, enter the following at the command prompt:
`%MGR_DIST%\Tools> generateAppiqKeystore.bat del`
The original certificate is deleted.
3. To create a new certificate containing the DNS name of the management server, enter the following at the command prompt:
`%MGR_DIST%\Tools> generateAppiqKeystore.bat`
4. If the program is unable to detect a DNS name, enter the following at the command prompt:
`%MGR_DIST%\Tools> generateAppiqKeystore.bat mycomputername`
where mycomputername is the DNS name of the computer
5. To copy the new certificate to the management server, enter the following at the command prompt:
`%MGR_DIST%\Tools> generateAppiqKeystore.bat copy`
The new certificate is copied to the correct location.

Stopping the Service for the Management Server

By default, the software runs as a service from the time the management server has restarted. If you must stop the service, it is recommended that you stop the service from the Services window, rather than from the command prompt window.

Important: The service must be running for users to access the management server.

To restart the AppStorManager service on a management server:

On Microsoft Windows:

1. Go to the **Services** window, usually accessible from the Control Panel.
 2. Right-click **AppStorManager**.
 3. Select **Stop** from the drop-down menu.
-

4. To start the management server, right-click **AppStorManager** and select **Start** from the drop-down menu.

On Sun Solaris:

1. Open a command prompt window.
2. Enter the following at the command prompt to stop the management server:
`/etc/init.d/appstormanager stop`
3. To start the management server, enter the following at the command prompt:
`/etc/init.d/appstormanager start`

Signing Out of the Management Server

Sign out of the management server to prevent unauthorized users from accessing the SAN.

To sign out, click the **Sign Out** link in the upper-left corner of the page.

Discovering NAS Devices, Tape Libraries, Switches and Storage Systems

Before you can use the management server, you must make the software aware of the elements on your network. An element is anything on the network that can be detected by the management server, such as a switch. This is done through the discovery process. Discovery obtains a list of discovered elements and information about their management interface and dependencies. The management server can discover only elements with a suitable management interface. Refer to the support matrix for supported hardware.

This chapter describes the following:

- Discovery Steps on page 22
 - Overview of Discovery Features on page 25
 - Step 1 - Discover Switches on page 34
 - Step 2 - Discover Storage Systems, NAS Devices and Tape Libraries on page 62
 - Step 3 - Build the Topology on page 88
 - Step 4 - Get Details on page 93
 - Troubleshooting Mode on page 98
 - Managing McDATA and EMC Connectrix Switches on page 99
 - Assigning a File Extension in Netscape 7 on page 102
 - Filtering Discovery Groups on page 103
 - Moving Elements to Another Discovery Group on page 103
 - Placing an Element in Quarantine on page 104
-

- Removing an Element from Quarantine on page 105
- Updating the Database with Element Changes on page 105
- Notifying the Software of a New Element on page 106

Discovery Steps

Discovery for switches, storage systems, tape libraries and NAS devices consists of several steps:

1. Discover your switches. See “Step 1 - Discover Switches” on page 34.
2. Discover your storage systems, tape libraries, and NAS devices. See “Step 2 - Discover Storage Systems, NAS Devices and Tape Libraries” on page 62.
3. (Optional) If you want to view the topology quickly in System Explorer, obtain the topology as described in “Step 3 - Build the Topology” on page 88. Keep in mind this step only gathers the information necessary for displaying the topology.
4. Perform Get Details. Get Details is required to obtain detailed information from the elements you discovered. You must run Get Details to obtain provisioning information. This step takes some time to complete. Run Get Details when the network is not busy. See “Step 4 - Get Details” on page 93 for more information.

Important: Before performing any provisioning operations that involve a Brocade switch you must perform Discovery Step 3 for any subset of elements that includes this Brocade switch.

Overall Discovery Tasks

Keep in mind the following:

- To save time, make sure the user names and passwords are correct. When credentials are not supplied, the default user names and passwords are tried for the element.
 - After you discover an EMC Connectrix or McDATA switch, the IP address displayed next to the name of the switch is actually the IP address of the service processor for the switch in the Get Details screens. To find the IP address of the switch, click the link for the switch in one of the following screens:
 - Topology screen (**Discovery > Topology**).
-

- Get Details screen (**Discovery > Details**).

Then, click the **Properties** tab. The Properties tab can also be accessed by double clicking the switch in System Explorer. Complete the steps in this chapter before you try to find the IP address of the switch.

- If you are having a problem with discovering an element, try enabling Troubleshooting Mode. See “Troubleshooting Mode” on page 98 for more information. If you are still having a problem, see Chapter 20, “Troubleshooting” on page 785.
- The management server does not display additional information about excluded elements in the user interface; however, the IP addresses of excluded elements appear in the following locations:
 - Discovery screen (**Discovery > Setup**).
 - Topology screen (**Discovery > Topology**).
 - Get Details screen (**Discovery > Details**).

The management server, however, does mention in the logs (**Discovery > View Logs**) that a provider instance has been created for an excluded element. You can ignore this message that appears in the logs.

- To obtain information about the storage area network (SAN), include in the discovery the IP addresses for the following:
 - **Fibre channel switch** The fibre channel switch contains a list of all elements within the fabric. The management server obtains a detailed listing of all elements connected to the switch fabric.
 - **A host containing a Host Bus Adapter (HBA)** All fibre channel host adapters look for available elements attached to the HBA. This information is gathered by CIM Extensions and sent to the management server. Since you have not installed CIM Extensions yet, the management server obtains limited information on the hosts when you perform discovery this time around.
 - **A proxy connected to the SAN** - Include a proxy that has a direct connection or a SAN connection to the management server. An example of a proxy is the EMC Solutions Enabler or Hitachi HiCommand Device Manager. Engenio storage systems do not require a proxy, as they can be accessed directly. Make sure the proxy service has started. On a computer running Windows, this can be determined by looking in the Services window. EMC Solutions Enabler version 5.1 requires additional steps for discovery. See “Discovering EMC Solutions Enabler 5.1” on page 66 for more information.

Note: Before performing any provisioning operations that involve a Brocade switch you must perform Discovery Step 3 for any subset of elements that includes this Brocade switch.

To make the management server aware of elements on your network follow the steps in the following table.

Table 2-1: Discovery Steps for Switches, NAS Devices, and Storage Systems

Step	Description
1	If you have several switches and storage systems that use the same password and user name, set that password and user name as the default. See “Setting Default User Names and Passwords” on page 25 for more information.
2	Discover your switches. See Table 2-2, “Discovery Requirements for Switches,” on page 35 for information on how to discover the types of switches in your network. One way to detect multiple IP addresses at once is to add an IP range for scanning. The management server scans the IP range for elements and populates the discovery list with element it could contact. You can then discover those elements. See “Adding an IP Range for Scanning” on page 28 for more information.

Table 2-1: Discovery Steps for Switches, NAS Devices, and Storage Systems (Continued)

Step	Description
3	Discover your NAS devices and storage systems. See Table 2-7, “Discovery Requirements for Storage Systems, Tape Libraries & NAS Devices,” on page 63.
4	Get Details (Discovery > Details) is required to obtain information from your discovered elements. Obtaining details takes some time. You might want to perform this process when the network and the managed elements are not busy. You can obtain a picture of device connectivity quickly, by clicking the Get Topology button on the Topology tab. See the following topics for more information: ■ “Step 3 - Build the Topology” on page 88 ■ “Step 4 - Get Details” on page 93

Overview of Discovery Features

With discovery, you can:

- Provide up to three default user name and passwords for discovery.
- Import pre-existing discovery lists, so you do not need to re-enter discovery information.
- Save your existing discovery list.
- Modify a discovery entry.
- Remove elements from a discovery list.

This section describes the following:

- Setting Default User Names and Passwords on page 25
 - Adding an IP Range for Scanning on page 28
 - Adding a Single IP Address or DNS Name for Discovery on page 29
 - Modifying a Single IP Address Entry for Discovery on page 31
 - Removing Elements from the Addresses to Discover List on page 32
-

Setting Default User Names and Passwords

You can specify up to three default user names and passwords. If several of the elements in the same domain use the same user name and password, assign the user name and password as the default. The management server uses the default user names and passwords if a user name and password are not assigned to an element in the **Setup** screen.

For example, if you have several hosts using the same user name and password, you could enter the default user name and password. If one of the hosts is connected to a storage system with another user name and password, you would also enter this user name and password. Do not specify the user name and password for the storage system in the individual range because that overrides the default user name and password.

To access a Windows-based device, prefix the user name with `domain_name\`, as shown in the following example. This is required by the Windows login mechanism.

```
domain_name\user_name
```

where

- `domain_name` is the domain name of the element
- `user_name` is the name of the account used to access that element

To save time, make sure the user names and passwords are correct. The software tries each of the default user names and passwords whenever it finds an element.

To add the default user name and passwords:

1. Click **Discovery**, then click **Setup** in the upper-right pane of the window.
 2. Under Discovery Setup, select **Step 1** at the top of the screen.
 3. Click the link, **Set Default User Name and Password**.
 4. In the **User Name** field, type the user name for one or more elements.
 5. In the **Password** field, type the corresponding password for the user name typed in the previous step.
 6. In the **Verify Password** field, retype the password.
 7. Repeat steps 4 through 6 for other default user names and passwords you want to add.
 8. Click **Add System**.
-

Setting User Names and Passwords

You can specify up to three user names and passwords. These user names and passwords are used during discovery if your IP Address does not have a user name and password specified.

If you are specifying a user name for a Windows host, prepend the user name with the Windows domain name.

For example: **mydomain\user**

User Name:

Password:

Verify Password:

User Name:

Password:

Verify Password:

User Name:

Password:

Verify Password:

OK Cancel Help

Figure 2-1: Setting Default User Names and Passwords

Adding an IP Range for Scanning

Instead of adding each IP address individually for scanning, the management server can detect a range of IP addresses, automatically populating the list of elements to be discovered.

Keep in mind the following:

- Include in the scanning, the proxy server that has a direct connection or a SAN connection to the management server, such as the EMC Solutions Enabler. Make sure the proxy service has started. On a computer running Microsoft Windows, this can be determined by looking in the **Services** window.
- You cannot scan an IP range to discover an instance of HiCommand Device Manager that listens on a port other than 2001. The management server does not allow port numbers in the scanning of IP ranges, and thus, you are not able to specify the port. See “Discovering HDS Storage Systems” on page 71 for more information.
- Enter a range within the same subnet. The management server cannot scan IP ranges across subnets.

To add an IP address range to scan:

1. Click **Discovery**, then click **Setup** in the upper-right pane of the window.
 2. Click the **IP Ranges** tab.
The IP ranges already added are listed.
 3. Click the **Add Range** button.
 4. In the **From IP Address** field, type a lowest IP address in the range to be scanned.
 5. In the **To IP Address** field, type the highest IP address in the range to be scanned.
 6. In the **User Name (Optional)** field, type a common user name for elements in the IP range.
 7. In the **Password (Optional)** field, type a common password for elements in the IP range.
 8. In the **Verify Password** field, retype the password.
 9. In the **Comment** field, type a brief description of the servers. For example, “Servers in Marketing.”
 10. Click **OK**.
-

Add Range for Scanning

If you are specifying a user name for a Windows host, you can prepend the user name with the Windows domain name.

For example, **mydomain\user**

From IP Address:*

To IP Address:*

User Name:

Password:

Verify Password:

Comment:

* required fields

OK Cancel Help

Figure 2-2: Adding an IP Range for Scanning

11. Click the **Start Scanning** button on the **IP Ranges** tab.
The management server scans the IP range and populates the **Addresses to Discover** table on the **IP Addresses** tab.

Adding a Single IP Address or DNS Name for Discovery

Keep in mind the following for discovering Engenio storage systems:

- Discover both controllers for the Engenio storage system. Each controller has its own IP address. In Step 1 of discovery, specify all the IP addresses for all the controllers (usually two). The management server discovers these controllers as one single storage system.

- The following steps provide general information on how to discover an element. See Table 2-2, “Discovery Requirements for Switches,” on page 35, Table 2-7, “Discovery Requirements for Storage Systems, Tape Libraries & NAS Devices,” on page 63, and Table 4-1, “Making the Management Server Aware of Hosts,” on page 119.
- To obtain drive-related statistics, install a proxy host. Ensure the proxy host has at least one LUN rendered by each controller of the array. See “Obtaining Disk Drive Statistics from Engenio Storage Systems” on page 150 for more information.
- The **Do Not Authenticate** option is for Engenio storage systems. The Engenio storage systems do not require a password for Get Details. If you want do not want to use the management server for provisioning on Engenio storage systems, you can leave the password field blank and select the **Do Not Authenticate** option. The management server will still monitor the Engenio storage system; however, you will not be able to do provisioning tasks.

To add a single IP address or DNS name to discover:

1. Click **Discovery**, then click **Setup** in the upper-right pane of the window.
 2. Under Discovery Setup, select **Step 1** at the top of the screen.
 3. On the **IP Addresses** tab, click the **Add Address** button.
 4. In the **IP Address/DNS Name** field, type the IP address or DNS name of the device you want to discover.
 5. In the **User Name (Optional)** field, type the user name. If you are discovering an Engenio storage system, leave this field blank.
This field can also be left blank if the element's user name and password are one of the default user names and passwords.
 6. Do one of the following for the password:
 - ☐ If you do not want to do provisioning on a storage system, leave the **Password** field blank. For Engenio storage systems, you must also select the **Do Not Authenticate** option.
 - ☐ If you want to do provisioning on a storage system, type the corresponding password for controller or proxy and make sure the **Do Not Authenticate** option is not selected.
 - ☐ For all other elements other than storage systems, provide the password if it is necessary for authentication. If the element does not require a password, leave the **Password** field blank.
 7. If you typed a password in the previous step, retype the password in the **Verify Password** field.
 8. In the **Comment** field (optional), type a comment for additional information. The information typed into this field is displayed under the Comment column in the Addresses to Discover list (**Discovery > Setup**).
-

9. Click **OK**.
10. To start discovering elements on the network, click the **Start Discovery** button on the **IP Addresses** tab.

Modifying a Single IP Address Entry for Discovery

You can change the user name and password the software uses to access an element. Whenever a user name and/or password has changed on an element the management server monitors, the management server must be made aware of the change. For example, assume the password for a host was changed. You would need to update the management server database with the new password.

Important: These steps only change the user name and password stored in the database. It does not change the device's user name and password.

To modify a user name or password for discovery:

1. Click **Discovery**, then click **Setup** in the upper-right pane of the window.
2. Click the  button for the element you want to modify the user name and/or password.
3. To change the user name, type the new user name in the **User Name** field.
4. To add or change a comment, type a comment in the **Comment** field.
5. To change the password:
 - a. Click the **Change password** button.
 - b. Type the new password in the **New Password** field.
 - c. Type the password again in the **Verify Password** field.
 - d. Click **OK** in the Change Password page.
6. Click **OK** in the Edit Address for Discovery page.
7. Select the option, **Step 2 - Topology: Select the discovered elements and build the topology view**.
8. Select the element for which you changed the user name and/or password.
9. Click the **Get Topology** button.

The software updates its database with the new user name and/or password.

Removing Elements from the Addresses to Discover List

When you remove IP addresses and/or ranges from the Addresses to Discover list, the elements associated with those IP addresses are not removed from the management server. Only the information that was used to discover them is removed.

To remove items from the Discovery list:

1. Click the **Discovery** icon in the upper-right pane of the home page.
2. Click **Setup**.
3. Select **Step 1** at the top of the page.
4. Do one of the following:
 - ☐ Select the IP addresses and/or IP ranges you want to remove from the list. Then, click the **Delete** button.
 - ☐ Click the  button corresponding to the elements you want to remove from the Addresses to Discover list.

Important: The elements associated with these addresses are not removed from the management server. See “Deleting Elements from the Product” on page 91 for information about how to remove an element from the management server.

Importing Discovery Settings from a File

If you have a previous discovery list you can import it, rather than re-entering the information.

The import discovery settings feature lets you import the following information to the **Discovery** list:

- IP addresses to discover
- Default user names and passwords, which are encrypted
- Discovery information for applications

To prevent re-entering the information for each management server instance, you can import the same file for multiple management server instances.

Important: When you import a file, your previous settings are overwritten.

If you are shown an error message when you try to import the discovery settings, verify you are using the right password. If you are using the correct password, there could be a possibility that the file is corrupt.

To import discovery settings from a file:

1. Click **Discovery**, then click **Setup** in the upper-right pane of the window.
2. Click the **Import Settings from File** link.
3. In the **Import Settings from File** window, do one of the following:
 - ☐ Click the **Browse** button to find the file.
 - ☐ In the **Filename** field, type a complete path to the file.
4. In the **Password** field, type the password for the management server.
5. Click **OK**. The information on the following tabs is updated:
 - ☐ IP Addresses
 - ☐ IP Ranges
 - ☐ Applications
 - ☐ Windows Proxy tab

Saving Discovery Settings to a File

After you have discovered your elements, save the discovery settings of the elements in your discovery list.

The **Save Settings to File** link on the **Discovery Targets** tab lets you save the following:

- IP addresses to discover
- Default user names and passwords, which are encrypted
- Oracle TNS Listener ports
- Microsoft Exchange configuration

To prevent re-entering the information for each instance of the management server, you can import the file for multiple instances.

To save the discovery settings to a file:

1. Click **Discovery**, then click **Setup** in the upper-right pane of the window.
2. Click **Setup** in the upper-right corner.
3. Click the **Save Settings to File** link.
4. In the **Password** field, type the password for the management server.
5. In the **Verify Password** field, type the password from the previous step. Then, click **OK**.
6. When you are asked if you want to open or save the file, do one of the following:
 - ☐ **Microsoft Internet Explorer** - Click the **Save** button.
 - ☐ **Netscape 7** - Select the **Save this file to disk** option.The Downloading window appears.
7. Type a name for the *.xml file and select the directory to which you want to save the file. The name of the file is `DiscoverySettings.xml` by default.

Important: Netscape 7 assumes the file is an HTML file. If you are running Netscape, make sure the file type is selected to **All Files** and the file extension is xml. You can make Netscape 7 recognize the xml file next time, by clicking the **Advanced** button See the topic, “Assigning a File Extension in Netscape 7” on page 102 for more information.

8. Leave the **Password** field blank, a password is not required. If you enter a password, that password is required later when you import the file. Leaving the password field blank allows the file to be imported without a password.
9. Click the **Save** button in the Save As window. The file is saved.

Step 1 - Discover Switches

This section describes the following:

- “SMI-S Switches Must Be Removed and Rediscovered After Upgrading” on page 37
 - “Discovering Brocade Switches” on page 37
 - “Discovering CNT Switches” on page 42
 - “Discovering Cisco Switches” on page 43
-

- “Discovering Sun StorEdge and QLogic Switches” on page 45
- “Changing the SNMP Trap Listener Port for Sun StorEdge Switches” on page 46
- “Discovering McDATA and EMC Connectrix Switches” on page 47
- “Excluding McDATA and EMC Connectrix Switches from Discovery” on page 58

The following table provides an overview of the discovery requirements for switches.

Table 2-2: Discovery Requirements for Switches

Element	Discovery Requirements	Additional Information
Brocade switches (Fabric Access API)	Enter the IP address/DNS name, user name and password of the Brocade switch to discover it. The user name (default admin) and password must be for the Admin Account.	See “Discovering Brocade Switches” on page 37.
Brocade switches (SMI-S*)	Enter the IP address/DNS name, user name and password of the Brocade SMI-S proxy server to discover it. The user name (default admin) and password must be for the Admin Account. Note that some SMI-S providers require you to specify the server using its IP address: <code>http://IPADDRESS</code> .	See “Discovering Brocade Switches” on page 37.
CNT switches	Enter the IP address and the port number for the InVsn Software that manages the switch as well as the user name and password.	See “Discovering CNT Switches” on page 42.
Cisco switches (SMI-S)*	Enter the IP address/DNS name of the Cisco switch as well as the user name and password of the switch. All SMI-S switches require a user name and password.	See “Discovering Cisco Switches” on page 43.

Table 2-2: Discovery Requirements for Switches (Continued)

Element	Discovery Requirements	Additional Information
Cisco switches (SNMP)	Enter the IP address/DNS name of the Cisco switch. Enter the SNMP read-only community string as the user name. You do not need to enter a password.	See “Discovering Cisco Switches” on page 43.
QLogic switches (SMI-S)*	Enter the IP address/DNS name of the QLogic SMI-S switch as well as the user name and password of the switch. All SMI-S switches require a user name and password.	See “Discovering Sun StorEdge and QLogic Switches” on page 45.
Sun StorEdge and QLogic switches (SNMP)	Enter the IP address/DNS name of the Sun StorEdge or QLogic switch. Enter the SNMP read-only community string as the user name. You do not need to enter a password.	See “Discovering Sun StorEdge and QLogic Switches” on page 45.
McDATA and EMC Connectrix switches	Additional steps are required for discovering these switches, and the steps vary according to your network configuration.	See “Discovering McDATA and EMC Connectrix Switches” on page 47.

* SMI-S switches must be removed and rediscovered after upgrading from builds earlier than 4.2. See “SMI-S Switches Must Be Removed and Rediscovered After Upgrading” on page 37 for important information. See “Important Information About Upgrading and Brocade Switches” on page 39 for specific details related to Brocade switches.

Important: Make sure pop-up blocking software is disabled. If your Web browser has an option for blocking pop-ups, disable it. The management server uses pop-ups for dialog boxes.

SMI-S Switches Must Be Removed and Rediscovered After Upgrading

You must remove and rediscover Brocade, Cisco, QLogic, or CNT switches supported through SMI-S after you upgrade the management server.

1. Remove the Brocade, Cisco, QLogic, or CNT switches supported through SMI-S . See “Deleting Elements from the Product” on page 91.
2. Rediscover the Brocade, Cisco, QLogic, or CNT switches. See “Step 1 - Discover Switches” on page 34.
3. Perform Get Details. See “Step 4 - Get Details” on page 93.

Discovering Brocade Switches

If you are upgrading to build 5.1 from any build prior to 4.2 of the management server your Brocade switches are discovered using the Brocade Fabric Access API after the upgrade. See “Important Information About Upgrading and Brocade Switches” on page 39 for upgrade details. If this is a new installation of the management server software, your Brocade switches are discovered using the Brocade SMI-S (Storage Management Initiative Specification) provider and you must download and install the Brocade SMI-S provider software from the following FTP site: ftp://ftp.compaq.com/pub/products/storageworks/smisproviders/brocade_provider.pdf.

Table 2-3: Brocade Discovery Methods

Management Server Installation Status:	Default Brocade Discovery Method:
New installation of Build 5.1	SMI-S (You must download the Brocade SMI-S provider from the following site: ftp://ftp.compaq.com/pub/products/storageworks/smisproviders/brocade_provider.pdf .)
Upgrade from Build 4.2	Fabric Access API (The Brocade Fabric Access API provider installed with build 4.2 of the management server is used.)

Important: When using the Brocade Fabric Access API, Brocade recommends that the switch with the most recent version of Brocade API firmware (newer than v.2.6.x) is configured as the proxy switch.

When discovering Brocade SMI-S hosts the IP address of the Brocade SMI-S proxy server must be entered for Discovery Step 1. Be sure to see the Support Matrix for the latest details on firmware requirements.

Brocade SMI-S Provider Installation Requirements (New Installations of the Management Server)

If this is a new installation of the management server, the management server is configured to discover Brocade switches through SMI-S, by default, as noted earlier. You must first download and install the Brocade SMI-S provider software from the following site:

```
ftp://ftp.compaq.com/pub/products/storageworks/smisproviders/  
brocade_provider.pdf
```

Be sure to check this FTP site periodically to be sure you are running a current version of the Brocade SMI-S provider. See the Support Matrix for details.

Keep in mind the following for Brocade SMI-S switches:

- You must verify that the Rapid program on the switch is set to 1. Rapid must be set to 1 so that the management server or any Brocade SMI-S agent can communicate with the switch. Be sure to see “Verifying Brocade Rapid Program Is Set to 1” on page 39 for more information. Note that this requirement also applies to Brocade switches that are configured to use the Fabric Access API provider.
 - Before performing any provisioning operations that involve a Brocade switch after an upgrade to build 4.2 or 5.1 from any build prior to 4.2, you must perform Discovery Data Collection/Get Details for any subset of elements that includes this Brocade switch. See “Step 4 - Get Details” on page 93 for details on running Discovery Data Collection/Get Details.
-

Verifying Brocade Rapid Program Is Set to 1

If you are discovering Brocade switches or Brocade SMI-S agents, verify that the Rapid program on the switch is set to 1.

1. (Optional) Set the command prompt window so that it displays many rows.
While completing the following steps, the command prompt window displays a large amount of data. You might want to expand the size and buffer of the command prompt window. To do this in Microsoft Windows 2000, click the upper-right corner of the command prompt window, click the **Layout** tab, and then modify the options under Screen Buffer Size and Window Size.
2. Access the Brocade switch or SMI-S agent by using the telnet option. For example,
`telnet`
`open 10.1.213.228`
where 10.1.213.228 is the IP address of the switch.
3. When prompted for the user name and password, supply them.
4. Type the following to see what is supported on the switch or Brocade SMI-S agent:
`configshow`
The output is displayed in a page-by-page format.
5. Select all of the output.
6. Paste the output in a text editor, for example Notepad. Use the Find command to search for `rpc.rapid`.
7. Verify Rapid is set to one, as displayed below:
`rpc.rapid: 1`

About Brocade Discovery (Specifying Brocade Discovery Using Fabric Access API or SMI-S)

For new installations, SMI-S is the default Brocade discovery protocol with this release of the management server as noted earlier. For existing installations, Fabric Access API is the default Brocade discovery protocol with this release of the management server. You can change from SMI-S to Fabric Access API by following the steps below:

To change the Brocade discovery settings to Fabric Access API on the management server:

1. Click **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree.

2. In the **Custom Properties** field type `cimom.brocade.useApi=true` to use Fabric Access API (or you can type: `cimom.brocade.useApi=false` to use SMI-S). If the `cimom.brocade.useApi` value is not specified the management server uses SMI-S.
3. When you are done, click **Save**.
4. Restart the service for the management server for your changes to take effect. While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements. See the help for your operating system platform for details on how to verify that the service is running.

Important: When using the Brocade API, Brocade recommends that the switch with the most recent version of Brocade API firmware (newer than v.2.6.x) is configured as the proxy switch.

When discovering Brocade SMI-S hosts the IP address of the Brocade SMI-S proxy server must be entered for Discovery Step 1. Be sure to see the Support Matrix for the latest details on firmware requirements.

To discover Brocade switches:

1. Click **Discovery**, then click **Setup** in the upper-right pane of the window.
 2. Select **Step 1** at the top of the page.
 3. Click the **IP Addresses** tab.
 4. Click the **Add Address** button.
 5. In the **IP Address/DNS Name** field, do one of the following depending on the provider:
 - ☐ **SMI-S**
Type the IP address of the proxy server that is running the SMI-S agent. (Some proxy servers require the following format: `http://IPADDRESS`.)
 - ☐ **Fabric Access API**
Type the IP address or DNS name of the Brocade switch you want to discover.
 6. In the **User Name** field, do one of the following depending on the provider:
 - ☐ **SMI-S**
Type the user name for the SMI-S proxy server.
This field can be left blank if one or more of the following conditions are fulfilled:
-

- The element's user name and password are one of the default user names and passwords.
 - The element does not require authentication.
- ☐ **Fabric Access API**
- This field can be left blank if one or more of the following conditions are fulfilled:
- The element's user name and password are one of the default user names and passwords.
 - The element does not require authentication.
7. In the **Password** field, do one of the following:
- ☐ **SMI-S**
- Type the password for the SMI-S proxy server.
- This field can be left blank if one or more of the following conditions exists:
- The proxy server's user name and password are one of the default user names and passwords.
 - The proxy server does not require authentication.
- ☐ **Fabric Access API**
- Type the corresponding password for the switch.
- This field can be left blank if one or more of the following conditions are fulfilled:
- The element's user name and password are one of the default user names and passwords.
 - The element does not require authentication.
8. If you typed a password in the previous step, retype the password in the **Verify Password** field.
9. In the **Comment** field (optional), type a comment for additional information. The information typed into this field is displayed under the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. To start discovering elements on the network, click the **Start Discovery** button on the **IP Addresses** tab.

Discovering CNT Switches

The management server uses the CNT SMI-S provider to discover CNT switches. A provider is a small software program that is used by the management server to communicate with a device, such as a switch.

This provider communicates with CNT InVsn Enterprise Manager to obtain information about the switch. The provider requires a certain version of InVsn depending on the switch model. See the following table for more information.

Table 2-4: Required Switch Models and InVsn Versions for Discovery

Switch Model	InVsn Software Version
FC/9000	9.0 or later
UMD	9.5 or later

Important: The InVsn credentials are used by the SMI-S provider. Make sure the SMI-S provider is enabled as described in the steps in this section.

Keep in mind the following for CNT switches:

- SNMP is not supported for CNT switches. If you discovered CNT switches in builds earlier than 4.0, you will need to rediscover the CNT switches by using the steps in this section.
- CNT InVsn Enterprise Manager must be running for the management server to discover it.
- The management server does not support provisioning for CNT switches. Only the active zone set and its zone members are reported.
- No ports are reported for uninstalled blades or GBICs.

To discover CNT switches

1. Before you can discover a CNT switch, you must do the following in the CNT InVsn Enterprise Manager software:
 - a. Open the file `ProductInfo.ini` in a text editor, such as Notepad. If the software was installed in the default directory, this file should be in the following directory:
`\Program Files\CNT\inVSN_EM`
 - b. Change the following entry in the file:
-

```
cimomenabled=TRUE
```

- c. Save the file, then restart the InVsn software.
2. In the **IP Address/DNS Name** field, type the primary IP address of the host running the InVsn software you want to discover followed by its namespace and port number, as shown in the following example:

```
192.168.10.76//root/cntfabric:5989
```

where

 - 192.168.10.76 is the IP address of the host running the InVsn software
 - //root/cntfabric is the namespace
 - 5989 is the port number
3. In the **User Name** field, type the user name for the login to the InVsn software.
4. In the **Password** field type the password for the login to the InVsn software.
5. In the **Verify Password** field type the password you provided previously.
6. Click **Start Discovery**.

Discovering Cisco Switches

The management server discovers Cisco switches through SNMP (Simple Network Management Protocol) and SMI-S (Storage Management Initiative Specification) connections—depending on the model of the Cisco switch. See the Support Matrix for details on supported switch models and firmware revisions.

Note: The Cisco MDS 9020 Multilayer Fabric Switch does not support VSANs.

Keep in mind the following for Cisco SNMP switches:

- When you discover a Cisco SNMP switch, you do not need to provide a password. .
- You can view zones, zone sets and zone aliases on a Cisco switch; however, you cannot use the management server to create, modify or remove them from a Cisco switch.
- The management server gathers information about the Cisco inactive database during Get Details . You can change the amount of information that is collected by modifying a property. See “Changing the Amount of Information Collected from the Inactive Zone Database (Cisco Switches)” on page 424.

- The management server groups active zone sets in all Virtual SANs (VSANs) in a fabric into a zone set called “ACTIVE”, and the “ACTIVE” zone set is shown associated with the physical fabric. The members of the “ACTIVE” zone set (zones, zone sets, zone aliases) have the name of the VSAN prefixed to their name. For example, an active zone named “ZONE1” from a VSAN named “VSAN1” is displayed as a zone on the physical fabric with name “VSAN1:CISCO1:ZONE1”.

Keep in mind the following for Cisco SMI-S switches:

- When you discover a Cisco SMI-S switch you need to provide a user name and password. All SMI-S switches require a user name and password.
- Cisco switches on the fabric are displayed without their names — each Cisco SMI-S switch name on the fabric is replaced with a generic name such as: Switch_1401 after running Get Details.

To view the correct names for the Cisco SMI-S switches after running Get Details, manually enter the correct name for the Cisco SMI-S switches.

To discover a Cisco switch:

1. Click **Discovery**, then click **Setup** in the upper-right pane of the window.
 2. Select **Step 1** at the top of the page.
 3. Click the **IP Addresses** tab.
 4. Click the **Add Address** button.
 5. In the **IP Address/DNS Name** field, type the DNS name or primary IP address of the Cisco switch you want to discover.
 6. Do one of the following:
 - For **Cisco** switches with **SNMP** connections:
In the **User Name** field, type the user name for the switch. This is the public community SNMP string for the switch. This field can be left blank if the element's user name and password are one of the default user names and passwords.
 - For **Cisco** switches with **SMI-S** connections:
In the **User Name** field, type the user name for this SMI-S switch.
 7. Do one of the following:
 - For **Cisco** switches with **SNMP** connections:
Leave the **Password** field blank.
 - For **Cisco** switches with **SMI-S** connections:
In the **Password** field, type the password for this SMI-S switch.
 8. Do one of the following:
 - For **Cisco** switches with **SNMP** connections:
-

Leave the **Verify Password** field blank.

- For **Cisco** switches with **SMI-S** connections:
In the **Verify Password** field, type the password of the SMI-S switch again.

Discovering Sun StorEdge and QLogic Switches

The management server discovers Sun StorEdge switches through an SNMP connection and QLogic switches—depending on the model—are discovered through SNMP or SMI-S. See the Support Matrix for details on supported switch models and firmware revisions.

Keep in mind the following for Sun StorEdge and QLogic SNMP switches:

- When you discover a Sun StorEdge or QLogic SNMP switch, you do not need to provide a password.
- The management server does not support provisioning for Sun StorEdge and QLogic switches. Only the active zone set and its zone members are reported.
- To manage a fabric of Sun StorEdge and/or QLogic switches, every switch in the fabric must be included in the discovery list. If a switch is not included in the discovery list, it may show up as a generic host system.
- No ports are reported for uninstalled blades or GBICs.
- The default SNMP trap listener port for all Sun StorEdge switches is 162. To change this port, see “Changing the SNMP Trap Listener Port for Sun StorEdge Switches” on page 46.
- To receive events from Sun StorEdge switches, verify the SNMP trap community string is set to public in SANbox Manager or via telnet. Also, make sure the SNMP traps are configured to be sent to the management server. See “About Event Manager” on page 536 for more information.

Keep in mind the following for QLogic SMI-S switches:

- A user name and password are required to discover any SMI-S switch.
- You must perform Get Details to obtain all available information from QLogic SMI-S switches—otherwise, attributes such as vendor, fabric, and port information will be missing for the QLogic SMI-S switches.

To discover Sun StorEdge or QLogic switches:

1. Click **Discovery > Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.

4. Click the **Add Address** button.
5. In the **IP Address/DNS Name** field, type the DNS name or primary IP address of the Sun StorEdge or QLogic switch you want to discover.
6. Do one of the following:
 - ☐ For **Sun StorEdge** and **QLogic** switches with **SNMP** connections:
In the **User Name** field, type the user name for the switch. This is the public community SNMP string for the switch. This field can be left blank if the element's user name and password are one of the default user names and passwords.
 - ☐ For **QLogic** switches with **SMI-S** connections:
In the **User Name** field, type the user name for this switch. All SMI-S switches require a user name and password.
7. Do one of the following:
 - ☐ For **Sun StorEdge** and **QLogic** switches with **SNMP** connections:
Leave the **Password** field blank.
 - ☐ For **QLogic** switches with **SMI-S** connections:
In the **Password** field, type the password for this switch.
8. Do one of the following:
 - ☐ For **Sun StorEdge** and **QLogic** switches with **SNMP** connections:
Leave the **Verify Password** field blank.
 - ☐ For **QLogic** switches with **SMI-S** connections:
In the **Verify Password** field, type the password of the switch again.

Changing the SNMP Trap Listener Port for Sun StorEdge Switches

The default SNMP trap listener port for all Sun StorEdge switches is 162. To change this port for all switches that are discovered through SNMP, modify the `cimom.snmpTrapListenerPort` property as described in the following steps:

1. Select **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree.
 2. Click **Show Default Properties** at the bottom of the page.
 3. Copy the `cimom.snmpTrapListenerPort` property. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
 4. Return to the Advanced page (**Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree).
-

5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Make your changes in the **Custom Properties** field. Make sure the property is not commented out by removing the hash (#) symbol in front of the property.
7. Set the `cimom.snmpTrapListenerPort` property to the port you want, as shown in the following example:
`cimom.snmpTrapListenerPort=162`
8. When you are done, click **Save**.
9. Restart the service for the management server for your changes to take effect. While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Discovering McDATA and EMC Connectrix Switches

McDATA and EMC Connectrix switches use the Fibre Channel Switch Application Programming Interface (SWAPI) to communicate with devices on the network. The management server can discover multiple instances of Enterprise Fabric Connectivity Manager.

Use one of the following techniques to discover McDATA and Connectrix switches:

Table 2-5: Discovery Settings for McDATA and Connectrix Switches

Discovery	SWAPI setting through a Proxy	SNMP setting Through a Proxy	Contacting the switch directly
Description	Use this option if you have Enterprise Fabric Connectivity (EFC) Manager. You will need to connect through the proxy instead of the switch. See “SWAPI Setting Through a Proxy” on page 49 for more information.	Contact the switch through a proxy. You can use this option with EMC Connectrix™ Manager and Enterprise Fabric Connectivity (EFC) Manager to contact the switch. See “SNMP Setting Through a Proxy” on page 53 for more information.	Contact the switch by its IP address or DNS name. This connection uses SNMP or SMI-S depending on the McDATA switch model. See the Support Matrix for details on McDATA switch models (Help > Documentation Center > Support Matrix). See “Contacting a McDATA or Connectrix Switch Directly” on page 55.
Provisioning Limitations	The SWAPI setting lets you activate a zone set, in addition to creating, editing, and deleting zones and zone sets. You cannot manage or view information about zone aliases.	This SNMP setting through a proxy does not let you manage or access information about zones, zone sets or zone aliases.	This SNMP setting provides view only access to the active zone set and its members. You cannot create, modify, and/or delete zone sets or its members. SMI-S provides active management of zones.

Keep in mind the following:

- If you change a discovery configuration from SNMP to SWAPI or vice versa, the user ID and password will no longer work. For this reason, it is recommended that you set this property before discovering any McDATA switches. If you must change the configuration, see “Changing the Discovery Settings” on page 57.
- After you discover a McDATA or Connectrix switch through a proxy, the IP address displayed next to the name of the switch is actually the IP address of the proxy for the switch in the Discovery, Topology, and Get Details screens. To find the IP address of the switch, click the link for the switch in the Topology screen (**Discovery > Topology**) or Get

Details screen (**Discovery > Details**) and then click the **Properties** tab. The **Properties** tab can also be accessed by double clicking the switch in System Explorer.

- If you want to add, remove, or replace McDATA or Connectrix switches after you have discovered the service processor, you must perform additional steps, see “About Managing McDATA and EMC Connectrix Switches” on page 100.
- If you have problems obtaining information from McDATA or Connectrix switches during discovery and/or Get Details, see “Step 2 - Discover Storage Systems, NAS Devices and Tape Libraries” on page 62.
- All McDATA switches in a fabric must be managed by the same EFC Manager. Do not have more than one EFC Manager to a fabric for McDATA switches. If you do use more than one EFC Manager in a fabric, you must use the same EFC Manager for your zoning. Do not use the other EFC Managers for zoning, as this will create zoning database problems.
- All Connectrix switches in a fabric must be managed by the same Connectrix Manager. Do not have more than one Connectrix Manager to a fabric for Connectrix switches. If you do use more than one Connectrix Manager in a fabric, you must use the same Connectrix Manager for your zoning. Do not use the other Connectrix Managers for zoning, as this will create zoning database problems.
- If you want the management server to receive SNMP events from Connectrix or McDATA switches, do one of the following:
 - If you discovered Connectrix Manager or EFC Manager, only enable SNMP trap forwarding to the management server on the Connectrix Manager or EFC Manager, not on the individual switches. Connectrix Manager or EFC Manager should be configured to forward SNMP traps to the IP address of the management server, and the community string should match the user ID you used to discover Connectrix Manager or EFC Manager.
 - If you discovered Connectrix or McDATA switches directly, enable SNMP trap forwarding on the switches, not on any other management software. The switches should be configured to forward SNMP traps to the IP address of the management server, and the community string should match the user ID you used to discover the Connectrix or McDATA switches.

SWAPI Setting Through a Proxy

With the SWAPI setting, the management server contacts a proxy to obtain information about the switches connected to it. Use Enterprise Fabric Connectivity (EFC) Manager for this option. If you do not have EFC Manager, see “SNMP Setting Through a Proxy” on page 53. EFC Manager versions 7.0, 1.3 and later can communicate with the management server and the switch. EFC Manager accesses the switch through a SWAPI connection. This configuration lets multiple

instances of the management server or other clients contact EFC Manager, which in turn provides information about the switch.

Important: EMC customers using the EMC Connectrix Manager (EMC's rebranded EFCM) cannot use the EMC Fibre Zone Bridge (EMC's rebranded Bridge Agent) to discover EMC switches using SWAPI. The McDATA SWAPI library is incompatible with EMC's Fibre Zone Bridge Agent.

If the Fibre Zone Bridge Agent is not installed or not needed, you can uninstall the Fibre Zone Bridge Agent and install McDATA's Bridge Agent. The McDATA Bridge Agent will work with EMC's Connectrix Manager, but it cannot co-exist with EMC's Fibre Zone Bridge Agent.

If you are running Connectrix Manager and you need to have the EMC Fibre Zone Bridge Agent running, you will not be able to discover EMC switches using SWAPI. You must discover them through the SNMP provider, either directly or through a proxy. See "SNMP Setting Through a Proxy" on page 53 for more information about using the SNMP provider to discover switches through a proxy. See "Contacting a McDATA or Connectrix Switch Directly" on page 55 for more information about discovering switches by their IP address.

Note: Neither McDATA nor EMC officially support running the EMC Connectrix Manager with the McDATA Bridge Agent. Although this configuration has been tested for discovering EMC switches using SWAPI, you should check with your EMC or McDATA representative to determine the implications of this configuration.

Step 1 - (McDATA Switches Only) Install the Bridge Agent

To communicate with EFC Manager, the management server requires the Bridge Agent. Refer to your McDATA representative for more information about the Bridge Agent.

Step 2 - Verify the Discovery Setting for McDATA and Connectrix Switches is Set to SWAPI

By default the discovery settings for McDATA and Connectrix switches is set to SWAPI. If you believe it has been changed to SNMP, you can perform the following steps to change it back to SWAPI.

1. Click **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the following property. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
`cimom.useSnmppMcDataProvider=TRUE`
4. Return to the Advanced page (**Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree).
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Uncomment the `cimom.useSnmppMcDataProvider` property by removing the number sign (#) in front of `cimom.useSnmppMcDataProvider`.
7. Change the `cimom.useSnmppMcDataProvider` property to false.
`cimom.useSnmppMcDataProvider=FALSE`
8. When you are done, click **Save**.
9. Restart the service for the management server for your changes to take effect. While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Step 3 - Discover the Proxy

To discover the proxy:

1. Click **Discovery**, then click **Setup** in the upper-right pane of the window.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click the **Add Address** button.
5. In the **IP Address/DNS Name** field, type the IP address or DNS name of the EFC Manager/Connectrix Manager you want to discover.

6. In the **User Name** field, type the user name for EFC Manager/Connectrix Manager. This field can be left blank if one or more of the following conditions are fulfilled:
 - ☐ The element's user name and password are one of the default user names and passwords.
 - ☐ The element does not require authentication.To access a Windows-based device, prepend the user name with the Windows domain name, as shown in the following example.

```
domain_name\user_name
```

where
 - ☐ `domain_name` is the domain name of the machine
 - ☐ `user_name` is the name of your network account
7. In the **Password** field, type the corresponding password for EFC Manager/Connectrix Manager. This field can be left blank if one or more of the following conditions are fulfilled:
 - ☐ The element's user name and password are one of the default user names and passwords.
 - ☐ The element does not require authentication.
8. If you typed a password in the previous step, retype the password in the **Verify Password** field.
9. In the **Comment** field (optional), type a comment for additional information. The information typed into this field is displayed under the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Click **OK**.
11. To start discovering elements on the network, click the **Start Discovery** button on the **IP Addresses** tab.
Discovery is complete when the software displays the DISCOVERY COMPLETED message in the Log Messages field.

CIM_ERR_FAILED When Trying to Activate a Zone Set Using McDATA SWAPI

When the user tries to activate a zone set using McDATA SWAPI, the operation may return CIM_ERR_FAILED with one of the following detailed messages:

```
Cannot activate zone set. SWAPI Handle is not valid for fabric
Cannot activate zone set. Active zone set information is out of date for fabric
```

```
There is no active SWAPI connection for fabric
Fabric is not in the cache
```

These error messages indicate that the SWAPI connection to the EFCM managing the fabric is no longer valid, or the active zone information was changed on the fabric without using the management server. The management server does not activate a zone set under these conditions.

(Discovery from HP SE only) To fix this problem, use the **Test** button on the discovery screen to check the status of the SWAPI connection. If necessary, re-discover the EFCM to re-establish the SWAPI connection.

Once the connection is working, the provisioning operation should succeed. If it continues to fail because the active zone set information is out of date, do a Get Details for this element to update the zoning information.

SNMP Setting Through a Proxy

This SNMP setting through a proxy does not let you manage or access information about zones, zone sets or zone aliases.

This option is required if you want to discover McDATA or Connectrix switches through a proxy using the SNMP provider. You can use this option with EMC Connectrix™ Manager and Enterprise Fabric Connectivity (EFC) Manager to contact the switch.

Step 1 - Change the Discovery Setting for Switches to SNMP

To change the discovery settings to SNMP:

1. Click **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the following command:
`#cimom.useSnmppMcDataProvider=true`
4. Return to the Advanced page.
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Uncomment the property by removing the hash (#) symbol in front of the property if needed:
`cimom.useSnmppMcDataProvider=true`

7. When you are done, click **Save**.
8. Restart the service for the management server.
9. Verify the following on the proxy and the switches accessible from the proxy:
 - ☐ The SNMP agent is enabled.
 - ☐ The read-only community string is configured.

Step 2 - Discover the Proxy

To discover the proxy:

1. Click **Discovery**, then click **Setup** in the upper-right pane of the window.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click the **Add Address** button.
5. In the **IP Address/DNS Name** field, type the IP address or DNS name of the proxy you want to discover.
6. In the **User Name** field, type the user name, which is the read-only community string of the EFC Manager or Connectrix Manager. The default community-string is "public" but this can be changed on the EFC Manager or Connectrix Manager.
7. Leave the **Password (Optional)** field blank. The password does not matter since the management server is not doing any configurations through SNMP.
8. You can leave the **Verify Password** field blank.
9. In the **Comment** field (optional), type a comment for additional information. The information typed into this field is displayed under the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. To start discovering elements on the network, click the **Start Discovery** button on the **IP Addresses** tab.
Discovery is complete when the software displays the DISCOVERY COMPLETED message in the Log Messages field.

Important: To obtain more information about the switch, you need to map the topology and obtain element details. See the topics, "Building the Topology View" on page 88 and "Get Details" on page 93.

Step 3 - Make Sure There Are No Port Conflicts for Receiving SNMP Traps

When the management server is configured to contact the proxy by SNMP, the management server receives events from the proxy in the form of SNMP traps. By default, the management server uses port 162 to receive SNMP traps. If another software package is using that port, the management server is unable to receive the traps. To change the port the management server uses:

1. Select **Product Health**. Then, click **Advanced** in the **Disk Space** tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the following command:
`cimom.snmpTrapListenerPort`
4. Return to the Advanced page.
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Set the `cimom.snmpTrapListenerPort` to another port, as shown in the following example:
`cimom.snmpTrapListenerPort=1234`
where 1234 is the new port
7. When you are done making your changes, click **Save**.
8. Restart the service for the management server for your changes to take effect. While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Step 4 - Set Up the Proxy to Send Traps to the Correct Port

When you are using the SNMP setting to discover a proxy, you must configure the SNMP agent on the proxy manager to send traps to the management server using the port you selected. This configuration sends traps from all switches managed by that proxy. Refer to your documentation for your proxy for more information.

Contacting a McDATA or Connectrix Switch Directly

The management server uses SWAPI to discover a McDATA or Connectrix switch by its IP address or DNS name. SWAPI is the default setting. If you want to discover McDATA or Connectrix switches by SNMP, you must change to SNMP before you begin the following steps. See “Changing the Discovery Settings” on page 57. If you are using a McDATA SMI-S compliant switch model, a direct connection to the switch is required and the switch must be contacted via its IP address. See the Support Matrix for McDATA switch details (**Help > Documentation Center > Support Matrix**).

To discover a McDATA or Connectrix switch directly:

1. Click **Discovery**, then click **Setup** in the upper-right pane of the window.
 2. Select **Step 1** at the top of the page.
 3. Click the **IP Addresses** tab.
 4. Click the **Add Address** button.
 5. In the **IP Address/DNS Name** field, type the IP address or DNS name of the switch you want to discover.
 6. In the **User Name** field, type the user name for accessing the switch. If you are using SNMP the user name is the read-only community string of the switch. The default community-string is "public" but this can be changed on the switch. If you are using SMI-S the user name is the user name of the admin login of the switch.
 7. If you are using SNMP leave the **Password (Optional)** field blank. The password does not matter since the management server is not doing any configurations through SNMP. If you are using SMI-S type the password of the admin account on the switch.
 8. In the **Verify Password** field type the same thing you typed in the password field.
 9. In the **Comment** field (optional), type a comment for additional information. The information typed into this field is displayed under the Comment column in the Addresses to Discover list (**Discovery > Setup**).
 10. Do not select the **Do Not Authenticate** option.
 11. Click **OK**.
 12. To start discovering elements on the network, click the **Start Discovery** button on the **IP Addresses** tab.
Discovery is complete when the software displays the DISCOVERY COMPLETED message in the Log Messages field.
-

Important: To obtain more information about the switch, you need to map the topology and obtain element details. See the topics “Building the Topology View” on page 88 and “Get Details” on page 93.

Make Sure There Are No Port Conflicts for Receiving SNMP Traps

When the management server is configured to contact a switch by SNMP, the management server receives events from the switch in the form of SNMP traps. By default, the management server uses port 162 to receive SNMP traps. If another software package is using that port, the management server is unable to receive the traps. To change the port the management server uses:

1. Select **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the following command:
`cimom.snmpTrapListenerPort`
4. Return to the Advanced page.
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Set the `cimom.snmpTrapListenerPort` to another port, as shown in the following example:
`cimom.snmpTrapListenerPort=1234`
where 1234 is the new port
7. When you are done making your changes, click **Save**.
8. Restart the service for the management server for your changes to take effect. While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Step 4 - Set Up the Proxy to Send Traps to the Correct Port

When you are using the SNMP setting to discover a proxy, you must configure the SNMP agent on the proxy manager to send traps to the management server using the port you selected. This configuration sends traps from all switches managed by that proxy. Refer to your documentation for your proxy for more information.

Changing the Discovery Settings

To change the discovery settings from SWAPI to SNMP or vice versa:

1. Delete all McDATA and Connectrix switches in the application by going to the Get Topology for Discovered Elements table (**Discovery > Topology**) and selecting the switches you want to delete. Then, click the **Delete** button.
2. Delete all McDATA and Connectrix switches listed in the Addresses To Discover table (**Discovery > Setup**) by selecting the switches you want to delete and clicking the Delete button.
3. Click **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree.
4. Click **Show Default Properties** at the bottom of the page.
5. Change the `cimom.useSnmpMcDataProvider` property as follows:
 - ☐ **SNMP setting** - Uncomment `cimom.useSnmpMcDataProvider` property by removing the number sign (#) in front of the `cimom.useSnmpMcDataProvider` property as follows: `cimom.useSnmpMcDataProvider=true`
 - ☐ **SWAPI setting** - Comment out the `cimom.useSnmpMcDataProvider` property by placing a number sign (#) in front of the `cimom.useSnmpMcDataProvider` property.
6. Click **Save**.
7. Restart the service for the management server.
8. Add new elements in the Discovery screen (**Discovery > Setup > Add Address**) by using the appropriate IP address and user name
 - ☐ **SWAPI connection** - Enter the IP address, user name and password for the proxy.
 - ☐ **SNMP connection** - Enter the IP address of the proxy. The default user name is "public" (the read-only community string). The password does not matter since the management server is not doing any configurations through SNMP.
9. Verify the following on the proxy and the switches accessible from the proxy:
 - ☐ The SNMP agent is enabled.
 - ☐ The read-only community string is configured.
10. Run Discovery.
11. Run Get Details.

Excluding McDATA and EMC Connectrix Switches from Discovery

Specific McDATA and Connectrix switches can be excluded from discovery by using system properties.

To exclude one or more switches from discovery, you must modify the `cimom.mcddata.exclude` property. Set the property `cimom.mcddata.exclude` to a comma separated list of Worldwide Names of the McDATA and Connectrix switches you want excluded, as shown in the following example:

```
cimom.mcddata.exclude=1000080088A07024,1000080088A0D0B6
```

The management server excludes the switches with one of the following Worldwide Names: 1000080088A07024 and 1000080088A0D0B6

If the `cimom.mcddata.exclude` property is not modified, the management server discovers and obtains details from all McDATA and Connectrix switches.

Important: The IP addresses of excluded elements appear in the discovery lists (**Discovery > Setup**), topology (**Discovery > Topology**), or Get Details lists (**Discovery > Details**). The management server does not display additional information about excluded elements in the user interface. The management server, however, does mention in the logs (**Discovery > View Logs**) that a provider instance has been created for an excluded element. You can ignore this message that appears in the logs.

To modify the `cimom.mcddata.exclude` property:

1. Select **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree.
 2. Click **Show Default Properties** at the bottom of the page.
 3. Copy the `cimom.mcddata.exclude` property. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
 4. Return to the Advanced page (**Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree).
-

5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Make your changes in the **Custom Properties** field. Make sure the property is not commented out by removing the hash (#) symbol in front of the property.
7. Add the Worldwide Names corresponding to the switches you want to exclude from discovery. Separate additional Worldwide Names with a comma, as shown by the following example:
`cimom.mcdata.exclude=1000080088A07024,1000080088A0D0B6`
where 1000080088A07024 and 1000080088A0D0B6 are the Worldwide Names for McDATA and Connectrix switches.
8. When you are done, click **Save**.
While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Viewing Log Messages

Use the **View Logs** page to obtain the status of the following:

- Discovery
- Building the Topology

During these operations, the management server displays its status at regular intervals.

To view logs for these operations:

1. Click **Discovery > View Logs**.
2. To obtain the latest status, click the **Get Latest Messages** button.

If the software is unable to discover or obtain information about a device, the log messages might provide some information as to where the problem occurred.

For example, if a host was not discovered, the log messages might indicate the provider configuration for that device was never created. This could mean the software was given the wrong user name and/or password for that host. As a result, the software logged onto the host with a guest account, which does not have enough permissions to start WMI.

Important: Look at Event Manager for additional information. See “About Event Manager” on page 536 for more information.

Viewing the Status of System Tasks

The Task Dashboard allows you to view the status of the tasks running on the management server. The dashboard provides the name of each task, its latest status, and the time the status was last reported.

To view the status of system tasks:

1. Click **Discovery > System Tasks**.
2. To obtain the latest status, click the **Get the Latest Status** button.

The following task statuses are provided by the Task Dashboard:

Table 2-6: Task Status Descriptions

Status	Description
Not Found	This task can not be found on this server.
Completed	This task has been completed successfully.
Failed	This task failed with an error.
Aborted	This task has been aborted by the user or other automated actions.
In Progress	This task is in progress. CPU and disk activities are active on this server.
Queued	This task is scheduled to be executed in the future.
Rejected	This task has been rejected by this server.

Duplicate Logs for Brocade Switches in Same Fabric

If you discover more than one Brocade switch in the same fabric, the discovery log displays duplicate listings for the Brocade switches. Each Brocade switch is listed multiple times with the IP address of the other switches and its own.

For example, assume you are discovering Brocade switches QBrocade2 and QBrocade5 in the same fabric, two duplicate entries are displayed in the log. QBrocade2 is listed twice, once with its own IP address, the other time with the IP address of QBrocade5, as shown below.

```
[Nov 27, 2002 8:45:05 AM] Discovered Switch: QBrocade2 at 192.168.10.22
[Nov 27, 2002 8:45:09 AM] Discovered Switch: QBrocade5 at 192.168.10.22
[Nov 27, 2002 8:45:09 AM] Enabling provider configuration:
APPIQ_BrocadeElementManagerConfig
[...]
[Nov 27, 2002 8:45:37 AM] Discovered Switch: QBrocade2 at 192.168.10.25
[Nov 27, 2002 8:45:42 AM] Discovered Switch: QBrocade5 at 192.168.10.25
[Nov 27, 2002 8:45:42 AM] Enabling provider configuration:
APPIQ_BrocadeElementManagerConfig
```

Note: On the **Topology** page, the software displays each Brocade switch (192.168.10.22 and 192.168.10.25) as elements:

```
192.168.10.22 Switch QBrocade2, QBrocade5 admin
192.168.10.25 Switch QBrocade2, QBrocade5 admin
```

Step 2 - Discover Storage Systems, NAS Devices and Tape Libraries

- “Discovering 3PAR Storage Systems” on page 65
 - “Discovering EMC Solutions Enabler 5.1” on page 66
 - “Excluding EMC Symmetrix Storage Systems from Discovery” on page 67
 - “Discovering EMC CLARiiON Storage Systems” on page 69
 - “Discovering Engenio Storage Systems” on page 69
-

- “Discovering HDS Storage Systems” on page 71
- “Excluding HDS Storage Systems from Discovery” on page 72
- “Discovering HP StorageWorks XP Arrays” on page 74
- “Discovering IBM Storage Systems” on page 77
- “Discovering Sun StorEdge 3510 Storage Systems” on page 79
- “Discovering Sun StorEdge 6920 and 6940 Storage Systems” on page 81
- “Discovering Sun StorEdge 6130 Storage Systems” on page 81
- “Discovering Xiotech Storage Systems” on page 82
- “Discovering HP NAS Devices on Windows” on page 83
- “Discovering HP NAS Devices on Linux” on page 84
- “Discovering NetApp NAS Devices” on page 85
- “Discovering Sun NAS Devices” on page 86
- “Discovering HP and IBM Tape Libraries” on page 87

Table 2-7: Discovery Requirements for Storage Systems, Tape Libraries & NAS Devices

Element	Discovery Requirements	Additional Information
3PAR storage systems	Discover the 3PAR storage system directly.	See “Discovering 3PAR Storage Systems” on page 65.
EMC CLARiiON storage systems	The EMC Navisphere CLI is required for the management server to communicate with the CLARiiON storage system.	See “Discovering EMC CLARiiON Storage Systems” on page 69 for more information.
EMC Symmetrix storage system (Including EMC Symmetrix DMX storage systems)	Discover the server running the EMC Solutions Enabler.	See “Discovering EMC Solutions Enabler 5.1” on page 66 for more information.

Table 2-7: Discovery Requirements for Storage Systems, Tape Libraries & NAS Devices

Element	Discovery Requirements	Additional Information
Engenio storage systems	Can be discovered two ways: ■ Entering the IP address/ DNS name, user name and password of a controller for an Engenio storage system. Discovers only the corresponding IP address of the controller. ■ Entering the IP address/ DNS name, user name and password of a proxy that is used to manage an Engenio storage system. Discovers all controllers known to the proxy.	See "Discovering Engenio Storage Systems" on page 69.
HDS storage systems	Discover the server running HiCommand Device Manager.	See "Discovering HDS Storage Systems" on page 71 for more information.
HP storage systems	Discover the server running the HP CIMOM.	See "Discovering HP StorageWorks XP Arrays" on page 74.
IBM Storage Systems	Discover the CIMOM that talks to the IBM storage systems you want to monitor.	See "Discovering IBM Storage Systems" on page 77.
Sun StorEdge 3510	Discovered through proxy software called Sun StorEdge™ Configuration Service. On the discovery page the user should enter the hostname or IP address of the computer running the Sun StorEdge 3510 SMI-S provider.	See "Discovering Sun StorEdge 3510 Storage Systems" on page 79.
Sun StorEdge 6920 and 6940	Discover the storage system directly.	See "Discovering Sun StorEdge 6920 and 6940 Storage Systems" on page 81.

Table 2-7: Discovery Requirements for Storage Systems, Tape Libraries & NAS Devices

Element	Discovery Requirements	Additional Information
Sun StorEdge 6130	Discover the storage system directly. The username does not matter. The password matters only for provisioning.	See “Discovering Sun StorEdge 6130 Storage Systems” on page 81.
Xiotech Storage Systems	Discover the storage system directly.	See “Discovering Xiotech Storage Systems” on page 82.
HP NAS Devices	Discover the device directly.	See “Discovering HP NAS Devices on Windows” on page 83 and “Discovering HP NAS Devices on Linux” on page 84.
NetApp Devices	Discover the device directly.	See “Discovering NetApp NAS Devices” on page 85.
Sun NAS Devices	Discover the server running the SMI-S provider for the Sun NAS Devices.	See “Discovering Sun NAS Devices” on page 86.
HP and IBM Tape Libraries	Provide the IP address, namespace, user name and password for the tape library.	See “Discovering HP and IBM Tape Libraries” on page 87

Discovering 3PAR Storage Systems

Note: You do not need to provide the interop namespace because the management server includes the interop namespace for 3PAR storage systems in its default list.

Important: To be able to discover a 3PAR storage system, the SMI-S server for the 3PAR storage system must be running. By default, the 3PAR SMI-S server is not started on the array. To start the SMI-S server, bring up the InForm CLI and run the following command:

```
startcim
```

This command starts the SMI-S server within a minute or so.

To discover a 3PAR storage system:

1. Click **Discovery > Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click the **Add Address** button.
5. In the **IP Address/DNS Name** field, the following for the 3PAR storage system you want to discover.
`<host>`
where `<host>` is the IP address or DNS name of the 3PAR storage system you want to discover.
6. Enter the **User Name** of the storage system.
7. Enter the **Password** of the storage system.
8. Retype the password in the **Verify Password** field.
9. In the **Comment** field (optional), type a comment for additional information. The information typed into this field is displayed under the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. To start discovering elements on the network, click the **Start Discovery** button on the IP Addresses tab.

Discovering EMC Solutions Enabler 5.1

EMC Solutions Enabler restricts access to itself through the nethost file. If present, the nethost file is located in the same directory as the netcnfg file. If you are using a nethost file, edit it to allow the management server to discover the Solutions Enabler and the Symmetrix storage systems that it manages.

Important: Use a nethost file unless you are running a version of the Solutions Enabler earlier than the 5.1 version. You must have the license installed for the Solutions Enabler. The nethost file provides access to the Solutions Enabler API.

Sometimes you can access an EMC Symmetrix storage system through several Solutions Enabler servers. In this case if you do not have access to a particular Solutions Enabler, you may still be able to access the Symmetrix storage system through another Solutions Enabler.

If you do not have a nethost file, you may need to create one. For example, assume you are running Solutions Enabler on a Solaris server, you would create a nethost file as described in the following steps. Refer to the documentation for Solutions Enabler for other operating systems.

1. Create a file called “nethost” in the `/opt/emc/API/symapi/config` directory.
2. Add the following lines to the nethost file:

Management Server Running on Microsoft Windows

```
[management server name] SYSTEM  
[management server IP] SYSTEM
```

Management Server Running on Sun Solaris

```
[management server name] root  
[management server IP] root
```

where

- [management server name] is the DNS name of the management server
- [management server name] is the IP address of the management server

3. Add the following line to the `/opt/emc/API/symapi/config/netcnfg` file:
`SYMAPI_SERVER - TCPIP <IP of SymAPI server> 2707`
4. Use the following command to start the daemon:
`/opt/emc/SYMCLI/V5.5.0/bin/symapisrv -service SYMAPI_SERVER start -background`
5. Use the following command to stop the daemon:
`/opt/emc/SYMCLI/V5.5.0/bin/symapisrv stop`
6. You may need to discover the Symmetrix arrays the SymAPI server can see by running the following command:
`/opt/emc/SYMCLI/V5.5.0/bin/symcfg discover`

Important: If error 214 is present in the discovery log and/or cimom.log during discovery, this means the SymAPI server is not licensed for remote connections. The end-user will have to acquire and install the license before discovery can occur.

Required Licenses

If you want to use all of the features of the management server, such as provisioning, with an EMC Symmetrix storage system, you must have licenses for the following products:

- BASE
- DeltaMark
- SERVER
- DevMasking
- Config Manager
- Mapping (SOLUTION_4)

Excluding EMC Symmetrix Storage Systems from Discovery

When multiple EMC Symmetrix storage systems are managed through a single Solutions Enabler, specific storage systems may be excluded from discovery by using system properties.

To exclude one or more Symmetrix storage systems from discovery, you must modify the `cimom.symmetrix.exclude` property. Set the property `cimom.symmetrix.exclude` to a comma separated list of serial numbers of the storage systems you want excluded, as shown in the following example:

```
cimom.symmetrix.exclude=000183500570,000183610580
```

The management server excludes the storage systems with one of the following serial numbers: 000183500570 and 000183610580.

If the `cimom.symmetrix.exclude` property, the management server discovers and obtains details from all EMC Symmetrix Storage Systems managed by discovered Solutions Enablers.

Important: The IP addresses of excluded elements appear in the discovery (**Discovery > Setup**), topology (**Discovery > Topology**), Discovery Data Collection (**Discovery > Run Discovery Data Collection**), or Get Details lists (**Discovery > Details**). The management server does not display additional information about excluded elements in the user interface. The management server, however, does mention in the logs that a provider instance has been created for an excluded element. You can ignore this message that appears in the logs.

To modify the `cimom.symmetrix.exclude` property:

1. Select **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the following command. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
`#cimom.symmetrix.exclude=000183500570,000183500575`
4. Return to the Advanced page.
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Make sure the property is not commented out by removing the hash (#) symbol in front of the property. Add the serial numbers corresponding to the Symmetrix storage systems you want to exclude from discovery. Separate additional serial numbers with a comma, as shown by the following example:
`cimom.symmetrix.exclude=000183500570,000183500575`
where 000183500570 and 000183500575 are serial numbers for Symmetrix storage systems.
7. When you are done, click **Save**.
8. Restart the service for the management server for your changes to take effect. While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Discovering EMC CLARiiON Storage Systems

The EMC Navisphere® CLI must be installed on the management server for the management server to communicate with the CLARiiON® storage system. At the time this documentation was created, EMC distributed the Navisphere CLI as part of the EMC Navisphere Software Suite. Contact your EMC representative for more information about obtaining the Navisphere CLI. Distribution rights for the Navisphere CLI belong to EMC. For Solaris, you must install the Navisphere Disk Array Management Tool CLI (NAVICLI).

Contact your EMC representative for more information about obtaining the Navisphere CLI. Distribution rights for the Navisphere CLI belong to EMC.

Important: Before you discover your CLARiiON storage systems, you must have already installed all required software components for your CLARiiON storage system, such as the Navisphere Host Agent. Refer to the documentation for your storage system for more information.

In Navisphere Manager add one of the following to the privilege user section:

```
SYSTEM@name_of_my_management_server
```

```
SYSTEM@IP_of_my_management_server
```

where

- `name_of_my_management_server` is the DNS name of the computer running the management server software
- `IP_of_my_management_server` is the IP address of the computer running the management server software

When you use the management server to discover the CLARiiON storage system, provide the IP address for the CLARiiON storage system and the user name and password used to log into Navisphere.

Discovering Engenio Storage Systems

Keep in mind the following when discovering an Engenio storage system:

- Discover all controllers on an Engenio storage system by entering the IP address of each controller.
- The management server must have the User Name field populated to discover the Engenio storage system. If your Engenio storage system does not have a user name set, you must enter something in the **User Name** field, even though the storage system has no user name.
- Discover both controllers for the Engenio storage system. Each controller has its own IP address. When you discover, specify all the IP addresses for all the controllers (usually two). The management server discovers these controllers as one single storage system.
- To obtain drive-related statistics, install a proxy host. Ensure the proxy host has at least one LUN rendered by each controller of the array. See the topic, “Obtaining Disk Drive Statistics from Engenio Storage Systems” on page 150 for more information.
- A license key is required for each storage system and that the key is obtained from the Web site specified on the Activation Card that shipped with your storage system.
- Engenio storage systems do not require a password for Get Details. If you want do not want to use the management server for provisioning on Engenio storage systems, select the **Do Not Authenticate** option. The management server will still monitor the Engenio storage system; however, you will not be able to do provisioning tasks.

Do the following to discover Engenio storage systems:

1. Click **Discovery > Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click the **Add Address** button.
5. In the **IP Address/DNS Name** field, type the IP address or DNS name of the controller or proxy you want to discover.
6. Enter the user name in the **User Name** field. If your Engenio storage system does not have a user name, you must enter something in the **User Name** field, even though the storage system has no user name.
7. Leave the **Password** field blank if you do not want to do provisioning on the Engenio storage system. If you want to do provisioning, type the corresponding password for controller or proxy.
8. If you typed a password in the previous step, retype the password in the **Verify Password** field.
9. In the **Comment** field (optional), type a comment for additional information. The information typed into this field is displayed under the Comment column in the Addresses to Discover list (**Discovery > Setup**).

10. If you do not plan to use provisioning in the product, select the **Do Not Authenticate** option.
11. Click **OK**.
12. To start discovering elements on the network, click the **Start Discovery** button on the **IP Addresses** tab.

Discovering HDS Storage Systems

HiCommand Device Manager is required for the management server to communicate with an HDS storage system. To discover an HDS storage system, enter the IP address, user name and password for the server running HiCommand Device Manager. Do not point to the disk array for the storage system.

To obtain information about HDS storage systems, the management server must be able to access the port HiCommand Device Manager uses to listen. By default, HiCommand Device Manager listens on port 2001, and the management server assumes this configuration at discovery time. If HiCommand Device Manager uses a different port, specify this other port when you discover HiCommand Device Manager.

Keep in mind the following:

- You cannot scan an IP range to discover an instance of HiCommand Device Manager that listens on a port other than 2001. The management server does not allow port numbers in the scanning of IP ranges, and thus, you are not able to specify the port.
- The management server communicates with HiCommand Device Manager through a nonsecure connection. If you want the management server to communicate with HiCommand Device Manager through a secure sockets layer (SSL) connection, you must modify an internal property or use HTTPS when you discover HiCommand Device Manager. See “Communicating with HiCommand Device Manager Over SSL” on page 825.

To discover an HDS storage system that listens on a port other than 2001:

1. Access the Discovery Setup page (**Discovery > Setup**).
2. Click the **Add Address** button.
3. In the **IP Address/DNS Name** field, type the name of the server and the port HiCommand Device Manager uses to listen separated by a colon, as shown in the following example:
`proxy2:1234`
where

- ☐ `proxy2` is the name of the server running HiCommand Device Manager
- ☐ `1234` is the port HiCommand Device Manager uses to listen
- 4. In the **User Name** field, type the user name for accessing HiCommand Device Manager.
- 5. In the **Password** field, type the password for accessing HiCommand Device Manager.
- 6. In the **Verify Password** field, retype the password for accessing HiCommand Device Manager.
- 7. In the **Comment** field (optional), type a comment for additional information. The information typed into this field is displayed under the Comment column in the Addresses to Discover list (**Discovery > Setup**).
- 8. Do not select the **Do Not Authenticate** option.
- 9. Click **OK**.

Excluding HDS Storage Systems from Discovery

When multiple HDS storage systems are managed through a single HiCommand Device Manager, specific storage systems may be excluded from discovery by using system properties.

To exclude one or more HDS storage systems from discovery, you must modify the `cimom.hds.exclude` property. Set the property `cimom.hds.exclude` to a comma separated list of serial numbers of the storage systems you want excluded, as shown in the following example:

```
cimom.hds.exclude=61038,61037
```

The management server excludes the storage systems with one of the following serial numbers: 61038 and 61037.

If the `cimom.hds.exclude` property is not specified, the management server discovers and obtains details from all HDS storage systems managed by the discovered HiCommand Device Manager.

The IP addresses of excluded elements appear in the discovery (**Discovery > Setup**), topology (**Discovery > Topology**) or Get Details list (**Discovery > Details**). The management server does not display additional information about excluded elements in the user interface. The management server, however, does mention in the logs (**Discovery > View Logs**) that a provider instance has been created for an excluded element. You can ignore this message that appears in the logs.

To modify the `cimom.hds.exclude` property:

1. Select **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree.

2. Click **Show Default Properties** at the bottom of the page.
3. Copy the following command. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
`#cimom.hds.exclude=61038,61037`
4. Return to the Advanced page.
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Make sure the property is not commented out by removing the hash (#) symbol in front of the property. Add the serial numbers corresponding to the HDS storage systems you want to exclude from discovery. Separate additional serial numbers with a comma, as shown by the following example:
`cimom.hds.exclude=61038,61037`
where 61038 and 61037 are serial numbers for HDS storage systems.
7. When you are done, click **Save**.
8. Restart the service for the management server for your changes to take effect. While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Discovering HP StorageWorks EVA or MSA Arrays

To discover HP StorageWorks EVA or MSA arrays, you must enter the following information for the instance of the HP CIMOM used to manage the storage system:

- user name and password used for accessing the HP CIMOM
- IP address of the server containing the HP CIMOM

The following should be installed on a server before you discover an HP EVA or MSA storage system:

- HP Storage Management Appliance software
 - HP OpenView Storage Operations Manager
 - HP StorageWorks Command View EVA or MSA
 - HP StorageWorks SMI-S EVA or MSA.
-

To determine provisioning support for HP StorageWorks Arrays, see Table 10-3, “Provisioning and Pool Support,” on page 427 and Table 10-4, “Volume and HSG Support,” on page 429.

To discover HP EVA or MSA storage systems:

1. Click **Discovery > Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click the **Add Address** button.
5. In the **IP Address/DNS Name** field, type the IP address or DNS name of the HP CIMOM you want to discover.
6. Enter the **User Name** used to access the HP CIMOM.
7. Enter the **Password** used to access the HP CIMOM.
8. If you typed a password in the previous step, retype the password in the **Verify Password** field.
9. In the **Comment** field (optional), type a comment for additional information. The information typed into this field is displayed under the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. To start discovering elements on the network, click the **Start Discovery** button on the **IP Addresses** tab.

Discovering HP StorageWorks XP Arrays

You can discover HP StorageWorks XP Arrays by using the following:

- Command View XP
- Command View XP Advanced Edition
- XP Provider

To determine provisioning support for HP StorageWorks Arrays, see Table 10-3, “Provisioning and Pool Support,” on page 427 and Table 10-4, “Volume and HSG Support,” on page 429.

Discovering HP XP Arrays by Using Command View XP

To discover HP StorageWorks XP arrays by using Command View XP, you must enter the following information for the instance of the HP CIMOM used to manage the storage system:

- user name and password used for accessing the HP CIMOM
- IP address of the server containing the HP CIMOM

The following should be installed on a server before you discover an HP XP storage system if you are discovering the storage system by using Command View XP:

- HP Storage Management Appliance software
- HP OpenView Storage Operations Manager
- HP StorageWorks Command View XP

To discover HP storage systems by using Command View XP:

1. Click **Discovery > Setup**.
 2. Select **Step 1** at the top of the page.
 3. Click the **IP Addresses** tab.
 4. Click the **Add Address** button.
 5. In the **IP Address/DNS Name** field, type the IP address or DNS name of the HP CIMOM you want to discover.
 6. Enter the **User Name** used to access the HP CIMOM.
 7. Enter the **Password** used to access the HP CIMOM.
If you have Command View version 2.0 or later, the default password is administrator. If you have Command View earlier than version 2.0, refer to the documentation that shipped with Command View for the default password.
 8. If you typed a password in the previous step, retype the password in the **Verify Password** field.
 9. In the **Comment** field (optional), type a comment for additional information. The information typed into this field is displayed under the Comment column in the Addresses to Discover list (**Discovery > Setup**).
 10. Do not select the **Do Not Authenticate** option.
 11. Click **OK**.
 12. To start discovering elements on the network, click the **Start Discovery** button on the **IP Addresses** tab.
-

Discovering HP XP Arrays by Using Command View XP Advanced Edition

To discover HP StorageWorks XP arrays by using Command View XP Advanced Edition, you must enter the following information:

- user name and password used for accessing Command View XP Advanced Edition
- IP address of the server running Command View XP Advanced Edition

The following should be installed on a server before you discover an HP XP storage system if you are discovering the storage system by using Command View XP Advanced Edition:

- HP Storage Management Appliance software
- HP OpenView Storage Operations Manager
- HP StorageWorks Command View XP Advanced Edition

To discover HP storage systems by using Command View XP Advanced Edition:

1. Click **Discovery > Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click the **Add Address** button.
5. In the **IP Address/DNS Name** field, type the IP address or DNS name of the server running Command View XP Advanced Edition.
6. Enter the **User Name** used to access Command View XP Advanced Edition.
7. Enter the **Password** used to access Command View XP Advanced Edition.
8. If you typed a password in the previous step, retype the password in the **Verify Password** field.
9. In the **Comment** field (optional), type a comment for additional information. The information typed into this field is displayed under the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. To start discovering elements on the network, click the **Start Discovery** button on the **IP Addresses** tab.

Discovering HP XP Arrays by using the XP Provider

To discover HP StorageWorks XP arrays by using XP provider, you must enter the following information for the XP array:

- user name and password used for accessing the XP array
- IP address of the XP array

To discover HP storage systems by using XP provider:

1. Click **Discovery > Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click the **Add Address** button.
5. In the **IP Address/DNS Name** field, type the IP address or DNS name of the XP storage system you want to discover.
6. Enter the **User Name** used to access the XP storage system.
7. Enter the **Password** used to access the XP storage system.
8. If you typed a password in the previous step, retype the password in the **Verify Password** field.
9. In the **Comment** field (optional), type a comment for additional information. The information typed into this field is displayed under the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. To start discovering elements on the network, click the **Start Discovery** button on the **IP Addresses** tab.

Discovering IBM Storage Systems

Before you can discover an IBM storage system, you must install the IBM CIM Agent. For Enterprise Storage Server (ESS) devices, the CIM Agent is called “CIM Agent for ESS”; for DS devices and mixed DS and ESS environments, use the “CIM Agent for DS Open (API)”. It is recommended not to install the IBM CIM Agent on the management server. Refer to the “CIM Agent for DS Open (API) - Installation and Configuration Guide” for details on configuring the CIM Agent. In short, this procedure entails:

- a. **Installing the software.** The installation checks for the existence of the ESSCLI. If the ESSCLI is not installed, installation of the CIM Agent cannot proceed. The ESSCLI is typically preinstalled on the ESS management server that was configured by the IBM field technician.
- b. **Configuring the protocol and ports used to communicate with the CIM Agent.** You can change the CIM Agent port value, protocol (HTTP/HTTPS), and enable or disable the debug option. Unless a secure connection is required between the management server and the CIM Agent, it is recommended to use port 5988 and protocol HTTP. You must change the default authentication method in order to discover the CIM Agent. Stop the IBM CIM Agent service. Then, edit the `cimom.properties` file in `C:\Program Files\IBM\cimagent` by default. Open the `cimom.properties` file and change the following property to false:
`DigestAuthentication=False`
- c. **Using the `setuser` command to configure a user to access the CIM Agent.** The user credentials specified here are used to access the CIM Agent and are specified in the Discovery Step 1. The credentials are not necessarily the same as those used to log in to the ESS Specialist management utility or the DS Storage Manager.
- d. **Using the `setdevice` command to configure the ESS and DS devices that are managed through the CIM Agent.** The `setdevice` command requires a valid user that has the necessary privileges to access and configure the ESS or DS storage systems.
- e. **Restarting the IBM CIM Agent service.** Restart the IBM CIM Agent service for your changes to take effect.
- f. **Verifying that the CIM Agent is able to communicate with the ESS devices.**

Note: You do not need to provide the interop namespace because the management server includes the interop namespace for IBM storage systems in its default list.

To discover an IBM storage system, you must discover its CIMOM, as described in the following steps:

1. Click **Discovery > Setup**.
 2. Select **Step 1** at the top of the page.
 3. Click the **IP Addresses** tab.
 4. Click the **Add Address** button.
 5. In the **IP Address/DNS Name** field, type the IP address or DNS name of the system running the IBM CIMOM you want to discover.
 6. Enter the **User Name** of the system running the IBM CIMOM.
-

7. Enter the **Password** of the system running the IBM CIMOM.
8. Retype the password in the **Verify Password** field.
9. In the **Comment** field (optional), type a comment for additional information. The information typed into this field is displayed under the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. To start discovering elements on the network, click the **Start Discovery** button on the IP Addresses tab.

Discovering Sun StorEdge 3510 Storage Systems

Before you can discover a Sun StorEdge 3510 storage system, you must set up a Sun StorEdge 3510 SMI-S provider and a Sun StorEdge™ Configuration Service. The provider cannot be installed on the same computer as the management server due to a port conflict.

The Sun StorEdge™ Configuration Service can be installed in one of the following locations:

- on the same computer as the Sun StorEdge 3510 SMI-S provider
- on the management server
- on a separate computer

To install the Sun StorEdge™ Configuration Service you must install the following packages:

- Sun StorEdge™ Configuration Service Console (SUNWscsu)
- Sun StorEdge™ Configuration Service Agent (SUNWscsd)
- Sun StorEdge™ Diagnostic Reporter Agent (SUNWscsa)

You must also install the following packages. Contact Sun technical support for information on how to obtain and configure these packages.

- WBEM Solutions J WBEM Server 1.0
 - Sun StorEdge™ CIM/WBEM Provider SDK (SUNWagsdk package) - A readme file is installed as part of SUNWagsdk package. Follow the instructions in that readme file.
 - Sun StorEdge™ 3510 SMI-S Provider (SUNW3x10a package) - A readme file is installed as part of SUNW3x10a package. Follow the instructions in that readme file.
-

To discover Sun StorEdge 3510 storage systems, you must discover the Sun StorEdge 3510 SMI-S provider. To discover a Sun StorEdge 3510 storage system, you must enter the following information for the instance of the Sun StorEdge 3510 SMI-S provider.

- user name and password used for the system running Sun StorEdge 3510 SMI-S provider
- IP address of the system running Sun StorEdge 3510 SMI-S provider

Important: The management server is unable to display logical volumes configured on Sun StorEdge 3510 storage systems. Any logical volumes as well as the logical drives that comprise them will not appear in the UI. There will be no indication that this happened.

To discover Sun StorEdge 3510 storage systems:

1. Click **Discovery > Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click the **Add Address** button.
5. In the **IP Address/DNS Name** field, type the IP address or DNS name of the system running the Sun StorEdge 3510 SMI-S provider you want to discover.
6. Enter the **User Name** of the system running the Sun StorEdge 3510 SMI-S provider.
7. Enter the **Password** of the system running the Sun StorEdge 3510 SMI-S provider.
8. Retype the password in the **Verify Password** field.
9. In the **Comment** field (optional), type a comment for additional information. The information typed into this field is displayed under the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. To start discovering elements on the network, click the **Start Discovery** button on the IP Addresses tab.

Discovering Sun StorEdge 6920 and 6940 Storage Systems

To discover Sun StorEdge 6920 and 6940 storage systems:

1. Click **Discovery > Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click the **Add Address** button.
5. In the **IP Address/DNS Name** field, type the IP address or DNS name of the storage system you want to discover.
6. Enter the **User Name** of the storage system.
7. Enter the **Password** used to access the storage system.
8. Retype the password in the **Verify Password** field.
9. In the **Comment** field (optional), type a comment for additional information. The information typed into this field is displayed under the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. To start discovering elements on the network, click the **Start Discovery** button on the **IP Addresses** tab.

Discovering Sun StorEdge 6130 Storage Systems

To discover Sun StorEdge 6130 storage systems:

1. Click **Discovery > Setup**.
 2. Select **Step 1** at the top of the page.
 3. Click the **IP Addresses** tab.
 4. Click the **Add Address** button.
 5. In the **IP Address/DNS Name** field, type the IP address or DNS name of the controller or proxy you want to discover.
 6. Leave the **User Name** field blank.
-

7. Leave the **Password** field blank if you do not want to do provisioning on the storage systems. If you want to do provisioning, type the corresponding password for controller or proxy.
8. If you typed a password in the previous step, retype the password in the **Verify Password** field.
9. In the **Comment** field (optional), type a comment for additional information. The information typed into this field is displayed under the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. If you do not plan to use provisioning in the product, select the **Do Not Authenticate** option.
11. Click **OK**.
12. To start discovering elements on the network, click the **Start Discovery** button on the **IP Addresses** tab.

Discovering Xiotech Storage Systems

Important: You must have Xiotech's Intelligent Control (ICON) software installed. If you do not have the software, contact your Xiotech representative.

To discover an Xiotech storage system:

1. Click **Discovery > Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click the **Add Address** button.
5. In the **IP Address/DNS Name** field, enter the IP address or DNS name for the storage system and its namespace. For example:
`<IP address/DNS name>/root/cimv2`
where
 - `<IP address/DNS name>` is the IP address or DNS name of the storage system.
 - `/root/cimv2` is its namespace.
6. You can leave the user name and password for the Xiotech storage system blank.

7. In the **Comment** field (optional), type a comment for additional information. The information typed into this field is displayed under the Comment column in the Addresses to Discover list (**Discovery > Setup**).
8. Do not select the **Do Not Authenticate** option.
9. Click **OK**.
10. To start discovering elements on the network, click the **Start Discovery** button on the IP Addresses tab.

Discovering HP NAS Devices on Windows

In order to discover an HP NAS device on Windows, you must first install a CIM extension on the device. Refer to the *Installation Guide* for information on how to install the CIM extension.

In order to enable NAS support, you must modify a property file on the NAS device. To enable NAS support:

1. Connect to the NAS device on which you have installed the CIM extension.
2. Browse to the installation directory and open the CIMextension/conf directory.
3. Copy the nas.properties-sample file and paste a copy into the same directory.
4. Rename the copied file to nas.properties.
5. Open the file and locate the following line:
`# Set to true to enable NAS data collection; "false" is the default
nas=false`
6. Change the value to `true` to enable NAS support, as shown in the following example:
`nas=true`
7. Save your changes and close the file.
8. Restart the CIM extension. Refer to the *Installation Guide* for information about starting CIM extensions.

To discover an HP NAS device on Windows:

1. Click **Discovery > Setup**.
 2. Select **Step 1** at the top of the page.
 3. Click the **IP Addresses** tab.
 4. Click the **Add Address** button.
 5. In the **IP Address/DNS Name** field, type the IP address or DNS name of the HP NAS device you want to discover.
-

6. Enter the **User Name** of the HP NAS device. You must provide a privileged login.
7. Enter the **Password** used to access the HP NAS device.
8. Retype the password in the **Verify Password** field.
9. In the **Comment** field (optional), type a comment for additional information. The information typed into this field is displayed under the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. To start discovering elements on the network, click the **Start Discovery** button on the **IP Addresses** tab.

Discovering HP NAS Devices on Linux

In order to discover an HP NAS device on Windows, you must first install a CIM extension on the device. Refer to the *Installation Guide* for information on how to install the CIM extension.

In order to enable NAS support, you must modify a property file on the NAS device. To enable NAS support:

1. Connect to the NAS device on which you have installed the CIM extension.
2. Browse to the installation directory and open the CIMextension/conf directory.
3. Copy the nas.properties-sample file and paste a copy into the same directory.
4. Rename the copied file to nas.properties.
5. Open the file and locate the following line:

```
# Set to true to enable NAS data collection; "false" is the default  
nas=false
```
6. Change the value to `true` to enable NAS support, as shown in the following example:

```
nas=true
```
7. Save your changes and close the file.
8. Restart the CIM extension. Refer to the *Installation Guide* for information about starting CIM extensions.

To discover an HP NAS device on Linux:

1. Click **Discovery > Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.

4. Click the **Add Address** button.
5. In the **IP Address/DNS Name** field, type the IP address or DNS name of the HP NAS device you want to discover.
6. Enter the **User Name** of the HP NAS device. You must provide a privileged login.
7. Enter the **Password** used to access the HP NAS device.
8. Retype the password in the **Verify Password** field.
9. In the **Comment** field (optional), type a comment for additional information. The information typed into this field is displayed under the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. To start discovering elements on the network, click the **Start Discovery** button on the **IP Addresses** tab.

Discovering NetApp NAS Devices

Keep in mind the following:

- SMNP must be enabled on the NetApp NAS device before it can be discovered.
- If you want to communicate with the NetApp NAS device via SSL you must set the `cimom.providers.netapp.useSSL` property to “true.” This property is located in the `jboss.properties` file located in the `%JBOSS4_DIST%\server\appiq\conf` directory on the management server. This is a global setting and will cause all NetApp NAS devices to communicate using SSL.
- If you want the management server to be able to receive events from a NetApp NAS device, you must add the IP address of the management server to the NetApp configuration. The management server runs on the same computer running the management server by default.
- You must provide a privileged login, which is one of the following:
 - ☐ the root user
 - ☐ a user belonging to the “Administrators” group. This is a predefined group by NetApp.
 - ☐ a user belonging to a group that has the following roles: `api-*`, `cli-*`, `login-http-admin`, and at least one of the following: `login-console`, `login-telnet`, `login-rsh`, or `login-ssh`
- Administrative HTTP access to the device can be restricted through the `httpd.access` and `httpd.admin.access` options. If that is the case, then the management server needs to be registered with the device. This is done by adding the IP addresses of the management

server to the `httpd.admin.access` option. More information related to this option is available in the NetApp documentation.

To discover a NetApp NAS device:

1. Click **Discovery > Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click the **Add Address** button.
5. In the **IP Address/DNS Name** field, type the IP address or DNS name of the NetApp NAS device you want to discover.
6. Enter the **User Name** of the NetApp NAS device. You must provide a privileged login.
7. Enter the **Password** used to access the NetApp NAS device.
8. Retype the password in the **Verify Password** field.
9. In the **Comment** field (optional), type a comment for additional information. The information typed into this field is displayed under the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. To start discovering elements on the network, click the **Start Discovery** button on the **IP Addresses** tab.

Discovering Sun NAS Devices

Note: You do not need to provide the interop namespace because it is included in the management servers list of default namespaces.

To discover a Sun NAS Device:

1. Click **Discovery > Setup**.
 2. Select **Step 1** at the top of the page.
 3. Click the **IP Addresses** tab.
 4. Click the **Add Address** button.
-

5. In the **IP Address/DNS Name** field, type the IP address or DNS name of the server running the SMI-S provider for the Sun NAS Devices you want to discover.
6. Enter the **User Name** of the CIMOM/provider for the Sun NAS Devices you want to discover. You must provide a privileged login.
7. Enter the **Password** used to access the CIMOM/provider for the Sun NAS Devices you want to discover.
8. Retype the password in the **Verify Password** field.
9. In the **Comment** field (optional), type a comment for additional information. The information typed into this field is displayed under the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. To start discovering elements on the network, click the **Start Discovery** button on the **IP Addresses** tab.

Discovering HP and IBM Tape Libraries

To discover an HP or IBM tape library:

1. Click **Discovery > Setup**.
 2. Select **Step 1** at the top of the page.
 3. Click the **IP Addresses** tab.
 4. Click the **Add Address** button.
 5. In the **IP Address/DNS Name** field, enter the IP address or DNS name for the tape library. Provide the namespace for the tape library after the IP address or DNS name of the tape library. For example: 192.168.1.1:root/ibm, where 192.168.1.1 is the IP address of the tape library and root/ibm is its namespace.
 6. Enter the **User Name** of the system running the tape library.
 7. Enter the **Password** of the system running the tape library.
 8. Retype the password in the **Verify Password** field.
 9. In the **Comment** field (optional), type a comment for additional information. The information typed into this field is displayed under the Comment column in the Addresses to Discover list (**Discovery > Setup**).
 10. Do not select the **Do Not Authenticate** option.
 11. Click **OK**.
-

12. To start discovering elements on the network, click the **Start Discovery** button on the IP Addresses tab.

Step 3 - Build the Topology

This section describes the following:

- “Building the Topology View” on page 88
- “Modifying the Properties of a Discovered Address” on page 90
- “Deleting Elements from the Product” on page 91

Building the Topology View

After you discover elements, the management server requires you to build a topology view, which is a graphical representation of port-level connectivity information.

If a switch has more than one connection to an element, the number of connections is displayed above the line linking the switch and the element. For example, assume the number two is shown between a switch and a storage system. This means the elements have two connections to each other. To view the port details for the connection, right-click the element and select **Show Port Details** from the drop-down menu.

If the topology changes, you can update how the element is viewed in the topology by selecting the element and clicking the Get Topology for Selected button in the Get Topology for discovered elements page (click **Discovery > Topology**). The management server obtains just enough information about where the element is connected in the topology, for example a switch connected to a host.

The management server marks an element with a question mark in the topology if the management server detects an element but it cannot obtain additional information about it. To learn more about fixing detected and/or disconnected elements, see the topic, “Troubleshooting Topology Issues” on page 813.

Important: The user interface may load slowly while the topology is being recalculated. It may also take more time to log into the management server during a topology recalculation.

To obtain enough information to display the topology in System Explorer:

1. Click the **Discovery** menu in the upper-right corner.
2. Click **Topology** in the upper-right corner.
The discovered elements are selected.
3. Select the discovery group from which you want to obtain the topology. If you are obtaining the topology for the first time, make sure **All Discovery Groups** is selected.
You can use discovery groups to break up getting the topology or getting details. For example, instead of getting the topology for all of the discovered elements, you could specify that the management server gets the topology for only the elements in Discovery Group 1, thus, saving you time. You add an element to a discovery group by modifying the properties used to discover the element. See “Modifying the Properties of a Discovered Address” on page 90.

4. Click the **Get Topology** button.
The management server obtains the topology for selected elements.

The management server displays the Log Message page. After the management server builds the topology, a link appears to take you to System Explorer so you can verify the topology view. You can also access System Explorer by clicking **System Explorer** in the left pane.

5. If you see errors in the topology, look at the log messages, which can provide an indication of what went wrong. See the topic, “Viewing Log Messages” on page 60 for more information about the messages that appear in this tab. Look at Event Manager for additional information. Access Event Manager by clicking the **Event Manager** button in the left pane. To obtain troubleshooting information, see the topic, “Troubleshooting Topology Issues” on page 813.

If the topology for an element in your network changes, select the element and click **Get Topology (Discovery > Topology)** to updated the information.

The software obtains just enough information about where the element is connected in the topology, for example a switch connected to a host.

6. See the topics in Chapter 20, “Troubleshooting” on page 785 for more information on troubleshooting problems with discovery and building the topology.
-

Modifying the Properties of a Discovered Address

All elements are initially placed in the Default discovery group. You can then move elements from the Default discovery group to other discovery groups. You can use discovery groups to break up Get Details. For example, you could specify that the management server get Get Details for only the elements in Discovery Group 1, thus, saving you time. This feature is sometimes referred to as segmented replication because you can specify getting Get Details for a segment of the discovered elements.

You can modify the following properties for discovering an device:

- **User name and password** - You can change the user name and password the management server uses to access a device. Whenever a user name and/or password has changed on a device the management server monitors, the management server must be made aware of the change. For example, assume the password for a host was changed. You would need to update the management server database with the new password. See “Modifying a Single IP Address Entry for Discovery” on page 31.
- **Discovery group** - All elements are initially placed in the Default discovery group. You can then move elements from the Default discovery group to other discovery groups. You can use discovery groups to break up getting the topology or Get Details. For example, you could specify that the management server gets the topology or Get Details for only the elements in Discovery Group 1, thus, saving you time. This feature is sometimes referred to as segmented replication because you can specify getting the topology or Get Details for a segment of the discovered elements.

Keep in mind the following:

- You can use this window to change the user name and password stored in the management server's database. It does not change the device's user name and password.
- Discovery groups cannot be renamed or created. You must use the existing discovery groups.
- You can also use the **Move to Discovery Group** button to move multiple elements to another discovery group. See “Moving Elements to Another Discovery Group” on page 103 for more information.

To change the discovery properties of an element:

1. Click **Discovery > Topology** or **Discovery > Details** in the upper-right pane.
2. Click the  button corresponding with the element you want to modify.
3. To move an element to another discovery group, select its new discovery group from the **Discovery Group** drop-down menu.

4. Click **OK** in the Edit Discovered Element window.

Deleting Elements from the Product

When you delete an element, all of its information is removed from the management server. This includes asset information, zoning, events, statistics, and fabrics assigned to switches.

To completely delete an element from the management server you must remove the elements, such as a switch or proxy that were used to discover the element. If you do not delete all switches and proxies that were used to discover the element, the element may reappear the next time you Get Details.

For example, assume you want to delete Switch_A. Switch_B and Switch_C were used to discover Switch_A. If you delete only Switch_B and Switch_A, Switch_A will most likely reappear when you Get Details because it is still accessible by Switch_C.

You can delete an element within the following tools:

- **System Explorer or Chargeback** - Gives you the option of deleting just the element or deleting the element and the elements that use the same switches and proxies for access.
- **Discovery Step 2 (Topology)** - Gives you the option of deleting multiple elements at a time. You are not given a detailed list of other elements you must delete; however, you can use the table on the Discovery screen to determine which switches and proxies provided access.

Deleting an Element Using System Explorer or Chargeback

To delete an element using System Explorer or Chargeback:

1. Do one of the following:
 - **In System Explorer** - Right-click an element and select **Delete Element** from the drop-down menu. Right-click an element and select **Delete Element** from the drop-down menu.
If you are blocking pop-ups and you use the right-click menu to delete an element from System Explorer, the Delete window is blocked and you are unable to delete the element. You must disable the popup blocker before you can delete the element.
 - **In Chargeback** - Click the  button for the element you want to delete.

2. If the element has multiple access points, you are asked which want to delete. Do one of the following:
 - **Delete the element and its access points.** This option lists not only the switch you want to delete, but also the other elements that use the same switches and proxies as the element you want to delete. For example, assume you want to delete Switch_A. Switch_B was used to discover Switch_A. Let's assume Switch_B is also the only path to Switch_D. If you delete Switch_B, you will no longer have access to Switch_D. This option would list Switch_D as one of the other elements that need to be deleted. An access point is the intersection of the IP address and the provider that discovered the IP address. A provider is software that is used to gather information about an element.
 - **Delete the element.** The element may reappear the next time you obtain element details. This is because not all switches and proxies connected to the element have not been removed. For example, assume you want to delete Switch_A. Switch_B is connected to Switch_A. If you do not delete Switch_B, the next time you obtain element details Switch_B will most likely find Switch_A again.
3. Click **OK**.

Deleting Elements Using Discovery Step 2 (Topology)

To delete multiple elements using Discovery Step 2 (Topology):

1. Click **Discovery > Topology**.
2. Determine the access points for the element you want to delete. In the following figure QBrocade2 is accessed by two switches: 192.168.10.25 and 198.168.10.22. You must delete both access points to completely remove the element. As a result, the QBrocade5 switch will also be removed because it has the same access points as QBrocade2.

192.168.10.25	Switch	QBrocade2 , QBrocade5	admin		
192.168.10.21	Switch	QBrocade1	admin		
192.168.10.22	Switch	QBrocade2 , QBrocade5	admin		
192.168.10.24	Switch	QBrocade3 , QBrocade4	admin		

Figure 2-3: Deleting Elements from the Management Server

3. Select all of the access points for the element you want to delete. Then, click the **Delete** button just above the table.

For example, assume you want to delete QBrocade2 in the previous figure. You would select the two listings for QBrocade2 on the Discovered Elements tab and click the **Delete** button in the **Get Topology for Discovered Elements** table. If you delete only one of the listings, QBrocade2 and QBrocade5 still appear in the topology since they are still accessible from one of the switches.

When you are asked if you want to remove the access points and its associated elements, keep in mind these elements will not be deleted if they are accessible from an access point not listed in the Delete Access Points window. For example, assume you selected access point 192.168.10.25 to be deleted. You are then told that switch1 will be deleted along with the access point. Let's assume switch1 is accessible from another access point, 192.168.10.29. When you remove access point 192.168.10.25, switch1 will still be accessible because it can be accessed from another access point that has not been removed.

4. Click **OK** if you want to remove the access points listed in the Delete Access Points window.
The access points are removed. If the elements listed have no other access points, they are no longer accessible from the management server.

Step 4 - Get Details

This section describes the following:

- Get Details on page 93
- Stopping the Gathering of Details on page 95
- Excluding EMC Symmetrix Storage Systems from Force Device Manager Refresh on page 96
- Excluding HDS Storage Systems from Force Device Manager Refresh on page 97

Get Details

Get Details is required to obtain detailed information from discovered elements. Get Details must be performed before you can do provisioning and/or obtain provisioning information, such as data about zone sets and LUN numbers.

Keep in mind the following:

- Get Details takes some time. You might want to perform this process when the network and the managed elements are not busy. To obtain a picture of device connectivity quickly, click the **Get Topology** button on the **Topology** tab.
- During Get Details the topology in System Explorer is recalculated. While the topology is being recalculated, the loading of the user interface may be slow. It may also take more time to log into the management server during a topology recalculation.
- If you have problems obtaining information from Connectrix and McDATA switches during Get Details, see the topic, “Step 2 - Discover Storage Systems, NAS Devices and Tape Libraries” on page 62.
- You can use discovery groups to break up getting the topology or Get Details. For example, instead of Get Details for all of the elements, you could specify that the management server Get Details for only the elements in Discovery Group 1, thus, saving you time. You add an element to a discovery group by modifying the properties used to discover the element. See “Moving Elements to Another Discovery Group” on page 103 for information on how to move one or more multiple elements to a discovery group. You can also move an element to another discovery group when you modify its discovery properties. See “Modifying the Properties of a Discovered Address” on page 90.
- When an element in a given discovery group is updated, its dependent elements are also updated. For example, assume Host_A is the only element in Discovery Group 1. Host_A is connected through a switch and storage system. When you Get Details for Discovery Group 1, you also obtain details from the switch and storage system.
- You can quarantine elements to exclude them from Get Details. See “Placing an Element in Quarantine” on page 104 for more information. Let us assume you want to discover all the elements in a discovery group, except for one. Perhaps the element you want to quarantine is being taken off the network for maintenance. You can use the quarantine feature to exclude one or more elements from discovery.
- If you want to receive status reports about Get Details, see “Configuring E-mail Notification for Get Details” on page 804 for information about how to configure this option.
- If the management server unable to obtain information from a UNIX host during Get Details as a result of a CIM Extension hanging, the management server places the access point where the CIM Extension is located in quarantine. The management server then moves onto getting details for the next element in the Get Details table. These UNIX hosts appear as missing until they are removed from quarantine. See “Removing an Element from Quarantine” on page 105 for information on how to remove an element from quarantine.

To obtain details about the devices on the network:

1. Click **Discovery > Details**.

2. Select **Include infrastructure details**, which gathers information about SAN details.
3. The management server obtains most of its information from device managers for storage systems with external databases, such as HP, HDS, and EMC storage systems. Select **Force Device Manager Refresh** if you want the management server to tell the device managers for your storage systems to obtain the latest information. If you do not select **Force Device Manager Refresh**, the management server gathers information from the external databases with the assumption the information in the external database is up to date. See the following topics for more information: “Excluding EMC Symmetrix Storage Systems from Force Device Manager Refresh” on page 96 and “Excluding HDS Storage Systems from Force Device Manager Refresh” on page 97.
4. Select the discovery group from which you want to obtain Get Details. If you are obtaining Get Details for the first time, make sure **All Discovery Groups** is selected.
5. Click the **Get Details** button.
During Get Details, the software changes its status light from green to red. You can view the progress of gathering details by clicking **Discovery > View Logs**.

When the software completes getting all elements details, it displays “GETTING ALL DETAILS COMPLETED” on the **View Logs** page.
6. See the “Adding a Discovery Schedule” on page 209 for information about automating the gathering of all element details.

Stopping the Gathering of Details

Obtaining details takes some time. If the network and managed elements are busy, you might need to stop the gathering of details and reschedule it for another time.

Important: If you stop the gathering of details, you should reschedule it. This type of collection obtains detailed information of devices in the network.

To stop the gathering of details:

1. Click **Discovery > View Logs**.
 2. On the **View Logs** page, click the “Click here” portion of the following message:
Click here if you wish to stop getting details.
 3. When you are asked if you are sure you want to stop Get Details, click **OK**.
The management server stops gathering details.
-

4. Schedule a time to resume getting details.

Excluding EMC Symmetrix Storage Systems from Force Device Manager Refresh

The management server obtains most of its information about Symmetrix storage systems from the EMC Solutions Enabler (proxy server) it discovered. If the EMC Solutions Enabler does not have the latest information, the management server also displays the outdated information.

To make the management server aware of any changes, make sure the Solutions Enabler it discovered has the latest information. This can be done by forcing the Solutions Enabler to refresh its data. The management server is then made aware of these changes.

When the **Force Device Manager Refresh** option is selected, the management server refreshes discovered EMC Solutions Enabler (proxy server), unless specified. If you do not want an EMC Solutions Enabler to be refreshed, you must assign the Symmetrix storage systems that use the Solutions Enabler to the `cimom.emc.skipRefresh` property, as described in the steps in this section.

To exclude EMC Symmetrix storage systems from a forced refresh:

1. Select **Configuration > Product Health > Advanced**.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the following command. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
`#cimom.emc.skipRefresh=000183500570,000183500575`
4. Return to the Advanced page (**Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree).
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Make sure the property is not commented out by removing the hash (#) symbol in front of the property. Add the serial numbers corresponding to the Symmetrix storage systems you want the refresh to skip. Separate additional serial numbers with a comma, as shown by the following example:

```
cimom.emc.skipRefresh=000183500570,000183500575
```

where 000183500570 and 000183500575 are serial numbers for Symmetrix storage systems. One of the ways to find the serial number is to double-click the storage system in System Explorer. Then, click the **Properties** tab.

7. When you are done, click **Save**.
8. Restart the service for the management server for your changes to take effect:
 - a. Go to the Services window on the management server.
 - b. Right-click **AppStorManager**.
 - c. Select **Restart** from the drop-down menu.While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.
9. To perform the forced refresh, select the **Force Device Manager Refresh** option on the Get Details page.
10. Click **Get Details**.

Excluding HDS Storage Systems from Force Device Manager Refresh

The management server obtains most of its information about the HDS storage systems from the HiCommand Device Manager (proxy server) it discovered. If HiCommand Device Manager, does not have the latest information, the management server also displays the outdated information.

To make the management server aware of any changes, make sure the HiCommand Device Manager it discovered has the latest information. This can be done by forcing the HiCommand Device Manager to refresh its data. The management server is then made aware of these changes.

When the **Force Device Manager Refresh** option is selected, the management server refreshes discovered HiCommand Device Manager (proxy server), unless specified. If you do not want a HiCommand Device Manager to be refreshed, you must assign the HDS storage systems that use HiCommand Device Manager to the `cimom.HdsSkipRefresh` property, as described in the steps in this section.

Important: Before performing any provisioning operations, you should perform a forced refresh.

To exclude HDS storage systems from a forced refresh:

1. Select **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the following command. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
`# cimom.HdsSkipRefresh=61038,61037`
4. Return to the Advanced page (**Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree).
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Make sure the property is not commented out by removing the hash (#) symbol in front of the property. Add the serial numbers corresponding to the HDS storage systems you want the refresh to skip. Separate additional serial numbers with a comma, as shown by the following example:
`cimom.HdsSkipRefresh=61038,61037`
where 61038 and 61037 are serial numbers for HDS storage systems. One of the ways to find the serial number is to double-click the storage system in System Explorer. Then, click the **Properties** tab.
7. When you are done, click **Save**.
8. Restart the service for the management server for your changes to take effect:
 - a. Go to the Services window on the management server.
 - b. Right-click **AppStorManager**.
 - c. Select **Restart** from the drop-down menu.While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.
9. To perform the forced refresh, select the **Force Device Manager Refresh** option on the Get Details or page.
10. Click **Get Details**.

Troubleshooting Mode

Troubleshooting Mode can be used to assist you in identifying and resolving host configuration issues during discovery, as described in the following steps:

1. If errors occur during discovery, an error message will display at the top of the screen below the discovery step where the errors occurred. If you receive an error message, enable Troubleshooting Mode by selecting the **Enable Troubleshooting Mode** check box located near the top of the page for each discovery step.
2. A red icon will display in the **Problems** column for each host for which a problem was detected. Clicking this icon for a particular host will cause a list of troubleshooting tips to display below the **Enable Troubleshooting Mode** check box. Use these tips to assist in the resolution of configuration problems for that host.
3. You can also enter Troubleshooting Mode by clicking the link located in the error message for one of the discovery steps. For example, if you are on discovery step 3, you can click the “Discovery->Setup in Troubleshooting mode” link located in the step 1 error message. Clicking this link will bring you to the step 1 page with Troubleshooting Mode enabled.

Note: Step 3 Troubleshooting Mode only identifies quarantined hosts.

When Troubleshooting Mode is enabled during discovery step 2 or 3, the following additional information is provided to assist in the identification of configuration issues:

- Host OS
- CIM Extension Version
- HBA (Driver Version)
- Multipathing

Managing McDATA and EMC Connectrix Switches

This section describes the following:

- “About Managing McDATA and EMC Connectrix Switches” on page 100
-

- “Adding McDATA and EMC Connectrix Switches” on page 100
- “Removing McDATA and EMC Connectrix Switches” on page 101
- “Swapping McDATA and EMC Connectrix Switches” on page 101

About Managing McDATA and EMC Connectrix Switches

Whenever you add, remove or replace McDATA or EMC Connectrix switches in an already discovered service processor, you must make the management server aware of those changes. After you add these switches to the service processor, you must perform Get Details. The management server obtains information about the new switches from the service processor. See the topic, “Adding McDATA and EMC Connectrix Switches” on page 100 for more information about adding switches.

When you remove switches from the service processor, you must remove them from the management server as well. See the topic, “Removing McDATA and EMC Connectrix Switches” on page 101 for more information about removing switches. When you replace McDATA or EMC Connectrix switches, you add and remove the switches as described previously. See the topic, “Swapping McDATA and EMC Connectrix Switches” on page 101 for more information.

Adding McDATA and EMC Connectrix Switches

After you add switches to an existing service processor, you must perform Get Details, as described in the following steps. If you are adding switches to a service processor that has not been discovered yet, see the topic, “Discovering McDATA and EMC Connectrix Switches” on page 47.

Important: Obtaining details takes some time. You might want to perform this process when the network and the managed elements are not busy.

To Get Details:

1. Click **Discovery > Details**.
 2. Click the **Get Details** button.
-

During Get Details, the software changes its status light from green to red. You can view the progress of gathering details by accessing the logs. See “Viewing Log Messages” on page 60.

Removing McDATA and EMC Connectrix Switches

After removing switches from a service processor, perform the following steps to remove the switches from the management server database.

1. Deleted the switches from the user interface by doing the following. These should be the same switches you removed from the service processor.
 - a. Click **System Explorer** in the left pane.
 - b. Right-click the switch you want to delete.
 - c. Select **Delete Element** from the drop-down menu.
 - d. Select the following option:
Just delete Switch [switch_name]. It may reappear the next time you get topology information or element details.
 - e. Repeat Steps a through d for each switch you want to delete.
2. Verify the switches have been removed from the Elements List in Discovery Steps 2 and 3 by doing the following:
 - a. To verify the switches have been removed from the element list in Discovery Step 3, click **Discovery > Details**.
 - b. To verify the switches have been removed from the element list in Discovery Step 2, click **Discovery > Topology**.

Swapping McDATA and EMC Connectrix Switches

After replacing switches in the service processor, you must make the management server aware of your changes by removing the old switches from the user interface and then performing Get Details so the management server can discover the new switches. If you are adding switches to a service processor that has not been discovered yet, see the topic, “Discovering McDATA and EMC Connectrix Switches” on page 47.

Important: Get Details takes some time. You might want to perform this process when the network and the managed elements are not busy.

Perform the following steps in the management server to swap the switches:

1. Deleted the switches from the user interface by doing the following. These should be the same switches you removed from the service processor.
 - a. Click **System Explorer** in the left pane.
 - b. Right-click the switch you want to delete.
 - c. Select **Delete Element** from the drop-down menu.
 - d. Select the following option:
Just delete Switch [switch_name]. It may reappear the next time you get topology information or element details.
 - e. Repeat Steps a through d for each switch you want to delete.
2. Verify the switches have been removed from the Elements List in Discovery Steps 2 and 3 by doing the following:
 - a. To verify the switches have been removed from the element list in Discovery Step 2, click **Discovery > Topology**.
 - b. To verify the switches have been removed from the element list in Discovery Step 3, click **Discovery > Details**.
3. Click **Discovery > Details**.
4. Click the **Get Details** button.

While Get Details, the software changes its status light from green to red. You can view the progress of gathering details by clicking **Discovery > View Logs**.

When the software completes Get Details, it displays a message saying Get Details is complete on the **View Logs** page.

Assigning a File Extension in Netscape 7

Netscape 7 automatically assigns unknown files an HTML extension. To make Netscape 7 recognize the type of file, you must assign a file extension.

To assign a MIME type:

1. Click the download file link or button in the software.
2. Click the **Advanced** button in the lower-left corner.
3. In the **Description of type** field, delete the existing text and type a description of the file.
4. In the **File extension** field, delete the existing text and type the file extension.
5. Click **OK**.
The next time Netscape 7 sees the associated MIME type, it will assign the extension you typed in the **File Extension** field.
6. Click **OK**.

Filtering Discovery Groups

You can determine which discovery groups are displayed on the Topology (**Discovery > Topology**) and Discovery (**Discovery > Details**) pages by modifying the discovery filter, as described in the following steps:

1. Access the Topology (**Discovery > Topology**) or Get Details (**Discovery > Details**) page.
2. Click the **Custom** button.
3. Select the discovery groups you want to include in Get Details. Deselect the discovery groups you do not want to be included in Get Details.
4. Click **OK**.
Elements in the selected discovery groups are selected on the Get Details page for Discovery. The management server obtains information from the selected elements during Get Details. To learn how to add an element to a different discovery group, see “Modifying the Properties of a Discovered Address” on page 90.

Moving Elements to Another Discovery Group

All elements are initially placed in the Default discovery group. You can then move elements from the Default discovery group to other discovery groups. You can use discovery groups to break up

Get Details and getting the topology. For example, you could specify that the management server gets Get Details for only the elements in Discovery Group 1, thus, saving you time. This feature is sometimes referred to as segmented replication because you can specify getting Get Details for a segment of the discovered elements.

Keep in mind the following:

- Discovery groups cannot be renamed or created. You must use the existing discovery groups.
- You can also use move an element to another discovery group when you modify its discovery properties. See “Modifying the Properties of a Discovered Address” on page 90 for more information.

To move an element to another discovery group:

1. Select the check boxes for the elements you want to move in the Get Details page.
2. Click the **Move to Discovery Group** button.
3. In the Select Discovery Group window, select the new discovery group for the selected elements.
4. Click **OK**.
The elements are moved to the new discovery group.

Placing an Element in Quarantine

When you click the **Get Details** button on the Get Details page, the management server automatically obtains details for the elements in the selected discovery group. Let us assume you want to discover all the elements in a discovery group, except for one. Perhaps the element you want to quarantine is being taken off the network for maintenance. You can use the quarantine feature to exclude one or more elements from discovery.

Note: After you perform Get Details for the discovery group containing the quarantined elements, the quarantined elements appear as missing throughout the product. The management server marks the quarantined elements as missing because it cannot obtain details from the quarantined element.

To quarantine an element:

1. Select the check boxes for the elements you want to quarantine on the Get Details page.
2. Click the **Set Quarantine** button.
3. When you are asked if you want to quarantine the selected elements, click **OK**.

The elements you quarantine appear with a flag () in the Quarantined column on the Get Details page.

The elements are excluded from discovery until you clear them from quarantine.

Removing an Element from Quarantine

To remove an element from quarantine:

1. Select the check boxes for the elements you want to remove from quarantine on the Get Details page.

Quarantined elements appear with a flag () in the Quarantined column on the Get Details page.

2. Click the **Clear Quarantine** button.
3. When you are asked if you want to remove the selected elements from quarantine, click **OK**.

The next time you perform Get Details for the element, the management server gathers data from the element.

Updating the Database with Element Changes

After you have initially discovered the elements, information about them might change. To update database with these changes, perform the steps described in this section.

Keep in mind the following:

- If you change the password of a host after you discover it, you must change the password for the host in the discovery list. Then, you must stop and restart the CIM Extension running on that host before you run a discovery.
- If you are adding, removing or replacing McDATA or Connectrix switches, you must perform different steps. See the topics, “Adding McDATA and EMC Connectrix Switches” on page 100, “Removing McDATA and EMC Connectrix Switches” on page 101, and “Swapping McDATA and EMC Connectrix Switches” on page 101.
- Get Details takes some time. You might want to perform this process when the network and the managed elements are not busy. To obtain a picture of device connectivity quickly, click the **Get Topology** button on the **Topology** tab.

To update the database:

1. Click **Discovery > Details**.
2. Select **Include infrastructure details**, which gathers information about SAN details. **Include backup details** is used for gathering information for Protection Explorer. You do not need to select **Include backup details** unless you have already discovered hosts running backup applications and installed CIM Extensions on those hosts. See “Step 1 - Discovering Your Hosts” on page 117 for more information about discovering master backup servers. The **Include backup details** option gathers the latest information about your backup applications.
3. The management server obtains most of its information from device managers for storage systems with external databases, such as HP, HDS, and EMC storage systems. Select **Force Device Manager Refresh** if you want the management server to tell the device managers for your storage systems to obtain the latest information. If you do not select **Force Device Manager Refresh**, the management server gathers information from the external databases with the assumption the information in the external database is up to date. See the following topics for more information: “Excluding EMC Symmetrix Storage Systems from Force Device Manager Refresh” on page 96 and “Excluding EMC Symmetrix Storage Systems from Force Device Manager Refresh” on page 96.
4. Click the **Get Details** button on the Get Details page.
5. View the status of the gathering of element details by looking in the **View Logs** page. See the topic, “Viewing Log Messages” on page 60 for more information about the messages viewed in this tab.
6. Verify the topology is displayed correctly by accessing System Explorer. Access System Explorer clicking its button in the left pane.

Notifying the Software of a New Element

When you add a new element to the network, such as a host, perform discovery to make the management server aware of the new element.

Keep in mind the following while performing discovery:

- If you change the password of a host after you discover it, you must change the password for the host in the discovery list. Then, you must stop and restart the CIM Extension running on that host.
- If you started a CIM Extension on a Sun Solaris host with the `./start -users` command, a user name provided in the command must be used to discover the host. For example, if you use `./start -users myname:yourname` (where `myname` and `yourname` are valid UNIX accounts) to start the CIM Extension, `myname` or `yourname` and its password must be used to discover the host.
- If you have Brocade switches, download and install the Brocade SMI-S provider software from the following FTP site if this is a new installation of the management server.
`ftp://ftp.compaq.com/pub/products/storageworks/smisproviders/brocade_provider.pdf`

If you are upgrading to build 5.1 from any build prior to 4.2 of the management server your Brocade switches are discovered using the Brocade Fabric Access API after the upgrade. See “Important Information About Upgrading and Brocade Switches” on page 39.

You must also verify that the Rapid program on the switch is set to 1. Rapid must be set to 1 so that the management server can communicate with the switch. See the topic, “Verifying Brocade Rapid Program Is Set to 1” on page 39 for more information.

- Additional steps are required for discovering McDATA switches, and the steps vary according to your network configuration. See the topic, “Discovering McDATA and EMC Connectrix Switches” on page 47 for more information.
- EMC CLARiiON storage systems require additional steps for discovery. See the topic, “Discovering EMC CLARiiON Storage Systems” on page 69 for more information.
- After you discover a McDATA or EMC Connectrix switch, the IP address displayed next to the name of the switch is actually the IP address of the service processor for the switch in the Get Details screens. To find the IP address of the switch, click the link for the switch in the Topology screen (**Discovery > Topology**) or Get Details screen (**Discovery > Details**) and then click the **Properties** tab. The **Properties** tab can also be accessed by double clicking the switch in System Explorer.

For additional information about adding switches, NAS devices, and storage systems, see:

- Table 2-1, “Discovery Steps for Switches, NAS Devices, and Storage Systems,” on page 24
- Chapter 4, “Discovering Applications and Hosts” on page 117

Installing and Discovering the Windows Proxy

This chapter describes the following:

- Installing the Windows Proxy on page 112
- Discovering the Windows Proxy on page 113
- Configuring Windows Proxy Authentication on page 114
- Decreasing the Maximum Java Heap Size on page 116
- Removing the Windows Proxy on page 116

The Windows Proxy is required for the management server on Solaris to communicate with Microsoft Windows hosts using WMI. Microsoft Windows hosts managed via CXWS can be discovered directly and the Windows Proxy is not necessary. The CIM Extensions included with version 5.0 of the management server manage the host via CXWS.

Note: If a host with no agent or an old Windows agent is upgraded to use a CXWS agent, the host must be rediscovered.

After you install and discover the Windows Proxy, discover the hosts on which you have already installed the CIM Extension. See Chapter 4, “Discovering Applications and Hosts” on page 117 for more information.

Keep in mind the following:

- File Server SRM will not work if the hosts behind the Windows proxy are on a private network. If you want to use File Server SRM and your license lets you use this functionality, the Windows hosts cannot be on a private network.
- File Server SRM will also not work if the Windows proxy and the management server do not have network connectivity.
- The management server is unable to discover a database on a Windows host if the host is on a private network behind a Windows proxy. The management server can discover the Windows host through the Windows proxy, but the management server is not able to detect the database.
- If you run into problems with starting the Windows proxy, decrease the maximum Java heap size, as described in “Decreasing the Maximum Java Heap Size” on page 116.
- When the Windows proxy is installed on a new server, the Windows hosts must be re-discovered.

Installing the Windows Proxy

Important: If you are upgrading the Windows proxy, you can install the latest version of the Windows Proxy over the previous version.

To install the Windows proxy:

1. Insert the CD-ROM for the CIM Extensions, go to the Windows directory and then double-click **InstallWindowsProxy.exe**.
 2. When you see the introduction screen, click **Next**.
 3. When you are asked for an installation directory, you can select the default or choose your own. To choose your own directory, click the **Choose** button. You can always display the default directory by clicking the **Restore Default Folder** button. When you are done, click **Next**.
 4. Read the important notes. Then, click **Next**.
 5. Check the pre-installation summary. You are shown the following:
 - ☐ Product Name
 - ☐ Installation Folder
-

- ☐ Disk Space Required
 - ☐ Disk Space Available
6. Do one of the following:
 - ☐ Click **Install** if you agree with the pre-installation summary.
 - ☐ Click **Previous** if you want to modify your selections.The Windows Proxy is installed.
 7. When you have been told the installation has been successful, click **Done** to quit the installation.

Important: Keep in mind that the Windows Proxy automatically starts when the system is restarted. The management server on Solaris can only obtain information from the Windows hosts when the Windows Proxy (AppStorWinProxy service) is running.

8. If the Windows host running the Windows proxy has a private and a public network interface, you must modify the winproxy.conf file.
9. Discover the Windows proxy as described in the topic, “Discovering the Windows Proxy” on page 113.

Discovering the Windows Proxy

Important: Install the Windows proxy before you try the following steps.

Keep in mind the following:

- Install the Windows proxy before you try the following steps.
- The recommended workaround for entering an IP address into the discovery list as well as the Windows Proxy list is to use IP address in one user interface and DNS name in the other.

To discover a Windows proxy:

1. Select **Discovery > Setup** on the management server.
-

2. Click the **Windows Proxy** tab.
3. Enter the following information for the Windows proxy:

Important: A primary key violation error is displayed when you have the same IP address or DNS name listed in both the Discovery list (**Discovery > Setup**) and in the Windows Proxy list. If you have already entered the IP address for a host into the discovery list (**Discovery > Setup**), provide its DNS name in the Windows Proxy list. Likewise, if the DNS name for a host is listed in the Discovery list, provide its IP address in the Windows Proxy list.

- ☐ **IP Address/DNS Name** - The IP address or DNS name used to access the host running the Windows proxy.
 - ☐ **User Name** - The user name of an account used to access the host running the Windows proxy.
 - ☐ **Password** - The password of an account used to access the host running the Windows proxy.
 - ☐ **Verify Password**
4. Click **OK**.
 5. Click the **IP Addresses** tab.
 6. Add the hosts and applications as described in the topic, Chapter 4, "Discovering Applications and Hosts" on page 117.
 7. Click **Start Discovery** if you have already added your hosts and applications for discovery.

Configuring Windows Proxy Authentication

To discover the Windows proxy, the management server requires by default the password and user name of the administrator's account of the host. If you do not want to use the administrator's password for discovery, you can modify the `winproxy.conf` file so that another user name and password can be used. The following options are available to you:

- **Create another Windows account for the host** - You can provide a user name and password other than the administrator's for discovery. Just create a Windows account for
-

the host. You must then set the following properties in the

[install_directory]\WindowsProxy\winproxy.conf file to true:
 winproxy.allowAllWindowsUsers and winproxy.authenticateWindowsUsers.
 After you modify the winproxy.conf file, you must restart the AppStorWinProxy service,
 which is the service for the Windows proxy. Refer to the following example:
 wrapper.java.additional.7=-Dwinproxy.authenticateWindowsUsers=true
 wrapper.java.additional.#=-Dwinproxy.allowAllWindowsUsers=true
 where # is the next consecutive number in the list of properties, for example
 wrapper.java.additional.7. This number can change based on the number of
 properties under # Java Additional Parameters in the winproxy.conf file.

- **Create a user name and password in the winproxy.conf file** - If you do not want to use Windows authentication to create another user account, you can set a user name and password in the winproxy.conf file. Although this user name and password can be used to discover the Windows proxy, it cannot be used to log into the host running the Windows proxy. See the following steps for more information on how to set a user name and password in the winproxy.conf file.

To set a user name and password in the winproxy.conf file:

1. Open the [install_directory]\WindowsProxy\winproxy.conf file in a text editor, such as Notepad.
2. Add the following underlined examples after the last line in put in the application parameters as follows:

```
# Application parameters. Add parameters as needed starting from 1
wrapper.app.parameter.1=com.appiq.cxws.main.WmiMain
wrapper.app.parameter.2=-reloading
wrapper.app.parameter.3=-u
wrapper.app.parameter.4=username
wrapper.app.parameter.5=-p
wrapper.app.parameter.6=password
```

where

- username is the name of the user account
- password is the password for the user account

The numbering must be consecutive. For example, if the last line in # Application Parameters ends at 2 you must number the code as follows:

```
wrapper.app.parameter.3=-u
wrapper.app.parameter.4=username
wrapper.app.parameter.5=-p
wrapper.app.parameter.6=password
```

where

- ☐ `username` is the name of the user account
 - ☐ `password` is the password for the user account
3. Restart the AppStorWinProxy service, which is the service for the Windows proxy.

Decreasing the Maximum Java Heap Size

If you run into problems with starting the Windows proxy on Windows XP, decrease the maximum Java heap size for the Windows proxy as follows:

1. Open the `[install_directory]\WindowsProxy\winproxy.conf` in a text editor, such as Notepad.
2. Change the value of the `wrapper.java.maxmemory` property from 1024 to 512 MB, as shown in the following example:
`wrapper.java.maxmemory=512`
3. Save the `winproxy.conf` file.
4. Restart the AppStorWinProxy service, which is the service for the Windows proxy.

Removing the Windows Proxy

To remove the Windows proxy:

1. Go to the Control Panel in Microsoft Windows.
 2. Double-click **Add or Remove Programs**.
 3. From the Currently installed programs list, select **HDS Windows Proxy**.
 4. Click the **Change/Remove** button.
 5. When you are told the product is about to be uninstalled, click **Uninstall**.
 6. When the program is done with removing the product, click **Done**.
 7. It is highly recommended you reboot the host.
-

Discovering Applications and Hosts

This chapter describes the following:

- “Step 1 - Discovering Your Hosts” on page 117
- “Step 2 - Setting Up Discovery for Applications” on page 125
- “Step 3 - Discovering Applications” on page 143
- “Changing the Oracle TNS Listener Port” on page 147
- “Adding/Modifying Microsoft Exchange Domain Controller Access” on page 148
- “Changing the Password for the Managed Database Account” on page 149
- “Obtaining Disk Drive Statistics from Engenio Storage Systems” on page 150

Step 1 - Discovering Your Hosts

Before you can discover your applications, you must discover their hosts. You discover hosts in the same way you discovered your switches and storage systems. You provide the host's IP address, user name and password. The user name and password must have administrative privileges. Unlike switches and storage systems, you must have installed CIM Extension on the host if you want to obtain detailed information about the host.

The management server automatically detects files servers on hosts through discovery. Before you map the topology (Step 2 in Discovery Setup), make sure the option for File Server SRM is selected, as described in “Step B - Build the Topology” on page 123.

Keep in mind the following:

- If you change the password of a host after you discover it, you must change the password for the host in the discovery list. Then, you must stop and restart the CIM Extension running on that host.
- If your license lets you discover UNIX and/or Linux hosts, the **Test** button for discovery reports SUCCESS from any UNIX and/or Linux hosts on which the management server can detect a CIM Extension. The CIM Extension must be running. The management server reports “SUCCESS” even if your license restricts you from discovering certain types of hosts. For example, assume your license lets you discover Solaris hosts but not AIX hosts. If you click the **Test** button, the management server reports “SUCCESS” for the AIX hosts. You will not be able to discover the AIX hosts. The IP address is not discoverable, because of the license limitation.
- You should have already installed a CIM Extension on the host you want to discover.
- If the management server is running on Sun Solaris, install and configure the Windows Proxy before starting the steps in this section.
- If you want to receive status reports about Get Details, see “Configuring E-mail Notification for Get Details” on page 804 for information about how to configure this option.
- Depending on your license, you may not be able to access File Server SRM and/or monitor certain applications may not be available. See the List of Features to determine if you have access to File Server SRM and/or are able to monitor the other applications. The List of Features is accessible from the Documentation Center (**Help > Documentation Center**). To learn more about File Server SRM, refer to the File Servers Guide, which is also available from the Documentation Center.
- If you are unable to discover a UNIX host because of DNS or routing issues, see “Unable to Discover a UNIX Host Because of DNS or Routing Issues” on page 826.

Discovery of hosts consists of three steps:

- **Setting up** - Finding the elements on the network. See “Step A - Set Up Discovery for Hosts” on page 119.
 - **Topology** - Mapping the elements in the topology. See “Step B - Build the Topology” on page 123.
 - **Details** - Obtaining detailed element information. See “Step D - Obtain Details” on page 124.
-

Step A - Set Up Discovery for Hosts

Note: The steps listed in the following table provide an overview and do not directly correspond to the steps listed in this section.

1. Discover your hosts as described in this section.
2. Build the topology. See “Step B - Build the Topology” on page 123.
3. If you want to view the topology, you can access System Explorer. See “(Optional) Step C - View the Topology” on page 123.
4. Perform Get Details. See “Step D - Obtain Details” on page 124.

Table 4-1: Making the Management Server Aware of Hosts

Step	Description
1	If several of the elements in the same domain use the same name and password, click the Set Default User name and Password link and provide up to three user names and passwords.
2	To enter an IP address range, click the IP Ranges tab and then the Add Range button. If necessary, enter a user name and password for the devices in the range. During a scan of an IP range, the software uses ping; however, ping is not guaranteed to return a response from all devices. If the management server is not able to detect the device during a scan of the IP range, discover the device directly as described in Step 4.
3	Click the Start Scanning button. The management server detects targets it can ping and adds them to the IP Address tab automatically. Ping does not guarantee a response. When you start the discovery process, these addresses are included.
4	If a device was not found during an IP scan, enter an IP address by clicking the IP Addresses tab and then the Add Address button. If necessary, enter a user name and password for the device.

Table 4-1: Making the Management Server Aware of Hosts (Continued)

Step	Description
5	Click the Start Discovery button.
6	Run Get Details. See Get Details takes some time. You might want to perform this process when the network and the managed elements are not busy. Click Discovery > Details .

Keep in mind the following:

- If you started a CIM Extension on a Sun Solaris host with the `./start -users` command, the user name provided in the command must be used to discover the host. For example, if you use `./start -users myname:yourname` (where `myname` and `yourname` are valid UNIX accounts) to start the CIM Extension, `myname` or `yourname` and its password must be used to discover the host.
- If you try to discover a Solaris host with multiple IP address, the management server picks only one IP address for discovery.
- To access a Windows-based device, prefix the user name with `domain_name\`, as shown in the following example. This is required by the Windows login mechanism.
`domain_name\username`
where
 - `domain_name` is the domain name of the element
 - `username` is the name of the account used to access that element
- You can configure the management server to obtain information about your master backup servers at a set interval. See the topic, “Scheduling Backup Collection for Master Servers” in the HP SE User Guide for more information about collectors.

To make the software aware of the devices on the network:

1. Click **Discovery > Setup**.
2. If several of the elements in the same domain use the same name and password, click the **Set Default User Name and Password** link. Provide up to three user names and passwords.
The management server tries the default user names and passwords for elements during discovery. For example, if you have a several hosts using the same user name and password, add their user name and password to the list of default user names and passwords. If one of the hosts is connected to a storage system with another user name and password, you would also add this user name and password to the list. Do not specify the user name and password for the storage system in the individual range because that overrides the default user name and password.

To access a Windows-based device, prefix the user name with `domain_name\`, as shown in the following example. This is required by the Windows login mechanism.

`domain_name\username`

where

- `domain_name` is the domain name of the element
- `username` is the name of the account used to access that element

3. To add an IP address range to scan:
 - a. Click the **IP Ranges** tab.
 - b. Click the **Add Range** button.
 - c. In the **From IP Address** field, type the lowest IP address in the range of the elements you want to discover.
 - d. In the **To IP Address** field, type the highest IP address in the range of the elements you want to discover.
 - e. In the **User Name (Optional)** field, type the user name.
To access a Windows-based device, prefix the user name with the Windows domain name, as shown in the following example. It is required by the Windows login mechanism.

`domain_name\username`

where

- `domain_name` is the domain name of the element
 - `username` is the name of the account used to access that element
- f. In the **Password (Optional)** field, type the password corresponding to the user name entered in the **User Name** field.
 - g. Type the password from the previous step in the **Verify Password** field.
 - h. In the **Comment** field, type a brief description of the servers. For example, Servers in Marketing.
 - i. Click **OK**.
 - j. Repeat steps b through i until all of the IP ranges have been entered.
 - k. Click the **Start Scanning** button.
The elements the management server detects during the scan are added to the **Addresses to Discover** list on the **IP Addresses** tab.
4. To add a single IP address or DNS name to discover:
 - a. Click the **IP Address** tab.
 - b. Click the **Add Address** button.

- c. In the **IP Address/DNS Name** field, type the IP address or DNS name of the device you want to discover.
- d. In the **User Name (Optional)** field, type the user name.
This field can be left blank if one or more of the following conditions are fulfilled:

- The element's user name and password are one of the default user names and passwords.
- The element does not require authentication.

To access a Windows-based device, prefix the user name with the Windows domain name, as shown in the following example.

```
domain_name\username  
where
```

- `domain_name` is the domain name of the machine
- `username` is the name of your network account

- e. In the **Password (Optional)** field, type the corresponding password for the user name entered in the previous step.
This field can be left blank if one or more of the following conditions are fulfilled:
- The element's user name and password are one of the default user names and passwords.
 - The element does not require authentication.
- f. If you typed a password in the previous step, retype the password in the **Verify Password** field.
- g. In the **Comment** field, type a brief description of the server. For example, Server Used for Nightly Backups.
- h. Click **OK**.
5. To start discovering elements on the network, click the **Start Discovery** button on the IP Addresses tab.
The software discovers the IP addresses selected.

During discovery, the following occurs:

- ☐ The software changes the status light from green to orange.
- ☐ You are shown the Log Messages page. To view the status of discovery, click **Discovery > View Logs**.

Discovery is complete when the software displays the DISCOVERY COMPLETED message in the Log Messages field.

Step B - Build the Topology

After you discover elements, the management server requires you build a topology view, which is a graphical representation of port-level connectivity information.

Important: The user interface may load slowly while the topology is being recalculated. It may also take more time to log into the management server during a topology recalculation.

To make the software aware of the devices on the network:

1. Click **Discovery > Topology**.
The discovered elements are selected.
2. Click the **Get Topology** button.
The management server obtains the topology for selected elements.

The management server displays the Log Message page. After the management server builds the topology, a link appears to take you to System Explorer so you can verify the topology view. You can also access System Explorer by clicking System Explorer in the left pane.

3. If you see errors in the topology, look at the log messages, which can provide an indication of what went wrong. Look at Event Manager for additional information. Access Event Manager by clicking the Event Manager button in the left pane. To obtain troubleshooting information, see the “Troubleshooting Topology Issues” on page 813.
If the topology for an element in your network changes, select the element and click **Get Topology** in **Discovery > Topology** to update the information.

The software obtains just enough information about where the element is connected in the topology, for example a switch connected to a host.

(Optional) Step C - View the Topology

Verify the topology is displayed correctly by accessing System Explorer.

To access System Explorer:

1. Click **System Explorer** button in the left pane.
 2. When you are asked if you want to trust the signed applet, click **Always**.
-

The **Always** option prevents this message from being displayed every time you access System Explorer, Capacity Explorer, and Performance Explorer.

The elements are shown connected to each other in the topology.

If you see a question mark above a host, the management server cannot obtain additional information about that element.

If a switch has more than one connection to an element, the number of connections is displayed above the line linking the switch and the element. For example, assume the number two is shown between a switch and a storage system. This means the elements have two connections to each other. To view the port details for the connection, right-click the element and select Show Port Details from the drop-down menu. If the topology changes, you can update how the element is viewed in the topology by selecting the element and clicking the Get Topology for Selected button in the Get Topology for discovered elements page (**Discovery > Topology**). The management server obtains just enough information about where the element is connected in the topology, for example a switch connected to a host.

The management server marks an element as “discovered” in the topology if the management server discovers an element but it cannot obtain addition information about it. To learn more about fixing discovered and/or disconnected elements, see the topic, “Troubleshooting Topology Issues” on page 813.

Step D - Obtain Details

After you obtain the topology of the network, you should obtained detailed information from the discovered elements. Get Details must be performed before you can do provisioning and/or obtain provisioning information, such as data about zone sets and LUN numbers.

Keep in mind the following:

- Obtaining details takes some time. You might want to perform this process when the network and the managed elements are not busy. To obtain a picture of device connectivity quickly, click the **Get Topology** button on the Topology tab.
 - During Get Details the topology in System Explorer is recalculated. While the topology is being recalculated, the loading of the user interface may be slow. It may also take more time to log into the management server during a topology recalculation.
 - If Get Details includes an AIX host, the system log displays three SCSI errors (2 FSCSI error and 1 FCS error) per IBM adapter port. You can ignore these errors.
-

- You can quarantine elements to exclude them from Get Details. See “Placing an Element in Quarantine” on page 105 for more information. Let us assume you want to discover all the elements in a discovery group, except for one. Perhaps the element you want to quarantine is being taken off the network for maintenance. You can use the quarantine feature to exclude one or more elements from discovery.
- When an element in a given discovery group is updated, its dependent elements are also updated. For example, assume Host_A is the only element in Discovery Group 1. Host_A is connected through a switch and storage system. When you Get Details for Discovery Group 1, you also obtain details from the switch and storage system.
- If the management server unable to obtain information from a UNIX host during Get Details as a result of a CIM Extension hanging, the management server places the access point where the CIM Extension is located in quarantine. The management server then moves onto getting details for the next element in the Get Details table. These UNIX hosts appear as missing until they are removed from quarantine. See “Removing an Element from Quarantine” on page 106 for information on how to remove an element from quarantine.

To obtain details:

1. Click the **Discovery > Details** in the upper-right corner.
 2. Click the **Get Details** button.
- During Get Details, the software changes its status light from green to red. You can view the progress of gathering details by clicking **Discovery > View Logs**.

When the software completes getting all elements details, it displays GETTING ALL DETAILS COMPLETED on the View Logs page.

For additional information, see “Step 1 - Discover Switches” on page 34 for information on how to automate the gathering of all element details. If you run into problems with discovery, see “Troubleshooting” on page 785.

Step 2 - Setting Up Discovery for Applications

Keep in mind the following when discovering applications:

- Make a list of the applications you want to monitor. Configure your applications first as described in this section and then run discovery.
-

- You should have already installed a CIM Extension on the hosts that have the applications you want to discover. After you installed the CIM Extension, you should have already discovered the host. See “Step 1 - Discovering Your Hosts” on page 117.

You can configure the management server to monitor hosts and applications, such as Oracle, Microsoft Exchange server, and Sybase Adaptive Server Enterprise, in addition to Microsoft SQL servers and file servers. If you want to obtain detailed information about the host and its applications, you must install a CIM Extension on the host, as described in the previous chapters.

The following is an overview of what you need to do. It is assumed you have already discovered the hosts running your applications. See “Step 1 - Discovering Your Hosts” on page 117.

Then, set up the configurations for your applications on the management server. Some applications may require you to provide additional discovery information about the application. Finally, perform discovery, (discovery from SE only) map the elements in the topology, and then Get Details. Get Details takes some time. Perform this step when the network is not busy. More details about the steps mentioned above are provided later.

See the following topics for more information:

- “Monitoring Oracle” on page 126
- “Monitoring Microsoft SQL Server” on page 132
- “Monitoring Sybase Adaptive Server Enterprise” on page 138
- “Monitoring Microsoft Exchange” on page 142

Monitoring Oracle

To monitor and manage Oracle, you must do the following:

- “Step A - Create the APPIQ_USER Account for Oracle” on page 127
- “Step B - Provide the TNS Listener Port” on page 130
- “Step C - Set up Discovery for Oracle 10g” on page 130

After you complete these steps, you must discover Oracle, and perform Get Details. See “Step 3 - Discovering Applications” on page 143.

Important: Before you begin these steps, make sure you purchased the module that lets you monitor Oracle. Contact your customer support if you are unsure if you purchased this module.

Step A - Create the APPIQ_USER Account for Oracle

The management server accesses Oracle through the APPIQ_USER account. This account is created when you run the `CreateOracleAct.bat` script on Microsoft Windows or `CreateOracleAct.sh` on UNIX on the computer running the Oracle database you want to monitor. The account has create session and select dictionary privileges to be used with the management server.

Keep in mind the following:

- The `CreateOracleAct.bat` script must run under SYS user.
- Create APPIQ_USER account on Oracle Database you want to monitor, not on the management server.
- You should have already installed the database for the management server.
- Verify that the instance TNS (Transparent Name Substrate) listener is running so that the management server can find the Oracle installation and its instances. For example on Microsoft Windows 2000, you can determine if the instance TNS listener is running by looking in the Services window for OracleOraHome92TNSListener. The name of the TNS listener might vary according to your version of Oracle. Refer to the Oracle documentation for information about verifying if the instance TNS listener is running. You can also verify the listener is running by entering the following at the command prompt: `lsnrctl status`. If the listener is not running you can start it by typing `lsnrctl start` on command line.
- Make sure you have all the necessary information before you begin the installation. Read through the following steps before you begin.

To create the Oracle user for management server:

1. Do one of the following:
 - **To run the script on IBM AIX, SGI IRIX, HP-UX, Linux or Sun Solaris**, log into an account that has administrative privileges, mount the CIM Extensions CD-ROM (if not auto-mounted), and go to the `/DBIQ/Oracle/unix` directory by typing the following:

```
# cd /cdrom/cdrom0/DBIQ/Oracle/unix
```

where `/cdrom/cdrom0` is the name of the CD-ROM drive

- ❑ **To run the script on Microsoft Windows**, go to the `DBIQ\Oracle\win` directory on the CIM Extensions CD-ROM.

Important: You must complete the following steps.

2. Verify you have the password to the SYS user account.
You are prompted for the password for this user account when you run the script.
3. Run the `CreateOracleAct.bat` script on Microsoft Windows or `CreateOracleAct.sh` script on a UNIX operating system on the computer with the Oracle database.
The script creates a user with create session and select dictionary privilege on a managed Oracle instance.

Note: You can use a remote Oracle client to run this script.

4. Specify the Oracle instance name, which must be visible to the client, as the first input when running the script. The script prompts you for the name of the Oracle instance on which to create user for Oracle management packages and the password of the SYS account.
You are asked to specify the default and temporary tablespaces for APPIQ_USER during the installation. You can enter users as default and temp as temporary if these tablespaces exist in the Oracle Instance.
5. Repeat the previous step for each Oracle instance you want to manage.
This script does the following in order:
 - ❑ Creates the APPIQ_USER account.
 - ❑ Grant create session and select on dictionary tables privileges to APPIQ_USER enabling management server to view statistics for the Oracle instances.

Removing the APPIQ_USER Account for Oracle

If you no longer want the management server to monitor an Oracle instance, you can remove the APPIQ_USER account for that Oracle instance by running the `UninstallOracleAct.bat` script on Windows or `UninstallOracleAct.sh` script on the Solaris platform.

Keep in mind the following:

- Before you remove the APPIQ_USER account for an Oracle instance, make sure no processes are running APPIQ_USER for that Oracle instance. The management server uses APPIQ_USER to obtain information about the Oracle database. For example, a process would be using APPIQ_USER if someone was using Performance Explorer to view monitoring statistics about that Oracle instance. One of the ways to make sure APPIQ_USER is not being used is to temporarily remove the host running Oracle (**Discovery > Topology**). After you removed the APPIQ_USER account for Oracle, discover and perform Get Details for the host if you want to continue monitoring it.
- If you receive a message about not being able to drop a user that is currently connected while you are removing the APPIQ_USER account for Oracle, re-run the script for removing APPIQ_USER.

To remove the APPIQ_USER account for that Oracle instance:

1. If you plan to remove the management software for Oracle from a Solaris host, do the following:
 - a. Log into an account that has administrative privileges.
 - b. Mount the CIM Extensions CD-ROM (if not auto-mounted).
 - c. Go to the `/DBIQ/Oracle/unix` directory by typing the following:

```
# cd /cdrom/cdrom0/DBIQ/Oracle/unix
```

where `/cdrom/cdrom0` is the name of the CD-ROM drive
2. If you plan to remove the management software for Oracle from a computer running Windows, go to the `\DBIQ\Oracle\win` directory on the CD-ROM.
3. Verify you have the password to the SYS user account.
You are prompted for the password for this user account when you run the script.
4. Run the `UninstallOracleAct.bat` for Windows or `UninstallOracleAct.sh` script for UNIX platform on the computer with the Oracle database.
5. This script removes the management software for the specified Oracle instance.

Note: You can use a remote Oracle client to run this script.

6. When you are asked for the Oracle instance name, enter the name of the Oracle instance you do not want the management server to monitor. The name must be visible to the client.
 7. Provide the password for the SYS user account.
The APPIQ_USER account for the specified Oracle instance is removed. The management server can no longer monitor that Oracle instance.
-

Step B - Provide the TNS Listener Port

If your Oracle instances use a different TNS Listener Port than 1521, change the port as described in the following steps:

1. Select **Discovery > Setup**. Then, click the **Applications** tab.
The TNS Listener Port setting applies to all Oracle instances you monitor.
2. To assign a new port, click the **Create** button for the Oracle Information table.
3. Type the new port number and click **OK**.
4. If necessary, click the  button to remove the old port number.

Important: Monitoring Oracle 10g and Oracle clusters require an additional step. If you are not monitoring Oracle 10g and Oracle clusters, see “Step 3 - Discovering Applications” on page 143.

Step C - Set up Discovery for Oracle 10g

Note: If you are discovering an Oracle cluster, see “Discovering Oracle Clusters” on page 131.

To monitor Oracle 10g, provide additional information as described in the following steps:

1. Select **Discovery > Setup**. Then, click the **Applications** tab.
 2. Click the **Create** button for the Database Information table.
 3. In the **Host IP/DNS Name** field, type the IP address or DNS name of the host running Oracle.
The **Management IP/DNS Name** field is optional.
 4. In the **Server Name** field, type the Oracle System Identifier (SID) of the Oracle database you want to monitor.
 5. In the **Port Number** field, type the monitored port.
If you are not sure of the monitored port, check the listener.ora file of the monitored database application. You can find the listener.ora file in the following directory on the host of the monitored database. Do not look for the listener.ora file on the management server for this information.
-

```
%ORA_HOME%\network\admin\listener.ora
```

The port can be found in the following code:

```
LISTENER =  
  (DESCRIPTION_LIST =  
    (DESCRIPTION =  
      (ADDRESS_LIST =  
        (ADDRESS = (PROTOCOL = TCP) (HOST = localhost) (PORT = 1521))  
        (ADDRESS = (PROTOCOL = IPC) (KEY = EXTPROC0))  
      )  
    )  
  )
```

6. Select **ORACLE** from the Database Type menu.
7. Click **OK**.

Discovering Oracle Clusters

Perform the following steps for each node in the cluster:

1. Install the CIM Extension on each node in the cluster.
2. Create the appiq_user account on each node in the cluster. See “Step A - Create the APPIQ_USER Account for Oracle” on page 127.
3. Click **Discovery > Setup**. Then, discover the host for the first node by clicking the **Add Address** button and providing the appropriate information for discovering the host, as described in “Adding a Single IP Address or DNS Name for Discovery” on page 29.
4. Discover first Oracle node as follows:
 - a. Select **Discovery > Setup**. Then, click the **Applications** tab.
 - b. Click the **Create** button for the Database Information table.
 - c. In the **Host IP/DNS Name** field, type the IP address or DNS name of the host running Oracle.
In the **Management IP/DNS Name** field, type the IP address the listener is listening on for the Oracle instance. The IP address can be a virtual IP or a host IP. You can find the IP address in the listener.ora file for the monitored database. You can find the listener.ora file in the following directory on the host of the monitored database. Do not look for the listener.ora file on the management server for this information.

```
%ORA_HOME%\network\admin\listener.ora
```
 - d. In the **Server Name** field, type the Oracle System Identifier (SID) of the Oracle database you want to monitor.

- e. In the **Port Number** field, type the monitored port.
If you are not sure of the monitored port, check the listener.ora file of the monitored database application. You can find the listener.ora file in the following directory on the host of the monitored database. Do not look for the listener.ora file on the management server for this information.

Microsoft Windows:

%ORA_HOME%\network\admin\listener.ora

Sun Solaris:

\$ORACLE_HOME/network/admin/listener.ora

The port can be found in the following code:

```
LISTENER =  
  (DESCRIPTION_LIST =  
    (DESCRIPTION =  
      (ADDRESS_LIST =  
        (ADDRESS = (PROTOCOL = TCP) (HOST = localhost) (PORT =  
1521))  
        (ADDRESS = (PROTOCOL = IPC) (KEY = EXTPROC0))  
      )  
    )  
  )
```

- f. Select **ORACLE** from the Database Type menu.
g. Click **OK**.
5. Repeat Steps 3 and 4 for each node in the cluster.

Monitoring Microsoft SQL Server

To manage and monitor Microsoft SQL Servers, you must do the following:

- “Step A - Create the APPIQ_USER for the SQL Server” on page 133
- “Step B - Provide the Microsoft SQL Server Name and Port Number” on page 135

Important: Make sure the Microsoft SQL server database is in “Mixed Mode authentication.” To switch to mixed mode authentication, see “Switching to Mixed Mode Authentication” on page 133.

Switching to Mixed Mode Authentication

Important: Do not make security changes to your Microsoft SQL Server installation unless you are familiar with the security requirements of your site.

Microsoft SQL Server must be running in Mixed Mode Authentication. You can switch to Mixed Mode Authentication as follows:

1. Open SQL Server Enterprise Manager (**Start Menu > Programs > Microsoft SQL Server > Enterprise Manager**).
2. Expand the tree-control until you can see your server.
3. Right-click the server name.
The SQL Server Properties (Configure) window appears.
4. Click the **Security** tab.
5. For “Authentication,” select **SQL Server and Windows**.
6. For “Startup service account,” select **System Account**

Step A - Create the APPIQ_USER for the SQL Server

The management server accesses SQL Server through the APPIQ_USER account. This account is created when you run the `CreateSQLServerAct.bat` script on Microsoft Windows on the computer running the SQL Server database you want to monitor. The account has create session and select dictionary privileges to be used with the management server.

Keep in mind the following:

- The script must run under SA user. To verify that the SA account is enabled, launch SQL Server’s Query Analyzer tool and attempt to connect to the database as SA with the SA user’s password.
- Obtain the SQL Server name before you run the script
- Create APPIQ_USER account on SQL Server database you want to monitor.
- You should have already installed the database for the management server.
- Make sure you have all the necessary information before you begin the installation. Read through the following steps before you begin.

To create the APPIQ_USER account for SQL Server:

1. To run the script on Microsoft Windows, go to the `DBIQ\sqlserver\win` directory on the CIM Extensions CD-ROM.

Important: You must complete the following steps.

2. Verify you have the password to the SA user account.
You are prompted for the password for this user account when you run the script.
3. Run the `CreateSQLServerAct.bat` script on Microsoft Windows on the computer with the SQL Server database.
The script creates a user with login to master and select privilege on data dictionary tables on a managed SQL Server instance.

Note: You can use a remote SQL Server `isql` to run this script.

4. Type the SQL Server instance name, which must be visible to the client, as the first input when running the script. The script prompts you for the name of the SQL Server on which to create user for SQL Server management packages and the password of the SA account. This script does the following in order:
 - ☐ Creates the `APPIQ_USER` account.
 - ☐ Grant create session and select on dictionary tables privileges to `APPIQ_USER` enabling management server to view statistics for the SQL Server.
 5. To determine if the `appiq_user` user was added correctly to your SQL server:
 - a. Open SQL Server Enterprise Manager.
 - b. Expand the user interface for SQL Server Enterprise Manager. Then, expand the specific SQL Server and select **Security**.
 - c. Double-click **Logins** and view the list of users authorized to access the SQL Server.
 - d. If `appiq_user` is not listed, the management server is not able to discover the database.
 6. To determine if the SQL Server is ready to accept connections from the management server:
 - a. Connect to the SQL Server installation through Query Analyzer using the account “`appiq_user`” and the password “password.”
 - b. Create a sample ODBC datasource for the SQL Server installation using the `appiq_user` account.
-

- c. Click the **Test** button to test the datasource.
7. Repeat steps 4 through 6 for each SQL Server you want to manage.

Step B - Provide the Microsoft SQL Server Name and Port Number

The server name for the Microsoft SQL server and port number for managing a SQL database must be provided in the following steps:

Important: You must provide the host name in the Host IP/DNS Name field. You cannot use localhost or parenthesis.

If you have name resolutions issues, your server may be discovered; however, your applications will not be discovered. You can address the name resolution issues by adding entries within the hosts file on the management server for the systems in question.

When configuring the System Application Discovery Settings for SQL servers, the following needs to be specified as described in the steps within this section:

- **Host IP/DNS Name:** <IP Address>
- **Database Server:** <SQL Server Name>
- **Port Number:** <SQL Port #>
- **Database Type:** SQLSERVER

To add information for discovering a SQL server:

1. Select **Discovery > Setup**. Then, click the **Applications** tab.
 2. Click the **Create** button for the Database Information table.
 3. In the **Host IP/DNS Name** field, type the IP address or DNS name of the host running SQL Server.
 4. You can leave the **Management IP/DNS Name** field blank. This field is for Oracle clusters. When you leave the **Management IP/DNS Name** field blank the management server automatically lists the DNS name or IP address of the host under the **Host IP/DNS Name** column and **Management IP/DNS Name** column.
 5. In the Database Server field, type the SQL database server name you want to monitor.
-

The SQL Server name is either the Windows system name (default) or the name specified when the SQL server was installed.

SQL Enterprise Manager displays the SQL server name as one of the following:

- ☐ As the name specified at the time the SQL server was installed.
- ☐ As the Windows system name (Windows 2000)
- ☐ As the local name (Windows 2003)

For example, if a Windows 2003 server called SQLTEST has an IP address of 192.168.2.10 with the default SQL port (1433) and shows the name of (local) within SQL Enterprise Manager, the correct system application discovery settings on the management server would be the following:

- ☐ **Host IP/DNS Name:** 192.168.2.10
 - ☐ **Database Server:** SQLTEST
 - ☐ **Port Number:** 1433
 - ☐ **Database Type:** SQLSERVER
6. In the **Port Number** field, type the port that SQL is using. If you do not provide a port number, the management server assumes the SQL server is using port 1433 (default). To determine the correct SQL Port Number that the SQL Server is using, follow these steps:
 - a. Open SQL Server Enterprise Manager.
 - b. Expand the user interface for SQL Server Enterprise Manager, and then select the specific SQL server. Right-click and then select **Properties** from the drop-down menu.
 - c. Click the **Network Configurations** button. On the General Tab, select the TCP/IP entry under the Enabled Protocols section. Then, click the **Properties** button.
 - d. The resulting window shows you the TCP/IP port your SQL server uses. Provide this port number in the **Port Number** field on the management server.
 7. Select **SQLSERVER** from the Database Type menu.
 8. Click **OK**.

Important: Perform Get Details for your inputs to take effect. See “Step 3 - Discovering Applications” on page 143.

Removing the APPIQ_USER Account for SQL Server

Important: Before you remove the APPIQ_USER account for the SQL Server databases on a host, make sure no processes are running APPIQ_USER for that SQL Server database. The management server uses APPIQ_USER to obtain information about a SQL Server database. One of the ways to make sure APPIQ_USER is not being used is to temporarily remove the host running SQL Server (**Discovery > Topology**). After you removed the APPIQ_USER account for SQL Server, discover and perform Get Details for the host if you want to continue monitoring it.

To remove the APPIQ_USER account for the SQL Server databases on a host:

1. To run the script on Microsoft Windows, go to the `DBIQ\sqlserver\win` directory on the CIM Extensions CD-ROM.

Important: You must complete the following steps.

2. Verify you have the password to the server administrator user account. You are prompted for the password for this user account when you run the script.
3. Run the `UninstallSQLServerAct.bat` script on Microsoft Windows on the computer with the SQL Server database.
4. Type the name of the SQL Server server.
5. Type the password for the server administrator account.
The account for APPIQ_USER is removed. The management server can no longer monitor the SQL Server databases on this host.

Deleting SQL Server Information

If you do not want the management server to monitor a SQL Server instance, you can remove its information, as described in the following steps:

1. Select **Discovery > Setup**. Then, click the **Applications** tab.
 2. In the Database Information table, click the  button, corresponding to the SQL Server instance you do not want the management server to monitor.
-

3. Perform Get Details to make the management server aware of your changes.

Monitoring Sybase Adaptive Server Enterprise

If you want to monitor Sybase Adaptive Server Enterprise you must:

- Create APPIQ_USER account on the database for Sybase
- Provide the database server name and port number
- Discover the application.

The required drivers for Sybase Adapter Server Enterprise were automatically installed along with the management server.

Important: Before you begin these steps, make sure you purchased Sybase IQ, which is the module that lets you monitor Sybase Adaptive Server Enterprise. Contact your customer support if you are unsure if you purchased this module.

Step A - Create the APPIQ_USER account for Sybase

The management server accesses Sybase through the APPIQ_USER account. This account is created when you run the CreateSybaseAct.bat script on Microsoft Windows or CreateSybaseAct.sh on UNIX on the computer running the Sybase database you want to monitor. The account has create session and select dictionary privileges to be used with the management server.

Keep in mind the following:

- The script must run under SA user.
 - Obtain the Sybase server name before you run the script
 - Create APPIQ_USER account on Sybase Database you want to monitor.
 - You should have already installed the database for the management server.
 - Make sure you have all the necessary information before you begin the installation. Read through the following steps before you begin.
-

To create the APPIQ_USER account for the Sybase server:

1. Do one of the following:
 - ❑ **To run the script on IBM AIX, SGI IRIX, or Sun Solaris**, log into an account that has administrative privileges, mount the CIM Extensions CD-ROM (if not auto-mounted), and go to the /DBIQ/sybase/unix directory by typing the following:

```
# cd /cdrom/cdrom0/DBIQ/sybase/unix
```

where /cdrom/cdrom0 is the name of the CD-ROM drive
 - ❑ **To run the script on Microsoft Windows**, go to the \DBIQ\sybase\win directory on the CIM Extensions CD-ROM.

Important: You must complete the following steps.

2. Verify you have the password to the SA user account.
You are prompted for the password for this user account when you run the script.
3. Run the CreateSybaseAct.bat script on Microsoft Windows or CreateSybaseAct.sh script on the UNIX operating system on the computer with the Sybase database.
The script creates a user with login to master and select privilege on data dictionary tables on a managed Sybase instance.

Note: You can use a remote Sybase isql to run this script.

4. Type the Sybase instance name, which must be visible to the client, as the first input when running the script. The script prompts you for the name of the sybase server on which to create user for Sybase management packages and the password of the SA account.
5. Repeat the previous step for each Sybase server you want to manage.
This script does the following in order:
 - ❑ Creates the APPIQ_USER account.
 - ❑ Grant create session and select on dictionary tables privileges to APPIQ_USER enabling management server to view statistics for the Sybase server.

Removing the APPIQ_USER Account for Sybase

Important: Before you remove the APPIQ_USER account for the Sybase databases on a host, make sure no processes are running APPIQ_USER for that Sybase database. The management server uses APPIQ_USER to obtain information about a Sybase database. One of the ways to make sure APPIQ_USER is not being used is to temporarily remove the host running Sybase (**Discovery > Topology**). After you removed the APPIQ_USER account for Sybase, discover and perform Get Details for the host if you want to continue monitoring it.

To remove the APPIQ_USER account for the Sybase databases on a host:

1. Do one of the following:
 - ☐ To run the script on IBM AIX, SGI IRIX, or Sun Solaris, log into an account that has administrative privileges, mount the CD-ROM (if not auto-mounted), and go to the /DBIQ/sybase/unix directory by typing the following:

```
# cd /cdrom/cdrom0/DBIQ/sybase/unix
```

where /cdrom/cdrom0 is the name of the CD-ROM drive
 - ☐ To run the script on Microsoft Windows, go to the \DBIQ\sybase\win directory on the CD-ROM.

Important: You must complete the following steps.

2. Verify you have the password to the SA user account.
You are prompted for the password for this user account when you run the script.
 3. Run the UninstallSybaseAct.bat script on Microsoft Windows or UninstallSybaseAct.sh script on the UNIX operating system on the computer with the Sybase database.
 4. Type the name of the Sybase server.
 5. Type the password for the SA account.
The account for APPIQ_USER is removed. The management server can no longer monitor the Sybase databases on this host.
-

Step B - Provide the Sybase Server Name and Port Number

You must provide the Sybase server name and port number for managing the Sybase database in the following steps:

To add information for discovering Sybase Adaptive Server Enterprise:

1. Select **Discovery > Setup**. Then, click the **Applications** tab.
2. Click the **Create** button for the Database Information table.
3. In the **Host IP/DNS Name** field, type the IP address or DNS name of the host running Sybase.
4. You can leave the **Management IP/DNS Name** field blank. This field is for Oracle clusters. When you leave the **Management IP/DNS Name** field blank the management server automatically lists the DNS name or IP address of the host under the **Host IP/DNS Name** column and **Management IP/DNS Name** column.
5. In the **Server Name** field, type the Sybase database you want to monitor.
6. In the **Port Number** field, type the port that Sybase is using.
7. Select **SYBASE** from the Database Type menu.
8. Click **OK**.

Important: Perform Get Details for your inputs to take effect. See “Step 3 - Discovering Applications” on page 143.

Deleting Sybase Information

If you do not want the management server to monitor a Sybase instance, you can remove its information, as described in the following steps:

1. Select **Discovery > Setup**. Then, click the **Applications** tab.
2. In the Database Information table, click the  button, corresponding to the Sybase instance you do not want the management server to monitor.
3. Perform Get Details to make the management server aware of your changes.

Monitoring Microsoft Exchange

To monitor Microsoft Exchange, you must make the management server aware of domain controller access. After information for controller access has been added, discover Microsoft Exchange, map the topology and perform Get Details. To save time, delay these steps until you have added the configurations for your other applications and hosts.

To monitor Microsoft Exchange, you must:

- Add information for Microsoft Exchange Domain Controller Access
- Discover the application.

Adding Microsoft Exchange Domain Controller Access

To obtain information about your Microsoft Exchange servers, you must provide the user name and password for at least a primary domain controller, in addition to a DNS name, as described in the following steps.

Important: The hosts should recognize the management server by name, as a reverse look-up is required by operating system security as well as Microsoft Exchange.

To provide information about the Microsoft Exchange servers:

1. Select **Discovery > Setup**. Then, click the **Applications** tab.
The information you provide for the primary domain controller and backup domain controller apply to all Microsoft Exchange servers you discover.
 2. In the Microsoft Exchange Configuration section, click the **Edit** button.
 3. Under the Primary Domain Controller section, perform the following steps:
 - a. In the **Host Name** field, type the fully qualified DNS name for the domain controller.
 - b. In the **User Name** field, type the user name for accessing the Microsoft Exchange server.
 - c. In the **Domain Password** field, type the corresponding password for accessing the Microsoft Exchange server.
 - d. In the **Verify Password** field, re-type the password for verification.
 4. Under the Backup Domain Controller section, perform the following steps:
-

- a. In the **Host Name** field, type the fully qualified DNS name for the domain controller.
 - b. In the **User Name** field, type the user name for accessing the Microsoft Exchange server.
 - c. In the **Domain Password** field, type the corresponding password for accessing the domain controller.
 - d. In the **Verify Password** field, re-type the password for verification.
5. Click the **OK** button.

Important: You must discover the host running Microsoft Exchange. See “Step 3 - Discovering Applications” on page 143.

Deleting a Microsoft Exchange Domain Controller

To delete a Microsoft Exchange domain controller:

1. Select **Discovery > Setup**. Then, click the **Applications** tab.
2. Click the  button, corresponding to the domain controller you want to remove.
3. Perform Get Details for your changes to take effect.

Step 3 - Discovering Applications

This step assumes you have already discovered your hosts and provided discovery information for your applications. To discover an application, do the following;

- Detect the application (“Step A - Detect Your Applications” on page 144)
- Obtain topology information about the application (“Step B - Obtain the Topology” on page 145)
- Perform Get Details (“Step C - Obtain Get Details” on page 145)

Keep in mind the following:

- This section assumes you have already set up the discovery configurations for your applications as described in “Step 2 - Setting Up Discovery for Applications” on page 125.
-

- If DNS records for your Microsoft Exchange servers are outdated or missing, the discovery of Microsoft Exchange may fail because Microsoft Exchange is dependant on Active Directory, which is dependant on DNS. Since Active Directory is dependant on DNS, Active Directory replication and Active Directory lookups may fail or contain errors if DNS records are not accurate.
- The management server is unable to discover Oracle on a Windows host if the host is on a private network behind a Windows proxy. The management server can discover the Windows host through the Windows proxy, but the management server is not able to detect Oracle.

Discovery consists of three steps:

- **Setting up** - Finding the elements on the network.
- **Topology** - Mapping the elements in the topology.
- **Details** - Obtaining detailed element information.

Step A - Detect Your Applications

To make the software aware of the applications on the network:

1. Click the **Discovery > Setup**.
2. To start discovering elements on the network, click the **Start Discovery** button on the IP Addresses tab.
The software discovers the IP addresses selected.

During discovery, the following occurs:

- ☐ The software changes the status light from green to orange.
- ☐ You are shown the Log Messages page. To view the status of discovery, click **Discovery > View Logs**.

Discovery is complete when the software displays the DISCOVERY COMPLETED message in the Log Messages field.

Keep in mind the following:

- ☐ If DNS records for your Microsoft Exchange Servers are outdated or missing, the discovery of Microsoft Exchange may fail because Microsoft Exchange is dependant on Active Directory, which is dependant on DNS. Since Active Directory is dependant on DNS, Active Directory replication and Active Directory lookups may fail or contain errors if DNS records are not accurate.
-

- If you are having problems discovering an element, see “Troubleshooting Discovery and Get Details” on page 802.

Step B - Obtain the Topology

The user interface may load slowly while the topology is being recalculated. It may also take more time to log into the management server during a topology recalculation.

To obtain the topology:

1. Click **Discovery > Topology**.
The discovered elements are selected.
2. Click the **Get Topology** button.
The management server obtains the topology for selected elements.
3. Select the discovery group from which you want to obtain the topology. If you are obtaining the topology for hosts for the first time, make sure **All Discovery Groups** is selected. You can use discovery groups to break up getting the topology or getting details. For example, instead of obtaining the topology for all of the elements, you could specify that the management server gets the topology for only the elements in Discovery Group 1, thus, saving you time. You add an element to a discovery group by modifying the properties used to discover the element. See “Modifying the Properties of a Discovered Address” on page 91.
4. If you see errors in the topology, look at the log messages, which can provide an indication of what went wrong. Look at Event Manager for additional information. Access Event Manager by clicking the Event Manager button in the left pane. To obtain troubleshooting information, see the “Troubleshooting Topology Issues” on page 813.
If the topology for an element in your network changes, select the element and click **Get Topology** in **Discovery > Topology** to update the information.

The software obtains just enough information about where the element is connected in the topology, for example a switch connected to a host.

Step C - Obtain Get Details

Obtain detailed information from the discovered applications as described in this section.

Keep in mind the following:

- Get Details takes some time. You might want to perform this process when the network and the managed elements are not busy.
- During Get Details the topology in System Explorer is recalculated. While the topology is being recalculated, the loading of the user interface may be slow. It may also take more time to log into the management server during a topology recalculation.
- To obtain a picture of device connectivity quickly, click the **Get Topology** button on the Topology tab.
- When you do Get Details that includes an AIX host, the system log displays three SCSI errors (2 FSCSI error and 1 FCS error) per IBM adapter port. You can ignore these errors.
- You can quarantine elements to exclude them from Get Details. See “Placing an Element in Quarantine” on page 105 for more information. Let us assume you want to discover all the elements in a discovery group, except for one. Perhaps the element you want to quarantine is being taken off the network for maintenance. You can use the quarantine feature to exclude one or more elements from discovery.
- If the management server unable to obtain information from an element during Get Details as a result of a CIM Extension hanging, the management server places the access point where the CIM Extension is located in quarantine. The management server then moves onto getting details for the next element in the Get Details table. These elements appear as missing until they are removed from quarantine. See “Removing an Element from Quarantine” on page 106 for information on how to remove an element from quarantine.

To obtain details:

1. Select **Discovery > Details**.
2. Select the discovery group from which you want to Get Details. If you are obtaining Get Details for hosts for the first time, make sure **All Discovery Groups** is selected.
You can use discovery groups to break up getting the topology or Get Details. For example, instead of Get Details for all of the elements, you could specify that the management server gets the element details for only the elements in Discovery Group 1, thus, saving you time. You add an element to a discovery group by modifying the properties used to discover the element. See “Modifying the Properties of a Discovered Address” on page 91.
3. Click the **Get Details** button.
During Get Details, the software changes its status light from green to red. You can view the progress of gathering details by clicking **Discovery > View Logs**.

When the software completes getting all elements details, it displays GETTING ALL DETAILS COMPLETED on the View Logs page.

Important: If the management server cannot communicate with an application, it labels the application as “Discovered”. The management server could find the application, but it could not obtain additional information about it.

4. Refer to the topic, for information about automating the gathering of Get Details. If you run into problems with discovery, see “Troubleshooting” on page 785.

Changing the Oracle TNS Listener Port

The software uses port 1521 by default to communicate with the TNS Listener service on the Oracle server. If your port is different or you use multiple ports, you can assign a new port number.

Important: The hosts should recognize the management server by name, as a reverse look-up is required by operating system security as well as the Oracle Transparent Name Substrate (TNS).

To change this port number or to add ports:

1. Select **Discovery > Setup**. Then, click the **Applications** tab.
2. To assign a new port, click the **Create** for the **Oracle Information** table.
3. Type the new port number and click **OK**.
4. If necessary, click the  button to remove the old port number.
5. Verify all elements have been discovered by clicking the **Start Discovery** button.

See “Troubleshooting Discovery and Get Details” on page 802 for more information.

Adding/Modifying Microsoft Exchange Domain Controller Access

To obtain information about your Microsoft Exchange servers, you must provide the user name and password for at least a primary domain controller, in addition to a DNS name, as described in the following steps.

Keep in mind the following:

- The hosts should recognize the management server by name, as a reverse look-up is required by operating system security as well as Microsoft Exchange.
- Make sure you provide just the user name and not the domain name when discovering Microsoft Exchange servers.

To provide information about the Microsoft Exchange servers:

1. Select **Discovery > Setup**. Then, click the **Applications** tab.
2. In the Microsoft Exchange Configuration section, click the **Edit** button.
3. Under the **Primary Domain Controller** section, perform the following steps:
 - a. In the **Host Name** field, type the fully qualified DNS name for the domain controller.
 - b. In the **User Name** field, type the user name for accessing the Microsoft Exchange server.
 - c. In the **Domain Password** field, type the corresponding password for accessing the Microsoft Exchange server.
 - d. In the **Verify Password** field, re-type the password for verification.
4. Under the **Backup Domain Controller** section, perform the following steps:
 - a. In the **Host Name** field, type the fully qualified DNS name for the domain controller.
 - b. In the **User Name** field, type the user name for accessing the Microsoft Exchange server.
 - c. In the **Domain Password** field, type the corresponding password for accessing the Microsoft Exchange server.
 - d. In the **Verify Password** field, retype the password for verification.
5. Click the **OK** button.
6. Verify all elements have been discovered by clicking the **Start Discovery** button.
7. Update the database with element changes. See “Step 1 - Discover Switches” on page 34.

See “Troubleshooting Discovery and Get Details” on page 802 for more information.

Changing the Password for the Managed Database Account

The management server connects to database applications through the use of the APPIQ_USER account, an unprivileged account with read-only privileges. You can change the password the management server uses to connect to database applications, such as Oracle and Sybase. When you change the password of APPIQ_USER, you must change the password of all database applications.

Keep in mind the following:

- Change the password in all database applications before you change the password through the user interface. The passwords must also match.
- You must enter a password in the **Password** and **Verify Password** fields.

To change the password:

1. Select **Discovery > Setup**. Then, click the **Applications** tab.
 2. Click the **Change Password** button.
 3. Verify you have already changed the password of the databases listed on this page.
 4. Type a new password in the **Password** field.
The management server requires the password to have the following characteristics:
 - ☐ a minimum of three characters
 - ☐ starts with a letter
 - ☐ contains only letters, numbers and underscores (_)
 - ☐ does not start or end with an underscore (_)
 5. Retype the password in the **Verify Password** field.
 6. Click **OK**.
 7. Verify that the management server can access the database applications by clicking the **Test** button for each database application.
-

Obtaining Disk Drive Statistics from Engenio Storage Systems

Important: Depending on your license, the ability to obtain disk drives statistics from Engenio storage systems may not be available. See the “List of Features” to determine if you have access to the additional statistics. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center**).

To obtain information about disk drive statistics from Engenio storage systems, you must install a CIM Extension on a host that can access the Engenio storage system. Ensure the proxy host has at least one LUN rendered by each controller of the array. Then, you must make the management server aware of that host, as described in the following steps:

1. Install the CIM Extension on a host that has access to the Engenio storage system.
2. Discover that host.
3. Select **Configuration > Performance > Data Collection**.
4. Verify the **Data Collection** tab is displayed.
5. Click the **Start** button corresponding with the disk drive statistics for an Engenio storage system.
6. Set the date and time. See “Setting the Date and Time for Performance Collectors” on page 259.
7. Type a repeat interval and then select a unit of measurement from the drop-down menu. The repeat interval determines how often the collectors gather the data.
8. Select a proxy host by clicking the **Browse** button.
9. Select a proxy host from the drop-down menu and then click **OK**.
The management server displays in the drop-down menu only hosts that are running a CIM Extension version 3.5 or later and have access to the corresponding Engenio storage system.

Note: You can always change the proxy host by returning to this page or by going to the **Properties** tab for an Engenio storage system. Double-click the Engenio storage system in System Explorer. Click the **Properties** tab. Then, click the **Browse** button on the Properties tab.

10. Saving the proxy host may take time. When you are asked if you want to continue, click **OK**.
11. Click **OK** again to set the time for starting the collector.
12. If you do not see any hosts displayed verify you have the latest CIM Extension installed and running on a host that can access the Engenio storage system.

Important: Depending on your license, role-based security may not be available. See the “List of Features” to determine if you have access to role-based security. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center**).

This chapter describes the following:

- About the Security for the Management Server on page 154
 - Managing User Accounts on page 161
 - Managing Roles on page 170
 - Managing Organizations on page 173
 - Changing the Password of System Accounts on page 181
 - Using Active Directory/LDAP for Authentication on page 182
-

About the Security for the Management Server

The management server offers security based on roles and organizations. Role-based security determines access to certain functionality depending on the user account assigned to a role. Organizations determine if you can modify an element type, such as hosts. The management server ships with the Everything organization, which lets you modify all element types.

See the following topics for more information:

- “About Roles” on page 154
- “About Organizations” on page 157
- “Planning Your Hierarchy” on page 160
- “Naming Organizations” on page 161

About Roles

The management server ships with several predefined roles that are listed in the following table. These roles determine which components of the software a user can access.

For example, users assigned to the Help Desk role have access to Application Explorer and Event Manager, but not to System Explorer, Provisioning, Policy Manager and Reporter. Likewise, users assigned to the domain administrator role have access to all of the features, as shown in the following table.

Table 5-1: Default Role Privileges

Feature	CIO (Chief Information Officer)	Domain Admini- strator	Storage Admini- strator	Server Admin- istrator	Applic- ation Admin- istrator	Help Desk
Application Explorer	X	X			X	X
System Explorer	X	X	X	X	X	
Event Manager		X	X	X	X	X
Provisioning		X	X			
Provisioning Administration		X	X			
Capacity Explorer	X	X	X	X	X	
Policy Manager		X	X			
Chargeback	X	X	X			
Business Tools	X	X	X			
Reporter	X	X	X	X	X	
Global Reporter	X	X	X			
File Server SRM		X		X		
Performance Explorer	X	X	X	X	X	
Access CLI		X	X			
Custom Commands		X	X			
System Configuration		X				

Keep in mind the following:

- Users with access to Global Reporter can view all the elements throughout the enterprise, including those on the server running Global Reporter. Grant access to Global Reporter only to those, who should be allowed to view all elements. Users, who had privileges to

Reporter in builds earlier than 3.5, are automatically given access to Global Reporter and thus they can see all elements. You may want to disable this functionality for some users.

- Only users belonging to the Domain Administrators role can add, modify, and delete users, roles, and organizations.
- The domain administrator can only edit active organizations.
- Domain administrators can change the user names and roles of other domain administrators, but they cannot modify their own user name and roles while logged into the management server. Domain administrators can also edit their full name, e-mail, phone, and other details, as well assign and un-assign any organization.
- If the System Configuration option is selected for a role, all users assigned to that role will have administration capabilities, as shown in the following list. If you do not want users belonging to that role to have those capabilities, do not assign the System Configuration option:
 - Schedule discovery
 - Find the CIM log level
 - Save log files, e-mail log files
 - Save the database, backup the database, and schedule a database backup
 - Configure Event Manager, File Server SRM and Performance Explorer
 - Configure reports and traps
 - Set up the management server to send e-mail

Roles also restrict access to element properties, element records, and Provisioning, as shown in the following table.

Table 5-2: Default Role Privileges with Elements

Role	Application	Host	Switch	Storage System	Tape Library	Others
CIO	View	View	View	View	View	View
Domain Administrator	Full Control	Full Control	Full Control	Full Control	Full Control	Full Control
Storage Administrator	View	View	Full Control	Full Control	Full Control	Full Control
Server Administrator	View	Full Control	View	View	View	View

Table 5-2: Default Role Privileges with Elements (Continued)

Role	Application	Host	Switch	Storage System	Tape Library	Others
Application Administrator	Full Control	View	View	View	View	View
Help Desk	View	View	View	View	View	View

By selecting one of the following options, users belonging to that role are restricted access:

- **Full Control** - Lets you view and modify the record for the element (Asset Management tab) and perform provisioning if applicable.
- **Element Control** - Lets you view and modify the record for the element (Asset Management tab). Provisioning cannot be performed.
- **View** - Lets you only view element properties.

For example, if a user belongs to a role that only lets you view the element properties on storage systems, that user would not be allowed to perform provisioning on storage systems because their role does not have the **Full Control** option selected for storage systems. That same role could also have the **Full Control** option selected for switches, allowing the user to perform provisioning for switches. Thus, the user would not be able to provision storage systems, but the user would be able to provision switches.

You can modify roles and/or create new ones. For example, you can modify the Help Desk role so that the users assigned to this role can also view Reporter and modify servers.

About Organizations

You can specify which elements users can access. For example, you can specify that some users have only access to certain switches and hosts. However, these users must already be assigned to roles that allow them to see switches and hosts.

Users only assigned to the organization can see just the elements that belong to the organization. If users are assigned to more than one organization, they see all elements that belong to the organizations to which they are assigned. For example, assume you created two organizations: one called OnlyHosts that allowed access to only hosts and another called OnlySwitches that allowed

access to only switches. If you assigned a user to OnlyHosts and OnlySwitches, they would have access to hosts and switches because those elements are listed in at least one of the organizations.

Organizations can also contain other organizations. An organization contained within another is called a child. The organization containing a child organization is called a parent. In the following figure, the BostonWebHosts organization contains two child organizations, BostonWebHost_Windows and BostonWebHost_Solaris. BostonWebHosts is a parent because it contains two organizations.




Figure 5-1: Parent-Child Hierarchy for Organizations

If a child contains organizations, it is also a parent. Let's assume you add two organizations called BostonWebMarketing and BostonWebProduction to BostonWebHost_Windows. BostonWebHost_Windows would become a parent because it now contains two organizations. It would also be a child because it is contained in BostonWebHosts.

Parent organizations allow access to all elements listed in their child organizations. For example, users assigned to the organization BostonWebHosts can access not only the elements in BostonWebHost_Windows, but also those in BostonWebHost_Solaris. This is because BostonWebHosts is a parent of the two child organizations.

The parent-child hierarchy for organizations saves you time when you add new elements. You need to add a new element only once. The change ripples through the hierarchy. For example, assume you add an element to BostonWebHost_Windows. Users not only assigned to BostonWebHost_Windows would see this addition, but also users assigned to any of the parent organizations containing BostonWebHost_Windows. For example, users assigned to BostonWebHosts would also see the addition because it contains BostonWebHost_Windows. Users, however, assigned to only BostonWebHost_Solaris would not see the addition.

A child organization can be in multiple parent organizations. As shown in the following figure BostonWebHosts and NYWebHosts are not only children of the WebHosts organization, but they are also children of the US East Coast organization. Assume you have a user that oversees all Web

hosts in the company, you could assign them to the WebHosts organization. Users managing hosts and storage systems on the east coast would be assigned to the US East Coast organization, which is a parent of BostonWebHosts, NYWebHosts, and StorageSystems organizations. Assume an element is added to NYWebHost_Solaris, users assigned to one or more of the following organizations would see the addition:

- NYWebHost_Solaris
- NYWebHosts
- WebHosts
- US East Coast




Figure 5-2: Children in Multiple Organizations

When you remove an element from an organization, users belonging to that organization or to one of its parents can no longer access that element if it is not a member of any other organization. For example, assume an element named *MyHost* was not only a member of *BostonWebHost_Solaris*, but also had mistakenly become a member of *BostonWebHost_Windows*. If you remove *MyHost* from *BostonWebHost_Solaris*, users belonging to *BostonWebHost_Solaris* can no longer access the element. Users belonging to the following organizations would still see the element because the element is still a member of *BostonWebHost_Windows*.

- BostonWebHosts
- WebHosts

- US East Coast

Keep in mind the following:

- You cannot edit the Everything organization.
- Users can only view all elements in the Discovery pages. In all other pages, only the members of the active organization are available.
- Discovery lists (**Discovery** tab) are not filtered. Users can see all elements in the discovery lists regardless of their affiliation with an organization.
- Event Manager displays events from all elements regardless of the user's organization.
- Domain administrators can change the user names and roles of other domain administrators, but they cannot modify their own user name and roles while logged into the management server. Domain administrators can also edit their full name, e-mail, phone, and other details, as well assign and un-assign any organization.
- Reports only display elements assigned to the user's organization, including child organizations. For example, if you attempt to view a Host Summary report and you do not have permission to access hosts through your organization, you are not given information about the hosts in the report. This is also true for e-mailing reports. Let's assume again you do not have permission to access hosts. The reports you e-mail will not contain information about hosts, including the host specific reports. If the users receiving your reports want to be able to view information about hosts, one of the following must happen:
 - The hosts in question must be added to your organization.
 - Someone else, who has the hosts in question already in their organization, must send the reports.

Planning Your Hierarchy

Before you begin creating organizations, plan your hierarchy. Do you want the hierarchy to be based on location, departments, hardware, software or tasks? Perhaps you want a combination of these options.

To help you with your task, create a table of users who manage elements on the network and the elements they must access to do their job. You might start seeing groups of users who oversee the same or similar elements. This table may help you in assigning users to the appropriate organizations.

Once you are done with planning your hierarchy, draw the hierarchy in an graphics illustration program, so you can keep track of which organizations are parents and children.

Create the child organizations first, then their parents. See the topic, “Adding an Organization” on page 173 for more information.

Naming Organizations

When you create an organization, give it a name that reflects its members. For example, you might want to use one or more of the following as a guideline:

- Type of elements that are members of the organization, such as switches, Sun Solaris hosts
- Location of the elements, such as San Jose
- Task, such as backup machines

You may find that it is easy to forget which containers are parents and children. When you name an organization, you might want to include a portion of the name of the dominant parent organization. For example, assume you have two types of Web hosts in Boston: Microsoft Windows and Sun Solaris. You might name the two children organizations `BostonWebHost_Windows` and `BostonWebHost_Solaris` and their parent, `BostonWebHosts`.

Managing User Accounts

This section discusses the following topics:

- Adding Users on page 161
 - Editing a User Account on page 163
 - Changing the Password for a User Account on page 164
 - Changing Your Password on page 165
 - Deleting Users on page 165
 - Modifying Your User Profile on page 165
 - Modifying Your User Preferences on page 166
 - Viewing the Properties of a Role on page 169
 - Viewing the Properties of an Organization on page 169
-

Adding Users

To access the management server, users must enter a user name and password. Only users belonging to the Domain Administrator role can add users.

Keep in mind the following:

- Only users belonging to the Domain Administrator role can add users.
- The user name and password should be alpha-numeric. They cannot exceed more than 256 characters. The user name cannot begin with a number.

To create an account:

1. Click **Security > Users**.
2. Click the **New User** button.
3. In the **Login Name** field, type a name for the user account, for example: jsmith
This name becomes the user name for the account.
4. (Optional) In the **Full Name** field, type a full name for the account.
This information is used to provide a correlation between an account name and a user.
The full name can contain spaces, but it cannot be longer than 512 characters.
5. Assign the user account to a pre-existing role by selecting a role from the **Role** drop-down menu. See “About the Security for the Management Server” on page 154 for more information about roles.
6. (Optional) In the **E-mail** field, type the user's e-mail address.
7. (Optional) In the **Phone** field, type the user's phone number.
8. (Optional) In the **Notes** field, provide additional information about the user.
9. (Optional) In the **Password** field, type a password for the user account.

Note: If you do not want to require the user to enter a password or the user will be using a password stored in Active Directory/LDAP, leave this field blank.

10. (Optional) In the **Verify Password** field, type the password you entered previously.
 11. Assign the user account to one or more organizations.
The organizations determine which elements the user can manage. To assign a user account to an organization, select the organizations from the table. See “About the Security for the Management Server” on page 154 for more information about roles and organizations, including the parent-child hierarchy.
-

12. Click **OK**.

Editing a User Account

Keep in mind the following:

- Only a user belonging to the Domain Administrator role is allowed to edit user accounts.
- The “admin” account acts differently than the other accounts. You cannot add or remove organizations from the “admin” account. You cannot remove the Everything organization from the “admin” account. New organizations are automatically added to the “admin” account when they are created.
- Although the domain administrator can see organization, user and role pages, the domain administrator can only edit active organizations.
- Domain administrators can change the user names and roles of other domain administrators, but they cannot modify their own user name and roles while logged into the management server. Domain administrators can also edit their full name, e-mail, phone, and other details, as well assign and un-assign any organization.
- This change takes effect immediately, even if the user is logged into the management server.
- You cannot change the password for a user account that has been authenticated against Active Directory/LDAP. To change the password for the user account, use Active Directory/LDAP. See “Step 3 - Add Users to the Management Server” on page 193.

If you want to change your password, follow the steps in “Changing Your Password” on page 165.

To modify a user account:

1. Click **Security > Users**.
2. Click the **Edit** button () corresponding to the user account you want to modify.
3. To change the account name, type a new name for the user account in the **Name** field, for example: jsmith
This name becomes the user name for the account.
4. To change the name assigned to the user account, type a new full name for the account in the **Full Name** field.
This information is used to provide a correlation between an account name and a user.
5. To change the role assign the user account, select a new role from the **Role** drop-down menu.

6. To change the e-mail address listed, type a new e-mail address in the **E-mail** field.
7. To change the phone number listed, type the user's new phone number in the **Phone** field.
8. Change or remove information from the **Notes** field if necessary.
9. To change the password:
 - a. Click the **Change Password** button.
 - b. Type a new password in the **Password** field.
 - c. Type the password again in the **Verify Password** field.
 - d. Click **OK**.
10. To change the organizations to which the user belongs, select or deselect the organizations from the table in the user interface.
The Everything organization is the default organization that lets users access all current and future elements.
11. Click **OK**.

Changing the Password for a User Account

To change the password for accessing the management server:

Keep in mind the following:

- Only a user belonging to the Domain Administrator role is allowed to change the password of another user.
- This change takes effect immediately, even if the user is logged into the management server.
- If a user account has been authenticated against Active Directory/LDAP, you cannot use the management server to change that user's password. You must use Active Directory/LDAP to change the password instead.

To modify a password:

1. Click **Security > Users**.
 2. Click **Users** from the drop-down menu.
 3. Click the **Edit** button () corresponding to the user account you want to modify.
 4. Click the **Change Password** button.
 5. Type a new password in the **New Password** field.
 6. Type the password again in the **Verify Password** field.
-

7. Click **OK**.

Changing Your Password

Note: You cannot use the management server to change your password if your user name has been authenticated against Active Directory/LDAP. See “Step 3 - Add Users to the Management Server” on page 193 for more information.

To change your password used for accessing the management server:

1. Click the name of your account in the upper-left corner.
2. On the **User Profile** tab, click the **Change Password** button.
3. Type a new password in the **New Password** field.
4. Type the password again in the **Verify Password** field.
5. Click **OK**.
6. Click the **Save Changes** button on the **User Profile** tab.
Your password used to access the management server is changed immediately.

Deleting Users

Keep in mind the following:

- You cannot delete the admin account.
- Only users belonging to the Domain Administrator role can delete users.

To delete a user account:

1. Click **Security > Users**.
2. Click the corresponding **Delete** button ()
The user account is deleted.

Modifying Your User Profile

While you are logged into the management server, you can change the following aspects of your user profile:

- Full Name
- E-mail address
- Phone number
- Password

However, you are not allowed to modify the following information:

- Login Name
- Role
- Organization affiliation

If you want this information modified, contact your domain administrator. Your domain administrator makes these changes.

To modify your user profile, do the following:

1. Click the name of your account in the upper-left corner.



Figure 5-3: Changing Your User Profile

2. On the **User Profile** tab, modify one or more of the following:
 - ☐ **Full Name**
 - ☐ **E-mail address**
 - ☐ **Phone number**
 - ☐ **Password** - To change the password, click the **Change Password** button. See “Changing Your Password” on page 165. This feature is not available if your user name has been authenticated against Active Directory or LDAP. Use Active Directory/ LDAP to change your password instead.
 3. When you are done with your modifications, click the **Save Changes** button.
-

Modifying Your User Preferences

Use the **User Preference** tab to modify your user preferences for System Explorer, Element Topology, and Event Manager. The **User Preference** tab controls what is displayed for your user account.

To access the **User Preferences** tab:

1. Click the name of your account in the upper-left corner.



Figure 5-4: Accessing the User Preferences Tab

2. Click the **User Preferences** tab.

Note: You can also access the **User Preference** tab by clicking the **Preferences** link in Event Manager.

System Explorer and Element Topology Preferences

To change the severity icons you view in System Explorer and in the element topology, select a severity level from the **Display Severity icons with this severity level or higher drop-down** menu. See the topic, “Viewing Event Status in the Topology” on page 341 for more information about event status.

If you want events refreshed within a time period, select the **Refresh events automatically** field. Then, enter in minutes how often you want the event information on the screen updated. If this option is set to every five minutes, the management server refreshes the severity icons displayed in System Explorer and the element topology every five minutes.

Event Manager Preferences

Preferences for Event Manager can also be accessed from the **Preferences** link in Event Manager.

Use the following table as a guideline for changing your user preferences for Event Manager.

Table 5-3: Changing User Preferences for Event Manager

If you want...	Do the following...
To be reminded whenever you change a filter	Select the option, Always remind me to apply filters when I change them .
Events refreshed automatically	Select the option, Refresh events automatically . Then, enter how often in minutes you want events refreshed.
Change the number of events displayed on each page	Select the number of events to appear on a page from the Number of Events combo box.
Change the severities to be included	Select a severity level you want displayed in Event Manager from the Severities to be Included drop-down menu. See the topic, “Selecting a Severity Level” on page 551. If you want to customize the filter for the severity levels, click the Custom button. See the topic, “Customizing the Severity Level Filter” on page 552 for more information.
Change the element types to be included	Select the element types you want to be included from the Element types to be included drop-down menu. Events from these elements are displayed in Event Manager. See the topic, “Selecting an Element Type for Events” on page 553. If you want to customize the filter for the element types, click the Custom button. See “Customizing the Element Type Filter for Events” on page 554.

Warnings for Slow Systems Operations

By default, the management server warns you when it encounters issues with handling large amounts of data from storage systems, such as long load times.

If you do not want to be warned, clear the **Warn about slow storage system operations** option on the **User Preferences** tab. See “Modifying Your User Preferences” on page 166 for information on how to access the **User Preferences** tab.

Viewing the Properties of a Role

You can quickly determine which components a user can access by viewing the properties of the user's role. Only users belonging to the Domain Administrator role can view the properties of roles.

To view the properties of a role:

1. Click **Security > Users**.
2. In the **Role** column, click the name of the role.

This page displays the following information:

- **Role Name** - The name of the role. This name appears in the users table (**Security > Users**)
- **Role Description** - A description of the role.
- **Access Level** - Determines how much access the user has to a type of element, such as hosts, storage systems, switches, and applications. See “About the Security for the Management Server” on page 154 for more information.
- **Access to the <product name>** - Determines which components in the management server the user can access, where <product name> is the name of your product.

To learn how to edit a role, see the topic, “Editing Roles” on page 171.

Viewing the Properties of an Organization

You can quickly determine which elements a user can access by viewing the properties of the user's organization. Only users belonging to the Domain Administrator can view the properties of an organization.

1. Click **Security > Users**.
2. In the **Organization** column, click the name of a organization.
3. To determine which elements are in a child organization, click the link of the child organization.
4. To learn more about an element, click the element's link.

This page displays the following information:

- **Name** - The name of the organization. This name appears in the users table (**Security > Users**)
- **Description** - A description of the organization
- **Organization Members** - Determines which elements the user can access. See “About the Security for the Management Server” on page 154 for more information.

To learn how to edit an organization, see the topic, “Editing Organizations” on page 175.

Managing Roles

This section discusses the following topics:

- Adding Roles on page 170
- Editing Roles on page 171
- Deleting Roles on page 172

Adding Roles

The management server ships with several roles. You can add roles to accommodate your organization. For example, you might want to add a role for quality assurance. See the topic, “About the Security for the Management Server” on page 154 for more information about roles and organizations.

Keep in mind the following:

- The **Role Name** and **Description** fields do not accept special characters, except spaces and the following characters: \$, -, ^, ., and _
- Only users belonging to the Domain Administrator role can add roles.

To add a role:

1. Select **Security > Roles**.
 2. Click the **New Role** button.
 3. In the **Role Name** field, type a name for the role. For example: Quality Assurance. The name can contain spaces, but it cannot be longer than 256 characters.
-

4. In the **Description** field, type a description for the role. For example: Role for those in quality assurance.
You cannot type more than 1024 characters in the **Description** field.
5. Select an access level for each element type:
 - ☐ **Full Control** - Lets you view and modify the record for the element (Asset Management tab) and perform provisioning.
 - ☐ **Element Control** - Lets you view and modify the record for the element (Asset Management tab).
 - ☐ **View** - Lets you view element properties.
For example, if a user belongs to a role that only lets you view the element properties on storage systems, that user would not be allowed to perform provisioning on storage systems because their role does not have the **Full Control** option selected for storage systems. That same role could also have the **Full Control** option selected for switches, allowing the user to perform provisioning for switches. Thus, the user would not be able to provision storage systems, but the user would be able to provision switches.
6. Select the features you want a user to be able to access. For example, if you want a user to have access to System Explorer, select System Explorer from the list.
See “Management Server Components” on page 4 for more information about these features.
7. Click **OK**.

Editing Roles

The software lets you modify the default roles and/or the roles you have created. See the topic, “About the Security for the Management Server” on page 154 for more information about roles and organizations.

Keep in mind the following:

- Only users belonging to the Domain Administrator role can modify roles.
- Domain administrators can change the user names and roles of other domain administrators, but they cannot modify their own user name and roles while logged into the management server. Domain administrators can also edit their full name, e-mail, phone, and other details, as well assign and un-assign any organization.
- After you click the **OK** button in the Edit Role window, any users assigned to the role you edited are logged out of the management server. Users see the changes when they log back into the management server.

- The **Role Name** field does not accept special characters, except spaces and the following characters: \$, -, ^, ., and _

To edit a role:

1. Select **Security > Roles**.
2. Click the **Edit** button ()
3. To edit the name of the role, change the name in the **Role Name** field.
The name can contain spaces, but it cannot be longer than 256 characters.
4. To edit the description of the role, change the description in the **Description** field.
You cannot type more than 1024 characters in the **Description** field.
5. To change the access level, change the options selected in the table.
 - ☐ **Full Control** - Lets you view and modify the record for the element (Asset Management tab) and perform provisioning.
 - ☐ **Element Control** - Lets you view and modify the record for the element (Asset Management tab).
 - ☐ **View** - Lets you view element properties.

For example, if a user belongs to a role that only lets you view the element properties on storage systems, that user would not be allowed to perform provisioning on storage systems because their role does not have the **Full Control** option selected for storage systems. That same role could also have the **Full Control** option selected for switches, allowing the user to perform provisioning for switches. Thus, the user would not be able to provision storage systems, but the user would be able to provision switches.
6. Select the features you want a user to be able to access. For example, if you want a user to have access to System Explorer, select System Explorer from the list.
See "Management Server Components" on page 4 for more information about these features.
7. Click **OK**.

Deleting Roles

Keep in mind the following:

- A role cannot be deleted if it contains a user.
- Only users belonging to the Domain Administrator role can delete roles.

To delete a role:

1. Select **Security > Roles**.
2. Click **Roles** from the drop-down menu.
3. Click the corresponding **Delete** button ().
The role is deleted.

Managing Organizations

This section discusses the following topics:

- Adding an Organization on page 173
- Viewing Organizations on page 175
- Editing Organizations on page 175
- Deleting an Organization on page 177
- Removing Members from an Organization on page 177
- Filtering Organizations on page 178

Adding an Organization

You can create new organizations to restrict access to certain elements. For example, assume you do not want the help desk to have access to elements belonging to a certain group. You could create an organization that does not allow access to those elements. Once you assign users to that organization, they would only be able to access the elements you specified.

See the topic, “About the Security for the Management Server” on page 154 for more information about roles and organizations.

Keep in mind the following:

- Create child organizations first, then their parents.
 - Event Manager displays events from all elements regardless of the user's organization.
 - Only users belonging to the Domain Administrator role can add organizations. Only active organizations can be edited.
 - All discovered elements are accessible in Business Tools, regardless of a user's restrictions. For example, assume your account belongs to an organization that has only
-

hosts as members. If you run the business tool Switch Risk Analysis, the management server still provides information about whether the switches are a risk in your environment.

To add an organization:

1. Select **Security > Organizations**.
2. Click the **New Organizations** button.
3. In the **Name** field, type a name for the organization.
The name of an organization has the following requirements:
 - ☐ Can contain spaces.
 - ☐ Can add digits to the beginning of an organization's name.
 - ☐ Cannot be longer than 256 characters.
 - ☐ Cannot contain the carot (^) symbol—currently the system allows the carot symbol to be entered, but the carot symbol should not be included in an organization's name.
4. In the **Description** field, type a description for the organization.
You cannot type more than 1024 characters in the **Description** field.
5. Click the **Add or Remove Members** button to determine which elements the user will see.
6. To add elements, expand the Element Types node in the tree. Select the element type that you would like to add. Then, in the right-hand pane, select the elements you would like to add by clicking the appropriate check boxes. Next, click the **Add** button and the selected elements are added to the Organization Members pane. If you would like to add storage volumes to the organization, see the topic “Adding Storage Volumes to an Organization” on page 174.
7. To add organizations, select the Organizations node. Then, in the right-hand pane, select the organizations you would like to add by clicking the appropriate check boxes. Next, click the **Add** button and the selected organizations are added to the Organization Members pane. The organizations in the Organization Members pane are listed as child organizations because they are now contained within the organization you are creating. See the topic, “About the Security for the Management Server” on page 154 for more information.
8. Once you are done adding the elements and organizations, click **OK**.

Adding Storage Volumes to an Organization

Only users belonging to the Domain Administrator role can add storage volumes to an organization.

To add storage volumes to an organization:

1. Click the **Add or Remove Members** button.
2. Expand the Element Types node in the tree and select the Storage Systems node.
3. In the right-hand pane, click the **Storage Volumes** tab and select a storage system from the **Showing Volumes for Storage System** drop-down menu.
4. If you want to filter the list of volumes for a storage system, click the **Show Volume Filter** link, select the appropriate filter criteria, and click the **Submit Query** button.
5. Select the storage volumes you want to add to the organization. Click the **+Ports** link in the **Ports** column if you want to see a list of the ports associated with a particular volume.
6. When you are finished selecting volumes, click the **Add** button located at the top of the pane. The selected volumes are added to the Organization Members pane. Click the **OK** button.

Viewing Organizations

The Setup Organizations page lists the organizations and their descriptions, in addition to the number of top-level elements, users and child organizations assigned to each organization.

Only users belonging to the Domain Administrator role can view organizations.

The No. of Top Level Elements column provides the total number of elements assigned to an organization, not including those within the child organization. An organization containing only child organizations displays 0 under the No. of Top Level Elements column; however, users assigned to that organization would have access to the elements assigned to its child organizations.

Assume an organization contains only two child organizations. As a result, 0 would be displayed under the No. of Top Level Elements column. Users assigned to that organization can access the elements assigned to the two child organizations.

Access the Setup Organizations page by clicking **Security > Organizations**.

You can access information about child organizations by clicking their link under the Child Organization column.

Editing Organizations

See the topic, “About the Security for the Management Server” on page 154 for more information about roles and organizations.

When elements are removed from an organization, users belonging only to that organization are no longer able to access the removed elements.

Keep in mind the following:

- Depending on your license, role-based security may not be available. See the “List of Features” to determine if you have access to role-based security. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center**).
- Only users belonging to the Domain Administrator role can edit organizations. Only active organizations can be edited.
- The **Name** and **Description** fields in the Edit Organization window do not accept special characters, except spaces and the following characters: \$, -, ., and _
- The organization name can contain spaces, but it cannot be longer than 256 characters.
- You cannot edit the Everything organization.

To edit an organization:

1. Select **Security > Organizations**.
 2. Click the  button.
 3. To change the name of the organization, type a new name in the **Name** field. The name of an organization has the following requirements:
 - ☐ Can contain spaces.
 - ☐ Can add digits to the beginning of an organization’s name.
 - ☐ Cannot be longer than 256 characters.
 - ☐ Cannot contain the carot (^) symbol—currently the system allows the carot symbol to be entered, but the carot symbol should not be included in an organization’s name.
 4. To change the description of the organization, type a new description in the **Description** field.
You cannot type more than 1024 characters in the **Description** field.
 5. Click the **Add or Remove Members** button.
 6. Add and remove elements as described in the topics, “Adding an Organization” on page 173 and “Removing Members from an Organization” on page 177.
-

7. Once you are done adding or removing elements, click **OK** in the Add Organization page.
8. In the Edit Organization page, click **OK**.

Deleting an Organization

When an organization is removed, users assigned only to that organization are no longer able to access the elements in the removed organization. For example, assume you belong to two organizations, onlyHosts and onlySwitchesandHosts. The organization onlyHosts contains only hosts, and the organization onlySwitchesandHosts contains only switches and hosts. If you delete the onlySwitchesandHosts organization, you will still have access to hosts because you still belong to the onlyHosts organization.

Keep in mind the following:

- You cannot remove the Everything organization, which is the default organization.
- You cannot delete an organization that contains a user, who belongs to no other organizations. For example, assume you create an organization named Org1 that contains two users: User1 and User2. User1 belongs to two other organizations, while User2 only belongs to the organization you just created. You will not be able to delete Org1 because the organization contains User2, who only belongs to the organization you are trying to delete.
- Only users belonging to the Domain Administrator role can delete organizations.

Depending on your license, role-based security may not be available. See the “List of Features” to determine if you have access to role-based security. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center**).

To delete an organization:

1. Click **Security > Organizations**.
2. Click the  button corresponding to the organization you want to remove. The software removes the organization.

Removing Members from an Organization

When you remove an element from an organization, users belonging to that organization or to one of its parents can no longer access that element if it is not a member of any other organization. For example, assume an element named MyHost was not only a member of BostonWebHost_Solaris,

but also had mistakenly become a member of BostonWebHost_Windows. If you remove MyHost from BostonWebHost_Solaris, users belonging to BostonWebHost_Solaris can no longer access the element. Users belonging to the BostonWebHost_Windows organization or to its parent would still see the element.

Important: Depending on your license, role-based security may not be available. See the “List of Features” to determine if you have access to role-based security. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center**).

Use one of the following methods to remove an element from an organization:

- In the Edit Organization window, click the  button corresponding to the element or child organization you want to remove from the organization.
- In the Add or Remove Organization Members window, select the element or child organization you want to remove by clicking the appropriate check box. Next, click the **Remove** button.
- Only users belonging to the Domain Administrator role can remove members from an organization. Only active organizations can be edited.

Filtering Organizations

The management server provides a filtering feature that lets you designate which organizations are active in your view. For example, assume you belong to an organization name Hosts and this organization contains two organizations: “Windows Hosts” and “Solaris Hosts.” If you want to view elements only in “Windows Hosts” and not in “Solaris Hosts” organizations, you could use the filtering feature to activate only the “Windows Hosts” organization.

Keep in mind the following:

- Users assigned to the “admin” account cannot filter organizations because the “admin” account belongs to the Everything organization by default. As a result, these users do not have access to filtering feature for organizations.
 - If you do not want to view an element, deselect all child organizations containing the element. You must also deselect all parent organizations containing the child organization that has that element. For example, assume you do not want to view all Solaris hosts and all Solaris hosts are in the “Solaris Hosts” organization. The “Solaris Hosts” organization is
-

contained in the Hosts organization. You must deselect the “Solaris Hosts” organization and the Hosts organization if you do not want to see Solaris hosts.

- The filter for organizations does not appear in Event Manager. Event Manager displays events from all elements regardless of the user’s organization.
- If you do not select any organizations for filtering, you do not see any elements in the topology.

To filter organizations:

1. Click the  button at the top of the screen, or click the link listing the organizations you can view.




Figure 5-5: Clicking the Organization Link

2. Deselect the organizations containing the elements you do not want to obtain information about. Assume you want to view only the elements in the “Windows Hosts” organization, you would select only “Windows Hosts.” Let’s assume you have a parent organization named “Hosts” that contains “Solaris Hosts” and “Windows Hosts.” You would need to deselect “Solaris Hosts” and “Hosts.” “Hosts” would need to be deselected because it contains organizations other than “Windows Hosts.”
Links are displays for the organizations if you belong to the Domain Administrator role. To learn more about the contents of an organization, click its link.




Figure 5-6: Filtering Organizations

3. Click **OK**.
You can now only obtain information about elements in the active organizations. The active organizations are listed in the link next to the filter button, as shown in the following figure.



Figure 5-7: Active Organization

Changing the Password of System Accounts

The management server uses the following accounts to access and manage the database for the management server. You should change the passwords to these accounts to prevent unauthorized access.

- **SYS** - Used for the management server database creation and upgrade. Default password: `change_on_install`
- **SYSTEM** - Used for management server database creation and upgrade, in addition to database import, export and re-initialization. Default password: `manager`
- **RMAN_USER** - Used for RMAN backup and restore. This user has sys privilege. Default password: `backup`
- **DB_SYSTEM_USER** - Used for all the database activity, including establishing a connection to the management server database. Default password: `password`

You must change the passwords of the SYS, SYSTEM, RMAN_USER, and DB_SYSTEM_USER accounts by using the `dbadmin.bat` tool, so the management server is aware of the changes. Do not change the password for one of these accounts by using Oracle. Make sure you keep the new passwords in a safe location, as it is your responsibility to remember the Oracle passwords.

The management server requires the password to have the following characteristics:

- a minimum of three characters
- starts with a letter
- contains only letters, numbers and underscores (`_`)
- does not start or end with an underscore (`_`)

To change the password of a system account:

1. Stop the AppStorManager service.
2. To access the database utility on Solaris, do the following on the management server:

- a. Set the display if you are accessing the dbAdmin tool remotely.
- b. The dbAdmin tool uses Perl. To set Perl in your path, enter the following command at the command prompt:

```
eval `/opt/productname/install/usersvars.sh`
```

where `/opt/productname` is the directory containing the software. It is defined by `$APPIQ_DIST`.

Important: You must include the back quotes around the full path to `usersvars.sh` in the command.

- c. Go to the `$APPIQ_DIST/Tools/dbAdmin` directory and then enter the following at the command prompt:

```
perl dbAdmin.pl
```

The full path to Perl is the following:

```
$APPIQ_DIST/JBossandJetty/server/appiq/remoteScripts/perl/bin/perl
```

3. To access the database utility on Windows, go to the `%MGR_DIST%\Tools\dbAdmin` directory on the management server and double-click `dbAdmin.bat`, where `[Install_Dir]` is the directory into which you installed the management server.
4. Click **Change Passwords** in the left pane.
5. Select an account name from the **User Name** combo box.
6. Type the current password in the **Old Password** field.
7. Type the new password in the **New Password** field.
8. Retype the password in the **Confirm Password** field.
9. Click **Change**.

The Database Admin Utility changes the password for the specified account.

Using Active Directory/LDAP for Authentication

The management server supports external authentication through Active Directory (AD) and Lightweight Directory Access Protocol (LDAP) directory services. When you configure the management server to use external authentication, user credentials are no longer stored in the management server database. This configuration centralizes all security related requirements to the

enterprise AD/LDAP infrastructure, such as password expiration, resets, and complexity requirements.

When a user attempts to log into the management server, the management server authenticates the user name and password against AD/LDAP for credential verification. If AD/LDAP verifies this user has the correct credentials, the management server allows this user access to the application.

Keep in mind the following:

- The `login-handler.xml` file contains configuration information for Active Directory and LDAP. It is important you enable either Active Directory or LDAP. You cannot enable both.
- If you want to go back and forth between internal and external (AD/LDAP) authentication, rename the `login-handler.xml` file before you modify it. This way you can easily switch back to internal authentication by changing the file name back to `login-handler.xml`.

To use AD/LDAP to authenticate your users, complete the following sections:

- “Step 1 - Configure the Management Server to Use Active Directory or LDAP” on page 183
- “Step 2 - Restart the AppStorManager Service and Login as the Designated Admin Account” on page 192
- “Step 3 - Add Users to the Management Server” on page 193
- “Step 4 - Provide Login Information to Your Users” on page 193

Step 1 - Configure the Management Server to Use Active Directory or LDAP

You must modify the `login-handler.xml` file if you want to use Active Directory/LDAP. How you modify the `login-handler.xml` file depends on whether you plan to use LDAP or Active Directory.

See the one of the following sections depending on whether you want to use Active Directory or LDAP:

- “Active Directory” on page 184
- “LDAP” on page 188

Active Directory

Active Directory allows by default connections with domain\username, instead of the distinguished name (DN) used by a generic LDAP server. However, you can use the generic LDAP server setup to authenticate with Active Directory, as described in “LDAP” on page 188.

To specify the management server to use Active Directory, do the following:

1. Before switching to Active Directory (AD) authentication mode, the management server needs to be configured with a designated Active Directory user and other AD specific credentials. At startup, the designated Active Directory user is mapped to the built-in “admin” user and overrides it with the Active Directory user information.

Important: Make sure the administrator account has already been created in Active Directory before you add it to the `login-handler.xml` file.

- a. On the management server look in one of the following locations:

- **Windows:** %MGR_DIST\Data\Configuration
- **Solaris:** \$MGR_DIST\Data\Configuration

Note: If you want to go back and forth between internal and external (AD/LDAP) authentication, rename the `login-handler.xml` file before you modify it. This way you can easily switch back to internal authentication by changing the file name back to `login-handler.xml`.

- b. In the `login-handler.xml` file, change the value of the `<AdminAccountName>` tag to the name of a user account in Active Directory, as shown in the following example:

```
<AdminAccountName>domain\PrimaryUser</AdminAccountName>
```

where `PrimaryUser` is the name of the user account that is designated as the primary user in Active Directory.

Keep in mind the following:

- ☐ For security reasons, it is recommended that the designated user not be the AD Domain Administrator
 - ☐ If you are using Active Directory, prefix the user name with the domain name, for example: `domain\PrimaryUser`
-

2. In the `login-handler.xml` file, comment out the section that contains `com.appiq.security.server.BasicLoginhandler`, which enables internal authentication mode. Only one login handler is allowed at a time.

```
<!--LoginHandlerClass>com.appiq.security.server.BasicLoginHandler</LoginHandlerClass-->
```
3. Comment out the `<LoginHandlerType>Default</LoginHandlerType>` tag as follows:

```
<!--LoginHandlerType>Default</LoginHandlerType-->
```
4. Uncomment the line containing the class name and login handler type so that it appears as follows:

```
<LoginHandlerClass>com.appiq.security.server.ActiveDirectoryLoginHandler</LoginHandlerClass>
<LoginHandlerType>ActiveDirectory</LoginHandlerType>
```
5. Replace `directory.hp.com` with the IP address or the fully qualified DNS name of your primary Domain Controller server in the `login-handler.xml` file, as shown in the following example:

```
<PrimaryServer port="389">192.168.10.1</PrimaryServer>
```

where

 - 192.168.10.1 is the IP address of the primary Domain Controller server running Active Directory.
 - 389 is the port on which Active Directory is running on the server.
6. Replace `directory2.hp.com` with the IP address or the fully qualified DNS name of your secondary Domain Controller server, if available.

```
<SecondaryServer>192.168.10.2</SecondaryServer>
```

where 192.168.10.2 is the IP address of the secondary Domain Controller server running Active Directory.
7. If you want the password to be saved in the management server database, change the value of the `<ShadowPassword>` tags to true, as shown in the following example:

```
<ShadowPassword>true</ShadowPassword>
```

Saving the passwords in the management server database lets a user still log into the management server if the management server is changed back to local mode. This, however, is not recommended as it defeats the purpose of externalizing a user's credentials.

The `login-handler.xml` file contains two sets of `<ShadowPassword>` tags: one for Active Directory and one for LDAP. Make sure you change the value of the `<ShadowPassword>` tags that are children of the `<ActiveDirectory>` tag.
8. If you want the user name to be case sensitive, change the value of the `<CaseSensitiveUserName>` tag to true, as shown in the following example:

```
<CaseSensitiveUserName>true</CaseSensitiveUserName>
```

If you change the value of `<CaseSensitiveUserName>` to true, the management server becomes case-sensitive to user names. The management server sees MyUserName and myusername as different users.

Important: AD servers are not case sensitive for user names so changing this tag to “true” for AD authentication is not recommended.

The `login-handler.xml` file contains two sets of `<CaseSensitiveUserName>` tags: one for Active Directory and one for LDAP. Make sure you change the value of the `<CaseSensitiveUserName>` tags that are children of the `<ActiveDirectory>` tag.

9. Provide the Active Directory search base in which you want the management server to look up AD/LDAP user attributes. Allow no spaces between commas and put in all components of fully qualified domain name, for example, `hds.usa.com` would be `DC=hds,DC=usa,DC=com`.

The search base is used to specify the starting point for the search. It points to a distinguished name of an entry in the directory hierarchy.

```
<SearchBase> dc=MyCompanyName,dc=COM</SearchBase>
```

10. Save the `login-handler.xml` file with your changes. The following is an example of a modified `login-handler.xml` file for use with AD server authentication. Underlined text is information that was modified:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<LoginHandler>
  <AdminAccountName>domain\primaryuser</AdminAccountName>
  <!-- for the default, using database for authentication -->
  <!--LoginHandlerClass>com.appiq.security.server.BasicLoginHandler</
LoginHandlerClass-->
  <!--LoginHandlerType>Default</LoginHandlerType-->
  <!-- uncomment the following to enable Active Directory login-->
  <LoginHandlerClass>com.appiq.security.server.ActiveDirectoryLoginHandler</
LoginHandlerClass>
  <LoginHandlerType>ActiveDirectory</LoginHandlerType>

  <ActiveDirectory>
    <PrimaryServer port="389">IP address of Primary Domain Controller</
    PrimaryServer>
    <SecondaryServer>IP Address of Secondary Domain Controller</
    SecondaryServer>
    <ssl>>false</ssl>
    <ShadowPassword>>false</ShadowPassword>
```

```

<CaseSensitiveUserName>>false</CaseSensitiveUserName>
<!-- provide SearchBase if full name and email attribute are to be
synchronized
between ActiveDirectory and the database.-->
<SearchBase>DC=domain extension1,DC=domain extension2,DC=COM</SearchBase>
<FullNameAttribute>displayName</FullNameAttribute>
<EmailAttribute>mail</EmailAttribute>
</ActiveDirectory>
<!-- uncomment the following for generic LDAP login
<LoginHandlerClass>com.appiq.security.server.LdapLoginHandler
</LoginHandlerClass>
<LoginHandlerType>LDAP</LoginHandlerType>
-->
<LDAP>
<!-- same as java.naming.provider.url ldap://ldap.companyname.com:389
-->
<Server port="389">IP address of LDAP server</Server>
<!-- LDAP env can be added, an example is shown below...
<LDAPEnv
name="java.naming.factory.initial">com.sun.jndi.ldap.LdapCtxFactory</
LDAPEnv>
-->
<ssl>>false</ssl>
<ShadowPassword>>false</ShadowPassword>
<CaseSensitiveUserName>>false</CaseSensitiveUserName>
<!-- multiple DN entries are allowed, they will be tried one at a time -->
<DN>CN=$NAME$,OU=Engineering,DC=HP,OU=US,DC=COM</DN>
<!-- provide FullNameAttribute and EmailAttribute if full name and email
attribute
are to be synchronized between LDAP and the database -->
<FullNameAttribute>displayName</FullNameAttribute>
<EmailAttribute>mail</EmailAttribute>
</LDAP>
</LoginHandler>

```

When you are done with your changes, the login-handler.xml file, may resemble the following:

```

<LoginHandler>
  <AdminAccountName>domain\primaryuser</AdminAccountName>
  <LoginHandlerClass>
    com.appiq.security.server.ActiveDirectoryLoginHandler
  </LoginHandlerClass>
  <LoginHandlerType>ActiveDirectory</LoginHandlerType>
  <ActiveDirectory>
    <PrimaryServer>IP address of primary domain controller</PrimaryServer>
    <SecondaryServer>IP address of secondary domain controller</
SecondaryServer>
  </ActiveDirectory>
</LoginHandler>

```

```
<ssl>false</ssl>
<ShadowPassword>false</ShadowPassword>
<CaseSensitiveUserName>false</CaseSensitiveUserName>
<SearchBase>DC=MyCompanyName,DC=COM</SearchBase>
  <FullNameAttribute>displayName</FullNameAttribute>
  <EmailAttribute>mail</EmailAttribute>
</ActiveDirectory>
</LoginHandler>
```

LDAP

The LDAP server requires a distinguished name (DN) and credentials. The DN can be configured, allowing name substitution and support for multiple DN configurations.

To configure the management server to use LDAP:

1. Before switching to LDAP authentication mode, the management server needs to be configured with a designated LDAP user through the `<AdminAccountName>` tag. At startup, the designated LDAP user is mapped to the built-in “admin” user and overrides it with the LDAP user information.

Important: Make sure the administrator account has already been created in LDAP before you add it to the `login-handler.xml` file.

- a. On the management server look in one of the following locations:
 - **Windows:** %MGR_DIST\Data\Configuration
 - **Solaris:** \$MGR_DIST\Data\Configuration

Note: If you want to go back and forth between internal and external (AD/LDAP) authentication, rename the `login-handler.xml` file before you modify it. This way you can easily switch back to internal authentication by changing the file name back to `login-handler.xml`.

- b. In the `login-handler.xml` file, change the value of the `<AdminAccountName>` tag to the name of a user account in LDAP, as shown in the following example:
`<AdminAccountName>Administrator</AdminAccountName>`
-

where Administrator is the name of a user account in LDAP.

2. In the `login-handler.xml` file, comment out the section that contains `com.appiq.security.server.BasicLoginhandler`, which enables internal authentication mode. Only one login handler is allowed at a time.
`<!--LoginHandlerClass>com.appiq.security.server.BasicLoginHandler</LoginHandlerClass-->`
3. Comment out the `<LoginHandlerType>Default</LoginHandlerType>` tag as follows:
`<!--LoginHandlerType>Default</LoginHandlerType-->`
4. Uncomment the line containing the class name and login handler type so that it appears as follows:

```
<LoginHandlerClass>com.appiq.security.server.LdapLoginHandler</LoginHandlerClass>
<LoginHandlerType>LDAP</LoginHandlerType>
```

5. Replace `directory.hp.com` with the IP address or the fully qualified name of your LDAP server in the `login-handler.xml` file, as shown in the following example:

```
<Server port="389">192.168.10.1</Server>
```

where

- 192.168.10.1 is the IP address of the server running LDAP.
 - 389 is the port on which LDAP is running on the server.
6. If you want the password to be saved in the management server database, change the value of the `<ShadowPassword>` tags to `true`, as shown in the following example:
`<ShadowPassword>true</ShadowPassword>`

Saving the passwords in the management server database lets a user still log into the management server if the management server is changed back to local mode. This, however, is not recommended as it defeats the purpose of externalizing a user's credentials.

The `login-handler.xml` file contains two sets of `<ShadowPassword>` tags: one for Active Directory and one for LDAP. Make sure you change the value of the `<ShadowPassword>` tags that are children of the `<LDAP>` tags.

7. If you want the user name to be case sensitive, change the value of the `<CaseSensitiveUserName>` tag to `true`, as shown in the following example:
`<CaseSensitiveUserName>true</CaseSensitiveUserName>`

If you change the value of `<CaseSensitiveUserName>` to `true`, the management server becomes case-sensitive to user names. The management server sees `MyUserName` and `myusername` as different users.

The `login-handler.xml` file contains two sets of `<CaseSensitiveUserName>` tags: one for Active Directory and one for LDAP. Make sure you change the value of the `<CaseSensitiveUserName>` tags that are children of the `<LDAP>` tags.

8. Provide the LDAP search base in which you want the management server to look up AD/LDAP user attributes. Allow no spaces between commas and put in all components of fully qualified domain name, for example, `hds.usa.com` would be `DC=hds,DC=usa,DC=com`. The search base is used to specify the starting point for the search. It points to a distinguished name of an entry in the directory hierarchy.

```
<SearchBase>CN=$NAME$,dc=MyCompanyName,dc=COM</SearchBase>
```

or:

```
<SearchBase>CN=$NAME$,OU=NetworkAdministration,
dc=MyCompanyName,ou=US,dc=COM</SearchBase>
```

The management server searches only those users in the company who are part of the NetworkAdministration organization (OU=NetworkAdministration) and in the United States (ou=US).

Important: Different LDAP implementations may be using different keynames for CN. The appropriate key should be mentioned in `login-handler.xml`. Refer to the documentation for your LDAP server to determine how to obtain the appropriate keyname. Your keyname may start with uid instead of CN, for example, `: uid=$NAME$,ou=<Optional org unit if applicable>,dc=windows,dc=hp,dc=com`

9. Save the `login-handler.xml` file.

The following is an example of a modified `login-handler.xml` file for use with an LDAP server. Underlined text is information that was modified:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<LoginHandler>
  <AdminAccountName>domain\admin</AdminAccountName>
  <!-- for the default, using database for authentication -->
  <!--LoginHandlerClass>com.appiq.security.server.BasicLoginHandler</
LoginHandlerClass-->
  <!--LoginHandlerType>Default</LoginHandlerType-->
  <!-- uncomment the following to enable Active Directory login>
  <LoginHandlerClass>com.appiq.security.server.ActiveDirectoryLoginHandler</
LoginHandlerClass>
  <LoginHandlerType>ActiveDirectory</LoginHandlerType-->

  <ActiveDirectory>
    <PrimaryServer port="389">IP address of Primary Domain Controller</
PrimaryServer>
    <SecondaryServer>IP Address of Secondary Domain Controller</
SecondaryServer>
```

```

<ssl>>false</ssl>
<ShadowPassword>>false</ShadowPassword>
<CaseSensitiveUserName>>false</CaseSensitiveUserName>
<!-- provide SearchBase if full name and email attribute are to be
synchronized
between ActiveDirectory and the database.-->
<SearchBase>DC=domain extension1,DC=domain extension2,DC=COM</SearchBase>
<FullNameAttribute>displayName</FullNameAttribute>
<EmailAttribute>mail</EmailAttribute>
</ActiveDirectory>
<!-- uncomment the following for generic LDAP login-->
<LoginHandlerClass>com.appiq.security.server.LdapLoginHandler</
LoginHandlerClass>
<LoginHandlerType>LDAP</LoginHandlerType>
<LDAP>
<!-- same as java.naming.provider.url ldap://ldap.companyname.com:389 -
->
<Server port="389">IP address or DNS name of LDAP server</Server>
<!-- LDAP env can be added, an example is shown below...
<LDAPEnv
name="java.naming.factory.initial">com.sun.jndi.ldap.LdapCtxFactory</
LDAPEnv>
-->
<ssl>>false</ssl>
<ShadowPassword>>false</ShadowPassword>
<CaseSensitiveUserName>>false</CaseSensitiveUserName>
<!-- multiple DN entries are allowed, they will be tried one at a time -->
<DN>CN=$NAME$,OU=Engineering,DC=mycompanyname,OU=US,DC=COM</DN>
<!-- provide FullNameAttribute and EmailAttribute if full name and email
attribute
are to be synchronized between LDAP and the database -->
<FullNameAttribute>displayName</FullNameAttribute>
<EmailAttribute>mail</EmailAttribute>
</LDAP>
</LoginHandler>

```

When you are done with your changes, the login-handler.xml file, may resemble the following:

```

<LoginHandler>
  <AdminAccountName>Administrator</AdminAccountName>
  <LoginHandlerClass>
    com.appiq.security.server.LdapLoginHandler
  </LoginHandlerClass>
  <LoginHandlerType>LDAP</LoginHandlerType>
  <LDAP>
    <Server port="389">IP address of LDAP server</Server>
  </LDAP>
</LoginHandler>

```

```
<ssl>false</ssl>
<ShadowPassword>false</ShadowPassword>
<CaseSensitiveUserName>false</CaseSensitiveUserName>
<DN>CN=$NAME$,OU=Engineering,DC=HP,OU=US,DC=COM</DN>
<FullNameAttribute>displayName</FullNameAttribute>
<EmailAttribute>mail</EmailAttribute>
</LDAP>
</LoginHandler>
```

Step 2 - Restart the AppStorManager Service and Login as the Designated Admin Account

In this section, you will restart the AppStorManager service and login as the designated Admin account.

1. After you modify the `login-handler.xml` file, you must restart the AppStorManager service, which is the service for the management server for your changes to take effect.

Important: The service must be running for users to access the management server.

On Microsoft Windows:

- a. Go to the **Services** window, usually accessible from the Control Panel.
- b. Right-click **AppStorManager**.
- c. Select **Stop** from the drop-down menu.
- d. To start the management server, right-click **AppStorManager** and select **Start** from the drop-down menu.

On Sun Solaris:

- a. Open a command prompt window.
 - b. Enter the following at the command prompt to stop the management server:
`/etc/init.d/appstormanager stop`
 - c. To start the management server, enter the following at the command prompt:
`/etc/init.d/appstormanager start`
2. Login as the designated administrator account you specified in “Step 1 - Configure the Management Server to Use Active Directory or LDAP” on page 183.
For example, the user name would be the following:

- ❑ **Active Directory** - domain\PrimaryUser

- ❑ **LDAP** - PrimaryUser

where `PrimaryUser` is the name of the user account in LDAP or is the designated primary user in Active Directory.

The password would be the following: [NTdomainpassword].

Step 3 - Add Users to the Management Server

You must now add your Active Directory/LDAP users to the management server. This step is required to prevent accidental access to the management server from other AD/LDAP users. Until the user is authenticated against Active Directory/LDAP, the management server views the user as an internal user, whose password can be changed within the management server.

Once a user is authenticated against Active Directory/LDAP, the user is tagged as an external user and the user's password must be managed through Active Directory/LDAP.

1. Log onto the management server by using the designated Admin account specified in "Step 1 - Configure the Management Server to Use Active Directory or LDAP" on page 183.
2. Create your users as described in "Adding Users" on page 161.
Keep in mind the following:
 - ❑ **(Active Directory)** Prefix the user name with the domain name, for example: domain\newuser.
 - ❑ The user names you create by using the management server must match the user names in Active Directory/LDAP.
 - ❑ It is not necessary to create a password since passwords used for login are those already configured on either the AD or LDAP server.

Step 4 - Provide Login Information to Your Users

Users are now able to log into the management server with the user name and password specified in Active Directory/LDAP.

Important: Tell your users not to give the password they use to access the management server to anyone. Since user credentials are now stored in Active Directory/LDAP, the password used to access the management server may also be used to access other accounts. In some instances, it may be their network user name and password.

To learn more about licensing, see the following topics:

- About the License on page 195
- Importing a License File on page 200
- Viewing Cumulative Licenses on page 200
- Viewing a Specific License on page 201
- Deleting a License on page 202
- Import License File Window Appears Too Small on page 202

About the License

The management server restricts the number elements it manages through its license. It is important you keep your license up to date with the requirements of your network. The management server has several different types of license restrictions, as shown in Table 6-1, “License Requirements,” on page 196.

Table 6-1: License Requirements

Type of Restriction	Description	Unit of Measurement
MAPs	The management software restricts the number of hardware elements it manages through the use of managed access points (MAPs) for hardware. A MAP is the sum of all storage access ports of all hardware elements that the management server manages. See Table 6-2, "Determining Managed Access Points," on page 198 for more information.	Number of MAPs
Raw NetApp Capacity	The Raw NetApp Capacity is the total disk capacity (unformatted capacity) of all discovered NetApp filers.	Terabytes (TB)
Managed Exchange Instances	The management server determines licensing for Microsoft Exchange instances by counting the number of instance of Microsoft Exchange it manages.	Number of instances of Microsoft Exchange the software manages

Table 6-1: License Requirements (Continued)

Type of Restriction	Description	Unit of Measurement
Managed Database Instances	The total number of instances of the following database the software manages: ■ Microsoft SQL Server ■ Oracle ■ Sybase Adaptive Server Enterprise This total is broken down by each type of database in the table.	Number of managed databases
For File Server SRM	The management server determines licensing for File Server SRM through terabytes (TB). When you purchased File Server SRM, you were given a number of TB you were allowed by the management server to monitor. The management server detects the number of TB that are being monitored on file servers and verifies that number is at or below the purchased amount.	Terabytes (TB)

Important: The management server updates the Current Usage Summary table every 24 hours after the first update. Elements the management server has discovered before the update are not reflected in the Current Usage Summary table. The time for the update is determined when the management server was first started. For example, the first update of the Current Usage Summary table occurs six hours after the management server was first started. The following updates occur every 24 hours. If the management server was started for the first time at noon, the first update of the Current Usage Summary table would occur at 6 p.m. All following updates would always occur at 6 p.m.

MAPs are determined as described in the following table.

Table 6-2: Determining Managed Access Points

Element	Managed Access Point
Hosts	The managed access points (MAPs) are the number of Fibre Channel ports with a minimum of one MAP. If a host has no Fibre Channel ports, the software assumes one MAP. The software does count direct attached storage, provided it is supported by the management server.
Switches	All ports on a switch are counted as MAPs.
Storage systems	The sum of all front facing ports. Storage systems with FA ports the software does not support, such as mainframe attached FICON, are still counted as MAPs. However, the management server does not count MAPs from storage systems it does not support. See the release notes for information about supported storage systems.

Example 1:

Assume you have the following environment:

- Brocade (two switches of 12 ports each, one switch of 16 ports) - Total 40 ports
- McDATA (one switch of 64 ports) - Total 64 ports
- Windows 2000 and Solaris Hosts (10 hosts with two Fibre Channel connection each) - Total 20 ports
- EMC Subsystem (one subsystem with 16 Fibre Channel ports) - Total 16 ports

The software calculates 140 MAPs in this environment.

Example 2:

Assume you have the same configuration above, and you added several devices to your network that the management server does not support. There are still 140 MAPs in this environment since the management server does not count the ports from devices it does not support.

Example 3:

Assume you have the same configuration as the first example with two Windows 2000 hosts that are directly attached to storage systems with no Fibre Channel (FC) connections with a total of 0 FC ports, as shown in the following figure:




Figure 6-1: An Example of Direct Attached Storage

The software calculates four MAPs (see the figure), since we assume one MAP for each host even though it has no Fibre Channel ports. The storage systems are counted since they are supported by the management server. If you include the MAPs from the first example (140 MAPs), it brings the total to 144 MAPs.

Importing a License File

If you cannot find the license file you want to import or if you are interested in expanding your license for managing additional elements, contact your software representative.

To import a license file:

1. Click the **Security** menu.
2. Click **Licenses** from the drop-down menu.
3. Click the **Import License File** button.
4. Click the **Browse** button.
You are shown the file system of the computer being used to access the management server.
5. Select the license file.
6. Click **OK**.

Viewing Cumulative Licenses

Use view cumulative license feature to view the complete number of elements the management server supports at the current time. The software adds up the number of licensed components from the licenses and takes into account the expiration date. See Table 6-1, "License Requirements," on page 196 for more information about the licensing capacities displayed.

Important: You cannot modify the license file since it is encrypted. If you want to increase the number of elements the management server is allowed to manage, contact technical support.

To view cumulative licenses:

1. Click the **Security** menu.
 2. Click **Licenses** from the drop-down menu.
 3. Click the **View Cumulative Licenses** button.
The properties for the cumulative licenses are displayed:
-

Notice in the **Cumulative License** window each feature has a property that is either set to true or false. If a value for a property is set to true, your license lets you access that feature. Likewise, if the value is set to false, the license does not let you access that feature.

You can determine how many elements your licenses supports by looking at the **Current Usage Summary** table at the bottom of the page. The management software displays the cumulative number for each type of licensed capacity in this table.

Viewing a Specific License

Important: Do not manually edit the license. If you want to increase the number of elements the management server is allowed to manage, contact technical support.

To view a license for a type of user:

1. Click the **Security** menu.
2. Click **Licenses** from the drop-down menu.
3. Click the  button corresponding to the license you want to view. The license name and file name is listed, along with its properties.

You can determine how many MAPs and/or MALs this license supports by looking at the properties in the license file. However, that can be misleading if you have other licenses that also provide support for MAPs and MALs. It is suggested you look at the cumulative licenses to obtain a total of the MAPs and MALs that are supported. See the topic, "Viewing Cumulative Licenses" on page 200 for more information about viewing cumulative licenses.

The MALs are split into three properties, LICENSE_FSRM_SIZE_TB, LICENSE_MAL_DATABASE, LICENSE_MAL_EXCHANGE. See the following list for the properties used for tracking MAPs and MALs:

- LICENSE_FSRM_SIZE_TB - The amount of space in Terabytes you are allowed for File Server SRM.
 - LICENSE_MAL_DATABASE - The number of database application instances the management server is allowed to monitor, such as Oracle and Sybase Adaptive Server Enterprise.
-

- ❑ LICENSE_MAL_EXCHANGE - The number of Microsoft Exchange instances the management server is allowed to monitor.
- ❑ LICENSE_MAPS - The number of MAPs the management software can manage.

Deleting a License

Important: Before you delete a license, make sure you have made a copy of it. If you delete the wrong license, you may lose access to certain features and/or access to the product. The management server saves the license files in the `%MGR_DIST%\Data\Licenses` directory.

To delete a license:

1. Click the **Security** menu.
2. Click **Licenses** from the drop-down menu.
3. Click the  button corresponding to the license you want to delete.

Import License File Window Appears Too Small

If you have the DPI (dots per inch) setting on your monitor set too high, scroll bars may appear on the Import License File window because not all of the text and/or buttons on the Import License File window can fit.

The higher your DPI setting, the higher the resolution, which means text and buttons occupy more space in the user interface. For example on Windows operating systems, a button that takes up 1 inch of square space with a DPI setting of 96 will take up 1.25 inches of square space when the DPI is set to 120 (125 normal size).

To view the Import License File window properly, you must reset your DPI to a lower setting. Refer to the documentation for your operating system for more information.

To reset the DPI on Windows XP:

1. Right-click your desktop, and select **Properties**.
2. Click the **Settings** tab. Then, click the **Advanced** button.
3. Click the **General** tab.
4. Select **Normal size (96 DPI)** from the **DPI Settings** drop-down menu.

Configuring the Management Server

This chapter contains the following topics:

- Trap Generation on page 206
 - Setting Up E-mail Notification on page 207
 - Managing Getting Discovery Details on page 208
 - Managing Product Health on page 211
 - Managing Logging on page 215
 - Managing the Display of Events on page 221
 - Managing File Server SRM on page 226
 - Managing Reports on page 232
 - Managing Performance Collection on page 255
 - Editing the Locale and Currency Settings on page 262
 - Process Names on page 263
-

Trap Generation

You can configure the software so that events received by the system generate SNMP traps, which the software can send to another event monitoring system, such as Micromuse™ Netcool® Solutions or HP OpenView. The software allows up to five SNMP trap destinations. The software can send either SNMPv1 or SNMPv2 traps. Whichever SNMP version you select will be used for all trap destinations. The default is SNMPv1. To change the default to SNMPv2, see “Changing the Default to SNMPv2” on page 206.

The software provides an SNMP MIB for each SNMP version that you can compile into your existing enterprise framework. This MIB contains trap definitions so your enterprise framework can understand the traps. The MIB can be found in ApplQ-Traps-v1.mib for SNMPv1 or ApplQ-Traps-v2.mib for SNMPv2 located in the %MGR_DIST%\Tools directory and in the Tools directory on the CIM Extension CD-ROM. You should only compile one of the two MIBs into your enterprise framework. Choose the MIB file that corresponds to the SNMP version you are using.

To configure trap forwarding:

1. Access the management server, as described in “Accessing the Management Server” on page 14.
2. Click **Configuration > Traps**.
3. Select the **Enable Trap Generation** option.
4. In the **SNMP Community String** field, type the SNMP community string, which is used for filtering.
Important: If you enter a value in the **SNMP Community String** field, the SNMP agent must know the SNMP community string entered in the field to receive the SNMP trap.
5. Click the **Save** button.
6. In the **New SNMP Destination** field, type the IP address of the server running an SNMP agent.
7. Click the **Add** button.

Changing the Default to SNMPv2

To change the system to send SNMPv2 traps:

1. Click **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree.
 2. Click **Advanced** in the Disk Space tree.
-

3. Click **Show Default Properties** at the bottom of the page.
4. Copy the following command. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.

```
outgoingSnmpTrapVersion=v2
```
5. Return to the Advanced page (**Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree).
6. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser.
7. Make sure the property is not commented out by removing the hash (#) symbol in front of the property.
8. When you are done, click **Save**.
9. Restart the service for the management server for your changes to take effect.
While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Setting Up E-mail Notification

The management server provides e-mail notification for reports and policies. For example, you can set up the management server to notify you by e-mail when the amount of free space on a host becomes too low.

You must assign an SMTP server from which the management server can send its e-mail notifications.

To configure e-mail notification:

1. Access the management server, as described in “Accessing the Management Server” on page 14.
 2. Select the **Configuration > E-mail Server** option in the upper-right corner.
 3. (Required) Select **Enabled** to enable e-mail notification.
 4. (Required) In the **Name** field, type the DNS name or the IP address of the Simple Mail Transfer Protocol (SMTP) server you want to use to send the e-mail notification.
-

5. (Required) In the **Port** field, type the port of the SMTP server you want to use to send the e-mail notification.
6. In the **User Name** field, enter a user name for the SMTP server.
7. In the **Password** field, enter a password for the SMTP server.
8. In the **Verify Password** field, re-type the password you entered previously.
9. (Required) In the **Sender** field, type the e-mail address of the sender. This address is displayed in the **From** field in the e-mail.
10. If you want the replies to go to an e-mail address other than the e-mail address specified in the **Sender** field, type the e-mail address you want to receive replies in the **Reply to** field.
11. Click the **Save** button.

Important: You should try sending a test e-mail.

To send a test e-mail:

1. In the **To** field, type an e-mail address.
The software verifies the address entered has a correct form.
2. In the **Subject** field, type a subject to distinguish from notification of a real event, for example:
`Testing E-mail Notification`
3. In the **Message** field, type a message, for example:
`I'm just testing e-mail notification.`
4. Click the **Send Test Message** button.

Managing Getting Discovery Details

You can schedule the management server to obtain discovery details at a specified interval. To learn more about scheduling details, see the following topics:

- Adding a Discovery Schedule on page 209
 - Disabling a Schedule on page 210
 - Editing a Schedule on page 211
 - Removing a Schedule on page 211
-

Adding a Discovery Schedule

Schedule getting details when the network is not busy.

Keep in mind the following:

- All collectors are stopped during include infrastructure details. This means that during the gathering of infrastructure details, information about the SAN, such as for Performance Explorer, is not updated.
- If you are creating multiple Discovery schedules, care must be taken to avoid scheduling conflicts—concurrently scheduled Discovery tasks—and that each scheduled task has enough time to start and finish before the next Discovery task is scheduled to start. See “Do Not Run Overlapping Discovery Schedules” on page 812.

To schedule discovery details:

1. Access the Discovery page by clicking **Configuration > Discovery**.
2. Click the **Schedule** tab.
3. Click the **New Schedule** button.
4. In the **Name** field, type a name for your discovery schedule.
5. In the **Description** field, type a description for your discovery schedule.
6. Select one or more of the following:
 - ☐ **Include infrastructure details** - During include infrastructure details, the management server gathers detailed information about the SAN infrastructure and the process can be network intensive.
 - ☐ **Force Device Manager Refresh** - If you want the device managers for HDS and EMC Symmetrix storage systems to obtain the latest information whenever getting discovery details occurs. The management server obtains most of its information for HDS and EMC Symmetrix storage systems from their device managers. If the device managers do not have the latest information, the management server also displays the outdated information. See the following topics for more information: “Excluding EMC Symmetrix Storage Systems from Force Device Manager Refresh” on page 97.
7. Select the **Enable** option.
8. Click the calendar icon, .
9. In the **Time** field, type the time in 24-hour format with the hour and minutes separate by a colon, for example, 22:15. Click the date on which you want the next time getting discovery details to run. Today's date is highlighted in pink. Click **Set**.

The date and time appear in the **Next Scheduled Run** field in the yyyy-mm-dd hh:min format.

If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.

10. In the **Repeat Interval** field, type an interval. Select one of the following units from the drop-down menu:
 - ☐ **Second(s)**
 - ☐ **Minute(s)**
 - ☐ **Hour(s)**
 - ☐ **Day(s)**
 - ☐ **Week(s)**
11. Click **Next**.
12. Select the discovery groups you want included in the discovery.

Important: Only the elements in the discovery groups you select are included in discovery.

13. Click **Finish** when done.
You can always disable a schedule. To learn more about how to remove the schedule, see the topic “Disabling a Schedule” on page 210.

Disabling a Schedule

To disable a schedule for getting SAN topology details:

1. Access the Discovery page by clicking **Configuration > Discovery**.
 2. Click the  button corresponding to the discovery schedule you want to disable.
 3. Deselect the **Enable** option.
 4. Click **Next**.
 5. Verify that the correct discovery groups are selected.
 6. Click **Finish**.
The schedule is disabled.
-

Editing a Schedule

To edit a schedule:

1. Access the Discovery page by clicking **Configuration > Discovery**.
2. Click the  button corresponding to the discovery schedule you want to modify.
3. If necessary, change the following properties:
 - ☐ Name
 - ☐ Description
 - ☐ Type of discovery
 - ☐ Schedule of the discovery
4. Click **Next**.
5. If necessary, change the discovery group you want assigned to the schedule.
6. Click **Finish**.

Removing a Schedule

To remove a schedule:

1. Access the Discovery page by clicking **Configuration > Discovery**.
2. Click the  button corresponding to the discovery schedule you want to remove.
The schedule is removed.

Managing Product Health

To obtain information from Product Health, you must do the following. .

1. Added the management server to your discovery list.
 - a. Click **Discovery > Setup**.
 - b. Click the **Monitoring Product Health** link next to Step 1.
 - c. Click **Add**.
 2. Discover the management server and include it in Get Details.
-

The Product Health menu option helps you to monitor and manage the management server. At installation, a CIM Extension was automatically installed on the management server so you could monitor the management server just as you would any other host. Product Health does the following:

- **Disk Space Monitoring** - This feature keeps track of the management server's use of disk space. See "Enabling Disk Space Monitoring" on page 212.
- **Database Alert Log** - The Database Alert Log scans the management server for critical errors at a specified interval and displays the information in its own chart. This setup frees up Event Manager for monitoring other elements. See "Enabling the Scanning of Critical Events of the Management Server Database" on page 220.
- **Log Files** - You can view and download product logs, as described in
- **Scheduled RMAN backups** - This feature lets you schedule RMAN backups. See "Scheduling RMAN Hot Backups of the Database" on page 275. If you see the buttons on this page disabled, the product is set to no archive mode. See "Setting the Archive Mode" on page 284 for more information about changing the archive mode.
- **Advanced** - This feature lets you modify advanced settings so you can configure the product to run optimally in your environment. See "Modifying Java Memory Settings" on page 214 and "Customizing Properties" on page 214 for more information.

Enabling Disk Space Monitoring

You can configure the management server to monitor itself just as it would any other element. This feature lets you monitor the amount of free space the management server has left. The management server uses disk space for many of its operations, such as when it collects performance data, gathers element properties, generates events, and creates a backup.

Important: To obtain information from Product Health, you must have already discovered the management server and obtained element details from it. See "Step 1 - Discovering Your Hosts" on page 117 for information on how to discover a host.

1. Make sure you have already performed Get Details for the management server. Discover the management server in the same manner as you would discover a host.
 2. Click **Configuration > Product Health**.
 3. Click the calendar icon, .
-

4. In the **Time** field, type the time in 24-hour format with the hour and minutes separate by a colon, for example, 22:15. Click the date on which you want to run the monitoring. Today's date is highlighted in pink. Click **Set**.

The date and time appear in the **Next Scheduled Run** field in the yyyy-mm-dd hh:min format.

If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.

5. In the **Repeat Interval** field, type an interval. Select one of the following units from the drop-down menu:

- ☐ **Second(s)**
- ☐ **Minute(s)**
- ☐ **Hour(s)**
- ☐ **Day(s)**
- ☐ **Week(s)**

6. Click **Save Changes**.

7. To view the results of the monitoring, click the **Result** tab periodically.

Viewing the Results of Disk Space Monitoring

To make sure your management server does not run out of space, you should check the results for disk space monitoring.

Important: The Results tab appears empty if the management server has not been included in Get Details. To obtain information from Product Health, you must have already discovered the management server and obtained element details from it. See “Step 1 - Discovering Your Hosts” on page 117 for information on how to discover a host.

To access the results for disk space monitoring:

1. Click **Configuration > Product Health**.
2. Select **Disk Space** in the tree.
3. Click the **Results** tab in the Disk Space window.

The following information is displayed:

- ☐ Data/Time
- ☐ Disk Capacity
- ☐ Free Space
- ☐ Database Files
- ☐ Archive Files
- ☐ RMAN Files
- ☐ Temp Tablespace

Advanced Settings

This section describes the following:

- Modifying Java Memory Settings on page 214
- Customizing Properties on page 214

Modifying Java Memory Settings

Do not modify the Java memory settings on Advanced page (**Configuration > Product Health > Advanced**) unless instructed to do so by technical support. Incorrectly changing these settings could adversely impact the performance of the software.

Customizing Properties

The management server lets you modify its properties. These properties control a variety of functionality, such as the ability to specify the number of time-outs for switches. You can always view the default setting of the properties by accessing the Default Properties page. To customize properties:

1. Click **Configuration > Product Health**.
 2. Click **Advanced** in the Disk Space tree.
 3. Click **Show Default Properties** at the bottom of the page.
 4. Copy the commands you want to modify. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
-

5. Return to the Advanced page (**Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree).
6. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
7. Make your changes in the **Custom Properties** field. Make sure the property is not commented out by removing the hash (#) symbol in front of the property.
8. When you are done, click **Save**.
9. Restart the service for the management server for your changes to take effect. While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Managing Logging

This section provides information about the following:

- Accessing the Log Files on page 215
- Downloading Logs to a File on page 216
- Downloading the User Audit Log on page 217
- Downloading the Discovery Summary Log on page 217
- Displaying a Log File in a Command Prompt Window on page 218
- Changing the Provider Log Level on page 218
- Enabling the Scanning of Critical Events of the Management Server Database on page 220
- Viewing the Results of Logging on page 221

Accessing the Log Files

You can obtain information about the software's and CIMOM's transactions in the log files, which are in the %MGR_DIST%\logs directory. CIMOM is a component in the CIM management infrastructure that handles the interaction between management applications and providers. The log files may contain information that is difficult to understand.

The following listing provides the logs files that would be the most useful to you. It is not a complete list of all the logs. Some logs are for internal use only.

- `appiq.log` - Provides information about the transactions in the software including web messages, EJB information, and general exceptions.
- `cimom.log` - Provides information about threads with CIMOM, such as provider transactions.

If you want to view all log files, you can save them in a zip file as described in the topic, “Downloading Logs to a File” on page 216.

Downloading Logs to a File

If you run into problems with the management server, use the “download logs to a file” feature to track the problem. This feature saves all the log files in a zip file, which is then stamped with the date and time (24-hour clock).

Note: Some of the log files are generated only when you run certain features. For example the `reports.log` is only generated when you run reports.

To save all logs to a file:

1. Click **Configuration > Product Health**.
 2. Select **Log Files** in the Product Health tree.
 3. Click the **Download Logs** button.
 4. When you are asked if you want to open or save the file, do one of the following:
 - ☐ **Microsoft Internet Explorer** - Click the **Save** button.
 - ☐ **Netscape 7** - Select the **Save this file to disk** option.
The Downloading window appears.
 5. Type a name for the *.zip file and select the directory to which you want to save the file. Keep in mind the following:
 - ☐ Netscape 7 assumes the file is an HTML file. If you are running Netscape, make sure the file type is selected to **All Files** and the file extension is zip. You can make Netscape 7 recognize the zip file next time, by clicking the **Advanced** button.
-

- ☐ Make sure the .zip file is saved to a location other than the local disk drives of the management server.
- 6. Click the **Save** button in the Save As window.

Downloading the User Audit Log

You can determine who has been accessing your machine by viewing the user audit log.

To access the user audit log:

1. Click **Configuration > Product Health**.
2. Click **Download User Audit Log**.
3. Save the zip file.
4. Unzip the zip file.
5. Open the text file in a text editor. If you are using NotePad, you may need to select **Format > Word Wrap**.

You may notice that information is displayed as follows:

```
[2005-05-09 09:22:24] INFO [admin/1000/computername.companyname.com] login succeeded
```

where

- [2005-05-09 09:22:24] - The time and date the action occurred.
- INFO - Level of warning
- [admin/1000/computername.companyname.com] - The user name and DNS name of the computer used to log into the management server. In this case, the user logged in as admin from computername.companyname.com.
- login succeeded - the action that occurred.

Downloading the Discovery Summary Log

You can view status information from Get Details by viewing the discovery summary logs, as described in the following steps:

1. Click **Configuration > Product Health**.
2. Click **Download Discovery Summary**.

3. Save the zip file.
4. Unzip the zip file.
5. Open the `GAEDSummary.log` file in a text editor. If you are using NotePad, you may need to select **Format > Word Wrap**.

Displaying a Log File in a Command Prompt Window

The software ships with `tail.exe`, which can display and update a log file in a command prompt window. This utility is extremely helpful if you do not want to enable the option that lets the management server service interact with the desktop.

To display a log file in a command prompt window:

1. Open a command prompt window.
2. Go to the `%MGR_DIST%\logs` directory by entering the following at the command prompt.
`c:\>cd %MGR_DIST%\logs`

3. Enter the following at the command prompt window:

```
%MGR_DIST%\logs>tail -f appiq.log
```

where `appiq.log` is the log file you want displayed in the command prompt window.

The `tail.exe` utility checks the file for updates and appends them to text in the command prompt window.

`Tail.exe` is a program distributed under the GNU General Public License. See <http://www.gnu.org> for more information.

Changing the Provider Log Level

The management server obtains information from its discovered elements through providers, which communicate with the hardware interface. These providers provide by default superficial logging to the `%JBoss4_DIST%\server\appiq\logs\appiq.log` file. You can change the level of logging provided by selecting the new level of logging from the **Provider Log Setting** combo box on the Product Health page (**Configuration > Product Health**), selecting **Log Files** and then clicking **Apply**. Only one logging level option can be selected.

Use the table below as a guideline for the different options. Several of the options mention providers. A provider is software that gathers information from an element, such as a switch.

Table 7-1: Logging Levels

Log Level Option	Description	Use When You
Default Logging	Provides superficial logging	Do not want additional logging.
Detailed tracing of Brocade provider	Provides detailed logging of the providers used to gather information from the Brocade switch.	Are having difficulty obtaining information from a Brocade switch.
Detailed tracing of CIM Object Manager	Provides detailed logging of the infrastructure that handles the interaction between management applications and providers. The CIM Object Manager supports services such as event notification, remote access, and query processing.	Are having difficulty obtaining information from the CIM Object Manager. You are having difficulty obtaining information from event notification, remote access, and query processing.
Detailed tracing of CLARiiON provider	Provides detailed logging of the providers used to gather information from CLARiiON storage systems.	Are having difficulty obtaining information from CLARiiON storage systems.
Detailed tracing of HDS provider	Provides detailed logging of the providers used to gather information from HDS storage systems.	Are having difficulty obtaining information from HDS storage systems.
Detailed tracing of HOST/SERVER provider	Provides detailed logging of the providers used to gather information from hosts/servers.	Are having difficulty obtaining information from host/servers.

Table 7-1: Logging Levels (Continued)

Log Level Option	Description	Use When You
Detailed tracing of all providers	Provides detailed logging of the providers, which gather information from the elements.	Are having difficulty obtaining information from more than one type of element.
Detailed tracing of SYMMETRIX provider	Provides detailed logging of the providers used to gather information from EMC Symmetrix storage systems	Are having difficulty obtaining information from EMC Symmetrix storage systems.

Enabling the Scanning of Critical Events of the Management Server Database

You can configure the management server to scan for only critical events occurring with the database for the management server at a specified time interval. The management server displays the results of these scans under **Configuration > Product Health > Log Files**. To enable this type of scanning, do the following:

1. Click **Configuration > Product Health**.
2. Select **Log Files** in the Product Health tree.
3. Click the calendar icon, .
4. In the **Time** field, type the time in 24-hour format with the hour and minutes separate by a colon, for example, 22:15. Click the date on which you want to run the next backup of the database. Today's date is highlighted in pink. Click **Set**.
The date and time appear in the **Next Scheduled Run** field in the yyyy-mm-dd hh:min format.

If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.
5. In the **Repeat Interval** field, type an interval. Select one of the following units from the drop-down menu:
 - ☐ **Second(s)**
 - ☐ **Minute(s)**

- ☐ **Hour(s)**
- ☐ **Day(s)**
- ☐ **Week(s)**

Note: The minimal interval you can schedule is one day. If you select **Hour(s)**, **Minute(s)** or **Second(s)**, you must enter an interval that equals more than a day. For example, if you select **Hour(s)**, you must enter 24 or more. Just as if you select **Minute(s)**, you must enter 1440.

6. Click the **Enable** option.
7. Click the **Save Schedule** button.
8. To view the results of the scanning, click the **Result** tab periodically.

Viewing the Results of Logging

You can view when an error occurred and at what time it was discovered by accessing the Results tab for logging.

To access the Results tab:

1. Click **Configuration > Product Health**.
2. Select Log Files in the tree.
3. Click the Results tab in the Log Files pane.
The following is displayed:
 - ☐ Scan Date/Time
 - ☐ Error Occurred Time
 - ☐ Error Description

Managing the Display of Events

This section describes the following:

- Controlling the Display of Cleared and Deleted Events on page 222
-

- Modifying the Clearing and Deletion Frequency on page 222
- Configuring the Clearing of Events on page 224
- Configuring the Deletion of Events on page 225

Controlling the Display of Cleared and Deleted Events

You can control how the management server displays events by modifying one or more of the following:

- **The Clearing and Deletion Frequency** - The frequency table determines how often the user interface in Event Manager removes events and marks events as “cleared.” Events always display as they occur in the user interface.
Important: Events are only removed or marked “cleared” when their automatic delay time is completed. See the topics, “Configuring the Clearing of Events” on page 224 and “Configuring the Deletion of Events” on page 225.

Assume you set the clearing and deletion frequency to every 15 minutes with the initial time at 11:00 a.m. So every 15 minutes the management server checks for events marked for deletion and for clearing, and it updated the user interface accordingly. Let's assume Informational events are set to be cleared every hour, and an Informational event occurs at 12:20 p.m. Exactly after an hour has passed, the management server marks this event as “cleared” but the user interface is not updated because the frequency update of the user interface is every 15 minutes. If you looked at the Event Manager at 1:35 p.m. the event would be marked cleared. See “Modifying the Clearing and Deletion Frequency” on page 222 for more information.

- **The Clearing of Events** - You can determine how often events are marked “cleared.” See “Configuring the Clearing of Events” on page 224.
- **The Deletion of Events** - You can determine how often events are deleted. Events are deleted every two weeks by default. See “Configuring the Deletion of Events” on page 225.

Modifying the Clearing and Deletion Frequency

You can modify how often the user interface in Event Manager removes events and marks events as “cleared.” Events are still displayed as they occur in the user interface.

Important: Events are only removed or marked “cleared” when their automatic delay time is completed. See the topics, “Configuring the Clearing of Events” on page 224 and “Configuring the Deletion of Events” on page 225.

Assume you set the clearing and deletion frequency to every 15 minutes with the initial time at 11:00 a.m. So every 15 minutes the management server checks for events marked for deletion and for clearing, and it updated the user interface accordingly. Let's assume Informational events are set to be cleared every hour, and an Informational event occurs at 12:20 p.m. Exactly after an hour has passed, the management server marks this event as “cleared” but the user interface is not updated because the frequency update of the user interface is every 15 minutes. If you looked at the Event Manager at 1:35 p.m. the event would be marked cleared. See “Modifying the Clearing and Deletion Frequency” on page 222 for more information.

To modify the clearing and deletion frequency:

1. Access the Events page by clicking **Configuration > Events**.
2. Click the calendar icon, .
3. In the **Time** field, type the time in 24-hour format with the hour and minutes separate by a colon, for example, 22:15. Click the date on which you want to run the clearing and deletion frequency. Today's date is highlighted in pink. Click **Set**.
The date and time appear in the **Next Scheduled Run** field in the yyyy-mm-dd hh:min format.

If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.
4. In the **Repeat Interval** field, type an interval. Select one of the following units from the drop-down menu:
 - ☐ **Second(s)**
 - ☐ **Minute(s)**
 - ☐ **Hour(s)**
 - ☐ **Day(s)**
 - ☐ **Week(s)**

Note: The higher the interval, the more demand there is on the management server.

5. Select the **Enable** option.
6. Click the **Save Changes** button.

Configuring the Clearing of Events

Depending on the severity of an event, the management server may mark the event as “clear” after 60 minutes. Events designated as Major and Critical are never marked as “clear.” You can change the time delay in clearing an event, and you can specify that the management server never marks an event as “clear”.

To help you in filtering events, you may want to have unimportant events marked as cleared rather than automatically deleted. Depending on how you have configured the deletion of events, you can view the “cleared” events at a later time.

See the following table for the default settings for clearing events.

Table 7-2: Default Settings for Clearing Events

Severity Level	Default Time Delay to Clear the Event (Hours)
Unknown	1
Informational	1
Warning	1
Minor	1
Major	Never
Critical	Never

To change the default time delay to clear an event, do the following:

1. Access the Events page by clicking **Configuration > Events**.
 2. Do one of the following:
-

- ☐ If you never want an event of the specified severity level marked as “cleared,” select the **Never** option next to the severity level.
 - ☐ If you want to change the delay time in clearing an event, select one of the following units of measurement from the combo box and type the number in the adjacent field:
 - **Minutes**
 - **Hours**
 - **Days**
 - **Weeks**
3. Click the **Save Changes** button.

Configuring the Deletion of Events

The management server automatically deletes events after two weeks by default. You can specify for each severity level different time periods for deleting events. For example, you could modify the management server to delete events with the Information severity level every two days. You could also specify the management server to never delete events with the Critical severity level.

To change the default time delay to delete an event, do the following:

1. Access the Events page by clicking **Configuration > Events**.
2. Do one of the following:
 - ☐ If you never want an event of the specified severity level automatically deleted, select the **Never** option under the Automatic Delete Delay column.
 - ☐ If you want to change the delay time in deleting an event, select one of the following units of measurement from the combo box and type the number in the adjacent field:
 - **Minutes**
 - **Hours**
 - **Days**
 - **Weeks**

For example, if you want events that are a week old deleted, you would type 1 and then select **Weeks** in the combo box in the Automatic Delete Delay column.

3. Click the **Save Changes** button.

Managing File Server SRM

To learn more about managing File Server SRM, see the following topics:

- Managing Collectors for File Server SRM on page 226
- Adding File Servers on page 227
- Setting Data Collection Properties on page 228
- File Server SRM Scan Status on page 230
- Performing a Test File Server SRM Scan on page 231

Managing Collectors for File Server SRM

Important: All collectors are stopped during Get Details. This means that during Get Details data, such as for File Server SRM, is not updated.

Note: It is not possible to run data collectors on quarantined elements. Attempting to run data collectors on quarantined elements will result in exceptions in the appiq.log file.

The management server uses File Server SRM data collectors to obtain information from a file server. File Server SRM provides information about a file server, such as the number of its directories. While File Server SRM data collectors are obtaining information from a file server, a second File Server SRM data collection cannot begin.

To manage collectors for File Server SRM:

1. Click **Configuration > FSRM**.
 2. Verify the **Data Collection** tab is displayed.
-

Table 7-3: About the Collectors for File Server SRM

Column Heading	Description
File Server	The name of the file server from which this collector gathers information.
Volumes	The volumes selected for scanning on a file server. Volumes selected for “Scan now” are not listed.
Next Scheduled Run	Displays the date and time when the collector is scheduled to run.
Interval	The interval in minutes the collector runs.
Running?	Gives the status of the collector. Collectors that are scheduled to run display a check mark in this column.
Start/Edit	If no schedule has been created for a file server, this column displays a Start button. Click the Start button to create a schedule. See “Setting Data Collection Properties” on page 228 for more information. If a schedule has been created, this column displays an Edit () button. Click the Edit button to edit the schedule for running a collector. See “Setting Data Collection Properties” on page 228 for more information.
Test Scan	Verifies that the management server can access the file server and perform a high level File Server SRM scan.
Delete	To remove the file server application associated with a particular host, click the Delete () button. You can also use the Delete button located above the table to delete multiple selected file servers.
Stop Collectors	Use the Stop Collectors button located above the table to stop selected collectors.

Adding File Servers

To select hosts to be marked as file servers:

1. Click **Configuration > FSRM**.
2. Verify the **Data Collection** tab is displayed.
3. Click the **Add File Server** button.
4. Click the check boxes for the hosts that you would like to mark as file servers. Only hosts that support File Server SRM are listed.
5. Click the **OK** button. The hosts are marked as file servers and you are returned to the **Data Collection** tab.

Setting Data Collection Properties

The management server obtains information from collectors that gather information from elements at specified times. If you want the management server to monitor file servers, you must start the collector for the specified file server.

For example, assume you installed File Server SRM on a host and you want to view reports about file management, such as partition consumption. You must start the collector for that host.

Keep in mind the following:

- If you want to scan network mounted drives on a host running Windows, you must make sure a user is logged onto the Windows host. The user must have permission to mount the network drives.
- If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.

To set scanning properties:

1. Click **Configuration > FSRM**.
 2. Verify you are looking at the **Data Collection** page.
 3. Do one of the following:
 - ☐ Click the **Start** button corresponding to the file server you want to monitor.
 - ☐ Click the **Edit** button () corresponding to the collector you want to modify.
 4. Under **Local Volumes**, select the local volumes you want to monitor.
 5. Under **External Volumes**, select the external volumes you want to monitor.
 6. Under **Network Volumes**, select the network volume you want to monitor.
-

7. Click the calendar icon, .
8. Enter the time in the Time field. Make sure the time resembles a 24-hour clock, for example 22:00 for 10 p.m.
9. Click the date.
The date is highlighted in pink.

You can navigate the calendar as follows:

 - Displays the same month in the previous year

 - Displays the previous month

 - Displays the next month

 - Displays the same month in the following year

10. When you are done, click the **Set** button.
The management server sets the time.

In the following figure, the calendar sets the date and time at 1:56 p.m. February 13, 2003.




Figure 7-1: Selecting the Date and Time from the Calendar

11. If you want to change the interval at which the collector will run, edit the number in the Interval field.

12. If you want to exclude a directory from the scan, enter it in the Exclude Directory List field. Additional directories must be separated by a comma. A default list of excluded directories can be provided by using the jboss property “fileSrm.excludeDirectories.” The default value for this property is “.snapshot, ~snapshot, RECYCLE.”

Note: Only first level directories are excluded from the scan. For example, if the directory “test” exists at C:\test and C:\\WINNT\system32\test, and you exclude the “test” directory, only C:\test will be excluded. C:\\WINNT\system32\test will still be scanned.

Note: Directory paths are case sensitive.

13. If you want to collect ownership information, click the Collect Ownership Information checkbox (for Windows File Servers only). If you select the option to collect ownership information, you are provided with the option to collect user’s usage information based on home directories or owners. Click the radio button that corresponds to your choice.

Note: This feature overrides the global setting available previously through the jboss property “fileSrm.collectWindowsSID”

14. Click **OK**.

File Server SRM Scan Status

You can review the status of the most recent (or current) scan that took place on each of the volumes on which a scan was started. To check the status of a scan:

1. Click **Configuration > FSRM**.
 2. Click the **Scan Status** tab.
-

Table 7-4: About File Server SRM Scan Status

Column Heading	Description
Server	The name of the file server from which this collector gathers information.
Volume	The volume selected for scanning on a file server.
Scan Start	The time that the scan began.
Last Status Polled	If the scan is still in progress, the latest time status is provided.
Scan End	If the scan is completed, the completion time is provided.
No. Of Files	The number of files that have been scanned.
Current Directory	The directory that is currently being scanned.
Scan Status	The following scan statuses are possible: ■ No Data ■ Scan in Progress ■ Scan Done ■ Scan Failed ■ Transferring Data ■ Transfer Completed ■ Uploading Data ■ Upload Completed ■ Error

Performing a Test File Server SRM Scan

If you have changed your network configuration and you are not sure the management server can still scan the server running File Server SRM, you can perform a test File Server SRM scan. During this scan, the management server verifies that it can access the server and perform a high level File Server SRM scan. It also verifies whether or not it can scan network volumes with host discovery credentials.

Keep in mind the following:

- This feature only works with CIM Extensions builds 5.1 and later.
- If you want to scan network mounted drives on a host running Windows, you must make sure a user is logged onto the Windows host. The user must have permission to mount the network drives.

To perform a test File Server SRM scan:

1. Click **Configuration > FSRM**.
2. Click the **Test** button in the Test Scan column for the corresponding file server. After you click the **Test** button, the management server attempts to access the server and perform a test File Server SRM scan.

Managing Reports

To learn more about managing reports, see one of the following topics:

- Architectural Overview of Report Views and Report Cache Refresh on page 232
- Managing Collectors for Reports on page 236
- Setting the Date and Time for Report Collectors on page 239
- Viewing Scheduled E-mail Deliveries for Reports on page 240
- Editing E-mail Schedules for Reports on page 242
- Viewing Data Aging Statistics for Reports on page 243
- Scheduling Report Cleanup on page 245
- Refreshing the Report Cache on page 245
- Setting Up Global Reporter on page 248

Architectural Overview of Report Views and Report Cache Refresh

Management server reports are based on report views (Materialized views) which are data snapshots of the management server schema at a certain time. The management server has a refresh process to refresh the report views. You can change the frequency of how often the report

views are refreshed, as described in “Refreshing the Report Cache” on page 245. The overall report architecture is displayed in the following figure.




Figure 7-2: Report Views and Report Cache Refresh

Suggestion to Schedule the Report Cache Refresh:

The report cache (report views) is a snapshot data of management server schema to a point in time. The report cache refresh is time and resource consuming. The following are guidelines for scheduling report cache refresh:

- To ensure the report cache has the latest data from Get Details, schedule the report cache refresh after Get Details finishes.
- Make sure to schedule the report cache refresh during off peak hours.

See “Refreshing the Report Cache” on page 245 for information on how to refresh the report cache.

Keep in mind the following:

- After an initial install, make sure you have already run Get Details. Then, refresh the report cache to populate it with the data from Get Details. If you do not refresh the report cache,

the cache is still in its initial state, which is empty. If the cache is empty, so too will the reports since they obtain their information from the report cache.

- During the report cache refresh, the report views are truncated and repopulated with latest data. If a report is run during report cache refresh and the corresponding view is getting refreshed, the report displays with no data.

Report Refresh Status

The management server has two types of views for its reports. During a report cache refresh these views are updated. You can check the status of these views as described in this section:

- **MVIEWCORE_STATUS** - This table keeps track of the refresh status of the core views. The core views are the views starting with `mvc`, `mvca` and `mvcs` as shown in Table 13-3, “Description of the Report Views,” on page 595. To obtain detailed information of **MVIEWCORE_STATUS**, see Table 13-47, “**MVIEWCORE_STATUS**,” on page 626.
- **MVIEW_STATUS** - This table keeps track of the refresh status of the regular views, which are views starting with `mv_`, as shown in Table 13-3, “Description of the Report Views,” on page 595. To obtain detailed information of the **MVIEW_STATUS**, see Table 13-48, “**MVIEW_STATUS**,” on page 626.

To query the **MVIEW_STATUS** table:

On Windows:

1. Enter the following from the command prompt:
`Sqlplus appiq_system/password`
where `password` is the password for the `appiq_system` account.
 2. Enter the following at the command prompt:
`Sql>col lastRefresh format a30`
This command ensures the data is displayed in a readable format.
 3. Enter the following at the command prompt on one line with a space between the closing parenthesis and `lastRefresh`:
`Sql> select mviewname, to_char(last_refresh_time,'mm/dd/yyyy hh24:mi:ss') lastRefresh,`
 4. Enter the following at the command prompt:
`status`
 5. Enter the following at the command prompt:
`from mview_status`
 6. Enter the following at the command prompt:
`order by 2;`
-

On Solaris:

1. From the command line change to the `$MGR_DIST/install` directory.
2. Run the `uservars` script to set the variables (``eval uservars.sh``)
3. Enter the following from the command prompt:
`Sqlplus appiq_system/password`
where `password` is the password for the `appiq_system` account.
4. Enter the following at the command prompt:
`Sql>col lastRefresh format a30`
This command ensures the data is displayed in a readable format.
5. Enter the following at the command prompt on one line with a space between the closing parenthesis and `lastRefresh`:
`Sql> select mviewname, to_char(last_refresh_time,'mm/dd/yyyy hh24:mi:ss') lastRefresh,`
6. Enter the following at the command prompt:
`status`
7. Enter the following at the command prompt:
`from mview_status`
8. Enter the following at the command prompt:
`order by 2;`

To query the `MVIEWCORE_STATUS` table:

On Windows:

1. Enter the following from the command prompt:
`Sqlplus appiq_system/password`
where `password` is the password for the `appiq_system` account.
2. Enter the following at the command prompt:
`Sql>col lastRefresh format a30`
This command ensures the data is displayed in a readable format.
3. Enter the following at the command prompt on one line with a space between the closing parenthesis and `lastRefresh`:
`Sql> select mviewname, to_char(last_refresh_time,'mm/dd/yyyy hh24:mi:ss') lastRefresh,`
4. Enter the following at the command prompt:
`status`
5. Enter the following at the command prompt:
`from mviewcore_status`

6. Enter the following at the command prompt:
`order by 2;`

On Solaris:

1. From the command line change to the `$MGR_DIST/install` directory.
2. Run the `uservars` script to set the variables (``eval uservars.sh``)
3. Enter the following from the command prompt:
`Sqlplus appiq_system/password`
where `password` is the password for the `appiq_system` account.
4. Enter the following at the command prompt:
`Sql>col lastRefresh format a30`
This command ensures the data is displayed in a readable format.
5. Enter the following at the command prompt on one line with a space between the closing parenthesis and `lastRefresh`:
`Sql> select mviewname, to_char(last_refresh_time,'mm/dd/yyyy hh24:mi:ss') lastRefresh,`
6. Enter the following at the command prompt:
`status`
7. Enter the following at the command prompt:
`from mviewcore_status`
8. Enter the following at the command prompt:
`order by 2;`

Managing Collectors for Reports

The management server uses statistics collectors to gather information for reports. To view a report you must have its corresponding collector running.

Important: All collectors are stopped during Get Details. This means that during Get Details data, such as for reports is not updated.

Note: It is not possible to run data collectors on quarantined elements. Attempting to run data collectors on quarantined elements will result in exceptions in the `appiq.log` file.

To view collectors for reports, click **Configuration > Reports**. Then, click the **Data Collection** tab.

Table 7-5: About Collectors for Reports

Column Heading	Description
Element	The name of the element from which this collector gathers information.
Element Type	The type of element from which the collector gathers information.
Reports	The reports the collector is responsible for providing information.
Next Scheduled Run	Displays the date and time when the collector is scheduled to run.
Interval	The interval in minutes the collector runs.
Running?	Gives the status of the collector. Collectors that are running display a check mark in this column.
Edit	To edit the schedule for running a collector, click the Edit  button. Then set the date and time. See “Setting the Date and Time for Report Collectors” on page 239 for more information. Important: If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the first day of the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.
Action	Displays one of the following buttons: ■ Stop - Stops the collector. The corresponding reports display information only gathered previously. See “Stopping Collectors” on page 238. ■ Start - Starts the collector. When you start a collector, it begins gathering information for its corresponding reports. See “Starting Collectors” on page 238.

Starting Collectors

Important: After you click **OK**, the collector starts immediately if the date and time you set has not passed. If the set time has passed, the collector starts two minutes after you click **OK**.

To start a collector:

1. Click **Configuration > Reports**. Then, click the **Data Collection** tab.
2. Click the **Start** button corresponding to the collector you want to start.
3. Click the calendar icon, .
4. Set the date and time as described in the topic, "Setting the Date and Time for Report Collectors" on page 239.
5. When you are done with setting the date and time, enter an interval for the collector to gather information in minutes.

Important: If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.

6. Click **OK**.
7. To start more than one collector at once, select more than one collector and then click the **Start Collectors** button.

Stopping Collectors

When you stop a collector, the management server stops gathering information for which the collector is responsible. For example, if a collector is not running, its corresponding reports are no longer receiving information to display. One of the following occurs:

- If there was originally no information gathered for the report, no data appears in the report.
- If information was previously gathered for the report, old data appears in the report.

To stop a collector.

1. Click **Configuration > Reports**. Then, click the **Data Collection** tab.
-

2. Click the **Stop** button corresponding to the collector you want to stop.
The collector stops gathering information for its corresponding reports.

To stop more than one collector at once, select more than one collector and then click the **Stop Collectors** button.

Setting the Date and Time for Report Collectors

Important: After you click **OK**, the collector starts immediately if the date and time you set has not passed. If the set time has passed, the collector starts two minutes after you click **OK**.

To set the date and time, do the following:

1. Click **Configuration > Reports**. Then, click the **Data Collection** tab.
2. Click the **Stop** button corresponding to the collector you want to stop.
3. Click the **Edit** button, .
4. Click the calendar icon, .
5. Enter the time in the time field. Make sure the time resembles a 24-hour clock, for example 22:00 for 10 p.m.
6. Click the date.
The date is highlighted in pink.

You can navigate the calendar as follows:



- Displays the same month in the previous year



- Displays the previous month



- Displays the next month



- Displays the same month in the following year

7. When you are done, click the **Set** button.
The management server displays the date and time.

Important: If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.




Figure 7-3: Selecting the Date and Time from the Calendar

In the previous figure, the calendar sets the date and time at 1:56 p.m. February 13, 2003.

Viewing Scheduled E-mail Deliveries for Reports

The Scheduled Deliveries tab displays all e-mail schedules for reports. An e-mail schedule instructs the management server to send a particular report to one or more recipients at a regular interval. For example, you could create an e-mail schedule that sends a host utilization report to your boss on a weekly basis. To learn more about creating e-mail schedules, see “Adding an E-mail Schedule for a Report” on page 571.

To view all e-mail schedules:

1. Click **Configuration > Reports**.
2. Click the **Scheduled Deliveries** tab at the top of the screen.
All e-mail schedules for reports are displayed.

Table 7-6: Viewing System Wide E-mail Schedules

Column Name	Description
Recipient	The person who receives the report.
Subject	The subject of the e-mail.
Format	The format of the report sent: ■ PDF ■ Microsoft EXCEL ■ XML
Scheduled By	The user who scheduled the report.
Edit	Click the  button to edit a schedule. See “Editing an E-mail Schedule for a Report” on page 574 for information about the options displayed in this window.
Delete	Click the  button to remove the corresponding schedule. To delete multiple schedules, select the schedules you want to delete in the far left column. To quickly select all schedules, select the box to the left of the Recipient column. Then, click the Delete button above the table.

Editing E-mail Schedules for Reports

Important: Schedule to send your reports soon after a report cache refresh. The reports display data that is in the report cache. If the report cache contains old data, the reports you send by e-mail will also show old data. The reports are refreshed every six hours by default. For example, assume you added an e-mail schedule that sends a report daily at 7 a.m. Also, assume you scheduled your report cache refreshes to take place daily at 8 a.m. Your reports will most likely show outdated data. It would make more sense to schedule your report cache refresh at 7 a.m. and then schedule to send your reports soon afterwards. See “Scheduling a Report Cache Refresh” on page 246.

To edit an e-mail schedule:

1. Click **Configuration > Reports**.
2. Click **Scheduled Deliveries** at the top of the screen.
3. Under the Edit column, click the **Edit** () button.
4. In the **To** field, change the recipient's e-mail address.
5. In the **Subject** field, change the subject of the e-mail.
6. In the **Message** field, change a message describing the report.
If you are e-mailing reports in bulk, you might want to let users know the e-mail is being sent by an automated process. You might also want to provide an e-mail address for users to provide feedback, for example:

This e-mail and its attached report are generated automatically. If you would like to change how often the report is sent to you or you want to be taken off the list, please contact `username@companyname.com`.

7. From the **Format** drop-down menu, select one of the following formats:
 - ☐ **PDF** - Requires the use of Adobe Acrobat, which can be downloaded for free from <http://www.adobe.com>.
 - ☐ **Excel** - Requires the use of Microsoft Excel.
 - ☐ **XML** - Requires the user has an understanding of XML.
 8. In the **Time to Run** field, type the time you want to send the report. This time must be entered in the 24-hour format. For example, if you want a report sent at 2:15 p.m., you would type 14:15 in **Time to Run** field.
 9. Select one of the following options to determine how frequently you want to send the report.
-

- ☐ **Daily** - If you selected daily, select how frequently you want the management server to send the report.
 - **Everyday** - The report is sent everyday.
 - **Weekday** - The report is sent only Monday through Friday.
 - **Everyday for a specified number of days** - Fill in the number of days you want the report to be sent daily. After the specified number of days, the report is no longer sent. For example, you could use this feature to send reports to someone's replacement while the person is away on vacation.
 - ☐ **Weekly** - If you selected weekly, use the **Frequency** drop-down menu to select the day of the week on which you want the report sent.
 - ☐ **Monthly** - If you selected monthly, select the time during the month you want the report sent.
10. Click the **OK** button.

Viewing Data Aging Statistics for Reports

Data Aging includes Data Rollup and Garbage Collection. Data Rollup controls how often a set of data is summarized. For example, hourly data is rolled into the daily table periodically. Garbage Collection refers to how long a set of data is preserved before it is permanently removed from the database.

The settings on this page control data aging for both reports and performance statistics.

Important: Do not modify the data on this page unless instructed by customer support. Changing them incorrectly can adversely affect the management server.

To view data aging statistics:

1. Click **Configuration > Reports**.
2. Click the **Data Aging** tab at the top of the screen.
Data aging statistics are displayed in the table.

Important: Perform the following steps **ONLY** if customer support has instructed you to modify one of the collectors on the page:

3. Click the **Edit** () button.
4. In the **Time** field, type the time in 24-hour format with the hour and minutes separate by a colon, for example, 22:15. Today's date is highlighted in pink. Click **Set**.
The date and time appear in the **Next Scheduled Run** field in the yyyy-mm-dd hh:min format.

If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.

5. Click the **Enable** option.

Note: If you are not allowed to disable the collector, the Enable option is unavailable.

6. To change the repeat interval, In the **Repeat Interval** field, type an interval. Select one of the following units from the drop-down menu:
 - ☐ **Second(s)**
 - ☐ **Minute(s)**
 - ☐ **Hour(s)**
 - ☐ **Day(s)**
 - ☐ **Week(s)**
 7. (Garbage Collection only) To change how long the data is preserved, type an interval in the **Preserve** field, and then select one of the following from the drop-down menu to the right of the field:
 - ☐ **Second(s)**
 - ☐ **Minutes**
 - ☐ **Hours**
 - ☐ **Days**
 - ☐ **Weeks**
 8. Click the **OK** button.
-

Scheduling Report Cleanup

Temporary files are created when you use Reporter to create reports. The report cleanup does not delete any data. It is used primarily for deleting temp files related to report management, such as when a report is viewed in the PDF or Microsoft Excel format.

You can schedule how often these files are removed, as described in the following steps:

1. Click **Configuration > Reports**.
2. Click the **Report Cleanup** tab at the top of the screen.
3. Click the calendar icon, .
4. In the **Time** field, type the time in 24-hour format with the hour and minutes separate by a colon, for example, 22:15. Click the date on which you want to run the next report cleanup. Today's date is highlighted in pink. Click **Set**.
The date and time appear in the **Next Scheduled Run** field in the yyyy-mm-dd hh:min format.

If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.
5. In the **Repeat Interval** field, type an interval. Select one of the following units from the drop-down menu:
 - ☐ **Second(s)**
 - ☐ **Minute(s)**
 - ☐ **Hour(s)**
 - ☐ **Day(s)**
 - ☐ **Week(s)**
6. Click the **Enable** option.
7. Click the **Save Changes** button.

Refreshing the Report Cache

The management server gathers information for reports from the database every six hours by default. The management server stores this information in its report cache and displays it when a report is requested. If you are seeing outdated information in the report, you can refresh the report using one of the following methods:

- **Refresh the report cache now.** - See “Refreshing the Report Cache Immediately” on page 246.
- **Schedule the report cache to be refreshed.** - See “Scheduling a Report Cache Refresh” on page 246.

Refreshing the Report Cache Immediately

The **Refresh Now** button lets you update the report cache immediately.

Keep in mind the following:

- If Get Details is occurring, wait for it to finish before clicking the **Refresh Now** button. This technique ensures the database is completely updated and thus, your reports will be as accurate as possible. Get Details collects the latest data. When you refresh the report cache, the management server transfers the information collected from Get Details and transfers it to the report cache.
- While the report cache is being refreshed, reports display no data.

To refresh the report cache now:

1. Click **Configuration > Reports**.
2. Click the **Report Cache** tab at the top of the screen.
3. Click **Refresh Now**.

Scheduling a Report Cache Refresh

The management server lets you schedule the refreshing of the report cache.

Keep in mind the following:

- The higher the frequency of the report cache interval, the more stress you put on the management server. A very frequent report cache interval, such as every 10 minutes, could hurt the response time of the management server to perform other tasks.
 - Make sure you schedule refreshing the report cache after Get Details. This technique ensures the database is completely updated and thus, your reports will be as accurate as possible. Get Details collects the latest data. When you refresh the report cache, the management server transfers the information collected from Get Details and transfers it to the report.
 - While the report cache is being refreshed, reports display no data.
-

If you find you are still viewing old information regarding elements on the network, you may need to perform Get Details. It is best to perform Get Details at regular intervals. See “Adding a Discovery Schedule” on page 209.

To schedule a report cache refresh:

1. Click **Configuration > Reports**.
2. Click the **Report Cache** tab at the top of the screen.
3. Click the calendar icon, .
4. In the **Time** field, type the time in 24-hour format with the hour and minutes separate by a colon, for example, 22:15. Click the date on which you want to run the next report cache refresh. Today's date is highlighted in pink. Click **Set**.
The date and time appear in the **Next Scheduled Run** field in the yyyy-mm-dd hh:min format.

If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.
5. In the **Repeat Interval** field, type an interval. Select one of the following units from the drop-down menu:
 - ☐ **Second(s)**
 - ☐ **Minute(s)**
 - ☐ **Hour(s)**
 - ☐ **Day(s)**
 - ☐ **Week(s)**
6. Click the **Enable** option.
7. Click the **Save Changes** button.

Setting Up Global Reporter

Important: Depending on your license, Global Reporter may not be available. See the “List of Features” to determine if you have access to Global Reporter. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center**).

Global Reporter lets you gather data for global reports from multiple management servers. For example, lets assume you have three management servers: one in London, one in Tokyo, and one in New York City. You are located in New York City and you want to obtain data in your global reports from the management servers in London and Tokyo, as well as from your own management server in New York City. Through the use of Global Reporter, you can gather data from all three sites.

When you set up Global Reporter, the management server pulls the data from the local database views at these sites through the database link. The flow of the data is shown by the arrows in the following figure. A global reporting view on the server running Global Reporter contains information in the database that can be used for local and global reports. The management server in New York is referred to as a Global Reporter server because it has Global Reporter enabled.




Figure 7-4: An Example of Global Reporting

Table 7-7: Description of “An Example of Global Reporting”

Item	Description
1	Management server for the SAN in Tokyo.
2	Database containing the local materialized views for the management server in Tokyo.
3	Management Server for the SAN in New York City.
4	Database for the Management Server in New York City. It contains the local materialized views for New York City, in addition to the global views containing local data from Tokyo, London, and New York.
5	Management server for the SAN in London
6	Database containing the local materialized views for the Management Server in London.

Keep in mind the following when setting up global reporting:

- **Multiple Global Reporter Servers** - Global reporting can be set up at multiple sites. Your associates in London and Tokyo can also set up global reporting at their site so they obtain data gathered by management server in New York City.
 - **Multiple Sites** - If you do have multiple Global Reporter servers, each Global Reporter server must contact each site that has the data the Global Reporter server wants to include in global reports. For example, assume you set up a Global Reporter server in Tokyo that collects data from a site in London. You also have another Global Reporter server in New York City. The Global Reporter server in New York City must contact the Global Reporter server in Tokyo and the site in London to obtain data from Tokyo and London. If the Global Reporter server in New York City contacts just the Global Reporter server in Tokyo, it does not obtain data from the site in London.
 - **Data Gathering** - The Global Reporter server gathers data from sites as scheduled or when the **Refresh Now** button on the Global Reporter tab is clicked. The management server pulls data from a site's local materialized views, which are snapshots of relevant data from the site's local database. During a refresh, the Global Reporter server contacts each site and attempts to pull data from its materialized views. If a site's local view is refreshing, the Global Reporter server skips pulling the data from that view and proceeds pulling data from the remaining views. The Global Reporter server then attempts to pull data from the skipped views.
-

- **Firewalls** - If a site is behind a firewall, the port for the Oracle TNS listener must be open. For example, if the site uses port 1521 for its TNS listener, this port must be open on the firewall.
- **Contact Status** - The Contact Status column on the Global Reporter tab only verifies if the management server was able to contact the site. It does not specify if the refresh was successful. For example, if the Global Reporter server is able to contact the site, but the refresh fails. The Contact Status column still reports “SUCCESS” because the site was able to be contacted. If you have trouble contacting the site, try pinging it to verify the network connection is working. If ping works, verify the management server on the site is running.
- **Unable to Contact Site** - If the management server is unable to contact one of the sites in the Global Reporting list, the refresh process will not start. You can verify if a site can still be contacted by clicking the **Test** button for the site. You may want to try pinging the machine and verify that management server on the remote server is running. If the site cannot be contacted, remove the obsolete site from the list by clicking the  button.
- **Security** - Users, whose role allows them access to view global reporting, can view all the elements throughout the enterprise. Grant access to viewing global reports only to those, who should be allowed to view all elements. Users belonging only to the following roles are given default privileges to view Global Reporter: CIO, Domain Admin, and Storage Administrator.
- **No data is displayed** - The Global Reporter server gathers information the site has collected. If the site has not collected any data, the Global Reporter server cannot obtain information from that site. Also, the Global Reporter server cannot display data until it has refreshed its view. The views are created at installation time and are empty until the view is refreshed. You can perform an immediate refresh by clicking the **Refresh Now** button on the Global Reporter server.
- **Data is outdated** - The Global Reporter server is as up to date as the sites. If the views on the sites are outdated, the Global Reporter server gathers this outdated information. If the management server is unable to contact the site, the Global Reporter server uses data from its last refresh.
- **Version of Sites** - The Global Reporter server can contact only sites running the same build. For example, if you upgrade one site, you must upgrade the rest.
- **Removing Sites** - To delete a site, click the corresponding  button for the site. Then, click the **Refresh Now** button near the bottom of the page. If you do not click the **Refresh Now** button and you attempt to create a new site with the same IP address as the deleted site, the management server tells you the site already exists. Also, when you delete a site, site names are not removed from the `tnsnames.ora` file. Names of deleted sites left in the `tnsnames.ora` file do not impact the user interface or the performance of the product.

See “Remote Sites Are Not Removed from the tnsnames.ora File” on page 254 for more information.

- **Editing Sites** - You can edit the port number and site name of a site, but you cannot edit its IP address.

To set up global reporting:

1. The Global Reporter server contacts sites through the IP address/DNS name and Oracle TNS listener port. Gather the IP address and port number of the Oracle TNS listener port that each site uses.
2. On the Global Reporter server, go to the Global Reporter tab by clicking **Configuration > Reports > Data Collection**.
3. Click the **Global Reporter** tab.
4. Click the **New Site** button.
5. Enter the information you gathered for a site in the first step.
 - **IP Address** - The IP address or DNS name of a server.
 - **Port number (optional)** - The Oracle TNS listener port the site uses. If this field is left blank, the management server assumes the database on the site uses port 1521.
 - **Site Name** - A name that includes the location of the site. For example, London1. Since data in the global reports is grouped by the site name, it is recommended you provide a unique site name to differentiate the sites.

Keep in mind the following:

- The Global Reporter server can contact only sites running the same build. For example, if you upgrade one site, you must upgrade the rest. After you upgrade your remote sites, click the **Refresh Now** button at the bottom of the page for Global Reporter.
 - The remote site is not required to have global reporting enabled in its license.
 - If you want data from the Global Reporter server included in global reports, add the local management server. Enter “localhost” as an IP Address/DNS name for your local management server.
6. Click **OK**.
 7. When you are done, click **OK**.

The management server verifies it can contact the site and it checks the version of the management server the site is running. The management server then adds the site to its list for global reporting.
 8. Repeat steps 4 and 5 for each site you want to add.
 9. Enable global reporting by selecting the **Enable** option on the Global Reporter tab.
-

10. Set the time you want the refresh to start by doing the following:
 - a. Click the  button. In the **Time** field, type a time in the 24-hour format with hours and minutes separated by a colon. For example, 23:15, for 11:15 p.m.
 - b. Select a date you want the job to start.
The date is selected.
 - c. Click the **Set** button.
If you click the Set button after the time has passed, you must reset the time.

The time and date you selected are displayed in the **Next Schedule Run** field.
11. Set up a repeat interval by typing a number in the **Repeat Interval** field and selecting a unit of measurement. For example, if you want the Global Reporter server to check the views of sites daily, you would type 1 in the **Repeat Interval** field and select **day(s)** as a unit.

Important: You must select a repeat interval that is 12 hours or more.

12. When you are done with your changes, click the **Save Changes** button.
Your changes are saved in case the browser is closed.
13. To obtain information immediately, click the **Refresh Now** button (**Configuration > Reports > Data Collection**).
14. When the Global Reporter server is done with updating its views, you can view the global reports. To view global reports, click **Reporter** (). Expand the Global node in the tree to view the reports.
You can filter the data viewed in the reports. See “Filtering Data in Global Reports” on page 569.

Editing Remote Site Information

To modify the information for your remote servers:

1. Click **Configuration > Reports > Data Collection**.
2. Click the **Global Reporter** tab.
3. Click  corresponding to the remote site you want to modify.
4. Modify the following information for your remote servers that are running the management server:

- ❑ **IP Address** - The IP address or DNS name of a server. If you change the IP address, you must modify the `listener.ora` file on the remote server, as described in the following step.
 - ❑ **Port number (optional)** - The Oracle TNS listener port the remote server uses. If this field is left blank, the management server assumes the database on the remote server uses port 1521.
 - ❑ **Site Name** - A name that includes the location of the site. For example, London1. Since data in the enterprise reports is grouped by the site name, it is recommended you provide a unique site name to differentiate the sites.
5. Click **OK** when done.
 6. If you changed the IP address of the remote server, you must edit the `listener.ora` file on that remote server as described in the following steps:
 - a. Log onto the remote site.
 - b. Stop the service for the management server running at the remote site.
 - c. Stop the listener service for Oracle (OracleOraHome92TNSListener) at the remote site.
 - d. Stop OracleServiceApplIQ at the remote site.
 - e. Open the following file in a text editor on the computer at the remote site:
Microsoft Windows:

```
%ORA_HOME%\network\admin\listener.ora
```


Sun Solaris:

```
$ORACLE_HOME/network/admin/listener.ora
```
 - f. Edit the IP address as shown in the following line.

```
(ADDRESS = (PROTOCOL = TCP) (HOST = 192.168.10.1) (PORT = 1521))
```
 - g. Save the file and exit.
 - h. Start the listener service for Oracle (OracleOraHome92TNSListener).
 - i. Start AppStorManager.

Remote Sites Are Not Removed from the tnsnames.ora File

Each time you add a site for global reporting, its contact information is added to the `tnsnames.ora` file. When you delete a site, site names are not removed from the `tnsnames.ora` file. Names of deleted sites left in the `tnsnames.ora` file do not impact the user interface or the performance of

the product. For example, assume you added the site “remotesiteA” for global reporting. It would appear in the `tnsnames.ora` file as the following:

```
remotesiteA =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP) (HOST = 192.168.1.2) (PORT = 1521))
    )
    (CONNECT_DATA =
      (SID = appiq)
    )
  )
```

Now, let's assume you removed “remotesiteA” from the user interface by clicking the corresponding



button for the site and then the **Refresh Now** button near the bottom of the Global Reporter page (**Configuration > Reports > Data Collection > Global Reporter**). Even though the remote site has been removed from the user interface, it still appears in the `tnsnames.ora` file. If you add “remotesiteA” as a remote site for global reporting again, another listing will be added to the `tnsnames.ora` file. These dual listings do not negatively impact the product.

Managing Performance Collection

This section describes the following:

- Managing Performance Collectors on page 255
 - Starting Performance Collectors on page 257
 - Stopping Performance Collectors on page 258
 - Setting the Date and Time for Performance Collectors on page 259
 - Viewing Data Aging Statistics for Performance on page 260
-

Managing Performance Collectors

The management server uses performance collectors to gather information for Capacity Explorer and Performance Explorer charts, as well as for monitoring. Use Data Collection page for Performance collectors to stop and start collectors as well as schedule when they run. Each row in the table corresponds to a collector. The Element Type column displays which element the collector gathers data for, and the Statistics column displays which statistics for that element the collector is tasked to gather.

For example, assume you see host_A in the Element Type column and Processor Utilization in the Statistics column. That collector gathers data about processor utilization for host_A.

You can have multiple collectors for an element. Let's go back to the example of host_A. Assume this host has another collector. Underneath the first row there is another row for the second collector. The Element Type is host_A, but this time the following statistics are listed: Physical Memory Used, Free Physical Memory, Virtual Memory Used, Free Virtual Memory. The second collector gathers information about all of those statistics listed for host_A.

You can see the result of these statistics in Capacity Explorer and Performance Explorer. When you select one of the following options from the Period combo-box, you can view the information these collectors gathered. Real time statistics are not affected by these collectors

- **Last Hour***
- **Last 24 Hours***
- **Last 7 Days**
- **Last Month**
- **Last Year**

*Only available in Performance Explorer

Important: All collectors are stopped during Get Details. This means that during Get Details data, such as for Performance Explorer, is not updated. Historical collectors, such as those available from the Configuration tab, are restarted when they are stopped during Get Details. Any charts that were active in Performance Explorer when Get Details was started are not restarted.

Note: It is not possible to run data collectors on quarantined elements. Attempting to run data collectors on quarantined elements will result in exceptions in the appiq.log file.

To manage performance collectors, click **Configuration > Performance > Data Collection**.

Table 7-8: About Performance Collectors

Column Heading	Description
Element	The name of the element from which this collector gathers information.
Element Type	The type of element from which the collector gathers information.
Statistics	The statistics that the collector is responsible for providing information.
Next Scheduled Run	Displays the date and time when the collector is scheduled to run.
Interval	The interval in minutes the collector runs.
Running?	Gives the status of the collector. Collectors that are running display a check mark in this column.

Table 7-8: About Performance Collectors (Continued)

Column Heading	Description
Edit	To edit the schedule for running a collector, click the Edit () button. Then set the date and time. See “Setting the Date and Time for Performance Collectors” on page 259 for more information. Important: If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the first day of the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.
Action	Displays one of the following buttons: ■ Stop - Stops the collector. The corresponding reports display information only gathered previously. See “Stopping Performance Collectors” on page 258. ■ Start - Starts the collector. When you start a collector, it begins gathering information for its corresponding reports. See “Starting Performance Collectors” on page 257.

Starting Performance Collectors

To start a collector:

1. Access the page for performance collectors page (**Configuration > Performance > Data Collection**).
2. Click the **Start** button corresponding to the collector you want to start.
3. Click the calendar icon, .
4. Set the date and time as described in the topic, “Setting the Date and Time for Performance Collectors” on page 259.
5. Type a repeat interval and then select a unit of measurement from the drop-down menu. The repeat interval determines how often the collectors gather the data.
6. If you are asked to provide a proxy host, do the following:
 - a. Select a proxy host by clicking the **Browse** button.

- b. Select a proxy host from the drop-down menu and then click **OK**.
 - c. Click **OK** again to set the time for starting the collector.
 - d. If you do not see any hosts displayed verify you have the latest CIM Extension version installed and running on a host that can access the Engenio storage system.
7. Click **OK**.
8. To start more than one collector at once, select more than one collector on the **Data Collection** tab and then click the **Start Collectors** button.

Stopping Performance Collectors

When you stop a collector, the management server stops gathering information for which the collector is responsible. For example, if a performance collector is not running, its corresponding statistics are no longer receiving information to display. One of the following occurs:

- If there was originally no information gathered for the statistic, no data appears for that statistic.
- If information was previously gathered for the statistic, old data appears for that statistic.

To stop a collector:

1. Access the page for performance collectors (**Configuration > Performance > Data Collection**).
2. Click the **Stop** button corresponding to the collector you want to stop.
The collector stops gathering information for its corresponding reports.

To stop more than one collector at once, select more than one collector and then click the **Stop Collectors** button.

Setting the Date and Time for Performance Collectors

To set the date and time, do the following:

1. Click **Configuration > Performance > Data Collection**.
2. Click the **Edit** button, .
3. Click the calendar icon, .

4. Enter the time in the time field. Make sure the time resembles a 24-hour clock, for example 22:00 for 10 p.m.
5. Click the date.
The date is highlighted in pink.

You can navigate the calendar as follows:



- Displays the same month in the previous year



- Displays the previous month



- Displays the next month



- Displays the same month in the following year

6. When you are done, click the **Set** button.
The management server displays the date and time.

Important: If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.



Figure 7-5: Selecting the Date and Time from the Calendar

In the previous figure, the calendar sets the date and time at 1:56 p.m. February 13, 2003.

Viewing Data Aging Statistics for Performance

Data Aging includes Data Rollup and Garbage Collection. Data Rollup controls how often a set of data is summarized. For example, hourly data is rolled into the daily table periodically. Garbage Collection refers to how long a set of data is preserved before it is permanently removed from the database.

The settings on this page control data aging for both reports and performance statistics.

Important: Do not modify the data on this page unless instructed by customer support. Changing them incorrectly can adversely affect the management server.

To view data aging statistics:

1. Click **Configuration > Performance > Data Collection**.
2. Click the **Data Aging** tab at the top of the screen.

Important: Perform the following steps **ONLY** if customer support has instructed you to modify one of the collectors on the page:

3. Click the **Edit** () button.
4. In the **Time** field, type the time in 24-hour format with the hour and minutes separate by a colon, for example, 22:15. Click the date on which you want to run the collector to run the next time. Today's date is highlighted in pink. Click **Set**.
The date and time appear in the **Next Scheduled Run** field in the yyyy-mm-dd hh:min format.

If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.

5. Click the **Enable** option.

Note: If you are not allowed to disable the collector, the Enable option is unavailable.

6. To change the repeat interval, In the **Repeat Interval** field, type an interval. Select one of the following units from the drop-down menu:
 - ☐ **Second(s)**
 - ☐ **Minute(s)**
 - ☐ **Hour(s)**
 - ☐ **Day(s)**
 - ☐ **Week(s)**
 7. (Garbage Collection only) To change how long the data is preserved, type an interval in the **Preserve** field, and then select one of the following from the drop-down menu to the right of the field:
 - ☐ **Second(s)**
 - ☐ **Minute(s)**
 - ☐ **Hour(s)**
 - ☐ **Day(s)**
 - ☐ **Week(s)**
 8. Click the **OK** button.
-

Editing the Locale and Currency Settings

The management server determines which languages and currency to display by looking at the language and currency settings for the operating system. You can override the management server's default locale and currency settings.

For example, assume the user interface for the management server is displayed in English, but you want to view it in Japanese. You could change the locale in the management server to Japanese without changing the locale setting of the computer running the management server.

When you change the locale and currency settings, the following occurs:

- The text in the product changes to the locale you set. For example, if you change the locale setting to French, the text in the product changes to French. Keep in mind the online help and PDFs stay in English regardless of the locale setting.
- The currency you selected is displayed in the product.

Note: The following in the online help looks at the localization of the operating system instead of the localization settings of the management server: Contents tab, Index tab, Search tab, Search text, output messages, text for images, and Related Topics heading. If you set the locale to Korean and your client computer is set to Japanese, the help components listed above appear in Japanese.

Table 7-9: Currency Settings

Abbreviation	Full Name
CNY	Chinese Yuan
EUR	Euro
GBP	Great Britain Pound
JPY	Japanese Yen
KRW	Korean Won
USD	U.S. Dollar

To change the locale and currency settings:

1. Click **Configuration > Locales**.
2. Choose a locale from the **Select Locale** drop-down menu.
3. Choose a currency from the **Select Currency** drop-down menu.
See Table 7-9, “Currency Settings,” on page 262 if you do not know the corresponding names for the currency abbreviations.
4. Click **Change Locale**.
5. Restart the management server.

Important: You must restart the management server for your changes to take effect.

Process Names

The following process names are displayed in the Windows Task Manager on the management server:

Table 7-10: Process Names on Windows

Process	Executable name
Application Server (JBoss)	storApplicationServer.exe
CIMOM for Default Discovery Group	storCimomDefault.exe
CIMOM for Discovery Group 1	storCimom1.exe
CIMOM for Discovery Group 2	storCimom2.exe

The executable names for the CIMOM's corresponding to the other discovery groups follow the same pattern as those described above.

The following process names are displayed in Solaris by running the `ps -ef` command:

Table 7-11: Process Names on Solaris

Process	Executable name
JBoss	storApplicationServer.sh
CIMOM for Default Discovery Group	storCimomDefault.sh
CIMOM for Discovery Group 1	storCimom1.sh
CIMOM for Discovery Group 2	storCimom2.sh

The executable names for the CIMOM's corresponding to the other discovery groups follow the same pattern as those described above.

Note: If you are using the prstat utility, all of the processes will be named `java.exe`.

Database Maintenance and Management

This chapter contains the following topics:

- Database Maintenance Window on page 267
- Overview of Backups on page 269
- Database Mode on page 271
- Architectural Overview of RMAN Backups on page 272
- Performing an RMAN Hot Backup on page 275
- Scheduling RMAN Hot Backups of the Database on page 275
- Viewing Results from RMAN Backup on page 277
- About the Database Admin Utility on page 277
- Database Password on page 290
- Troubleshooting Listener and Database Connection Problems on page 291

Database Maintenance Window

Schedule a maintenance window of two to four hours weekly during off peak hours of operation to do the following database maintenance operation:

1. Stop the service for the management server. See “Stopping the Service for the Management Server” on page 19.
2. Access the Database Admin Utility. See “Accessing the Database Admin Utility” on page 278 for more information.
3. Using the Database Admin Utility, verify the database is in an open state and the listener is running. See “Checking the Database and Listener Status” on page 279.
If the database state is not OPEN and shows error get the following logs and then contact technical support:
 - ☐ Log file for the Database Admin Utility (%MGR_DIST%/logs/dbAdmin.log)
 - ☐ Database alert log which can be found in one of the following locations:
 - **Windows:** \oracle\admin\APPIQ\bdump
 - **Solaris:** \$ORACLE_BASE\admin\APPIQ\bdump
4. Reset the temporary tablespace, as described in “Resetting the Temp and Undo Tablespace” on page 283.
5. Export the database, as described in “Exporting the Database” on page 280.
6. If the database is running in archive mode, set the database to no-archive mode, as described in “Setting the Archive Mode” on page 284. Then, clean the \oracle\oradata\APPIQ\archive on Microsoft Windows and \$ORACLE_HOME/oradata/APPIQ/archive on Sun Solaris.

Note: The archive directory (\oracle\oradata\APPIQ\archive on Microsoft Windows and \$ORACLE_HOME/oradata/APPIQ/archive on Sun Solaris) only exists if you have previously set the management server to archive mode.

7. Return the database to archive mode, as described in “Setting the Archive Mode” on page 284.
 8. Start the service for the management server.
 9. If database is in archive mode, take a current RMAN backup by clicking the **Backup Now** button on the **Configuration > Product Health > Disk Space > RMAN Backup**. See “Performing an RMAN Hot Backup” on page 275 and “Architectural Overview of RMAN Backups” on page 272. On successful completion of RMAN backup, the backup is saved to the following directory:
 - ☐ **On Microsoft Windows:**
%ORA_HOME%\rman\current
 - ☐ **On Sun Solaris:**
\$ORACLE_HOME/rman/current
-

10. Clean the following folders:

- ❑ **On Microsoft Windows:**
%ORA_HOME%\rman\backup1
%ORA_HOME%\rman\backup2
- ❑ **On Sun Solaris:**
\$ORACLE_HOME/rman/backup1
\$ORACLE_HOME/rman/backup2

Overview of Backups

The management server provides the following backups:

Important: Export and RMAN backups should be done regularly in combination.

Table 8-1: Description of Backups

Backup Type	Description	Files Backed Up	Database Mode
Export backup	Done through the Database Admin Utility. See “Exporting the Database” on page 280 for more information about exporting the database.	Database Schema, Oracle network Configuration files (tnsnames.ora, listener.ora), CIM repository, File Server SRM	Does not matter
RMAN HOT backup	The backup is referred to as being “hot” because the management server is still running while the backup is occurring. You can configure the RMAN backup to run by default. See “Scheduling RMAN Hot Backups of the Database” on page 275.	Database Files, Control files, Redo files, Archive Files, Oracle network Configuration files (tnsnames.ora, listener.ora), CIM Repository, File Server SRM	To do an RMAN hot backup, the management server must be set to archive mode. See “Setting the Archive Mode” on page 284.
RMAN Cold backup	Done through the Database Admin Utility.	Same files as an RMAN HOT backup.	If the management server is set to no-archive mode, users can perform an RMAN cold backup. See “Setting the Archive Mode” on page 284.

Backup Destination and Operation for RMAN: The management server has three backup points available for RMAN backups. RMAN backup destinations are:

- **On Microsoft Windows:**
 - %ORA_HOME%\rman\current
 - %ORA_HOME%\rman\backup1
 - %ORA_HOME%\rman\backup2
- **On Sun Solaris:**


```
$ORACLE_HOME/rman/current  
$ORACLE_HOME/rman/backup1  
$ORACLE_HOME/rman/backup2
```

The schedule backup writes in backup1 and backup2 folder in rotation and save now backup from the management server keeps overwriting in current folder. See “Architectural Overview of RMAN Backups” on page 272 for more information about RMAN backups.

Database Mode

To facilitate the HOT RMAN Backup, you must change the database mode to archive mode, as described in “Setting the Archive Mode” on page 284. The default database archive destination is `\oracle\oradata\APPIQ\archive` on Microsoft Windows and `$ORACLE_HOME/oradata/APPIQ/archive` on Sun Solaris. This destination can be modified as described in “Changing the Archive Destination” on page 287. Depending on the input/output of the data, archiving can be in the range of 2.5 GB to 10 GB per day. The archive folder gets cleaned during on a scheduled RMAN backup.

You can change the database mode to run in no archive mode with RMAN backup disabled, as described in “Setting the Archive Mode” on page 284. When you change the database to no archive mode, you reset the logs SCN, set the archive parameter in the database parameter file and disable the RMAN backup scheduler.

Important: Export the database after switching to no archive mode. See “Exporting the Database” on page 280 for more information.

Keep in mind the following implications if you do decide to change the database to no archive mode.

- If you set the management server to **Disable Database Archive Mode and RMAN Backup**, it is up to you to back up the management server manually. If you forget to back up your management server and your management server fails, you will not have a database to import. To learn more about manual backups, see “Running a Cold Backup” on page 286. Export the database after you change the database mode to no archive. See “Exporting the Database” on page 280.
 - Scheduled RMAN backup jobs do not run. If you change the database mode to no archive during an RMAN backup, the RMAN backup will error out.
-

- If the database fails as a result of a corrupt data file, the database can only be restored to the last export backup available. This requires recreating the database along with the import.

When you change the database to archive mode, you reset the logs SCN, set the archive parameter in the database parameter file and enable the RMAN backup scheduler. Take a current RMAN backup after switching to archive mode, as described in “Performing an RMAN Hot Backup” on page 275.

Architectural Overview of RMAN Backups

By default the management server does not backup the database automatically. If you enabled the database archive mode and RMAN backup as described in “Setting the Archive Mode” on page 284, the management server backs up the Oracle instance for the management server every three days and saves the backup for two weeks. The management server alternates the directories in which it saves the backup. When the management server first performs a scheduled backup, the backup is saved in `$ORACLE_HOME/rman/backup1` on Solaris and `%ORA_HOME%\rman\backup1` on Microsoft Windows. The next time the management server performs a scheduled backup, it is saved `$ORACLE_HOME/rman/backup2` on Solaris and `%ORA_HOME%\rman\backup2` on Microsoft Windows. The management server saves the backup in alternating directories (`backup1` and `backup2`), so you always have a copy of the last backup and the previous backup. To learn how to change the frequency of the scheduled backups, see “Scheduling RMAN Hot Backups of the Database” on page 275.

You can back up the database at any time by clicking the **Backup Now** button on the Database tab. When you back up the database using this technique, the backup is saved only in `$ORACLE_HOME/rman/current` on Solaris and `%ORA_HOME%\rman\current` on Microsoft Windows. To recover the database, contact customer support.

**Table 8-2: Backup Directories in
%ORA_HOME%\rman**

Directory	Contains
current	The backup when the Database Server Backup button was clicked.
log	A log of when the backup was done.
backup1	Information from the automatic backup (alternating day).
backup2	Information from the automatic backup (alternating day).

Let's assume you recently installed the management server and you have not done any backups. You have schedule the backups to take place every three days. You performed a backup, and it is stored in the `backup1` folder. The next scheduled backup occurs on day 4 and it is saved in the `backup2` folder. If your database fails, you can restore the database from day 1 or day 4. Let's assume you have a scheduled backup on day 7, which is saved to the `backup1` folder. This backup replaces the backup from day 1. The available backups are from day 4 and 7. Assume you click the **Backup Now** button on day 8, the backup is saved in the `CURRENT` folder because the backup is recording the current state of the database. If your database fails, you can restore the database from day 4, 7 or 8, as described in the following table:

Table 8-3: Backup Example

Day	Backup Type	Backup 1	Backup 2	Current	Available Backup
Day 1	Scheduled	Day 1 backup	-----	-----	Day 1 Backup
Day 4	Scheduled	-----	Day 4 backup	-----	Backups from Days 1 and 4
Day 7	Scheduled	Day 7 backup	-----	-----	Backups from Days 4 and 7

Table 8-3: Backup Example (Continued)

Day	Backup Type	Backup 1	Backup 2	Current	Available Backup
Day 8	Save Now	-----	-----	Day 8	Backups from Days 4, 7, and 8
Day 10	Scheduled	-----	Day 10 backup	-----	Backups from Days 7, 8, and 10

Keep in mind the following:

- Only one user at a time can back up the database.
- The management server archives files for the backup in a separate directory. Do not modify the files in this directory (`\oracle\oradata\APPIQ\archive` on Microsoft Windows and `$ORACLE_HOME/oradata/APPIQ/archive` on Sun Solaris) . For average database activity, the management server requires at least 100 GB of disk space for archive files. If there is higher database activity than average, more disk space may be required.

The following is saved during the backup:

- **Management server RMAN backup files** - These files contain information about the elements your management server monitors.
- **Oracle Network Configuration Files** -The configuration files are `tnsnames.ora` and `listener.ora`.
- **CIM Repository**
- **Property files, such as `jboss.properties`**
- **`rmanbackup.log` file**
- **`spfileappiq.ora` file**

To back up the database manually:

1. Before you can back up the database manually, you must enable database archive mode and RMAN backup as described in “Setting the Archive Mode” on page 284.
2. In the management server, click **Configuration > Product Health**.
3. Select **RMAN Backup** in the Product Health tree.
4. Click the **Backup Now** button.
The database is backed up.

Performing an RMAN Hot Backup

You can perform an RMAN hot backup instantly. The backup is referred to as being “hot” because the management server is still running. When you perform an RMAN hot backup, the following files are backed up: database files, control files, redo files, archive files, Oracle network configuration files (tnsnames.ora, listener.ora), CIM repository, and File Server SRM.

Note: The buttons on the RMAN backup page appear disabled if the database archive mode is disabled. See “Setting the Archive Mode” on page 284 for more information about changing the archive mode.

To perform an RMAN Hot Backup:

1. Verified you have enabled database archive mode and RMAN backup as described in “Setting the Archive Mode” on page 284.
2. Click **Configuration > Product Health**.
3. Select **RMAN Backup** in the Product Health tree.
4. Click **Backup Now**.
The management server performs an RMAN backup.

When the backup is complete, the Setup tab is refreshed with the status of the manual hot backup. The Results tab is also updated with the status of the RMAN hot backup. The Results tab displays the status of the previous RMAN hot backups (manual and scheduled).

Scheduling RMAN Hot Backups of the Database

The management server lets you schedule an RMAN hot back up of the database, as described in the following steps. The backup is referred to as being “hot” because the management server is still running.

Important: Before you can schedule the RMAN backup, you must enable database archive mode and RMAN backup as described in “Setting the Archive Mode” on page 284. The buttons on the RMAN backup page appear disabled if the database archive mode is disabled.

To learn more about backing up the database, see “About the Database Admin Utility” on page 277.

1. Verified you have enabled database archive mode and RMAN backup as described in “Setting the Archive Mode” on page 284.
2. Click **Configuration > Product Health**.
3. Select **RMAN Backup** in the Product Health tree.
4. Click the calendar icon, .
5. In the **Time** field, type the time in 24-hour format with the hour and minutes separate by a colon, for example, 22:15. Click the date on which you want to run the next backup of the database. Today's date is highlighted in pink. Click **Set**.
The date and time appear in the **Next Scheduled Run** field in the yyyy-mm-dd format.

If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.

6. In the **Repeat Interval** field, type an interval. Select one of the following units from the drop-down menu:
 - ☐ **Second(s)**
 - ☐ **Minute(s)**
 - ☐ **Hour(s)**
 - ☐ **Day(s)**
 - ☐ **Week(s)**

Note: The minimal interval you can schedule is one day. If you select **Hour(s)**, **Minute(s)** or **Second(s)**, you must enter an interval that equals more than a day. For example, if you select **Hour(s)**, you must enter 24 or more. Just as if you select **Minute(s)**, you must enter 1440.

7. Click the **Enable** option.
-

8. Click the **Save Schedule** button.

You can always disable the schedule by clearing the **Enable** option.

When the scheduled RMAN hot backup is complete, the Results tab is updated with the status of the RMAN hot backup. The Results tab displays the status of the previous RMAN hot backups (manual and scheduled).

Viewing Results from RMAN Backup

You can determine the results of the RMAN backups. The Results tab for RMAN backup provides the data and time of the RMAN backup, in addition to its status and backup folder.

To view the results from RMAN backup:

1. Click **Configuration > Product Health**.
2. Select **RMAN Backup** in the Product Health tree.
3. Click the **Results** tab in the RMAN Backup window.
The following information is displayed:

- ☐ Data/Time of the backup
- ☐ Status of the backup
- ☐ Backup folder

About the Database Admin Utility

The Database Admin Utility lets you easily manage your database from restoring it from a cold backup to resetting the temp tablespace. This tool provides flexible importing and exporting features, which let you save time. The Database Admin Utility requires you to stop the AppStorManager service so that it can manage the database optimally.

To learn more about the Database Admin Utility, see the following topics:

- Accessing the Database Admin Utility on page 278
- Refreshing the Database Admin Utility on page 279

- Checking the Database and Listener Status on page 279
- Changing the Password of System Accounts on page 279
- Exporting the Database on page 280
- Importing the Database on page 281
- Re-initializing the Database on page 282
- Resetting the Temp and Undo Tablespace on page 283
- Restarting the Database on page 283
- Restoring a Cold Backup on page 284
- Setting the Archive Mode on page 284
- RMAN Restore on page 285
- Running a Cold Backup on page 286
- Changing the Archive Destination on page 287
- Viewing Logs from the Database Admin Utility on page 288
- Changing the Oracle Listener Password on page 288
- Resetting the Admin Password for the Management Server on page 289
- Resetting the Log on page 289
- Warning Messages During Reinitializing the Database on page 290

Accessing the Database Admin Utility

To access the Database Admin Utility:

1. Stop the AppStorManager service.
2. To access the database utility on Sun Solaris, do the following on the management server:
 - a. Set the display if you are accessing the Database Admin Utility remotely.
The Database Admin Utility uses Perl. To set Perl in your path, enter the following command at the command prompt:

```
eval `/opt/productname/install/usersvars.sh`
```


where `/opt/productname` is the directory containing the software. It is defined by `$APPIQ_DIST`.

Important: You must include the back quotes around the full path to `usersvars.sh` in the command.
 - b. Go to the `$APPIQ_DIST/Tools/dbAdmin` directory and then enter the following at the command prompt:


```
perl dbAdmin.pl
```

The full path to Perl is the following:

\$APPIQ_DIST/JBossandJetty/server/appiq/remoteScripts/perl/bin/perl

3. To access the database utility on Windows, go to the %MGR_DIST%\Tools\dbAdmin directory on the management server and double-click dbAdmin.bat.

Refreshing the Database Admin Utility

The Database Admin Utility requires the following are not in use:

- JBoss
- CIMOM

If any of these components are in use, the Database Admin Utility will not work. If you are shown an error message when you start the Database Admin Utility, stop the AppStorManager service. Then, click the **Refresh** button.

Checking the Database and Listener Status

To find the database and listener status:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 278.
2. Click **Check Database Status** in the left pane.
3. Type the SYS password in the **Password for User SYS** field.
4. Click **Check Database Status**.
The database and listener status is shown in the lower pane.
5. To clear the pane, click **Clear All**.

Changing the Password of System Accounts

The management server uses the following accounts to access and manage the database for the management server. You should change the passwords to these accounts to prevent unauthorized access.

- **SYS** - Used for the management server database creation and upgrade. Default password: `change_on_install`
- **SYSTEM** - Used for management server database creation and upgrade, in addition to database import, export and re-initialization. Default password: `manager`
- **RMAN_USER** - Used for RMAN backup and restore. This user has sys privilege. Default password: `backup`
- **DB_SYSTEM_USER** - Used for all the database activity, including establishing a connection to the management server database. Default password: `password`

You must change the passwords of the SYS, SYSTEM, RMAN_USER, and DB_SYSTEM_USER accounts by using the Database Admin Utility, so the management server is aware of the changes. Do not change the password for one of these accounts by using Oracle. Make sure you keep the new passwords in a safe location, as it is your responsibility to remember the Oracle passwords.

The management server requires the password to have the following characteristics:

- a minimum of three characters
- starts with a letter
- contains only letters, numbers and underscores (`_`)
- does not start or end with an underscore (`_`)

To change the password of a system account:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 278.
2. Click **Change Passwords** in the left pane.
3. Select an account name from the **User Name** combo box.
4. Type the current password in the **Old Password** field.
5. Type the new password in the **New Password** field.
6. Retype the password in the **Confirm Password** field.
7. Click **Change**.
The Database Admin Utility changes the password for the specified account.

Exporting the Database

Use the Database Admin Utility to save the database in a format that can be imported.

Important: The export database feature should not be used for backing up the server. If you are interested in backing up the database for disaster recovery, use the Database Server Backup feature. See “About the Database Admin Utility” on page 277 for more information.

To export the database:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 278.
2. Click **Export Database** in the left pane.
3. Click the **Browse** button to select a file path. Then, type a file name in the **File name** field. Then, click the **Open** button.

Important: Select a directory outside of the directory tree of the management server. This way if you remove the management server, you will not lose the saved database.

The file name and its path is displayed in the Database Admin Utility. The .zip file extension is automatically added to the file name.

4. Type the password of the SYSTEM account. The default password of the SYSTEM account is the following: `manager`
You are notified when the database process is complete.
5. Select **Clear Report Cache** if you do not want the report cache to be included with the database you are exporting. When a user imports this database, their report cache will be empty until the report cache is refreshed (**Configuration > Reports > Report Cache**). This option may save you time with exporting the database if your database includes a large amount of data for reports.
6. Click **Export Database**.

Importing the Database

You can revert to an earlier configuration by uploading a file (*.zip) containing the database information.

The software stores a snapshot of the data in its database. Since this file is a snapshot of the network at a certain time, it may not contain your most current network configuration. If you want to view an up-to-date network configuration and the latest information about the elements, perform discovery and Get Details.

To import a database:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 278.
2. Click **Import Database** in the left pane.
3. Click the **Browse** button. Select a database file with a .zip extension to import. Then, click the **Open** button.
The file name is displayed in the Database Admin Utility.
4. Type the password of the SYSTEM account. Then, click **OK**. The default password of the SYSTEM account is the following: `manager`
You are notified when the database process is complete.
5. Select **Populate Report Cache** if you want to refresh the reports cache during the import. Keep in mind that the amount of time to import the database may increase if you select this option and the database has a large amount of data for reports. Alternatively you can refresh the report cache from the management server (**Configuration > Reports > Report Cache**).
6. Click the **Import Database** button.

Re-initializing the Database

Keep in mind the following:

- Re-initializing the database removes everything from the database. This is not recommended unless you are sure about what you are doing. It is strongly suggested you export the database before you re-initialize it. See “Exporting the Database” on page 280 for more information on how to save the database.
- When you re-initialize the database, all users are logged out of the management server.
- Ignore the warning messages in the command prompt window that pops up when the `dbAdmin.bat` file is double-clicked. See “Warning Messages During Reinitializing the Database” on page 290 for more information.

To clear the database:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 278.
2. Click **Re-initialize Database** in the left pane.
3. Type the password of the SYSTEM account. Then, click **OK**. The default password of the SYSTEM account is the following: `manager`
You are notified when the re-initialization is complete.
4. Click the **Re-initialize Database** button.

Resetting the Temp and Undo Tablespace

The temporary and undo tablespace may grow large due to high database activity. You should regularly reset the temp and undo tablespace to its initial value, as described in this section.

To reset the tablespace:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 278.
2. Click **Reset Temp Tablespace** in the left pane.
3. Type the SYS password in the **SYS Password** field.
4. Click the **Reset Temp Tablespace** button.

Restarting the Database

You may sometimes need to restart the ApplQ instance of the database. Use this feature in the Database Admin Utility when the database is down or you need to shutdown and then start the database.

To restart the database:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 278.
2. Click **Restart Database Server** in the left pane.
3. Type the SYS password in the **SYS Password** field.
4. Click the **Restart Database Server** button.

Restoring a Cold Backup

If you performed an RMAN cold backup, follow the steps in this section to restore the RMAN cold backup. You can only perform a cold backup if you selected **Disable Database Archive Mode and RMAN Backup**. See “Setting the Archive Mode” on page 284 for information about changing the archive mode. The backup is referred to as being “cold” because the management server is not running while the backup is occurring.

To restore a cold backup:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 278.
2. Click **Restore Cold Backup** in the left pane.
3. Click the **Browse** button.
4. In the **File Name** field, provide a directory path containing the cold backup (may automatically be populated in some Web browsers), but do not provide a file name. In this release the cold backup is saved in the COLDBACKUP directory in the path specified by the person who did the cold backup. The RMAN backup consists of several files. The Database Admin Utility automatically notices the backup files in the directory provided.
5. Type the password of the SYS account.
6. Click **Run Cold Backup**.

Setting the Archive Mode

By default the management server database runs in no-archive mode, which requires you to backup the database manually by using a cold RMAN backup. A cold RMAN backup is an RMAN backup without the management server running. If you want to take an RMAN hot backup of the database, switch the database to archive mode. An RMAN hot backup is with the database running while the backup is occurring.

Keep in mind the following implications if you decide to leave the database in no archive mode.

- If you leave the database in no archive mode (**Disable Database Archive Mode and RMAN Backup**), it is up to you to back up the management server manually. If you forget to back up your management server and your management server fails, you will not have a database to import. To learn more about manual backups, see “Running a Cold Backup” on page 286. Export the database after you change the database mode to no archive. See “Exporting the Database” on page 280.
-

- Scheduled RMAN backup jobs do not run. If you change the database mode to no archive during an RMAN backup, the RMAN backup will error out.
- If the database fails as a result of a corrupt data file, the database can only be restored to the last export backup available. This requires recreating the database along with the import.

To change the archive mode:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 278.
2. Click **Archive Mode Switch** in the left pane.
3. Select one of the following:
 - **Enable Database Archive Mode and RMAN Backup** - Select this option if you plan to run database in archive mode and automated backup while the management server is running. If database is already in archive mode, you can use this option to clean up archive files. See “About the Database Admin Utility” on page 277 for more information about the automatic backups. When you change the database to archive mode, you reset the logs SCN, set the archive parameter in the database parameter file and enable the RMAN backup scheduler. Take a current RMAN backup after switching to archive mode (**Configuration > Product Health > Disk Space > RMAN Backup > Backup Now**) as described in “Performing an RMAN Hot Backup” on page 275.
 - **Disable Database Archive Mode and RMAN Backup** - Select this option if and only if you always shut down the management server prior to a backup. With database running in no-archive mode, you can only cold back up your database. When you change the database to no archive mode, you reset the logs SCN, set the archive parameter in the database parameter file and disable the RMAN backup scheduler. See “Running a Cold Backup” on page 286 for more information. Export the database after switching to no archive mode. See “Exporting the Database” on page 280.
4. Type the SYS password in the **SYS Password** field.
5. Click the **Change Settings** button.

RMAN Restore

The Database Admin Utility lets you restore the management server database from a previously scheduled (hot) RMAN backup, which are stored in the following directories by default: backup1, backup2, and current. You may not have a previously scheduled (hot) RMAN backup if you selected **Disable Database Archive Mode and RMAN Backup**.

To restore the database:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 278.
2. Click **RMAN Restore** in the left pane.
3. Select one of the following directories to restore:
 - ☐ **Current** - The last restore from when the **Database Server Backup** button on the Database tab was clicked.
 - ☐ **Backup1***
 - ☐ **Backup2***

*Information from the automatic backup (alternating weeks). See “About the Database Admin Utility” on page 277 for more information about the automatic backup.
4. Type the SYS password in the **SYS Password** field.
5. Click **Restore RMAN**.

The Database Admin Utility restores the selected database.

Running a Cold Backup

If you are running the database in no-archival mode, you should perform a cold backup frequently. The backup is referred to as being “cold” because the management server is not running while the backup is occurring.

Note: This feature does not run if the **Enable Database Archive Mode and RMAN Backup** option is selected. The **Enable Database Archive Mode and RMAN Backup** option performs an RMAN backup periodically. See “About the Database Admin Utility” on page 277 and “Setting the Archive Mode” on page 284 for more information.

The following is saved during a cold RMAN backup:

- **Management server RMAN backup files** - These files contain information about the elements your management server monitors.
 - **Oracle Network Configuration Files** -The configuration files are `tnsnames.ora` and `listener.ora`.
 - **CIM Repository**
 - **Property files, such as `jboss.properties`**
-

- **rmanbackup.log** file
- **spfileappiq.ora** file

To run a cold backup:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 278.
2. Click **Run Cold Backup** in the left pane.
3. Click the **Browse** button to select a file path.
4. In the **File Name** field, provide a directory path (may automatically be populated in some Web browsers), but do not provide a file name.

Important: The management server saves the backup in a directory called COLDBACKUP in the path you specified. Any pre-existing content in this directory, such as previous cold backups, is removed.

5. Type the password of the SYS account.
6. Click **Run Cold Backup**.

Changing the Archive Destination

The current archive directory by default is `\oracle\oradata\APPIQ\archive` on Microsoft Windows and `$ORACLE_HOME/oradata/APPIQ/archive` on Sun Solaris. Over time your database will grow. If you feel you are running out of space, you can add a new volume and change the archive destination to a new volume, as described in the following steps:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 278.
 2. Click **Set Archive Destination** in the left pane.
 3. Click the **Browse** button to select a file path.
 4. Type the password of the SYS account.
 5. Click **Set Archive Destination**.
-

Viewing Logs from the Database Admin Utility

You can determine the status of a process executed by the Database Admin Utility by looking at its logs. To view the log from the Database Admin Utility, click **View Log** in the Database Admin Utility.

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 278.
2. Click **View Log**.
A separate window appears displaying the logs.
3. You can refresh the logs by clicking the **Refresh** button in the separate window displaying the logs. To clear this window, **click Reset Log** in the Database Admin Utility window.

Changing the Oracle Listener Password

By default, Oracle provides a blank Oracle Listener Password. To change the Oracle listener password:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 278.
2. Click **Set Oracle Listener Password** in the left pane.
3. Type the new Oracle listener password in the **New Listener Password** field in the right pane.
4. Type the new Oracle listener password again in the **Confirm New Listener Password** field.
If you want to clear the **New Listener Password** and **Confirm New Listener Password** fields, click the **Clear All** button.

Important: Only the following are allowed in the password: letters (a-z, A-Z), numbers(0-9) and underscores(_).

5. Click the **Set Oracle Listener Password** button.
The Oracle Listener Password is changed.
-

Resetting the Admin Password for the Management Server

If you no longer know the password for the admin user account and you need the admin user account to access the management server, you can reset its password to the default, which is **password**.

Important: Once you have reset the password, you should change the admin password to prevent unauthorized access. The admin user account provides full access to the management server. See “Changing Your Password” on page 165 for more information.

To reset the admin password:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 278.
2. Click **Reset Product Admin Password** in the left pane.
3. Type the password for the SYS user account in the **Password for User SYS** field. The default password for the SYS user account is `change_on_install`.

Note: If you want to clear the **Password for User SYS** field, click the **Clear All** button.

4. Click the **Reset Product Admin Password** button. The admin password is reset to **password**.

Resetting the Log

When you click the **View Logs** button, a separate window appears displaying the logs. You can clear this separate window by clicking **Reset Log**. When you are asked if you want to reset the log, click **Yes**.

Warning Messages During Reinitializing the Database

When you use `dbAdmin.bat` to re-initialize the database, warning messages resembling the following appear in the command prompt window that pops up when the `dbAdmin.bat` file is double-clicked. You can ignore these messages.

```
Connected.  
  
Creating FSRM DATA tablespace  
Creating FSRM INDX tablespace  
Connected.  
  
Warning: View created with compilation errors.  
Warning: View created with compilation errors.  
Warning: View created with compilation errors.
```

Database Password

The management server uses the following accounts to access and manage the database for the management server. You should change the passwords to these accounts to prevent unauthorized access.

Table 8-4: Default Passwords for Database Accounts

Account Name	Default Password
SYS (database admin user)	change_on_install
SYSTEM (database admin user)	manager
APPIQ_SYSTEM (management server admin user)	password

You should change the default password after the initial installation to prevent a security breach. See “Changing the Password of System Accounts” on page 279.

Important: Do not change any Database Password manually or from the command line. If a password is changed manually, DB-Admin tool will not work and export, import, rman backup will start failing. See “Changing the Password of System Accounts” on page 279 for information about changing database passwords.

How to change the Database Password for APPIQ_USER account on Managed Database Application:

To discover and manage the database application from management server, APPIQ_USER (GLOBAL) account is used. On all the managed applications the APPIQ_USER account with the default password of “password” was created when you configured the management server before you discovered the database application.

You can change the password for the APPIQ_USER account but this should be done for all database applications. See “Changing the Password for the Managed Database Account” on page 149.

For example: If the management server is managing two Oracle and two Sybase database applications, ask the database administrator to change the password for APPIQ_USER on all the Oracle and Sybase managed databases to a single password. Then, change the password for APPIQ_USER on the management server as described in “Changing the Password for the Managed Database Account” on page 149.

Troubleshooting Listener and Database Connection Problems

If there is any connection problem from management server or you see a JBoss connection problem in appiq.log, first verify that the listener and database are running:

Checking Listener Status:

If you are not able to start the listener contact technical support with the error message and network files (tnsnames.ora, listener.ora) from the following directory:

Microsoft Windows:

%ORA_HOME%\network\admin\listener.ora

Sun Solaris:

\$ORACLE_HOME/network/admin/listener.ora

■ On Windows:

- From the command line:

lsnrctl status

This will show you the status of listener. If the listener is not up try to start as:

lsnrctl start

- From the Database Admin Utility. See “Checking the Database and Listener Status” on page 279.

■ On Solaris: From command line su – oracle and as a Oracle user

- From the command line:

lsnrctl status

This will show you the status of listener. If the listener is not up try to start as:

lsnrctl start

- From the Database Admin Utility. See “Checking the Database and Listener Status” on page 279.

Checking Database Status:

These steps restart the database. If you receive an error message, contact technical support.

■ On Windows:

- From the command line:

Sqlplus /nolog

Sql>connect sys/change_on_install@appiq as sysdba

Sql> startup force;

- From the Database Admin Utility. See “Checking the Database and Listener Status” on page 279.

These steps restart the database. If you receive an error message, contact technical support.

- **On Solaris:** From command line su – oracle and as a Oracle user
 - From the command line:
Sqlplus /nolog
Sql>connect sys/change_on_install@appiq as sysdba
Sql> startup force;
 - From the Database Admin Utility. See “Checking the Database and Listener Status” on page 279.

Viewing Element Topology and Properties

This chapter describes the following:

- About System Explorer on page 296
 - Accessing System Explorer on page 298
 - About Cisco Switches and VSANs in System Explorer on page 299
 - About the User Interface on page 299
 - Viewing Storage Elements on page 324
 - Setting Up Custom Commands on page 353
 - Using External Tools on page 366
 - About the Navigation Tab on page 367
 - Viewing Element Properties on page 374
 - Viewing Element Topology on page 377
 - Creating a Virtual Application on page 390
 - About the Provisioning Tab on page 391
 - About the Events Tab on page 391
 - Asset Attributes of an Element on page 392
 - About the Collectors Tab on page 394
 - About the Monitoring Tab on page 395
 - About the Policies Tab on page 396
-

- Determining If a Host Belongs to a File System on page 396
- About the Data from CXFS File Systems on page 397

About System Explorer

System Explorer is the gateway to many features that let you view details about the discovered elements. System Explorer provides a topology that lets you view how the devices in your network are connected. For example, direct attached storage connections are displayed by a dotted line.

Note: To view direct attach storage, you must enable the  button. See Table 9-1, “Description of Toolbar in System Explorer,” on page 301 for more information.

Use the utilities provided in the toolbar to modify the topology. For example, you can filter out fabrics and change the placement of elements in the topology through drag and drop functionality. See “About the Toolbar in System Explorer” on page 300 for more information.

The following tabs provided in the middle pane provide additional information:

- **List** - Provides information about the elements by fabric and domain. See “About the List Tab” on page 304 for more information.
- **Access** - Provides information about zone entries, persistent bindings, and storage system LUN masking. You can also manage zone, zone aliases, and zone sets from this tab. See “About the Access Tab” on page 307 for more information.
- **Path** - Provides information about an element's path. See “About the Path Tab” on page 315 for more information.

When you right-click an element in the topology or in the List, Access, or Path tab, a drop-down menu appears. This menu provides additional functionality depending on the type of element clicked, such as telnet and the creation of zone sets. See the topics, “About the Right-Click Menu Options” on page 317.

If a switch has more than one connection to a host or storage system, the number of connections is displayed on the line connecting the elements. If there is only one connection, no number is displayed since the line indicates a connection exists.

Keep in mind the following:

- If your Java plug-in control panel cache setting is set to 50 MB, it is recommended you increase this setting to 100 MB or more. Increasing this setting improves the reloading performance of System Explorer.
- Individual virtual SANs (VSANs) for Cisco switches are not displayed in the topology or fabric tree. The switches in a VSAN are displayed under the fabric to which their VSAN belongs. See “About Cisco Switches and VSANs in System Explorer” on page 299 for more information.
- ISL connections are not shown as connected ports between two McDATA or Connectrix switches that are not both managed by EFC Manager or Connectrix Manager in a fabric.
- The user interface may load slowly while the topology is being recalculated.
- NAS stitching to Windows XP hosts does not appear in the topology.
- If multiple management servers are using the same Windows Proxy, hosts discovered by other management servers appear in the topology. The Windows Proxy tracks all Windows hosts discovered. If you have more than one management server using the same Windows Proxy, the management servers are aware of all hosts discovered by the Windows Proxy. For example, assume you have two management servers that are named A and B and are using the same Windows Proxy. You use management server A to discover host1, host2, and host3. You then use management server B to discover only host3. You will see host1, host2, and host3 in the topology for management server B, although host3 only appears in the lists for discovery, topology, and Get Details.
- The following IBM HBA appears as QLogic HBAs in the Navigation and Properties pages, in addition to reports:
 - IBM MSJ
 - FaStT FC-2/2-133

By double-clicking an element in the topology you have access to the following features:

- **Navigation** - The Navigation tab provides information about an element and how it relates to other elements in its path. See “About the Navigation Tab” on page 367 for more information.
- **Properties** - The Properties tab provides a detailed status of the element. See “Viewing Element Properties” on page 374 for more information.
- **Topology** - The Topology tab provides a graphical representation of an element's path. It displays additional information not found in System Explorer, such as adapters, slots, and fibre channel ports. See “Viewing Element Topology” on page 377 for more information.

- **Chargeback** - The Asset Management tab lets you keep track of asset attributes, such as contact information for the element's owner. See "Asset Attributes of an Element" on page 392 for more information.
- **Collectors** - The **Collectors** tab lets you start a collector for a report and view the collector's corresponding reports once the information has been gathered. See "About the Collectors Tab" on page 394 for more information.
- **Provisioning** - The Provisioning tab lets you manage zones, zone sets, and zone aliases, in addition to pools, volumes, LUNs and LUN mappings. See the topics, "About the Provisioning Tab" on page 391 and "About Provisioning" on page 399 for more information.
- **Events** - The Events tab lets you view events for an element. See the topics, "About Event Manager" on page 536 and "About the Events Tab" on page 391 for more information.
- **Policies** - The Policies tab lets you create utilization policies, which can send an e-mail, generate an event, or run a custom script when a set threshold for an element is triggered. See "About Policy Manager" on page 682 for more information.

Grey Screen When Attempting to Access System Explorer

Errors may occur if the client computer you use to access the management server has software that blocks JavaScript and/or pop-ups. You may be shown a grey screen when attempting to access System Explorer. Other errors include not being able to get past the login screen, view topology, or perform many other functions. Set your blocking software appropriately to allow the user interface to function properly.

Accessing System Explorer

To access System Explorer, click **System Explorer** ()

Keep in mind the following:

- ☐ If you are unable to access System Explorer, make sure your Web browser is set to allow JavaScript and cookies.
 - ☐ Java 2 Runtime Environment is required to access several features in the management server, such as System Explorer. If you are accessing the management
-

server and you do not have the Java 2 Runtime environment, you are asked to install it if your Web browser is on Windows. If your Web browser is on a Solaris system, you must manually install the Java plug-in. See “Installing the Java Plug-in” on page 16 for more information.

- When you are asked if you want to trust the signed applet for the software, click **Always**. The **Always** option prevents this message from being displayed every time you access System Explorer.

About Cisco Switches and VSANs in System Explorer

The management server does not display individual VSANs in its topology or fabric tree. The switches in a VSAN are displayed under the fabric to which their VSAN belongs. For example, assume switch_A belongs to VSAN1 and switch_B belongs to VSAN_2. VSAN_1 and VSAN_2 belong to the same fabric. The management server displays switch_A and switch_B under the same fabric without their VSAN listed in the tree.

To determine the VSAN to which a port on a Cisco switch belongs, access the Properties page for the port. The Properties page can be accessed by double-clicking the switch in System Explorer and then clicking the **Properties** tab. Then, click the hyperlink for the port in the Properties page for the switch. You can also view information about Cisco port types, such as TE ports by accessing the Properties page for a port.

Note: If a TE port belongs to multiple VSANs, the management server mentions only the primary VSAN.

About the User Interface

This section describes the following:

- About the User Interface for System Explorer on page 300
- About the Toolbar in System Explorer on page 300
- Description of the Icons Displayed in the Topology on page 304
- About the List Tab on page 304
- About the Access Tab on page 307
- About the Path Tab on page 315
- About the Right-Click Menu Options on page 317

About the User Interface for System Explorer

System Explorer displays an easy to use interface, which provides the following:

- **Toolbar** - Provides buttons and menus to help you modify the topology in System Explorer. See “About the Toolbar in System Explorer” on page 300.
- **Tabs** - Provides information about individual elements. The following tabs are provided:
 - **List** - Provides information about the elements by type and by fabric and domain. See “About the List Tab” on page 304.
 - **Access** - Provides access to tools that let you provision and view information about provisioning. See “About the Access Tab” on page 307. Provisioning may not be available in your kit. To determine if you can access the Provisioning feature, access the Feature List, which is accessible from the Documentation Center.
 - **Path** - Provides information about an element's path. See “About the Path Tab” on page 315.
- **Right-Click Menu** - Provides features you can use to manage that element. See “About the Right-Click Menu Options” on page 317.

About the Toolbar in System Explorer

The following table provides a brief description of the buttons and menus in System Explorer toolbar.

Table 9-1: Description of Toolbar in System Explorer



Icon	Description
	Prints the topology. See “Printing Elements in System Explorer” on page 330 for more information.
	Exports the topology to an XML file that can be viewed in Microsoft Visio. See “Exporting the Topology to Visio” on page 332 for more information.
	Magnifies the view
	Decreases the magnification
	Lets you set the magnification to a percentage of the default magnification
	Opens a smaller pane, which provides a global view of the topology. This lets you position the main view to a certain section of the topology. For more information, see “Using the Global View” on page 330.
	Fits the topology to the window, so you can see the entire topology.
	Lets you move an element in the topology. See “Arranging Elements in the Topology” on page 327.
	Lets you move multiple elements at once. This button is not accessible from the Topology tab. See “Arranging Elements in the Topology” on page 327.
	Lets you move the entire topology at once. Click the Pan button () and then click any place in the topology. Drag the mouse to a new location.
	Updates the layout of the topology and removes the last saved layout from the database. Elements that have been manually moved might revert to their original position. This button is not accessible from the Topology tab.

Table 9-1: Description of Toolbar in System Explorer (Continued)

Icon	Description
	Saves the current topology, so that when you return to System Explorer, the saved layout is restored. This option can be especially useful if you have moved elements in the topology and you want to keep their current location. This button is not accessible from the Topology tab. When you click the button, you are asked if you want the layout to apply to all users. ■ Yes - All users who log into the management server see the topology you created. Only users with system configuration capability can save their layout for all other users ■ No - Other users cannot view the topology you saved. The saved topology appears the next time you log into the management server.
	Opens a new window containing the topology. This feature lets you view different domains of the topology at once. This button is not accessible from the Topology tab. See “About the New Window Option” on page 386 for more information.
	Lets you view only selected fabrics in the topology. This button may not be accessible from the Topology tab. See “Filtering Fabrics” on page 340 for more information.
	Used in Capacity Explorer to hide the lower pane.
	Lets you find an element by name or by Worldwide Name in the topology. You can type part of the information, and the management server highlights the elements that match. After you populate the search field, click the  button or press ENTER. To expand the Search field, close the left pane. See “Opening and Closing the Left Pane” on page 9 for more information.

Table 9-1: Description of Toolbar in System Explorer (Continued)

Icon	Description
	Displays the event severity icons for the elements displayed in the topology. See “Viewing Event Status in the Topology” on page 341 for more information. Disabled for Performance Explorer and Capacity Explorer.
	This button calculates the topology paths. Clicking this button also allows you to view direct attached storage in System Explorer. Direct attached storage is indicated by dotted lines. If any of the paths are not fully calculated, a pop-up dialog box displays a list of all the hosts with partially calculated paths.

Description of the Icons Displayed in the Topology

The following icons are displayed in the System Explorer topology.

Table 9-2: Elements Displayed in the Topology

Icon	Description
	Indicates an application.
	Indicates a file server.
	Indicates a host. This particular icon is for a host running Microsoft Windows.

Table 9-2: Elements Displayed in the Topology (Continued)

Icon	Description
	Indicates a storage system.
	Indicates a switch. This particular icon is for a Brocade switch.
	Indicates a filer. This particular icon indicates a NAS filer.
	Indicates a tape library.

About the List Tab

The List tab provides information about the elements by type and by fabric and domain.

To find the fabrics in a domain, expand the domain node. You can see the elements in each fabric by expanding the fabric node, as shown in the following figure.




Figure 9-1: Expanding the Fabric Node

The “unknown” Fabric lists elements that have a Fibre Channel port connected to an undiscovered Fabric or the Fibre Channel port remains not connected.

When you click a fabric name in the tree, its members are highlighted in the right pane, as shown in the following figure.




Figure 9-2: Fabric Members Are Highlighted

When you right-click an element in the List tab, a drop-down menu is displayed. The options displayed depend on the type of element you clicked. See About the Right-Click Menu Options on page 317 for an explanation of the options in the drop-down menu.

Viewing Elements by Type

You can view elements by type under the **All Elements** node on the **List** tab. This is especially helpful in determining how many elements you have of a specified type, such as the number of storage systems.

When you expand the element type node, all elements of that type are listed. If you select the element type node, all elements of that type are selected in the topology. For example, assume you want to determine the number of applications that the management server monitors. When you expand the **Applications** node, the applications are listed. When you select the **Applications** node, the applications are highlighted in the topology, as shown in the following figure.




Figure 9-3: Highlighting the Applications in the Topology

If you select an element in the left pane, the element is highlighted in the topology. You also have access to additional functionality by right clicking the element. See About the Right-Click Menu Options on page 317 for more information.

About the Access Tab

The Access tab provides information about the following:

- Zone Entries
 - Host Bindings
-

- Storage System LUN Masking
- You can also manage zone, zone aliases, and zone sets from this tab by right-clicking an element. See “About the Right-Click Menu Options” on page 317.

Obtaining Information About Zone Entries

To view the zone entries in a domain, expand **domain** > **Fabric name** > **Zone Set** nodes. Click the Zone Set node to see the members of the zone set highlighted in the right pane, as shown in the following figure.




Figure 9-4: Members of a Zone Set

The  icon is displayed next to the name of the active zone set. The  icon is displayed next to the inactive zone sets.

To view members of a zone, do one or more of the following:

- Expand the node of the zone in the tree. The software displays the zone members underneath the node of the zone.
- Click the node of the zone in the tree. The software highlights the zone members in the right pane.




Figure 9-5: Displaying Zone Members

To view the relationship of the zone member to the switch, click the zone member in the tree. The software highlights the zone member and its switch.




Figure 9-6: Zone Member to Switch

To view information about a zone member's port, expand the zone member node, as shown in the following figure. Notice that when you select the zone member node in the tree, it appears highlighted in the right pane.




Figure 9-7: Obtaining Information About a Zone Member's Adapter

Obtaining Information About Host Bindings

To view the elements that have host bindings, click the Host Bindings node in the tree. The software highlights the elements that have host bindings in the right pane, as shown in the following figure.




Figure 9-8: Highlighting Elements with Host Bindings

To view a host binding, click the HBA node in the tree. The HBA node is under the element node, as shown in the following figure.

When you click the HBA node, the host and the element to which it has the binding are highlighted. A green line between the two elements indicates they have a binding.




Figure 9-9: Displaying Host Bindings

To view information about the ports on an HBA card, expand the HBA node in the tree, as shown in the following figure:



Figure 9-10: HBA Port Properties

Obtaining Information About Storage System LUN Masking

To obtain information about a storage system LUN masking, expand the fabric node. Then, expand the storage system node and click the Fibre Channel port. The values of the WWNs are displayed under the node, and the storage system is highlighted in the right pane, as shown in the following figure.




Figure 9-11: WWN Properties

The software displays properties of the WWN. If the LUN has a LUN masking, expand the LUN node to obtain information about the LUN masking.




Figure 9-12: WWN Properties

To view a LUN masking, expand a LUN node.

About the Path Tab

The Path tab provides information about an element's path. By clicking a host's node, you can determine the host's path in the application.

Note: When you expand a domain node, if any of the paths for hosts are not fully calculated, a pop-up dialog box displays a list of all the hosts with partially calculated paths. In addition, the current state of the path calculation is appended to the node name.

When you click a host node in the tree, the elements in the host's path appear highlighted in the right pane, as shown in the following figure.




Figure 9-13: Obtaining Path Information

You can also determine the elements in a hosts path by expanding the **Application Path** and **Path** nodes under the host node in the tree.




Figure 9-14: Path Information Visible in the Tree

When you right-click an element in the List tab, a drop-down menu is displayed. The options displayed depend on the type of element you clicked. See “About the Right-Click Menu Options” on page 317 for an explanation of the options in the drop-down menu.

About the Right-Click Menu Options

When you right-click an element in the topology pane or in the List or Access, or Path tabs, you see a drop-down menu. The options displayed in the menu depend on the type of element clicked.

See the following table for an explanation of the options displayed for elements right-clicked in the topology or right-clicked in the List or Path tab.

Note: Right-click menu options are not available to undiscovered fabrics.

Table 9-3: Drop-Down Menu Options Accessible from the Topology*

Menu Option	Description
Go to Navigation Details	This menu option redirects you to the Navigation page. If the element is labeled “Discovered,” you are shown the Properties page. Elements are labeled “Discovered” when the management server has discovered the element, but it cannot obtain additional information about it. See “About the Navigation Tab” on page 367.
Go to Element Topology	Displays a graphical representation of the path of an element. This also includes multipathing. See “Viewing Element Topology” on page 377.
Show Events	Displays the events for an element. See “About the Events Tab” on page 391.
Show Policies	Displays the Policy tab for the element. You can then view, add, modify and delete policies assigned to the element. See “About the Policies Tab” on page 396 for more information.
Update Element Data	The management server gathers new and changed details from the element and then redraws the topology with the updated information. The Update Element Data functionality does not detect element components that have been removed, such as ports and LUNs. For example, assume you removed several LUNs from an array. If you right-click the storage system and select Update Element Data, the LUNs still appear in the user interface. You must perform Get Details for the deleted LUNs to be removed from the user interface. See “Step 4 - Get Details” on page 94.
Show Impact	Highlights the elements that are impacted. See “Showing the Impact of an Element” on page 336.
Show Port Details	Lets you determine the use of each port for all elements in the network. See “Viewing Ports” on page 335 for more information.

Table 9-3: Drop-Down Menu Options Accessible from the Topology* (Continued)

Menu Option	Description
External Tools	Provides several ways to access an element: ■ Telnet - Lets you access a host or a switch through the telnet utility. The Telnet feature is only accessible to Web browsers on Microsoft Windows operating systems. ■ Browse - Lets you access the main Web page for a host or a switch. ■ Set Up External Tools - Lets you add URLs for accessing the management tools for the storage system. In some instances, the management tool for the storage system is accessible from this menu. For example, HiCommand for HDS storage systems and Command View for HP XP storage systems are accessible from the External Tools menu. See “Using External Tools” on page 366 for more information.
Discovered Elements	Lets you group unnamed generic hosts. It provides the following options. See “About Hiding Generic Hosts” on page 350. ■ Hide Generic Hosts for the Switch - Hides unnamed generic hosts that are connected to the switch. See “Hiding Generic Hosts for a Switch” on page 351. ■ Expand Generic Hosts for the Switch - Displays hidden unnamed generic hosts that are connected to the switch. See “Expanding Generic Hosts for a Switch” on page 352. ■ Hide Generic Hosts for All Switches - Hides unnamed generic hosts connected to all switches. See “Hiding Generic Hosts for All Switches” on page 352. ■ Expand Generic Hosts for All Switches - Displays hidden unnamed generic hosts that are connected to the switches. See “Expanding Generic Hosts for All Switches” on page 353.

Table 9-3: Drop-Down Menu Options Accessible from the Topology* (Continued)

Menu Option	Description
Provision	Provides provisioning tools for switches and storage systems: Switches - Lets you activate and deactivate zone sets, in addition to managing the following: ■ Zone Aliases (Not applicable to McDATA switches) ■ Zones ■ Zone Sets (The delete zone set option is disabled for active zone sets). Notes: ■ When McDATA or Connectrix switches are discovered through a proxy by using SNMP, you cannot view or perform any provisioning operations for those switches. For example, you cannot view zone sets, zones, and/or zone alias. ■ When McDATA or Connectrix switches are discovered by their IP address by using SNMP, you can only view the active zone set and its members. You cannot create, modify, and/or delete zone sets or its members. Storage Systems - Lets you manage the following: ■ Storage Pools ■ Volumes ■ Host Security Groups *Offered only when supported by storage system See “About the Provisioning Tab” on page 391, “SAN Zoning Overview” on page 401 and “Setting Up Storage Partitioning” on page 427 and for more information. Provisioning wizards may not be available in your kit. To determine if you can access Provisioning wizards, access the Feature List, which is accessible from the Documentation Center.
Add Virtual Application	Lets you add a virtual application so you can monitor it. A virtual application is a placeholder you create for an application. For example, you could create a virtual application for an application that was created just for your company. See “Adding a Virtual Application” on page 325.

Table 9-3: Drop-Down Menu Options Accessible from the Topology* (Continued)

Menu Option	Description
Set Business Cost (Applications only)	Lets you assign a business cost to an application. See "Assigning a Business Cost to an Application" on page 338.
Delete Element	Removes an element and its discovery instance from the system. It also removes other elements discovered through the removed element. See "Deleting Elements from the Product" on page 92. Important: If you are blocking pop-ups in a Netscape or Mozilla Web browser and you use the right-click menu to delete an element from System Explorer, the Delete window is blocked and you are unable to delete the element. You must disable the popup blocker before you can delete the element.
Adding a Custom Command	Lets you run a custom command on an element, for example to start an executable or a script. See "Adding a Custom Command" on page 355.
Group together with other elements	Lets you group "Discovered" hosts and storage systems. See one of the following topics: ■ "Grouping "Discovered" Hosts" on page 345 ■ "Grouping "Discovered" Storage Systems" on page 347
Ungroup into multiple elements	Lets you ungroup "Discovered" hosts and storage systems. See one of the following topics: ■ "Ungrouping "Discovered" Hosts" on page 346 ■ "Ungrouping "Discovered" Storage Systems" on page 348
Recalculate Topology	Lets you know about topology changes. The management server contacts the elements on the topology list (Discovery > Topology) to determine topology changes. The management server uses this updated information to redraw the topology.
Change Fabric Name (Only Available from the List Tab)	Lets you change the name of the fabric. See "Changing the Fabric Name" on page 349 for more information.

*Additional menu items may appear for types of Automators and Advisors, such as Reachable Storage. The descriptions of the right-click menu options for a fabric appear in the following section, since these menu options are only accessible from the List tab.

When you right-click a fabric in the List tab, you are shown the following options:

- **Change Fabric Name** - Type the new fabric name in the Change Fabric Name window. Then, click **OK**.
- **Go to Properties** - Displays the properties of the fabric. See “Viewing Element Properties” on page 374 tab for more information. It also provides access to the Events and Provisioning tabs. The Events tab displays events occurring within the fabric. The Provisioning tab lets you set up and manage zone provisioning. Provisioning wizards may not be available in your kit. To determine if you can access Provisioning wizards, access the Feature List, which is accessible from the Documentation Center.
- **Delete This Fabric** - Deletes a fabric. When you are asked if you want to delete the fabric, click **Yes** if you do not mind waiting for the management server to recalculate the topology. If the elements in the deleted fabric do not belong to another fabric, they are moved to the **unknown** node on the List tab.

Provisioning features are available from the right-click menu in the Access tab. When you right-click a zone, zone alias, or zone set in the Access tab, a drop-down menu is displayed. The provisioning options displayed in this menu depend on the type of element clicked.

Important: If you do not see the provisioning features from the right-click menu, your license does not allow you to access provisioning.

The following table describes the menu options. To learn more about each task and its required steps, see Chapter 10, “Provisioning” on page 399.

Table 9-4: Drop-Down Menu Options on the Access Tab





Task	To Perform Task, Right-Click...
1Activate a zone set 1Modify a zone set 1,3Delete a zone set	A zone set under the Zone Entries node. 
1Create a zone	A zone with the blue folder icon under a fabric node. 
1Delete a zone 1Modify a zone 1,2Show zone details	A zone under the Zone Set node. 
1Create a zone alias	A zone alias with a blue folder under the fabric node. 
1Delete a zone alias 1Modify a zone alias 1Show zone alias details	A zone alias under the Zone Aliases node. 

Table 9-4: Drop-Down Menu Options on the Access Tab (Continued)

Task	To Perform Task, Right-Click...
Provision Storage Pool	A storage system under the Storage System LUN Access node. 
Provision Volume	
Provision Host Security Groups	

1When McDATA and Connectrix switches are discovered through a proxy by SNMP, you cannot view or perform any provisioning operations for those switches.

2Only these options are accessible when McDATA and Connectrix switches are discovered by a switch IP address directly through SNMP. This setting provides view only access to the active zone set and its members. You cannot create, modify, and/or delete zone sets or its members.

3The delete zone set option is disabled for active zone sets.

Viewing Storage Elements

System Explorer has a wide range of features to help you in viewing your storage elements in the topology. For example, you can filter fabrics, arrange elements in the topology, and search for elements in the topology. You can even view the status of elements in the topology and find the impact of removing an element.

See the following topics in this section for more information.

- Adding a Virtual Application on page 325
- Adding Information for “Discovered” Hosts on page 326
- Arranging Elements in the Topology on page 327
- Closing Topology Windows on page 329
- Using the Global View on page 330
- Printing Elements in System Explorer on page 330

- Exporting the Topology to Visio on page 332
- Updating Element Data on page 335
- Viewing Ports on page 335
- Showing the Impact of an Element on page 336
- Showing the Impact of an Element on page 336
- Assigning a Business Cost to an Application on page 338
- Expanding the Topology Pane on page 340
- Filtering Fabrics on page 340
- Viewing Event Status in the Topology on page 341
- Custom Name for Switch Truncated in Topology on page 344
- Managing Groups on page 344
- Managing Fabrics on page 349
- Hiding and Showing Generic Hosts on page 350

Adding a Virtual Application

The management server lets you monitor applications not listed in the support matrix. For example, assume your company has created an internal application, and you want to be able to use the management server to monitor that application. You can create a virtual application for that product. A virtual application is a placeholder you create for an application.

Note: Only a user belonging to a role that has System Configuration selected on the **Edit Role** page (such as the Domain Administrator role) is allowed to create a virtual application.

Once you create the virtual application, it will appear connected to a host in your topology.

1. Select a host in **System Explorer**.
 2. Right-click and select **Add Virtual Application**.
 3. Type the following information for the virtual application.
 - ☐ **Name**
 - ☐ **Product**
 - ☐ **Description**
 - ☐ **Vendor**
-

□ **Version**

4. Select the storage volume for the application.

Note: You can view the properties of a volume by clicking its link.

5. Click **OK**.
The virtual application appears connected to the selected host.

Adding Information for “Discovered” Hosts

The software labels a host as “Discovered” when it cannot obtain additional information about a host it has discovered. To learn why the software was unable to obtain information about the element, see “Troubleshooting Discovery and Get Details” on page 802.

If you have more than one “Discovered” host, it can be difficult to differentiate them. To make them easier to identify, you may want to add information about the host, such as the following:

- Custom Name
- IP Address
- DNS Name
- Operating System
- Version of the operating system

Important: Do not add information for generic elements during Get Topology or Get Details. You can determine if the management server is getting the topology or all element details by looking at label near the status button.

1. Click **System Explorer** ()
 2. Double-click a “Discovered” host in the right pane.
The Properties tab is displayed.
 3. In the **Custom Name** field, type a name for the element.
Keep in mind the following:
 - The name must contain 1 to 64 characters.
-

- ☐ The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or one of the following symbols: period (.), dollar sign (\$), caret (^), hyphen (-), an underscore (_) or a space.
 - ☐ The name is case sensitive, for example, “Element1” and “element1” are different elements.
4. In the **IP Address** field, type an IP address for the element.
 5. In the **DNS Name** field, type a DNS name for the element.
 6. In the **Version** field, type the version of the operating system.
 7. In the **Operating System** field, type one of the following operating systems:
 - ☐ **AIX** - corresponds to IBM AIX®
 - ☐ **HP-UX** - corresponds to all versions of HP-UX™
 - ☐ **IRIX** - corresponds to SGI IRIX®
 - ☐ **Linux**
 - ☐ **Windows** - corresponds to Microsoft Windows®
 - ☐ **Solaris** - corresponds to Sun Solaris™
 8. Click **Save**.
When you access System Explorer, the information you entered appears in the topology.

Arranging Elements in the Topology

To improve usability, arrange the topology so it suits your environment. For example, if you plan to filter various fabrics, you might want to arrange the topology so that elements are arranged by fabric. This way when you filter the fabrics, large gaps do not appear in the topology. You can arrange elements individually or in groups, as described in the following sections.

Note: The topology displays direct attached connections as a dotted line from the host to the storage system. To view direct attach storage, you must enable the  button. See Table 9-1, “Description of Toolbar in System Explorer,” on page 301 for more information.

To arrange elements individually:

1. Click the  button.
 2. Click the element you want to move and drag it to a new location.
-

3. Repeat the previous step for each element you want to move.
The management server provides buttons to help you with viewing and arranging the topology. To learn more about those buttons, see “About the Toolbar in System Explorer” on page 300.
4. Once you have finished arranging the topology, click the  button to save it.
5. To learn more about filtering fabrics, see “Filtering Fabrics” on page 340.

To arrange elements in a group:

1. Click the  button.
2. Holding down the mouse button, drag the mouse diagonally across the elements you want to move. Then, lift your finger off the mouse button when you are done.
A square encloses the elements, as shown in the following figure. If you want to redo the square, just click outside of the square and retry.



Figure 9-15: Enclosing the Elements

3. To move the elements within the square, click within the square. Holding down the mouse button, drag the elements to the new location.




Figure 9-16: Dragging Multiple Elements to Their New Location

Closing Topology Windows

Whenever you select a new topology view, the software creates a pane for that view.

To lessen the number of panes open, do the following:

1. Right-click the tab of one of the views.
2. Select one of the following from the drop-down menu:
 - ☐ **Close** - Closes the current topology pane in System Explorer.
 - ☐ **Close All** - Closes all of the topology panes in System Explorer.
 - ☐ **Close All But Current** - Closes all of the topology panes in System Explorer, except the current one.

Using the Global View

If you have a large storage network, navigating it can be daunting, especially if it cannot fit in the pane. The global view () provides “a bird's eye” view of the network. With this view, you can move the viewing area to a certain section of the network.

1. Click the  button at the top of the screen.
A smaller pane displaying “a bird's eye” view of topology appears.
2. Move the brackets so they enclose the area of the network you want to view in the main pane.

Printing Elements in System Explorer

You can print the topology displayed in System Explorer.

To print the topology:

1. Click **System Explorer** (.
2. If the topology spans more than one screen, arrange the elements so they are closer together, preferably on one screen. To move an element, click the  button and then click the element you want to move. Drag the element to its new location. Moving elements closer together prevents the printout from appearing too stretched.
3. Click the  button.
4. Use the fields on the Paper tab to modify the setup of the page. When you are done, click **Apply**. If you want the default settings, click **Default**.
A preview of the printout is displayed in the right pane.

Important: Before you change the margins, decide on a unit of measurement.

- ☐ **Paper format** - Select the paper size from the drop-down menu.
 - ☐ **Unit** - Select cm (centimeters) or inch for the margins.
 - ☐ **Paper width** - Displays the width of the paper. You can modify the measurement in this field when you select the **Custom** option in the Paper format drop-down menu.
-

- ☐ **Paper height** - Displays the height of the paper. You can modify the measurement in this field when you select the **Custom** option in the Paper format drop-down menu.
 - ☐ **Top margin** - Type a measurement for the top margin.
 - ☐ **Bottom margin** - Type a measurement for the bottom margin.
 - ☐ **Left margin** - Type a measurement for the left margin.
 - ☐ **Right margin** - Type a measurement for the right margin.
 - ☐ **Orientation** - Click an orientation for the printout.
5. Click the **View Selection** tab to modify how the printout will appear on the page. You can modify the following. When you are done, click **Apply**. If you want the default settings, click **Default**.
A preview of the printout is displayed in the right pane.

Important: Before you change the margins, decide on a unit of measurement.

- ☐ **Start x** - Determines the horizontal placement of the printout on the page with zero being the closest to the right margin. For example, if the value is 50 for **Start x**, the printing starts at 50 inches or centimeters (depending on what you selected) from the right margin. You can also enter negative numbers. Anything more than zero expands the printout to another page.
 - ☐ **Start y** - Determines the vertical placement of the printout on the page with zero being the closest to the bottom margin. For example, if the value is 50 for **Start y**, the printing starts at 50 inches or centimeters (depending on what you selected) from the bottom. You can also enter negative numbers.
 - ☐ **Width** - Determines the width of the printout.
 - ☐ **Height** - Determines the height of the printout.
- To remove extra space around the topology, click the **Trimmed** button.
6. To change how many pages the printout will use, select one of the following. When you are done, click **Apply**. If you want the default settings, click **Default**.
A preview of the printout is displayed in the right pane.

Important: Before you change the margins, decide on a unit of measurement.

- ☐ **Unit** - Select cm (centimeters) or inch for the margins.
-

- **Position/Size** - Lets you change the position and size of the printout so that it spans several pages:
 - **Start x** - Determines the horizontal placement of the printout on the page with zero being the closest to the right margin. For example, if the value is 50 for **Start x**, the printing starts at 50 inches or centimeters (depending on what you selected) from the right margin. Determines the horizontal placement of the printout. Anything more than zero expands the printout to another page.
 - **Start y** - Determines the vertical placement of the printout on the page with zero being the closest to the bottom margin. For example, if the value is 50 for **Start y**, the printing starts at 50 inches or centimeters (depending on what you selected) from the bottom.
 - **Width** - Determines the width of the printout. If the width entered does not fit on the page, the printout wraps around to another page.
 - **Height** - Determines the height of the printout. If the height entered does not fit on the page, the printout wraps around to another page.
 - **Resolution (pixel/unit)** - Lets you change the resolution so that the printout spans several pages.
 - **Page** - Lets you expand the printout so it prints on several pages without modifying the graphic.
7. To preview your pages, click the Preview tab. Then click the page you want to preview. The page appears in the right pane.
 8. When you are ready to print, click **Print**.
 9. Click **Close**.

Note: To revert back to all of the original settings, click the **Default** button next to the **Print** button.

Exporting the Topology to Visio

You can export the topology to an XML file that can be viewed in Microsoft Visio. To export the topology:

1. Click **System Explorer** ()
2. Click the **Export to Visio** button.
3. Name the file and select the directory where you want the file to be saved.

4. Click the **Save** button. The XML file is saved to the directory that you chose.

Note: Backup topology can only be exported through Protection Explorer and not through System Explorer.

Viewing the Topology in Microsoft Visio

Once you have exported the topology to an XML file, you can view it in Microsoft Visio. To view the topology in Visio:

1. Install the necessary components as described in “Installing Storage Planner” on page 333.
2. Configure Visio as described in “Configuring Visio to View Exported Topology” on page 334.
3. Start Visio and select **Storage Planner > Import XML File**.
4. Browse to the XML file that you would like to view, and click the **Open** button. The selected topology displays in Visio.
5. Right-click any element and select **Properties**. The Custom Properties window opens and displays additional information about each element.
6. Each fabric is displayed in a separate layer. Select **View > Layer Properties** to display the Layer Properties window. This allows you to customize your view of the various fabrics. For example, you can change the color or visibility settings for each fabric.

Installing Storage Planner

Note: Microsoft Visio must be installed before installing Storage Planner.

Follow these steps to install Storage Planner:

1. Go to the Windows directory on the Cim Extensions CD-ROM and run StoragePlanner.exe.
 2. The Welcome to the Storage Planner Setup Wizard window displays. Click the **Next** button.
 3. The Select Destination Location windows displays. Click the **Next** button.
 4. The Select Components window displays. Click the **Next** button
-

5. Click the **Install** button. The Storage Planner component is installed.
6. If you haven't installed Microsoft XML 6.0 Parser, select the checkbox.
7. Click the **Finish** button.
8. If you are installing Microsoft XML 6.0 Parser, the MSXML 6.0 Parser Setup wizard displays.
9. Select the **I Accept** radio button and click the **Next** button.
10. Enter your name and company information. Click the **Next** button.
11. Click the **Install** button. Microsoft XML 6.0 Parser is installed.
12. Click the **Finish** button.

Configuring Visio to View Exported Topology

Before you can view the exported topology, some Visio settings must be changed. Follow these steps to correctly configure Visio:

1. Open Microsoft Visio.
2. Select **Tools > Options**.
3. Click the **Security** tab.
4. Click the **Macro Security** button.
5. Click the **Medium** radio button.
6. Click the **OK** button. You are returned to the previous window.
7. Click the **Advanced** tab.
8. Click the **File Paths** button.
9. Click the button to the right of the **Add-ons** field. Browse to your Visio installation directory and select the 1033/Solutions directory.
10. Click the button to the right of the **Start-up** field. Browse to your Visio installation directory and select the 1033/Solutions directory.
11. Click the **OK** button. You are returned to the previous window.
12. Click the **OK** button.
13. Restart Visio and select the **Enable Macros** button when prompted.

Note: If the Storage Planner menu item does not appear when you restart Visio, go to **Tools > Add-Ons > Storage Planner** and select **Storage Planner**.

Updating Element Data

System Explorer lets you update data about elements directly from this screen. When you update element data, the management server updates infrastructure details from the element and then redraws the topology with the updated information.

Keep in mind the following:

- Do not update element data during Get Topology or Get Details. You can determine if the management server is getting the topology or all element details by looking at the label near the status button.
- Updating element data does not detect element components that have been removed, such as ports and LUNs. For example, assume you removed several LUNs from an array. If you right-click the storage system and select **Update Element Data**, the LUNs still appear in the user interface. You must perform Get Details for the deleted LUNs to be removed from the user interface. See “Step 4 - Get Details” on page 94.
- Update Element Data does not correctly update switch ISL and topology changes. Get Details must be performed to obtain switch ISL and topology changes.

To update an element:

1. Click **System Explorer** ()
2. Right-click the element you want to update.
3. Select **Update Element Data** from the drop-down menu.
The software begins to update its database with the updated infrastructure details from the element.

During this process the status button appears red and “Getting Details” appears next to it.

When the process is complete, the status button returns to green.

Viewing Ports

When looking at an element on the network, such as a switch, it can be difficult to determine how the ports are used. System Explorer provides a view that lets you determine the use of each port for all elements in the network.

To view the ports:

1. Do one of the following:
-

- **Access System Explorer** - Click **System Explorer** ()
- **Access the Topology page (for an element)** - Do one of the following:
 - Double-click an element in System Explorer, and then click the **Topology** tab
 - Right-click an element, and then select **Show Element Topology** from the drop-down menu.
- 2. Right-click an element in the topology.
- 3. From the drop-down menu, select **Show Port Details**.
The ports are displayed.

Showing the Impact of an Element

You can display an element's impact. For example, assume you want to replace a switch. You can use this feature to determine which elements in the network would be impacted by taking the switch off the network.

To find an element's impact:

1. Right-click the element from which you want to obtain impact information.
2. Select **Show Impact** from the drop-down menu.
The other elements in the path of the element you right clicked are highlighted in yellow. For example, assume you right-clicked the Oracle instance klu2e (shown in the following figure) and selected **Show Impact** from the drop-down menu. The elements on which klu2e is dependent are highlighted in yellow. This means that if any of these highlighted elements are removed from the network, klu2e may have difficulty functioning.




Figure 9-17: Showing the Impact of an Application

However, the show impact feature not only shows the other elements on which an element is dependent, but it also displays the other elements dependent on it. For example, assume you right-clicked a switch and selected **Show Impact** from the drop-down menu. The highlighted elements would include its dependent elements, such as the hosts, applications, and storage systems connected to it. These elements might have difficulty communicating with one another if the switch was removed.

Likewise, if you decided to show the impact of a host, the highlighted elements would not only include its dependent elements, such as its applications, but also the elements on which it is dependent, such as switches.

Use the following table as a guideline to help you in determining whether the highlighted elements are dependent or required.

Table 9-5: Show Impact Results

If you select "Show Impact" for	Highlighted Elements
An Application (virtual or real)	Elements required by the application, such as its host and a switch.
A Host	■ Elements dependent on the host, such as its applications. ■ Elements required by the host, such as switches.
A Switch	Elements dependent on the switch, such as hosts and storage systems
Storage System	Elements dependent on the storage system, such as hosts.

Assigning a Business Cost to an Application

The management server lets you assign a business cost to an application, including virtual applications. This information is used in Event Manager for ranking events from elements. Event Manager determines the rank of an event by taking into account the business cost of the application and the severity of the event. You can sort events by rank in Event Manager by clicking the Rank column.

For example, assume you assigned a business cost of \$40 to one application and a business cost of \$30 to another application. If an event occurs with the same severity level occurs from both applications. The \$40 application has a higher rank because it has a higher "business cost."

The events from the elements in the path of the applications also inherit a business cost from the applications that use it. For example, assume a host has an application assigned a cost of \$30. The host would have a business cost of \$30. If the host has two applications (both valued at \$30), the host would be valued at \$60 because the two applications are using it. Likewise the switch

connected to the host would also have a value of \$60 because the two \$30 applications use it. If a switch has a \$40 application on one host and two \$30 applications on another host, the switch has a value of \$100, see the following figure.

The cost of the storage system is determined by the applications in its path. Two storage systems connected to a switch can have different business costs based on the applications in their path. For example, a storage system would have a value of \$60 if two \$30 applications are in its path, as shown in the following figure.




Figure 9-18: Determining Business Cost

To assign a business cost to an application:

1. Do one of the following:
 - ☐ Right-click an application in System Explorer, and then select **Set Business Cost** from the drop-down menu.
 - ☐ Double-click an application in System Explorer. Click the **Properties** tab. Then click the **Change** button next to the Business Cost field.
 - ☐ Click an application in Application Explorer. Click the **Properties** tab. Then click the **Change** button next to the Business Cost field.
 2. In the **Business Cost** field, type an amount, for example 35.25.
-

3. Click **OK**.

Expanding the Topology Pane

To increase screen space for viewing the topology, hide the List, Access, and Path tabs by clicking the arrow pointing left on the border between the pane containing the tabs and the main pane in System Explorer.




Figure 9-19: Expanding the Topology Pane

Note: To obtain more screen space, you might also want to close the left pane, as described in the topic, “Opening and Closing the Left Pane” on page 9.

To display these tabs, click the arrow pointing right on the border for the left pane.

Filtering Fabrics

To view a specified fabric in the topology, do the following:

1. Click **System Explorer** ().
2. Click the  button near the top of the screen.
3. Deselect the fabrics you do not want to view, as shown in the following figure:



Figure 9-20: Filtering Fabrics

- 4. Click the **Apply Filter** button.
System Explorer displays the selected fabrics.

Viewing Event Status in the Topology

You can obtain a status of the events occurring on the elements displayed in System Explorer by clicking the  button located on the toolbar. Elements with events that have occurred within the last five minutes have a triangle displayed next to them. The color of the triangle corresponds to the severity of the event.

Note: The Event Status button () is disabled in Capacity Explorer and Performance Explorer.

Table 9-6: Legend for Severity Level Icons

Icon	Description
	The event has a critical impact.
	The event has a major impact.

Table 9-6: Legend for Severity Level Icons (Continued)

Icon	Description
	The event has a minor impact.
	The event is providing a warning.
	The event is providing information.
	The severity of the event is not known.

The triangle corresponding to the highest severity is shown. For example, the management server displays an icon for a critical event next to an element that might also have minor events.

Since the severity level for an element is set by the manufacturer, the meanings of the severity levels vary. It is best to view the description of the event.

Use the Severity drop-down menu to filter which type of events you want to view. It displays the severity icons with the selected severity level or higher. For example, you can be notified of just critical and major events by selecting **Major** from the Severity drop-down menu and clicking the  button.

Note: If you select a severity, click the  button, and then leave System Explorer, System Explorer remembers your selection.

See the following table for a description of each severity level.

**Table 9-7: Descriptions of the Severity Levels Presented
in the Severity Menu**

Severity Level	Description
Critical	Denotes elements that have a critical severity level. For example, Brocade switches that have a failed firmware download and the failure reason code for each respective switch.
Major	Denotes elements that have a major severity level. It also displays icons for the critical severity level. For example, one or more physical fabric objects (device port, switch, or fabric) have disappeared.
Minor	Denotes elements that have a minor severity level. It displays icons for the following severity levels: ■ Critical ■ Major An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
Warning	Denotes elements that have a warning severity level. It displays icons for the following severity levels: ■ Critical ■ Major ■ Minor An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.

Table 9-7: Descriptions of the Severity Levels Presented in the Severity Menu (Continued)

Severity Level	Description
Informational	Denotes elements that have an informational severity level. It displays icons for the following severity levels: ■ Critical ■ Major ■ Minor ■ Warning An example of an informational event is a progress report event for firmware download operation currently in progress.
Unknown	Denotes elements that have unknown severity level. It displays icons for the following severity levels: ■ Critical ■ Major ■ Minor ■ Information

Custom Name for Switch Truncated in Topology

If a custom name for a switch is long, its name may appear truncated in the topology. The full name appears once the cursor is positioned over the switch with the custom name.

Managing Groups

This section describes the following:

- About Groups on page 345
- Grouping "Discovered" Hosts on page 345
- Ungrouping "Discovered" Hosts on page 346

- Grouping "Discovered" Storage Systems on page 347
- Ungrouping "Discovered" Storage Systems on page 348

About Groups

System Explorer lets you group together hosts and storage systems that have been labeled "Discovered" so the management server sees them as one element in the topology. The management server labels an element as "Discovered" when it has discovered the element, but it cannot obtain additional information about it. Grouped elements preserve space in the topology since only one element is displayed to represent the group. It also provides a way to keep track of all your "Discovered" hosts and storage systems.

Grouping "Discovered" Hosts

If you have several "Discovered" hosts, you might want to group them together so the management server sees them as one element in the topology. The management server labels a host as "Discovered" it when it has found the host but it cannot obtain additional information about it. Grouped elements preserve space in the topology since only one "Discovered" host is displayed to represent the group. Grouping also provides a way to keep track of your "Discovered" hosts.

Keep in mind the following:

- A user's role must include an access level of Element Control or Full Control for hosts. See the topic, "Editing Roles" on page 171 for more information about the access level of a role.
- Grouped elements are still seen as "Discovered" so the management server is unable to monitor or manage them.
- Do not create groups during Get Topology or Get Details. You can determine if the management server is getting the topology or all element details because the status button appears red during both operations.
- You can tell if a host is generic by double-clicking the host in System Explorer and then clicking the Properties tab. If a host is generic, it is listed as "Generic Host" for its description.

To group "Discovered" hosts:

1. Access System Explorer, as described in "Accessing System Explorer" on page 298.
 2. Right-click a "Discovered" host and select **Group together with other hosts** from the drop-down menu.
-

3. In the **Custom Name** field, type a custom name for the group.
4. In the **IP Address** field, type an IP address for the group.
5. In the **DNS Name** field, type the DNS name for the group.
6. In the **Version** field, type a version number for the group.
7. In the **Operating System** field, type the operating system for the hosts in the group.
8. Select the hosts you want to be a part of the group and click the button with the greater than sign (>).
The hosts are added to the group.

You can sort the hosts by:

- ☐ **Name** - To sort hosts by name, click the Hosts column heading.
- ☐ **Port** - To sort hosts by port, click the Ports column heading.
- ☐ **Connected Switches** - To sort hosts by connected switches, click the Connected Switches column heading.

An arrow appears next to the column heading that sorts the hosts. For example, if the hosts are being sorted by name, an arrow appears next to the Hosts column heading. If the arrow next to the column heading is pointing up, the hosts are sorted in ascending numerical and alphabetical order. If the arrow is pointing down, the hosts are sorted in descending numerical and alphabetical order. Click the column heading consecutively to change the direction of the arrow.

9. To remove hosts from the group, click the button with the less than sign (<).
10. Click **OK**.
The management server no longer displays the grouped elements in the topology individually. A host icon with the group name on the bottom represents the group. The group cannot be monitored or managed.

Ungrouping "Discovered" Hosts

You might want to ungroup "Discovered" hosts, if one of the hosts in the group is going to change. For example, the host will be taken off line.

Keep in mind the following:

- Do not ungroup elements during Get Topology or Get Details. You can determine if the management server is getting the topology or all element details because the status button appears red during both operations.
- A user's role must include an access level of Element Control or Full Control for hosts. See the topic, "Editing Roles" on page 171 for more information about the access level of a role.

To ungroup multiple elements:

1. Access System Explorer, as described in “Accessing System Explorer” on page 298.
2. Right-click host icon for a group and select **Ungroup into multiple hosts** from the drop-down menu.
3. When you are asked if you want to ungroup the elements, click **OK**.
The elements are ungrouped.

Grouping "Discovered" Storage Systems

If you have several "Discovered" storage systems, you might want to group them together so the management server sees them as one element in the topology. The management server labels a storage system as "Discovered" when it has found the storage system, but it cannot obtain additional information about it. Grouped elements preserve space in the topology since only one "Discovered" element is displayed to represent the group. It also provides a way to keep track of all your "Discovered" storage systems.

Keep the following in mind:

- A user's role must include an access level of Element Control or Full Control for storage systems. See “Editing Roles” on page 171 for more information about the access level of a role.
- Grouped elements are still seen as "Discovered", so the management server is unable to monitor or manage them.
- Do not group storage systems during Get Topology or Get Details. You can determine if the management server is getting the topology or all element details by looking at label near the status button.
- You can tell if a storage system is generic by double-clicking the storage system in System Explorer and then clicking the Properties tab. Look for its description.

To group "Discovered" storage systems:

1. Access System Explorer, as described in “Accessing System Explorer” on page 298.
 2. Right-click a "Discovered" storage system and select **Group together with other** from the drop-down menu.
 3. In the **Custom Name** field, type a custom name for the group.
 4. In the **Vendor** field, type the vendor names for the storage systems in the group.
 5. Select **Tape Library** if you want the entire group to be considered a tape library. Keep in mind this tape library will be considered as discovered, meaning it will not be managed or monitored by the management server.
-

6. Select the storage systems you want to be a part of the group and click the button with the greater than sign (>).

The storage systems are added to the group.

You can sort the storage systems by:

- ☐ **Name** - To sort storage systems by name, click the Storage Systems column heading.
- ☐ **Port** - To sort storage systems by port, click the Ports column heading.
- ☐ **Connected Switches** - To sort storage systems by connected switches, click the Connected Switches column heading.

An arrow appears next to the column heading that sorts the storage systems. For example, if the storage systems are being sorted by name, an arrow appears next to the Storage Systems column heading. If the arrow next to the column heading is pointing up, the storage systems are sorted in ascending numerical and alphabetical order. If the arrow is pointing down, the storage systems are sorted in descending numerical and alphabetical order. Click the column heading consecutively to change the direction of the arrow.

7. To remove storage systems from the group, click the button with the less than sign (<).
8. Click **OK**.

The management server no longer displays the grouped elements in the topology individually. A storage system icon with the group name on the bottom represents the group. The group cannot be monitored or managed.

Ungrouping "Discovered" Storage Systems

You might want to ungroup "Discovered" storage systems, if one of the storage systems in the group is going to change. For example, the storage system will be taken off line.

Keep in mind the following:

- A user's role must include an access level of Element Control or Full Control for storage systems. See the topic, "Editing Roles" on page 171 for more information about the access level of a role.
- Do not ungroup elements during Get Topology or Get Details. You can determine if the management server is getting the topology or all element details by looking at label near the status button.

To ungroup multiple elements:

1. Access System Explorer, as described in "Accessing System Explorer" on page 298.

2. Right-click a storage system icon for a group and select **Ungroup into multiple storage** from the drop-down menu.
3. When you are asked if you want to ungroup the elements, click **OK**. The elements are ungrouped.

Managing Fabrics

This section describes the following:

- Changing the Fabric Name on page 349
- Deleting Fabrics on page 349

Changing the Fabric Name

To change a fabric name:

1. Access System Explorer.
2. Click the **List** tab.
3. Right-click a fabric name.
4. Select **Change Fabric Name** from the drop-down menu.
5. In the **Enter a Fabric Name** field, type a new fabric name.
6. Click **OK**.

Deleting Fabrics

When you delete a fabric, the elements in the fabric are not removed. After you delete the fabric, the management server recalculates the entire topology. The recalculation may take some time, especially if you have a large topology.

To delete a fabric:

1. Access System Explorer.
 2. Click the **List** tab.
 3. Right-click a fabric name.
 4. Select the **Delete This Fabric** option from the drop-down menu.
 5. When you are asked if you want to delete the fabric, click **Yes**.
-

The management server recalculates the topology. If the elements in the deleted fabric do not belong to another fabric, they are moved to the **unknown** node on the List tab.

Hiding and Showing Generic Hosts

This section describes the following:

- About Hiding Generic Hosts on page 350
- Hiding Generic Hosts for a Switch on page 351
- Expanding Generic Hosts for a Switch on page 352
- Hiding Generic Hosts for All Switches on page 352
- Expanding Generic Hosts for All Switches on page 353

About Hiding Generic Hosts

You can reduce the amount of time it takes to arrange your topology, by using the hide unnamed generic hosts feature to hide hosts that a switch has detected. An element is considered to be generic if the management server can detect the element but it cannot obtain additional information about the element during Getting the Topology or Get Details.

When you use the show/hide feature, your changes persist to the next time you log into the management server. If you log in as another user, you will not see your changes. This feature allows each user to arrange the topology as he or she wishes.

Note: This feature only pertains to unnamed generic hosts. If you name a generic host, you cannot use this feature to hide the named host. The hide generic element feature also does not work for grouped unnamed generic hosts and missing elements. To learn how to give a custom name to an unnamed generic host, see “Assigning a Custom Name” on page 377. To learn more about groups, see the topic, “About Groups” on page 345.

The management server provides two variations of this feature:

- **Hiding Generic Hosts:** This feature hides unnamed generic hosts detected by a switch. The management server detects an element by looking at the ports on a switch. If it cannot find additional information about the element, it marks it as generic with a question mark over the element.
-

For example, assume you have a switch with 10 "discovered" elements. The management server detected these elements by looking at the ports on the switch and determined the type of element connected. "Discovered" elements appear with a question mark above their icon in the topology. The question mark indicates the management server has detected the element, but it cannot obtain additional information. To learn how to use this feature, see "Hiding Generic Hosts for a Switch" on page 351.

- **Hide Generic Hosts for the Switch:** This feature hides unnamed generic hosts within a domain. To learn how to use this feature, see "Hiding Generic Hosts for All Switches" on page 352.

Hiding Generic Hosts for a Switch

Simplify your topology by hiding unnamed generic hosts connected to a switch. If you have an unnamed generic host connected to more than one switch and you want to hide the generic element, you must repeat the following steps for each switch connected to the generic host. You can hide all unnamed generic hosts at once by using the **Hide Generic Hosts for All Switches** feature. See "Hiding Generic Hosts for All Switches" on page 352 for more information.

Note: This feature only pertains to unnamed generic hosts. If you name a generic host, you cannot use this feature to hide the named host. The hide generic element feature also does not work for grouped unnamed generic hosts and missing elements. To learn how to give a custom name to an unnamed generic host, see "Assigning a Custom Name" on page 377. To learn more about groups, see the topic, "About Groups" on page 345.

To hide generic hosts connected to a switch:

1. Right-click the switch.
2. Select **Discovered Element > Hide Generic Hosts for the Switch** from the drop-down menu.
A "+" icon is added to icon of the switch you right-clicked to indicate it has hidden generic hosts.

Expanding Generic Hosts for a Switch

Note: This feature only pertains to unnamed generic hosts. If you name a generic host, you cannot use this feature to hide the named host. The hide generic element feature also does not work for grouped unnamed generic hosts and missing elements. To learn how to give a custom name to an unnamed generic host, see “Assigning a Custom Name” on page 377. To learn more about groups, see the topic, “About Groups” on page 345.

Use the **Expand Generic Hosts for the Switch** feature to display hidden generic hosts. The hidden generic hosts are displayed in the upper-right corner of the topology if you select **Discovered Element > Expand Generic Hosts for the Switch** after you have saved the topology.

To display hidden generic hosts connected to a switch:

1. Right-click a switch with a "+" icon. This "+" icon indicates the switch has hidden generic hosts.
2. Select **Discovered Element > Expand Generic Hosts for the Switch** from the drop-down menu.
The hidden elements for the switch appear in the upper-right corner of the topology.

Hiding Generic Hosts for All Switches

Note: This feature only pertains to unnamed generic hosts. If you name a generic host, you cannot use this feature to hide the named host. The hide generic element feature also does not work for grouped unnamed generic hosts and missing elements. To learn how to give a custom name to an unnamed generic host, see “Assigning a Custom Name” on page 377. To learn more about groups, see the topic, “About Groups” on page 345.

You can hide all unnamed generic hosts and unnamed generic Cisco switches by doing the following:

1. Right-click a switch.
 2. Select **Discovered Element > Hide Generic Hosts for All Switches** from the drop-down menu. All unnamed generic hosts are hidden. A "+" icon is added to the icon of all switches that have generic hosts that have been hidden.
-

Expanding Generic Hosts for All Switches

Note: This feature only pertains to unnamed generic hosts. If you name a generic host, you cannot use this feature to hide the named host. The hide generic element feature also does not work for grouped unnamed generic hosts and missing elements. To learn how to give a custom name to an unnamed generic host, see “Assigning a Custom Name” on page 377. To learn more about groups, see the topic, “About Groups” on page 345.

Use the **Expand Generic Hosts for All Switches** feature to display hidden generic hosts. The hidden generic hosts are displayed in the upper-right corner of the topology if you select **Discovered Element > Expand Generic Hosts for All Switches** after you have saved the topology.

To display hidden generic hosts for a domain:

1. Right-click a switch with a "+" icon. This "+" icon indicates the switch has hidden generic hosts.
2. Select **Discovered Element > Expand Generic Hosts for All Switches** from the drop-down menu.
The hidden elements for the domain appear in the topology.

Setting Up Custom Commands

This section describes the following:

- About Custom Commands on page 354
 - Adding a Custom Command on page 355
 - Editing a Custom Command on page 358
 - Deleting a Custom Command on page 358
 - Software Environment Variables for Scripting on page 358
 - Using the Remote Console on page 362
-

About Custom Commands

Custom commands let you run a command you created on the management server. The command could point to an executable or a script that does not use the graphical user interface. For example, assume you have already created a script that backs up a storage system. You could run that script from System Explorer.

You can also use environment variables in your scripts. For example, you could use the variables to obtain information about a host, such as its total physical memory and the number of processors.

Keep in mind the following:

- The custom command always runs on the management server unless you are running the telnet utility. You can obtain information about the element on which you right-clicked by using the software's environment variables. See “Software Environment Variables for Scripting” on page 358.
- Run scripts at your own risk. The management server lets you run any script including those that can disable the management server.
- Custom commands only supports executables and scripts that do not use the graphical user interface.
- (Management Servers on Windows Only) If you leave the Remote Console (`cmd /k`) open after running a script, users can traverse the directory structure of the management server.
- If you want a Perl script to run as a custom command on Sun Solaris, you must prefix the script with the Perl executable, for example, `perl myscript.pl`, where `myscript.pl` is the script you want to run. A best practice is to prefix the script with the path to Perl and the Perl executable, for example: `perl/bin/perl myscript.pl`, where `perl/bin/` is the directory containing the Perl executable, `perl` is the executable and `myscript.pl` is the script you want to run.
- If you want a Perl script to run as a custom command on Microsoft Windows, you must prefix the script name with the complete path to Perl. The management server already has a directory containing the perl executable inside the folder, `%JBoss4_DIST%\server\appiq\remoteScripts\perl\bin`. You would prefix the script name as follows:
`.\perl\bin\perl myscript.pl`
where `.\perl\bin\` is the directory containing the Perl executable in the RemoteScripts directory, `perl` is the executable and `myscript.pl` is the script you want to run.

Adding a Custom Command

Use the custom command functionality in the software to run a command you created on the management server. The command could point to an executable or a script. For example, assume you have already created a script that backs up a storage system. You could run that script from System Explorer, as described in the following steps.

Keep in mind the following:

- The custom command always runs on the management server unless you are running the telnet utility. You can obtain information about the element on which you right-clicked by using the software's environment variables. See the topic, "Software Environment Variables for Scripting" on page 358.
- Run scripts at your own risk. The management server lets you run any script including those that can disable the management server.
- (Management Servers on Windows Only) If you leave the Remote Console (`cmd /k`) open after running a script, users can traverse the directory structure of the management server.
- If you want a Perl script to run as a custom command on Sun Solaris, you must prefix the script name with the Perl executable, for example, `perl myscript.pl`, where `myscript.pl` is the script you want to run. A best practice is to prefix the script name with the path to Perl and the Perl executable, for example: `perl/bin/perl myscript.pl`, where `perl/bin/` is the directory containing the Perl executable, `perl` is the executable and `myscript.pl` is the script you want to run.
- If you want a Perl script to run as a custom command on Microsoft Windows, you must prefix the script name with the complete path to perl. The management server already has a directory containing the perl executable inside the folder, `%JBOS4_DIST%\server\appiq\remoteScripts\perl\bin`. You would prefix the script name as follows:
`.\perl\bin\perl myscript.pl`
where `.\perl\bin\` is the directory containing the Perl executable in the RemoteScripts directory, `perl` is the executable and `myscript.pl` is the script you want to run.

To add a custom command:

1. Right-click an element in System Explorer.
 2. Select **Custom Commands > Set Up Custom Commands** from the drop-down menu.
 3. (Optional) If you plan to use a command to activate a file, such as a script, the file must be uploaded to the management server, as described below:
 - a. In the Custom Command Setup window, click the **Browse** button to find the file containing the custom command.
-

- b. Click the **Open** button, then click the **Upload to server** button.
The file is saved on the management server.
4. Click the **Add Command** button in the upper-right corner of the window.
5. In the Add Custom Command window, type a name for the command in the **Name** field, for example, **backup command**.
6. In the **Description** field, type a description, for example, **This command activates a script that backs up an element**.
7. In the **Command Line** field type a command, which will run on the management server. This could be a command required to start a script, for example the following:

```
myscript.bat
```

The remote console automatically becomes inactive once the command finishes.

(Windows Only) If you want the remote console to stay open, prefix the command with the following:

```
cmd /k
```

For example:

```
cmd /k dir
```

Important: When the command in the above example runs, it leaves the remote console window open, so you can enter additional commands. Users can traverse the directory structure of the management server when the remote console is left open.

The file is appended to the command line.(Optional) If you plan to use a file in the command, select the file from the **Files** drop-down menu. Then, click the **Append To Command Line** button.

If the file is missing, repeat step 3.

8. Select one of the following options to determine which elements you want the command to be visible. For example if you select the **All Elements** option, the command is visible in the drop-down menu when you right-click any element.
 - ☐ **Name of the Element** - Select the name of the element if you want the command to be visible in the drop-down menu when you right-click only this element.
 - ☐ **All Elements** - Select this option if you want the command to be visible in the drop-down menu when any element is right-clicked.
 - ☐ **Selected element types and filter criteria** - Select this option if you want to narrow the filtering criteria for an element type. For example, you could specify that the
-

command is only in the drop-down menu when a Brocade switch is right-clicked. The options are as follows:

- **Applications** - (Optional) If you want the command to be visible in the drop-down menu when a particular application is right-clicked, type the name of the product in the **Product Name** field. To make sure you type the correct product name, type the product name displayed in the Product Name field on the Properties tab, accessible by double-clicking the application in System Explorer and clicking the Properties tab.
- **Hosts** - (Optional) If you want the command to be visible in the drop-down menu when a particular host is right-clicked, type the name of the operating system in the **OS Name** field. To make sure you type the correct operating system, type the operating system displayed in the Target Operating System field on the Properties tab, accessible by double-clicking the host in System Explorer and clicking the Properties tab.
- **Switches** - (Optional) If you want the command to be visible in the drop-down menu when a switch from a particular vendor is right-clicked, type the name of the vendor in the **Vendor Name** field. To make sure you type the correct vendor name, type the vendor name displayed on the Properties tab, accessible by double-clicking the switch in System Explorer and clicking the Properties tab.
- **Storage Systems** - (Optional) If you want the command to be visible in the drop-down menu when a storage system from a particular vendor is right-clicked, type the name of the vendor in the **Vendor Name** field. To make sure you type the correct vendor name, type the vendor name displayed on the Properties tab, accessible by double-clicking the storage system in System Explorer and clicking the Properties tab.

9. Click **OK**.

10. When you want to run a command, do the following:

- a. Right-click an element.
- b. Select **Custom Commands** from the drop-down menu.
- c. Select the command from the **Custom Commands** drop-down menu.
A remote console displays the result of the command.

You can stop a command by clicking the **Stop** button in the remote console. Once a command has been executed the console becomes inactive. The software assumes you are in the `%MGR_DIST%\JBossandJetty\server\appiq\remotescripts` directory on the management server when the script is executed.

Editing a Custom Command

To edit a custom command:

1. Right-click an element in System Explorer.
2. Select **Custom Commands > Set Up Custom Commands** from the drop-down menu.
3. Click the  button corresponding to the custom command you want to edit.
4. Make the appropriate changes in the **Edit Custom Command** window.
5. Click **OK**.
The custom command is modified.

Deleting a Custom Command

To delete a custom command:

1. Right-click an element in System Explorer.
2. Select **Custom Commands > Set Up Custom Commands** from the drop-down menu.
3. Click the  button corresponding to the custom command you want to delete.
The custom command is deleted.

Software Environment Variables for Scripting

Keep in mind the following:

- Run scripts at your own risk. The management server lets you run any script including those that can disable the management server.
- (Management Servers on Windows Only) If you leave the Remote Console (`cmd /k`) open after running a script, users can traverse the directory structure of the management server.

The software provides the following environment variables for you to put in your scripts. For example, assume you have a script that backups a host. You could use the variables to obtain information about the host.

The software gathers information about the element you right-clicked. For example, if you use the variable, `APPIQ_ELEMENT_ID`, the management server obtains the element ID of the element you right-clicked.

The following table lists the variables that can be used to gather information for all elements.

Table 9-8: Variables for All Elements*

Variable	Value
APPIQ_ELEMENT_ID	The identifier of an element.
APPIQ_ELEMENT_NAME	The name of the element.
APPIQ_ELEMENT_STATUS	The status of the element, such as the following: ■ Managed ■ Generic ■ Missing ■ Virtual Application ■ Asset
APPIQ_ELEMENT_DESCRIPTION	The description for the element.
APPIQ_ELEMENT_MANAGEMENT_IP_ADDRESS	Contains the IP address of the first access point used to discover the element.
APPIQ_ELEMENT_VENDOR	The vendor for the element.
APPIQ_ELEMENT_TYPE_NAME	The type of element, for example, switch, application, or host.

*If an application resides on the host, the variables in this table provide information about the application. See the last table for variables that return information about the host.

The following table lists variables that can be used to gather information for only storage systems, switches, and hosts.

Table 9-9: Variables for Only Storage Systems, Switches, and Hosts*

Variable	Value
APPIQ_ELEMENT_IP_ADDRESS	The IP address of the element.
APPIQ_ELEMENT_DNS_NAME	The DNS name of the element.
APPIQ_ELEMENT_MODEL	The model of the element.
APPIQ_ELEMENT_VERSION	The version of the element.

*If an application resides on the host, the variables in this table provide information about the application. See the last table for variables that return information about the host.

The following table lists variables that can be used to gather information for only switches.

Table 9-10: Variables for Only Switches

Variable	Value
APPIQ_ELEMENT_SWITCH_ID	The identifier for the switch.
APPIQ_ELEMENT_IP_GATEWAY	The IP gateway of the switch.
APPIQ_ELEMENT_IP_NETWORK_MASK	The IP network mask for the switch.
APPIQ_ELEMENT_SWITCH_STATUS	The status of the switch.
APPIQ_ELEMENT_DOMAIN_ID	The domain identifier of the switch.

The following table lists variables that can be used to gather information for only hosts.

Table 9-11: Variables for Only Hosts*

Variable	Value
APPIQ_ELEMENT_OPERATING_SYSTEM	The operating system of the host.
APPIQ_ELEMENT_NUMBER_OF_PROCESSORS	The number of processors used by the host.
APPIQ_ELEMENT_TOTAL_PHYSICAL_MEMORY	The total physical memory of the host.
APPIQ_ELEMENT_DOMAIN	The domain of the host.

*If an application resides on the host, the variables in this table provide information about the application. See the last table for variables that return information about the host.

The following table lists variables that can be used to gather information for applications and hosts. Use the variables with the "APPIQ_HOST" prefix when you are using variables from the first table to gather information about the application. For example, if you are running a script containing APPIQ_ELEMENT_STATUS on a host, it would obtain information about the status of the application. You would need to run APPIQ_HOST_STATUS to obtain information about the status of the host on which the application resides.

Table 9-12: Variables for Only Applications

Variable	Value
APPIQ_ELEMENT_PRODUCT_NAME	The name of the application.
APPIQ_HOST_NAME	The name of the host on which the application resides
APPIQ_HOST_ID	The identifier of a host on which the application resides.
APPIQ_HOST_STATUS	The status of the host on which the application resides.
APPIQ_HOST_DESCRIPTION	The description of the host on which the application resides.
APPIQ_HOST_VENDOR	The vendor of the host on which the application resides.
APPIQ_HOST_TYPE_NAME	The type name of the host on which the application resides.
APPIQ_HOST_IP_ADDRESS	The IP address of the host on which the application resides.
APPIQ_HOST_DNS_NAME	The DNS name of the host on which the application resides.
APPIQ_HOST_MODEL	The model of the host on which the application resides.
APPIQ_HOST_VERSION	The version of the host on which the application resides.
APPIQ_HOST_OPERATING_SYSTEM	The operating system of the host on which the application resides.
APPIQ_HOST_NUMBER_OF_PROCESSORS	The number of processors on the host on which the application resides.
APPIQ_HOST_TOTAL_PHYSICAL_MEMORY	The total physical memory of the host on which the application resides.
APPIQ_HOST_DOMAIN	The domain of the host on which the application resides.

Using the Remote Console

This section describes the following:

- About the Remote Console on page 362
- Keeping the Remote Console Active on page 363
- About the Buttons on the Remote Console on page 364
- About the Drop-Down Menu Options on page 365
- Copying Text from the Remote Console on page 366

About the Remote Console

Whenever you run a custom command on the management server, the remote console appears. The remote console displays the result of a custom command. For example, you can use the remote console to start a remote command prompt on the management server.




Figure 9-21: Remote Console

Keeping the Remote Console Active

Note: This section is for only management servers running on Microsoft Windows.

Keep in mind the following:

- The remote console become inactive when the custom command finishes its execution. To use the menus and buttons in the remote console, the remote console must be kept active.
- If you leave the Remote Console (`cmd /k`) open after running a script, users can traverse the directory structure of the management server.

To keep the remote console window active, create a remote command prompt:

1. Right-click an element in System Explorer.
 2. Select **Custom Commands > Set Up Custom Commands** from the drop-down menu.
-

3. Click the **Add Command** button in the upper-right corner of the window.
4. In the **Add Custom Command** window, type a name for the command in the **Name** field, for example, **prompt**.
5. In the **Description** field, type a description, for example, **Accesses the remote console**.
6. In the **Command Line** field type the following command, which will run on the management server:
`cmd /k`
7. Select the **All elements** option.
8. Click **OK**.
9. When you want to run the remote command prompt, do the following:
 - a. Right-click an element from which you want to obtain information.
 - b. Select **Custom Commands** and select the command from the drop-down menu.
The software displays the remote console on the management server.
10. To enter a command in the remote console, type the command in the **Commands** field. Then, press ENTER.
You can stop a command by clicking the **Stop** button in the remote console.

Keep in mind the following:

- You can quickly access information about the element you right-clicked by typing the following at the command prompt:
`set appiq`
- The software ships with a utility called plink. To view the commands for plink, type the following in the **Commands** field and then press ENTER:
`plink`

About the Buttons on the Remote Console

The remote console provides the **Stop** and **Help** buttons, as described in the following table.

Table 9-13: Buttons on the Remote Console

Button	Description
Stop	Stops a command. Once a command has been executed the console becomes inactive.
Help	Provides the following information about the remote console: ■ Clearing the remote console - Type <code>CLS</code> in the Commands field of the remote console. ■ Copying text to the Commands field - Place the cursor at the end of the line in the window below the Commands field, as shown in the following figure. Then, press ENTER. The command is copied to the Commands field. Note: If you are viewing the remote console on Microsoft Windows, you can copy text by using CTRL + C, then use CTRL + P to paste it.

About the Drop-Down Menu Options

The remote console also provides the following drop-down menu options.

Table 9-14: Drop-Down Menu Options

Option	Description
CRLF	(Default setting) Provides a carriage return and a linefeed. Important: Do not use this option when you are using telnet to access another computer. You must select the CR option after you enter a user name to be able to enter a password. If you leave the setting at CRLF, the software enters a carriage return and a line feed when you click OK. As a result, no value is entered for the password.
CR	Provides a carriage return.
LF	Provides a linefeed.

Copying Text from the Remote Console

To copy text from the remote console:

1. Select the text in the remote console.
2. Right-click the top frame in the remote console.
3. Select **Copy** from the drop-down menu.
The text is stored in the buffer of your computer to be pasted elsewhere.

Using External Tools

This section describes the following:

- About the External Tools Feature on page 366
- Set Up External Tools - Lets you add a URL for accessing management software, such as Hitachi HiCommand Device Manager and EMC ControlCenter™ Navisphere. See “Setting up External Tools” on page 367.

About the External Tools Feature

The management server ships with an external tools feature that lets you:

- **Browse the Element** - Access a host or a switch through its main Web page. The software assumes the host or switch has a Web page at `http://hsIPAddress`, where `hsIPAddress` is the IP address of the host or switch. To access the main Web page of the host or switch, right-click the element in System Explorer and select **External Tools > Browse to 192.168.1.2**, where 192.168.1.2 is the IP address of the host or switch.
 - **Telnet to the Element** - Access a host or a switch through the telnet utility. Telnet must be already enabled on the element. The command uses `telnet://hsIPAddress`, where `hsIPAddress` is the IP address of the host or switch. To telnet to a host or switch, right-click the element in System Explorer and select **External Tools > Telnet to 192.168.1.2**, where 192.168.1.2 is the IP address of the host or switch.
 - **Set Up External Tools** - Lets you add a URL for accessing management software, such as Hitachi HiCommand Device Manager and EMC ControlCenter™ Navisphere. See “Setting up External Tools” on page 367 for more information.
-

- **Access the Management Tool for the Storage System** - In some instances, the management tool for the storage system is accessible from this menu. For example, HiCommand for HDS storage systems and Command View for HP XP storage systems are accessible from the External Tools menu.

Setting up External Tools

You can add URLs used for accessing external tools used for managing an element, such as Hitachi HiCommand Device Manager and EMC ControlCenter™ Navisphere for storage systems.

Note: When you add a URL, it only applies to the element you originally right-clicked.

To add a URL for accessing external tools:

1. Access System Explorer.
2. Right-click the element and select **External Tools > Set Up External Tools**.
3. Click the **Add New Management URL** button.
4. In the **Description** field, type the name of the product you plan to access.
5. In the **URL** field, type the URL that is used to access the product.
6. Click **OK**.

When you right-click the element and select **External Tools**, the external tool is listed.

To delete the URL for an external tool, click the corresponding  button in the External Tools window.

About the Navigation Tab

The Navigation tab not only provides information about an element, but it also illustrates how the element relates to other elements in its path. For example, the Navigation page displays logical and physical components, such as ports, zone sets, zones and zone aliases. It also displays the dependencies for switches, as shown in the following figure.

Keep in mind the following:

- When McDATA or Connectrix switches are discovered through a proxy by using SNMP, you cannot view or perform any provisioning operations for those switches. For example, you cannot view zone sets, zones, and/or zone alias.
- When McDATA or Connectrix switches are discovered by their IP address by using SNMP, you can only view the active zone set and its members. You cannot create, modify, and/or delete zone sets or its members.
- If you see a message that zone aliases are not supported on a Brocade switch, perform Get Details. The management server does not gather provisioning information from a fabric until Get Details is performed.




Figure 9-22: Obtaining Information About a Host

You can learn more about a component, by clicking it in the Navigation page. For example, assume you brought up the Navigation page shown in the previous figure, and you want to learn which hosts are dependent. You would click Hosts in the page. You are shown information about the dependent hosts, as in the following figure:




Figure 9-23: Details of a Host Connected to a Switch

The following table provides an overview of the information presented for each type of element.

Table 9-15: Information Available from the Navigation Page

Element	Dependencies	Front Physical	Back Physical	Logical	Physical
Applications	✓				
*Hosts	✓				
Switches	✓			✓	✓
Storage Systems	✓	✓	✓	✓	

*The management server displays cxfs for SGI IRIX computers if it detects cxfs on the cluster. On individual IRIX computers "cxfs" is not displayed when you type the following at the command prompt:

```
df -k
```

Data may be missing from the Navigation tab for a McDATA or a Connectrix switch if the switch was discovered by using SNMP and one of the following techniques:

- Its IP address
- Enterprise Fabric Connectivity (EFC) Manager or EMC Enterprise Connectrix Manager (ECM)

If the McDATA and/or Connectrix switches are discovered by their IP address (SNMP connection), the following fields are empty:

- IP Gateway
- Switch ID
- FC Net Address
- FC Net Mask

If the McDATA and/or Connectrix switches are discovered by SNMP and by Enterprise Fabric Connectivity (EFC) Manager or by EMC Enterprise Connectrix Manager (ECM), the following fields are empty:

- IP Gateway
- Switch ID
- DNS Name

- IP Address
- IP Net Mask
- FC Net Address
- FC Net Mask

To learn how to access the Navigation tab, see “Accessing the Navigation Tab” on page 373 for more information.

Finding the Status of a Port on a Switch

The management server can detect the status of a switch. This can be especially useful if a port is being a problem. To find the status of a port on a switch:

1. Access System Explorer as described in “Accessing System Explorer” on page 298.
2. Double-click a switch in the right pane.
3. Click the **Ports** button under the Physical column in the Navigation tab, as shown in the following figure.




Figure 9-24: Finding the Status of a Port

4. Under the Name column in the Ports table, click the port that you want to obtain the status.
5. On the Properties page, the status of the port is displayed in the right column.

The status of the port can be online, off line or unknown.

Table 9-16: Port Status Definitions

Status	Definition
Online	Port is physically installed with node connections
Offline	Port is physically installed, but without node connections Brocade switches also display this status if the port is not physically installed (Gigabit Interface Converter (GIBIC) is not installed).
*Not Installed (McDATA SWAPI connections only)	Port is not physically installed (Gigabit Interface Converter (GIBIC) is not installed).
*Unknown (McDATA SNMP connections only)	Port is not installed.

*An ES4500 switch displays its status differently when a port is not installed:

- **SWAPI connection** - Unknown status
- **SNMP connection** - Offline

Accessing the Navigation Tab

To access the **Navigation** tab:

1. Access the management server.
2. To access the **Navigation** tab, do one of the following:
 - ☐ Click an element in Application Explorer.
 - ☐ Double-click an element in Capacity Explorer, Performance Explorer or System Explorer.
3. Click the **Navigation** tab.

Viewing Element Properties

This section describes the following:

- About the Properties Tab on page 374
- Accessing the Properties Tab on page 375
- Viewing Fabric Properties on page 376
- Assigning a Custom Name on page 377

About the Properties Tab

The Properties tab provides detailed information about an element. Since the information obtained from each type of element varies, the Properties tab displays information only relevant for that type of element. For example, the Properties tab for fabrics lists the zones, zone sets, switches, and zone aliases, as compared to the Properties tab for a host, which lists the processors, cards, applications, and storage volumes the host uses.

See “Viewing Fabric Properties” on page 376 for more information about the Properties tab for fabrics.

The Properties tab usually provides the following, although this list does vary from element to element:

- ***Assign a Custom Name** - To make it easier to identify the element, assign the element a Custom Name. See the topic, “Assigning a Custom Name” on page 377. This option is not available to all elements.
 - ***IP Address (Generic Hosts Only)** - Type an IP address for a generic host.
 - ***DNS Name (Generic Hosts Only)** - Type a DNS name for a generic host.
 - ***Version (Generic Hosts Only)** - Type a version number for the generic host.
 - ***Operating System (Generic Hosts Only)** - Type an operating system for a generic host.
 - **Vendor** - The vendor name.
 - **View element properties** - The element properties viewed on the type of device. The properties provided vary according to the element. The following information is usually provided:
 - **Record Creation** - The first time the software contacted this element.
 - **Discovery Status** - The status of the discovery of the element, for example “Contacted.”
-

- ☐ **Vendor** - The name of the vendor.
- ☐ **IP Address** - The IP address of the element.
- ☐ **DNS Name** - The element's domain name system (DNS) name.
- ☐ **Provider Name** - The name of the provider.
- ☐ **Model** - The model of the element.
- **Update Element Data** - To update the displayed properties, click the **Update Element Data** button at the bottom of the screen. The management server gathers new and changed details from the element and then redraws the topology with the updated information.
Important:
 - ☐ Do not update element data during Get Topology or Get Details. You can determine if the management server is getting the topology or all element details by looking at the label near the status button.
 - ☐ The Update Element Data functionality does not detect element components that have been removed, such as ports and LUNs. For example, assume you removed several LUNs from an array. If you perform **Update Element Data** for the storage system, the LUNs still appear in the user interface. You must perform Get Details for the deleted LUNs to be removed from the user interface. See the topic, "Step 1 - Discover Switches" on page 34 for more information about Get Details.

*Click the **Save** button for your changes to take effect.

Accessing the Properties Tab

To access the **Properties** tab:

1. Access the management server.
 2. To access the **Properties** tab for an element, do one of the following:
 - ☐ Click an element, except a file server, in Application Explorer.
 - ☐ Double-click an element in System Explorer.
 3. To access the **Properties** tab for a fabric:
 - a. In System Explorer, click the **List** tab.
 - b. Right-click a fabric name in the **List** tab, for example 100000606930260d.
 - c. Select **Go to Properties** from the right-click menu.
 4. (Not applicable to fabrics) Click the **Properties** tab.
-

Viewing Fabric Properties

You can view the following properties of a fabric:

- **Assign a Custom Name** - To make it easier to identify the element, assign the element a Custom Name. See “Assigning a Custom Name” on page 377. This option is not available to all elements.
- **Vendor** - The vendor name.
- **Created** - The first time the software contacted this element.
- **Discovery Status** - The status of the discovery of the element, for example “Contacted.”
- **Install Date** - Not applicable
- **Name Detected** - The name of the fabric detected.
- **OID** - Not applicable
- **Description** - Provides information about the fabric.
- **WWN** - The Worldwide Name of the fabric.
- **Zones** - The zones in the fabric. To learn more about a zone, click its link.
- **Zone Sets** - The zone sets in the fabric. To learn more about a zone set, click its link.
- **Switches** - The switches in the fabric. To learn more about a switch, click its link.
- **Zone Aliases** - The zone aliases in the fabric.
- **Update Element Data** - To update the displayed properties, click the **Update Element Data** button at the bottom of the screen. The management server gathers new and changed details from the element and then redraws the topology with the updated information.

Keep in mind the following:

- Do not update element data during Get Topology or Get Details. You can determine if the management server is getting the topology or all element details by looking at the label near the status button.
- The Update Element Data functionality does not detect element components that have been removed, such as ports and LUNs. For example, assume you removed several LUNs from an array. If you perform **Update Element Data** for the storage system, the LUNs still appear in the user interface. You must perform Get Details for the deleted LUNs to be removed from the user interface. See the topic, “Step 1 - Discover Switches” on page 34 for more information about Get Details.

To learn how to access the Properties tab for a fabric, see the topic, “Accessing the Properties Tab” on page 375.

Assigning a Custom Name

To make it easier to identify an element instance in the system, assign the instance a Custom Name. The Custom Name also appears in Chargeback.

Important: Do not assign a custom name during Get Topology or Get Details. You can determine if the management server is getting the topology or all element details by looking at label near the status button.

Note: Since all users query the same database, this name is displayed to others using the software. As a result, you might want to make them aware of the name.

1. Access the **Custom Name** field by double-clicking the element in System Explorer and then clicking the **Properties** tab.
2. In the **Custom Name** field, type a name.
Keep in mind the following:
 - ☐ The name must contain 1 to 64 characters.
 - ☐ The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or one of the following symbols: period (.), dollar sign (\$), caret (^), hyphen (-), an underscore (_) or a space.
 - ☐ The name is case sensitive, for example, “Element1” and “element1” are different elements.
3. Click **Save**.

Viewing Element Topology

This section describes the following:

- About the Topology Tab on page 378
 - About the New Window Option on page 386
-

- Printing the Topology on page 387

About the Topology Tab

The Topology tab provides a graphical representation of an element's path. It displays additional information not found in System Explorer, such as adapters, slots, and fibre channel ports.

For example, assume you want to view the topology of a server called QASERVER02, and it contains seven fixed local disks. If you double-click the server in System Explorer and then click the **Topology** tab, you can see the path of the server. The Topology tab also displays the drives of the server's fixed local disks, as well as the adapter used to connect the server with the switch, as shown in the following figure. According to the following figure, the server can access three storage systems: Engenio, EMC, and HDS.




Figure 9-25: Topology of a Server

The topology extends the length of the screen. The second portion of the topology is provided by the following figure.




Figure 9-26: Topology of a Server (Continued)

Note: If any of the paths are not fully calculated, a pop-up dialog box displays a list of all the hosts with partially calculated paths. In addition, the current state of the path calculation is appended to the node name in the left pane.

Multipathing

Multipathing is the process of providing a server more than one path to a storage system, so that in the case of an emergency, the server will have continuous access to the storage system.

Multipathing can be done many ways. For example, you can provide redundant switches for a host to access a storage system. Another example of multipathing is providing redundant paths from the host to the switch. To determine if your multipathing software is supported, refer to the release notes.

Keep in mind the following:

- "SANtricity" Manager Utilities" must be installed on the host running RDAC for the management server to obtain RDAC information.
- HDLM on Sun Solaris requires the storage array to be included in discovery in order to report the correct information back to the bindings page. For more information about HDLM issues, see "Troubleshooting HDLM."
- The software supports VERITAS Volume Manager without VxDMP, but VxDMP is required to do multipathing.
- Microsoft Windows 2003 hosts with Service Pack 1 and IBM TotalStorage DS6800 arrays do not stitch properly as a result of the Subsystem Device Driver (SDD) appearing on the same disk. The multipathing page returns the Windows SDD path as something similar to "600507630efe01a80000000000001104:c0t0d0p3". This makes it difficult to match it up to your SDD path names.

The following figure displays how the software detects multipathing for a server running VERITAS Volume Manager. "MP" is displayed on the path of the redundant volumes.




Figure 9-27: Multipathing Displayed in the Topology

The topology extends the length of the screen. The second portion of the topology is provided by the following figure.




Figure 9-28: Multipathing Displayed in the Topology (Continued)

Keep in mind the following:

- If you do not see all of the elements in the path displayed, verify they have been discovered and details have been obtained from them. See Chapter 2, “Discovering NAS Devices, Tape Libraries, Switches and Storage Systems” on page 21.
- See “Known Device Issues” on page 831 for multipathing issues regarding certain devices.
- The management server displays only the active path for an RDAC host. It also displays only the active path when PowerPath is running on a host connected to a CLARiiON storage system. The management server also does not support RDAC configurations for monitoring disk statistics.

Direct Attached Storage

To view direct attach storage, you must enable the  button. See Table 9-1, “Description of Toolbar in System Explorer,” on page 301 for more information.

Once the  button is enabled, the management server displays the link between the storage system port and the port to the host as a dotted line, as shown in the following figure:




Figure 9-29: Direct Attached Storage in the Topology

In this figure, Slot A Port 1 belongs to the storage system and Adapter 1 Port 0 belongs to the host. The dotted line indicates the storage system is directly attached to the host.

Filers

Element topology for a filer shows the connection from a host to the filer going through an IP cloud, which represents the IP network.

Accessing the Topology

To access the Topology tab:

- 1. Access the management server.
- 2. Access the Topology tab by doing one of the following:
 - ❑ Selecting an element in Application Explorer, and then clicking the **Topology** tab.
 - ❑ Double-clicking an element in Capacity Explorer or System Explorer, and then clicking the **Topology** tab.
 - ❑ Right-clicking an element in System Explorer, and then selecting **Go To Element Topology** from the drop-down menu.

Table 9-17: About the Toolbar in the Topology Tab



Icon	Description
	Prints the topology. See the topic, “Printing the Topology” on page 387 for more information.
	Magnifies the view
	Decreases the magnification
	Lets you set the magnification to a percentage of the default magnification
	Opens a smaller pane, which provides a global view of the topology. This lets you position the main view to a certain section of the topology. For more information, see “Using the Global View” on page 330.
	Lets you drag an element in the topology.
	Lets you move the entire topology at once. Click the Pan button () and then click any place in the topology. Drag the mouse to a new location.

Table 9-17: About the Toolbar in the Topology Tab (Continued)

Icon	Description
	Lets you find an element by name or by Worldwide Name in the topology. You can type part of the information, and the management server highlights the elements that match. After you populate the search field, click the  button or press ENTER. To expand the Search field, close the left pane. See the topic, “Opening and Closing the Left Pane” on page 9 for more information.
	Displays the event severity icons for the elements displayed in the topology. See “Viewing Event Status in the Topology” on page 341 for more information. Disabled for Performance Explorer and Capacity Explorer.

About the New Window Option

The New Window Option in System Explorer lets you view several sections of the topology at once. Click the  button. A new window pops up. Use this pane to view another section of the topology.




Figure 9-30: New Window Option

Printing the Topology

The software lets you print the topology. This option is extremely helpful when you want to show someone the layout of the network, such as in a presentation.

To print the topology:

1. Access the management server.
2. Access the Topology tab by doing one of the following:
 - ☐ Selecting an element in Application Explorer, and then clicking the **Topology** tab.
 - ☐ Double-clicking an element in Capacity Explorer or System Explorer, and then clicking the **Topology** tab.

3. If the topology spans more than one screen, arrange the elements so they are closer together, preferably on one screen. To move an element, click the  button and then the element you want to move. Drag the element to its new location. Moving elements closer together prevents the printout from appearing too stretched.
4. Click the  button.
5. Use the fields on the Paper tab to modify the setup of the page. When you are done, click **Apply**. If you want the default settings, click **Default**.
A preview of the printout is displayed in the right pane.

Important: Before you change the margins, decide on a unit of measurement.

- ☐ **Paper format** - Select the paper size from the drop-down menu.
 - ☐ **Unit** - Select cm (centimeters) or inch for the margins.
 - ☐ **Paper width** - Displays the width of the paper. You can modify the measurement in this field when you select the **Custom** option in the Paper format drop-down menu.
 - ☐ **Paper height** - Displays the height of the paper. You can modify the measurement in this field when you select the **Custom** option in the Paper format drop-down menu.
 - ☐ **Top margin** - Type a measurement for the top margin.
 - ☐ **Bottom margin** - Type a measurement for the bottom margin.
 - ☐ **Left margin** - Type a measurement for the left margin.
 - ☐ **Right margin** - Type a measurement for the right margin.
 - ☐ **Orientation** - Click an orientation for the printout.
6. Click the **View Selection** tab to modify how the printout will appear on the page. You can modify the following. When you are done, click **Apply**. If you want the default settings, click **Default**.
A preview of the printout is displayed in the right pane.

Important: Before you change the margins, decide on a unit of measurement.

- ☐ **Start x** - Determines the horizontal placement of the printout on the page with zero being the closest to the right margin. For example, if the value is 50 for **Start x**, the printing starts at 50 inches or centimeters (depending on what you selected) from the right margin. You can also enter negative numbers. Anything more than zero expands the printout to another page.
-

- ☐ **Start y** - Determines the vertical placement of the printout on the page with zero being the closest to the bottom margin. For example, if the value is 50 for **Start y**, the printing starts at 50 inches or centimeters (depending on what you selected) from the bottom. You can also enter negative numbers.
 - ☐ **Width** - Determines the width of the printout.
 - ☐ **Height** - Determines the height of the printout.
- To remove extra space around the topology, click the **Trimmed** button.

7. To change how many pages the printout will use, select one of the following. When you are done, click **Apply**. If you want the default settings, click **Default**.
A preview of the printout is displayed in the right pane.

Important: Before you change the margins, decide on a unit of measurement.

- ☐ **Unit** - Select cm (centimeters) or inch for the margins.
 - ☐ **Position/Size** - Lets you change the position and size of the printout so that it spans several pages:
 - **Start x** - Determines the horizontal placement of the printout. Anything more than zero expands the printout to another page.
 - **Start y** - Determines the vertical placement of the printout.
 - **Width** - Determines the width of the printout. If the width entered does not fit on the page, the printout wraps around to another page.
 - **Height** - Determines the height of the printout. If the height entered does not fit on the page, the printout wraps around to another page.
 - ☐ **Resolution (pixel/unit)** - Lets you change the resolution so that the printout spans several pages.
 - ☐ **Page** - Lets you expand the printout so it prints on several pages without modifying the graphic.
8. To preview your pages, click the Preview tab. Then click the page you want to preview. The page appears in the right pane.
9. When you are ready to print, click **Print**.
10. Click **Close**.

Note: To revert back to all of the original settings, click the **Default** button next to the **Print** button.

Creating a Virtual Application

The management server lets you keep track of unsupported applications. For example, assume your company has created an internal application, and you want to be able to use the software to keep track of that application. You can create a virtual application for that product. A virtual application is a placeholder you create for an application.

Once you create the virtual application, it will appear connected to a host in your topology.

1. Access System Explorer by clicking the **System Explorer** button in the left pane.
2. Right-click the host that contains the application you want to monitor.
If the host is not in the topology, verify you have discovered the element and obtained element details. See Chapter 2, “Discovering NAS Devices, Tape Libraries, Switches and Storage Systems” on page 21.
3. Select **Add Virtual Application** from the drop-down menu.
4. Type the following information for the virtual application.
 - ☐ **Name**
 - ☐ **Product**
 - ☐ **Description**
 - ☐ **Vendor**
 - ☐ **Version**
5. Click **Next**.
6. Select a storage volume containing the application for which you are creating the virtual application.

Note: You can view the properties of a volume by clicking its link.

7. If applicable, choose a disk partition by clicking the **Disk Partitions** tab or the **Next** button and then selecting a disk partition.
 8. Click **Finish**.
-

About the Provisioning Tab

The provisioning tab provides different functionality depending on the type of element you double-clicked in System Explorer or clicked in the Provisioning pages. You can also access the provisioning table by right-clicking a fabric, selecting the **Go to Properties** option, and clicking the Provisioning tab.

If you selected a switch or a fabric, you are shown zone provisioning tools that let you manages zones, zone aliases, and zone sets. These tools provide a wide range of functionality, such as the following:

- “Creating Zone Aliases” on page 409
- “Creating a Zone in a Fabric” on page 413
- “Creating a Zone Set” on page 417
- “Activating a Zone Set” on page 421

For more information about setting up zones, see “SAN Zoning Overview” on page 401.

If you double-clicked a storage system, you are shown storage provisioning tools that let you create storage pools, volumes, and host security groups. These tools provide a wide range of functionality, such as the following:

- “Managing Storage Pools” on page 434
- “Managing Volumes” on page 438
- “Rules for Creating Host Security Groups” on page 448
- “Managing Host Security Groups” on page 452

About the Events Tab

The Events tab lets you view, clear, sort and filter events for an element. An event can be anything that occurs on the element, such as a device connected to a Brocade switch has gone off-line. It provides the following information about the events:

- **ID**- The identification number assigned to the event
 - **Severity** - Provides the severity level
 - **Time** - The time the event was recorded.
-

- **Summary Text** - A brief explanation of the event. When you click the summary text, the details of the event are displayed.

The Events tab lets you use Event Manager to:

- **View Event Details** - See “Viewing Event Details” on page 541.
- **Clear Events** - See “Clearing Events” on page 542.
- **Unclear Events** - See “Unclearing Events” on page 543.
- **Delete Events** - See “Deleting Events” on page 545.
- **Sort Events** - See “Sorting Events” on page 545.
- **Select a Severity for Filtering** - See “About Filtering Events” on page 550 and “Selecting a Severity Level” on page 551.
- **Customize the Severity Level Filter** - See “Customizing the Severity Level Filter” on page 552 and “About Filtering Events” on page 550.

To view all events, click the **Event Manager** button in the left pane. See “About Event Manager” on page 536.

Asset Attributes of an Element

Important: Depending on your license, Chargeback may not be available. See the "List of Features" to determine if you have access to Chargeback. The "List of Features" is accessible from the Documentation Center (**Help > Documentation Center**).

Access asset information for an element by doing the following:

- Clicking an element (except for a file server) in Application Explorer, and then clicking the **Asset Management** tab.
- Double-clicking an element in Capacity Explorer, Performance Explorer or System Explorer, and then clicking the **Asset Management** tab.
- Clicking an element in Chargeback.

Chargeback provides a handy way for you to keep track of your asset information for an element. You can easily store warranty and licensing information, as well as contact information for the

element. For example, assume a switch on the network is having some problems, and you want to contact the person in charge of that switch. You can use the element's asset record to not only find the contact information for that switch, but also the location of that switch.

The Asset Management tab displays general asset information about an element. It also provides access to other screens that provide additional asset information, such as staff, geographic, licensing, and warranty information. You can access these other screens by expanding the Asset Record node and clicking one of its children, as shown in the following figure. To learn more about these other screens, see the following topics:

- “Adding Asset Information” on page 724
- “Adding General Information” on page 727
- “Adding Staff Information” on page 728
- “Adding Geographic Information” on page 729
- “Adding Licensing and Warranty Information” on page 729
- “Adding Custom Information” on page 729




Figure 9-31: Viewing Asset Records

To set up chargeback, expand the Chargeback node and click Asset-based or Storage-based. To learn more about each type of Chargeback, see “Setting Up Asset-Based Chargeback” on page 733 and “Setting Up Storage-Based Chargeback” on page 739.

This section describes the fields on the Asset Management tab. When you are done with adding information on this page, click the **Save Changes** button at the bottom of the page.

Note: The fields that accept input cannot contain more than 250 characters.

- **Custom Name** - A name you assigned to the element. See “Assigning a Custom Name” on page 377 for more information.
- **Date Created** - Date the element was discovered.
- **Date Last Modified** - Date the record was last modified.
- **Description** - A description of the element. This description cannot be more than 250 characters.
- **Status** - The current status of the element. If the status of the element has changed, select the new status from the Status drop-down menu.
 - ☐ **New** - Default category for all detected elements.
 - ☐ **Missing** - No longer detectable through discovery
 - ☐ **Repaired** - The element is being repaired. The software does not automatically select this status.
 - ☐ **In Use** - The element is in use.
- **Vendor** - The vendor for the element.
- **Model** - The model of the element.
- **Serial Number** - Serial number of the element.
- **Barcode Number** - The barcode on the device.
- **Asset Code** - The asset code assigned to the element.
- **Asset Type** - The asset type assigned to the element.
- **Asset Tag** - The asset tag assigned to the element.
- **Asset Category** - The asset category assigned to the element.
- **Geographic Location** - The location of the element, for example, Boston, Massachusetts.
- **(Storage Systems Only) Storage Tier Classification** - Click the **Set Storage Tier Cost** link to set up storage tiers. See “Defining Storage Tiers” on page 721 for more information.

About the Collectors Tab

The management server uses collectors to gather information. The **Collectors** tab provides information about the collectors for a particular element.

To start collectors and view reports for an element:

1. Access the **Collectors** page by doing one of the following:
 - ❑ Clicking an element in Application Explorer, and then clicking the **Collectors** tab. (Click the the **Scan Schedule** tab for file servers).
 - ❑ Double-clicking an element in Capacity Explorer, Performance Explorer or System Explorer, and then clicking the **Collectors** tab.
 - ❑ Clicking an element in Chargeback, and then clicking the **Collectors** tab.
2. To change a collector's start time, modify the time and date entered in the **Next Scheduled Run** field. If you decide to change the start time, make sure the date is in the yyyy-mm-dd format with the time resembling a 24-hour clock. There should be a space between the date and the time, as shown below:
`2005-06-26 09:41`
After the collector runs, the value in this column is updated to the next time the collector will run.
3. To change how often the collector runs, type the number of minutes in the **Interval** field. **Important:** Do not make the interval too short. Running a collector too frequently uses up space on the management server and impacts its performance.
4. To enable the collector, click the **Start** button.
5. To stop a collector, click the **Stop** button.
6. To view a report, click its link. See “Viewing Reports” on page 563 for more information.

About the Monitoring Tab

You can easily access performance information about an element by doing the following:

1. Double clicking the element in System Explorer or Application Explorer.
 2. Click the Monitoring tab.
The element appears highlighted in Performance Explorer.
 3. Select one of the monitoring options in the lower pane to view specific performance data about the element.
 4. See Chapter 14, “Viewing Performance Data” on page 629 for more information about Performance Explorer.
-

About the Policies Tab

The Policies tab lets you view the utilization policies for an element. Utilization policies can automatically send an e-mail, generate an event, or run a custom script when an element is being overused. If the policy table is unpopulated, no policies exist for the element.

The Policies tab lets you use Policy Manager to do the following. See “About Policy Manager” on page 682 for more information:

- **Add Policies**
- **Test Policies**
- **Edit Policies**
- **Delete Policies**

You can access the Policies tab by doing one of the following:

- Double-clicking an element in Capacity Explorer, Performance Explorer or System Explorer and then clicking the **Policies** tab.
- Right-clicking an element in Capacity Explorer, Performance Explorer or System Explorer and then selecting **Show Policies** from the drop-down menu.

To access utilization policies for other elements and to create other types of policies, click the **Policy Manager** button in the left pane.

Determining If a Host Belongs to a File System

You can determine if a host is a member of a file system, such as CXFS™ on the Navigation tab and/or in Capacity Explorer.

To find the information on the Navigation page:

1. Access System Explorer as described in “Accessing System Explorer” on page 298.
 2. Double-click the host you want to determine is a member of a file system.
 3. Click the Navigation tab.
 4. Click the **Storage Volumes** button.
-

5. Look under the File System Type column. The system type, such as CXFS, is listed. Additional information about the storage volume is provided:
 - ☐ Name of the storage volume
 - ☐ Description of a storage volume
 - ☐ Drive Type

To find the information in Capacity Explorer:

1. Access Capacity Explorer as described in “Accessing Capacity Explorer” on page 662.
2. Select the host you want to determine is a member of a file system.
3. Scroll to the bottom of the page.
4. If a storage volume is a member of a shared file system, such as CXFS or XFS, it is displayed in the Storage Volume column.
You may need to expand the Storage Volume column if the volume names are long.

About the Data from CXFS File Systems

The management server can only monitor CXFS file systems from the host generating the input/output. For example, assume the elements in the following figure are part of a CXFS file system. When you generate input/out into the metadata server into /folder, only the metadata server is able to monitor the file system. For example, assume the metadata server generates 100 KB write, the management server displays 0 KB write for /folder on the metadata client.

The information in /folder on the metadata server is actually being mirrored to the /folder on the metadata client. The management server, however, does not detect the changes being mirrored to /folder on the metadata client.




Figure 9-32: CXFS File System

Important: Depending on your license, Provisioning may not be available. See the “List of Features” to determine if you have access to Provisioning wizards. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center**).

This chapter describes the following:

- About Provisioning on page 399
- Managing Zones on page 401
- Managing Storage on page 426

About Provisioning

The software provides the following tools to assist you in provisioning your storage. These tools are accessible by clicking **Provisioning** ().

- **Path Provisioning tool** - Lets you schedule provisioning tasks to take place when the network traffic is light. See “About Path Provisioning” on page 474 for more information.
- **SAN Zoning tool** - Lets you create and modify zones, zone aliases and zone sets. Click the **Provisioning** button next to the fabric on which you want to do provisioning. You can also view the properties of a fabric or switch by clicking its link in the table. See “SAN Zoning Overview” on page 401 for more information.
- **Storage System Provisioning tool** - Lets you manage storage pools, volumes, and host security groups. Click the **Provisioning** button next to the storage system on which you want to do provisioning. You can also view the properties of a storage system by clicking its link in the table. See “Setting Up Storage Partitioning” on page 427.

Important: Ports designated as an Initiator on a storage system belonging to the HDS Freedom Storage™ Lightning 9900™ Series or Freedom Storage Lightning 9900V Series cannot be used for provisioning. If you select one of these ports, you receive a message saying that provisioning failed because the HiCommand Database was not refreshed. The management server does not support provisioning for ports designated as Initiators on these storage systems.

Once you have become adept at provisioning, you might want to try accessing the provisioning screens by doing one of the following:

- Double-clicking a storage system or switch in System Explorer and then clicking the **Provisioning** tab.
- Right-clicking a storage system or switch in the Access tab in System Explorer. See “About the Access Tab” on page 307 for more information.

About Provisioning Brocade Switches After Upgrading

After you upgrade the management server, perform Get Details for any subset of elements that includes the Brocade switch before performing any provisioning operations that involve that Brocade switch. See “Step 4 - Get Details” on page 94 for details.

Managing Zones

This section describes the following:

- SAN Zoning Overview on page 401
- Accessing Information About Zone Aliases on page 408
- Creating Zone Aliases on page 409
- Modifying a Zone Alias on page 410
- Deleting Zone Aliases on page 411
- Accessing Information About Zoning on page 411
- Creating a Zone in a Fabric on page 413
- Adding and Removing Zone Members on page 415
- Deleting Zones on page 415
- Accessing Information About Zone Sets on page 416
- Creating a Zone Set on page 417
- Modifying a Zone Set on page 419
- Deleting Zone Sets on page 419
- Copying a Zone Set on page 420
- Activating a Zone Set on page 421
- Zones and Zone Sets Are Sometimes Listed Twice on page 423
- Changing the Amount of Information Collected from the Inactive Zone Database (Cisco Switches) on page 424

SAN Zoning Overview

Important: Depending on your license, Provisioning may not be available. See the “List of Features” to determine if you have access to Provisioning. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center**).

Use SAN zoning to control what can be seen in the storage area network (SAN). SAN zoning lets you group elements into zones, which can then be grouped into active and inactive zone sets. Only

elements in an active zone set can be seen. A switch fabric can have multiple zone sets, but only one zone set can be active.

Zones are an excellent way to split hardware resources because they work by exclusion. For example, you can set up your switch ports so that elements connected to some of the ports appear in one zone and the rest appear in another zone. Members of a zone can only communicate with other members in the zone. If two elements are not within the same zone, they cannot communicate.

Zones are usually created for a particular task, such as controlling access between devices or groups. You might create zones based on an application or an operating system. For example, some network administrators prefer to put all of the Microsoft Windows computers in one zone and all of the Sun Solaris computers in another. As mentioned previously, you can create zones according to an application. For example, you might want to create a zone for production and another zone for finance. This way the users in the finance department are not even aware of the disks and ports available for production and vice versa.

Only elements in an active zone set can communicate with each other. When a zone set is not active, it does not have any effect. If we do not want users in the Production and Finance zones to have access to the same storage, these two zones must be in two different zone sets, which must be both active. Since you can only have one active zone set to a fabric, the Production zone belongs to a zone set in one fabric and the Finance zone belongs to another zone set in another fabric.

A zone can be in more than one zone set. For example, the Finance zone could be in an active and inactive zone set. This allows for more flexibility. For example assume the Finance zone is a member of an active zone set named Zone Set One and a member of an inactive zone set named Zone Set Two in the same fabric. Zone Set Two contains additional zones. Assume you want to make users aware of those elements in the additional zones. When you activate Zone Set Two, users would still have access to the elements in Finance zone because it is also a member of Zone Two.

Note: Create zone aliases to easily keep track of your zones. Instead of remembering a port's name, you can assign a meaningful name. As a best practice, a zone should contain either zone aliases or ports, but not both.

The SAN Zoning tool is able to manage the two types of zoning:

- **Switch Port Zoning (also known as hard zoning)** - A hard zone is created by assigning a domain/port to a zone. Any device attached to the port is automatically in the zone.
-

- **WWN Zoning (also known as soft zoning)** - A soft zone is created by assigning a world wide name (WWN) of a device port to a zone.

The following figure provides an example of hard zoning. Ports 1 through 5 on the switch are assigned to a zone for production and ports 4 through 8 are assigned to the zone for the finance department. Users in the finance department can access storage systems B and C but not storage system A. Likewise, users in the production department can access storage systems A and C, but not storage system B. The following figure is an example of switch port zoning because the zones are assigned to ports on a switch.




Figure 10-1: Resources in Two Zones

The following diagram provides an overview of zoning:




Figure 10-2: Overview of Zoning Capabilities

Description

- **Zone Sets** - A zone set is a collection of zones. You can have only one zone set active at a time in a fabric; however, you can have a zone in more than one zone set. Zones sets are usually created for a particular task.
- **Zone** - A collection of zone aliases and ports.
- **Ports** - The Worldwide Name of the port to which an element is connected. A port in a zone can be of two types:
 - switch port WWN
 - WWN of the connected element

Use the following table as a guideline for setting up zoning.

Table 10-1: Setting Up Zoning

Step	Description	Where to Find
1	(Optional) Create a zone alias. Zone aliases are used to give meaningful names to switch ports, HBA ports, or storage system ports.	"Creating Zone Aliases" on page 409
2	Create zones. Zoning is the primary tool to constrain groups of SAN members. A zone defines a logical SAN that contains limited element membership. The only elements visible to members of a zone are other members of that zone.	"Creating a Zone in a Fabric" on page 413
3	Create zone sets. A zone set contains multiple zones.	"Creating a Zone Set" on page 417
4	Activate a zone set. A switch fabric can have multiple zone sets defined, but only one zone set can be active.	"Activating a Zone Set" on page 421

Keep in mind the following:

- If you use another product to make zoning changes, such as add a zone, you must perform Get Details for the management server to be made aware of these changes.
- The management server creates a zone by finding the port or Worldwide Name of discovered elements. The management server cannot create a zone by Fibre Channel addresses. An active zone will appear empty in the user interface of the management server if you use third-party software to create the zone by Fibre Channel addresses.

Use the following table for more information about what is supported for each type of switch.

Table 10-2: Zoning Support

Switch Type	View Active Zones	View Inactive Zones	View Aliases	Zone/ Zone Set Provisioning¹	Zone Set Copying	Port Statistics
Brocade	Y	Y	Y	Y	N	Y
² McDATA SWAPI to EFCM	Y	Y	N	Y	Y	Y
² McDATA SNMP through proxy	N	N	N	N	N	Y
² McDATA SNMP to switches	Y	N	N	N	N	Y
Cisco SNMP	Y	N	Y	N	N	Y
Cisco SMI-S ⁵	Y	Y	Y	Y	N	Y
CNT SMI-S	Y	Y	N	N	N	Y
⁴ QLogic SNMP	Y ³	N	N	N	N	Y
⁴ QLogic SMI-S	Y	Y	Y	Y	N	Y

¹The ability to create, modify, and remove zone aliases, zones, and zone sets.

²Also applies to EMC Connectrix switches.

³Also applies to Sun StorEdge SNMP switches.

⁴SANbox-2 only

⁵Soft zoning is only support for Cisco SMI-S switches. See “soft zone” on page 841 in the Glossary for more information. Hard zoning is not supported on Cisco SMI-S switches.

Issues for Sun StorEdge and QLogic SNMP (Simple Network Management Protocol) switches:

- Active zoning is not reported for Sun StorEdge and QLogic switches that have the SNMPv1 Agent. Switches that do report active zoning, do not supply information about inactive zones to the management server.
- If a zone alias for a Sun StorEdge or QLogic switch is a member of an active zone, the zone alias is not shown in the user interface, but its members are displayed as belonging to the active zone.

Issues for QLogic SMI-S switches:

- The management server prefixes the names of zone aliases, zones, and zone sets on QLogic switches with the name of the switch. For example, assume you have the name `qlogic1:DeviceName` displayed under the **Name** column.
 - `qlogic1` is the name of the switch containing the zone alias, zone or zone set.
 - `DeviceName` is the name the user assigned to the zone alias, zone or zone set.

Issues for Cisco SMI-S (Storage Management Initiative Specification) switches:

- Cisco SMI-S switches use virtual storage area networks (VSANs). VSANs are virtual fabrics that are subsets of physical Fibre Channel switch networks.
- The management server prefixes the names of zone aliases, zones, and zone sets on Cisco switches with the name of the Virtual SAN and the switch the they sets reside on. For example, assume you have the name `VSANNAME:cisco1:DeviceName` displayed under the **Name** column.
 - `VSANNAME` is the name of the virtual SAN in which the zone alias, zone, or zone set was created
 - `cisco1` is the name of the switch containing the zone alias, zone or zone set.
 - `DeviceName` is the name the user assigned to the zone alias, zone or zone set.

Issues for McDATA and Connectrix switches:

- Only one client at a time can provision on a McDATA or Connectrix fabric. Since each fabric has a separate lock, you can perform simultaneous provisioning on two different fabrics. For example, you could perform provisioning by using the user interface and the CLI at the same time on two different fabrics. Simultaneous provisioning on the same fabric is not supported.

- The management server does not support enabled default zones on McDATA or Connectrix switches. When a default zone is enabled on a McDATA or Connectrix switch, it is not listed as part of the active zone set.

Accessing Information About Zone Aliases

The software provides a listing of zone aliases in a fabric. You can view the properties of the zone alias and its port from this page.

Keep in mind the following:

- See Table 10-2, “Zoning Support,” on page 406 for more information about which zoning features are supported for your switches.
- (Cisco SIM-S Switch Connections Only) Zone aliases that are part of an active zone are listed multiple times for each switch in the virtual SAN. When a zone set is activated, the zone alias is copied to each switch in the virtual SAN. The zone alias is then listed twice (active and non-active versions) for each switch.

To obtain information about zone aliases in a fabric:

1. Click **Provisioning** ().
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric on which you want to do provisioning.
4. Click **Step 1 Zone Alias**.
This page lists the zone aliases and their ports under the following columns:

- **Name** - Click the name of the zone alias to view its properties.
Keep in mind the following:
 - (Cisco SMI-S switches) The management server prefixes the names of zone aliases on Cisco switches with the name of the Virtual SAN and the switch the zone aliases reside on. For example, assume you have the name VSANNAME:cisco1:Alias35 displayed under the **Name** column. VSANNAME is the name of the virtual SAN in which the zone alias was created, cisco1 is the name of the switch containing the zone alias, and Alias35 is the name the user assigned to the zone alias.
 - (QLogic SMI-S switches) The management server prefixes the names of zone aliases on QLogic switches with the name of the switch. For example, assume

you have the name Qlogic1:Alias displayed under the **Name** column. Qlogic1 is the name of the switch, and Alias is the name the user assigned to the zone alias.

- ❑ **Ports** - In some instance, you may be able to click the link of a port to view its properties.
- ❑ **Active** - A check mark appears in the Active column if the zone alias is included in an active zone set.
- ❑ **Edit** - Click the **Edit** button () to edit an alias. See “Modifying a Zone Alias” on page 410 for more information.
- ❑ **Delete** - Click the  button corresponding to the zone alias you want to delete. See “Deleting Zone Aliases” on page 411 for more information.

To create a zone alias, click the **New Zone Alias** button. See “Creating Zone Aliases” on page 409 for more information.

Creating Zone Aliases

Zone aliases are used to give meaningful names to switch ports, HBA ports, or storage system ports.

To avoid remembering a port's Worldwide Name (WWN), add a port to a zone alias as described in the following steps:

1. Click **Provisioning** ()
 2. In the right pane, click the **SAN Zoning** tab.
 3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to create a zone alias.
 4. Click **Step 1 Zone Alias**.
 5. Click the **New Zone Alias** button.
 6. (Cisco SIM-S Switch Connections Only) Select the virtual SAN in which you want to create the zone alias from the **VSAN** drop-down menu.
Only the ports in the VSAN you selected are displayed in the Potential Ports pane.
 7. (Cisco SIM-S and QLogic SMI-S Switch Connections Only) Select the switch on which you want to create the zone alias from the **Switch** drop-down menu.
 8. In the **Zone Alias Name** field, type a name for the zone alias.
- Naming Conventions for Brocade Switches:**
- ❑ The name must contain 1 to 64 characters.

- ❑ The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or an underscore (_).
- ❑ The name is case sensitive. For example, “ZoneAlias1” and “zonealias1” are different zone aliases.
- ❑ You cannot create a zone alias with the same name as an existing zone, zone alias or zone set. For example, if you create a zone alias named “new”, you cannot give a zone, zone alias, or zone set the same name.
- ❑ The following characters are invalid for Brocade switches: caret (^), dash (-), and dollar sign (\$).

Naming Conventions for McDATA and Connectrix Switches:

- ❑ The name can have a maximum of 64 characters.
 - ❑ The first character of a zone alias name must be a letter (A-Z, AZ).
 - ❑ A zone alias name cannot contain spaces.
 - ❑ Valid characters are a-a, AA, 0-9, caret (^), dash (-), underscore (_), and dollar sign (\$).
 - ❑ All names must be unique and may not differ by case. For example, myzonealias and MyZoneAlias are considered to be the same zone alias.
9. Add ports to the zone alias by selecting a port in the **Potential Ports** pane.

A port is not in the virtual SAN if the  icon is next to it.
 10. Remove ports from the zone by selecting them in the **Ports in the Zone Alias** pane and clicking the **Remove From Zone** button.
 11. Click **OK**.
The zone alias is created.

Modifying a Zone Alias

Zone aliases are used to give meaningful names to switch ports, HBA ports, or storage system ports.

To modify a zone alias:

1. Click **Provisioning** ().
2. In the right pane, click the **SAN Zoning** tab.

3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to modify the ports.
4. Click **Step 1 Zone Alias**.
5. Click the **Edit** button () .
6. Add ports to the zone alias by selecting a port in the **Potential Ports** pane.
7. Remove ports from the zone by selecting them in the **Ports in the Zone Alias** pane and clicking the **Remove From Zone** button.

Note: To select all of the ports, select the check box next to the Port heading.

8. Click **OK**.

Deleting Zone Aliases

You cannot delete a zone alias if it is the only member in a zone.

To delete a zone alias:

1. Click **Provisioning** () .
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to delete a zone alias.
4. Click **Step 1 Zone Alias**.
5. Click the  button corresponding to the zone alias you want to delete.
6. When you are asked if you want to delete the zone alias, click **OK**.

Accessing Information About Zoning

Keep in mind the following:

- See Table 10-2, “Zoning Support,” on page 406 for more information about which zoning features are supported for your switches.

- (Cisco SIM-S Switch Connections Only) Zone that are part of an active zone set are listed multiple times for each switch in the virtual SAN. When a zone set is activated, the zone is copied to each switch in the virtual SAN. The zone is then listed twice (active and non-active versions) for each switch.

To access information about zones and to be able to manage them, do the following steps:

1. Click **Provisioning** ().
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to access information about zoning.
4. Click **Step 2 Zone**.

This page lists the zones, their aliases and ports under the following columns:

- **Name** - Click the name of the zone to view its properties.
Keep in mind the following:
Keep in mind the following:
 - (Cisco SMI-S Switches) The management server prefixes the names of zones on Cisco switches with the name of the Virtual SAN and the switch the zones reside on. For example, assume you have the name VSANNAME:cisco1:ZoneName displayed under the **Name** column. VSANNAME is the name of the virtual SAN in which the zone was created, cisco1 is the name of the switch containing the zone, and ZoneName is the name the user assigned to the zone.
 - (QLogic SMI-S switches) The management server prefixes the names of zones on QLogic switches with the name of the switch. For example, assume you have the name Qlogic1:zone_name displayed under the **Name** column. Qlogic1 is the name of the switch, and zone_name is the name the user assigned to the zone.
- **Zone Aliases** - Click the name of the zone alias to view its properties.
- **Ports** - In some instance, you may be able to click the link of a port to view its properties.
- **Active** - A check mark appears in the Active column if the zone is included in an active zone set.
- **Edit** - Click the **Edit** button () to edit a zone. See “Adding and Removing Zone Members” on page 415 for more information.
- **Delete** - Click the  button corresponding to the zone you want to delete. See “Deleting Zones” on page 415 for more information.

To create a zone, click the **New Zone** button. See “Creating a Zone in a Fabric” on page 413 for more information.

Creating a Zone in a Fabric

To learn why zones are so important, see “SAN Zoning Overview” on page 401. A zone must have at least one member.

Keep in mind the following:

- Soft zoning is only support for Cisco SMI-S switches. See “soft zone” on page 841 in the Glossary for more information. Hard zoning is not supported on Cisco SMI-S switches.
- Depending on your license, Provisioning may not be available. See the “List of Features” to determine if you have access to Provisioning. The “List of Features” is accessible from the Documentation Center.
- See Table 10-2, “Zoning Support,” on page 406 for more information about which zoning features are supported for your switches.

To create a zone:

1. Click **Provisioning** ().
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to create a zone.
4. Click **Step 2 Zone**.
5. Click the **New Zone** button.
6. (Cisco SIM-S Switch Connections Only) Select the virtual SAN in which you want to create the zone from the **VSAN** drop-down menu.
Only the ports in the VSAN you selected are displayed in the Potential Members pane.
7. (Cisco SIM-S and QLogic SMI-S Switch Connections Only) Select the switch on which you want to create the zone from the **Switch** drop-down menu.
8. In the **Zone Name** field, type a name for the zone.
Naming Conventions for Brocade Switches:
 - ☐ The name must contain 1 to 64 characters.
 - ☐ The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or an underscore (_).

- ❑ The name is case sensitive. For example, “Zone1” and “zone1” are different zones.
- ❑ You cannot create a zone with the same name as an existing zone, zone alias or zone set. For example, if you create a zone named “new”, you cannot give a zone, zone alias, or zone set the same name.
- ❑ The following characters are invalid for Brocade switches: caret (^), dash (-), and dollar sign (\$).

Naming Conventions for McDATA and Connectrix Switches:

- ❑ The name can have a maximum of 64 characters.
- ❑ The first character of a zone name must be a letter (A-Z, AZ).
- ❑ A zone name cannot contain spaces.
- ❑ Valid characters are a-a, AA, 0-9, caret (^), dash (-), underscore (_), and dollar sign (\$).
- ❑ All names must be unique and may not differ by case. For example, myzone and MyZone are considered to be the same zone.

Naming Conventions for QLogic Switches

- ❑ Only alphanumeric characters and caret (^), dash (-), underscore (_), and dollar sign (\$) are allowed for zone sets, zones and alias names.
 - ❑ The first character of a zone set, zone, or alias must be alphanumeric (a-z or A-Z).
 - ❑ Zone sets, zones, and alias names are case sensitive.
 - ❑ A zone, zone set, and alias name can include up to 64 alphanumeric characters including caret (^), underscore (_), and dollar sign (\$).
9. Add members to the zone by selecting a member in the **Potential Members** pane. A zone member can be a port attached to a switch, a Worldwide Name (WWN) or a zone alias. Keep in mind the following:
- ❑ As a best practice, a zone should contain zone aliases only, and there should be a zone alias for each port/WWN.
 - ❑ You cannot create a zone with an existing name.

A port is not in the virtual SAN if the  icon is next to it.

10. Remove members from the zone by selecting them in the **Zone Members** pane and clicking the **Remove From Zone** button.
11. Click **OK**.
-

Adding and Removing Zone Members

Important: A zone must have at least one member.

To add and remove zone members:

1. Click **Provisioning** ().
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to modify a zone.
4. Click **Step 2 Zone**.
5. Click the **Edit** button ().
6. Add members to the zone by selecting a member in the **Potential Members** pane. A zone member can be a port attached to a switch, a Worldwide Name (WWN) or a zone alias.
7. Remove members from the zone by selecting them in the **Zone Members** pane and clicking the **Remove From Zone** button.
8. Click **OK**.

Deleting Zones

You cannot delete a zone if it is the only member in one of the zone sets or if it is a member of an active zone set. If you want to delete a zone in an active zone set, first move the zone to an inactive zone set. Then, delete it.

If you are using EFC Manager to delete zones, see “Changes in EFC Manager Requiring Get Details” on page 830.

To delete a zone:

1. Click **Provisioning** ().
 2. In the right pane, click the **SAN Zoning** tab.
 3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to delete a zone.
-

4. Click **Step 2 Zone**.
5. Click the  button corresponding to the zone you want to delete.
6. When you are asked if you want to delete the zone, click **OK**.

Accessing Information About Zone Sets

Keep in mind the following:

- See Table 10-2, “Zoning Support,” on page 406 for more information about which zoning features are supported for your switches.
- (Cisco SIM-S Switch Connections Only) Active zone sets are listed multiple times for each switch in the virtual SAN. When a zone set is activated, the zone set is copied to each switch in the virtual SAN. The zone set is then listed twice (active and non-active versions) for each switch.

To access information about zone sets and to be able to manage them, do the following steps:

1. Click **Provisioning** ().
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to access information about a zone set.
4. Click **Step 3 Zone**.
5. (Cisco SIM-S Switch Connections Only) Select the virtual SAN in which you want to view zone sets from the **VSAN** drop-down menu.

This page lists information about zone sets under the following columns:

- **Name** - Click the name of the zone set to view its properties.
Keep in mind the following:
 - (Cisco SMI-S switches) The management server prefixes the names of zone sets on Cisco switches with the name of the Virtual SAN and the switch the zone sets reside on. For example, assume you have the name VSANNAME:cisco1:ZoneSetName displayed under the **Name** column. VSANNAME is the name of the virtual SAN in which the zone set was created, cisco1 is the name of the switch containing the zone set, and ZoneSetName is the name the user assigned to the zone set.
 - (QLogic SMI-S switches) The management server prefixes the names of zone sets on QLogic switches with the name of the switch. For example, assume you have the name Qlogic1:Zone_set_name displayed under the **Name** column.

Qlogic1 is the name of the switch, and Zone_set_name is the name the user assigned to the zone set.

- ❑ **Zones** - Click the name of the zone to view its properties.
- ❑ **Active** - To make a zone set active, select its corresponding **Active** option. When you select a zone set, you make elements outside of the zone set inaccessible. See “Activating a Zone Set” on page 421 for more information.
- ❑ **Edit** - Click the **Edit** button () to edit a zone set. See “Modifying a Zone Set” on page 419 for more information.
- ❑ **Delete** - Click the  button corresponding to the zone set you want to delete. See “Deleting Zone Sets” on page 419 for more information.

To create a zone set, click the **New Zone Set** button. See “Creating a Zone Set” on page 417 for more information.

Creating a Zone Set

To learn why zone sets are so important, see “SAN Zoning Overview” on page 401.

To create a zone set in a fabric:

1. Click **Provisioning** ().
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to create a zone set.
4. Click **Step 3 Zone Set**.
5. (Cisco SIM-S Switch Connections Only) Select the virtual SAN in which you want to view zone sets from the **VSAN** drop-down menu.
6. Click the **New Zone Set** button.
7. (Cisco SIM-S Switch Connections Only) Select the virtual SAN in which you want to create the zone set from the **VSAN** drop-down menu.
Only the ports in the VSAN you selected are displayed in the Potential Members pane.
8. (Cisco SIM-S and QLogic SMI-S Switch Connections Only) Select the switch on which you want to create the zone set from the **Switch** drop-down menu.
9. In the **Zone Set Name** field, type a unique name for the new zone set.
Naming Conventions for Brocade Switches:
 - ❑ The name must contain 1 to 64 characters.

- ☐ The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or an underscore (_).
- ☐ The name is case sensitive. For example, "ZoneSet1" and "zoneset1" are different zone sets.
- ☐ You cannot create a zone set with the same name as an existing zone, zone alias or zone set. For example, if you create a zone set named "new", you cannot give a zone, zone alias, or zone set the same name.
- ☐ The following characters are invalid for Brocade switches: caret (^), dash (-), and dollar sign (\$).

Naming Conventions for McDATA and Connectrix Switches:

- ☐ The name can have a maximum of 64 characters.
 - ☐ The first character of a zone set name must be a letter (A-Z, AZ).
 - ☐ A zone set name cannot contain spaces.
 - ☐ Valid characters are a-a, AA, 0-9, caret (^), dash (-), underscore (_), and dollar sign (\$).
 - ☐ All names must be unique and may not differ by case. For example, myzoneset and MyZoneSet are both valid individually, but they are not considered to be unique.
10. If you want to make the zone set active, select the option **Activate this Zone**. Keep in mind the following:
- ☐ Only one zone set can be active at a time. When you make a zone set active, the previous active zone set becomes inactive.
 - ☐ You cannot create a zone set with the same name as the name of an existing zone set.
11. Add zones to the zone set by selecting a zone in the **Zones Not in this Zone Set** and clicking the greater than sign (>).

Note: Zones can be in multiple zone sets. For example, zone1 could be in several zone sets.

12. Remove zones from the zone set by selecting a zone in the **Zones in this Zone Set** and clicking the less than sign (<).
13. Click **OK**.
-

Modifying a Zone Set

To modify a zone set:

1. Click **Provisioning** ().
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to modify a zone set.
4. Do one of the following:
 - ☐ Click **Step 3 Zone Set**.
 - ☐ Click **Step 4 Activate this Zone Set**.
5. (Cisco SIM-S Switch Connections Only) Select the virtual SAN in which you want to view zone sets from the **VSAN** drop-down menu.
6. Click the **Edit** button ().
7. Add zones to the zone set by selecting a zone in the **Zones Not in this Zone Set** and clicking the greater than sign (>).
Note: Zones can be in multiple zone sets. For example, zone1 could be in several zone sets. Zone sets are usually created for a particular task.
8. Remove zones from the zone set by selecting a zone in the **Zones in this Zone Set** and clicking the less than sign (<).
9. Click **OK**.

Deleting Zone Sets

The software does not display all elements in a zone set, such as quick loop and fabric assist elements. When you delete a zone set, all elements, including quick loop and fabric assist, which are not viewable in the software, are deleted.

Only the zone set is deleted, not the zones contained in the zone set. For example, assume Zone A is contained in two zone sets: one named Zone_Set_One and another named Zone_Set_Two. If you delete Zone_Set_One, the zone has not been deleted so it is still in Zone_Set_Two.

If you are using EFC Manager to delete zone sets, see “Changes in EFC Manager Requiring Get Details” on page 830.

To delete a zone set:

1. Click **Provisioning** ().
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to delete a zone set.
4. Do one of the following:
 - ☐ Click **Step 3 Zone Set**.
 - ☐ Click **Step 4 Activate Zone Set**.
5. (Cisco SIM-S Switch Connections Only) Select the virtual SAN in which you want to view zone sets from the **VSAN** drop-down menu.
6. Click the  button corresponding to the zone set you want to delete.
7. When you are asked if you want to delete the zone set, click **OK**.

Copying a Zone Set

This feature copies a zone set and all of its members, such as zones and zone aliases. You can use this feature to copy inactive and active zone sets. The newly created zone set is inactive.

The management server stops the copying process of an active zone set if the management server finds one of the following:

- An inactive zone set with the same name as the name entered for the copy.
- An inactive zone with the same name as an active zone, but they do not have the same content.

Active zones in a zone set have corresponding inactive zones for redundancy. If you attempt to copy an active zone set containing a zone that does not have a corresponding inactive zone, the management server creates an inactive zone with the same name as the active zone. The inactive zone is used as a backup for the active zone.

Important: This feature is supported only for switches that support zone set copying. Refer to Table 10-2, “Zoning Support,” on page 406 for information on which switches support zone set copying. Aliases in the zone set are not copied over for McDATA switches.

To copy a zone set:

1. Click **Provisioning** ().
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the Provision button corresponding to the fabric in which you want to copy a zone set.
4. Click **Step 3 Zone Set**.
5. Click the corresponding  button for the zone set you want to copy.
6. Enter the name of the new zone set. If you are copying an active zone set, make sure you do not enter the name of a pre-existing inactive zone set.

Naming Conventions for McDATA and Connectrix Switches:

- ☐ The name can have a maximum of 64 characters.
 - ☐ The first character of a zone set name must be a letter (A-Z, AZ).
 - ☐ A zone set name cannot contain spaces.
 - ☐ Valid characters are a-a, AA, 0-9, ^, -, _, and \$.
 - ☐ All names must be unique and may not differ by case. For example, myzoneset and MyZoneSet are both valid individually, but they are not considered to be unique.
 - ☐ (Brocade Switches) You cannot create a zone set with the same name as an existing zone, zone alias or zone set. For example, if you create a zone set named “new”, you cannot give a zone, zone alias, or zone set the same name.
7. Click **OK**.
The zone is copied.

Activating a Zone Set

You can only have one zone set in a fabric active at a time; however, you could have a zone in more than one zone set.

Keep in mind the following:

- (Cisco SIM-S Switch Connections Only) Active zone sets are listed multiple times for each switch in the virtual SAN. When a zone set is activated, the zone set is copied to each switch in the virtual SAN. The zone set is then listed twice (active and non-active versions) for each switch.
- See Table 10-2, “Zoning Support,” on page 406 for more information about which zoning features are supported for your switches.

To make a zone set active:

1. Click **Provisioning** ().
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to activate a zone set.
4. Click **Step 4 Activate Zone Set**.
5. (Cisco SIM-S Switch Connections Only) Select the virtual SAN in which you want to view zone sets from the **VSAN** drop-down menu.
6. Select the corresponding **Active** option.
7. (McDATA and Connectrix switches) The management server lets you create a backup copy of the zone set you want to activate. To create a backup of the zone set that will become active, do the following:
 - a. Select the option, **Make a backup copy of the active zone set after activation**.
 - b. (Optional) In the **Name** field, modify the name that has been assigned to the backup zone set. The management server assigns the name by appending the date and time of the zone set you have selected to become active, as shown in the following example:

zone_name_2005-05-17_13-41-05
where

- zone_name is the name of the zone you are making active.
- 2005-05-17 is the date you made the zone active. The date is formatted as year-month-day. In the example, the date is May 17, 2005.
- 13-41-05 is the time the copy was made. The time is formatted as hour-minute-second, and it uses the 24-hour notation. In the example, the time is 1:41:05 p.m.

The name of the backup zone set must follow the following naming conventions for McDATA and Connectrix switches:

- The name can have a maximum of 64 characters.
 - The first character of a zone set name must be a letter (A-Z, AZ).
 - A zone set name cannot contain spaces.
 - Valid characters are a-a, AA, 0-9, caret (^), dash (-), underscore (_), and dollar sign (\$).
 - All names must be unique and may not differ by case. For example, myzoneset and MyZoneSet are both valid individually, but they are not considered to be unique.
-

Note: The management server truncates the name of the backup zone set if it is more than 44 characters long. For example, assume you have a zone set named `McDATA_Switches_Burlington_Massachusetts_United_States`. The management server truncates the name to `McDATA_Switches_Burlington_Massachusetts_Un_2005-05-17_13-41-05` when it creates the backup zone set because the original name was more than 44 characters long. The management server truncates the original name so that it can fit the date and time into the zone set name, which cannot be more than 64 characters.

8. To activate the zone, click **OK**.

Important: Only one zone set can be active at a time. When you make a zone set active, the previous active zone set becomes inactive.

Zones and Zone Sets Are Sometimes Listed Twice

Sometimes the Navigation tab and the Provisioning pages list the zones and zone sets twice for McDATA and Connectrix switches. EFC Manager and Connectrix Manager contain an off line Zoning Library. This Zoning Library holds all zone sets, zones and zone members. When you activate a zone set, the zone set along with its zones and zone members is copied to the McDATA and/or Connectrix switches and activated. This creates an active copy, in addition to the saved copy that already exists in the Zoning Library. If you edit the saved copy of the zone or zone set in the Zoning Library, you have to reactivate the saved copy of the active zone set.

For example, assume you have the following information in the EFC Zoning Library:

```
ZoneSet A
    Zone A1
        ZoneMember A1a

ZoneSet B
    Zone B1
        ZoneMember B1b
```

If a user activates ZoneSetB, the existing information in ZoneSetB is copied to the switch and activated. The Zoning Library, however, still contains the older information. The following is what you would see in the Zoning Library:

```
ZoneSets:
    ZoneSet A
    ZoneSet B    (this is the inactive ZoneSet in the Zoning Library)
    ZoneSet B    (this is the active ZoneSet which is not the same as
the one above)
Zones:
    Zone A1
    Zone B1    (this is the inactive one)
    Zone B1    (this is the active one)
ZoneMembers:
    ZoneMember A1a
    ZoneMember B1b    (inactive)
    ZoneMember B1b    (active)
```

Notice ZoneSetB and its members are listed twice in the previous example. They are also displayed twice on the **Navigation** tabs and Provisioning pages. If you click Zone A1 in EFC Manager or Connectrix Manager, you are shown it has member ZoneMember A1a. The other zone, Zone B1, has ZoneMember B1b. The name of an item in the provider cannot be changed, because the management server is required to report the actual name of the element.

If you remove a zone from the Zoning library, the zone is shown only once in the user interface of the management server. The zone is not displayed at all in the management server if it is an inactive zone. Each time you make changes in the Zoning Library using EFC Manager or Connectrix Manager, you must perform Get Details for the management server to obtain the latest information.

Changing the Amount of Information Collected from the Inactive Zone Database (Cisco Switches)

You can change how much information the management server collects during Get Details from the inactive database for Cisco switches discovered through SNMP. The zone database is stored on

each switch. If the switches are configured to distribute zone database changes to all switches in the fabric, it is only necessary to get zone database information from one switch in each VSAN.

The following options are supported:

- **all** - Show all zone database information from all switches. This may produce duplicate information and result in large zone tables if zone databases are distributed and/or large.
- **primary (default)** - Show only zone database information from the primary switch in each VSAN. This option is the recommended setting if zone information is distributed to all switches in the fabric or if only the primary switch zone database is used. Note that the primary switch can be different for different VSANs.
- **none** - Do not display any zone database information. This option is not recommended.

To change the amount of information collected:

1. Click **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the following command. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
`cimom.cisco.displayZoneDatabase=primary`
4. Return to the Advanced page (**Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree).
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Change the value of `cimom.cisco.displayZoneDatabase`. Make your changes in the **Custom Properties** field. For example, to set the value of `cimom.cisco.displayZoneDatabase` to all, you would change the value from primary to all as shown in the following example:
`cimom.cisco.displayZoneDatabase=all`
7. When you are done, click **Save**.
8. Restart the service for the management server for your changes to take effect. While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

About the Messages Displayed in the Brocade Console

If you see the following messages in the Brocade Console, you can ignore them. It is just the management server talking to the Brocade switch and subscribing to the event service.

```
wwn: 10:0:0:60:69:12:1a:5, cnt: 1, ipclnts: 0
```

```
total number of api event proxy servers: 0
```

```
total number of api event proxy servers: 0
```

```
total number of api event proxy servers: 1
```

```
proxy server 0 Did: 16776194
```

```
wwn: 10:0:0:60:69:10:56:fb, cnt: 1, ipclnts: 0
```

```
total number of api event proxy servers: 1
```

```
proxy server 0 Did: 16776194
```

```
wwn: 10:0:0:60:69:10:56:fb, cnt: 1, ipclnts: 0
```

```
wwn: 10:0:0:60:69:12:1a:5, cnt: 1, ipclnts: 0
```

The Brocade Console is accessible by typing the user name and password for the switch in a Telnet window. Then, type `switchshow` at the prompt to view the status of the switch.

Managing Storage

This section describes the following:

- Setting Up Storage Partitioning on page 427
-

- Modifying the Cache Settings (Engenio and Sun 6130) on page 433
- Changing the Owner of a Volume (Engenio, CLARiiON and Sun 6130) on page 434
- Managing Storage Pools on page 434
- Managing Volumes on page 438
- Rules for Creating Host Security Groups on page 448
- Managing Host Security Groups on page 452
- General Provisioning Issues on page 461
- Provisioning Issues by Vendor on page 462

Setting Up Storage Partitioning

Each storage vendor treats storage partitioning differently. For example, Hitachi and EMC ship their storage systems with the volumes already created. Other storage vendors, such as Engenio ship their storage system as an empty array.

Despite the differences among storage systems, you can still use this product to manage your provisioning. Some tasks, such as volume creation, might create different results according to the type of storage system. To learn how host security groups are created on your storage systems, see “Rules for Creating Host Security Groups” on page 448.

To learn which storage systems are supported, see Table 10-3, “Provisioning and Pool Support,” on page 427 and Table 10-4, “Volume and HSG Support,” on page 429.

Table 10-3: Provisioning and Pool Support

Storage System	Storage Provisioning	Create/Delete Pool	Create Pool Using Settings³
3PAR	N	N	N
Engenio and Sun 6130	Y	Y	Y
CLARiiON	Y, ¹	Y	Y
Symmetrix	Y	N	N

Table 10-3: Provisioning and Pool Support (Continued)

Storage System	Storage Provisioning	Create/Delete Pool	Create Pool Using Settings ³
HDS	Y	N	N
Sun 35xx	N	Y	Y
Sun 6920 and 6940	Y	N	N
HP XP with Command View XP	N	N	N
HP XP with Command View XP Advanced Edition	Y	N	N
HP XP with XP Provider for SE	Y	N	N
HP MSA	Y	N	N
HP EVA	Y	N	N
IBM DS and ESS ²	Y	N	N
Xiotech	Y	N	N

Note from table:

¹The EMC Navisphere® CLI is required to communicate with the CLARiiON storage system. The CLARiiON storage system must be configured to recognize the management server as a privileged user. See “EMC Navisphere CLI Is Required” on page 463 for more information.

²IBM and other storage systems that use external providers show raw capacity for a storage pool, instead of formatted capacity.

³The “Create Pool Using Settings” column refers to the functionality that lets you choose the type of pool, usually RAID level.

See the following table for information regarding support for volumes and host security groups.

Table 10-4: Volume and HSG Support

Storage System	Create/ Delete Volume	Create Volume Using Settings⁵	Create/ Delete Meta Volume	HSG Provisioning Supported
3PAR	N	N	N	N
Engenio and Sun 6130	Y	Y	N	Y
CLARiiON	Y	Y ⁴	Y	Y
Symmetrix	Y ¹	N	Y	Y ²
HDS	Y ¹	N	Y	Y
Sun 35xx	N	N	N	Y
Sun 6920 and 6940 ⁶	Y	N	N	Y
HP XP with Command View XP	N	N	N	N
HP XP with Command View XP Advance Edition	Y	N	Y	Y
HP XP with XP Provider for SE	Y ¹	N	Y	Y
HP MSA	Y	Y	N	Y
HP EVA	Y	Y	N	Y
IBM DS ⁷	Y	N	N	Y
IBM ESS ⁷	Y ³	N	N	Y
Xiotech	Y	Y	N	Y

Notes from table:

¹Creating volume means marking existing device as accessible to HSG management, deleting volume means returning device to free device pool.

²Solutions Enabler 5.0 or later is required.

³Volumes created on IBM storage systems cannot be deleted.

⁴RAID level can be specified for first volume in a pool, subsequent volumes inherit this setting

⁵The “Create Volume Using Settings” column refers to the provisioning capability on the management server that lets you create different kinds of volumes from a pool depending on whether you want it optimized, such as for streaming, random access or high availability.

⁶Unlike the native tool for the Sun 6130 storage system, the management server refers to virtual disks as storage pools.

⁷Volume groups are not represented in the provisioning or properties of IBM storage systems. The HSGs are based on a host connection, which does not give any indication as to which volume groups each of them belongs. Volume group can have many host connections, but only one host connection can be part of a volume group.

See the following topics for more information about how storage provisioning works on your storage system.

- Issues Specific to CLARiiON Storage Systems on page 462
 - Issues Specific to EMC Symmetrix Storage Systems on page 463
 - Issues Specific to HDS Storage Systems on page 465
 - Issues Specific to Engenio Storage Systems on page 470
-

The following table provides an overview of the required steps for setting up storage partitioning.

Table 10-5: Setting Up Storage Partitioning

Step	Description	Where to Find
1	(Engenio, CLARiiON, Sun 6130 and Sun 35xx) Create a storage pool. Sometimes referred to as a volume group or RAID group. A storage pool is a group of disks associated together through a RAID configuration. The pool's capabilities define the level of protection for the associated volumes and LUNs. EMC Symmetrix and HDS storage systems have storage pools that are predefined. Engenio and CLARiiON storage systems require a storage pool to be created and volumes to be allocated from a storage pool.	"Creating a Storage Pool (Engenio, CLARiiON, Sun 6130 and Sun 35xx)" on page 434

Table 10-5: Setting Up Storage Partitioning (Continued)

Step	Description	Where to Find
2	Create a volume. A volume is a virtual disk. Volumes are created in sizes that are desirable for being shown as a LUN. On CLARiiON storage systems, a volume is owned by one of the storage processors. Creating a volume also creates a LUN for this volume and each port of the storage processor that owns this volume. A volume can be associated with more than one fibre channel port, creating multiple LUNs corresponding to the same volume. The defining characteristics of a LUN are the volume, port, and LUN number. On CLARiiON storage systems, a LUN mapped to a port is visible to all the ports on that controller. Mapping a volume to a port on a CLARiiON storage system also maps that volume to all ports that reside on the same storage processor as the selected port. It also causes the volume to be unmapped from all the ports of the other storage processor. Some storage systems have their volumes fully configured during install. For these storage systems, Users can concatenate multiple volumes together to create a new volume.	“Creating a Storage Volume” on page 441
3	Create a Host Security Group. Host Security Groups define which initiators (HBA ports) have access to specified storage volumes. They are associated with a fibre-channel port and contain a list of HBA port initiators and the volumes they can see.	“Creating Host Security Groups” on page 454 and “Rules for Creating Host Security Groups” on page 448

Modifying the Cache Settings (Engenio and Sun 6130)

Important: Depending on your license, Provisioning may not be available. See the “List of Features” to determine if you have access to Provisioning. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center**).

To modify the cache settings:

1. Click the  button corresponding to the volume you want to modify.
2. Type the cache read ahead multiplier (0 to 65535 bytes) in the **Cache read-ahead multiplier** field.
A cache read ahead multiplier copies additional data blocks into the cache while it is reading and copying host-requested data blocks from disk to cache. To disable this option, type 0.
3. (Optional) Select **Read Caching**.
When this option is enabled, the host's operations are stored in controller cache memory.
4. (Optional) Select **Write Caching**.
When this option is selected, data is written to the cache memory of a controller
5. (Optional) Select **Write Caching with Mirroring**.
Use this option to preserve data if a controller or the cache fails. When this option is enabled, the data is written to two redundant controllers of the same cache size. This configuration provides redundancy in case a controller fails. One controller performs uncompleted write operations when the other controller fails.

Note: See “Changing the Owner of a Volume (Engenio, CLARiiON and Sun 6130)” on page 434 for information about changing the owner assigned to the volume.

Changing the Owner of a Volume (Engenio, CLARiiON and Sun 6130)

When a volume is created, the management server automatically assigns a controller to be the owner of the volume. You can change the owning controller if you want to use a different one for LUN masking.

Note: If the owner becomes unreachable as a result of a network failure or the owner itself fails, the other controller in the pair automatically becomes the owner of the volume.

To change the owner:

1. Click the  button corresponding to the volume you want to modify.
2. Select the owner for the volume from the **Current Owner** drop-down menu.
3. Click **OK**.

Managing Storage Pools

This section describes the following:

- Creating a Storage Pool (Engenio, CLARiiON, Sun 6130 and Sun 35xx) on page 434
- Accessing Information About Storage Pools on page 436
- Deleting a Storage Pool (Engenio and CLARiiON Only) on page 437

Creating a Storage Pool (Engenio, CLARiiON, Sun 6130 and Sun 35xx)

A storage pool is a group of disks associated together through a RAID configuration. The pool's capabilities define the level of protection for the associated volumes and LUNs. Create at least one storage pool before provisioning a volume.

To create a storage pool:

1. Click **Provisioning** ().
-

2. In the right pane, click the **Storage Systems** tab.
3. In the right pane, click the **Provision** button corresponding to the storage system in which you want to create a storage pool.
4. Click **Step 1 Storage Pool**.
5. Click the **New Storage Pool** button.
6. Select a setting and size for the storage pool.
The sizes displayed depends on the RAID level you want. For example, RAID 0 does not require additional drives so you can assign more space to the pool.




Figure 10-3: Selecting a Setting and Size for a Storage Pool

7. Click **OK**.
The storage pool is created.

Note: When you create a pool on an Engenio storage system, a placeholder volume is created inside the new volume group. The name of the placeholder volume starts with “Required - do not delete.” The placeholder volume is required because the storage pool cannot not exist without it. The management server does not display the placeholder volume, but other monitoring products may display this volume.

Accessing Information About Storage Pools

If you use another product to make provisioning changes, you must perform Get Details for the management server to be made aware of these changes.

To access information about storage pools:

1. Click **Provisioning** ().
2. In the right pane, click the **Storage Systems** tab.
3. In the right pane, click the **Provision** button corresponding to the storage system in which you want to access information about storage pools.
4. Click **Step 1 Storage Pool**.

This page lists information about the storage pools under the following columns:

- ☐ **Pool Name** - Click the name of the storage pool to view its properties.
- ☐ **Size** - Displays the amount of space assigned to the storage pool.
- ☐ **Available** - Displays the amount of space available in the storage pool.
- ☐ **Used** - Displays the amount of space used in the storage pool.
- ☐ **Volumes** - Click the name of the volume to view its properties. If the storage system has a large number of volumes, not all the volumes are displayed. To display all the volumes, select the **Show All** option.
- ☐ **Capabilities** - Click the RAID level or name to view its properties.
- ☐ ***Mainframe** - Displays whether the storage pools have volumes that are on a mainframe.
- ☐ ****Delete** - Click the  button corresponding to the storage pool you want to delete. See the topic, “Deleting a Storage Pool (Engenio and CLARiiON Only)” on page 437 for more information.

*This option is available to only HDS storage systems.

**These options are available to only Engenio, CLARiiON, Sun 6130, and Sun 35xx storage systems.

5. To create a storage pool, click the **New Storage Pool** button in the upper-right corner of the page. See the topic, “Creating a Storage Pool (Engenio, CLARiiON, Sun 6130 and Sun 35xx)” on page 434 for more information.
6. If the table contains more than 10 entries, navigation buttons are displayed. To navigate through the table, click one of the following:
 - ☐  - Goes to the first page.
 - ☐  - Goes one page back.
 - ☐  - Goes one page forward.
 - ☐  - Goes to the last page.

Deleting a Storage Pool (Engenio and CLARiiON Only)

Important: When you delete a storage pool on an Engenio storage system, all the volumes for the volume group are deleted, including the placeholder volume.

To delete a storage pool:

1. Click **Provisioning** ().
 2. In the right pane, click the **Storage Systems** tab.
 3. In the right pane, click the **Provision** button corresponding to the storage system in which you want to delete a storage pool.
 4. Click **Step 1 Storage Pool**.
 5. Click the  button corresponding to the storage pool you want to delete.
When you delete a storage pool, you are warned about the other volumes that will be deleted.
 6. Click **OK**.
 7. The storage pool and its volumes are deleted.
-

Managing Volumes

This section describes the following:

- Accessing Information About Volumes on page 438
- Filtering Volumes on page 440
- Creating a Storage Volume on page 441
- Deleting a Storage Volume on page 445
- Changing the Cache Block Size for a Storage System on page 446
- Modifying the Cache Settings (Engenio and Sun 6130 Only) on page 447

Accessing Information About Volumes

Important: Some storage vendors require a password to access the storage system. If the same password is not entered, an authentication error message is displayed. Refer to Get Details to determine the user account that was used to access the storage system during discovery.

If you use another product to make provisioning changes, you must perform Get Details for the management server to be made aware of these changes.

To access information about volumes:

1. Click **Provisioning** ().
2. In the right pane, click the **Storage Systems** tab.
3. Click the **Provision** button for the storage system you want to provision.
4. Click **Step 2 Volume**.
All volumes are displayed

To filter the volumes displayed, see “Filtering Volumes” on page 440.

The following columns list information about the volumes:

- ☐ **Volume** - Click the name of the volume to view its properties.
 - ☐ **Size** - Displays the size of the volume in megabytes (MB).
 - ☐ **Ports** - Click a ports link to view its properties.
-

- **Pool** - Click the name of the pool to view its properties.
- ***Mainframe** - Displays whether the storage pools have volumes that are on a mainframe.
- ****Default Owner** - The controller which owns the storage system when it is rebooted.
- ****Current Owner** - The controller which currently owns the volume.
- ****Segment Size** - Displays the amount of space assigned to a volume in megabytes (MB).
- ****Read ahead** - Displays the cache read ahead multiplier.
- ****Edit** - Click the  button corresponding to the volume you want to edit. See the topic, “Modifying the Cache Settings (Engenio and Sun 6130)” on page 433 for more information.
- **Delete** - Click the  button corresponding to the volume you want to delete. See the topic, “Deleting a Storage Volume” on page 445 for more information.

*This option is available to only HDS storage systems.

**Not accessible to all storage systems.

To create a volume, click the **New Volume** button in the upper-right corner of the page. To delete several volumes at once, select the volumes you want to delete and then click the **Delete Selected Volumes** button.

Note: See “About Volumes on HDS Storage Systems” on page 440 if you have an HDS storage system.

5. If the table contains more than 10 entries, navigation buttons are displayed. To navigate through the table, click one of the following:
-  - Goes to the first page.
 -  - Goes one page back.
 -  - Goes one page forward.
 -  - Goes to the last page.

Filtering Volumes

You can filter the volumes displayed.

To filter the list of volumes:

1. Click the **Show Volume Filter** link to display the filtering options.
You see the **Hide Volume Filter** link instead of the Show Volume Filter link if the volume filter is already displayed.
2. Enter all or part of a volume name in the in the **Name Contains** field.
3. Enter a size in megabytes in the **Size (>=) MB** field.
Volumes that are greater than or equal to the size specified appear on the page.
4. Select one of the following:
 - ☐ **All Volumes** - All volumes are displayed.
 - ☐ **Unmapped Volumes** - Only unmapped volumes are displayed.
 - ☐ **Mapped Volumes** - Only mapped volumes are displayed.
5. Select the port you want displayed from the **Ports** drop-down menu.
If you want to display all ports on the page, select **All** from the **Ports** drop-down menu.
6. Select the storage pools you want displayed from the **Storage Pools** drop-down menu.
7. Click the **Filter** button.
The table updates and displays only the elements that meet the filter criteria.

Note: If you want to reset the filter criteria, click the **Reset** button.

About Volumes on HDS Storage Systems

Volumes from single LDEVs are shown as `LDEV:0` on HDS storage systems. When volumes made up of multiple LDEVs are first created, they are not mapped to a target port on the storage system. The software remembers these LDEVs constitute a single volume, but it does not make changes to the storage system until the volume is mapped to a port. As a result, they are referred to as Groups, for example:

```
Group:0 (LDEV:0, LDEV:1)
```

where 0 in `Group:0` is the volume identifier and `LDEV:0` and `LDEV:1` are the LDEVs that make up this volume.

After you create a storage volume on an HDS storage system, you must map the volume to a target port on the storage system. You can map the volume to the target port by using the storage system Provisioning tool. In the tool, click **Step 3, LUNs**.

Once the volume is mapped, it is displayed as a logical unit size expansion (LUSE), as shown below:

```
LUSE:0 (LDEV:0, LDEV:1)
```

where 0 in `Group:0` is the volume identifier and `LDEV:0` and `LDEV:1` are the LDEVs that make up this volume.

Creating a Storage Volume

Important: Some storage vendors require a password to access the storage system. If the same password is not entered, an authentication error message is displayed. Refer to **Discovery > Details** to determine the user account that was used to access the storage system during discovery.

When you create a storage volume, you can set its size, volume capabilities, and storage pool.

Volumes are shipped already created on HDS and Symmetrix storage systems. When you create a volume in the management server on these storage systems, you are defining the volume as being allocated.

Keep in mind the following:

- Some vendor's tools for HDS might round off the volume size, so that a 6.87-GB volume appears as 7 GB (7168 MB) in the tool. The management server displays the size of the volume without rounding. For example, assume you want to create a 14-GB (14336 MB) LUSE volume, and according to the storage tool, you have two 7-GB LDEVs, which are really 6.87 GB (7034.88 MB). If you look at the native tool, it would be logical to assume only two LDEVs would be used to create the 14-GB LUSE volume. The management server would use three LDEVs because each LDEV is really 6.87 GB.
 - If you are creating a volume on an HP EVA storage system, its external SMI-S provider may round the specified number of megabytes to the nearest whole gigabyte.
 - You cannot create volumes in existing storage pools on HP MSA and EVA storage systems.
-

- PvLinks based on HP disk partitions are not supported. Any volumes created on such PvLink meta devices are shown as local. If you partition a regular (non-PvLink) external disk and create volumes based on it, then volumes are recognized as external volumes.

Keep in mind the following for HDS storage systems:

- A LUSE volume on an HDS storage system is not created until you map that volume to a target array port. In the management server, the create volume and LUN creation tasks are two different operations. So if you want to create a LUSE volume and perform LUN creation on the HDS box, it is a two-step process. First, use the management server to create LUSE volumes. The management server puts volumes in the repository as grouped LDEVs. Then, create a LUN and map a volume to a target port, by creating a host security group. See the topic, “Creating Host Security Groups” on page 454 and “Rules for Creating Host Security Groups” on page 448.
- HDS ships some of its storage systems with volumes already created. When the software first discovers an HDS storage system, it detects the volumes created by HDS. When you use the software's “create a volume” feature, you are assigning the already created volume. See the topic, “About Provisioning on HDS Storage Systems” on page 466 for more information.
- LUSE made up of volumes from different RAID levels are not supported. You cannot use this product to provision this type of LUSE. Existing LUSE of this type maybe incorrectly reported.

Keep in mind the following for Engenio storage systems:

- Make sure you select a volume group that can accommodate the requested size for the new volume.
- You can create volumes from existing free extent areas within a volume group.
- No volume-to-LUN masking is done by default, except for CLARiiON storage systems. See the topic, “Creating Host Security Groups” on page 454 and “Rules for Creating Host Security Groups” on page 448.

To create a storage volume:

1. Access the Create Storage Volume wizard by doing the following:
 - a. Click **Provisioning** ().
 - b. In the right pane, click the **Storage Systems** tab.
 - c. In the right pane, click the **Provision** button corresponding to the storage system in which you want to access information about volumes.
 - d. Click **Step 2 Volume**.

- e. If you want to create LUSE volumes on an HDS storage system, click the **Unmapped Volumes** folder in the left pane. Then, select the desired number of LDEVs for the LUSE volume then click the **Delete Selected Volumes** button. Take note of the array group from which you deleted the LDEVs. You need this information to create the LUSE volume.
2. Click the **New Volume** button.

Note: You can also access the Create Storage Volume wizard from the Navigation tab in System Explorer. To access the wizard from the Navigation tab, click the Volumes link for a storage system, then click the **New Volume** button at the bottom of the screen.

3. (Engenio, CLARiiON, and Sun 6130) In the **Volume Name** field, type a name for the volume. If you do not provide a name, the software assigns one.
If you type a volume name, keep in mind the following:
 - ☐ It cannot be more than 30 characters.
 - ☐ The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or one of the following symbols: dollar sign (\$), caret (^), or an underscore (_).
 - ☐ The name is case sensitive, for example, "StorageVolume1" and "storagevolume1" are different storage volumes.
 - ☐ The name must be different from any other volume name on an Engenio storage system.
4. In the **Size** field, type the size of the volume in megabytes (MB) or gigabytes (GB). Select the appropriate unit of measurement from the drop-down menu to the right of the **Size** field.

Important: The management server creates a volume of at least the size specified. For example, assume you requested that the management server create a 15-MB volume, and you have only three free extents: 10 MB, 17 MB, and 100 MB. A 17-MB volume will be created instead of a 15-MB volume because that is the closest size you have for a free extent. Keep in mind that although the management server tries to find free extents that make the volume size as close to the requested size, it does not guarantee it will pick the optimal combination of free extents.

5. Select a storage pool for the volume.
-

Keep in mind the following:

- ❑ If you do not see space available in the storage pool, you must delete volumes. For example, assume you want to create an 8-MB volume, but you do not have space available. Each volume is made up of 4 MB. You must delete two volumes from that storage pool. Make sure those volumes you delete are not being used. See “Deleting a Storage Volume” on page 445 for more information.
 - ❑ If you do not see a storage pool, verify that you have obtained all element details from the storage system. See “Step 1 - Discover Switches” on page 34 for more information.
6. Click **Next**.
 7. Select a volume capability. The volume capabilities listed depend on the type of storage system. For example, if an EMC Symmetrix storage system is selected, **Pool default settings** are displayed in the field. If an Engenio storage system is selected, the following is displayed.
 - ❑ **<Default>** - Provides the default cache read ahead multiplier and the default segment size for the storage system.
 - ❑ **File System (Typical)** - Provides a cache read ahead multiplier of 1 with a segment size of 64 KB.
 - ❑ **Database** - Provides a cache read ahead multiplier of 0 with a segment size of 64 KB.
 - ❑ **Multimedia** - Provides a cache read ahead multiplier of 8 with a segment size of 128 KB.
 - ❑ **Custom** - Lets you customize the cache read ahead multiplier and the segment size.
-

Note: (HDS only) Under the “Volume Capabilities Tab”, leave the default selection “<default>” selected and click the “Finish” button. Once the settings have been made, you see a new volume called a Group Volume in the list of unmapped volumes. Technically this is not a LUSE yet as it has not been assigned to a port; it is a logical grouping within the management server. Think of it as a place holder. From this point you can select the new group volume and assign it to a port. Once the volume has been assigned to a port, the management server makes the changes to the array and creates the LUSE. See “Creating Host Security Groups” on page 454 and “Rules for Creating Host Security Groups” on page 448 for information on how to create a LUN.

The HP-EVA specific volume capabilities are the following: default, RAID 0, RAID 1, and RAID 5.

For information about the volume capability, refer to the documentation accompanying the storage system.

8. (Engenio only) If you selected the **Custom** option, do the following:
 - a. Type the cache read ahead multiplier (0 to 65535 bytes) in the **Cache read ahead multiplier** field.
A cache read ahead multiplier copies additional data blocks into the cache while it is reading and copying host-requested data blocks from disk to cache. Select the multiplier that maximizes performance for the way the volume will be utilized.
 - b. Select a segment size from the drop-down menu.
9. Click **Finish**.

Deleting a Storage Volume

When you delete a storage volume on an HDS or Symmetrix storage system, the software marks the deleted volume as hidden in the CIM repository, thus, making it unassigned, instead of being deleted. The software keeps track of the “deleted volumes.”

Keep in mind the following:

- Some storage vendors require a password to access the storage system. If the same password is not entered, an authentication error message is displayed. Refer to **Discovery > Details** to determine the user account that was used to access the storage system during discovery.
- If you remove volumes from host storage groups that are command devices or are pair volumes on HDS storage systems, later modification of the pair volumes may be disabled.

To delete a storage volume:

1. Click **Provisioning** ().
2. In the right pane, click the **Storage Systems** tab.
3. In the right pane, click the **Provision** button corresponding to the storage system in which you want to access information about volumes.
4. Click **Step 2 Volume**.
5. To display only unmapped volumes, click the unmapped volumes node in the left pane.
6. Click the  button corresponding to the volume you want to delete.
7. When you are asked if you want to delete the volume, click **OK**.

8. To delete several storage volumes at once, select the storage volumes you want to delete and then click the **Delete Selected Volumes** button.




Figure 10-4: Deleting Several Volumes at Once

Note: To select all volumes, select the check box next to the Volume heading, as shown in the following figure:




Figure 10-5: Selecting All Volumes

Changing the Cache Block Size for a Storage System

Engenio and Sun 6130 storage systems let you change the cache block size on them.

To change the cache block size of a storage system:

1. Do one of the following:
 - Click **Provisioning** (). Then, click the link for the storage system.
 - Click **System Explorer** (). Then, double-click the storage system displayed in the right pane.
2. Scroll to the bottom of the Navigation page.
3. Click the **Change** button.
4. Select the cache block size from the drop-down menu. Set a higher cache size for applications that requires a lot of input and output, such as multimedia.
5. Click **OK**.

Modifying the Cache Settings (Engenio and Sun 6130 Only)

To modify the cache settings:

1. Click **Provisioning** ().
2. In the right pane, click the **Storage Systems** tab.
3. In the right pane, click the **Provision** button corresponding to the storage system in which you want to access information about volumes.
4. Click **Step 2 Volume**.
5. Click the **Edit** button () corresponding to the volume you want to modify.
6. Type the cache read ahead multiplier (0 to 65535 bytes) in the **Cache read-ahead multiplier** field.

A cache read ahead multiplier copies additional data blocks into the cache while it is reading and copying host-requested data blocks from disk to cache. To disable this option, type 0.
7. (Optional) Select **Read Caching**.

When this option is enabled, the host's operations are stored in controller cache memory.
8. (Optional) Select **Write Caching**.

When this option is selected, data is written to the cache memory of a controller
9. (Optional) Select **Write Caching with Mirroring**.

Use this option to preserve data if a controller or the cache fails. When this option is enabled, the data is written to two redundant controllers of the same cache size. This

configuration provides redundancy in case a controller fails. One controller performs uncompleted write operations when the other controller fails.

10. To change the current owner, select a new owner from the **Current Owner** drop-down menu.
11. Click **OK**.

Rules for Creating Host Security Groups

The management server now uses host security groups instead of LUN masking and LUN mapping. With the introduction of host security groups, the management server has a new definition of mapped for this release. Mapped is capacity that is accessible by one or more hosts external to the array (aggregated capacity of volumes that are accessible from hosts external to the subsystem).

The number of initiators allowed in the HSG depends on the type of storage system:

Table 10-6: Allowed Initiators in Host Security Groups

Storage System	Allowed Initiators in HSGs
EMC CLARiiON	0 or more initiators
Engenio (LSI)	0 or more initiators
HDS	0 or more initiators
HP EVA and MSA	0 or more initiators
IBM	only one initiator per host security group
Sun 35xx	One initiator per host security group
Sun 6920 and 6940	only one initiator per host security group
Sun 6130	0 or more initiators
Symmetrix	one initiator for host security masking
Xiotech	1 initiator

Note: For the “Volume Creation and LUN Security” option in Path Provisioning, the All Ports node is not shown because volumes cannot be placed inside HSGs for All Ports.

Each storage system treats host security groups differently. See the following sections for more information.

EMC CLARiiON and Sun 6130 Storage Systems

Upon the creation of volumes, by default the volume is assigned to one of the two controllers. Even though this volume is mapped to a controller, it is not visible from the outside world by a host. The management server reports this volume as unmapped since it is not visible by a host initiator.

- Volumes can be only on SPA or SPB because CLARiiON is active/passive storage, which means it can have only one active path to a volume. Addition of initiators to any of the ports on a storage processor is listed for all ports of that storage processor.
- The HSG is created on all ports of the processor you select unless you select an initiator that uses a different processor and the initiator does not belong to a host security group. For example, assume you select processor SPA. Then, you select an initiator that belongs to SPB but it does not belong to a host security group. The host security group is created for all ports on SPB.
- HSGs can consist of initiators (WWN) only. You do not need to specify volumes. The initiator is shown in both host security groups SPA and SPB.
- HSGs can consist of volumes (LUNs) only. You do not need to specify initiators.
- When you select an initiator for the host security group, the initiator has to be registered with the CLARiiON storage system.
- You can have more than one initiator in a security group if you have the proper multipathing software installed on that particular host where the initiator is located.

Engenio Storage Systems

- When you create the HSG by using the management server, the HSG appears as a host with its volumes and ports displayed underneath the tree in SANtricity.
- When you create the HSG by using the management server, you can specify the controller, but not the port for the HSG.
- An initiator is equivalent to a host port in SANtricity.
- You can have multiple volumes and initiators in the HSG.
- If you create a volume on a host and you have no multipathing on the host, make sure the volume is on the preferred path. You can make sure the volume is on the preferred path by using SANtricity.
- The ID for a host security group changes when you rename the host security group.
- When creating a host security group, if you provide a volume, but not an initiator, the host security group is created, but the volume is part of the Default Group, not the group you

created. This occurs because until a volume is assigned a host security group with an initiator, it is visible to all initiators and is assigned to the Default Group. To assign a volume to a specific host security group, you must add an initiator to that host security group.

EMC Symmetrix Storage Systems

If no LUN security is turned on for an FA port, all volumes assigned to the FC port are visible by hosts that are on the SAN and have been zoned by the SAN. All volumes assigned to the FC port appear in the mapped category.

When you create the HSG on a Symmetrix storage system, you are creating LUN mapping and masking in one step. In the native tools for Symmetrix storage systems, you will not see the HSG you created by using the management server. Instead you will see a volume binded to a port and a masked LUN to a host in the native tools.

- HSG is associated with individual ports
- HSGs only allow one initiator for host security masking.
- To create a host security group, you must specify a port, initiator and a volume.
- Every port has a LUN HSG, even if no LUNs are defined for that port. To bind a LUN to a port, edit the host security group and add the desired LUN to a port.
- You can also add LUNs to a Mask HSG. To add initiators, you need to create the HSG.

HDS Storage Systems

- FC port contains only volumes but no initiators (HBA WWN) assignment, the management server displays these volumes as unmapped since no external host can see these volumes yet.
- You can have 0 to multiple initiators in a host security group
- You can have 0 to multiple volumes in a host security group.
- Every HSG is only on one port on the array. You can have host security groups with the same name, as long as they are on different ports.
- Host security groups appear in the native tool for HDS storage systems in the logical view, listed by LDEV and in the physical view, listed by port.
- In the native tool for HDS storage systems, host security groups are referred to as a host security domain.
- See “Host Security Group Names on HDS Storage Systems” on page 466 for more information about HSG names.

HP MSA and EVA Storage Systems

- You can have one initiator per host security group
-

- You can have 0 to multiple volumes in a host security group.
- A host security group spans all ports on the array.

IBM Storage Systems

- To assign a host mode to a host security group, you must modify a property, as described in “Setting the Host Mode for IBM Storage Systems” on page 460.
- You can name a host security group on IBM storage systems. The host security group will be given the name of the initiator you select for the host security group.
- The management server can read the names of HSGs created by the native tool.
- Only one initiator per host security group.
- The volume can be assigned to more than one initiator.
- You can select any number of ports from one to all when creating the HSG.
- You can create the HSG with or without LUNs.
- You can add mapped and unmapped volumes to a host security group, but they should have the same host mode.
- The default is all ports are in the host security group. If no ports are selected, the default is used.

Sun 6920 and 6940 Storage Systems

- You can have two or more initiators assigned to the same volume, but you can have only one initiator to a host security group.
- You can select any number of ports from one to all when creating the HSG.
- You must create the HSG with one initiator and at least one LUN.
- When the management server creates the HSG, the name of the HSG is the initiator WWN.
- The management server can read the names of HSGs created by the native tool.
- Creation of a new HSG with an initiator that already exists in another HSG adds the volumes to the previously created HSG.
- The management server displays the ports connected to the initiator you selected for creating an HSG. If the initiator is not connected to any ports, no ports are displayed.
- If you select a number of ports, only the selected initiators will be part of the host security group.
- Only the storage ports connected to the initiator in the host security group will be included in the host security group, even if the default is all ports. For example, if an initiator is not connected to a storage system, no ports will be in the host security group.

Xiotech Storage Systems

- The initiators used in a host security group must be registered with the array.
- Any given HSG is assigned to only one storage port.
- Host security groups must have one initiator and at least one volume

Managing Host Security Groups

This section describes the following:

- Accessing Information About Host Security Groups on page 452
- Creating Host Security Groups on page 454
- Editing Host Security Groups on page 457
- Deleting Host Security Groups on page 460

Accessing Information About Host Security Groups

Host Security Groups define which initiators (HBA ports) have access to specified storage volumes. They are associated with a fibre-channel port and contain a list of HBA port initiators and the volumes they can detect.

Keep in mind the following:

- Each type of storage system treats host security groups differently. See “General Provisioning Issues” on page 461.
- Not all HDS storage systems support host security groups. Refer to the documentation accompanying the HDS storage system.

If you use another product to make provisioning changes, you must perform Get Details (**Discovery > Details**) for the management server to be made aware of these changes.

You can access information about host security groups from the Provisioning wizard or from System Explorer:

To access host security groups from the Provisioning wizard:

1. Click **Provisioning** ().
2. In the right pane, click the **Storage Systems** tab.

3. Click the **Provision** button corresponding to the storage system in which you want to access information about host security groups.
4. Click **Step 3 Host Security Group**.

To access host security groups from System Explorer:

1. Double-click a storage system in System Explorer.
2. Click **Host Security Groups** in the Navigation tab.

This page lists the following information about host security groups:

Note: The following two features are only available when accessed from the Provisioning wizard.

- **View all host security groups** - Click the **All** category in the tree. All the host security groups appear in the right pane.
- **View only host security groups assigned to a certain port** - Click a port in the tree. The host security group assigned to the port appears in the right pane.

This page lists information about host security groups under the following columns:

- **Port** (Displayed if applicable) - Lists the port associated with the host security group.

Note: You can filter the list of ports so that the information for only one port is displayed. To filter the list of ports: Click the **Show Filters** button, select the port that you would like to view, and click the **Apply** button. Only the information for the selected port is displayed in the table below.

- **Name** - The name of the host security group.
- **Initiators** - Displays one of the following:
 - The caption for the discovered port if the port has been discovered by the management server (for example Columbia:Adapter0 Port 0). A tool tip over the caption gives the full WWN.
 - The WWN if the port's not been discovered.
- **Volumes** - The volumes in the host security group.
- **Host Mode** - Displays the port settings for your operational environment. The settings for the host mode vary by the model of the HDS storage system. With some hardware, you must select a special host mode on the port for the storage system to enable certain

servers and HBAs to “see” the LUNs on the port. Refer to your documentation for the HDS storage system.

- **Host Mode 2** (Displayed if applicable) - Optional settings on the port that describe how the host accesses the port. Multiple options exist. Refer to your documentation for HDS storage system.
3. If the table contains more than 10 entries, navigation buttons are displayed. To navigate through the table, click one of the following:
- ☐  - Goes to the first page.
 - ☐  - Goes one page back.
 - ☐  - Goes one page forward.
 - ☐  - Goes to the last page.

You can also create, edit and delete host security groups from this page. See the following topics for more information:

- “Creating Host Security Groups” on page 454
- “Editing Host Security Groups” on page 457
- “Deleting Host Security Groups” on page 460
- “Setting the Host Mode for IBM Storage Systems” on page 460

Creating Host Security Groups

Host Security Groups define which initiators (HBA ports) have access to specified storage volumes. They are associated with a fibre-channel port and contain a list of HBA port initiators and the volumes they can see.

Keep in mind the following:

- If you are running the management server on Solaris and a Leadville driver looks like it has become corrupted (for example c4t4849544143484920523430314545364530303436pö d0), change the host mode from “standard mode” to “Solaris mode.”
 - You cannot use the management server to add a host to a “Host Group”. For example, you cannot have nested host groups.
 - In releases previous to build 4.0, host security groups were only supported for HDS storage systems. In this build, host security groups are now available to all storage systems that
-

support provisioning. Each storage system treats host security groups differently, see “Rules for Creating Host Security Groups” on page 448 for more information.

- Each type of storage system treats host security groups differently. See “General Provisioning Issues” on page 461.

To create a host security group:

1. Click **Provisioning** ().
2. In the right pane, click the **Storage Systems** tab.
3. In the right pane, click the **Provision** button corresponding to the storage system in which you want to access information about host security groups.
4. Click **Step 3 Host Security Group**.
5. Click the **New Host Security Group** button in the upper-right corner of the screen.

Step 1 - Add Details for the Host

1. Type a unique name for the host security group in the **Name** field.

Note: You cannot name a host security group on IBM storage systems. The host security group will be given the name of the initiator you select for the host security group.

Keep in mind the following:

- The name must contain the following number of characters. If you enter no characters, you are given the option of using a default name.
 - **CLARiiON storage systems** - 1 to 64 characters
 - **Engenio storage systems** - 1 to 30 characters
 - **Sun 6130 storage systems** - 1 to 30 characters
 - **All other storage systems** - 1 to 50 characters
- The first and last letter cannot be spaces
- You cannot have the following characters in the name:
 - <
 - >
 - ;
 - :
 - ,
 - |

- /
- *
- ?
- \
- \\
- \t
- \n
- \b

2. (Displayed if applicable) Select the port you want associated with the host security group. This port should contain your LUNs.
-

Note: Each type of storage system handles ports for HSGs differently. See “Rules for Creating Host Security Groups” on page 448.

3. (HDS storage system) Click the **Options** button to the right of the **Host Mode** field. Select a host mode resembling the port settings for your environment and then click **OK**. If your host mode is not listed, type it in the **Host Mode** field.
4. (HDS storage systems) If your storage system supports a second host mode, type the second host mode in the **Second Host Mode** field.
A second host mode is an optional settings on the port that describe how the host accesses the port (not applicable to all storage systems).
5. (IBM storage systems) You cannot assign the host mode for an IBM storage system in the user interface. You must modify an internal property to set the host mode. See “Setting the Host Mode for IBM Storage Systems” on page 460 for more information.
6. Click **Next**.

Step 2 - Add Initiators to the Host Security Group

1. To add an initiator to the host security group, click the **Add** button in the upper-right corner.
 2. Do one of the following:
 - ☐ Type the WWN of the port you want to add to the host security group
 - ☐ Select the initiator you want to add to the host security group.Notice that when the mouse hovers over the port, you are shown additional information, such as the name and WWN of the port on the switch that the host uses.
 3. Click the **Add** button at the bottom of the window.
 4. When you are done with adding initiators, click the **Close** button.
-

5. To remove an initiator from the host security group, click the  button. To remove multiple HBA initiators from the list, select the HBA ports you want to remove and then click the **Remove Selected** button.
6. Click **Next**.

Step 3 - Add Volumes to the Host Security Group

1. To add a volume to the host security group, click the **Add** button in the upper-right corner of the window.
2. Select a volume. Then, do one of the following:
 - ☐ If you want the unit number to be selected automatically by the server, leave the **Auto-Select** option selected.
 - ☐ If you want to choose a unit number, deselect the **Auto-Select** option and enter the unit number in the **Unit Number** field at the top of the window.

Note: Keep in mind that for Engenio storage systems LUN numbers cannot be duplicated, and that the management server can use an existing LUN number if the access mode is “No Access” for the created LUN.

3. Click **Add** at the bottom of the window.
The volume is added to host security group.
4. When you are done adding volumes, click the **Close** button.
5. To remove a volume from the host security group, click the  button. To remove multiple volumes from the list, select the volumes you want to remove and then click the **Remove Selected** button.
6. Click **Finish**.

Editing Host Security Groups

To edit a host security group:

1. Click **Provisioning** ().
 2. In the right pane, click the **Storage Systems** tab.
 3. In the right pane, click the **Provision** button corresponding to the storage system in which you want to access information about host security groups.
-

4. Click **Step 3 Host Security Group**.
5. Click the  button corresponding to the host security group you want to edit.

Step 1 - Edit Details for the Host

1. Modify the name of the host security group.

Note: The ID for a host security group changes when you rename the host security group.

2. Change the port you want associated with the host security group. This port should contain your LUNs.

Note: Each type of storage system handles ports for HSGs differently. See “Rules for Creating Host Security Groups” on page 448.

3. (Only HDS storage system) Click the **Options** button to the right of the **Host Mode** field. Select a host mode resembling the port settings for your environment and then click **OK**. If your host mode is not listed, type it in the **Host Mode** field.
4. (Only HDS storage systems) If your storage system supports a second host mode, type the second host mode in the **Second Host Mode** field.
A second host mode is an optional settings on the port that describe how the host accesses the port (not applicable to all storage systems).
5. (IBM storage systems) You cannot assign the host mode for an IBM storage system in the user interface. You must modify an internal property to set the host mode. See “Setting the Host Mode for IBM Storage Systems” on page 460 for more information.
6. Click **Next**.

Step 2 - Change the Initiators Assigned to the Host Security Group

1. Change the initiators assigned to the host security group, by doing one or more of the following:

☐ **Add an Initiator** - Click the **Add** button. Then, do one of the following:

- Type the WWN of the port you want to add to the host security group
- Select the initiator you want to add to the host security group.

Notice that when the mouse hovers over the port, you are shown additional information, such as the name and WWN of the port on the switch that the host uses.

Then, click **Add**. Click **Close** to exit the window.

- ☐ **Delete an initiator** - Click the  corresponding to the initiator you want to remove.
- ☐ **Delete multiple initiators** - To remove multiple initiators from the host security group, select the initiators you want to remove and then click **Remove Selected**.

Caution: Removing an HBA can cause hosts that are using it to lose access to their storage. This may result in data loss.

2. Click **Next**.

Step 3 - Change the Volumes Assigned to the Host Security Group

Important: You cannot delete the default host security group.

1. Change the volumes assigned to the host security group, by doing one or more of the following:
 - ☐ **Add a volume** - Click the **Add** button, select a volume and then, do one of the following:
 - If you want the unit number to be selected automatically by the server, leave the **Auto-Select** option selected.
 - If you want to choose a unit number, deselect the **Auto-Select** option and enter the unit number in the **Unit Number** field at the top of the window.Then, click **Add**. Click **Close** to exit the window.

Note: Keep in mind that for Engenio storage systems LUN numbers cannot be duplicated, and that the management server can use an existing LUN number if the access mode is “No Access” for the created LUN.

Caution: Removing an HBA can cause hosts that are using it to lose access to their storage. This may result in data loss.

- ☐ **Remove a volume** - Click  corresponding to the volume you want to remove from
-

the host security group.

- **Remove multiple volumes** - To remove multiple volumes from the host security group, select the volumes you want to remove and then click **Remove Selected**.
2. Click **Finish**.

Deleting Host Security Groups

Important: You cannot delete the default host security group.

To delete a host security group:

1. Click **Provisioning** ().
2. In the right pane, click the **Storage Systems** tab.
3. In the right pane, click the **Provision** button corresponding to the storage system in which you want to access information about host security groups.
4. Click **Step 3 Host Security Group**.
5. Click the  button corresponding to the host security group you want to delete. The host security group is removed.
6. To remove multiple host security groups, select the host security groups you want to remove. Then click **Delete Selected**.

Setting the Host Mode for IBM Storage Systems

The host mode for an IBM storage system cannot be set in the user interface. You can set the host mode by assigning the host mode to the `smi.ProvisioningIbmEss.hostConnectionProfile` property, as described in the following steps:

1. Click **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree.
 1. Click **Show Default Properties** at the bottom of the page.
 2. Copy `smi.ProvisioningIbmEss.hostConnectionProfile`. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
-

3. Return to the Advanced page (**Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree).
4. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
5. Assign a host mode to `smi.ProvisioningIbmEss.hostConnectionProfile` in the **Custom Properties** field, as shown in the following example:
`smi.ProvisioningIbmEss.hostConnectionProfile=AIX`
where AIX is the host mode.
6. When you are done, click **Save**.
7. Restart the service for the management server for your changes to take effect.
8. While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

General Provisioning Issues

This section describes the following:

- Provisioning Can Make a Device Inaccessible on page 461
- Provisioning Does Not Make an Operating System Aware of a Device on page 461

Provisioning Can Make a Device Inaccessible

Keep in mind that provisioning can break a connection between an array and a host. When you rezone a device, make sure no users or applications are using the device, as it will become unavailable. For example, assume a port for a disk drive is a member of zone set A, which is active. If you make zone set A inactive and this port is not a member of the new active zone set, the disk drive will become unavailable.

Provisioning Does Not Make an Operating System Aware of a Device

When a port in a zone set becomes active, you must take the appropriate steps to make it available to the operating system. For example, assume a port for a disk drive is a member of zone set A. If

you make the zone set active, the host will not automatically recognize the disk drive. You will need to configure the operating system, so that it becomes aware of the device. Refer to the documentation that accompanies the operating system for the host.

Provisioning Issues by Vendor

This section describes the following:

- Issues Specific to CLARiiON Storage Systems on page 462
- Issues Specific to EMC Symmetrix Storage Systems on page 463
- Issues Specific to HDS Storage Systems on page 465
- Issues Specific to HP Storage Systems on page 469
- Issues Specific to Engenio Storage Systems on page 470

Issues Specific to CLARiiON Storage Systems

The following issues are specific to CLARiiON® storage systems:

- Make the Management Server a Privileged User for CLARiiON on page 462
- EMC Navisphere CLI Is Required on page 463

Make the Management Server a Privileged User for CLARiiON

Before you can provision a CLARiiON storage system, you must configure it to recognize the management server as a privileged user. This task can be completed by using NaviCli as follows:

```
C:\>navicli -h 192.168.1.249 remoteconfig -setconfig -o -adduser  
SYSTEM@hostname
```

- where
 - 192.168.1.249 is the URL of the CLARiiON storage system
 - SYSTEM@hostname is the name of the computer running the management server
 - This example tells the CLARiiON storage system at 192.168.1.249 to accept configuration change requests from management server running on the computer named hostname.
-

EMC Navisphere CLI Is Required

The EMC Navisphere® CLI is required for the management server to communicate with the CLARiiON® storage system. At the time this documentation was created, EMC distributed the Navisphere CLI as part of the EMC Navisphere Software Suite. Contact your EMC representative for more information about obtaining the Navisphere CLI. Distribution rights for the Navisphere CLI belong to EMC. EMC Navisphere CLI Is Required

In Navisphere Manager add one of the following to the privilege user section:

Management server on Solaris:

```
SYSTEM@root
```

```
SYSTEM@root
```

where

- `name_of_my_management_server` is the DNS name of the computer running the management server software
- `IP_of_my_management_server` is the IP address of the computer running the management server software

Management server on Windows:

```
SYSTEM@name_of_my_management_server
```

```
SYSTEM@IP_of_my_management_server
```

where

- `name_of_my_management_server` is the DNS name of the computer running the management server software
- `IP_of_my_management_server` is the IP address of the computer running the management server software

When you discover the CLARiiON storage system, provide the IP address for the CLARiiON storage system and the user name and password used to log into Navisphere.

Issues Specific to EMC Symmetrix Storage Systems

The following issues are specific to EMC Symmetrix Storage Systems:

- About Provisioning on EMC Symmetrix Storage Systems on page 464

- LUN Provisioning Is Not Supported with Solutions Enabler 4.3 on page 464
- “Process Has an Exclusive Lock” Message on page 464
- Some EMC Volumes, Their LUNs and LUN Maskings Are Hidden on page 465

About Provisioning on EMC Symmetrix Storage Systems

EMC ships its Symmetrix storage system with volumes already created. When the software first discovers an EMC Symmetrix storage system, it assumes the devices on the Symmetrix storage system are volumes.

This software refers to the term “device” to define a piece of hardware in the storage network. EMC uses the term “device” to refer to a volume on one of its storage systems. In this section, the term “device” is used in the context of EMC storage systems.

When you use the software's “create a volume” feature, you are assigning the already created volume. If necessary the software will create a meta device, which is a device that is a concatenation of several devices.

The software does not delete the volumes created by EMC. When you use the software's “delete a volume” feature, the software marks the volume as hidden in its repository. These “hidden volumes” are stored in the “Free Device” list.

If you use a device in the “Free Device” list when you create a volume, that device is removed from the “Free Device” list. The device was moved from the list because it is now assigned.

LUN Provisioning Is Not Supported with Solutions Enabler 4.3

LUN mapping and LUN Masking works only when SYMAPI server version 5.0 or later is used. SYMAPI version 4.3 does not allow LUN provisioning, even though it shows existing LUNs. It also does not show existing LUN maskings.

“Process Has an Exclusive Lock” Message

You will receive a message resembling the one shown below if a process has already locked the EMC Symmetrix storage system and you attempt a process that requires a lock on the Symmetrix storage system. The Symmetrix storage system can become locked for many reasons. For example, the storage system becomes locked when it performs LUN mapping, LUN masking or Get

Details. The Symmetrix storage system may also remain locked after a provisioning operation has failed.

“SYMAPI routine SymDevMaskSessionStart failed with error code 188: The operation failed because another process has an exclusive lock on the local Symmetrix.”

After the management server has detected the lock on the Symmetrix storage system, it tries to access the storage system for 15 minutes and logs the errors.

If you receive the error message, determine if someone is performing an operation that requires a lock, such as LUN mapping, LUN masking or Get Details. This also applies even if one of the processes is being used by a third-party product, such as for LUN masking. If so, wait until the process is complete. Only manually remove the lock if you are certain that no other processes are occurring on the storage system. To learn how to remove the lock, refer to the documentation for the Symmetrix storage system.

If a provisioning failure has caused the Symmetrix storage system to remain locked, you are alerted to this situation in Event Manager and on the Properties tab. You may receive a message resembling the following:

```
Unable to end device masking session. Symmetrix '000001835005700' may be
locked.
```

Some EMC Volumes, Their LUNs and LUN Maskings Are Hidden

EMC volumes, their LUNs and LUN maskings for the volumes that play a special role (such as holding device masking information, attached BCVs, gate-keeper devices) are hidden. On the host side, the software shows LUN maskings for these volumes. So you may see a LUN masking, but not its volume residing on the EMC storage system.

Issues Specific to HDS Storage Systems

This section describes issues specific to HDS storage systems.

- About Provisioning on HDS Storage Systems on page 466
- Host Security Group Names on HDS Storage Systems on page 466
- Unable to Provision When HDS CruiseControl Is Enabled on page 467
- Increasing the Wait Time for the Management Server on page 468
- Initiator Ports Cannot Be Used for Provisioning on page 469

- The Software Is Unable to Map a Host to Volumes in Different Host Storage Domains of the Same Port on HDS 9900V on page 469

About Provisioning on HDS Storage Systems

Important: The management server does not allow LUSE of different RAID levels.

HDS ships some of its storage systems with volumes already created. When the software first discovers an HDS storage system, it detects the volumes created by HDS. When you use the software's "create a volume" feature, you are assigning the already created volume.

The software does not delete volumes created by HDS. When you use the software's "delete a volume" feature, the software marks the volume as hidden in its repository. These "hidden volumes" are stored in the "Free LDEVs" list.

If you use an LDEV in the "Free LDEVs" list when you create a volume, the LDEV is removed from the "Free LDEVs" list. The LDEVs was moved from the list because it is now assigned.

HDS cannot create a LUSE volume (made up of multiple LDEVs) without mapping it to a target port (that is without creating a LUN). In the software, creating a volume and creating a LUN are two different operations. Therefore, the software keeps the volumes, made from multiple LDEVs in the "Grouped LDEVs" list in the repository. Once these volumes are mapped to the target port and a LUN is created for them, they are removed from the repository and a real LUSE volume is created on the HDS box.

For example, assume you have several 2-GB "Free LDEVs" and you want to create a 4-GB volume. Since the request is larger than an LDEV, two of the "Free LDEVs" will be used for the 4-GB volume.

Host Security Group Names on HDS Storage Systems

In releases previous to build 5.0, when you used the management server to create a host security group, the newly created host security group was displayed by the "name" attribute. The "name" attribute is stored only in the HiCommand database and it is not stored in the HSG itself on the device. This is why the "name" attribute does not appear in the native tools. The name attribute can be 50 characters in length. Hitachi storage arrays contain HSG records on the device. These are identified by the "DisplayName" attribute that is read only. The DisplayName is displayed in the native tools. On some HDS storage systems, an additional "nickname" attribute is available to be set

on the HSG and its value is stored on the device itself. The management server used the “name” attribute instead of the “nickname” attribute because not all HDS storage systems support setting the “nickname” attribute.

In build 5.0, the management server uses the “nickname” attribute when it’s available. If you want to update the HSG names on the management server so they match the “nickname” attribute on the devices, you can use the Host Security Group Name Upgrader tool. The tool should be run once before upgrade or before the first full sync after an upgrade. If it is run afterwards, the management server will need to be resynchronized so that the database is updated with the renamed HSG’s. Follow these steps to use the Host Security Group Name Upgrader tool:

1. Go to the `[Installation_Directory]\Tools` directory on the management server.
2. On Windows, run `HsgNameUpgrader.bat`. On Solaris, run `HsgNameUpgrader.sh`. The Host Security Group Name Upgrader tool opens.
3. Enter a host name and credentials in the Server section.
4. Click the **Load** button. The storage system data is added to the tool and the new nickname is shown.
5. Edit the new nickname if necessary.
6. Click the **Purge old names** check box if you want to remove the old names from the management server database.
7. Click the **Process Select Array** button if you want to process a single array, or click the **Process All Arrays** button if you want to process all of the arrays.

Unable to Provision When HDS CruiseControl Is Enabled

When HDS CruiseControl is enabled on an HDS array, such as an HDS Lightning 9980V, you are unable to do provisioning. You might also receive the following error message:

An error was encountered during this operation. Some of the operation may have been applied to the storage subsystem. A refresh of the storage subsystem is recommended. "The LDEV is HIHSM reserved; cannot be used in a LUSE".

To use the Provisioning tool, disable HDS CruiseControl. Refer to the HDS CruiseControl product documentation for more information.

Increasing the Wait Time for the Management Server

By default, the management server waits 20 minutes for a response from HiCommand Device Manager after sending a provisioning command. If the management server does not receive a response from HiCommand Device Manager after 20 minutes, the management server assumes the provisioning command did not go through. It then tries to contact HiCommand Device Manager again while the previous command is still proceeding. After two retries, the management server stops attempting to contact HiCommand Device Manager.

For example, assume you initiated a provisioning command to add the HSG. The management server waits 20 minutes for a response. If it receives no response during the time, it sends another “Add” command and waits 20 minutes for a response. If no response is received, it sends another command. After the second retry, the management server stops attempting to contact HiCommand Device Manager. Multiple HSGs with the same configuration of LUNs and Worldwide Names (WWNs) were created when the management server attempted to contact HiCommand Device Manager.

The management server has a similar behavior when a delete command is initiated. It sends the delete command to HiCommand Device Manager, waits 20 minutes, and sends another delete command if it receives no response. The command tries to delete the same HSG, but the target HSG is deleted after the first command is completed. The second command returns an error.

If you need more time for HiCommand Device Manager to respond, you can increase the amount of time the management server waits, by modifying the

`cimom.provider.hds.ProvisioningTimeout` property as described in the following steps:

1. Select **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree on the management server.
 2. Click **Show Default Properties** at the bottom of the page.
 3. Copy the following text.
`cimom.provider.hds.ProvisioningTimeout`
 4. Return to the Advanced page (**Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree).
 5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
 6. Make your changes in the **Custom Properties** field. Change the value of the `cimom.provider.hds.ProvisioningTimeout` property. The value is in milliseconds. For example, assume you want the management server to wait an hour. You would assign
-

3600000 to `cimom.provider.hds.ProvisioningTimeout`, since 3600000 milliseconds is one hour.

```
cimom.provider.hds.ProvisioningTimeout=3600000
```

7. When you are done, click **Save**.
8. Restart the service for the management server for your changes to take effect. While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Initiator Ports Cannot Be Used for Provisioning

Ports designated as an Initiator on a storage system belonging to the HDS Freedom Storage™ Lightning 9900™ Series or Freedom Storage Lightning 9900V Series cannot be used for provisioning. If you select one of these ports, you receive a message saying that provisioning failed because the HiCommand Database was not refreshed. The management server does not support provisioning for ports designated as Initiators on these storage systems.

The Software Is Unable to Map a Host to Volumes in Different Host Storage Domains of the Same Port on HDS 9900V

On HDS 9900V storage systems, if a host is already mapped to a volume and you try to map the same host to a volume in another host storage domain, corresponding to the same port, it will fail. However for HDS 9900, the host can be mapped to a volume in another host storage domain corresponding to the same port.

A Default LUN Number Is Used Instead of a User-Specified One

When you create a LUN, the user-specified LUN is ignored. The LUN is created with the next available default number.

Issues Specific to HP Storage Systems

- Size Reported as 0 on EVA Array Volumes of 1 GB or More (Correct Size is Reported After Running Get All Details on EVA Array on page 470)

- Cannot Always Delete Selected Volume on MSA, MSA Allows Deletion of Only Topmost or Last Volume on page 470

Size Reported as 0 on EVA Array Volumes of 1 GB or More (Correct Size is Reported After Running Get All Details on EVA Array)

If you create a volume of 1 GB or more on an EVA, the size reported is 0 until a **Get All Details** is run on the EVA array. Smaller volume sizes do not exhibit this problem. This is a limitation in the EVA array software, and ApplIQ is working with HP to resolve this issue.

Cannot Always Delete Selected Volume on MSA, MSA Allows Deletion of Only Topmost or Last Volume

You can delete only the last or topmost volume on each of the MSA array's controllers (modeled as pools). Attempting to delete other volumes will return a generic error code 4. This is a limitation in the MSA array software and ApplIQ is working with HP to resolve the issues.

Issues Specific to Engenio Storage Systems

This section describes the following for Engenio:

- Creating and Deleting Storage Pools on page 470
- Creating and Deleting Storage Volumes on page 470

Creating and Deleting Storage Pools

For Engenio, a storage pool is the same as a volume group. Create at least one storage pool before provisioning a volume. When you delete a storage pool on an Engenio storage system, all the volumes for the volume group are deleted, including the placeholder volume.

Creating and Deleting Storage Volumes

Keep in mind the following when you create a volume on an Engenio storage system:

- Make sure you select a volume group that can accommodate the requested size for the new volume.
- You can create volumes from existing free extent areas within a volume group.
- The volume capabilities, their cache read ahead multiplier and segment size are shown in the following table:

Table 10-7: Volume Usage

Volume Capability	Cache Read Ahead Multiplier	Segment Size
File and Default	1	64 KB
Database	0	64 KB
Multimedia	8	128 KB
Custom	0 to 65,535	8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB

- No volume-to-LUN masking is done by default.

Important: The management server creates a placeholder volume when a storage pool is created. This placeholder volume is not viewable in the management server, but it might be viewable in other storage tools. Do not delete this placeholder volume.

Important: Depending on your license, Path Provisioning may not be available. See the “List of Features” to determine if you have access to Path Provisioning. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center**).

This chapter describes the following:

- About Path Provisioning on page 474
 - How Path Provisioning Works on page 476
 - How to Use Path Provisioning on page 477
 - Default System Action Templates on page 479
 - Creating a System Action Template on page 517
 - Modifying a System Action Template on page 518
 - Adding a Host on page 519
 - Creating a Host Security Group on page 520
 - Scheduling Provisioning Jobs on page 521
 - Executing Provisioning Jobs on page 523
 - Monitoring Provisioning Jobs on page 524
 - Deleting Multiple Jobs on page 525
-

- Naming Conventions for Zones on page 525
- Using Multipathing with Path Provisioning on page 526
- Customizing Path Provisioning on page 527
- Assigning Templates to a Role on page 533

About Path Provisioning

Path Provisioning lets you schedule provisioning tasks to take place at a later time when the network traffic is light. For example, you could use Path Provisioning to schedule multiple provisioning tasks to take place at 1 a.m. When you come in the later that morning, you can view the status of the provisioning tasks.

You can also use Path Provisioning to identify host/storage dependencies so you can make informed decisions when deciding where new volumes, zones or LUN security is needed.

To view the latest provisioning information in Path Provisioning, click the **Refresh** button. The **Refresh** button updates the following:

- The Path Provisioning screen with changes made in Provisioning. For example, assume you used the wizards in Provisioning to create a host security group. When you access Path Provisioning, your changes are not shown.
 - The Path Provisioning screen with changes from executed jobs. After a job is executed in Path Provisioning, the Path Provisioning screen is not updated until you click the **Refresh** button or you exit and re-enter Path Provisioning.
 - Adding a volume to an existing host security group on an EMC Symmetrix or DMX array only performs the masking operation, but it does not map the port on the array. The native tools for the array show that the masking to the host initiator took place, but the volume is still not mapped to a storage port.
 - Other parts of the product, such as:
 - ☐ Application Explorer
 - ☐ Capacity Explorer
 - ☐ Performance Explorer
 - ☐ Provisioning
 - ☐ System Explorer
-

Keep in mind the following:

- Path Provisioning runs within a Java applet. If you receive “out of memory” messages when you view Path Provisioning, you may need to increase the amount of memory assigned to the Java plug-in on the client computer.
- If you select a direct attached host, the storage ports appear in the LUN pane with a  next to them, indicating they are unreachable. You can still select these storage ports that have a  next to their icon and schedule the job. These storage ports are shown as unreachable in the user interface because the user interface uses switches to display the association between a host and a storage system. If the management server cannot detect a switch, as with a DAS connection, the user interface assumes the storage ports are unreachable.
- Ports designated as an Initiator on a storage system belonging to the HDS Freedom Storage™ Lightning 9900™ Series or Freedom Storage Lightning 9900V Series cannot be used for provisioning. If you select one of these ports, you receive a message saying that provisioning failed because the HiCommand Database was not refreshed. The management server does not support provisioning for ports designated as Initiators on these storage systems.
- The HBAs displayed may not have a connection to the selected storage system. This provides flexibility. You can select a disconnected HBA for a job to take place when it will be connected to the storage system.

Keep in mind the following for McDATA and Connectrix switches:

- When McDATA or Connectrix switches are discovered through a proxy by using SNMP, you cannot view or perform any provisioning operations for those switches. For example, you cannot view zone sets, zones, and/or zone alias.
- When McDATA or Connectrix switches are discovered by their IP address by using SNMP, you can only view the active zone set and its members. You cannot create, modify, and/or delete zone sets or its members.
- Zone aliases are not supported for McDATA or Connectrix switches.
- You can view zones, zone sets and zone aliases on a Cisco switch; however, you cannot use the management server to create, modify or remove them from a Cisco switch.
- Path Provisioning looks for the names of the active zone set and of the active zones and all of their saved counterparts in the zoning library in EFC Manager. The provisioning job only occurs if those names match.

How Path Provisioning Works

When you select a storage system in Path Provisioning, the management server displays all the information relevant for provisioning with respect to the storage system selected. This information includes:

- Mapped, unmapped and unmasked (mapped to one or more storage system ports but not associated to any host initiator port) volumes of the storage system,
- Already masked LUNs of the storage system and the front end ports of the storage system (possible candidates for LUN mapping step)
- Hosts that are reachable (hosts belonging to the same fabric as the storage system) along with their HBA and host initiator ports. Single multipathing functionality is supported.
- Existing zones to which the ports of the storage system belong are displayed. If the Host and the Storage System belong to multiple fabrics, zones of all those fabrics are displayed.

The above information is displayed in several different panes (storage system, host, volume, LUN and zone panes). You can select relevant information from each of these panes (such as, which volumes to be used for the LUN Mapping Task). Each of these selections potentially can create a provisioning task to be performed. You can create a job that contains one of the following:

- A single task of one type (such as, Map LUN)
- Set of tasks of the same type (such as, multiple LUN Mapping)
- Multiple LUN Masking tasks or a job consisting of several tasks of different types (such as, Map LUN, zone required ports, mask LUN, which go together as a combination).
- You can also create multiple jobs each consisting of multiple tasks of different types.

The jobs can be executed immediately or can be scheduled to start at a later time. Jobs with all its required details (such as, parameters to invoke Provider Service methods) are stored in a job queue. The scheduler, at the scheduled time, retrieves the job from the job queue and performs the tasks of the job using the details stored in the job.

The status of each of the job can be one of the following, and it is displayed in the State column in the Provision Job section of the screen (lower pane):

- **Created** - The job has been created, but it will not be executed. The job cannot be viewed by others and it is deleted when the Web browser is closed. See “Scheduling Provisioning Jobs” on page 521 for information about changing the state of the job from “created” to “scheduled”.
-

- **Scheduled** - The job has been tasked to execute at a specified time and date. Jobs are assigned a scheduled state after you select the job and click the **Execute Job** button ( button).
- **Started** - The job has started. You cannot delete a job once it has started.
- **Failed** - The job failed.
- **Ended** - The job has finished.

How to Use Path Provisioning

All the provisioning tasks are centralized on one screen. You select the provisioning actions you want to take by selecting a default system action template from the System Action combo-box in the upper-left corner or you can create your own system action template, as described in “Creating a System Action Template” on page 517.

The following are the default templates:

- **Volume Creations + LUN Security + Zone Operation**- First you create a meta volume. Second, you map a meta volume to a Fibre Channel port and host HBAs (HSG). Third, you create a zone. See “Volume Creation, LUN Security, and Zone Operation” on page 479 for more information.
 - **Meta Volume Creation** - Creates a meta volume. See “Meta Volume Creation” on page 486 for more information.
 - **LUN Security** - Map a meta volume to Fibre Channel port and host HBAs (HSG). See the topic “LUN Security” on page 488 for more information.
 - **Zone Operation** - Lets you perform just a zone operation. See “Zone Operation” on page 493 for more information.
 - **Volume Creation + LUN Security** - First you create a meta volume. Second you map a meta volume to a Fibre Channel port). See “Volume Creation and LUN Security” on page 499 for more information.
 - **LUN Security and Zone Operation** - Lets you create a host security group (HSG) with the host HBA WWN along with zoning operations,. See “LUN Security and Zone Operation” on page 504.
 - **Volume Assignment** - Lets you assign a volume to existing host security groups. See “Volume Assignment” on page 511.
-

After you select a system action, complete the steps in the various panes. If a step is disabled after all data has been loaded, that step is not required.

When you have completed your steps, schedule the task as described in the topic, “Scheduling Provisioning Jobs” on page 521.

Note: You can customize which templates users can access. See “Assigning Templates to a Role” on page 533 for more information.

The following table provides an overview of the steps required for Path Provisioning.

Table 11-1: Overview for Path Provisioning

Step	Description	Where to Find Additional Information
1	Select a system action.	See the topics: ■ “Volume Creation, LUN Security, and Zone Operation” on page 479 ■ “Meta Volume Creation” on page 486 ■ “LUN Security” on page 488 ■ “Zone Operation” on page 493 ■ “Volume Creation and LUN Security” on page 499 ■ “LUN Security and Zone Operation” on page 504 ■ “Volume Assignment” on page 511
2	Complete the steps.	

Table 11-1: Overview for Path Provisioning (Continued)

Step	Description	Where to Find Additional Information
3	(optional) Schedule a provisioning job if you want it to take place at a later time.	See “Scheduling Provisioning Jobs” on page 521.
4	Execute the job. The job does not run until the Execute Job  button is clicked. Once this button is clicked the job is saved in the management server database.	See “Executing Provisioning Jobs” on page 523.

Default System Action Templates

This section describes the following default system action templates:

- Volume Creation, LUN Security, and Zone Operation on page 479
- Meta Volume Creation on page 486
- LUN Security on page 488
- Zone Operation on page 493
- Volume Creation and LUN Security on page 499
- LUN Security and Zone Operation on page 504
- Volume Assignment on page 511
- Providing a LUN Number on page 516
- Adding a Host on page 519

Volume Creation, LUN Security, and Zone Operation

You can use Path Provisioning to create meta volumes. Then, map that meta volume to a Fibre Channel port and host HBA (HSG). Finally, you could designate that volume to appear in a pre-existing zone or create your own.

Keep in mind the following:

- If you have options still selected from a previous job, just clear the options you do not want in your next job. For example, assume you created a volume. Now you want to create a new volume on the same host used previously. Just clear the Volume pane. To clear a pane, click the  button.
- If you want to clear all the steps, except for the Step 1 (storage systems), select another option from the System Action combo-box.
- (HDS only) Before you can create a volume. You must delete some unmapped LDEVs using the standard provisioning tool. To delete LDEVs, click **Provisioning** (). Then, click the Storage Systems tab, then the **Provision** button for the storage system. Then, click **Step 2 Volume**. Click the **Unmapped Volumes** folder in the left pane. Select the desired number of LDEVs for the LUSE volume then click the **Delete Selected Volumes** button. Take note of the array group from which you deleted the LDEVs. You need this information to create the LUSE volume.

To access Path Provisioning:

1. Click **Provisioning** ().
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. Select the following from the **System Action** combo-box: Volume Creation + LUN Security + Zone Operation

Step 1 - Select Storage System

1. Wait for the management server to load the storage systems into the Storage System panel.
2. Select a storage system on which you want to create the metavolume.

Note: The S heading in the Storage Systems pane means that only a single selection is allowed.

(HDS only) Select the storage system from which you want to create the LUSE volume.

3. Click the **Step 1** button underneath the pane. Information about the selected storage system, such as its model and IP address is displayed below the Storage System pane. The Host pane becomes populated. Notice in
-

the following pane that some hosts have a red X over their icon. This means the host is not accessible.




Figure 11-1: Selecting a Storage System

Step 2 - Select a Host

1. Wait for all data to be loaded. You can tell if all data has been loaded when you see the following messages:
 - ☐ Loading host data finished.
 - ☐ Loading volume data finished.
 - ☐ Loading HSG data finished.

- Loading zone data finished.

The Step 2 button appears disabled until data has been loaded

2. Select a host that is reachable.

Keep in mind the following:

- To create a provisioning job for multipathing, click the  button above the Host pane. Then, select the option for multipathing. When you select this option, you must select the same host for both paths. See “Host Customize Dialog” on page 528 for more information.
- To automatically create a zone if a zone does not meet a preset criteria, click the  button above the Zone pane. Then, select the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**. Select one of the options listed. See “Zone Customize Dialog” on page 530 for more information about setting the criteria. After you set the criteria, select a host and click **Step 2**. If the management server finds a zone that meets the criteria, it selects the zone in the Zone pane. If the management server does not find a zone that meets the criteria, it asks for a zone name. See “Naming Conventions for Zones” on page 525 for more information about the naming requirements for a zone. After you type a zone name, the new zone is displayed in the Zone pane, but it has not been created yet. It is created when the job runs.
- To configure zoning manually, click the  button above the Zone pane. Then, deselect the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**. In the zone pane, you must select a zone or create one manually.

If you see the  icon next to a host, it means the host is unreachable.

3. Click the **Step 2** button.

Information about the selected port, such as its Worldwide Name, is displayed below the Host pane. The volumes for that host are displayed in the Volume pane.

Step 3 - Select a Volume

To select a volume:

1. In the Volume pane select mapped and unmapped volumes. You can select multiple volumes on Windows computers by pressing the Ctrl key as you select the volumes.
 - **Mapped** - There are two types of mapped volumes:
 - **Masked** - The volume is exposed to the storage port and to the host.
-

- **Unmasked** - The volume is exposed to the storage port, but not to the host.
- **Unmapped** – The volume is not exposed to the storage port. The management server puts all unmapped volumes in this category when it first discovers an array.
- **Free Extents** – Available free extents that can be used to create a meta volume. You can create meta volumes on EMC Symmetrix and LUSE on HDS storage systems. To create a meta volume or LUSE, select multiple free extents under the Free Extents node in the Volume pane. Select multiple LDEVS from the Free Extents menu by holding down the shift key on your keyboard and selecting free LDEVS. When you select free extents, they must be of the same type. For example, on Symmetrix, you cannot select a mirrored volume and a BCV (business continuous volume) to create a meta volume.

When you first discover a storage system, no free extents are displayed. This is because the management server puts all unmapped volumes into the “unmapped” category for the list of volumes by default. To move a volume to the free extent node, delete the unmapped volume. When the volume is deleted, it is moved to the free extents node. The free extents category is used internally by the management server. See “Deleting a Storage Volume” on page 445 for more information.

Important: Make sure the free extents you select are not being used. Data on the free extents becomes unusable when a meta volume is created.

Keep in mind the following:

- You can narrow the type of volumes displayed in the Volumes pane by using the Volume Customize Dialog () located above the Volumes pane. See “Host Customize Dialog” on page 528 for more information.
 - If the LUN has already been selected and Step 4 is clicked, skip this step or click the  button.
2. Click **Step 3**.
 3. If you are asked to specify a LUN number, provide a LUN for each volume displayed. See “Providing a LUN Number” on page 516 for information about numbering LUNs.

Step 4 - Select a Host Security Group

1. Select a host security group. See “Creating a Host Security Group” on page 520 for information on how to create a host security group. See “General Provisioning Issues” on page 461 for information on how your storage system handles host security groups.
2. Click **Step 4**.

Step 5 - Select a Zone

Note: If the zone has already been selected and Step 5 is clicked, skip this step or click the  button to clear the selection.

If you do not have a zone selected or you do not have one created yet, most likely the option, **Automatically Configure Zoning**, is not selected in the Zone Customize Dialog (). The management server assumes you want to select a pre-existing zone or create one manually when you deselect the **Automatically Configure Zoning** option.

- **To reuse a zone** - Select a zone in the Zone pane and then click **Step 5**. Expand a fabric node to view its zones.
- **To create a zone** - Select a fabric in the zone pane, click the  button, type a zone name. See the following topic for more information.

If you see the  icon, the fabric cannot be reached.

Table 11-2: Zone Icons

Icon	Description
	Its location on the page determines its use: ■ Displayed above Zone pane - Button for creating zones. ■ In Zone pane - Icon for a zone.
	Zone Alias
	Port

Important: (McDATA switches only) Path Provisioning looks for the names of the active zone set and of the active zones and all of their saved counterparts in the zoning library in EFC Manager. The provisioning job only occurs if those names match.

Creating a Zone

To create a zone:

1. Select a fabric in the zone pane.




Figure 11-2: Selecting a Fabric

2. Click the  button located above the Zone pane.
3. Type a zone name in the dialog box. The following restrictions apply:

Naming Conventions for Brocade Switches:

- ☐ The name must contain 1 to 64 characters.
- ☐ The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or an underscore (_).
- ☐ The name is case sensitive. For example, “Zone1” and “zone1” are different zones.
- ☐ You cannot create a zone with the same name as an existing zone, zone alias or zone set. For example, if you create a zone named “new”, you cannot give a zone, zone alias, or zone set the same name.
- ☐ The following characters are invalid for Brocade switches: caret (^), dash (-), and dollar sign (\$).

Naming Conventions for McDATA and Connectrix Switches:

- ☐ The name can have a maximum of 64 characters.
- ☐ The first character of a zone name must be a letter (A-Z, AZ).
- ☐ A zone name cannot contain spaces.
- ☐ Valid characters are a-a, AA, 0-9, caret (^), dash (-), underscore (_), and dollar sign (\$).

- ☐ All names must be unique and may not differ by case. For example, myzone and MyZone are considered to be the same zone.
- 4. Click **OK**.
The new zone is added to the Zone pane.
- 5. Click the **Create Job** button.
The job is added to the bottom pane.
- 6. Do one of the following:
 - ☐ If you want the job to execute now, click the Execute Job () button
 - ☐ If you want the job to execute at a later time, schedule the job as described in the topic, “Scheduling Provisioning Jobs” on page 521.

Meta Volume Creation

Keep in mind the following when creating meta volumes:

- If you have options still selected from a previous job, just clear the options you do not want in your next job. For example, assume you created a volume. Now you want to create a new volume on the same host used previously. Just clear the Volume pane. To clear a pane, click the  button.
- If you want to clear all the steps, except for the Step 1 (storage systems), select another option from the System Action combo-box.
- (HDS only) Before you can create a volume. You must delete some unmapped LDEVs using the standard provisioning tool. To delete LDEVs, click **Provisioning** (). Then, click the Storage Systems tab, then the **Provision** button for the storage system. Then, click **Step 2 Volume**. Click the **Unmapped Volumes** folder in the left pane. Select the desired number of LDEVs for the LUSE volume then click the **Delete Selected Volumes** button. Take note of the array group from which you deleted the LDEVs. You need this information to create the LUSE volume.

You can use Path Provisioning to create meta volumes, as described in the following steps:

1. Click **Provisioning** ().
 2. In the right pane, click **Start Here** on the Path Provisioning tab.
 3. Select the following from the System Action combo-box: **Meta Volume Creation**
-

4. Wait for the management server to load the storage systems into the Storage System panel.
5. Select the storage system on which you want to create the metavolume.
(HDS only) Select the storage system from which you want to create the LUSE volume.

Note: The S heading in the Storage Systems pane means that only a single selection is allowed.

6. Click the **Step 1** button underneath the pane.
Information about the selected storage system, such as its model and IP address is displayed below the Storage System pane.
7. In the Volume pane, select multiple free extents under the Free Extents node in the Volume pane.
The meta volume containing the selected free extents is created when the job runs.
 - ☐ **Mapped** - There are two types of mapped volumes:
 - **Masked** - The volume is exposed to the storage port and to the host.
 - **Unmasked** - The volume is exposed to the storage port, but not to the host.
 - ☐ **Unmapped** – The volume is not exposed to the storage port. The management server puts all unmapped volumes in this category when it first discovers an array.
 - ☐ **Free Extents** – Available free extents that can be used to create a meta volume. You can create meta volumes on EMC Symmetrix and LUSE on HDS storage systems. To create a meta volume or LUSE, select multiple free extents under the Free Extents node in the Volume pane. Select multiple LDEVs from the Free Extents menu by holding down the shift key on your keyboard and selecting free LDEVs. When you select free extents, they must be of the same type. For example, on Symmetrix, you cannot select a mirrored volume and a BCV (business continuous volume) to create a meta volume.
When you first discover a storage system, no free extents are displayed. This is because the management server puts all unmapped volumes into the “unmapped” category for the list of volumes by default. To move a volume to the free extent node, delete the unmapped volume. When the volume is deleted, it is moved to the free extents node. The free extents category is used internally by the management server. See “Deleting a Storage Volume” on page 445 for more information.

Keep in mind the following:

- ☐ Make sure the free extents you select are not being used. Data on the free extents becomes unusable when a meta volume is created.
-

- ☐ You can narrow the type of volumes displayed in the Volumes pane by using the Volume Customize Dialog () , located above the Volumes pane. See “Host Customize Dialog” on page 528.
 - ☐ If the LUN has already been selected and Step 4 is clicked, skip this step or click the  button.
8. Click **Step 3**.
 9. Click the **Create Job** button.
The job is added to the bottom pane.
 10. Do one of the following:
 - ☐ If you want the job to execute now, select the job and then click the Execute Job () button.
 - ☐ If you want the job to execute at a later time, schedule the job as described in the topic, “Scheduling Provisioning Jobs” on page 521.

LUN Security

Keep in mind the following:

- If you have options still selected from a previous job, just clear the options you do not want in your next job. For example, assume you just mapped a volume to a port. Now you want to map the a different volume on the same host to another port. Just clear the Volume and LUN panes. To clear a pane, click the  button.
- If you want to clear all the steps, except for the Step 1 (storage systems), select another option from the System Action combo-box.

You can use Path Provisioning to designate subsystem LUN security, as described in the following steps:

1. Click **Provisioning** ().
 2. In the right pane, click **Start Here** on the Path Provisioning tab.
 3. Select the following from the System Action combo-box: LUN Security
-

Step 1 - Select Storage System

1. Wait for the management server to load the storage systems into the Storage System panel.
2. Select the storage system on which you want to create the metavolume.

Note: The S heading in the Storage Systems pane means that only a single selection is allowed.

3. Click the **Step 1** button underneath the pane.
Information about the selected storage system, such as its model and IP address is displayed below the Storage System pane. The Host pane becomes populated. Notice in the following pane that some hosts have a red X over their icon. This means the host is not accessible.




Figure 11-3: Selecting a Storage System

Step 2 - Select a Host

1. Wait for all data to be loaded. You can tell if all data has been loaded when you see the following messages:
 - ☐ Loading host data finished.
 - ☐ Loading volume data finished.
 - ☐ Loading HSG data finished.The Step 2 button appears disabled until data has been loaded

2. Do one of the following:

- Select a host that is reachable.
- Add a host that is not currently connected to the network by clicking the  button. See “Adding a Host” on page 491.

Keep in mind the following:

- To create a provisioning job for multipathing, click the  button above the Host pane. Then, select the option for multipathing. When you select this option, you must select the same host for both paths. See “Host Customize Dialog” on page 528 for more information.
- To automatically create a zone if a zone does not meet a preset criteria, click the  button above the Zone pane. Then, select the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**. Select one of the options listed. See “Zone Customize Dialog” on page 530 for more information about setting the criteria. After you set the criteria, select a host and click **Step 2**. If the management server finds a zone that meets the criteria, it selects the zone in the Zone pane. If the management server does not find a zone that meets the criteria, it asks for a zone name. See “Naming Conventions for Zones” on page 525 for more information about the naming requirements for a zone. After you type a zone name, the new zone is displayed in the Zone pane, but it has not been created yet. It is created when the job runs.
- To configure zoning manually, click the  button above the Zone pane. Then, deselect the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**. In the zone pane, you must select a zone or create one manually.

If you see the  icon next to a host, it means the host is unreachable.

3. Click the **Step 2** button.

Information about the selected port, such as its Worldwide Name, is displayed below the Host pane. The volumes for that host are displayed in the Volume pane.

Adding a Host

The management server lets you add hosts that are not currently connected to the network. While you are creating a job, add the host.

Important: Make sure the added host is physically connected to the network before the scheduled job runs.

1. Click the  button.
2. Type the name for the host in the **Host Name** field.
3. Type the port name of the host in the **Port WWN** field.
4. Click the **Add** button.
5. Repeat Steps 2 and 3 for multiple ports.
6. If you want to remove the host, click the  button.
7. When you are done with your changes, click **OK**.
The host is added to the list of hosts.
8. Physically connect the host to the network before the job runs.

Step 3 - Select a Volume

To select a volume:

1. In the Volume pane select mapped and unmapped volumes. You can select multiple volumes on Windows computers by pressing the Ctrl key as you select the volumes.
 - ☐ **Mapped** - There are two types of mapped volumes:
 - **Masked** - The volume is exposed to the storage port and to the host.
 - **Unmasked** - The volume is exposed to the storage port, but not to the host.
 - ☐ **Unmapped** – The volume is not exposed to the storage port. The management server puts all unmapped volumes in this category when it first discovers an array.
 - ☐ **Free Extents** – Available free extents that can be used to create a meta volume. You can create meta volumes on EMC Symmetrix and LUSE on HDS storage systems. To create a meta volume or LUSE, select multiple free extents under the Free Extents node in the Volume pane. When you select free extents, they must of the same type. For example, on Symmetrix, you cannot select a mirrored volume and a BCV (business continuous volume) to create a meta volume.
When you first discover a storage system, no free extents are displayed. This is because the management server puts all unmapped volumes into the “unmapped” category for the list of volumes by default. To move a volume to the free extent node, delete the unmapped volume. When the volume is deleted, it is moved to the free

extents node. The free extents category is used internally by the management server. See “Deleting a Storage Volume” on page 445 for more information.

Important: Make sure the free extents you select are not being used. Data on the free extents becomes unusable when a meta volume is created.

2. Click **Step 3**.
3. If you are asked to specify a LUN number, provide a LUN for each volume displayed. See “Providing a LUN Number” on page 516 for information about numbering LUNs. Keep in mind the following:
 - You can narrow the type of volumes displayed in the Volumes pane by using the Volume Customize Dialog () located above the Volumes pane. See “Host Customize Dialog” on page 528.
 - If the LUN has already been selected and Step 4 is clicked, skip this step or click the  button.

Step 4 - Select a Host Security Group

1. Select a host security group. See “Creating a Host Security Group” on page 520 for information on how to create a host security group. See “General Provisioning Issues” on page 461 for information on how your storage system handles host security groups.
2. Click **Step 4**.
3. Click the **Create Job** button.
The job is added to the bottom pane.
4. Do one of the following:
 - ☐ If you want the job to execute now, click the Execute Job () button
 - ☐ If you want the job to execute at a later time, schedule the job as described in the topic, “Scheduling Provisioning Jobs” on page 521.

Zone Operation

Keep in mind the following:

- If you have options still selected from a previous job, just clear the options you do not want in your next job. For example, assume you created a zone. Now you want to create a new

zone that includes the same host used previously. Just clear the Zone pane. To clear a pane, click the  button.

- If you want to clear all the steps, except for the Step 1 (storage systems), select another option from the System Action combo-box.

You can use Path Provisioning to perform just zoning operations, as described in the following steps:

1. Click **Provisioning** ().
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. Select the following from the System Action combo-box: Zone Operation

Step 1 - Select Storage System

1. Wait for the management server to load the storage systems into the Storage System panel.
2. Select a storage system.

Note: The S heading in the Storage Systems pane means that only a single selection is allowed.

3. Click the **Step 1** button underneath the pane.
Information about the selected storage system, such as its model and IP address is displayed below the Storage System pane. The Host pane becomes populated. Notice in the following pane that some hosts have a red X over their icon. This means the host is not accessible.
-




Figure 11-4: Selecting a Storage System

Step 2 - Select a Host

1. Wait for all data to be loaded. You can tell if all data has been loaded when you see the following messages:
 - ☐ Loading host data finished.
 - ☐ Loading HSG data finished.
 - ☐ Loading zone data finished.The Step 2 button appears disabled until data has been loaded
2. Select a host that is reachable.

Keep in mind the following:

- ❑ To create a provisioning job for multipathing, click the  button above the Host pane. Then, select the option for multipathing. When you select this option, you must select the same host for both paths. See “Host Customize Dialog” on page 528 for more information.
- ❑ To automatically create a zone if a zone does not meet a preset criteria, click the  button above the Zone pane. Then, select the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**. Select one of the options listed. See “Zone Customize Dialog” on page 530 for more information about setting the criteria. After you set the criteria, select a host and click **Step 2**. If the management server finds a zone that meets the criteria, it selects the zone in the Zone pane. If the management server does not find a zone that meets the criteria, it asks for a zone name. See “Naming Conventions for Zones” on page 525 for more information about the naming requirements for a zone. After you type a zone name, the new zone is displayed in the Zone pane, but it has not been created yet. It is created when the job runs.
- ❑ To configure zoning manually, click the  button above the Zone pane. Then, deselect the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**. In the zone pane, you must select a zone or create one manually.

If you see the  icon next to a host, it means the host is unreachable.

3. Click the **Step 2** button.
Information about the selected port, such as its Worldwide Name, is displayed below the Host pane. The volumes for that host are displayed in the Volume pane.
 4. Select a host security group in the LUN pane. See “Creating a Host Security Group” on page 520 for information on how to create a host security group. See “General Provisioning Issues” on page 461 for information on how your storage system handles host security groups.
 5. Click **Step 4**.
-

Step 3 - Select a Zone

Note: If the zone has already been selected and Step 5 is clicked, skip this step or click the  button to clear the selection.

If you do not have a zone selected or you do not have one created yet, most likely the option, **Automatically Configure Zoning**, is not selected in the Zone Customize Dialog (). The management server assumes you want to select a pre-existing zone or create one manually when you deselect the **Automatically Configure Zoning** option.

- **To reuse a zone** - Select a zone in the Zone pane and then click **Step 5**. Expand a fabric node to view its zones.
- **To create a zone** - Select a fabric in the zone pane, click the  button, type a zone name. See the following topic for more information.

If you see the  icon, the fabric cannot be reached.

Table 11-3: Zone Icons

Icon	Description
	Its location on the page determines its use: ■ Displayed above Zone pane - Button for creating zones. ■ In Zone pane - Icon for a zone.
	Zone Alias
	Port

Important: (McDATA switches only) Path Provisioning looks for the names of the active zone set and of the active zones and all of their saved counterparts in the zoning library in EFC Manager. The provisioning job only occurs if those names match.

Creating a Zone

To create a zone:

1. Select a fabric in the zone pane.




Figure 11-5: Selecting a Fabric

2. Click the  button located above the Zone pane.
3. Type a zone name in the dialog box. The following restrictions apply:

Naming Conventions for Brocade Switches:

- ☐ The name must contain 1 to 64 characters.
- ☐ The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or an underscore (_).
- ☐ The name is case sensitive. For example, "Zone1" and "zone1" are different zones.
- ☐ You cannot create a zone with the same name as an existing zone, zone alias or zone set. For example, if you create a zone named "new", you cannot give a zone, zone alias, or zone set the same name.
- ☐ The following characters are invalid for Brocade switches: caret (^), dash (-), and dollar sign (\$).

Naming Conventions for McDATA and Connectrix Switches:

- ☐ The name can have a maximum of 64 characters.
- ☐ The first character of a zone name must be a letter (A-Z, AZ).
- ☐ A zone name cannot contain spaces.
- ☐ Valid characters are a-a, AA, 0-9, caret (^), dash (-), underscore (_), and dollar sign (\$).

- ☐ All names must be unique and may not differ by case. For example, myzone and MyZone are considered to be the same zone.
- 4. Click **OK**.
The new zone is added to the Zone pane.
- 5. Click the **Create Job** button.
The job is added to the bottom pane.
- 6. Do one of the following:
 - ☐ If you want the job to execute now, click the Execute Job () button
 - ☐ If you want the job to execute at a later time, schedule the job as described in the topic, “Scheduling Provisioning Jobs” on page 521.

Volume Creation and LUN Security

Keep in mind the following:

- If you have options still selected from a previous job, just clear the options you do not want in your next job. For example, assume you created a volume. Now you want to create a new volume on the same host used previously. Just clear the Volume pane. To clear a pane, click the  button.
- If you want to clear all the steps, except for the Step 1 (storage systems), select another option from the System Action combo-box.
- (HDS only) Before you can create a volume. You must delete some unmapped LDEVs using the standard provisioning tool. To delete LDEVs, click **Provisioning** (). Then, click the Storage Systems tab, then the **Provision** button for the storage system. Then, click **Step 2 Volume**. Click the **Unmapped Volumes** folder in the left pane. Select the desired number of LDEVs for the LUSE volume then click the **Delete Selected Volumes** button. Take note of the array group from which you deleted the LDEVs. You need this information to create the LUSE volume.

You can create a meta volume and designate LUN security, as described in the following steps.

1. Click **Provisioning** ().
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. Select the following from the System Action combo-box: Volume Creation + LUN Security

Step 1 - Select Storage System

1. Wait for the management server to load the storage systems into the Storage System panel.
 2. Select the storage system on which you want to create the metavolume.
-

Note: The S heading in the Storage Systems pane means that only a single selection is allowed.

(HDS only) Select the storage system from which you want to create the LUSE volume.

3. Click the **Step 1** button underneath the pane.
Information about the selected storage system, such as its model and IP address is displayed below the Storage System pane.




Figure 11-6: Selecting a Storage System

Step 2 - Select a Volume

To select a volume:

1. In the Volume pane select mapped and unmapped volumes. You can select multiple volumes on Windows computers by pressing the Ctrl key as you select the volumes.
 - **Mapped** - There are two types of mapped volumes:
 - **Masked** - The volume is exposed to the storage port and to the host.
 - **Unmasked** - The volume is exposed to the storage port, but not to the host.

- ❑ **Unmapped** – The volume is not exposed to the storage port. The management server puts all unmapped volumes in this category when it first discovers an array.
- ❑ **Free Extents** – Available free extents that can be used to create a meta volume. You can create meta volumes on EMC Symmetrix and LUSE on HDS storage systems. To create a meta volume or LUSE, select multiple free extents under the Free Extents node in the Volume pane. Select multiple LDEVs from the Free Extents menu by holding down the shift key on your keyboard and selecting free LDEVs. When you select free extents, they must be of the same type. For example, on Symmetrix, you cannot select a mirrored volume and a BCV (business continuous volume) to create a meta volume.

When you first discover a storage system, no free extents are displayed. This is because the management server puts all unmapped volumes into the “unmapped” category for the list of volumes by default. To move a volume to the free extent node, delete the unmapped volume. When the volume is deleted, it is moved to the free extents node. The free extents category is used internally by the management server. See “Deleting a Storage Volume” on page 445 for more information.

Important: Make sure the free extents you select are not being used. Data on the free extents becomes unusable when a meta volume is created

Keep in mind the following:

- ❑ You can narrow the type of volumes displayed in the Volumes pane by using the Volume Customize Dialog () located above the Volumes pane. See “Host Customize Dialog” on page 528.
- ❑ If the LUN has already been selected and Step 4 is clicked, skip this step or click the  button.




Figure 11-7: Selecting a Volume

2. Click **Step 3**.
3. If you are asked to specify a LUN number, provide a LUN for each volume displayed. See "Providing a LUN Number" on page 516 for information about numbering LUNs.

Step 3 - Select a Host Security Group

1. Select a host security group in the LUN pane. See “Creating a Host Security Group” on page 520 for information on how to create a host security group. See “General Provisioning Issues” on page 461 for information on how your storage system handles host security groups.
2. Click **Step 4**.
3. Click the **Create Job** button.
The job is added to the bottom pane.
4. Do one of the following:
 - ☐ If you want the job to execute now, click the Execute Job () button
 - ☐ If you want the job to execute at a later time, schedule the job as described in the topic, “Scheduling Provisioning Jobs” on page 521.

LUN Security and Zone Operation

Keep in mind the following:

- If you have options still selected from a previous job, just clear the options you do not want in your next job. For example, assume you created a volume. Now you want to create a new volume on the same host used previously. Just clear the Volume pane. To clear a pane, click the  button.
- If you want to clear all the steps, except for the Step 1 (storage systems), select another option from the System Action combo-box.
- This template does not create volumes nor associate the volumes to HSG.

You can use Path Provisioning to create a host security group (HSG) with the host HBA WWN along with zoning operations, as described in the following steps:

1. Click **Provisioning** ()
 2. In the right pane, click **Start Here** on the Path Provisioning tab.
 3. Select the following from the System Action combo-box: LUN Security and Zone Operation
-

Step 1 - Select Storage System

1. Wait for the management server to load the storage systems into the Storage System panel.
2. Select the storage system on which you want to create the metavolume.

Note: The S heading in the Storage Systems pane means that only a single selection is allowed.

3. Click the **Step 1** button underneath the pane.
Information about the selected storage system, such as its model and IP address is displayed below the Storage System pane. The Host pane becomes populated. Notice in the following pane that some hosts have a red X over their icon. This means the host is not accessible.




Figure 11-8: Selecting a Storage System

Step 2 - Select a Host

1. Wait for all data to be loaded. You can tell if all data has been loaded when you see the following messages:
 - ☐ Loading host data finished.
 - ☐ Loading volume data finished.
 - ☐ Loading HSG data finished.The Step 2 button appears disabled until data has been loaded

2. Do one of the following:

- Select a host that is reachable.
- Add a host that is not currently connected to the network by clicking the  button. See “Adding a Host” on page 491.

Keep in mind the following:

- To create a provisioning job for multipathing, click the  button above the Host pane. Then, select the option for multipathing. When you select this option, you must select the same host for both paths. See “Host Customize Dialog” on page 528 for more information.
- To automatically create a zone if a zone does not meet a preset criteria, click the  button above the Zone pane. Then, select the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**. Select one of the options listed. See “Zone Customize Dialog” on page 530 for more information about setting the criteria. After you set the criteria, select a host and click **Step 2**. If the management server finds a zone that meets the criteria, it selects the zone in the Zone pane. If the management server does not find a zone that meets the criteria, it asks for a zone name. See “Naming Conventions for Zones” on page 525 for more information about the naming requirements for a zone. After you type a zone name, the new zone is displayed in the Zone pane, but it has not been created yet. It is created when the job runs.
- To configure zoning manually, click the  button above the Zone pane. Then, deselect the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**. In the zone pane, you must select a zone or create one manually.

If you see the  icon next to a host, it means the host is unreachable.

3. Click the **Step 2** button.

Information about the selected port, such as its Worldwide Name, is displayed below the Host pane. The volumes for that host are displayed in the Volume pane.

Adding a Host

The management server lets you add hosts that are not currently connected to the network. While you are creating a job, add the host.

Important: Make sure the added host is physically connected to the network before the scheduled job runs.

1. Click the  button.
2. Type the name for the host in the **Host Name** field.
3. Type the port name of the host in the **Port WWN** field.
4. Click the **Add** button.
5. Repeat Steps 2 and 3 for multiple ports.
6. If you want to remove the host, click the  button.
7. When you are done with your changes, click **OK**.
The host is added to the list of hosts.
8. Physically connect the host to the network before the job runs.

Step 3 - Select a Host Security Group

1. Select a host security group. See “Creating a Host Security Group” on page 520 for information on how to create a host security group. See “General Provisioning Issues” on page 461 for information on how your storage system handles host security groups.
 2. Click **Step 4**.
 3. Click the **Create Job** button.
The job is added to the bottom pane.
 4. Do one of the following:
 - ☐ If you want the job to execute now, click the Execute Job () button
 - ☐ If you want the job to execute at a later time, schedule the job as described in the topic, “Scheduling Provisioning Jobs” on page 521.
-

Step 4 - Select a Zone

Note: If the zone has already been selected and Step 5 is clicked, skip this step or click the  button to clear the selection.

If you do not have a zone selected or you do not have one created yet, most likely the option, **Automatically Configure Zoning**, is not selected in the Zone Customize Dialog (). The management server assumes you want to select a pre-existing zone or create one manually when you deselect the **Automatically Configure Zoning** option.

- **To reuse a zone** - Select a zone in the Zone pane and then click **Step 5**. Expand a fabric node to view its zones.
- **To create a zone** - Select a fabric in the zone pane, click the  button, type a zone name. See the following topic for more information.

If you see the  icon, the fabric cannot be reached.

Table 11-4: Zone Icons

Icon	Description
	Its location on the page determines its use: ■ Displayed above Zone pane - Button for creating zones. ■ In Zone pane - Icon for a zone.
	Zone Alias
	Port

Important: (McDATA switches only) Path Provisioning looks for the names of the active zone set and of the active zones and all of their saved counterparts in the zoning library in EFC Manager. The provisioning job only occurs if those names match.

Creating a Zone

To create a zone:

1. Select a fabric in the zone pane.




Figure 11-9: Selecting a Fabric

2. Click the  button located above the Zone pane.
3. Type a zone name in the dialog box. The following restrictions apply:

Naming Conventions for Brocade Switches:

- ☐ The name must contain 1 to 64 characters.
- ☐ The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or an underscore (_).
- ☐ The name is case sensitive. For example, "Zone1" and "zone1" are different zones.
- ☐ You cannot create a zone with the same name as an existing zone, zone alias or zone set. For example, if you create a zone named "new", you cannot give a zone, zone alias, or zone set the same name.
- ☐ The following characters are invalid for Brocade switches: caret (^), dash (-), and dollar sign (\$).

Naming Conventions for McDATA and Connectrix Switches:

- ☐ The name can have a maximum of 64 characters.
- ☐ The first character of a zone name must be a letter (A-Z, AZ).
- ☐ A zone name cannot contain spaces.
- ☐ Valid characters are a-a, AA, 0-9, caret (^), dash (-), underscore (_), and dollar sign (\$).

- ☐ All names must be unique and may not differ by case. For example, myzone and MyZone are considered to be the same zone.
- 4. Click **OK**.
The new zone is added to the Zone pane.
- 5. Click the **Create Job** button.
The job is added to the bottom pane.
- 6. Do one of the following:
 - ☐ If you want the job to execute now, click the Execute Job () button
 - ☐ If you want the job to execute at a later time, schedule the job as described in the topic, "Scheduling Provisioning Jobs" on page 521.

Volume Assignment

Keep in mind the following:

- If you have options still selected from a previous job, just clear the options you do not want in your next job. For example, assume you created a volume. Now you want to create a new volume on the same host used previously. Just clear the Volume pane. To clear a pane, click the  button.
- If you want to clear all the steps, except for the Step 1 (storage systems), select another option from the System Action combo-box.
- (HDS only) Before you can create a volume. You must delete some unmapped LDEVs using the standard provisioning tool. To delete LDEVs, click **Provisioning** (). Then, click the Storage Systems tab, then the **Provision** button for the storage system. Then, click **Step 2 Volume**. Click the **Unmapped Volumes** folder in the left pane. Select the desired number of LDEVs for the LUSE volume then click the **Delete Selected Volumes** button. Take note of the array group from which you deleted the LDEVs. You need this information to create the LUSE volume.

You can assign a volume to existing host security groups, as described in the following steps.

1. Click **Provisioning** (.
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. Select the following from the System Action combo-box: Volume Assignment

Step 1 - Select Storage System

1. Wait for the management server to load the storage systems into the Storage System panel.
 2. Select the storage system on which you want to create the metavolume.
-

Note: The S heading in the Storage Systems pane means that only a single selection is allowed.

(HDS only) Select the storage system from which you want to create the LUSE volume.

3. Click the **Step 1** button underneath the pane.
Information about the selected storage system, such as its model and IP address is displayed below the Storage System pane.




Figure 11-10: Selecting a Storage System

Step 2 - Select a Volume

To select a volume:

1. In the Volume pane select mapped and unmapped volumes. You can select multiple volumes on Windows computers by pressing the Ctrl key as you select the volumes.
 - **Mapped** - There are two types of mapped volumes:
 - **Masked** - The volume is exposed to the storage port and to the host.
 - **Unmasked** - The volume is exposed to the storage port, but not to the host.

- ❑ **Unmapped** – The volume is not exposed to the storage port. The management server puts all unmapped volumes in this category when it first discovers an array.
- ❑ **Free Extents** – Available free extents that can be used to create a meta volume. You can create meta volumes on EMC Symmetrix and LUSE on HDS storage systems. To create a meta volume or LUSE, select multiple free extents under the Free Extents node in the Volume pane. Select multiple LDEVs from the Free Extents menu by holding down the shift key on your keyboard and selecting free LDEVs. When you select free extents, they must of the same type. For example, on Symmetrix, you cannot select a mirrored volume and a BCV (business continuous volume) to create a meta volume.

When you first discover a storage system, no free extents are displayed. This is because the management server puts all unmapped volumes into the “unmapped” category for the list of volumes by default. To move a volume to the free extent node, delete the unmapped volume. When the volume is deleted, it is moved to the free extents node. The free extents category is used internally by the management server. See “Deleting a Storage Volume” on page 445 for more information.

Important: Make sure the free extents you select are not being used. Data on the free extents becomes unusable when a meta volume is created

Keep in mind the following:

- ❑ You can narrow the type of volumes displayed in the Volumes pane by using the Volume Customize Dialog () located above the Volumes pane. See “Host Customize Dialog” on page 528.
- ❑ If the LUN has already been selected and Step 4 is clicked, skip this step or click the  button.




Figure 11-11: Selecting a Volume

2. Click **Step 3**.
3. If you are asked to specify a LUN number, provide a LUN for each volume displayed. See "Providing a LUN Number" on page 516 for information about numbering LUNs.

Step 3 - Select a Host Security Group

1. Select a host security group in the LUN pane. See “Creating a Host Security Group” on page 520 for information on how to create a host security group. See “General Provisioning Issues” on page 461 for information on how your storage system handles host security groups.
2. Click **Step 4**.
3. Click the **Create Job** button.
The job is added to the bottom pane.
4. Do one of the following:
 - ☐ If you want the job to execute now, click the Execute Job () button
 - ☐ If you want the job to execute at a later time, schedule the job as described in the topic, “Scheduling Provisioning Jobs” on page 521.

Providing a LUN Number

Keep in mind the following when providing a LUN number:

- **Unique numbers:** LUN numbers must be unique.
- **HDS 9200 series storage systems:** LUN numbers must be between 0 and 128.
- **HDS 9500V series storage systems:** LUN numbers must be between 0 and 512.
- **Engenio storage systems:** LUN numbers must be between 0 and 31.
- **Symmetrix storage systems:** LUN numbers must be between 1 and 8190.

You can enter a LUN number for a volume by placing the cursor under the LUN Number column, as shown in the following figure. Click **OK** when you are done. The window displayed in the following figure is displayed if your storage system requires you to provide a LUN number.




Figure 11-12: Specifying a LUN Number

Creating a System Action Template

You can create your own system action template, as described in the following steps.

1. Click **Provisioning** ().
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. Click the **Configure Templates** button at the top of the screen.
4. Click the **New Template** button in the Provisioning Template Configuration window.
5. If you want to change the name that was assigned to the new template, enter the name you want for the template in the **Template Name** field. Then, click **Apply**.
6. Select a master template from which you want to base your new template. For more information about the default master templates, see "Default System Action Templates" on page 479.
7. Click the following tabs and select the options you want for your template.

Note: Not all tabs are available for some default templates.

- ☐ **Storage Options** - See “Storage System Customize Dialog” on page 528 for more information for these options.
 - ☐ **Host Options** - See “Host Customize Dialog” on page 528 for more information about these options.
 - ☐ **Volume Options** - See “Volume Customize Dialog” on page 528 for more information about these options.
 - ☐ **Host Security Options** - See “LUN Customize Dialog” on page 529 for more information about these options.
 - ☐ **Zone Options** - See “Zone Customize Dialog” on page 530 for more information about these options.
8. When you are done, do one of the following:
- ☐ Click **Apply** if you want to apply your changes and keep the Provisioning Template Configuration window open.
 - ☐ Click **OK** if you want to apply your changes and leave the Provisioning Template Configuration window.
 - ☐ Click **Cancel** if you want to leave the Provisioning Template window without saving your changes. You will lose your changes unless you had previously clicked **Apply**.
- When you want to use your new template, select your new template from the **System Actions** drop-down menu.

Modifying a System Action Template

You can modify only system action templates you have created.

To modify a system action template:

1. Click **Provisioning** ().
 2. In the right pane, click **Start Here** on the Path Provisioning tab.
 3. Click the **Configure Templates** button at the top of the screen.
 4. Select the template you want to modify in the **Provisioning Templates** panel.
 5. Make the necessary changes.
 6. Click **Apply**.
 7. When you are done, do one of the following:
-

- ❑ Click **Apply** if you want to apply your changes and keep the Provisioning Template Configuration window open.
- ❑ Click **OK** if you want to apply your changes and leave the Provisioning Template Configuration window.
- ❑ Click **Cancel** if you want to leave the Provisioning Template window without saving your changes. You will lose your changes unless you had previously clicked **Apply**.

Adding a Host

Important: Make sure the added host is physically connected to the network before the scheduled job runs.

The management server lets you add hosts that are not currently connected to the network. While you are creating a job, add the host. This feature is only available when you select **LUN Security** from the System Action menu.

1. Click the  button.
 2. Type the name for the host in the **Host Name** field.
 3. Type the port WWN of the host in the **Port WWN** field.
 4. Click the **Add** button.
 5. Repeat Steps 3 and 4 for multiple ports.
 6. If you want to remove the host, click the  button.
 7. When you are done with your changes, click **OK**.
The host is added to the list of hosts.
 8. Physically connect the host to the network before the job runs.
-

Creating a Host Security Group

You cannot assign the host mode for an IBM storage system in the user interface. You must modify an internal property to set the host mode. See “Setting the Host Mode for IBM Storage Systems” on page 460 for more information.

Keep in mind the following:

- Each storage system handles host security groups differently. See “General Provisioning Issues” on page 461.
- When the **Volume Creation and LUN Security** option is selected on the System Action menu, you cannot add a host security group to an Engenio storage system or to any storage system that Engenio resells under a different brand.

To create a host security group:

1. Select a storage system in the Storage System pane.
2. Click **Step 1**.
3. Select a host in the Host pane.
4. Click **Step 2**.
5. Select a port in the LUN pane.



Figure 11-13: Selecting a Port

6. Click  at the top of the LUN pane.
7. When you are asked to provide a name for the new host security group, type a unique name for the host security group.

Note: You are not asked for the name of the host security group for Symmetrix storage systems.

Keep in mind the following:

- ☐ The name must contain 1 to 50 characters. If you enter no characters, you are given the option of using a default name.
 - ☐ The first and last letter cannot be spaces
 - ☐ You cannot have the following characters in the name:
 - <
 - >
 - ;
 - :
 - ,
 - |
 - /
 - *
 - ?
 - \
 - \\
 - \t
 - \n
 - \b
8. (HDS) Select the host mode for the host security group.
 9. (HDS) Provide a second host mode if applicable.
 10. (HDS) Click the **Create Host Security Group** button.
 11. For non-HDS storage systems, click **OK**.
The host security group is created in the LUN pane.

Scheduling Provisioning Jobs

Keep in mind the following:

- You must have already created a provisioning job. See “About Path Provisioning” on page 474 for more information.
 - You cannot delete a job once it has started. A job has started if it has a state of “started.”
-

- You can deselect jobs by clicking the check marks or by clicking the Clear Selection () button in the Provision Jobs pane.

You can determine the status of a job by looking in the State column:

- **Created** - The job has been created, but it will not be executed. The job cannot be viewed by others and it is deleted when the Web browser is closed. See “Scheduling Provisioning Jobs” on page 521 for information about changing the state of the job from “created” to “scheduled”.
- **Scheduled** - The job has been tasked to execute at a specified time and date. Jobs are assigned a scheduled state after you select the job and click the **Execute Job** button () button.
- **Started** - The job has started. You cannot delete a job once it has started.
- **Failed** - The job failed.
- **Ended** - The job has finished.

To schedule a provisioning job:

1. Click the **Create Job** button in the lower pane.
The job is given the “created” status. The job, however, is not executed.

Keep in mind the following:

- When you close the Web browser window, all jobs with a status of “created” are erased.
- Other users cannot see a job with a status of “created”.
- Change the schedule of job only if its status is “created”. Once you click the Execute Job () icon, the job is saved in the database for the management server and its status changes from “created” to “scheduled”. Thus, it should not be modified.

If you are unable to click the **Create Job** button, verify all required Step buttons have been clicked. If you are still unable to click the **Create Job** button, verify if the Host Customize Dialog () is selected for multipathing. If the Host Customize Dialog is set for multipathing, select the second path from the Path combo box and repeat the provisioning steps, except the steps for selecting a system action, storage system and host. See “Host Customize Dialog” on page 528.

2. Schedule the job by selecting the job and then clicking the  button in the Provision Jobs pane.
-

- a. (Optional) In the **Time** field, change the time displayed. The management server automatically displays a time five minutes from when you clicked the . Type a time in 24-hour format with hours and minutes separated by a colon. For example, 23:15, for 11:15 p.m.
 - b. Select a date you want the job to start.
The date is selected.
 - c. Click the **Set** button.
If you click the **Set** button after the time has passed, you must reset the time.
3. To execute the job, select the job and then click the Execute Job () button. The job is executed at the scheduled time.
You can obtain more information about a schedule job by selecting it. When you select a scheduled job in the Provision Jobs pane, detailed information appears in the Job Console pane, located in the lower-right corner. The Message Console tab, located also in the lower-right corner, provides information such as if the job has ended or failed. If a job has failed, the reason for the failure is provided on the Message Console tab.

Keep in mind the following:

- ☐ Jobs are executed according to the time set on the management server, since they are centrally saved in the database for the management server.
- ☐ When the management server is determining which job to perform first, it looks first for jobs requiring volume and/or zone creation. It does not perform the jobs in the order suggested by the Provision Job ID nor according to the order of the jobs appearing in the Provision Jobs table.
- ☐ The management server can detect when the required volume and/or zone has already been created. For example, assume you created a job to create a volume. The next few jobs use this volume. These jobs will determine if the volume has been created, before creating the volume.

Executing Provisioning Jobs

If you want to save and execute a job, you must click the **Execute Job** button () button. When you click the **Execute Job** button () button, the job is saved on the management server. Other users can now see the job.

When the management server is determining which job to perform first, it looks first for jobs requiring volume and/or zone creation. It does not perform the jobs in the order suggested by the Provision Job ID nor according to the order of the jobs appearing in the Provision Jobs table.

The management server can detect when the required volume and/or zone has already been created. For example, assume you created a job to create a volume. The next few jobs use this volume. These jobs will determine if the volume has been created, before creating the volume.

You can schedule a job to start now or at a pre-determined time in the future:

- **To start a job now** - Click the **Execute Job** button () button before the job is scheduled.
- **To start a job at a pre-determined time in the future** - Schedule the job first. Then, click the **Execute Job** button () button. The job is executed according to the time set on the management server.

To view the latest information in Path Provisioning, click the **Refresh** button. The **Refresh** button updates the following:

- The Path Provisioning screen with changes made in Provisioning. For example, assume you used the wizards in Provisioning to create a host security group. When you access Path Provisioning, your changes are not shown.
- The Path Provisioning screen with changes from executed jobs. After a job is executed in Path Provisioning, the Path Provisioning screen is not updated until you click the **Refresh** button or you exit and re-enter Path Provisioning.
- Other parts of the product, such as:
 - ☐ Application Explorer
 - ☐ Capacity Explorer
 - ☐ Performance Explorer
 - ☐ Provisioning
 - ☐ System Explorer

Monitoring Provisioning Jobs

If you want to view the latest status of the provisioning jobs listed, click the  button in the Provision Jobs pane. The management server gathers information about the provisioning jobs listed to

determine their latest status. You might want to try this feature when you are not sure if a job has ended.

Deleting Multiple Jobs

Important: You cannot delete a job once it has started. A job has started if it has a state of “started.”

To delete jobs, select the jobs in the M column and then click the  icon until all jobs have been removed, as shown in the following figure.

You can deselect jobs by clicking the check marks or by clicking the Clear Selection () button in the Provision Jobs pane.




Figure 11-14: Deleting Multiple Provisioning Jobs

Naming Conventions for Zones

The following are the naming conventions for zones:

Naming Conventions for Brocade Switches:

- The name must contain 1 to 64 characters.
- The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or an underscore (_).
- The name is case sensitive. For example, “Zone1” and “zone1” are different zones.
- You cannot create a zone with the same name as an existing zone, zone alias or zone set. For example, if you create a zone named “new”, you cannot give a zone, zone alias, or zone set the same name.
- The following characters are invalid for Brocade switches: caret (^), dash (-), and dollar sign (\$).

Naming Conventions for McDATA and Connectrix Switches:

- The name can have a maximum of 64 characters.
- The first character of a zone name must be a letter (A-Z, AZ).
- A zone name cannot contain spaces.
- Valid characters are a-a, AA, 0-9, caret (^), dash (-), underscore (_), and dollar sign (\$).
- All names must be unique and may not differ by case. For example, myzone and MyZone are considered to be the same zone.

Using Multipathing with Path Provisioning

You can use provisioning with multipathing. Set the Host Customize Dialog to the multipathing option. Then, repeat the provisioning steps for each path, as described in the following steps:

1. Select one of the following system actions:
 - **Volume Creation + LUN Security + Zoning** - First you create a meta volume. Second, you map a meta volume to a Fibre Channel port and host HBAs (HSG). Third, you create a zone. See “Volume Creation, LUN Security, and Zone Operation” on page 479 for more information.
 - **LUN Security** - Map meta volume to Fibre Channel port and host HBAs (HSG). See the topic “LUN Security” on page 488 for more information.
 - **Zone Operation** - Create a new zone. See “Zone Operation” on page 493 for more information.
 2. Select a storage system, and then click the **Step 1** button.
 3. Click the  button above the Host pane.
-

-
- The screenshot shows the 'Provision Jobs' window. At the top, there is a toolbar with icons for a folder, calendar, mail, floppy disk, and trash. Below the toolbar is a table with columns: M, ID, Description, Path 1, and Path 2. The first row is selected, and a context menu is open over it. The context menu has two options: 'Path 1' and 'Path 2'. The 'Path 1' option is currently selected. The 'Create Job' button is visible in the top right corner.
- | M | ID | Description | Path 1 | Path 2 |
|-------------------------------------|----|--|--------|----------------|
| ☒ | 0 | Path 1 Created Lun Mapping Task: StorageVolumeld = 16676 | Path 1 | 04/02/25 17:20 |
| ☐ | | | | |
| ☐ | | | | |
| ☐ | | | | |
| ☐ | | | | |

8. Repeat Step 6.
You do not need to select a system action, storage system, and host. Furthermore, you must select the same host used for the first path.
9. Set the schedule for the job as described in the topic, “Scheduling Provisioning Jobs” on page 521.

Customizing Path Provisioning

- Storage System Customize Dialog on page 528
- Host Customize Dialog on page 528
- Volume Customize Dialog on page 528

- LUN Customize Dialog on page 529
- Zone Customize Dialog on page 530

Storage System Customize Dialog

The Storage System Customize Dialog lets you specify the type of storage systems displayed in the Storage System pane. Select one of the following:

- **Show all available storage systems** - All storage systems are displayed.
- **Show storage system with available raw storage** - Only storage systems with available raw storage are displayed.
- **Show storage system with available unmapped storage** - Only storage systems with available unmapped storage are displayed.
- **Show storage system with the following characteristics** - Select one or more of the storage tiers. Only storage systems belonging to the selected storage tiers are displayed. Storage systems are assigned to storage tiers in Chargeback. This menu is blank if no storage systems have been assigned to storage tiers. See “Adding Asset Information” on page 724 and “Adding General Information” on page 727 for information about assigning storage systems to a storage tier.

Click **OK** when you are done with your selections.

The Storage System pane is updated.

Host Customize Dialog

The Host Customize Dialog lets you use multipathing with Path Provisioning. Select one of the following options:

- **Single path: Select one port of a host.** Select this option if you do not have multipathing or you do not want to use multipathing with Path Provisioning.
 - **Multipath: Select more than one port within a single server.** Select this option if you want to use multipathing with Path Provisioning. See “Using Multipathing with Path Provisioning” on page 526 for more information on how to use multipathing with Path Provisioning.
-

Volume Customize Dialog

1. If you want to view all volumes, select the **Show All Volumes** option. Once you have selected the **Show All Volumes** option, you have the option to select the **Hide mapped volumes** option, which hides mapped volumes.
2. Select one of the following for metavolumes:
 - ☐ **Concatenating** - Only concatenating metavolumes are displayed.
 - ☐ **Striped** - Only striped metavolumes are displayed. (Applies only to EMC storage systems)
3. Click **OK** when you are done with your selections.
The Volume pane is updated.

LUN Customize Dialog

1. Specify how LUN are mapped by selecting one or more of the following options in the LUN customize dialog:
 - ☐ **Perform number of path verification based on host selection** - The path verification is based on the host you selected.
 - ☐ **Automatically Assign Volume to Storage Port based on** - Select one of the following:
 - ***Most unused ports (the number of LUNs assigned)** - Assigns a volume to a FA port based on how often a port is used. Unused ports have a better chance of having a volume assigned to them as compared to frequently used ports. This option helps you spread out the traffic.
 - ***Linked port if there is any** - A linked port is more likely to be assigned a volume. The management server looks for any storage system port that is zoned to the selected host bus adapter. If the management server cannot find a storage system port zoned to the selected HBA, the management server selects a port with the least amount of LUNs.

*You can only access these options if **Automatically Assign Volume to Storage Port Based on** is selected.

 - ☐ **Verify that in a multipath configuration, storage system ports do not connect to the same switch** - This option makes sure the multipath configuration is preserved. A storage system ports should not get connected to the same switch.
 - ☐ **Assign a LUN number automatically** - Do not select this option if you want to assign a LUN number manually.
2. Click **OK** when you are done with your selections.

Zone Customize Dialog

When the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**, is selected in the zone customize dialog () , the management server automatically selects a zone that meets its criteria. If the management server cannot find a zone that meets its criteria, it creates a zone on the fly. You set the criteria for automatically configuring zones by selecting one of the following options:

- ***Create new zone if there is no zone containing ONLY the selected zone members (HBA and storage system ports)** - The management server checks if an identical zone exists. An identical zone contains only same HBA and storage system ports you selected. If the zone contains additional members, it is not considered to be identical.
 - ☐ If an identical zone exists, it is selected in the Zone pane.
 - ☐ If an identical zone does not exist, the management server asks you to name a zone that will be created. The new zone appears in the Zone pane, even though it is not created until the job runs.
- ***Create new zone if there is no zone containing AT LEAST the selected zone members** - The management server tries to find a zone that contains the HBA and storage system ports you selected.
 - ☐ If a zone contains additional members, the management server selects that zone in the Zone pane.
 - ☐ If the management cannot find a zone containing the HBA and storage system ports you selected, the management server asks you to name a zone that will be created. The new zone appears in the Zone pane, even though it is not created until the job runs.

*To enable one of these options, the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**, must be selected.

Important: (McDATA switches only) Path Provisioning looks for the names of the active zone set and of the active zones and all of their saved counterparts in the zoning library in EFC Manager. The provisioning job only occurs if those names match.

About the Use Switch Port Zoning Mode Option

By default Path Provisioning creates zones through Worldwide Names, which means it looks for a port on a host and a port on a storage system to create the zone. Zoning through WWNs is not

dependent upon the switch. This means you could change switches and as long as the host and storage system are able to access each other through the network.

You can create zones through switch port zoning if you select the **Use Switch Port Zoning Mode** option. When you select the **Use Switch Port Zoning Mode** option, you select ports on a switch to create your zone. For example, ports 1 and 2 on a switch could be designated for a zone. This means that whatever is connected to either of those ports will be a part of the zone.

Automatically Configure Zoning

To automatically configure zoning:

1. Click the  button above the Zone pane before you select a host.
2. Select the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**.
3. Select one of the options mentioned in the previous bulleted list. Then, click **Apply**.
4. Select a storage system and click **Step 1**.
5. Select a host and click **Step 2**.
6. One of the following occurs:
 - ☐ If the management server cannot find a zone that meets the criteria set in Step 3, it asks for a zone name.
 - ☐ If the management server finds a zone that meets the criteria set in Step 3, it selects the zone in the Zone pane.
7. If you are asked for a zone name, type the zone name and then click **OK**.
See “Naming Conventions for Zones” on page 525 for restrictions on naming zones.

The new zone is displayed in the Zone pane, but it has not been created yet. It will be created when the job runs.

Manually Configure Zoning

The management server assumes you want to select a pre-existing zone or create one manually when you deselect the **Automatically create new zone if no existing zone containing HBA and storage system ports is detected** option.

- **To reuse a zone** - Select a zone in the Zone pane and then click **Step 5**. Expand a fabric node to view its zones.

- **To create a zone** - Select a fabric in the zone pane, click the  button, type a zone name. See the following topic for more information.

If you see the  icon, the fabric cannot be reached.

Table 11-5: Zone Icons

Icon	Description
	Its location on the page determines its use: ■ Displayed above Zone pane - Button for creating zones. ■ In Zone pane - Icon for a zone.
	Zone Alias
	Port

To create a zone:

1. Select a fabric in the zone pane.




Figure 11-16: Selecting a Fabric

2. Click the  button located above the Zone pane.
3. Type a zone name in the dialog box. The following restrictions apply:

Naming Conventions for Brocade Switches:

 - The name must contain 1 to 64 characters.

- ❑ The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or an underscore (_).
- ❑ The name is case sensitive. For example, "Zone1" and "zone1" are different zones.
- ❑ You cannot create a zone with the same name as an existing zone, zone alias or zone set. For example, if you create a zone named "new", you cannot give a zone, zone alias, or zone set the same name.
- ❑ The following characters are invalid for Brocade switches: caret (^), dash (-), and dollar sign (\$).

Naming Conventions for McDATA and Connectrix Switches:

- ❑ The name can have a maximum of 64 characters.
 - ❑ The first character of a zone name must be a letter (A-Z, AZ).
 - ❑ A zone name cannot contain spaces.
 - ❑ Valid characters are a-a, AA, 0-9, caret (^), dash (-), underscore (_), and dollar sign (\$).
 - ❑ All names must be unique and may not differ by case. For example, myzone and MyZone are considered to be the same zone.
4. Click **OK**.
The new zone is added to the Zone pane.

Assigning Templates to a Role

You can assign templates to a role to restrict a user's access to all templates. For example, you could specify that only users belonging to a certain role can only access the Volume Assignment template.

Important: You must belong to a role that has "Provisioning Administration" privileges to be able to assign templates to a role. The Domain Administrator and Storage Administrators roles have these privileges by default. To determine if your role has "Provisioning Administration" privileges, your domain administrator will need to go to the **Security > Users** page. Then, the domain administrator must click the name of the role to see which privileges are assigned to your role.

1. Click **Provisioning** ().
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. Click **Assign Templates**.
4. Select the role from the **Select Role** drop-down menu you want to have access to the Path Provisioning template.

For a role to appear in the **Select Role** drop-down menu, it must have the following attributes:

- ☐ Has Provisioning privileges, meaning it has enough privileges to access Path Provisioning.
 - ☐ Does not have Provisioning Administration privileges. Roles that have Provisioning Administration privileges do not appear in the **Select Role** drop-down menu.
5. In the Provisioning Templates panel select the templates you want users belonging to the role you selected in Step 4 to be able to access. Then, click **Add**.
The templates the users belonging to the selected role will be able to access appear in Assigned Provisioning Templates panel.

You can also click the following buttons:

- ☐ **Add All >>** -- Makes all the templates available to users assigned to the selected role.
 - ☐ **< Remove** - Select a template in the Assigned Provisioning Templates panel and click the **< Remove** button to make a template not available to users assigned to the selected role.
 - ☐ **<< Remove All** - Makes all the templates unavailable to users assigned to the selected role.
6. When you are done, do one of the following:
 - ☐ Click **Apply** if you want to apply your changes and keep the Provisioning Template Configuration window open.
 - ☐ Click **OK** if you want to apply your changes and leave the Provisioning Template Configuration window.
 - ☐ Click **Cancel** if you want to leave the Provisioning Template window without saving your changes. You will lose your changes unless you had previously clicked **Apply**.
-

This chapter describes the following:

- About Event Manager on page 536
 - Viewing Event Details on page 540
 - Clearing Events on page 542
 - Configuring the Clearing of Events on page 543
 - Configuring the Deletion of Events on page 544
 - Deleting Events on page 545
 - Sorting Events on page 545
 - Adding Journal Entries on page 546
 - Changing the CLARiiON Event Polling Interval on page 546
 - Brocade Events on page 548
 - Filtering Events on page 550
 - Enabling the Rank Column in Event Manager on page 556
-

About Event Manager

Event Manager lets you view, clear, sort and filter events from managed elements. An event can be anything that occurs on the element, such as a device connected to a Brocade switch has gone off-line. It provides the following information about the events:

- **ID** - The identification number assigned to the event
- **Element** - The source of the event. An element can be a switch, host, application, fabric or anything else on the network.
- **Severity** - Provides the severity level
- **Time** - The time the event was recorded.
- **Summary Text** - A brief explanation of the event. When you click the summary text, the details of the event are displayed.
- **Element Type** - Specifies whether the source of this event is an application, a host, etc.
- **Rank** - The cost implication. This column is hidden until you enable it, as described in “Enabling the Rank Column in Event Manager” on page 556.

The top of the Event Manager pane provides a subtotal of the severity of the events and the number of events from the different element types. For example, in the following figure, zero critical events were found, but 23 events were found to have a severity level of minor and one of “Warning”. The definition for each severity level varies according to the type of element.




Figure 12-1: Event Manager Summary

Note: The name for the management server is displayed for <Product Name>.

Access Event Manager by doing one of the following:

- Click **Event Manager** (⚠) in the left pane.
- To view events from a specific element:
 - Double-click the element in Capacity Explorer, Performance Explorer or System Explorer and then click the **Events** tab. Only events from the element that was double-clicked are displayed. You cannot access the **Element Type** filter through this method.

- Right-click the element in Capacity Explorer, Performance Explorer or System Explorer and then select the **Show Events** option in the menu.

You can change your user preferences for Event Manager by doing one of the following:

- Clicking the Preferences link in Event Manager. See “Modifying Your User Preferences” on page 167 for more information.
- Clicking **Configuration > Events**. See “Controlling the Display of Cleared and Deleted Events” on page 222.

The following icons are displayed in Event Manager.

Table 12-1: About the Icons in Event Manager

Icon	Description
	Event was marked cleared. See the topic, “Clearing Events” on page 542 for more information.
	The severity of the event is not known.
	The event is informational.
	The event might have some impact.
	The event has a minor impact.
	The event has a major impact.
	The event has a critical impact.

Event Manager does not support events from all discovered elements. See the following table for more information about which elements Event Manager supports.

Table 12-2: Supported Hardware

Hardware	Events Supported?
Brocade switches	Y, ¹
⁴ QLogic SNMP switches	Y, ³

Table 12-2: Supported Hardware (Continued)

Hardware	Events Supported?
⁵ McDATA switches SWAPI to EFCM	Y, ²
⁵ McDATA switches SNMP through proxy	Y, ³
⁵ McDATA switches SNMP to switches	Y, ³
Cisco switches SNMP	Y, ³
CNT switches SMI	N
Engenio and Sun 6130 storage systems	Y
Tape Libraries	N
HP NAS Filers	N
NetApp Filers ⁶	Y
Sun NAS Devices	N
CLARiiON storage systems	Y
Symmetrix storage systems	Y
HDS storage systems	Y
Sun 35xx storage systems	N
Sun 6920 and 6940 storage systems	N
HP XP with Command View XP	N
HP XP with Command View XP AE	Y
HP XP with XP Provider for SE	Y
HP MSA storage systems	N
HP EVA storage systems	Y
IBM ESS storage systems	N
Xiotech storage systems	N

Notes from the table:

¹Not all events that show up in the webtool appears in Event Manager.

²Not all events that show up in EFCM will appear in Event Manager.

³Need to configure the switch or proxy to send traps to the management server.

⁴Also applies to Sun StorEdge switches.

⁵Also applies to EMC Connectrix switches.

⁶To receive events from NetApp filers, you must add the management server address as an SNMP trap host on the NetApp filer.

Keep in mind the following:

- By default the management server displays events from all of the elements, except those from itself. If you want to view events from the management server
 - Select the **All** or **<Product Name of the Management Server>** option from the **Show Element Type** drop-down menu. See “Selecting an Element Type for Events” on page 553 for more information about the menu options.
 - Click the **Customize** button next to the **Show Element Type** drop-down menu in Event Manager. Select the management server and then click **OK**.

When you are asked if you want to apply your changes, click **Yes**. Then, click the **Apply Filters** button.

- The management server provides separate event notification for every event that is reported from the devices it is monitoring. Excessive notification could delay provisioning, as the providers are busy notifying the management server of the events. If you do not want the management server to be notified of every event, modify the event threshold of the devices to filter out some of the events. Refer to the documentation accompanying the device for more information about setting the threshold.
- To receive events from Sun StorEdge or QLogic switches, verify the SNMP trap community string is set to public in SANbox Manager or through telnet. Also, make sure the SNMP traps are configured to be sent to the management server.
- The management server does not receive SNMP v1 traps from Sun StorEdge switches that have the SNMPv1 agent.
- If you want the management server to be able to receive events from a NetApp Filer, you must add the IP address of the management server CIMOM to the NetApp configuration. The management server CIMOM runs on the same computer running the management server by default.
- Event Manager displays events from all elements regardless of the user's organization.
- Event Manager displays events it receives from CNT InVsn Enterprise Manager. As of version 9.5 of CNT InVsn Enterprise Manager, InVsn Enterprise Manager does not provide

events to the management server. As future versions of CNT InVsn Enterprise Manager provide event support, the management server will be able to provide information about those events.

Issues with Brocade Switches

- Event Manager does not display events from Brocade switches with the firmware version 3.0. This firmware version is not supported by Event Manager. You can, however, specify Event Manager not to display events from additional firmware versions. See “Filtering Events” on page 550.
- Events from Brocade Fabric Watch have “Fabric watch:” in their heading, as shown in the following example:

```
Fabric watch: threshold is below low boundary, Value:30 Class: Environment  
ClassArea: Temperature Index: 1 EventType: Below EventState: Faulty
```

Issues with McDATA and Connectrix Switches

- If you discovered McDATA and/or Connectrix switches through SWAPI, Event Manager does not report events for switch hardware failures, except for those regarding switch ports (port off-line/port online).
- If you are not receiving events from McDATA and Connectrix switches discovered through SNMP, make sure you have the correct port set for receiving SNMP traps. You must also configure the proxy or each switch (depending on your configuration for discovery) to send the SNMP traps to the correct port on the management server. See “Discovering McDATA and EMC Connectrix Switches” on page 47 for more information.
- If you want the management server to receive SNMP events from Connectrix or McDATA switches, do one of the following:
 - If you discovered Connectrix Manager or EFC Manager, only enable SNMP trap forwarding to the management server on the Connectrix Manager or EFC Manager, not on the individual switches. Connectrix Manager or EFC Manager should be configured to forward SNMP traps to the IP address of the management server, and the community string should match the user ID you used to discover Connectrix Manager or EFC Manager.
 - If you discovered Connectrix or McDATA switches directly, enable SNMP trap forwarding on the switches, not on any other management software. The switches should be configured to forward SNMP traps to the IP address of the management server, and the community string should match the user ID you used to discover the Connectrix or McDATA switches.

Viewing Event Details

To access event details:

1. Access Event Manager as described in “About Event Manager” on page 536.
2. Do one of the following in Event Manager:
 - Right-click an event, and select **Event Details** from the drop down menu, as shown in the following figure:




<u>Element</u>	<u>Severity</u>	<u>Time</u>
QBrocade2		2004/04/20 2
QBrocade2		2

Context menu options: Event Details, Mark as Cleared

Figure 12-2: Right-clicking Event Details

- In Event Manager click the event summary, as shown in the following figure.



☐	<u>ID</u>	<u>Element</u>	<u>Severity</u>	<u>Time</u>	<u>Summary Text</u>
☐	4219	QBrocade2		2004/04/20 21:10	2004/04/20 17:29 Logout

Figure 12-3: Accessing Event Details

The event details are displayed.

The Event Details pane provides information on one or more of the following:

- **Serial Number** - The number assigned to the event
- **Summary Text** - A brief explanation of the event
- **Element** - It is the source of the event. An element can be a switch, host, application, fabric or anything else on the network. If this field is blank, the event did not come from an element.
- **Source** - If the event came from an element, the name of the element is provided.
- **Time Reported** - The time the event was reported to the management server.
- **Probable Cause** - Provides an explanation of a probable cause.

- **Recommended Actions** - Provides recommendations.
- **Event Type** - Provides a description of the event.
- **Correlated Indications** - Information obtained from the provider.
- **Severity** - Provides the severity level, which can be one of the following:
 - ☐ **Clear**
 - ☐ **Unknown**
 - ☐ **Informational**
 - ☐ **Warning**
 - ☐ **Minor**
 - ☐ **Major**
 - ☐ **Critical**
- **Journal Entries** - Use this field to enter additional information and then click **Add Journal Entry**. This field is limited to 4,000 characters.

Note: Events listed in Event Manager may not be attributed to the correct source until Get Details has completed.

Clearing Events

You might want to mark an event as “cleared” if you have already reviewed it. Use one of the following techniques to mark an event as “cleared”:

- Click the **Clear this Event** button in the Event Details pane, which is accessible by clicking the summary text link.
- Right-click the event listed in Event Manager and select **Mark as Cleared** from the drop-down menu. This menu does not appear if you right-click the summary text link.
- Select the check box to the left of the event and click the **Clear** button in Event Manager.

Note: You can clear more than one event by selecting the check box to the left of the ID column header and then clicking the **Clear** button in Event Manager.

Unclearing Events

You can unclear events by doing one of the following:

- Click the **Set this Event to Not Cleared** button in the Event Details pane
- Select the check box to the left of the event and click the **Undo Clear** button in Event Manager. You can unclear multiple events at once by select them and then clicking the **Undo Clear** button.
- Right-click an event in Event Manager and select the **Mark as NOT Cleared** option from the drop-down menu. This menu does not appear if you right-click the summary text link.

Configuring the Clearing of Events

Depending on the severity of an event, the management server may mark the event as “clear” after 60 minutes. Events designated as Major and Critical are never marked as “clear.” You can change the time delay in clearing an event, and you can specify that the management server never marks an event as “clear”.

To help you in filtering events, you may want to have unimportant events marked as cleared rather than automatically deleted. Depending on how you have configured the deletion of events, you can view the “cleared” events at a later time.

See the following table for the default settings for clearing events.

Table 12-3: Default Settings for Clearing Events

Severity Level	Default Time Delay to Clear the Event (Hours)
Unknown	1
Informational	1
Warning	1
Minor	1
Major	Never
Critical	Never

To change the default time delay to clear an event, do the following:

1. Access the Events page by clicking **Configuration > Events**.
2. Do one of the following:
 - ☐ If you never want an event of the specified severity level marked as “cleared,” select the **Never** option next to the severity level.
 - ☐ If you want to change the delay time in clearing an event, select one of the following units of measurement from the combo box and type the number in the adjacent field:
 - **Minutes**
 - **Hours**
 - **Days**
 - **Weeks**
3. Click the **Save Changes** button.

Configuring the Deletion of Events

The management server automatically deletes events after two weeks by default. You can specify for each severity level different time periods for deleting events. For example, you could modify the management server to delete events with the Information severity level every two days. You could also specify the management server to never delete events with the Critical severity level.

To change the default time delay to delete an event, do the following:

1. Access the Events page by clicking **Configuration > Events**.
2. Do one of the following:
 - ☐ If you never want an event of the specified severity level automatically deleted, select the **Never** option under the Automatic Delete Delay column.
 - ☐ If you want to change the delay time in deleting an event, select one of the following units of measurement from the combo box and type the number in the adjacent field:
 - **Minutes**
 - **Hours**
 - **Days**
 - **Weeks**

For example, if you want events that are a week old deleted, you would type 1 and then select **Weeks** in the combo box in the Automatic Delete Delay column.

3. Click the **Save Changes** button.

Deleting Events

To delete an event:

1. (Optional) Select the events you want to delete in Event Manager.
2. Click the **Delete** button at the top of the screen.
3. Select one of the following options in the Select Events to Delete window:
 - ☐ **Delete selected events** - All selected elements are deleted.
 - ☐ **Delete all events** - All elements are deleted, regardless of your filter settings.
 - ☐ **Delete all events earlier than** - Elements earlier than the date specified are deleted, regardless of your filter settings. Click the calendar icon, .
4. Click **OK**.

Sorting Events

In Event Manager, you can sort events. For example, if you want to see the most severe event on a page, click the Severity column header link. Click it again to sort events in the reverse order.

To sort the events:

1. Access Event Manager as described in “About Event Manager” on page 536.
 2. In the Event Manager table, click the column title corresponding to the attribute you want to sort the event. The following is a listing of the attributes:
 - ☐ **ID** - The identification number assigned to the event
 - ☐ **Element** - The source of the event. An element can be a switch, host, application, fabric or anything else on the network.
 - ☐ **Severity** - Provides the severity level
 - ☐ **Time** - The time the event was recorded.
-

- ❑ **Summary Text** - A brief explanation of the event. When you click the summary text, the details of the event are displayed.
- ❑ **Element Type** - Specifies whether the source of this event is an application, a host, etc.
- ❑ **Rank** - The cost implication. This column is hidden until you enable it, as described in “Enabling the Rank Column in Event Manager” on page 556.

Adding Journal Entries

While you are tracking an event, add journal entries to make others aware of what you are doing and to prevent others from repeating your steps. For example, assume a host went down. You could use journal entries as a way to track your steps. Others would know what you did to get the host running. They could use this information to solve problems with other hosts.

To add a journal entry to an event:

1. Access Event Manager as described in “About Event Manager” on page 536.
2. Access the Event Details page, by doing one of the following:
 - ❑ Clicking a link in the **Summary Text** column.
 - ❑ Right-clicking an event listing and selecting **Event Entries** from the drop-down menu.
3. In the **Journal Entries** field, type the entry for the event. This field is limited to 4,000 characters.
4. Click the **Add Journal Entry** button.
The entry is added with the user's account name and the date and time it was added.

Changing the CLARiiON Event Polling Interval

You can change how frequently the management server polls the CLARiiON storage systems by modifying the `cimom.ClariionEventPollInterval` property. You may want to change this interval if you are receiving many “information” messages from the CLARiiON storage system.

Important: Do not set a very long time interval. The management server does not become aware of events occurring on CLARiiON storage system until it polls the storage system. For example, assume you set the polling interval to every two days. A serious issue could occur on the first day, but you would not know about it until the second day because you set a long time interval.

To change the polling interval:

1. Click **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the `cimom.ClariionEventPollInterval` property. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
4. Return to the Advanced page (**Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree).
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Make your changes in the **Custom Properties** field. Make sure the property is not commented out by removing the hash (#) symbol in front of the property.
7. Change the value assigned to the `cimom.ClariionEventPollInterval` property. The value is in milliseconds. In the example below, the polling interval is set to 5 minutes.
`cimom.ClariionEventPollInterval=300000`
8. When you are done, click **Save**.
9. Restart the service for the management server for your changes to take effect. While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Important: While the AppStorManager service is stopped, the following occurs:

- Users are not be able to access the management server.
 - The management server is unable to monitor elements at this time.
-

Brocade Events

This section describes the following:

- Brocade Switch Events on page 548
- Supported Brocade Events on page 549

Brocade Switch Events

When a Brocade switch generates an event, it assigns a code to the event instead of an event severity level. So, the software assigns an event severity level to the event according to the event's code. This lets you filter Brocade switch events by severity level in Event Manager, as described in the following table.

Note: Events regarding firmware downloads are removed from the following table since the management server cannot be made aware of those events.

Table 12-4: Brocade Switch Events

*Code	Event Severity Level	*Name	Description
0	Minor	EV_OBJ_CHANGED	A physical fabric object (switch port or fabric) has changed state.
1	Major	EV_OBJ_DELETE	One or more physical fabric objects (device port, switch, or fabric) have disappeared.
2	Warning	EV_OBJ_CREATE	One or more new physical fabric objects (device port, switch, or fabric) have appeared.

Table 12-4: Brocade Switch Events (Continued)

*Code	Event Severity Level	*Name	Description
3	Critical	EV_CONNECTED_OBJECT_OFFLINE	A device connected to a switch has gone offline.
4	Major	EV_CONNECTED_OBJECT_ONLINE	A device connected to a switch has come online.
5	Info	EV_RSCN	For those RSCN events not covered by EV_OBJ_xxx codes listed above. Examples include: fabric, domain, connected area state unknown, connected device state unknown.
7	Minor	EV_API_HEART_CONDITION	Monitoring proxy switch.
11	Major	EV_STATE_CHANGE	State changes such as: login, logout, login failed, configuration change, track on/off, port up/down, fabric segmentation, security violations, zone change.
12	Major	EV_PLATFORM_CHANGE	The platform database has changed.

*This information does not appear in the event description, but it appears here for additional information.

Supported Brocade Events

The Event Manager displays the following events from Brocade switches:

- **RSCN events:** Events about the state of the switch, such as it being offline.
- **Fabric Watch events:** Events about the threshold of the switch, such as threshold below.

- **Fabric Access library events:** Events about proxy switch health.
- **Zoning events:** Events about zoning, such as zone-related state change notification.

Filtering Events

This section describes the following:

- About Filtering Events on page 550
- Selecting a Severity Level on page 551
- Customizing the Severity Level Filter on page 552
- Selecting an Element Type for Events on page 553
- Customizing the Element Type Filter for Events on page 554
- Filtering Events by Summary Text or Element Name on page 555
- Hiding Filters in Event Manager on page 556

About Filtering Events

The management server provides several types of event filters to specify which events you want Event Manager to display.

You can use all the event filters at once or you can use just one of them. You can filter events by:

- **Severity level**
- **Element type**

or element name For example, assume you are having problems with a host named WickedWitch, and you want to see only events from WickedWitch and you only want to see events with a severity level of Critical and Major.

You could do the following:

- Set the severity level to critical and major by clicking the **Custom** button next to the **Show Severity** field.
 - Type WickedWitch in the **Summary Text or Element Name** field.
-

Once you set all of your filters, you would click the **Apply Filters** button.

Selecting a Severity Level

You can filter the events, so you only view events that are of a specified severity level. For example, you can specify that Event Manager displays only events with a Critical severity level.

To select a severity level:

1. Select an option from the **Show Severity** drop-down menu in Event Manager. A dialog box displays and allows you to select when you want to apply the new filter settings.
2. If you want to apply the new filters immediately, click the **Yes** button. If you want to apply the new filter settings later, click the **No** button.
3. If you clicked the **No** button in the previous dialog box, click the **Apply Filters** button. The definition for each severity level varies according to the type of element.

Important: Events marked as “Clear” are not displayed if you select a severity level other than **Clear** or **All**.

Table 12-5: Severity Levels

Severity Level	Description
Clear	The event has been marked as cleared.
Unknown	The event does not fall into the other categories and further information can not be obtained from it.
Informational	Provides informational data. For example, for Brocade switches it could be a list of switches that have successfully completed firmware download.
Warning	Provides warning data. For example, for a Brocade switch one or more new physical fabric objects (device port, switch, or fabric) have appeared.

Table 12-5: Severity Levels (Continued)

Severity Level	Description
Minor	Provides a message to indicate a minor problem. For example, for a Brocade switch a physical fabric object (switch port or fabric) has changed state.
Major	Provides a message to indicate a major problem. For example, for a Brocade switch one or more physical fabric objects (device port, switch, or fabric) have disappeared.
Critical	Provides a message to indicate a major problem. For example, for a Brocade switch a device connected to the switch has gone off line.
All	Displays all events, regardless of severity level.
All But Clear	Displays all events, except for those marked cleared.

Customizing the Severity Level Filter

You can customize the severity level in your filter by clicking the **Customize** button next to the **Show Severity** drop-down menu in Event Manager.

For example, you can specify you want Critical and Major events displayed in Event Manager.

Important: Events marked as “Clear” are not displayed if you select a severity level other than **Clear**.

To select more than one severity level for filtering:

1. Click the **Custom** button next to the **Show Severity** drop-down menu in Event Manager.
-

2. In the Event Filters window, make sure the severity levels you want to appear in Event Manager are selected.
For example, in the following figure the filters are selected so that Event Manager displays only events that are critical or major.




Figure 12-4: Selecting Event Filters

Notice in the previous figure that the number of events meeting a criteria, such as Minor, is displayed next to the options.

3. Click **OK**.
4. Click the **Apply Filters** button.
Event Manager displays the events that meet the criteria selected in the Event Filters pane.

Selecting an Element Type for Events

You can filter by element type, so only the events from a certain type of element are displayed. For example, you can specify that only events from hosts are displayed.

To filter by element type, select an option from the **Show Element Type** drop-down menu in Event Manager. When you are asked if you want to apply your changes, click **Yes**. Then, click the **Apply Filters** button.

Note: The **Show Element Type** drop-down menu is not available if you accessed Event Manager by double-clicking an element in Capacity Explorer, Performance Explorer or System Explorer and then clicking the **Events** tab.

Table 12-6: Element Types

Element Type	Description
Applications	Lists only events from applications, such as Microsoft Exchange and Oracle.
Host	Lists only events from hosts.
Switch	Lists only events from switches.
Storage System	Lists only events from storage systems.
Fabric	Lists only events from fabrics.
Other	Lists only events from elements that do not fit in the previous categories.
Management Server*	Lists only events from the management server.
All	Lists all events.
Tape Library	Lists only events from tape libraries.
All But <Product Name of Management Server>	Lists all events, except those from the management server.

*May display the product name instead of “management server.”

Customizing the Element Type Filter for Events

You can customize the element type for your filter by clicking the **Customize** button next to the **Show Element Type** drop-down menu in Event Manager.

For example, you can specify you want Critical and Major events displayed in Event Manager.

Note: The **Show Element Type** drop-down menu is not available if you accessed Event Manager by double-clicking an element in Capacity Explorer, Performance Explorer or System Explorer and then clicking the **Events** tab.

To select more than one severity level for filtering:

1. Click the **Custom** button next to the **Show Element Type** drop-down menu in Event Manager.
2. In the Event Filters window, select the element types from which you want obtain event messages.
3. Click **OK**.
4. When you are asked if you want to apply your changes, click **Yes**.
5. Click the **Apply Filters** button.
Event Manager displays the events that meet the criteria selected in the Event Filters pane.

Filtering Events by Summary Text or Element Name

You can filter events by summary text or element name. For example, assume you are having problems with a host named WickedWitch. You could specify that only events from WickedWitch are displayed in Event Manager.

Note: If you accessed Event Manager by double-clicking an element in Capacity Explorer, Performance Explorer or System Explorer and then clicking the **Events** tab, you can filter only events by summary text.

To filter elements by summary text or element name:

1. In the **Summary Text or Element Name** field in Event Manager, type the summary text or element name you want to be used in the filter.
2. To filter events by element type, do one of the following:
 - ☐ Select a filter from the **Show Element Type** drop-down menu. See “Selecting an Element Type for Events” on page 553 for more information about the menu options.

- ☐ Click the **Customize** button next to the **Show Element Type** drop-down menu in Event Manager. Select the elements you want to be included in the filter and then click **OK**.
- 3. When you are asked if you want to apply your changes, click **Yes**.
- 4. Click the **Apply Filters** button.
Only events with the specified text are displayed in Event Manager.

Hiding Filters in Event Manager

Hide the filters for additional screen space. When you hide the filters, the following features are hidden:

- Show Severity
- Show Element Type
- Summary Text or Element Name

To hide the filters, click the **Hide Filters** link in the upper-left corner of Event Manager.

To display the filters, click the **Show Filters** link in the upper-left corner of Event Manager.

Enabling the Rank Column in Event Manager

By default, the Rank column in Event Manager is hidden. The Rank column displays the cost implication of an event, which is determined by the business cost of the application that uses the element and the severity of the event. The business cost of an application is set in System Explorer, as described in “Assigning a Business Cost to an Application” on page 338.

To enable the Rank column:

1. Click **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree.
 2. Click **Show Default Properties** at the bottom of the page.
 3. Copy the `calculateBusinessCost` property. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
-

4. Return to the Advanced page (**Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree).
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Change the value assigned to the `calculateBusinessCost` property to true in the **Custom Properties** field.
`calculateBusinessCost=true`
7. When you are done, click **Save**.
8. Restart the service for the management server for your changes to take effect.

Important: Depending on your license, Reporter may not be available. See the “List of Features” to determine if you have access to Reporter. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center**).

This chapter describes the following:

- About Reporter on page 560
 - Accessing Reporter on page 563
 - Viewing Reports on page 563
 - Viewing Report Collectors for an Element on page 565
 - Refreshing a Report on page 565
 - Changing the Formatting of a Report on page 566
 - Opening a Report in a New Window on page 566
 - Maximizing the Screen Space for a Report on page 567
 - Filtering Data in Global Reports on page 569
 - Sending a Report by E-mail on page 570
 - Managing E-mail Schedules for Reports on page 571
 - Creating Custom Reports on page 579
-

About Reporter

Reporter provides a variety of detailed reports, such as dependency, event, and utilization reports for discovered elements. To view a report, click a report name in the tree in Reporter. The report appears in the right pane.

Important: If you are shown the message “The report does not contain any data,” verify the collector for the report is running. You may also want to verify a collector is running if you believe the report is not displaying the latest information. Data collectors gather information for reports. If you stop a data collector, its reports are not updated. See “Managing Collectors for Reports” on page 236.

The management server provides reports that display performance information in a variety of formats:

- **HTML (Default)** - The software displays the report in a Web page by default.
- **PDF** - The software displays the report in Adobe Acrobat, a good option if you need to print the report. The software assumes you already have Adobe Acrobat Reader installed on your computer. To obtain a copy of Adobe Acrobat Reader, go to <http://www.adobe.com>.
- **Excel** - The software displays the report in Microsoft Excel, providing you have a copy of Microsoft Excel already installed.
- **XML** - The software displays the report in the XML format.

It also provides some reports with pie charts. The  icon next to the report name indicates the report displays a graphic.

The software provides reports for the following:

- **Global** - These global-wide reports provides data gathered from multiple management servers. For example, lets assume you have three management servers: one in London, one in Tokyo, and one in New York City. You can gather data from all three management servers. To view information for these reports, you must set up global reporting. To learn more, see “Setting Up Global Reporter” on page 248 for more information.
 - **Asset Management** - These reports provide information based on assets and ownership.
 - **Chargeback** - These reports provide cost information about the management and storage usage of an element. To populate these reports, enter information for Chargeback, as described in the topic, “Setting Up Chargeback” on page 716.
-

- **System** - These reports are enterprise wide and they collect information about the following:
 - **Application** - Data about applications the management server monitors, such as reports on application utilization and dependencies.
 - **Events** - Data about events occurring on the elements the management server monitors, such as summary reports on events.
 - **Fabric** - Data about fabrics, such as SAN components not zoned and world wide names that appear in zones but not in SANs.
 - **File Server** - Data about the file servers the management server monitors, such as reports on groups and users by server. This information is provided only if you have purchased the license for File Server SRM.
 - **HBA** - A summary report on the host bust adapters (HBAs) the management server detects.
 - **Host** - Data about the hosts in the management server monitors, such as reports on host storage allocation and total host utilization.
 - **NAS** - Data about NAS storage devices, such as reports on volume and aggregate usage.
 - **Performance** - Historical performance data for devices, such as reports on I/O performance.
 - **Storage System** - Data about storage systems the management server monitors, such as reports on storage system capacity and storage system utilization.
 - **Switch** - Data about switches the management server monitors, such as reports on switch port traffic and port utilization by connection type.
- **Applications** - These reports provide information about an application, such as Oracle or Microsoft Exchange.
- **Hosts** - These reports provide information about a host.
- **NAS** - These reports provide information about NAS storage devices.
- **Storage Systems** - These reports provide information about a storage system.
- **Switches** - These reports provide information about a switch.
- **Tape Libraries** - These reports provide information about a tape library.
- **Recent** - Lists the last 10 reports viewed. This option is not displayed when you first access Reporter.

Troubleshooting Reporter

- Reports in general contain no data or incomplete/incorrect data while the report cache is being refreshed. See “Refreshing a Report” on page 565.

- Some reports display data trends for several days in the future. For example, assume you run a report to gather information from the last three days. The report may display data trends for the next several days based on the current information. Please keep in mind the data trends are just assumptions and should not be treated as fact.
- The following IBM HBA appears as QLogic HBAs in the Navigation and Properties pages, in addition to reports:
 - IBM MSJ
 - FaStT FC-2/2-133
- Certain reports display elements assigned to the user's organization, including child organizations. For example, if you attempt to view an Assets by Department report and you do not have permission to access hosts through your organization, you are not given information about the hosts in the report. This is also true for e-mailing certain reports. Let's assume again you do not have permission to access hosts. If you e-mail an Assets by Department report, your e-mail will not contain information about hosts. If the users receiving your reports want to be able to view information about hosts, one of the following must happen:
 - The hosts in question must be added to your organization.
 - Someone else, who has the hosts in question already in their organization, must send the reports.

The same is true for organization filters. Let's assume you belong to two organizations: OnlyHosts and OnlySwitches. If you set your organization filter to display only the elements in OnlySwitches and not in OnlyHosts, your certain reports display only the elements in OnlySwitches. If you send an Assets by Department report by e-mail, the information in the report displays information about the elements that are currently allowed through the organization filter. You can set the organization filter when you create an e-mail schedule. See "Sending a Report by E-mail" on page 570.

The following reports display all information regardless of a user's organizations:

- Event Reports
- Application Reports
- Hosts Reports*
- Storage System Reports*
- Switch Reports*

*Reports that display information for only one element, such as Asset Summary, Details, Events, and Utilization reports.

Accessing Reporter

To access Reporter, click **Reporter** ()

Viewing Reports

Collectors gather information for reports. If you stop a collector, its reports are not updated. You might want to verify a collector is running if you are having difficulty viewing a report or you believe the report is not displaying the latest information. See “Viewing Report Collectors for an Element” on page 565 for more information.

Reports in general contain no data or incomplete/incorrect data while reports are being materialized (materialized views are being populated).

Important: The elements you see in the report are based on your organizations and which organizations are selected in the organization filter at the top of the page.

To view a report:

1. Access Reporter as described in “Accessing Reporter” on page 563.
2. In the middle pane, expand the tree and click the type of report you want.
The report is shown in the right pane. A list of the available filters displays at the top of pane. You can hide the list of available filters by clicking the **Hide Filters** button located in the upper-left corner of the pane.

Note: The list of available filters varies depending on the type of report you have selected. The options described below are available for all reports. For information about additional parameters that are available for some reports, see “Report Parameters” on page 564

3. To change the format of the report, select one of the following from the **Format** drop-down menu and then click the **Run Report** button:
 - ☐ **HTML (Default)** - The software displays the report in a Web page by default.
 - ☐ **PDF** - The software displays the report in Adobe Acrobat, a good option if you need to print the report. The software assumes you already have Adobe Acrobat Reader installed on your computer.
To obtain a copy of Adobe Acrobat Reader, go to <http://www.adobe.com>.
 - ☐ **Excel** - The software displays the report in Microsoft Excel, providing you have a copy of Microsoft Excel already installed.
 - ☐ **XML** - The software display the report in the XML format.
4. To view a report in a new window, click the Open in new window option located next to the drop-down menu. Then, click the **Run Report** button.
5. Some reports provide the option to view data within a time period. If your report offers this option, select the time range for the **Start Date** and **End Date** fields, by clicking the  icon. Then, click the **Run Report** button.
6. To send a report by e-mail, use the **Scheduled Deliveries** tab. See “Sending a Report by E-mail” on page 570 for more information.

Report Parameters

Many reports offer an additional set of parameters for filtering information. If additional parameters are available, they will display at the top of the pane along with the default formatting filter. The following are examples of the parameters available (depending on the type of report):

- Organization
- Vendor
- OS Type
- Array Type
- Start Time
- End Time
- Rollup Interval (hourly, daily, monthly, etc.)

To apply parameters to a report:

1. Select the values that you would like to use to filter the information displayed in the report.
 2. Click the **Run Report** button. The selected parameters are applied to the generated report.
-

Viewing Report Collectors for an Element

The management server uses collectors to gather information. The Collectors tab provides information about the collectors for a particular element.

To start collectors and view reports for an element:

1. Access the **Collectors** page by doing one of the following:
 - ☐ Clicking an element in Application Explorer, and then clicking the **Collectors** tab. (Click the the **Scan Schedule** tab for file servers).
 - ☐ Double-clicking an element in System Explorer, and then clicking the **Collectors** tab.
 - ☐ Clicking an element in Chargeback, and then clicking the **Collectors** tab.
2. To change a collector's start time, modify the time and date entered in the **Next Scheduled Run** field. If you decide to change the start time, make sure the date is in the yyyy-mm-dd format with the time resembling a 24-hour clock. There should be a space between the date and the time, as shown below:

2003-08-20 09:41

After the collector runs, the value in this column is updated to the next time the collector will run.

3. To change how often the collector runs, type the number of minutes in the **Interval** field.

Important: Do not make the interval too short. Running a collector too frequently uses up space on the management server and impacts its performance.

4. To enable the collector, click the **Start** button.
5. To stop a collector, click the **Stop** button.
6. To view a report, click its link. See "Viewing Reports" on page 563 for more information.

Refreshing a Report

If you want to view the latest information in a report, click the **Refresh Now** button on the Report Cache tab. To access the Report Cache tab, select **Configuration > Reports**. Then, click the

Report Cache tab. The management server gathers the latest information from the database and makes this information available to the reports.

The reports are refreshed every six hours by default. You can change the schedule for refreshing a report. See “Refreshing the Report Cache” on page 245.

If you are still seeing old information after you click the **Refresh Now** button, verify your database is being updated within the appropriate time frame for your organization. Some of the information contained in the database depend on collectors. Verify that the collectors are running at the appropriate intervals for your organization. Click the **Data Collection** tab to access the configuration pages for the many types of collectors.

If you find you are still viewing old information regarding elements on the network, you may need to perform Get Details. It is best to perform Get Details at regular intervals. See “Adding a Discovery Schedule” on page 209.

Changing the Formatting of a Report

Reports are displayed in HTML format by default. You can change the formatting of the report by selecting one of the following options from the **Format** drop-down menu in Reporter and then clicking the **Run Report** button:

- **HTML (Default)** - The software displays the report in a Web page by default.
- **PDF** - The software displays the report in Adobe Acrobat, a good option if you need to print the report. The software assumes you already have Adobe Acrobat Reader installed on your computer. To obtain a copy of Adobe Acrobat Reader, go to <http://www.adobe.com>.
- **Excel** - The software displays the report in Microsoft Excel, providing you have a copy of Microsoft Excel already installed.
- **XML** - The software display the report in the XML format.

Opening a Report in a New Window

Use this feature to view two or more reports simultaneously.

To view a report in a new window:

1. Select the **Open in new window** option.
2. Click the **Run Report** button.
A new window opens, and it displays just the report.
3. You can arrange the windows on the screen so you can view one report in one corner and another report in a different corner. You can also toggle between the two reports, depending on your operating system.

Maximizing the Screen Space for a Report

In some instances you might need to maximize the screen space to view a report. This can be done by hiding the middle and left pane, as described below:

1. To hide the middle pane, click the border between the tree listing the reports and the main pane and drag it to the left side of the page, as shown in the following figure.



Figure 13-1: Hiding the Middle Pane

2. To display the middle pane, drag the border to the right.
3. To close the left pane, click the section between the arrows at the upper-right border for the left pane.



Figure 13-2: Closing the Left Pane in Reporter

The buttons in the left pane are moved to the top of the page. See the following table for an explanation of the buttons.

Table 13-1: Buttons Displayed When the Left Pane Is Closed

Button	Provides Access to...
	The Home page
	Application Explorer
	System Explorer
	Capacity Explorer
	Performance Explorer
	Event Manager
	Provisioning
	Policy Manager
	Business Tools
	Chargeback
	Reporter

4. To open the left pane, click the section between the arrows in the upper-left border.

Filtering Data in Global Reports

Important: Depending on your license, Global Reporter may not be available. See the “List of Features” to determine if you have access to global reports. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center**). If your license lets you access global reports but you cannot access them, contact your system administrator to make sure your role lets you view and/or set up global reports.

You can filter the data in global reports so you only see the data gathered from certain sites and assigned to certain organizations. For example, assume your Global Reporter server gathers data from 10 sites throughout the world. You can filter the Global Reporter reports so you see data gathered from just one site.

If you do not see information in your reports, verify you have global reporting set up correctly. See “Setting Up Global Reporter” on page 248.

To filter data global reports:

1. Access Reporter as described in “Accessing Reporter” on page 563.
 2. Expand the **Global** node.
 3. Select a global report.
 4. In the right pane, click the **Filter Data** button.
 5. Expand the **Sites and Organizations** node.
 6. In the Filtering window, select the sites you want to view.
 7. Expand the node of each selected site to view its organizations.
 8. Select the organizations containing the elements you want to view. To view all organizations at a site, select the **Everything** option.
 9. Click **OK**.
The filter applies to all global reports.
-

Sending a Report by E-mail

You can send by e-mail an attached report in PDF, XML or Microsoft Excel format. If you want to send reports by e-mail on a regular basis, set up an e-mail schedule for the report, as described in “Adding an E-mail Schedule for a Report” on page 571.

Keep in mind the following:

- Before you can e-mail a report, you must set up e-mail notification, as described in the topic, “Setting Up E-mail Notification” on page 207.
- The elements in the report you send are based on your organizations and which organizations are selected in the organization filter at the top of the page.
- Send your reports soon after a report cache refresh. The reports display data that is in the report cache. If the report cache contains old data, the reports you send by e-mail will also show old data. The reports are refreshed every six hours by default. See “Scheduling a Report Cache Refresh” on page 246.

To send a report by e-mail:

1. Access Reporter as described in “Accessing Reporter” on page 563.
 2. Expand the tree in the middle pane, and click the report you want to send by e-mail.
 3. Click the **E-mail Report** button.
You are told the E-mail server is not enabled if you have not set up e-mail notification. You must set up e-mail notification before you can send e-mail a report. See “Setting Up E-mail Notification” on page 207.
 4. In the top field, type the recipient's e-mail address.
The software verifies the address entered has a correct form. To send multiple addresses, separate each address with a comma (,), for example:
`john.example@appiq.com, jerry.example@appiq.com`
 5. From the **Format** drop-down menu, select one of the following formats:
 - ☐ **PDF** - Requires the use of Adobe Acrobat, which can be downloaded for free from <http://www.adobe.com>.
 - ☐ **Excel** - Requires the use of Microsoft Excel.
 - ☐ **XML** - Requires the user has an understanding of XML.
 6. (Optional) Modify the subject and message.
 7. Click the **OK** button.
The report is sent.
-

Managing E-mail Schedules for Reports

This section describes the following:

- Adding an E-mail Schedule for a Report on page 571
- Editing an E-mail Schedule for a Report on page 574
- Deleting E-mail Schedules for a Report on page 577
- Viewing E-mail Schedules for a Report on page 578

Adding an E-mail Schedule for a Report

You can add an e-mail schedule so that a user receives an attached report on a regular basis. The report can be in the form a PDF, XML or Microsoft Excel document.

Keep in mind the following:

- Before you can add an e-mail schedule, you must set up e-mail notification, as described in “Setting Up E-mail Notification” on page 207.
- Schedule to send your reports soon after a report cache refresh. The reports display data that is in the report cache. If the report cache contains old data, the reports you send by e-mail will also show old data. The reports are refreshed every six hours by default. For example, assume you added an e-mail schedule that sends a report daily at 7 a.m. Also, assume you scheduled your report cache refreshes to take place daily at 8 a.m. Your reports will most likely show outdated data. It would make more sense to schedule your report cache refresh at 7 a.m. and then schedule to send your reports soon afterwards. See “Scheduling a Report Cache Refresh” on page 246.
- The management server service must be running for users to receive e-mail notification.
- Only the e-mail schedules created by the current user are listed. To view the e-mail schedules for all reports, click **Configuration > Reports**. Then, click the **Scheduled Deliveries** tab at the top of the screen.
- The elements in the report you send are based the organizations selected from the Organization Filters tab. For example assume you belong to two organizations: OnlyHosts and OnlySwitches. If you select OnlySwitches and not OnlyHosts in the Organization Filters tab (accessible by clicking **Add E-mail Schedule**), the user only receives information about the elements in the OnlySwitches organization. This is still true even if you change your organization filtering at the top of the page (.

To add an e-mail schedule:

1. Access Reporter as described in “Accessing Reporter” on page 563.
 2. Expand the tree in the middle pane, and click the report you want to send at a scheduled time.
 3. When the report is displayed in the right pane, click the **Scheduled Deliveries** tab in the right pane.
 4. Click the **Add E-mail Schedule** button.
 5. Verify the **Properties** tab is displayed.
 6. In the top field, type the recipient's e-mail address.
The software verifies the address entered has a correct form. To send multiple addresses, separate each address with a comma (,), for example:
`john.example@appiq.com,jerry.example@appiq.com`
 7. In the **Subject** field, type a subject for the e-mail messages you plan to send.
-

Note: Provide the name of the report in the subject field so users can distinguish the message from others.

8. In the **Message** field, type a message describing the report.
If you are e-mailing reports in bulk, you might want to let users know the e-mail is being sent by an automated process. You might also want to provide an e-mail address for users to provide feedback, for example:

`This e-mail and its attached report are generated automatically. If you would like to change how often the report is sent to you or you want to be taken off the list, please contact username@companyname.com.`
 9. From the **Format** drop-down menu, select one of the following formats:
 - ☐ **PDF** - Requires the use of Adobe Acrobat, which can be downloaded for free from <http://www.adobe.com>.
 - ☐ **Excel** - Requires the use of Microsoft Excel.
 - ☐ **XML** - Requires the user has an understanding of XML.
 10. In the **Time to Run** field, type the time you want to send the report. This time must be entered in the 24-hour format. For example, if you want a report sent at 2:15 p.m., you would type 14:15 in the **Time to Run** field.
 11. Select one of the following options to determine how frequently you want to send the report.
 - ☐ **Daily** - If you selected daily, select how frequently you want the management server to send the report.
 - **Everyday** - The report is sent everyday.
-

- **Weekday** - The report is sent only Monday through Friday.
 - **Every x days** - Fill in the frequency you want the report to be sent. For example, if you enter 15, the report is sent every 15 days.
 - **Weekly** - If you selected weekly, use the **Frequency** drop-down menu to select the day of the week on which you want the report sent.
 - **Monthly** - If you selected monthly, select the time during the month you want the report sent.
 - To send the report on the first or last day of the month, select the first option. Then, select **First** or **Last** from the drop-down menu.
 - To send the report on a specified day during the month, select the second option. Then, type the day on which you want the report sent. If you type a day that is not in the month, for example 30 for February, the report is sent on the last day of the month.
12. Click **Next**.
The Organization Filters tab is not available for all reports. If an Organization Filters tab is not available, you see an **OK** button instead of a **Next** button. Click the **OK** button and skip the rest of the steps in this procedure.
13. Select the organizations containing the elements you want used in the report. If you find all organizations selected, deselect the organizations containing the elements you do not want displayed in the report.
Organizations you belong to and their children are displayed. Only elements belonging to selected organizations are displayed in the report you send. For example, assume you selected OnlyHosts, as shown in the following figure. The user receiving the report would only see data about elements in OnlyHosts. Information about elements in other organizations would not be displayed.



Figure 13-3: Selecting Organizations Used in This Report

14. Click **Finish**.
The schedule is created.

Editing an E-mail Schedule for a Report

Important: Only the e-mail schedules created by the current user are listed. To view the e-mail schedules for all reports, click **Configuration > Reports**. Then, click the **Scheduled Deliveries** tab.

To edit an e-mail schedule for a report:

1. Access Reporter as described in “Accessing Reporter” on page 563.
2. Expand the tree in the middle pane, and click the report you want to send at a scheduled time.
3. When the report is displayed in the right pane, click the **Scheduled Deliveries** tab in the right pane.
4. Under the Edit column, click the **Edit** () button.
5. Verify the **Properties** tab is displayed.
6. In the top field, type the recipient's e-mail address.
The software verifies the address entered has a correct form. To send multiple addresses, separate each address with a comma (,), for example:

`john.example@appiq.com,jerry.example@appiq.com`

7. In the **Subject** field, change the subject of the e-mail.
8. In the **Message** field, change a message describing the report.
If you are e-mailing reports in bulk, you might want to let users know the e-mail is being sent by an automated process. You might also want to provide an e-mail address for users to provide feedback, for example:

This e-mail and its attached report are generated automatically. If you would like to change how often the report is sent to you or you want to be taken off the list, please contact `username@companyname.com`.

9. From the **Format** drop-down menu, select one of the following formats:
 - ☐ **PDF** - Requires the use of Adobe Acrobat, which can be downloaded for free from <http://www.adobe.com>.
 - ☐ **Excel** - Requires the use of Microsoft Excel.
 - ☐ **XML** - Requires the user has an understanding of XML.
10. In the **Time to Run** field, type the time you want to send the report. This time must be entered in the 24-hour format. For example, if you want a report sent at 2:15 p.m., you would type 14:15 in **Time to Run** field.
11. Select one of the following options to determine how frequently you want to send the report.
 - ☐ **Daily** - If you selected daily, select how frequently you want the management server to send the report.
 - **Everyday** - The report is sent everyday.
 - **Weekday** - The report is sent only Monday through Friday.
 - **Every x days** - Fill in the frequency you want the report to be sent. For example, if you enter 15, the report is sent every 15 days.

- ❑ **Weekly** - If you selected weekly, use the **Frequency** drop-down menu to select the day of the week on which you want the report sent.
 - ❑ **Monthly** - If you selected monthly, select the time during the month you want the report sent.
 - To send the report on the first or last day of the month, select the first option. Then, select **First** or **Last** from the drop-down menu.
 - To send the report on a specified day during the month, select the second option. Then, type the day on which you want the report sent. If you type a day that is not in the month, for example 30 for February, the report is sent on the last day of the month.
12. Click **Next**.
The Organization Filters tab is not available for all reports. If an Organization Filters tab is not available, you see an **OK** button instead of a **Next** button. Click the **OK** button and skip the rest of the steps in this procedure.
13. Select the organizations containing the elements you want used in the report. If you find all organizations selected, deselect the organizations containing the elements you do not want displayed in the report.
Organizations you belong to and their children are displayed. Only elements belonging to selected organizations are displayed in the report you send. For example, assume you selected OnlyHosts, as shown in the following figure. The user receiving the report would only see data about elements in OnlyHosts. Information about elements in other organizations would not be displayed.
-



Figure 13-4: Selecting Organizations Used in This Report

14. Click **Finish**.

Deleting E-mail Schedules for a Report

Important: Only the e-mail schedules created by the current user are listed. To view the e-mail schedules for all reports, click **Configuration > Reports**. Then, click the **Scheduled Deliveries** tab.

To delete an e-mail schedule:

1. Access Reporter as described in “Accessing Reporter” on page 563.

- 2. Expand the tree in the middle pane, and click the report corresponding to the e-mail schedule you want to delete.
- 3. When the report is displayed in the right pane, click the **Scheduled Deliveries** tab in the right pane.
- 4. Click the  button corresponding to the e-mail schedule you want to remove.

Viewing E-mail Schedules for a Report

Important: Only the e-mail schedules created by the current user are listed. To view the e-mail schedules for all reports, click **Configuration > Reports**. Then, click the **Scheduled Deliveries** tab.

To view the E-mail schedules assigned to a report.

- 1. Access Reporter as described in “Accessing Reporter” on page 563.
- 2. Expand the tree in the middle pane, and click the report corresponding to the e-mail schedules you want to view.
- 3. When the report is displayed in the right pane, click the **Scheduled Deliveries** tab in the right pane.
Information about the e-mail schedules for that report are displayed.

Table 13-2: Viewing E-mail Schedules for a Report

Column Name	Description
Recipient	The person who receives the report.
Subject	The subject of the e-mail, brief summary of what it is about.
Format	The format of the report sent: ■ PDF ■ Microsoft EXCEL ■ XML

Table 13-2: Viewing E-mail Schedules for a Report (Continued)

Column Name	Description
Edit	Click the  button to edit a schedule of the report. See “Adding an E-mail Schedule for a Report” on page 571 and “Editing an E-mail Schedule for a Report” on page 574 for information about the options displayed in this window.
Delete	Click the  button to remove the corresponding schedule.

Creating Custom Reports

This section describes the following:

- About Creating Custom Reports on page 579
- Configuring Report Designer to Work with the Management Server on page 581
- Designing Custom Reports on page 583
- Integrating Custom Reports on page 592
- Detailed Schema Information on page 595
- Views from Previous Releases on page 626
- Implementing Custom Reports on Sun Solaris on page 628

Important: You must install Report Designer before you can create custom reports. Obtain a copy of Report Designer from your sales professional. Follow the installation instructions that accompany it.

About Creating Custom Reports

To create customized reports, you need a program for creating reports, such as Report Designer. Use Report Designer to create customized reports on the management server. Report Designer links to the database of the management server, so you can view real-time data in your customized

Configuring Report Designer to Work with the Management Server

Important: The steps in this section are for configuring Report Designer version 6.0.

You can use data from the management server to create your custom reports. These steps assume you have installed Report Designer on a Microsoft Windows computer that can access the `C:\oracle\ora92\jdbc\lib` directory on the management server.

To use the management server with Report Designer:

1. Add the following class path to Report Designer. Refer to the documentation accompanying Report Designer for more information.
`C:\oracle\ora92\jdbc\lib\classes12.jar`
where `C:\oracle\ora92` is the directory containing `classes12.jar`.
2. Select **File > New Catalog**.
A catalog is a repository for reports. The catalog and the reports that are based on it must be in the same directory for the report to run. This is because the catalog contains the object definitions that are used by the reports in the catalog.
3. In the name field, type a name for the catalog, such as `custom.cat`.
4. Save the catalog in a directory especially for it. This way you can work on the reports remotely and then easily move them when you are ready to integrate them with the product.
5. Click **OK**.
When you are done with creating the catalog, the Catalog Browser window appears. If you do not see the Catalog Browser window, open the catalog by selecting **File > Open Catalog**. Select the catalog you want to open, and then click the **Open** button.
6. In the Catalog Browser window, expand the Default node in the tree. Then, right-click **Connection** and select **New Connection**.
7. Verify that the OracleOraHome92TNSListener service is running on the management server. If you want to view live data in your custom reports, the management server does not need to be running; however, the Oracle database for the management server does need to be running.

8. In the Get JDBC Connection Information window do the following:
 - a. Deselect the Use ODBC Data Source option.
 - b. Select the JDBC Driver option.
 - c. Type the following for the JDBC driver: `oracle.jdbc.driver.OracleDriver`
 - d. Type the following in the JDBC URL field:
`jdbc:oracle:thin:@HostIP/DNS:1521:APPIQ`
where HostIP/DNS is the host IP address or DNS name of the host running the management server

If Report Designer is running on the same computer as the management server, you can use localhost for the DNS name, as shown in the following example:

```
jdbc:oracle:thin:@localhost:1521:APPIQ
```

- e. Type the following in the User Name field: `Report_User`
This is the user name that is used to access the schema view in the management server database. This user has read privileges only for the schema views.
 - f. Type the following in the Password field: `appiq`
 9. Click **OK**.
Report Designer searches for JDBC driver.

You may have entered incorrect path information if Report Designer cannot find the JDBC drivers.
 10. Select `APPIQ_SYSTEM` under the schemas section.
The Tables pane becomes populated. These are all the tables you can use to create the reports. It is best to select as many tables as possible rather than too few. If you are not too sure as to which tables you may be using, you may want to select them all.
 11. Select the tables and then click **Add**. Click the **Done** button when you have finished adding tables.
The tables populate the Catalog Browser window.
 12. Click **File > Save Catalog** to save the catalog.
-

Designing Custom Reports

Important: This section assumes you have already installed and configured Report Designer and integrated it with the management server.

These steps assume you are running Report Designer 6.0, and you have the online help for Report Designer installed. The instructions provide general information about using Report Designer. For additional information, refer to the online help for Report Designer. The management server only supports Report Designer running on Microsoft Windows.

Note: You cannot create custom reports for NetApp filers.

Creating Standard Reports

The following steps are for Microsoft Windows and for Report Designer 6.0:

1. Open Report Designer.
 2. If the Choose Report screen does not appear, select **File > New** in Report Designer.
 3. Click the **Standard Report** icon and then click **Create**.
-




Figure 13-6: Choosing a Standard Report

4. Select the data you want in your report by selecting the corresponding materialized views (tables) displayed in the Data tab. Use the Report Categories table in “About Creating Custom Reports” on page 579 as a guideline. For example, in the following figure, the MV_APPLICATIONVW table has been selected. System application data will be made available to the report, according to the Report Categories table. You can, however, specify you do not want all data displayed in the report. When you are done, click **Next**. To find a definition of the listings in a table, see “Detailed Schema Information” on page 595.




Figure 13-7: Adding Tables for a Standard Report

5. When you are asked if Report Designer will create a new query, click **OK**.
6. If you selected more than one materialized views (tables) in the Data tab, you need to link common search criteria, such as the `DEVICE_ID` in one table to the `DEVICE_ID` of another table. Sometimes the search criteria will have different terminology. For example, in the following figure, `MV_EVENTVW_DOMAINNAME` is linked with `DOMAIN NAME`. If a search criteria is cut off in a table, you can expand the size of the table by clicking a table border and dragging it to the appropriate position.

To find a definition of a search criteria, see “Detailed Schema Information” on page 595.

Refer to the online help for Report Designer for more information. When you are done, click **Next**.




Figure 13-8: Linking Common Data in Tables for a Standard Report

7. Enter search criteria that will be used to generate the report. For example, let's assume you want the report to display information only about Oracle applications. You would enter a search criteria that tells Report Designer to display data from Oracle applications. When you are done, click **Next**.




Figure 13-9: Creating Search Criteria for Standard Reports

8. Select the data you want displayed in the report. Click the data source in the left and then click the arrow pointing right. Its report field appears in the Report fields pane. When you are done, click **Next**.
The order of the report fields in the Report fields pane determines their sequence in the report. For example, a report field at the top of the list in the Report fields pane will appear in the far left column in the report. Likewise, a report field at the bottom of the list in the Report fields pane will appear in the far right column in the report. The column heading for

the data is determined by the text in the AutoLabel column in the Report fields pane. To find a definition of the listings in a table, see “Detailed Schema Information” on page 595.




Figure 13-10: Deciding Which Data Should Appear in the Report

9. Select the fields in the left pane in the order you want them sorted in your report. Then, click the arrow pointing right. When you are done, click **Next**. For example, in the following figure, information in the report will first be sorted by an application name. Let's assume you then selected APPLICATIONID. Applications would first be sorted by application name and then by their application ID.




Figure 13-11: Sorting Information in the Report

10. Use the Style tab to determine the layout of the report. When you are done, click **Finish**.




Figure 13-12: Selecting the Layout of the Report

The report template is displayed. You will not see any data reported, only placeholders, as shown in the following figure.

APPLICATIONNAME	ELEMENT_ID	SUMMARY_TEXT	TIME_REPORTED	SEVERITY
XXXXXXXXXX	XXXXXXXXXX	##### DD	XXXXXXXXXX	##### DD
XXXXXXXXXX	##### DD	MM/dd/yy	##### DD	

Figure 13-13: Report Template Displayed

11. To view the report with its data, click the **View** tab.
12. The database for the management server must be running on the management server to be able to view active data in the report. Verify that the OracleOraHome92TNSListener service is running on the management server.
If you do not see any text, verify that the management server has collected this data. See which tables in your custom report map to pre-existing reports. Use the table in “About Creating Custom Reports” on page 579 as a guideline. Then, access Reporter on the management server and verify the corresponding reports are displaying information.

If you are running Report Designer remotely from the management server and you cannot view active data in your reports even with the Oracle database running, verify you did the following when you installed Report Designer.
 - ❑ Save classes12.jar from the management server on your computer and pointed this file in the class path when you installed Report Designer. If so, the file is listed in [Report Designer installation directory]\jreport.bat.
 - ❑ Verify the connection information for the catalog is correct. Open the catalog and expand the Connection node. Verify the IP address/DNS name listed is correct.
 If you are still having problems, verify the network from your computer to the management server is stable. Try accessing the management server console from your computer as a test.

13. If you want to view live data in your custom reports, the management server does not need to be running; however, the Oracle database for the management server does need to be running.
14. Click **File > Save Template** to save the report.
15. Refer to the online help for Report Designer for information on how to designer the reports.




Figure 13-14: Clicking the View Tab

Integrating Custom Reports

After you are satisfied with your custom report, you must integrate it with the management server so that other users can access the report. You must first deploy the custom report to the management server. Then, you must integrate the report so that it is accessible from Reporter. Custom reports appear with their own icon in the reports tree.

Important: You can have only one catalog (*.cat) file to a directory. If you have more than one catalog, create a subdirectory under `%JBOSS4_DIST%\server\appiq\reports\custom`. For example, assume you have two reports that use a different catalog, you would create two subdirectories under `reports\custom`. The support files and catalog for one report would go into one subdirectory (`reports\custom\subdirectory1`), and the support files and catalog for the other report would go into another subdirectory (`reports\custom\subdirectory2`).

To integrate custom reports:

1. Deploy the custom report files (CLS, CAT and all other files) from Report Designer. Use Report Designer's deploy catalog mechanism to transfer the files. Create a directory called `custom` under `%JBOSS4_DIST%\server\appiq\reports`. Then, deploy the catalog to the following directory:
`%JBOSS4_DIST%\server\appiq\reports\custom`
Refer to the documentation accompanying Report Designer for more information.
 2. Create a node for your custom reports. This node appears in the tree in Reporter.
 - a. Create the directory
`%JBOSS4_DIST%\server\appiq\reports\customTreeNodees` if it doesn't already exist
 - b. Create an XML file in that directory for your tree node, for example, `custom.xml`. The management server uses this file to determine where to put the node for custom reports in the tree.
 - c. Use a text editor, such as Notepad, to open the XML file you created in the previous step. Enter the following into the XML file:
`<NODE NAME="CUSTOM" LABEL="Custom" PARENT="SYSTEM"/>`
where
 - `NAME` is how reports and other tree nodes refer to this node.
 - `LABEL` is the label that appears in the user interface for the node.
 - `PARENT` is the name of the parent node in the tree (optional). In this case, the Custom report appears under the SYSTEM node in the tree in Reporter.
 3. To populate the tree in Reporter with your custom reports:
-

- a. Create the directory
%JBOSS4_DIST%\server\appiq\reports\definitions\custom if it doesn't already exist
- b. Create an XML file in that directory for your report, for example:
CustomHostReport.xml
- c. Open the file you created in the previous step in a text editor, such as Notepad. Add the following information to the file, modifying it for your settings. Refer to files in %JBOSS4_DIST%\server\appiq\reports\definitions if you need additional samples.
- d. Make sure you use an ID that is not used by any of the existing reports. You have to specify a title to appear in the tree and the file names for the CLS and CAT files, as shown in the following example.

```
<REPORT ID = "9098"
  TITLE = "Custom Host Connectivity"
  FILE_NAME = "//definitions/custom/Host_Connectivity.cls"
  CATALOG_FILE = "//definitions/custom/host.cat"
  SUPPORTS_ORGANIZATION_FILTERS = "true">
  <TREE_NODE NAME = "CUSTOM" />
  <ELEMENT_FUNCTION_CALLBACK NAME = "isHost" />
</REPORT>
```

where

- REPORT ID is the unique ID for the report.
- TITLE is the name you want to appear for the report in the tree in Reporter.
- FILE_NAME is the file name of the CLS file for the report and its path. In the example, the CLS file is in the
%JBOSS4_DIST%\server\appiq\reports\definitions\custom directory.
- CATALOG_FILE is the file name of the CAT file for the report and its path. In the example, the CAT file is in the
%JBOSS4_DIST%\server\appiq\reports\definitions\custom directory.
- SUPPORTS_ORGANIZATION_FILTERS if this property is set to true, the report supports organization filtering. If this property is set to false, the Organization Filters tab does not appear. If you leave omit SUPPORTS_ORGANIZATION_FILTERS the management server assumes you want the report to support organization filtering.
- TREE_NODE NAME is the name of the tree node you want the report to appear under. Use the node name of the file you created in

%JBOS4_DIST%\server\appiq\reports\customTreeNodees, not its file name.

- `ELEMENT_FUNCTION_CALLBACK NAME` refers to the type of information you want to obtain from the report. For example if you want to obtain information from all hosts, enter `ishosts`. See the following table for a general listing of the element function callback names you should use. If you do not see your element listed, refer to the existing report definitions for the element located in %JBOS4_DIST%\server\appiq\reports\definitions.

4. Restart the AppStorManager service.

The custom reports are displayed in the reports tree with a  icon.

Detailed Schema Information

This section provides information about the current materialized views. If you created reports in previous releases, you are most likely using old views. The old views are still supported; however, it is strongly suggested you use the new views in this section when creating reports. See “Views from Previous Releases” on page 626 for a mapping between the old and new views.

Table 13-3: Description of the Report Views

Materialized View (Tables)	Description
MVC_ORGANIZATIONVW	Provides information about organizations. See Table 13-19, “MVC_ORGANIZATIONVW,” on page 611.
MVC_ORGRELATIONVW	Provides information about the relationships in an organization, such as the parent and child identifiers in an organization. See Table 13-20, “MVC_ORGRELATIONVW,” on page 611.
MVC_OPTIONALTABLEVW	Provides information about the names and values of the optional value. See Table 13-25, “MVC_OPTIONALTABLEVW,” on page 613.
MVC_ASSETSUMMARYVW	Provides summary information about assets.

Table 13-3: Description of the Report Views (Continued)

Materialized View (Tables)	Description
MVC_HOSTSUMMARYVW	Provides summary information about hosts. See Table 13-4, "MVC_HOSTSUMMARYVW," on page 600.
MVC_APPLICATIONSUMMARYVW	Provides summary information about applications. See Table 13-33, "MVC_APPLICATIONSUMMARYVW," on page 619.
MVC_UNITACCESSVW	Provides information about unit access, such as access mode, host group, host group name, and host group modes. See Table 13-34, "MVC_UNITACCESSVW," on page 620.
MVCA_DBAPPCAPACITYVW	Provides capacity information for a supported database application. See Table 13-35, "MVCA_DBAPPCAPACITYVW," on page 621.
MVCA_EXCHANGEAPPCAPACITYVW	Provides capacity information for Microsoft Exchange. See Table 13-36, "MVCA_EXCHAPPCAPACITYVW," on page 621.
MVCA_VIRTUALAPPCAPACITYVW	Provides capacity information for a virtual application. See Table 13-37, "MVCA_VIRTUALAPPCAPACITYVW," on page 622.
MVC_EVENTSVW	Provides event information, such as description and time reported. See Table 13-18, "MVC_EVENTSVW," on page 610.
MVCA_FSRM_VOLUMESUMMARYVW	Provides summary information about File Server SRM volumes. Table 13-38, "MVCA_FSRM_VOLUMESUMMARYVW," on page 622.
MVCA_FSRM_AGESUMMARYVW	Provides summary information about ages in File Server SRM. See Table 13-39, "MVCA_FSRM_AGESUMMARYVW," on page 623.

Table 13-3: Description of the Report Views (Continued)

Materialized View (Tables)	Description
MVCA_FSRM_EXTDETAILSUMMARYVW	Provides summary information about extent details. See Table 13-40, "MVCA_FSRM_EXTDETAILSUMMARYVW," on page 623.
MVCA_FSRM_DIRDETAILSUMMARYVW	Provides information about directories in File Server SRM. See Table 13-41, "MVCA_FSRM_DIRDETAILSUMMARYVW," on page 623.
MVCA_FSRM_USERSUMMARYVW	Provides information about users in File Server SRM. See Table 13-42, "MVCA_FSRM_USERSUMMARYVW," on page 624.
MVCA_FSRM_GROUPSUMMARYVW	Provides information about groups in File Server SRM. See Table 13-43, "MVCA_FSRM_GROUPSUMMARYVW," on page 624.
MVCA_FSRM_TOPNFILES	
MVCA_FSRM_AGEDFILEDETAILS	Provides information about the age properties of files in File Server SRM. See Table 13-45, "MVCA_FSRM_AGEDFILEDETAILS," on page 625.
MVCA_FSRM_LARGEDIRINFO	
MVC_STORAGEPOOLSUMMARYVW	Provides summary information about storage pools. See Table 13-9, "MVC_STORGAEPOOLSUMMARYVW," on page 603.
MVC_STORAGESYSTEMSUMMARYVW	Provides summary information about a storage system. See Table 13-8, "MVC_STORAGESYSTEMSUMMARYVW," on page 602.

Table 13-3: Description of the Report Views (Continued)

Materialized View (Tables)	Description
MVC_STORAGEVOLUMESUMMARYVW	Provides summary information about a storage volume. See Table 13-10, "MVC_STORAGEVOLUMESUMMARYVW," on page 604.
MVC_HOSTDISKDRIVEVW	Provides information about host disk drives. See Table 13-7, "MVC_HOSTDISKDRIVEVW," on page 602.
MVC_HOSTVOLUMESUMMARYVW	Provides summary information about host volumes. See Table 13-6, "MVC_HOSTVOLUMESUMMARYVW (logical volumes)," on page 601.
MVC_DISKEXTENTSUMMARYVW	Provides summary information about disk extents. See Table 13-26, "MVC_DISKEXTENTSUMMARYVW," on page 613.
MVC_STORAGEVOLUMEPORTS	Provides information about storage volume ports. See Table 13-27, "MVC_STORAGEVOLUMEPORTS," on page 614.
MVC_VOLUMEDISKDRIVEVW	Provides information about volume disk drives. See Table 13-28, "MVC_VOLUMEDISKDRIVEVW," on page 615.
MVC_STORAGEPROCESSORSUMMARYVW	Provides information about storage processors. Table 13-29, "MVC_STORAGEPROCESSORSUMMARYVW," on page 615.
MVC_DISKDRIVESUMMARYVW	Provides summary information about disk drives. See Table 13-30, "MVC_DISKDRIVESUMMARYVW," on page 616.
MVC_DISK_EXTENTVW	Provides information about disk extents. See Table 13-31, "MVC_DISK_EXTENTVW," on page 617.

Table 13-3: Description of the Report Views (Continued)

Materialized View (Tables)	Description
MVC_SWITCHSUMMARYVW	Provides summary information about switches. See Table 13-11, “MVC_SWITCHSUMMARYVW,” on page 605.
MVC_PORTSUMMARYVW	Provides summary information about ports. See Table 13-12, “MVC_PORTSUMMARYVW,” on page 606.
MVC_ZONESUMMARYVW	Provides summary information about zones. See Table 13-13, “MVC_ZONESUMMARYVW,” on page 607.
MVC_ZONEVW	Provides information about zones. See Table 13-14, “MVC_ZONEVW,” on page 608.
MVC_PATHVW	Provides path information. See Table 13-15, “MVC_PATHVW,” on page 608.
MVC_SUBPATHVW	Provides subpath information. See Table 13-16, “MVC_SUBPATHVW,” on page 609.
MVC_MULTIPATHVW	Provides multipath information. See Table 13-17, “MVC_MULTIPATHVW,” on page 610.
MVC_SWITCHCONFIGVW	Provides switch configuration information. See Table 13-24, “MVC_SWITCHCONFIGVW,” on page 613.
MVC_HOSTCAPACITYVW	Provides host capacity information. See Table 13-21, “MVC_HOSTCAPACITYVW,” on page 611.
MVC_STORAGESYSTEMCONFIGVW	Provides storage system configuration information. See Table 13-22, “MVC_STORAGESYSTEMCONFIGVW,” on page 612.
MVC_STORAGEPOOLCONFIGVW	Provides information about storage pool configurations. See Table 13-23, “MVC_STORAGEPOOLCONFIGVW,” on page 612.

Table 13-3: Description of the Report Views (Continued)

Materialized View (Tables)	Description
MVIEWCORE_STATUS	Provides information about the core views. The core views are the views starting with mvc, mvca, and mvcs. See Table 13-47, “MVIEWCORE_STATUS,” on page 626.
MVIEW_STATUS	Show the status of the materialized views for reports. Table 13-48, “MVIEW_STATUS,” on page 626.

The following tables provide information about each report view:

Table 13-4: MVC_HOSTSUMMARYVW

Column Name	Type	Description
HOSTID	NUMBER(38)	HostID
HOSTNAME	VARCHAR2(256)	Host Name
DOMAINID	NUMBER(38)	DomainID
VENDOR	VARCHAR2(256)	Host Vendor
DESCRIPTION	VARCHAR2(1024)	Host Description
STATUS	NUMBER(38)	Operation status (provide map here)
IP	VARCHAR2(16)	Host IP
DNS	VARCHAR2(50)	Host DNS Name
Model	VARCHAR2(256)	Host Model
Version	VARCHAR2(256)	Host Version number
OS	VARCHAR2(24)	Host Operating System
TOTALPHYSICALMEM	NUMBER(38)	Total physical memory
NUMBERPROCESSOR	Number	Number of processors
SUPPORTFLAG	NUMBER(38)	Support flag (unused now)
BASETABLENAME	CHAR(4)	Name of the base table for optional values

Table 13-5: MVC_CARDSUMMARYVW

Column Name	Type	Description
CardID	NUMBER(38)	CardID
CardName	VARCHAR2(256)	Card Name
ContainerID	NUMBER(38)	Container ID
CardType	VARCHAR2(7)	Card Type (HBA, SCSI)
DomainID	NUMBER(38)	Domain ID
Vendor	VARCHAR2(256)	Card Vendor
Description	VARCHAR2(1024)	Card Description
Status	NUMBER(38)	Operational status
WWN	VARCHAR2(256)	Node WWN
Model	VARCHAR2(256)	Card model
SerialNumber	VARCHAR2(256)	Card Serial Number
Version	VARCHAR2(256)	Card Version
Firmware	VARCHAR2(256)	Firmware version
DriverVersion	VARCHAR2(256)	Driver version
BASETABLENAME	CHAR(4)	Name of the base table for optional values

Table 13-6: MVC_HOSTVOLUMESUMMARYVW (logical volumes)

Column Name	Type	Description
LogicalVolumeID	NUMBER(38)	Storage Volume ID
LogicalVolumeName	VARCHAR2(256)	Name of the logical volume
DomainID	NUMBER(38)	Domain ID
Description	VARCHAR2(1204)	Description
HostID	NUMBER(38)	Container Host ID
DeviceID	VARCHAR2(254)	Logical Device ID
FileSystemType	VARCHAR2(254)	File System Type

Table 13-6: MVC_HOSTVOLUMESUMMARYVW (logical volumes) (Continued)

Column Name	Type	Description
Blocksize	NUMBER(38)	These 3 fields may not be needed for Host
NumberOfBlocks	NUMBER(38)	Logical Volumes
ConsumableBlocks	NUMBER(38)	
BASETABLENAME	CHAR(14)	Name of the base table for optional values

Table 13-7: MVC_HOSTDISKDRIVEVW

Column Name	Type	Description
HostID	NUMBER(38)	Host ID
DiskDriveID	NUMBER(38)	Disk Drive ID
ExtentID	NUMBER(38)	Disk Partition ID
DiskDrive	VARCHAR2(256)	Disk Drive Name
DiskPartition	VARCHAR2(256)	Disk Partition Name
DiskPartitionDescription	VARCHAR2(1024)	Description of the partition
DiskPartitionSPace	Number	Capacity of the disk partition in megabytes

Table 13-8: MVC_STORAGESYSTEMSUMMARYVW

Column Name	Type	Description
StorageSystemID	NUMBER(38)	Storage system ID
StorageSystemName	VARCHAR2(256)	Storage system Name
DomainID	NUMBER(38)	Domain ID
Vendor	VARCHAR2(254)	Vendor
Description	VARCHAR2(1024)	Description of the Storage System
Status	NUMBER(38)	Operational Status (provide map here)
IP	VARCHAR2(16)	Not used

Table 13-8: MVC_STORAGESYSTEMSUMMARYVW (Continued)

Column Name	Type	Description
Model	VARCHAR2(254)	Model
SerialNumber	VARCHAR2(254)	Serial Number
Version	VARCHAR2(254)	Version
StorageSystemStatus	VARCHAR2(254)	Intrinsic status of the system
ResetCapability	VARCHAR2(254)	Indicating reset capability
ProvisionCapabilities	NUMBER(38)	Provide map here
SupportFlag	NUMBER(38)	Provide map here
BASETABLENAME	Varchar	Name of the base table for optional values

Table 13-9: MVC_STORGAEPOOLSUMMARYVW

Column Name	Type	Description
StoragePoolID	NUMBER(38)	Storage Pool ID
StoragePoolName	VARCHAR2(256)	Pool Name
StoragePoolDescription	VARCHAR2(1024)	Description of the storage pool
Status	NUMBER(38)	Operational status (provide map)
StorageSystemID	NUMBER(38)	ID of storage system to which pool belongs
PoolSettingID	NUMBER(38)	Storage capabilities ID
ParentPoolID	NUMBER(38)	Parent pool ID, reference to Storage Pool ID
TotalAvailableSpace	NUMBER(38)	Total available space in bytes
CIMPoolID	VARCHAR2(254)	Reserved
PoolType	NUMBER(38)	Pool Type (provide map here)
StorageCapabilityInternalName	VARCHAR2(254)	Internal Name of the capability
NoSinglePTOfFailure	NUMBER(1)	No single point of failure indication
DefaultNoSinglePtOfFailure	NUMBER(1)	Default no single point of failure indication

Table 13-9: MVC_STORGAEPOOLSUMMARYVW (Continued)

Column Name	Type	Description
MinDataRedundancy	NUMBER(18)	Minimum data redundancy indication
MaxDataRedundancy	NUMBER(18)	Maximum data redundancy indication
DefaultDataRedundancy	NUMBER(18)	Default data redundancy indication
MinSpindleRedundancy	NUMBER(18)	
MaxSpindleRedundancy	NUMBER(18)	
DefaultSpindleRedundancy	NUMBER(18)	
MinDeltaReservation	NUMBER(18)	
MaxDeltaReservation	NUMBER(18)	
DefaultDeltaReservation	NUMBER(18)	
StorageCapabilityCommonName	VARCHAR2(256)	Name of the pool capability
StorageCapabilityDescription	VARCHAR2(1024)	Description of the pool capability

Table 13-10: MVC_STORAGEVOLUMESUMMARYVW

Column Name	Type	Description
StorageVolumeID	NUMBER(38)	StorageVolume ID
StorageVolumeName	VARCHAR2(256)	StorageVolume Name
DomainID	NUMBER(38)	Domain ID
OID	VARCHAR2(254)	Reserved
Status	NUMBER(38)	Operational status (provide map here)
StorageSystemID	NUMBER(38)	ID of the storage system that contains this volume
StorageCapabilityID	NUMBER(38)	Storage Capability ID
VolumeDeviceID	VARCHAR2(254)	Device ID
AccessType	VARCHAR2(254)	Volume access type
Blocksize	NUMBER(38)	Size per block in bytes
NumberOfBlocks	NUMBER(38)	Total number of blocks in the volume
ConsumableBlocks	NUMBER(38)	Total number of consumable blocks
SeqAccess	NUMBER(1)	Sequential access

Table 13-10: MVC_STORAGEVOLUMESUMMARYVW (Continued)

Column Name	Type	Description
Availability	VARCHAR2(254)	Availability indication
StatusInfo	VARCHAR2(254)	Status of the volume
PoolID	NUMBER(38)	ID of the Storage Pool that contains this volume
VolumeType	NUMBER(38)	Type of volume
BASETABLENAME	CHAR(14)	Name of the base table for optional values

Table 13-11: MVC_SWITCHSUMMARYVW

Column Name	Type	Description
SwitchID	NUMBER(38)	Switch ID
SwitchName	VARCHAR2(256)	Switch Name
DomainID	NUMBER(38)	Domain ID
Vendor	VARCHAR2(254)	Switch Vendor
Description	VARCHAR2(1024)	Description of the Switch
Status	NUMBER(38)	Operational status (provide map here)
IP	VARCHAR2(16)	Switch IP
DNS	VARCHAR2(50)	DNS of the Switch
WWN	VARCHAR2(254)	WWN of the Switch
Model	VARCHAR2(254)	Switch Model
SerialNumber	VARCHAR2(254)	Serial Number of the Switch
Version	VARCHAR2(254)	Switch's hardware version
LoginName	VARCHAR2(254)	Login name for this Switch
LoginPwd	VARCHAR2(254)	Login password for this Switch
HardZoningCapability	VARCHAR2(254)	
SoftZoningCapability	VARCHAR2(254)	
ZoningInstalled	NUMBER(1)	
MaxModuleNumber	NUMBER(38)	
CurrentZoningEnforcement	VARCHAR2(254)	

Table 13-11: MVC_SWITCHSUMMARYVW (Continued)

Column Name	Type	Description
SwitchDomainID	NUMBER(38)	
SwitchStatus	VARCHAR2(254)	
SwitchState	VARCHAR2(254)	
IPGateway	VARCHAR2(254)	
IPNetwork	VARCHAR2(16)	
FCAddress	VARCHAR2(254)	
FCNetmask	VARCHAR2(16)	
SwitchRole	VARCHAR2(254)	
ProvisionSupportFlag	NUMBER(1)	
FabricWWN	Varchar	
FabricID	NUMBER(38)	Fabric ID
BASETABLENAME	CHAR(6)	Name of the base table for optional values

Table 13-12: MVC_PORTSUMMARYVW

Column Name	Type	Description
PortID	NUMBER(38)	Port ID
PortName	VARCHAR2(256)	Port Name
DomainID	NUMBER(38)	Domain ID
Description	VARCHAR2(1024)	Description
Status	NUMBER(38)	Operational status (provide map)
WWN	VARCHAR2(32)	Port WWN
ContainerID	NUMBER(38)	Container ID of this port
TargetPort	NUMBER(38)	Target Port ID (Never populated)
Connected_To_WWN	VARCHAR2(32)	WWN of connected port
Device_ID	VARCHAR2(254)	Port Device ID
PortState	VARCHAR2(254)	
PortStatus	VARCHAR2(254)	

Table 13-12: MVC_PORTSUMMARYVW (Continued)

Column Name	Type	Description
Physical_State	VARCHAR2(254)	
Port_ID	NUMBER(38)	
Port_Speed	NUMBER(38)	
Max_Speed	NUMBER(38)	Port Max Speed, bit/second
PortNumber	NUMBER(18)	
SCSIPort	NUMBER(18)	
ConnectedToNodeWWN	VARCHAR2(32)	
PortType	VARCHAR2(254)	
BASETABLENAME	CHAR(4)	Name of the base table for optional values

Table 13-13: MVC_ZONESUMMARYVW

Column Name	Type	Description
ZoneID	NUMBER(38)	Zone ID
ZoneName	VARCHAR2(254)	Zone Name
DominalID	NUMBER(38)	Domain ID (currently only one domain)
CimClassName	VARCHAR2(28)	
Status	NUMBER(38)	ApplQ status
ActiveZone	VARCHAR2(3)	
ZoneType	VARCHAR2(254)	
ProtocolType	VARCHAR2(254)	
ReadOnly	NUMBER(1)	
FabricID	NUMBER(38)	Fabric ID
FabricWWN	VARCHAR2(254)	Fabric WWN
FabricCName	VARCHAR2(256)	
ZoneCapID	NUMBER(38)	Zone capability ID
ZoneCapabilitiesName	VARCHAR2(254)	Name of the zone capabilities
ZC_MaxName_length	NUMBER(18)	Name length limit

Table 13-13: MVC_ZONESUMMARYVW (Continued)

Column Name	Type	Description
MaxZoneSets	NUMBER(18)	Number of maximum zone sets
MaxZones	NUMBER(18)	
MaxZoneMembers	NUMBER(18)	
MaxZonePerZoneSet	NUMBER(18)	
MaxZoneAliases	NUMBER(18)	
EnhencdZoning	NUMBER(1)	

Table 13-14: MVC_ZONEVW

Column Name	Type	Description
ZoneSetID	NUMBER(38)	Zone set ID
ZoneSetName	VARCHAR2(256)	Zone set Name
ZoneID	NUMBER(38)	Zone ID
FabricID	NUMBER(38)	ID of the fabric which the zone belongs
ZoneMemberID	NUMBER(38)	Zone member ID
ZoneMemberName	VARCHAR2(254)	Name of the zone member
ZoneMemberType	VARCHAR2(254)	Type of the zone member
ZoneMemberInFabric	NUMBER(1)	
ZonePortWWN	VARCHAR2(32)	WWN of the zone port
ZoneAlias	VARCHAR2(256)	

Table 13-15: MVC_PATHVW

Column Name	Type	Description
PathID	NUMBER(38)	Path ID
HostID	NUMBER(38)	Host ID on this path
LogicalVolumeID	NUMBER(38)	Logical Volume ID
DiskPartitionID	NUMBER(38)	Disk Partition ID if raw partition is present

Table 13-15: MVC_PATHVW (Continued)

Column Name	Type	Description
IsMounted	VARCHAR2(5)	(TRUE, FALSE)
ContainerPathID	NUMBER(38)	Currently not in use
ApplicationID	NUMBER(38)	Application ID for the application file
ApplicationFileID	NUMBER(38)	Application File ID
FilePath	VARCHAR2(256)	Application file path
FileName	VARCHAR2(128)	Application file name

Table 13-16: MVC_SUBPATHVW

Column Name	Type	Description
SubPathID	NUMBER(38)	SubPath ID
PathID	NUMBER(38)	Parent Path ID
HostID	NUMBER(38)	Host ID
DiskDriveID	NUMBER(38)	Disk Drive ID
HBACardID	NUMBER(38)	HBA Card ID
HBAPortID	NUMBER(38)	HBA Port ID
HostSwitchPortID	NUMBER(38)	ID of Host Switch Port
SystemSwitchPortID	NUMBER(38)	ID of the system switch port
StorageSystemPortID	NUMBER(38)	
StorageVolumeID	NUMBER(38)	
StorageSystemID	NUMBER(38)	
LUN	NUMBER(38)	
IsLocal	VARCHAR2(6)	
FabricID	NUMBER(38)	
MultipathDeviceID	NUMBER(38)	
PathSwitchID	NUMBER(38)	

Table 13-17: MVC_MULTIPATHVW

Column Name	Type	Description
PathID	NUMBER(38)	Path ID
MultiPathDiskExtentID	NUMBER(38)	ID of multipath disk extent
VolumeManagerVolumeID	NUMBER(38)	Volume Manager Volume ID
VxvmDiskPartitionID	NUMBER(38)	ID of volume manager volume disk partition
multipathDeviceID	NUMBER(38)	

Table 13-18: MVC_EVENTSVW

Column Name	Type	Description
DomainID	NUMBER(38)	
ID	NUMBER(38)	
Element_ID	NUMBER(38)	
Device_ID	NUMBER(38)	
Summary_Text	VARCHAR2(4000)	
Time_Reported	Date	
Severity	NUMBER(38)	
Cleared	NUMBER(1)	
Source	VARCHAR2(254)	
Type	NUMBER(38)	
SubType	NUMBER(38)	
Probable_Cause_Description	VARCHAR2(4000)	
Recommanded_Actions	VARCHAR2(254)	
Correlated_Indications	VARCHAR2(254)	
Type_Description	VARCHAR2(254)	
Element_type	NUMBER(38)	
Name	VARCHAR2(256)	

Table 13-18: MVC_EVENTSVM (Continued)

Column Name	Type	Description
Vendor	VARCHAR2(254)	
Model	VARCHAR2(254)	

Table 13-19: MVC_ORGANIZATIONVM

Column Name	Type	Description
ORGID	NUMBER(38)	ID from Organization table
ORGNAME	VARCHAR2(256)	Organization name
ORGDESCRIPTION	VARCHAR2(4000)	Description of this Organization
ELEMENTID	NUMBER(38)	Element ID mapped to this Organization
DOMAINID	NUMBER(38)	Domain ID

Table 13-20: MVC_ORGRELATIONVM

Column Name	Type	Description
ParentOrgID	NUMBER(38)	Parent Organization ID
ChildOrgID	NUMBER(38)	Child Organization ID
DOMAINID	NUMBER(38)	Domain ID

Table 13-21: MVC_HOSTCAPACITYVM

Column Name	Type	Description
HostID	NUMBER(38)	Host ID
VolumeID	NUMBER(38)	Logical Volume ID
TimeStamp	TIMESTAMP	Time of data collection
Total	NUMBER(38)	Total capacity in megabytes

Table 13-21: MVC_HOSTCAPACITYVW (Continued)

Column Name	Type	Description
Used	NUMBER(38)	Used capacity in megabytes
Free	NUMBER(38)	Free capacity in megabytes

Table 13-22: MVC_STORAGESYSTEMCONFIGVW

Column Name	Type	Description
StorageSystemID	NUMBER(38)	Storage System ID
CollectionTime	TIMESTAMP(6)	Configuration statistics collection time
MaskedMB	NUMBER	Masked storage in megabytes
ExportedMB	NUMBER	Storage exposed in megabytes
UnexportedMB	NUMBER	Volume unexposed in megabytes
AvailableMB	NUMBER	Available storage left on storage pool
ProvisionedMB	NUMBER	Provisioned storage in megabytes
RawStorageMB	NUMBER	Unconfigured storage in megabytes
TotalMB	NUMBER	Total storage in megabytes
MainframeStorageMB	NUMBER	Mainframe storage in megabytes
AvailablePorts	NUMBER(38)	Number of available ports
ConnectedPorts	NUMBER(38)	Number of used ports
TotalPorts	NUMBER(38)	Total system ports

Table 13-23: MVC_STORAGEPOOLCONFIGVW

Column Name	Type	Description
StoragePoolID	NUMBER(38)	Storage Pool ID
CollectionTime	TIMESTAMP (6)	Configuration stats collection time
MaskedMB	NUMBER	Masked Storage in megabytes
ExportedMB	NUMBER	Storage exposed in megabytes
UnexportedMB	NUMBER	Volume unexposed in megabytes

Table 13-23: MVC_STORAGEPOOLCONFIGVW (Continued)

Column Name	Type	Description
AvailableMB	NUMBER	Available storage left on storage pool
ProvisionedMB	NUMBER	Provisioned storage in megabytes
TotalMB	NUMBER	Total storage in megabytes
StorageCapability	VARCHAR2(255)	Storage pool capability
MainframePool	VARCHAR2(13)	Indicate whether reserved for mainframe: MAINFRAMEPOOL

Table 13-24: MVC_SWITCHCONFIGVW

Column Name	Type	Description
SwitchID	NUMBER(38)	Switch ID
AvailablePorts	NUMBER(38)	Number of available ports
ConnectedPorts	NUMBER(38)	Number of used ports
TotalPorts	NUMBER(38)	Number of total ports of this switch

Table 13-25: MVC_OPTIONALTABLEVW

Column Name	Type	Description
BaseTableName	VARCHAR2(30)	Name of the base table
BaseTableID	NUMBER(38)	ID of the base table record
OptionalName	VARCHAR2(50)	Name of the optional value
OptionalValue	VARCHAR2(4000)	Value of the optional value

Table 13-26: MVC_DISKEXTENTSUMMARYVW

Column Name	Type	Description
DiskExtentID	NUMBER(38)	Disk Extent ID
DiskEnxtentName	VARCHAR2(256)	Name of the extent

Table 13-26: MVC_DISKEXTENTSUMMARYVW (Continued)

Column Name	Type	Description
DomainID	NUMBER(38)	Domain ID
DiskExtentDescription	VARCHAR2(1024)	Description of the extent
Status	NUMBER(38)	Operational status
Access_Type	VARCHAR2(254)	Access type
BlockSize	NUMBER(38)	Block Size
Number_Of_Blocks	NUMBER(38)	Number of total blocks
Consumable_Blocks	NUMBER(38)	Number of usable blocks
StoragePoolID	NUMBER(38)	Storage Pool ID
SystemID	NUMBER(38)	Container System ID
BASETABLENAME	CHAR(11)	Name of the base table

Table 13-27: MVC_STORAGEVOLUMEPORTS

Column Name	Type	Description
ID	NUMBER(38)	
Storage_Volume_ID	NUMBER(38)	Storage Volume ID
Port_ID	NUMBER(38)	
LUN_ID	NUMBER(38)	
Access_Mode	Varchar(254)	
Access_State	Varchar(254)	
Host_Group	Varchar(254)	
Host_Group_Name	Varchar(254)	
Host_Group_Modes	Varchar(1024)	

Table 13-28: MVC_VOLUMEDISKDRIVEVW

Column Name	Type	Description
VolumeID	NUMBER(38)	Storage Volume ID
DiskDriveID	NUMBER(38)	Disk Drive ID
ExtentID	NUMBER(38)	Disk Extent ID

Table 13-29: MVC_STORAGEPROCESSORSUMMARYVW

Column Name	Type	Description
SystemProcessorID	NUMBER(38)	Storage System Processor ID
SystemProcessorName	VARCHAR2(256)	Name of the system processor
DomainID	NUMBER(38)	Domain ID
Vendor	VARCHAR2(254)	Vendor
Description	VARCHAR2(1024)	
Status	NUMBER(38)	Operational status
IP	VARCHAR2(16)	
DNS	VARCHAR2	
WWN	VARCHAR(16)	
Model	VARCHAR2(254)	
PowerManagement	VARCHAR2(254)	
SerialNumber	VARCHAR2(254)	
Version	VARCHAR2(254)	
ContainerID	NUMBER(38)	Container system ID
ProcessorStatus	NUMBER	Status of the processor
ResetCapability	VARCHAR2(254)	Reset Capability
Roles	VARCHAR2(254)	Roles
ProviderTag	VARCHAR2(254)	Provider name tag
SupportFlags	NUMBER(38)	Support flags
BASETABLENAME	CHAR(14)	Name of the base table

Table 13-30: MVC_DISKDRIVESUMMARYVW

Column Name	Type	Description
DiskDriveID	NUMBER(38)	Disk Drive ID
DiskDriveName	VARCHAR2(256)	Name of the Disk Drive
DomainID	NUMBER(38)	Domain ID
OID	VARCHAR2(254)	Object ID of the Disk Drive
Vendor	VARCHAR2(254)	Vendor
Description	VARCHAR2(1024)	Description
Status	NUMBER(38)	Operational status
Model	VARCHAR2(254)	Disk Drive Model
Name	VARCHAR2(254)	Name coming from disk drive
CardID	NUMBER(38)	Card ID
DiskDriveStatus	VARCHAR2(254)	Disk Drive Status
SCSIBus	NUMBER(38)	SCSI Bus
SCSILUN	NUMBER(38)	SCSI LUN
SCSITargetID	NUMBER(38)	SCSI target ID
SCSIPort	NUMBER(38)	SCSI PORT ID
SystemID	NUMBER(38)	Container System ID
MaxMediaSize	NUMBER(38)	Maximum media size
MaxBlockSize	NUMBER(38)	Maximum block size
MinBlockSize	NUMBER(38)	Minimum block size
EnableStatus	VARCHAR2(254)	
Availability	VARCHAR2(254)	
BASETABLENAME	CHAR(11)	Name of the base table

Table 13-31: MVC_DISK_EXTENTVW

Column Name	Type	Description
ExtentID	NUMBER	Disk Extent ID
ContainerExtentID	NUMBER	Container Extent ID
DiskID	NUMBER	Disk Drive ID

Table 13-32: MVC_ASSETSUMMARY

Column Name	Type	Description
DOMAINID	NUMBER(38)	Domain ID
ASSETID	NUMBER(38)	ID of the Asset depending on Assetclass for example if the assetclass is Host then this is host.id
ASSETCLASS	VARCHAR2(13)	The Asset class are "HOST", "APPLICATION", "STORAGESYSTEM" or "SWITCH"
NAME	VARCHAR2(256)	Name of the Host for example in case of Host this host name.
DATECREATED	DATE	Creation Date
DATELASTMODIFIED	DATE	Date last modified.
DESCRIPTION	VARCHAR2(255)	Asset description
STATUS	VARCHAR2(8)	Asset status. Can be 'NEW', 'MISSING', 'REPAIRED', 'IN USE'
VENDOR	VARCHAR2(254)	Asset Vendor
MODEL	VARCHAR2(254)	Asset Model
SERIALNUMBER	VARCHAR2(254)	Asset serial number
BARCODE	VARCHAR2(255)	Asset bar code
ASSETCODE	VARCHAR2(255)	Asset code
ASSETTYPE	VARCHAR2(255)	Asset type
ASSETTAG	VARCHAR2(255)	Asset Tag

Table 13-32: MVC_ASSETSUMMARY (Continued)

Column Name	Type	Description
ASSETCATEGORY	VARCHAR2(255)	Asset Category
GEOGRAPHICLOCATION	VARCHAR2(255)	Asset location
STORAGETIERCLASSIFICATION	VARCHAR2(255)	Asset storage Tier Name for example "Ultra High Availability"
STORAGETIERCOSTPERGB	NUMBER(36,2)	Asset storage Tier cost
DEPARTMENTNO	VARCHAR2(255)	Asset department no
DEPARTMENTNAME	VARCHAR2(255)	Asset department Name
PERCENTAGEOWNED	NUMBER(5, 2)	Percentage owned by Department
ADMINISTRATOR	VARCHAR2(255)	Asset maintained by
STAFFNAME	VARCHAR2(255)	Staff Name
STAFFPHONENUMBER	VARCHAR2(255)	Staff PH#
STAFFDEPARTMENT	VARCHAR2(255)	Staff Department
STAFFEMAIL	VARCHAR2(255)	Staff E-mail
RACKNUMBER	VARCHAR2(100)	
FLOOR	VARCHAR2(100)	
DATACENTER	VARCHAR2(100)	
CITY	VARCHAR2(100)	
REGION	VARCHAR2(100)	
COUNTRY	VARCHAR2(50)	
CONTINENT	VARCHAR2(20)	
ADDRESS	VARCHAR2(1024)	
ZIPCODE	VARCHAR2(16)	
LICENSE	VARCHAR2(4000)	
PURCHASEORDERNUMBER	VARCHAR2(255)	Asset Purchase Order no
DATEPURCHASED	DATE	Asset Date purchased
COST	NUMBER(36,2)	Asset cost
VALUE	NUMBER(36,2)	

Table 13-32: MVC_ASSETSUMMARY (Continued)

Column Name	Type	Description
SALVAGECOST	NUMBER(36,2)	Asset Deprecated Salvage Cost
DEPRECITIONPERIOD	NUMBER(18)	Asset Deprecation Period
DEPRECITIONMETHOD	NUMBER(16)	Asset Deprecation Method ('Straight line', 'Fixed declining', 'Double declining')
DEPRECITEDVALUE	NUMBER(36,2)	Asset Depreciated Value
RESELLER	VARCHAR2(255)	Asset Reseller
COMMENTS	VARCHAR2(4000)	Comments
ASSETFIXCOSTTAXPERDEPTPERYEAR	NUMBER	Asset Fixed cost tax per department per Year
STORAGFIXCOSTTAXPERDEPTPERYEAR	NUMBER	Storage Fixed cost tax per department per Year

Table 13-33: MVC_APPLICATIONSUMMARYVW

Column Name	Type	Description
APPLICATIONID	NUMBER(38)	Application ID
APPLICATIONNAME	VARCHAR2(256)	Application Name
DOMAINID	NUMBER(38)	Domain ID
APPLICATIONTYPE	NUMBER(38)	Application type
VENDOR	VARCHAR2(254)	Vendor
DESCRIPTION	VARCHAR2(1024)	Description of the application
STATUS	NUMBER(38)	Operational status
BUILDNUMBER	VARCHAR2(254)	Software build number
SERIALNUMBER	VARCHAR2(254)	Software serial number
TARGETOS	VARCHAR2(254)	Target operating system
PRODUCTNAME	VARCHAR2(254)	Product name
VERSION	VARCHAR2(254)	Software version
IDENTIFICATIONCODE	VARCHAR2(254)	Software identification code

Table 13-33: MVC_APPLICATIONSUMMARYVW (Continued)

Column Name	Type	Description
SOFTWAREELEMENTID	VARCHAR2(254)	Software element ID
HOSTID	NUMBER(38)	ID of host where the application is running from

Table 13-34: MVC_UNITACCESSVW

Column Name	Type	Description
ID	NUMBER(38)	
STORAGE_VOLUME_ID	NUMBER(38)	
STORAGE_SYSTEM_PORT_ID	NUMBER(38)	
HBA_PORT_ID	NUMBER(38)	
ACCESS_MODE	VARCHAR(254)	
INITIATOR	VARCHAR(254)	
INITIATOR_FORMAT	NUMBER(38)	
UNIT_NUMBER	NUMBER(38)	
HOST_GROUP	VARCHAR(254)	
HOST_GROUP_NAME	VARCHAR(254)	
HOST_GROUP_MODES	VARCHAR(1024)	

Application Core Views**Table 13-35: MVCA_DBAPPCAPACITYVW**

Column Name	Type	Description
DBApplicationID	NUMBER(38)	
HostID	NUMBER(38)	
CapacityType	Varchar2(7)	Type of capacity data
Timestamp	Timestamp(10)	
DBInstanceName	Varchar2(255)	
DBLogicalName	Varchar2(255)	
DBPhysicalName	Varchar2(512)	
TotalMB	NUMBER	
FreeMB	NUMBER	
UsedMB	NUMBER	

Table 13-36: MVCA_EXCHAPPCAPACITYVW

Column Name	Type	Description
ExchangeAppID	NUMBER(38)	
HostID	NUMBER(38)	
CapacityType	Varchar2(7)	
Timestamp	Timestamp	
StorageGroupID	NUMBER(38)	
ExchangeFilePath	Varchar(512)	
TotalMB	NUMBER	
FreeMB	NUMBER	
UsedMB	NUMBER	

Table 13-37: MVCA_VIRTUALAPPCAPACITYVW

Column Name	Type	Description
VirtualAppID	NUMBER(38)	
HostID	NUMBER(38)	
Timestamp	Date	
VirtualPath	Varchar2(512)	
TotalIMB	NUMBER	
FreeMB	NUMBER	
UsedMB	NUMBER	

Table 13-38: MVCA_FSRM_VOLUMESUMMARYVW

Column Name	Type	Description
VolumeID	NUMBER(38)	
VolumeName	Varchar2(256)	
FSID	NUMBER(38)	
TotalDirectories	NUMBER(38)	
TotalFiles	NUMBER(38)	
DomainID	NUMBER(38)	
Timestamp	Timestamp(6)	

Table 13-39: MVCA_FSRM_AGESUMMARYVW

Column Name	Type	Description
AgeID	NUMBER(38)	
VolumeID	NUMBER(38)	
TotalFiles	NUMBER(38)	
TotalSize	NUMBER(38)	
Timestamp	Timestamp(6)	

Table 13-40: MVCA_FSRM_EXTDETAILSUMMARYVW

Column Name	Type	Description
ExtName	Varchar2(254)	
VolumeID	NUMBER(38)	
TotalFiles	NUMBER(38)	
TotalSize	NUMBER(38)	
Timestamp	Timestamp(6)	

Table 13-41: MVCA_FSRM_DIRDETAILSUMMARYVW

Column Name	Type	Description
DirKey	NUMBER(38)	
ParentKey	NUMBER(38)	
DirName	Varchar2(254)	
DirLSevel	NUMBER(38)	
DirSize	NUMBER(38)	
TotalSubDirectories	NUMBER(38)	
TotalFiles	NUMBER(38)	

Table 13-41: MVCA_FSRM_DIRDETAILSUMMARYVW (Continued)

Column Name	Type	Description
VolumeID	NUMBER(38)	
Timestamp	Timestamp(6)	

Table 13-42: MVCA_FSRM_USERSUMMARYVW

Column Name	Type	Description
UserID	NUMBER(38)	
FSID	NUMBER(38)	
UserProvidederID	Varchar2(254)	
UserName	Varchar2(254)	
DirName	Varchar2(254)	
Department	Varchar2(254)	
Email	Varchar2(254)	
Quota	NUMBER(38)	
DomainID	NUMBER(38)	

Table 13-43: MVCA_FSRM_GROUPSUMMARYVW

Column Name	Type	Description
GroupID	NUMBER(38)	
FSID	NUMBER(38)	
GroupName	Varchar2(254)	
DirName	Varchar2(254)	
Contact	Varchar2(254)	
Department	Varchar2(254)	
Email	Varchar2(254)	
Quota	NUMBER(38)	
DomainID	NUMBER(38)	

Table 13-44: MVCA_FSRM_TOPNFILES

Column Name	Type	Description
GroupID	NUMBER(38)	
FSID	NUMBER(38)	
GroupName	Varchar2(254)	
DirName	Varchar2(254)	
Contact	Varchar2(254)	
Department	Varchar2(254)	
Email	Varchar2(254)	
Quota	NUMBER(38)	
DomainID	NUMBER(38)	

Table 13-45: MVCA_FSRM_AGEDFILEDETAILS

Column Name	Type	Description
VolumeID	NUMBER(38)	
FileName	Varchar2(254)	
FileSize	NUMBER(38)	
FileAge	NUMBER(38)	
Timestamp	Timestamp(6)	
DomainID	NUMBER(38)	

Table 13-46: MVCA_FSRM_LARGEDIRINFO

Column Name	Type	Description
VolumeID	NUMBER(38)	
DirName	Varchar2(256)	
DirSize	NUMBER(38)	

Table 13-46: MVCA_FSRM_LARGEDIRINFO (Continued)

Column Name	Type	Description
TotalFiles	NUMBER(38)	
TotalDirs	NUMBER(38)	
Timestamp	Timestamp(6)	
DomainID	NUMBER(38)	

Table 13-47: MVIEWCORE_STATUS

Name	Type
MVIEWNAME	NOT NULL VARCHAR2(30)
LAST_REFRESH_TIME	DATE
TOTALREFRESHTIME	VARCHAR2(32)
STATUS	VARCHAR2(10)

Table 13-48: MVIEW_STATUS

Name	Type
MVIEWNAME	NOT NULL VARCHAR2(30)
LAST_REFRESH_TIME	DATE
TOTALREFRESHTIME	VARCHAR2(32)
STATUS	VARCHAR2(10)

Views from Previous Releases

In this release, the materialized views were renamed, revised and in some cases removed. The following views were dropped from this release:

- MV_STORAGESYSTEMCAPSUMMARYVW
- MV_HOSTDETAILVW

■ MV_UNITACCESSVW

The following table provides the views from earlier releases. Although you can still use these views for your reports created in previous versions, it is strongly recommended you use the new views for any new report development. See the following tables for the corresponding new views.

Table 13-49: Views from Previous Releases

View Name	Corresponding New View
MV_ASSETCOUNTVW	No replacement
MV_ASSETSUMMARYVW	MVC_ASSETSUMMARYVW
MV_ASSETSUMMARYVW	MVC_ASSETSUMMARYVW
MV_HOSTSUMMARYVW , MV_HOSTCONNECTIVITYVW, MV_HOSTDETAILVW	MVC_HOSTSUMMARYVW
MV_APPLICATIONVW	Information about applications is now spread across several materialized views.
MV_EVENTVW	MVC_EVENTS_VW
MV_FILESERVERVW	Information about file servers is now spread across several materialized views.
MV_HOSTSTORAGEALLLOCATIONVW	Information spread across several materialized views, such as MVC_HOSTSUMMARYVW.
MV_HOSTSTORAGE_SUMMARYVW	Information about storage on hosts is now spread across several reports.
MV_HOSTHOSTSTORAGELOCATIONVW	No replacement. Information in several views now.
MV_HOSTSTORAGE_SUMMARYVW	MVC_HOSTSUMMARYVW
MV_STORAGE_SYSTEMPORTUTILVW	No replacement. Information in several views, such as MVC_STORAGE_SYSTEMCONFIGVW.
MV_SWITCHDETAILVW	MVC_SWITCHSUMMARYVW (In this release of the product, MV_SWITCHDETAILVW does not contain zone data. Use MVC_SWITCHSUMMARYVW to get the zone data.)
MV_ASSETSUMMARYVW	MVC_ASSETSUMMARY

Table 13-49: Views from Previous Releases (Continued)

View Name	Corresponding New View
MV_STORAGESYSTEMC APSUMMARYVW	STORAGEVOLUMESUMMARYVW

Implementing Custom Reports on Sun Solaris

If your management server runs only on Solaris, you can still take advantage of customized reports even though Report Designer only runs on Microsoft Windows:

1. Once Report Designer is installed on a Windows server, copy the report directory from the Sun Solaris server to the Microsoft Windows server.
This directory is located in the `JBossandJetty/server/appiq/reports` subdirectory on the Solaris server. Any empty directory can be used on the Windows server.

Note: You can use binary FTP to copy the files.

2. Customize your reports as described in this chapter.
3. Replace the customized files.
4. Once your customizations are complete, copy the reports directory back to the Solaris server. Copy the entire `JBossandJetty/server/appiq/reports` subdirectory.
5. Restart AppStorManager.
6. Restart the management server processes on the Solaris server.
Your customized reports should now be available in the product.

Viewing Performance Data

Important: Depending on your license, Performance Explorer may not be available. See the “List of Features” to determine if you have access to Performance Explorer. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center**).

This chapter describes the following:

- About Performance Explorer on page 630
 - Accessing Performance Explorer on page 631
 - Creating Performance Charts on page 631
 - About the Toolbars in Performance Explorer on page 632
 - Comparing the Performance of Different Elements on page 636
 - Viewing Summary Charts on page 637
 - Viewing Trending Information for Performance on page 638
 - Removing Performance Data from a Graph on page 639
 - Setting a Custom Period on page 639
 - About the Monitoring Options on page 641
 - Task Already Running on Element on page 652
-

- Monitoring with Direct Attached Storage on page 653
- Supported Host Configurations for Monitoring on page 653
- Sudden Dips Displayed in Certain Charts in Performance Explorer on page 656
- Values Continue to Increase in Charts for Aggregated Drives and Aggregate Volumes on page 657

About Performance Explorer

Important: Depending on your license, Performance Explorer may not be available. See the “List of Features” to determine if you have access to Performance Explorer. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center**).

Performance Explorer provides a graphical representation of the performance history of a managed element, such as bytes transmitted per second for a switch.

You can manipulate the charts, so they show a different reporting period and frequency. For example, you could show the performance of an element over the past 24 hours with an hourly monitoring frequency.

Keep in mind the following:

- If you see “There is not enough data to produce a chart [chart_title] at this time.” message, lessen the frequency option, so that the amount of time listed in the **Frequency** drop-down menu has passed before you view the chart again. For example, if the **Frequency** menu displays hourly, you need to wait an hour for data to appear in the chart. If the frequency option is **Default**, the information is displayed according to the interval of the performance collector that gathers information for the graph. For example, if you set the collector for the graph to every 15 minutes and the period set to **Last Hour**, the default setting displays four data points in the graph. See “Managing Performance Collection” on page 255 for more information about setting the performance collectors.
 - Verify the performance collector for that element is enabled (**Configuration > Performance**). See “Managing Performance Collection” on page 255 for more information about enabling performance collectors.
-

- If you select the **Default** option from the Frequency combo box in Performance Explorer and the maximum number of data points is more than 20, the management server displays statistics averaged over a three-minute interval. All data points collected within the three-minute period are averaged together to display one data point. For example, assume you are viewing the performance of the read bandwidth on a storage system and you selected the **Default** option from the Frequency combo box in Performance Explorer. If there are more than 20 data points within the selected time period, Performance Explorer employs an averaging algorithm to return an averaged representation of data in a three-minute period.
- Direct attached storage ports are not displayed in the storage tree in the bottom left pane.
- Performance Explorer is not available to file servers.
- To learn more about the buttons in the toolbar, see “About the Toolbars in Performance Explorer” on page 632.
- The Aggregated volume and aggregated drive filters are no longer shown for Engenio 5884 storage systems or any of its controllers. These filters were available in builds earlier than 4.0 of the management server.

Important: All collectors are stopped during Get Details. This means that during Get Details data for Performance Explorer is not updated. Historical collectors, such as those available from the Configuration tab, are restarted when they are stopped during Get Details. Any charts that were active in Performance Explorer when Get Details was started are not restarted.

Accessing Performance Explorer

To access Performance Explorer, click **PerfExplorer** ().

Creating Performance Charts

To create a performance chart for an element:

1. Access Performance Explorer as described in “Accessing Performance Explorer” on page 631.
2. Select the element you want to monitor.
3. Under the Monitoring tab in the lower-left pane, select the element again. In some instances, you may need to select an element's port, such as a switch.
4. In the pane under the tree, select a monitoring option.
See “About the Monitoring Options” on page 641 for more information.
5. Use the Chart, Unit and Frequency combo box to modify the chart. When you are done with your selections, click the  button in the lower pane. To learn more about these features, see “About the Toolbars in Performance Explorer” on page 632
6. To monitor more than one element in a chart, see “Comparing the Performance of Different Elements” on page 636.

About the Toolbars in Performance Explorer

Performance Explorer provides two toolbars. One in the upper pane and another in the lower pane. The toolbar in the upper pane is the same as the one in System Explorer. See “About the Toolbar in System Explorer” on page 300 for information about the toolbar in the upper pane.

The toolbar in the lower pane provides the following information:

Table 14-1: Toolbar in Lower Pane of Performance Explorer


Icon	Description
	Saves the graph in the cache. If you close the Web browser window, the graph is no longer saved.
	Lets you print a graph.
	Clears the graph of the elements you have selected.
	Lets you determine the type of graph displayed. Select one of the following options and then click the  button for your selection to take effect: ■ Bar - Displays each data point as a bar. The data for the different elements is displayed side by side. ■ Line - Displays each data point as a dot with a line connected to the previous data points. The data for the different elements for a specific point in time is displayed in the same column. ■ Summary - Displays a single line that summarizes the values for a single statistic. Multiple statistics can be shown with multiple lines. See “Viewing Summary Charts” on page 637.

Table 14-1: Toolbar in Lower Pane of Performance Explorer (Continued)


Icon	Description
	Lets you determine the unit of the measurement in the graph. Select one of the following options and then click the  button for your selection to take effect: ■ Default - Displays the data in its default unit, which is usually megabytes. ■ GB (Binary) - Displays the data in gibibytes, which is a binary unit of measurement. The computer handles the data in binary format. A gibibyte is equal to 1,073,741,824 bytes. ■ GB (Decimal) - Displays the data in gigabytes. A gigabyte is equal to 1,000,000,000 bytes.

Table 14-1: Toolbar in Lower Pane of Performance Explorer (Continued)


Icon	Description
	Lets you format the graph to provide data within the time period specified. Select the option to the left of the Period combo box. Select one of the following options and then click the  button for your selection to take effect: ■ Last Hour - Information collected in the last hour is reported. If you select Last Hour, the only frequency available is Default. ■ Last 24 Hours - Information collected in the last 24 hours is reported. ■ Last 7 Days - Information collected in the seven days is reported. ■ Last Month - Information collected in the last month is reported. ■ Last Year - Information collected in the last year is reported. ■ Custom - Lets you created custom settings for the graph. ■ Real Time - Displays the data as it is currently being gathered by the management server. The summary chart is the default chart for all real time performance display. If you select Real Time, the only frequency and unit available is Default. To set a custom period, select the Custom option. See “Setting a Custom Period” on page 639 for more information.

Table 14-1: Toolbar in Lower Pane of Performance Explorer (Continued)



Icon	Description
	Lets you change the display frequency. Select one of the following options and then click the  button for your selection to take effect: ■ Default - The information is displayed according to the interval of the performance collector that gathers information for the graph. For example, if you set the collector for the graph to every 15 minutes and the period set to Last Hour, the default setting displays four data points in the graph. See “Managing Performance Collection” on page 255 for more information about setting the performance collectors. See “Creating Performance Charts” on page 631 for more information. ■ Hourly - The information displayed is in hourly increments. ■ Daily - The information displayed is in daily increments. ■ Weekly - The information displayed is in weekly increments. ■ Monthly - The information displayed is in monthly increments.
	Click the  button after you have modified one or more of the following combo-boxes for your changes to take effect: ■ Chart ■ Unit ■ Period ■ Frequency
	If the chart contains more than one series of data, you can filter out the additional data series by using this Filter button.

Table 14-1: Toolbar in Lower Pane of Performance Explorer (Continued)

Icon	Description
	Lets you modify the performance data displayed in the graph and change graph settings.
	Creates a chart window.
	Lets you arrange the chart windows as follows: ■ Tiled - Displays the windows within their own tile. ■ Cascade - Displays the windows on top each other with the active window on top. ■ Minimize All - Minimizes all windows. ■ Restore All - Opens all windows that have been minimized.

Comparing the Performance of Different Elements

Use Performance Explorer to compare the performance of different elements. Let's assume you want to compare the physical memory used on different hosts, you would do the following:

1. Access Performance Explorer as described in "Accessing Performance Explorer" on page 631.
2. Click the element in the topology you want to see its performance.
In this example, let's assume you clicked a host named Cortez.
3. Select a component of the element.
In this example, let's assume you selected Cortez---->Monitoring under the **Application Path** node.
4. Select the graph you want to see on the Monitoring tab, located in the lower pane.
Let's assume you selected Physical Memory Used (%).
5. Scroll to the bottom of the storage tree in the lower-left pane.
6. Expand the **All Elements** node by clicking the (+) symbol next to the node name.

7. Expand the node for the host you want to compare.
For example, if you want to compare a host name HostA against Cortez, you would expand the following nodes in order: **All Elements** node, **Hosts** node, **HostA** node, **Application Path** node.
8. Select the same type of component you selected previously for Cortez.
9. Select the same type of graph in the left pane.
For example, if you selected “Physical Memory Used (%)” for the first element, you must select the same option for the other elements. You cannot select “Physical Memory Used (%)” for the first element and then “Processor Utilization (%)” for the second element. This is because these two options measure different types of data.

Performance Explorer displays information for the different elements in the same graph.
10. Repeat steps 7 through 9 for each host you want to compare against Cortez.

Viewing Summary Charts

Performance Explorer provides summary charts which display a single line that summarized the values for a single statistic. Multiple statistics can be shown with multiple lines.

The characteristic of the line has the following attributes:

- The vertical length of the line indicate the minimum and maximum value of the statistic within the selected data time frame.
- The green marker marks the value obtained from the latest time stamp (latest polled value).
- The purple marker marks the average value.

Summary charts are the default type for real time performance display.

To view a summary chart:

1. Access Performance Explorer as described in “Accessing Performance Explorer” on page 631.
 2. Create a chart as described in “Creating Performance Charts” on page 631.
 3. Select **Summary** from the Chart combo box in the lower pane.
 4. Click the  button in the lower pane for the change to take effect.
-

Viewing Trending Information for Performance

The management server can display trending information in its reports. For example, you can configure Performance Explorer to display trending information for the next week. This information can give you an indication of an element's future performance based on past performance.

Keep in mind the following:

- An element's performance can drastically change in the future. Please keep in mind the data trends are just assumptions and should not be treated as fact.
- Trending requires at least two sets of data gathered within the frequency specified.

To view trending information:

1. Access Performance Explorer as described in "Accessing Performance Explorer" on page 631.
 2. Click the element for which you want to see its performance.
 3. Under the Monitoring tab in the lower-left pane, select the element again. In some instances, you may need to select an element's port, such as a switch.
 4. In the lower-left pane on the Monitoring tab, click a performance monitoring option. The performance monitoring options listed in this figure vary according to the type of element. The monitoring buttons in the figure are for a switch.
 5. Click the **Edit Chart Property** () button in the lower pane.
 6. In the Performance Graphic Editing Dialog window, type a number in the **Trend** field. The number corresponds to the number of frequency intervals for which the trending information will be provided. For example, if you type 5 in the **Trend** field, the chart provides trending information for five frequency intervals, such as five weeks if weeks was selected from the **Frequency** drop-down combo box menu.
 7. Click **OK**.
The trending information is displayed in the chart.
-

Note: If there is not enough data to display, Performance Explorer does not display the chart. For example, if you selected the weekly option from **Frequency** drop-down combo box menu and you only have two days of data, a chart is not displayed regardless of the value in the **Trend** field. Performance Explorer does not display a chart if there is not enough data and the trending number is ignored.

Removing Performance Data from a Graph

To remove multiple data from a performance graph:

1. Access Performance Explorer as described in “Accessing Performance Explorer” on page 631.
2. In Performance Explorer, click the graph you want to modify in the bottom-right pane.
3. Click the Edit Chart Property () button in the lower pane.
4. In the Performance Graphic Editing Dialog window, click the  button corresponding with the data you want to remove from the graph.
5. When you are asked if you want to remove the data, click **Yes**.
6. Click **OK**.
7. Click the  button in the lower pane.

Setting a Custom Period

You can format the graph to provide data within a custom time period. This feature can be extremely useful for pinpointing performance changes. For example, assumed you changed the firmware of a switch two weeks ago and you want to compare the performance of that switch before and after you changed its firmware. You could create a graph that provides performance data two weeks before you changed the firmware. You could then create another graph that provides performance data two weeks after you changed the firmware.

Important: If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003/11/31, the software assumes the date is 2003/12/01.

To set a custom period:

1. Access Performance Explorer as described in “Accessing Performance Explorer” on page 631.
2. Click the element for which you want to see its performance.
3. Under the Monitoring tab in the lower-left pane, select the element again. In some instances, you may need to select an element's port, such as a switch.
4. In the lower-left pane on the Monitoring tab, click a performance monitoring option. The performance monitoring options listed in this figure vary according to the type of element.
5. Select **Custom** from the Period combo box in the lower pane. The Performance Graphic Editing Dialog window appears when you select the Custom option.
6. In the Performance Graphic Editing Dialog window, select the Custom option near the lower-left corner.
7. Click the calendar  button to the right of the **Start** field.
8. Enter the time in the time field. Make sure the time resembles a 24-hour clock, for example 22:00 for 10 p.m.
9. Click the date.
The date is highlighted in pink.

You can navigate the calendar as follows:



- Displays the same month in the previous year



- Displays the previous month



- Displays the next month



- Displays the same month in the following year

10. When you are done, click the **Set** button.
The start time and date are displayed in the **Start** field.

- 11. To set the end date, click the calendar  button to the right of the End field. Repeat steps 8 through 10.
The ending time and date are displayed in the **End** field.
- 12. Click **OK**.
- 13. Click the  button.

About the Monitoring Options

See the following table for more information about some of the monitoring options.

Table 14-2: About the Monitoring Options

Available Monitoring Options	Available to Which Elements	Description
Active Client Logons	Microsoft Exchange Stores	The number of active client logons.
Average Delivery Time (Messages/Sec)	Microsoft Exchange Stores	The average delivery time for messages per second.
Average IO Size (Bytes/Sec)	Engenio Storage Systems	The average input/out size (Bytes/Sec)

Table 14-2: About the Monitoring Options (Continued)

Available Monitoring Options	Available to Which Elements	Description
Buffer Cache Hits Count per Second	NAS Filers (System)	The number of buffer cache hits per second.
Buffer Cache Misses Count per Second	NAS Filers (System)	The number of buffer cache misses per second.
Buffer Hit Ratio	Oracle	The buffer cache hit ratio calculates how often a requested block has been found in the buffer cache, without requiring disk access. If the ratio is less than 90 percent investigate further. You may need to increase cache size by increasing db_block_buffer(8i) or db_cache_size(9i).
Bytes Transferred	Engenio Storage Systems	Bytes transferred on a drive.
Bytes Transmitted (MB/Sec)	■ Storage Systems ■ Host (port for HBA card) ■ NAS Filer (IP Port) ■ Switch Port	Number of bytes transmitted by the port per second

Table 14-2: About the Monitoring Options (Continued)

Available Monitoring Options	Available to Which Elements	Description
Bytes Received (MB/Sec)	■ Storage Systems ■ Host (port for HBA card) ■ NAS Filer (IP Port) ■ Switch Port	Number of bytes received by the port per second
CPU Usage Percentage (%)	Sybase	Percentage of the Sybase CPU being used. For example, if this monitoring appears pegged at 100 percent, this means one or more Sybase databases on the host are using 100 percent of the Sybase CPU, not the host CPU. The host CPU could be pegged at 60 percent, while the Sybase CPU is pegged at 100 percent.
CRC Errors (Errors/Sec)	■ Storage Systems ■ Host (port for HBA card) ■ Switch Port	The number of cyclic redundancy checking (CRC) errors per second.
Data Received (rate/sec)	NAS Filers	The data received per second by an IP port on a filer.
Data Sent (rate/sec)	NAS Filers	The data sent per second by an IP port on a filer.

Table 14-2: About the Monitoring Options (Continued)

Available Monitoring Options	Available to Which Elements	Description
Dictionary Hit Ratio	Oracle	Use this ratio to determine shared pool is large enough to store dictionary cache data adequately. If the ratio is less then 95 percent investigate further. You may need to increase shared_pool_size.
*Disk Read (KB/second)	Disk Drives	The speed at which the disk is read. To receive this data from a 64-bit AIX host, the bos.perf.libperfstat file must be installed on the host. See the Prerequisites section in the chapter for installing the CIM Extension on AIX in the Installation Guide.
Disk Total (KB/Sec)	Disk Drives	Total speed at which the disk is read and written for HP-UX hosts. To receive this data from a 64-bit AIX host, the bos.perf.libperfstat file must be installed on the host. See the Prerequisites section in the chapter for installing the CIM Extension on AIX in the Installation Guide.
Disk Utilization (%)	Disk Drives	The percentage of space used on the disk. To receive this data from a 64-bit AIX host, the bos.perf.libperfstat file must be installed on the host. See the Prerequisites section in the chapter for installing the CIM Extension on AIX in the Installation Guide.

Table 14-2: About the Monitoring Options (Continued)

Available Monitoring Options	Available to Which Elements	Description
*Disk Write (KB/second)	Disk Drives	The speed to which the disk is written. To receive this data from a 64-bit AIX host, the bos.perf.libperfstat file must be installed on the host. See the Prerequisites section in the chapter for installing the CIM Extension on AIX in the Installation Guide.
Exchange Services	Microsoft Exchange	Services Microsoft Exchange depends on to operate. A red line indicates the services have stopped, a green line indicates the services are running.
File Read Percent	Oracle	Percentage of “reads” for the file against the total “reads” in the database.
File Total I/O Percent	Oracle	This gives Datafile IO percentage against Total IO. Percentage of read+write for the file against total read+write in the DB.
File Write Percent	Oracle	Percentage of “write” for the file against the total “writes” in the database.
Final Destination Currently Unreachable Queue Size	Microsoft Exchange	The number of messages that cannot currently be sent to their final recipient.
Free Physical Memory (KB)	Host	The amount of free physical memory on the host. To receive this data from a 64-bit AIX host, the bos.perf.libperfstat file must be installed on the host. See the Prerequisites section in the chapter for installing the CIM Extension on AIX in the Installation Guide.

Table 14-2: About the Monitoring Options (Continued)

Available Monitoring Options	Available to Which Elements	Description
Free Virtual Memory (KB)	Host	The amount of free virtual memory on the host. To receive this data from an AIX host, the bos.perf.libperfstat file must be installed on the host.
In Memory Sort Ratio	Oracle	This ratio shows the proportion of sorts that are performed in memory. Optimally, most sorts should be performed in memory. If the ratio is less than 90 percent investigate further. You may need to increase sort_area_size or pga_aggregate_target (for Oracle 9i if the workarea_size_policy is auto.)
Inode Cache Hits Count per Second	NAS Filers (System)	The number of inode cache hits per second.
Inode Cache Misses Count per Second	NAS Filers (System)	The number of inode cache misses per second.
Invalid CRC Errors (errors/second)	Storage Systems	The speed at which invalid cycle redundancy checking (CRC) errors are found.
Library Cache Hit Ratio	Oracle	This gives an indication of how many times parsed SQL or PL/SQL statements needed to be reloaded, which requires a reload/rebuild of the statement. If the ratio is less than 95% investigate further. You may need to increase shared_pool_size.
Link Failures (failures/second)	■ Storage Systems ■ Host (port for HBA card) ■ Switch Port	The number of link failures per second.

Table 14-2: About the Monitoring Options (Continued)

Available Monitoring Options	Available to Which Elements	Description
Mail Box Count	Microsoft Exchange Stores	The number of mail boxes for a store.
Memory Usage Percentage (%)	Sybase	Percentage of the Sybase memory being used, not the host memory.
Messages Awaiting Directory Lookup Queue Size	Microsoft Exchange	The number of messages waiting for the recipient to be resolved in the global catalog, so the message can be sent.
Messages to be Routed Queue Size	Microsoft Exchange	The number of messages waiting to be routed to the recipient.
Name Cache Hits per Second	NAS Filers	The number of cache hits per second.
Name Cache Misses per Second	NAS Filers	The number of name cached misses per second on a NAS filer.
Packets Received (rate/second)	NAS Filers (IP Port)	The number of packets received per second by the IP ports on a filer.

Table 14-2: About the Monitoring Options (Continued)

Available Monitoring Options	Available to Which Elements	Description
Packets Sent (rate/sec)	NAS Filers (IP Port)	The packets sent per second by the IP ports on a filer.
Parse CPU to Total CPU Ratio	Oracle	Ratio closer to 0 percent is good. A high ratio means system is performing too many parses and indicates that shared pool is configured poorly or application is using SQL and so SQL cannot be shared. Setting Cursor_sharing=Force or SIMILAR might help in this case
Pre-submission Queue Size	Microsoft Exchange	The number of messages being held by the Pre-submission Queue, which holds messages waiting to be submitted to the Messages to be Routed Queue.
Processor Utilization (%)	Hosts	The percentage of the processor being used. To receive this data from a 64-bit AIX host, the bos.perf.libperfstat file must be installed on the host. See the Prerequisites section in the chapter for installing the CIM Extension on AIX in the Installation Guide.
Physical I/O Percentage (%)	Sybase	Percentage of the Sybase physical/ input output, not the host physical/input output.

Table 14-2: About the Monitoring Options (Continued)

Available Monitoring Options	Available to Which Elements	Description
Physical Memory Used (%)	Hosts	The percentage of physical memory used on the host. To receive this data from a 64-bit AIX host, the bos.perf.libperfstat file must be installed on the host. See the Prerequisites section in the chapter for installing the CIM Extension on AIX in the Installation Guide.
Percent Read (%)	Engenio Storage Systems	The percentage read.
Processor Busy Time	NAS Filers (processor)	The amount of time the processor is busy.
Processor Elapsed Time	NAS Filers (processor)	The amount of time that has passed since the NAS filer was rebooted.
Read IO Rate (Reads/Sec)	Engenio Storage Systems	The input/output of the read rate.
Read Operations	EMC Storage Systems	Read operations (bytes/second).
Read Requests	EMC Storage Systems	Read requests (bytes/second).
Receive Queue Size	Microsoft Exchange Stores	The size of the receive queue.
Recovered Errors	Engenio Storage Systems	The number of recovered errors on the drive.

Table 14-2: About the Monitoring Options (Continued)

Available Monitoring Options	Available to Which Elements	Description
Redo Buffer Allocation Retries Ratio	Oracle	It is number of times DSP had to wait to write to log buffer. If the ratio is more then 1 percent investigate further and consider increasing redo log buffer size.
Redo Logspace Request Ratio	Oracle	The number of times lgwr had to wait for writing to redo log file. If the ratio is more then .0002 (1 in 5000) investigate further and consider increasing redo log buffer size.
Requests Serviced	EMC Storage Systems	Requests serviced (bytes/second).
Reserved Inodes	NAS Filer (File System)	The number of reserved inode files on a NAS file system.
Retried Requests	Engenio Storage Systems	The number of retried requests for a drive.
Send Queue Size	Microsoft Exchange Stores	The size of the send queue.
SMTP Local Delivery Queue Size	Microsoft Exchange	The number of messages in the SMTP Local Delivery Queue, the queue hold mail box of local recipients.

Table 14-2: About the Monitoring Options (Continued)

Available Monitoring Options	Available to Which Elements	Description
SMTP Server Queues Summary	Microsoft Exchange	Provides a count of the total number of messages in the following queues: ■ Final Destination Currently Unreachable ■ Messages Awaiting Directory Lookup ■ Messages to Be Routed ■ Pre-submission ■ SMTP Local Delivery
Storage Group Size	Microsoft Exchange Storage Groups	The size of the storage group in megabytes.
Store size (MB)	Microsoft Exchange Stores	The size of the store in megabytes.
System Event Time Waited (ms)	Oracle	Displays the delta of time waited for the system wait events.
Tablespace Read Percent	Oracle	Percentage of “reads” for the tablespace against the total “reads” in the database.
Tablespace Write Percent	Oracle	Percentage of “writes” for the tablespace against the total “writes” in the database.
Tablespace Total I/O Percent	Oracle	This gives the tablespace input/output percentage against the total input/output. Percentage of read+write for the tablespace against total read+write in the database.
Timeouts	Engenio Storage Systems	The number of timeouts on the drive.
Total Bandwidth (Bytes/Sect)	Engenio Storage Systems	The Total bandwidth.

Table 14-2: About the Monitoring Options (Continued)

Available Monitoring Options	Available to Which Elements	Description
Total Inodes	NAS Filers (File systems)	The total number of inode files for a NAS file system.
Total IOs	Engenio Storage Systems	The total input/output.
Total IO Rate (IOs/Sec)	Engenio Storage Systems	The total input/output rate.
Unrecovered Errors	Engenio Storage Systems	The number of unrecovered errors on a drive.
Used Inodes	NAS Filers (File Systems)	The number of unused inode files on a NAS file system.
Virtual Memory Used (%)	Hosts	The percentage of virtual memory used on the host. To receive this data from an AIX host, the bos.perf.libperfstat file must be installed on the host.
Write IO Rate (Writes/Sec)	Engenio Storage Systems	The input/output of the write rate.
Write Operations	EMC Storage Systems	Write operations (bytes/second).

*These monitoring options are not available on HP-UX hosts because HP-UX hosts do not return read/write data separately.

Task Already Running on Element

If the software is trying to collect information for monitoring or statistic collection and the element is not returning the information in time, you are shown a message resembling the following:

```
Task xxxx is already running on element xxxx
```

The software cannot obtain the information in a timely manner because one or more of the following:

- The element might be inherently slow.
- The element might be busy with other tasks.
- You are trying to collect too much information at once from the element.

If you think you might be trying to collect too much information from the element, you might want to narrow down the collection. For example, if you are trying to collect monitoring information for three disks drives on a server, you might want to try collecting information for one disk drive.

The software stops trying to gather information after several unsuccessful attempts.

Monitoring with Direct Attached Storage

A port on a storage system directly attached to a host, does not appear in the left panel for monitoring. If you want to monitor the port, connect the port to a switch.

Supported Host Configurations for Monitoring

Refer to the following table to determine which host configurations the management server can monitor. The management server supports configurations that have a Y next to them. Nonsupported configurations have an N next to them.

Note: If a configuration listed in the following table is not supported, you can still obtain processor and memory statistics from the host. The exception, however, is Windows NT 4, which does not provide any monitoring information to the management server.

If the host has several configurations listed in the table below and one of them cannot be monitored, monitoring is not supported for all of the configurations on the host. For example, assume you have

a host with Solaris 9 Sun Foundation Suite Leadville with MPXIO and Solstice Disksuite/Volume Manager. Even though the management server supports monitoring for Solstice Disksuite/Volume Manager, both of those devices, Sun Foundation Suite Leadville with MPXIO and Solstice Disksuite/Volume Manager, cannot be monitored because Solaris 9 Sun Foundation Suite Leadville with MPXIO is not supported, as shown with the following formula:

A monitorable configuration (Y) + a unmonitorable configuration (N) =
unmonitorable configuration (N)

Keep in mind the following:

- In all configurations, you cannot monitor a VCM database device.
- The management server only monitors the top or bottom layer of Solstice Disksuite/Volume Manager. For example assume you have a normal configuration for Solstice Disksuite/Volume Manager (/folder <- d1 <- d2 <- d3 <- d4 <- cxydzs#). The management server reports on the folder and the highest layer, which is d1. Assume you have soft partitioning (/folder <- d1 <- d2 <- d3 <- d4 <- cxydzs#). The management server reports on the folder and the lowest layer, which is d4. The management server skips all layers between the d# devices. You are not able to monitor the middle layers, which are d2 and d3 in the previous example.

The following table is a sample of possible configurations, not a complete list of all products in the marketplace.

Table 14-3: Host Monitoring Support

Host Configuration	Monitoring Supported?
Solaris 8,9,10	Y
Solaris 8,9,10 VXVM	Y
Solaris 8,9 VXVM dmp	Y
Solaris 8,9 PowerPath	Y
Solaris 8,9 HDLM	N
Solaris 8,9 RDAC	N
Solaris 8,9 Sun Foundation Suite Leadville	Y
Solaris 8,9 DAS	N
Solaris 8,9 Sun Foundation Suite Leadville + MPXIO	N

Table 14-3: Host Monitoring Support (Continued)

Host Configuration	Monitoring Supported?
Solaris 8,9 Solstice DiskSuite/Volume Manager	Y
Solaris 8,9 Sun Foundation Suite Leadville + MPXIO + Solstice DiskSuite/Volume Manager	N
Windows 2003 HDLM	N
Windows 2003 Volume Manager	N
Windows 2003	Y
Windows NT 4	N
Windows 2000	Y
Windows 2000 PowerPath	N
Windows 2000 HDLM	N
Windows 2000 RDAC	N
Windows 2000 DAS	N
Windows 2000 Volume Manager	N
Windows 2000 Volume Manager and PowerPath	N
Windows 2000 Volume Manager and HDLM	N
SGI ProPack 3.0	Y
SGI ProPack 3.0 XVM	Y
SGI ProPack 3.0 CXFS	Y (only on node sending I/O)
AIX 5.1, 5.2 LVM	Y
AIX 5.1 PowerPath	Y
AIX 5.1 HDLM	Y
AIX 5.3	Y
AIX 5.3 SDD	N
AIX 5.3 SDDPCM	N
Irix 6.5.x	Y
Irix 6.5.x XVM	Y
Irix 6.5.x CXFS	Y (only on node sending I/O)
Redhat 2.1	Y
Redhat 3.0 Sistina LVM	Y

Table 14-3: Host Monitoring Support (Continued)

Host Configuration	Monitoring Supported?
Redhat 2.1 HDLM	N
Redhat 3.0	Y
Redhat 3.0 HDLM	N
Redhat 4.0	Y
Redhat 4.0 Itanium	Y
HP-UX 11.x Itanium	Y
HP-UX 11i	Y
HP-UX 11.0	Y
HP-UX 11i LVM	Y
HP-UX 11.0 LVM	Y
HP-UX 11i PV Link	Y
HP-UX 11.0 PV Link	Y
HP-UX 11i PV Link Volumes	Y
HP-UX 11.0 PV Link Volumes	Y
HP-UX 11.i with PowerPath	Y
HP-UX 11.0 with PowerPath	Y
HP-UX 11.0 with HDLM	Y
HP-UX 11.i with HDLM	Y

Sudden Dips Displayed in Certain Charts in Performance Explorer

In Performance Explorer and on the Monitoring tab, charts that display data gathered by certain volume and drive counters display their charts with the results increasing to the maximum value and then decreasing rapidly to a very low number that starts rising again.

The sudden dip in the charts occur because the counters holding the values displayed in the charts can contain at most 231 (2147483648) integers. The counters revert back to zero after reaching 231 (2147483648) integers. Then, they continue to go up again. The counters do not usually display zero in a chart because they are quickly gathering data again.

The following charts for individual drives are impacted: ReadIOs, WriteIOs, TotalIOs, Bytes Transferred, Unrecovered Errors, Recovered Errors, Timeouts, Retried Requests

The following charts for individual volumes are impacted: Bytes Read, Bytes Read Large, Bytes Written, Bytes Written Large, ReadIOs Large, WriteIOs Large, ReadIOs, WriteIOs, TotalIOs, Bytes Transferred, ReadHitIOs

Values Continue to Increase in Charts for Aggregated Drives and Aggregate Volumes

Values in charts for aggregate drives and aggregate volumes continue to rise smoothly in Performance Explorer and on the Monitoring tab. The only time the values in the charts come back down to the values returned by the array, is when the service for the management server restarts. Then, the charts display the values returned by the array. The values in the charts continue to rise until the service for the management server is stopped.

The values in the charts continue to rise because the values do not revert back to zero when an individual drive or volume counter that makes up the aggregation contain 231 (2147483648) integers. Instead, the values continue to rise smoothly.

The following charts for aggregate drives are impacted and charts for aggregated drives are also provided for arrays, controllers, and volume groups: ReadIOs, WriteIOs, TotalIOs, Bytes Transferred, Unrecovered Errors, Recovered Errors, Timeouts, Retried Requests

The following charts for aggregate volumes are impacted, and charts for aggregated volumes are also provided for arrays, controllers, and volume groups: Bytes Read, Bytes Read Large, Bytes Written, Bytes Written Large, ReadIOs Large, WriteIOs Large, ReadIOs, WriteIOs, TotalIOs, Bytes Transferred, ReadHitIOs

Finding an Element's Storage Capacity

Important: Depending on your license, Capacity Explorer may not be available. See the “List of Features” to determine if you have access to Capacity Explorer. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center**).

This chapter describes the following:

- About Capacity Explorer on page 660
 - Accessing Capacity Explorer on page 662
 - About the Toolbars in Capacity Explorer on page 662
 - Finding the Capacity of an Element on page 667
 - Obtaining Utilization Reports on page 673
 - Printing Elements in Capacity Explorer on page 673
 - Viewing Capacity Charts on page 676
 - Viewing Trending Information for Storage Capacity on page 678
 - “Different Results for the df -k Command and Capacity Explorer” on page 679
-

About Capacity Explorer

Important: Depending on your license, Capacity Explorer may not be available. See the “List of Features” to determine if you have access to Capacity Explorer. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center**).

Capacity Explorer provides graphical representation of an element's storage capacity and utilization in the storage network. It provides the following tabs:

- **List**
- **Path**
- **Utilization**
- **Capacity Data**
- **Capacity Chart**

Capacity Explorer provides a different set of information for NetApp NAS devices. For more information, see “Capacity Information for NetApp NAS Devices” on page 670.

Keep in mind the following:

- The Capacity Explorer displays the total capacity of an application, including the network drives. If you look at the capacity of the application in Chargeback, the capacity differs. Chargeback provides only network capacity. See “Capacity in Chargeback and Capacity Explorer Differ” on page 744 for more information.
- Volume names from ambiguous automounts on Solaris hosts are not displayed in Capacity Explorer. See “Volume Names from Ambiguous Automounts Are Not Displayed” on page 798.
- Capacity Explorer takes extra time to load the first time you access it after restarting the management server. The extra time is required for the management server to calculate the element capacity data. Capacity Explorer loads faster during subsequent times because the element capacity data has already been calculated.
- After discovering new elements, the capacity data for those elements will show up as null until the data collectors have run. Data collectors are set to run every 15 minutes by default.

About the List Tab

The List tab lets you quickly access an element. For example, to quickly access a host in the topology, expand the List View Tree node, then select your host in the tree by expanding the All Elements and Hosts nodes. When you select an element in the tree, it appears highlighted in the topology.

About the Path Tab

The Path tab provides information about an element's path. By clicking a host's node, you can determine the host's path in the application. When you click a host node in the tree, the elements in the host's path appear highlighted in the right pane.

About the Utilization Tab

Provides host and switch reports about certain aspects of utilization and storage, for example:

- **Remote vs. Local Utilization**
- **Remote Storage**
- **Local Storage**
- **Total Port Utilization**
- **Port Utilization by Fabric**

See “Obtaining Utilization Reports” on page 673 for more information.

About the Capacity Data Tab

The Capacity Data Tab provides information about the capacity of an element. You can quickly determine the general capacity of an element by its colors, as explain in the following legend.

Figure 15-1: Legend for Capacity Explorer

Color	Description
	Description varies according to element type: ■ For storage systems - The space is mapped. ■ For all elements except storage systems - The space is used.
	Unallocated

Figure 15-1: Legend for Capacity Explorer

	Unused Raw (storage systems only)
	Unmapped (storage systems only)

For example, the element displayed in the following figure is about 75 percent available. The rest of it is being used.



Figure 15-2: Capacity of an Element

You can obtain more detailed information about an element by clicking it in the right pane or in the Capacity Explorer tree, as explained in the topic “Finding the Capacity of an Element” on page 667.

About the Capacity Chart Tab

The Capacity Chart tab lets create bar or line charts to view your capacity data. You can use these charts to display trending information. See “Viewing Capacity Charts” on page 676 and “Viewing Trending Information for Storage Capacity” on page 678 for more information.

Accessing Capacity Explorer

To access Capacity Explorer,click **Capacity Explorer** ().

About the Toolbars in Capacity Explorer

Capacity Explorer provides two toolbars. One in the upper pane and another in the lower pane. The toolbar in the upper pane is the same as the one in System Explorer. See “About the Toolbar in System Explorer” on page 300 for information about the toolbar in the upper pane.

When the Capacity Chart tab is active, the toolbar in the lower pane provides the following information:

Note: Capacity Explorer provides additional tabs for NetApp NAS devices. The toolbars are the same.

Table 15-1: Toolbar in Lower Pane of Capacity Explorer

Icon	Description
	Lets you switch between the following: ■ Bar - Displays each data point as a bar. The data for the different elements is displayed side by side. ■ Line - Displays each data point as a dot with a line connected to the previous data points. The data for the different elements for a specific point in time is displayed in the same column.

Table 15-1: Toolbar in Lower Pane of Capacity Explorer (Continued)


Icon	Description
	Lets you format the graph to provide data within the time period specified. Select the option to the left of the Period combo box. Select one of the following from the drop-down combo box. Then, click the  button to update the chart. ■ Last 24 Hours - Information collected in the last 24 hours is reported. This option is only available to hosts and applications. ■ Last 7 Days - Information collected in the last seven days is reported. ■ Last Month - Information collected in the last month is reported. ■ Last Year - Information collected in the last year is reported.

Table 15-1: Toolbar in Lower Pane of Capacity Explorer (Continued)


Icon	Description
	Lets you format the graph to provide data within the starting and ending time specified. Select the option to the left of the Start field. Click the  icon. In the Time field, type the time you want the graph to start. Use the 24-hour clock. Select a date and click Set. Repeat the steps for the End field. Then, click the  button to update the chart. Important: If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the first day of the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.
	(Applications and Hosts only) Lets you change the display frequency. The options are the following: ■ Hourly - The information is displayed in hourly increments. ■ Daily - The information is displayed in daily increments. ■ Weekly - The information is displayed in weekly increments. ■ Monthly - The information is displayed in monthly increments.

Table 15-1: Toolbar in Lower Pane of Capacity Explorer (Continued)

Icon	Description
	Lets you set trending information. In the pane displaying the chart, type an integer in the Trend field. The integer corresponds to the number of frequency intervals for which the trending information will be provided. For example, if you type 5 in the Trend field, the chart provides trending information for five frequency intervals, such as five weeks if weeks was selected from the Frequency drop-down combo box menu. Click the  button to update the chart. When a switch or storage system is selected, the frequency field is set to hourly. See “Viewing Trending Information for Storage Capacity” on page 678. Keep in mind the following: ■ An element's performance can drastically change in the future. Please keep in mind the data trends are just assumptions and should not be treated as fact. ■ Trending requires at least two sets of data gathered within the frequency specified.
	Applies period, frequency and trending information.

Table 15-1: Toolbar in Lower Pane of Capacity Explorer (Continued)

Icon	Description
	If the chart contains more than one series of data, you can filter out the additional data series by using this Filter button.
	Lets you print a graph.

Finding the Capacity of an Element

Note: Capacity Explorer rounds the data it displays. This rounding could result in discrepancies in the totals you add for a property against the data displayed in the Summary column. For example, if you add the total capacity from each data pool and compare that total against the number for Total Capacity displayed under the Summary heading, you will most likely find the results differ.

To find the capacity of an element:

1. Access Capacity Explorer as described in “Accessing Capacity Explorer” on page 662.
2. (Optional) To quickly view the capacity of all of the elements in a fabric or application path, click the fabric or application path displayed in the tree for Capacity Explorer, as shown in the following figure.




Figure 15-3: Viewing the Capacity of Elements in a Fabric

See “About Capacity Explorer” on page 660 for an explanation of the colors displayed.

3. Do one of the following:

- ☐ Click an element in the right pane.
- ☐ Click an element in the tree for Capacity Explorer.

When an element is clicked, a pane appears along the bottom of the page and displays the capacity information.

Capacity Information for Applications

Capacity Explorer displays the following information under the **Capacity Summary** and **Capacity** columns:

- **Total Capacity**
- **Total Used**
- **Available**
- **Percentage Used** - The percentage used compared to the total capacity of the storage groups or database files.

The following additional information is displayed for each storage groups (Microsoft Exchange) or database file (Oracle):

- **Total Capacity**

- **Total Used**
- **Available**
- **Percentage Used** - The percentage used compared to the total capacity of the storage group or database file.

Capacity Information for Hosts

Capacity Explorer displays the following information under the **Capacity Summary** and **Capacity** columns:

- **Total Capacity**
- **Total Used**
- **Available**
- **Unmounted Volume** - The amount in megabytes of unmounted storage.

Note: This field automatically displays 0 MB if you discovered the host but not the storage system connected to it. This may occur if you did not enter the IP address of the storage system when performing discovery and/or your license does not allow you to discover a particular storage system. See the Supported Elements section in the “List of Features” to determine which storage systems you can discover. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center**).

- **Percentage Used** - The percentage used compared to the total capacity of the storage volumes.

The following additional information is displayed for each volume:

- **Total Capacity**
 - **Total Used**
 - **Available**
 - **Percentage Used** - The percentage used compared to the total capacity of the storage volume.
-

Capacity Information for NetApp NAS Devices

Capacity Explorer displays the following under the **Capacity Summary** and **Capacity** columns on the **Volume Data** tab:

- Total Aggregate Available
- Total Aggregate Used
- Total Volume Maximum
- Total Volume Used
- Percentage Total Volume Used

Capacity Explorer displays the following information about each volume:

- Volume Name
- Aggregate Available
- Aggregate Used
- Volume Maximum
- Volume Used
- Percentage Volume Used

Capacity Explorer displays the following information under the **Capacity Summary** and **Capacity** columns on the **Aggregate Data** tab:

- Aggregate Name
- Total Allocated Capacity
- Total Available Capacity
- Total Used Capacity
- Total Reserved Capacity
- Percentage Used

Capacity Explorer displays the following information about each aggregate:

- Allocated Capacity
 - Available Capacity
 - Used Capacity
 - Reserved Capacity
 - Percentage Used
-

Capacity Explorer displays the following information under the **Capacity Summary** and **Capacity** columns on the **Quota Data** tab:

- Total Disk Limit
- Total Disk Used
- Percentage Disk Used
- Total File Limit
- Total File Used
- Percentage File Used

Capacity Explorer displays the following information about each disk or file:

- Quota Name
- Quota Type - There are two types of quotas: disk (space) and file (count).
- Quota Limit - The amount of disk space or the number of files that is reserved for the target.
- Threshold - The amount of disk space that would have to be exceeded before a message is logged.
- Quota Soft Limit - The amount of disk space or the number of files that would have to be exceeded before a message is logged and and SNMP trap is generated.
- Quota Used
- Percentage Used

Capacity Information for Storage Systems

Capacity Explorer displays the following information under the **Capacity Summary** and **Capacity** columns:

- **Total Capacity**
- **Mapped**
- **Total Internal (HDS arrays only)**
- **Unmapped**
- **Total Allocated**
- **Unallocated**
- **Unused Raw**
- **Percentage Allocated**

Capacity Explorer displays additional information for each storage pool (Volume Group):

- **Total Capacity** - Total allocated, unallocated and unused raw.
- **Mapped** - Provisioned volumes exported to a system port. Mapped volumes are accessible.
- **Unmapped** - Provisioned volumes that are not exported to a system port. Unmapped volumes are not accessible.
- **Total allocated** - Provisioned volume.
- **Unallocated** - Available space left on the storage pool
- **Percentage Allocated** - The percentage allocated compared to the total capacity of the storage system.

The following table describes how the properties are calculated for HDS array groups.

**Table 15-2: Explanation of the Properties of the Capacity Levels
for HDS Array Groups**

Property Displayed for an Array Group (CIM_StoragePool)	How It Is Calculated	Explanation
Total Capacity	sum of volume.size + pool.totalRemainingSpace	sum of sizes of all LDEVs in the array group plus the total free space
Total Internal	(sum of volume.size - sum of external volume.size) + pool.totalRemainingSpace	sum of sizes of all internal LDEVs
Mapped	sum of volume.size for each volume that has a LUN	LDEVs that have LUNs
Unmapped	sum of volume.size for each volume that doesn't have a LUN	LDEVs that don't have LUNs but are not on the management server's "free" list

**Table 15-2: Explanation of the Properties of the Capacity Levels
for HDS Array Groups (Continued)**

Property Displayed for an Array Group (CIM_StoragePool)	How It Is Calculated	Explanation
Total Allocated	Sum of volume.size for all volumes	All LDEVs that are not on the management server's “free” list
Unallocated	pool.totalRemainingSpace	LDEVs on the management server's “free” list plus the total free space

Obtaining Utilization Reports

The software provides the following utilization reports to help you determine how much of your storage is being used:

- **Host Storage Utilization Reports**
- **Switch Utilization Reports**
- **Subsystem Utilization Reports**

To view a utilization report:

1. Access Capacity Explorer as described in “Accessing Capacity Explorer” on page 662.
2. Click the **Utilization** tab.
3. To view a utilization report, click one of the reports listed in the tree, as shown in the following figure.




Figure 15-4: Viewing a Utilization Report

The report appears in the right pane.

Printing Elements in Capacity Explorer

The software lets you print the topology in Capacity Explorer. For example, you can provide a printout to upper management that shows not only the topology of the network, but also the capacity of each element.

To print the elements in Capacity Explorer:

1. Access Capacity Explorer as described in “Accessing Capacity Explorer” on page 662.
2. If the topology spans more than one screen, arrange the elements so they are closer together, preferably on one screen. To move an element, click the  button and then the element you want to move. Drag the element to its new location. Moving elements closer together prevents the printout from appearing too stretched.
3. Click the  button.
4. Use the fields on the Paper tab to modify the setup of the page. When you are done, click **Apply**. If you want the default settings, click **Default**. A preview of the printout is displayed in the right pane.

Important: Before you change the margins, decide on a unit of measurement.

- ☐ **Paper format** - Select the paper size from the drop-down menu.

- ☐ **Unit** - Select cm (centimeters) or inch for the margins.
 - ☐ **Paper width** - Displays the width of the paper. You can modify the measurement in this field when you select the **Custom** option in the Paper format drop-down menu.
 - ☐ **Paper height** - Displays the height of the paper. You can modify the measurement in this field when you select the **Custom** option in the Paper format drop-down menu.
 - ☐ **Top margin** - Type a measurement for the top margin.
 - ☐ **Bottom margin** - Type a measurement for the bottom margin.
 - ☐ **Left margin** - Type a measurement for the left margin.
 - ☐ **Right margin** - Type a measurement for the right margin.
 - ☐ **Orientation** - Click an orientation for the printout.
5. Click the **View Selection** tab to modify how the printout will appear on the page. You can modify the following. When you are done, click **Apply**. If you want the default settings, click **Default**.
A preview of the printout is displayed in the right pane.

Important: Before you change the margins, decide on a unit of measurement.

- ☐ **Start x** - Determines the horizontal placement of the printout on the page with zero being the closest to the right margin. For example, if the value is 50 for **Start x**, the printing starts at 50 inches or centimeters (depending on what you selected) from the right margin. You can also enter negative numbers. Anything more than zero expands the printout to another page.
 - ☐ **Start y** - Determines the vertical placement of the printout on the page with zero being the closest to the bottom margin. For example, if the value is 50 for **Start y**, the printing starts at 50 inches or centimeters (depending on what you selected) from the bottom. You can also enter negative numbers.
 - ☐ **Width** - Determines the width of the printout.
 - ☐ **Height** - Determines the height of the printout.
- To remove extra space around the topology, click the **Trimmed** button.
6. To change how many pages the printout will use, select one of the following. When you are done, click **Apply**. If you want the default settings, click **Default**.
A preview of the printout is displayed in the right pane.

Important: Before you change the margins, decide on a unit of measurement.

- ☐ **Unit** - Select cm (centimeters) or inch for the margins.
 - ☐ **Position/Size** - Lets you change the position and size of the printout so that it spans several pages:
 - **Start x** - Determines the horizontal placement of the printout. Anything more than zero expands the printout to another page.
 - **Start y** - Determines the vertical placement of the printout.
 - **Width** - Determines the width of the printout. If the width entered does not fit on the page, the printout wraps around to another page.
 - **Height** - Determines the height of the printout. If the height entered does not fit on the page, the printout wraps around to another page.
 - ☐ **Resolution (pixel/unit)** - Lets you change the resolution so that the printout spans several pages.
 - ☐ **Page** - Lets you expand the printout so it prints on several pages without modifying the chart.
7. To preview your pages, click the Preview tab. Then click the page you want to preview. The page appears in the right pane.
 8. When you are ready to print, click **Print**.
 9. Click **Close**.

Note: To revert back to all of the original settings, click the **Default** button next to the **Print** button.

Viewing Capacity Charts

Capacity Explorer provides a graphical representation of the capacity history of an element, such as port summary information for switches.

You can manipulate the charts, so they show a different reporting period and frequency. For example, you could show the capacity of a host over the past 24 hours with an hourly monitoring frequency.

Keep in mind the following:

- Verify the performance collector for the element is enabled (**Configuration > Performance > Data Collection**). See “Managing Performance Collection” on page 255 for more information about enabling performance collectors.
- Switches and storage systems display data from the last time Get Details was performed. See “Step 1 - Discover Switches” on page 34 for more information.
- If you see “There is not enough data to produce a chart [chart_title] at this time.” message, lessen the frequency option, so that the amount of time listed in the **Frequency** drop-down menu has passed before you view the chart again. For example, if the **Frequency** menu displays hourly, you need to wait an hour for data to appear in the chart.

To find the performance of an element:

1. Access Capacity Explorer as described in “Accessing Capacity Explorer” on page 662.
2. Click the element for which you want to see a history of its capacity.
3. In the bottom split pane, click the **Chart** tab.
4. In the lower-middle split pane, click a report title.
The chart for the monitoring option is displayed in the lower-right pane.

Note: You can change the orientation of the chart by holding down the mouse button when you click the chart and moving the mouse.

5. To change the reporting period, do one of the following. Then, click the  button to update the chart.
 - **Display data within a time period** - You can format the graph to provide data within the time period specified. Select the option to the left of the **Period** combo box. Select one of the following from the drop-down combo box. Then, click the  button to update the chart.
 - **Last 24 hours** - Information collected in the last 24 hours is reported. This option is not available to switches and storage systems.
 - **Last 7 Days** - Information collected in the seven days is reported.
 - **Last Month** - Information collected in the last month is reported.
 - **Last Year** - Information collected in the last year is reported.
 - **Display data within a starting and ending time** - You can format the graph to provide data within the starting and ending time specified. Select the option to the left of the **Start** field. Click the  icon. In the **Time** field, type the time you want the graph

to start. Use the 24-hour clock. Select a date and click **Set**. Repeat the steps for the

End field. Then, click the  button to update the chart.

Important: If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the first day of the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.

6. (Applications and Hosts only) To change the display frequency, select one of the following from the **Frequency** drop-down combo box.
 - ☐ **Hourly** - The information is displayed in hourly increments.
 - ☐ **Daily** - The information is displayed in daily increments.
 - ☐ **Weekly** - The information is displayed in weekly increments.
 - ☐ **Monthly** - The information is displayed in monthly increments.
7. If the chart contains more than one series of data, you can filter out the additional data series by clicking the  button.
8. To add trending information, type an integer in the **Trend** field.
The integer corresponds to the number of frequency intervals for which the trending information will be provided. For example, if you type 5 in the **Trend** field, the chart provides trending information for five frequency intervals, such as five weeks if weeks was selected from the **Frequency** drop-down combo box.

See "Viewing Trending Information for Storage Capacity" on page 678.
9. Click the  button to update the chart.
10. To print the chart, click the  button displayed in the same pane as the chart.

Viewing Trending Information for Storage Capacity

The management server can display trending information in its reports. For example, you can configure Capacity Explorer to display trending information for the next week. This information can give you an indication of an element's future capacity need based on its past capacity utilization.

Keep in mind the following:

- An element's capacity can drastically change in the future. Please keep in mind the data trends are just assumptions and should not be treated as fact.
- Trending requires at least two sets of data gathered within the frequency specified.

To view trending information:

1. Access Capacity Explorer as described in “Accessing Capacity Explorer” on page 662.
2. Click the element for which you want to see a history of its performance.
3. In the bottom pane, click the **Chart** tab.
4. In the lower-middle split pane, click a monitoring option.
5. In the pane displaying the chart, type a number in the **Trend** field.
The number corresponds to the number of frequency intervals for which the trending information will be provided. For example, if you type 5 in the **Trend** field, the chart provides trending information for five frequency intervals, such as five weeks if weeks was selected from the **Frequency** drop-down menu.
6. Click the  button.
The trending information is displayed.

Note: If there is not enough data to display, Capacity Explorer does not display the chart. For example, if you selected the weekly option from **Frequency** drop-down menu and you only have two days of data, a chart is not displayed regardless of the value in the **Trend** field. Capacity Explorer does not display a chart if there is not enough data and the trending number is ignored.

Different Results for the df -k Command and Capacity Explorer

If you run the `df -k` command on UNIX for a storage system, you may notice the total capacity displayed does not match the total capacity in Capacity Explorer. This difference occurs because Capacity Explorer calculates the total capacity differently than the `df -k` command. The `df -k` command calculates the total capacity as follows:

```
used capacity + available capacity + reserved capacity
```

Capacity Explorer calculates the capacity differently, as shown below:

`used capacity + available capacity`

The difference between the two calculations is the capacity reserved for superuser. If a file system has a reserved capacity, the total capacity from the `df -k` command and Capacity Explorer will differ.

For example, assume you ran the `df -k` command for the file system, `/dev/dsk/c0t0d0s0`. After you run the `df -k` command, you notice that the total capacity displayed is 6688076 KB. When you look at Capacity Explorer, the total capacity displayed is 6621196 KB. Actually, Capacity Explorer displays results in megabytes, but for this example, it is easier to have the totals using the same units.

The totals differ. How did this happen? When you ran the `df -k` command, the computer ran the equation mentioned earlier (`used capacity + available capacity + reserved capacity`):

`1904031 KB + 4717165 KB + 66880 KB = 6688076 KB`

where

- 1904031 KB is the used capacity
- 4717165 KB is the available capacity
- 66880 KB is the capacity reserved for the superuser. The percentage of the reserved capacity can be set when using the `newfs -m` command.

Capacity Explorer calculated the total capacity by using the equation discussed previously (`used capacity + available capacity`) and displayed the result in megabytes:

`1904031 KB + 4717165 KB = 6621196 KB`

where

- 1904031 KB is the used capacity
- 4717165 KB is the available capacity

Because Capacity Explorer does not include the reserved capacity in its calculations, the difference between the two calculations is the capacity reserved for the superuser, which is 66880 KB.

Important: Depending on your license, Policy Manager may not be available. See the “List of Features” to determine if you have access to Policy Manager. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center**).

This chapter describes the following:

- About Policy Manager on page 682
 - Accessing Policy Manager on page 682
 - Creating Policies on page 683
 - Modifying Policies on page 697
 - Viewing Policies on page 706
 - Deactivating a Policy on page 707
 - Deleting Policies on page 707
 - Providing E-mail Notification for a Policy on page 708
 - Providing Event Generation for a Policy on page 708
 - Providing a Custom Command for a Policy on page 710
-

About Policy Manager

Policy Manager can automatically send an e-mail, generate an event, or run a custom script when an element is being overused or when one of the following occurs:

- A new element is discovered
- Successful provisioning occurred
- An event occurred on one or more specified elements

Important: Policies are not triggered when a collector is running. For example, assume you configured a policy to be triggered when a host has a critical event. If a collector is running when a critical event occurs, the policy is not triggered. Collectors are used to gather information for reports, monitoring and File Server SRM. You can modify a collector's schedule, by going to **Configuration > Reports, Configuration > Performance** or **Configuration > File SRM**.

if the collector for a report or Performance Explorer is running, the policy is not triggered.

Policy Manager lets you create the following types of policies:

- **Utilization policies** - Monitors the utilization of an element. Options provided depend on the type of element.
For example, you can configure Policy Manager so you receive an e-mail message when the amount of free space on a server decreases to a specified level.
- **Infrastructure policies** - Monitors the following:
 - ☐ A new element is discovered
 - ☐ Successful provisioning occurred
 - ☐ An event occurred on one or more specified elements

Accessing Policy Manager

This topic provides the various techniques for accessing Policy Manager. You have only one way to access policies for discovery, provisioning, and events; however, you have multiple ways to access

policies. The flexibility in accessing utilization policies lets you easily create and manage policies without interrupting your work flow.

Policy Manager provides the following options for accessing policies:

- **To access from System Explorer** - Double-click an element in System Explorer and then click the **Policies** tab. This method displays the utilization policies for just the element that was double-clicked.
- **To access policies for file servers** - Access the Policy tab in Application Explorer. Expand the **Applications** and **File Servers** nodes in the tree in Application Explorer. Click the file server name in the Application Explorer tree. Click the **Policies** tab in the right pane.

To access all policies in Policy Manager:

1. Click **Policy Manager** ().
2. Click a policy in the Policy Manager tree.

Creating Policies

This section describes the following:

- Creating a Utilization Policy on page 683
- Creating Policies for Discovery on page 690
- Creating Policies for Provisioning on page 692
- Creating Policies for Events on page 694
- Testing a Utilization Policy on page 697

Creating a Utilization Policy

You can create a utilization policy that generates an event, sends an e-mail, or runs a custom command when an element is being overused. For example, you can configure Policy Manager so you receive an e-mail message when the amount of free space on a server decreases to a specified level.

Keep in mind the following:

- If you plan to use e-mail notification with your policy, first assign an SMTP server from which the management server can send its e-mail notifications. See “Setting Up E-mail Notification” on page 207.
- Policies that are triggered for virtual applications are also triggered for file servers.

To create a utilization policy:

1. Access Policy Manager as described in the topic, “Accessing Policy Manager” on page 682.
2. In the left pane, select an element or element type you want the policy to apply to.
3. In the right pane, click the **Add** button.
4. Select a policy. Then, click **Next**. See Table 16-1, “Policy Templates,” on page 687 for more details on available policies.
5. In the **Name** field on the Policy Properties tab, modify the name of the policy or keep the default.
6. In the **Description** field, modify the description for the policy or keep the default.
7. In the **Re-arm Period** field, specify a re-arm period in minutes.
The re-arm period is the amount of time after the policy executes before it can execute again. It is useful for limiting the number of times the same actions will execute.

Important: Specify shorter periods for important applications.

8. In the **Condition** drop-down menu, specify a comparison operator.
 9. To the right of the comparison operator drop-down menu, do the following depending on the field displayed:
 - ☐ Type a percentage.
 - ☐ Type an amount in gigabytes (GB).
 10. For trending policies, enter the number of days in the **Historic period** field (min=7; max=180) and **Projection period** field (min=1; max=180).
 11. Select an action to occur when the policy condition is fulfilled by clicking one of the following buttons (more than one action can be assigned to a policy):
 - ☐ **Send E-mail** - Policy Manager sends an e-mail when the condition is fulfilled. Enter a comma-separated list of e-mail addresses, and then click **OK**.
 - ☐ **Generate Event** - Policy Manager generates an event of the specified event type. The event appears in Event Manager. Select one of the following event types.
-

Important: Since the severity level for an element is set by the manufacturer, the meanings of the severity levels vary. It is best to view the description of the event.

- **Unknown** - The severity level is not known.
- **Informational** - An example of an informational event is a progress report event for firmware download operation currently in progress.
- **Warning** - An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
- **Minor** - An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
- **Major** - An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
- **Critical** - An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.

After you select an event level, click **OK**.

- **Execute a Custom Command** - Policy Manager executes a custom command on the management server when the condition is fulfilled. Type a command that will execute the script in the field. Then, click **OK**. The software assumes you are in the `%JBOS4_DIST%\server\appiq\remotescripts` directory on the management server when the script is executed. You can use environment variables in your script, such as `POLICY_NAME` and `POLICY_DESCRIPTION`. `POLICY_NAME` provides the policy name and `POLICY_DESCRIPTION` provides the policy description. See “Software Environment Variables for Scripting” on page 358 for more information. Prefix the command with “start” if the custom command triggers a user interface component, such as a program that uses the user interface (Example: Microsoft Internet Explorer) or a command prompt window.

For example, assume you want the custom command to open a command prompt window and list the contents in the directory. You would prefix the command with start as shown below:

```
start dir
```

12. Click **Finish**.
13. To test a policy, click the **Test** button in the Utilization Policy table.
The management server fires a test for all utilization policies associated with that element.

Utilization Policies

Delete ✓

☐	<u>Name</u>	<u>Description</u>	<u>Active</u>	Test	Edit	Delete
☐	Application Free Space	Free space for each Application	✓	Test		

Figure 16-1: Testing a Newly Created Utilization Policy

About the Policy Templates

The following table provides descriptions of the policy templates.

Table 16-1: Policy Templates

Policy Type	Description	Applies to
Free Space	Amount of free space on one of the following: ■ A host ■ A Microsoft Exchange instance ■ A database instance, such as Microsoft SQL Server, Sybase, or Oracle. ■ File Servers ■ Virtual Applications	■ Hosts ■ Microsoft Exchange ■ Microsoft SQL Server ■ Oracle ■ Sybase ■ File Servers ■ Virtual Applications
Free Space Trending	Forecast of amount of free space on one of the following: ■ A host ■ A database instance, such as Microsoft SQL Server, Sybase, or Oracle. ■ File Servers ■ Virtual Applications	■ Hosts ■ Microsoft SQL Server ■ Oracle ■ Sybase ■ File Servers ■ Virtual Applications

Table 16-1: Policy Templates (Continued)

Policy Type	Description	Applies to
Percent Used	Description varies according to element type: ■ Hosts - Percent of storage used for a host. ■ Microsoft Exchange instances - Percent of storage used for a Microsoft Exchange instance. ■ Database instances - Percent of storage used for a database instance, such as Microsoft SQL Server, Sybase, or Oracle. ■ Switches - Percent of ports used for a switch. ■ Storage Systems - Percent of ports used for a storage system ■ File Servers - Percent of disk space used. ■ File Server Users - Percent of storage used for a file server user. ■ Virtual Applications - Percent of storage used for a virtual application.	■ Hosts ■ Microsoft Exchange ■ Microsoft SQL Server ■ Oracle ■ Sybase ■ Storage Systems ■ Switches ■ File Servers ■ File Server Users ■ Virtual Applications

Table 16-1: Policy Templates (Continued)

Policy Type	Description	Applies to
Trending for Percent Used	Forecast of percent used. Description varies according to element type: ■ Hosts - Percent of storage used for a host. ■ Database instances - Percent of storage used for a database instance, such as Microsoft SQL Server, Sybase, or Oracle. ■ Switches - Percent of ports used for a switch. ■ File Servers - Percent of disk space used. ■ Virtual Applications - Percent of storage used for a virtual application.	■ Hosts ■ Microsoft SQL Server ■ Oracle ■ Sybase ■ Switches ■ File Servers ■ Virtual Applications
Unmapped Storage	Space not assigned to any mapped volume for the storage system	Storage systems
Unmapped Storage Trending	Forecast of space not assigned to any mapped volume for the storage system.	Storage systems
Unmapped Storage Percent	Percent of total space not assigned to any mapped volume for the storage system	Storage systems
Unmapped Storage Percent Trending	Forecast of percent of total space not assigned to any mapped volume for the storage system.	Storage systems
Unused Raw Storage	Space not assigned to any storage pool for a storage system.	Storage systems
Unused Raw Storage Trending	Forecast of space not assigned to any storage pool for a storage system.	Storage systems

Table 16-1: Policy Templates (Continued)

Policy Type	Description	Applies to
Unused Raw Storage Percent	Percent of total space not assigned to any storage pool for a storage system.	Storage systems
Unused Raw Storage Percent Trending	Forecast of percent of total space not assigned to any storage pool.	Storage systems
Unused Ports	Number of unused ports for a switch or storage system.	■ Switches ■ Storage systems
Unused Ports Trending	Forecast of unused ports for a switch.	Switches

Creating Policies for Discovery

You can create an infrastructure policy that generates an event, sends an e-mail, or runs a custom command when an element is discovered.

Important: If you plan to use e-mail notification with your policy, first assign an SMTP server from which the management server can send its e-mail notifications. See “Setting Up E-mail Notification” on page 207.

To create a policy for discovery:

1. Access Policy Manager as described in the topic, “Accessing Policy Manager” on page 682.
 2. In the Policy Manager tree, expand the node, Infrastructure Policies. Then, select **New Element Discovery**.
 3. Click the **Add** button in the right pane.
-

4. In the **Name** field, type a name for the policy.
5. In the **Description** field, type a description for the policy.
6. Select one or more element types.
When a condition is fulfilled on a select element, Policy Manager generates an event, sends an e-mail, or runs a custom command.
7. Select **Fire when event is cleared** if you want the policy to fire when the event is cleared, Otherwise the event fires when the event is received.
8. Select one of the following from the **Summary Text** drop-down menu:
 - ☐ **Is anything** - Regardless of the contents of the event's summary text, Policy Manager sends an e-mail, generates an event, or runs a custom command.
 - ☐ **Contains** - If the event's summary text contains the specified text, Policy Manager sends an e-mail, generates an event, or runs a custom command. Type the specified text in the field that appears to the right of the **Summary Text** drop-down menu.
 - ☐ **Matches Regular Expression** - If the event's summary text matches the specified expression, Policy Manager sends an e-mail, generates an event, or runs a custom command. Type the specified text in the field that appears to the right of the **Summary Text** drop-down menu.
9. Select an action to occur when the policy condition is fulfilled by clicking one of the following buttons (more than one action can be assigned to a policy):
 - ☐ **Send E-mail** - Policy Manager sends an e-mail when the condition is fulfilled. Enter a comma-separated list of e-mail addresses, and then click **OK**.
 - ☐ **Generate Event** - Policy Manager generates an event of the specified event type. The event appears in Event Manager. Select one of the following event types.

Important: Since the severity level for an element is set by the manufacturer, the meanings of the severity levels vary. It is best to view the description of the event.

 - **Unknown** - The severity level is not known.
 - **Informational** - An example of an informational event is a progress report event for firmware download operation currently in progress.
 - **Warning** - An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
 - **Minor** - An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
 - **Major** - An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.

- **Critical** - An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.

After you select an event level, click **OK**.

- **Execute a Custom Command** - Policy Manager executes a custom command on the management server when the condition is fulfilled. Type a command that will execute the script in the field. Then, click **OK**. The software assumes you are in the `%JBOSSE4_DIST%\server\appiq\remotescripts` directory on the management server when the script is executed. You can use environment variables in your script, such as `POLICY_NAME` and `POLICY_DESCRIPTION`. `POLICY_NAME` provides the policy name and `POLICY_DESCRIPTION` provides the policy description. See “Software Environment Variables for Scripting” on page 358 for more information. Prefix the command with “start” if the custom command triggers a user interface component, such as a program that uses the user interface (Example: Microsoft Internet Explorer) or a command prompt window.

For example, assume you want the custom command to open a command prompt window and list the contents in the directory. You would prefix the command with start as shown below:

```
start dir
```

10. Click **OK**.

Creating Policies for Provisioning

You can create an infrastructure policy that generates an event, sends an e-mail, or runs a custom command when successful provisioning occurred.

Important: If you plan to use e-mail notification with your policy, first assign an SMTP server from which the management server can send its e-mail notifications. See “Setting Up E-mail Notification” on page 207.

To create a policy for provisioning:

1. Access Policy Manager as described in the topic, “Accessing Policy Manager” on page 682.
 2. In the Policy Manager tree in the middle pane, expand the node, Infrastructure Policies. Then, click **Provisioning**.
-

3. Click the **Add** button.
4. In the **Name** field, type a name for the policy.
5. In the **Description** field, type a description for the policy.
6. Select one or more element types.
When a condition is fulfilled on a select element, Policy Manager generates an event, sends an e-mail, or runs a custom command.
7. Select **Fire when event is cleared** if you want the policy to fire when the event is cleared, Otherwise the event fires when the event is received.
8. Select an action to occur when the policy condition is fulfilled by clicking one of the following buttons (more than one action can be assigned to a policy):
 - ☐ **Send E-mail** - Policy Manager sends an e-mail when the condition is fulfilled. Enter a comma-separated list of e-mail addresses, and then click **OK**.
 - ☐ **Generate Event** - Policy Manager generates an event of the specified event type. The event appears in Event Manager. Select one of the following event types.

Important: Since the severity level for an element is set by the manufacturer, the meanings of the severity levels vary. It is best to view the description of the event.

- **Unknown** - The severity level is not known.
- **Informational** - An example of an informational event is a progress report event for firmware download operation currently in progress.
- **Warning** - An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
- **Minor** - An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
- **Major** - An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
- **Critical** - An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.

After you select an event level, click **OK**.

- ☐ **Execute a Custom Command** - Policy Manager executes a custom command on the management server when the condition is fulfilled. Type a command that will execute the script in the field. Then, click **OK**. The software assumes you are in the `%JBOSS4_DIST%\server\appiq\remotescripts` directory on the management server when the script is executed. You can use environment variables in your script, such as `POLICY_NAME` and `POLICY_DESCRIPTION`. `POLICY_NAME`. `POLICY_NAME`

provides the policy name and `POLICY_DESCRIPTION` provides the policy description. See “Setting Up E-mail Notification” on page 207 for more information.

Prefix the command with “start” if the custom command triggers a user interface component, such as a program that uses the user interface (Example: Microsoft Internet Explorer) or a command prompt window.

For example, assume you want the custom command to open a command prompt window and list the contents in the directory. You would prefix the command with start as shown below:

```
start dir
```

9. Click **OK**.

Creating Policies for Events

You can create a policy that generates an event, sends an e-mail, or runs a custom command when an event occurred on one or more specified elements

For example, you can create a policy that sends an e-mail when a new element generates a critical event.

If you plan to use e-mail notification with your policy, first assign an SMTP server from which the management server can send its e-mail notifications. See “Setting Up E-mail Notification” on page 207.

To create a policy for events:

1. Access Policy Manager as described in the topic, “Accessing Policy Manager” on page 682.
 2. In the Policy Manager tree in the middle pane, expand the node, Infrastructure Policies. Then, click **Events**.
 3. Click the **Add** button.
 4. In the **Name** field, type a name for the policy.
 5. In the **Description** field, type a description for the policy.
 6. In the **Re-arm Period** field, specify a re-arm period in minutes.
The rearm period is the amount of time after the policy executes before it can execute again. It is useful for limiting the number of times the same actions will execute.
-

Important: Specify shorter periods for important applications.

7. Select one or more element types.
When a condition is fulfilled on a select element, Policy Manager generates an event, sends an e-mail, or runs a custom command.
8. Select **Fire when event is cleared** if you want the policy to fire when the event is cleared, Otherwise the event fires when the event is received.
9. In the **Severity** drop-down menu, specify a comparison operator.
10. To the right of the **Severity** drop-down menu, select one of the following severity levels:

Important: Since the severity level for an element is set by the manufacture, the meanings of the severity levels vary. It is best to view the description of the event.

- ☐ **Unknown** - The severity level is not known.
 - ☐ **Informational** - An example of an informational event is a progress report event for firmware download operation currently in progress.
 - ☐ **Warning** - An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
 - ☐ **Minor** - An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
 - ☐ **Major** - An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
 - ☐ **Critical** - An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.
11. Select one of the following from the **Summary Text** drop-down menu:
 - ☐ **Is anything** - Regardless of the contents of the event's summary text, Policy Manager sends an e-mail, generates an event, or runs a custom command.
 - ☐ **Contains** - If the event's summary text contains the specified text, Policy Manager sends an e-mail, generates an event, or runs a custom command. Type the specified text in the field that appears to the right of the **Summary Text** drop-down menu.
 - ☐ **Matches Regular Expression** - If the event's summary text matches the specified expression, Policy Manager sends an e-mail, generates an event, or runs a custom command. Type the specified text in the field that appears to the right of the **Summary Text** drop-down menu.
-

12. Select an action to occur when the policy condition is fulfilled by clicking one of the following buttons (more than one action can be assigned to a policy):

- ☐ **Send E-mail** - Policy Manager sends an e-mail when the condition is fulfilled. Enter a comma-separated list of e-mail addresses, and then click **OK**.
- ☐ **Generate Event** - Policy Manager generates an event of the specified event type. The event appears in Event Manager. Select one of the following event types.

Important: Since the severity level for an element is set by the manufacturer, the meanings of the severity levels vary. It is best to view the description of the event.

- **Unknown** - The severity level is not known.
- **Informational** - An example of an informational event is a progress report event for firmware download operation currently in progress.
- **Warning** - An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
- **Minor** - An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
- **Major** - An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
- **Critical** - An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.

After you select an event level, click **OK**.

- ☐ **Execute a Custom Command** - Policy Manager executes a custom command on the management server when the condition is fulfilled. Type a command that will execute the script in the field. Then, click **OK**. The software assumes you are in the %JBOS4_DIST%\server\appiq\remotescripts directory on the management server when the script is executed. You can use environment variables in your script, such as POLICY_NAME and POLICY_DESCRIPTION. POLICY_NAME provides the policy name and POLICY_DESCRIPTION provides the policy description. See "Software Environment Variables for Scripting" on page 358 for more information. Prefix the command with "start" if the custom command triggers a user interface component, such as a program that uses the user interface (Example: Microsoft Internet Explorer) or a command prompt window.

For example, assume you want the custom command to open a command prompt window and list the contents in the directory. You would prefix the command with start as shown below:

```
start dir
```

13. Click **OK**.

Testing a Utilization Policy

After you create or modify a utilization policy, test it to verify it provides the results you are anticipating. To test a policy, click the **Test** button in the Utilization Policy table. The management server fires a test for all utilization policies associated with that element.

Keep in mind the following:

- If you want to run the **Test** functionality more than once, set the Re-arm period to zero before clicking the **Test** button a second time.
- Policies that are triggered for virtual applications are also triggered for file servers.

Utilization Policies						
Delete ✓						
☐	<u>Name</u>	<u>Description</u>	<u>Active</u>	<u>Test</u>	<u>Edit</u>	<u>Delete</u>
☐	Application Free Space	Free space for each Application	✓	Test		

Figure 16-2: Testing a Utilization Policy

Modifying Policies

This section describes the following:

- Modifying Utilization Policies on page 698
- Modifying Discovery Policies on page 700
- Modifying Provisioning Policies on page 702
- Modifying Policies for Events on page 703

Modifying Utilization Policies

To modify a utilization policy:

1. Access Policy Manager as described in the topic, “Accessing Policy Manager” on page 682.
2. Click the  button corresponding to the policy you want to modify.
3. In the **Name** field, change the name for the policy.
4. In the **Description** field, change the description for the policy.
5. Select or deselect the **Active** option.
When the **Active** option is selected, the policy is active.
6. In the **Re-arm Period** field, change the re-arm period in minutes.
The rearm period is the amount of time after the policy executes before it can execute again. It is useful for limiting the number of times the same actions will execute.

Important: Specify shorter periods for important applications.

7. In the **Condition** drop-down menu, change the conditions of the policy.
8. For trending policies, change the number of days in the **Historic period** field (min=7; max=180) or **Projection period** field (min=1; max=180).
9. To change the action to occur when the policy condition is fulfilled click one of the following buttons (more than one action can be assigned to a policy):
 - ☐ **Send E-mail** - Policy Manager sends an e-mail when the condition is fulfilled. Enter a comma-separated list of e-mail addresses, and then click **OK**.
 - ☐ **Generate Event** - Policy Manager generates an event of the specified event type. The event appears in Event Manager. Select one of the following event types.

Important: Since the severity level for an element is set by the manufacturer, the meanings of the severity levels vary. It is best to view the description of the event.

- **Unknown** - The severity level is not known.
 - **Informational** - An example of an informational event is a progress report event for firmware download operation currently in progress.
 - **Warning** - An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
-

- **Minor** - An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
- **Major** - An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
- **Critical** - An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.

After you select an event level, click **OK**.

- **Execute a Custom Command** - Policy Manager executes a custom command on the management server when the condition is fulfilled. Type a command that will execute the script in the field. Then, click **OK**. The software assumes you are in the `%JBOS4_DIST%\server\appiq\remotescripts` directory on the management server when the script is executed. You can use environment variables in your script, such as `POLICY_NAME` and `POLICY_DESCRIPTION`. `POLICY_NAME` provides the policy name and `POLICY_DESCRIPTION` provides the policy description. See “Software Environment Variables for Scripting” on page 358 for more information. Prefix the command with “start” if the custom command triggers a user interface component, such as a program that uses the user interface (Example: Microsoft Internet Explorer) or a command prompt window.

For example, assume you want the custom command to open a command prompt window and list the contents in the directory. You would prefix the command with start as shown below:

```
start dir
```

10. Click **OK**.
11. To test a policy, click the **Test** button in the Utilization Policy table.
The management server fires a test for all utilization policies associated with that element.




Figure 16-3: Testing a Modified Utilization Policy

Modifying Discovery Policies

To modify a policy for discovery.

1. Access Policy Manager as described in the topic, “Accessing Policy Manager” on page 682.
2. In the Policy Manager tree in the middle pane, expand the node, Infrastructure Policies. Then, click **New Element Discovery**.
3. Click the button corresponding to the policy you want to modify.
4. In the **Name** field, change the name for the policy.
5. In the **Description** field, change the description of the policy.
6. Select or deselect one or more element types.
When a condition is fulfilled on a select element, Policy Manager generates an event, sends an e-mail, or runs a custom command.
7. Select **Fire when event is cleared** if you want the policy to fire when the event is cleared, Otherwise the event fires when the event is received.
8. Change the action to occur when the policy condition is fulfilled by clicking one of the following buttons (more than one action can be assigned to a policy):
 - ☐ **Send E-mail** - Policy Manager sends an e-mail when the condition is fulfilled. Enter a comma-separated list of e-mail addresses, and then click **OK**.
 - ☐ **Generate Event** - Policy Manager generates an event of the specified event type. The event appears in Event Manager. Select one of the following event types.

Important: Since the severity level for an element is set by the manufacturer, the meanings of the severity levels vary. It is best to view the description of the event.

- **Unknown** - The severity level is not known.
- **Informational** - An example of an informational event is a progress report event for firmware download operation currently in progress.
- **Warning** - An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
- **Minor** - An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
- **Major** - An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
- **Critical** - An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.

After you select an event level, click **OK**.

- **Execute a Custom Command** - Policy Manager executes a custom command on the management server when the condition is fulfilled. Type a command that will execute the script in the field. Then, click **OK**. The software assumes you are in the %JBOS4_DIST%\server\appiq\remotescripts directory on the management server when the script is executed. You can use environment variables in your script, such as POLICY_NAME and POLICY_DESCRIPTION. POLICY_NAME provides the policy name and POLICY_DESCRIPTION provides the policy description. See “Software Environment Variables for Scripting” on page 358 for more information. Prefix the command with “start” if the custom command triggers a user interface component, such as a program that uses the user interface (Example: Microsoft Internet Explorer) or a command prompt window.

For example, assume you want the custom command to open a command prompt window and list the contents in the directory. You would prefix the command with start as shown below:

```
start dir
```

9. Click **OK**.

Modifying Provisioning Policies

You can create an infrastructure policy that generates an event, sends an e-mail, or runs a custom command when successful provisioning occurred.

If you plan to use e-mail notification with your policy, first assign an SMTP server from which the management server can send its e-mail notifications. See “Setting Up E-mail Notification” on page 207.

To modify a policy for provisioning:

1. Access Policy Manager as described in the topic, “Accessing Policy Manager” on page 682.
2. In the Policy Manager tree in the middle pane, expand the node, Infrastructure Policies. Then, click **Provisioning**.
3. Click the  button corresponding to the policy you want to modify.
4. In the **Name** field, change the name for the policy.
5. In the **Description** field, change the description for the policy.
6. Select one or more element types.
When a condition is fulfilled on a select element, Policy Manager generates an event, sends an e-mail, or runs a custom command.
7. Select **Fire when event is cleared** if you want the policy to fire when the event is cleared, Otherwise the event fires when the event is received.
8. Change the action to occur when the policy condition is fulfilled by clicking one of the following buttons (more than one action can be assigned to a policy):
 - ☐ **Send E-mail** - Policy Manager sends an e-mail when the condition is fulfilled. Enter a comma-separated list of e-mail addresses, and then click **OK**.
 - ☐ **Generate Event** - Policy Manager generates an event of the specified event type. The event appears in Event Manager. Select one of the following event types.

Important: Since the severity level for an element is set by the manufacturer, the meanings of the severity levels vary. It is best to view the description of the event.

- **Unknown** - The severity level is not known.
-

- **Informational** - An example of an informational event is a progress report event for firmware download operation currently in progress.
- **Warning** - An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
- **Minor** - An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
- **Major** - An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
- **Critical** - An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.

After you select an event level, click **OK**.

- **Execute a Custom Command** - Policy Manager executes a custom command on the management server when the condition is fulfilled. Type a command that will execute the script in the field. Then, click **OK**. The software assumes you are in the %JBOSS4_DIST%\server\appiq\remotescripts directory on the management server when the script is executed. You can use environment variables in your script, such as POLICY_NAME and POLICY_DESCRIPTION. POLICY_NAME provides the policy name and POLICY_DESCRIPTION provides the policy description. See “Software Environment Variables for Scripting” on page 358 for more information. Prefix the command with “start” if the custom command triggers a user interface component, such as a program that uses the user interface (Example: Microsoft Internet Explorer) or a command prompt window.

For example, assume you want the custom command to open a command prompt window and list the contents in the directory. You would prefix the command with start as shown below:

```
start dir
```

9. Click **OK**.

Modifying Policies for Events

To modify a policy for events:

1. Access Policy Manager as described in the topic, “Accessing Policy Manager” on page 682.
2. In the Policy Manager tree in the middle pane, expand the node, Infrastructure Policies. Then, click **Events**.

3. Click the  button corresponding to the policy you want to modify.
4. In the **Name** field, change the name for the policy.
5. In the **Description** field, type the description for the policy.
6. In the **Re-arm Period** field, change the re-arm period in minutes.
The rearm period is the amount of time after the policy executes before it can execute again. It is useful for limiting the number of times the same actions will execute.

Important: Specify shorter periods for important applications.

7. Select or deselect one or more element types.
When a condition is fulfilled on a select element, Policy Manager generates an event, sends an e-mail, or runs a custom command.
8. Select **Fire when event is cleared** if you want the policy to fire when the event is cleared, Otherwise the event fires when the event is received.
9. In the **Severity** drop-down menu, specify a comparison operator.
10. To the right of the **Severity** drop-down menu, you can change the severity level:

Important: Since the severity level for an element is set by the manufacture, the meanings of the severity levels vary. It is best to view the description of the event.

- ☐ **Unknown** - The severity level is not known.
 - ☐ **Informational** - An example of an informational event is a progress report event for firmware download operation currently in progress.
 - ☐ **Warning** - An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
 - ☐ **Minor** - An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
 - ☐ **Major** - An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
 - ☐ **Critical** - An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.
11. To change how Policy Manager scans the summary text, select one of the following from the **Summary Text** drop-down menu. Policy Manager scans the summary text and responds according to one of the following actions selected:

- ☐ **Is anything** - Regardless of the contents of the event's summary text, Policy Manager sends an e-mail, generates an event, or runs a custom command.
 - ☐ **Contains** - If the event's summary text contains the specified text, Policy Manager sends an e-mail, generates an event, or runs a custom command. Type the specified text in the field that appears to the right of the **Summary Text** drop-down menu.
 - ☐ **Matches Regular Expression** - If the event's summary text matches the specified expression, Policy Manager sends an e-mail, generates an event, or runs a custom command. Type the specified text in the field that appears to the right of the **Summary Text** drop-down menu.
12. Change the action to occur when the policy condition is fulfilled by clicking one of the following buttons (more than one action can be assigned to a policy):
- ☐ **Send E-mail** - Policy Manager sends an e-mail when the condition is fulfilled. Enter a comma-separated list of e-mail addresses, and then click **OK**.
 - ☐ **Generate Event** - Policy Manager generates an event of the specified event type. The event appears in Event Manager. Select one of the following event types.

Important: Since the severity level for an element is set by the manufacturer, the meanings of the severity levels vary. It is best to view the description of the event.

- **Unknown** - The severity level is not known.
- **Informational** - An example of an informational event is a progress report event for firmware download operation currently in progress.
- **Warning** - An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
- **Minor** - An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
- **Major** - An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
- **Critical** - An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.

After you select an event level, click **OK**.

- ☐ **Execute a Custom Command** - Policy Manager executes a custom command on the management server when the condition is fulfilled. Type a command that will execute the script in the field. Then, click **OK**. The software assumes you are in the `%JBOSS4_DIST%\server\appiq\remotescripts` directory on the management
-

server when the script is executed. You can use environment variables in your script, such as `POLICY_NAME` and `POLICY_DESCRIPTION`. `POLICY_NAME` provides the policy name and `POLICY_DESCRIPTION` provides the policy description. See “Software Environment Variables for Scripting” on page 358 for more information. Prefix the command with “start” if the custom command triggers a user interface component, such as a program that uses the user interface (Example: Microsoft Internet Explorer) or a command prompt window.

For example, assume you want the custom command to open a command prompt window and list the contents in the directory. You would prefix the command with start as shown below:

```
start dir
```

13. Click **OK**.

Viewing Policies

To view policies, access Policy Manager as described in the topic, “Accessing Policy Manager” on page 682.

In the right pane, the policies are listed in a table.

Table 16-2: Policy Table Description

Heading	Description
Name	The name of the policy.
Description	A description of the policy.
Active	If there is a check mark in this column, the policy is active.
Edit	Click the  button to edit a policy, for example to change its Active status.
Delete	Click the  button to remove a policy.

Deactivating a Policy

Policies are activated when they are created. You can deactivate a policy, but still keep it stored in the management server. For example, assume you created a policy that sends an e-mail whenever an event of type minor is generated for a server. You might want to deactivate this policy before you upgrade the server.

To deactivate a policy:

1. Access Policy Manager as described in the topic, “Accessing Policy Manager” on page 682.
2. Click the  button corresponding to the policy you want to modify.
3. Deselect the **Active** option.
4. Click the **Finish** button.
The policy is deactivated.

Deleting Policies

To delete a policy:

1. Access Policy Manager as described in the topic, “Accessing Policy Manager” on page 682.
 2. In the Policy Manager tree in the middle pane, click the element or infrastructure to view its policies.
 3. Do one of the following:
 - ☐ **Delete a policy** - If you want to delete just one policy, click the  button corresponding to the policy you want to delete.
 - ☐ **Delete several policies at once** - If you want to delete several policies at once, select the check boxes next to the policies you want to delete. To select all of the policies, select the check box next to the **Name** heading. Then, click the **Delete** button.The policies are deleted.
-

Providing E-mail Notification for a Policy

You can configure Policy Manager to provide e-mail notification when an element is being overused or when any of the following occurs:

- A new element is discovered
- Successful provisioning occurred
- An event occurred on one or more specified elements

Keep in mind the following:

- First assign an SMTP server from which the management server can send its e-mail notifications. See “Setting Up E-mail Notification” on page 207 for more information.
- The following instructions assume you have already created a policy. If you have not created a policy, see the following topics:
 - “Creating a Utilization Policy” on page 683
 - “Creating Policies for Discovery” on page 690
 - “Creating Policies for Provisioning” on page 692
 - “Creating Policies for Events” on page 694

To set up e-mail notification for a policy:

1. Access Policy Manager as described in the topic, “Accessing Policy Manager” on page 682.
2. Click the  button corresponding to the policy you want to modify.
3. Click the **Send E-mail** button.

Policy Manager sends an e-mail when the condition is fulfilled. The software verifies the address entered has a correct form.
4. Enter a comma-separated list of e-mail addresses, and then click **OK**.
5. Click the **Finish** button.

Providing Event Generation for a Policy

You can configure Policy Manager to generate an event when an element is being overused or when any of the following occurs:

- A new element is discovered
- Successful provisioning occurred
- An event occurred on one or more specified elements

The following instructions assume you have already created a policy. If you have not created a policy, see the following topics:

- “Creating a Utilization Policy” on page 683
- “Creating Policies for Discovery” on page 690
- “Creating Policies for Provisioning” on page 692
- “Creating Policies for Events” on page 694

To set up event generation for a policy:

1. Access Policy Manager as described in the topic, “Accessing Policy Manager” on page 682.
2. Click the  button corresponding to the policy you want to modify.
3. Click the **Generate Event** button.
Policy Manager generates an event of the specified event type. The event appears in Event Manager. Select one of the following event types.

Important: Since the severity level for an element is set by the manufacturer, the meanings of the severity levels vary. It is best to view the description of the event.

- ☐ **Unknown** - The severity level is not known.
 - ☐ **Informational** - An example of an informational event is a progress report event for firmware download operation currently in progress.
 - ☐ **Warning** - An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
 - ☐ **Minor** - An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
 - ☐ **Major** - An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
 - ☐ **Critical** - An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.
4. After you select an event level, click **OK**.
-

5. Click the **Finish** button.

Providing a Custom Command for a Policy

You can configure Policy Manager to run a custom command on the management server when an element is being overused or when any of the following occurs:

- A new element is discovered
- Successful provisioning occurred
- An event occurred on one or more specified elements

The following instructions assume you have already created a policy. If you have not created a policy, see the following topics:

- “Creating a Utilization Policy” on page 683
- “Creating Policies for Discovery” on page 690
- “Creating Policies for Provisioning” on page 692
- “Creating Policies for Events” on page 694

To set up a custom script for a policy:

1. Access Policy Manager as described in “Accessing Policy Manager” on page 682.
 2. Click the  button corresponding to the policy you want to modify.
 3. Click the **Execute Custom Command** button.
 4. Type a command that will execute the script in the field. Then, click **OK**. The software assumes you are in the %JBOSS4_DIST%\server\appiq\remotescripts directory on the management server when the script is executed. You can use environment variables in your script, such as `POLICY_NAME` and `POLICY_DESCRIPTION`. `POLICY_NAME` provides the policy name and `POLICY_DESCRIPTION` provides the policy description. See “Software Environment Variables for Scripting” on page 358 for more information.
Prefix the command with “start” if the custom command triggers a user component, such as a program that uses the user interface (Example: Microsoft Internet Explorer) or a command prompt window.
-

For example, assume you want the custom command to open a command prompt window and list the contents in the directory. You would prefix the command with `start` as shown below:

```
start dir
```

5. Click the **Finish** button.
Policy Manager executes a remote script on the management server when the condition is fulfilled.

Important: Depending on your license, Chargeback may not be available. See the “List of Features” to determine if you have access to Chargeback. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center**).

This chapter describes the following:

- About Chargeback on page 714
 - Setting Up Chargeback on page 716
 - Accessing Chargeback on page 716
 - Creating an Asset Record on page 717
 - Changing the Status of an Element on page 718
 - Saving Chargeback Information on page 719
 - Viewing Assets on page 720
 - Defining Storage Tiers on page 721
 - Adding Asset Information on page 724
 - Managing Departments on page 730
 - Setting the Infrastructure Cost on page 731
 - Setting Up Asset and Storage Based Chargeback on page 733
-

- Viewing Chargeback on page 752
- Chargeback Reports on page 756
- Filtering Assets on page 765

About Chargeback

Chargeback helps you manage departmental ownership, track cost and assemble business reports making inquiries, such as audits and inventory reviews, easier.

The management server provides the following types of chargeback.

- **Asset-based** - Asset-based chargeback calculates chargeback based on the departmental ownership percentages and the depreciated value of the assets. Each piece of equipment is owned by a department or a set of departments. Each department has a percentage ownership of the equipment.
- **Storage-based** - Storage-based chargeback calculates charges based on the actual amount of storage used by an application, the type of storage it is using and the ownership percentage assigned to each department. The chargeback number is further refined by an additional fixed infrastructure tax on a per department basis.

After you add information about all of your assets, back up the database by using the Database Admin Utility. Backing up the database saves your chargeback information. If the database fails, your asset information is restored when you restore the database. See “Performing an RMAN Hot Backup” on page 275 for more information.

First set up your chargeback as described in the topic, “Setting Up Chargeback” on page 716. When you are done with adding your chargeback information, you can view chargeback as following:

- **By element** - Displays chargeback for a single element. See “Viewing Chargeback by Element” on page 752 for more information.
- **By department** - Displays chargeback for a department. See “Viewing Chargeback by Department” on page 753 for more information.
- **By owner** - Displays chargeback for by owner. See “Viewing Chargeback by Owner” on page 754 for more information.

Chargeback also provides the following reports. See “Viewing Chargeback Reports” on page 756 for more information.

- **Array Based Assets** - Displays asset information from storage arrays: host name, department, HBA port, HBA port WWN, storage volume, volume size, and cost.
- **Asset Based** - Displays the following asset-based chargeback for each department owning elements: department, asset name, ownership ratio, and chargeback amount. Total asset-based cost per month is also displayed.
- **Storage Based** - Displays storage-based chargeback for each department owning applications: department, application, ownership ratio, total capacity, and chargeback amount. Total storage-based cost per month is also displayed.
- **Storage System Based by Tier** - Displays storage-based chargeback by tier.

After viewing the reports, you can e-mail them to co-workers. You can even set up a schedule for a chargeback report to be mailed at regular intervals to a co-worker. See the topics, “E-mailing a Chargeback Report” on page 758 and “Adding an E-mail Schedule for a Chargeback Report” on page 759 for more information.

Chargeback helps you track of the status of your elements. Elements that have recently been discovered are automatically given the status of New. You can then change the status of elements to in use, missing or repaired. Since the management server cannot determine what you plan to do with an element, you must change the status manually. However, you can easily change the status of a group of elements at once. See “Changing the Status of an Element” on page 718 for more information.

Chargeback also provides a brief listing of your assets by name, status, serial number, vendor/model, and description. You can filter elements by status and/or type for easy navigation. You can even create custom filters. See “About Filtering Assets” on page 766.

Setting Up Chargeback

To be able to view chargeback, you must first complete the steps in the following table.

Table 17-1: Setting Up Chargeback

Step	Description	Where to Find More Info
1	Create an asset record if it does not exist. You can use Chargeback for applications and hardware the management server does not detect. Just create an asset record for that element.	“Creating an Asset Record” on page 717
2	Set the status of the asset.	“Changing the Status of an Element” on page 718
3	(Optional) Add asset information for asset management.	“Adding Asset Information” on page 724
4	Add departments.	“Adding Departments” on page 730
5	Set up Chargeback.	■ “Setting Up Asset-Based Chargeback” on page 733 ■ “Setting Up Storage-Based Chargeback” on page 739

Accessing Chargeback

To access Chargeback, click **Chargeback** ().

Creating an Asset Record

Important: Only a user belonging to a role that has System Configuration selected on the Edit Role page (such as the Domain Administrator role) is allowed to create a record.

You can use Chargeback for any application or hardware, even those the management server does not detect. Just create an asset record for that application or hardware the management server does not monitor. Then, follow the steps for setting up Chargeback, as described in the topic, “Setting Up Chargeback” on page 716.

After you create a record, the element you created the record for is treated as a discovered element. A discovered element is an element that has been detected by the management server, but the management server cannot obtain detailed information about the element. If you create a record for an application, that application is treated as a virtual application.

Important: You can easily remove an element's record by clicking the  button. When you remove an element's record, the management server no longer monitors that element. See “Deleting Elements from the Product” on page 92 for more information.

To create a record:

1. Access Chargeback, as described in “Accessing Chargeback” on page 716.
 2. Click the **New** button.
 3. In the **Add Asset Record** window, type the following:
 - ☐ Name
 - ☐ Vendor
 - ☐ Model
 4. Select one of the following for the type of element:
 - ☐ **Hardware** - Host
 - ☐ **Hardware** - Storage System
 - ☐ **Hardware** - Switch
 - ☐ **Software** - Application
 5. Click **OK**.
-

Changing the Status of an Element

Chargeback helps you track of the status of your elements. Elements that have recently been discovered are automatically given the status of New. You can then change the status of elements to in use, missing or repaired. Since the management server cannot determine what you plan to do with an element, you must change the status manually. However, you can easily change the status of a group of elements at once.

Important: Once you are done with change the status of your elements, save your settings. See “Saving Chargeback Information” on page 719.

To change the status of an element:

1. Access Chargeback, as described in “Accessing Chargeback” on page 716.
 2. Under the Status column, select the status of the element:
 - ☐ **New (Default)** - The status of the element has not been set yet.
 - ☐ **Missing** - The element cannot be found. It may have been taken off line.
 - ☐ **Repaired** - The element is repaired.
 - ☐ **In Use** - The element is running.

The status settings are set manually. For example, if the status of an element changes from being “In Use” to being “Repaired,” you must change this status manually. Refer to the Topology and Event Manager for the latest status of an element.
 3. To change the status of multiple elements at once:
 - a. Select the elements you want to modify.
 - b. Click the **Set Status** button.
 - c. Select the new status for the elements you selected from the Asset Status drop-down menu.
 - d. Click **OK**.
-



Figure 17-1: Selecting an Element

Saving Chargeback Information

After you change the status of your elements, save your settings by clicking the **Save Listing to File** link. The following information is saved as comma-separated values, which can be viewed using a text editor, such as Notepad, or a spreadsheet program, such as Microsoft Excel.

- ID
- Name
- Status
- Category
- Serial Number
- Vendor
- Model

Viewing Assets

To obtain asset information about an element:

1. Access Chargeback, as described in “Accessing Chargeback” on page 716. The following is displayed in the right pane:
 - ❑ **Name** - The name of the element.
 - ❑ **Status** - The state of the element. An element is automatically assigned the status of “new” when it is first discovered. You can change the status of an element to in use, repaired or missing. See “Changing the Status of an Element” on page 718 for more information.
 - ❑ **Chargeback** - Click the  icon to view chargeback for an element. You must first set up Chargeback before you can view its calculations. See “Setting Up Chargeback” on page 716 for more information.
 - ❑ **Vendor/Model** - The vendor and/or model of an element.
 - ❑ **Type** - The type of element, such as an application. See Table 17-2, “Element Type Icons,” on page 721 for an explanation of each icon.
 - ❑ **Serial Number** - The serial number of the element.
 - ❑ **Description** - Description of the element.
-

Table 17-2: Element Type Icons

Graphic	Element Type
	Application
	Host
	Switch
	Storage System

2. To remove an asset record, click the  button corresponding to the record you want to remove.

Defining Storage Tiers

Storage-based chargeback lets you charge the application owners based on the amount of storage allocated to them. Each storage system is designated a storage tier classification. You can specify a charge for each storage tier. This charge is referred to as storage tier cost. The management server determines the storage cost of the application by multiplying the storage tier cost by the allotted storage. If the application uses more than one storage system, the storage cost from each storage system is added for the total storage cost.

Note: All of the elements in an application path (from the application to the volume) must belong to the user's organization in order for the cost to be calculated.

Storage tiers for storage-based information can have any name. The following default storage tier names are provided:

- **Ultra High Availability** is usually assigned to the ultra high availability storage. This tier contains the premium storage in your network, usually the most expensive.

- **High Availability** is usually assigned to the high availability storage. This tier contains storage that is not as expensive as the storage assigned to the Ultra High Availability tier.
- **IDE Based Storage** is usually assigned to IDE-based storage. This tier contains storage that is comparatively inexpensive.

Creating a New Storage Tier

You can create you own storage tiers (up to maximum of 64).

Follow these steps to create a new storage tier:

1. Access Chargeback, as described in “Accessing Chargeback” on page 716.
2. Click the **Ownership** tab in the right pane.
3. Click the **Chargeback Method** drop-down menu and select **Storage-based**.
4. Click the **Set Storage Tier Cost** button. The **Storage Tiers** window displays.
5. Click the **Create New Storage Tier** button.
6. Enter the necessary information in the **Name**, **Monthly Cost/GB** and **Description** fields.
7. Click the **OK** button. The information you entered is saved and you are returned to the **Storage Tiers** window.

Adding Elements to a Storage Tier

It is possible to add any combination of storage systems, volumes, and pools to a storage tier. A storage element can only be assigned to one tier.

Follow these steps to add elements to a storage tier:

1. Access Chargeback, as described in “Accessing Chargeback” on page 716.
 2. Click the **Ownership** tab in the right pane.
 3. Click the **Chargeback Method** drop-down menu and select **Storage-based**.
 4. Click the **Set Storage Tier Cost** button. The **Storage Tiers** window displays.
 5. Click the **Add/Remove Elements** button for the storage tier you want to add elements to. The **Add or Remove Storage Elements from Tier** window displays.
 6. Click the **Storage Systems** tab and select the storage systems you want to add to the tier.
 7. Click the **Storage Volumes** tab and select a storage system from the **Showing Volumes for Storage System** drop-down menu.
-

8. If you want to filter the list of volumes for a storage system, click the **Show Volume Filter** link, select the appropriate filter criteria, and click the **Filter** button.
9. Select the storage volumes you want to add to the storage tier. Click the **+Ports** link in the **Ports** column if you would like to see a list of the ports associated with a particular volume.
10. Click the **Storage Pools** tab and select a storage system from the **Showing Pools for Storage System** drop-down menu.
11. Select the storage pools you want to add to the storage tier. Click the **+Volumes** link in the **Volumes** column if you would like to see a list of the volumes associated with a particular pool.
12. When you are finished selecting the storage systems, volumes, and pools you want to add to the storage tier, click the **Add Selected Elements to Storage Tier** button. The selected elements are added to the **Assets for Tier** section.
13. Click the **OK** button. You will be returned to the **Storage Tiers** window.

Removing Elements from a Storage Tier

Follow these steps to remove elements from a storage tier:

1. Access the **Add or Remove Storage Elements from Tier** window, as described in “Adding Elements to a Storage Tier” on page 722.
2. In the **Assets for Tier** section, select the elements you would like to remove from the storage tier.
3. Click the **Remove Selected Elements from Storage Tier** button. The selected elements are removed from the storage tier.
4. Click the **OK** button. You will be returned to the **Storage Tiers** window.

Editing a Storage Tier

Follow these steps to edit a storage tier:

1. Access the **Storage Tiers** window, as described in “Adding Elements to a Storage Tier” on page 722.
2. Click the **Edit** () button for the storage tier you want to edit.
3. Update the information for the storage tier.
4. Click the **OK** button. You will be returned to the **Storage Tiers** window.

Deleting a Storage Tier

Follow these steps to edit a storage tier:

1. Access the **Storage Tiers** window, as described in “Adding Elements to a Storage Tier” on page 722.
2. Click the **Delete** button () for the storage tier you want to delete. The storage tier is removed from the system.

Note: You can't delete a storage tier if it has one or more assets assigned to it.

3. Click the **OK** button. You will be returned to the **Storage Tiers** window.

Adding Asset Information

This section describes the following:

- Adding Asset Information on page 724
- Adding General Information on page 727
- Adding Staff Information on page 728
- Adding Geographic Information on page 729
- Adding Licensing and Warranty Information on page 729
- Adding Custom Information on page 729

Adding Asset Information

Chargeback provides a handy way for you to keep track of your asset information for an element. You can easily store warranty and licensing information, as well as contact information for the element. For example, assume a switch on the network is having some problems, and you want to contact the person in charge of that switch. You can use the element's asset record to not only find the contact information for that switch, but also the location of that switch.

Important: After you add information about all of your assets, back up the database. Backing up the database saves your chargeback information. If the database fails, your asset information is restored when you restore the database. See “Performing an RMAN Hot Backup” on page 275 for more information.

You can view and add asset information by doing the following:

1. Access Chargeback, as described in “Accessing Chargeback” on page 716.
2. Do one of the following:
 - ☐ Click the  icon corresponding to the element.
 - ☐ Click the link for the element in the right pane, as shown in the following figure.




Figure 17-2: Accessing an Element's Asset Information

3. To access the different types of asset information, click the Asset Record node or the nodes under it, as shown in the following figure. To view general information about an element, click the Asset Record node in the tree. To view specific asset information such as ownership, click the Staff node.




Figure 17-3: Asset Record Node

You can also access the tree shown in this figure from Application Explorer and System Explorer:

- ☐ To access the tree from Application Explorer, click the name of an application in the Application Explorer tree. In the right pane, click the **Asset Management** tab.
- ☐ To access the tree from System Explorer, double-click an element in the topology. In the right pane, click the **Asset Management** tab.

Note: The nodes under Chargeback in the tree are for creating and viewing reports. The Storage-based node is available only for applications. See “Chargeback Reports” on page 756 for more information.

Adding General Information

The management server provides a page for you to enter the following general information about an element. When you are done with adding information on this page, click the **Save Changes** button at the bottom of the page. To learn more about Chargeback, see “About Chargeback” on page 714. and “Setting Up Chargeback” on page 716.

Note: This page enforces the maximum number of characters you can enter in a field. When you can no longer add additional characters to a field, you have reached the maximum number of characters that can be entered for that field.

- **Custom Name** - A name you assigned to the element. See “Assigning a Custom Name” on page 377 for more information.
 - **Date Created** - Date the element was discovered.
 - **Date Last Modified** - Date the record was last modified.
 - **Description** - A description of the element.
 - **Status** - The current status of the element. If the status of the element has changed, select the new status from the Status drop-down menu.
 - ☐ **New** - Default category for all detected elements.
 - ☐ **Missing** - No longer detectable through discovery
-

- ☐ **Repaired** - The element is being repaired. The software does not automatically select this status.
- ☐ **In Use** - The element is in use.
- **Vendor** - The vendor for the element.
- **Model** - The model of the element.
- **Serial Number** - Serial number of the element.
- **Barcode Number** - The barcode on the device.
- **Asset Code** - The asset code assigned to the element.
- **Asset Type** - The asset type assigned to the element.
- **Asset Tag** - The asset tag assigned to the element.
- **Asset Category** - The asset category assigned to the element.
- **Geographic Location** - The location of the element, for example, Boston, Massachusetts.
- **(Storage Systems Only) Storage Tier Classification** - Click the **Set Storage Tier Cost** link to set up storage tiers. See “Defining Storage Tiers” on page 721 for more information.

Adding Staff Information

This page provides contact information about the element.

Important: Keep this information up to date. Other users will use this information to contact you about the element, for example, if it is having problems.

- **Administrator** - The person or department that maintains the element.
 - **Staff Name** - The name of the person who maintains the element.
 - **Staff Phone Number** - A phone number for the person who maintains the element.
 - **Staff Department** - The department that maintains the element.
 - **Staff E-Mail** - An e-mail address of the person who maintains the element.
 - **Staff #2 Name** - The name of an additional person who maintains the element.
 - **Staff #2 Phone Number** - A phone number for an additional person who maintains the element.
 - **Staff #2 Department** - An additional department that maintains the element.
 - **Staff #2 E-Mail** - An e-mail address of an additional person who maintains the element.
-

Adding Geographic Information

Use this page to add geographic information about the element. This page is helpful in keeping track where all of your elements are located, especially if you have more than 100 elements. For example, assume you are told one of your servers is having problems and you need physical access to it. You can use this page to find where the server is located.

- **Rack Number** - The number of the rack that holds the element.
- **Floor** - The floor on which the element is located, for example third floor.
- **Data Center** - The name of the data center where the element is located.
- **Address** - The address where the element is located.
- **City** - The city where the element is located, for example, Boston.
- **Region** - The region where the element is located, for example, New England.
- **Country** - The country where the element is located, for example, the United States.
- **Continent** - The continent where the element is located, for example, North America.
- **Zip Code** - The zip code for the town where the element is located. For example, if the element is located in Burlington, Massachusetts, the zip code would be 01803. If your country does not use zip codes, you can leave this field blank.

Adding Licensing and Warranty Information

Use this page to provide licensing and warranty information.

- **License (maximum of 4000 characters)** - The license of the element.
- **Warranty Information (maximum of 4096 characters)** - Information about the warranty. In this field, you probably want to enter information such as how long the warranty lasts and what it covers.
- **Comments (maximum of 4000 characters)** - Any financial information you might want to add about the element

Adding Custom Information

You can provide up to six custom properties.

1. In the **Name** field, assign a name for the field. The name cannot be more than 50 characters. For example, Backup Contact.

2. In the **Value** field, provide the information for the field. Do not provide more than 255 characters. For example, Joe Smith.
3. Repeat the previous step for each custom property you want to add.
4. When you are done, click **Save Changes**.

Managing Departments

This section describes the following:

- Adding Departments on page 730
- Editing a Department on page 731
- Removing a Department from Chargeback on page 731

Adding Departments

Before you can assign a department to an element, you must add it to the list of departments, as described in the following steps.

To add a department:

1. Access Chargeback, as described in “Accessing Chargeback” on page 716.
 2. Click the **Departments** tab above the table.
 3. Click the **New** button.
 4. In the **Add Department** window, provide the following information:
 - ☐ Department Name (Required)
 - ☐ Department Number (Required)
 - ☐ E-mail
 - ☐ Phone
 5. Click **OK**.
The new department is added.
-

Editing a Department

To edit a department:

1. Access Chargeback, as described in “Accessing Chargeback” on page 716.
2. Click the **Departments** tab above the table.
3. Click the  button corresponding with the department you want to edit.
4. In the **Edit Department** window, you can edit all fields except the department number.
5. Click **OK**.

Removing a Department from Chargeback

Over time, departments in your company may merge and others may be dissolved. To keep up with these changes, you may need to remove obsolete departments from your list. If an element is only assigned to this department, it no longer has an owner. Likewise, if an element is assigned to this department and several others, it continues to be assigned to the other departments.

For example, assume you want to delete a department called TooSmall. The TooSmall department owns 50 percent of a host and the Server department owns 50 percent of the host. When you remove TooSmall, the host is only owned by the Server department by 50 percent.

To remove a department from the list:

1. Access Chargeback, as described in “Accessing Chargeback” on page 716.
2. Click the **Departments** tab above the table.
3. Click the **Delete** button () corresponding to the department you want to remove.

Setting the Infrastructure Cost

For a more detailed analysis of Chargeback information you can optionally specify infrastructure cost. There are two types of monthly infrastructure costs that you can set up: Asset-based and Storage-based. Each can be set up to calculate infrastructure cost based on ownership using the Ownership tab on the Chargeback page. Asset and storage-based infrastructure costs that you specify are calculated as part of your monthly Chargeback results. See About Asset-based

Infrastructure Cost on page 732 and About Storage-based Infrastructure Cost on page 732 for more information on asset and storage-based infrastructure cost. Follow these steps to set up infrastructure costs on the management server:

1. Access **Chargeback**, as described in “Accessing Chargeback” on page 716.
2. Click the **Ownership** tab in the right pane.
3. Click the **Chargeback Method** drop-down menu and select either **Asset-based** or **Storage-based**.
4. Click the **Set Infrastructure Cost** button.
5. Type the monthly infrastructure charge for calculating asset-based or storage-based infrastructure costs depending on your current selection.
6. Click **OK**.

About Asset-based Infrastructure Cost

Asset-based infrastructure cost is an optional chargeback calculation that you can specify on the Ownership page. Setting an asset-based infrastructure cost calculates a monthly infrastructure charge which is identical for each department and is applied once each month on top of their total ownership cost. Modifying the infrastructure charge impacts the asset-based chargeback result for all department owners. Asset-based infrastructure cost, when optionally set up, is added to the total Asset-based chargeback calculation results.

About Storage-based Infrastructure Cost

Storage-based infrastructure cost is an optional chargeback calculation that you can specify on the Ownership page. Setting a storage-based infrastructure cost calculates a monthly infrastructure charge which is identical for each department and is applied once each month on top of their total storage cost. Modifying the infrastructure charge will impact the storage-based chargeback result for all department owners. The same infrastructure charge is assigned to a department regardless of it owning one or 100 elements. Storage-based infrastructure cost, when optionally set up, is added to the total storage-based chargeback calculation results.

For example, the infrastructure cost is not included when you view the chargeback for individual elements; however, the infrastructure cost is added to the Total Cost/Month value in the storage-based Chargeback report. It is also displayed when you view chargeback per department. The infrastructure cost is added to each department, regardless of whether the department owns one or 100 elements.

Setting Up Asset and Storage Based Chargeback

This section describes the following:

- Setting Up Asset-Based Chargeback on page 733
- Setting Up Storage-Based Chargeback on page 739
- Editing Percentage of Ownership on page 743
- Removing Department Ownership of an Element on page 744
- How a Depreciation Method Is Calculated on page 745
- Capacity in Chargeback and Capacity Explorer Differ on page 744

Setting Up Asset-Based Chargeback

Asset-based chargeback calculates chargeback based on the departmental ownership percentages and the depreciated value of the assets. Each piece of equipment is owned by a department or a set of departments. Each department has a percentage ownership of the equipment.

The management server calculates monthly chargeback from the financial information provided. You can then use these monthly calculations to determine the cost impact on your enterprise on a monthly basis. You can even break the cost down by department. If you have a infrastructure cost, you can add that into the calculations as well.

To set up asset-based chargeback:

- Specify Financial information.
- Assign a Departmental Ownership Percentage.
- Review the asset-based chargeback result.


Important: You must have already added your departments, as described in the topic, “Editing a Department” on page 731.

To provide information for asset-based chargeback:

1. Access Chargeback, as described in “Accessing Chargeback” on page 716.
-

2. Do one of the following:

- ❑ Click the  icon corresponding to the element.
- ❑ Click the link for an element in the right pane, as shown in the following figure.



New...		Set Status... ✓	
☐	<u>Name</u>	<u>Status</u>	
☐	<u>babyibm2</u>	In Use ▼	
☐	<u>000183500570 (Symm48:□)</u>	New ▼	

Figure 17-4: Accessing an Element's Asset Information

3. Click the **Asset-based** node under the Chargeback node, as shown in the following figure.




Figure 17-5: Clicking the Asset-based Node

Note: You can also access the tree shown in this figure from Application Explorer and System Explorer.

- ❑ To access the tree from Application Explorer, click the name of an application in the Application Explorer tree. In the right pane, click the **Asset Management** tab.
- ❑ To access the tree from System Explorer, double-click an element in the topology. In the right pane, click the **Asset Management** tab.

Step 1 - Specify Financial information

1. Verify the option, **Step 1 - Specify Financial information**, is selected at the top of the page.
2. Provide the following financial information:
 - ❑ **Purchase Order Number** - The purchase order of the element.
 - ❑ ***Date Purchased** - The date the element was purchased. To select the date, Click the calendar icon, . If you select a future date, the purchase date is set to today when calculating depreciation. The management server only supports dates within the years 1900 and 3000.
 - ❑ **Reseller** - The company that directly sold you the element.
 - ❑ ***Purchase Price** - The price of the element when it was bought.
 - ❑ **Salvage Value** - The amount of money an item is worth for salvage value. You cannot go below this number when depreciating an item.
 - ❑ ***Depreciation Period (months)** - This time period in which you plan to keep the element.
 - ❑ ***Depreciation Method** - Select how the depreciation is calculated:
 - **Straight Line** - The device loses the same amount of value in each period. To learn more about how the management server calculates straight line depreciation, see “Calculating Straight Line Depreciation” on page 746.
 - **Fixed Declining Balance** - The Fixed Declining Balance method calculates depreciation based on the value of the asset each month instead of a fixed rate like straight line depreciation. To learn more about how the management server calculates fixed declining balance, see “Calculating Fixed Declining Balance” on page 747.
 - **Double Declining Balance** - This depreciation method doubles the calculation of the Fixed Declining Balance method. Thus, it doubles the speed at which a device depreciates. To learn more about how the management server calculates double declining balance, see “Calculating Double Declining Balance” on page 749.
 - ❑ **Value as of** - The value of the element as of the end of the previous month. For example, assume you entered and/or viewed this chargeback information in the

middle of March. The value would be for the month of February, but not for March. The value is calculated from the following fields:

- Date Purchased
- Original Cost (\$)
- Depreciation Salvage Value (\$)
- Depreciation Period
- Depreciation Method

*Required

3. Click **Save Changes**.

Step 2 - Assign Departmental Ownership Percentage

Assign the percentage of ownership to an element and the monthly infrastructure charge by doing the following:

1. Select the option, **Step 2 - Assign Departmental Ownership Percentage.**, at the top of the page.
 2. Click the **Add Ownership** button.
 3. Select a department from the **Department** drop-down menu.
 4. Type the percentage of ownership in the **Ownership %** field. If you do not see a department listed, add it to the list as described in the topic, "Adding Departments" on page 730. Click the **Manage Departments** link. After you have added the department, close the window you used to add the department and then refresh the page.
 5. Click **OK**.
The department with its percentage of ownership is added to the table.
 6. If multiple departments own the element repeat the steps above for each department. You can have departments owning more than 100 percent of the element.
 7. (Optional) Specify a monthly infrastructure charge for when asset-based calculation is being done. This charge is identical for each department and is applied once each month on top of their total ownership cost. Modifying the infrastructure charge impacts the asset-based chargeback result for all department owners.
For example, the infrastructure cost is not included when you view the chargeback for individual elements; however, the infrastructure cost is added to the Total Cost/Month value in the Asset-based Chargeback report. It is also displayed when you view chargeback per department. The infrastructure cost is added to each department, regardless of whether the department owns one or 100 elements.
-

- a. Click the **Set Infrastructure Cost** button.
 - b. Type the monthly infrastructure charge for when asset-based calculation is performed.
 - c. Click **OK**.
8. When you are done with assigning the element to a department, click the **Save Changes** button.

Step 3 - Review Asset-based Chargeback Result

Important: The management server displays chargeback information up to the end of the previous month. For example, assume you view chargeback information in the middle of March. The calculations for chargeback would include the month of February, but not March.

To view the result of Asset-based Chargeback:

1. Select the option, **Step 3 - Review Asset-based Chargeback Result.**, at the top of the page.
2. If you see empty values, verify you have provided the required values in the previous steps.

The ownership cost is determined by the following formula:

$$(\text{Depreciation}) \times (\text{Ownership \%}) = \text{Ownership Cost}$$

Ownership Cost is how much owning the element will cost a department. The depreciation is determined by the depreciation method you selected in Step 1 - Specify Financial information.

Important: The infrastructure cost is not included in ownership cost because the information displayed on this page is per asset. The asset-based infrastructure cost is a monthly charge that is applied to each departmental owner in addition to any ownership charges. The infrastructure cost is not included when you view the chargeback for individual elements; however, the infrastructure cost is added to the Total Cost/Month value in the Asset-based Chargeback report. It is also displayed when you view chargeback per department. The infrastructure cost is added to each department, regardless of whether the department owns one or 100 elements.

Oracle Instance cortezi
 Host CORTEZ

☐ Step 1 - Specify financial information.
☐ Step 2 - Assign departmental ownership percentage.
☒ Step 3 - Review asset-based chargeback result.

Asset-based Chargeback result for the month ending: 2004/02/28

Purchase Date: 2003/02/10
 Period Ending: 2004/02/28 Months passed since purchased: 13

Purchase Price: \$2,500.00
 Salvage Value: \$500.00 Depreciable Amount: \$2,000.00

Depreciation Period: 30 months
 Depreciation Method: Straight Line Value as of 2004/02/01: \$1,700.00
 Value as of 2004/02/28: \$1,633.33
 One Month Depreciation: \$66.67

Total Asset-based Chargeback = Sum of Ownership Cost
 (Depreciation) x (Ownership %) = Ownership Cost

<u>Department No.</u>	<u>Department Owner</u>	<u>Depreciation</u>	<u>Ownership %</u>	<u>Ownership Cost</u>
1234	Engineering	\$66.67	100%	\$66.67

Save Changes

Figure 17-6: Ownership Cost

Setting Up Storage-Based Chargeback

Storage-based Chargeback calculates charges based on the actual amount of storage used by an application on the storage system, the type of storage it is using and the ownership percentage assigned to each department. The chargeback number is further refined by an additional fixed infrastructure tax on a per department basis.

To obtain this storage-based chargeback, you must follow these steps:

- Step 1 - Assign Departmental Ownership Percentage.
- Step 2 - Review Storage Tier Cost.
- Step 3 - Review Storage Dependency and Cost.
- Step 4 - Review storage-based chargeback result.

Keep in mind the following:

- You must have already added your departments, as described in the topic, “Adding Departments” on page 730.
- Not all applications use storage on storage systems. Storage-based chargeback is applicable only for those applications that use storage on storage systems.
- You must have access to the storage system the application uses. Verify your organization and roles allow you access. See “About the Security for the Management Server” on page 154 for more information.
- Chargeback provides only network capacity. If you look at the capacity of an application in Capacity Explorer, the capacity differs. Capacity Explorer displays the total capacity of an application, including the network drives. See the topic, “Capacity in Chargeback and Capacity Explorer Differ” on page 744 for more information.

To provide information for storage-based Chargeback:

1. Access Chargeback.
2. Do one of the following:
 - ☐ Click the  icon corresponding to the element.
 - ☐ Click the link for a host running an application in the right pane, as shown in the following figure.




Figure 17-7: Accessing an Element's Asset Information

3. Click the **Storage-based** node under the Chargeback node, as shown in the following figure.




Figure 17-8: Clicking Storage-Based Node

Note: You can also access the tree shown in this figure from Application Explorer and System Explorer.

- To access the tree from Application Explorer, click the name of an application in the Application Explorer tree. In the right pane, click the **Asset Management** tab.

- To access the tree from System Explorer, double-click an element in the topology. In the right pane, click the **Asset Management** tab.

Step 1 - Assign Departmental Ownership Percentage

1. Select the option, **Step 1 - Assign Departmental Ownership Percentage.**, at the top of the page.
2. Click the **Add Ownership** button.
3. Select department from the **Department** combo-box.

Note: If you do not see a department listed, add it to the list as described in the topic, “Adding Departments” on page 730. Click the **Manage Departments** link. After you have added the department, close the window you used to add the department and then refresh the page.

4. Type the percentage of ownership in the **Ownership %** field.
 5. Click **OK**.
The department with its percentage of ownership is added to the table.
 6. If multiple departments own the element repeat the steps above for each department. You can have departments owning more than 100 percent of the element.
 7. (Optional) Specify a monthly infrastructure charge for when storage-based calculation is being done. This charge is identical for each department and is applied once each month on top of their total storage cost. Modifying the infrastructure charge will impact the storage-based chargeback result for all department owners. The same infrastructure charge is assigned to a department regardless of it owning one or 100 elements. For example, the infrastructure cost is not included when you view the chargeback for individual elements; however, the infrastructure cost is added to the Total Cost/Month value in the Storage-based Chargeback report. It is also displayed when you view chargeback per department. The infrastructure cost is added to each department, regardless of whether the department owns one or 100 elements.
 - a. Click the **Set Infrastructure Cost** button.
 - b. Type the monthly infrastructure charge for when storage-based calculation is performed.
 - c. Click **OK**.
 8. When you are done with assigning the element to a department, click the **Save Changes** button.
-

Step 2 - Review Storage Tier Cost

Storage-based chargeback lets you charge the application owners based on the amount of storage allocated to them. Each storage system is designated a storage tier classification. You can specify a charge for each storage tier. This charge is referred to as storage tier cost. The management server determines the storage cost of the application by multiplying the storage tier cost by the allotted storage. If the application uses more than one storage system, the storage cost from each storage system is added for the total storage cost.

Storage tiers for storage-based information can have any name. The following default storage tier names are provided:

- **Ultra High Availability** - usually assigned to the ultra high availability storage. This tier contains the premium storage in your network, usually the most expensive.
- **High Availability** - usually assigned to the high availability storage. This tier contains storage that is not as expensive as the storage assigned to the Ultra High Availability tier.
- **IDE Based Storage** - usually assigned to IDE-based storage. This tier contains storage that is comparatively inexpensive.

Click the **Storage Tiers** link to review your storage tier costs. See “Defining Storage Tiers” on page 721 for more information about setting up storage tiers.

Step 3 - Review Storage Dependency and Cost

Important: The management server displays chargeback information up to the end of the previous month. For example, assume you view chargeback information in the middle of February. The calculations for chargeback would include the month of January, but not February.

Click the **Examine storage/volume dependency and cost details** link to view the details.

The Storage Dependency for Application table is displayed if all of the following conditions apply:

- The application depends on a storage system.
 - The organizations to which you belong allow you to view the storage system.
-

If the table is empty and you know the application is dependent on a storage system, verify you have access to the storage system; otherwise, data cannot be calculated for this report.

The details are provided in a tree table. Expand the various nodes to drill down into the application cost and examine how the storage systems, storage volumes, storage pools, and the assigned tiers are contributing to the total cost.

Step 4 - Review Storage-Based Chargeback Result

Important: The management server displays chargeback information up to the end of the previous month. For example, assume you view chargeback information in the middle of February. The calculations for chargeback would include the month of January, but not February.

The total storage-based monthly chargeback is displayed. This number is calculated as follows:

$$(\text{Total Storage Cost}) \times (\text{Ownership \%}) = \text{Ownership Cost}$$

where

- **Total Storage Cost** is the total Storage Cost from (Step 3 - Review Storage Dependency and Cost).
- **Ownership %** is the percentage of ownership.

Editing Percentage of Ownership

To edit the department ownership of an asset:

1. Access Chargeback.
2. Click the link for the element in the right pane.
3. Click Asset-based under the Chargeback node in the tree.
4. Verify the option, "Step 2 - Assign Departmental Ownership Percentage" is displayed in the right pane.
5. Click the **Edit** () button corresponding to the percentage of ownership you want to modify.
6. In the **Ownership %** field, type a new percentage of ownership.

7. Click **Save Changes**.

Removing Department Ownership of an Element

Sometimes you may need to remove ownership from an element, such in the case of an element being moved from one department to another. When department ownership is removed from an element, the department is still accessible from the list of departments. If you want to make the department inaccessible to all elements, remove from the list of departments as described in the topic, "Removing a Department from Chargeback" on page 731.

To remove ownership:

1. Access Chargeback.
2. Click the link for the element in the right pane.
3. Click Asset-based under the Chargeback node in the tree.
4. Verify the option, "Step 2 - Assign Departmental Ownership Percentage" is displayed in the right pane.
5. Click the Delete () button corresponding to the department you want to remove.
6. Click **Save Changes**.
The department is removed.

Capacity in Chargeback and Capacity Explorer Differ

The capacity displayed for an application in Chargeback differs from the capacity displayed in Capacity Explorer. The management server uses only network storage when calculating chargeback capacity. Local capacity is not counted. The following figure shows the chargeback capacity for an Oracle instance named RETAIL. Noticed the chargeback capacity is .89 GB.

<u>Storage System</u>	<u>Mounted Storage</u>	<u>Unmounted Storage</u>	<u>Total Storage</u>
000183500570 (Symm48:3830)	0.89 GB	0 GB	0.89 GB

Figure 17-9: Chargeback Capacity

If you were to view the Oracle instance RETAIL in Capacity Explorer, you would be shown the local and network capacity, which is a total of 1,042 MB, as shown in the following figure. Of the 1,042 MB, 133 MB is on a local drive. The rest (909 MB) is on a network drive. Convert 909 MB to gigabytes (0.887 GB) and round the output (0.89 GB), the capacity in Capacity Explorer matches the number in Chargeback.

RETAIL	Database Instance Name	Total Capacity
	RETAIL	1,042 MB
	Database Files	Total Capacity
	INDX	58 MB
	RBS	520 MB
	RETAILSPACE	5 MB
	SYSTEM	264 MB
	TEMP	72 MB
	TOOLS	12 MB
	USERS	108 MB
	RedoGroup 1	1 MB
	RedoGroup 2	1 MB
	RedoGroup 3	1 MB

Figure 17-10: Capacity in Capacity Explorer

How a Depreciation Method Is Calculated

This section describes the following:

- Calculating Straight Line Depreciation on page 746

- Calculating Fixed Declining Balance on page 747
- Calculating Double Declining Balance on page 749

Calculating Straight Line Depreciation

When the management server calculates straight line depreciation, it calculates depreciation based on a fixed rate. These instructions describe how the management server performs the straight line depreciation calculation. An example is provided for each step so that you can try the calculations for yourself.

The following is how the management server calculates straight line depreciation:

1. The management server rolls back the purchase date to the beginning of the purchase month. If the purchase date is later than today (for example, a future purchase), then the purchase date is roll back to today.
Example: Assume the purchase date of an element is January 15, 2003. The management server adjusts the purchased date to January 1, 2003 when calculating months to depreciate.
 2. It determines the period ending date. This is equivalent to the last day of the previous full month.
Example: Assume today's date is January 9, 2004. The management server sets the period ending to December 31, 2003.
 3. The management server calculates the delta between purchase date and the period ending. This determines how many months worth of depreciation amount the management server need to take into account.
Example: Using the examples from the previous two steps, the delta is 12 months (January 1, 2003 - December 31, 2003).
 4. It subtracts the salvage value from the purchase price. This is the depreciable amount.
Example: Assume the purchase price for the element is \$2500 and the Salvage Value is \$100. The depreciable amount is \$2400, which was calculated by subtracting the Salvage Value (\$100) from the purchase price (\$2500).
 5. It takes the depreciable amount and divides it by the depreciation period (the number of months it takes for the asset to fully depreciated to either 0 or salvage value). This gives us the depreciation for a single month.
Example: Let's use the depreciable amount (\$2400) calculated in the previous step. Let's assume the depreciation period is 24 months. Divide \$2400 by 24. The result is \$100, which is the one month depreciation.
-

6. It multiplies the depreciation for a single month by the delta from step 3. This is the total depreciation.
Example: To find the total depreciation, multiply the one month depreciation from the previous step (\$100) by the delta (12 months), which was calculated in Step 3. The result of $\$100 \times 12$ months is \$1,200, which is the total depreciation.
7. To determine the value as of end of last month, the management server simply subtract the total depreciation from the purchase price.
Example: Subtract the total depreciation (\$1200), which was calculated in the last step, from the purchase price (\$2500), which was provided in Step 4. The value as of the end of last month is \$1300.

Calculating Fixed Declining Balance

The Fixed Declining Balance method calculates depreciation based on the value of the asset each month instead of a fixed rate like straight line depreciation.

These instructions describe how the management server performs the fixed declining balance calculation. An example is provided for each step so that you can try the calculations for yourself.

The following is how the management server calculates fixed declining balance:

Example: Assume the purchase date of an element is January 15, 2003. The management server adjusts the purchased date to January 1, 2003 when calculating months to depreciate. The management server rolls back the purchase date to the beginning of the purchase month. If the purchase date is later than today (for example, a future purchase), then the purchase date is roll back to today.

1. It determines the period ending date. This is equivalent to the last day of the previous full month.
Example: Assume today's date is January 9, 2004. The management server sets the period ending to December 31, 2003.
2. The management server calculates the delta between purchase date and the period ending. This determines how many months worth of depreciation amount the management server need to take into account.
Example: Using the examples from the previous two steps, the delta is 12 months (January 1, 2003 - December 31, 2003).
3. The management server takes the user-specified depreciation period and use it as the life of the asset.

Example: Let's assume the depreciation period is 24 months and that it is also the life of the asset.

4. The management server calculates the declining ratio using this formula: $(1.0 / \text{life})$. This determines the rate at which depreciation should occur each month.

Example: Use the example from step 3 (24 months) and use it in the following formula to find the rate of depreciate per month:

$$1.0 / 24$$

The depreciation ratio is 0.042.

5. For each month identified by delta from Step 2, the management server calculates the following:

The example for the following steps can be found at the end of these instructions.

- a. Determine the "would-be" depreciation for the month. This means multiplying the asset value for the month by the declining ratio from step 4.
- b. Subtract the depreciation for the month from the asset value for the month. If the result is less than the salvage value, it means the asset value after depreciation would be less than the salvage. In this case, the management server simply depreciate the asset to the salvage value. Once the management server depreciates an asset down to its salvage value, the depreciation for that asset stops.
- c. If the management server subtracts the depreciation for the month from the asset value and the result is greater than the salvage value, then the management server know it is safe to depreciate the asset by the depreciation amount calculated in step a. The depreciated asset value for the month would be asset value minus depreciation. The new asset value will be used to compute the depreciation for next month. This process continues until one of the following occurs:
 - The management server has depreciated the asset value for the number of months equal to delta.
 - The asset value has depreciated down to the salvage cost. If no salvage cost is specified, then the asset value has depreciated down to 0.

Example: For Step 5, let's complete Steps a through c for the first month and then repeat these steps for the second month.

Step 5a - Let's assume the asset value of the element is \$2500. Calculate the "would-be" depreciation of the month by multiplying the asset value by the declining ratio from Step 4 (0.042):

$$\$2500 \times .042 = \$105$$

Step 5b - Assume the salvage value is \$100. Determine if the asset value after depreciation is less than the salvage value by using the following formula.

Asset value of the month (\$2500) - Depreciation for the month (\$105) =
\$2395

Since \$2395 (the depreciated asset value) is greater than the salvage value (\$100), the asset value after depreciation is \$2395. Go to Step 6c.

Step 5c - The new asset value (\$2395) is used to calculate the depreciation for the next month. Let's go through the calculations for the next month.

Step 5a - Assume the asset value of the element is \$2395. Calculate the "would-be" depreciation of the month by multiplying the asset value by the declining ratio from Step 4 (0.042):

$\$2395 \times .042 = \$100.59$

Step 5b - Assume the salvage value is \$100. Determine if the asset value after depreciation is less than the salvage value by using the following formula:

Asset value of the month (\$2395) - Depreciation for the month (\$100.59) =
\$2294.41

Since the \$2294.41 (the depreciated asset value) is greater than the salvage value (\$100), the asset value for the month is \$2294.41. Go to Step 5c. The management server repeats Steps 5a through 5c for 12 months (the delta from Step 2), unless the depreciated asset value reaches the salvage value or 0 if the salvage value is not specified.

Calculating Double Declining Balance

The Double Declining Balance method and the Fixed Declining Balance are very similar. The difference is that instead of using the depreciation ratio determined by $(1.0 / \text{life})$, the management server doubles the ratio to increase the rate of depreciation. This provides for a more realistic depreciation when your asset tends to lose its value in the early part of its life. For instance, a new car's blue book value decreases dramatically once it is sold and driven off the lot of the car dealership.

These instructions describe how the management server performs the double declining balance calculation. An example is provided for each step so that you can try the calculations for yourself:

1. The management server rolls back the purchase date to the beginning of the purchase month. If the purchase date is later than today (for example, a future purchase), then the purchase date is roll back to today.
Example: Assume the purchase date of an element is January 15, 2003. The management server adjusts the purchased date to January 1, 2003 when calculating months to depreciate.

2. It determines the period ending date. This is equivalent to the last day of the previous full month.
Example: Assume today's date is January 9, 2004. The management server sets the period ending to December 31, 2003.
 3. The management server calculates the delta between purchase date and the period ending. This determines how many months worth of depreciation amount the management server need to take into account.
Example: Using the examples from the previous two steps, the delta is 12 months (January 1, 2003 - December 31, 2003).
 4. The management server takes the user-specified depreciation period and use it as the life of the asset.
Example: Let's assume the depreciation period is 24 months and that it is also the life of the asset.
 5. The management server calculates the declining ratio using this formula: $(1.0 / \text{life}) * 2$. This determines the rate at which depreciation should occur each month.
Example: Use the example from step 4 (24 months) and use it in the following formula to find the rate of depreciate per month:
$$(1.0 / 24) * 2$$

The depreciation ratio is 0.084.
 6. For each month identified by delta from Step 3, the management server calculates the following:
The example for the following steps can be found at the end of these instructions.
 - a. Determine the "would-be" depreciation for the month. This means multiplying the asset value for the month by the declining ratio from step 5.
 - b. Subtract the depreciation for the month from the asset value for the month. If the result is less than the salvage value, it means the asset value after depreciation would be less than the salvage. In this case, the management server simply depreciate the asset to the salvage value. Once the management server depreciates an asset down to its salvage value, the depreciation for that asset stops.
 - c. If the management server subtracts the depreciation for the month from the asset value and the result is greater than the salvage value, then the management server know it is safe to depreciate the asset by the depreciation amount calculated in step a. The depreciated asset value for the month would be asset value minus depreciation. The new asset value will be used to compute the depreciation for next month. This process continues until one of the following occurs:
 - The management server has depreciated the asset value for the number of months equal to delta.
-

- The asset value has depreciated down to the salvage cost. If no salvage cost is specified, then the asset value has depreciated down to 0.

Example: For Step 6, let's complete Steps a through c for the first month and then repeat these steps for the second month.

Step 6a - Let's assume the asset value of the element is \$2500. Calculate the "would-be" depreciation of the month by multiplying the asset value by the declining ratio from Step 5 (0.084):

$$\$2500 \times 0.084 = \$210$$

Step 6b - Assume the salvage value is \$100. Determine if the asset value after depreciation is less than the salvage value by using the following formula.

$$\text{Asset value of the month } (\$2500) - \text{Depreciation for the month } (\$210) = \$2290$$

Since \$2290 (the depreciated asset value) is greater than the salvage value (\$100), the asset value after depreciation is \$2290. Go to Step 6c.

Step 6c - The new asset value (\$2290) is used to calculate the depreciation for the next month. Let's go through the calculations for the next month.

Step 6a - Assume the asset value of the element is \$2290. Calculate the "would-be" depreciation of the month by multiplying the asset value by the declining ratio from Step 5 (0.084):

$$\$2290 \times .084 = \$192.36$$

Step 6b - Assume the salvage value is \$100. Determine if the asset value after depreciation is less than the salvage value by using the following formula:

$$\text{Asset value of the month } (\$2290) - \text{Depreciation for the month } (\$192.36) = \$2097.64$$

Since the \$2097.64 (the depreciated asset value) is greater than the salvage value (\$100), the asset value for the month is \$2097.64. Go to Step 6c. The management server repeats Steps 6a through 6c for 12 months (the delta from Step 3), unless the depreciated asset value reaches the salvage value or 0 if the salvage value is not specified.

Viewing Chargeback

This section describes the following:

- Viewing Chargeback by Element on page 752
- Viewing Chargeback by Department on page 753
- Viewing Chargeback by Owner on page 754

Viewing Chargeback by Element

You can view chargeback for an element by clicking the  icon next to the element listed on the Asset tab. If you see empty values, make sure chargeback has been set up, as described in “Setting Up Asset-Based Chargeback” on page 733 and “Setting Up Storage-Based Chargeback” on page 739.

Important: The management server displays chargeback information up to the end of the previous month. For example, assume you view chargeback information in the middle of February. The calculations for chargeback would include the month of January, but not February.

To view chargeback by element:

1. Access Chargeback as described in “Accessing Chargeback” on page 716.
 2. Click the  icon next corresponding to the element you want to view chargeback. Asset-based chargeback is displayed.
 3. (Applications only) To view storage-based chargeback, click **Storage-based** under the Chargeback node to the left of the Asset-based chargeback information.
-




Figure 17-11: Accessing Storage-Based Chargeback

Chargeback information for the element is displayed. If you see missing values, set up chargeback as described in the topics, “Setting Up Asset-Based Chargeback” on page 733 and “Setting Up Storage-Based Chargeback” on page 739.

Viewing Chargeback by Department

You can determine how much a department is being charged for equipment use by viewing chargeback by department. This feature lets you view the monthly costs associated with using hardware and applications. If you see empty values, make sure chargeback has been set up, as described in “Setting Up Asset-Based Chargeback” on page 733 and “Setting Up Storage-Based Chargeback” on page 739.

Important: The management server displays chargeback information up to the end of the previous month. For example, assume you view chargeback information in the middle of February. The calculations for chargeback would include the month of January, but not February.

To view chargeback by department:

1. Access Chargeback as described in “Accessing Chargeback” on page 716.
 2. Click the **Departments** tab in the right pane.
-

3. Click the  icon next corresponding to the element you want to view chargeback. You are shown the following information:
- ☐ **Department Name** - Provided when the department was added.
 - ☐ **Department Number** - Provided when the department was added.
 - ☐ **E-mail** - May be blank if information was not provided.
 - ☐ **Phone** - May be blank if information was not provided.
 - ☐ **Monthly Infrastructure Cost for Asset** - How much it costs to operate the element on a monthly basis.
 - ☐ **Ownership Cost** - How much it costs the department in operating the element
 - ☐ **Total Asset-based Chargeback** - How much it costs the department in operating the element. This number is based on the following formula:

$$(\text{Monthly Infrastructure cost}) + ((\text{Depreciation}) \times (\text{Ownership \%}))$$
 - ☐ **Monthly Infrastructure Cost for Storage-based Chargeback** - How much it costs for an application to use a certain amount of storage.
 - ☐ **Total Storage-based Chargeback** - How much it costs the department for an application to use a certain amount of storage:

$$\text{Monthly Infrastructure Cost} + \text{Ownership Cost}$$
 where Ownership Cost is $(\text{Ownership \%}) \times (\text{Storage Cost})$
- This page also displays two tables:
- ☐ Table for Asset-Based Chargeback lists the asset, depreciation, ownership percentage, and ownership cost.
 - ☐ Table for Storage-Based Chargeback lists the application, storage allotted, storage cost, ownership %, and ownership cost. The storage allotted value includes mounted and unmounted storage. Any volumes the application can access are included in the storage calculations.

Viewing Chargeback by Owner

You can view Chargeback for all elements by using the **Ownership** tab. The **Ownership** tab shows the ownership distribution across different departments and helps you with quickly identifying the assets without a department owner.

To view Chargeback by owner:

1. Access Chargeback.
 2. Click the **Ownership** tab in the right pane.
-

3. Select one of the following from the Chargeback Method drop-down menu:
 - ☐ **Asset-based** - Displays chargeback information for assets.
 - ☐ **Storage-based** - Displays chargeback information for storage (applications only).

The management server displays asset-based or storage-based chargeback information based on your selection. The management server displays chargeback information from the previous month. For example, assume you view chargeback information in the middle of February. The calculations for chargeback would include the month of January, but not February.

If you see empty values, make sure chargeback has been set up, as described in “Setting Up Asset-Based Chargeback” on page 733 and “Setting Up Storage-Based Chargeback” on page 739.

Note: You can sort elements according to a column heading. Just click a column heading in the table to sort the data. The arrow next to column heading displays whether the items are being sorted in ascending or descending order. If the arrow is pointing up, items are sorted in ascending and alphabetical order. If the arrow is pointing down, items are sorted in descending and reserved alphabetical order.

The following information is displayed:

- **Department Owner** - The department that owns the element. This information was provided when the department was added.
- **Department Number** - The number of the department that owns the element. This information was provided when the department was added.
- **Asset Name** - The name of the element.
- **Vendor** - The company that supplied the element.
- **Serial Number** - The serial number of the element.
- **Ownership Percentage** - The percentage the department owns of the element.
- **Storage (GB)**** - The amount in gigabytes the application uses. This value includes mounted and unmounted storage. Any volumes the application can access are included in the storage calculations.
- **Storage Cost**** - (Storage-based only) How much it costs to run the storage that the application uses.
- **Ownership Cost** - How much it costs the department to use the asset.

**Applies to only storage-based chargeback.

Chargeback Reports

This section describes the following:

- Viewing Chargeback Reports on page 756
- E-mailing a Chargeback Report on page 758
- Managing E-mail Schedules for Chargeback Reports on page 759

Viewing Chargeback Reports

You can access Chargeback reports one of two ways:

- From Chargeback
 - From Reporter
-

Keep in mind the following:

- To populate Chargeback reports, enter information for Chargeback, as described in the topic, “Setting Up Chargeback” on page 716.
- If you want to view the latest information in a report, select **Configuration > Reports** and then click the **Refresh Now** button on the Report Cache tab. See “Refreshing the Report Cache” on page 245 for more information.

To view Chargeback reports:

1. Do one of the following:
 - To access the reports through Chargeback, click the **Reports** tab in Chargeback and expand the **Chargeback Reports** node in the tree in the middle pane.
 - To access the reports through Reporter, click the **Reporter** button. Then, expand the **Reporter > Chargeback** nodes in the tree in the middle pane.
2. Click one of the default reports:
 - **Array-Based Chargeback** - Displays asset information from storage arrays: host name, department, HBA port, HBA port WWN, storage volume, volume size, and cost.
 - **Asset Based** - Displays the following asset-based chargeback for each department owning elements: department, asset name, ownership ratio, and chargeback amount. Total asset-based cost per month is also displayed.
 - **Storage Based** - Displays storage-based chargeback for each department owning applications: department, application, ownership ratio, total capacity, and capacity amount. Total storage-based cost per month is also displayed.
 - **Storage System Based by Tier** - Displays storage-based chargeback by tier. The report is displayed as a Web page in the right pane. See the topic, “Chargeback Reports” on page 756 for more information about each of the reports.
3. Select one of the following options for a different output:
 - **PDF**
 - **Excel**
 - **XML**
4. To view the report in a new window, select the **Open in new window** option. Then, click the **Run Report** button.

E-mailing a Chargeback Report

You can e-mail a Chargeback report in PDF, XML or Microsoft Excel format. If you want to e-mail reports by on a regular basis, set up an e-mail schedule for the report, as described in “Adding an E-mail Schedule for a Chargeback Report” on page 759.

Important: Before you can e-mail a report, you must set up e-mail notification, as described in the topic, “Setting Up E-mail Notification” on page 207.

To e-mail a report:

1. Do one of the following:
 - ☐ To access the reports through Reporter, click the **Reporter** button. Then, expand the **Reporter > Chargeback** nodes in the tree in the middle pane.
 - ☐ To access the reports through Chargeback, click the **Reports** tab in Chargeback and expand the **Chargeback Reports** node in the tree in the middle pane.
 2. Expand the tree in the middle pane, and click the report for which you want to send by e-mail.
 3. When the report is displayed in the right pane, click the **E-mail Report** button in the upper-right corner of the right pane.
 4. In the **To** field, type the recipient's e-mail address.
The software verifies the address entered has a correct form. To send multiple addresses, separate each address with a semicolon (;), for example:
`john.example@appiq.com;jerry.example@appiq.com`
 5. From the **Format** drop-down menu, select one of the following formats:
 - ☐ **PDF** - Requires the use of Adobe Acrobat, which can be downloaded for free from <http://www.adobe.com>.
 - ☐ **Excel** - Requires the use of Microsoft Excel.
 - ☐ **XML** - Requires the user has an understanding of XML.
 6. (Optional) Modify the subject and message.
 7. Click the **OK** button.
The report is sent.
-

Managing E-mail Schedules for Chargeback Reports

This section describes the following:

- Adding an E-mail Schedule for a Chargeback Report on page 759
- Editing an E-mail Schedule for a Chargeback Report on page 761
- Deleting E-mail Schedules for a Chargeback Report on page 763
- Viewing E-mail Schedules for a Chargeback Report on page 763
- Viewing the History of an E-mail Chargeback Schedule on page 765

Adding an E-mail Schedule for a Chargeback Report

You can add an e-mail schedule so that a user receives an attached report on a regular basis. The report can be in the form a PDF, XML or Microsoft Excel document.

Keep in mind the following:

- Before you can add an e-mail schedule, you must set up e-mail notification, as described in the topic, “Setting Up E-mail Notification” on page 207.
- The management server service must be running for users to receive e-mail notification.
- Only the e-mail schedules created by the current user are listed. To view the e-mail schedules for all reports, click **Configuration > Reports** in the upper-right corner of the screen. Then, click the **Scheduled Deliveries** tab at the top of the screen.

To add an e-mail schedule:

1. Access Chargeback, as described in “Accessing Chargeback” on page 716.
2. Do one of the following:
 - To access the reports through Chargeback, click the **Reports** tab in Chargeback and expand the **Chargeback Reports** node in the tree in the left pane.
 - To access the reports through Reporter, click the **Reporter** button. Then, expand the **Reporter > Chargeback** nodes in the tree in the middle pane.
3. Click the report for which you want to set up an e-mail schedule.
The report is displayed as a Web page in the right pane. See the topic, “Chargeback Reports” on page 756 for more information about each of the reports.
4. When the report is displayed in the right pane, click the **Scheduled Deliveries** tab in the right pane.
5. Click the **Add E-mail Schedule** button.

6. In the **To** field, type the recipient's e-mail address.
The software verifies the address entered has a correct form. To send multiple addresses, separate each address with a semicolon (;), for example:

`john.example@appiq.com;jerry.example@appiq.com`

7. In the **Subject** field, type a subject for the e-mail messages you plan to send.
-

Note: Provide the name of the report in the subject field so users can distinguish the message from others.

8. In the **Message** field, type a message describing the report.
If you are e-mailing reports in bulk, you might want to let users know the e-mail is being sent by an automated process. You might also want to provide an e-mail address for users to provide feedback, for example:

This e-mail and its attached report are generated automatically. If you would like to change how often the report is sent to you or you want to be taken off the list, please contact `username@companyname.com`.

9. From the **Format** drop-down menu, select one of the following formats:
- ☐ **PDF** - Requires the use of Adobe Acrobat, which can be downloaded for free from <http://www.adobe.com>.
 - ☐ **Excel** - Requires the use of Microsoft Excel.
 - ☐ **XML** - Requires the user has an understanding of XML.
10. In the **Time to Run** field, type the time you want to send the report. This time must be entered in the 24-hour format. For example, if you want a report sent at 2:15 p.m., you would type 14:15 in **Time to Run** field.
11. Select one of the following options to determine how frequently you want to send the report.
- ☐ **Daily** - If you selected daily, select how frequently you want the management server to send the report.
 - **Everyday** - The report is sent everyday.
 - **Weekday** - The report is sent only Monday through Friday.
 - **Everyday for a specified number of days** - Fill in the number of days you want the report to be sent daily. After the specified number of days, the report is no longer sent. For example, you could use this feature to send reports to someone's replacement while the person is away on vacation.
 - ☐ **Weekly** - If you selected weekly, use the **Frequency** drop-down menu to select the day of the week on which you want the report sent.
-

- **Monthly** - If you selected monthly, select the time during the month you want the report sent.
 - To send the report on the first or last day of the month, select the first option. Then, select **First** or **Last** from the drop-down menu.
 - To send the report on a specified day during the month, select the second option. Then, type the day on which you want the report sent. Keep in mind that the number of days in a month vary. So if you type 30 in this field, users will not receive a report in February. Also, if you type 29 in this field, users do not receive the report in February during non-leap years.
12. Click the **OK** button.
The schedule is created.

Editing an E-mail Schedule for a Chargeback Report

Important: Only the e-mail schedules created by the current user are listed. To view the e-mail schedules for all reports, click **Configuration > Reports** in the upper-right corner of the screen. Then, click the **Scheduled Deliveries** tab at the top of the screen.

To edit an e-mail schedule for a report:

1. Access Chargeback, as described in "Accessing Chargeback" on page 716.
2. Do one of the following:
 - To access the reports through Chargeback, click the **Reports** tab in Chargeback and expand the **Chargeback Reports** node in the tree in the middle pane.
 - To access the reports through Reporter, click the **Reporter** button. Then, expand the **Reporter > Chargeback** nodes in the tree in the middle pane.
3. Click the report for which you want to edit an e-mail schedule.
The report is displayed as a Web page in the right pane. See the topic, "Chargeback Reports" on page 756," for more information about each of the reports.
4. When the report is displayed in the right pane, click the **Scheduled Deliveries** tab in the right pane.
5. Under the Edit column, click the **Edit** () button.
6. In the **To** field, change the recipient's e-mail address.
The software verifies the address entered has a correct form. To send multiple addresses, separate each address with a semicolon (;), for example:

john.example@appiq.com;jerry.example@appiq.com

7. In the **Subject** field, change the subject of the e-mail.
8. In the **Message** field, change a message describing the report.
If you are e-mailing reports in bulk, you might want to let users know the e-mail is being sent by an automated process. You might also want to provide an e-mail address for users to provide feedback, for example:

This e-mail and its attached report are generated automatically. If you would like to change how often the report is sent to you or you want to be taken off the list, please contact `username@companyname.com`.

9. From the **Format** drop-down menu, select one of the following formats:
 - ☐ **PDF** - Requires the use of Adobe Acrobat, which can be downloaded for free from <http://www.adobe.com>.
 - ☐ **Excel** - Requires the use of Microsoft Excel.
 - ☐ **XML** - Requires the user has an understanding of XML.
 10. In the **Time to Run** field, type the time you want to send the report. This time must be entered in the 24-hour format. For example, if you want a report sent at 2:15 p.m., you would type 14:15 in **Time to Run** field.
 11. Select one of the following options to determine how frequently you want to send the report.
 - ☐ **Daily** - If you selected daily, select how frequently you want the management server to send the report.
 - **Everyday** - The report is sent everyday.
 - **Weekday** - The report is sent only Monday through Friday.
 - **Everyday for a specified number of days** - Fill in the number of days you want the report to be sent daily. After the specified number of days, the report is no longer sent. For example, you could use this feature to send reports to someone's replacement while the person is away on vacation.
 - ☐ **Weekly** - If you selected weekly, use the **Frequency** drop-down menu to select the day of the week on which you want the report sent.
 - ☐ **Monthly** - If you selected monthly, select the time during the month you want the report sent.
 - To send the report on the first or last day of the month, select the first option. Then, select **First** or **Last** from the drop-down menu.
 - To send the report on a specified day during the month, select the second option. Then, type the day on which you want the report sent. Keep in mind that the number of days in a month vary. So if you type 30 in this field, users will not receive a report in February. Also, if you type 29 in this field, users do not receive the report in February during non-leap years.
-

12. Click the **OK** button.

Deleting E-mail Schedules for a Chargeback Report

Important: Only the e-mail schedules created by the current user are listed. To view the e-mail schedules for all reports, click **Configuration > Reports** in the upper-right corner of the screen. Then, click the **Scheduled Deliveries** tab at the top of the screen.

To delete an e-mail schedule:

1. Access Chargeback, as described in “Accessing Chargeback” on page 716.
2. Do one of the following:
 - ☐ To access the reports through Chargeback, click the **Reports** tab in Chargeback and expand the **Chargeback Reports** node in the tree in the middle pane.
 - ☐ To access the reports through Reporter, click the **Reporter** button. Then, expand the **Reporter > Chargeback** nodes in the tree in the middle pane.
3. Click the report for which you want to delete an e-mail schedule.
The report is displayed as a Web page in the right pane. See the topic, “Chargeback Reports” on page 756 for more information about each of the reports.
4. Click the **Scheduled Deliveries** tab.
5. Click the  button corresponding to the e-mail schedule you want to remove.

Viewing E-mail Schedules for a Chargeback Report

Important: Only the e-mail schedules created by the current user are listed. To view the e-mail schedules for all reports, click **Configuration > Reports** in the upper-right corner of the screen. Then, click the **Scheduled Deliveries** tab at the top of the screen.

To view the E-mail schedules assigned to a report.

1. Access Chargeback as described in “Accessing Chargeback” on page 716.
 2. Do one of the following:
-

- ❑ To access the reports through Chargeback, click the **Reports** tab in Chargeback and expand the **Chargeback Reports** node in the tree in the middle pane.
 - ❑ To access the reports through Reporter, click the **Reporter** button. Then, expand the **Reporter > Chargeback** nodes in the tree in the middle pane.
3. Click the report for which you want to view e-mail schedules.
 4. When the report is displayed in the right pane, click the **Scheduled Deliveries** tab in the right pane.
Information about the e-mail schedules for that report are displayed.

Table 17-3: Viewing E-mail Schedules for a Chargeback Report

Column Name	Description
Recipient	The person who receives the report.
Subject	The subject of the e-mail, brief summary of what it is about.
Format	The format of the report sent: ■ PDF ■ Microsoft EXCEL ■ XML
Last Delivered	The time the last report was sent to the recipient.
History	Click the View button to display the times and dates when the report was sent. You can also delete a historical entry by clicking the  button for the corresponding entry.
Edit	Click the  button to edit a schedule of the report. See the topic, “Adding an E-mail Schedule for a Chargeback Report” on page 759 for information about the options displayed in this window.
Delete	Click the  button to remove the corresponding schedule.

Viewing the History of an E-mail Chargeback Schedule

You can display the times and dates when the report was sent. You can also delete a historical entry by clicking the  button for the corresponding entry.

To view the history of an e-mail schedule:

1. Access Chargeback, as described in “Accessing Chargeback” on page 716.
2. Do one of the following:
 - To access the reports through Chargeback, click the **Reports** tab in Chargeback and expand the **Chargeback Reports** node in the tree in the middle pane.
 - To access the reports through Reporter, click the **Reporter** button. Then, expand the **Reporter > Chargeback** nodes in the tree in the middle pane.
3. Click the report for which you want to view established e-mail schedules.
4. When the report is displayed in the right pane, click the **Scheduled Deliveries** tab in the right pane.
5. Under the History column, click the **View** button.
You are shown when the report was sent.
6. To remove a historical entry, click the  button.

Filtering Assets

This section describes the following:

- About Filtering Assets on page 766
 - Selecting an Element Type for Chargeback on page 766
 - Customizing the Element Type Filter for Chargeback on page 767
 - Filtering Assets by Status on page 767
 - Customizing the Asset Status Filter for Chargeback on page 768
 - Hiding Filters in Chargeback on page 768
-

About Filtering Assets

The management server provides several types of filters to specify which assets you want Chargeback to display.

You can use all the filters at once or you can use just one of them. You can filter assets by:

- **Status**
- **Element type**

For example, assume you need a host to install an application, but you are not too sure which hosts are in use. You could set the filters so that only hosts with a status of "in use" are displayed. You could then click the element to find contact information for the owner.

You could also use the filters to find out which elements are missing or repaired. For example, assume you want to find out which hosts are missing or repaired.

You could do the following:

- Set the filter to display only hosts by selecting the **Host** option from the Show Element Type combo box.
- Click the **Custom** button next to the Show Status combo box. Verify that only Missing and Repaired are selected. Click **OK**.

Once you set all of your filters, you would click the **Apply Filters** button.

Selecting an Element Type for Chargeback

You can filter by element type, so only certain types of elements are displayed. For example, you can specify that only hosts are displayed.

To filter by element type, select an option from the **Show Element Type** drop-down menu in Chargeback. When you are asked if you want to apply your changes, click **Yes**. If you want to apply your changes at a later time you can click **No** and then, click the **Apply Filters** button when you are ready for your changes to take effect. Chargeback displays only the elements you specified in your filter.

Table 17-4: Element Types

Element Type	Description
Applications	Displays only applications, such as Microsoft Exchange and Oracle.
Host	Displays only hosts.
Switch	Displays only switches.
Storage System	Displays only storage systems.
All	Lists all elements.

Customizing the Element Type Filter for Chargeback

You can customize the element type for your filter by clicking the **Customize** button next to the **Show Element Type** drop-down menu in Chargeback.

For example, you can specify you want only hosts and switches displayed in Chargeback.

To select more than one element for filtering:

1. Click the **Custom** button next to the **Show Element Type** drop-down menu in Chargeback.
2. Select the element types you want displayed in Chargeback.
3. Click **OK**.
4. When you are asked if you want to apply your changes, click **Yes**. if you want them to take effect immediately. If not, click **No**. Then, click the **Apply Filters** button when you want your changes to take effect.
Chargeback displays the elements that meet the criteria selected in the filter.

Filtering Assets by Status

You can filter an asset by status, so only certain assets of a specified status are displayed. For example, you can specify that only assets in use are displayed.

To filter by asset status:

1. Select the following options from the **Show Element Type** drop-down menu in Chargeback.
 - ☐ **All (Default)** - All assets are displayed.
 - ☐ **New** - Only assets with the status of New are displayed.
 - ☐ **Missing** - Only assets with the status of Missing are displayed.
 - ☐ **In Use** - Only assets with the status of In Use are displayed.
2. When you are asked if you want to apply your changes, click **Yes**. If you want to apply your changes at a later time you can click **No** and then, click the **Apply Filters** button when you are ready for your changes to take effect. Chargeback displays only the assets you specified in your filter.

Customizing the Asset Status Filter for Chargeback

You can filter multiple assets by clicking the **Customize** button next to the **Show Status** drop-down menu in Chargeback.

For example, you can specify you want only assets that are missing displayed in Chargeback.

To select more than one assets for filtering:

1. Click the **Custom** button next to the **Show Status** drop-down menu in Chargeback.
2. Select the statuses you want displayed in Chargeback.
3. Click **OK**.
4. When you are asked if you want to apply your changes, click **Yes**. if you want them to take effect immediately. If not, click **No**. Then, click the **Apply Filters** button when you want your changes to take effect.
Chargeback displays the assets that meet the criteria selected in the filter.

Hiding Filters in Chargeback

Hide the filters for additional screen space. When you hide the filters, the following features are hidden:

- Show Element Type
 - Show Status
-

To hide the filters, click the **Hide Filters** link in the upper-left corner of Chargeback.

To display the filters, click the **Show Filters** link in the upper-left corner of Chargeback.

Important: Depending on your license, Business Tools may not be available. See the “List of Features” to determine if you have access to Business Tools. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center**).

This chapter describes the following:

- About the Business Tools on page 772
 - Using the HBA Replacement Automator on page 773
 - Using the Brocade Firmware Update Utility on page 774
 - Setting up Risk Analysis on page 777
 - Global Change Management on page 778
-

About the Business Tools

Important: Depending on your license, Business Tools may not be available. See the “List of Features” to determine if you have access to Business Tools. The “List of Features” is accessible from the Documentation Center (**Help** -> **Documentation Center**).

Business Tools provides the following functionality to help you manage the business aspect of your network.

- **Advisors** - Provides detailed information for you to make informed decisions about your network, such as non-compliant HBA firmware. You may have access to the following advisors depending on your release:
 - **Reachable Storage** - Provides information about the storage accessible from a selected host, such as the following:
 - Free volumes on current storage systems
 - LUNs mapped to host, but not mounted with file systems
 - Free volumes on other storage systems in host fabrics
 - Free volumes on all other storage systems
 - **HBA Risk Analysis** - It examines whether the HBAs are at risk. Configure the `hba_risk_analysis.conf` file before you run this tool. See “Setting up Risk Analysis” on page 777 for more information.
 - **Switch Risk Analysis** - It examines whether the switches are at risk. Configure the `switch_risk_analysis.conf` file before you run this tool. See “Setting up Risk Analysis” on page 777 for more information.
 - **Global Change Management** - It lets you record the properties and connections of managed elements. See “Global Change Management” on page 778.
- **Automators** - Lets you automate particular tasks. You may have access to the following automators depending on your release:
 - **Brocade Firmware Update Utility** - Lets you update Brocade firmware remotely.
 - **HBA Replacement Automator** - Makes the management server aware of a replaced HBA card so that the latest information for zoning and LUN masking is still available. See “Using the HBA Replacement Automator” on page 773 before you use this tool.

Keep in mind the following:

- After you use an automator, refresh the element or perform Get Details so the management server can obtain new information and update the user interface accordingly.
-

Only Get Details removes elements that are no longer there from the user interface. For example, removed ports could appear in Port Details but not in the topology for the host if you refresh the element, instead of Get Details.

- All discovered elements are accessible in Business Tools, regardless of a user's restrictions. For example, assume your account belongs to an organization that has only hosts as members. If you run Switch Risk Analysis, the management server still provides information about whether the switches are a risk in your environment.

Using the HBA Replacement Automator

The HBA Replacement Automator makes the management server aware of a replaced HBA card so that the latest information for zoning and LUN masking is still available. You must complete all steps to prevent LUN masking from being lost after running the HBA Replacement Automator.

Important: After installing the new HBA, you must run Get Details before you run the HBA Replacement Automator script; otherwise, you cannot change zone and zone aliases to accommodate the new Worldwide Name (WWN).

You have two ways to replace the HBA. You can install the new HBA with the old HBA or you can install the new HBA by itself.

Installing New HBA with Old HBA

To install the new HBA with the old HBA:

1. Install the new HBA with the old HBA.
 2. Run Get Details.
 3. Run the HBA Replacement Automator (**Business Tools > HBA Replacement Automator**).
 4. When you run the Automator, select the WWNs of the new and old HBA.
-

Installing New HBA by Itself

To install the new HBA by itself:

1. Write down the WWN of the HBA you want to replace. You can find the WWN name of the HBA by doing the following:
 - a. Double-click the host of the HBA in System Explorer.
 - b. In the Navigation pane, click the **Host Bus Adapters** button at the top of the pane. The host bus adapters for the host are listed and their WWNs. Write down the WWN.
2. Install the new HBA by itself.
3. Run Get Details.
4. Run the HBA Replacement Automator (**Business Tools > HBA Replacement Automator**).
5. When you run the Automator, enter the WWN of the old HBA.

Using the Brocade Firmware Update Utility

You can update the firmware of a Brocade switch remotely by using the Brocade Firmware Update Utility on the management server. The management server must obtain some information about the switch and the location of the firmware before it can upgrade the firmware on the switch. This information can be entered in a configuration file or it can be entered when you run the Brocade Firmware Update Utility.

The directory containing the new firmware must be:

- in a shared network folder
 - on an server with FTP capabilities
-

Updating the Firmware Without Modifying the Configuration File

To update the firmware without modifying the configuration file:

1. Click **Business Tools** ().
2. Click **Brocade Firmware Update** in the right pane under Automators.
3. Select the switch you want to upgrade, and then click **OK**.
If you have not provided in the configuration file the user name and password of the switch, you are shown the following message:

```
One or both of the following keys are missing or contain invalid values:  
'switch_user'  
'switch_pass'
```
4. In the **Commands** field, type the user name of the switch. Press ENTER.
5. In the **Commands** field, enter the password of the switch.
If you have not provided in the configuration file the user name and password of the server that has the firmware, you are shown the following message:

```
One or both of the following keys are missing or contain invalid values:  
'ftp_user'  
'ftp_pass'
```
6. In the **Commands** field, enter the user name of the server that has the firmware.
7. In the **Commands** field, enter the password of the server that has the firmware.
If you have not provided in the configuration file the IP address of the server that has the firmware, you are shown the following message:

```
""is not a valid value for key 'ftp_server'
```
8. In the **Commands** field, enter the IP address of the server that has the firmware.
If you have not provided the path to the firmware, you are shown the following message:

```
"" is not a valid value for key 'firmware'.
```
9. Enter the complete path to the firmware. Include the file name in the path, but not the drive letter. Use slashes (/) even if the firmware is on a Microsoft Windows computer, for example: `/myshare/firmware_directory/firmware_filename`.
The directory containing the new firmware must be:
 - ☐ in a shared network folder
 - ☐ on an server with FTP capabilities
10. When you are asked if you want to upgrade, enter `Yes`.

11. After the firmware is loaded, reboot the switch.

Updating the Firmware With Modifying the Configuration File

You can enter the information for the firmware upgrade in the `update_Brocade.conf` file located in `%JBOS4_DIST%\server\appiq\remotescripts\automators`.

1. Open the `update_Brocade.conf` file in a text editor, such as Notepad.
2. Acceptable fields are the following:
 - ☐ `switch_user` - username for switch
 - ☐ `switch_pass` - password for switch username
 - ☐ `ftp_server` - IP address of FTP server with firmware file
 - ☐ `ftp_user` - username for FTP server
 - ☐ `ftp_pass` - password for FTP username
 - ☐ `firmware` - absolute path to firmware file on FTP server
3. Assigned the following values. Verify the number sign (#) is removed from in front of the field. The program will prompt for these fields if they are not specified. All fields must be tab delimited.

```
switch_user admin
switch_pass switchpassword
ftp_server 192.168.1.1
ftp_user root
ftp_pass password
firmware /myshare/firmware_directory/firmware_file
where
```

- ☐ `admin` is the user name for the switch
 - ☐ `switchpassword` is the password for the switch
 - ☐ `192.168.1.1` the IP address of the server that has the new firmware
 - ☐ `root` is the user name of the server that has the new firmware
 - ☐ `password` is the password of the server that has the new firmware
 - ☐ `/myshare/firmware_directory/firmware_file` is the full path to the firmware file.
4. Save the `update_Brocade.conf` file.
-

5. Click **Business Tools** ().
6. Click **Brocade Firmware Update** in the right pane under Automators.
7. Select the switch you want to upgrade, and then click **OK**.
8. When you are asked if you want to upgrade the firmware, enter **Yes**.
9. After the firmware is updated, reboot the switch.

Setting up Risk Analysis

The risk analysis tools flag which HBAs and/or switches are a risk in your environment. The HBA Risk Analysis and Switch Risk Analysis determine which elements are at risk by checking them against predefined profiles you created. Before you run either of these tools, create profiles for your environment, as described in the following steps.

1. Go to the %JBOS4_DIST%\server\appiq\remotescripts\advisors directory and open one of the following files in a text editor, such as Notepad:
 - ☐ **HBA Risk Analysis** - hba_risk_analysis.conf
 - ☐ **Switch Risk Analysis** - switch_risk_analysis.conf
2. In the configuration file, do one of the following:
 - ☐ If you are editing hba_risk_analysis.conf, define a profile for each type of HBA in your environment.
 - ☐ If you are editing switch_risk_analysis.conf, define a profile for each type of switch in your environment.

The profile name can be any name you want, and the number and type of fields in the profile is defined by you. For example, assume you want to check that all HBAs in Solaris hosts are a certain model (LP-9000). The profile would be in hba_risk_analysis.conf and it would resemble the following:

```
any profile name
the 'OS' field = Solaris
the 'Model' field = LP-9000
```

Let's expand that analysis to verifying that those HBA's also have a certain driver version and a certain firmware level, as shown in the following example profile:

```
any profile name
the 'OS' field = Solaris
the 'Model' field = LP-9000
```

```
the 'DriverVersion' field = whatever
the 'FirmwareVersion' field = whatever
```

Important: The use of 'special' characters in the field values are interpreted as regular expressions (or wildcards) by the search. For example, "Model LP900[02]" will match the values "LP-9000" or "LP-9002". "FirmwareVersion 3.90A[0-9]" will match any firmware version 3.90A0 through 3.90A9.

3. Once you have your profiles configured, run HBA Risk Analysis or Switch Risk Analysis. The tool compares every HBA or switch against the profiles and flag any HBA or switch that does not match at least one profile.

Global Change Management

Global Change Management lets you save the current configuration and/or compare changes with a previous configuration.

Accessing Global Change Management

1. Access Business Tools.
2. Click **Global Change Management Tools**.

Saving the Current Configuration

To save the current configuration, use option 0. The current configuration is saved to a DAT file in the following directory. This file can be opened by using a text editor, such as Notepad:

```
%JBoss4_DIST%\server\appiq\remotescripts\advisors\saved-configurations
```

The DAT file contains the data for the elements discovered in the configuration. The following is an example of a portion of a DAT file for a saved configuration:

```
$VAR1 = 'conf3';
$VAR2 = 1117783473;
__DATA__
$VAR1 = {
    '1002' => {
```

```
'PROPERTIES' => {  
  'ID' => '1002',  
  'HostType' => 'Default',  
  'DnsName' => 'QA67',  
  'SupportFlags' => '7',  
  ....  
}
```

where

- \$VAR1 is the configuration name, which is conf3 in this case.
- 1002 is the element ID.

Comparing a Previous Configuration by Using Global Change Management

Note: Global Change Management requires you to provide the name of the current configuration the first time you run the tool. Global Change Management then assigns the saved configuration name to a number.

To compare a previous configuration, enter the number corresponding with the previous saved configuration. While the script is determining the changes, it lists the elements it is analyzing. Once it is done, it lists the changes under the heading “CHANGED PROPERTIES” on the screen.

Storage Terminology Table

The software uses the storage terminology defined by the Storage Networking Industry Association (SNIA). However, each storage company has its own vocabulary. If you are not familiar with SNIA's terminology, use the table below to map the unfamiliar terms to the terms you know.

For more information about storage networking terms, see the glossary in this user guide and refer to the Storage Networking Industry Association's (SNIA) Web site. SNIA has compiled a list of storage terms at <http://www.snia.org/education/dictionary/>.

Storage Terminology Table

The software uses the storage terminology defined by the Storage Networking Industry Association (SNIA). However, each storage company has its own vocabulary. If you are not familiar with SNIA's terminology, use the table below to map the unfamiliar terms to the terms you know.

For more information about storage networking terms, see the glossary in this user guide and refer to the Storage Networking Industry Association's (SNIA) Web site. SNIA has compiled a list of storage terms at <http://www.snia.org/education/dictionary/>.

Table 19-1: Storage Terminology Table

SNIA	EMC² Symmetrix	EMC² CLARiiON	Hitachi Data Systems	Engenio	Network Appliance
Back-end SCSI Controller	Disk Director (DA)	NA	NA	Fibre Drive Channel	NA
Composite Disk Extent	Device or Metadevice	LUN	Logical Device (LDEV)/LUN Size Expansion(LU SE)	Volume or Free Capacity	LUN
Disk Drive	Disk Drive	Disk	Physical Device (PDEV)	Drive	Disk Drive
Disk Extent	Hyper	Part of a LUN that resides on the same physical disk.	Part of LDEV that resides on a physical disk	Port of volume or free capacity that resides on the disk drive	Part of a LUN that resides on the same physical disk.
Front-end SCSI Controller	Fibre Channel Director/ Interface Adapter (FA)	SP Port	Cluster	Fibre channel host port	Fibre channel target adapter port
Front-end Fibre Channel Port	FA port	SP Port	Port	Fibre channel host port	Fibre channel target adapter port
Front-end Storage System Processor	FA	SP	Port Controller	Controller	CPU Mainboard

Table 19-1: Storage Terminology Table (Continued)

SNIA	EMC² Symmetrix	EMC² CLARiiON	Hitachi Data Systems	Engenio	Network Appliance
Storage Pool	Collection of devices with the same protection level	RAID Group	Array Group	Volume group	Volume
Storage System	Storage Array	Storage Array	Storage Array	Storage array	Filer
Volume	Device or Metadevice	LUN	LDEV/LUSE	Volume	LUN

This chapter describes the following:

- “Data is late or an error occurred” Message on page 786
 - appiq.log Filled with Connection Exceptions on page 786
 - Receiving “HTTP ERROR: 503” When Accessing the Management Server on page 787
 - Permanently Changing the Port a CIM Extension Uses (UNIX Only) on page 789
 - Configuring UNIX CIM Extensions to Run Behind Firewalls on page 790
 - Volume Names from Ambiguous Automounts Are Not Displayed on page 798
 - Solaris Management Server Suddenly Restarts on page 798
 - Installing the Software Security Certificate on page 799
 - Troubleshooting After Upgrading on page 801
 - Troubleshooting Discovery and Get Details on page 802
 - Troubleshooting Topology Issues on page 813
 - Troubleshooting the Java Plug-in on page 827
 - Troubleshooting Provisioning on page 828
 - Troubleshooting Hardware on page 830
-

“Data is late or an error occurred” Message

If you see the message “Data is late or an error occurred” when you try to obtain information from a UNIX host, verify you were logged in as root when you started the CIM Extension (`./start`). You must be logged in as root if you want to use the `./start` command, even if you are using the `./start -users username` command, where `username` is a valid UNIX account.

The CIM Extension only provides the information within the privileges of the user account that started the CIM Extension. This is why you must use root to start the CIM Extension. Only root has enough privileges to provide the information the management server needs.

If you continue to see the message, contact customer support.

appiq.log Filled with Connection Exceptions

When an Oracle REDO log becomes corrupt, the management server is unable to connect to the database. Whenever the management server is unable to connect to the Oracle database, it writes to the `appiq.log` file. Many exceptions may cause the Application Log on Windows to become full.

To fix the problem, stop the management server and Oracle. Then, remove the corrupted REDO log, as described in the following steps:

1. Stop the AppStorManager service, which is the service the management server uses.

Note: While the service is stopped, the management server cannot monitor elements and users cannot access the management server.

2. To find the corrupt log file, look in the `alert_APPIQ.log` file, which can be found in one of the following locations:
 - ❑ **Windows:** `\oracle\admin\APPIQ\bdump`
 - ❑ **Solaris:** `$ORACLE_BASE\admin\APPIQ\bdump`
-

You can verify if the REDO log listed in the alert_APPIQ.log file is corrupt by looking for a “redo block corruption” error in the REDO log.

3. On the management server, enter the following at the command prompt:
`Sqlplus /nolog`
4. Enter the following:
`Sql> connect sys/change_on_install as sysdba`
5. Enter the following:
`Sql> startup mount;`
6. Enter the following:
`Sql> ALTER DATABASE CLEAR UNARCHIVED LOGFILE
'C:\ORACLE\ORADATA\APPIQ\REDO02.LOG';`
where C:\ORACLE\ORADATA\APPIQ\REDO02.LOG is the corrupted log file and its path.
7. Enter the following:
`Sql> alter database open`
8. Enter the following:
`Sql> shutdown immediate;`
9. Enter the following:
`Sql> startup`

Receiving “HTTP ERROR: 503” When Accessing the Management Server

If you receive a message resembling the following when you try to access the management server, make sure your database for the management server is running. If it is not, start the database.

```
Receiving HTTP ERROR: 503 javax.ejb.EJBException: null;
```

Refer to the following sections for more information about how to start database for the management server.

Windows

Access the Services window to make sure the OracleOraHome92TNSListener service has started and is set to automatic. Refer to the Windows documentation for information on how to access the Services window.

If the OracleOraHome92TNSListener service has not started but the AppStorManager service has started, start the OracleOraHome92TNSListener service and then restart AppStorManager.

Solaris

To verify the Oracle service has started, enter the following at the command prompt:

```
# ps -ef | grep ora
```

Output resembling the following is displayed if the service has started:

```
/opt/oracle/product/9.2.0.1.0/bin/tnslsnr LISTENER -inherit
./appstormservice /opt/productname/ManagerData/conf/solaris-wrapper.
oracle   356      1  0   Jul 30 ?          0:01 ora_pmon_APPIQ
oracle   358      1  0   Jul 30 ?          0:26 ora_dbw0_APPIQ
oracle   360      1  0   Jul 30 ?          1:13 ora_lgwr_APPIQ
oracle   362      1  0   Jul 30 ?          0:39 ora_ckpt_APPIQ
oracle   364      1  0   Jul 30 ?          0:10 ora_smon_APPIQ
oracle   366      1  0   Jul 30 ?          0:00 ora_reco_APPIQ
oracle   368      1  0   Jul 30 ?
```

If you find your service for the Oracle has not started, you can start the service by entering the following at the command prompt:

```
# /etc/rc3.d/S98dbora start
```

If you need to stop the service for Oracle, enter the following at the command prompt:

```
# /etc/rc3.d/S98dbora stop
```

Important: If you are starting the services manually, start the Oracle service before the service for the management server.

Errors in the Logs

If you access the logs, you are shown messages resembling the following. The complete text has been shorten as a result of space constraints:

```
Aug 04 2004 11:59:07] INFO
[com.appiq.service.policyManager.policyService.PolicyService] Creating
[Aug 04 2004 11:59:07] INFO
[com.appiq.service.policyManager.policyService.PolicyService] Created
[Aug 04 2004 11:59:07] INFO
[com.appiq.service.policyManager.policyService.PolicyService] Starting
[Aug 04 2004 11:59:07] INFO
[com.appiq.service.policyManager.policyService.PolicyService] Starting
Policy Factory
[Aug 04 2004 11:59:11] ERROR
[com.appiq.security.DatabaseSecurityManager] DatabaseSecurityManager
Error:

org.jboss.util.NestedSQLException: Could not create connection; -
nested throwable: (java.sql.SQLException: ORA-01033: ORACLE
initialization or shutdown in progress

); - nested throwable: (org.jboss.resource.ResourceException: Could
not create connection; - nested throwable: (java.sql.SQLException:
ORA-01033: ORACLE initialization or shutdown in progress

))
```

Permanently Changing the Port a CIM Extension Uses (UNIX Only)

CIM Extensions on UNIX use port 4673 by default. You can start a CIM Extension on another port by entering `./start -port 1234`, where 1234 is the new port. With this method, you must always remember to provide the nondefault port when starting the CIM Extension.

You can configure a CIM Extension to remember the nondefault port, so you only need to enter `./start` to start the CIM Extension:

1. Go to the `/opt/APPQcime/conf` directory.
2. Open the `cim.extension.parameters` file in a text editor, and provide the following line:
`-credentials username:password`
`-port 1234`

Important: The values for `-credentials` and `-port` must be on separate lines, as shown in the example.

where

- ☐ `username` is the user that is used to discover the CIM Extension. You will need to provide this user name and its password when you discover the host.
 - ☐ `password` is the password of `username`.
 - ☐ `1234` the new port for the CIM Extension
3. Save the file.
 4. Restart the CIM Extension for your changes to take effect.

Note: The CIM Extension looks for parameters in the `cim.extension.parameters` file whenever it starts, such as when it is started manually or when the host is rebooted.

5. The management server assumes the CIM Extension is running on port 4673. The management server also listens on port 17000 for CIM Extensions from build 4.0. If you change the port number, you must make the management server aware of the new port number.

In the IP Address/DNS Name field in the Add Address for Discovery page (**Discovery > Setup > Add Address**), type a colon and then the port number after the IP address or DNS name, as shown in the following example:

```
192.168.1.2:1234
```

where

- ☐ `192.168.1.2` is the IP address of the host
- ☐ `1234` is the new port number

If you have already added the host to the discovery list (**Discovery > Setup**) on the management server, you must remove it and then re-add it. You cannot have more than one listing of the host with different ports.

Configuring UNIX CIM Extensions to Run Behind Firewalls

In some instances you will need to discover a host behind a firewall. Use the following table as a guideline. Assume the management server wants to discover HostA, which has three network interface cards on three separate networks with three separate IPs: 10.250.250.10, 172.31.250.10, and 192.168.250.10. In the following table different configurations are presented:

- The “Manual Start Parameters for CIM Extensions” column provides what you would enter to start the CIM Extension manually on the host. See the Installation Guide for more information on how to start a CIM Extension manually.
 - The “If Mentioned in `cim.extension.parameters`” column provides information on how you would modify the `cim.extension.parameters` file. See “Permanently Changing the Port a CIM Extension Uses (UNIX Only)” on page 789.
 - The “Step 1 Discovery (**Discovery** > **Setup**) and RMI Registry Port” column - Provides information about what IP addresses are required for the discovery list. The RMI Registry port is the port the CIM Extension uses. Keep in mind that when a port other than 4673 is used for the CIM Extension, the port must be included in the discovery IP. For example, 192.168.1.1:1234, where 192.168.1.1 is the IP for the host and 1234 is the port the CIM Extension uses.
-

Table 20-1: Troubleshooting Firewalls

Configuration	Manual Start Parameters for CIM Extension	If Mentioned in cim.extension.parameters	Step 1 Discovery and RMI Registry Port
Firewall port 4673 opened between host and management server	start		10.250.250.10 OR 172.31.250.10 OR 192.168.250.10 Communication Port: 4673
Firewall port 1234 opened between host and management server, but specific port	start -port 1234	-port 1234	10.250.250.10 :1234 OR 172.31.250.10 :1234 OR 192.168.250.10 :1234 Communication Port: 1234
Firewall port 4673 opened between host and management server on the 172.31.250.x subnet	start -on 172.31.250.10	-on 172.31.250.10	172.31.250.10 Communication Port: 4673

Table 20-1: Troubleshooting Firewalls (Continued)

Configur- ation	Manual Start Parameters for CIM Extension	If Mentioned in cim.extension.paramet ers	Step 1 Discovery and RMI Registry Port
Firewall port 1234 opened between host and management server on the 192.168.250.x subnet	start -on 192.168.250.10:1234	-on 172.31.250.10:1234	172.31.250.10 :1234 Communication Port: 1234
With 3 firewall ports opened on different ports respectively 1234, 5678, 9012.	start -on 10.250.250.10:1234 -on 172.31.250.10: 5678 -on 192.168.250.10: 9012	-on 10.250.250.10:1234 -on 172.31.250.10:5678 -on 192.168.250.10: 9012	10.250.250.10 :1234 OR 172.31.250.10 :5678 OR 192.168.250.1 0:9012 Communication Port: 1234, 5678, 9012

Table 20-1: Troubleshooting Firewalls (Continued)

Configur- ation	Manual Start Parameters for CIM Extension	If Mentioned in cim.extension.paramet ers	Step 1 Discovery and RMI Registry Port
With firewall port 4673 opened between host and management server. NAT environment where 10.250.250.10 subnet is translated to 172.16.10.10 when it reaches other side of the firewall	start		172.16.10.10 Communication Port: 17001

Table 20-1: Troubleshooting Firewalls (Continued)

Configur- ation	Manual Start Parameters for CIM Extension	If Mentioned in cim.extension.paramet ers	Step 1 Discovery and RMI Registry Port
With firewall port 1234 opened between a host and management server. NAT environment where 10.250.250.10 subnet is translated to 172.16.10.10 when it reaches other side of the firewall	start -port 1234	-port 1234	172.16.10.10 Communication Port: 17001
With 3 firewall ports opened on different ports respectively 1234, 5678, 9012. NAT environment where all 3 NICs are translated to different 172.16.x.x subnets	start -on 10.250.250.10:1234 -on 172.31.250.10:5678 -on 192.168.250.10:9012	-on 10.250.250.10:1234 -on 172.31.250.10:5678 -on 192.168.250.10:9012	172.16.10.10: 1234 OR 172.16.20.20: 5678 OR 172.16.30.30: 9012 Communication Port: 1234, 5678, 9012

Table 20-1: Troubleshooting Firewalls (Continued)

Configuration	Manual Start Parameters for CIM Extension	If Mentioned in cim.extension.parameters	Step 1 Discovery and RMI Registry Port
False DNS or IP is slow to resolve		jboss.properties, stop and restart service cimom.Dcxws.agency.firstwait=200000 cimom.Dcxws.agency.timeout=200000	Any IP that is reachable Communication Port: 4673
No DNS, never resolve		jboss.properties, stop and restart service cimom.Dcxws.agency.firstwait=200000 cimom.Dcxws.agency.timeout=200000	Any IP that is reachable Communication Port: 4673

Table 20-1: Troubleshooting Firewalls (Continued)

Configur- ation	Manual Start Parameters for CIM Extension	If Mentioned in cim.extension.paramet ers	Step 1 Discovery and RMI Registry Port
No firewall. Don't want to use root credentials. Want to discover with a non-existent user.	start -credentials abcuser:passwd	-credentials abcuser:passwd	Specify abcuser and password in the discovery list. Communication Port: 4673
With 3 firewall ports opened on different ports respectively 1234, 5678, 9012. Don't want to use root credentials. Want to discover with a non-existent user.	start -on 10.250.250.10:1234 -on 172.31.250.10:5678 -on 192.168.250.10:9012 -credentials abcuser:passwd	-on 10.250.250.10:1234 -on 172.31.250.10:5678 -on 192.168.250.10:9012 -credentials abcuser:passwd	10.250.250.10 :1234 OR 172.31.250.10 :5678 OR 192.168.250.1 0:9012. Then, specify abcuser and passwd in the discovery list. Communication Port: 1234, 5678, 9012

Volume Names from Ambiguous Automounts Are Not Displayed

Volume names from ambiguous automounts on Solaris hosts are not displayed on the Storage Volumes page and in Capacity Explorer. Some Solaris hosts have autofs and NFS mounted through an automounter. The management server cannot display volume names from ambiguous automounts because it cannot determine if the comma separate strings that are part of the mounted volume name are host names or part of the name of a remote volume.

The following example is a comma separate string that is part of a mounted volume name. The management server cannot tell whether `test` and `three` are host names or part of the name of a remote volume. As a result, the management server does not display the volume name.

```
VolumeName = two:/ntlocal2,two:/comma,test,three,one:/ntlocal
```

Solaris Management Server Suddenly Restarts

When the memory usage for management server Java process grows considerably, users may experience sudden restart of management server on machines with low physical memory. The restart occurs because Java Virtual Machine (JVM) for management server exits when it is not able to expand heap during Garbage Collection. This is a known JVM issue.

Work around:

1. Increase the swap size on solaris server.
 2. Set `-Xms` and `-Xmx` to the same value and `-XX:PermSize` and `-XX:MaxPermSize` to the same value so that no heap expansion takes place during Garbage Collection. These variables can be set using the Advanced option under the Product Health menu.
-

Installing the Software Security Certificate

To stop receiving a Security Alert message each time you use the HTTPS logon, install the software security certificate, as described in the following steps.

Important: Enter the DNS name of the computer in the URL instead of localhost. If you use `https://localhost` to access the management server, you are shown a “Hostname Mismatch” error.

Installing the Certificate by Using Microsoft Explorer 6.0

1. Access the management server by typing the following:
`https://machinename`
where `machinename` is the name of the management server.
2. When the security alert message appears, click **OK**.
If you do not want the Web browser to warn you about a secure connection at any Web site, select the **In the future, do not show this warning** option.
3. When you are told there is a problem with the site's security certificate, click the **View Certificate** button.
4. When you are shown the certificate information, click the **Install Certificate** button at the bottom of the screen.
5. When you are shown the Certificate Import Wizard, click **Next** to continue the installation process.
6. Select one of the following:
 - ☐ **Automatically select the certificate store based on the type of certificate** - This option places the certificate automatically in the appropriate location.
 - ☐ **Place all certificates in the following store** - This option lets you pick the store where the certificate will be stored.
7. Click **Finish**.
8. When you are asked if you want to install the certificate, click **Yes**.
You are shown the following message when the certificate is installed.

Installing the Certificate by Using Netscape Navigator 7

1. Access the management server by typing the following:
`https://machinename`
where `machinename` is the name of the management server.
2. When the security alert message appears, click the **Always** button.
3. When you are told you are requesting an encrypted page, click **OK**.
4. Click the **Always** button when you are asked if you want to accept the certificate.
5. When asked if you wanted to trust the signed applet, click the **Always** button.

Changing the Security Certificate to Match the Name of the Server

If your users are shown a Security Alert window with the following message, you might want to modify the security certificate so users feel more comfortable with installing the certificate:

“The name of the security certificate is invalid or does not match the name of the site.”

You can change the security certificate so that users receive the following message instead:

“The security certificate has a valid name matching the name of the page you are trying to view.”

When you change the certificate, you must use the `generateAppiqKeystore.bat` program to delete the original certificate. Then, use the `generateAppiqKeystore.bat` program to create a new certificate and to copy the new certificate to the management server.

To change the certificate:

1. (Solaris management servers only) Go to the `[Install_Dir]/Tools` directory and run the following command:

```
eval `./usersvars.sh`
```

Important: The quotes in the example must be entered as left single quotes.

2. Go to the following directory:
-

□ **Solaris** - [Install_Dir]/Tools

□ **Windows** - [Install_Dir]\Tools

where [Install_Dir] is the directory into which you installed the management server.

3. To delete the original certificate, enter the following at the command prompt:

```
perl generateAppIQKeyStore.pl del
```

The original certificate is deleted.

Note: If you see an error message when you enter this command, a previous certificate may not have been created. You can ignore the error message.

4. To create a new certificate containing the DNS name of the management server, enter the following at the command prompt:

```
perl generateAppIQKeyStore.pl
```

5. If the program is unable to detect a DNS name, enter the following at the command prompt:

```
perl generateAppIQKeyStore.pl create mycomputername
```

where mycomputername is the DNS name of the computer

6. To copy the new certificate to the management server, enter the following at the command prompt:

```
perl generateAppIQKeyStore.pl copy
```

The new certificate is copied to the correct location.

Troubleshooting After Upgrading

This section provides information on troubleshooting after upgrading. It includes the following topics:

SMI-S Switches Must Be Removed and Rediscovered After Upgrading on page 802

SMI-S Switches Must Be Removed and Rediscovered After Upgrading

You must remove and rediscover Cisco, QLogic, or CNT switches supported through SMI-S after you upgrade the management server.

1. Remove the Cisco, QLogic, or CNT switches supported through SMI-S. See “Deleting Elements from the Product” on page 92.
2. Rediscover the Cisco, QLogic, or CNT switches. See “Step 1 - Discover Switches” on page 34.
3. Perform Get Details. See “Step 4 - Get Details” on page 94.

Troubleshooting Discovery and Get Details

This section describes the following:

- Names are Changed After Running Get Details for Cisco SMI-S Switches on page 803
 - Configuring E-mail Notification for Get Details on page 804
 - Increasing the Time-out Period and Number of Retries for Switches in Progress on page 805
 - “Connection to the Database Server Failed” Error on page 807
 - Using the Test Button to Troubleshoot Discovery on page 807
 - DCOM Unable to Communicate with Computer on page 809
 - Duplicate Listings for Brocade Switches in Same Fabric on page 810
 - Element Logs Authentication Errors During Discovery on page 810
 - EMC Device Masking Database Does Not Appear in Topology (AIX Only) on page 810
 - Management Server Does Not Discover Another Management Server's Database on page 811
 - Microsoft Exchange Drive Shown as a Local Drive on page 811
 - Unable to Discover Microsoft Exchange Servers on page 811
 - Nonexistent Oracle Instance Is Displayed on page 811
-

- Requirements for Discovering Oracle on page 811
- “Do Not Run Overlapping Discovery Schedules” on page 812
- “This storage system uses unsupported firmware. ManagementClassName: ???” Message on page 812
- Troubleshooting Topology Issues on page 813
- Incorrect Topology Sometimes Displayed for CNT Switches on page 819
- Unable to Find Elements on the Network on page 819
- Unable to See Path Information on page 819
- Device Locking Mechanism for Brocade Element Manager Query/Reconfiguration on page 819
- A Discovered Sun StorEdge A5000 JBOD Does Not Display Its WWN Properly on page 820
- Unable to Monitor McDATA Switches on page 820
- Unable to Detect a Host Bus Adapter on page 820
- Navigation Tab Displays Removed Drives as Disk Drives on page 821
- Unable to Obtain Information from a CLARiiON Storage System on page 821
- Discovery Fails Too Slowly for a Nonexistent IP Address on page 821
- “CIM_ERR_FAILED” Message on page 822
- Communicating with HiCommand Device Manager Over SSL on page 825
- Unable to Discover a UNIX Host Because of DNS or Routing Issues on page 826
- Troubleshooting Hardware on page 830

Names are Changed After Running Get Details for Cisco SMI-S Switches

Cisco switches on the fabric are displayed without their names—each Cisco SMI-S switch name on the fabric is replaced with a generic name such as: Switch_1401 after running Get Details.

To view the correct names for the Cisco SMI-S switches after running Get Details, manually enter the correct name for the Cisco SMI-S switches.

Configuring E-mail Notification for Get Details

The management server lets you send status reports about Get Details to users. The status reports that are sent to users can also be found in the `GAEDSummary.log` file in the `[Install_DIR]\logs` directory on the management server.

To configure the management server to send status reports on Get Details to your e-mail account:

1. Enable e-mail notification for the management server. Refer to the User Guide for more information.
2. Add or edit the e-mail address for the Admin account.

The status reports for Get Details automatically go the e-mail account provided for the Admin user. To add or edit an e-mail address for the Admin account, log in as Admin and then follow the steps in “Modifying Your User Profile” on page 166.
3. If you want additional users to receive the status reports for Get Details, do the following:
 - a. Click **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree.
 - b. Click **Show Default Properties** at the bottom of the page.
 - c. Copy the `gaedemail` property. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
 - d. Return to the Advanced page (**Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree).
 - e. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
 - f. Assign the e-mail accounts you want to receive the report to the `gaedemail` property. For example, if you want `user1@appiq.com` and `user2@appiq.com` to receive these status reports, modify the `gaedemail` property in the **Custom Properties** field as follows:
`gaedemail=user1@appiq.com;user2@appiq.com`
Note: Make sure the harsh (#) symbol is removed from the `gaedmail` property.
 - g. When you are done, click **Save**.
 - h. Restart the service for the management server for your changes to take effect.

Increasing the Time-out Period and Number of Retries for Switches in Progress

If you are having difficulty obtaining information from switches with SNMP connections during Get Details, you may need to increase the time-out period and the number of retries. By default, the management server gives a switches five seconds to respond to its requests for information during Get Details. If the switch does not respond the first time, the management server tries again. The management server says it cannot contact the switch if it does not receive a response from the switch a second time.

To change the time-out period and number of retries for switches, modify the properties specified Table 20-2, "Time-out Properties," on page 806 and Table 20-3, "Retry Properties," on page 806 as described in the following steps:

1. Access the management server.
2. Select **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree.
3. Click **Show Default Properties** at the bottom of the page.
4. Copy the commands specified in Table 20-2, "Time-out Properties," on page 806. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
5. Return to the Advanced page (**Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree).
6. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
7. Make sure the property is not commented out by removing the hash (#) symbol in front of the property. To modify the time-out period, set the corresponding property for your switch in the following table to the number of millisecond you want. The default is 5000 ms. For example, to change the time-out period to 30000 ms for a McDATA switch, you would set the `cimom.McData.Snmp.Timeout` property to 30000, as shown in the following example:

```
cimom.McData.Snmp.Timeout=30000
```

Table 20-2: Time-out Properties

Switch	Property
McDATA/Connectrix discovered through SNMP	<code>cimom.McData.Snmp.Timeout</code>
Cisco	<code>cimom.Cisco.Snmp.Timeout</code>
Other switches discovered through SNMP: ■ Sun StorEdge ■ QLogic	<code>cimom.snmp.switch.timeout</code>

8. To modify the number of retries, repeat Steps 4 through 6 by copying and pasting the property specified in Table 20-3, “Retry Properties,” on page 806. Set the corresponding property for your switch in the following table to the number of retries you want. The default is two retries. For example, to change the number of retries to five for a McDATA switch, set the `cimom.McData.Snmp.Retries` properties as shown in the following example:
- ```
cimom.McData.Snmp.Retries=5
```

**Table 20-3: Retry Properties**

| Switch                                                                                                                              | Property                               |
|-------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------|
| McDATA/Connectrix<br>discovered through<br>SNMP                                                                                     | <code>cimom.McData.Snmp.Retries</code> |
| Cisco                                                                                                                               | <code>cimom.Cisco.Snmp.Retries</code>  |
| Other switches<br>discovered through<br>SNMP:<br><ul style="list-style-type: none"> <li>■ Sun StorEdge</li> <li>■ QLogic</li> </ul> | <code>cimom.snmp.switch.retries</code> |

9. When you are done, click **Save**.
10. Restart the service for the management server for your changes to take effect.



While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

## “Connection to the Database Server Failed” Error

If you received an error message resembling the following after getting all element details, verify that the database instance is running:

```
The connection to the database server failed. Check that the Oracle
instance 'OIQ3 on host '192.168.1.162:1521 is running correctly and has
the management software for Oracle installed correctly.
```

Assume you received the error message listed above. You would want to verify the following:

- Oracle instance OIQ3 on host 192.168.1.162 port 1521 is running.
- The management software for Oracle is installed on the server running the Oracle instance. One of the installation's tasks is to create an APPIQ\_USER user account with enough privileges for the software to view statistics from the database.

Once you have verified the items listed above, run Get Details again. If you continue to see the error message, contact customer support.

## Using the Test Button to Troubleshoot Discovery

If you are having problems discovering an element, click the **Test** button on the Discovery setup page (**Discovery > Setup**). When you click the **Test** button, the management server attempts to ping the element. Then, it runs a series of device-specific connectivity tests. The output of these tests can be viewed in the discovery log window.

The management server uses a provider to communicate with an element. A provider is software that communicates with the element and the management server. When you click the **Test** button, it checks every available provider against the element to see which one works. When this test is being performed, you may notice messages such as “Test provider not supported,” “Connection Refused” or “Failed to Establish Connection.” This means a provider was tested against the element and the provider was not the correct one.

When the correct provider is found, a message is displayed, such as “ExampleComputer responds to a Wind32 system” or “Connection accepted,” as shown below:

```
Testing provider APPIQ_Win32Provider for: 192.168.1.2
ExampleComputer responds as a Win32 system with CIM Extensions 3.0.0.129
```

The success messages are intertwined with the other messages, so you need to scroll through the log messages. For example, the success message shown previously appeared in the middle of the log messages, as shown in the following example. The success message is underlined in the following example.

To make it easier to view the log messages, copy and paste the log messages from the log window to a text editor.

LOG MESSAGES

```
[2004/01/15 09:10] Test Discovery Started
[2004/01/15 09:10] Successfully pinged 192.168.1.2
[2004/01/15 09:10]
Testing provider APPIQ_SolarisProvider for: 192.168.1.2
Connection refused to host: 192.168.1.2; nested exception is:
java.net.ConnectException: Connection refused: connect
Testing provider APPIQ_CimProxyProvider for: 192.168.1.2
Test provider functionality not supported for APPIQ_CimProxyProvider
Testing provider APPIQ_McDataProvider for: 192.168.1.2
Can't connect.
No current SWAPI connection to host 192.168.1.2. Cannot establish
connection
Testing provider APPIQ_AltixProvider for: 192.168.1.2
Connection refused to host: 192.168.1.2; nested exception is:
java.net.ConnectException: Connection refused: connect
Testing provider APPIQ_IrixProvider for: 192.168.1.2
Connection refused to host: 192.168.1.2; nested exception is:
java.net.ConnectException: Connection refused: connect
Testing provider APPIQ_Win32Provider for: 192.168.1.2
ExampleComputer responds as a Win32 system with CIM Extensions 3.0.0.129
Windows host does not support remote testing
VERITAS Volume Manager not available
HDLN Multipathing Software not available
Powerpath Multipathing Software not available
RDAC Multipathing Software not available
Testing provider APPIQ_EmcProvider for: 192.168.1
Can't connect
```

---

```

appiqSymInitialize() failed with error code 510
Testing provider APPIQ_AixProvider for: 192.168.1.2
Connection refused to host: 192.168.1.2; nested exception is:
java.net.ConnectException: Connection refused: connect
Testing provider APPIQ_HdsProvider for: 192.168.1.2
Cannot connect to Proxy
Cannot connect to Proxy
Testing provider APPIQ_BrocadeElementManager for: 192.168.1.2
Cannot connect
Cannot connect
Testing provider EngenioSSI_Provider for: 192.168.1.2
Failed to establish connection.
Testing provider APPIQ_ClariionProvider for: 192.168.1.2
NaviCLI not installed
No such file: C:\Program Files\EMC\Navisphere CLI\NaviCLI.exe
[2004/01/15 09:10] Test Discovery Completed
TEST DISCOVERY COMPLETED in 5 seconds

```

---

**Note:** By design the **Test** button is not available when any of the discovery steps are occurring.

---

## DCOM Unable to Communicate with Computer

Sometimes the following error message appears in the event log of the management server when the software is monitoring a Brocade switch:

```

DCOM was unable to communicate with the computer 192.168.10.21 using
any of the configured protocols

```

where 192.168.10.21 is the IP address of the Brocade switch.

Ignore this error message.

---

## Duplicate Listings for Brocade Switches in Same Fabric

If you discover more than one Brocade switch in the same fabric, the **Targets** tab displays duplicate listings for the Brocade switches. Each Brocade switch is listed multiple times with the IP address of the other switches and its own.

For example, assume you discovered Brocade switches QBrocade2 and QBrocade5 in the same fabric, the switches are listed twice on the **Targets** tab. QBrocade2 is listed twice, once with its own IP address, the other time with the IP address of QBrocade5, as shown below:

```
192.168.10.22 Switch QBrocade2, QBrocade5 admin
192.168.10.25 Switch QBrocade2, QBrocade5 admin
```

## Element Logs Authentication Errors During Discovery

During discovery, you may see SNMP authentication errors on the element you are trying to discover. The management server is probing the element with an SNMP request. If the element does not know the management server, it logs authentication errors.

## EMC Device Masking Database Does Not Appear in Topology (AIX Only)

An EMC device masking database attached to an AIX host does not appear in the Topology tree under the **Application Path - Unmounted** node on the **Topology** tab in System Explorer.

If the EMC device masking database is attached to a host running Microsoft Windows or Sun Solaris, the masking database appears under the **Application Path - Unmounted** node.

---

## Management Server Does Not Discover Another Management Server's Database

In some situations the management server may not discover another management server's database. Make sure that the Oracle monitoring software (CreateOracleAct.bat for Microsoft Windows or CreateOracleAct.sh for UNIX) is installed on the management server to be discovered and that the Oracle instance is added to the discovery list.

## Microsoft Exchange Drive Shown as a Local Drive

Microsoft Exchange Servers have a drive M. The software displays this drive as a local fixed disk, instead of a Microsoft Exchange Server special drive.

## Unable to Discover Microsoft Exchange Servers

If DNS records for your Microsoft Exchange servers are outdated or missing, the discovery of Microsoft Exchange may fail because Microsoft Exchange is dependant on Active Directory, which is dependant on DNS. Since Active Directory is dependant on DNS, Active Directory replication and Active Directory lookups may fail or contain errors if DNS records are not accurate.

## Nonexistent Oracle Instance Is Displayed

The software uses the Oracle Transparent Name Substrate (TNS) listener port to detect Oracle instances on a server. Sometimes an Oracle instance is removed from the server, but not from the TNS listener port. This results in the software detecting the nonexistent Oracle instance and displaying it in the topology. Refer to Oracle documentation for information on how to remove the deleted Oracle instance from the TNS listener port.

## Requirements for Discovering Oracle

To discover Oracle:

- The management software for Oracle must be installed. For information about installing the management software for Oracle, refer to the *Installation Guide*.
- By default, the software sets the TNS Listener Port to 1521. If you use another port, you can change the port number on the Discovery Targets tab.
- Oracle discovery relies on the TNS networking substrate on which Oracle is built (TNS is Oracle's proprietary protocol). The software does not use TNS listener password. If you have set a TNS Listener password, the software is not able to discover the Oracle instances serviced by the listener.

## Do Not Run Overlapping Discovery Schedules

If you are creating multiple Discovery schedules, care must be taken to avoid scheduling conflicts—concurrently scheduled Discovery tasks—and that each scheduled task has enough time to start and finish before the next Discovery task is scheduled to start. For example, if a scheduled Discovery is still in progress when another scheduled Discovery attempts to start, the Discovery task that attempts to start will not start because the first Discovery is still running. The Discovery that is unable to start is re-scheduled according to its recurring rule. If the Discovery task is scheduled to run on a daily basis for example, then the Discovery will start again on the next day. To check the status of scheduled Discovery tasks, view the appiq.log file in the following directory:

```
[Install_Dir]\jbossandjetty\server\appiq\logs
```

## "This storage system uses unsupported firmware. ManagementClassName: ???" Message

The following message is displayed when an Engenio storage system is discovered, and is running unsupported firmware:

**"This storage system uses unsupported firmware. ManagementClassName: ???"**

The management class name for the unsupported array is displayed in the message.

New releases of storage system firmware are supported with each new release of this software. See the Support Matrix for the latest information on supported firmware.

---

---

## Troubleshooting Topology Issues

This section describes the following:

- About the Topology on page 813
- Undiscovered Hosts Display as Storage Systems on page 817
- Solaris Machines Appear to Have Extra QLogic HBAs on page 818
- No Stitching for Brocade Switches with Firmware 3.2.0 on page 818
- Link Between a Brocade Switch and a Host Disappears from the Topology on page 818
- Incorrect Topology Sometimes Displayed for CNT Switches on page 819
- Unable to Find Elements on the Network on page 819
- Unable to See Path Information on page 819
- Device Locking Mechanism for Brocade Element Manager Query/Reconfiguration on page 819
- A Discovered Sun StorEdge A5000 JBOD Does Not Display Its WWN Properly on page 820
- Unable to Monitor McDATA Switches on page 820
- Unable to Detect a Host Bus Adapter on page 820
- Navigation Tab Displays Removed Drives as Disk Drives on page 821
- Unable to Obtain Information from a CLARiiON Storage System on page 821
- Discovery Fails Too Slowly for a Nonexistent IP Address on page 821
- “CIM\_ERR\_FAILED” Message on page 822
- Communicating with HiCommand Device Manager Over SSL on page 825
- Unable to Discover a UNIX Host Because of DNS or Routing Issues on page 826

## About the Topology

The software determines the topology by looking at the following:

- **Fibre channel switch** The fibre channel switch contains a list of all elements within the fabric. The software obtains a detailed listing of all elements connected to the switch fabric.
  - **A host containing a Host Bus Adapter (HBA)** All fibre channel host adapters look for available elements attached to the HBA. This information is gathered by CIM Extensions and sent to the management server.
-

- **A proxy connected to the SAN** - Include a proxy that has a direct connection or a SAN connection to the management server. An example of a proxy is the EMC Solutions Enabler or Hitachi HiCommand Device Manager. Engenio storage systems do not require a proxy, as they can be accessed directly. Make sure the proxy service has started. On a computer running Windows, this can be determined by looking in the **Services** window.

Table 20-4: Troubleshooting Discovery and Get Details

| Scenario                                                                                                                                                          | Description                                                                              | What to do                                                                                                                    |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <div><p>The host appears discovered and it is connected to the switch.</p></div> | The software is aware of the host, but it cannot obtain additional information about it. | <p>Verify that a CIM Extension is installed on the host.</p> <p>Try discovering the element again. Then, run Get Details.</p> |



Table 20-4: Troubleshooting Discovery and Get Details (Continued)

| Scenario                                                                                                                                                                                                                                                                        | Description                                                                                                                                                                                                                  | What to do                                                                                                                                                                                                                                             |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  <p>Host_3017</p>  <p>QBrocade1</p> <p>Host appears discovered and it is not connected to the switch.</p>     | <p>The switch was previously made aware of the host, but it can no longer contact it.</p> <p>If the steps provided do not work, see “Link Between a Brocade Switch and a Host Disappears from the Topology” on page 818.</p> | <p>Verify that the host is on and the network cables are connected to it.</p> <p>Try discovering the element again. Then, run Get Details.</p>                                                                                                         |
|  <p>Host_3017</p>  <p>QBrocade1</p> <p>The host appears managed, but it is not connected to the switch.</p> | <p>There is a problem with Get Details from the host.</p> <p>If the steps provided do not work, see “Link Between a Brocade Switch and a Host Disappears from the Topology” on page 818.</p>                                 | <p>Try getting the topology again.</p> <ol style="list-style-type: none"> <li>1. Click the <b>Discovery</b> menu. Then, click the <b>Topology</b> tab.</li> <li>2. Verify the element is selected and click the <b>Get Topology</b> button.</li> </ol> |

Table 20-4: Troubleshooting Discovery and Get Details (Continued)

| Scenario                                                                                                                                                                                    | Description                         | What to do                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <div><p>Host_3017</p></div> <p>The element appears discovered, but a connected switch does not appear.</p> | The switch has not been discovered. | <p>Try discovering the switch again.</p> <ol style="list-style-type: none"><li>1. Click the <b>Discovery</b> menu. Click the <b>Setup</b> tab and the <b>Add Address</b> button on the <b>IP Addresses</b> tab.</li><li>2. Type the IP address or DNS Name of the switch. Then, type its user name and password. Click <b>OK</b>.</li><li>3. Verify the element is selected.</li><li>4. Click the <b>Start Discovery</b> button.</li><li>5. After discovery has completed, click the <b>Topology</b> tab.</li><li>6. Verify the element is selected and click the <b>Get Topology</b> button.</li></ol> |

Table 20-4: Troubleshooting Discovery and Get Details (Continued)

| Scenario                                                                                                                                                                                                                                                                                                                                                                                                     | Description                         | What to do                                                                         |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|------------------------------------------------------------------------------------|
| <p>When discovering a Windows-based host, the correct IP address is entered, but the host does not appear in the topology.</p> <p>The following can be seen on the host:</p> <ul style="list-style-type: none"> <li>■ *In Windows Event Manager the <code>WinMgmt.exe</code> process is not running. This process starts WMI.</li> <li>■ In the Windows Event Log, DCOM error messages are shown.</li> </ul> | An invalid user account was entered | Enter a valid user account that has administrative privileges so it can start WMI. |

\*The CIM Extension for Microsoft Windows enhances Windows Management Instrumentation (WMI) so that it can gather information from host bus adapters and make the information available to management server.

**Important:** One way to determine what is happening is to look at the log messages during discovery and getting element details. See “Viewing Log Messages” on page 60 for more information.

## Undiscovered Hosts Display as Storage Systems

On rare occasions, undiscovered hosts will display as storage systems in System Explorer. Specifying the host’s WWN in the `jboss.properties` file will cause the host to display correctly. Follow these steps to modify the `jboss.properties` file:

1. Determine the host's WWN. This information is available on the IEEE Standards Association web site at <http://standards.ieee.org/regauth/oui/oui.txt>.
2. Go to the %JBOSS4\_DIST%\server\appiq\conf directory on the management server.
3. Open the jboss.properties file in a text editor and locate the `hostPortWWNs` property.
4. Uncomment the property and enter the host's WWN in hex format. Multiple WWNs can be entered as a comma-separated list. For example:  
`hostPortWWNs=00-01-C9,00-01-C8`
5. Restart the service for the management server.

## Solaris Machines Appear to Have Extra QLogic HBAs

Solaris machines using fibre channel drives internally will always appear to have extra QLogic HBAs. After discovering a Solaris machine, internal fiber channel drives will show an extra QLogic adapter on the host adapters page.

## No Stitching for Brocade Switches with Firmware 3.2.0

Stitching does not appear for hosts attached to Brocade switches running firmware 3.2.0. There is no stitching when the PID format is 0. The port setting must be the same for all Brocade switches in the fabric or the fabric will become segmented. The PID format should be set to 1 for all Brocade switches running firmware later than 2.6.0 and 3.0. The PID=0 setting is a legacy Port ID format that does not support the numbers of ports beyond 16.

## Link Between a Brocade Switch and a Host Disappears from the Topology

If a link that used to work between a Brocade switch and a host disappears from the topology, you may need to rediscover the Brocade switch and the host. Also, confirm that both are online and there are no network connection issues. As a last resort, you may need to reboot the switch. In some instances, the API of the Brocade switch has been known to hang. Rebooting the switch clears the switch of the API hang.

---

## Incorrect Topology Sometimes Displayed for CNT Switches

The CNT SMI-S provider for CNT switches does not return the correct topology information when more than one fabric is managed by the same InVSN™ Storage Network Manager. McDATA, which completed its acquisition of CNT in the summer of 2005, has been made aware of this issue.

## Unable to Find Elements on the Network

The management server uses ping to find the devices on the network enabled for IP. Ping is a program that lets you verify that a particular IP address exists. Ping is not guaranteed to return a response from all devices. If Discovery is not able to find a device automatically, enter the IP address for the device on the Discovery Targets tab, which can be accessed by clicking the Discovery button at the top of the screen in the management server. Sometimes ping cannot find the device if one of the following conditions occur:

- Network configuration does not support ping, including data center security (firewalls).
- Device has the ping responder turned off.
- Device does not support ping.

## Unable to See Path Information

You will not be able to see path information if LUN masking information is missing. To view LUN masking information, follow the steps described in the section, “Accessing Information About Host Security Groups” on page 452.

## Device Locking Mechanism for Brocade Element Manager Query/Reconfiguration

Please keep in mind that the configuration for Brocade switches is locked while getting all details for elements in a zones. The software ensures that each CIM query locks out any reconfiguration. For example, if you are getting details for elements in all zones, you cannot add a new Brocade switch while you are doing it (the discovery or configuration process waits until the collection of details is finished before proceeding). However, simultaneous CIM queries do not lock each other out.

## A Discovered Sun StorEdge A5000 JBOD Does Not Display Its WWN Properly

Although full monitoring and management support is available only to those devices for which there is a provider, the software's topology displays other devices found on your storage area network (SAN) to give you a more complete view. However, because these devices do not have a provider, only basic information is returned. In some cases, as with the Sun StorEdge A5000 JBOD (Just a Bunch of Disks), the Worldwide Name (WWN) presented and reported to the management server may be different from the official WWN of the device, as the management server reports the WWN of the port connected to the fabric.

## Unable to Monitor McDATA Switches

McDATA switches use the Fibre Channel Switch Application Programming Interface (SWAPI) to communicate with devices on the network. The McDATA switches allow only one SWAPI connection at a time. For example, if the management server discovers the IP address of the McDATA switch, other management servers and third-party software are not able to communicate with the switch using SWAPI.

Use Enterprise Fabric Connectivity (EFC) Manager to communicate with the McDATA switch. EFC Manager versions 7.0 and later can communicate with the management server and the switch. This configuration lets multiple instances of the management server or other clients contact EFC Manager, which in turn provides information about the switch. To communicate with the EFC Manager, discover the McDATA switches as described in Chapter 2, “Discovering NAS Devices, Tape Libraries, Switches and Storage Systems” on page 21.

---

**Important:** EFC Manager uses the SWAPI connection, preventing other third-party software from contacting the switch.

---

## Unable to Detect a Host Bus Adapter

The software is unable to detect a host bus adapter if you install its driver before you have completed installing the Solaris operating system for the first time. For example, you installed the

---

HBA drives too early when you used JumpStart to install Solaris. The best way to install the HBA driver is to install it after Solaris has been installed and is running.

## Navigation Tab Displays Removed Drives as Disk Drives

If you remove an internal disk from a Solaris host and do not enter the `cfgadm` command, the Navigation tab displays the empty slot as `DiskDrives_XXXXXX` after getting element details. The `cfgadmn` command makes the software realize the drive has been removed. Refer to the documentation that shipped with the Solaris operating system for more information about the `cfgadm` command.

## Unable to Obtain Information from a CLARiiON Storage System

If you are having difficulty obtaining topology information or element details from a CLARiiON storage system, the NaviCLI might have timed out as a result of the service processor being under a heavy load. The management server uses the NaviCLI to communicate with the CLARiiON storage system. This situation has been seen in the field when the service processor is running more than 35,000 IOPS (IOs/Sec).

Try obtaining the topology and/or Get Details from a CLARiiON storage system when the service processor is not under such a heavy load.

## Discovery Fails Too Slowly for a Nonexistent IP Address

If you enter a nonexistent IP address, the management server times out by default after 20 seconds on Windows and after three minutes and after 45 seconds on Solaris. If you want to shorten the time-out period, modify the `cimom.CimXmlClientHttpConnectTimeout` property as described in this section.

---

**Note:** The management server does not accept a period longer than its default setting. If you set `cimom.CimXmlClientHttpConnectTimeout` property to more than 20 seconds on Windows or three minutes and 45 seconds on Solaris, the management server ignores the values of this property and reverts back to the default settings.

---

To modify the default time-out:

1. Access the management server.
  2. Click **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree.
  3. Click **Show Default Properties** at the bottom of the page.
  4. Copy the `cimom.CimXmlClientHttpConnectTimeout` property you want to modify. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
  5. Return to the Advanced page (**Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree).
  6. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
  7. Make your changes in the **Custom Properties** field. Make sure the property is not commented out by removing the hash (#) symbol in front of the property.
  8. To modify the time-out period, set the `cimom.CimXmlClientHttpConnectTimeout` property to the number of millisecond you want. For example, to change the time-out period to 200 ms, set the `cimom.CimXmlClientHttpConnectTimeout` property, as shown in the following example:  
`cimom.CimXmlClientHttpConnectTimeout=200`
  9. When you are done, click **Save**.
  10. Restart the service for the management server for your changes to take effect. While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.
-



## “CIM\_ERR\_FAILED” Message

If you are in a McDATA environment where the EFC Manager Service Processor is managing multiple switches, it is possible that the management server will send SWAPI requests faster than the EFC Manager Service Processor can handle them. The management server may detect this as a failed connection and take corrective action. When this happens, you are shown a “CIM\_ERR\_FAILED” message whenever the management server tried to access the McDATA switches and directors.

The management server then attempts to reconnect to the EFCM by creating a new SWAPI connection. EFCM versions 8.x and later have five SWAPI connections available. EFCM versions 7.1.3 and later but before version 8.x have three SWAPI connections available. If the management server reconnects successfully, a reconnect event is generated and no further action is necessary.

If the management server cannot reconnect to the EFCM, another event is generated with a severity of major. If this happens, any operation the management server is performing (Get Details) involving switches on that EFCM fails.

To prevent the “CIM\_ERR\_FAILED” messages, increase the delay between the management server’s SWAPI calls to EFCM, as described in the following steps:

1. Click **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy `cimom.mcData.swapIThrottle=200`. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
4. Return to the Advanced page (**Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree).
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Make your changes in the **Custom Properties** field by changing the value of `cimom.mcData.swapIThrottle`. For example, the default is 200 milliseconds. To change the value to 800 milliseconds, change the xxx value to 800, as shown in the following example:

```
cimom.mcData.swapIThrottle=800
```

---

**Note:** If you want no delay, change the value to 0 for 0 milliseconds. The maximum delay you can have is 1,000 milliseconds (`cimom.mcData.swapIThrottle=1000`),

---

7. When you are done, click **Save**.
8. Restart the service for the management server for your changes to take effect. While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.
9. Verify if you can re-establish communication with EFCM by following the steps in “Re-establishing Communication with EFCM” on page 824. You may need to change the value of the `cimom.mcData.swapIThrottle` property if you cannot re-establish communication with EFCM after following the steps in that section.

## Re-establishing Communication with EFCM

To re-establish communication with EFCM, perform the following steps:

1. To check the status of the connection, click the **Test** button on the Discovery Setup screen. If the McDATA provider reports that it can connect to EFCM, the connection has been restored. A provider is a component of the management server that is used to gather information about an element. In this case, the McDATA provider gathers information about McDATA switches for the management server. To ensure the management server does not have corrupt data as a result of the loss of communication, perform Get Details to obtain the latest information from the element.
  2. If the ping results to EFCM fails, there is a network problem that must be resolved. Once network connectivity is restored, use the **Test** button to verify the McDATA provider can communicate with EFCM, then do a Get Details.
  3. If the **Test** button results from the management server indicates that it still cannot communicate with EFCM, wait approximately three minutes for the lost SWAPI connection to time out, then click the **Test** button again. If it works, do a Get Details.
  4. If the **Test** button results continue to indicate a lost connection after three minutes, perform the following steps to restore the connection. Please note these steps involve restarting services on the EFCM server. Any other applications using SWAPI to communicate with EFCM are affected by these actions.
-

- a. Open the EFCM client. Make sure that the EFCM is still actively managing at least one switch. If there are no switches under management, you will not be able to connect to this EFCM.
- b. On the EFCM server, stop and restart the Bridge Agent service. Repeat Steps 1 through 3. If the connection is still down, proceed to Step c.
- c. On the EFCM server, stop and restart the EFCM services. On Windows, use the McDATA EFCM Manager options in the **Start > Programs** menu. Repeat Step 1 through 3. If the connection is still down, proceed to Step d.
- d. Reboot the EFCM server. Repeat Step 1 through 3. If the connection is still down, proceed to Step e.
- e. Stop and restart the service for the management server. Repeat Step 1 through 3. If the connection is still down, proceed to Step f.
- f. Reboot the management server. Repeat Step 1 through 3. If the connection is still down, proceed to Step g.
- g. If none of the above steps have restored the connection, refer to the support matrix to determine if the EFCM and switch versions are all supported. Contact technical support for further information.

## Communicating with HiCommand Device Manager Over SSL

By default the management server communicates with HiCommand Device Manager through a nonsecure connection. You can configure the management server so that it communicates with HiCommand Device Manager over a secure socket layer (SSL) connection by doing one of the following:

- **Use HTTPS in the discovery address** - Prepend `https://` to the discovery address to force the connection to HTTPS mode, for example, `https://192.168.1.1`, where 192.168.1.1 is the IP address of the host running HiCommand Device Manager. Use this option if you have one HiCommand Device Manager you want to communicate through a secure connection (SSL) and another you want to communicate through a nonsecure connection.
- **Modify an internal property** - Change the value of the `cimom.provider.hds.useSecureConnection` to `true`, as described in the steps in this section. Use this option if you want all connections to HiCommand Device Manager to be secure (SSL).

To set all connections with HiCommand Device Manager to SSL:

1. Click **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the `cimom.provider.hds.useSecureConnection` property. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
4. Return to the Advanced page (**Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree).
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Make your changes in the **Custom Properties** field. Make sure the property is not commented out by removing the hash (#) symbol in front of the property.
7. Change the value assigned to the `cimom.provider.hds.useSecureConnection` property to true, as shown in the following example:  
`cimom.provider.hds.useSecureConnection=true`
8. When you are done, click **Save**.  
If you want to connect to another instance of HiCommand Device Manager by using a nonsecure connection, prepend `http://` to the discovery address to force the connection to nonsecure mode, for example, `http://192.168.1.1`, where 192.168.1.1 is the IP address of the host running HiCommand Device Manager.
9. Restart the service for the management server for your changes to take effect. While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

**Important:** While the AppStorManager service is stopped, the following occurs:

- Users are not be able to access the management server.
- The management server is unable to monitor elements at this time.

## Unable to Discover a UNIX Host Because of DNS or Routing Issues

If the management server is unable to discover a UNIX host because of a DNS or routing issues, you will need increase the amount of time that passes before the management server times out for that CIM Extension. By default, the management server waits 1,000 ms before it times out. It is recommended you increasing the time before the management server times out to 200000 ms (3.33

---

minutes), as described in the following steps. If you continue to see time out issues, you can still increase the time before the management server times out, but keep in mind that it will lengthen discovery.

To increase the time out period:

1. Select **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree.
2. Paste the following text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.

```
cimom.cxws.agency.firstwait=200000
```

```
cimom.cxws.agency.timeout=200000
```

where

- ☐ `cimom.cxws.agency.firstwait` - The `firstwait` property controls the amount of time required for the management server to wait after it first contacts the CIM Extension on the host before the management server attempts to proceed with a username and password. The default value is 1,000 ms. You are modifying it to wait 20,000 ms or 3.33 minutes.
  - ☐ `cimom.cxws.agency.timeout` - The `timeout` property controls the allowable interval of silence before either the CIM Extension or the management server start to question whether its partner is still alive. If an entity (management server or extension) has not received a message from the other during the interval set by the `timeout` property, it will send an "are you there" message. If that message is not acknowledged during the interval set by the `timeout` property, the entity will conclude that the connection is no longer functioning. The CIM Extension will stop attempting to make a connection. When this occurs on the side of the management server, the management server will attempt to re-connect (and it will keep re-attempting until the host becomes available). The default value is 1,000 ms. You are modifying it to wait 20,000 ms or 3.33 minutes.
3. Click **Save**.
  4. Restart the service for the management server for your changes to take effect. While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

---

## Troubleshooting the Java Plug-in

This section describes the following:

- Java Applet Has Data from a Different Version of Management Server Software on page 827
- OutOfMemoryException Messages on page 828
- Improving Reload Performance in System Explorer on page 828

### Java Applet Has Data from a Different Version of Management Server Software

If you attempt to monitor a host with old JAR (Java Archive) files, you might be unable to monitor the host and you might see the following error message:

```
The Java applet has data from a different version of the management
server. Please close and re-start your browser.
```

The reason for this error message is that the client still has JARs from the previous build in its Java Plug-in cache. To remove the old JARs, clear the cache for the Java plug-in.

### OutOfMemoryException Messages

In some rare cases it may be necessary to increase the amount of memory for the Java plug-in on the client computer. This should only be done if you are seeing `OutOfMemoryException` messages in the Java console on the client side.

### Unable to View System Explorer After Upgrade

System Explorer might not display if the Java applet plug-in for the Web browser is configured to use a proxy. This issue has been seen after the management server has been upgraded and the Web browser has cached Java class files. Clearing the cache does not correct this issue. The only known work around is to disable the proxy.

---

## Improving Reload Performance in System Explorer

If your Java plug-in control panel cache setting is set at 50 MB, it is recommended you increase this setting to 150 MB or more. Increasing this setting improves the reloading performance of System Explorer.

---

## Troubleshooting Provisioning

This section describes the following:

- Cannot Access a Resource Owned by Another Controller on page 829
- Error -56 on page 829
- “Can’t delete this zone” Message on page 829
- Changes in EFC Manager Requiring Get Details on page 829

### Cannot Access a Resource Owned by Another Controller

You received the message about not being able to access a resource owned by another controller because you tried to access a controller that has not been discovered. You should discover all controllers on the Engenio storage system.

For example, assume you discovered only one of the controllers on an Engenio storage system with two controllers. If you want to change a volume, such as add or delete a LUN, you will not be able to make the change to the volume associated with the controller that has not been discovered.

See Chapter 2, “Discovering NAS Devices, Tape Libraries, Switches and Storage Systems” on page 21 for more information on how to discover a controller.

---

## Error -56

If you see `error -56`, the switch has network connection failures/problems. To solve the problem, make sure the switch is physically connected to the network. Then, redo the task you were originally trying to complete.

If you now see `-21 (OBJECT_NOT_FOUND)` errors, the switch needs to be rediscovered.

## “Can't delete this zone” Message

If you are shown the following message when you try to delete a zone, move the zone to an inactive zone set. Then, delete the zone.

```
Can't delete this zone, it is member of an Active Zoneset
```

## Changes in EFC Manager Requiring Get Details

If you use EFC Manager to delete zones or zone sets, perform Get Details on the management server afterwards. The changes are not reflected by the management server until Get Details is done.

---

## Troubleshooting Hardware

This section describes the following:

- About Swapping Host Bus Adapters on page 830
  - “Fork Function Failed” Message on AIX Hosts on page 830
  - Known Driver Issues on page 830
  - Known Device Issues on page 831
  - “mailbox command 17 failure status FFF7” Message on page 834
  - “Process Has an Exclusive Lock” Message on page 835
-



## About Swapping Host Bus Adapters

Swapping brands of host bus adapters (HBA) on a Microsoft Windows 2000 host may have undesirable side effects. For example, after swapping out one brand of an HBA for another (including driver installation), `WinMgmt.exe` might crash repeatedly and appear to be associated with an error in the Windows Event Log about being unable to retrieve data from the `PerfLib` subkey in the Registry. To solve this problem, reinstall the operating system.

## “Fork Function Failed” Message on AIX Hosts

If a CIM Extension running on AIX detects low physical or virtual memory while starting, a “Fork Function Failed” message appears. A CIM Extension on AIX uses additional memory and CPU resources at start time. If the resources on the AIX machine is already low, you may see the “Fork Function Failed” message. Depending on the AIX operating system or hardware, the host may crash after you see this message.

## Known Driver Issues

If you are having problems with a driver, keep in mind the following:

- The software requires the driver to have a compliant SNIA HBA API. Emulex driver version 4.21e does not support the SNIA HBA API.
- If the driver has a compliant SNIA HBA API, make sure the driver is installed correctly.

## Known Device Issues

The following table provides a description of the known device issues. You can find the latest information about device issues in the release notes.

Table 20-5: Known Device Issues

| Device                               | Software | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------------|----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AIX host                             | NA       | <p>If you are receiving replication errors for an AIX host, the provider may be trying to connect to the host using the 0.0.0.0 IP address instead of the real host IP address. If this situation is occurring, you would see a message containing the following when you start the CIM Extension:</p> <pre>CXWS 3.1.0.144 on 0.0.0.0/0.0.0.0 now accepting connections</pre> <p>To fix this situation, add the following line to the <code>/opt/APPQcime/tools/start</code> file on the AIX host:</p> <pre>export NSORDER=local,bind</pre> |
| AIX host using an IBM Storage System | NA       | <p>If you have an AIX host using an IBM storage system, not all bindings may be displayed on the bindings page on the Navigation tab. For example, assume diskA on host123 has six paths. All six bindings may not be displayed.</p>                                                                                                                                                                                                                                                                                                        |

Table 20-5: Known Device Issues (Continued)

| Device                                          | Software                                           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------------------------------------|----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hosts running SGI IRIX version 6.5.22 or 6.5.24 | NA                                                 | If a host running SGI IRIX version 6.5.22 or 6.5.24, the HBA port page on the Navigation tab in System Explorer displays 0 GB/s for HBA ports.                                                                                                                                                                                                                                                                                                                                                                                                    |
| SGI IRIX host                                   | CXFS file systems                                  | <p>The management server can only monitor CXFS file systems from the host generating the input/output. For example, assume the elements are part of a CXFS file system. When you generate input/output into the metadata server into /folder, only the metadata server is able to monitor the file system. For example, assume the metadata server generates 100 KB write, the management server displays 0 KB write for /folder on the metadata client.</p> <p>See “About the Data from CXFS File Systems” on page 397 for more information.</p> |
| Solaris host                                    | Sun SAN Foundation Suite driver (Leadville driver) | The bindings page reports a SCSI number that comes from the HBA API. This number cannot be seen by the user. For example SCSI target 267008 does not correlate to anything.                                                                                                                                                                                                                                                                                                                                                                       |
| Solaris host                                    | HDLM                                               | <p>If you sync the Solaris host by itself without the switches and storage, the storage volume page reports all drive types as local.</p> <p>Once you discover the host with the switches and storage, it reports its drives as being external. It was the same result with Active-Active and Active-Standby.</p>                                                                                                                                                                                                                                 |
| Solaris host                                    | HDLM                                               | Solaris HDLM disks cannot be monitored. If you try monitoring them, the management server displays a message saying “data is late or an error occurred.”                                                                                                                                                                                                                                                                                                                                                                                          |

Table 20-5: Known Device Issues (Continued)

| Device       | Software | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------|----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Solaris host | HDLM     | <p>Do a Get Details for the host by itself. In the bindings page, the controller number are displayed as c-1. For example c-1t0d58.</p> <p>Perform Get Details on the host with storage and switches. The controller numbers are displayed correctly.</p>                                                                                                                                                                                                                                        |
| Solaris host | VxVM     | <p>If you discover a host with any typical SAN disk groups off line, the storage volume page shows SAN mount points as local instead of external. These disks, however, are not accessible.</p> <p>When you perform Get Details with all disk groups online, disks on the SAN are shown as external. Hosts connected directly to a storage system are shown as local, except for hosts connected by fibre. Hosts connected directly to a storage system through fibre are shown as external.</p> |
| Windows host | VxVM     | <p>When a Windows host with VxVM is used, the SCSI bus number is always reported to be one in the SCSI bus column of the Disk Drives page.</p>                                                                                                                                                                                                                                                                                                                                                   |

Table 20-5: Known Device Issues (Continued)

| Device              | Software                                                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------|----------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Any host            | NA                                                       | The Unmounted Volume field under Capacity Summary automatically displays 0 MB if you discovered the host but not the storage system connected to it. This may occur if you did not enter the IP address of the storage system when performing discovery and/or your license does not allow you to discover a particular storage system. See the Supported Elements section in the “List of Features” to determine which storage systems you can discover. The “List of Features” is accessible from the Documentation Center ( <b>Help &gt; Documentation Center</b> ). |
| IBM Storage Systems | Subsystem Device Driver (SDD) or MPIO (Multi Pathing IO) | If you discovery an IBM storage system without SDD, incorrect stitching is displayed in System Explorer for the storage system. You are shown only one path if the storage system is using MPIO instead of SDD.                                                                                                                                                                                                                                                                                                                                                         |

## “mailbox command 17 failure status FFF7” Message

If one or more of your Microsoft Windows hosts are using an Emulex HBA driver, you may see the following message in Windows Event Viewer:

“mailbox command 17 failure status FFF7”

This message can be safely ignored. The HBA API is being used to access data in the FLASH memory of the adapter that does not exist and this is causing the event to be logged. This issue has been seen with version 5.2.2 of the driver.

## **”Process Has an Exclusive Lock” Message**

You will receive a message resembling the one shown below if a process has already locked the EMC Symmetrix storage system and you attempt a process that requires a lock on the Symmetrix storage system. The Symmetrix storage system can become locked for many reasons. For example, the storage system becomes locked when it performs LUN mapping, LUN masking or Get Details. The Symmetrix storage system may also remain locked after a provisioning operation has failed.

“SYMAPI routine SymDevMaskSessionStart failed with error code 188: The operation failed because another process has an exclusive lock on the local Symmetrix.”

After the management server has detected the lock on the Symmetrix storage system, it tries to access the storage system for 15 minutes and logs the errors.

If you receive the error message, determine if someone is performing an operation that requires a lock, such as LUN mapping, LUN masking or Get Details. This also applies even if one of the processes is being used by a third-party product, such as for LUN masking. If so, wait until the process is complete. Only manually remove the lock if you are certain that no other processes are occurring on the storage system. To learn how to remove the lock, refer to the documentation for the Symmetrix storage system.

If a provisioning failure has caused the Symmetrix storage system to remain locked, you are alerted to this situation in Event Manager and on the Properties tab. You may receive a message resembling the following:

```
Unable to end device masking session. Symmetrix '000001835005700' may be
locked.
```

---

## **A**

### **access point**

It is the intersection of the IP address and the provider that discovered the IP address. It is displayed on the screens for discovery. A provider is software that is used to gather information about an element.

### **active zone set**

An active zone set is the zone set in use. You can have only one zone set active at a time; however, you can have a zone in more than one zone set. Zones sets are usually created for a particular task.

Zones work by exclusion. Members of a zone can communicate with other members in the zone. If two devices are not within the same zone, they cannot communicate. Only elements in active zones can communicate with each other. When a zone is not active, it does not have any effect.

## **C**

### **Common Information Model (CIM)**

The Common Information Model is a common data model of an implementation-neutral schema for describing overall management information in a network/enterprise environment. CIM is comprised of a specification and a schema. The specification defines the details for integration

---

with other management models (i.e. SNMP's MIBs or the DMTF's MIFs) while the schema provides the actual model descriptions.

### **Common Information Model Object Manager (CIM Object Manager)**

A component in the CIM management infrastructure that handles the interaction between management applications and providers. The CIM Object Manager supports services such as event notification, remote access, and query processing. The CIM Object Manager also grants access to the CIM Object Manager repository.

## **F**

### **File Server Storage Resource Management (SRM)**

File Server Storage Resource Management (SRM) does a recursive lookup on the file system and stores the information in an embedded database. File Server SRM can scan files very quickly because of its structure in the database and because it uses a multi-threaded process. More than one process can be used at a time to scan the files.

## **D**

### **device**

This documentation set defines a device as a piece of hardware in the storage network.

EMC uses the term device to refer to a volume on one of its storage systems.

## **E**

### **element**

An element is anything on the network that can be detected by the management server, such as hosts and switches.

### **element created in Chargeback**

An element created in Chargeback is a type of generic element. When you create a record for an element in Chargeback, the element appears as a generic element in the topology.

## **G**

### **generic element**

An element is considered to be generic if the management server can detect the element but it cannot obtain additional information about the element during Get the Topology or Get Details.

### **global reporting view**

A global reporting view contains information in the database that can be used for global reports.

---



**Global Reporter server**

A management server that has global reporting enabled.

**H****hard zone**

A hard zone is created by assigning a domain/port to a zone. Any device attached to the port is automatically in the zone.

**host persistent binding**

A system SCSI target ID assigned permanently to an element. The host binding is implemented on the host bus adapter (HBA), resulting in the HBA being tied to a certain LUN.

**I****initiator WWN**

The Worldwide Name (WWN) of a host bus adapter's port. The WWN differentiates the port from others.

**inode file**

An inode file stores information about a file, excluding the file's data.

**M****Managed Application Licenses (MALs)**

Managed application licenses (MALs) are the number of detected instances of Microsoft Exchange, Oracle and Sybase Adaptive Server Enterprise.

**managed object**

A hardware or software system component that is represented as an instance of the CIM class. Information about managed objects is supplied by data and event providers, as well as by the CIM Object Manager.

**Managed Access Points (MAPs)**

Manage access points (MAPs) are the sum of all storage access ports of all hardware elements that the management server manages.

**materialized view**

A materialized view is a snapshot of data, from the database, created from a query.

---

**mapped**

Mapped is capacity that is accessible by one or more hosts external to the array (aggregated capacity of volumes that are accessible from hosts external to the subsystem).

**meta device**

This term is used by EMC. A meta device is a device that is a concatenation of several devices.

**metavolume**

Metavolumes are created from a disk, slice, stripe, or other metavolumes. Metavolumes are extremely useful because they can expand their storage capacity, such as to mainframe volume sizes. Also referred to “LDEVs” for HDS storage systems.

**missing element**

The management server was able to discover the element, but it lost contact with the element before more information could be gathered during “Get the Topology” or Get Details. A missing element can be managed if the management server lost contact with the element after Get Details was performed.

**multipathing**

The process of providing a server more than one path to a storage system. So that in the case of an emergency, the server will have continuous access to the storage system. Multipathing can be done many ways. For example, you can provide redundant switches for a server to access a storage system. Another example of multipathing is providing redundant paths from the server to the switch.

## P

**provider**

A provider is software that is used to gather information about an element.

**proxy server**

A device, such as a host, connected to a storage system. It is sometimes referred to as a storage system proxy or an API proxy. An example of a proxy server is the EMC Solutions Enabler or Hitachi HiCommand Device Manager. Engenio storage systems do not require a proxy, as they can be accessed directly.

## S

**SAN**

A Storage Area Network (SAN) is a high speed network configuration that is dedicated to transporting storage data among network devices, such as storage systems, hosts (servers), switches, and tape libraries to end users. In addition to connecting local elements to storage

---

arrays, it may also be extended to off site or remote locations for the purposes of backup, archival or acting as a hot site in the event of a disaster.

A SAN can communicate via current technologies such as ESCON (mainframe), fibre channel, or newer technology such as iSCSI. SAN's can support several configurations such as disk mirroring, RAID 5, backup/restore, and data migration, as well as being able to incorporate Network Attached Storage (NAS).

### **SMTP**

Simple Mail Transfer Protocol (SMTP) is a TCP/IP protocol used in sending and receiving e-mail.

### **soft zone**

A soft zone is created by assigning a world wide name (WWN) of a device port to a zone.

### **Storage Management Infrastructure Specification**

A Storage Networking Industry Association (SNIA) standard for implementing data storage management using the Common Information Model (CIM).

### **storage pool**

It is a group of volumes. Also known as volume group.

### **switch port zoning**

A type of zoning in which the port of the switch is physically in the zone. Any device attached to the port is automatically in the zone.

## **V**

### **Virtual Storage Area Networks (VSAN)**

See VSAN.

### **Virtual Application**

A placeholder you create for an unsupported application. For example, assume your company has created an internal application, and you want to be able to use the software to keep track of that application. You can create a virtual application for that product.

### **VSAN (Virtual Storage Area Networks)**

Virtual Storage Area Networks (VSANs) are logical fabrics formed as subsets of physical FC (Fibre Channel) switch networks. A VSAN is defined as a set of Fx, E and TE ports-entry/exit, traditional ISL (Inter-switch Link) ports, and trunked ISL ports. TE ports are ISLs that may be shared between a named set of VSANs. All ports other than TE ports must be members of exactly one VSAN. On Cisco switches, there is a default VSAN (VSAN 1) that initially includes

all ports, and an isolated VSAN (VSAN 4094) where ports end up if their owning VSAN is deleted.

## W

### **Web-Based Enterprise Management (WBEM)**

Web-Based Enterprise Management (WBEM) is an initiative based on a set of management and Internet standard technologies developed to unify the management of enterprise computing environments. WBEM provides the ability for the industry to deliver a well-integrated set of standard-based management tools leveraging the emerging technologies such as CIM and XML.

### **Windows Management Instrumentation (WMI)**

Microsoft created WMI as its implementation of Web-based Enterprise Management (WBEM). For more information about WMI, refer to the Microsoft Web site at <http://www.microsoft.com>.

The Windows CIM Extension enhances Windows Management Instrumentation (WMI) so that it can gather information from host bus adapters and make the information available to the management server.

### **Worldwide Name (WWN) zoning**

A type of zoning in which the port is assigned to a Worldwide Name of a host or a storage system. It is not dependent on the switch.

### **WWN (Worldwide Name)**

Fibre-channel Worldwide Name. Usually formatted as 16-hexadecimal digits. This name is globally unique, and it identifies the connection or set of connections to the network.

## Z

### **zone**

A collection of zone aliases and ports.

### **zone alias**

To avoid remembering a port's Worldwide Name (WWN), assign the port to a zone alias.

### **zone member**

A zone member is a port attached to a switch, a Worldwide Name (WWN) or a zone alias.

As a best practice, a zone should contain either zone aliases or ports, but not both.

---

**zone set**

A zone set is a collection of zones. You can have only one zone set active in a fabric; however, you can have a zone in more than one zone set. Zones sets are usually created for a particular task.

---



---

# *Index*

---

## **Numerics**

3PAR storage systems 65

## **A**

about

Access tab 307

asset attributes 392

Business Tools 772

buttons in System Explorer 300

Capacity Explorer 660

Chargeback reports 756

custom commands 354

custom reports 579

CXFS 397

Event Manager 536

Events tab 391

filtering events 550

groups 345

Home page 6

List tab 304

monitoring options 642

Navigation screen 367

Path tab 315

Performance Explorer 632

Policies tab 396

Policy Manager 682

provisioning 399

Provisioning tab 391

remote console 362

Reporter 560

Reports tab 394

security 154

storage tiers 716

System Explorer 296

Topology Screen 378

virtual applications 841

zone aliases 842

zone sets 843

access point 837

Access tab 300, 307

accessing

asset attributes 720

Database Admin Utility 278

domain controller 125

---

- e-mail schedules 578
  - events 555
  - host security groups 452
  - logs 216
  - management server 14
  - Navigation screen 373
  - Policy Manager 682
  - properties 374
  - storage pools 436
  - System Explorer 298
  - trending information 678
  - volume information 438
  - zone sets 416
  - zoning information 411
  - accessing information
    - zone aliases 408
  - account
    - password 163
  - Accounters 772
  - accounts
    - users 162
  - activating
    - a zone set 421
    - zone sets 416
  - active
    - zone sets 416
  - Active Client Logons 642
  - Active Directory 811
  - active zone set 837
  - adding
    - asset information 724
    - asset record 717
    - custom command 355
    - custom information 729
    - departments 730
    - discovery policies 690
    - discovery schedule 209
    - domain controller 125, 148
    - elements 173, 176
    - e-mail schedule 571
    - event policies 694
    - general information 727
    - geographic information 729
    - host security groups 454
    - hosts 491, 507, 519
    - IP address 29
    - IP range 28
    - journal entries 546
    - license 200
    - licensing information 729
    - new elements 108
    - organizations 173
    - provisioning policies 692
    - roles 170
    - staff information 728
    - storage pool 434
    - storage volume 441
    - switches 101
    - TNS Listener Port 147
    - user accounts 162
    - utilization policies 683
    - virtual application 317
    - virtual applications 390
    - warranty information 729
    - zone alias members 409
    - zone members 415
    - zone sets 416, 417
    - zones 413
  - Advisors 772
  - aging statistics 243
  - AIX 810
  - API proxy 840
  - appiq.log 215, 218
  - APPIQ\_OWNER account 125
  - APPIQ\_USER 149
  - application
    - virtual 317
  - Application Administrator role 154
  - applications
    - business costs 338
    - discovering 125
    - virtual 390
  - archive destination 287
  - archive mode
    - setting 284
  - arranging
    - elements 327
  - array group 781
  - asset attributes
    - elements 392
-



- viewing 720
- asset based
  - reports 756
- asset information
  - adding 724
- Asset Management 772
- asset record
  - creating 717
- asset-based chargeback
  - setting up 733
- assigning
  - business costs 338
  - costs 716
  - Custom Name 377
  - e-mail schedules 571
  - MIME types 103
- authentication errors
  - SNMP 810
- automatically
  - clearing events 222
  - deleting events 222
- Automators 772
- available pools 436
- Average Blocks Requested 642
- Average Delivery Time 642

**B**

- back-end SCSI controller 781
- backing up
  - database 275, 277
- bar chart 632
- binding
  - persistent 839
- block size
  - cache 470
- Bridge Agent 47
- Brocade
  - switches 548
  - tracing 218
- Brocade Rapid program 89
- Brocade switches 89
  - discovering 37
  - events 550
- browsing
  - hosts 366

- switches 366
- Buffer Hit Ratio 642
- building
  - topology 34, 89
- business cost
  - setting 317
- business costs
  - assigning 338
- Business Tools 772
  - risk Analysis 777
- buttons
  - remote console 362, 364
  - right 667
  - System Explorer 300
- Bytes Received 642
- Bytes Transmitted 642

**C**

- cache
  - refreshing reports 245
- cache block size
  - changing 470
- cache read-ahead multiplier 433
- cache settings
  - modifying 433, 470
- calculating
  - fixed declining balance 747
  - straight line depreciation 746
- calculations for
  - double declining balance 749
- capacity charts 676
- Capacity Explorer
  - about 660
  - capacity charts 676
  - Capacity tab 660
  - finding capacity 667
  - global view 663
  - HDS 667
  - legend 660
  - magnifying 663
  - Path tab 660
  - printing 663, 674
  - toolbar 663
  - trending 678
  - utilization reports 673

- Utilization tab 660
  - Capacity tab 660
  - categorizations
    - storage 716
  - certificate
    - installing 17
  - changing
    - archive destination 287
    - cache block size 470
    - cache settings 433
    - custom commands 358
    - database 34
    - database passwords 279
    - date 239, 259
    - discovery policies 700
    - domain controller 125, 148
    - DPI settings 202
    - e-mail address 166
    - e-mail schedules 242
    - event policies 703
    - external tools 366
    - fabric name 317, 349
    - frequency 632
    - full name 166
    - logging 218
    - login name 166
    - number of retries 62, 805
    - Oracle Listener Password 288
    - organizations 176
    - password 91, 149, 164, 165, 166
    - phone number 166
    - properties 214
    - provisioning policies 702
    - roles 171
    - SNMP trap listener 46
    - status 718
    - time 239, 259
    - time-out period 62, 805
    - TNS Listener Port 147
    - tools 366
    - topology 663
    - user account 163
    - user name 91
    - user preferences 167
    - user profile 166
    - utilization policies 698
    - zone set 419
  - changing ports
    - zone alias 410
  - Chargeback 772
    - about 714
    - asset information 724
    - asset record 717
    - asset-based 714, 733
    - custom information 729
    - custom properties 729
    - customizing filters 767
    - department views 753
    - double declining balance 749
    - e-mail schedule 759
    - e-mailing reports 758
    - filtering assets 767, 768
    - fixed declining balance 747
    - general information 727
    - geographic information 729
    - hiding filters 768
    - infrastructure cost 731
    - licensing information 729
    - reports 756
    - saving 719
    - schedules editing 761
    - selecting element type 766
    - setting up 716, 733, 739
    - staff information 728
    - storage subsystem reports 756
    - storage tiers 721
    - storage-based 714, 739
    - straight line depreciation 746
    - viewing by element 752
    - viewing by enterprise 754
    - viewing report schedules 763
    - warranty information 729
  - chart
    - bar 632
    - clearing 632
    - filtering data 632
    - line 632
    - printing 632
    - saving 632
    - summary 632
-

- time period 632
- charts
  - custom periods 640
  - removing data 640
  - summary 638
  - thresholds 640
  - viewing 638
- checker
  - scan 230
- checking
  - database 279
- child organizations 154
- choosing
  - fabrics 340
- CIM
  - log levels 218
- CIM Extension
  - port 789
- CIM Extensions 1
- CIM Object Manager 838
  - tracing 218
- cimom.CimXmlClientHttpConnectTimeout 821
- cimom.ClariionEventPollInterval 546
- cimom.emc.skipRefresh 97
- cimom.hds.exclude 73
- cimom.log 215
- cimom.provider.hds.hiCommandTimeout 465
- cimom.symmetrix.exclude 68
- CIO role 154
- Cisco SMI-S switches
  - troubleshooting 803
- CISCO switches
  - topology 299
  - VSAN 299
- CLARiiON 427, 462
  - polling interval 546
- CLARiiON storage systems
  - no data 821
- cleanup
  - reports 245
- clearing
  - elements 32
  - events 222, 224, 391, 542, 543
  - graph 632
  - clearing events 222
- closing
  - Topology windows 329
- cluster 781
- CNT
  - switches 43
- cold backup 286
- collectors
  - for reports 394
  - managing 256
  - performance 259
  - reports 236, 239
  - setting time 239
  - starting 226, 236, 565
  - stopping 226, 236
- command window
  - log file 218
- commands
  - custom 355, 358
  - remote 317
- Common Information Model (CIM) 837
- comparing
  - element performance 637
- components 3
- configuring
  - business cost 317
  - Chargeback 716, 733, 739
  - e-mail notification 207
- console
  - buttons 362
  - remote 362
- controller
  - removing 125
- conventions
  - naming zones 525
- copying
  - zone sets 420
- copying text 366
- costs
  - assigning 716
  - business 338
- CR 365
- CRC Errors 642

- creating
    - asset record 717
    - custom commands 355
    - custom name 377
    - custom reports 583
    - departments 730
    - discovery policies 690, 692
    - e-mail schedule 571
    - event policies 694
    - global reports 248
    - host security groups 454
    - journal entries 546
    - LUSE volumes 465
    - new password 165
    - organizations 173
    - roles 170
    - storage pool 434
    - storage volume 441
    - topology 21
    - user accounts 162
    - utilization policies 683
    - virtual application 317
    - virtual applications 390
    - volumes 462
    - zone alias members 409
    - zone in fabric 413
    - zone sets 416, 417
  - CRLF 365
  - cumulative licenses 200
  - Current View combo box 663
  - custom
    - periods 640
  - custom commands
    - about 354
    - adding 355
    - deleting 358
    - editing 358
    - setting up 317
    - stopping 355
  - custom name
    - truncated 344
  - custom properties 729
  - custom reports 579
    - integrating 592
  - customer information
    - adding 729
  - Customizing 214
  - customizing
    - element filters 554
    - filters 552, 554
    - properties 214
    - severity filters 552
  - cut off
    - custom name 344
  - CXFS 396, 397
  - D**
  - DA 781
  - data
    - backing up 277
    - performance 632
    - trending 678
  - data aging statistics 243
  - data collection 226
  - database
    - AIX 810
    - backing up 275, 277
    - changing passwords 279
    - checking status 279
    - cold backup 284, 286
    - Database Admin Utility 281
    - export 280
    - reinitializing 282, 290
    - restarting 283
    - RMAN restore 285
    - updating 34
  - Database Admin Utility
    - accessing 278
    - archive destination 287
    - archive mode 284
    - changing passwords 279
    - cold backup 286
    - database passwords 279
    - exporting database 280
    - refreshing 279, 288
    - reinitializing database 282
    - resetting logs 289
    - resetting tablespace 283
    - restarting database 283
    - restoring cold backup 284
-

- RMAN restore 285
- temp tablespace 283
- viewing logs 288
- database connection failed
  - error 807
- date
  - setting 239, 259
- DCOM
  - unable to communicate 809
- deactivating
  - policies 707
  - zone sets 416
- definitions
  - ports 372
- deleting
  - custom commands 358
  - departments 731
  - domain controller 125
  - element 317
  - elements 32, 92, 336
  - e-mail schedules 577, 763
  - events 225, 391, 544, 545
  - fabrics 317, 349
  - host security groups 460
  - license 202
  - multiple jobs 525
  - organizations 177
  - performance data 640
  - policies 707
  - roles 172
  - storage pool 437
  - storage pools 436
  - switches 102
  - TNS Listener Port 147
  - user accounts 165
  - volumes 445, 462
  - zone aliases 411
  - zone sets 416, 419, 830
  - zones 415, 830
- deleting events 222
- department ownership 731
  - removing 744
- departments
  - adding 730
  - deleting 731
  - editing 730, 731
  - removing 731
- designing
  - custom reports 583
- detailed
  - tracing 218
- detailed schema 595
- details
  - event 541
  - obtaining 94
- detecting
  - IP range 28
  - McDATA switches 101
  - switches 101
- device 781
  - about 838
  - difference in meaning 781
  - finding 461
  - inaccessible 461
- device issues 831
- devices
  - deleting 92, 336
  - updating 335
- different
  - Java applet 828
- different version
  - Java applet 828
- direct attached storage 654
- directory lookup 642
- disabling
  - discovery schedule 210
  - schedule 210
- discovered
  - hosts 326, 345
  - storage systems 347
- discovered address
  - modifying 91
- Discovered Elements
  - deleting elements 336
- discovered elements
  - deleting elements 92
- discovering
  - applications 125
  - Brocade switches 37, 89
  - CNT switches 43

- DNS Name 29
  - EMC Solutions Enabler 66
  - event details 541
  - HDS storage systems 72
  - HDS systems 73
  - HP XP storage systems 74, 75
  - IBM storage systems 78
  - IP address 29
  - McDATA switches 47
  - Microsoft Exchange 125, 142, 811
  - NetApp filers 86
  - new elements 108
  - Oracle 125, 126
  - Oracle clusters 126
  - passwords 26
  - policies 700
  - SQL servers 132
  - storage system 21
  - storage systems 70, 80
  - Sun StorEdge storage systems 80, 82
  - Sun StorEdge switches 45
  - switches 21, 37
  - Sybase 125, 138
  - Symmetrix systems 68
  - troubleshooting 811, 813, 819, 835, 836
  - user names 26
  - discovery
    - authentication errors 810
    - quarantine 105, 106
    - time-out 821
    - troubleshooting 807
    - Windows proxy 114
  - discovery groups 91
  - discovery policies
    - creating 690
  - discovery requirements
    - Oracle 811
  - discovery schedule
    - adding 209
    - disabling 210
    - editing 211
    - removing 211
  - discovery settings
    - importing 32
    - saving 33
  - disk director 781
  - disk drive 150, 214, 781, 821
  - disk extent 781
  - Disk Read 642
  - disk space monitoring
    - results of 213
  - Disk Utilization 642
  - Disk Write 642
  - displaying
    - deleted Oracle instances 811
    - element home page 317
    - element impact 336
    - element properties 374
    - filters 556
    - log file 218
    - port details 317
  - DNS 811
  - documentation
    - finding liv
  - Domain Administrator role 154
  - domain controller
    - access 148
    - accessing 125, 148
    - removing 125
  - domain controller access 125, 148
  - double declining
    - calculating 749
  - DPI settings 202
  - dragging
    - elements 663
  - drivers
    - fixing 831
  - drives
    - Microsoft Exchange 811
    - uninitialized 821
  - duplication
    - zone sets 423
    - zones 423
  - E**
  - edit
    - event policies 703
  - editing
    - cache settings 433
-

- custom commands 358
- date 239
- departments 730, 731
- discovery policies 700
- discovery schedule 211
- e-mail address 166
- e-mail schedule 574
- e-mail schedules 242, 761
- fabric name 317, 349
- frequency 632
- full name 166
- host security groups 457
- logging 218
- login name 166
- organizations 176, 177
- ownership percentage 743
- password 164, 165, 166
- phone number 166
- provisioning policies 702
- roles 171
- status 718
- time 239
- user account 163
- user preferences 167
- user profile 166
- utilization policies 698
- zone set 419
- zone sets 416
- editing ports
  - zone alias 410
- EFC Manager 47, 830
- element
  - about 838
  - asset attributes 392
  - deleting 317
  - impact 336
  - performance 632
  - refresh 317
- element created in Chargeback
  - about 838
- element details
  - obtaining 94
- element home page
  - viewing 317
- element name
  - filtering 555
- element performance
  - comparing 637
- element properties
  - viewing 374
- element topology 317
- element type
  - filtering 554
  - selecting 553
- elements
  - adding 173, 176
  - arranging 327
  - capacity 667
  - deleting 32, 92, 336
  - dragging 663
  - getting details 34
  - managing 176
  - modifying 91
  - organization 176
  - removing 177
  - topology 89, 327
  - unable to find 813, 819
  - unmanaged 838
  - updating 335
- e-mail
  - notification 708
  - reports 570
  - schedules 242, 570
- e-mail address
  - changing 166
- e-mail notification
  - configuring 207
- e-mail schedule
  - adding 571
  - editing 574
  - organizations 571
  - viewing 763
- e-mail schedules 579
  - deleting 577, 763
  - editing 761
  - viewing 578
- e-mailing
  - chargeback reports 758
  - reports 759
- EMC CLARiiON 70

- EMC Navisphere Agent 462
  - EMC Solutions Enabler 66
  - EMC Symmetrix 427, 463
  - Empty Chart message 676
  - entries
    - journal 546
  - error
    - database connection failed 807
    - error -56 830
  - Error 503 787
  - error message
    - exclusive lock 835, 836
  - errors
    - authentication 810
  - event
    - details 541
    - icons 341
  - Event Manager
    - about 536
    - clearing 542
    - clearing events 222
    - deleting events 222
    - events 542, 545, 548, 552
    - filtering 552
    - journal entries 546
    - SNMP traps 536
    - sorting 545
    - switches 548
  - event policies
    - creating 694
  - event severity
    - setting 708
  - event types
    - all 553
    - all but management server 553
    - applications 553
    - fabrics 553
    - hosts 553
    - management servers 553
    - other 553
    - selecting multiple 554
    - storage systems 553
    - switches 553
  - events
    - clearing 222, 224, 391, 542, 543
    - deleting 225, 391, 544, 545
    - excluding 550
    - filtering 550, 552, 555
    - filters 391
    - managing 550
    - policies 703
    - polling 546
    - removing 542
    - severity 391
    - severity levels 551
    - sorting 391, 545
    - switches 548
    - unclearing 391, 543
    - viewing 391
  - Events tab
    - about 391
  - exceptions 828
  - excluding
    - events 550
    - HDS systems 73
    - switches 59
    - Symmetrix systems 68
  - exclusive lock
    - error message 835, 836
  - export
    - database 280
  - external tools 366
    - setting up 366
  - F**
  - FA 781
    - port 781
  - fabric
    - name 349
    - properties 376
  - fabrics
    - changing name 317
    - filtering 340, 663
    - removing 317, 349
    - viewing 376
  - fibre channel director 781
  - file extension
    - assigning 103
  - File Server SRM 838
    - properties 228
-



- scan checker 230
- File SRM scan 231
- file systems 226
  - about 397
- files
  - all logs 216
- filtering
  - assets 767
  - element type 554
  - events 550, 552, 555
  - fabrics 340, 663
  - global reports 569
  - organizations 178
  - volumes 440
- filtering data 632
- filters
  - Chargeback 767
  - hiding 556, 768
  - host dependency 528
  - LUN dependency 529
  - showing 556
  - storage system dependency 528
  - volume dependency 528, 529
  - zone dependency 530
- Final Destination 642
- finding
  - applications 125
  - capacity 667
  - devices 461
  - documentation liv
  - element impact 336
  - event details 541
  - hosts 125
  - IP address 29
  - IP range 28
  - new elements 108
  - performance data 632
  - storage systems 21
  - switches 21
  - trending information 678
- first time
  - users 3
- fixed declining
  - balance 747
- fixing

- drivers 831
- formatting
  - reports 566
- forwarding
  - SNMP traps 536
- Free Physical Memory 642
- frequency
  - changing 632
- front-end fibre channel port 781
- front-end SCSI controller 781
- front-end storage system
  - processor 781
- full name
  - changing 166

## **G**

- general information
  - adding 727
- generic
  - hosts 326
- geographic information
  - adding 729
- Get Details
  - email notification 804
- getting
  - element details 94
- getting details 34, 94
  - applications 125
  - hosts 125
- Global Reporter server
  - about 839
- global reporting view 838
- global reports
  - filtering 569
  - setting up 248
  - tnsnames.ora file 254
- global view 663
  - using 330, 386
- graph
  - bar 632
  - clearing 632
  - filtering data 632
  - line 632
  - printing 632
  - saving 632

- time period 632
- graphs
  - saving 632
- grey screen 298
- grouping
  - hosts 345
  - storage systems 347
- groups
  - about 345
- H**
- hard zone 839
- HBA Risk Analysis 777
- HBAs
  - replacing 773
  - swapping 831
- HDS 427
  - capacity definitions 667
  - timeout 465
- HDS 9200 465
- HDS 9900 465
- HDS 9900V 465
- HDS storage systems
  - discovering 72
- HdsSkipRefresh 98
- Help button
  - remote console 364
- Help Desk role 154
- hiCommandTimeout 465
- hiding
  - filters 556, 768
- hiding tabs 340
- hierarchy
  - organizations 154
- host
  - not in topology 813, 819
- host bus adapter
  - unable to detect 820
- host bust adapters
  - replacing 773
- host dependency
  - host 528
- host persistent binding 839
- host security groups
  - accessing 452

- creating 454
- deleting 460
- editing 457
- hosts
  - adding 491, 507, 519
  - browsing 366
  - discovered 326
  - discovering 125
  - generic 326
  - grouping 345
  - removing 32
  - telnet to 366
  - ungrouping 346
- hot-swapped
  - drives 821
- HP XP storage systems 74, 75
- HTTP Error 503 787
- HTTPS 17
- hyper 781

- I**
- IBM storage systems
  - discovering 78
- icons
  - event 341
- impact
  - element 336
  - showing 317
- importing
  - database 281
  - discovery settings 32
  - license 200
- importing database 281
- inaccessible
  - device 461, 829
- increasing
  - memory 828
- information
  - obtaining element 94
  - performance 632
- infrastructure cost
  - setting 731
- initiator WWN 839
- In-memory Sort Ratio 642
- inode file

---

- about 839
- installing
  - Java plug-in 16
  - security certificate 17
- integrating
  - custom reports 592
- interface adapter 781
- internal
  - drives 821
- intervals
  - events polling of 546
- Invalid CRC Errors 642
- IP range 28
- issues
  - devices 831
- J**
- Java applet
  - different version 828
- Java plug-in 828
  - installing 16
- jboss.properties
  - modifying 214
- jobs
  - provisioning 524
- journal entries 546
- K**
- keeping active
  - remote console 363
- killing
  - commands 355
  - custom commands 355
- L**
- layout
  - arranging 327
- LDEV 781
- levels
  - event 548
  - severity 551
- LF 365
- Library Cache Hit Ratio 642
- license 195
  - deleting 202

- importing 200
  - viewing 200, 201
- licensing information
  - adding 729
- line chart 632
- Link Failures 642
- List tab 300, 304
- local drives 811
- locating
  - storage systems 21
  - switches 21
- log file
  - displaying 218
- log messages
  - viewing 62
- logging
  - changing 218
  - viewing 221
- logging off 20
- logical device 781
- Logical Unit Size Expansion 781
- login name
  - modifying 166
- logs
  - accessing 216
  - refreshing 288
  - resetting 289
  - saving 216
  - viewing 288
- LUN numbering
  - requirements 516
- LUN Security
  - Path Provisioning 488
- LUSE 781
- LUSE volumes
  - cannot create 465
- M**
- magnifying
  - Capacity Explorer 663
  - topology 663
- MALs 195, 839
- Managed Access Points
  - about 839
- managed access points 195, 839

- managed application license 195, 839
  - managed object 839
  - management server
    - about 1
    - accessing 14
    - assigning a name 377
    - backup 277
    - components 3
    - security 154
    - stopping 19
    - variables 358
  - managing
    - collectors 236, 259
    - elements 173, 176, 177
    - events 550
    - performance collectors 256
    - switches 101
  - mapped 840
  - MAPs 195, 839
    - about 839
  - materialized view 839
    - about 839
  - materialized views
    - previous 626
  - maximizing
    - screen space 567
  - maximum thresholds 640
  - McDATA switches 820
    - adding 101
    - discovering 47
  - members
    - zone sets 416
  - memory
    - increasing 828
  - messages
    - data is late 786
    - OutOfMemoryException 828
  - messages waiting 642
  - meta device 840
  - meta volume 840
  - metadata client 397
  - metadata server 397
  - metadevice 781
  - Microsoft Exchange
    - discovering 125, 142, 811
    - drive M 811
  - Microsoft Exchange Services 642
  - MIME types 103
  - minimum
    - thresholds 640
  - missing element 840
  - mixed mode authentication 133
  - modifying
    - cache block size 470
    - cache settings 433, 470
    - Chargeback 716
    - custom command 358
    - database 34
    - database passwords 279
    - date 239, 259
    - departments 730
    - discovered address 91
    - discovery IP address 31
    - discovery policies 700
    - DNS name for discovery 31
    - domain controller 125, 148
    - DPI settings 202
    - elements 91
    - e-mail address 166
    - e-mail schedule 574
    - e-mail schedules 242
    - event policies 703
    - external tools 366
    - fabric name 317, 349
    - filters 767, 768
    - frequency 632
    - full name 166
    - jboss.properties 214
    - logging 218
    - login name 166
    - Oracle Listener Password 288
    - organizations 176
    - ownership percentage 743
    - password 91, 149, 164, 165, 166
    - phone number 166
    - properties 214
    - provisioning policies 702
    - roles 171
    - schedules 240
    - SNMP trap listener 46
-

- status 718
- time 239, 259
- TNS Listener Port 147
- topology 663
- user account 163
- user name 91
- user preferences 167
- user profile 166
- utilization policies 698
- zone members 415
- zone set 419
- zone sets 416
- modifying ports
  - zone alias 410
- monitoring
  - provisioning jobs 524
- monitoring options
  - about 642
- moving
  - elements 327, 663
  - groups 327
  - topology 663
- multipathing 378, 840

## **N**

- name
  - fabric 349
  - filtering 555
  - storage pool 436
  - storage tiers 716
  - zone sets 416
- naming
  - storage tier 721
- naming conventions
  - for zones 525
- naming organizations 154
- navigating elements 367
- navigation details 317
- Navigation screen
  - about 367
  - accessing 373
- Navigation tab
  - duplication 423
- Navisphere 366
- Navisphere Agent 462

- NetApp filers
  - discovering 86
- netcnfg 66
- nethost 66
- new elements
  - adding 108
- new password 165
- new window 566
- New Window option 386
- no data
  - CLARiiON storage systems 821
- nonexistent IP addresses 821
- nonexistent Oracle instances 811
- notification
  - e-mail 207, 708
- number of retries
  - changing 62, 805

## **O**

- obtaining
  - performance data 632
  - security certificate 17
  - topology information 89
  - utilization reports 673
- old reports 245
- opening
  - new window 386, 663
- opening reports
  - new window 566
- Oracle
  - deleted instances 811
  - discovering 125, 126
  - discovery requirements 811
- Oracle Listener Password 288
- Oracle TNS Listener Port 147
- organizations
  - about 154
  - adding 173
  - deleting 177
  - editing 176, 177
  - elements 173, 176, 177
  - e-mailing reports 571
  - filtering 178
  - properties 169
  - users 169

- viewing 175
- organizing
  - topology 345
- outdated reports 245
- OutOfMemoryException 828
- overview
  - SAN zoning 401, 405
- owners
  - adding 730
  - editing 730
  - removing 731
- ownership
  - percentage 743
  - removing 731

## P

- Pan button 663
- parent organizations 154
- Parse CPU 642
- password
  - changing 91, 149, 163, 164, 165, 166
- path information
  - unable to find 819
- Path Provisioning
  - adding hosts 491, 507, 519
  - deleting jobs 525
  - executing 523
  - filters 529
  - LUN Security 488
  - monitoring jobs 524
  - storage system dependency
    - filter 528
  - Volume Creation 486
  - Volume Creation and LUN Security 499
  - Volume Creation, LUN Security, and Zone Operation 479, 516
  - volume dependency filter 528, 529
  - zone dependency filter 530
  - Zone Operation 493
- Path tab 300, 315, 660
- paths 378
- PDEV 781
- percentage
  - ownership 743
- performance
  - collectors 259
  - finding 632
  - statistics 261
- performance collectors
  - managing 256
  - starting 258
  - stopping 259
- Performance Explorer
  - about 632
  - accessing 632
  - chart format 632
  - charts 632
  - clearing graph 632
  - comparing 637
  - custom periods 640
  - filtering data 632
  - frequency 632
  - Java plug-in 16
  - multiple elements 632
  - printing 632
  - removing data 640
  - saving graph 632
  - summary charts 638
  - thresholds 640
  - time period 632
  - trending 632
- periods
  - custom 640
- persistitent binding 839
- phone number
  - editing 166
- physical device 781
- Physical Memory Used 642
- planning organizations 154
- points
  - managed access 195
- policies
  - creating 683
  - deactivating 707
  - deleting 707
  - discovery 690, 700
  - event 694, 703
  - provisioning 692, 702
  - testing 683, 697

---

- utilization 683, 698
- viewing 682, 706
- Policies tab 396
- Policy Manager
  - about 682
  - accessing 682
  - creating policies 682
  - deactivating policies 682, 707
  - deleting policies 682, 707
  - modifying policies 682
  - viewing policies 682, 706
- polling
  - events 546
- pool
  - information 436
  - storage 434
- pools
  - accessing 436
  - deleting 436
- port 781
  - CIM Extension 789
- port controller 781
- port details
  - showing 317
- ports
  - definitions 372
  - viewing 335
- predicting 632
  - performance 632
  - storage capacity 678
- Pre-submission Queue Size 642
- previous materialized views 626
- printing
  - Capacity Explorer 663, 674
  - graphs 632
  - System Explorer 330
  - topology 387, 663
- privileges
  - roles 154
- problems
  - drivers 831
- process
  - exclusive lock 835, 836
- Processor Utilization 642
- profile

- user 166
- properties
  - customizing 214
  - fabric 376
  - File Server SRM 228
  - organizations 169
  - roles 169
  - viewing 374
- provider 840
- provisioning
  - about 399
  - adding hosts 491, 507, 519
  - copying zone sets 420
  - LUN numbering 516
  - monitoring jobs 524
  - policies 702
  - troubleshooting 461, 462, 463, 829, 830, 835, 836
- provisioning jobs
  - scheduling 521
- provisioning policies
  - creating 692
- Provisioning tab 391
- proxy server 840

## Q

- quarantine
  - adding elements 105
  - clearing elements 106
- queue size 642

## R

- RAID level 436
- Rapid program 89
- read caching 433
- Read Operations 642
- Read Requests 642
- recalculating topology 6
- Receive Queue Size 642
- receiving
  - SNMP traps 536
- refresh
  - element 317
- refresh view 565
- refreshing

- Database Admin Utility 279
  - element properties 374
  - logs 288
  - report cache 245
  - reports 565
  - Symmetrix systems 97
  - reinitializing
    - database 282, 290
    - warnings 290
  - remote
    - commands 317
    - console 362
  - remote console
    - buttons 364
    - copying text 366
    - keeping active 363
    - menu options 365
  - remote drives 811
  - remote scripts
    - running 710
  - remote sites
    - filtering 569
  - removing
    - department ownership 744
    - departments 731
    - discovery schedule 211
    - domain controller 125
    - element 317
    - elements 32, 92, 176, 177, 336
    - e-mail schedules 577, 763
    - events 225, 391, 542, 544, 545
    - fabrics 317, 349
    - host security groups 460
    - license 202
    - multiple jobs 525
    - organizations 177
    - ownership 731
    - performance data 640
    - policies 707
    - roles 172
    - schedules 240
    - storage pool 437
    - storage pools 436
    - switches 102
    - TNS Listener Port 147
    - user accounts 165
    - volumes 445
    - zone aliases 411
    - zone members 415
    - zone sets 419, 830
    - zones 830
  - removing ports
    - zone alias 410
  - renaming
    - storage tiers 716
  - replacing
    - fabric name 349
    - HBAs 773
    - switches 102
  - replication 34
  - Reporter
    - about 560
    - custom 579
  - reports
    - asset based 756
    - Chargeback 759
    - cleanup 245
    - collectors 236
    - custom 579
    - data aging statistics 243
    - deleting e-mail schedules 577
    - designing 583
    - detailed schema 595
    - e-mail 570
    - e-mail schedule 579
    - e-mail schedules 240, 242, 571, 578
    - e-mailing 758
    - formatting 566
    - global 248, 569
    - HTML 566
    - integrating 592
    - Microsoft Excel 566
    - new window 566
    - organizations 571
    - PDF 566
    - refreshing 565
    - refreshing cache 245
    - screen space 567
    - sending 570
    - starting collectors 565
-



- storage based 756
- storage subsystem 756
- utilization 673
- XML 566
- Reports tab
  - about 394
- Requests Serviced 642
- required steps liii
- resetting
  - tablespace 283
- restarting
  - database 283
- restoring
  - cold backup 284
  - database 279
  - RMAN 285
- right button 667
- risk analysis
  - setting up 777
- RMAN
  - destination 287
- RMAN backup
  - viewing results 277
- RMAN backups 275
- RMAN restore 285
- roles
  - about 154
  - adding 170
  - Application Administrator 154
  - CIO 154
  - deleting 172
  - Domain Administrator 154
  - editing 171
  - Element Control privilege 154
  - Full Control privilege 154
  - Help Desk 154
  - privileges 154
  - properties 169
  - Server Administrator 154
  - Storage Administrator 154
  - users 169
  - View privilege 154
- running
  - cold backup 286
  - remote scripts 710

Rx Bytes 642

## **S**

- SAN 840
- SAN Zoning Overview 401
- saving
  - chargeback information 719
  - discovery settings 33
  - graph 632
  - graphs 632
  - logs 216
  - settings to a file 33
  - topology 300, 663
- scan checker 230
- scan status 230
- scanning
  - IP range 28
- schedules
  - deleting 577
  - disabling 210
  - editing 761
  - e-mail 242, 570, 578
  - modifying 240
  - removing 240
- scheduling
  - database backup 275
  - discovery 209, 210, 211
  - File Server SRM scan checker 230
  - provisioning jobs 521
  - report cleanup 245
  - reports 759
- screen space
  - maximizing 567
- scripting
  - variables 358
- scripts
  - remote 710
  - stopping 355
- searching
  - topology 663
- security
  - Management server 154
  - roles 170, 171
- security certificate
  - installing 17

- seeing
    - license 200, 201
  - selecting
    - element type 553
    - fabrics 340
    - severity levels 551
  - Send Queue Size 642
  - sending
    - reports 570
  - server
    - accessing 14
    - backup 277
  - Server Administrator role 154
  - services
    - stopping 19
  - sets
    - zone 416
  - setting
    - archive mode 284
    - business cost 317
    - date 239, 259
    - discovery passwords 26
    - discovery user name 26
    - event severity 708
    - infrastructure cost 731
    - report collectors 394
    - storage cost 721
    - storage tier 721
    - thresholds 640
    - time 239, 259
    - time period 632
  - setting up
    - Chargeback 716, 733, 739
    - e-mail notification 207
    - external tools 366
    - global reports 248
    - remote commands 317
  - settings
    - cache 470
  - severity
    - events 391
  - Severity combo box 663
  - severity levels 341, 548
    - all 551
    - all but clear 551
    - clear 551
    - critical 551
    - informational 551
    - major 551
    - minor 551
    - unknown 551
    - warning 551
  - showing
    - element impact 336
    - filters 556
    - impact 317
    - port details 317
  - signing out 20
  - size
    - cache block 470
    - storage pool 436
  - SMI-S support
    - troubleshooting 802
  - SMI-S switches
    - see the support matrix 37, 802
  - SMTP 207, 841
  - SNIA specification 841
  - SNMP
    - authentication errors 810
    - trap forwarding 536
  - SNMP trap destinations 206
  - SNMP trap listener
    - changing 46
  - soft zone 841
  - sorting
    - events 391, 545
  - specifying
    - custom periods 640
  - SQL Server
    - authentication modes 133
  - SQL servers
    - discovering 132
  - staff information
    - adding 728
  - Standard XML Connector 1
  - starting
    - collectors 226, 236, 565
    - performance collectors 258
  - Statistics 214
    - data aging 243
-

- performance 261
- statistics 150
- status
  - changing 718
  - scan 230
- stop button
  - remote console 364
- stopping
  - collectors 226, 236
  - performance collectors 259
  - SAN details 96
  - scripts 355
  - services 19
- storage
  - assigning costs 716
- Storage Administrator role 154
- storage area network 840
- storage array 781
- storage based
  - reports 756
- storage capacity
  - predicting 678
- storage cost
  - setting 721
- storage elements
  - viewing 324
- storage pool 781, 841
  - creating 434
  - deleting 437
- storage pools
  - accessing 436
  - deleting 436
- storage subsystem based
  - reports 756
- storage system 781
- storage system proxy 840
- storage systems 80, 150, 214
  - discovering 21, 80
  - grouping 347
  - polling interval 546
  - removing 32
  - ungrouping 348
- storage tier
  - setting 721
- storage tiers 716
- storage-based
  - Chargeback 714
- storage-based chargeback
  - setting up 739
- Store Group Size 642
- Store Size 642
- straight line depreciation
  - calculating 746
- suggested topics 3
- summary text
  - filtering 555
- Sun StorEdge
  - SNMP trap listener 46
- Sun StorEdge A5000 820
- Sun StorEdge storage systems 80, 82
- Sun StorEdge switches 45
- swapped
  - drives 821
- swapping
  - switches 102
- swapping HBAs 831
- switch
  - events 548
- switch port zoning 841
- Switch Risk Analysis 777
- switches
  - adding 101
  - browsing 366
  - CISCO 299
  - discovering 21, 37
  - duplication 423
  - excluding 59
  - managing 101
  - McDATA 47, 101, 820
  - name truncated 344
  - number of retries 62, 805
  - port status 372
  - removing 32, 102
  - replacing 102
  - telnet to 366
  - time-out period 62, 805
  - unable to monitor 820
- Sybase
  - discovering 125, 138
- System Event Times 642

- System Explorer
    - about 296
    - accessing 298
    - buttons 300
    - can't access 828
    - Cisco switches 299
    - deleting elements 92, 336
    - Discovered hosts 326
    - element impact 336
    - event icons 341
    - global views 330, 386
    - groups 345
    - hiding tabs 340
    - Java plug-in 16
    - menu options 317
    - printing 330
    - scripting variables 358
    - storage elements 324
    - Topology windows 329
    - update 335
    - user interface 300
    - viewing ports 335
    - virtual applications 390
  - T**
  - table
    - terms 781
  - tablespace
    - resetting 283
  - tabs
    - Access 307
    - hiding 340
    - List 304
    - Path 315
  - task
    - running on element 653
  - telnet to
    - hosts 366
    - switches 366
  - terms
    - table 781
  - testing
    - policies 683, 697
  - text
    - filtering 555
    - text copying 366
  - threads 228
  - thresholds 640
  - tiers
    - storage 716
  - time
    - setting 239, 259
  - timeout
    - HDS 465
  - time-out period
    - changing 62
  - TNS Listener Port
    - changing 147
  - tnsnames.ora file 254
  - toolbar
    - Capacity Explorer 663
    - System Explorer 300
  - tools
    - external 366
  - topology
    - AIX 810
    - building 89
    - changing 663
    - Cisco switches 299
    - filtering 663
    - host not appearing 813, 819
    - moving 663
    - new window option 663
    - organizing 345
    - printing 387, 663
    - saving 663
    - updating 6
    - VSANs 299
  - topology issues 813
  - Topology tab
    - about 378
  - Total Cost/Month for Each Departments 756
  - tracing
    - detailed 218
  - trap forwarding 536
  - trap generation 206
  - traps
    - SNMP 536
  - trending
-

- Capacity Explorer 678
  - reports 560
- triangle
  - icons 341
- troubleshooting
  - Cisco SMI-S switches
    - switch names 803
  - CLARiiON storage systems 462
  - discovery 807
  - discovery and getting element
    - details 807, 809, 811, 813, 819, 835, 836
  - EMC storage systems 463
  - EMC Symmetrix storage systems 463
  - Engenio storage systems 470
  - HDS storage systems 465
  - JavaScript blockers 298
  - Microsoft Exchange 811
  - popups 298
  - provisioning 461, 462, 463, 829, 830, 835, 836
  - SMI-S switches 37, 802
- truncated
  - custom name 344
- Tx Bytes 642
- types of
  - monitoring 642
- U**
  - unable to
    - discover 807
  - Unable to access resource 829
  - unable to detect
    - host bus adapter 820
  - unable to find
    - elements 819
    - path information 819
  - unable to retrieve data 831
  - unclearing
    - events 391, 543
  - ungrouping
    - hosts 346
    - storage systems 348
  - uninitialized
    - drives 821
  - updating
    - database 34
    - element properties 374
    - elements 335
    - license 200
  - updating topology 6
  - used pools 436
  - user accounts
    - creating 162
    - deleting 165
  - user audit log 217
  - user guide
    - about li
  - user interface
    - System Explorer 300
  - user name
    - changing 91
  - user preferences
    - changing 167
  - user profile
    - modifying 166
- users
  - about 154
  - adding 162
  - first time 3
  - organizations 169
  - roles 169, 170, 171
- using
  - global view 330, 386
  - journal entries 546
- utilization
  - policies 697, 698
- utilization policies
  - creating 683
- Utilization tab 660
- V**
  - variables
    - scripting 358
  - viewing
    - all logs 216
    - asset attributes 720
    - capacity 667
    - capacity charts 676

- chargeback by element 752
- charts 638
- cumulative licenses 200
- data aging statistics 261
- database status 279
- department views 753
- element asset attributes 392
- element home page 317
- element impact 336
- element properties 374
- e-mail schedules 578, 763
- event details 541
- events 391
- fabrics 340, 376
- log file 218
- log messages 62
- logging 221
- logs 288
- organization properties 169
- organizations 175
- performance data 632
- policies 682, 706
- ports 335
- report refresh 565
- RMAN backup results 277
- schedule history 765
- scheduled e-mail deliveries 240
- security certificate 17
- specific license 201
- storage elements 324
- topology 21
- trending information 678
- utilization reports 673
- volume information 438
- zone aliases 408
- views
  - global 330, 386
- virtual application
  - about 841
  - adding 317, 390
  - creating 390
- Virtual Memory Used 642
- volume 781
  - meta 840
- volumes

- creating 462
- deleting 445, 462
- filtering 440
- unable to create 465
- viewing information 438
- VSAN 299

## **W**

- warnings
  - reinitializing 290
- warranty information
  - adding 729
- WBEM 842
  - about 842
- window
  - opening reports 566
- Windows Management
  - Instrumentation 842
- Windows proxy
  - discovery 114
- WinMgmt.exe 813
- WMI 842
  - about 842
- Worldwide Name 842
- WorldWide Name zoning 841
- Worldwide Name zoning 842
  - about 842
- write caching 433
- Write Operations 642
- WWN 842

## **X**

- XFS 396
- Xiotech storage systems 83

## **Z**

- zone
  - hard 842
  - soft 842
- zone alias
  - about 842
  - creating 409
  - deleting 411
  - modifying ports 410
- zone hard

---

- about 842
- zone member 842
- zone members
  - adding 415
  - removing 415
- zone set
  - creating 417
- zone sets
  - accessing 416
  - activating 421
  - copying 420
  - creating 416, 417
  - deleting 419, 830
  - duplication 423
  - editing 416
  - members 416
  - modifying 419
  - removing 416
- zones
  - creating 413
  - deleting 415, 830
  - duplication 423
  - naming conventions 525
  - removing 415
  - viewing aliases 408
- Zoning Library
  - duplication 423

---