



# Sun StorageTek™ Operations Manager 6.0 User Guide

---

Sun Microsystems, Inc.  
[www.sun.com](http://www.sun.com)

Part No. 817-7923-16  
January 2008, Revision A

Submit comments about this document at: <http://www.sun.com/hwdocs/feedback>

Copyright 2002-2007 Hewlett-Packard Development Company, L.P., 3000 Hanover Street, Palo Alto, California 94304, U.S.A. All rights reserved.  
Copyright 2002-2007 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries.

This document and the product to which it pertains are distributed under licenses restricting their use, copying, distribution, and decompilation. No part of the product or of this document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd. Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, Sun StorageTek, Java, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and in other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and in other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

U.S. Government Rights—Commercial use. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

Microsoft and Windows are registered trademarks of Microsoft Corporation. Oracle is a registered trademark of Oracle Corporation. AIX and IBM are registered trademarks of International Business Machines Corporation in the United States, other countries or both. McDATA is a registered trademark of McDATA Corporation. Engenio is a registered trademark of Engenio Corporation. CLARiiON is a registered trademark of EMC Corporation. 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. 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. HBAware is a trademark of Emulex Corporation.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

---

Copyright 2002-2007 Hewlett-Packard Development Company, L.P., 3000 Hanover Street, Palo Alto, Californie 94304, Etats-Unis. Tous droits réservés.  
Copyright 2002-2007 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, Californie 95054, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. a les droits de propriété intellectuelle relatants à la technologie qui est décrit dans ce document. En particulier, et sans la limitation, ces droits de propriété intellectuelle peuvent inclure un ou plus des brevets américains énumérés à <http://www.sun.com/patents> et un ou les brevets plus supplémentaires ou les applications de brevet en attente dans les Etats-Unis et dans les autres pays.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a.

Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, Sun StorageTek, Java, et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciées de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

Microsoft et Windows sont des marques déposées de Microsoft Corporation. Oracle est la marque déposée de Oracle Corporation. AIX et IBM des marques déposées de International Business Machines Corporation aux Etats-Unis et dans d'autres pays. McDATA est la marque déposée de McDATA Corporation. Engenio est la marque déposée de Engenio Corporation. CLARiiON est la marque déposée de EMC Corporation. SGI et IRIX des marques déposées de Silicon Graphics, Inc. Netscape est la marque déposée de Netscape Communications Corporation aux Etats-Unis et dans d'autres pays. HP, HP-UX, et OpenVMS, Tru64 UNIX des marques déposées de Hewlett-Packard Development Company. QLogic est une marque déposée de QLogic Corporation. Emulex est la marque déposée d'Emulex Corporation. HBAware est une marque déposée d'Emulex Corporation.

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.



Adobe PostScript

# Contents

---

## **1. Overview 1**

About This Product 1

Suggested Topics for First-Time Users 2

Product Components 3

    Management Server Components 3

The User Interface 5

    The Top Pane 5

    The Left Pane 7

        Opening and Closing the Left Pane 8

    The Home Page 9

Accessing the Management Server 11

Installing the Java Plug-in 12

Installing the Software Security Certificate 14

    Installing the Certificate by Using Microsoft Explorer 6.0 15

    Installing the Certificate by Using Firefox 1.5 16

    Changing the Security Certificate to Match the Name of the Server 16

Restarting the Service for the Management Server 18

Signing Out of the Management Server 19

<b>2. Discovering Switches, Storage Systems, NAS Devices, and Tape Libraries</b>	<b>21</b>
Discovery Steps	22
Overall Discovery Tasks	22
Overview of Discovery Features	24
Setting Default User Names and Passwords	24
Adding an IP Range for Scanning	27
Adding a Single IP Address or DNS Name for Discovery	29
Modifying a Single IP Address Entry for Discovery	30
Removing Elements from the Addresses to Discover List	31
Importing Discovery Settings from a File	31
Importing a File	32
Re-discovering the Management Server	32
Saving Discovery Settings to a File	33
Discover Switches	34
Discovering Brocade Switches	35
Migrating Brocade API Switches to SMI-S After Upgrading	35
To Discover Brocade SMI-S Switches	38
Discovering CNT Switches	39
Discovering Cisco Switches	40
Discovering Sun StorEdge, QLogic and HP StorageWorks M-Series for p-Class BladeSystems	42
Discovering McDATA and EMC Connectrix Switches	44
Discovering McDATA and Connectrix switches with SMI-S	45
Discovering McDATA and Connectrix Switches through a Proxy with SWAPI	47
Discovering McDATA and Connectrix Switches through a Proxy with SNMP	49
Discovering McDATA and Connectrix Switches through a Direct Connection and SNMP	51



Changing the Discovery Settings	52
Excluding McDATA and EMC Connectrix Switches from Discovery	53
Managing McDATA and EMC Connectrix Switches	54
Changing the SNMP Trap Listener Port and Community String for Switches Discovered with SNMP	57
Discover Storage Systems, NAS Devices and Tape Libraries	58
Discovering 3PAR Storage Systems	59
Discovering EMC Solutions Enabler	60
Excluding EMC Symmetrix Storage Systems from Discovery	61
Excluding EMC Symmetrix Storage Systems from Force Device Manager Refresh	62
Discovering EMC CLARiiON Storage Systems	63
Discovering LSI Storage Systems	64
Discovering HDS Storage Systems	65
Excluding HDS Storage Systems from Discovery	66
Excluding HDS Storage Systems from Force Device Manager Refresh	68
Discovering HP StorageWorks MSA Arrays	69
Discovering HP StorageWorks EVA Arrays	70
Obtaining SNMP Traps using Command View EVA	71
Discovering HP StorageWorks XP Arrays	74
Discovering HP XP Arrays by Using Command View XP and SMI-S	74
Discovering HP XP Arrays Using Command View XP Advanced Edition	75
Discovering HP XP Arrays by using the built-in XP Provider	76
Discovering IBM Storage Systems	77
Discovering Sun StorEdge 3510 Storage Systems	79
Discovering Sun StorEdge 6920 and 6940 Storage Systems	80
Discovering Sun StorEdge 6130 Storage Systems	81
Discovering Xiotech Storage Systems	82

Discovering HP NAS Devices on Windows	83
Discovering HP NAS Devices on Linux	84
Discovering NetApp NAS Devices	85
Enabling SSL Communication with a NetApp NAS Device	86
Discovering Sun NAS Devices	87
Discovering HP and IBM Tape Libraries	88
Building the Topology	88
Building the Topology View	89
Modifying the Properties of a Discovered Address	90
Get Details	91
About Get Details	91
Running Get Details	92
Stopping the Gathering of Details	93
Using Discovery Groups	93
Creating Custom Discovery Lists	94
Managing Discovery Groups	95
Moving Elements Between Discovery Groups	96
Deleting Elements from the Product	97
Deleting an Element Using System Explorer or Chargeback	97
Deleting Elements Using Discovery Step 2 (Topology)	98
Working with Quarantined Elements	99
Placing an Element in Quarantine	100
Removing an Element from Quarantine	100
Updating the Database with Element Changes	101
Notifying the Software of a New Element	102
Viewing Log Messages	103
Viewing the Status of System Tasks	103

### **3. Discovering Applications, Backup Hosts and Hosts 105**

Step 1 — Discovering Your Hosts and Backup Manager Hosts	105
Step A — Set Up Discovery for Hosts	107
Step B — Build the Topology	110
(Optional) Step C — View the Topology	111
Step D — Get Details	111
Step 2 — Setting Up Discovery for Applications	113
Creating Custom Passwords on Managed Database Instances	114
Monitoring Oracle	115
Step A — Create the APPIQ_USER Account for Oracle	115
Removing the APPIQ_USER Account for Oracle	117
Step B — Provide the TNS Listener Port	119
Step C — Set up Discovery for Oracle 10g	119
Discovering Oracle Real Application Clusters (RAC)	120
Deleting Oracle Application Information	123
Monitoring Microsoft SQL Server	123
Switching to Mixed Mode Authentication	124
Step A — Create the appiq_user Account for the Microsoft SQL Server	125
Step B — Provide the Microsoft SQL Server Name and Port Number	128
Removing the appiq_user Account for Microsoft SQL Server	130
Deleting Microsoft SQL Server Information	131
Monitoring Microsoft SQL Server Clusters	132
Monitoring Sybase Adaptive Server Enterprise	134
Step A — Create the APPIQ_USER account for Sybase	135
Removing the APPIQ_USER Account for Sybase	136
Step B — Provide the Sybase Server Name and Port Number	137
Deleting Sybase Information	138
Monitoring Microsoft Exchange	138
Adding Microsoft Exchange Domain Controller Access	138

Editing a Microsoft Exchange Domain Controller	140
Deleting a Microsoft Exchange Domain Controller	140
Monitoring Microsoft Exchange Failover Clusters	140
Monitoring Caché	141
Step A — Import the Wrapper Class Definitions into the Caché Instance	141
Step B — Create APPIQ_USER Account on the Caché Instance	146
Removing the APPIQ_USER Account from the Caché Instance	149
Step C — Provide the Caché Instance Name and Port Number	151
Deleting Caché Information	152
Step 3 — Discovering Applications	152
Step A — Detect Your Applications	153
Step B — Obtain the Topology	153
Step C — Run Get Details	154
Changing the Oracle TNS Listener Port	155
Changing the Password for the Managed Database Account	156
<b>4. Installing and Discovering the Windows Proxy</b>	<b>159</b>
Installing the Windows Proxy	160
Discovering the Windows Proxy	161
Configuring Windows Proxy Authentication	162
Decreasing the Maximum Java Heap Size	163
Removing the Windows Proxy	164
<b>5. Host and Application Clustering</b>	<b>165</b>
About Clustering	165
Discovering Clusters	165
Automatic Discovery of Host Clusters	166
Manual Discovery of Host Clusters	167
Filtering Hosts	168

File Servers and Clusters	169
Clustering in System Explorer	169
Clustering in Topology	171
Clustering in Capacity Manager	173
<b>6. Managing Security</b>	<b>175</b>
About Security for the Management Server	175
About Roles	176
About Organizations	179
Planning Your Hierarchy	181
Naming Organizations	182
Managing User Accounts	182
Adding Users	182
Editing a User Account	184
Changing the Password for a User Account	185
Changing Your Password	186
Deleting Users	186
Modifying Your User Profile	186
Modifying Your User Preferences	187
System, Capacity and Performance Manager Preferences	187
System Explorer and Element Topology Preferences	188
Warnings for Slow Systems Operations	188
Viewing the Properties of a Role	188
Viewing the Properties of an Organization	189
Managing Roles	189
Adding Roles	190
Editing Roles	190
Deleting Roles	191
Managing Organizations	192

Adding an Organization	192
Adding Storage Volumes to an Organization	194
Viewing Organizations	194
Editing an Organization	195
Removing an Organization	196
Removing Members from an Organization	196
Filtering Organizations	197
Changing the Password of System Accounts	198
Using Active Directory/LDAP for Authentication	200
Step 1 — Configure the Management Server to Use AD or LDAP	200
Configuring the Management Server to Use Active Directory	201
Configuring the Management Server to Use LDAP	205
Step 2 — Restart the AppStorManager Service and Login as the Designated Admin Account	209
Step 3 — Add Users to the Management Server	210
Step 4 — Provide Login Information to Your Users	210

## **7. Managing Licenses 213**

Modifying the License Summary Page	217
------------------------------------	-----

## **8. Configuring the Management Server 219**

Trap Generation	219
Changing the Default to SNMPv2	220
Setting Up E-mail Notification	221
Configuring Print Settings	222
Setting the Date and Time for Scheduled Tasks	225
Managing Getting Discovery Details	225
Adding a Discovery Schedule	226
Disabling a Schedule	227
Editing a Schedule	228

Removing a Schedule	228
Modifying Collector Settings for Newly Discovered Elements	229
Managing Product Health	230
Enabling Disk Space Monitoring	231
Viewing the Results of Disk Space Monitoring	231
Advanced Settings	232
Modifying Java Memory Settings	232
Customizing Properties	232
Managing Logging	233
Accessing the Log Files	233
About Log Files	234
Downloading Logs to a File Using the Download Logs Feature	235
Downloading Logs to a File Using the Log Download Utility	236
Downloading the User Audit Log	237
Downloading the Discovery Summary Log	237
Displaying a Log File in a Command Prompt Window	238
Changing the Provider Log Level	238
Enabling the Scanning of Critical Events of the Management Server Database	240
Viewing the Results of Logging	240
Managing the Display of Events	241
Controlling the Display of Cleared and Deleted Events	241
Modifying the Clearing and Deletion Frequency	241
Configuring the Clearing of Events	242
Configuring the Deletion of Events	243
Managing File Server SRM	244
Managing Backup Collection	244
Viewing Collectors for Backup Servers	244
Scheduling Backup Collection for Backup Managers	245

Editing the Schedule of Backup Collection	246
Setting the Backup Sessions Retention Period	246
Session Monitoring	246
Drive Monitoring	247
Viewing the Status of Backup Collection	247
Managing Reports	248
Architectural Overview of Report Views and Report Cache Refresh	248
Suggestion for Scheduling the Report Cache Refresh:	249
Report Refresh Status	250
Managing Collectors for Reports	252
Starting Collectors	254
Stopping Collectors	255
Viewing Scheduled E-mail Deliveries for Reports	255
Editing E-mail Schedules for Reports	256
Viewing Data Aging Statistics for Reports	257
Scheduling Report Cleanup	258
Refreshing the Report Cache	259
Refreshing the Report Cache Immediately	259
Scheduling a Report Cache Refresh	260
Setting Up Global Reporter	260
Editing Remote Site Information	265
Remote Sites Are Not Removed from the tnsnames.ora File	266
Managing Custom Reports (Importing and Deleting)	267
Importing Custom Reports Error Messages	269
Deleting Custom Reports	269
Managing Performance Collection	270
Managing Performance Collectors	270
Starting Performance Collectors	272



Stopping Performance Collectors	273
Viewing Data Aging Statistics for Performance	273
Editing the Locale and Currency Settings	274
Process Names	275
Process names on Windows	276
Process Names on Unix Systems	276
Editing a Collector Schedule	277
<b>9. Database Maintenance and Management</b>	<b>279</b>
Database Maintenance Window	279
Overview of Backups	281
Database Mode	282
Archive Mode	283
No-Archive Mode	283
Architectural Overview of RMAN Backups	284
Data Saved During a Backup	285
Backing up the Database Manually	285
Performing an RMAN Hot Backup	286
Scheduling RMAN Hot Backups	287
Viewing Results from RMAN Backup	288
About the Database Admin Utility	289
Accessing the Database Admin Utility	289
Refreshing the Database Admin Utility	290
Checking the Database and Listener Status	291
Use Only the Database Admin Utility to Change the Password of System Accounts	291
Exporting the Database	292
Importing the Database	293
Re-initializing the Database	294

Resetting the Temp and Undo Tablespace	295
Restarting the Database	295
Clearing Archives	295
Restoring a Cold Backup	296
Changing the Archive Mode	296
Restore RMAN Backup	297
Running a Cold Backup	298
Changing the Archive Destination	299
Downloading Log Files	299
Viewing the Database Admin Utility Log File	300
Resetting/Clearing the Database Admin Log File	300
Changing the Oracle Listener Password	300
Resetting the Admin Password for the Management Server	301
Warning Messages During Reinitializing the Database	302
About Database Passwords	302
Generating a Support Database	303
About Importing a Customer Support Database	304
Troubleshooting Listener and Database Connection Problems	305
<b>10. Viewing Element Topology and Properties</b>	<b>307</b>
About System Explorer	307
Grey Screen When Attempting to Access System Explorer	310
Accessing System Explorer	310
About Cisco Switches and VSANs in System Explorer	311
About the User Interface for System Explorer	311
About the User Interface	311
The Toolbar in System Explorer	312

Icons Displayed in the Topology	314
The List Tab	315
Viewing Clustered Elements	316
Viewing Elements by Type	316
The Access Tab	317
Obtaining Information About Zone Entries	317
Obtaining Information About Host Bindings	320
Obtaining Information About Storage System LUN Masking	323
About the Path Tab	324
About the Right-Click Menu Options	326
Viewing Storage Elements	332
Adding a Virtual Application	333
Adding Information for Discovered Hosts	334
Arranging Elements in the Topology	335
Closing Topology Windows	337
Using the Global View	337
Printing the Topology	338
Exporting the Topology to Microsoft Visio	340
Viewing the Topology in Microsoft Visio	340
Installing Storage Planner	341
Configuring Visio to View Exported Topology	342
Updating Element Data	343
Viewing Ports	343
Showing the Impact of an Element	344
Assigning a Business Cost to an Application	346
Expanding the Topology Pane	347
Filtering Fabrics	348
Viewing Event Status in the Topology	349

Custom Name for a Switch Truncated in the Topology	350
Managing Groups	350
About Groups	351
Grouping Discovered Hosts	351
Ungrouping Discovered Hosts	352
Grouping Discovered Storage Systems	353
Ungrouping Discovered Storage Systems	354
Managing Fabrics	355
Changing the Fabric Name	355
Deleting Fabrics	355
Hiding and Showing Generic Hosts	356
About Hiding Generic Hosts	356
Hiding Generic Hosts for a Switch	357
Expanding Generic Hosts for a Switch	357
Hiding Generic Hosts for All Switches	357
Expanding Generic Hosts for All Switches	358
Setting Up Custom Commands	358
About Custom Commands	358
Important Considerations	358
Adding a Custom Command	359
Editing a Custom Command	361
Deleting a Custom Command	362
Software Environment Variables for Scripting	362
Using the Remote Console	365
About the Remote Console	366
Keeping the Remote Console Active	366
Buttons on the Remote Console	368
Menu Options	368

Copying Text from the Remote Console	368
Using External Tools	369
The External Tools Feature	369
Setting up External Tools	369
About the Navigation Tab	370
Finding the Status of a Port on a Switch	373
Accessing the Navigation Tab	375
Viewing Element Properties	375
About the Properties Tab	375
Accessing the Properties Tab	377
Viewing Fabric Properties	377
Assigning a Custom Name	378
Viewing Element Topology	379
The Topology Tab	379
Multipathing	382
Direct Attached Storage	385
Filers	385
Accessing the Topology	386
About the New Window Option	387
Printing the Topology	388
Creating a Virtual Application	390
The Provisioning Tab	391
About the Events Tab	392
Asset Attributes of an Element	393
About the Collectors Tab	395
About the Monitoring Tab	396
About the Policies Tab	396
Determining If a Host Belongs to a File System	397

About the Data from CXFS File Systems 397

## **11. Provisioning 399**

About Provisioning 399

About Provisioning Brocade Switches After Upgrading 400

Managing Zones 400

SAN Zoning Overview 401

Accessing Information About Zone Aliases 406

Creating a Zone Alias 407

Zone Naming Conventions 407

Modifying a Zone Alias 408

Deleting a Zone Alias 409

Accessing Information About Zoning 409

Creating a Zone in a Fabric 410

Adding and Removing Zone Members 411

Deleting a Zone 412

Accessing Information About Zone Sets 412

Creating a Zone Set 413

Modifying a Zone Set 414

Deleting a Zone Set 415

Copying a Zone Set 416

Activating a Zone Set 417

Zones and Zone Sets Listed Twice 418

Changing the Amount of Information Collected  
from the Inactive Zone Database (Cisco Switches) 419

About the Messages Displayed in the Brocade Console 420

Managing Storage 421

Setting Up Storage Partitioning 422

How to Set Up Storage Partitioning 426

Modifying the Cache Settings (LSI and Sun 6130)	427
Changing the Owner of a Volume (LSI, CLARiiON and Sun 6130)	428
Managing Storage Pools	428
Creating a Storage Pool (LSI, CLARiiON, Sun 6130 and Sun 35xx)	429
Accessing Information About Storage Pools	430
Deleting a Storage Pool (LSI and CLARiiON Only)	431
Managing Volumes	431
Accessing Information About Volumes	432
Filtering Volumes	433
Creating a Storage Volume	434
Deleting a Storage Volume	438
Changing the Cache Block Size for a Storage System (LSI and Sun 6130)	440
Modifying the Cache Settings (LSI and Sun 6130 Only)	440
Rules for Creating Host Security Groups	441
Managing Host Security Groups	446
Accessing Information About Host Security Groups	446
Creating Host Security Groups	448
Editing a Host Security Group	450
Deleting a Host Security Group	453
Setting the Host Mode for IBM Storage Systems	453
General Provisioning Issues	454
Provisioning Can Make a Device Inaccessible	454
Provisioning Does Not Make an Operating System Aware of a Device	454
Provisioning Issues by Vendor	454
Issues Specific to CLARiiON Storage Systems	455
Issues Specific to EMC Symmetrix Storage Systems	456
Issues Specific to HDS Storage Systems	457

Issues Specific to HP Storage Systems	461
Issues Specific to LSI Storage Systems	462
<b>12. Path Provisioning</b>	<b>463</b>
About Path Provisioning	463
How Path Provisioning Works	465
How to Use Path Provisioning	466
About the User Interface	468
Default System Action Templates	469
Volume Creation, LUN Security, and Zone Operation	469
Step 1 - Select Storage System	470
Step 2 - Select a Host	471
Step 3 - Select a Volume	473
Step 4 - Select a Host Security Group	474
Step 5 - Select a Zone	474
Creating a Meta Volume	476
LUN Security	478
Step 1 - Select Storage System	478
Step 2 - Select a Host	479
Step 3 - Select a Volume	481
Step 4 - Select a Host Security Group	482
Zone Operation	482
Step 1 - Select Storage System	483
Step 2 - Select a Host	484
Step 3 - Select a Zone	486
Volume Creation and LUN Security	487
Step 1 - Select Storage System	488
Step 2 - Select a Volume	489
Step 3 - Select a Host Security Group	491



LUN Security and Zone Operation	492
Step 1 - Select Storage System	492
Step 2 - Select a Host	493
Step 3 - Select a Host Security Group	495
Step 4 - Select a Zone	496
Volume Assignment	497
Step 1 - Select Storage System	498
Step 2 - Select a Volume	499
Step 3 - Select a Host Security Group	501
Providing a LUN Number	502
Creating a System Action Template	502
Modifying a System Action Template	504
Adding a Host	504
Creating a Host Security Group	505
Scheduling Provisioning Jobs	506
Executing Provisioning Jobs	508
Monitoring Provisioning Jobs	509
Deleting Multiple Jobs	509
Naming Conventions	509
Using Multipathing with Path Provisioning	510
Customizing Path Provisioning	511
Storage System Customize Dialog Box	511
Host Customize Dialog Box	512
Customize Volume Options Dialog Box	512
Customize HSG Options	513
Customize Zone Options Dialog Box	513
Automatically Configure Zoning	515
Manually Configure Zoning	515

Assigning a Template to a Role 516

### **13. Running Reports 519**

About Reporter 519

Available Reports 520

Troubleshooting Reporter 521

Accessing Reporter 523

Viewing Reports 523

Report Parameters 525

Viewing Report Collectors for an Element 525

Refreshing a Report 526

Changing the Formatting of a Report 527

Opening a Report in a New Window 527

Maximizing the Screen Space for a Report 527

Filtering Data in Global Reports 529

Sending a Report by E-mail 530

Managing E-mail Schedules for Reports 531

Adding an E-mail Schedule for a Report 531

Editing an E-mail Schedule for a Report 534

Deleting an E-mail Schedule for a Report 535

Viewing E-mail Schedules for a Report 535

Creating Custom Reports 536

About Creating Custom Reports 537

Configuring Report Designer to Work with the Management Server 538

Designing Custom Reports 540

Creating Standard Reports 540

Managing Custom Reports (Importing and Deleting) 549

Importing Custom Reports Error Messages 551

Deleting Custom Reports 551

	Integrating Custom Reports	552
	Detailed Schema Information	554
	Views from Previous Releases	598
	Implementing Custom Reports on Sun Solaris	604
<b>14.</b>	<b>Event Management</b>	<b>605</b>
	About Event Manager	605
	Accessing Event Manager	607
	Event Manager Icons	608
	Viewing Event Details	611
	Clearing Events	613
	Configuring the Clearing of Events	613
	Configuring the Deletion of Events	614
	Deleting Events	615
	Sorting Events	615
	Adding Journal Entries	616
	Changing the CLARiiON Event Polling Interval	617
	Brocade Events	618
	Brocade Switch Events	618
	Supported Brocade Events	619
	Filtering Events	619
	Setting up a filter	620
	Selecting a custom time period	622
	Resetting a filter	625
	Setting up advanced filtering	626
	Clearing Advance Filtering Options	630
<b>15.</b>	<b>Viewing Performance Data</b>	<b>631</b>
	About Performance Explorer and Array Performance Pack	632

Array Performance Pack Requirements	632
Licensing Requirements and Setup	632
Software Requirements	633
EVAPerf Data Collector Requirements	633
Specifying Data Collectors	634
EVA Array Discovery	637
EVA Metrics and Considerations	638
EVAPerf Considerations	638
General Considerations for Performance Explorer	639
Accessing Performance Explorer	640
Creating Performance Charts	640
The Toolbars in Performance Explorer	641
Comparing the Performance of Different Elements	644
Viewing Summary Charts	645
Viewing Trending Information for Performance	645
Removing Performance Data from a Graph	646
Setting a Custom Period	647
Monitoring Options	649
Managing Late Data or Errors	658
Monitoring with Direct Attached Storage	658
Supported Host Configurations for Monitoring	659
Sudden Dips Displayed in Certain Charts in Performance Explorer	661
Values Continue to Increase in Charts for Aggregated Drives and Aggregate Volumes	662
<b>16. Finding an Element's Storage Capacity</b>	<b>663</b>
About Capacity Explorer	663
Accessing Capacity Explorer	666
The Toolbars in Capacity Explorer	666

Finding the Capacity of an Element	668
Capacity Information for Applications	669
Capacity Information for Hosts	670
Capacity Information for NetApp NAS Devices	670
Capacity Information for Storage Systems	672
Viewing the raw capacity of a storage system	672
Viewing post-RAID information	673
Obtaining Utilization Reports	675
Printing Elements in Capacity Explorer	676
Viewing Capacity Charts	676
Viewing Trending Information for Storage Capacity	678
Different Results for the df -k Command and Capacity Explorer	679
<b>17. Managing Policies</b>	<b>681</b>
About Policy Manager	681
Accessing Policy Manager	682
Creating Policies	683
Actions Available for When a Policy Condition is Fulfilled	683
Severity Levels	684
Creating a Utilization or Backup Policy	684
About the Policy Templates	686
Creating Policies for Discovery	689
Creating Policies for Provisioning	690
Creating Policies for Events	691
Testing a Utilization Policy	692
Modifying Policies	693
Modifying Utilization and Backup Policies	693
Modifying Discovery Policies	694
Modifying Provisioning Policies	695

Modifying Policies for Events	696
Viewing Policies	697
Deactivating a Policy	698
Deleting Policies	698
Providing E-mail Notification for a Policy	699
Providing Event Generation for a Policy	699
Providing a Custom Command for a Policy	700
<b>18. Chargeback</b>	<b>703</b>
About Chargeback	703
Setting Up Chargeback	705
Accessing Chargeback	705
Creating an Asset Record	706
Changing the Status of an Element	707
Saving Chargeback Information	708
Viewing Assets	708
Creating a New Storage Tier	710
Adding Elements to a Storage Tier	710
Removing Elements from a Storage Tier	711
Editing a Storage Tier	712
Deleting a Storage Tier	712
Adding Asset Information	712
Adding Asset Information	713
Adding General Information	714
Adding Staff Information	715
Adding Geographic Information	716
Adding Licensing and Warranty Information	716
Adding Custom Information	716
Managing Departments	717

Adding Departments	717
Editing a Department	717
Removing a Department from Chargeback	718
Setting the Infrastructure Cost	718
About Asset-based and Storage-based Infrastructure Cost	719
Setting Up Asset and Storage Based Chargeback	719
Setting Up Asset-Based Chargeback	720
Step 1 - Specify Financial information	722
Step 2 - Assign Departmental Ownership Percentage	723
Step 3 - Review Asset-based Chargeback Result	723
Setting Up Storage-Based Chargeback	726
Step 1 - Assign Departmental Ownership Percentage	728
Step 2 - Review Storage Tier Cost	728
Step 3 - Review Storage Dependency and Cost	729
Step 4 - Review Storage-Based Chargeback Result	729
Editing Percentage of Ownership	730
Removing Department Ownership of an Element	730
How Capacity Differs in Chargeback and Capacity Explorer	731
How a Depreciation Method Is Calculated	732
Calculating Straight Line Depreciation	732
Calculating Fixed Declining Balance	733
Calculating Double Declining Balance	735
Viewing Chargeback	737
Viewing Chargeback by Element	738
Viewing Chargeback by Department	739
Viewing Chargeback by Owner	740
Chargeback Reports	741
Viewing Chargeback Reports	741

E-mailing a Chargeback Report	742
Managing E-mail Schedules for Chargeback Reports	743
Adding an E-mail Schedule for a Chargeback Report	743
Editing an E-mail Schedule for a Chargeback Report	745
Deleting E-mail Schedules for a Chargeback Report	747
Viewing E-mail Schedules for a Chargeback Report	747
Viewing the History of an E-mail Chargeback Schedule	748
Filtering Assets	749
About Filtering Assets	749
Selecting an Element Type for Chargeback	750
Customizing the Element Type Filter for Chargeback	750
Filtering Assets by Status	751
Customizing the Asset Status Filter for Chargeback	751
Hiding Filters in Chargeback	752
<b>19. Managing Backups</b>	<b>753</b>
About Protection Explorer	753
Requirements for Using Protection Explorer	754
Determining if You Have Enough Media to Run a Backup	755
Accessing Protection Explorer	757
Viewing Sessions that are Running	757
Determining if the Last Scheduled Backup was Successful	758
Viewing the Summary Backup Charts	758
Viewing Backup Results for a Backup Manager Host	759
Viewing Backup Results for a Client	759
Viewing Backup Information for a Client	760
Viewing Backup Reports	760
About the User Interface	761
About the Topology Icons in Protection Explorer	762



	About the Toolbars in Protection Explorer	764
	Changing the Topology Settings	766
	Exporting the Topology to Visio	767
	Right-Click Menu Options on the Topology Tab	768
	About the Summary Backup Charts	772
	About the Tabs in the Topology Lower Pane	773
	Sorting Information in the Lower Pane	776
	Modifying Summary Backup Charts	777
	Viewing Charts for a Backup Manager Host	778
<b>20.</b>	<b>Business Tools</b>	<b>781</b>
	About the Business Tools	781
	Using Business Tools in Remote AD/LDAP Authentication Mode	782
	Using the HBA Replacement Automator	783
	Installing New HBA with Old HBA	783
	Installing New HBA by Itself	784
	Setting up Risk Analysis	784
	Global Change Management	785
<b>21.</b>	<b>Troubleshooting</b>	<b>793</b>
	Troubleshooting Installation/Upgrade	793
	If Your Installation or Upgrade Failed, Capture the Logs	794
	Checking Installation Log Files	795
	“The environment variable ‘perl5lib’ is set.” Message	795
	“SEVERE: OUI-10029...” Message	796
	Brocade API Switches Displaying Stale Data	796
	Troubleshooting the Oracle Database (Windows)	796
	Use Only the Installation Wizard (or Unix Scripts) to Install/Upgrade Oracle	796
	Cancelling an Installation or Upgrade Before Completion	797

Uninstalling Oracle Using the Oracle Scripts	797
Re-installing the Management Server	798
Existing Oracle Database Is Detected	798
Configuring the Java Console	798
“Data is late or an error occurred” Message	799
appstorm.<timestamp>.log Filled with Connection Exceptions	799
Receiving HTTP ERROR: 503 When Accessing the Management Server	800
Windows	801
Unix systems	801
Errors in the Logs	802
Permanently Changing the Port a CIM Extension Uses (UNIX Only)	803
Configuring UNIX CIM Extensions to Run Behind Firewalls	804
Volume Names from Ambiguous Automounts Are Not Displayed	807
Solaris Management Server Suddenly Restarts	808
Installing the Software Security Certificate	808
Installing the Certificate by Using Microsoft Internet Explorer 6.0	809
Changing the Security Certificate to Match the Name of the Server	810
Windows	810
Sun Solaris and Linux	811
Troubleshooting Discovery and Get Details	812
Troubleshooting Mode	813
Unable to discover Emulex host bus adapters	813
CIMOM Service Not Starting After Trying to Discover Sybase or SQL Server Applications	814
Configuring E-mail Notification for Get Details	814
Increasing the Time-out Period and Number of Retries for Switches in Progress	815
“Connection to the Database Server Failed” Error	817
Using the Test Button to Troubleshoot Discovery	817

DCOM Unable to Communicate with Computer	820
Duplicate Listings/Logs for Brocade Switches in Same Fabric	820
Duplicate listings: Targets tab	820
Duplicate Logs	820
Duplicate entries for the same element on the Get Details page	821
Element Logs Authentication Errors During Discovery	821
EMC Device Masking Database Does Not Appear in Topology (AIX Only)	822
Management Server Does Not Discover Another Management Server's Database	822
Microsoft Exchange Drive Shown as a Local Drive	822
Unable to Discover Microsoft Exchange Servers	822
Nonexistent Oracle Instance Is Displayed	822
Requirements for Discovering Oracle	823
Do Not Run Overlapping Discovery Schedules	823
"This storage system uses unsupported firmware. ManagementClassName: class_name" Message	823
Troubleshooting Topology Issues	824
About the Topology	824
Undiscovered Hosts Display as Storage Systems	827
Solaris Machines Appear to Have Extra QLogic HBAs	828
No Stitching for Brocade Switches with Firmware 3.2.0	828
Link Between a Brocade Switch and a Host Disappears from the Topology	828
Incorrect Topology Sometimes Displayed for CNT Switches	828
Unable to Find Elements on the Network	829
Unable to See Path Information	829
Device Locking Mechanism for Brocade Element Manager Query/Reconfiguration	829
A Discovered Sun StorEdge A5000 JBOD Does Not Display Its WWN Properly	829

Sun 6920 Storage Systems: “ReplicatorSQLException: Database create error” During Get Details	830
Mirrored Volumes Cannot Be Provisioned on Sun 6920 Storage Systems	830
Unable to Monitor McDATA Switches	830
Unable to Detect a Host Bus Adapter	831
Navigation Tab Displays Removed Drives as Disk Drives	831
Unable to Obtain Information from a CLARiiON Storage System	831
Discovery Fails Too Slowly for a Nonexistent IP Address	832
“CIM_ERR_FAILED” Message	833
Re-establishing Communication with EFCM	834
CIM_ERR_FAILED When Trying to Activate a Zone Set Using McDATA SWAPI	835
Communicating with HiCommand Device Manager Over SSL	835
Unable to Discover a UNIX Host Because of DNS or Routing Issues	836
ERROR replicating APPIQ_EVAStorageVolume during Get Details for an EVA array	837
Recalculating the Topology	838
Troubleshooting the Java Plug-in	838
Java Applet Has Data from a Different Build of Management Server Software	838
OutOfMemoryException Messages	839
Unable to View System Explorer After Upgrade	839
Improving Reload Performance in System Explorer	839
Troubleshooting Provisioning	839
Cannot Access a Resource Owned by Another Controller	839
Error -56	840
“Can't delete this zone” Message	840
Changes in EFC Manager Requiring Get Details	840
Troubleshooting Hardware	840
About Swapping Host Bus Adapters	841

“Fork Function Failed” Message on AIX Hosts	841
Known Driver Issues	841
Known Device Issues	841
“mailbox command 17 failure status FFF7” Message	845
“Process Has an Exclusive Lock” Message	845



# Figures

---

FIGURE 1-1	Status Light	6
FIGURE 1-2	Closing the Left Pane	8
FIGURE 1-3	Opening the Left Pane	9
FIGURE 2-1	Setting Default User Names and Passwords	26
FIGURE 2-2	Adding an IP Range for Scanning	28
FIGURE 2-3	Deleting Elements from the Management Server	99
FIGURE 3-1	Selecting Import from Disk	142
FIGURE 3-2	Enabling Compile Class	143
FIGURE 3-3	Selecting appiq.cls	144
FIGURE 3-4	Importing Wrapper Class Definitions	146
FIGURE 5-1	System Explorer Cluster Representation	170
FIGURE 5-2	Cluster Element Topology Representation	172
FIGURE 5-3	Capacity Manager Cluster Representation	174
FIGURE 6-1	Parent-Child Hierarchy for Organizations	179
FIGURE 6-2	Children in Multiple Organizations	180
FIGURE 6-3	Clicking the Name of Your User Account	187
FIGURE 6-4	Clicking the Organization Link	197
FIGURE 6-5	Filtering Organizations	198
FIGURE 6-6	Active Organization	198
FIGURE 7-1	An Example of Direct Attached Storage	216

FIGURE 8-1	227
FIGURE 8-2	Report Views and Report Cache Refresh 249
FIGURE 8-3	An Example of Global Reporting 261
FIGURE 8-4	Manage Custom Reports Screen 268
FIGURE 8-5	Screen Displays Custom Reports Available 269
FIGURE 10-1	Expanding the Fabric Node 315
FIGURE 10-2	Highlighting a Fabric's Members in the Topology 316
FIGURE 10-3	Highlighting the Applications in the Topology 317
FIGURE 10-4	Members of a Zone Set 318
FIGURE 10-5	Displaying a Zone Member and its Switch 319
FIGURE 10-6	Zone Member to Switch 320
FIGURE 10-7	Obtaining Information About a Zone Member's Adapter 320
FIGURE 10-8	Highlighting Elements with Host Bindings 321
FIGURE 10-9	Displaying Host Bindings 322
FIGURE 10-10	HBA Port Properties 322
FIGURE 10-11	WWN Properties 323
FIGURE 10-12	WWN Properties 324
FIGURE 10-13	Obtaining Path Information 325
FIGURE 10-14	Path Information Visible in the Tree 326
FIGURE 10-15	Enclosing the Elements 336
FIGURE 10-16	Dragging Multiple Elements to Their New Location 337
FIGURE 10-17	Showing the Impact of an Element 345
FIGURE 10-18	Determining Business Cost 347
FIGURE 10-19	Expanding the Topology Pane 348
FIGURE 10-20	Filtering Fabrics 348
FIGURE 10-21	Remote Console 366
FIGURE 10-22	Obtaining Information About a Host 371
FIGURE 10-23	Details of a Host Connected to a Switch 372
FIGURE 10-24	Finding the Status of a Port 374
FIGURE 10-25	Topology of a Server 380



FIGURE 10-26	Topology of a Server (Continued)	381
FIGURE 10-27	Multipathing Displayed in the Topology	383
FIGURE 10-28	Multipathing Displayed in the Topology (Continued)	384
FIGURE 10-29	Direct Attached Storage in the Topology	385
FIGURE 10-30	New Window Option	388
FIGURE 10-31	Viewing Asset Records	394
FIGURE 10-32	CXFS File System	398
FIGURE 11-1	Resources in Two Zones	402
FIGURE 11-2	Overview of Zoning Structure	403
FIGURE 11-3	Selecting a Setting and Size for a Storage Pool	429
FIGURE 11-4	Deleting Several Volumes at Once	439
FIGURE 11-5	Selecting All Volumes	439
FIGURE 12-1	Selecting a Storage System	471
FIGURE 12-2	Selecting a Fabric	475
FIGURE 12-3	Selecting a Storage System	479
FIGURE 12-4	Selecting a Storage System	484
FIGURE 12-5	Selecting a Fabric	487
FIGURE 12-6	Selecting a Storage System	489
FIGURE 12-7	Selecting a Volume	491
FIGURE 12-8	Selecting a Storage System	493
FIGURE 12-9	Selecting a Fabric	497
FIGURE 12-10	Selecting a Storage System	499
FIGURE 12-11	Selecting a Volume	501
FIGURE 12-12	Specifying a LUN Number	502
FIGURE 12-13	Selecting a Port	505
FIGURE 12-14	Selecting the Second Path	511
FIGURE 12-15	Selecting a Fabric	516
FIGURE 13-1	Hiding the Middle Pane	528
FIGURE 13-2	Closing the Left Pane in Reporter	528
FIGURE 13-3	Selecting Organizations Used in This Report	534

FIGURE 13-4	Report Architecture	537
FIGURE 13-5	Choosing a Standard Report	541
FIGURE 13-6	Adding Tables for a Standard Report	542
FIGURE 13-7	Linking Common Data in Tables for a Standard Report	543
FIGURE 13-8	Creating Search Criteria for Standard Reports	544
FIGURE 13-9	Deciding Which Data Should Appear in the Report	545
FIGURE 13-10	Sorting Information in the Report	546
FIGURE 13-11	Selecting the Layout of the Report	547
FIGURE 13-12	Report Template Displayed	548
FIGURE 13-13	Result of Clicking the View Tab	549
FIGURE 13-14	Manage Custom Reports Screen	550
FIGURE 13-15	Screen Displays Custom Reports Available	551
FIGURE 14-1	Accessing summary information	606
FIGURE 14-2	Accessing Event Details	612
FIGURE 14-3	Clearing an event	613
FIGURE 14-4	Accessing the filter feature	620
FIGURE 14-5	The Filter feature in Event Manager	620
FIGURE 14-6	Selecting a time period	621
FIGURE 14-7	Custom Time Period	623
FIGURE 14-8	Selecting a start date for filtering	623
FIGURE 14-9	Selecting a date	624
FIGURE 14-10	Selecting a date	625
FIGURE 14-11	Accessing the filter feature	626
FIGURE 14-12	The Filter feature in Event Manager	626
FIGURE 14-13	Accessing the filter feature	627
FIGURE 14-14	The Filter feature in Event Manager	627
FIGURE 14-15	Accessing advanced options for filtering	628
FIGURE 14-16	Clicking the Add button	628
FIGURE 14-17	Listing of elements	629
FIGURE 14-18	Advance filtering options	630

FIGURE 15-1	Data Collector Selection	635
FIGURE 15-2	Expanded View of Available Metrics in Performance Manager	636
FIGURE 15-3	EVA Array Host Port Data Rates	637
FIGURE 16-1	Capacity of an Element	665
FIGURE 16-2	Viewing the Capacity of Elements in a Fabric	669
FIGURE 16-3	Post-RAID tab	674
FIGURE 16-4	Viewing a Utilization Report	675
FIGURE 17-1	Testing a Newly Created Utilization Policy	686
FIGURE 17-2	Testing a Utilization Policy	693
FIGURE 17-3	Testing a Modified Utilization Policy	694
FIGURE 18-1	Selecting an Element	708
FIGURE 18-2	Accessing an Element's Asset Information	713
FIGURE 18-3	Asset Record Node	714
FIGURE 18-4	Accessing an Element's Asset Information	721
FIGURE 18-5	Clicking the Asset-based Node	721
FIGURE 18-6	Ownership Cost	725
FIGURE 18-7	Accessing an Element's Asset Information	727
FIGURE 18-8	Clicking Storage-Based Node	727
FIGURE 18-9	Chargeback Capacity	731
FIGURE 18-10	Capacity in Capacity Explorer	731
FIGURE 18-11	Accessing Storage-Based Chargeback	738
FIGURE 19-1	Summary Settings Page for Protection Explorer Charts	777
FIGURE 21-1	Perl5lib environment variable message	795



# Tables

---

TABLE 1-1	Status Light Settings	6
TABLE 1-2	Buttons Displayed When the Left Pane Is Closed	8
TABLE 1-3	Icons on the Home Page	9
TABLE 2-1	Discovery Requirements for Switches	34
TABLE 2-2	Discovery Settings for McDATA and Connectrix Switches	44
TABLE 2-3	Discovery Requirements for Storage Systems, Tape Libraries & NAS Devices	58
TABLE 2-4	Discovery Group Ports	94
TABLE 2-5	Task Status Descriptions	104
TABLE 3-1	Script Names for Managed Databases	114
TABLE 6-1	Default Role Privileges	176
TABLE 6-2	Default Role Privileges by Elements	178
TABLE 7-1	License Restrictions	213
TABLE 7-2	Determining Managed Access Points	215
TABLE 8-1	Logging Levels	239
TABLE 8-2	Default Settings for Clearing Events	243
TABLE 8-3	About Collectors for Reports	253
TABLE 8-4	Viewing System-Wide E-mail Schedules	255
TABLE 8-5	Description of “An Example of Global Reporting”	261
TABLE 8-6	About Performance Collectors	271
TABLE 8-7	Currency Settings	275

TABLE 8-8	Process Names on Windows	276
TABLE 8-9	Process Names on UNIX systems	276
TABLE 9-1	Backup Directories in %ORA_HOME%\rman	281
TABLE 9-2	Description of Backups	281
TABLE 9-3	Sample Backup Example	284
TABLE 9-4	Default Passwords for Database Accounts	302
TABLE 10-1	Feature of the Toolbar in System Explorer	312
TABLE 10-2	Icons Displayed in the Topology	314
TABLE 10-3	Menu Options Accessible from the Topology*	327
TABLE 10-4	Menu Options on the Access Tab	331
TABLE 10-5	Show Impact Results	346
TABLE 10-6	Severity Levels	349
TABLE 10-7	Variables for All Elements	362
TABLE 10-8	Variables for Storage Systems, Switches, and Hosts Only	363
TABLE 10-9	Variables for Switches Only	364
TABLE 10-10	Variables for Hosts Only	364
TABLE 10-11	Variables for Applications Only	365
TABLE 10-12	Buttons on the Remote Console	368
TABLE 10-13	Menu Options	368
TABLE 10-14	Information Available from the Navigation Page	372
TABLE 10-15	Port Status Definitions	374
TABLE 10-16	The Toolbar in the Topology Tab	386
TABLE 11-1	Setting Up Zoning	403
TABLE 11-2	Zoning Support	404
TABLE 11-3	Provisioning and Pool Support	422
TABLE 11-4	Volume and Host Security Group Support	423
TABLE 11-5	Allowed Initiators in Host Security Groups	441
TABLE 11-6	Volume Usage	462
TABLE 12-1	Overview for Path Provisioning	468
TABLE 12-2	Feature Toolbar	469

TABLE 12-3	Zone Icons	474
TABLE 12-4	Zone Icons	486
TABLE 12-5	Zone Icons	496
TABLE 12-6	Zone Icons	515
TABLE 13-1	Buttons Displayed When the Left Pane Is Closed	528
TABLE 13-2	Viewing E-mail Schedules for a Report	536
TABLE 13-3	Description of the Report Views	554
TABLE 13-4	MVC_HOSTSUMMARYVW	560
TABLE 13-5	MVC_CARDSUMMARYVW	561
TABLE 13-6	MVC_HOSTVOLUMESUMMARYVW (logical volumes)	561
TABLE 13-7	MVC_HOSTDISKDRIVEVW	562
TABLE 13-8	MVC_STORAGESYSTEMSUMMARYVW	562
TABLE 13-9	MVC_STORAGEPOOLSUMMARYVW	562
TABLE 13-10	MVC_STORAGEVOLUMESUMMARYVW	563
TABLE 13-11	MVC_SWITCHSUMMARYVW	564
TABLE 13-12	MVC_PORTSUMMARYVW	565
TABLE 13-13	MVC_ZONESUMMARYVW	566
TABLE 13-14	MVC_ZONEVW	567
TABLE 13-15	MVC_PATHVW	567
TABLE 13-16	MVC_SUBPATHVW	567
TABLE 13-17	MVC_MULTIPATHVW	568
TABLE 13-18	MVC_EVENTSVW	568
TABLE 13-19	MVC_ORGANIZATIONVW	570
TABLE 13-20	MVC_ORGRELATIONVW	570
TABLE 13-21	MVC_HOSTCAPACITYVW	571
TABLE 13-22	MVC_STORAGESYSTEMCONFIGVW	571
TABLE 13-23	MVC_STORAGEPOOLCONFIGVW	571
TABLE 13-24	MVC_SWITCHCONFIGVW	572
TABLE 13-25	MVC_OPTIONALTABLEVW	572
TABLE 13-26	MVC_DISKEXTENTSUMMARYVW	572

TABLE 13-27	MVC_STORAGEVOLUMEPORTS	573
TABLE 13-28	MVC_VOLUMEDISKDRIVEVW	573
TABLE 13-29	MVC_STORAGEPROCESSORSUMMARYVW	574
TABLE 13-30	MVC_DISKDRIVESUMMARYVW	574
TABLE 13-31	MVC_DISK_EXTENTVW	575
TABLE 13-32	MVC_ASSETSUMMARY	575
TABLE 13-33	MVC_APPLICATIONSUMMARYVW	577
TABLE 13-34	MVC_UNITACCESSVW	578
TABLE 13-35	MVCA_DBAPPCAPACITYVW	578
TABLE 13-36	MVCA_EXCHAPPCAPACITYVW	579
TABLE 13-37	MVCA_VIRTUALAPPCAPACITYVW	579
TABLE 13-38	MVCA_FSRM_VOLUMESUMMARYVW	579
TABLE 13-39	MVCA_FSRM_AGESUMMARYVW	580
TABLE 13-40	MVCA_FSRM_EXTDETAILSUMMARYVW	580
TABLE 13-41	MVCA_FSRM_DIRDETAILSUMMARYVW	580
TABLE 13-42	MVCA_FSRM_USERSUMMARYVW	581
TABLE 13-43	MVCA_FSRM_TOPNFILES	581
TABLE 13-44	MVCA_FSRM_AGEDFILEDETAILS	581
TABLE 13-45	MVCA_FSRM_LARGEDIRINFO	582
TABLE 13-46	MVCA_BU_MASTERSERVERSUMMARY	582
TABLE 13-47	MVCA_BU_MEDIASERVERSUMMARY	583
TABLE 13-48	MVCA_BU_CLIENTSUMMARY	583
TABLE 13-49	MVCA_BU_MEDIASUMMARY	583
TABLE 13-50	MVCA_BU_JOBSUMMARY	585
TABLE 13-51	MVCA_BU_LIBRARYSUMMARY	585
TABLE 13-52	MVIEWCORE_STATUS	586
TABLE 13-53	MVIEW_STATUS	586
TABLE 13-54	MVC_DISCOVERYDETAILSVW	586
TABLE 13-55	MVC_HOSTRELATIONVW	587
TABLE 13-56	MVC_APPLICATIONRELATIONVW	587



TABLE 13-57	MVC_STORAGETIERDETAILVW	588
TABLE 13-58	MVCA_BU_OPTIONALTABLEVW	588
TABLE 13-59	MVCA_BU_DRIVESTATVW	589
TABLE 13-60	MVCA_EXCHMAILBOXDETAILVW	589
TABLE 13-61	MVCA_EXCHPUBLICFOLDERDETAILVW	590
TABLE 13-62	MVCA_EXCHGESTORESUMMARYVW	590
TABLE 13-63	MVCA_EXCHSTORGROUPSUMMARYVW	591
TABLE 13-64	MVCA_FSRM_FILEREREPORTDATAVW	591
TABLE 13-65	MVCA_FSRM_DIRREPORTDATAVW	591
TABLE 13-66	MVCA_FSRM_REPORTRULEVW	592
TABLE 13-67	MVCS_HOSTMEMORYSTATSVW	592
TABLE 13-68	MVCS_HOSTCPUSTATSVW	593
TABLE 13-69	MVCS_EVACTRLSTATSVW	593
TABLE 13-70	MVCS_EVADISKSTATSVW	594
TABLE 13-71	MVCS_EVAHOSTFCPORTSTATSVW	595
TABLE 13-72	MVCS_EVASPAGVOLUMESTATSVW	596
TABLE 13-73	MVCS_EVASTORAGESYSTEMSTATSVW	597
TABLE 13-74	MVCS_EVAVOLUMESTATSVW	597
TABLE 13-75	Views from Previous Releases	598
TABLE 14-1	Severity Levels	607
TABLE 14-2	Icons in Event Manager	608
TABLE 14-3	Supported Hardware for Events	608
TABLE 14-4	Default Settings for Clearing Events	614
TABLE 14-5	Brocade Switch Events	618
TABLE 15-1	Toolbar in Lower Pane of Performance Explorer	641
TABLE 15-2	About the Monitoring Options	649
TABLE 15-3	Host Monitoring Support	660
TABLE 16-1	Color Coding for Capacity Explorer	665
TABLE 16-2	Toolbar in Lower Pane of Capacity Explorer	666
TABLE 16-3	Explanation of the Properties of the Capacity Levels for HDS Array Groups	674

TABLE 17-1	Policy Templates	686
TABLE 17-2	Policy Table Description	697
TABLE 18-1	Setting Up Chargeback	705
TABLE 18-2	Element Type Icons	709
TABLE 18-3	Viewing E-mail Schedules for a Chargeback Report	748
TABLE 18-4	Element Types	750
TABLE 19-1	Topology Icons in Protection Explorer	762
TABLE 19-2	The Toolbar in Protection Explorer	764
TABLE 19-3	Toolbar for Charts	766
TABLE 19-4	Right-Click Menu Options on the Topology Tab	768
TABLE 19-5	Right-Click Menu Options on the Summary Tab	770
TABLE 19-6	To Obtain Additional Information from a Chart on the Summary Tab	771
TABLE 19-7	Show Details for Tabs on the Lower Pane of the Topology Tab	771
TABLE 19-8	Protection Explorer Summary Charts	772
TABLE 19-9	Tabs in the Lower Pane of Protection Explorer Topology	774
TABLE 19-10	Buttons on the Summary Settings Page	777
TABLE 21-1	Troubleshooting Firewalls	805
TABLE 21-2	Time-out Properties	816
TABLE 21-3	Retry Properties	817
TABLE 21-4	Troubleshooting Discovery and Get Details	825
TABLE 21-5	Known Device Issues	842

# Revision History

---

Short Name	Part Number	Dash	Date	Comments
USER GUIDE	817-7923-16	-05	January 2008	

---



# Preface

---

*This document describes how to use Sun StorageTek™ Operations Manager 6.0 and its components.*

*This document assumes you have a basic understanding of the following:*

- *Networking*
- *Storage Area Networks (SANs)*
- *The Common Information Model (CIM)*

---

## Before You Read This Book

In order to fully use the information in this document, you must have thorough knowledge of the topics discussed in the *Sun StorageTek™ Operations Manager 6.0 Installation Guide*.

---

# Using UNIX Commands

This document might not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices. See the following for this information:

- Software documentation that you received with your system
- Solaris™ operating environment documentation, which is at

<http://docs.sun.com>

---

## Related Documentation

Application	Title	Part Number
Installation	Release Notes	-----
Operations Manager	<i>Sun StorageTek™ Operations Manager 6.0 Installation Guide</i>	817-7922-16
Resource Manager	<i>Sun StorageTek™ Resource Manager 6.0 Guide</i>	817-7925-16
CLI	<i>Sun StorageTek™ Operations Manager 6.0 CLI Guide</i>	817-7924-16
Application Module	<i>Sun StorageTek™ Application Module 6.0 Guide</i>	817-7926-16

*Table listing other documents that are related to this book or product.*

---

## Accessing Sun Documentation

You can view, print, or purchase a broad selection of Sun documentation, including localized versions, at:

<http://www.sun.com/documentation>

---

## Contacting Sun Technical Support

If you have technical questions about this product that are not answered in this document, go to:

<http://www.sun.com/service/contacting>

---

## Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can submit your comments by going to:

<http://www.sun.com/hwdocs/feedback>

Please include the title and part number of your document with your feedback:

*Sun StorageTek™ Operations Manager 6.0 User Guide* , part number 817-7923-16





## Overview

---

This chapter contains the following topics:

- “About This Product” on page 1
- “Suggested Topics for First-Time Users” on page 2
- “Product Components” on page 3
- “The User Interface” on page 5
- “Accessing the Management Server” on page 11
- “Installing the Java Plug-in” on page 12
- “Installing the Software Security Certificate” on page 14
- “Restarting 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.

This product gives 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.

### **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

### **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.

---

## **Suggested Topics for First-Time Users**

As a first-time user, you should first become familiar with 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 run 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:

- “Discovering Switches, Storage Systems, NAS Devices, and Tape Libraries” on page 21
- “Discovering Applications, Backup Hosts and Hosts” on page 105

Once you have discovered and obtained details about the devices in the network, begin adding users and adding them to roles and organizations. See “Managing Security” on page 175 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:

- “Provisioning” on page 399
- “Event Management” on page 605
- “Running Reports” on page 519

- “Viewing Performance Data” on page 631

---

## 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 3 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. See 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.
- **Optional: Module for managing Caché** — Management software for Caché lets the administrator actively manage the data storage requirements for Caché.

## 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, or 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 provide 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 “Business Tools” on page 781 for more information.
- **Capacity Explorer** — Capacity Explorer provides a graphical representation of an element's storage capacity in the storage network. See “Finding an Element's Storage Capacity” on page 663 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 “Chargeback” on page 703 for more information.
- **Command Line Interface** — The command line interface (CLI) provides an alternate way for you to manage elements that 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, for example, a device connected to a Brocade switch has gone off-line. See “Event Management” on page 605 for more information.
- **Path Provisioning** — Path Provisioning lets you schedule a provisioning task, such as creating zones, to run at a later time. See “Path Provisioning” on page 463 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 “Viewing Performance Data” on page 631 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 occurs
  - An event occurs on one or more specified elements
 See “Managing Policies” on page 681 for more information.
- **Protection Explorer** — Protection Explorer helps you to keep track of element backups. See “Managing Backups” on page 753 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 “Provisioning” on page 399.

- **Reporter** — Reporter provides a variety of detailed reports, such as dependency, event, and utilization reports for discovered elements. See “Running Reports” on page 519 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 “Viewing Element Topology and Properties” on page 307.

---

## 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** — The status light and the buttons for the management server's utilities are displayed in the left pane.
- **Right pane** — The output of a feature, such as the topology in System Explorer are displayed in the right pane.

## The Top Pane

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

---

**Caution** – 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 “Managing Security” on page 175 for more information.

---

- **Configuration** — This menu provides the tools for you to manage the management server, such as saving the database. See “Configuring the Management Server” on page 219.
- **Security** — This menu lets you manage users, organizations, roles, and licenses. See “Managing Security” on page 175.
- **Discovery** — This menu provides the tools for the management server to discover and obtain information from elements in your network. See “Discovering Switches, Storage Systems, NAS Devices, and Tape Libraries” on page 21 and “Discovering Applications, Backup Hosts and Hosts” on page 105.

- **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 is usually green under normal conditions, as shown in the following figure:



**FIGURE 1-1** Status Light

Table 1-1, “Status Light Settings,” on page 6 shows the possible displays for the status light.

**TABLE 1-1** Status Light Settings

Scenario	Status Light
Normal	
Discovery	
Getting Topology	
Backup Topology Details	
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**).

---

# The Left Pane

The buttons used to access the management server's main utilities are displayed in the left pane.

---


**Caution** — 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, or Reporter.

---

- **Application Explorer** — Application Explorer lets you monitor and display information obtained from discovered databases.
- **Protection Explorer** — Protection Explorer lets you monitor the overall status of the backup process, and visualize the backup configuration and recoverability of a file, directory, volume, or server. See “About Protection Explorer” on page 753.
- **System Explorer** — System Explorer lets you access systems, and view assets by fabric and logical path. See “About System Explorer” on page 307.
- **Capacity Explorer** — Capacity Explorer provides a graphical representation of an element's storage capacity in the storage network. See “About Capacity Explorer” on page 663.
- **Performance Explorer** — Performance Explorer provides a graphical representation of the results obtained from monitoring your elements. See “Creating Performance Charts” on page 640.
- **Event Manager** — Event Manager lets you view, clear, sort, and filter API-generated events. See “About Event Manager” on page 605.
- **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 681.
- **Business Tools** — The management server provides advisors and automators. Advisors provide 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 781.
- **Chargeback** — Chargeback helps you manage ownership by department, track costs, and assemble business reports. You can view data gathered by Chargeback by element, department, or the entire enterprise. You can show upper management reports with the data that Chargeback gathered. These reports can be e-mailed on a regular schedule. See “About Chargeback” on page 703 for more information.
- **Reporter** — Reporter provides detailed reporting on the infrastructure, such as statistics and usage trends. See “About Reporter” on page 519.

## 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:



**FIGURE 1-2** Closing the Left Pane



The buttons in the left pane are moved to the top of the page. See Table 1-2, “Buttons Displayed When the Left Pane Is Closed,” on page 8 for an explanation of the buttons.

**TABLE 1-2** Buttons Displayed When the Left Pane Is Closed

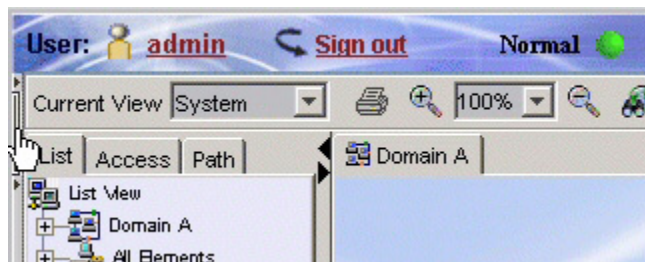
Button	Provides Access to...
	The Home page
	Application Explorer
	Protection Explorer
	System Explorer
	Capacity Explorer
	Performance Explorer
	Event Manager
	Provisioning
	Policy Manager
	Business Tools



**TABLE 1-2** Buttons Displayed When the Left Pane Is Closed (*Continued*)

Button	Provides Access to...
	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








## 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 send you email, 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.	"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
 Protection Explorer	Protection Explorer lets you keep track of element backups.	"Managing Backups" on page 753
 System Explorer	System Explorer shows you the topology of your SAN and gives you the ability to explore details about each element.	"Viewing Element Topology and Properties" on page 307
 Performance Explorer	Performance Explorer provides detailed performance-management capabilities, enabling you to visualize what you have and how it is performing.	"Viewing Performance Data" on page 631
 Capacity Explorer	Capacity Explorer provides a graphical representation of an element's storage capacity and utilization in the storage network.	"Finding an Element's Storage Capacity" on page 663
 Event Manager	Event Manager keeps you informed of what is happening with your managed elements. Its filter and report format allows you to easily view, clear, and sort the events you are interested in.	"Event Management" on page 605
 Provisioning	Provisioning simplifies your SAN-zoning and storage management tasks. The format minimizes errors by giving you easy-to-follow instructions and step-by-step screens.	"Provisioning" on page 399
 Chargeback	Chargeback helps you manage departmental ownership, track costs, and assemble business reports, making inquires, such as audits and inventory reviews, easier.	"Chargeback" on page 703

**TABLE 1-3** Icons on the Home Page (*Continued*)

Feature	Description	Where to Find More Information
 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.	"Business Tools" on page 781
 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.	"Running Reports" on page 519

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 Security for the Management Server" on page 175 for more information.

## Accessing the Management Server

Keep in mind the following:

- 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 12 for more information.
- It may take more time to log into the management server during a topology recalculation.

To access the management server:

1. Access the Sun Web Console by entering the following URL in a Web browser:

`https://mgmt-station-hostname.domain:6789`

where

- `mgmt-station-hostname.domain` is the management station host name
- 6789 is the default port number for the Sun Web Console software

2. Enter the user name and password that is used to access the host.
3. Click **Log In**.
4. On the Sun Web Console, click **Sun StorEdge™ Enterprise Storage Manager**.
5. 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 800: Receiving HTTP ERROR: 503 javax.ejb.EJBException: null; CausedByException is: Unexpected Error; nested exception is: java.lang.NoClassDefFoundError

---

6. In the management server login page, enter `admin` in the **Name** box, password in the **Password** box, and click **Login**.
7. Click **Security > Licenses** in the upper-right corner.
8. Select the tools and enter the number of MAPs, MALs, instances, and terabytes you are licensed to use.  
  
Contact customer support if you are uncertain of which products you purchased and for how many MAPs, instances, and terabytes.
9. Click **Save Changes** to certify that you are authorized to use the components selected.
10. When you are shown the license agreement, accept the license if you agree with its terms.

Your changes take effect.

---

## 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 or Linux, you must manually install the Java plug-in.

To install the Java plug-in on Solaris:

1. Go to the following URL and download the installation file for the Sun JRE:

`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.

To install the Java plug-in on Linux:

1. Go to the following URL and download the installation file for the Sun JRE when asked:

`http://<management_server>/servlet.html?page=JavaPluginLinux`

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/i386/ns7/libjavaplugin_oji.so
```

where `$JRE_HOME` is the directory containing the JRE installation.

3. Set the executable permission of the downloaded file:

```
chmod +x downloaded_file_name
```

4. In a terminal window, go to the `$HOME/.mozilla/plugins` directory. Create a `plugins` directory if it does not exist.
5. Remove any existing links to the Java plug-in in this directory. You may use the `rm libjavaplugin_oji.so` command in a terminal window to remove an existing symbolic link to the Java plug-in.
6. Create a symbolic link to the Java plug-in by using the following command:

```
ln -s $JRE_HOME/plugin/i386/ns7/libjavaplugin_oji.so .
```

---

**Note** – Remember the dot at the end of the command.

---

---

**Note** – If you create this symbolic link in any directory other than `$HOME/.mozilla/plugins`, your browser will not be able to use this new Java plug-in.

---

7. 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 to the Java plug-in in the `plugins` directory under the browser's installation directory.

---

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

---

8. 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 15
- “Installing the Certificate by Using Firefox 1.5” on page 16

- “Changing the Security Certificate to Match the Name of the Server” on page 16

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.

Management server installation on Linux requires a hardcoded IP of the server’s IP address to start the `appstormanager` service. (Linux requires the Fully Qualified Domain Name and the IP address on separate lines on `/etc/hosts` for the management server to start. This is the OS default.) Refer to the Installation Guide for additional details.

- To receive a trusted certificate, you need to purchase a certificate from a trusted entity. (Most browsers have trust relationships set up for Verisign, Entrust, and Thawte, among others.) Set the Common Name (CN) to the name of your management server. 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 entering 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 **View Certificate**.
4. When you are shown the certificate information, click **Install Certificate** 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**.

## Installing the Certificate by Using Firefox 1.5

1. Access the management server by entering the following:  
`https://machinename`  
where `machinename` is the name of the management server.
2. When the security alert message appears, click the **Accept this certificate permanently** radio button.
3. Click **OK**.

## 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` program to delete the original certificate, then use the `generateAppiqKeystore` program to create a new certificate and to copy the new certificate to the management server.

### On Microsoft Windows:

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.

### On Sun Solaris and Linux:

1. Go to the <Install\_Dir>/install directory and run the following command:

```
eval `./usersvars.sh`
```

---

**Caution** – The quotes in the example must be entered as left single quotes.

---

2. Go to the following directory:

```
<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.

---

## Restarting 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 restart the service, it is recommended that you restart the service from the Services window, rather than from the command prompt window.

---

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

### On Sun Solaris and Linux:

---

**Caution** – Linux management servers require a fixed IP address for starting the appstormanager service.

---

1. Open a command prompt window.
2. To stop the management server, enter the following at the command prompt:  

```
/etc/init.d/appstormanager stop
```
3. To start the management server, enter the following at the command prompt:  

```
/etc/init.d/appstormanager start
```

4. To see the status of the management server, enter the following at the command prompt:

```
/etc/init.d/appstormanager status
```

---

## 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 Switches, Storage Systems, NAS Devices, and Tape Libraries

---

Before you can use the management server, you must execute the discovery process to make the software aware of the elements on your network, such as switches, storage systems, NAS devices, and tape libraries. Discovery obtains a list of discovered elements and information about their management interface and dependencies.

---

**Note** – The management server can discover only elements with a suitable management interface. Refer to the support matrix for information about supported hardware.

---

This chapter contains the following information:

- “Discovery Steps” on page 22
- “Overview of Discovery Features” on page 24
- “Discover Switches” on page 34
- “Discover Storage Systems, NAS Devices and Tape Libraries” on page 57
- “Building the Topology” on page 88
- “Get Details” on page 90
- “Using Discovery Groups” on page 92
- “Deleting Elements from the Product” on page 96
- “Working with Quarantined Elements” on page 98
- “Updating the Database with Element Changes” on page 99
- “Notifying the Software of a New Element” on page 101
- “Viewing Log Messages” on page 101
- “Viewing the Status of System Tasks” on page 102

---

# Discovery Steps

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

1. Discover your switches. See “Discover Switches” on page 34.
2. Discover your storage systems, tape libraries, and NAS devices. See “Discover Storage Systems, NAS Devices and Tape Libraries” on page 57.
3. If you want to view the topology quickly in System Explorer, obtain the topology as described in “Building the Topology” on page 88 (Optional). 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, including provisioning information. See “Get Details” on page 90.

---

**Note** – Running Get Details takes time. You might want to perform this process when the network and the managed elements are not busy. See “Get Details” on page 90 for more information.

---

## Overall Discovery Tasks

Before you begin the discovery process, note the following:

- Get Details does not default to an automatic schedule. In most cases, we recommend running Get Details once a day during off-peak hours. For more information, see “Get Details” on page 90.
- Make sure the credentials you enter are correct. When credentials are not supplied, the default user names and passwords are tried for the element.
- For elements that support multiple discovery protocols (for example, SNMP and SMI-S), only one protocol at a time is supported for a given element. If you want to change the protocol used to discover an element that has already been discovered, delete the element before attempting to rediscover it with a different protocol. For more information, see “Deleting Elements from the Product” on page 96.
- Elements discovered through SMI-S and hosts discovered with CIM extensions from Build 5.1 and later of the product cannot be added to discovery groups. These elements are listed separately and can be placed independently into scheduled Get Details tasks without being part of a discovery group. This allows you greater flexibility when gathering discovery data. (For more information, see “Creating Custom Discovery Lists” on page 93).

- If you have a problem discovering an element, try enabling Troubleshooting Mode. For more information, see “Troubleshooting Mode” on page 813.
- 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.

---

**Note** – Since CIM extensions have not been installed yet, the management server will not be able to obtain this data when you perform discovery for elements. For more information, see “Deploying and Managing CIM Extensions” on page 155 and “Discovering Applications, Backup Hosts and Hosts” on page 105.

---

- 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. LSI 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 requires additional steps for discovery. See “Discovering EMC Solutions Enabler” on page 60 for more information.

The process for making the management server aware of the elements on your network consists of four stages:

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. For more information, see “Setting Default User Names and Passwords” on page 24.
2. Discover your switches. For information on how to discover the types of switches in your network, see “Discover Switches” on page 34.
3. Discover your storage systems, NAS devices and tape libraries. For more information, see Table 2-3, “Discovery Requirements for Storage Systems, Tape Libraries & NAS Devices,” on page 57.
4. Perform Get Details (**Discovery > Details**), which is required to obtain information from your discovered elements.

---

**Note** – Running Get Details takes time. You might want to perform this process when the network and the managed elements are not busy. See “Get Details” on page 90 for more information.

---

---

# Overview of Discovery Features

Discovery features allow you to:

- 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.
- Import or save discovery settings to a file.

This section contains the following topics:

- “Setting Default User Names and Passwords” on page 24
- “Adding an IP Range for Scanning” on page 27
- “Adding a Single IP Address or DNS Name for Discovery” on page 29
- “Modifying a Single IP Address Entry for Discovery” on page 30
- “Removing Elements from the Addresses to Discover List” on page 31
- “Importing Discovery Settings from a File” on page 31
- “Saving Discovery Settings to a File” on page 33

## 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 that 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.

---

**Caution** – 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, before you begin, 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 **Set Default User Name and Password**.

The Set Default User Name and Password pane appears (Figure 2-1, “Setting Default User Names and Passwords,” on page 26).

**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:

**FIGURE 2-1** Setting Default User Names and Passwords

4. In the User Name box, enter the user name for one or more elements.
5. In the Password box, enter the corresponding password for the user name entered in the previous step.

6. In the Verify Password box, re-enter the password.
7. Repeat steps 4 through 6 for other default user names and passwords you want to add.
8. Click **Add System**.

## Adding an IP Range for Scanning

The management server can be set up so that when scanning, instead of adding each IP address individually the 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 a 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. For Microsoft Windows systems, the check the proxy service status 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 port 2001. The management server does not allow port numbers in the scanning of IP ranges, so you are not able to specify the port. For more information, see “Discovering HDS Storage Systems” on page 65.
- Enter a range within the same subnet. The management server cannot scan IP ranges across subnets.
- One way to detect multiple IP addresses at one time is to add an IP range for scanning. The management server scans the IP range for elements and populates the discovery list with the elements it could contact. You can then discover those elements.

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 **Add Range**.

The Add Range for Scanning pane appears (Figure 2-2, “Adding an IP Range for Scanning,” on page 28).

### 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:*	<input type="text" value="192.168.1.2"/>
To IP Address:*	<input type="text" value="192.168.1.95"/>
User Name:	<input type="text" value="admin"/>
Password:	<input type="password" value="••••"/>
Verify Password:	<input type="password" value="••••"/>
Comment:	<input type="text" value="Servers in Marketing"/>

\* required fields

**FIGURE 2-2** Adding an IP Range for Scanning

4. In the From IP Address box, enter a lowest IP address in the range to be scanned.
5. In the To IP Address box, enter the highest IP address in the range to be scanned.
6. In the User Name box (optional), enter a common user name for elements in the IP range.
7. In the Password box (optional), enter a common password for elements in the IP range.
8. In the Verify Password box, re-enter the password.
9. In the Comment box, enter a brief description of the servers; for example, "Servers in Marketing."
10. Click **OK** to close the Add Range for Scanning pane.
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

The following steps provide general information on how to discover an element. For more information, see Table 2-1, “Discovery Requirements for Switches,” on page 34, Table 2-3, “Discovery Requirements for Storage Systems, Tape Libraries & NAS Devices,” on page 57.

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 **Add Address**.
4. In the IP Address/DNS Name box, enter the IP address or DNS name of the device you want to discover.
5. In the User Name box (optional), enter the user name. This box can be left blank if you are discovering an LSI storage system or if the element's user name and password are one of the default user names and passwords.
6. To set the password, take one of the following actions:
  - If you do not want to do provisioning on a storage system, leave the Password box blank. For LSI storage systems, you must also select the **Do Not Authenticate** option.
  - If you want to do provisioning on a storage system, enter the corresponding password for controller or proxy and make sure the **Do Not Authenticate** option is not selected.
  - For all elements other than storage systems, provide the password if it is required for authentication. If the element does not require a password, leave the Password box blank.
7. If you entered a password in the previous step, re-enter the password in the Verify Password box.
8. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
9. Click **OK**.
10. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

# 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, if the password for a host was changed, you would need to update the management server database with the new password.

---

**Caution** – 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 **Edit** () button for the element whose user name and/or password you want to modify.
3. To change the user name, enter the new user name in the User Name box.
4. To add or change a comment, enter a comment in the Comment box.
5. To change the password:
  - a. Click **Change password**.
  - b. Enter the new password in the New Password box.
  - c. Enter the password again in the Verify Password box.
  - 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 **Get Topology**.

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, and then click **Delete**.
  - Click the **Delete** () button corresponding to the elements you want to remove from the Addresses to Discover list.

---

**Caution** – The elements associated with these addresses are not removed from the management server. For information about how to remove an element from the management server, see “Deleting Elements from the Product” on page 96.

---

## 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 allows you to import the following information to the Discovery list:

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

Note the following:

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

---

**Caution** – When you import a file, your previous settings are overwritten.

---

- If you receive an error message when you try to import the discovery settings, verify that you are using the right password. If you are using the correct password, there is a possibility that the file is corrupt.
- When you save the discovery settings to a file, the management server is not included in the list and you must rediscover the management server. For instructions, see “Importing a File” on page 32 and “Re-discovering the Management Server” on page 32.

## Importing 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 **Browse** to find the file.
  - In the Filename box, enter a complete path to the file.
4. In the Password box, enter the password for the discovery list. If the discovery list did not have a password assign to it, leave this field blank.
5. Click **OK**. The information on the following tabs is updated:
  - IP Addresses
  - IP Ranges
  - Applications
  - Windows Proxy tab

## Re-discovering the Management Server

1. Click **Discovery**, and then click **Setup** in the upper-right pane of the window.
2. Click the **Monitoring Product Health** link.

The Monitoring Product Health window appears.

3. Click **Add**.

The Discovery Setup, Step 1 - Setup page shows the management server as localhost.

4. Select the check box next to localhost and click **Start Discovery**.

When Step 1 discovery is finished, the management server is put into the default discovery group.



5. Select **Discovery > Details**.
6. Run **Get Details** for the discovery group that contains the `localhost` entry.

## 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 information:

- 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 box, enter the password for the management server.
5. In the Verify Password box, enter the password from the previous step, and then click **OK**.
6. When you are asked if you want to open or save the file, choose **Save**.

The Downloading window appears.

7. Enter a name for the \*.xml file and select the directory to which you want to save the file. The default name of the file is `DiscoverySettings.xml`.
8. In the Password box, provide a password for the discovery list.

---

**Note** – This password is required later when you import the file. Choose a password you will remember.

---

9. Click the **Save** button in the Save As window. The file is saved.

---

# Discover Switches

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

**TABLE 2-1** Discovery Requirements for Switches

Element	Discovery Requirements	Additional Information
Brocade switches (SMI-S)	IP address or DNS name, and the user name and password from the Brocade SMI Agent security setup.	See “Discovering Brocade Switches” on page 35.
CNT switches	IP address and the port number for the InVsn Software that manages the switch and the user name and password.	See “Discovering CNT Switches” on page 39.
Cisco switches (SMI-S)	IP address/DNS name of the Cisco switch and the user name and password of the switch.	See “Discovering Cisco Switches” on page 40.
Cisco switches (SNMP)	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 40.
QLogic, and HP M-Series switches (SMI-S)	Enter the IP address/DNS name of the SMI-S switch as well as the user name and password of the switch.	See “Discovering Sun StorEdge, QLogic and HP StorageWorks M-Series for p-Class BladeSystems” on page 42.
Sun StorEdge, HP M-Series, and QLogic switches (SNMP)	IP address/DNS name of the Sun StorEdge, QLogic, or HP M-Series switch. Enter the SNMP read-only community string as the user name. You do not need to enter a password.	See “Discovering Sun StorEdge, QLogic and HP StorageWorks M-Series for p-Class BladeSystems” on page 42.
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 44.

---

**Note** – If your Web browser has an option for blocking pop-ups, disable it. The management server uses pop-ups for dialog boxes.

---

## Discovering Brocade Switches

The management server uses the Brocade SMI-S Provider (also known as the Brocade SMI Agent) to discover Brocade switches. Before you can discover Brocade switches with SMI-S, however, you must first download and install the Brocade SMI Agent software. You can download the Brocade SMI Agent and documentation from the Brocade web site. For more information on Brocade SMI Agent versions, see the support matrix.

---

**Caution** – With this release, discovery of Brocade switches through the Fabric Access API is not supported. For information on migrating existing switches to SMI-S, see “Migrating Brocade API Switches to SMI-S After Upgrading” on page 35.

---

## Migrating Brocade API Switches to SMI-S After Upgrading

After successfully upgrading the management server, any Brocade switches that use the Brocade Fabric Access API provider must be migrated to the Brocade SMI-A provider. The management server will prompt you to migrate your Brocade switches the first time you log on to the management server after the upgrade and will display the Brocade API switches that need to be migrated.

Until you migrate your Brocade switches to SMI-A, data such as topology and zoning from the Brocade switch will be stale and you will be unable to use the Brocade switch to perform provisioning or gather port performance statistics through the Brocade switch. The Brocade Fabric Access API switches are quarantined and you will have the option to migrate to the Brocade SMI-A provider at your discretion in case your SAN policy requires that you validate the new Brocade SMI Agent provider before migrating your Brocade switches.

The quarantined API-managed Brocade switches retain their historical data and that data remains intact during the migration to the SMI-A provider.

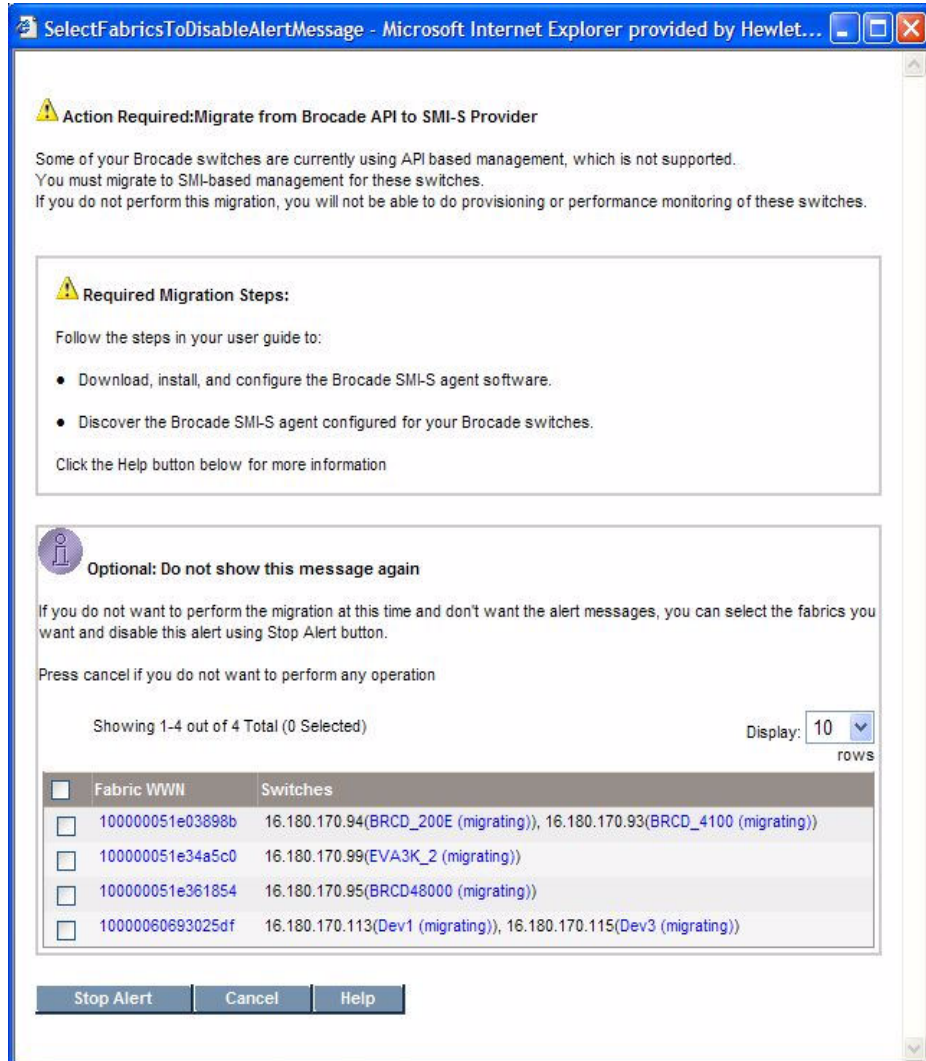
However, new data will not be collected for the quarantined Brocade switches until you migrate the switches to the SMI-A provider.

After migrating the Brocade switches to SMI-A, the Brocade SMI-A proxy server is placed in its own discovery group. This new discovery group is not part of any Get Details schedule. If the Brocade switches were part of a Get Details schedule prior to migration, you must manually adjust those schedules to run Get Details for the migrated Brocade switches. If the schedules are not adjusted manually, Get Details will not run for the migrated switches as per pre-migration schedules.

Follow these steps to migrate your Brocade switches to the Brocade SMI-A provider:

1. Download the Brocade SMI Agent v120.6.0a provider software and its Installation Guide from the Brocade website:  
<http://www.brocade.com/support/SMIAGENT.jsp>  
See the support matrix for your edition for details on the latest supported version for the management server.
2. Install the Brocade SMI Agent with a minimum version of 120.6.0a and configure the proxy servers on the server with which you will manage your Brocade access points following the installation and configuration instructions included in the Brocade v120.6.0a Installation Guide. Refer to the Brocade document for SMI-A requirements.

3. Log on to the management server. The management server alerts you to migrate your Brocade Fabric Access API switches when you first log on.




Your Brocade switches are quarantined until you migrate to the SMI-A provider. The migration message is displayed each time you log on to the management server until each Brocade switch is migrated to the new Brocade SMI-A provider or you choose to disable the message.

4. Run Discovery Step 1 for the Brocade proxy server. See the chapter, "Discovering Switches, Storage Systems, NAS Devices, and Tape Libraries."

5. Run Get Details. See the chapter, “Discovering Switches, Storage Systems, NAS Devices, and Tape Libraries.”

The Brocade switches are migrated to the SMI-A provider.

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

6. If you were using discovery schedules to collect details for Brocade switches prior to migrating them to SMI-A, add the new discovery group for the Brocade proxy server to your pre-existing Get Details schedules as described in the following steps:
  - a. Access the Discovery page by selecting **Configuration > Discovery**.
  - b. Click the **Edit** () button corresponding to the discovery schedule you want to modify.
  - c. Click the **Discovery Groups** tab.
  - d. Select the Brocade proxy under the list of discovery groups.
  - e. Click **Add Selected Groups To Schedule**.
  - f. Click **Finish**.

## To Discover Brocade SMI-S Switches

1. Click **Discovery**, and 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 **Add Address**.
5. In the IP Address/DNS Name box, enter the IP address of the proxy server that is running the SMI-S agent. (Some proxy servers require the following format:  
`http://IPADDRESS.`)
6. In the User Name box, enter the user name for the SMI-S proxy server.

This box 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 box, enter the password for the SMI-S proxy server.

This box 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.
8. If you entered a password in the previous step, re-enter the password in the Verify Password box.
  9. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
  10. Do not select the **Do Not Authenticate** option.
  11. Click **OK**.
  12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

## Discovering CNT Switches

The management server uses the CNT SMI-S provider to discover CNT switches. 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 support matrix for the required InVSN version for your switch model.

---

**Caution** – 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.

---

When discovering CNT switches, note the following:

- SNMP is not supported for CNT switches.
- 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. Take the following steps 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. Make the following entry in the file:  
`cimomenabled=TRUE`
  - c. Save the file, and then restart the InVsn software.
2. In the IP Address/DNS Name box, enter 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 box, enter the user name for the login to the InVsn software.
4. In the Password box, enter the password for the login to the InVsn software.
5. In the Verify Password box, enter the password you provided in step 4.
6. Click **Start Discovery**.

## Discovering Cisco Switches

The management server discovers Cisco switches through SNMP and SMI-S connections depending on the switch model. See the support matrix for details on supported switch models and firmware revisions.

Note the following when discovering Cisco switches with SNMP:

- 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. For more information, see “Changing the Amount of Information Collected from the Inactive Zone Database (Cisco Switches)” on page 419.
- The management server groups active zone sets in all Virtual SANs (VSANs) in a fabric into a zone set called ACTIVE, which 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.
- No ports are reported for uninstalled blades or GBICs.



- To receive events from Cisco switches, verify that the SNMP trap community string is set to match the community string defined in the custom properties (the default is `public`), and make sure the SNMP traps are configured to be sent to the management server. For more information, see “Changing the SNMP Trap Listener Port and Community String for Switches Discovered with SNMP” on page 57.

Note the following when discovering Cisco switches with SMI-S:

- Before you can discover Cisco switches with SMI-S, you must download and install the Cisco cimserver software. See your Cisco documentation for more information.
- Enable the CIM Server for Cisco switches discovered through the SMI-S provider.
  - a. On the Cisco switch, enter the following command to display the Common Information Models (CIM) configurations and settings:  
`cisco_switch# show cimserver`
  - b. To enter configuration mode, enter the following:  
`cisco_switch# config`
  - c. To enable access to the server, enter the following:  
`cisco_switch# cimserver enableHttps`  
 And/or  
`cisco_switch# cimserver enableHttp`
  - d. To enable the CIM Server, enter the following:  
`cisco_switch(config)# cimserver enable`
  - e. To exit configuration mode, enter the following:  
`cisco_switch(config)# exit`
- When you discover a Cisco SMI-S switch you need to provide a user name and password.
- If you are using the SMI-S provider, discover all Cisco switches in a fabric. If you discover only one switch, inactive zones and zone sets residing on other switches are not displayed on the management server.

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 **Add Address**.
5. In the IP Address/DNS Name box, enter the DNS name or primary IP address of the Cisco switch you want to discover.
6. Take one of the following actions:

- For **Cisco** switches with SNMP connections:  
In the User Name box, enter the user name for the switch. This is the public community SNMP string for the switch. This box 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 box, enter the switch user name.
7. Take one of the following actions:
- For **Cisco** switches with SNMP connections:  
Leave the Password box blank.
  - For **Cisco** switches with SMI-S connections:  
In the Password box, enter the switch password.
8. Take one of the following actions:
- For **Cisco** switches with SNMP connections:  
Leave the Verify Password box blank.
  - For **Cisco** switches with SMI-S connections:  
In the Verify Password box, enter the switch password again.

## Discovering Sun StorEdge, QLogic and HP StorageWorks M-Series for p-Class BladeSystems

The management server discovers Sun StorEdge switches through an SNMP connection and QLogic and HP M-Series switches are discovered through SNMP or SMI-S. See the support matrix for details on supported switch models and firmware revisions.

Note the following when discovering these switches with SNMP:

- When you discover these switches, you do not need to provide a password.
- The management server does not support provisioning for Sun StorEdge, QLogic, and HP M-Series switches. Only the active zone set and its zone members are reported.
- To manage a fabric of Sun StorEdge, QLogic, or HP M-Series 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 switches is 162. To change this port, see "Changing the SNMP Trap Listener Port and Community String for Switches Discovered with SNMP" on page 57.
- To receive events from Sun StorEdge, QLogic, and HP M-Series switches, verify that the SNMP trap community string is set to match the community string defined in the custom properties (the default is `public`), and make sure the

SNMP traps are configured to be sent to the management server. For more information, see “Changing the SNMP Trap Listener Port and Community String for Switches Discovered with SNMP” on page 57.

Note the following when discovering these switches with SMI-S:

- Before you can discover these switches with SMI-S, you must download and install the cimserver software. See your switch documentation for more information.
- 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.

---

**Note** – You may see an error replicating the switch fabric name for QLogic-based switches. This error can be ignored.

---

To discover Sun StorEdge or QLogic switches:

1. Select **Discovery > Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the DNS name or primary IP address of the switch you want to discover.
6. Take one of the following actions:
  - For switches with SNMP connections:  
In the User Name box, enter the user name for the switch. This is the public community SNMP string for the switch. This box can be left blank if the element's user name and password are one of the default user names and passwords.
  - For switches with SMI-S connections:  
In the User Name box, enter the user name for this switch. All SMI-S switches require a user name and password.
7. Take one of the following actions:
  - For switches with SNMP connections:  
Leave the Password box blank.
  - For switches with SMI-S connections:  
In the Password box, enter the password for this switch.
8. Take one of the following actions:
  - For switches with SNMP connections:  
Leave the Verify Password box blank.

- For switches with SMI-S connections:  
In the Verify Password box, enter the password of the switch again.

## Discovering McDATA and EMC Connectrix Switches

McDATA and EMC Connectrix switches use SMI-S, the Fibre Channel Switch Application Programming Interface (SWAPI), or SNMP to communicate with devices on the network. The management server can discover multiple instances of Enterprise Fabric Connectivity (EFC) Manager. Use one of the following methods to discover McDATA and Connectrix switches:

**TABLE 2-2** Discovery Settings for McDATA and Connectrix Switches

Method	Description
<b>SMI-S Discovery</b>	SMI-S is the default discovery method for new installations. The SMI-S 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 and nicknames are not supported.
<b>SWAPI setting through a Proxy</b>	You will need to connect through the proxy instead of the switch. For more information, see “Discovering McDATA and Connectrix Switches through a Proxy with SWAPI” on page 47. 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.
<b>SNMP setting Through a Proxy</b>	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. For more information, see “Discovering McDATA and Connectrix Switches through a Proxy with SNMP” on page 49. This SNMP setting through a proxy does not let you manage or access information about zones, zone sets or zone aliases.
<b>Contacting the switch directly</b>	Contact the switch by its IP address or DNS name. This connection uses SNMP. See the support matrix for details on switch models ( <b>Help &gt; Documentation Center</b> ). For more information, see “Discovering McDATA and Connectrix Switches through a Direct Connection and SNMP” on page 50. 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.

Keep in mind the following:

- SMI-S is the default method for discovering McDATA and Connectrix switches. If you need to migrate to SMI-S or change the discovery settings, see “Changing the Discovery Settings” on page 52.

- You can only choose one discovery method for McDATA and Connectrix switches. For example, if you use SMI-S, you cannot discover additional McDATA and Connectrix switches with SWAPI or SNMP.
- If you use EFC Manager or Connectrix Manager, see the support matrix to verify the version requirements.
- Brocade 5000ni switches running in McDATA mode are managed by the Brocade SMI Agent and not by McDATA SMI-S. For more information, see “Discovering Brocade Switches” on page 35.
- If you change the discovery settings, the user ID and password will no longer work. For this reason, set this property before discovering any McDATA or Connectrix switches. If you must change the configuration, see “Changing the Discovery Settings” on page 52.
- After you discover a McDATA or Connectrix switch through a proxy, the IP address displayed next to the name of the switch is 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 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 “Managing McDATA and EMC Connectrix Switches” on page 54.
- All McDATA or Connectrix switches in a fabric must be managed by the same EFC Manager or Connectrix Manager. Do not have more than one EFC Manager or Connectrix Manager to a fabric for McDATA or Connectrix switches.
- If you want the management server to receive SNMP traps 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 only on the Connectrix Manager or EFC Manager, not on the individual switches.
  - If you discovered Connectrix or McDATA switches directly, enable SNMP trap forwarding on the switches, not in any other management software.
- For more information about the SNMP port and community string, see “Changing the SNMP Trap Listener Port and Community String for Switches Discovered with SNMP” on page 57.

## Discovering McDATA and Connectrix switches with SMI-S

Before you can discover McDATA and Connectrix switches with SMI-S, you must download and install the McDATA SMI-S provider software. See your switch documentation for more information.

Note the following when discovering these switches with SMI-S:

- Before attempting to discover your switches, ensure that EFC Manager or Connectrix Manager is installed and configured or add your switches to the SMI-S provider.

- For upgrades only: To migrate your existing switches to SMI-S, follow the procedure in “Changing the Discovery Settings” on page 52.
- Discovering McDATA and Connectrix switches with SMI-S is the default setting. To view or change the discovery settings, see “Changing the Discovery Settings” on page 52.
- You can install only one instance of the SMI-S provider on the management station.
- Installation of the McDATA SMI-S provider is not supported on Linux systems.
- A McDATA or Connectrix switch cannot be managed by more than one SMI-S provider.
- When you install the SMI-S provider, there are two modes:
  - In coexist mode the SMI-S provider communicates with EFC Manager or Connectrix Manager and adds all the switches in the managed list of EFC Manager or Connectrix Manager.
  - In direct mode, you must add each switch to the SMI-S provider with its IP address, credentials and switch type. You can use a McDATA’s `manageswitch.bat` file to manage the addition and deletion of switches.
- If you selected direct mode during the SMI-S provider installation, when you add switches, you must enter the switch type based on the McDATA model number even if your switch is an OEM model. For more information about the switch type, see your McDATA documentation.
- The SMI-S provider can be installed on the same server as EFC Manager or Connectrix Manager.
- If you selected coexist mode during the SMI-S provider installation you can have only one EFC Manager or Connectrix Manager server.
- If you are using EFC Manager or Connectrix Manager you cannot add managed switches in direct mode. To add switches in direct mode you must remove them from EFC Manager or Connectrix Manager first.
- If the SMI-S provider is installed on a machine other than the management server, network links between them must pass http traffic on port 5988 (default) or https on port 5989. The port used by the SMI-S provider can be configured. See your switch documentation for more information.

To discover the proxy:

1. Select **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 **Add Address**.
5. In the IP Address/DNS Name box, enter the IP address or DNS name of the proxy you want to discover.
6. In the User Name box, enter the user name.
7. In the Password box, enter the password.

---

**Note** – The user name and password are defined during the SMI-S provider installation. These credentials might be different from the EFC Manager credentials.

---

8. Re-enter the password in the Verify Password box.
9. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.  
  
Discovery is complete when the software displays the DISCOVERY COMPLETED message in the Log Messages box.

---

**Caution** – 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 “About Get Details” on page 90.

---

## Discovering McDATA and Connectrix Switches through a Proxy with SWAPI

With the SWAPI setting, the management server contacts a proxy to obtain information about the switches connected to it. Use EFC Manager or Connectrix Manager for this option. If you do not have EFC Manager or Connectrix Manager, see “Discovering McDATA and Connectrix Switches through a Direct Connection and SNMP” on page 50.

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.

---

**Caution** – EMC customers using the EMC Connectrix Manager (EMC’s rebranded EFC Manager) 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 it

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 cannot discover EMC Connectrix switches using SWAPI. You must discover them through the SNMP provider, either directly or through a proxy. For more information, see "Discovering McDATA and Connectrix Switches through a Proxy with SNMP" on page 49 and "Discovering McDATA and Connectrix Switches through a Direct Connection and SNMP" on page 50 .

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 Connectrix switches using SWAPI, you should check with your EMC or McDATA representative to determine the implications of this configuration.

---

1. For McDATA switches only, install the McDATA Bridge Agent. To communicate with EFC Manager, the management server requires the Bridge Agent. Consult your McDATA representative for more information about the Bridge Agent.
2. Change the discovery setting for McDATA and Connectrix switches to SWAPI following the steps in "Changing the Discovery Settings" on page 52.
3. Discover the Proxy:
  - a. Click **Discovery**, then click **Setup** in the upper-right pane of the window.
  - b. Select **Step 1** at the top of the page.
  - c. Click the **IP Addresses** tab.
  - d. Click **Add Address**.
  - e. In the IP Address/DNS Name box, enter the IP address or DNS name of the EFC Manager or Connectrix Manager you want to discover.
  - f. In the User Name box, enter the user name for EFC Manager or Connectrix Manager.

This box 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.

In the Password box, enter the corresponding password for EFC Manager or Connectrix Manager.

This box can be left blank if one or more of the following conditions is fulfilled:



The element's user name and password are one of the default user names and passwords.

The element does not require authentication.

- g. If you entered a password in the previous step, re-enter the password in the Verify Password box.
- h. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery > Setup**).
- i. Click **OK**.
- j. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

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

## Discovering McDATA and Connectrix Switches through a Proxy with SNMP

---

**Note** – Discovering McDATA or Connectrix switches through a proxy using the SNMP protocol does not let you manage or access information about zones, zone sets or zone aliases.

---

You can use this option with EMC Connectrix™ Manager and Enterprise Fabric Connectivity (EFC) Manager to contact the switch.

1. Change the discovery setting for McDATA and Connectrix switches to SNMP following the steps in “Changing the Discovery Settings” on page 52.
2. Discover the Proxy:
  - a. Select **Discovery**, then click **Setup** in the upper-right pane of the window.
  - b. Select **Step 1** at the top of the page.
  - c. Click the **IP Addresses** tab.
  - d. Click **Add Address**.
  - e. In the IP Address/DNS Name box, enter the IP address or DNS name of the proxy you want to discover.

- f. In the User Name box, enter 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.
- g. Leave the Password and Verify Password boxes blank. The password does not matter since the management server is not doing any configurations through SNMP.
- h. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
- i. Do not select the **Do Not Authenticate** option.
- j. Click **OK**.
- k. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

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

---

**Caution** – 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 “About Get Details” on page 90.

---

- 3. Make sure there are no port conflicts for receiving SNMP traps. When the management server is configured to contact the proxy by SNMP, it 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. For information about changing the port, see “Changing the SNMP Trap Listener Port and Community String for Switches Discovered with SNMP” on page 57.
- 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 from all switches managed by the proxy to the management server using the port you selected. For more information, see the documentation for your proxy.

## Discovering McDATA and Connectrix Switches through a Direct Connection and SNMP

The management server uses SMI-S or SWAPI to discover a McDATA or Connectrix switch through a proxy. If you want to discover McDATA or Connectrix switches directly, you must change the discovery settings to SNMP before you begin the following steps. See “Changing the Discovery Settings” on page 52. See the support matrix for McDATA switch details (**Help > Documentation Center**).

To discover a McDATA or Connectrix switch directly:

1. Make sure there are no port conflicts for receiving SNMP traps. When the management server is configured to contact the proxy by SNMP, it 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. For information about changing the port, see “Changing the SNMP Trap Listener Port and Community String for Switches Discovered with SNMP” on page 57.
2. 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 then sends traps from all switches managed by that proxy. See the proxy documentation for more information.
3. Select **Discovery**, then click **Setup** in the upper-right pane of the window.
4. Select **Step 1** at the top of the page.
5. Click the **IP Addresses** tab.
6. Click **Add Address**.
7. In the IP Address/DNS Name box, enter the IP address or DNS name of the switch you want to discover.
8. In the User Name box, enter 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.
9. If you are using SNMP leave the Password box (optional) blank. The password does not matter since the management server is not doing any configurations through SNMP. If you are using SMI-S enter the password of the admin account on the switch.
10. In the Verify Password box enter the same thing you entered in the password box.

11. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery > Setup**).
12. Do not select the **Do Not Authenticate** option.
13. Click **OK**.
14. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.  
  
Discovery is complete when the software displays the `DISCOVERY COMPLETED` message in the Log Messages box.

---

**Caution** – 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 “About Get Details” on page 90.

---

## Changing the Discovery Settings

To change the discovery settings for McDATA and Connectrix switches:

1. If you have already discovered your switches, 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, and then click **Delete**.
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 **Delete**.
3. Select **Configuration > Product Health**, and then click **Advanced** in the **Disk Space** tree.
4. Click **Show Default Properties** at the bottom of the page.

To enable SNMP:

- a. Uncomment the `cimom.useSnmpMcDataProvider` property by removing the pound sign (#) in front of it.
- b. Change the `cimom.mcdata.dontUseSmis` property as follows:  
`cimom.mcdata.dontUseSmis=true`

---

**Note** – The `cimom.mcdata.dontUseSmis` property exists only in upgrade installations. If the property does not exist on your system, enter it manually.

---

To enable SWAPI:

- a. Comment out the `cimom.useSnmppMcDataProvider` property by placing a pound sign (#) in front of it.
- b. Change the `cimom.mcdata.dontUseSmis` property as follows:  
`cimom.mcdata.dontUseSmis=true`

---

**Note** – The `cimom.mcdata.dontUseSmis` property exists only in upgrade installations. If the property does not exist on your system, enter it manually.

---

To enable SMI-S:

- a. Comment out the `cimom.useSnmppMcDataProvider` property by placing a pound sign (#) in front of it.
  - b. Change the `cimom.mcdata.dontUseSmis` property as follows:  
`cimom.mcdata.dontUseSmis=false`.
5. Click **Save**.
  6. Discover the switch. For instructions, see “Discovering McDATA and EMC Connectrix Switches” on page 44.

---

**Note** – If you change the discovery settings, when you discover the switch with the new method, make sure you enter the correct credentials. For example, if you change from SNMP to SMI-S, the required credentials are different. See the section for the specific discovery method for information on the credentials to enter.

---

## 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, modify the `cimom.mcdata.exclude` property. Set the property `cimom.mcdata.exclude` to a comma-separated list of Worldwide Names (WWN) of the McDATA and Connectrix switches you want excluded, as shown in the following example:

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

The management server excludes the switches with the following WWNs: 1000080088A07024 and 1000080088A0D0B6

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

---

**Caution** – 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**) when a provider instance has been created for an excluded element. You can ignore this log message.

---

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

1. Select **Configuration > Product Health**, and then click **Advanced** in the **Disk Space** tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the `cimom.mdata.exclude` property.
4. Return to the Advanced page (**Configuration > Product Health**, and then click **Advanced** in the Disk Space tree).
5. Paste the copied text into the Custom Properties box.
6. Make your changes to the text in the Custom Properties box. Remove the pound (#) symbol in front of the property to make sure it is not commented out.
7. Add the WWNs corresponding to the switches you want to exclude from discovery. Separate additional WWNs with a comma, as shown by the following example:  

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

where 1000080088A07024 and 1000080088A0D0B6 are the WWN for McDATA and Connectrix switches.
8. When you are done, click **Save**.

## 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 by performing Get Details to obtain information about the new switches from the service processor. For more information about adding switches, see, “Adding McDATA and EMC Connectrix Switches” on page 54.

When you remove switches from the service processor, you must remove them from the management server. For more information about removing switches, see “Removing McDATA and EMC Connectrix Switches” on page 55.

When you replace McDATA or EMC Connectrix switches, you add and remove the switches as described previously. For more information, see “Replacing McDATA and EMC Connectrix Switches” on page 56.

### *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 44.

---

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

---

To run Get Details:

1. Select **Discovery > Details**.
2. Click **Get Details**.

During Get Details, the software status light changes from green to red. You can view the progress of gathering details by accessing the logs. For more information, see “Viewing Log Messages” on page 101.

### *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. Delete 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 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 that the switches have been removed from the elements list in Discovery Steps 2 and 3 by taking the following steps:
  - a. To verify that the switches have been removed from the element list in Discovery Step 3, select **Discovery > Details**.
  - b. To verify that the switches have been removed from the element list in Discovery Step 2, select **Discovery > Topology**.

### *Replacing 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 44.

To swap the switches, perform the following steps on the management server:

1. Delete the switches from the user interface by taking the following steps (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 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 that the switches have been removed from the element list in Discovery Steps 2 and 3 by doing the following:
  - a. To verify that the switches have been removed from the element list in Discovery Step 2, select **Discovery > Topology**.
  - b. To verify that the switches have been removed from the element list in Discovery Step 3, select **Discovery > Details**.
3. Select **Discovery > Details**.



4. Click **Get Details**.

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

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

## Changing the SNMP Trap Listener Port and Community String for Switches Discovered with SNMP

The default SNMP trap listener port for all switches is 162. To change this port for all switches that are discovered through SNMP, modify the `cimom.snmpTrapListenerPort` property.

The default SNMP trap community string is public. To change this port for all switches that are discovered through SNMP, modify the `cimom.snmpTrapListenerCommunityString` property.

1. Select **Manage Product Health**, and then click **Advanced** in the Disk Space tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Do one of the following:
  - Copy the `cimom.snmpTrapListenerPort` property.
  - Copy the `cimom.snmpTrapListenerCommunityString` property.
4. Click **Close** to return to the Advanced page.
5. Paste the copied text into the Custom Properties box.
6. Make your SNMP trap listener port or SNMP trap community string change in the Custom Properties box. Remove the pound (#) symbol in front of the property to make sure it is not commented out. For example:  
`cimom.snmpTrapListenerPort=162.`
7. Click **Save**.

# Discover Storage Systems, NAS Devices and Tape Libraries

The following table provides an overview of the discovery requirements for storage systems, NAS devices and tape libraries.

**TABLE 2-3** 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 59.
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 63 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” on page 60 for more information.
LSI storage systems	Can be discovered two ways: <ul style="list-style-type: none"><li>• Entering the IP address/DNS name, user name and password of a controller for an LSI storage system. Discovers only the corresponding IP address of the controller.</li><li>• Entering the IP address/DNS name, user name and password of a proxy that is used to manage an LSI storage system. Discovers all controllers known to the proxy.</li></ul>	See “Discovering LSI Storage Systems” on page 64.
HDS storage systems	Discover the server running HiCommand Device Manager.	See “Discovering HDS Storage Systems” on page 65 for more information.
HP MSA storage systems	Discover the server running the MSA SMI-S provider.	See “Discovering HP StorageWorks MSA Arrays” on page 68.
HP EVA storage systems	Discover the server running Command View EVA.	See “Discovering HP StorageWorks EVA Arrays” on page 69.
HP XP storage systems	Discover the server running the SMI-S provider or the built-in provider.	See “Discovering HP StorageWorks XP Arrays” on page 73.

**TABLE 2-3** Discovery Requirements for Storage Systems, Tape Libraries & NAS Devices

Element	Discovery Requirements	Additional Information
IBM Storage Systems	Discover the CIMOM that talks to the IBM storage systems you want to monitor.	See “Discovering IBM Storage Systems” on page 76.
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 78.
Sun StorEdge 6920 and 6940	Discover the storage system directly.	See “Discovering Sun StorEdge 6920 and 6940 Storage Systems” on page 80.
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 80.
Xiotech Storage Systems	Discover the storage system directly.	See “Discovering Xiotech Storage Systems” on page 81.
HP NAS Devices	Discover the device directly.	See “Discovering HP NAS Devices on Windows” on page 82 and “Discovering HP NAS Devices on Linux” on page 83.
NetApp Devices	Discover the device directly.	See “Discovering NetApp NAS Devices” on page 84.
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	Discover the server running the SMI-S provider for the tape library.	See “Discovering HP and IBM Tape Libraries” on page 87

## Discovering 3PAR Storage Systems

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, start the InForm CLI and run the following command:

```
startcim
```

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

---

**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.

---

To discover a 3PAR storage system:

1. Select **Discovery > Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter 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. Re-enter the password in the Verify Password box.
9. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

## Discovering EMC Solutions Enabler

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. See the EMC documentation for details.

To discover Symmetrix storage systems, you must create and configure a VCM volume on the storage system. The VCM database on the Solutions Enabler host must also be configured. For more information, see the *EMC Solutions Enabler Symmetrix CLI Command Reference*.

---

**Caution** – 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. You 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)

## Using Only One Subnet

To allow Solutions Enabler to respond correctly, limit the management server to a single subnet. If your management server is on two or more subnets, discovering a storage array through Solutions Enabler might not work. Limiting the management server to a single subnet allows Solutions Enabler to respond correctly.

## Excluding EMC Symmetrix Storage Systems from Discovery

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

To exclude one or more Symmetrix storage systems from discovery, 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 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.

---

**Caution** – The IP addresses of excluded elements appear in the discovery (**Discovery > Setup**), topology (**Discovery > Topology**) and 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 when 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**, and 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.symmetrix.exclude=000183500570,000183500575
```
4. Click **Close** to return to the Advanced page.
5. Paste the copied text into the Custom Properties box.
6. Remove the pound (#) symbol in front of the property to make sure it is not commented out. 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**.

## 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 the 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.

```
#cimom.emc.skipRefresh=000183500570,000183500575
```
4. Click **Close** to return to the Advanced page.
5. Paste the copied text into the Custom Properties box.
6. Remove the pound (#) symbol in front of the property to make sure it is not commented out. 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, and then click the **Properties** tab.
7. When you are done, click **Save**.
8. To force a refresh for elements that are not configured to skip the refresh, select the **Force Device Manager Refresh** option on the Get Details page.
9. Click **Get Details**.

## 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. 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.

---

**Caution** – Before you discover a CLARiiON storage system, you must have already installed all required software components for that CLARiiON storage system, such as the Navisphere Host Agent. See 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 LSI Storage Systems

When discovering LSI storage systems, note the following:

- Discover all controllers on an LSI storage system by entering the IP address of each controller. The management server discovers these controllers as one single storage system.
- The management server must have the User Name box populated to discover the LSI storage system. Even if your LSI storage system does not have a user name set, you must enter something in the User Name box.
- To obtain drive-related statistics, install a proxy host. Ensure that the proxy host has at least one LUN rendered by each controller of the array.
- 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.
- LSI storage systems do not require a password for Get Details. If you want do not want to use the management server for provisioning on LSI storage systems, select the **Do Not Authenticate** option. The management server will still monitor the LSI storage system; however, you will not be able to do provisioning tasks.

Do the following to discover LSI storage systems:

1. Select **Discovery > Setup**.
2. Select **Step 1** at the top of the page.



3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the IP address or DNS name of the controller or proxy you want to discover.
6. Enter the user name in the User Name box. If your LSI storage system does not have a user name, you must enter something in the User Name box, even though the storage system has no user name.
7. Leave the Password box blank if you do not want to do provisioning on the LSI storage system. If you want to do provisioning, enter the corresponding password for controller or proxy.
8. If you entered a password in the previous step, re-enter the password in the Verify Password box.
9. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in 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. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

## 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 port 2001. The management server does not allow port numbers in the scanning of IP ranges, so 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 835.

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

1. Access the Discovery Setup page (**Discovery > Setup**).
2. Click **Add Address**.
3. In the IP Address/DNS Name box, enter 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 box, enter the user name for accessing HiCommand Device Manager.
5. In the Password box, enter the password for accessing HiCommand Device Manager.
6. In the Verify Password box, re-enter the password for accessing HiCommand Device Manager.
7. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in 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 can 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**) when 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**, and 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.hds.exclude=61038,61037
```
4. Click **Close** to return to the Advanced page.
5. Paste the copied text into the Custom Properties box.
6. Remove the pound (#) symbol in front of the property to make sure it is not commented out. 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**.

## 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.

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.

---

**Caution –** 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**, and 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.HdsSkipRefresh=61038,61037
```

4. Click **Close** to return to the Advanced page.

5. Paste the copied text into the Custom Properties box.

6. Remove the pound (#) symbol in front of the property to make sure it is not commented out. 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.

---

**Note –** To find the serial number, double-click the storage system in System Explorer, and then click the **Properties** tab.

---

7. When you are done, click **Save**.
8. To force a refresh for elements that are not configured to skip the refresh, select the **Force Device Manager Refresh** option on the Get Details page.
9. Click **Get Details**.

## Discovering HP StorageWorks MSA Arrays

Before you can discover MSA arrays, you must download and install the HP MSA SMI-S Provider software. See your array documentation for more information. Keep in mind the following:

- To determine provisioning support for HP StorageWorks Arrays, see Table 11-3, “Provisioning and Pool Support,” on page 422 and Table 11-4, “Volume and Host Security Group Support,” on page 423.
- The Array Configuration Utility (ACU) application should not be running when the management server is using the MSA provider.
- The management URL on the Properties page for the MSA can be used only if the ACU is installed on the same host as the SMI-S provider and the Execution Mode is set to Remote Service. See the ACU *Readme* file for information about execution modes and how to change them.
- Selective Storage Presentation (SSP) for the array must be enabled for provisioning to work.
- MSA volumes must be deleted in the reverse order of their creation. For example, if you have six volumes, and you want to delete the second one you created, you must delete the volumes one at a time, starting with the volume created sixth and continuing with the fifth, fourth, third, and then the second.
- The MSA SMI-S provider updates its cache every four minutes. If the array is managed by an application other than the management server, changes to the array configuration might not be reflected by a Get Details task that ran before the cache update.

To discover HP MSA storage systems:

1. Select **Discovery > Setup** in the upper-right pane of the home page window.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the IP address or DNS name of the HP CIMOM you want to discover.
6. Enter the user name used to access the MSA SMI-S provider.
7. Enter the password used to access the MSA SMI-S provider.

8. If you entered a password in the previous step, re-enter the password in the Verify Password box.
9. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

## Discovering HP StorageWorks EVA Arrays

The management server uses the built-in EVA provider. Before discovering EVA arrays, note the following:

- HP StorageWorks Command View EVA must be installed on a server before you can discover an HP EVA storage system.
- If you have both active and standby Command View EVA proxy machines, you can discover both the proxy machine that is actively managing the array, and the proxy machine that is not actively managing the array. If you discover only the proxy machine that is not actively managing the array, then only top level array information is collected.

If both proxy machines are discovered, keep them in the same discovery group. They can be moved to other discovery groups, but they must be moved together to the same group at the same time. When discovering the proxy machines separately, the machine that has already been discovered must be in the Default discovery group. For more information about discovery groups, see “Using Discovery Groups” on page 92.

- To determine provisioning support for HP StorageWorks Arrays, see Table 11-3, “Provisioning and Pool Support,” on page 422 and Table 11-4, “Volume and Host Security Group Support,” on page 423.
- EVA arrays can only be provisioned if they are actively managed by the Command View server that they are discovered through.
- When an EVA is discovered by the built-in EVA provider, a cache is created and populated with the current array configuration. Each subsequent cache refresh will start 30 minutes after completion of the previous cache refresh. The time between cache refresh starts depends on factors such as the EVA configuration, model, and SAN traffic.

When you perform a provisioning operation (creating, deleting, or modifying a pool or volume), the cache information about provisioning is immediately updated. If you provision an EVA using Command View EVA or a different management station, the cached information about the EVA will not be accurate until the cache is refreshed.

To discover an EVA array:

1. Select **Discovery > Setup** in the upper-right pane of the management server's home page 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 box, enter the IP address of the Command View server.
6. Enter the user name used to access the Command View server.
7. Enter the password used to access the Command View server.
8. If you entered a password in the previous step, re-enter the password in the Verify Password box.
9. (Optional) In the Comment box, enter any additional information. The information entered in this box is displayed in the Comment column in the Addresses to Discover list.
10. Do not select the Do Not Authenticate option.
11. Click **OK**.
12. To start discovering elements on the network, check the check box next to the elements you want to discover, and click **Start Discovery** on the IP Addresses tab.

## Obtaining SNMP Traps using Command View EVA

You must configure Command View EVA so it can send SNMP traps from the EVA to the management server. When the management server receives these SNMP traps, it converts them to WBEM Indications for display in its Event Manager.

## *Community String Requirements*

- The default community string for Command View EVA 6.x is `Public` and the default community string for is `public`. The community strings must be a case-sensitive match, so if you are using the default values in the management server and Command View EVA 6.x, you must change the community strings to a case-sensitive match.
- If you are using the default community strings for Command View EVA 7.x and the management server, no changes to the community strings are needed. If you change the community strings to non-default values, then they must be a case-sensitive match.

---

**Caution** – Other applications may be using the default community strings to communicate with Command View EVA. If you change the community string in Command View EVA, you might break Command View EVA’s connection to other applications. If a change is needed, we recommend changing the community string on the management server to match the string in Command View EVA.

---

## *Obtaining SNMP traps from Command View*

To obtain SNMP traps from Command View EVA:

1. Verify that the community strings follow the rules in “Community String Requirements” on page 71. For information on viewing or changing community strings, see “Viewing or Changing the Community String” on page 72, “Viewing or Changing the Community String in Command View EVA 6.x” on page 72, or “Viewing or Changing the Community String in Command View EVA 7.x” on page 72.
2. Configure event and host notification. For instructions, see “Configuring event and host notification in Command View EVA” on page 73.

## *Viewing or Changing the Community String*

To view or change the community string:

1. Select .
2. Click **Advanced** in the Disk Space tree.
3. Click **Show Default Properties** at the bottom of the page.
4. Copy the `cimom.snmpTrapListenerCommunityString` variable.  
The management server uses the value that is listed last, so be sure to search to the end of the page to locate the latest build.



5. Click **Close** to return to the Advanced page.
6. Paste the copied text into the Custom Properties box.
7. Change the value by entering  
`cimom.snmpTrapListenerCommunityString=<value>` where <value> is the desired community string value.
8. Click **Save**.

### *Viewing or Changing the Community String in Command View EVA 6.x*

To view or change the community string:

1. Open the `c:\hsvmafiles\nsaserver.ini` file in a text editor on the Command View EVA server.
2. Find the line `Authority=Public`  
This example shows the Command View EVA 6.x default: `Public`.
3. Change the value to the desired community string. For example, if you want to change the community string to `public`, enter `Authority=public`
4. Restart the service for Command View EVA.

### *Viewing or Changing the Community String in Command View EVA 7.x*

To view or change the community string:

1. Open the `C:\Program Files\Hewlett-Packard\Sanworks\Element Manager for StorageWorks HSV\config\cveva.cfg` file in a text editor on the Command View EVA server.
2. Find the following command lines:  

```
# Authority. Default = Public
authority Public
```
3. Change the community string to the desired value. For example, if you want to change the community string to `public`, enter `authority public`
4. Restart the service for Command View EVA.

## *Configuring event and host notification in Command View EVA*

See the HP StorageWorks Command View EVA user guide for instructions on configuring Command View EVA event notification.

## Discovering HP StorageWorks XP Arrays

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

- “Discovering HP XP Arrays by Using Command View XP and SMI-S” on page 73
- “Discovering HP XP Arrays Using Command View XP Advanced Edition” on page 74
- “Discovering HP XP Arrays by using the built-in XP Provider” on page 75

---

**Note** – To determine provisioning support for HP StorageWorks Arrays, see Table 11-3, “Provisioning and Pool Support,” on page 422 and Table 11-4, “Volume and Host Security Group Support,” on page 423.

---

---

**Note** – HP StorageWorks Command View XP should be installed on a server before you discover an HP XP storage system.

---

## Discovering HP XP Arrays by Using Command View XP and SMI-S

Before you can discover XP arrays, you must download and install the XP SMI-S Provider software. See the support matrix for details.

---

**Caution** – The Command View XP SMI-S provider does not return information related to external storage available to the HP XP storage arrays, including the external LDEVs. As a result, that information is not available in the management server user interface or reports.

---

To discover an HP XP array using Command View XP and SMI-S:

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

5. In the IP Address/DNS Name box, enter the IP address or DNS name of the Command View XP server you want to discover.
6. Enter the user name for accessing the XP SMI-S provider.
7. Enter the password for accessing the XP SMI-S provider.  
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 it for the default password.
8. If you entered a password in the previous step, re-enter the password in the Verify Password box.
9. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

## Discovering HP XP Arrays Using Command View XP Advanced Edition

HP StorageWorks Command View XP Advanced Edition must be installed on a server before you discover an HP XP storage system.

To discover an HP XP array using Command View XP Advanced Edition:

1. Select **Discovery > Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter 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 entered a password in the previous step, re-enter the password in the Verify Password box.

9. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

## Discovering HP XP Arrays by using the built-in XP Provider

To discover an HP XP array using the built-in XP Provider:

1. Select **Discovery > Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter 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.

---

**Note** – The account must be a Partition Storage Administrator account.

---

8. If you entered a password in the previous step, re-enter the password in the Verify Password box.
9. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

# 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 IBM 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 best not to install the IBM CIM Agent on the management server. For more information, see the *CIM Agent for DS Open (API) - Installation and Configuration Guide* for details on configuring the CIM Agent. Briefly, this procedure entails:

1. Installing the software (ESS devices only).

The installation checks for the existence of the ESSCLI. If the ESSCLI is not installed, installation of the CIM Agent cannot proceed.

2. 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 best to use port 5988 and protocol HTTP.

3. Changing the default authentication method in order to discover the CIM Agent.

- a. Stop the IBM CIM Agent service, and then edit the `cimom.properties` file in `C:\Program Files\IBM\cimagent`.

- b. Open the `cimom.properties` file and change the following property to `false`:

```
DigestAuthentication=False
```

4. 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. The credentials are not necessarily the same as those used to login to the ESS Specialist management utility or the DS Storage Manager.

5. 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 with the necessary privileges to access and configure the ESS or DS storage systems.

- a. Navigate to `\Program Files\ibm\cimagent\setdevice`.

- b. Do one of the following:

For ESS devices, enter `cmd address <ipaddress> <username> <password>` where `ipaddress` is the IP address of the management console server of the ESS device and `username` and `password` are the management console credentials.

For DS devices enter `cmd addressserver <ipaddress> <username> <password>` where `ipaddress` is the IP address of the management console server of the DS device and `username` and `password` are the management console credentials.

6. Restarting the IBM CIM Agent service.
7. Verifying that the CIM Agent is able to communicate with the storage devices.  
Enter the following command to verify communication:  
`verifyconfig -u username -p password` where `username` and `password` are the credentials to access CIM Agent and were created by `setuser`.

---

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

---

**Note** – The IBM CIMOM user name and password are defined with the `setuser` command.

---

8. Re-enter the password in the Verify Password box.
9. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.

12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

## 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. The packages can be found on the Sun Enterprise Storage Manager Accessory CD-ROM. Refer to the readme file on the Sun StorEdge™ ESM Accessory CD-ROM for information about configuring these three 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

---

**Caution** – 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. Select **Discovery > Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter 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. Re-enter the password in the Verify Password box.
9. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

## Discovering Sun StorEdge 6920 and 6940 Storage Systems

To discover Sun StorEdge 6920 and 6940 storage systems:

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



5. In the IP Address/DNS Name box, enter 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. Re-enter the password in the Verify Password box.
9. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

## Discovering Sun StorEdge 6130 Storage Systems

To discover Sun StorEdge 6130 storage systems:

1. Select **Discovery** > **Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the IP address or DNS name of the controller or proxy you want to discover.
6. Leave the User Name box blank.
7. If you do not want to do provisioning on the storage systems, leave the password box blank. If you want to do provisioning, enter the corresponding password for controller or proxy.
8. If you entered a password in the previous step, re-enter the password in the Verify Password box.
9. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in 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. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

## Discovering Xiotech Storage Systems

---

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

---

To discover a Xiotech storage system:

1. Select **Discovery > Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, 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. A user name and password are required. Enter anything for the user name and password.
7. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery > Setup**).
8. Select the **Do Not Authenticate** option.
9. Click **OK**.
10. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

# 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 and then modify one of its properties files. See the *Installation Guide* for information on how to install the CIM extension.

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 APPQTime/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. See the *Installation Guide* for information about starting CIM extensions.

To discover an HP NAS device on Windows:

1. Select **Discovery > Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter 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. Re-enter the password in the Verify Password box.

9. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

## Discovering HP NAS Devices on Linux

In order to discover an HP NAS device on Linux, you must first install a CIM extension on the device and then modify one of its properties files. See the *Installation Guide* for information about starting CIM extensions.

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 `/opt/APPQCime/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 then close the file.
8. Restart the CIM extension. See the *Installation Guide* for information about starting CIM extensions.

To discover an HP NAS device on Linux:

1. Select **Discovery** > **Setup**.
2. Select **Step 1** at the top of the page.

3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter 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. Re-enter the password in the Verify Password box.
9. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

## 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 through SSL you must set the `cimom.providers.netapp.useSSL` property to `true`. This is a global setting and will cause all NetApp NAS devices to communicate using SSL. For more information, see “Enabling SSL Communication with a NetApp NAS Device” on page 85.
- 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.
- 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 you are restricting Administrative HTTP access, 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. For more information, see the NetApp NAS device documentation.

To discover a NetApp NAS device:

1. Select **Discovery > Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter 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. Re-enter the password in the Verify Password box.
9. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

## Enabling SSL Communication with a NetApp NAS Device

To enable SSL communication with a NetApp NAS device:

1. Select **Configuration > Product Health**, and then click **Advanced** in the **Disk Space** tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the following property:  
`#cimom.providers.netapp.useSSL=true`
4. Click **Close** to return to the Advanced page.
5. Paste the copied text into the Custom Properties box.

6. Uncomment the `cimom.providers.netapp.useSSL` property by removing the pound symbol (#) in front of `cimom.providers.netapp.useSSL`.
7. When you are done, click **Save**.

## 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. Select **Discovery > Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter 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. Re-enter the password in the Verify Password box.
9. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

# Discovering HP and IBM Tape Libraries

Before you can discover an HP or IBM tape library, you must download and install the corresponding SMI-S provider software. Refer to the support matrix and your tape library documentation for more information.

To discover an HP or IBM tape library:

1. Select **Discovery > Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the IP address or DNS name of the SMI-S provider for the tape library.
6. Enter the user name and password of the provider running the tape library. The user name and password are the provider's user name and password, not the credentials for the operating system's user name. The default user name/password for IBM is cimuser/cimpass and for HP it's administrator/administrator unless you've made changes.
7. Enter the **Password** of the system running the tape library.
8. Re-enter the password in the Verify Password box.
9. (Optional) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery > Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

---

## Building the Topology

This section contains the following topics:

- "Building the Topology View" on page 88
- "Modifying the Properties of a Discovered Address" on page 89
- "Deleting Elements from the Product" on page 96



# 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, if the number two is shown between a switch and a storage system, it means that 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 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 (select **Discovery > Topology**). The management server obtains enough information about where the element is connected in the topology; for example, showing where a switch connected to a host.

If the management server detects an element but it cannot obtain additional information about it, it marks the element with a question mark in the topology. To learn more about fixing detected and/or disconnected elements, see “Troubleshooting Topology Issues” on page 824.

---

**Caution** – 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. See “Recalculating the Topology” on page 838 for more information.

---

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

1. Click the **Discovery** menu in the upper-right corner of the home page.
2. Click **Topology** in the upper-right corner.  
The discovered elements are selected.
3. Select **All Discovery Groups** or click **Specified Discovery Groups** to specify a customized list. If you are obtaining the topology for the first time, select **All Discovery Groups**.

---

**Note** – For information on selecting a custom discovery list, see “Creating Custom Discovery Lists” on page 93.

---

4. Click **Get Topology**.

The management server obtains the topology for selected elements and 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.

---

**Note** – You can also access System Explorer by clicking **System Explorer** in the left pane.

---

5. Review the topology for errors and/or changes.
  - 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. For more information, see “Viewing Log Messages” on page 101 and “Troubleshooting Topology Issues” on page 824.
  - If the topology for an element in your network changes, select the element and click **Get Topology (Discovery > Topology)** to updated the information.

## Modifying the Properties of a Discovered Address


You can modify the user name and password the management server uses to access a device. However, 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, if the password for a host was changed, you would need to update the management server database with the new password. For more information, see “Modifying a Single IP Address Entry for Discovery” on page 30.

---

**Note** – If you 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.

---

To change the discovery properties of an element:

1. Select **Discovery > Topology** or **Discovery > Details** in the upper-right pane.
2. Click the **Edit** () 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** menu.
4. Click **OK** in the Edit Discovered Element window.

---

# Get Details

This section contains the following topics:

- “About Get Details” on page 90
- “Running Get Details” on page 91
- “Stopping the Gathering of Details” on page 92

## About 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:

- Running Get Details takes 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 **Get Topology** on the Topology tab.
- Reports show data from the last successful Get Details and report cache update. When a scheduled Get Details finishes, the report cache refreshes automatically. If you run Get Details manually, the report cache updates every 6 hours. For information about refreshing the report cache, see “Refreshing a Report” on page 526.
- During Get Details the data you see in the user interface is not updated until the data collection is finished.
- 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.
- You can use discovery groups to break up Get Details. For example, instead of running Get Details for all elements, you could specify only the elements in Discovery Group 1. For more information, see “Using Discovery Groups” on page 92.
- When an element in a discovery group is updated, its dependent elements are also updated.
- You can quarantine elements to exclude them from Get Details. For example, if you want to get information about all the elements in a discovery group except for one, you can quarantine that element. For more information, see “Placing an Element in Quarantine” on page 98.
- If a problem occurs with a host or SMI-S element during Get Details, the host or element is automatically quarantined. To remove the element from quarantine, see “Removing an Element from Quarantine” on page 99.

- If you want to receive status reports about Get Details, see “Configuring E-mail Notification for Get Details” on page 814 for information about how to configure this option.
- If an element changes and you run Get Details while the provider cache is updating, an error might occur or the gathered details might be inconsistent with the actual element status.

## Running Get Details

To obtain details about the elements on the network:

1. Select **Discovery > Details**.
2. Select **Include infrastructure details**, which gathers the latest information about SAN details. 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. For information about discovering master backup servers, see “Step 1 — Discovering Your Hosts and Backup Manager Hosts” on page 105.
3. 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 such as HP, HDS, and EMC storage systems with the assumption that 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 62 and “Excluding HDS Storage Systems from Force Device Manager Refresh” on page 67.
4. Select **All Discovery Groups** or click **Specified Discovery Groups** to specify a customized list. If you are running Get Details for the first time, select **All Discovery Groups**.

---

**Note** – For information on selecting a custom discovery list, see “Creating Custom Discovery Lists” on page 93.

---

5. Click **Get Details**.

During Get Details, the software changes its status light from green to red and the log opens and shows the progress of Get Details.

When the software finishes getting all element details, it displays `GETTING ALL DETAILS COMPLETED` on the View Logs page and the status light turns green.

6. See the “Adding a Discovery Schedule” on page 226 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.

---

**Caution** – If you stop the gathering of details, you should reschedule it. This type of collection obtains detailed information about elements in the network.

---

To stop the gathering of details:

1. Select **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.

---

**Note** – Existing operations will finish before the management server stops gathering details.

---

4. Schedule a time to resume getting details.

---

## Using Discovery Groups

The discovery groups feature is sometimes called *segmented replication* because it allows you to run Get Details/ for a segment of elements. Because The product runs more slowly when Get Details is in progress, it is helpful to break the process into segments which can then be run at night or on multiple days. For example, if Get Details for all elements takes twelve hours, you could break the elements into several small groups and schedule Get Details to run at night on multiple days.

---

**Note** – For more about data collection, see “About Get Details” on page 90.

---

When planning discovery groups, consider the following requirements and capabilities:

- By default, the product is configured with a default discovery group plus four additional groups.
- Discovery groups affect the amount of memory needed for the product. Before configuring discovery groups, check the support matrix and verify that your system meets the memory requirements for using discovery groups.
- Do not move elements between discovery groups when Get Details is running. If you do this, an error will occur when Get Details tries to discover elements that were moved.
- An element can be a member of one discovery group at a time.
- Elements discovered through SMI-S and hosts discovered with CIM extensions from Build 5.1 and later of the product cannot be added to discovery groups. These elements can, however, be placed independently into scheduled Get Details tasks without being part of a discovery group. This allows you greater flexibility when gathering discovery data. For more information, see “Creating Custom Discovery Lists” on page 93.
- When an element in a discovery group is updated, its dependent elements are also updated.
- Each discovery group communicates over a specific port. The defaults are:

**TABLE 2-4** Discovery Group Ports

Default	5986
Discovery Group 1	5984
Discovery Group 2	5982
Discovery Group 3	5980
Discovery Group 4	5978

## Creating Custom Discovery Lists

You can create a discovery list for Get Details or Get Topology, which will allow you to select a set of discovery groups to use the next time Get Details runs.

1. Select **Discovery > Details or Discovery > Topology**.
2. Click the **Specified Discovery Groups** link.
3. Select the check box next to each item you want to add to the discovery list.

Elements discovered through SMI-S and hosts discovered with CIM extensions from Build 5.1 and later of the product appear in the list individually. You can add individual elements, discovery groups, or both to the same discovery list.

4. Click **Add Selected Discovery Groups to Discovery List** to move them into the Discovery List.

---

**Caution** – Do not run Get Details for all discovery groups simultaneously.

---

5. Click **OK** to save and return to the previous window. The elements are selected in the elements table.
6. Click **Get Details** or **Get Topology**.


## Managing Discovery Groups

You can manage discovery groups from the Discovery Setup page.

---

**Note** – The Default discovery group cannot be edited.

---

1. Select **Discovery > Details** or **Discovery > Topology**.
2. Click **Manage Discovery Groups**.  
The Discovery Groups page shows a list of your discovery groups, including the name, Port Number, and included elements.
3. Click **Edit** .
4. To rename the group, enter a new name in the Name box.
5. To add a member, select the member from the Potential Members section, and then click the **Add Selected Discovery Groups to Discovery Group** button to move it into the Current Members section.
6. To remove a member, select the member from the Current Members section, and then click the **Remove Selected Discovery Groups from Discovery Group** button to move it into the Potential Members section.

---

**Note** – The path to the log file for the discovery group is listed at the top of the page.

---

7. Click **OK**.
8. Click **Back to Discovery Page**.

## Moving Elements Between Discovery Groups

All elements are initially placed in the Default discovery group. You can move elements between discovery groups.

---

**Caution** – Do not move elements between discovery groups when Get Details is running. If you do this, an error will occur when Get Details tries to discover elements that were moved.

---

### *Method 1: Select Discovery Group*

To select a new discovery group for an element:

1. Select Discovery Setup (**Discovery > Details**).

The Get Details page appears.

2. Select the check box for the element you want to move.

3. Click **Move to Discovery Group**.

The Select Discovery Group window appears.

4. Select the new discovery group for the selected element.

5. Click **OK**.

The management server notifies you that it can take a few minutes to move an element.


6. Click **OK**.

The elements are moved to the new discovery group.

### *Method 2: Edit a Discovered Element*

To edit a discovered element:

1. Select Discovery Setup (**Discovery > Details**). The Get Details page appears.

2. Click the **Edit** () button next to the element you want to modify.

3. Select a new discovery group in the **Discovery Group** menu.

4. Click **OK**.

The management server notifies you that it can take a few minutes to move an element.



5. Click **OK**.

The elements are moved to the new discovery group.

---

## 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 menu. Right-click an element and select **Delete Element** from the 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.







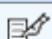

- **In Chargeback** - Click the **Delete** () button for the element you want to delete.
2. If the element has multiple access points, you are asked which want to delete. Do one of the following:
    - **Delete the element and its access points.** This option lists not only the switch you want to delete, 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. Select **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.

92.168.10.25	Switch	<a href="#">QBrocade2</a> , <a href="#">QBrocade5</a>	admin		
92.168.10.21	Switch	<a href="#">QBrocade1</a>	admin		
92.168.10.22	Switch	<a href="#">QBrocade2</a> , <a href="#">QBrocade5</a>	admin		
92.168.10.24	Switch	<a href="#">QBrocade3</a> , <a href="#">QBrocade4</a>	admin		

**FIGURE 2-3** Deleting Elements from the Management Server

3. Select all of the access points for the element you want to delete, and 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. Assume also that 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.

## Working with Quarantined Elements

When an element is quarantined, it is not included in the Get Details process until it is removed from quarantine. For more information, see “Removing an Element from Quarantine” on page 99. If a problem occurs with a host or SMI-S element during Get Details, the host or element is automatically quarantined.

## 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. Assume you want to discover all the elements in a discovery group, except for one, which is being taken off of the network for maintenance. You can use the quarantine feature to exclude this element 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 **Set Quarantine**.
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 **Clear Quarantine**.
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, and 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 use a different procedure. For more information, see “Managing McDATA and EMC Connectrix Switches” on page 54.
- Running Get Details takes 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. Select **Discovery > Details**.
2. Select **Include infrastructure details**, which gathers information about SAN details.

---

**Note – Include backup details** is used for gathering information for Backup Manager. You do not need to select it unless you have already discovered hosts running backup applications and installed CIM extensions on those hosts. For more information about discovering master backup servers, see “Step 1 — Discovering Your Hosts and Backup Manager Hosts” on page 105.

---

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 based on the assumption the information in the external database is up-to-date.

For more information, see the following topics: “Excluding EMC Symmetrix Storage Systems from Force Device Manager Refresh” on page 62 and “Excluding EMC Symmetrix Storage Systems from Force Device Manager Refresh” on page 62.

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 101 for more information about the messages viewed in this tab.
6. Verify the topology is displayed correctly by accessing System Explorer. Access System Explorer by 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, and 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, in the command, you must provide a user name to 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 this is a new installation of the management server and you have Brocade switches, download and install the Brocade SMI Agent software as described in the Brocade SMI Agent documentation.
- Additional steps are required for discovering McDATA and EMC Connectrix switches; the steps vary according to your network configuration. For more information, see “Discovering McDATA and EMC Connectrix Switches” on page 44.
- EMC CLARiiON storage systems require additional steps for discovery. For more information, see “Discovering EMC CLARiiON Storage Systems” on page 63 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.

---

# Viewing Log Messages

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

- Discovery
- Building the Topology
- Backup details

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

To view logs for these operations:

1. Select **Discovery > View Logs**.
2. To view the progress of Get Details, click the **Infrastructure** tab.
3. To view the progress of Backup Details, click the **Backup** tab.
4. To obtain the latest status, click **Get Latest Messages**.

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 that 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 Windows Management Instrumentation (WMI).

---

**Note** – The logs show data from the most recent discovery, test, or data collection task.

---

---

**Caution** – Look at Event Manager for additional information. See “About Event Manager” on page 605 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. Select **Discovery > System Tasks**.
2. To obtain the latest status, click **Get the Latest Status**.

The following task statuses are provided by the Task Dashboard:

**TABLE 2-5** 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.



## Discovering Applications, Backup Hosts and Hosts

---

This chapter describes the following:

- “Step 1 — Discovering Your Hosts and Backup Manager Hosts” on page 105
- “Step 2 — Setting Up Discovery for Applications” on page 112
- “Step 3 — Discovering Applications” on page 151
- “Changing the Oracle TNS Listener Port” on page 154
- “Changing the Password for the Managed Database Account” on page 155

---

### Step 1 — Discovering Your Hosts and Backup Manager 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 a CIM extension on the host if you want to obtain detailed information about the host and its applications, including those applications for backup. See the support matrix for information about which backup applications the management server supports.

For information about discovering clustered hosts, see “Host and Application Clustering” on page 165.

The management server automatically detects file 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 110.

The management server also detects the backup applications its supports, such as Veritas™ NetBackup™ or HP Data Protector. If you are licensed for Protection Explorer and you want to manage and monitor your backup applications, select **Include backup details** when you run Get Details, as described in “Step D — Get Details” on page 111.

Keep in mind the following:

- Elements discovered through SMI-S and hosts discovered with CIM extensions from Build 5.1 and later of the product cannot be added to discovery groups. These elements are listed separately and can be placed independently into scheduled Get Details tasks without being part of a discovery group. This allows you greater flexibility when gathering discovery data. (For more information, see “Creating Custom Discovery Lists” on page 94).

If you are upgrading from a previous build of the product, and you rediscover your hosts, they will be moved out of their existing discovery groups. Each rediscovered host would be placed in its own discovery group. If the original discovery groups containing these hosts were included in scheduled Get Details tasks, the schedules would be modified to contain the new discovery groups for rediscovered hosts.

- After installing the CIM extension on a DataProtector system on Windows, check the Logon account for the DataProtector CRS service and verify that it matches the AppStorWin32Agent service. To determine the Logon account for the DataProtector CRS service, go to **Control Panel > Administrative Tools > Services**, select the DataProtector CRS service, access its Properties page, and select the **Logon** tab. To determine the Logon account for the AppStorWin32Agent service, go to **Control Panel > Administrative Tools > Services**, select the AppStorWin32Agent service, access its Properties page, and select the **Logon** tab.
- If you change the password of a host after you discover it, stop and restart the CIM extension running on the host, and change the host password in the discovery list.
- 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.
- If you want to receive status reports about Get Details, see “Configuring E-mail Notification for Get Details” on page 814 for information about how to configure this option.
- Depending on your license, you may not be able to access Protection Explorer, File Server SRM and/or monitor certain applications may not be available. See the List of Features to determine if you have access to Protection Explorer, 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, see 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 836.
- Get Details can hang if obtaining information from an AIX host where SAN storage was previously available is no longer visible to the operating system. You may need to reboot the management server to resolve this issue.
- When discovering a Linux host from the management server, the operating system/server type is not available.
- If you started a CIM extension on a Sun Solaris host by using the `cim.extension.parameters` config file or 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.
- You can configure the management server to obtain information about your backup manager hosts at a set interval. See the topic, “Scheduling Backup Collection for Backup Managers” in the User Guide for more information about collectors.

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 107.
- **Topology** — Mapping the elements in the topology. See “Step B — Build the Topology” on page 110.
- “(Optional) Step C — View the Topology” on page 110
- **Details** — Obtaining detailed element information. See “Step D — Get Details” on page 111.

## Step A — Set Up Discovery for Hosts

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** box, enter the lowest IP address in the range of the elements you want to discover.
- d. In the **To IP Address** box, enter the highest IP address in the range of the elements you want to discover.
- e. In the **User Name (Optional)** box, enter 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)** box, enter the password corresponding to the user name entered in the **User Name** box.
- g. Enter the password from the previous step in the **Verify Password** box.
- h. In the **Comment** box, enter 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** box, enter the IP address or DNS name of the device you want to discover.
  - d. In the **User Name (Optional)** box, enter the user name.

This box 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)** box, enter the corresponding password for the user name entered in the previous step.

This box 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 entered a password in the previous step, entered the password in the **Verify Password** box.
  - g. In the **Comment** box, enter 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 DISCOVERY COMPLETED message is displayed in the Log Messages box.

## 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.

---

**Caution** – The management server’s 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 Log Message page is displayed by the management server. 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 824.

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 the **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 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 additional information about it. To learn more about fixing discovered and/or disconnected elements, see the topic, “Troubleshooting Topology Issues” on page 824.

## Step D — Get Details

After you obtain the topology of the network, you should obtain 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. Clusters won’t be recognized until Get Details is completed. Get Details must be run on all of the participating nodes of application clusters.

Keep in mind the following:

- Running Get Details takes 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 **Get Topology** on the Topology tab.

- Reports show data from the last successful Get Details and report cache update. When a scheduled Get Details finishes, the report cache refreshes automatically. If you run Get Details manually, the report cache updates every 6 hours. For information about refreshing the report cache, see “Refreshing a Report” on page 526.
- During Get Details the data you see in the user interface is not updated until the data collection is finished.
- 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.
- You can use discovery groups to break up Get Details. For example, instead of running Get Details for all elements, you could specify only the elements in Discovery Group 1. For more information, see “Using Discovery Groups” on page 93.
- When an element in a discovery group is updated, its dependent elements are also updated.
- If you want to monitor and manage backup servers, select **Include backup details**. If you also want to manage and monitor the host itself, select **Include infrastructure details**; otherwise, the host appears as a generic element in the topology in System Explorer.
- If Get Details includes an AIX host, three SCSI errors (2 FSCSI error and 1 FCS error) per IBM adapter port are displayed in the system log. You can ignore these errors.
- You can quarantine elements to exclude them from Get Details. For example, if you want to get information about all the elements in a discovery group except for one, you can quarantine that element. For more information, see “Placing an Element in Quarantine” on page 100.
- If a problem occurs with a host or SMI-S element during Get Details, the host or element is automatically quarantined. To remove the element from quarantine, see “Removing an Element from Quarantine” on page 100.
- If you want to receive status reports about Get Details, see “Configuring E-mail Notification for Get Details” on page 814 for information about how to configure this option.

To obtain details:

1. Click **Discovery > Details** in the upper-right corner.
2. Verify the **Include backup details** option is selected if you want to monitor and manage backup applications in Protection Explorer.
3. Verify the **Include infrastructure details** option is selected. This option is required to manage and monitor your elements not related to the backup infrastructure.
4. 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 Get Details is finished GETTING ALL DETAILS COMPLETED is displayed on the View Logs page.

---

## 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 and Backup Manager Hosts” on page 105.

You can configure the management server to monitor hosts and applications, such as Oracle, Microsoft Exchange server, Caché, 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 installation guide.

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 and Backup Manager Hosts” on page 105, 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, map the elements in the topology, and then run 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 114
- “Monitoring Microsoft SQL Server” on page 123
- “Monitoring Sybase Adaptive Server Enterprise” on page 134
- “Monitoring Microsoft Exchange” on page 137
- “Monitoring Caché” on page 140

# Creating Custom Passwords on Managed Database Instances

Depending on the password policy, SQL Server 2005 may require that passwords be alphanumeric. For this reason, a managed SQL Server 2005 database instance might not accept the default managed database password (password) during APPIQ\_USER creation. A script is provided to input an alphanumeric password for SQL Server 2005. For all other applications, this script is optional.

Because the management server uses a single password for managing all types of databases, the script for specifying a custom password is provided for each managed database type (SQLServer, Oracle, Sybase, and Caché). If the password is changed on any managed database instance, you should run the respective custom password scripts for each of the other managed database instances, and specify the same password.

The script names for each database type are as follows:

**TABLE 3-1** Script Names for Managed Databases

Database Type	With Default Password	With Custom Password
Oracle	CreateOracleAct.sh (or .bat) or CRACCT.COM (for OpenVMS)	CreateOracleActWithCustomPwd.sh (or .bat) or CUSTACCT.COM (for OpenVMS)
SQL Server	CreateSQLServerAct.bat	CreateSQLServerActCustomPwd.bat
Sybase	CreateSybaseAct.bat	CreateSybaseActCustomPwd.bat
Caché 5.0.20	createCacheDB50User.sh (or .bat)	createCacheDB50UserCustomPwd.sh (or .bat)
Caché 5.2 and 2007.1	createCacheDBUser.sh (or .bat) or CRUSER.COM (for OpenVMS)	createCacheDBUserCustomPwd.sh (or .bat) or CUSTUSER.COM (for OpenVMS)

After changing the password on all managed database instances, the password must be changed on the management server. To change the password on the management server:

1. Select **Discovery > Setup**.
2. Click the **Applications** tab.
3. Click **Change Password** in the Change Password for Managed Database Account section.
4. Enter the password that was used for creating APPIQ\_USER on the managed database instances.

# Monitoring Oracle

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

- “Step A — Create the APPIQ\_USER Account for Oracle” on page 115
- “Step B — Provide the TNS Listener Port” on page 118
- “Step C — Set up Discovery for Oracle 10g” on page 119

After you complete these steps, you must discover Oracle, and perform Get Details. See “Step 3 — Discovering Applications” on page 151.

Keep in mind the following:

- 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.
- By default discovery of Oracle is not supported through autoscan. To enable autoscan, add the line - "oracleautoscan=true" in the Custom Properties window from the Advanced page in **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree. Auto scans are only supported for Oracle 9i. To discover Oracle 10g instances, you must enter the application information described “Step C — Set up Discovery for Oracle 10g” on page 119.

## 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 platforms) or `CRACCT.COM` (on OpenVMS) 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.

---

**Note** – To create the APPIQ\_USER with a custom password, run `CreateOracleActWithCustomPwd.bat` (on Microsoft Windows) or `CreateOracleActWithCustomPwd.sh` (on UNIX platforms) or `CUSTACCT.COM` (on OpenVMS). For more information, see “Creating Custom Passwords on Managed Database Instances” on page 113.

---

Keep in mind the following:

- The `CreateOracleAct.bat` script must run under SYS user.
- Create the APPIQ\_USER account on the 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. See 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.

- When creating the APPIQ\_USER account on an Oracle Real Application Cluster (RAC) Database, this script should be run only once, on any one of the instances of the Oracle RAC Database. Since all the instances of an Oracle RAC access the same Database, it is sufficient to create the APPIQ\_USER account on any one of the instances.
- 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 the 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/DBIQ/oracle/unix
```

where `/cdrom` is the name of the directory where you mounted the CD-ROM.

- **To run the script on Microsoft Windows**, go to the `DBIQ\oracle\win` directory on the CIM extensions CD-ROM.
- **To run the script on OpenVMS:**

Log into an account that has administrative privileges.

Mount the CIM Extensions CD-ROM (if not auto-mounted) using the following command.

```
$ MOUNT /MEDIA=CDROM  
/UNDEFINED_FAT=STREAM:32767/OVERRIDE=IDENTIFICATION  
DQB0
```

where `DQB0` is the CD-ROM drive.

Go to the directory containing the Oracle agent creation script using the following command.

```
$ SET DEF DQB0:[OVMS.DBIQ.ORACLE]
```

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 UNIX platforms) or `CRACCT.COM` (on OpenVMS) on the computer with the Oracle database. On OpenVMS, run

CRACCT.COM on the host using the following command.

```
$ @CRACCT.COM
```

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 the 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.
- Grants create session and select on dictionary tables privileges to APPIQ\_USER, enabling the 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 UNIX platforms) or `RMACCT.COM` (on OpenVMS).

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.

- When removing the APPIQ\_USER account from an Oracle RAC Database, this script should be run only once, on any one of the instances of the Oracle RAC Database. Since all the instances of an Oracle RAC access the same Database, it is sufficient to remove the APPIQ\_USER account from any one of the instances.

To remove the APPIQ\_USER account for that Oracle instance:

1. If you plan to remove the management software for Oracle from a UNIX platform, 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/DBIQ/oracle/unix
```

where /cdrom is the name of the directory where you mounted the CD-ROM.
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. If you plan to remove the management software for Oracle from a computer running OpenVMS do the following:
  - a. Mount the CIM Extensions CD-ROM (if not auto-mounted) using the following command:

```
$ MOUNT /MEDIA=CDROM  
UNDEFINED_FAT=STREAM:32767/OVERRIDE=IDENTIFICATION  
DQB0
```

where DQB0 is the CD-ROM drive.
  - b. Go to the directory containing the Oracle agent creation script using the following command:

```
$ SET DEF DQB0:[OVMS.DBIQ.ORACLE]
```
4. 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.
5. Run UninstallOracleAct.bat (on Windows) or UninstallOracleAct.sh or RMACCT.COM ( on OpenVMS).
6. This script removes the management software for the specified Oracle instance.

---

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


---

7. 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.
8. 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. Enter the new port number and click **OK**.
4. If necessary, click the  button to remove the old port number.

---

**Caution** – Monitoring Oracle 10g or Oracle clusters requires an additional step. If you are not monitoring Oracle 10g or Oracle clusters, go to “Step 3 — Discovering Applications” on page 151.

---

## Step C — Set up Discovery for Oracle 10g

---

**Note** – If you are discovering an Oracle cluster, see “Discovering Oracle Real Application Clusters (RAC)” on page 120.

---

---

**Note** – By default discovery of Oracle is not supported through auto scan. To enable autoscan, add the line - "oracleautoscan=true" in the Custom Properties window from the Advanced page in **Configuration > Product Health**. Then, click **Advanced** in the **Disk Space** tree. Autoscan is only supported for Oracle 9i. To discover Oracle 10g instances, you must enter the application information described in the following procedure.

---

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** box, enter the IP address or DNS name of the host running Oracle.  
The **Management IP/DNS Name** box is optional.
4. In the **Server Name** box, enter the Oracle System Identifier (SID) of the Oracle database you want to monitor.
5. In the **Port Number** box, enter 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.

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 Real Application Clusters (RAC)

Since Oracle RAC is an active-active application cluster, one RAC instance can provide information for the whole RAC. Regardless of the instance through which the database is accessed, the same sets of tables are accessed. This includes the data dictionary tables that are used to understand the logical and physical storage organization of the Oracle RAC application.

### Discovery of Oracle RAC Instances Using One Instance



Because one RAC instance can provide information for the whole RAC, it is possible to identify and discover all the instances in the Oracle RAC cluster from any one of its instances. This means that you can enter the application setup information for one instance of the Oracle RAC, and the management server will automatically discover the other instances, subject to certain conditions. The conditions to be satisfied for discovering all the instances of Oracle RAC using application setup information from one of its instances are as follows:

- Only the Oracle RAC instances running on hosts already discovered and identified as part of the same cluster will be discovered as part of the Oracle RAC on the management server.
- The management server is able to contact the hosts running Oracle RAC instances using the short host name. The management server can be configured to access the hosts running Oracle RAC instances using the short name in the following ways:
  - On the management server, add entries for each host running an Oracle RAC instance in `/etc/hosts` (on UNIX platforms) or `%WINDIR%\system32\drivers\etc\hosts` (on Windows).
  - Add the domain of the host in the domain search list of the management server under the search option of `/etc/resolv.conf` (on UNIX platforms) or Append these DNS suffixes (in order) on the **Advanced TCP/IP Settings > DNS** tab (on Windows).
- The listener is configured on the same IP address that is used to discover the host. For example, on the Application Setup page, the management IP address for the application should be the same as the host IP address.
- Typically, all the instances of Oracle RAC will be listening on the same TNS port number. If this is not the case, the port numbers for the other instances should be specified in the default port list in the Application Setup page. For example, if SID1 is listening on TNS port LP1, and SID2 is listening on TNS port LP2, then it is possible to automatically discover SID2, provided that TNS port LP2 is part of the default port list in the Application Setup page.

To discover Oracle RAC:

1. Install the CIM extension on each node in the cluster.
2. If the cluster is not automatically discovered by the management server, create the cluster using Cluster Manager. For more information about Cluster Manager, see “Host and Application Clustering” on page 165.
3. Create the APPIQ\_USER account on any one node in the cluster. See “Step A — Create the APPIQ\_USER Account for Oracle” on page 115.
4. Click **Discovery > Setup**, and 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 an IP Range for Scanning” on page 27.
5. Discover the first Oracle node as follows:
  - a. Select **Discovery > Setup**, and click the **Applications** tab.

- b. Click the **Create** button for the Database Information table.
- c. In the **Host IP/DNS Name** box, enter the IP address or DNS name of the host running Oracle.

In the **Management IP/DNS Name** box, enter 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` (on Windows)

`$ORACLE_HOME/network/admin/listener.ora` (on UNIX platforms)

- d. In the **Server Name** box, enter the Oracle System Identifier (SID) of the Oracle database you want to monitor.
- e. In the **Port Number** box, enter 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`

#### Unix Platforms:

`$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 = EXTPROC))
    )
  )
)
```

- f. Select **ORACLE** from the Database Type menu.
- g. Click **OK**.

6. If the conditions described in the “Discovery of Oracle RAC Instances Using One Instance” section are satisfied, then all the other instances in the Oracle RAC will also be discovered, and the Oracle RAC application cluster will also be constructed by the management server.
7. If the other instances of the Oracle RAC are not discovered in the previous step, repeat steps 4 and 5 for each node in the cluster.


### About Discovery of an Oracle RAC Application Cluster on a Host Cluster Discovered Using Cluster Manager

When the underlying host cluster is not discovered, the management server will be “Oracle RAC safe,” but not fully “Oracle RAC aware.” Each instance will show up as a standalone Oracle application, and data will be collected for each instance separately (even though both instances will return identical capacity data). However, the management server does not explicitly identify and construct the Oracle RAC application cluster. Also, when the underlying host cluster is not discovered, other instances of the Oracle RAC cannot be discovered automatically as described in the Discovery of Oracle RAC Instances Using One Instance section.

However, if you create the host cluster at a later point in time, subsequent discovery of any instance in Oracle RAC will identify and construct the Oracle RAC application cluster. The management server will shift to “Oracle RAC aware” mode on top of the host cluster that you created.

## Deleting Oracle Application Information

If you do not want the management server to monitor an Oracle 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 Oracle Application 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 SQL Server

---

**Note** – If you are planning to monitor Microsoft SQL Server clusters, see “Monitoring Microsoft SQL Server Clusters” on page 131

---

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

- “Step A — Create the appiq\_user Account for the Microsoft SQL Server” on page 124
- “Step B — Provide the Microsoft SQL Server Name and Port Number” on page 128

---

**Caution** – 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 123.

---

## Switching to Mixed Mode Authentication

---

**Caution** – 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:

### Microsoft SQL Server 2000:

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 and select **Properties**.  
The SQL Server Properties (Configure) window appears.
4. Click the **Security** tab.
5. For “Authentication,” select **SQL Server and Windows**.
6. If the SQL instance is a clustered instance, make sure that the Startup Service Account is that of a Domain Administrator account. If the SQL instance is not clustered, make sure that the Startup Service Account is that of System Account.

### Microsoft SQL Server 2005:

1. Open SQL Server Management Studio (**Start Menu > Programs > Microsoft SQL Server 2005 > SQLServer Management Studio**).
2. Connect to the Microsoft SQL Server 2005 instance.
3. Right-click the server name and select **Properties**. The SQL Server Properties (Configure) window is displayed.

4. Select **Security**.
5. For “Server Authentication,” select **SQL Server and Windows Authentication Mode**, and then click **OK**. You may be prompted to restart the SQL server.
6. Open SQL Server Configuration Manager (**Start Menu > Programs > Microsoft SQL Server 2005 > SQLServer Configuration Manager**). Make sure that the SQL instance is logged on with a Domain Administrator account if it is a clustered instance and System Account if it is a non-clustered instance.

## Step A — Create the appiq\_user Account for the Microsoft SQL Server

### Microsoft SQL Server 2000:

The management server accesses Microsoft SQL Server through the appiq\_user account. This account is created when you run the `CreateSQLServerAct.bat` or `CreateSQLServerActCustom.bat` script on the computer running the Microsoft SQL Server database you want to monitor. This account has create session and select dictionary privileges, which allow the management server to view statistics for the Microsoft SQL Server.

---

**Note** – For more information about creating the appiq\_user account with a custom password, see “Creating Custom Passwords on Managed Database Instances” on page 113.

---

Keep in mind the following:

- The script must run under the SA user account. 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
- 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 Microsoft SQL Server:

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

---

**Caution** – 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. In a new command window, run the `CreateSQLServerAct.bat` script on the computer with the SQL Server database.

---

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

---

4. The script prompts you for the name of the Microsoft SQL Server on which to create the `appiq_user` account. If you are creating the account on a default instance, enter the host name if the instance is non-clustered and the `SQLNetwork Name` if the instance is clustered. If you are creating the account on a named instance, enter the host name and the instance name as follows:

**For a non-clustered instance:**

<Host Name>\<Instance Name>

**For a clustered instance:**

<SQL Network Name>\<Instance Name>

5. If you are running the `CreateSQLServerActCustom.bat` script, you will be prompted for a password for the `appiq_user` account. Provide a password that meets the password policy criteria described in “Creating Custom Passwords on Managed Database Instances” on page 113. If you are running the `CreateSQLServerAct.bat` script, the default password (`password`) is automatically used.
6. The script prompts you for the SA user password. Enter the password.  
The `appiq_user` account is created.
7. To determine if the `appiq_user` account was added correctly to your Microsoft 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. Click the refresh button in SQL Server Enterprise Manager. If the `appiq_user` is not listed, the management server is not able to discover the database.
8. 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.
9. Repeat these steps for each Microsoft SQL Server 2000 instance you want to manage.

### Microsoft SQL Server 2005:

The management server accesses Microsoft SQL Server through the appiq\_user account. To create this account, run the CreateSQLServerActCustomPwd.bat script on the computer running the Microsoft SQL Server database you want to monitor. This account has create session and select dictionary privileges, which allow the management server to view statistics for the Microsoft SQL Server.

Keep in mind the following:

- The script must run under the SA user account. 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
- 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 Microsoft SQL Server:

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

---

**Caution** – 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. In a new command window, run the CreateSQLServerActCustomPwd.bat script on the computer with the SQL Server database.

---

**Note** – You can use a remote SQL Server isql to run this script.

---

4. The script prompts you for the name of the SQL Server on which to create the appiq\_user account. If you are creating the account on a default instance, enter the host name if the instance is non-clustered and the SQLNetwork Name if the instance is clustered. If you are creating the account on a named instance, enter the host name and the instance name as follows:

**For a non-clustered instance:**

<Host Name>\<Instance Name>

**For a clustered instance:**

<SQL Network Name>\<Instance Name>

5. The script prompts you for the password for the appiq\_user account. Provide a password that meets the password policy criteria described in “Creating Custom Passwords on Managed Database Instances” on page 113.
6. The script prompts you for the SA user password. Enter the password.  
The appiq\_user account is created.
7. To determine if appiq\_user was added correctly to your SQL server:
  - a. Open SQL Server Management Studio.
  - b. Expand the user interface for SQL Server Management Studio, and 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. Click the Refresh button in SQL Server Management Studio. If the appiq\_user is not listed, the management server is not able to discover the database.
8. To determine if the SQL Server is ready to accept connections from the management server:
  - a. Connect to the SQL Server installation through SQL Server Management Studio using the appiq\_user account and the password specified earlier.
  - 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.
9. Repeat these steps for each Microsoft SQL Server 2005 instance 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:



---

**Caution** – 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** box, enter the IP address or DNS name of the host running Microsoft SQL Server. You must provide the host name. You cannot use localhost or parenthesis.
4. You can leave the **Management IP/DNS Name** box blank. This box is for Oracle clusters. When you leave the **Management IP/DNS Name** box 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 box, enter 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. It is one of the following:

- The name specified at the time the SQL server was installed
- The Windows system name (Windows 2000)
- 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/SQL Server Management Studio, 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** box, enter the port that SQL is using.

To determine the correct SQL Port Number that the SQL Server is using, follow these steps:

#### **Microsoft SQL Server 2000**

- 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 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** box on the management server.

#### **Microsoft SQL Server 2005**

- a. Open SQL Server Configuration Manager.
  - b. Select the specific SQL Server 2005 Network Configuration entry for the SQL Server 2005 instance.
  - c. Select the TCP/IP entry on the right pane, and then click the Properties right click menu.
  - d. From the IP Addresses tab, obtain the Port Number configured for the instance. Provide this port number in the Port Number box on the management server.
7. Select **SQLSERVER** from the Database Type menu.
  8. Click **OK**.

---

**Caution** – Perform Get Details for your inputs to take effect. See “Step 3 — Discovering Applications” on page 151.

---

## Removing the appiq\_user Account for Microsoft SQL Server

---

**Caution** – 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 from the Microsoft 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.

---


**Caution** – 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 DropSQLServerAct.bat script on Microsoft Windows on the computer with the SQL Server database.
4. Enter the name of the SQL Server server.
5. Enter 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 Microsoft SQL Server Information

If you do not want the management server to monitor a Microsoft 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 Microsoft SQL Server Clusters

---

**Caution** – Make sure the Microsoft SQL Server cluster database is in “Mixed Mode authentication.” To switch to mixed mode authentication, see “Switching to Mixed Mode Authentication” on page 123.

---

To monitor and manage Microsoft SQL Server clusters:

1. Install CIM Extensions on each of the participating nodes.
2. Create the appiq\_user account as described in “Step A — Create the appiq\_user Account for the Microsoft SQL Server” on page 124.

---

**Note** – This step needs to be run on any one of the participating host nodes of the Microsoft SQL Server cluster.

---

3. Enter the server name and port number as described in “Provide the Microsoft SQL Server Name and Port Number for a Cluster” on page 131.

### *Provide the Microsoft SQL Server Name and Port Number for a Cluster*

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

---

**Caution** – 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 Microsoft SQL Server cluster:

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** box, enter the IP address or DNS name of at least one of the participating host nodes running Microsoft SQL Server cluster. You must provide the host name. You cannot use localhost or parenthesis.
4. You can leave the Management IP/DNS Name box blank. When you leave the Management IP/DNS Name box 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 box, enter the SQL database server name you want to monitor.

The SQL Server name would be one of the following:

- The name specified at the time the SQL server was installed
- The Microsoft SQL Network Name (the default instance)

For example, if a Microsoft SQL Server cluster instance called SQLCLUSTER is running on a 2 node Windows 2003 cluster (individual host node IP address being 192.168.2.10 and 192.168.2.11) at the default SQL port (1433) and shows the name of Microsoft SQL Network Name within SQL Enterprise Manager / SQL Server Management Studio, the correct system application discovery settings on the management server would be either of the following:

- **Host IP/DNS Name:** 192.168.2.10
- **Database Server:** SQLCLUSTER
- **Port Number:** 1433
- **Database Type:** SQLSERVER

Or

- **Host IP/DNS Name:** 192.168.2.11
- **Database Server:** SQLCLUSTER
- **Port Number:** 1433
- **Database Type:** SQLSERVER

6. In the **Port Number** box, enter the port that SQL is using.

To determine the correct SQL Port Number that the SQL Server is using, follow these steps:

#### **Microsoft SQL Server 2000 Cluster**

- 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 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** box on the management server.

### Microsoft SQL Server 2005 Cluster

- a. Open SQL Server Configuration Manager.
  - b. Select the specific SQL Server 2005 Network Configuration entry for the SQL Server 2005 instance.
  - c. Select the TCP/IP entry on the right pane, and then click the Properties right click menu.
  - d. From the IP Addresses tab, obtain the Port Number configured for the instance. Provide this port number in the Port Number box on the management server. If Dynamic Ports are used, the Port Number is located under IPAll > TCP Dynamic Ports.
7. Select **SQLSERVER** from the Database Type menu.
  8. Click **OK**.

---

**Caution** – Perform Get Details for your inputs to take effect. See “Step 3 — Discovering Applications” on page 151.

---

## 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.

---

**Caution** – 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 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.

---

**Note** – To create the APPIQ\_USER with a custom password, run `CreateSybaseActCustomPwd.bat`. For more information, see “Creating Custom Passwords on Managed Database Instances” on page 113.

---

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.

---

**Caution** – 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 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. Enter 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

---

**Caution** – 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.



---

**Caution** – 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 the computer with the Sybase database.
4. Enter the name of the Sybase server.
5. Enter 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** box, enter the IP address or DNS name of the host running Sybase.
4. You can leave the **Management IP/DNS Name** box blank. This box is for Oracle clusters. When you leave the **Management IP/DNS Name** box 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** box, enter the Sybase database you want to monitor.
6. In the **Port Number** box, enter the port that Sybase is using.
7. Select **SYBASE** from the Database Type menu.
8. Click **OK**.


---

**Caution** – Perform Get Details for your inputs to take effect. See “Step 3 — Discovering Applications” on page 151.

---

## 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

---

**Note** – If you are planning to monitor Microsoft Exchange Clusters, see “Monitoring Microsoft Exchange Failover Clusters” on page 140.

---

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 (“Step 3 — Discovering Applications” on page 151).

## Adding Microsoft Exchange Domain Controller Access

Before adding a domain controller, note the following:

- The hosts should recognize the management server by name, because a reverse look-up is required by both operating system security and Microsoft Exchange. Make sure the domain controller, Exchange server host, and management server are accessible to one other using the host name and the fully-qualified domain name.

- The user name you provide must be the Common Name (CN) of the Active Directory User for accessing the Microsoft Exchange server. If you enter the Windows user name and it is different from the CN, the management server will not discover the Exchange instance.

To find the CN for a user on a domain controller server:

- a. Install the ADSIEdit MMC snap-in if it is not installed.
- b. Select **Start > Run** and enter `adsiedit.msc`.
- c. When the snap-in opens, expand the DOMAIN directory and navigate to the **CN=Users** folder to see the CN for each user in the Active Directory.

To provide information about your domain controllers:

1. Select **Discovery > Setup**, and then click the **Applications** tab.
2. In the Exchange Information section, click **Create**.
3. Click the **Add New Domain Controller** link.
  - a. In the Domain box, enter the domain name.
  - b. In the Domain Controller Name box, enter the fully qualified DNS name for the domain controller.
  - c. In the User Common Name box, enter the Common Name (CN) of the Active Directory User for accessing the Microsoft Exchange server.
  - d. In the Domain Password box, enter the corresponding password for accessing the Microsoft Exchange server.
  - e. In the Verify Password box, re-enter the password for verification.
4. Click **Add**.

The domain controller is added to the table.
5. Click **OK**.
6. Repeat these steps for each domain controller.
7. When all of your domain controllers are added, run `wmiadap /f` on the Exchange Server to refresh the Exchange data.

---

**Caution** – You must discover the host running Microsoft Exchange. See “Step 3 — Discovering Applications” on page 151.

---

## Editing a Microsoft Exchange Domain Controller

To provide information about your domain controllers:


1. Select **Discovery > Setup**, and then click the **Applications** tab.
2. Click the **Edit** button next to the Exchange domain controller you want to edit.
3. Enter a new User Common Name or Domain Password.
4. Click **Edit**.

The domain controller updates are added to the table.



Click **OK**.

## Deleting a Microsoft Exchange Domain Controller

To delete all of the domain controllers of a particular domain:

1. Select **Discovery > Setup**, then click the **Applications** tab.
2. Click the **Delete** () button corresponding to the domain you want to remove.
3. Run Get Details for your changes to take effect.

To delete a particular domain controller in a domain:

1. Select **Discovery > Setup**, then click the **Applications** tab.
2. Identify the domain for the domain controller you want to remove, and click the **Edit** () button corresponding to that domain.
3. In the Edit window, click the **Delete** () button corresponding to the domain controller you want to remove.
4. Run Get Details for your changes to take effect.

## Monitoring Microsoft Exchange Failover Clusters

To monitor and manage Microsoft Exchange Failover Clusters:

1. Install CIM Extensions on each of the participating nodes of Microsoft Exchange Failover Cluster.
2. Add information for Microsoft Exchange Domain Controller Access. See “Adding Microsoft Exchange Domain Controller Access” on page 138.
3. Perform Get Details on each of the participating nodes of the Exchange Cluster.

## Monitoring Caché

To monitor Caché, you must do the following:

- “Step A — Import the Wrapper Class Definitions into the Caché Instance” on page 140
- “Step B — Create APPIQ\_USER Account on the Caché Instance” on page 145
- “Step C — Provide the Caché Instance Name and Port Number” on page 150

After you complete these steps, you must discover Caché. See “Step 3 — Discovering Applications” on page 151.

---

**Note** – The required drivers for Caché were automatically installed along with the management server.

---

---

**Caution** – Before you begin these steps, make sure you purchased Caché IQ, which is the module that lets you monitor Caché. Contact your customer support if you are unsure if you purchased this module.


---

### Step A — Import the Wrapper Class Definitions into the Caché Instance

To import the wrapper classes:

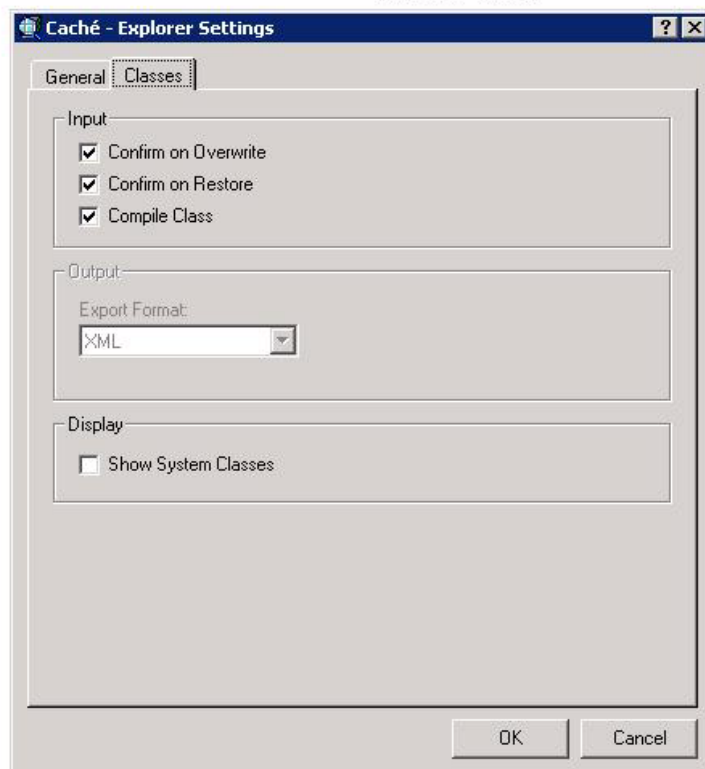
#### **For Caché 5.0 (5.0.20 onwards)**

1. Launch the Caché Explorer by right-clicking the Caché Cube icon in the system tray area of the Windows toolbar and selecting **Explorer**.
2. Right click the **Classes** folder located at **Namespaces > “%SYS” > Classes**.
3. Select **Import from disk**.



**FIGURE 3-1** Selecting Import from Disk

4. Browse the CIM Extension CD, select the wrapper xml file, and click **Open**.
  - On IBM AIX, Linux, or HP-UX, log into an account that has administrative privileges, and mount the CIM Extensions CD-ROM (if not auto-mounted). The wrapper file is  
`/cdrom/DBIQ/cachedb/unix/cachedb50_sqlprojs.xml`  
 where `/cdrom` is the name of the directory where you mounted the CD-ROM
  - On Microsoft Windows, the wrapper file on the CIM Extensions CD-ROM is  
`\DBIQ\cachedb\win\cachedb50_sqlprojs.xml`.
  - When the Import Classes window is displayed, click **Options**.
  - Select the **Classes** tab, enable the **Compile Class** checkbox, and click **OK**.



**FIGURE 3-2** Enabling Compile Class

5. In the Import Classes pop-up window, select `appiq.cls`, and click **Import**.

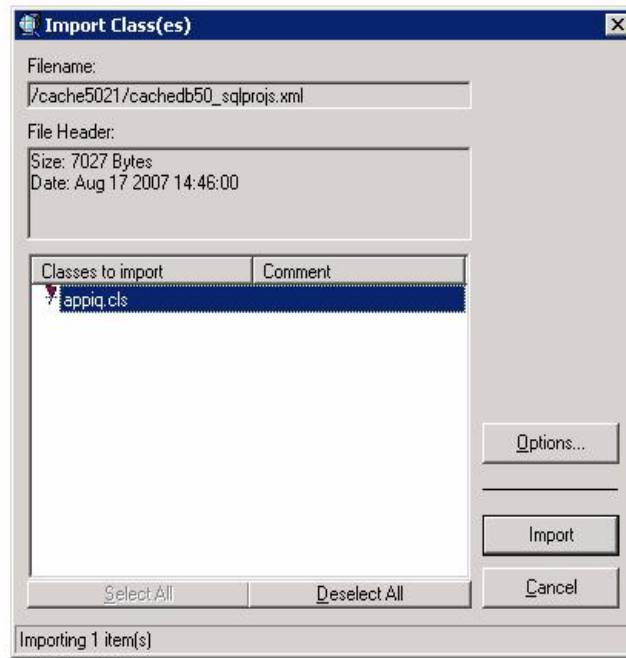


FIGURE 3-3 Selecting appiq.cls

### For Caché 5.2 and Caché 2007.1

1. Launch the Caché System Management Portal by right-clicking the Caché Cube icon in the system tray area of the Windows toolbar and selecting **System Management Portal**.
2. Click the **Classes** link under Data Management.
3. On the Classes page, select the **Namespaces** radio button, and then select **%SYS**.
4. Click **Import**.
5. Browse the CIM Extension CD, select the wrapper xml file, and click **Open**.
  - On IBM AIX, Linux, or HP-UX, log into an account that has administrative privileges, and mount the CIM Extensions CD-ROM (if not auto-mounted). The wrapper file is  
`/cdrom/DBIQ/cachedb/unix/cachedb_sqlprojs.xml`  
 where /cdrom is the name of the directory where you mounted the CD-ROM
  - On Microsoft Windows, the wrapper file on the CIM Extensions CD-ROM is  
`\DBIQ\cachedb\win\cachedb_sqlprojs.xml`.
  - On OpenVMS:



- a. Log in as system and mount the CIM Extensions CD-ROM.
- b. Copy the wrapper file (for example: `DQB0:[OVMS.DBIQ.CACHE]SQLPROJS.XML`), where `DQB0` is the CD-ROM drive, to any internal location on the OpenVMS host.

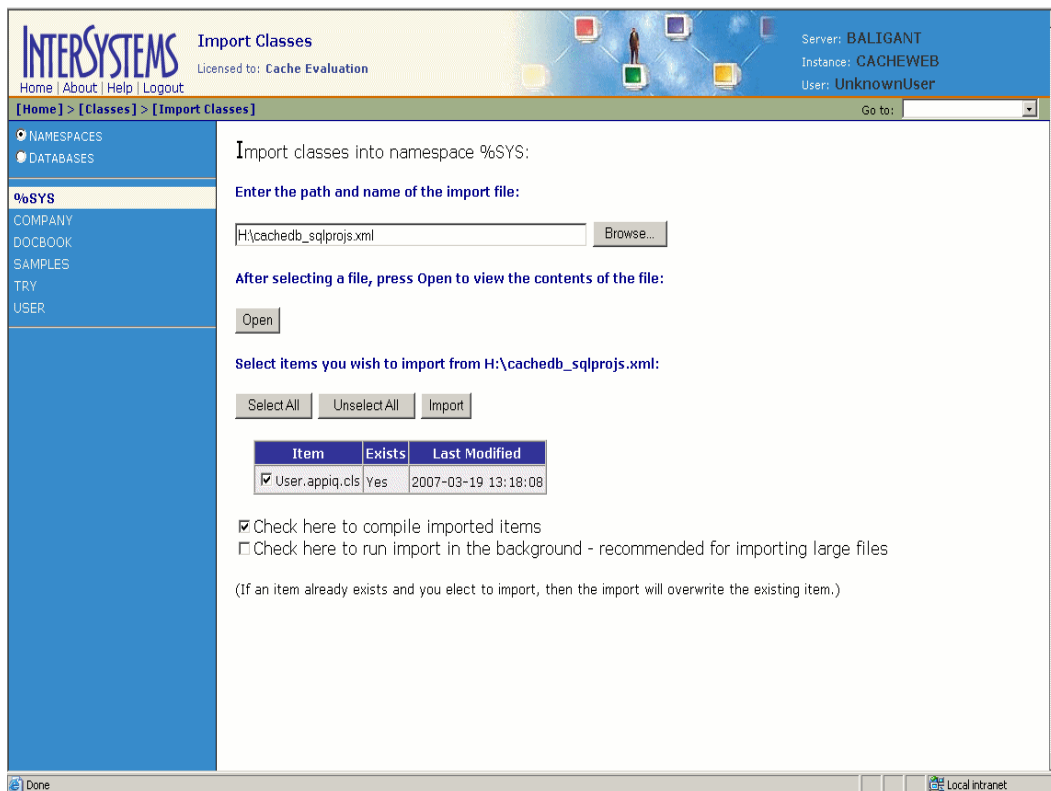
For example, copy `$DQB0:[OVMS.DBIQ.CACHE]SQLPROJS.XML`  
`$DKA0:[000000]SQLPROJS.XML`

where `DKA0` is a local drive on the OpenVMS host.

- c. Browse to `$DKA0` and specify `SQLPROJS.XML` within `$DKA0` as the import file.
6. After the file is opened, click **Select All**.
  7. Select **Check here to compile imported items**, and click **Import**.

The wrapper class definitions are imported into the Caché %SYS namespace.

The following image shows an example of importing the wrapper class definitions:



**FIGURE 3-4** Importing Wrapper Class Definitions

## Step B — Create APPIQ\_USER Account on the Caché Instance

The management server accesses Caché through the APPIQ\_USER account. This account is created when you run the appropriate script (described below) on the computer running the Caché database you want to monitor. You can execute these scripts from the management server also.

This script creates APPIQROLE with execute permissions for the SQL projections imported into the Caché managed instance, creates an APPIQ\_USER account, and assigns APPIQROLE to APPIQ\_USER.

The script must run as the \_SYSTEM user. You should enter the Caché server name, the Super Server port number, and the password of the \_SYSTEM user account as arguments for the script.

---

**Note** – If you are running Caché 5.2 or later, and the Caché instance was installed using “Locked Down” security mode, see “Locked Down Security Mode” on page 147 before creating the APPIQ\_USER account.

---

To create APPIQ\_USER for the Caché instance:

1. Do one of the following:

**To create APPIQ\_USER on the host:**

- To run the script on IBM AIX, HP\_UX, or Linux, log into an account that has administrative privileges, mount the CD-ROM (if not auto-mounted), and go to the  
/DBIQ/cachedb/unix directory by entering the following:

```
# cd /cdrom/DBIQ/cachedb/unix
```

where /cdrom is the name of the directory where you mounted the CD-ROM .

- To run the script on Microsoft Windows, go to the DBIQ\cachedb\win directory on the CD-ROM.
- To run the script on OpenVMS, log in as system, mount the CD-ROM drive, and go to the [OVMS.DBIQ.CACHE] directory by entering the following:  
SET DEF DQB0:[OVMS.DBIQ.CACHE]

Where DQB0 is the name of the CD-ROM drive.

**To remotely create APPIQ\_USER on the Caché instance from the management server:**

- To run the script on Linux or Solaris, go to the /opt/<product name>/install/cachedb/unix directory by entering the following:  
# cd opt/<product name>/install/cachedb/unix
- To run the script on Windows, go to the %MGR\_DIST%\install\cachedb\win directory

2. Verify you have the password to the \_SYSTEM user account.
3. For Caché 5.0: run createCacheDB50User.bat (on Windows) or createCacheDB50User.sh (on UNIX platforms) on the computer with the CacheDatabase. To specify a custom password for the APPIQ\_USER account, run createCacheDB50UserCustomPwd.bat (on Windows) or createCacheDB50UserCustomPwd.sh (on UNIX platforms) on the computer with the CacheDatabase.

For later versions of Caché: run createCacheDBUser.bat (on Windows) or createCacheDBUser.sh (on UNIX platforms) or CRUSER.COM (on OpenVMS) on the computer with the CacheDatabase. To specify a custom password for the

APPIQ\_USER account, run createCachedBUserCustomPwd.bat (on Windows) or createCachedBUserCustomPwd.sh (on UNIX platforms) or CUSTUSER.COM (on OpenVMS) on the computer with the CacheDatabase.

4. Enter the Caché server name, the Super Server port number and the password of the \_SYSTEM user account as arguments for the script. If you are running the custom password creation script, enter the custom password as the fourth argument.

When invoking the scripts on OpenVMS, enclose the arguments in double quotes:

```
$ @CRUSER.COM "<host name>" "<super server port>" "<password for  
_SYSTEM user>"
```

5. Repeat the previous step for each Caché instance you want to manage.

### *Locked Down Security Mode*

For Caché 5.2 and later versions, if the Caché instance was installed using “Locked Down” security mode, the following steps must be carried out before creating the APPIQ\_USER account:

1. Launch the System Management Portal.
2. Click the **Security Management** link under System Administration.
3. On the Security Management page, click **Services**.
4. Click **%Service\_Bindings** on the Services page.
5. On the Edit definition for Service %Service\_Bindings page:
  - a. Under Allowed Incoming Connections, click **Add** and enter the IP address of the management server in the Explorer User Prompt window.
  - b. If the create APPIQ\_USER scripts are being executed from the host on which Caché instance is running, add the IP address of the host.
  - c. Click the **Service Enabled** checkbox on the Edit definition for Service %Service\_Bindings page.
  - d. Click **Save**.
6. Click the **Security Management** link under System Administration in the System Management portal.
7. On the Security Management page, click the **Users** link .
8. Click the **Edit** link for \_SYSTEM user.

9. On the Edit Definition for User \_SYSTEM page, click the **User Enabled** checkbox and enter a password for the \_SYSTEM user in the Password and Confirm Password boxes.
10. Click the **Save** button.

Once the APPIQ\_USER has been created, the \_SYSTEM user can be disabled from the System Management portal.

## Removing the APPIQ\_USER Account from the Caché Instance

If you no longer want the management server to monitor a Caché instance, you can remove the APPIQ\_USER account and APPIQROLE for that Caché instance by running `dropCacheDBUser.bat` (on Windows) or `dropCacheDBUser.sh` (on UNIX platforms) or `DROPUSER.COM` (on OpenVMS).

Before you remove the APPIQ\_USER account from the Caché instances on a host, make sure no processes are running APPIQ\_USER for that Caché instance. The management server uses APPIQ\_USER to obtain information about a Caché instance. One of the ways to make sure APPIQ\_USER is not being used is to temporarily remove the host running Caché (**Discovery > Topology**). After you remove the APPIQ\_USER account for Caché, discover and perform Get Details for the host if you want to continue monitoring it.

For Caché 5.2 and later versions, if the Caché instance was installed using “Locked Down” security mode, ensure that the \_SYSTEM user has been enabled before trying to remove the APPIQ\_USER account. To ensure that the \_SYSTEM user has been enabled:

1. Launch the System Management Portal
2. Click the **Security Management** link under System Administration.
3. On the Security Management page, click the **Users** link.
4. Click the **Edit** link for \_SYSTEM user.
5. On the Edit Definition for User \_SYSTEM page, click the **User Enabled** checkbox and enter a password for the \_SYSTEM user in the Password and Confirm Password fields.
6. Click **Save**.

Once the APPIQ\_USER has been removed, the \_SYSTEM user can be disabled from the System Management portal. The %Service\_Bindings service that was enabled before creating the APPIQ\_USER can also be disabled.

To remove the APPIQ\_USER account:

1. Do one of the following:

**To remove the APPIQ\_USER account from the host:**

- To run the script on IBM AIX, HP\_UX, or Linux, log into an account that has administrative privileges, mount the CD-ROM (if not auto-mounted), and go to the  
/DBIQ/cachedb/unix directory by entering the following:

```
# cd /cdrom/DBIQ/cachedb/unix
```

where /cdrom is the name of the directory where you mounted the CD-ROM

- To run the script on Microsoft Windows, go to the DBIQ\cachedb\win directory on the CD-ROM.
- To run the script on OpenVMS, log in as system, mount the CD-ROM drive, and go to the [OVMS.DBIQ.CACHE] directory by entering the following :  
SET DEF DQB0:[OVMS.DBIQ.CACHE]

Where DQB0 is the name of the CD-ROM drive.

**To remotely remove the APPIQ\_USER account from the Caché instance from the management server:**

- To run the script on or Solaris, go to the /opt/<product name>/install/cachedb/unix directory by entering the following:  
# cd opt/<product name>/install/cachedb/unix
- To run the script on Windows, go to the %MGR\_DIST%\install\cachedb\win directory

2. Verify you have the password to the \_SYSTEM user account.
3. For Caché 5.0, run dropCachedB50User.bat (on Windows) or dropCachedB50User.sh (on UNIX platforms) on the computer with the CacheDatabase. For later versions of Caché, run dropCachedBUser.bat (on Windows) or dropCachedBUser.sh (on UNIX platforms), or DROPUSER.COM (on OpenVMS) on the computer with the CacheDatabase.
4. Enter the Caché server name, the Super Server port number and the password of the \_SYSTEM user account as arguments for the script.

When invoking the scripts on OpenVMS, enclose the arguments in double quotes:

```
$ @DROPUSER.COM "<host name>" "<super server port>" "<password for  
_SYSTEM user>"
```

5. Repeat the previous step for each Caché instance you want to manage.

After deleting the APPIQ\_USER account from the Caché instance, you can also delete the wrapper class definitions.

### For Caché 5.0 (5.0.20 onwards)

1. Launch the Caché Explorer.
2. Click the Classes folder located at **Namespaces > “%SYS” > Classes**. Right-click the `User.appiq` class, and select **Delete**.
3. The Confirm Deletion window displays. Click **Yes**.

### For Caché 5.2 and Caché 2007.1

1. Launch the Caché System Management Portal.
2. Click the **Classes** link under Data Management.
3. On the Classes page, select the Namespaces radio button, and then click **%SYS**.
4. Click **Delete**.
5. Enter `User.appiq.cls` in the Enter search mask box, and click **Search**.
6. Select `User.appiq.cls`, and click **Delete**.

## Step C — Provide the Caché Instance Name and Port Number

To provide the Caché instance name and SuperServer port number for managing the Caché instance:

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 box, enter the IP address or DNS name of the host running Caché.
4. You can leave the Management IP/DNS Name box blank. This box is for clusters. When you leave the Management IP/DNS Name box 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 box, enter the Caché instance name you want to monitor.
6. In the Port Number box, enter the SuperServer port that Caché is using.
7. Select **Cache** from the Database Type menu.
8. Click **OK**.


---

**Caution** – Perform Get Details for your inputs to take effect. See “Step 3 — Discovering Applications” on page 151.

---

## Deleting Caché Information

If you do not want the management server to monitor a Caché 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 Caché instance you do not want the management server to monitor.
3. Perform Get Details to make the management server aware of your changes.

---

## 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 152)
- Obtain topology information about the application (“Step B — Obtain the Topology” on page 152)
- Perform Get Details (“Step C — Run Get Details” on page 153)

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 112.
- If you used a custom password for the APPIQ\_USER account, you must change the password on the management server before performing Get Details. See “Creating Custom Passwords on Managed Database Instances” on page 113.
- 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 **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.
- The Log Messages page is displayed. To view the status of discovery, click **Discovery > View Logs**.

The DISCOVERY COMPLETED message is displayed in the Log Messages box when Discovery is complete.

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 812.

## 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 90.

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 824.

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 — Run 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 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, three SCSI errors (2 FSCSI error and 1 FCS error) per IBM adapter port are displayed in the system log. You can ignore these errors.
- You can quarantine elements to exclude them from Get Details. See “Placing an Element in Quarantine” on page 100 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 is unable to obtain information from an element during Get Details as a result of a CIM extension failure, 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 100 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 90.

3. Click **Get Details**.

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**.

The DISCOVERY COMPLETED message is displayed in the Log Messages box when Discovery is complete.

---

**Caution** – 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. See “Adding a Discovery Schedule” in the User Guide for information about automating the gathering of Get Details. If you run into problems with discovery, see “Troubleshooting” on page 793.

## 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.

---

**Caution** – 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** button for the **Oracle Information** table.
  3. Enter 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 812 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** boxes.

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. Enter a new password in the **Password** box.

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. Re-enter the password in the **Verify Password** box.
6. Click **OK**.
7. Verify that the management server can access the database applications by clicking the **Test** button for each database application.
- 8.



## Installing and Discovering the Windows Proxy

---

This chapter describes the following:

- “Installing the Windows Proxy” on page 160
- “Discovering the Windows Proxy” on page 161
- “Configuring Windows Proxy Authentication” on page 162
- “Decreasing the Maximum Java Heap Size” on page 163
- “Removing the Windows Proxy” on page 164

The Windows Proxy is required when the management server is on Sun Solaris or Linux and you want to obtain information from Microsoft Windows hosts that do not have a CIM extension installed. First, install the Windows Proxy as described in “Installing the Windows Proxy” on page 160. Then, discover the Windows Proxy as described in “Discovering the Windows Proxy” on page 161.

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 163.
- When the Windows proxy is installed on a new server, the Windows hosts must be re-discovered.

---

# Installing the Windows Proxy

---

**Caution** – 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 Utilities CD-ROM, 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.

---

**Caution** – Keep in mind that the Windows Proxy automatically starts when the system is restarted. The management server 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 161.



---

# Discovering the Windows Proxy

---

**Caution** – 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:

---

**Caution** – 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, “Discovering Applications, Backup Hosts and Hosts” on page 105.
  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 **SUN 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.

# Host and Application Clustering

---

This chapter contains the following topics:

- “About Clustering” on page 165
- “Discovering Clusters” on page 165
- “Clustering in System Explorer” on page 169
- “Clustering in Topology” on page 171
- “Clustering in Capacity Manager” on page 172

---

## About Clustering

The management server provides full support for managing clusters. Cluster support includes the following features:

- Clusters are recognized as managed elements.
- System Explorer supports clusters in all areas.
- The element topology shows which shared resources an application instance uses.
- Cluster capacity utilization is accurately reported.
- Capacity utilization trending is provided for applications running on clusters.

The management server supports automatic discovery of several popular cluster servers, and allows management of other clusters through Cluster Manager.

---

## Discovering Clusters

The following cluster services support automatic discovery:

- Microsoft Cluster Services (MSCS) on Windows 2003
- Veritas Clusters on Solaris

Cluster services that don't support automatic discovery can be discovered manually by using Cluster Manager. See "Manual Discovery of Host Clusters" on page 167.

The following application clusters are supported:

- Oracle Real Application Clusters (RAC)
- Microsoft Exchange 2000/2003 FailOver Clusters and 2007 Single Copy Cluster (SCC)
- Microsoft SQL Server 2000 and 2005

For information about discovering application clusters, see "Discovering Applications, Backup Hosts and Hosts" on page 105.

Refer to the support matrix for a complete list of supported configurations. The support matrix is accessible from the Documentation Center (**Help > Documentation Center**).

## Automatic Discovery of Host Clusters

MSCS on Windows 2003 and Veritas Clusters on Solaris support automatic discovery. To discover hosts using either of these cluster services:

8. Discover your hosts and applications as described in "Discovering Applications, Backup Hosts and Hosts" on page 105. The clusters are automatically recognized by the management server.

---

**Note** – The following optional steps describe how to select a preferred host from which shared resource capacity data will be collected.

---

9. *Optional:* Access Cluster Manager by right-clicking a cluster in System Explorer and selecting Edit Cluster. The Cluster Manager Overview page is displayed. Click **Next**.
10. *Optional:* Cluster Manager Step 2 (Select Preferred Host for Cluster Shared Resources) is displayed. Select a preferred host for each of the cluster shared resources. Keeping the default selection of "None" will result in shared resource capacity data being collected from an available active host that shares the resource. Choosing a particular active host results in the specified host being used for data collection. If the specified host becomes unavailable, an available active host is used for data collection.

Specify the preferred host for individual cluster shared resources. If a resource is not shared by the preferred host selection, the preferred host menu for that shared resource will continue to display the previous selection.

When you have finished specifying preferred hosts, click **Finish**.

# Manual Discovery of Host Clusters

If you are using a cluster service that doesn't support automatic discovery, you must manually create your clusters. For the list of cluster services that support automatic discovery, see "Discovering Clusters" on page 165.

---

**Note** – In some environments, using Cluster Manager to manually create a cluster with NetApp hosts may result in unsuccessful or incomplete cluster creation.

---

To manually discover clusters:

1. Discover your hosts and applications as described in "Discovering Applications, Backup Hosts and Hosts" on page 105.
2. Access Cluster Manager by right-clicking a host in System Explorer and selecting **Build Cluster**. The Cluster Manager Overview page is displayed. Click **Next**.
3. Cluster Manager Step 2 (Specify Cluster Properties and Cluster Members) is displayed. Follow these steps to specify the cluster properties and cluster members:
  - a. In the Cluster Properties section, specify the cluster name, cluster server, and cluster virtual IP.
  - b. In the Available Hosts section, select the hosts to add to the Cluster Members table. If desired, use the filter to assist in the selection of hosts. For details about the filtering functionality, see "Filtering Hosts" on page 168.

You may also use the Select Related Hosts button to facilitate the selection of hosts. Select a host in the table, and click **Select Related Hosts** to automatically select any related hosts.
  - c. After you have selected the hosts that you would like to add to the cluster, click **Add Selected Hosts to Cluster**. The selected hosts are added to the Cluster Members table.
  - d. Click **Next**.
4. Cluster Manager Step 3 (Specify Cluster Shared Resources) is displayed. Select **Automatic** or **Manual**. If you select Automatic, click **Display Cluster Shared Resources**, and the table at the bottom of the page is automatically populated.

If you select Manual discovery, follow these steps:

- a. Enter a name in the Cluster Shared Resource Name box.
- b. Select a resource type from the Resource Type menu. The menu includes the following resource types:

**Logical Disk**  
**Disk Partition**  
**Volume Manager Volume**  
**Disk Drive**

- c. Select the relevant resource for each cluster host, and click **Save Selections as Cluster Shared Resource**. The selections are added to the Cluster Shared Resources table.
  - d. Repeat steps a through c for each shared resource in the cluster.
  - e. Click **Next**.
5. Cluster Manager Step 4 (Select Preferred Hosts for Cluster Shared Resources) is displayed. Select a preferred host for each of the cluster shared resources. Shared resource capacity data will be collected from the specified node. Selecting "None" will result in no information being collected about the cluster shared resource.

Specify the preferred host for individual cluster shared resources. If a resource is not shared by the preferred host selection, the preferred host menu for that shared resource will continue to display the previous selection.

When you have finished specifying preferred hosts, click **Finish**.

## Filtering Hosts

The Available Hosts table on Cluster Manager Step 2 (Specify Cluster Properties and Cluster Members) allows you to filter the list of hosts displayed. To filter the list of hosts:

1. Click the **+ Filter** link to display the filtering options.  
If the volume filter is already displayed, the **- Filter** link is shown instead, which will collapse the filtering options.
2. Enter all or part of a volume name in the Name Contains box.
3. Select an operating system from the Operating System menu.
4. Enter all or part of a vendor name in the Vendor Contains box.
5. Enter a number in the Processors ( $\geq$ ) box.  
Hosts with at least as many processors as specified will display in the table.
6. Enter a number in the HBAs ( $\geq$ ) box.  
Hosts with at least as many HBAs as specified will display in the table.



7. Enter a number in the Ports ( $\geq$ ) box.

Hosts with at least as many ports as specified will display in the table.

8. Click **Filter**.

The table is updated to display only the elements that meet the filter criteria.

---

**Note** – To reset the filter criteria, click **Reset**.

---

## File Servers and Clusters

If you have marked a host as a file server and you move it into or out of a cluster, you must remove the file server data from the host and then re-mark it as a file server. To remove the file server data from the host and re-mark it as a file server:

1. Select **Configuration > File SRM**.

2. Verify that the **File Servers** tab is displayed.

3. Select the file servers you want to remove, and then click **Delete**.

4. Click **Add File Server**.

5. Click the check boxes for the hosts that you would like to mark as file servers.

6. Click **OK**.

The hosts are marked as file servers, and you are returned to the **File Servers** tab.

7. After removing the file server data from the host and then re-marking it as a file server, you must rescan the cluster member nodes and the cluster nodes. If a rescan is not completed, incorrect data may be displayed.

---

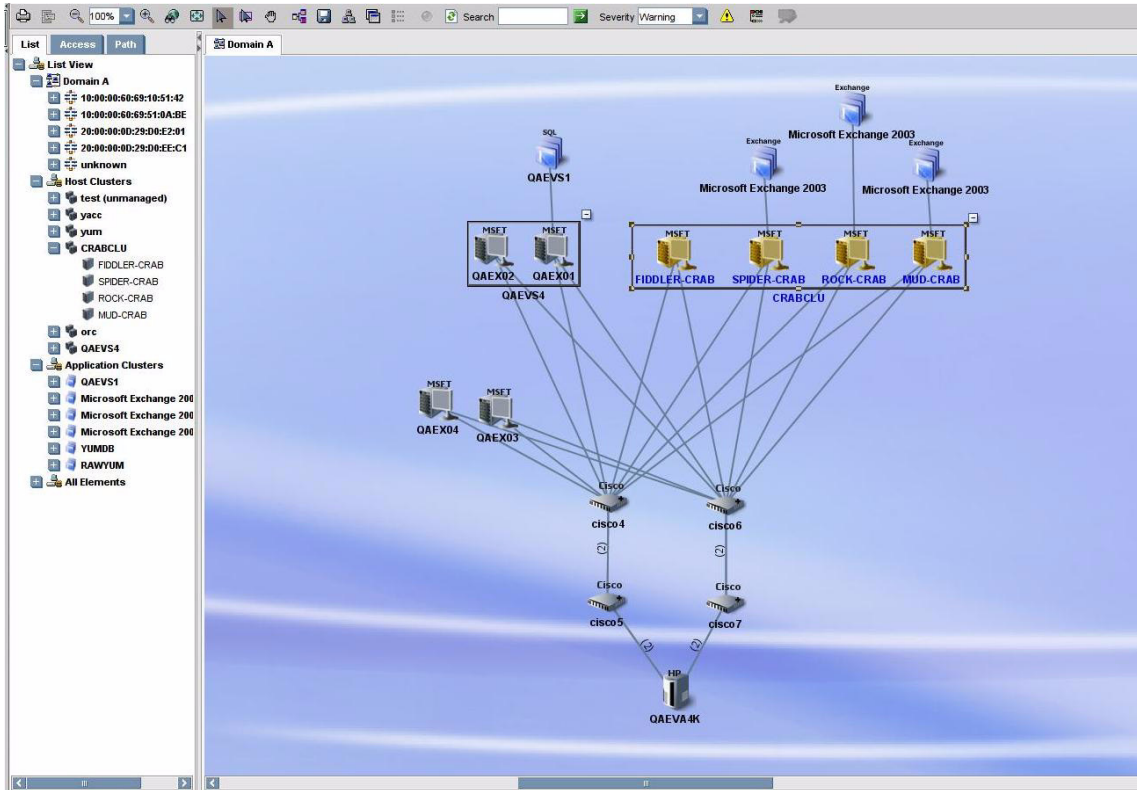
## Clustering in System Explorer

System Explorer has been enhanced to seamlessly support clusters in all areas. You can view connectivity information from all levels on a single canvas — from applications running on clusters, to the storage array spindles that share volumes for all the nodes of a cluster.

For detailed information about System Explorer, see “Viewing Element Topology and Properties” on page 307.

The following figure shows how clusters are displayed in System Explorer. Note that the tree nodes on the List tab reflect the structure of the clusters.

In this figure, the box on the left of the topology canvas shows a cluster with two hosts, and the box on the right shows a cluster with four hosts. Both clusters are in the expanded view mode, so all of the nodes are displayed. To minimize the view of a cluster, click the (-) button.



**FIGURE 5-1** System Explorer Cluster Representation

In the minimized view of a cluster, all of the nodes of the cluster are collapsed into a single box. To expand the display to show all of the nodes, click on the (+) button.

In the minimized view, a dotted line from an application to a cluster indicates that the application only runs on some of the clustered hosts. A solid line indicates that the application runs on all of the clustered hosts.

Double-click a cluster to open the Properties page for the cluster. Double-click an individual cluster node to open the Properties page for that node.

---

# Clustering in Topology

Element topology expands System Explorer's view to show exactly which shared resources a particular application instance uses. Individual paths from application nodes are listed in the path tree as well.

For detailed information about viewing element topology, see "Viewing Element Topology" on page 379.

In the following figure, individual instances of Microsoft Exchange Server 2003 share HP EVA virtual disk array group shared resources:




FIGURE 5-2 Cluster Element Topology Representation

---

# Clustering in Capacity Manager

In Capacity Explorer, it is possible to see the whole capacity utilization by the cluster. Clusters are represented as managed elements, and the capacity calculator intelligently avoids double counting of the capacity from individual nodes at the cluster level.

For detailed information about Capacity Explorer, see “Finding an Element’s Storage Capacity” on page 663.

You can drill down to various levels to see the following details of cluster capacity utilization:

- Whole cluster capacity
- Individual application instance capacity
- Individual cluster node capacity
- Capacity trending over a period of time
- Shared resources of individual nodes

The following figure shows an example of how clusters are represented in Capacity Manager:

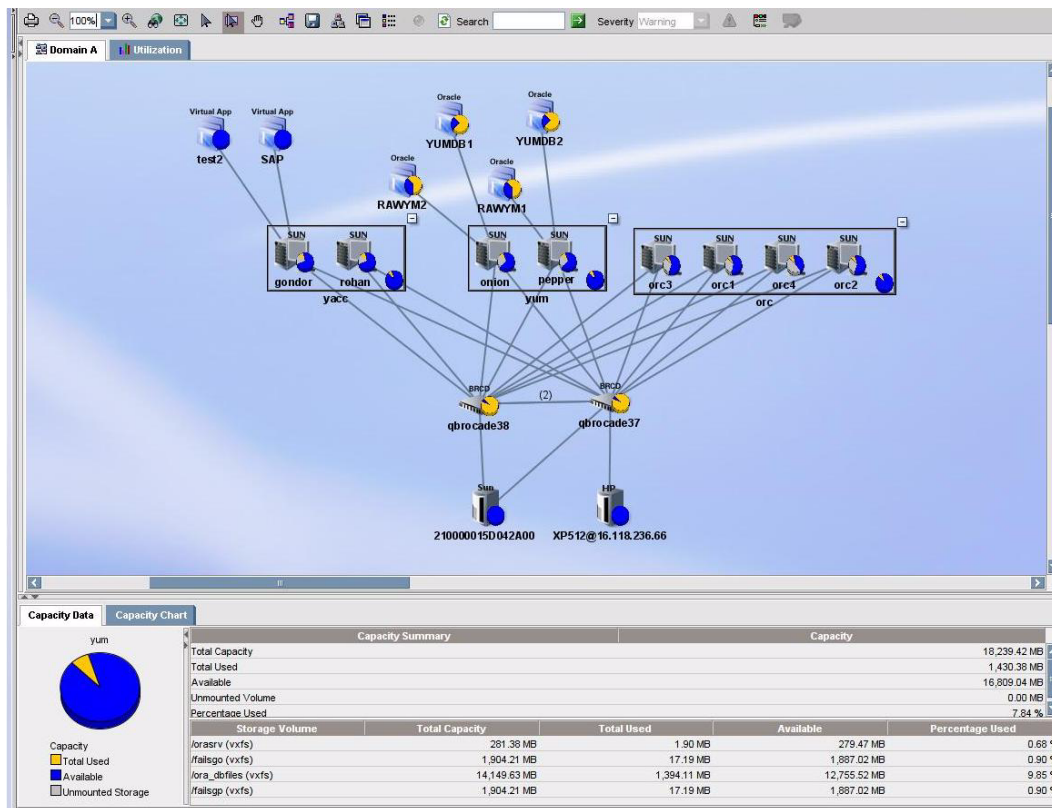


FIGURE 5-3 Capacity Manager Cluster Representation

## Managing Security

---

---

**Caution** – 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 contains the following topics:

- “About Security for the Management Server” on page 175
- “Managing User Accounts” on page 182
- “Managing Roles” on page 189
- “Managing Organizations” on page 192
- “Changing the Password of System Accounts” on page 198
- “Using Active Directory/LDAP for Authentication” on page 200

---

### About Security for the Management Server

The management server offers security based on the assignment of roles and organizations. Role-based security determines access to specific functionality depending on the user account assigned to a role. Organization-based security determines 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 176
- “About Organizations” on page 179
- “Planning Your Hierarchy” on page 181

- “Naming Organizations” on page 182

## About Roles

The management server ships with several predefined roles, which 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, Protection Explorer, and Reporter. Likewise, users assigned to the domain administrator role have access to all of the features, as shown in Table 6-1, “Default Role Privileges,” on page 176.

**TABLE 6-1** Default Role Privileges

Feature	Role					
	CIO	Domain Administrator	Storage Administrator	Server Administrator	Application Administrator	Help Desk
Application Explorer	X	X			X	X
System Explorer	X	X	X	X	X	
Event Manager		X	X	X	X	X
Protection Explorer	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				



## *Granting Global Reporter Access*

Users with access to Global Reporter can view all 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.

## *Domain Administrator Role Privileges*

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.

## *System Configuration Option*

If the System Configuration option is selected for a role, all users assigned to that role will have the administration capabilities shown in the following list:

- 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

If you do not want users belonging to that role to have those capabilities, do not assign the System Configuration option.

## *Roles Used to Restrict Access*

Roles also restrict access to element properties, element records, and Provisioning, as shown in Table 6-2, "Default Role Privileges by Elements," on page 178.

**TABLE 6-2** Default Role Privileges by Elements

Role	Element					
	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
Application Administrator	Full Control	View	View	View	View	View
Help Desk	View	View	View	View	View	View

### *Options for Restricting a Role*

In addition, you can assign one of the following options within a role to further allow or restrict access for a specific element:

- **Full Control** — Lets you view and modify the record for the element on the Asset Management tab, and perform provisioning if applicable.
- **Element Control** — Lets you view and modify the record for the element on the Asset Management tab. You cannot perform provisioning.
- **View** — Lets you only view element properties.

For example, if users belong to a role that only lets them view the element properties on storage systems, those users 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 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 use organizations to 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 assigned to an organization can see only the elements that belong to that 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. A user assigned to OnlyHosts and OnlySwitches 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. The figure below shows a parent-child hierarchy in which BostonWebHosts organization contains two child organizations, BostonWebHost\_Windows and BostonWebHost\_Solaris. BostonWebHosts is a parent because it contains two organizations.



**FIGURE 6-1** Parent-Child Hierarchy for Organizations


If a child contains organizations, it is also a parent. For example, if 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; for example, when you add a new element, you need to add it only once; the change ripples through the hierarchy. For example, if you add an element to BostonWebHost\_Windows, not only users 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 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. For example, if you have a user that oversees all Web hosts in the company, you could assign that user 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. For example, if 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 6-2** 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 view all elements only 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.
- Events from all elements regardless of the user's organization are displayed by Event Manager.
- 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 when you email reports. If you do not have permission to access hosts, the reports you email, including the host-specific reports, 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.

## 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? Or 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 a graphics illustration program, so you can keep track of which organizations are parents and children.

Create the child organizations first, then their parents. See “Adding an Organization” on page 192 for more information.

## Naming Organizations

When you create an organization, give it a name that reflects its members. 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 which are children. When you name an organization, you might want to include a portion of the name of the dominant parent organization. For example, if 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 contains the following topics:

- “Adding Users” on page 182
- “Editing a User Account” on page 184
- “Changing the Password for a User Account” on page 185
- “Changing Your Password” on page 186
- “Deleting Users” on page 186
- “Modifying Your User Profile” on page 186
- “Modifying Your User Preferences” on page 187
- “Viewing the Properties of a Role” on page 188
- “Viewing the Properties of an Organization” on page 189

## Adding Users

This section contains procedures for adding users and authorizing privileges. Only users belonging to the Domain Administrator role can add or modify users.

Keep in mind the following:

- On Windows and Sun Solaris systems — The user name and password must be alpha-numeric, and cannot exceed 256 characters. The user name cannot begin with a number.
- On Linux systems — The user name and password cannot exceed 256 characters.

To create an account:

1. Click **Security > Users**.
2. Click the **New User** button.
3. In the **Login Name** box, enter a name for the user account, for example: jsmith  
This name becomes the user name for the account.
4. (Optional) In the **Full Name** box, enter 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** menu. See “About Security for the Management Server” on page 175 for more information about roles.
6. (Optional) In the **E-mail** box, enter the user's e-mail address.
7. (Optional) In the **Phone** box, enter the user's phone number.
8. (Optional) In the **Notes** box, provide additional information about the user.
9. (Optional) In the **Password** box, enter 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 box blank.

---

10. (Optional) In the **Verify Password** box, enter 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 Security for the Management Server” on page 175 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.
- See “Domain Administrator Role Privileges” on page 177.
- User modifications take 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 210.

If you want to change your password, follow the steps in “Changing Your Password” on page 186.

To modify a user account:

1. Click **Security > Users**.
2. Click the **Edit** () button for the user account you want to modify.
3. To change the account name, enter a new name for the user account in the **Name** box; for example: jsmith  
This name becomes the user name for the account.
4. To change the name assigned to the user account, enter a new name for the account in the Full Name box.  
This information is used to provide a correlation between an account name and a user.
5. To change the role assigned to the user account, select a new role from the Role menu.
6. To change the e-mail address listed, enter a new e-mail address in the **E-mail** box.
7. To change the phone number listed, enter the user's new phone number in the **Phone** box.
8. Change or remove information from the **Notes** box if necessary.
9. To change the password:



- a. Click **Change Password**.
  - b. Enter a new password in the **Password** box.
  - c. Enter the password again In the **Verify Password** box.
  - 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.

---

**Note** – The Everything organization is the default organization that lets users access all current and future elements.

---

11. Click **OK**. The user account is updated.


## 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 menu.
3. Click the **Edit** button () corresponding to the user account you want to modify.
4. Click **Change Password**.
5. Enter a new password in the **New Password** box.
6. Enter the password again in the **Verify Password** box.
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 210 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. Enter a new password in the **New Password** box.
4. Enter the password again In the **Verify Password** box.
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, ask your Domain Administrator to make the changes.

To modify your user profile (other than name, role, and organization affiliation):

1. Click the name of your account in the upper-left corner.



**FIGURE 6-3** Clicking the Name of Your User Account

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 186. 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 **Save Changes**.

## Modifying Your User Preferences

Use the User Preference tab to modify your user preferences for System Explorer and Element Topology. 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.
2. Click the **User Preferences** tab.

## System, Capacity and Performance Manager Preferences

Select one of the following:

- **Load-on-Demand:** Does not populate the tree nodes or display elements in the topology when the page opens (Faster). Use this option for medium to large environments.
- **(Default) Automatic Loading:** Populate fabric tree nodes and display all elements in the topology when the page opens (Slower).

## 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 menu.

If you want events refreshed within a time period, select the **Refresh events automatically** box 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.

## Warnings for Slow Systems Operations

By default, the management server warns you when it encounters issues occurring when 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 187 for information on how to access the User Preferences tab.

## Viewing the Properties of a Role

If you are assigned the Domain Administrator role, you can determine which components a user can access by viewing the properties of the user's role.

To view the properties of a role:

1. Click **Security > Users**.
2. In the Role column, click the name of the role.

The following information for the selected role is displayed:

- 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 — How much access the user has to a type of element, such as hosts, storage systems, switches, and applications. See “About Security for the Management Server” on page 175 for more information.
- Access to the <product name> — 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 “Editing Roles” on page 190.

## Viewing the Properties of an Organization

If you are assigned the Domain Administrator role, you can determine which elements a user can access by viewing the properties of the user's organization

To view the properties of an organization:

1. Click **Security > Users**.
2. In the Organization column, click the name of a organization.
3. Take one of the following actions:
  - To determine which elements are in a child organization, click the link of the child organization.
  - To learn more about an element, click the element's link to display 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 Security for the Management Server” on page 175 for more information.

To learn how to edit an organization, see “Editing an Organization” on page 195.

---

## Managing Roles

This section contains the following topics:

- “Adding Roles” on page 190
- “Editing Roles” on page 190
- “Deleting Roles” on page 191

## 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 “About Security for the Management Server” on page 175 for more information about roles and organizations.

Keep in mind the following:

- The Role Name and Description boxes 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. Click **Security > Roles**.
2. Click **New Role**.
3. In the Role Name box, enter 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 box, enter a description for the role; for example: Role for those in quality assurance.  
The description cannot be more than 1024 characters.
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.  
See “Options for Restricting a Role” on page 178.
6. Select the features you want a user to be able to access.  
See “Management Server Components” on page 3 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 “About Security for the Management Server” on page 175 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.
- After you click **OK** 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 box does not accept special characters, except spaces and the following characters: \$, -, ^, ., and \_

To edit a role:

1. Click **Security > Roles**.
2. Click the **Edit** () button.
3. Make the desired changes:
  - To edit the name of the role, change the name in the Role Name box. The name can contain spaces, but it cannot be longer than 256 characters.
  - To edit the description of the role, change the description in the Description box. The description cannot be more than 1024 characters.
  - 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.

See “Options for Restricting a Role” on page 178.
4. Select the features you want a user to be able to access.
 

See “Management Server Components” on page 3 for more information about these features.
5. 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. Click **Security > Roles**.

2. Select **Roles** from the menu.
3. Click the corresponding **Delete** button ()  
The role is deleted.

---

## Managing Organizations

This section contains the following topics:

- “Adding an Organization” on page 192
- “Viewing Organizations” on page 194
- “Editing an Organization” on page 195
- “Removing an Organization” on page 196
- “Removing Members from an Organization” on page 196
- “Filtering Organizations” on page 197

### Adding an Organization

You can create new organizations to restrict access to certain elements. For example, if 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 will only be able to access the elements you specified.

See “About Security for the Management Server” on page 175 for more information about roles and organizations.

Keep in mind the following:

- Create child organizations first, then their parents.
- Events from all elements regardless of the user’s organization are displayed by Event Manager.
- 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.
- Moving a cluster from one organization to another moves all of the cluster’s nodes to the target organization.

To add an organization:



1. Click **Security > Organizations**.
2. Click the **New Organizations** button.
3. In the **Name** box, enter 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 caret (^) symbol—currently the system allows the caret symbol to be entered, but the caret symbol should not be included in an organization's name.
4. In the **Description** box, enter a description for the organization.

The Description box cannot have more than 1024 characters.
5. Click **Add or Remove Members** to determine which elements the user will see.
6. To add elements:
  - a. Expand the Element Types node in the tree, and select the element type that you would like to add.
  - b. In the right-hand pane, select the elements you would like to add by clicking the appropriate check boxes.
  - c. Click **Add**.
  - d. The selected elements are added to the Organization Members pane. To add storage volumes to the organization, see “Adding Storage Volumes to an Organization” on page 194.
7. To add organizations:
  - a. Click the **Organizations** node.
  - b. In the right-hand pane, select the elements you would like to add by clicking the appropriate check boxes.
  - c. Click **Add**. 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 “About Security for the Management Server” on page 175 for more information.
8. Click **OK** when you are done adding the elements and organizations.

## 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 **Add or Remove Members**.
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 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 **Submit Query**.
5. Select the storage volumes you want to add to the organization. Click the **+Ports** link in the Ports column 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.
7. Click **OK**. The selected volumes are added to the Organization Members pane.

## Viewing Organizations

The Setup Organizations page lists the organizations with their descriptions. The page also shows 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 directly to an organization. This number does not include those within the child organization. A zero (0) in the Elements column indicates that the organization contains only child organizations; 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**.

To access information about a child organization, click its link in the Child Organization column.

# Editing an Organization


When elements are removed from an organization, users belonging only to that organization are no longer able to access the removed elements.

See “About Security for the Management Server” on page 175 for more information about roles and organizations.

Keep in mind the following:

- Depending on your license, role-based security may not be available. See the List of Features accessible from the Documentation Center.
- Only users belonging to the Domain Administrator role can edit organizations.
- Only active organizations can be edited.
- You cannot edit the Everything organization.

To edit an organization:

1. Click **Security > Organizations**.
2. Click the Edit () button.
3. To change the name of the organization, enter a new name in the Name box.

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 include special characters, except spaces and the following characters: \$, -, ., and \_
  - Cannot contain the carot (^) symbol.
4. To change the description of the organization, enter a new description in the **Description** box.  
You cannot enter more than 1024 characters in the **Description** box.
  5. Click **Add or Remove Members**.
  6. Add or remove elements as described in “Adding an Organization” on page 192 and “Removing Members from an Organization” on page 196.
  7. Once you are done adding or removing elements, click **OK** in the Add Organization or Remove Organization page.
  8. In the Edit Organization page, click **OK**.


## Removing an Organization

When an organization is removed, users assigned only to that organization are no longer able to access its elements. 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.
- Only users belonging to the Domain Administrator role can delete organizations.
- 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.

To delete an organization:


1. Click **Security > Organizations**.
2. Click the Delete () 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.

Use one of the following methods to remove an element from an organization:

- In the Edit Organization window, click the Delete () 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, and then click **Remove**.

- Only users belonging to the Domain Administrator role can remove members from an organization.


## 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: WindowsHosts and SolarisHosts. If you want to view elements only in WindowsHosts and not in SolarisHosts organizations, you could use the filtering feature to activate only the WindowsHosts 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 the filtering feature for organizations.
- If you do not want to view an element, deselect all child organizations containing that 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 SolarisHosts organization. The SolarisHosts organization is contained in the Hosts organization. You must deselect the SolarisHosts organization and the Hosts organization if you do not want to see the Solaris hosts.
- The filter for organizations does not appear in Event Manager. Events from all elements regardless of the user's organization are displayed by Event Manager.
- 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 6-4** Clicking the Organization Link

2. Deselect the organizations that contain the elements you do not want to obtain information about. For example, if you want to view only the elements in the WindowsHosts organization, you would select only WindowsHosts. If you have a parent organization named Hosts that contains SolarisHosts and WindowsHosts, you would need to deselect SolarisHosts and Hosts. You would need to deselect Hosts because it contains organizations other than WindowsHosts.

If you belong to the Domain Administrator role, links are displayed for the organizations. To learn more about the contents of an organization, click its link.

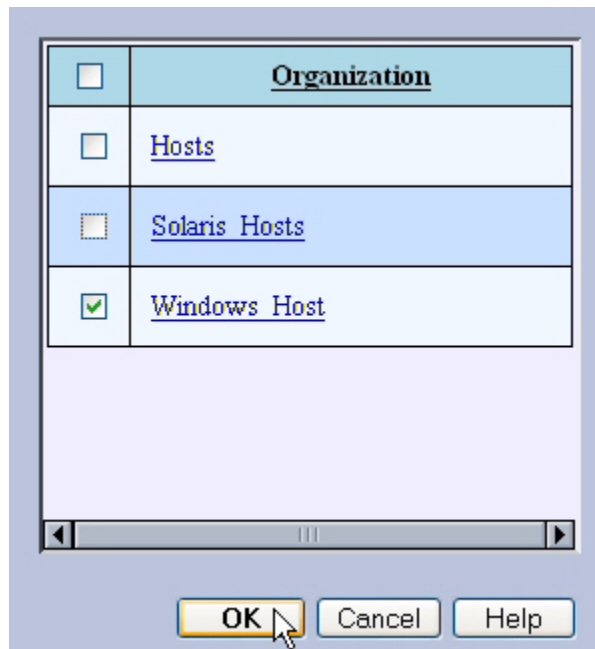


FIGURE 6-5 Filtering Organizations

3. Click **OK**.

You can now only obtain information about elements in the active organizations. These active organizations are listed in the link next to the filter button, as shown in the following figure.



FIGURE 6-6 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 to create and update the management server database. Default password: `change_on_install`
- **SYSTEM** — Used to create and upgrade, import, export and re-initialize the management server database. 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`

To change the passwords of the SYS, SYSTEM, RMAN\_USER, and DB\_SYSTEM\_USER accounts, you must use the Database Admin Utility, so the management server is aware of the changes. Do not change the password for any 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 password requirements for the management server are:

- Must have a minimum of three characters
- Must start with a letter
- May contain only letters, numbers and underscores (`_`)
- May 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 289.
2. Click **Change Passwords** in the left pane.
3. Select an account name from the User Name box.
4. Enter the current password in the Old Password box.
5. Enter the new password in the New Password box.
6. Re-enter the password in the Confirm Password box.
7. 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 that 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 both AD and LDAP. It is important to enable either AD 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`.
- Business Tools do not work when the management server is configured for AD/LDAP authentication.

To use AD/LDAP to authenticate your users, complete the following procedures:

- “Step 1 — Configure the Management Server to Use AD or LDAP” on page 200
- “Step 2 — Restart the AppStorManager Service and Login as the Designated Admin Account” on page 209
- “Step 3 — Add Users to the Management Server” on page 210
- “Step 4 — Provide Login Information to Your Users” on page 210

## Step 1 — Configure the Management Server to Use AD or LDAP

If you want to use AD/LDAP, you must modify the `login-handler.xml` file. How you modify the `login-handler.xml` file depends on whether you plan to use AD or LDAP.

To configure the management server:



- To use AD, see “Configuring the Management Server to Use Active Directory” on page 201
- To use LDAP, see “Configuring the Management Server to Use LDAP” on page 205

## Configuring the Management Server to Use Active Directory

By default, AD allows connections with `domain\username`, instead of with the distinguished name (DN) used by a generic LDAP server. However, you can use the generic LDAP server setup to authenticate with AD, as described in “Configuring the Management Server to Use LDAP” on page 205.

To specify the management server to use AD:

1. Before switching to AD authentication mode, the management server needs to be configured with a designated AD user and other AD-specific credentials. At startup, the designated AD user is mapped to the built-in Admin user and overrides it with the AD user information.

---

**Caution** – Make sure the administrator account has already been created in AD 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`

**UNIX systems:** `$MGR_DIST/Data/Configuration`

- b. In the `login-handler.xml` file, change the value of the `<AdminAccountName>` tag to the name of a user account in AD, 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 AD.

For security reasons, it is recommended that the designated user not be the AD Domain Administrator

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 AD.
- 389 is the port on which AD 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 AD.

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 allows a user to also 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 AD 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.

---

**Caution** – 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 AD and one for LDAP. Make sure you also change the value of the <CaseSensitiveUserName> tags that are children of the <ActiveDirectory> tag.

9. Provide the AD 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</LDAP
Env>
-->
<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>
  <ssl>false</ssl>

```

```

<ShadowPassword>false</ShadowPassword>
<CaseSensitiveUserName>false</CaseSensitiveUserName>
<SearchBase>DC=MyCompanyName,DC=COM</SearchBase>
    <FullNameAttribute>displayName</FullNameAttribute>
    <EmailAttribute>mail</EmailAttribute>
  </ActiveDirectory>
</LoginHandler>

```

## Configuring the Management Server to Use 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.

---

**Caution** – 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

**UNIX systems:** \$MGR\_DIST/Data/Configuration

- 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</Logi
nHandlerClass-->

```

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 allows a user to also 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 AD 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. For example, the management server sees `MyUserName` and `myusername` as different users.

The `login-handler.xml` file contains two sets of `<CaseSensitiveUserName>` tags: one for AD and one for LDAP. Make sure you also 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).

---

**Caution** – Different LDAP implementations may be using different keynames for CN. The appropriate keyname should be named 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>PreferredUser\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</Logi
nHandlerClass>
<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</LDA
PEnv>
-->
<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>
    <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.

---

**Caution** – 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 menu.
- d. To start the management server, right-click **AppStorManager** and select **Start** from the menu.

On UNIX systems:

- 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 AD or LDAP” on page 200.

For example, the user name would be the following:

- AD — domain\PrimaryUser
- LDAP — PrimaryUser

where `PrimaryUser` is the name of the user account in LDAP or is the designated primary user in AD.

The password would be the following: `[NTdomainpassword]`.

## Step 3 — Add Users to the Management Server

Once the management server is configured for Active Directory/LDAP, the users can be added to the management server. This is required to prevent accidental access to the management server from other AD/LDAP users. Until the user is authenticated against AD/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 AD/LDAP, the user is tagged as an external user and the user’s password must be managed through AD/LDAP.

To add a user to the management server:

1. Log onto the management server by using the designated Admin account specified in “Step 1 — Configure the Management Server to Use AD or LDAP” on page 200.
2. Create the users as described in “Adding Users” on page 182 observing the following rules:
  - AD: 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 AD/LDAP.
  - It is not necessary to create a password, since the passwords used for login are those already configured on either the AD or LDAP server.

## Step 4 — Provide Login Information to Your Users

Notify your users that they are now able to log into the management server, and provide them with the user name and password you have specified in Active Directory/LDAP

---

**Caution** – Remind your users not to give the password they use to access the management server to anyone. Since user credentials are now stored in AD/LDAP, the password used to access the management server may also be used to access other accounts. In some instances, it may even be their network user name and password.

---



## Managing Licenses

This chapter contains the following topics:

- “Modifying the License Summary Page” on page 217

The management server restricts the number of 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 7-1, “License Restrictions,” on page 213.

**TABLE 7-1** License Restrictions

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 7-2, “Determining Managed Access Points,” on page 215 for more information.	Number of MAPs
Backup Size	The management server determines licensing for Protection Explorer through gigabytes (GB). The management server compares the number of gigabytes for Protection Explorer with what you are backing up. If you are backing up more than your license allows, you are warned the next time you log onto the management server.	Gigabytes (GB)
Raw NetApp Capacity	The Raw NetApp Capacity is the total disk capacity (unformatted capacity) of all discovered NetApp filers.	Terabytes (TB)

**TABLE 7-1** License Restrictions

Type of Restriction	Description	Unit of Measurement
Managed Exchange Instances	The management server determines licensing for Microsoft Exchange instances by counting the number of instances of Microsoft Exchange it manages.	Number of instances of Microsoft Exchange the software manages
Managed Database Instances	The total number of instances of the following databases the software manages: <ul style="list-style-type: none"><li>• Microsoft SQL Server</li><li>• Oracle</li><li>• Sybase Adaptive Server Enterprise</li><li>• InterSystems Caché</li></ul> This total is broken down by each type of database in the table.	Number of managed databases
For File Server SRM	<p>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.</p> <p>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.</p> <p>You do not have to monitor everything associated with your file server. You can choose to manage only the mount points that are important to you. Only the files associated with these mount points are counted toward the file server TB.</p>	Terabytes (TB)

---

**Caution** – The management server Current Usage Summary is first updated six hours after the management server (AppStorManager) starts, and then the updates occur every 24 hours thereafter. 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 is first started. For example, the first update of the Current Usage Summary table occurs six hours after the management server is first started. The following updates occur every 24 hours.

If the management server is 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 Table 7-2, “Determining Managed Access Points,” on page 215.

**TABLE 7-2** 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 MAPs are 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 add 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 and with a total of 0 FC ports, as shown in the following figure:



**FIGURE 7-1** 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.

If we had a configuration which included a switch, two managed hosts, and several unmanaged hosts, the MAPs would not be used against the unmanaged hosts.

Some switches allow the user to turn off an unused GBIC. (Gigabit Interface Converter). If a GBIC is turned off, the port is not counted. However, if the GBIC is turned on, or if there is no GBIC, the port is counted.



---

# Modifying the License Summary Page

If you have purchased additional elements, you must modify the License Summary page. For example, assume you purchased an additional 200 MAPs, which lets you monitor 200 more devices, such as hosts, switches, and storage systems. To make the management server aware of these changes, you must enter the new total of MAPs you are licensed to use on the License Summary page.

---

**Caution** – Select only the applications you are licensed to access. Enter only the MAPs, terabytes and instances, you are authorized to use.

---

To modify the License Summary page:

1. Select **Security > Licenses**.
2. Select the applications you have recently purchased.
3. If you have added one or more of the following, add the amount you have purchased to the total listed.
  - MAPs
  - Gigabytes that will be backed up by Protection Explorer.
  - Terabytes that will be scanned by File SRM
  - Number of instances of each type of application you want to monitor.
4. Select **Save Changes**.
5. When you are shown the license agreement, accept the license if you agree with its terms.



# Configuring the Management Server

---

This chapter contains the following topics:

- “Trap Generation” on page 219
- “Setting Up E-mail Notification” on page 221
- “Configuring Print Settings” on page 222
- “Setting the Date and Time for Scheduled Tasks” on page 225
- “Managing Getting Discovery Details” on page 225
- “Modifying Collector Settings for Newly Discovered Elements” on page 229
- “Managing Product Health” on page 230
- “Managing Logging” on page 233
- “Managing the Display of Events” on page 241
- “Managing File Server SRM” on page 244
- “Managing Backup Collection” on page 244
- “Managing Reports” on page 248
- “Managing Performance Collection” on page 269
- “Editing the Locale and Currency Settings” on page 273
- “Process Names” on page 275
- “Editing a Collector Schedule” on page 276

---

## 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 220.

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 `AppIQ-Traps-v1.mib` for SNMPv1 or `AppIQ-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.

The software does not have the capability to forward traps received from other devices. It can take events from Event Manager and create SNMP traps from them. These traps are generated using an SNMP MIB.

The software does receive SNMP traps from some devices. These traps are translated into events in Event Manager. When they are sent out as SNMP traps, the information in the trap will be the same as the original device trap, but the format of the trap will be different. For example, the trap will contain the original severity and description information, but the Trap OID, fields, and codes will be different.

To configure trap forwarding:

1. Access the management server, as described in “Accessing the Management Server” on page 11.
2. Select **Configuration > Traps**.
3. Select the **Enable Trap Generation** option.
4. In the SNMP Community String box, enter the SNMP community string, which is used for filtering.  
**Important:** If you enter a value in the SNMP Community String box, the SNMP agent must know the SNMP community string entered in the box to receive the SNMP trap.
5. Click **Save**.
6. In the New SNMP Destination box, enter the IP address of the server running an SNMP agent.
7. Click **Add**.

## Changing the Default to SNMPv2

To change the system to send SNMPv2 traps:

1. Select **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.  

```
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 box. How you paste the text depends on your Web browser.
7. To make sure the property is not commented out, remove the hash (#) symbol in front of the property.
8. When you are done, click **Save**.

You do not need to restart the AppStorManager service for your changes to take effect.

---

## Setting Up E-mail Notification

---

**Caution** – Depending on your license, e-mail notification may not be available. See the List of Features to determine if you have access to e-mail notification. The List of Features is accessible from the Documentation Center (**Help > Documentation Center**).

---

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 11.
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 box, enter 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 box, enter the port of the SMTP server you want to use to send the e-mail notification.
6. In the User Name box, enter a user name for the SMTP server.
7. In the Password box, enter a password for the SMTP server.
8. In the Verify Password box, enter the password you entered previously.
9. *Required:* In the Sender box, enter the e-mail address of the sender.  
This address is displayed in the From box 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 box, enter the e-mail address you want to receive replies in the Reply to box.
11. Click **Save**.

---

**Caution** – You should try sending a test e-mail.

---

To send a test e-mail:

1. In the To box, enter an e-mail address.  
The software verifies that the address entered has a correct form.
2. In the Subject box, enter a subject to distinguish this e-mail from notification of a real event, for example:  
Testing E-mail Notification
3. In the Message box, enter a message, for example:  
I'm just testing e-mail notification.
4. Click **Send Test Message**.

---

## Configuring Print Settings

To configure print settings:

5. Click the  button.

6. 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.

---

**Caution** – 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.
7. 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.

---

**Caution** – 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.

8. 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.

---

**Caution** – 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.

9. To preview your pages, click the **Preview** tab. Then click the page you want to preview.

The page appears in the right pane.

10. When you are ready to print, click **Print**.

11. Click **Close**.

---

**Note** – To revert back to all of the original settings, click the **Default** button next to the **Print** button.


---



---

# Setting the Date and Time for Scheduled Tasks

While configuring the management server, there are several occasions when you will need to set the date and time for a scheduled task. To set the date and time for a scheduled task:

1. Click the calendar icon, .
2. In the Time box, enter 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 task to run. Today's date is highlighted in pink. Click **Set**.

The date and time appear in the Next Scheduled Run box in the yyyy-mm-dd hh:min format.

If you change the date in the box 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.

3. In the Repeat Interval box, enter an interval. Select one of the following units from the list:
  - **Second(s)**
  - **Minute(s)**
  - **Hour(s)**
  - **Day(s)**
  - **Week(s)**

---

## Managing Getting Discovery Details

You can schedule the management server to obtain discovery details at a specified interval. The management server provides several types of discovery details:

- **Include infrastructure details** - If the Include infrastructure details option is selected, the management server gathers detailed information about the SAN infrastructure. This process can be network intensive.
- **Include backup details** - To obtain the latest backup information, select the Include backup details option, and schedule the discovery to run nightly after you run your backup sessions. It is recommended you do not schedule the discovery of infrastructure details and backup details to run at the same time.

This section contains the following topics:

- “Adding a Discovery Schedule” on page 226
- “Disabling a Schedule” on page 227
- “Editing a Schedule” on page 228
- “Removing a Schedule” on page 228
- “Modifying Collector Settings for Newly Discovered Elements” on page 229

## Adding a Discovery Schedule

Schedule getting details when the network is not busy.

Keep in mind the following:

- All collectors are stopped during a discovery using the Include infrastructure details option. 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, take care to avoid scheduling conflicts—concurrently scheduled discovery tasks—and ensure 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 823.
- Do not run Get Details for all discovery groups simultaneously.
- Hosts discovered with CIM extensions from Build 5.1 and later of the the product cannot be added to discovery groups. These hosts can be placed independently into scheduled Get Details tasks without being part of a discovery group. This allows you greater flexibility when gathering discovery data.

If you are upgrading from a previous build of the product, and you rediscover your hosts, they will be moved out of their existing discovery groups. Each rediscovered host would be placed in its own discovery group. If the original discovery groups containing these hosts were included in scheduled Get Details tasks, the schedules would be modified to contain the new discovery groups for rediscovered hosts.

To schedule discovery details:

1. Access the Discovery page by selecting **Configuration > Discovery**.
2. Click the **Discovery Schedule** tab.
3. Click **New Schedule**.
4. In the Name box, enter a name for your discovery schedule.
5. In the Description box, enter a description for your discovery schedule.
6. Select one or more of the following:

- **Include infrastructure details** - If the Include infrastructure details option is selected, the management server gathers detailed information about the SAN infrastructure. This process can be network intensive.
- **Include backup details** - To obtain the latest backup information, select the Include backup details option, and schedule the discovery to run nightly after you run your backup sessions. It is recommended you do not schedule the discovery of infrastructure details and backup details to run at the same time.
- **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. 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. For more information, see “Excluding EMC Symmetrix Storage Systems from Force Device Manager Refresh” on page 62.

7. Select the **Enable** check box.

8. Set the date, time, and repeat interval for this task. For more information, see “Setting the Date and Time for Scheduled Tasks” on page 225.

9. Click **Next**.

10. Select the discovery groups you want included in the discovery:

---

**Caution** – Only the elements in the discovery groups you select are included in discovery.

---

d. Select the items in the Discovery Groups section and click the **Add Selected Items to Discovery Group** button to move them to the Custom Discovery Groups section.

e. Click **Finish**.


The scheduled Get Details operation appears in the list of scheduled discoveries

**FIGURE 8-1**

## Disabling a Schedule

To disable a schedule for getting SAN topology details:

1. Access the Discovery page by selecting **Configuration > Discovery**.


2. Click the **Edit** () button corresponding to the discovery schedule you want to disable.

3. Deselect the **Enable** option.
4. Click **Next**.
5. Click **Finish**.

The schedule is disabled.


## Editing a Schedule

To edit a schedule:

1. Access the Discovery page by selecting **Configuration > Discovery**.
2. Click the **Edit** () button corresponding to the discovery schedule you want to modify.
3. If necessary, change the following properties:
  - Name
  - Description
  - Type of discovery
  - Schedule
4. Click **Next**.
5. If necessary, change the discovery groups you want assigned to the schedule.
6. Click **Finish**.

## Removing a Schedule

To remove a schedule:

1. Access the Discovery page by selecting **Configuration > Discovery**.
2. Click the **Delete** () button corresponding to the discovery schedule you want to remove.
3. When prompted to confirm, click **OK**.

The schedule is removed.

---

# Modifying Collector Settings for Newly Discovered Elements

The management server is capable of collecting many different types of data. Instead of using a single large process, these data are gathered using many different collectors. You can decide whether all or some of these collector schedules should start or stop when a new element is discovered.


To review the list of collectors that are available for each element type, access the Discovery page by selecting **Configuration > Discovery**, and click the **Collector Settings** tab.

The Default Collector Settings page provides a comprehensive listing of what collectors are available for each element type and what category each collector is classified as.

To help you find your collectors quickly, this page offers a set of filters. The filter area is collapsed by default. To expand the filter area, click the **+** symbol. The following filters are supported:

- **Element Type** – Retrieves the list of collectors designed to support the specified element types.
- **Collector Category** – Specifies whether you want to see a listing of all the collectors, or only collectors designed to gather performance statistics, report data, or capacity data.

To apply the filter settings, click **Filter** to refresh the content of the page. To restore the filters to their default settings, click **Reset**, and refresh the collector page.

You can modify the default collector settings for all future discovery elements by clicking the **Edit** () button, or by selecting one or more collectors using the checkboxes and clicking **Edit Selected**.

The following properties are available when you modify default collector settings:

- **Enable selected collectors for newly discovered elements** – Each collector is designed to gather data for a specific type of elements. If you select this checkbox, the management server will start the selected collectors automatically whenever it discovers a system matching the element type supported by a selected collector. This means that the start time of the selected collectors for the newly discovered system is set to the same time as when the system is discovered.

If you do not select this checkbox, the collectors are still available to the discovered elements, but the collector schedules are stopped by default. You can

start any of the collectors after a system is discovered using one of the Data Collector Configuration pages. For example, **Configuration > Report**, **Configuration > Performance**, or **Configuration > Backup**.

- **Repeat Interval** – Specifies how often the collector should run on a recurring basis when the management server starts a collector schedule for an element type.

---

## Managing Product Health

To obtain information from Product Health:

1. Add the management server to your discovery list:
  - a. Select **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 is automatically installed on the management server so you can 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 231.
- **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 240.
- **Log files** - You can view and download product logs, as described in "Accessing the Log Files" on page 233 and "Downloading Logs to a File Using the Download Logs Feature" on page 235.
- **Scheduled RMAN backups** - This feature lets you schedule RMAN backups. See "Scheduling RMAN Hot Backups" on page 287. If the buttons on the RMAN Backup page are disabled, the product is set to No archive mode. See "Changing the Archive Mode" on page 296 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 232 and "Customizing Properties" on page 232 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.

---

**Caution** – To obtain information from Product Health, you must have already discovered the management server and obtained element details from it. For more information on how to discover a host, see “Step 1 — Discovering Your Hosts and Backup Manager Hosts” on page 105.

---

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. Select **Configuration > Product Health**.
3. Set the date, time, and repeat interval for this task. For more information, see “Setting the Date and Time for Scheduled Tasks” on page 225.
4. Click **Save Changes**.
5. To view the results of the monitoring, click the **Results** 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.

---

**Caution** – 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. For information on how to discover a host, see “Step 1 — Discovering Your Hosts and Backup Manager Hosts” on page 105.

---

To access the results for disk space monitoring:

1. Select **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:

- Date/Time
- Disk Capacity
- Free Space
- Database Files
- Archive Files
- RMAN Files
- Temp Tablespace

## Advanced Settings

This section contains the following topics:

- “Modifying Java Memory Settings” on page 232
- “Customizing Properties” on page 232

### Modifying Java Memory Settings

Do not modify the Java memory settings on the 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. Select **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.
5. Return to the Advanced page (step 1 and step 2).
6. Paste the copied text into the Custom Properties box. How you paste the text depends on your Web browser.



7. Make your changes in the Custom Properties box. To make sure the property is not commented out, remove the hash (#) symbol in front of the property.
8. When you are done, click **Save**.

---

## Managing Logging

This section contains the following topics:

- “Accessing the Log Files” on page 233
- “Downloading Logs to a File Using the Download Logs Feature” on page 235
- “Downloading Logs to a File Using the Log Download Utility” on page 236
- “Downloading the User Audit Log” on page 237
- “Downloading the Discovery Summary Log” on page 237
- “Displaying a Log File in a Command Prompt Window” on page 238
- “Changing the Provider Log Level” on page 238
- “Enabling the Scanning of Critical Events of the Management Server Database” on page 240
- “Viewing the Results of Logging” on page 240

## 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, and there is a trace for the XML received from a CIMOM. The log files may contain information that is difficult to understand.

Some log files will be more appropriate for your specific needs than others, and some will be more useful for troubleshooting or other support needs. Some are for internal use only. The following lists logs files that would likely be the most useful to you, although it is not a complete list of all the logs.

- appstorm.<timestamp>.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 “Downloading Logs to a File Using the Download Logs Feature” on page 235.

# About Log Files

On the management server, the following log files rollover on Startup, at the start of a new day (midnight), and by size:

- appstorm.<timestamp>.log
- AppstormProvisioning.log
- AppstormRemoteConsole.log
- Discovery.log
- GAEDSummary.log
- LicenseChanges.log
- userAudit.log
- All CIMOM logfiles

The following provides information about log file timestamps, sort criteria, configurable parameters, and adding a trace for the XML received from a CIMOM.

**Log file timestamp** - A timestamp (YYMMDD-HHMMSS) is inserted into the filename at its creation, making its origin more quickly identified. (i.e., appstorm.20071012-122025.log).

**Log file sort criteria** - Logfiles sort in order of their creation, based upon the timestamp in their filename.

**Log file configurable parameters** - Configurable parameters for all log files are these:

- Maximum size of the logfile before it rolls over (`MaxFileSize`). This parameter resides in `log4j.xml` and is used to limit the size of an individual log file.
- Maximum total amount of space the logfile can use (`MaxTotalSize`). This parameter resides in `log4j.xml` is is used to limit the total size of a set of log files (e.g., all appstorm logs).

Log file “appenders” manage the log file rollover when the `MaxFileSize` and `MaxTotalSize` parameters are reached. These parameters can be changed for any log file by using the appstorm.<timestamp>.log appenders at the following directory:

```
<management_server_install_directory>/JBossandJetty/server/appiq/conf/log4j.xml
```

At the `log4j.xml` directory indicated above, change the appender values to the new desired values:

```
<param name="MaxFileSize" value="100MB"/> <!--Max size of a file before it is rolled over -->
```

```
<param name="MaxTotalSize" value="900MB"/> <!--Max size of a all these log files, oldest is deleted when size is exceeded -->
```

**Example 1 (Log file rollover based on size)** - Assume the appstorm.<timestamp>.log file MaxFileSize=100MB and MaxTotalSize=900MB.

- If the size of the current appstorm.<timestamp>.log file exceeds 100MB before the next day starts, a new appstorm.<timestamp>.log file is created.
- If any rollover occurs, and the total size of all appstorm.<timestamp>.log files exceeds 900 MB, the oldest appstorm.<timestamp>.log files are deleted until the total size is below 900 MB.
- Whenever a time-based or size-based rollover occurs, a footer is appended to the current file, and a header is placed on the new file. These headers and footers describe why the rollover occurred and the logfile to, or from, which it is being rolled.

**Example 2 (Log file rollover based on time)** - Assume a new day occurred. The current logfile (appstorm.20071012-154625.log) would receive this footer:

```
****Log File Rollover due to Time****
```

```
****Next Log
```

```
File:<Installation_Directory>/logs/appstorm.20071013-000055.log****
```

The next logfile (appstorm.20071013-000055.log) would receive this header:

```
****Log File Rollover due to Time****
```

```
****Previous Log
```

```
File:C:/hp/StorageEssentials/logs/appstorm.2007-154625.log****
```

**Adding trace for XML received from CIMOM** - Traces are normally very large files. For that reason, the trace is turned “off” by default. To add a trace, go into the properties file in the following directory:

```
%/JBoss4_DIST%\server\appiq\conf
```

At the conf directory, uncomment the line shown below by deleting the pound sign (#):

```
#wbem.debug.sml=1
```

After uncommenting the line, set the level to at least 3 for the XML traces to be written. After they are written, a user can go to the /JBoss4\_DIST%\bin directory to view them.

## Downloading Logs to a File Using the Download Logs Feature

If you run into problems with the management server, use the Download Logs 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` file is only generated when you run reports.

---

To save all logs to a file:

1. Select **Configuration > Product Health**.
2. Select **Log Files** in the Product Health tree.
3. Click **Download Logs**.
4. When you are asked if you want to open or save the file, save the file.
5. Enter a name for the \*.zip file, and select the directory to which you want to save the file.

---

**Caution** – Make sure the zip file is saved to a location other than the local disk drives of the management server.

---

6. Click **Save** in the Save As window.

## Downloading Logs to a File Using the Log Download Utility

In addition to the Download Logs feature, there is also an automated process of gathering and downloading logs that is accessible from a command-line utility. This command-line utility is helpful in a situation when the user is unable to access the Download Logs feature from the user interface. For example, if the the management server is unable to start.

---

**Note** – The Log Download Utility does not trigger CIMOM thread dumps, Environment variable dumps, Port usage information, or the latest JBoss thread dump information. The Download Logs button on the Configuration screen can be used to trigger the items not triggered by the Log Download Utility.

---

To download logs to a file by using the Log Download Utility:

1. Open the Command Prompt window and go to the following directory:  
`%MGR_DIST%\Tools\logDownloader`
2. In the same Command Prompt window, enter the following:  
`logDownloader.bat [Target Directory]`

The Log Download Utility creates a zip file of all of the log files on the management server named:

AllLogsxxx-xx.zip

where xx-xx is the time stamp of the collection.

The zip file is copied to the following directory unless you specify a different location:

%MGR\_DIST%\Tools\logDownloader

## 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. Select **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.

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. Select **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.

## 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 appstorm.<timestamp>.log
```

where `appstorm.<timestamp>.log` is the log file you want displayed in the command prompt window and `<timestamp>` is the timestamp for the log file.

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\appstorm.<timestamp>.log` file. You can change the level of logging provided by selecting the new level of logging from the **Provider Log Setting** menu 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 “Logging Levels” on page 239 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 8-1** 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 and servers.	Are having difficulty obtaining information from hosts or servers.
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:

1. Select **Configuration > Product Health**.
2. Select **Log Files** in the Product Health tree.
3. Set the date, time, and repeat interval for this task. For more information, see “Setting the Date and Time for Scheduled Tasks” on page 225.

---

**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.

---

4. Click the **Enable** option.
5. Click **Save Schedule**.
6. To view the results of the scanning, click the **Results** 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. Select **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 contains the following topics:

- “Controlling the Display of Cleared and Deleted Events” on page 241
- “Modifying the Clearing and Deletion Frequency” on page 241
- “Configuring the Clearing of Events” on page 242
- “Configuring the Deletion of Events” on page 243

## 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.

---

**Caution** – Events are only removed or marked cleared when their automatic delay time is completed. See “Configuring the Clearing of Events” on page 242, and “Configuring the Deletion of Events” on page 243.

---

- **The clearing of events** - You can determine how often events are marked cleared. For more information, see “Configuring the Clearing of Events” on page 242.
- **The deletion of events** - You can determine how often events are deleted. By default, events are deleted every two weeks by. See “Configuring the Deletion of Events” on page 243.

## 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.

---

**Caution** – Events are removed or marked cleared only when their automatic delay time is completed. See the topics, “Configuring the Clearing of Events” on page 242 and “Configuring the Deletion of Events” on page 243.

---

Assume you set the clearing and deletion frequency to every 15 minutes, with the initial time at 11:00 a.m., so that every 15 minutes the management server checks for events marked for deletion and for clearing, and it updates the user interface accordingly. Now assume Informational events are set to be cleared every hour, and an Informational event occurs at 12:20 p.m. Exactly one hour later, 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 look at the Event Manager at 1:35 p.m., the event is be marked as cleared.

To modify the clearing and deletion frequency:

1. Access the Events page by selecting **Configuration > Events**.
2. Set the date, time, and repeat interval for this task. For more information, see “Setting the Date and Time for Scheduled Tasks” on page 225.

---

**Note** – The higher the interval, the more demand there is on the management server.

---

3. Select the **Enable** option.
4. Click **Save Changes**.

## Configuring the Clearing of Events

Depending on the severity of an event, the management server may mark the event as cleared after 60 minutes. Events designated as Major and Critical are never marked as cleared. You can change the time delay in clearing an event, and you can specify that the management server never marks an event as cleared.

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.

The following table shows the default settings for clearing events:

**TABLE 8-2** 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:

1. Access the Events page by selecting **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 enter the number in the adjacent box:
    - Minutes**
    - Hours**
    - Days**
    - Weeks**
3. Click **Save Changes**.

## Configuring the Deletion of Events

By default, the management server automatically deletes events after two weeks. However, you can specify for each severity level different time periods for deleting events. For example, you can modify the management server to delete events with the Information severity level every two days. You can also modify the management server to never delete events with the Critical severity level.

To change the default time delay to delete an event:

1. Access the Events page by selecting **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 in 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 list, and enter the number in the adjacent box:

**Minutes**

**Hours**

**Days**

**Weeks**

For example, if you want events that are a week old deleted, you enter 1 and then select **Weeks** in the combo box in the Automatic Delete Delay column.

3. Click **Save Changes**.

---

## Managing File Server SRM

For information about configuring File Server SRM, see the “Overview of File Server SRM” chapter in the File Servers Guide.

---

## Managing Backup Collection

This section contains the following topics:

- “Viewing Collectors for Backup Servers” on page 244
- “Scheduling Backup Collection for Backup Managers” on page 245
- “Editing the Schedule of Backup Collection” on page 246
- “Setting the Backup Sessions Retention Period” on page 246
- “Session Monitoring” on page 246
- “Drive Monitoring” on page 247
- “Viewing the Status of Backup Collection” on page 247

## Viewing Collectors for Backup Servers

The management server uses collectors to gather information for Protection Explorer. Protection Explorer provides information about backup protection, such as whether last night’s backup was successful.

To manage collectors for Protection Explorer:

1. Select **Configuration > Protection**.
2. Select one of the collection tabs. The three collection tabs are Image Collection, Sessions Collection, and Media Collection.

The following is displayed on the collection tabs:

- **Backup Manager** - Displays the names of the backup servers.
- **Next Scheduled Run** - Displays the next time the management server is scheduled to obtain image details from the backup server.
- **Interval in Minutes** - Displays how often the management server is scheduled to obtain image details.
- **Running** - A check mark means the collector is running.
- **Edit Schedule** - Lets you modify the collection details
- **Start Collectors** - Click this button to start the collectors.
- **Stop Collectors** - Click this button to stop the collectors.

## Scheduling Backup Collection for Backup Managers

You can configure the management server to obtain information about your master backup servers at a set interval. You can obtain image details, session details, or media details.

Keep in mind the following:

- Make sure these collectors run at least daily, so the latest backup information is displayed in Protection Explorer.
- All collectors are stopped during Get Details. This means that during Get Details, data such as for Protection Explorer is not updated.
- The process of Get Details takes time. It is recommended you perform this procedure when the network is not busy.
- It is not possible to run data collectors on quarantined elements. Attempting to do so will result in exceptions in the `appstorm.<timestamp>.log` file.

To obtain details for a backup server:


1. Select **Configuration > Protection**.
2. Select one of the collection tabs.
3. Select the management servers for which you want to obtain details.
4. Click the **Start Collectors** button.
5. Set the date, time, and repeat interval for this task. For more information, see “Setting the Date and Time for Scheduled Tasks” on page 225.

6. Click **OK**.

The management server gathers information about image, session, or media details.

## Editing the Schedule of Backup Collection

To change when the management server obtains details for a backup server:

1. Select **Configuration >Protection**.
2. Select one of the collection tabs.
3. Click the **Edit** () button corresponding to the Backup Manager you want to modify for its collection of details.
4. Click the **Start Collectors** button.
5. Set the date, time, and repeat interval for this task. For more information, see “Setting the Date and Time for Scheduled Tasks” on page 225.
6. Click **OK**.

The management server gathers information about image, session, or media details.

## Setting the Backup Sessions Retention Period

You can set the retention value for sessions to be stored in the database. To set the backup sessions retention period:

1. Select **Configuration >Protection**.
2. Click the **Retention Configuration** tab.
3. Enter the number of days you would like sessions to be retained. The retention period must be a minimum of 30 days to a maximum of 1098 days.
4. Click **Submit**.

The retention period is set.

## Session Monitoring

To view the current running session details for Backup Manager:

1. Select **Configuration >Protection**.
2. Click the **Session Monitoring** tab.

The following is displayed:

- **Backup Manager** - Displays the names of the backup servers.
- **Next Scheduled Run** - Displays the next time the management server is scheduled to obtain current running session details for the backup server.
- **Interval (Minutes)** - Displays how often the management server is scheduled to obtain current running session details.
- **Running** - A check mark means the collector is running.
- **Edit** - Lets you modify the collection details
- **Start Collectors** - Click this button to start the collectors.
- **Stop Collectors** - Click this button to stop the collectors.

## Drive Monitoring

To view drive monitoring details for Backup Manager:

1. Select **Configuration >Protection**.
2. Click the **Drive Monitoring** tab.

The following is displayed:

- **Backup Manager** - Displays the names of the backup servers.
- **Next Scheduled Run** - Displays the next time the management server is scheduled to obtain drive monitoring details for the backup server.
- **Interval (Minutes)** - Displays how often the management server is scheduled to obtain drive monitoring details.
- **Running** - A check mark means the collector is running.
- **Edit** - Lets you modify the collection details
- **Start Collectors** - Click this button to start the collectors.
- **Stop Collectors** - Click this button to stop the collectors.

## Viewing the Status of Backup Collection

The management server keeps track of the collections it has completed on the discovered Backup Managers on the Status tab. The Status tab displays the time when the collection started and ended on a Backup Manager, in addition to its status.

To access the Status tab:

1. Select **Configuration > Protection**.
  2. Click the **Status** tab.
- 

## Managing Reports

This section contains the following topics:

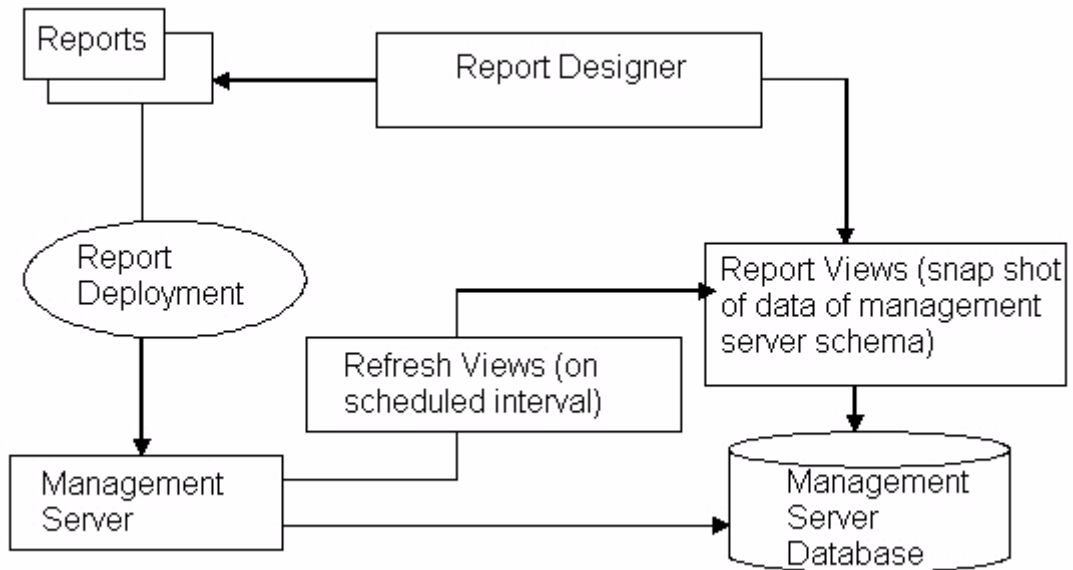
- “Architectural Overview of Report Views and Report Cache Refresh” on page 248
- “Managing Collectors for Reports” on page 252
- “Editing a Collector Schedule” on page 276
- “Viewing Scheduled E-mail Deliveries for Reports” on page 255
- “Editing E-mail Schedules for Reports” on page 256
- “Viewing Data Aging Statistics for Reports” on page 257
- “Scheduling Report Cleanup” on page 258
- “Refreshing the Report Cache” on page 259
- “Setting Up Global Reporter” on page 260
- “Managing Custom Reports (Importing and Deleting)” on page 267

## 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 259. The overall report architecture is displayed in Figure 8-2, “Report Views and Report Cache Refresh,” on page 249.



**FIGURE 8-2** Report Views and Report Cache Refresh

## Suggestion for Scheduling the Report Cache Refresh:

The report cache (report views) is snapshot data of a management server schema up 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 that 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.

For information on how to refresh the report cache, see “Refreshing the Report Cache” on page 259.

Keep in mind the following:

- After an initial installation, 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 be, since they obtain their information from the report cache.

- During the report cache refresh, the report views are truncated and repopulated with the latest data. If a report is run during report cache refresh and the corresponding view is refreshed, the report is displayed 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 the following 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 554. To obtain detailed information of **MVIEWCORE\_STATUS**, see Table 13-52, "MVIEWCORE\_STATUS," on page 586.
- **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 554. To obtain detailed information about the **MVIEW\_STATUS**, see Table 13-53, "MVIEW\_STATUS," on page 586.

To query the **MVIEW\_STATUS** table:

### On Windows:

1. Enter the following at 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 that 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 at 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 that 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 data collectors to gather information for reports. To view a report, you must have its corresponding collector running, and your report cache must be up-to-date. See “Refreshing the Report Cache” on page 259 for details.

To view collectors for reports, select **Configuration > Reports**, and click the **Data Collection** tab.

For each element, the Report Data Collectors page displays all the collectors that the management server uses to gather report data. You can decide whether data should be collected for each element by starting or stopping the collector schedule. If a collector is gathering data when you stop it, that collector will complete the collection. You are stopping all future runs of the collector; not the currently active process. The Enabled column reflects whether you've started your collector schedule or not.

The Report Data Collector page offers a set of filters to help you find your collectors quickly. The filter area is collapsed by default. To expand the filter area, click the **+** symbol. The following filters are supported:

- **Element Type** - Use this filter to see the collectors for a specific element type or for all elements.
- **Collector State** - Use this filter to see all schedules or only the schedules that have been started or stopped.
- **Element Name Contains** - Use this filter to retrieve all the elements whose name contains the specified string.

To apply the filter settings, click **Filter** to refresh the content of the Report Data Collector page. To restore the filters to their default settings, click **Reset**, and refresh the collector page.

In addition to changing the collector schedule after an element has been discovered, you may wish to decide whether a collector schedule should be started or not for future discovery elements. You can access the default collector schedule settings by clicking the Default Collector Settings link above the Filter area, or select **Configuration > Discovery**, and click the **Collector Settings** tab. For more information about the Collector Settings tab, see “Modifying Collector Settings for Newly Discovered Elements” on page 229.

---

**Caution –** All collectors are stopped during Get Details. This means that during Get Details data is not updated.

---



---

**Note –** It is not possible to run data collectors on quarantined elements. Attempting to do so will result in exceptions in the `appstorm.<timestamp>.log` file.


---

The following table provides details about the Data Collection tab:

**TABLE 8-3** About Collectors for Reports

Column Heading	Description
Element	Displays the name of the element from which this collector gathers information.

**TABLE 8-3** About Collectors for Reports (*Continued*)

Column Heading	Description
Element Type	Displays the type of element from which the collector gathers information.
Reports	Lists the reports for which the collector is responsible for providing information.
Enabled	Displays the status of the collector. Collectors that are running display a check mark in this column.
Interval (Minutes)	Displays the interval in minutes between collector runs.
Next Scheduled Run	Displays the date and time when the collector is scheduled to run.
Edit	To edit the schedule for running a collector, click the <b>Edit</b> (  ) button. For more information, see "Editing a Collector Schedule" on page 276.
Action	Displays one of the following buttons: <ul style="list-style-type: none"><li>• <b>Stop</b> - Stops the collector. The corresponding reports display only information gathered previously. See "Stopping Collectors" on page 255.</li><li>• <b>Start</b> - Starts the collector. When you start a collector, it begins gathering information for its corresponding reports. See "Starting Collectors" on page 254.</li></ul>

## Starting Collectors

---

**Caution** – After you click **OK**, if the date and time you set has not passed, the collector starts immediately. If the set time has passed, the collector starts two minutes after you click **OK**.

---

To start a collector:

1. Select **Configuration > Reports**. Then, click the **Data Collection** tab.
2. Click the **Start** button corresponding to the collector you want to start.

To start more than one collector at once, select more than one collector and then click **Start Collectors**.

3. Set the date, time, and repeat interval for this task. For more information, see "Editing a Collector Schedule" on page 276.
4. Click **OK**.

## Stopping Collectors

When you stop a collector, the management server stops gathering the 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. Select **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 **Stop Collectors**.

## 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 531.

To view all e-mail schedules:

1. Select **Configuration > Reports**.
2. Click the **Scheduled Deliveries** tab at the top of the screen.



All e-mail schedules for reports are displayed.

The following table provides details about the Scheduled Deliveries tab.

**TABLE 8-4** Viewing System-Wide E-mail Schedules

Column Name	Description
Recipient	The person who receives the report.
Subject	The subject of the e-mail.


**TABLE 8-4** Viewing System-Wide E-mail Schedules (*Continued*)

Column Name	Description
Format	The format of the report sent: <ul style="list-style-type: none"><li>• PDF</li><li>• Microsoft EXCEL</li><li>• XML</li></ul>
Scheduled By	The user who scheduled the report.
Edit	Click the <b>Edit</b> (  ) button to edit a schedule. For information about the options displayed in this window, see “Editing an E-mail Schedule for a Report” on page 534.
Delete	Click the <b>Delete</b> (  ) 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 <b>Delete</b> button above the table.

## Editing E-mail Schedules for Reports

**Caution** – Schedule your reports to be sent 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 add an e-mail schedule that sends a report daily at 7 a.m. Also, assume you schedule 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 your reports to be sent soon afterwards. See “Scheduling a Report Cache Refresh” on page 260.

To edit an e-mail schedule:

1. Select **Configuration > Reports**.
2. Click **Scheduled Deliveries** at the top of the screen.
3. In the Edit column, click the **Edit** () button.
4. In the To box, change the recipient's e-mail address.
5. In the Subject box, change the subject of the e-mail.



6. In the Message box, 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 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 that the user has an understanding of XML.

8. In the Time to Run box, enter the time you want to send the report in 24-hour format.

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 every day.

**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 a person's replacement while the person is away on vacation.

- **Weekly** - If you selected weekly, use the **Frequency** 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 **OK**.

## 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 the Data Aging page control data aging for both reports and performance statistics.

---

**Caution** – Do not modify the data on the Data Aging page unless instructed by customer support. Changing them incorrectly can adversely affect the management server.

---

To view data aging statistics:


1. Select **Configuration > Reports**.
2. Click the **Data Aging** tab at the top of the screen.

Data aging statistics are displayed in the table.

---

**Caution** – 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. Set the date, time, and repeat interval for this task. For more information, see “Setting the Date and Time for Scheduled Tasks” on page 225.
5. Click the **Enable** option.

---

**Note** – If you are not allowed to disable the collector, the Enable option is unavailable.

---

6. *Garbage Collection only*: To change how long the data is preserved, enter an interval in the Preserve box, and then select one of the following from the list to the right of the box:
  - **Second(s)**
  - **Minutes**
  - **Hours**
  - **Days**
  - **Weeks**
7. Click **OK**.

## 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 temporary 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. Select **Configuration > Reports**.
2. Click the **Report Cleanup** tab at the top of the screen.
3. Set the date, time, and repeat interval for this task. For more information, see “Setting the Date and Time for Scheduled Tasks” on page 225.
4. Click the **Enable** option.
5. Click **Save Changes**.

## Refreshing the Report Cache

The management server gathers information for reports from the database every six hours. 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 259.
- **Schedule the report cache to be refreshed.** - See “Scheduling a Report Cache Refresh” on page 260.

Keep in mind the following:

- If Get Details is occurring, wait for it to finish before clicking **Refresh Now**. 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.

## Refreshing the Report Cache Immediately

To refresh the report cache now:

1. Select **Configuration > Reports**.
2. Click the **Report Cache** tab at the top of the screen.
3. Click **Refresh Now**.

## Scheduling a Report Cache Refresh

When you schedule the refreshing of the report cache, keep in mind that 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.

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 226.

To schedule a report cache refresh:

1. Select **Configuration > Reports**.
2. Click the **Report Cache** tab at the top of the screen.
3. Set the date, time, and repeat interval for this task. For more information, see “Setting the Date and Time for Scheduled Tasks” on page 225.
4. Click the **Enable** option.
5. Click **Save Changes**.

## Setting Up Global Reporter

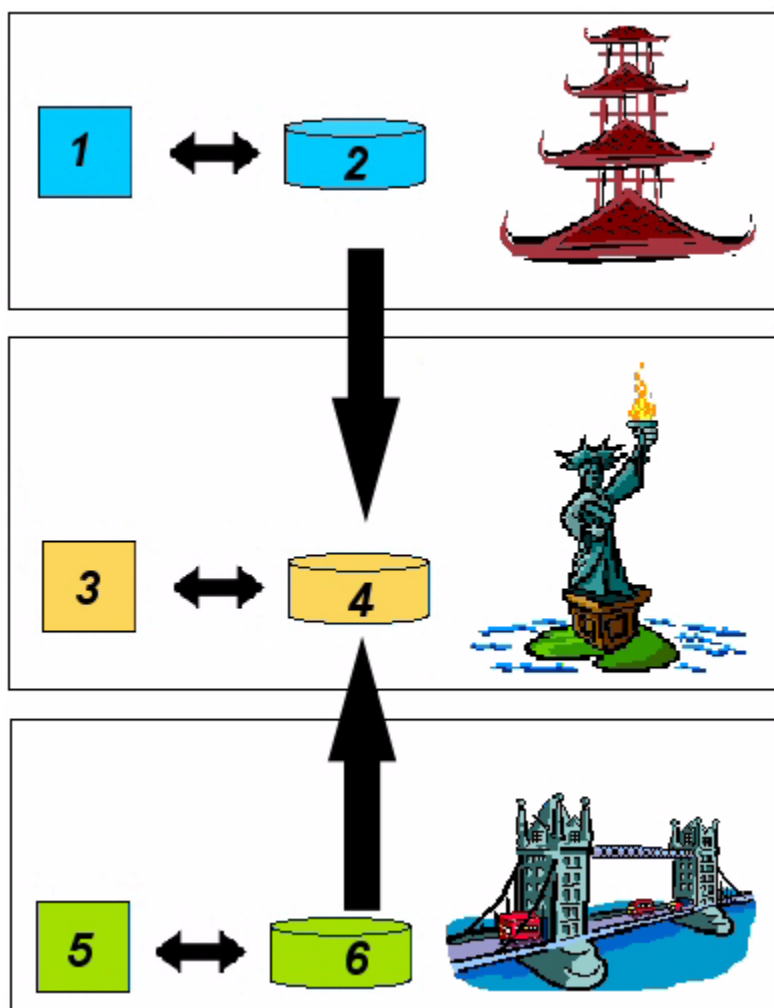
---

**Caution** – 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.

---

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 figure below. 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 8-3** An Example of Global Reporting

**TABLE 8-5** 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.

**TABLE 8-5** Description of “An Example of Global Reporting” (Continued)

Item	Description
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 the 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 to pull 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 that 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 **Delete** () 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. For more information, see “Remote Sites Are Not Removed from the `tnsnames.ora` File” on page 266.
- **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 selecting **Configuration > Reports > Data Collection**.
3. Click the **Global Reporter** tab.

4. Click **New Site**.

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 (optional)** - The Oracle TNS listener port the site uses. If this box 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 that it can contact the site and it checks the build 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 box, enter a time in 24-hour format.
- b. Select the date when you want the job to start.  
The date is selected.
- c. Click **Set**.

If you click **Set** after the time you specified has passed, you must reset the time.

The time and date you selected are displayed in the Next Scheduled Run box.




11. Set up a repeat interval by typing a number in the Repeat Interval box and selecting a unit of measurement. For example, if you want the Global Reporter server to check the views of sites daily, you would enter 1 in the Repeat Interval box and select **day(s)** as a unit.

---


**Caution** – You must select a repeat interval that is 12 hours or more.

---

12. When you are done, click **Save Changes**.
13. To obtain information immediately, click **Refresh Now** (**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 529.

## Editing Remote Site Information

To modify the information for your remote servers:

1. Select **Configuration > Reports > Data Collection**.
2. Click the **Global Reporter** tab.
3. Click **Edit** () button 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 (optional)** - The Oracle TNS listener port the remote server uses. If this box 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 OracleServiceAppIQ 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

**UNIX systems :**

\$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 Custom Reports (Importing and Deleting)

You can manage custom reports using the graphical user interface. This interface enables you to import and delete custom reports and is accessed from the **Reports** tab under the **Configuration** choice. The interface eliminates the need to manually deploy custom reports at customer locations.

To use this feature, click **Configuration** on the main screen. Then, click **Reports**. From the **Reports** choices, click **Manage Custom Reports**.

The following screen displays.

Scheduled Deliveries Data Collection Data Aging Report Cleanup Report Cache Global Reporter **Manage Custom Reports**

To import and view a custom report:

1. Package the report definition, report template and database into a zip file.
2. Use the Browse button below to specify the name/path of the zip file.
3. Use the Import button below to import the zip file.
4. View the report from the "System" link of the Reporter
5. For more information Please Click [Help](#)

Import Reports

Zip file name

Imported Reports

Customer/Report Name	Report File(s)	Delete

**FIGURE 8-4** Manage Custom Reports Screen

The screen lists these instructions to import and view a custom report:

1. Package the report definition, report template, and database into a zip file.
2. Use the **Browse** button in the display to specify the name/path of the zip file.
3. Use the **Import** button in the display to import the zip file.
4. View the report from the "system" link of the Reporter.

The instructions also remind you to refer to **Help** for additional information.

After you package the necessary files into a .zip file, use the **Browse** button to navigate to the desired .zip file. When ready to upload the selected file, click the **Import** button.

After importing, a screen display similar to the following shows the imported report files. The imported reports can then be viewed from the "system" link of the Reporter.

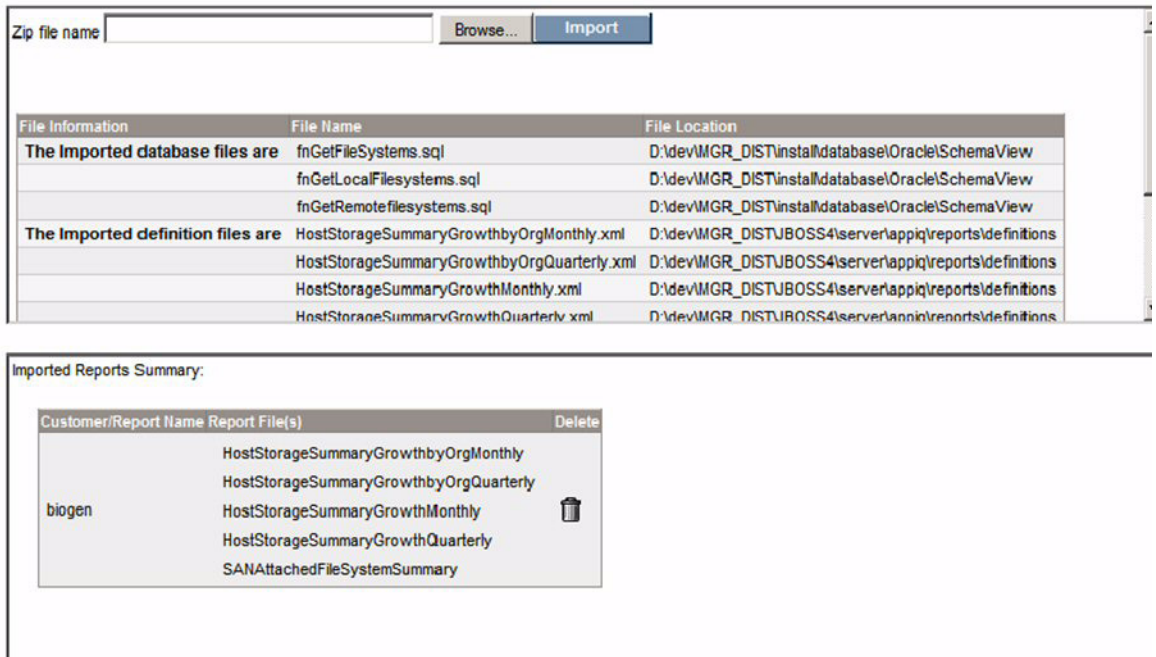


FIGURE 8-5 Screen Displays Custom Reports Available

## Importing Custom Reports Error Messages

If you attempt to load a file type other than .zip, the system will give an error message saying that it is not a valid file type.

If you try to import a report that already exists, the system gives an error message saying the report already exists. In this situation, you must delete the report that already exists, then import the new report.

## Deleting Custom Reports

If you want to delete imported custom reports by clicking **Delete**, the system displays a warning message. If you confirm that you want to delete by clicking **Yes**, the system deletes the reports and supporting files. You can confirm that the files are deleted by navigating to the **Reporter** tab and clicking the link.

---

# Managing Performance Collection

This section contains the following topics:

- “Managing Performance Collectors” on page 269
- “Starting Performance Collectors” on page 271
- “Stopping Performance Collectors” on page 272
- “Setting the Date and Time for Scheduled Tasks” on page 225
- “Viewing Data Aging Statistics for Performance” on page 272

## Managing Performance Collectors

The management server uses performance collectors to gather information for Capacity Explorer and Performance Explorer charts, as well as for monitoring.

To manage performance collectors, select **Configuration > Performance**, and click the Data Collection tab.

For each element, the Performance Data Collectors page displays all the collectors that the management server uses to gather data. You can decide whether data should be collected for each element by starting or stopping the collector schedule. If a collector is gathering data when you stop it, that collector will complete the collection. You are stopping all future runs of the collector; not the currently active process. The Enabled column reflects whether you’ve started your collector schedule or not.

If you are not sure whether you want to start or stop a collector schedule, click **+Statistics** to see the list of performance statistics that the collector is designed to collect. Once the statistics are collected, you can use Performance Explorer to review the data.

The Performance Data Collector page offers a set of filters to help you find your collectors quickly. The filter area is collapsed by default. To expand the filter area, click the **+** symbol. The following filters are supported:

- **Element Type** – Use this filter to see the collectors for a specific element type or for all elements. When you select a specific element type, the Collector Type filter is adjusted to show you the collector types that are meaningful to the selected element type.
- **Collector State** – Use this filter to see all schedules or only the schedules that have been started (Enabled) or stopped (Disabled).
- **Collector Type** – Use this filter to see all the collectors for a specific collector type or all collector types. This filter works in conjunction with the Element Type filter.

- **Element Name Contains** – Use this filter to retrieve all the elements whose name contains the specified string.

To apply the filter settings, click **Filter** to refresh the content of the Report Data Collector page. To restore the filters to their default settings, click **Reset**, and refresh the collector page.

In addition to changing the collector schedule after an element has been discovered, you may wish to decide whether a collector schedule should be started or not for future discovery elements. You can see the default collector schedule settings using the Default Collector Settings link above the Filter area, or select **Configuration > Discovery**, and click the **Collector Settings** tab. For more information about the Collector Settings tab, see “Modifying Collector Settings for Newly Discovered Elements” on page 229.

---

**Caution** – All collectors are stopped during Get Details. This means that during Get Details, data 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 do so will result in exceptions in the `appstorm.<timestamp>.log` file.

---

The following table provides details about the Data Collection tab:

**TABLE 8-6** About Performance Collectors

Column Heading	Description
Element	Displays the name of the element from which this collector gathers information.
Element Type	Displays the type of element from which the collector gathers information.
Collector Type	Displays the collector type. Click <b>+Statistics</b> to see the list of performance statistics that the collector is designed to collect
Enabled	Displays the status of the collector. Collectors that are running display a check mark in this column.
Interval (Minutes)	Displays the interval in minutes between collector runs.
Next Scheduled Run	Displays the date and time when the collector is scheduled to run.
Edit	To edit the schedule for running a collector, click the <b>Edit</b>  button. Then set the date and time. For more information, see “Editing a Collector Schedule” on page 276.

**TABLE 8-6** About Performance Collectors (*Continued*)

Column Heading	Description
Action	<p>Displays one of the following buttons:</p> <ul style="list-style-type: none"> <li>• <b>Stop</b> - Stops the collector. The corresponding reports display only information gathered previously. See “Stopping Performance Collectors” on page 272.</li> <li>• <b>Start</b> - Starts the collector. When you start a collector, it begins gathering information for its corresponding reports. See “Starting Performance Collectors” on page 271.</li> </ul>

## Starting Performance Collectors

To start a collector:

1. Access the page for performance collectors (**Configuration > Performance > Data Collection**).
2. Click the **Start** button corresponding to the collector you want to start.  
To start more than one collector at once, select more than one collector on the **Data Collection** tab, and then click **Start Collectors**.
3. Set the date, time, and repeat interval for this task. For more information, see “Editing a Collector Schedule” on page 276.
4. If you are asked to provide a proxy host, do the following:
  - a. Click **Browse**.
  - b. Select a proxy host from the 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 that you have the latest CIM extension version installed and running on a host that can access the LSI storage system.
5. Click **OK**.



# 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 **Stop Collectors**.

# 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 the Data Aging page control data aging for both reports and performance statistics.

---

**Caution** – Do not modify the data on the Data Aging page unless instructed by customer support. Changing them incorrectly can adversely affect the management server.

---


To view data aging statistics:

1. Select **Configuration > Performance > Data Collection**.
2. Click the **Data Aging** tab at the top of the screen.

---

**Caution** – 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. Set the date, time, and repeat interval for this task. For more information, see “Setting the Date and Time for Scheduled Tasks” on page 225.
5. Click the **Enable** option.

---

**Note** – If you are not allowed to disable the collector, the Enable option is unavailable.

---

6. *Garbage Collection only*: To change how long the data is preserved, enter an interval in the Preserve box, and then select one of the following from the list to the right of the box:
  - **Second(s)**
  - **Minute(s)**
  - **Hour(s)**
  - **Day(s)**
  - **Week(s)**
7. Click **OK**.

---

## 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 8-7**    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. Select **Configuration > Locales**.
2. Choose a locale from the **Select Locale** menu.
3. Choose a currency from the **Select Currency** menu.

See Table 8-7, “Currency Settings,” on page 274 if you do not know the corresponding names for the currency abbreviations.

4. Click **Change Locale**.
5. Restart the management server.

---

**Caution** – You must restart the management server for your changes to take effect.

---

---

# Process Names

This section describes the process names on Windows and Unix systems.

## Process names on Windows

The following process names are displayed in the Windows Task Manager on the management server:

**TABLE 8-8** 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.

## Process Names on Unix Systems

The following process names are displayed in UNIX systems by running the `ps -ef` command:

**TABLE 8-9** Process Names on UNIX systems

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`.

---

---

# Editing a Collector Schedule

There are several ways to modify a collector schedule:

- Locate the collector you want to edit in the collector table, and click the **Edit** () button in that collector's row.
- Locate the collector you want to edit in the collector table, and select the collector using the checkbox in the first column of the table. You can select one or more collectors this way. Click the **Edit Selected** button above the collector table to modify the selected collector's schedules at the same time.
- If you want to modify all of the collectors displayed on the same page of the collector table, click the checkbox in the first column header, and then click the **Edit Selected** button above the table.
- If you want to modify all of the collectors displayed on all of the pages, click the **Select All Pages** link above the collector table, and then click the **Edit Selected** button above the table.

When you edit a collector schedule, you can change the following schedule properties:

- **Next Scheduled Run** – Specifies a start time for when the collector should run.
- **Repeat Interval** – Specifies how often the collector should run on a recurring basis.
- **Spread Start Time** – This checkbox can only be selected when more than one collector is being edited. If you select this checkbox, the management server automatically adjusts the start time of the selected collectors and spreads the load of the collectors across the specified repeat interval.



# Database Maintenance and Management

---

This chapter contains information about backing up and restoring the database using the Database Admin Utility.

This chapter contains the following topics:

- “Database Maintenance Window” on page 279
- “Overview of Backups” on page 281
- “Database Mode” on page 282
- “Architectural Overview of RMAN Backups” on page 284
- “Performing an RMAN Hot Backup” on page 286
- “Scheduling RMAN Hot Backups” on page 287
- “Viewing Results from RMAN Backup” on page 288
- “About the Database Admin Utility” on page 289
- “About Database Passwords” on page 302
- “Troubleshooting Listener and Database Connection Problems” on page 305

---

## 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 “Restarting the Service for the Management Server” on page 18.
2. Access the Database Admin Utility. See “Accessing the Database Admin Utility” on page 289 for more information.
3. Using the Database Admin Utility, verify that the database is in an open state and the listener is running. See “Checking the Database and Listener Status” on page 291.

If the database state is not OPEN and shows an 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:

**On Microsoft Windows:** \oracle\admin\APPIQ\bdump

**On UNIX systems:** \$ORACLE\_BASE/admin/APPIQ/bdump

4. Reset the temporary tablespace, as described in “Resetting the Temp and Undo Tablespace” on page 295.
5. Export the database, as described in “Exporting the Database” on page 292.
6. If the database is running in archive mode, set the database to no-archive mode, as described in “Changing the Archive Mode” on page 296. Then, clean the \oracle\oradata\APPIQ\archive on Microsoft Windows and \$ORACLE\_HOME/oradata/APPIQ/archive on UNIX systems.

---

**Note** – The archive directory (\oracle\oradata\APPIQ\archive on Microsoft Windows and \$ORACLE\_HOME/oradata/APPIQ/archive on UNIX systems) only exists if you have previously set the management server to archive mode.

---

7. Return the database to archive mode, as described in “Changing the Archive Mode” on page 296.
8. Start the service for the management server.
9. If the 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 286 and “Architectural Overview of RMAN Backups” on page 284.

On successful completion of RMAN backup, the backup is saved to the following directory:

- **On Microsoft Windows:**

%ORA\_HOME%\rman\current

- **On UNIX systems:**

\$ORACLE\_HOME/rman/current

10. Clean the following folders:

- **On Microsoft Windows:**

%ORA\_HOME%\rman\backup1

%ORA\_HOME%\rman\backup2

- **On UNIX systems:**



\$ORACLE\_HOME/rman/backup1  
\$ORACLE\_HOME/rman/backup2

**TABLE 9-1** Backup Directories in %ORA\_HOME%\rman

Directory	Contains
current	The backup when the <b>Database Server Backup</b> 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).

## Overview of Backups

The management server provides the following backups:

---

**Caution** – Export and RMAN backups should be done regularly and in combination.

---

**TABLE 9-2** Description of Backups

Backup Type	Description	Files Backed Up	Database Mode
<b>Export backup</b>	Done through the Database Admin Utility. See “Exporting the Database” on page 292 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

**TABLE 9-2** Description of Backups (*Continued*)

Backup Type	Description	Files Backed Up	Database Mode
<b>RMAN HOT backup</b>	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” on page 287.	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 “Changing the Archive Mode” on page 296.
<b>RMAN Cold backup</b>	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 “Changing the Archive Mode” on page 296.

**Backup Destination and Operation for RMAN:** The management server has three backup points available for RMAN backups:

■ **On Microsoft Windows:**

```
%ORA_HOME%\rman\current
%ORA_HOME%\rman\backup1
%ORA_HOME%\rman\backup2
```

■ **On UNIX systems:**

```
$ORACLE_HOME/rman/current
$ORACLE_HOME/rman/backup1
$ORACLE_HOME/rman/backup2
```

The scheduled backup writes in the backup1 and backup2 folders, in rotation. The Backup Now backup from the management server keeps overwriting in the current folder. See “Architectural Overview of RMAN Backups” on page 284 for more information about RMAN backups.

## Database Mode

The database can be set to archive mode or no-archive mode. This section contains the following topics:

- “Archive Mode” on page 282
- “No-Archive Mode” on page 283

## Archive Mode

To facilitate the HOT RMAN backup, you must change the database mode to archive mode, as described in “Changing the Archive Mode” on page 296. The default database archive destination is

`\oracle\oradata\APPIQ\archive` on Microsoft Windows and  
`$ORACLE_HOME/oradata/APPIQ/archive` on UNIX systems. This destination can be modified as described in “Changing the Archive Destination” on page 299.

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.

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 286.

## No-Archive Mode

You can change the database mode to run in no-archive mode with RMAN backup disabled, as described in “Changing the Archive Mode” on page 296. 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.

---

**Caution** – Export the database after switching to no-archive mode. See “Exporting the Database” on page 292 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 **no-archive mode**, 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 298. Export the database after you change the database mode to no-archive. See “Exporting the Database” on page 292.
- Scheduled RMAN backup sessions 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.

---

# 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 “Changing the Archive Mode” on page 296, the management server backs up the Oracle instance for the management server every three days and saves the backup for two weeks. When the management server first performs a scheduled backup, the backup is saved in %ORA\_HOME%\rman\backup1 on Microsoft Windows and \$ORACLE\_HOME/rman/backup1 on UNIX systems. The next time the management server performs a scheduled backup, it is saved in %ORA\_HOME%\rman\backup2 on Microsoft Windows and as \$ORACLE\_HOME/rman/backup2 on UNIX systems. 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” on page 287.

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 %ORA\_HOME%\rman\current on Microsoft Windows and in \$ORACLE\_HOME/rman/current on UNIX systems. To recover the database, contact customer support. See Table 9-1, “Backup Directories in %ORA\_HOME%\rman,” on page 281.

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. If you have a scheduled backup on day 7, it is saved to the backup1 folder. This backup replaces the backup from day 1. The available backups are now from day 4 and 7. If 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 Table 9-3, “Sample Backup Example,” on page 284:

**TABLE 9-3** Sample Backup Example

Day	Backup Type	Backup 1	Backup 2	Current	Available Backup
Day 1	Scheduled	Day 1 backup	-----	-----	Day 1 Backup

**TABLE 9-3** Sample Backup Example (*Continued*)

Day	Backup Type	Backup 1	Backup 2	Current	Available 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
Day 8	Backup 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 UNIX systems).
- 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.

## Data Saved During a Backup

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 `appiq_jboss.properties`**
- **`rmanbackup.log` file**
- **`spfileappiq.ora` file**

## Backing up the Database Manually

To back up the database manually:

1. Before you can back up the database manually, enable the database archive mode and RMAN backup as described in “Changing the Archive Mode” on page 296.

2. In the management server, click **Configuration > Product Health**.
3. Select **RMAN Backup** in the Product Health tree.
4. Click **Backup Now**.

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
- File Server SRM.

---

**Note** – The buttons on the RMAN backup page appear disabled when the database archive mode is disabled. See “Changing the Archive Mode” on page 296 for more information about changing the archive mode.

---

To perform an RMAN hot backup:

1. Verify that you have enabled the database archive mode and RMAN backup as described in “Changing the Archive Mode” on page 296.
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 updated with the status of the RMAN hot backup and displays the status of the previous RMAN hot backups (manual and scheduled).

---

# Scheduling RMAN Hot Backups

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 “hot” because the management server is still running.

---

**Caution** – Before you can schedule the RMAN backup, you must enable the database archive mode and RMAN backup as described in “Changing the Archive Mode” on page 296. 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 289.

1. Verify that you have enabled database archive mode and RMAN backup as described in “Changing the Archive Mode” on page 296.
2. Click **Configuration > Product Health**.
3. Select **RMAN Backup** in the Product Health tree.

4. Click the calendar icon .

5. In the Time box, take the following actions:

- a. Enter the time in 24-hour format.
- b. Click the date on which you want to run the next backup of the database. Today’s date is highlighted in pink.
- c. Click **Set**.

The date and time appear in the Next Scheduled Run box in the yyyy-mm-dd format.

If you change the date in the box 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 box, enter an interval. Select one of the following units from the menu:
  - **Second(s)**
  - **Minute(s)**
  - **Hour(s)**
  - **Day(s)**

- **Week(s)**

---

**Note** – The minimum interval you can schedule is one day. Whether 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 or more.

---

7. Click the **Enable** option.

8. Click **Save Schedule**.

---

**Note** – 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 backup. The status of the previous RMAN hot backups (manual and scheduled) is displayed in the Results tab.

---

## Viewing Results from RMAN Backup

The results of an RMAN backup can be determined by checking the Results tab for the date, time, and status of the backup.

To view the results of 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:

- Date/Time of the backup
- Status of the backup
- Backup folder



---

# About the Database Admin Utility

The Database Admin Utility allows you to manage your database, from restoring it from a cold backup to resetting the temp tablespace. The tool provides flexible importing and exporting features, which let you save time.

To learn more about the Database Admin Utility, see the following topics:

- “Accessing the Database Admin Utility” on page 289
- “Refreshing the Database Admin Utility” on page 290
- “Checking the Database and Listener Status” on page 291
- “Use Only the Database Admin Utility to Change the Password of System Accounts” on page 291
- “Exporting the Database” on page 292
- “Importing the Database” on page 293
- “Re-initializing the Database” on page 294
- “Resetting the Temp and Undo Tablespace” on page 295
- “Restarting the Database” on page 295
- “Restoring a Cold Backup” on page 296
- “Changing the Archive Mode” on page 296
- “Restore RMAN Backup” on page 297
- “Running a Cold Backup” on page 298
- “Changing the Archive Destination” on page 299
- “Downloading Log Files” on page 299
- “Changing the Oracle Listener Password” on page 300
- “Resetting the Admin Password for the Management Server” on page 301
- “Resetting/Clearing the Database Admin Log File” on page 300
- “Warning Messages During Reinitializing the Database” on page 302
- “Generating a Support Database” on page 303
- “About Importing a Customer Support Database” on page 304
- “Troubleshooting Listener and Database Connection Problems” on page 305

## Accessing the Database Admin Utility

To access the Database Admin Utility:

1. Stop the AppStorManager service.
2. Access the database utility by doing the following on the management server:
  - On Solaris:
    - 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.

---

**Caution** – 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
```

- On Linux:

- 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:

```
. /opt/productname/install/setvars.sh
```

where /opt/productname is the directory containing the software. It is defined by \$APPIQ\_DIST.

- 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
```

- On Windows: Go to the %MGR\_DIST%\Tools\dbAdmin directory and double-click dbAdmin.bat.

## Refreshing the Database Admin Utility

The Database Admin Utility requires that JBoss and CIMOM not be in use when it is operating. If either of these components is in use, the Database Admin Utility will not work. If you are shown an error message when you start the utility, stop the AppStorManager service, and 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 289.
2. Click **Check Database Status** in the left pane.
3. Enter the SYS password in the Password for User SYS box.
4. Click **Check Database Status**.

The database and listener status is shown in the lower pane.

5. To clear the pane, click **Clear All**.

---

**Note** – The Check database status option is available even when the management server is running.

---

## Use Only the Database Admin Utility to Change 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 being sure to use only the Database Admin Utility to make the changes.

---

**Caution** – Do not use the Oracle tools to change the passwords.

---

- **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 289.
2. Click **Change Passwords** in the left pane.
3. Select an account name from the **User Name** box.
4. Type the current password in the **Old Password** box.
5. Type the new password in the **New Password** box.
6. Retype the password in the **Confirm Password** box.
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.

---

**Caution** – Do not use the export database feature for backing up the server. If you are interested in backing up the database for disaster recovery, see “Running a Cold Backup” on page 298.

---

To export the database:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 289.
2. Click **Export Database** in the left pane.

3. Click **Browse** to select a file path, enter a file name in the **File name** box, and click **Open**.

---

**Caution** – Select a directory outside of the directory tree of the management server. Then if you remove the management server, you will not lose the saved database.

---

The file name with its path is displayed in the Database Admin Utility. The .zip file extension is automatically added to the file name.

4. Enter the password of the SYSTEM account. The default password of the SYSTEM account is `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, the report cache will be empty until it is refreshed (**Configuration > Reports > Report Cache**). This option may save you time with exporting the database if your database includes a large amount of report data.

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 289.
2. Click **Import Database** in the left pane.
3. Click **Browse**, select a database file with a .zip extension to import, and click **Open**.

The file name is displayed in the Database Admin Utility.

4. Enter the password of the SYSTEM account, and click **OK**. The default password of the SYSTEM account is `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 when the database has a large amount of data for reports. You can instead refresh the report cache from the management server (**Configuration > Reports > Report Cache**).
6. (For non-production systems only) Select **Include Management DB** if you want to include management data for the management server host. This option, which is for support purposes only, provides data on the health of the product so you can determine if there is an issue with the management server itself.

---

**Caution** – Do not use this feature on production systems. Selecting this option affects product health reporting.

---

7. Click the **Import Database** button.

## Re-initializing the Database

---

**Caution** – 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 292 for more information on how to save the database.

---

Keep in mind the following:

- 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 pop up when the Database Admin Utility runs. See “Warning Messages During Reinitializing the Database” on page 302 for more information.

To re-initialize the database:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 289.
2. Click **Re-initialize Database** in the left pane.
3. Enter the password of the SYSTEM account, and click **OK**. The default password of the SYSTEM account is `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 289.
2. Click **Reset Temp Tablespace** in the left pane.
3. Enter the SYS password in the SYS Password box.
4. Click **Reset Temp Tablespace**.

## Restarting the Database

You may sometimes need to restart the AppIQ instance of the database. Use this feature in the Database Admin Utility when the database is down or when you need to shutdown and restart the database.

To restart the database:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 289.
2. Click **Restart Database Server** in the left pane.
3. Enter the SYS password in the SYS Password box.
4. Click **Restart Database Server**.

## Clearing Archives

Archive files can require large amounts of disk space and should be deleted periodically. Use the Clear Archive option to manage the disk space used by the Oracle database for the management server. Follow these steps:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 289.
2. Enter the password for the SYS user account and click **Clear Archives** in the left pane.

The archives are deleted.

## 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**. The backup is referred to as being “cold” because the management server is not running while the backup is occurring. For information about changing the archive mode, see “Changing the Archive Mode” on page 296.

To restore a cold backup:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 289.
2. Click **Restore Cold Backup** in the left pane.
3. Click **Browse**.
4. In the File Name box, provide the directory path containing the cold backup (which 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 Database Admin Utility automatically notices the backup files in the directory provided.
5. Enter the password of the SYS account.
6. Click **Run Cold Backup**.

## Changing the Archive Mode

By default the management server database runs in no-archive mode, which requires you to backup the database manually 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, change the database to archive mode. An RMAN hot backup is done with the database running while the backup is occurring.

If you decide to leave the database in no-archive mode, see “No-Archive Mode” on page 283 for additional important information.

To change the archive mode:

1. Access the Database Admin Utility, as described in “Accessing the Database Admin Utility” on page 289.
2. Click **Change Archive Mode** in the left pane.

A text message indicating the current status of the archive mode is displayed.



3. Do one of the following:

- The **Enable RMAN Backup** button is displayed if the archive mode is currently disabled. Select this option if you plan to run automated backups while the management server is running. If the database is already in archive mode, you can use this option to clean up archive files. See “About the Database Admin Utility” on page 289 for more information about the automatic backups.

Changing the database to archive mode resets the logs, sets the archive parameter in the database parameter file, and enables the RMAN backup scheduler. After switching to archive mode, take a current RMAN backup (**Configuration > Product Health > Disk Space > RMAN Backup > Backup Now**) as described in “Performing an RMAN Hot Backup” on page 286.

- The **Disable RMAN backup** button is displayed if the archive mode is currently enabled. Select this option if and only if you always shut down the management server prior to a backup. With the database running in no-archive mode, you can only cold back up your database. Changing the database to no-archive mode resets the logs SCN (System Change Number), sets the archiver parameter in the database parameter file, and disables the RMAN backup scheduler. See “Running a Cold Backup” on page 298 for more information. Be sure to export the database after switching to no-archive mode. See “Exporting the Database” on page 292.

4. Enter the SYS password in the SYS Password box.

5. Click **Change Settings**.

## Restore RMAN Backup

The Database Admin Utility lets you restore the management server database from a previously scheduled (hot) RMAN backup, which is 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 289.
2. Click **Restore RMAN Backup** 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 289 for more information about the automatic backup.

4. Type the SYS password in the **SYS Password** box.
5. Click **Restore RMAN Backup**.

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 during the backup.

---

**Note** – The cold backup does not run if archive mode, which runs an RMAN backup periodically, is enabled. See “About the Database Admin Utility” on page 289 and “Changing the Archive Mode” on page 296 for more information.

---

The following data 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**

To run a cold backup:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 289.
2. Click **Run Cold Backup** in the left pane.
3. Click **Browse** to select a file path.
4. In the File Name box, provide a directory path (may automatically be populated in some Web browsers), but do not provide a file name.

---

**Caution** – 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. Enter the password of the SYS account.
6. Click **Run Cold Backup**.

## Changing the Archive Destination

The default archive directory is `\oracle\oradata\APPIQ\archive` on Microsoft Windows and `$ORACLE_HOME/oradata/APPIQ/archive` on UNIX systems. 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 289.
2. Click **Set Archive Destination** in the left pane.
3. Click **Browse** to select a file path.
4. Enter the password of the SYS account.
5. Click **Set Archive Destination**.

## Downloading Log Files

Use the Download Logs option to view log files generated by the management server. Download Logs saves all the log files in a zip file, which is 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` file is only generated when you run reports.

---

Follow these steps to download all of the log files for the management server.

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 289.
2. In the Directory box at the top of the screen, enter the path to an existing directory or browse to a directory using the Browse button.

3. Click **Download Logs**. The log files are saved to a zip file which is copied to the directory you specified

---

**Note** – The Download Logs option is available when the management server is running.

---

## Viewing the Database Admin Utility Log File

You can view the Database Admin Utility log file in a separate window following these steps:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 289.
2. Click the **View Log** button at the bottom of the Database Admin Utility window.  
The logs are displayed in a separate window.

## Resetting/Clearing the Database Admin Log File

When you click the **View Log** button at the bottom of the Database Admin window, the Database Admin Utility log file is displayed. You can reset/clear or refresh the Database Admin log information by doing the following:

- To refresh the log information, click the **Refresh** button.
- To clear this window, click the **Reset Log** button at the bottom of the Database Admin Utility window and click **Yes** to confirm that you want to refresh the log information.

## 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 289.
2. Click **Set Oracle Listener Password** in the left pane.

3. Enter the new Oracle listener password in the **New Listener Password** box in the right pane.  
Only the following characters are allowed in the password: letters (a – z, A – Z), numbers (0 – 9) and underscores(\_).
4. Enter the new Oracle listener password again in the Confirm New Listener Password box.  
If you want to clear the New Listener Password and Confirm New Listener Password boxes, click **Clear All**.
5. Click **Set Oracle Listener Password**.  
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**.

---

**Caution** – 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 186 for more information.

---

To reset the admin password:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 289.
2. Click **Reset Product Admin Password** in the left pane.
3. Enter the password for the SYS user account in the **Password for User SYS** box.  
The default password for the SYS user account is `change_on_install`.

---

**Note** – If you want to clear the **Password for User SYS** box, click the **Clear All** button.

---

4. Click the **Reset Product Admin Password** button.  
The admin password is reset to **password**.

# Warning Messages During Reinitializing the Database

When you use the Database Admin Utility to re-initialize the database, warning messages similar to the following appear in the command prompt window that pops up when the utility runs. 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.
```

## About Database Passwords

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 using only the Database Admin Utility (do not use the command line or the Oracle tools to change the database password).

**TABLE 9-4** 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 “Use Only the Database Admin Utility to Change the Password of System Accounts” on page 291.

---

**Caution** – Do not change any database password manually or from the command line. If a password is changed manually, the Database Admin Utility will not work, and export, import, and RMAN backup will fail. See “Use Only the Database Admin Utility to Change the Password of System Accounts” on page 291 for information about changing database passwords.

---

**How to change the Database Password for the APPIQ\_USER account on the Managed Database Application:**

To discover and manage the database application from the management server, the 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 156.

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 156.

## Generating a Support Database

The Database Admin Utility allows you to generate a support database that can be used during support calls. This support database is only useful to customer support representatives for the management server.

---

**Caution** – Never import a support database to the management server. The Support Database is only intended for use by Customer Support.

---

Follow these steps to generate a support database for customer support use:

1. Access the Database Admin Utility as described in “Accessing the Database Admin Utility” on page 289.
2. Click **Generate Support Database** in the left pane.
3. Click Browse, select a path, and enter a file name in the File name box. Select a directory other than the directory where the management server resides. This way if you remove the management server, you will not lose the saved database.
4. Click the Open button. The file name and path is displayed in the Database Admin Utility. The .zip file extension is automatically added to the file name.
5. Type the password of the SYSTEM account. The default password of the SYSTEM account is: manager.

6. Select **Clear Report Cache** if you do not want the report cache to be included with the support database. When a customer support representative imports the support database, the report cache will be empty until the report cache is refreshed. Generating the support database is much faster if you omit the report data.
7. Optional. Select the **Include File Server Data** check box if you want to capture the File Server data. File Server SRM data files are excluded by default in the support zip file because this data might include confidential directory and file names.
8. Click **Generate Support Database**. The support database is created based on your selections. See “About Importing a Customer Support Database” on page 304 for more information.

## About Importing a Customer Support Database

---

**Caution** – Do not import a support database. This functionality is for only customer support.

---

The Import Support Database functionality is for customer support use only. Importing a support database to your management server, results in the following:

- Only the admin user will be able to log into the management server, no other user will be able to log in to the system. To enable other users to log in, the admin user would need to re-enter the passwords for all other users.
- The passwords for all discovered elements in the database — hosts, switches, and storage systems would become invalid and discovery would fail for all elements.

---

**Caution** – The **Include Management DB** option, which is for support purposes only, provides data on the health of the product so you can determine if there is an issue with the management server itself. Do not use this feature on production systems. Selecting this option affects product health reporting.

---

Use **Export Database** to generate a valid database backup for use on your management server. See “Exporting the Database” on page 292.



# Troubleshooting Listener and Database Connection Problems

If there is a problem connecting to the Oracle database from the management server or you see a JBoss connection problem in `appstorm.<timestamp>.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
```

### ■ UNIX systems:

```
$ORACLE_HOME/network/admin/listener.ora
```

To verify if the listener and database are running:

### ■ Windows:

- From the command line, enter:

```
lsnrctl status
```

This shows you the status of listener. If the listener is not running, enter:

```
lsnrctl start
```

- From the Database Admin Utility. See “Checking the Database and Listener Status” on page 291.

### ■ On UNIX systems:

- From the command line, first login as an Oracle user by entering `su - oracle`. Then, enter:

```
lsnrctl status
```

This shows you the status of listener. If the listener is not running, enter:

```
lsnrctl start
```

- From the Database Admin Utility. See “Checking the Database and Listener Status” on page 291.

## Checking Database Status:

These steps should restart the database. If you receive an error message, contact technical support.

### ■ On Windows:

- From the command line, enter the following commands:

```
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 291.
- **On UNIX systems:**
  - From the command line, first login as an Oracle user by entering `su - oracle`. Then, enter the following commands:

```
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 291.

### **Generating a Support Database for Customer Support:**

When contacting Customer Support, you may be asked to generate support database. See “Generating a Support Database” on page 303.

## Viewing Element Topology and Properties

---

This chapter contains the following topics:

- “About System Explorer” on page 307
- “Accessing System Explorer” on page 310
- “About Cisco Switches and VSANs in System Explorer” on page 310
- “About the User Interface for System Explorer” on page 311
- “Viewing Storage Elements” on page 332
- “Setting Up Custom Commands” on page 358
- “Using External Tools” on page 369
- “About the Navigation Tab” on page 370
- “Viewing Element Properties” on page 375
- “Viewing Element Topology” on page 379
- “Creating a Virtual Application” on page 390
- “The Provisioning Tab” on page 391
- “About the Events Tab” on page 392
- “Asset Attributes of an Element” on page 393
- “About the Collectors Tab” on page 395
- “About the Monitoring Tab” on page 396
- “About the Policies Tab” on page 396
- “Determining If a Host Belongs to a File System” on page 397
- “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.

Another example is the display of Inter-Switch Link (ISL) trunking for supported Brocade switches. The topology screens, as well as other related displays, show the ISLs between switches and indicate the total number of ISLs and how many of them are trunked (for supported switches). For example, 6(3 trunked) means 6 is the total number and 3 is how many of them are trunked. ISL trunking information for supported switches is also provided by switch port Properties, switch port Detail table on the Navigation page, and by Reporter in various predefined reports, .

---

**Note** – To view direct-attached storage, you must enable the  button. See Table 10-1, “Feature of the Toolbar in System Explorer,” on page 312 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 “The Toolbar in System Explorer” on page 312 for more information.

The following tabs, located in the middle pane, provide additional information:

- **List** - Provides information about the elements by fabric and domain. See “The List Tab” on page 315 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 “The Access Tab” on page 317 for more information.
- **Path** - Provides information about an element's path. See “About the Path Tab” on page 324 for more information.

When you right-click an element in the topology or in the List, Access, or Path tab, a menu appears. This menu provides additional functionality, depending on the type of element clicked, such as telnet or the creation of zone sets. See the topics, “About the Right-Click Menu Options” on page 326.

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 that a connection exists. For HP blade servers, loop based connections are represented by a solid line with the word LOOP next to it, and NPIV based connections are represented by a solid line with the word NPIV.

Keep in mind the following:

- If your Java plug-in control panel cache 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 310 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.
- On the Switch Navigation Ports page, a Brocade switch shows an L-Port as an FL-Port.
- The following IBM HBAs appear 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 370 for more information.
- **Properties** - The Properties tab provides a detailed status of the element. See “Viewing Element Properties” on page 375 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 379 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 393 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 395 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, “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 605 and “About the Events Tab” on page 392 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 681 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 12 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. (Double-click the switch in System Explorer, and click 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 for System Explorer

This section contains the following topics:

- “About the User Interface” on page 311
- “The Toolbar in System Explorer” on page 312
- “Icons Displayed in the Topology” on page 314
- “The List Tab” on page 315
- “The Access Tab” on page 317
- “About the Path Tab” on page 324
- “About the Right-Click Menu Options” on page 326

## About the User Interface





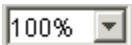


System Explorer displays an easy-to-use interface, which provides the following:

- **Toolbar** - Provides buttons and menus to help you modify the topology in System Explorer. See “The Toolbar in System Explorer” on page 312.
- **Tabs** - Provide information about individual elements. The following tabs are provided:
  - **List** - Provides information about the elements by type and by fabric and domain. See “The List Tab” on page 315.
  - **Access** - Provides access to tools that let you provision and view information about provisioning. See “The Access Tab” on page 317. Provisioning may not be available in your build of the product. To determine if you can access the Provisioning feature, access the List of Features, which is accessible from the Documentation Center (**Help > Documentation Center**).
  - **Path** - Provides information about an element's path. See “About the Path Tab” on page 324.
- **Right-click menu** - Provides features you can use to manage that element. See “About the Right-Click Menu Options” on page 326.

## The Toolbar in System Explorer












Table 10-1, “Feature of the Toolbar in System Explorer,” on page 312 provides a brief description of the buttons and menus in the System Explorer toolbar.

**TABLE 10-1** Feature of the Toolbar in System Explorer



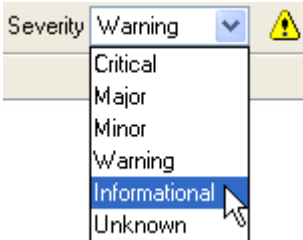


Button	Description
	Prints the topology. See “Printing the Topology” on page 338.
	Exports the topology to an XML file that can be viewed in Microsoft Visio. See “Exporting the Topology to Microsoft Visio” on page 340.
	Magnifies the view
	Decreases the magnification
	Lets you set the magnification to 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. See “Using the Global View” on page 337.
	Fits the topology to the window, so you can see the entire topology.



**TABLE 10-1** Feature of the Toolbar in System Explorer (*Continued*)

Button	Description
	Lets you move an element in the topology. See “Arranging Elements in the Topology” on page 335.
	Lets you move multiple elements at once. This button is not accessible from the Topology tab. See “Arranging Elements in the Topology” on page 335.
	Lets you move the entire topology at once. Click the <b>Pan</b> (  ) 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.
	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. <ul style="list-style-type: none"> <li>• <b>Yes</b> - All users who log in to the management server see the topology you created. Only users with system configuration capability can save their layout for all other users</li> <li>• <b>No</b> - Other users cannot view the topology you saved. The saved topology appears the next time you log into the management server.</li> </ul>
	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 387.
	Lets you view only selected fabrics in the topology. This button may not be accessible from the Topology tab. See “Filtering Fabrics” on page 348.
	In Capacity Explorer, hides the lower pane.
	Change Observer button - Monitors changes in the database status on the server. When changes are detected, the button turns amber. Click on the amber button and a pop-up window displays the elements that have changed on the server. When no changes are detected, the button is greyed out.
	Reloads the Change Observer button to display the latest changes to elements on the server.



**TABLE 10-1** Feature of the Toolbar in System Explorer (*Continued*)

Button	Description
	<p>Lets you find an element by name or by Worldwide Name in the topology. You can enter part of the information, and the management server highlights the elements that match.</p> <p>After you populate the search box, click  or press ENTER.</p> <p>To expand the Search box, close the left pane. See “Opening and Closing the Left Pane” on page 8 for more information.</p>
	<p>Displays the event severity icons for the elements displayed in the topology.</p> <p>See “Viewing Event Status in the Topology” on page 349. This feature is disabled for Performance Explorer and Capacity Explorer.</p>
	<p>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 window displays a list of all the hosts with partially calculated paths.</p>
	<p>Displays the backup topology. The backup topology is also displayed in Protection Explorer.</p> <p>See “About Protection Explorer” on page 753 for more information about the backup protection provided in this product.</p>






## Icons Displayed in the Topology

Table 10-2, “Icons Displayed in the Topology,” on page 314 provides a brief description of the icons displayed in the System Explorer topology.

**TABLE 10-2** Icons Displayed in the Topology

Icon	Description
	Indicates an application.
	Indicates a file server.

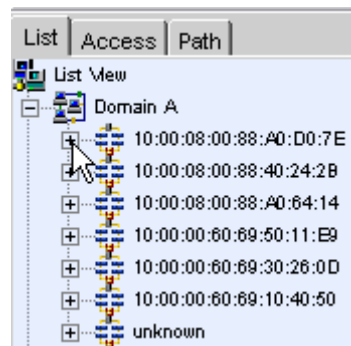
**TABLE 10-2** Icons Displayed in the Topology (*Continued*)

Icon	Description
	Indicates a host. This particular icon is for a host running Microsoft Windows.
	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.

## The List Tab

The List tab provides information about the elements by type, by cluster, or by fabric and domain.

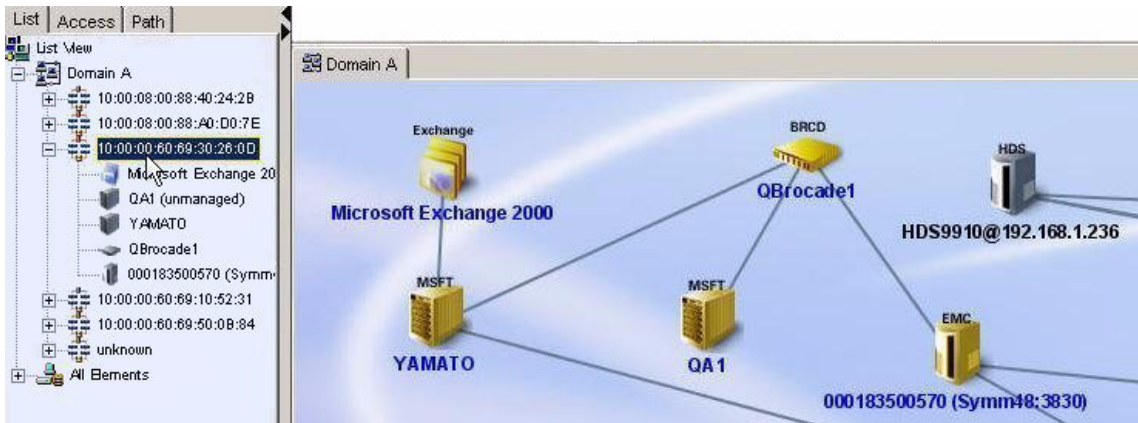
To find the fabrics in a domain, expand the domain node. You can see the elements in each fabric by expanding the fabric node, as shown in the following figure.



**FIGURE 10-1** Expanding the Fabric Node

The “unknown” Fabric lists elements that have a Fibre Channel port connected to an undiscovered Fabric or that have a Fibre Channel port that remains unconnected.

When you click a fabric name in the tree, its members are highlighted in the right pane, as shown in the following figure.



**FIGURE 10-2** Highlighting a Fabric's Members in the Topology

When you right-click an element in the List tab, a menu is displayed. The options displayed depend on the type of element you clicked. See “About the Right-Click Menu Options” on page 326 for an explanation of the options in the menu.

## Viewing Clustered Elements

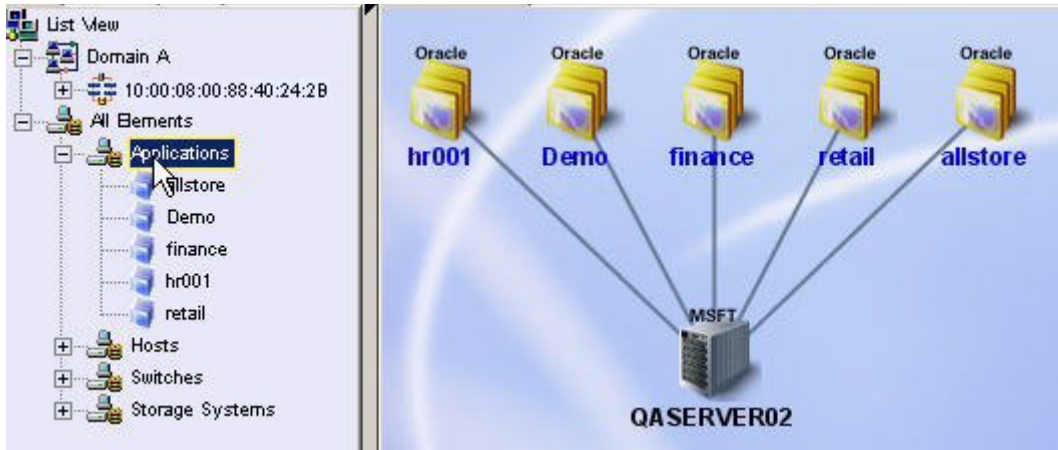
You can also view elements that are part of a host cluster or application cluster. To find the elements in a cluster, expand the Host Clusters or Application Clusters node. When you click a cluster name in the tree, its members are highlighted in the topology.

## 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 tree of 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 tree of 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 10-3** 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 326 for more information.

## The Access Tab

The Access tab provides information about the following:

- Zone entries
- Host bindings
- Storage system LUN masking

You can also manage zones, zone aliases, and zone sets from this tab by right-clicking an element. See “About the Right-Click Menu Options” on page 326.

## Obtaining Information About Zone Entries

To view the zone entries in a domain, expand the tree for the domain, fabric, and zone set. Select the zone set node to see the members of the zone set highlighted in the right pane, as shown in the following figure.

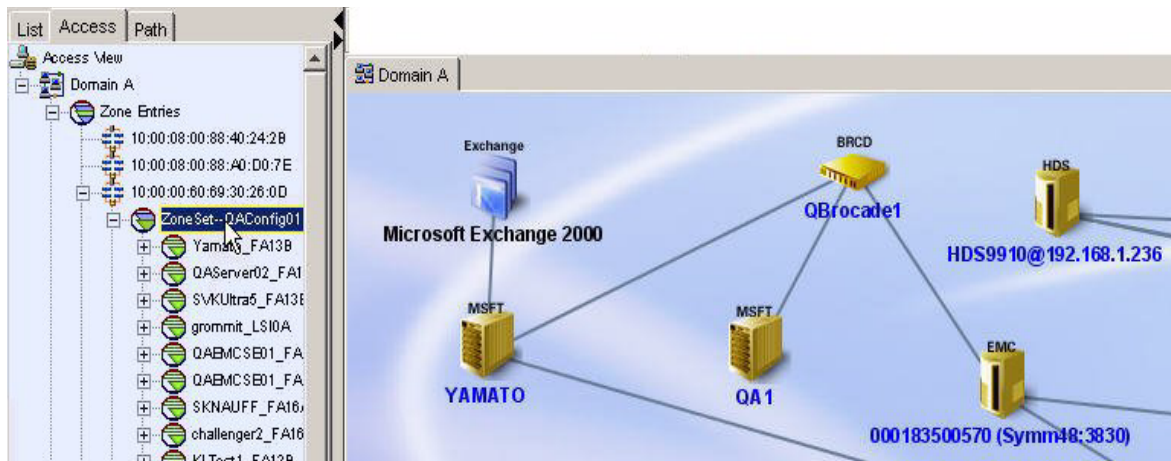


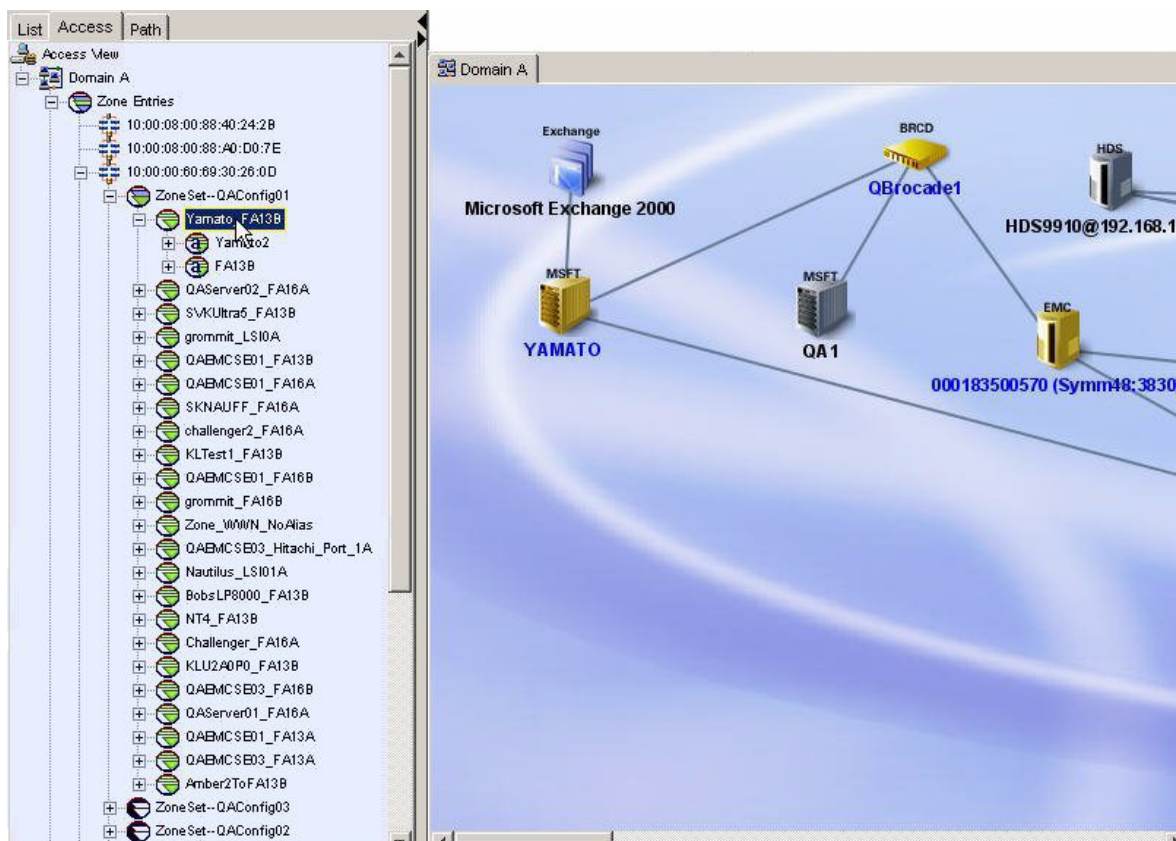


FIGURE 10-4 Members of a Zone Set

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 10-5** Displaying a Zone Member and its Switch

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.



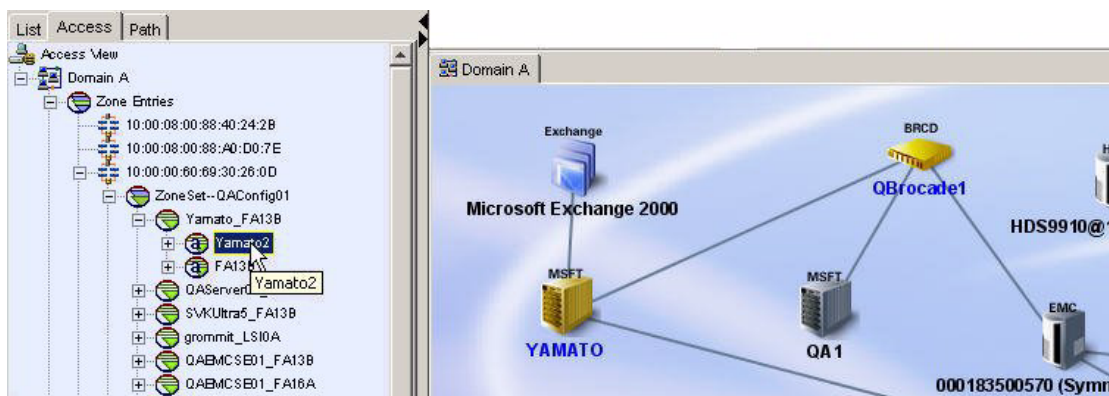


FIGURE 10-6 Zone Member to Switch

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.

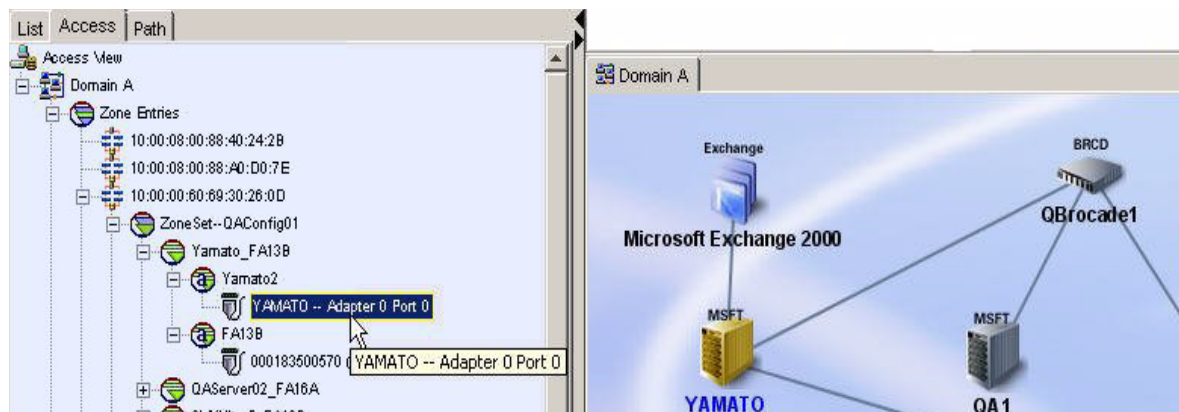
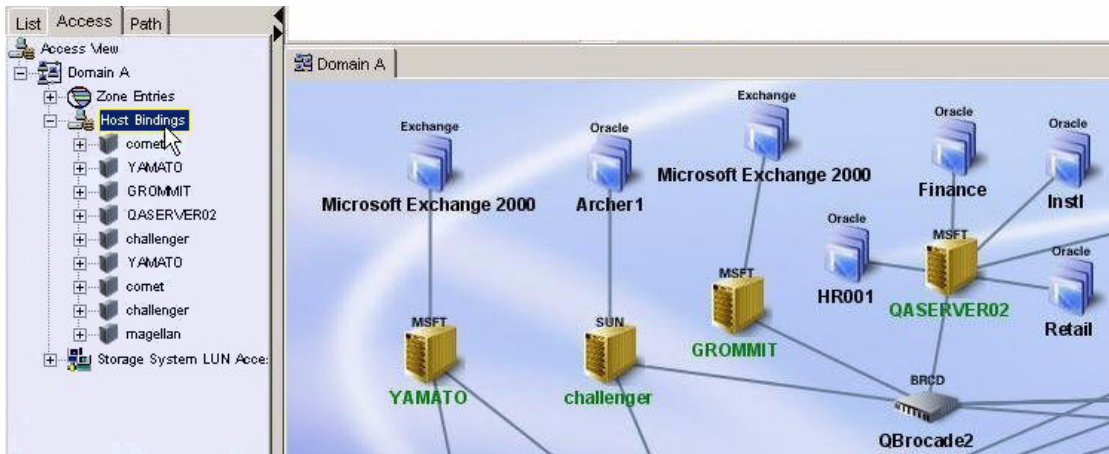


FIGURE 10-7 Obtaining Information About a Zone Member's Adapter

## 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 10-8** 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.

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, as shown in the following figure:



**FIGURE 10-9** 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 10-10** HBA Port Properties

## Obtaining Information About Storage System LUN Masking

To obtain information about a storage system's LUN masking, expand the fabric node, and 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.

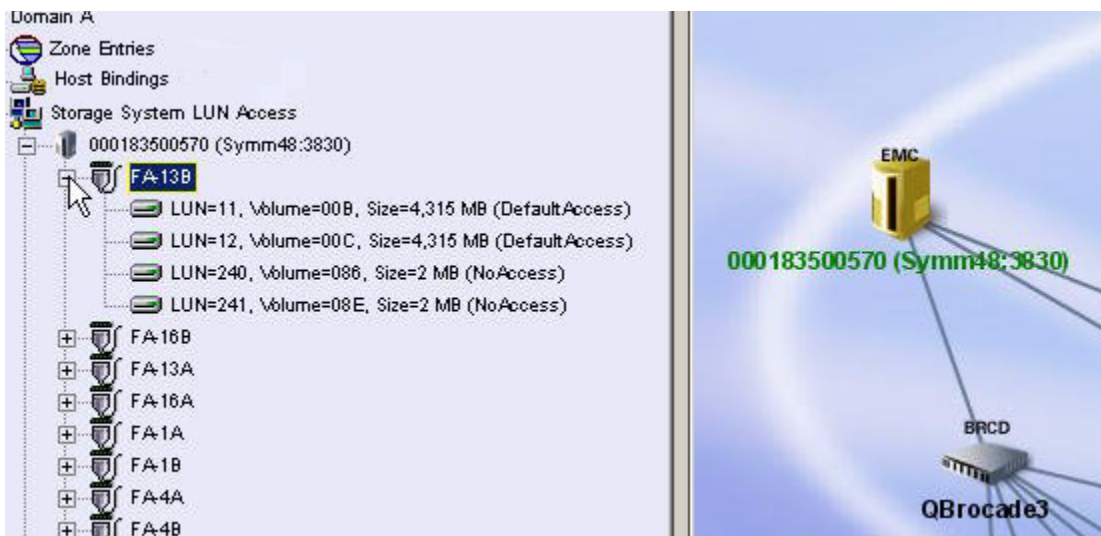
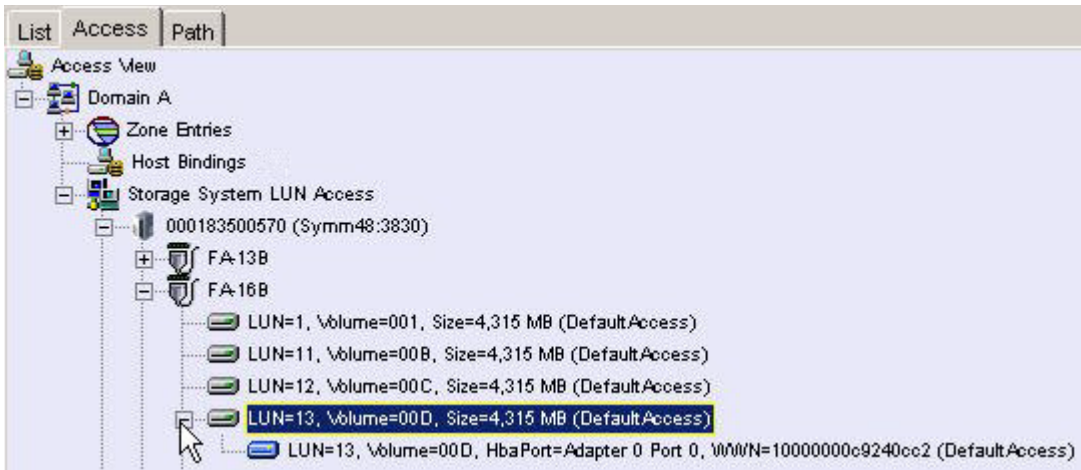


FIGURE 10-11 WWN Properties

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, as shown in the following figure.



**FIGURE 10-12** WWN Properties

To view a LUN masking, expand a LUN node.

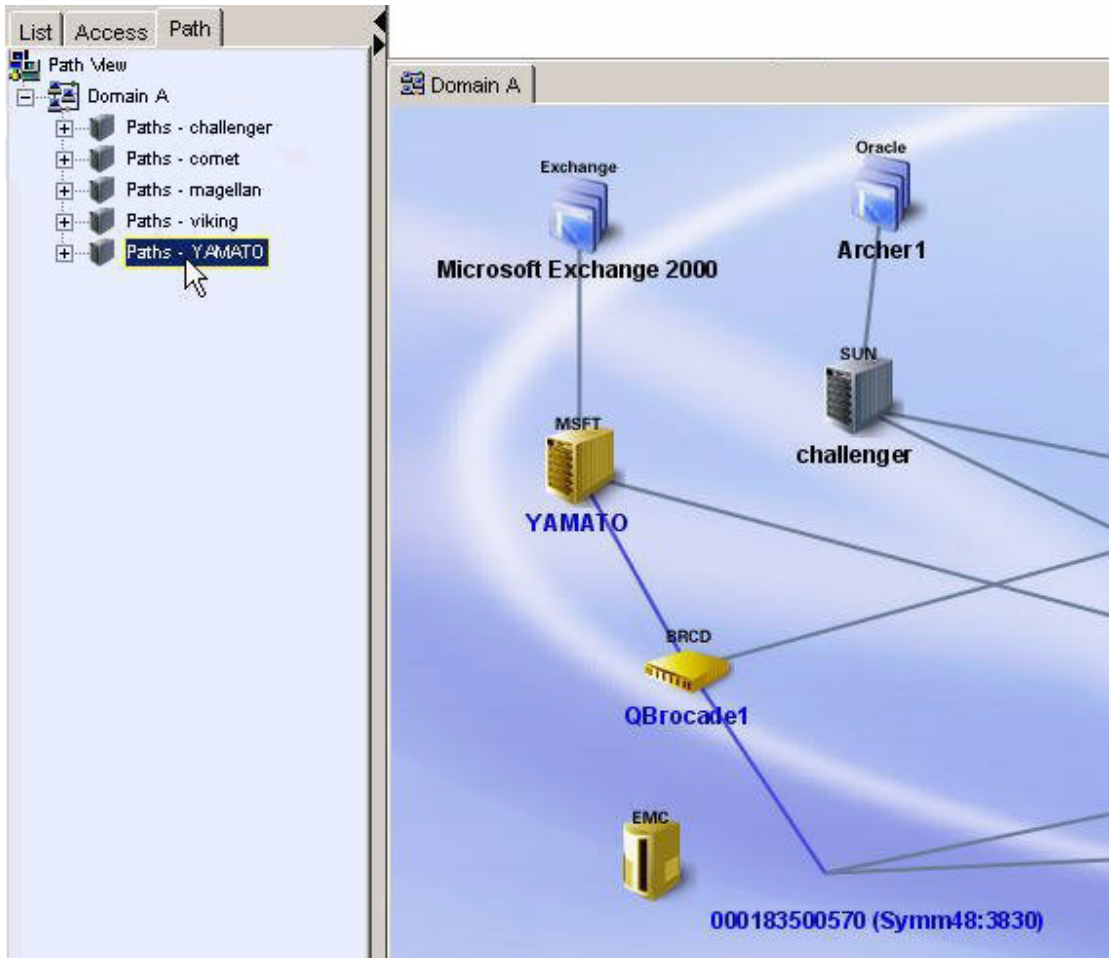
## About the Path Tab

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.

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.

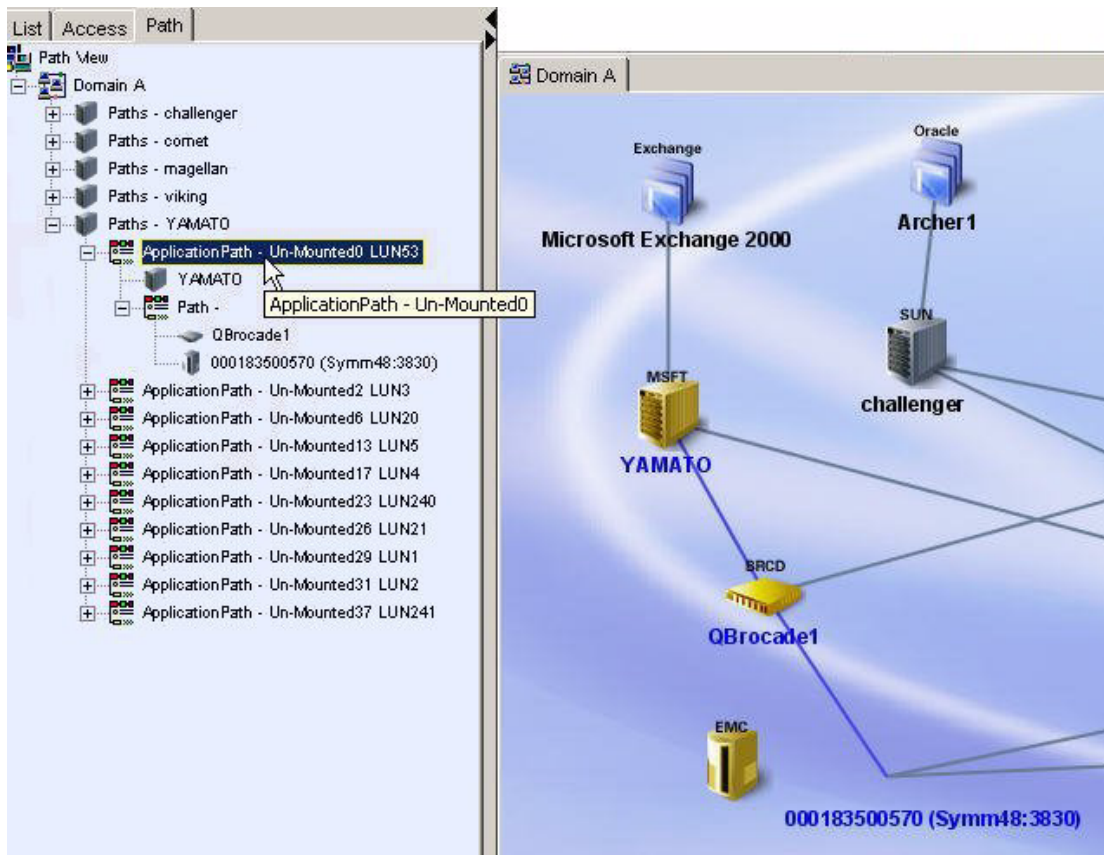
When you click a host node in the tree, the elements in the host's path appear highlighted in the right pane, as shown in the following figure.

**Caution** – If you are using the Load-on-Demand feature, you must load the path first by expanding the tree node. If a path tree node is not loaded, there are no elements under the path node and no highlighting occurs. For more information about the Load-on-Demand feature, see “System, Capacity and Performance Manager Preferences” on page 187.



**FIGURE 10-13** 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, as shown in the following figure.



**FIGURE 10-14** Path Information Visible in the Tree

When you right-click an element in the List tab, a menu is displayed. The options displayed depend on the type of element you clicked. See “About the Right-Click Menu Options” on page 326 for an explanation of the options in the menu.

## About the Right-Click Menu Options

When you right-click an element in the topology pane or in the List, Access, or Path tab, you see a 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 in the List or Path tab. For the options on the Access tab, see Table 10-4, “Menu Options on the Access Tab,” on page 331.

---

**Note** – Right-click menu options are not available to undiscovered fabrics.

---

**TABLE 10-3** Menu Options Accessible from the Topology\*

Menu Option	Description
<b>Go to Navigation Details</b>	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 cannot obtain additional information about it. See “About the Navigation Tab” on page 370.
<b>Go to Element Topology</b>	Displays a graphical representation of the path of an element. This also includes multipathing. See “Viewing Element Topology” on page 379.
<b>Show Events</b>	Displays the events for an element. See “About the Events Tab” on page 392.
<b>Show Policies</b>	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.
<b>Update Element Data</b>	<p>The management server gathers new and changed details from the element and then redraws the topology with the updated information.</p> <p>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 <b>Update Element Data</b>, 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 “Get Details” on page 91.</p>
<b>Show Impact</b>	Highlights the elements that are impacted. See “Showing the Impact of an Element” on page 344.
<b>Show Cluster Impact</b>	Highlights the elements that are impacted. See “Showing the Impact of an Element” on page 344.
<b>Show Port Details</b>	Lets you determine the use of each port for all elements in the network. See “Viewing Ports” on page 343 for more information.
<b>Build/Edit Cluster</b>	Lets you manually build or edit host and application clusters. See “Host and Application Clustering” on page 165.

**TABLE 10-3** Menu Options Accessible from the Topology\* (Continued)

Menu Option	Description
<b>External Tools</b>	<p>Provides several ways to access an element:</p> <ul style="list-style-type: none"> <li>• <b>Telnet</b> - Lets you access a host or a switch through the telnet utility. The Telnet feature is accessible only to Web browsers on Microsoft Windows operating systems.</li> <li>• <b>Browse</b> - Lets you access the main Web page for a host or a switch.</li> <li>• <b>Set Up External Tools</b> - 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.</li> </ul> <p>See "Using External Tools" on page 369 for more information.</p>
<b>Discovered Elements</b>	<p>Lets you group unnamed generic hosts. It provides the following options. See "About Hiding Generic Hosts" on page 356.</p> <ul style="list-style-type: none"> <li>• <b>Hide Generic Hosts for the Switch</b> - Hides unnamed generic hosts that are connected to the switch. See "Hiding Generic Hosts for a Switch" on page 356.</li> <li>• <b>Expand Generic Hosts for the Switch</b> - Displays hidden unnamed generic hosts that are connected to the switch. See "Expanding Generic Hosts for a Switch" on page 357.</li> <li>• <b>Hide Generic Hosts for All Switches</b> - Hides unnamed generic hosts connected to all switches. See "Hiding Generic Hosts for All Switches" on page 357.</li> <li>• <b>Expand Generic Hosts for All Switches</b> - Displays hidden unnamed generic hosts that are connected to the switches. See "Expanding Generic Hosts for All Switches" on page 357.</li> </ul>



**TABLE 10-3** Menu Options Accessible from the Topology\* (Continued)

Menu Option	Description
<b>Provision</b>	<p>Provides provisioning tools for switches and storage systems:</p> <p><b>Switches</b> - Lets you activate and deactivate zone sets, in addition to managing the following:</p> <ul style="list-style-type: none"> <li>• Zone aliases (Not applicable to McDATA switches)</li> <li>• Zones</li> <li>• Zone sets (The delete zone set option is disabled for active zone sets).</li> </ul> <p><b>Notes:</b></p> <ul style="list-style-type: none"> <li>• 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, or zone aliases.</li> <li>• 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, or delete a zone set or its members.</li> </ul> <p><b>Storage Systems</b> - Lets you manage the following:</p> <ul style="list-style-type: none"> <li>• Storage pools</li> <li>• Volumes</li> <li>• Host security groups</li> </ul> <p>These options are offered only when supported by the storage system.</p> <p>See “The Provisioning Tab” on page 391, “SAN Zoning Overview” on page 401 and “Setting Up Storage Partitioning” on page 422 and for more information.</p> <p>Provisioning wizards may not be available in your build of the product. To determine if you can access Provisioning wizards, access the List of Features, which is accessible from the Documentation Center.</p>
<b>Add Virtual Application</b>	<p>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 333.</p>
<b>Set Business Cost</b> (Applications only)	<p>Lets you assign a business cost to an application. See “Assigning a Business Cost to an Application” on page 346.</p>
<b>Delete Element</b>	<p>Removes an element and its discovery instance from the system. It also removes other elements discovered through the removed element.</p> <p><b>Important:</b> 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 pop-up blocker before you can delete the element.</p>
<b>Adding a Custom Command</b>	<p>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 359.</p>

**TABLE 10-3** Menu Options Accessible from the Topology\* (Continued)

Menu Option	Description
<b>Group together with other elements</b>	Lets you group “Discovered” hosts and storage systems. See one of the following topics: <ul style="list-style-type: none"> <li>• “Grouping Discovered Hosts” on page 351</li> <li>• “Grouping Discovered Storage Systems” on page 353</li> </ul>
<b>Ungroup into multiple elements</b>	Lets you ungroup “Discovered” hosts and storage systems. See one of the following topics: <ul style="list-style-type: none"> <li>• “Ungrouping Discovered Hosts” on page 352</li> <li>• “Ungrouping Discovered Storage Systems” on page 354</li> </ul>
<b>Recalculate Topology</b>	Lets you know about topology changes. The management server contacts the elements on the topology list ( <b>Discovery &gt; Topology</b> ) to determine topology changes. The management server uses this updated information to redraw the topology.
<b>Change Fabric Name</b> (Only Available from the List Tab)	Lets you change the name of the fabric. See “Changing the Fabric Name” on page 355 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:

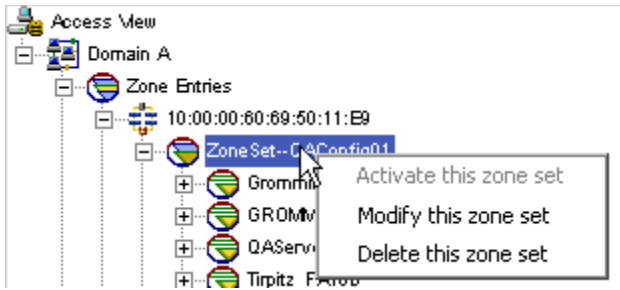

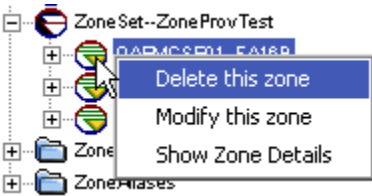

- **Change Fabric Name** - Enter the new fabric name in the Change Fabric Name window, and then click **OK**.
- **Go to Properties** - Displays the properties of the fabric. See “Viewing Element Properties” on page 375 for more information. It also provides access to the Events and Provisioning tabs. 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 List of Features, 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 menu is displayed. The provisioning options displayed in this menu depend on the type of element clicked.

**Caution** – 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 “Provisioning” on page 399.

**TABLE 10-4** Menu Options on the Access Tab

Task	To Perform Task, Right-Click...
Activate a zone set <sup>1</sup>	<p>A zone set under the Zone Entries node.</p> 
Modify a zone set <sup>1</sup>	
Delete a zone set <sup>1,3</sup>	
Create a zone <sup>1</sup>	<p>A zone with the blue folder icon under a fabric node.</p> 
Delete a zone <sup>1</sup>	<p>A zone under the Zone Set node.</p> 
Modify a zone <sup>1</sup>	
Show zone details <sup>1,2</sup>	
Create a zone alias <sup>1</sup>	<p>A zone alias with a blue folder under the fabric node.</p> 

**TABLE 10-4** Menu Options on the Access Tab (*Continued*)

Task	To Perform Task, Right-Click...
Delete a zone alias <sup>1</sup>	A zone alias under the Zone Aliases node.
Modify a zone alias <sup>1</sup>	
Show zone alias details <sup>1</sup>	
Provision a storage pool	A storage system under the Storage System LUN Access node.
Provision a volume	
Provision a host security group	



<sup>1</sup>When McDATA and Connectrix switches are discovered through a proxy by SNMP, you cannot view or perform any provisioning operations for those switches.

<sup>2</sup>Only 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, or delete zone sets or its members.

<sup>3</sup>The 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 333
- “Adding Information for Discovered Hosts” on page 334
- “Arranging Elements in the Topology” on page 335

- “Closing Topology Windows” on page 337
- “Using the Global View” on page 337
- “Printing the Topology” on page 338
- “Exporting the Topology to Microsoft Visio” on page 340
- “Updating Element Data” on page 342
- “Viewing Ports” on page 343
- “Showing the Impact of an Element” on page 344
- “Assigning a Business Cost to an Application” on page 346
- “Expanding the Topology Pane” on page 347
- “Filtering Fabrics” on page 348
- “Viewing Event Status in the Topology” on page 349
- “Custom Name for a Switch Truncated in the Topology” on page 350
- “Managing Groups” on page 350
- “Managing Fabrics” on page 354
- “Hiding and Showing Generic Hosts” on page 355

## 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 as connected to a host in your topology.

1. Select a host.
2. Right-click, and select **Add Virtual Application**.
3. Enter 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 812.


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

---

**Caution** – 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 the 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 box, enter a name for the element.  
Keep in mind the following:
  - The name must contain 1 to 64 characters.
  - The following characters and symbols are accepted: letters, numerals (0 to 9), ~, @, \*, \_ , +, ., < >, (), [ ], { }, |.
  - The name is case sensitive, for example, “Element1” and “element1” are different elements.
4. In the IP Address box, enter an IP address for the element.
5. In the DNS Name box, enter a DNS name for the element.


6. In the Version box, enter the version of the operating system.
7. In the Operating System box, enter 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™
  - NonStop
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. Thus, 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 10-1, “Feature of the Toolbar in System Explorer,” on page 312 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 “The Toolbar in System Explorer” on page 312.

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 348.

To arrange elements in a group:

1. Click the  button.
2. Holding down the mouse button, move the cursor diagonally across the elements you want to move.


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 10-15** 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 10-16** 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 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 high-level 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 high-level view of the topology appears.

2. Move the brackets so they enclose the area of the network you want to view in the main pane.

## Printing the Topology

To print the topology displayed in System Explorer:

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. This prevents the printout from appearing too stretched. To move an element, click the  button and then click the element you want to move. Drag the element to its new location.
3. Click the  button.
4. The Paper tab shows the page setup. If you want the default settings, click **Default**. You can modify the following settings:

---

**Caution** – Before you change the margins, decide on a unit of measurement.

---

- **Paper format** - Select the paper size from the menu.
- **Unit** - Select **cm** (centimeters) or **inch** for the margins.
- **Paper width** - To modify the width of the paper, select the **Custom** option in the Paper format menu.
- **Paper height** - To modify the measurement in this box, select the **Custom** option in the Paper format menu.
- **Top margin** - Enter a measurement.
- **Bottom margin** - Enter a measurement.
- **Left margin** - Enter a measurement.
- **Right margin** - Enter a measurement.
- **Orientation** - Click an orientation for the printout.

A preview of the printout is displayed in the right pane.

5. When you are done, click **Apply**.
6. To see how the printout will appear on the page, click the **View Selection** tab. If you want the default settings, click **Default**. You can modify the following settings:

- **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 **Trimmed**.

A preview of the printout is displayed in the right pane.

7. When you are done, click **Apply**.

8. The Pages tab shows how many pages the printout will use. If you want the default settings, click **Default**. You can modify the following settings:

---

**Caution** – 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** - Same as in step 6.
  - Start y** - Same as in step 6.
  - 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.

A preview of the printout is displayed in the right pane.

9. When you are done, click **Apply**.

10. To preview your pages, click the **Preview** tab. Then click the page you want to preview.

The page appears in the right pane.

11. When you are ready to print, click **Print**.

12. Click **Close**.


---

**Note** – To return to all of the original settings, click the **Default** button next to the **Print** button.

---

## Exporting the Topology to Microsoft 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 **Export to Visio**.
3. Name the file, and select the directory where you want the file to be saved.
4. Click **Save**. The XML file is saved to the directory that you chose.

---

**Note** – Backup topology can be exported only through Protection Explorer, 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:

1. Install the necessary components, as described in “Installing Storage Planner” on page 341.
2. Configure Visio, as described in “Configuring Visio to View Exported Topology” on page 341.
3. Start Visio, and select **Storage Planner > Import XML File**.
4. Browse to the XML file that you would like to view, and click **Open**. The selected topology is displayed in Visio.
5. Right-click any element, and select **Properties**. The Custom Properties window opens and displays additional information about each element.

- 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:

- From the Windows directory on the Utilities CD-ROM, run `StoragePlanner.exe`.  
The Welcome to the Storage Planner Setup Wizard window is displayed.
- Click **Next**.  
The Select Destination Location windows is displayed.
- Click **Next**.  
The Select Components window is displayed.
- Click **Next**.
- Click **Install**.  
The Storage Planner component is installed.
- If you have not installed Microsoft XML 6.0 Parser, select the check box.
- Click **Finish**.

---

**Note** – If you are installing Microsoft XML 6.0 Parser, the MSXML 6.0 Parser Setup wizard is displayed. If MSXML 6.0 parser was installed previously, you will not see the following steps.

---

- Select the **I Accept** option button, and click **Next**.
- Enter your name and company information. Click **Next**.
- Click **Install**. Microsoft XML 6.0 Parser is installed.
- Click **Finish**.

## 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 **Macro Security**.
5. Click the **Medium** radio button.
6. Click **OK**. You are returned to the previous window.
7. Click the **Advanced** tab.
8. Click **File Paths**.
9. Click the button to the right of the Add-ons box. Browse to your Visio installation directory, and select the following path:  
`<Visio_installation_directory>\1033\Solutions\StoragePlanner`  
where `<Visio_installation_directory>` is the installation directory for Visio, for example  
`C:\Program Files\Microsoft Office\Visio11`.
10. Click the button to the right of the Start-up box. Browse to your Visio installation directory, and select the following path:  
`<Visio_installation_directory>\1033\Solutions`  
where `<Visio_installation_directory>` is the installation directory for Visio
11. Click **OK**. You are returned to the previous window.
12. Click **OK**.
13. Restart Visio, and select **Enable Macros** 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. For the deleted LUNs to be removed from the user interface, you must perform Get Details. See “Get Details” on page 91.
- Update Element Data does not correctly update switch ISL and topology changes. To obtain switch ISL and topology changes, you must perform Get Details.

To update an element:

1. Click **System Explorer** ()
2. Right-click the element you want to update.
3. Select **Update Element Data** from the 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 you are 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 menu.

2. Right-click an element in the topology.
3. From the menu, select **Show Port Details**.

The ports are displayed.

## Showing the Impact of an Element

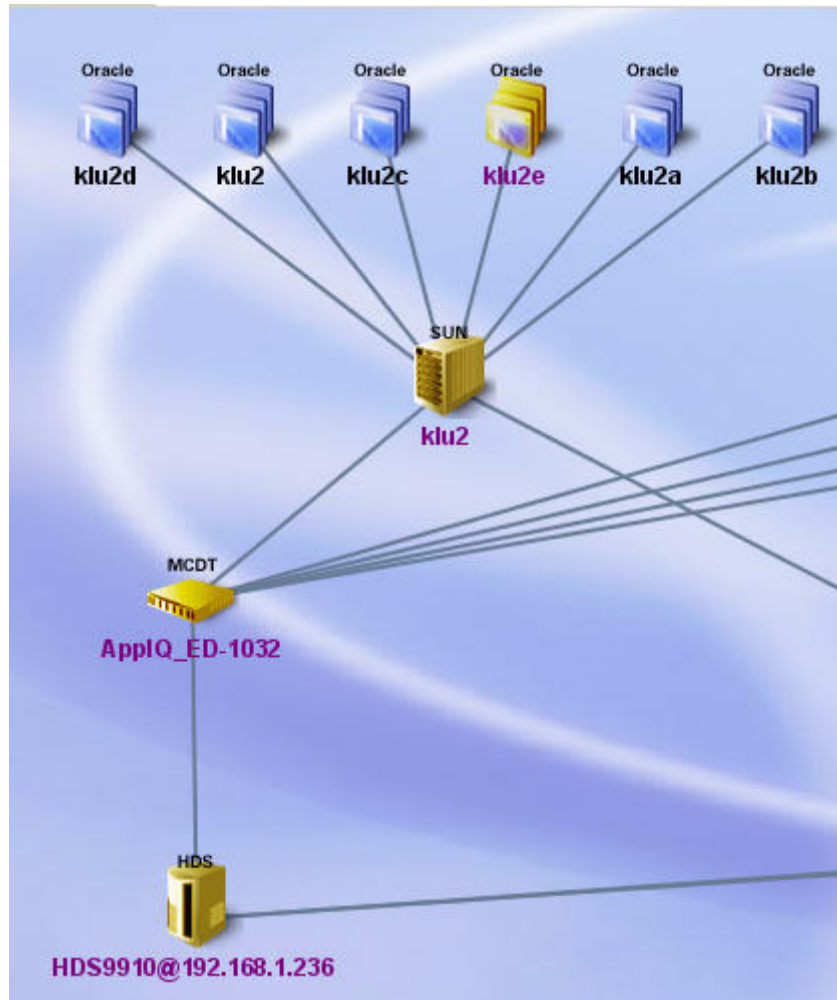
You can display an element's impact. For example, assume you want to replace a switch. You can use this feature to determine which elements in the network would be impacted by taking the switch off the network.

To find an element's impact:

1. Right-click the element from which you want to obtain impact information.
2. Select **Show Impact** (or **Show Cluster Impact** if you right-clicked a cluster) from the 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 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 10-17** Showing the Impact of an Element

However, the Show Impact feature not only displays the 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 menu. Each highlighted element 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, each highlighted element 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 10-5** Show Impact Results

If you select Show Impact for...	The software highlights...
An Application (virtual or real)	Elements required by the application, such as its host and a switch.
A Host	Elements that are dependent on the host, such as its applications.  Elements that are required by the host, such as switches.
A Switch	Elements dependent on the switch, such as hosts and storage systems
A 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 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.

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 has a value of \$60 if two \$30 applications are in its path, as shown in the following figure.




**FIGURE 10-18** 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 menu.
  - Double-click an application in System Explorer. Click the **Properties** tab. Then click the **Change** button next to the Business Cost box.
  - Click an application in Application Explorer. Click the **Properties** tab. Then click the **Change** button next to the Business Cost box.
2. In the Business Cost box, enter 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.



**FIGURE 10-19** 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 8.

---

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:

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 10-20** Filtering Fabrics


4. Click **Apply Filter**.

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 displayed next to them an icon that indicates the severity of the event. Table 10-6, “Severity Levels,” on page 349 shows the icons and describes their meaning.

---


**Note** – The Event Status button () is disabled in Capacity Explorer and Performance Explorer.

---

**TABLE 10-6** Severity Levels


Icon	Severity Level	Description
	The event has a critical impact.	Denotes elements that have a critical severity level. The elements may also have events of lower severity levels. Example: A Brocade switch has a failed firmware download. The failure reason code for each respective switch is displayed.
	The event has a major impact.	Denotes elements that have a major severity level. The elements may also have events of lower severity levels. Example: one or more physical fabric objects (device port, switch, or fabric) have disappeared.
	The event has a minor impact.	Denotes elements that have a minor severity level. The elements may also have events of lower severity levels. Example: A physical fabric object (switch port or fabric) has changed state.
	The event is providing a warning.	Denotes elements that have a warning severity level. The elements may also have events of lower severity levels. Example: One or more new physical fabric objects (device port, switch, or fabric) have appeared.
	The event is providing information.	Denotes elements that have an informational severity level. The elements may also have events of lower severity levels. Example: A progress report event for a firmware download operation is currently in progress.

**TABLE 10-6** Severity Levels


Icon	Severity Level	Description
	The severity of the event is not known.	Denotes elements that have an unknown severity level. It displays icons for the following severity levels: <ul style="list-style-type: none"> <li>• Critical</li> <li>• Major</li> <li>• Minor</li> <li>• Information</li> </ul>

The icon 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 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 only critical and major events by selecting **Major** from the Severity menu and clicking the  button.

---

**Note** – If you select a severity, click the  button, and then leave System Explorer, System Explorer remembers your selection.

---

## Custom Name for a Switch Truncated in the 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 contains the following topics:

- “About Groups” on page 351
- “Grouping Discovered Hosts” on page 351
- “Ungrouping Discovered Hosts” on page 352
- “Grouping Discovered Storage Systems” on page 353
- “Ungrouping Discovered Storage Systems” on page 354

## 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. 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 "Editing Roles" on page 190 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 either the topology or all element details, because the status button appears red during both operations.
- You can determine 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 310.
2. Right-click a Discovered host, and select **Group together with other hosts** from the menu.
3. In the Custom Name box, enter a custom name for the group.
4. In the IP Address box, enter an IP address for the group.
5. In the DNS Name box, enter the DNS name for the group.
6. In the Version box, enter a version number for the group.

7. In the Operating System box, enter 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 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

If one of the hosts in a group is going to change, you might want to ungroup Discovered hosts. An example of such a change would be when a 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 either 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 190 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 310.



2. Right-click the host icon for a group, and select **Ungroup into multiple hosts** from the menu.
3. When you are asked if you want to ungroup the elements, click **OK**.  
The elements are ungrouped.

## Grouping Discovered Storage Systems

As with hosts, 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. Grouping elements preserves 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 190 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 either the topology or all element details by looking at label near the status button.
- To determine if a storage system is generic, double-click the storage system in System Explorer, click the Properties tab, and look at the storage system's description.

To group Discovered storage systems:

1. Access System Explorer, as described in "Accessing System Explorer" on page 310.
2. Right-click a Discovered storage system, and select **Group together with other** from the menu.
3. In the Custom Name box, enter a custom name for the group.
4. In the Vendor box, enter 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 that 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 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 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

If one of the storage systems in a group is going to change, you might want to ungroup Discovered hosts. An example of such a change would be when a 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 190 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 310.
2. Right-click a storage system icon for a group, and select **Ungroup into multiple storage** from the menu.
3. When you are asked if you want to ungroup the elements, click **OK**.

The elements are ungrouped.

## Managing Fabrics

This section contains the following topics:

- “Changing the Fabric Name” on page 355
- “Deleting Fabrics” on page 355

### 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 menu.
5. In the Enter a Fabric Name box, enter 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 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 contains the following topics:

- “About Hiding Generic Hosts” on page 356
- “Hiding Generic Hosts for a Switch” on page 356
- “Expanding Generic Hosts for a Switch” on page 357
- “Hiding Generic Hosts for All Switches” on page 357
- “Expanding Generic Hosts for All Switches” on page 357

---

**Note** – The feature described in this section pertains only 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 378. To learn more about groups, see the topic, “About Groups” on page 351.

---

## 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.

The management server provides two variations of this feature:

- **Hiding Generic Hosts for One Switch:** 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 by displaying a question mark above its icon.

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 that 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 356.

- **Hiding Generic Hosts for All Switches:** 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 357.

## 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 357 for more information.

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 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

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 menu.  
The hidden elements for the switch appear in the upper-right corner of the topology.

## Hiding Generic Hosts for All Switches

To hide all unnamed generic hosts and unnamed generic Cisco switches:

1. Right-click a switch.
2. Select **Discovered Element > Hide Generic Hosts for All Switches** from the 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

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 menu.

The hidden elements for the domain appear in upper right corner of the topology.

---

## Setting Up Custom Commands

This section contains the following topics:

- "About Custom Commands" on page 358
- "Adding a Custom Command" on page 359
- "Editing a Custom Command" on page 361
- "Deleting a Custom Command" on page 362
- "Software Environment Variables for Scripting" on page 362
- "Using the Remote Console" on page 365

## About Custom Commands

Custom commands let you run a command that you create 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.

## Important Considerations

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.

- The custom command always runs on the management server unless you are running the telnet utility. You can obtain information about an element on which you right-click by using the software's environment variables. See “Software Environment Variables for Scripting” on page 362.
- 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 a UNIX system, 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 %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

myscript.pl is the script you want to run.

## Adding a Custom Command

Before adding a custom command, be sure you are aware of the considerations listed in “Important Considerations” on page 358.

To add a custom command:

1. Right-click an element in System Explorer.
2. Select **Custom Commands > Set Up Custom Commands** from the 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 follows:

- a. In the Custom Command Setup window, click **Browse** to find the file containing the custom command.
- b. Click **Open**, and then click **Upload to server**.

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, enter a name for the command in the Name box, for example, backup command.
6. In the Description box, enter a description, for example, This command activates a script that backs up an element.
7. In the Command Line box, enter a command.

This could be a command required to start a script, for example:

```
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
```

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** menu, and then click **Append To Command Line**.

If the file is missing, repeat step 3.

8. Select one of the following options to determine the elements for which you want the command to be visible. For example, if you select the **All Elements** option, the command is visible in the 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 menu only when you right-click this element.
  - **All Elements** - Select this option if you want the command to be visible in the menu when you right-click any element.
  - **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 menu when a Brocade switch is right-clicked. The options are as follows:

**Applications** - If you want the command to be visible in the menu when a particular application is right-clicked, enter the name of the product in the Product Name box. To make sure you enter the correct product name, enter



the product name displayed in the Product Name box on the Properties tab, accessible by double-clicking the application in System Explorer and then clicking the Properties tab.

**Hosts** - If you want the command to be visible in the menu when a particular host is right-clicked, enter the name of the operating system in the OS Name box. To make sure you enter the correct operating system, enter the operating system displayed in the Target Operating System box on the Properties tab, accessible by double-clicking the host in System Explorer and then clicking the Properties tab.

**Switches** - If you want the command to be visible in the menu when a switch from a particular vendor is right-clicked, enter the name of the vendor in the Vendor Name box. To make sure you enter the correct vendor name, enter the vendor name displayed on the Properties tab, accessible by double-clicking the switch in System Explorer and then clicking the Properties tab.

**Storage Systems** - If you want the command to be visible in the menu when a storage system from a particular vendor is right-clicked, enter the name of the vendor in the Vendor Name box. To make sure you enter the correct vendor name, enter the vendor name displayed on the Properties tab, accessible by double-clicking the storage system in System Explorer and then clicking the Properties tab.

9. Click **OK**.

To run a command:

1. Right-click an element.
2. Select **Custom Commands** from the menu.
3. Select the command from the **Custom Commands** 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 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 menu.
3. Click the  button corresponding to the custom command you want to delete.

The custom command is deleted.

## Software Environment Variables for Scripting

The software provides environment variables for you to put in your scripts. For example, assume you have a script that backs up a host. You could use variables to obtain information about the host.

The software gathers information about the element you right-click. For example, if you use the variable `APPIQ_ELEMENT_ID` the management server obtains the element ID of the element you right-click.

Table 10-7, “Variables for All Elements,” on page 362 lists the variables that can be used to gather information for all elements. If an application resides on the host, the variables in this table provide information about the application. See Table 10-11, “Variables for Applications Only,” on page 364 for variables that return information about the host.

**TABLE 10-7** Variables for All Elements

Variable	Value
<code>APPIQ_ELEMENT_ID</code>	The identifier of an element.
<code>APPIQ_ELEMENT_NAME</code>	The name of the element.

**TABLE 10-7** Variables for All Elements

Variable	Value
APPIQ_ELEMENT_STATUS	The following statuses are available: <ul style="list-style-type: none"> <li>• Managed</li> <li>• Generic</li> <li>• Missing</li> <li>• Virtual Application</li> <li>• Asset</li> </ul>
APPIQ_ELEMENT_DESCRIPTION	The description for the element.
APPIQ_ELEMENT_MANAGEMENT_IP_ADDRESS	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.

Table 10-8, “Variables for Storage Systems, Switches, and Hosts Only,” on page 363 lists variables that can be used to gather information for storage systems, switches, and hosts only. If an application resides on the host, the variables in this table provide information about the application. See Table 10-11, “Variables for Applications Only,” on page 364 for variables that return information about the host.

**TABLE 10-8** Variables for Storage Systems, Switches, and Hosts Only

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.

Table 10-9, “Variables for Switches Only,” on page 363 lists variables that can be used to gather information for switches only.

**TABLE 10-9** Variables for Switches Only

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.

Table 10-10, “Variables for Hosts Only,” on page 364 lists variables that can be used to gather information for hosts only. If an application resides on the host, the variables in this table provide information about the application. See Table 10-11, “Variables for Applications Only,” on page 364 for variables that return information about the host.

**TABLE 10-10** Variables for Hosts Only

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.

Table 10-11, “Variables for Applications Only,” on page 364 lists variables that can be used to gather information for applications. 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 10-11** Variables for Applications Only

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_ADDRESSES	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 contains the following topics:

- “About the Remote Console” on page 365
- “Keeping the Remote Console Active” on page 366
- “Buttons on the Remote Console” on page 367

- “Menu Options” on page 368
- “Copying Text from the Remote Console” on page 368

## About the Remote Console

Whenever you run a custom command on the management server, the remote console appears, as shown in the figure below. 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 10-21 Remote Console

## Keeping the Remote Console Active

---

**Note** – This section is for management servers running on Microsoft Windows only.

---

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 menu.
3. Click the **Add Command** button in the upper-right corner of the window.
4. In the Add Custom Command window, enter a name for the command in the Name box, for example, prompt.
5. In the Description box, enter a description, for example, Accesses the remote console.
6. In the Command Line box, enter the following command, which will run on the management server:

```
cmd /k
```

7. Select the **All elements** option.
8. Click **OK**.

To run the remote command prompt:

1. Right-click an element from which you want to obtain information.
2. Select **Custom Commands** and select the command from the menu.

The software displays the remote console on the management server.

To enter a command in the remote console:

1. Enter the command in the Commands box, and then
2. 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, enter the following in the Commands box and then press ENTER:

```
plink
```

## Buttons on the Remote Console

The remote console provides the Stop and Help buttons, as described in Table 10-12, “Buttons on the Remote Console,” on page 367.

**TABLE 10-12** Buttons on the Remote Console

Button	Description
<b>Stop</b>	Stops a command. Once a command has been executed the console becomes inactive.
<b>Help</b>	<p>Provides the following information about the remote console:</p> <ul style="list-style-type: none"><li>• <b>Clearing the remote console</b> - Enter CLS in the Commands box of the remote console.</li><li>• <b>Copying text to the Commands box</b> - Place the cursor at the end of the line in the window below the Commands box, and then press ENTER. The command is copied to the Commands box.</li></ul> <p><b>Note:</b> 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.</p>

## Menu Options

The remote console also provides the following menu options.

**TABLE 10-13** Menu Options

Option	Description
<b>CRLF</b>	<p>(Default setting) Provides a carriage return and a linefeed.</p> <p><b>Important:</b> 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 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.</p>
<b>CR</b>	Provides a carriage return.
<b>LF</b>	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 menu.



The text is stored in the buffer of your computer to be pasted elsewhere.

---

## Using External Tools

This section contains the following topics:

- “The External Tools Feature” on page 369
- “Setting up External Tools” on page 369.

### 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 369 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 for accessing external tools used for managing an element, such as Hitachi HiCommand Device Manager and EMC ControlCenter™ Navisphere for sentertorage systems.

---


**Note** – When you add a URL, it applies only 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 **Add New Management URL**.
4. In the Description box, enter the name of the product you plan to access.
5. In the URL box, enter 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.

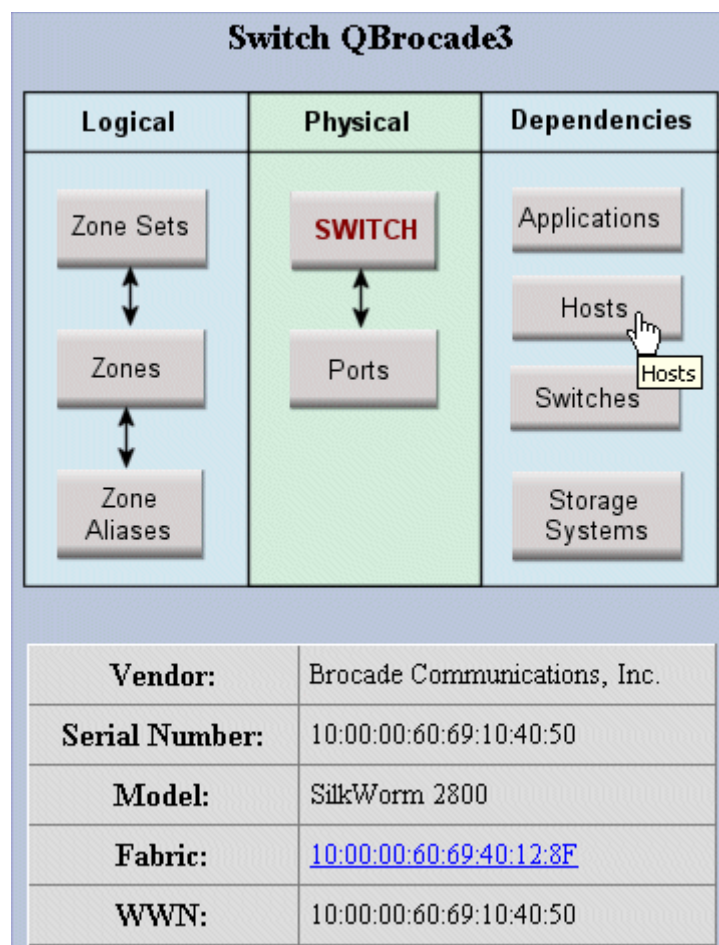
---

## 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 figure below.

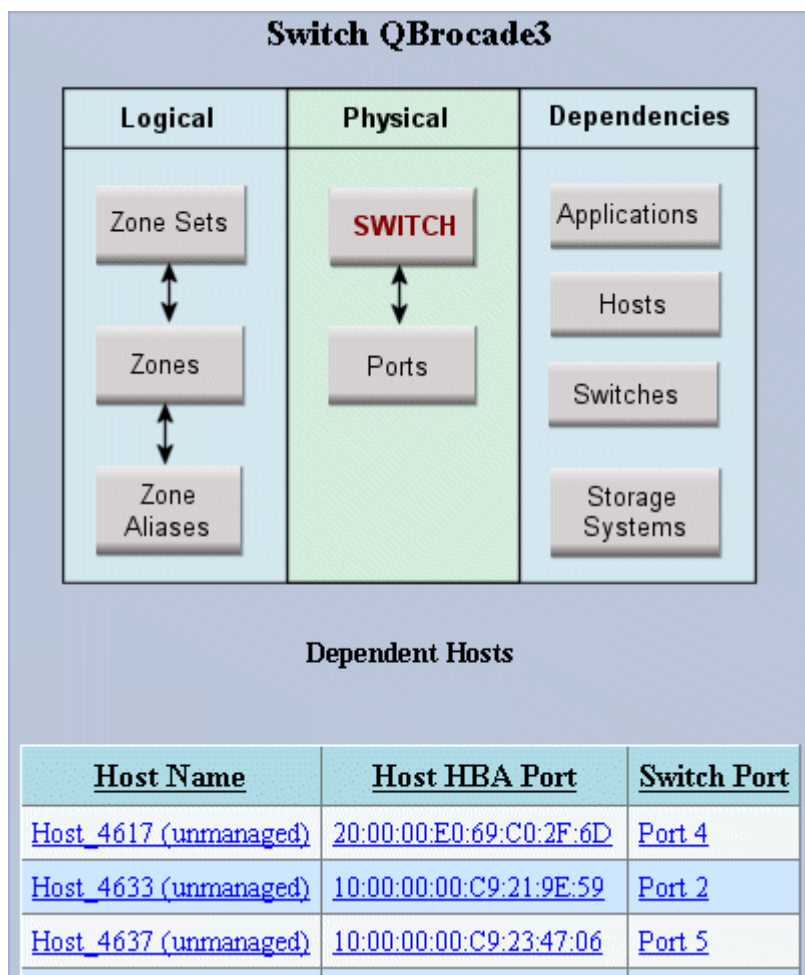
Keep in mind the following:

- 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, or zone aliases.
- 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, or delete a zone set 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 10-22** 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, and you want to learn which hosts are dependent. Click **Hosts** in the page. You are shown information about the dependent hosts, as shown in the following figure:



**FIGURE 10-23** Details of a Host Connected to a Switch

The following table provides an overview of the information presented for each type of element:

**TABLE 10-14** Information Available from the Navigation Page

Element	Dependencies	Front Physical	Back Physical	Logical	Physical
Applications	✓				
Hosts*	✓				

**TABLE 10-14** Information Available from the Navigation Page (*Continued*)

Element	Dependencies	Front Physical	Back Physical	Logical	Physical
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 enter 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 a McDATA or Connectrix switch is discovered by their IP address (SNMP connection), the following boxes are empty:

- IP Gateway
- Switch ID
- FC Net Address
- FC Net Mask

If a McDATA or Connectrix switch is discovered by SNMP and by Enterprise Fabric Connectivity (EFC) Manager or by EMC Enterprise Connectrix Manager (ECM), the following boxes are empty:

- IP Gateway
- Switch ID
- DNS Name
- IP Address
- IP Net Mask
- FC Net Address
- FC Net Mask

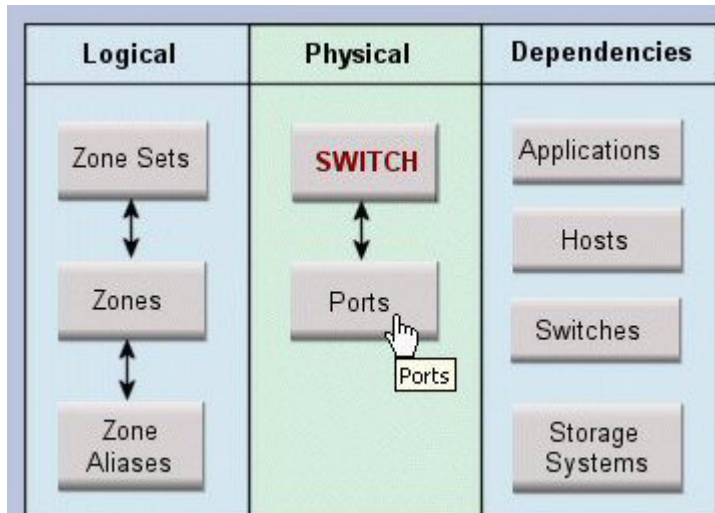
To learn how to access the Navigation tab, see “Accessing the Navigation Tab” on page 375 for more information.

## 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 310.

2. Double-click a switch in the right pane.
3. Click the **Ports** button in the Physical column in the Navigation tab, as shown in the following figure.



**FIGURE 10-24** Finding the Status of a Port

4. Under the Name column in the Ports table, click the port whose status you want to obtain.
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 10-15** 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.
    - Click one of the following elements in Protection Explorer, and then click **Navigation** in the lower-right corner.
      - Backup Client
      - Backup Library
      - Host
      - Master backup server
      - Master backup media
  3. Click the **Navigation** tab. (This is not necessary if you accessed the Navigation tab from Protection Explorer.)
- 

## Viewing Element Properties

This section contains the following topics:

- “About the Properties Tab” on page 375
- “Accessing the Properties Tab” on page 377
- “Viewing Fabric Properties” on page 377
- “Assigning a Custom Name” on page 378

## 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 only information relevant to 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. For supported Brocade switches, trunking ISL ports have Trunking State set to 2, and non-trunking ISL ports have Trunking State set to 1.

See “Viewing Fabric Properties” on page 377 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 378. This option is not available to all elements.
- **IP Address (Generic Hosts Only)** - Enter an IP address for a generic host.
- **DNS Name (Generic Hosts Only)** - Enter a DNS name for a generic host.
- **Version (Generic Hosts Only)** - Enter a version number for a generic host.
- **Operating System (Generic Hosts Only)** - Enter an operating system for a generic host.
- **Vendor** - Enter the vendor name.
- **View element properties** - Lets you view the element properties for 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 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.

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. For the deleted LUNs to be removed from the user interface, you must perform Get Details. See “Get Details” on page 91.



# 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

The Properties tab lets you view properties and take certain actions.

You can view the following properties of a fabric:

- **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** - 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.

You can take the following actions:

- **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 378. This option is not available to all elements.

- **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. For the deleted LUNs to be removed from the user interface, you must perform Get Details. See “Get Details” on page 91.

To learn how to access the Properties tab for a fabric, see the topic, “Accessing the Properties Tab” on page 377.

## 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.

---

**Caution** – 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.

---

---

**Note** – Since all users query the same database, this name is displayed to others using the software, so you might want to make them aware of the name.

---

1. Access the custom name box by double-clicking the element in System Explorer and then clicking the **Properties** tab.
2. In the custom name box, enter a name.

Keep in mind the following:

- The name must contain 1 to 64 characters.
- The following characters and symbols are accepted: letters, numerals (0 to 9), ~, @, \*, \_ , +, . , < > , ( ) , [ ] , { } , | .
- The name is case sensitive, for example, “Element1” and “element1” are different elements.

3. Click **Save**.

---

## Viewing Element Topology

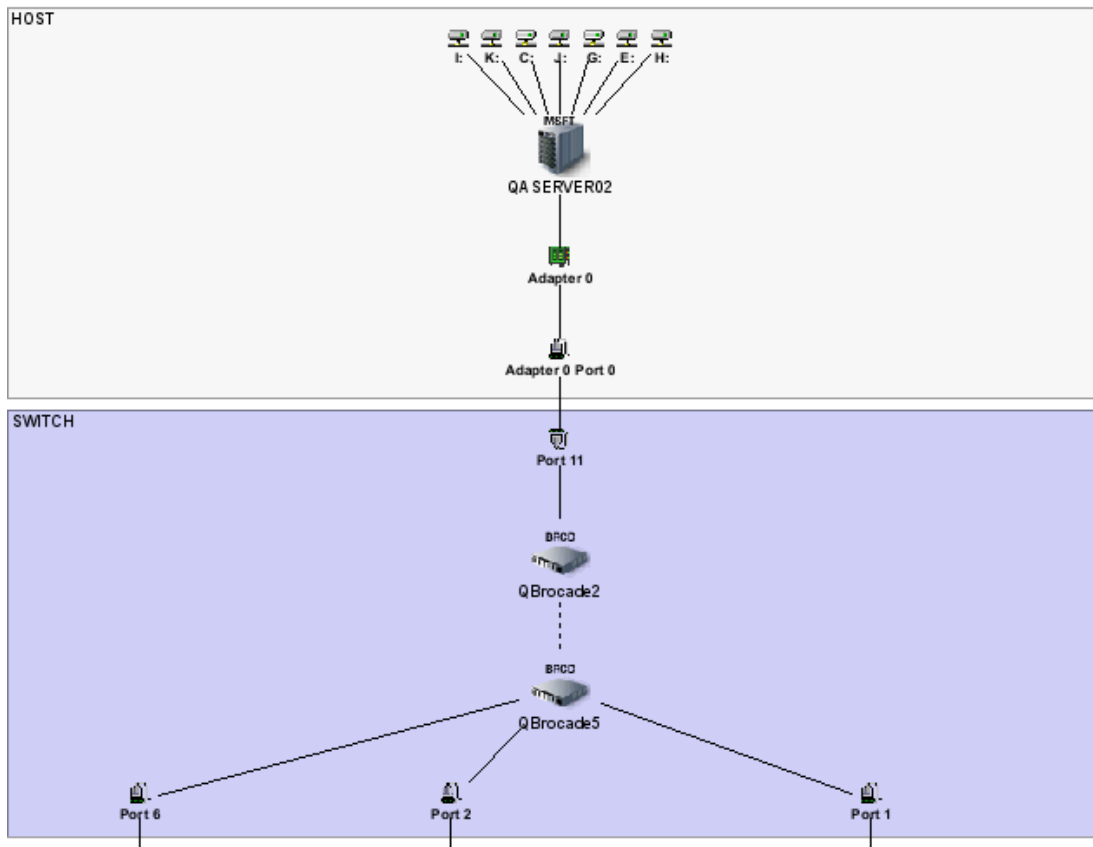
This section contains the following topics:

- “The Topology Tab” on page 379
- “About the New Window Option” on page 387
- “Printing the Topology” on page 388

### The Topology Tab

The Topology tab provides a graphical representation of an element's path. It displays 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: LSI, EMC, and HDS.



**FIGURE 10-25** Topology of a Server

The topology extends the length of the screen. The second portion of the topology is provided by the following figure.

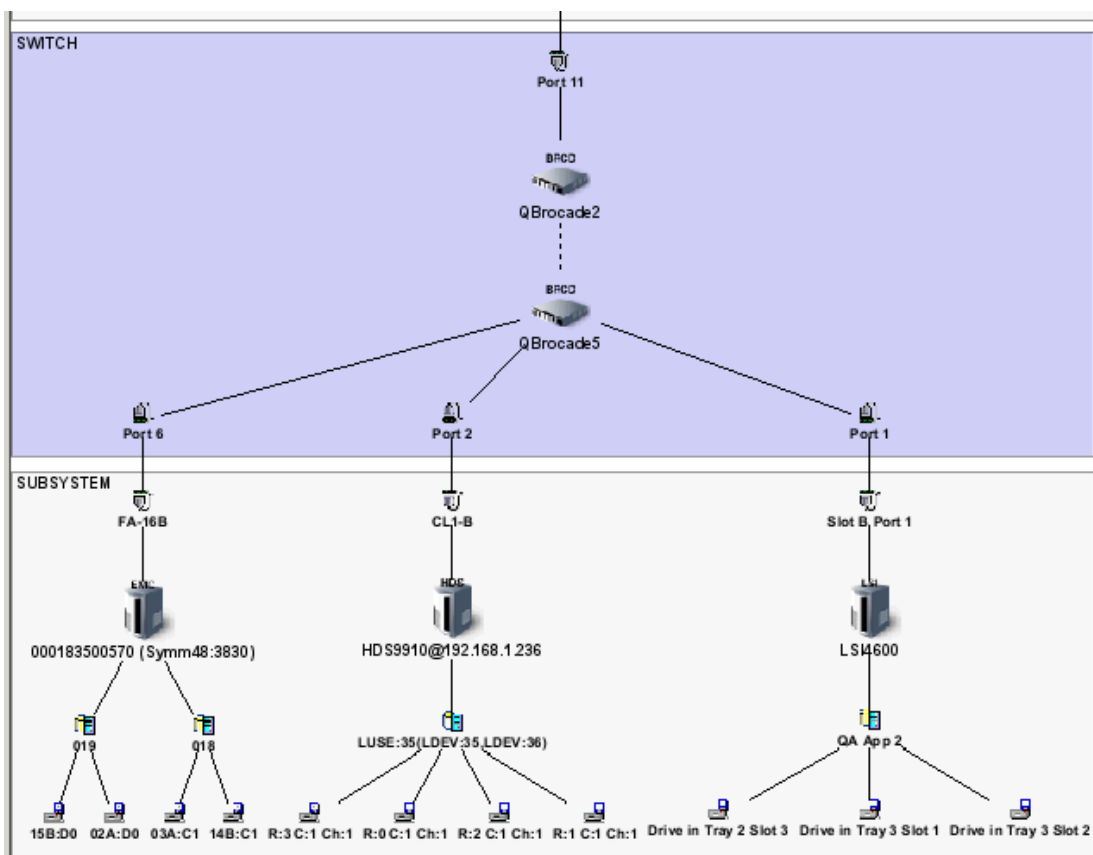


FIGURE 10-26 Topology of a Server (Continued)

**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.

**Note** – Drilling down into EVA VDisks mounted to a host will reveal that the VDisks comprise all disks on the EVA. This reflects the information provided to the management server by the EVA Provider.

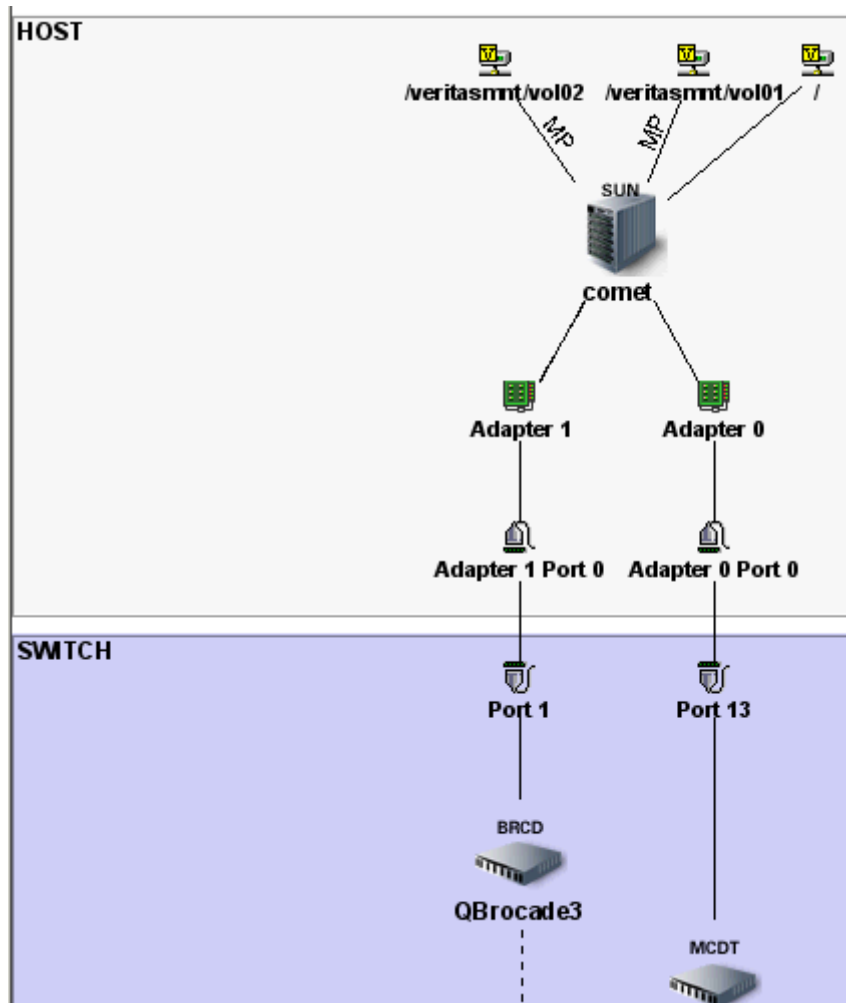
## Multipathing

Multipathing is the process of providing a server more than one path to a storage system, so that in case of an emergency, the server will have continuous access to the storage system. Multipathing can be done many ways. One example of multipathing is providing redundant switches for a host to access a storage system. Another is providing redundant paths from the host to the switch. To determine if your multipathing software is supported, see the support matrix, which is available from the Documentation Center.

Keep in mind the following:

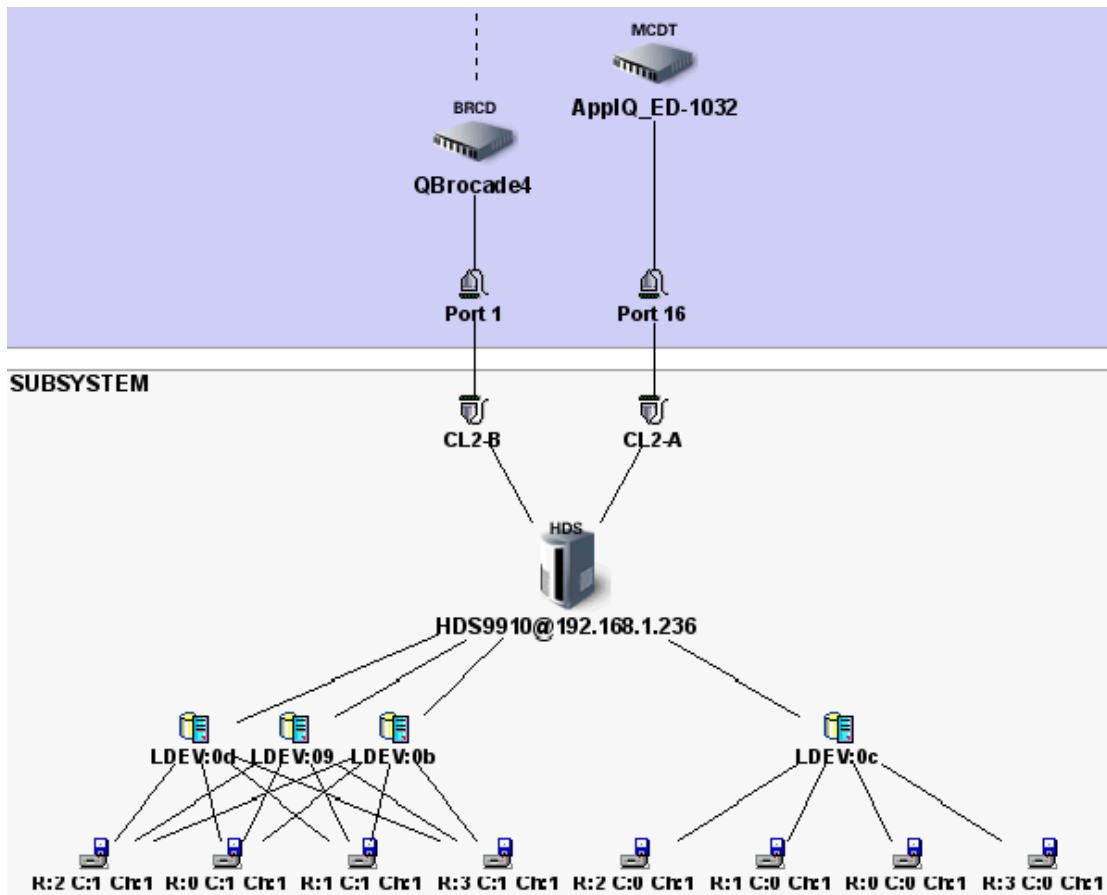
- SANtricity Manager Utilities must be installed on the host running RDAC for the management server to obtain RDAC information.
- HDLM on Sun Solaris requires the storage array to be included in discovery to report the correct information back to the bindings page. See “Known Device Issues” on page 841 for more information about 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 `600507630efe01a800000000000001104:c0t0d0p3`. This makes it difficult to match it up to your SDD path names.
- Elements managed by the QLogic F/O software display a blank value for **Type of MP** when you view multipathing information.

The following figure shows how the software detects multipathing for a server running VERITAS Volume Manager. MP is displayed on the path of the redundant volumes:



**FIGURE 10-27** Multipathing Displayed in the Topology

The topology extends the length of the screen. The following figure displays the second portion of the topology:




**FIGURE 10-28** 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 “Discovering Switches, Storage Systems, NAS Devices, and Tape Libraries” on page 21.
- For multipathing issues regarding certain devices, see “Known Device Issues” on page 841.
- 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 CLARiON 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 10-1, “Feature of the Toolbar in System Explorer,” on page 312 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 10-29** 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 that 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. To access the Topology tab, do one of the following:
  - Select an element in Application Explorer, and then click the **Topology** tab.

---




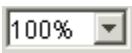




**Note** – The Topology tab is not available for clustered file servers.

---



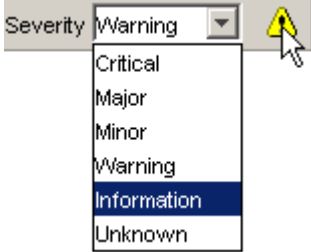
- Double-click an element in Capacity Explorer or System Explorer, and then click the **Topology** tab.

Right-click an element in System Explorer, and then select **Go To Element Topology** from the menu. Table 10-16, “The Toolbar in the Topology Tab,” on page 386 describes the icons on the toolbar.


**TABLE 10-16** The Toolbar in the Topology Tab

Icon	Description
	Prints the topology. See “Printing the Topology” on page 388 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. See “Using the Global View” on page 337.
	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 10-16** The Toolbar in the Topology Tab (*Continued*)

Icon	Description
	<p>Lets you find an element in the topology by name or by Worldwide Name. Begin entering the information; the management server highlights the elements that match.</p> <p>After you populate the search box, click the  button or press ENTER.</p> <p>To expand the Search box, close the left pane. See the topic, “Opening and Closing the Left Pane” on page 8 for more information.</p>
	<p>Displays the event severity icons for the elements displayed in the topology. This feature is disabled for Performance Explorer and Capacity Explorer.</p> <p>See “Viewing Event Status in the Topology” on page 349.</p>

## 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 opens. Use this window to view another section of the topology.

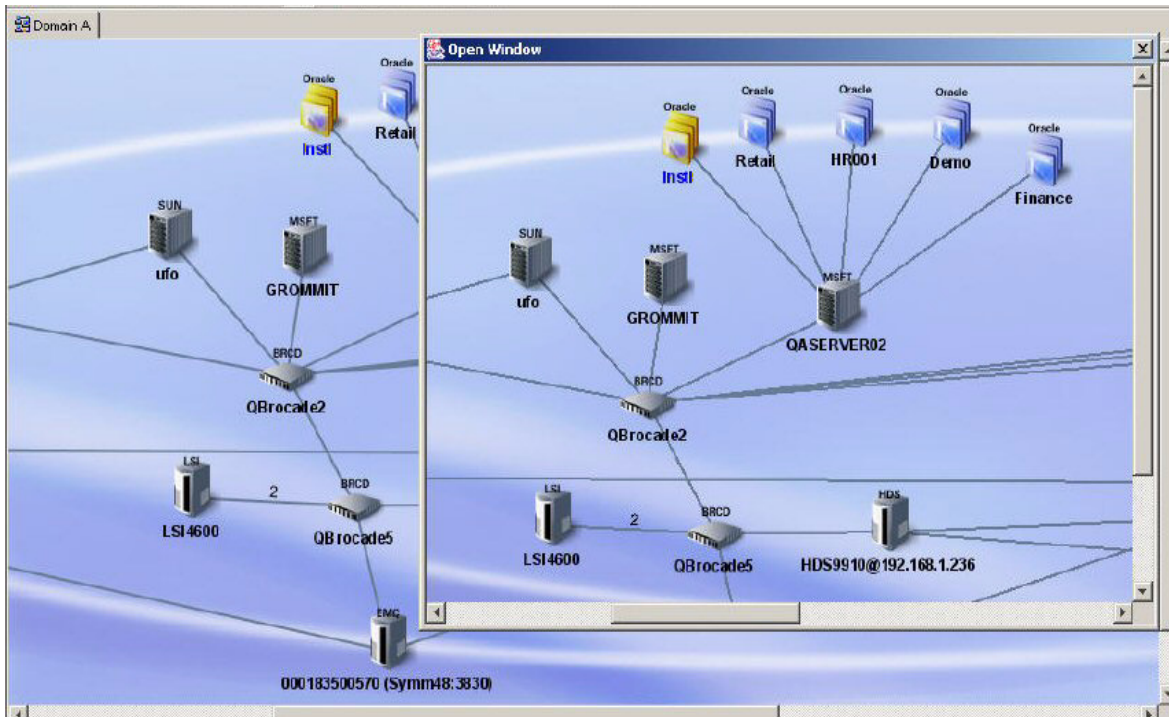




FIGURE 10-30 New Window Option

## 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. To access the Topology tab, do one of the following:
  - Select an element in Application Explorer, and then click the **Topology** tab.
  - Double-click an element in Capacity Explorer or System Explorer, and then click 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 provides a more compact printout.
4. Click the  button.

5. The Paper tab shows the page setup. If you want the default settings, click **Default**. You can modify the following settings:

---

**Caution** – Before you change the margins, decide on a unit of measurement.

---

- **Paper format** - Select the paper size from the menu.
- **Unit** - Select `cm` (centimeters) or `inch` for the margins.
- **Paper width** - To modify the width of the paper, select the Custom option in the Paper format menu.
- **Paper height** - To modify the measurement in this box, select the **Custom** option in the Paper format menu.
- **Top margin** - Enter a measurement.
- **Bottom margin** - Enter a measurement.
- **Left margin** - Enter a measurement.
- **Right margin** - Enter a measurement.
- **Orientation** - Click an orientation for the printout.

A preview of the printout is displayed in the right pane.

6. When you are done, click **Apply**.
7. To see how the printout will appear on the page, click the **View Selection** tab. If you want the default settings, click **Default**. You can modify the following settings:
- **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 **Trimmed**.

A preview of the printout is displayed in the right pane.

8. When you are done, click **Apply**.
9. The Pages tab shows how many pages the printout will use. If you want the default settings, click **Default**. You can modify the following settings:

---

**Caution** – 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** - Same as in step 6.

**Start y** - Same as in step 6.

**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.

A preview of the printout is displayed in the right pane.

10. When you are done, click **Apply**.

11. To preview your pages, click the **Preview** tab. Then click the page you want to preview.

The page appears in the right pane.

12. When you are ready to print, click **Print**.

13. Click **Close**.

---

**Note** – To return 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 “Get Details” on page 91.
3. Select **Add Virtual Application** from the menu.
4. Enter 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**.

---

## The Provisioning Tab

The provisioning tab provides different functionality, depending on the type of element you double-click in System Explorer or click 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 manage zones, zone aliases, and zone sets. These tools provide a wide range of functionality, such as the following:

- “Creating a Zone Alias” on page 407
- “Creating a Zone in a Fabric” on page 410
- “Creating a Zone Set” on page 413
- “Activating a Zone Set” on page 417

For more information about setting up zones, see “SAN Zoning Overview” on page 401.

If you double-click 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 428
- “Managing Volumes” on page 431
- “Rules for Creating Host Security Groups” on page 441
- “Managing Host Security Groups” on page 446

---

## 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, for example, a device connected to a Brocade switch has gone off-line. The Events tab provides the following information about the events:

- **ID**- The identification number assigned to the event
- **Severity** - 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 611.
- **Clear Events** - See “Clearing Events” on page 613.
- **Delete Events** - See “Deleting Events” on page 615.
- **Sort Events** - See “Sorting Events” on page 615.
- **Select a Severity for Filtering** - See “Setting up a filter” on page 620 and “” on page 607.

To view all events, click the **Event Manager** button in the left pane. See “About Event Manager” on page 605.



---

# Asset Attributes of an Element

---

**Caution** – Depending on your license, Chargeback may not be available. To determine if you have access to Chargeback, see the List of Features, which is accessible from the Documentation Center (**Help > Documentation Center**).

---

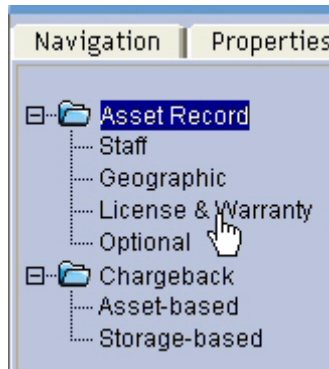
Chargeback provides a handy way for you to keep track of your asset information for an element. You can easily store warranty and licensing information, as well as contact information for the element. For example, assume a switch on the network is having some problems, and you want to contact the person in charge of that switch. You can use the element's asset record to find both the contact information for that switch and the location of the switch.

To access asset information for an element, do one of the following:

- Click an element (except a file server) in Application Explorer, and then click the **Asset Management** tab.
- Double-click an element in Capacity Explorer, Performance Explorer, or System Explorer, and then click the **Asset Management** tab.
- Click a discovered host in Protection Explorer, and then click **Chargeback** in the lower-left corner.
- Click an element in Chargeback.

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 713
- “Adding General Information” on page 714
- “Adding Staff Information” on page 715
- “Adding Geographic Information” on page 716
- “Adding Licensing and Warranty Information” on page 716
- “Adding Custom Information” on page 716



**FIGURE 10-31** 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 720 and “Setting Up Storage-Based Chargeback” on page 726.

The boxes on the Asset Management tab are as follows. When you are done with adding information on this page, click the **Save Changes** button at the bottom of the page.

---

**Note** – The boxes that accept input cannot contain more than 250 characters.

---

- **Custom Name** - A name you assign to the element. See “Assigning a Custom Name” on page 378 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 menu.
  - **New** - This is the default category for all detected elements.
  - **Missing** - The element is 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 This section contains the following topics:” on page 709 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. To access the **Collectors** page, do one of the following:
  - Click an element in Application Explorer, and then click the **Collectors** tab. (For file servers, click the **Scan Schedule** tab).
  - Double-click an element in Capacity Explorer, Performance Explorer, or System Explorer, and then click the **Collectors** tab.
  - Click an element in Chargeback, and then click the **Collectors** tab.
  - Click a discovered host in Protection Explorer, and then click **Collectors** in the lower-left corner.
2. To change a collector's start time, modify the time and date entered in the Next Scheduled Run box. If you decide to change the start time, make sure the date is in yyyy-mm-dd format and the time in 24-hour format. There should be a space between the date and the time, as shown:

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, enter the number of minutes in the Interval box.
 

**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 **Start**.
5. To stop a collector, click **Stop**.
6. To view a report, click its link. See “Viewing Reports” on page 523 for more information.

---

## About the Monitoring Tab

You can easily access performance information about an element:

1. Double-click 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.

See “Viewing Performance Data” on page 631 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 681 for more information:

- Add Policies
- Test Policies
- Edit Policies
- Delete Policies

To access the Policies tab, do one of the following:

- Double-click an element in Capacity Explorer, Performance Explorer, or System Explorer, and then click the **Policies** tab.
- Right-click an element in Capacity Explorer, Performance Explorer, or System Explorer, and then select **Show Policies** from the menu.
- Click a discovered host in Protection Explorer, and then click **Policies** in the lower-left corner.

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 or in Capacity Explorer.

To use the Navigation to determine if a host is part of a file system:

1. Access System Explorer as described in “Accessing System Explorer” on page 310.
2. Double-click the host.
3. Click the **Navigation** tab.
4. Click **Storage Volumes**.

The system type, such as CXFS, is listed in the File System Type column. The following information about the storage volume is also provided:

- Name of the storage volume
- Description of a storage volume
- Drive Type

To use Capacity Explorer to determine if a host is part of a file system:

1. Access Capacity Explorer as described in “Accessing Capacity Explorer” on page 666.
2. Select the host.
3. Scroll to the bottom of the page.

If a storage volume is a member of a shared file system, such as CXFS or XFS, it is listed 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/output into the

metadata server into `/folder`, only the metadata server is able to monitor the file system. For example, if the metadata server generates a 100-KB write, the management server displays a 0-KB write for the `/folder` on the metadata client.

The information in the `/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 the `/folder` on the metadata client.



**FIGURE 10-32** CXFS File System

# Provisioning

---

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 contains the following topics:

- “About Provisioning” on page 399
- “Managing Zones” on page 400
- “Managing Storage” on page 421

---

## 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. For more information, see “About Path Provisioning” on page 463.
- **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. For more information, see “SAN Zoning Overview” on page 401.

- **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. For more information, see “Setting Up Storage Partitioning” on page 422.

---

**Caution** – Ports designated as an Initiator on a storage system belonging to the HDS Freedom Storage™ Lightning 9900™ Series or the 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.

---

Once you have become adept at provisioning, you might want to try accessing the provisioning screens using one of the following methods:

- Double-click a storage system or switch in System Explorer and then click the **Provisioning** tab.
- Right-click a storage system or switch in the Access tab in System Explorer. For more information, see “The Access Tab” on page 317.

## 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 switch. For more information, see “Discovering Switches, Storage Systems, NAS Devices, and Tape Libraries” on page 21.

---

## Managing Zones

This section contains the following topics:

- “SAN Zoning Overview” on page 401
- “Accessing Information About Zone Aliases” on page 406
- “Creating a Zone Alias” on page 407
- “Modifying a Zone Alias” on page 408
- “Deleting a Zone Alias” on page 409
- “Accessing Information About Zoning” on page 409
- “Creating a Zone in a Fabric” on page 410
- “Adding and Removing Zone Members” on page 411
- “Deleting a Zone” on page 412
- “Accessing Information About Zone Sets” on page 412
- “Creating a Zone Set” on page 413
- “Modifying a Zone Set” on page 414
- “Deleting a Zone Set” on page 415
- “Copying a Zone Set” on page 416
- “Activating a Zone Set” on page 417



- “Zones and Zone Sets Listed Twice” on page 418
- “Changing the Amount of Information Collected from the Inactive Zone Database (Cisco Switches)” on page 419

## SAN Zoning Overview

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.

### *Uses of Zones*

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 of the same zone. If two elements are not in 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; or you might 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. If you do not want users in the Production and Finance zones to have access to the same storage, the two zones must be in two different zone sets, both of which must be 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, which allows for flexibility. For instance, in our example, the Finance zone could be in both an active zone set and an inactive zone set. Assume that the Finance zone is a member of an active zone set named Zone Set One and also a member of an inactive zone set named Zone Set Two, and that Zone Set Two contains additional zones. If you activate Zone Set Two, users can be aware of those additional elements and still have access to the Finance zone (because it is a member of Zone Set Two).

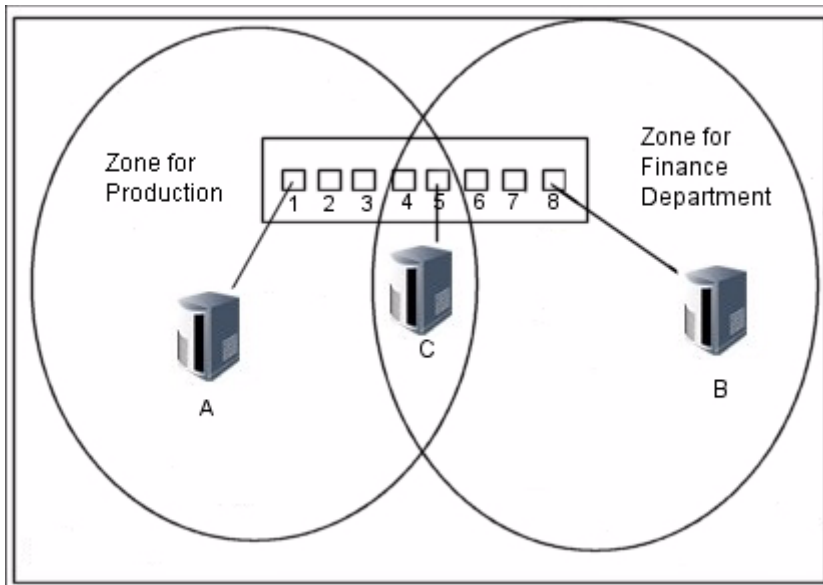
You can create zone aliases to keep track of your zones easily. Instead of having to remember a port's name, you can assign a name that is easy to remember. As a best practice, a zone should contain either zone aliases or ports, but not both.

### *Types of Zoning*

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 figure below 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 Finance. Users in Finance can access storage systems B and C but not storage system A. Likewise, users in Production can access storage systems A and C, but not storage system B.



**FIGURE 11-1** Resources in Two Zones

The following figure provides an overview of zoning structure:



**FIGURE 11-2** Overview of Zoning Structure

### *Zoning Structure*

- **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 WWN of the port to which an element is connected. The WWN of the port can be either the WWN of a switch port or the WWN of the connected element.

Use Table 11-1, “Setting Up Zoning,” on page 403 as a guideline for setting up zoning.

**TABLE 11-1** Setting Up Zoning

Step	Action	Description	For More Information
1	Create a zone alias	<i>Optional:</i> Zone aliases are used to give meaningful names to switch ports, HBA ports, or storage system ports.	“Creating a Zone Alias” on page 407
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 410
3	Create zone sets	A zone set contains multiple zones.	“Creating a Zone Set” on page 413
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 417

Keep in mind the following:

- If you use another product to make zoning changes (such as adding a zone) you must perform Get Details in order to make the management server aware of these changes.
- The management server creates a zone by finding the port or WWN of discovered elements. The management server cannot create a zone using Fibre Channel addresses. If you use a third-party software to create a zone using Fibre Channel addresses, the active zone will appear empty in the user interface of the management server.

### *Activities Supported by Zoning*

Table 11-2, “Zoning Support,” on page 404 contains information about the options supported for each type of switch . For additional information on specific switch types, refer to the paragraphs following this table.

**TABLE 11-2** Zoning Support

Switch Type	Capability					
	View Active Zones	View Inactive Zones	View Aliases	Zone/ Zone Set Provisioning <sup>1</sup>	Zone Set Copying	Port Statistics
Brocade	Y	Y	Y	Y	N	Y
McDATA SWAPI to EFCM <sup>2</sup>	Y	Y	N	Y	Y	Y
McDATA SNMP through proxy <sup>2</sup>	N	N	N	N	N	Y
McDATA SNMP to switches <sup>2</sup>	Y	N	N	N	N	Y
Cisco SNMP	Y	N	Y	N	N	Y
Cisco SMI-S <sup>5</sup>	Y	Y	Y	Y	N	Y
CNT SMI-S	Y	Y	N	N	N	Y
QLogic SNMP <sup>4</sup>	Y <sup>3</sup>	N	N	N	N	Y
QLogic SMI-S <sup>4</sup>	Y	Y	Y	Y	N	Y

<sup>1</sup>The ability to create, modify, and remove zone aliases, zones, and zone sets.

<sup>2</sup>Also applies to EMC Connectrix switches.

<sup>3</sup>Also applies to Sun StorEdge SNMP switches.

<sup>4</sup>Applies to SANbox-2 only.

<sup>5</sup>Soft zoning is only supported for Cisco SMI-S switches. For more information, see “soft zone” on page 851 in the Glossary. Hard zoning is not supported on Cisco SMI-S switches.

### *Issues for Sun StorEdge and QLogic SNMP 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 the name `qlogic1:DeviceName` is displayed in 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 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 they sets reside on. For example, assume you have the name `VSANNAME:cisc01:DeviceName` displayed in the Name column.
  - `VSANNAME` is the name of the virtual SAN in which the zone alias, zone, or zone set was created
  - `cisc01` 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. However, 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.

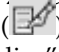

Keep in mind the following:

- For more information about which zoning features are supported for your switches, see Table 11-2, “Zoning Support,” on page 404.
- *Cisco SMI-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 access information about zone aliases in a fabric:

1. Click **Provisioning** (  ).
2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for 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.
- **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. For more information, see “Modifying a Zone Alias” on page 408.
- **Delete** - Click the **Delete** (  ) button for the zone alias you want to delete. For more information, see “Deleting a Zone Alias” on page 409.

To create a zone alias, click **New Zone Alias**. For more information, see “Creating a Zone Alias” on page 407.

# Creating a Zone Alias

Zone aliases are used to give meaningful names to switch ports, HBA ports, or storage system ports.

To create a zone alias:

1. Click **Provisioning** (  ).
2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for the fabric in which you want to create a zone alias.
4. Click **Step 1 Zone Alias**.
5. Click **New Zone Alias**.
6. *Cisco SMI-S Switch Connections Only:* Select the virtual SAN in which you want to create the zone alias from the **VSAN** menu.  
Only the ports in the VSAN you selected are displayed in the Potential Ports pane.
7. *Cisco SMI-S and QLogic SMI-S Switch Connections Only:* Select the switch on which you want to create the zone alias from the **Switch** menu.
8. In the Zone Alias Name box, enter a name for the zone alias. For more information, see “Zone Naming Conventions” on page 407.
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 **Remove From Zone**.
11. Click **OK**.

The zone alias is created.

## Zone Naming Conventions

The following naming conventions apply to zones, zone sets, and zone aliases:

### Naming Conventions for Brocade Switches:

- Names can have a maximum of 64 characters.
- Names must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or an underscore (\_).

- Names are case-sensitive. For example, Zone1 and zone1 are considered to be 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 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:**



- Names can have a maximum of 64 characters.
- Names must begin with a letter.
- Names 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.
- The first character must be alphanumeric (a-z or A-Z).
- Names are case-sensitive.
- Names can include up to 64 alphanumeric characters including caret (^), underscore (\_), and dollar sign (\$).

## Modifying a Zone Alias

To modify a zone alias:

1. Click **Provisioning** (  ).
2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for the fabric in which you want to modify the ports.
4. Click **Step 1 Zone Alias**.
5. Click the **Edit** (  ) button.
6. Take one of the following actions:
  - To add ports, select a port in the **Potential Ports** pane.
  - To remove ports, select the ports in the **Ports in the Zone Alias** pane you want to remove, and then click **Remove From Zone**.



---

**Note** – To select all of the ports, select the check box next to the Port heading.



---

7. Click **OK**.

## Deleting a Zone Alias

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. Click the **Provision** button for the fabric in which you want to delete a zone alias.
4. Click **Step 1 Zone Alias**.
5. Click the **Delete** (  ) button for 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:

- For more information about which zoning features are supported for your switches, see Table 11-2, “Zoning Support,” on page 404.
- *Cisco SMI-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 manage them:

1. Click **Provisioning** (  ).
2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for 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:

(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. For more information, see “Adding and Removing Zone Members” on page 411.
- **Delete** - Click the **Delete** () button for the zone you want to delete. For more information, see “Deleting a Zone” on page 412.

To create a zone, click **New Zone**. For more information, see “Creating a Zone in a Fabric” on page 410.

## 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 the only support for Cisco SMI-S switches. For more information, see “soft zone” on page 851 in the Glossary. Hard zoning is not supported on Cisco SMI-S switches.
- For more information about which zoning features are supported for your switches, see Table 11-2, “Zoning Support,” on page 404.

To create a zone:

1. Click **Provisioning** ()

2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for the fabric in which you want to create a zone.
4. Click **Step 2 Zone**.
5. Click **New Zone**.
6. *Cisco SMI-S Switch Connections only*: Select the virtual SAN in which you want to create the zone from the **VSAN** menu.  
Only the ports in the VSAN you selected are displayed in the Potential Members pane.
7. *Cisco SMI-S and QLogic SMI-S Switch Connections only*: Select the switch on which you want to create the zone from the **Switch** menu.
8. In the Zone Name box, enter a name for the zone. For more information, see “Zone Naming Conventions” on page 407.
9. Add members to the zone by selecting a member in the **Potential Members** pane.  
Keep in mind the following:
  - A zone member can be a port attached to a switch, a WWN or a zone alias.
  - 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. Click **OK**.


## Adding and Removing Zone Members


---

**Caution** – 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. Click the **Provision** button for the fabric in which you want to modify a zone.
4. Click **Step 2 Zone**.



5. Click the **Edit** () button.
6. Take one of the following actions:
  - To add a member to the zone, select a member in the Potential Members pane. A zone member can be a port attached to a switch, a WWN or a zone alias.
  - To remove from members from the zone, select the members in the Zone Members pane and click Remove From Zone.
7. Click **OK**.

## Deleting a Zone

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, and then delete it.

If you are using EFC Manager to delete zones, see “Changes in EFC Manager Requiring Get Details” on page 840.

To delete a zone:


1. Click **Provisioning** () .
2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for the fabric in which you want to delete a zone.
4. Click **Step 2 Zone**.
5. Click the **Delete** () button for 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:

- For more information about which zoning features are supported for your switches, see Table 11-2, “Zoning Support,” on page 404.
- *Cisco SMI-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:

1. Click **Provisioning** (  ).
2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for the fabric in which you want to access information about a zone set.
4. Click **Step 3 Zone**.
5. *Cisco SMI-S Switch Connections only:* Select the virtual SAN in which you want to view zone sets from the **VSAN** menu.

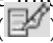

This page lists information about zone sets under the following columns:

- **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. For more information, see “Activating a Zone Set” on page 417.
- **Edit** - Click the **Edit** (  ) button to edit a zone set. For more information, see “Modifying a Zone Set” on page 414.
- **Delete** - Click the **Delete** (  ) button for the zone set you want to delete. For more information, see “Deleting a Zone Set” on page 415.

To create a zone set, click **New Zone Set**. For more information, see “Creating a Zone Set” on page 413.

## 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. Click the **Provision** button for the fabric in which you want to create a zone set.
4. Click **Step 3 Zone Set**.
5. *Cisco SMI-S Switch Connections only:* Select the virtual SAN in which you want to view zone sets from the **VSAN** menu.
6. Click **New Zone Set**.
7. *Cisco SMI-S Switch Connections only:* Select the virtual SAN in which you want to create the zone set from the **VSAN** menu.  
Only the ports in the VSAN you selected are displayed in the Potential Members pane
8. *Cisco SMI-S and QLogic SMI-S Switch Connections only:* Select the switch on which you want to create the zone set from the **Switch** menu.
9. In the Zone Set Name box, enter a unique name for the new zone set. For more information, see “Zone Naming Conventions” on page 407.
10. Take the desired action:
  - To make a zone set active, select the **Activate this Zone** option. Note that only one zone set can be active at a time; when you make a zone set active, the previous zone set becomes inactive.
  - To add zones to the zone set, select a zone in the **Zones Not in this Zone Set** list and click the greater than sign (>). Note that a zone can be in multiple zone sets.
  - To remove zones from the zone set, select a zone in the **Zones in this Zone Set** list and click the less than sign (<).
11. 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. Click the **Provision** button for 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 SMI-S Switch Connections only*: Select the virtual SAN in which you want to view zone sets from the **VSAN** menu.
  6. Click the **Edit** () button.
  7. Take the desired action:
    - To add zones to the zone set, select a zone in the **Zones Not in this Zone Set** list and clicking the greater than sign (>). Note that a zone can be in multiple zone sets.
    - To remove zones from the zone set, select a zone in the **Zones in this Zone Set** list, and click the less than sign (<).
  8. Click **OK**.


## Deleting a Zone Set


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 840.

To delete a zone set:

1. Click **Provisioning** ()
2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for the fabric which contains the zone set you want to delete.
4. Do one of the following:
  - Click **Step 3 Zone Set**.
  - Click **Step 4 Activate Zone Set**.
5. *Cisco SMI-S Switch Connections only*: Select the virtual SAN in which you want to view zone sets from the **VSAN** menu.

6. Click the **Delete** () button for 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 it finds one of the following:

- An inactive zone set that has 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.


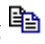
Active zones in a zone set have corresponding inactive zones for redundancy. If you attempt to copy an active zone set containing a zone that does not have a corresponding inactive zone, the management server creates an inactive zone with the same name as the active zone. The inactive zone is used as a backup for the active zone.

---

**Caution** – This feature is supported only for switches that support zone set copying. Refer to Table 11-2, “Zoning Support,” on page 404 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. Click the **Provision** button for 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. For more information, see “Zone Naming Conventions” on page 407.
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. When you make a zone set active, the previous active zone set becomes inactive. However, you could have a zone in more than one zone set.

Keep in mind the following:

- *Cisco SMI-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.
- For more information about which zoning features are supported for your switches, see Table 11-2, “Zoning Support,” on page 404.

To make a zone set active:

1. Click **Provisioning** (  ).
2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for the fabric in which you want to activate a zone set.
4. Click **Step 4 Activate Zone Set**.
5. *Cisco SMI-S Switch Connections only:* Select the virtual SAN in which you want to view zone sets from the **VSAN** 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:
  - a. Select the option, **Make a backup copy of the active zone set after activation**.
  - b. *Optional:* In the Name box, 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 in the format *yyyy-mm-dd*. In the example, the date is May 17, 2005.

13-41-05 is the time the copy was made in the format *hh-mm-ss*, using 24 hour notation. 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 naming conventions described in “Zone Naming Conventions” on page 407.

---

**Note** – The management server truncates the name of the backup zone set if it is more than 44 characters long so that it can fit the date and time into the zone set name, which cannot be more than 64 characters. For example, if 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.

---

8. To activate the zone, click **OK**.

## Zones and Zone Sets Listed Twice

Sometimes the Navigation tab and the Provisioning pages list the zones and zone sets twice for McDATA and Connectrix switches. This is because EFC Manager and Connectrix Manager contain an offline 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 must then 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)

```

Note that ZoneSetB and its members are listed twice in the previous example. They are also displayed twice on the Navigation tabs and Provisioning pages. For example, If you click **Zone A1** in EFC Manager or Connectrix Manager, you are shown that it has member ZoneMember A1a. 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)

For Cisco switches discovered through SNMP, you can change how much information the management server collects from the inactive database during Get Details. 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 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. Select **Configuration > Product Health**, and 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**, and then click **Advanced** in the **Disk Space** tree).
5. Paste the copied text into the Custom Properties box. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the box and select **Paste**.
6. In the Custom Properties box, change the value of `cimom.cisco.displayZoneDatabase`. 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**.

## About the Messages Displayed in the Brocade Console

You can ignore the following messages in the Brocade Console. 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
```

To view the status of the switch, access the Brocade Console by typing the user name and password for the switch in a Telnet window, and then enter `switchshow` at the prompt

---

## Managing Storage

This section contains the following topics:

- “Setting Up Storage Partitioning” on page 422
- “Modifying the Cache Settings (LSI and Sun 6130)” on page 427
- “Changing the Owner of a Volume (LSI, CLARiiON and Sun 6130)” on page 428
- “Managing Storage Pools” on page 428
- “Managing Volumes” on page 431
- “Rules for Creating Host Security Groups” on page 441
- “Managing Host Security Groups” on page 446
- “General Provisioning Issues” on page 454
- “Provisioning Issues by Vendor” on page 454

# 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 LSI 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 depending on 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 441.

For details about setting up storage partitioning, see “How to Set Up Storage Partitioning” on page 426

To learn which storage systems are supported, see Table 11-3, “Provisioning and Pool Support,” on page 422 and Table 11-4, “Volume and Host Security Group Support,” on page 423.

**TABLE 11-3** Provisioning and Pool Support

Storage System	Storage Provisioning	Create/Delete Pool	Create Pool Using Settings <sup>1</sup>	Additional Information
3PAR	N	N	N	
CLARiiON	Y <sup>1</sup>	Y	Y	See “Additional Information About CLARiiON Storage Systems” on page 425
LSI and Sun 6130	Y	Y	Y	
HDS	Y	N	N	
HP EVA	Y	Y	Y	
HP MSA	Y	N	N	Selective Storage Presentation (SSP) for the array must be enabled for provisioning to work.
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	Y	N	N	

**TABLE 11-3** Provisioning and Pool Support (*Continued*)

Storage System	Storage Provisioning	Create/Delete Pool	Create Pool Using Settings <sup>1</sup>	Additional Information
IBM DS and ESS	Y	N	N	See “Additional Information About IBM DS and ESS” on page 425
Sun 35xx	N	Y	Y	
Sun 6920 and 6940	Y	N	N	
Symmetrix	Y	N	N	
Xiotech	Y	N	N	

<sup>1</sup>The “Create Pool Using Settings” column refers to the functionality that lets you choose the type of pool, usually RAID level.

For information regarding support for volumes and host security groups, see Table 11-4, “Volume and Host Security Group Support,” on page 423.

**TABLE 11-4** Volume and Host Security Group Support

Storage System	Create/Delete Volume	Create Volume Using Settings <sup>1</sup>	Create/Delete Meta Volume	Host Security Group Provisioning Supported	Additional Information
3PAR	N	N	N	N	
CLARiiON	Y	Y	Y	Y	RAID level can be specified for first volume in a pool, subsequent volumes inherit this setting.
LSI and Sun 6130	Y	Y	N	Y	
HDS	Y	N	Y	Y	Creating a volume means marking an existing device as accessible to host security group management, deleting a volume means returning the device to the free device pool.
HP EVA	Y	Y	N	Y	
HP XP with Command View XP	N	N	N	N	

**TABLE 11-4** Volume and Host Security Group Support (*Continued*)

<b>Storage System</b>	<b>Create/ Delete Volume</b>	<b>Create Volume Using Settings<sup>1</sup></b>	<b>Create/ Delete Meta Volume</b>	<b>Host Security Group Provisioning Supported</b>	<b>Additional Information</b>
HP XP with Command View XP Advance Edition	Y	N	Y	Y	
HP XP with XP Provider	Y	N	Y	Y	Creating a volume means marking an existing device as accessible to host security group management, deleting a volume means returning the device to the free device pool.
HP MSA	Y	N	N	Y	
IBM DS	Y	N	N	Y	See “Additional Information About IBM DSS and IBM ESS” on page 426
IBM ESS	Y	N	N	Y	Volumes created on IBM storage systems cannot be deleted. See “Additional Information About IBM DSS and IBM ESS” on page 426
Sun 35xx	N	N	N	Y	
Sun 6920 and 6940 <sup>5</sup>	Y	N	N	Y	Unlike the native tool for the Sun 6130 storage system, the management server refers to virtual disks as storage pools.
Symmetrix	Y	N	Y	Y	Creating a volume means marking an existing device as accessible to host security group management, deleting a volume means returning the device to the free device pool.



**TABLE 11-4** Volume and Host Security Group Support (*Continued*)

Storage System	Create/Delete Volume	Create Volume Using Settings <sup>1</sup>	Create/Delete Meta Volume	Host Security Group Provisioning Supported	Additional Information
Xiotech	Y	Y	N	Y	See “Additional Information About Xiotech” on page 426

### *Additional Information About CLARiiON Storage Systems*

The 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. For more information, see “EMC Navisphere CLI Is Required” on page 455.

### *Additional Information About HP EVA Arrays*

- To determine provisioning support for HP StorageWorks Arrays, see Table 11-3, “Provisioning and Pool Support,” on page 422 and Table 11-4, “Volume and Host Security Group Support,” on page 423.
- EVA arrays can only be provisioned if they are actively managed by the Command View server that they are discovered through.

When an EVA is discovered by the built-in EVA provider, a cache is created and populated with the current array configuration. Each subsequent cache refresh will start 30 minutes after completion of the previous cache refresh. The time the cache refresh takes depends on factors such as the EVA configuration, model, and SAN traffic.

When you perform a provisioning operation, for example, create, delete, or modify a pool, or volume, the cache information about provisioning is immediately updated. If you provision an EVA using Command View EVA or a different management station, the cached information about the EVA will not be accurate until the cache is refreshed.

### *Additional Information About IBM DS and ESS*

IBM and other storage systems that use external providers show raw capacity for a storage pool, instead of formatted capacity.

<sup>1</sup>The “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.

## *Additional Information About IBM DSS and IBM ESS*

Volume groups are not represented in the provisioning or properties of IBM storage systems. The host security groups are based on a host connection, which does not give any indication as to which volume groups each of them belongs. A Volume group can have many host connections, but only one host connection can be part of a volume group.

If there is a failure in a Wizard Provisioning operation on the IBM ESS F20 array, the failure is not reported to the management server from the IBM ESS management software. The Wizard Provisioning UI hangs, appears busy, and becomes unresponsive.

## *Additional Information About Xiotech*

During a provisioning operating on Xiotech, the Provisioning interface remains busy and eventually times out. You can close the provisioning page after a few minutes and check on the array, which should be updated properly.

For more information about how storage provisioning works on your storage system, see the following topics .

- “Issues Specific to CLARiiON Storage Systems” on page 455
- “Issues Specific to EMC Symmetrix Storage Systems” on page 456
- “Issues Specific to HDS Storage Systems” on page 457
- “Issues Specific to LSI Storage Systems” on page 462

## How to Set Up Storage Partitioning

To set up storage partitioning:

1. *LSI, CLARiiON, Sun 6130 and Sun 35xx only:* Create a storage pool (sometimes referred to as a volume group or RAID group). For more information about storage pools, see “About Storage Pools” on page 426
2. Create a volume. For more information about volumes, especially in CLARiiON storage systems, see “About Volumes” on page 427.
3. Create a host security group. For more information about host security groups, see “About Host Security Groups” on page 427.

## *About Storage Pools*

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. LSI and CLARiiON storage systems require that a storage pool be created and volumes be allocated from the storage pool. For more information, see “Creating a Storage Pool (LSI, CLARiiON, Sun 6130 and Sun 35xx)” on page 429.

### *About Volumes*

A volume is a virtual disk. Volumes are created in sizes that are desirable for being shown as a LUN. 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 volume is owned by one of the storage processors. Creating a volume also creates a LUN for this volume and for each port of the storage processor that owns the volume. 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.

For more information, see “Creating a Storage Volume” on page 434.


### *About Host Security Groups*

Host security groups define which initiators (HBA ports) have access to specified storage volumes. Host security groups are associated with a fibre-channel port and contain a list of HBA port initiators and the volumes they can see.

For more information, see “Creating Host Security Groups” on page 448 and “Rules for Creating Host Security Groups” on page 441.

## Modifying the Cache Settings (LSI and Sun 6130)

To modify the cache settings:

1. Click the **Edit** () button for the volume you want to modify.
2. Enter the cache read ahead multiplier (0 to 65535 bytes) in the Cache read-ahead multiplier box. 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, enter 0.
3. *Optional:* Select **Read Caching**.

Use this option to store the host's operations in controller cache memory.

4. *Optional:* Select **Write Caching**.

Use this option to write data 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** – For information about changing the owner assigned to the volume, see “Changing the Owner of a Volume (LSI, CLARiiON and Sun 6130)” on page 428.

---

## Changing the Owner of a Volume (LSI, 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 controller 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 **Edit** () button for the volume you want to modify.
2. Select the owner for the volume from the **Current Owner** menu.
3. Click **OK**.

## Managing Storage Pools


This section contains the following topics:

- “Creating a Storage Pool (LSI, CLARiiON, Sun 6130 and Sun 35xx)” on page 429
- “Accessing Information About Storage Pools” on page 430
- “Deleting a Storage Pool (LSI and CLARiiON Only)” on page 431

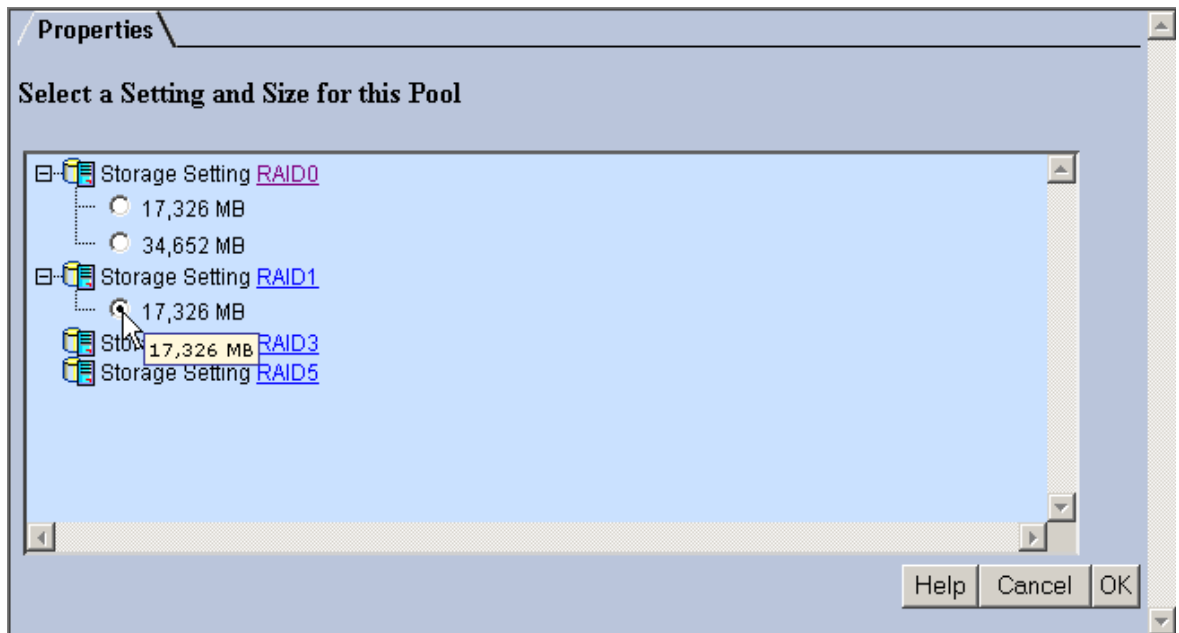
## Creating a Storage Pool (LSI, CLARiiON, Sun 6130 and Sun 35xx)

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. You should 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. Click the **Provision** button for the storage system in which you want to create a storage pool.
4. Click **Step 1 Storage Pool**.
5. Click **New Storage Pool**.
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 11-3** 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 LSI 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. Click the **Provision** button for the storage system in which you want to access information about storage pools.
4. Click **Step 1 Storage Pool**.

The following information about the storage pools is displayed:

- 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. (Available to only HDS storage systems).
  - Delete - Click the **Delete** (  ) button for the storage pool you want to delete. For more information, see “Deleting a Storage Pool (LSI and CLARiiON Only)” on page 431. (Available to only LSI, 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. For more information, see “Creating a Storage Pool (LSI, CLARiiON, Sun 6130 and Sun 35xx)” on page 429.

6. If the table contains more than 10 entries, the following buttons for navigating through the table are enabled:
-  - Move to the first page.
  -  - Move back one page.
  -  - Move forward one page.
  -  - Move to the last page.



## Deleting a Storage Pool (LSI and CLARiiON Only)

---

**Caution** – When you delete a storage pool on an LSI storage system, all the volumes in 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. Click the **Provision** button for the storage system in which you want to delete a storage pool.
4. Click **Step 1 Storage Pool**.
5. Click the **Delete** (  ) button for the storage pool you want to delete.  
A message warns you about the other volumes that will be deleted.
6. Click **OK**. The storage pool and its volumes are deleted.

## Managing Volumes

This section contains the following topics:

- “Accessing Information About Volumes” on page 432
- “Filtering Volumes” on page 433
- “Creating a Storage Volume” on page 434
- “Deleting a Storage Volume” on page 438
- “Changing the Cache Block Size for a Storage System (LSI and Sun 6130)” on page 440
- “Modifying the Cache Settings (LSI and Sun 6130 Only)” on page 440

## Accessing Information About Volumes


---

**Caution** – Some storage vendors require a password to access the storage system. If the correct password is not entered, an authentication error message is displayed. Refer to Get Details to determine which user account 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 433.

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. (This option is available to only HDS storage systems).
- \*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 **Edit** (  ) button for the volume you want to edit. For more information, see “Modifying the Cache Settings (LSI and Sun 6130)” on page 427.
- Delete - Click the **Delete** (  ) button for the volume you want to delete. For more information, see “Deleting a Storage Volume” on page 438.

\*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 **Delete Selected Volumes**.

---

**Note** – If you have an HDS storage system, see “About Volumes on HDS Storage Systems” on page 434.

---

5. If the table contains more than 10 entries, the following buttons for navigating through the table are enabled:
-  - Move to the first page.
  -  - Move back one page.
  -  - Move forward one page.
  -  - Move to the last page.

## Filtering Volumes

To filter the list of volumes displayed:

1. Click the **Show Volume Filter** link to display the filtering options.  
If the volume filter is turned on, the link appears as **Hide Volume Filter**.
2. Enter all or part of a volume name in the Name Contains box.
3. Enter a size in megabytes in the Size ( $\geq$ ) MB box.  
Volumes that are greater than or equal to the size specified appear on the page.
4. Select one of the following:
  - **All Volumes** to display all existing volumes.
  - **Unmapped Volumes** to display only unmapped volumes are displayed.
  - **Mapped Volumes** to display only mapped volumes are displayed.
5. Select the port you want to display from the **Ports** menu.  
If you want to display all ports, select **All** from the **Ports** menu.
6. Select the storage pools you want displayed from the **Storage Pools** menu.
7. Click **Filter**.  
The table is updated to display only the elements that meet the filter criteria.

---

**Note** – To reset the filter criteria, click **Reset**.

---

## *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 that 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, 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). For example:

```
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

---

**Caution** – Some storage vendors require a password to access the storage system. If the correct password is not entered, an authentication error message is displayed. Refer to **Discovery > Details** to determine which user account 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.

On HDS and Symmetrix storage systems, volumes are shipped already created. When you create a volume in the management server on these storage systems, you are defining the volume as being allocated.

## *Rounding Volume Size*

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 required to create the 14-GB LUSE volume. The management server would use three LDEVs because each LDEV is only 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.

### *Support for PvLinks*

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.

### *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 then perform LUN creation on the HDS box, it is a two-step process. First, use the management server to create LUSE volumes. Then create a LUN and map a volume to a target port, by creating a host security group. For more information, see “Creating Host Security Groups” on page 448 and “Rules for Creating Host Security Groups” on page 441.

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. For more information, see “About Provisioning on HDS Storage Systems” on page 458.

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 may be incorrectly reported.

It is not possible to extend LUSEs using the management server software. You must perform this operation using the native tools.

---

**Note** – In a Hitachi Freedom Storage™ Thunder 9500™ V Series storage system, it is not possible to create a LUSE from a group of RAID-0 LDEVs.

---

## *For LSI 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. For more information, see “Creating Host Security Groups” on page 448 and “Rules for Creating Host Security Groups” on page 441.

## *Creating a Storage Volume*

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. Click the **Provision** button for the storage system in which you want to access information about volumes.
  - d. Click **Step 2 Volume**.
  - e. Select the desired number of LDEVs for the LUSE volume then click **Delete Selected Volumes**. Take note of the array group from which you deleted the LDEVs. You need this information to create the LUSE volume.
2. Click **New Volume**.

---

**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. *LSI, CLARiiON, and Sun 6130 only:* In the Volume Name box, enter a name for the volume. If you do not provide a name, the software assigns one.

If you enter a volume name, observe the following conventions:

- 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 LSI storage system.
4. In the Size box, enter the size of the volume in megabytes (MB) or gigabytes (GB). Select the appropriate unit of measurement from the menu to the right of the Size box.

---

**Caution** – The management server creates a volume of at least the size specified. For example, assume you asked the management server to 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 available. Keep in mind that although the management server tries to find free extents that make the volume size as close to the requested size, there is no 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, therefore you must delete two volumes from that storage pool. Make sure those volumes you delete are not being used. For more information, see “Deleting a Storage Volume” on page 438.
- If you do not see any storage pools, verify that you have obtained all element details from the storage system. For more information, see “Discover Switches” on page 34.

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 box. If an LSI storage system is selected, the following information 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, keep the default selection and click **Finish**. Once the settings have been made, you will 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. For more information about how to create a LUN, see “Creating Host Security Groups” on page 448 and “Rules for Creating Host Security Groups” on page 441.

---

8. *LSI only:* If you selected the **Custom** option, do the following:

- a. enter the cache read ahead multiplier (0 to 65535 bytes) in the Cache read ahead multiplier box.

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 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, 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 correct password is not entered, an authentication error message is displayed. Refer to **Discovery > Details** to determine which user account 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. Click the **Provision** button for the storage system in which you want to access information about volumes.
- 4. Click **Step 2 Volume**.
- 5. To display only unmapped volumes, select the unmapped volumes node in the left pane.
- 6. Click the **Delete** (🗑️) button for 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 **Delete Selected Volumes**.

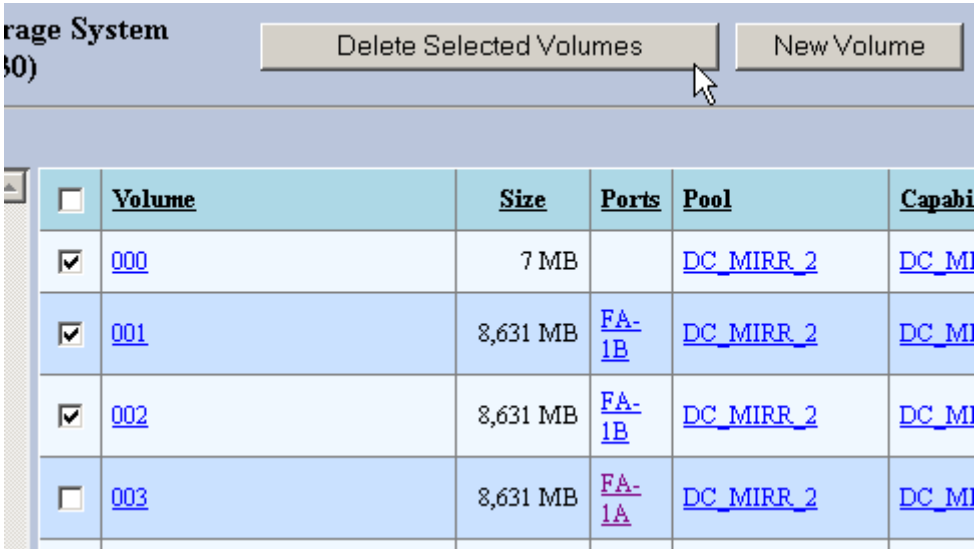


FIGURE 11-4 Deleting Several Volumes at Once

**Note –** To select all volumes, select the check box next to the Volume heading, as shown in the following figure:






FIGURE 11-5 Selecting All Volumes

## Changing the Cache Block Size for a Storage System (LSI and Sun 6130)



LSI and Sun 6130 storage systems let you change the cache block size.

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 **Change**.
4. Select the cache block size from the menu. Set a higher cache size for applications that require a lot of input and output, such as multimedia.
5. Click **OK**.

## Modifying the Cache Settings (LSI and Sun 6130 Only)

To modify the cache settings:

1. Click **Provisioning** () .
2. In the right pane, click the **Storage Systems** tab.
3. Click the **Provision** button for the storage system in which you want to access information about volumes.
4. Click **Step 2 Volume**.
5. Click the **Edit** () button for the volume you want to modify.
6. Enter the cache read ahead multiplier (0 to 65535 bytes) in the Cache read-ahead multiplier box. To disable this option, enter 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 controllers of the same cache size, providing redundancy, so that if one controller fails, the other controller performs the uncompleted write operations.

10. To change the current owner, select a new owner from the **Current Owner** menu.
11. Click **OK**.

## Rules for Creating Host Security Groups

This section contains the following topics:

- “Host Security Groups on EMC CLARiiON and Sun 6130 Storage Systems” on page 442
- “Host Security Groups on LSI Storage Systems” on page 443
- “Host Security Groups on EMC Symmetrix Storage Systems” on page 443
- “Host Security Groups on HDS Storage Systems” on page 444
- “Host Security Groups on HP MSA Storage Systems” on page 444
- “Host Security Groups on HP EVA Storage Systems” on page 444
- “Host Security Groups on IBM Storage Systems” on page 444
- “Host Security Groups on Sun 6920 and 6940 Storage Systems” on page 445
- “Host Security Groups on Xiotech Storage Systems” on page 445

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* refers to 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 host security group depends on the type of storage system:

**TABLE 11-5** Allowed Initiators in Host Security Groups

Storage System	Allowed Initiators in Host Security Groups
EMC CLARiiON	0 or more initiators
LSI	0 or more initiators
HDS	0 or more initiators
HP EVA	0 or more initiators
HP MSA	1 initiator
IBM	Only 1 initiator per host security group
Sun 35xx	1 initiator per host security group
Sun 6920 and 6940	Only 1 initiator per host security group

**TABLE 11-5** Allowed Initiators in Host Security Groups

Storage System	Allowed Initiators in Host Security Groups
Sun 6130	0 or more initiators
Symmetrix	1 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 host security groups for All Ports.

---

Each storage system treats host security groups differently. See the following sections for more information.

### *Host Security Groups on EMC CLARiiON and Sun 6130 Storage Systems*

Keep in mind the following rules for host security groups on EMC CLARiiON and Sun 6130 storage systems

- When a volume is created, by default it 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 host security group is created on all ports of the processor you select unless you select an initiator that uses a different processor and does not belong to a host security group. For example, assume you select processor SPA, and then you select an initiator that belongs to SPB but does not belong to a host security group. The host security group is created for all ports on SPB.
  - Host security groups can consist of initiators (WWNs) only. You do not need to specify volumes. The initiator is shown in both host security groups SPA and SPB.
  - Host security groups 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 the particular host where the initiator is located.

## *Host Security Groups on LSI Storage Systems*

Keep in mind the following rules for host security groups on LSI storage systems.

- When you create a host security group using the management server, the host security group appears as a host, with its volumes and ports displayed underneath the tree in SANtricity.
- When you create a host security group using the management server, you can specify the controller, but not the port, for the host security group.
- In SANtricity, an initiator is equivalent to a host port.
- You can have multiple volumes and initiators in the host security group.
- If you create a volume on a host that does not have multipathing, make sure the volume is on the preferred path. Use SANtricity to make sure the volume is on the preferred path.
- 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 to 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.
- If you create an host security group for a Unix host on an LSI array, you must use SANtricity to change the host mode of the host security group to Linux in order to be compatible with the Unix host. If you do not, the Unix host will show a large number of LUN numbers available. You must use SANtricity to change the host mode because it is not possible to select the host mode when creating an host security group on an LSI array. The host security group will be created with the Windows NT host mode by default if it is not manually changed.

## *Host Security Groups on EMC Symmetrix Storage Systems*

Keep in mind the following rules for host security groups on EMC Symmetrix Storage Systems.

- If LUN security is not turned on for an FA port, all volumes assigned to the FC port are visible to 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 a host security group 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 host security group you created by using the management server. Instead you will see a volume bound to a port and a masked LUN bound to a host in the native tools.
- Host security group is associated with individual ports
- Host security groups 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 host security group, 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 host security group. To add initiators, you need to create the host security group.

## *Host Security Groups on HDS Storage Systems*

Keep in mind the following rules for host security groups on 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 zero to multiple initiators in a host security group
- You can have zero to multiple volumes in a host security group.
- A host security group can be on only 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, the host security groups are listed by LDEV; in the physical view, they are listed by port.
- In the native tool for HDS storage systems, host security groups are referred to as a host security domain.
- For more information about host security group names, see “Host Security Group Names on HDS Storage Systems” on page 458.

## *Host Security Groups on HP MSA Storage Systems*

Keep in mind the following rules for host security groups on HP MSA storage systems.

- When you create a host security group on the MSA, the host mode is set to the value defined by the `smi.ProvisioningHpMsa.hostConnectionProfile` property (the default is Windows). To change the value, follow the instructions in “Customizing Properties” on page 232, and then restart the service.
- You can have one initiator per host security group.
- You can have one or more volumes in a host security group.
- A host security group spans all ports on the array.

## *Host Security Groups on HP EVA Storage Systems*

Keep in mind the following rules for host security groups on HP EVA storage systems.

- You can have multiple initiators per host security group.
- You can have zero to multiple volumes in a host security group.
- A host security group spans all ports on the array.

## *Host Security Groups on IBM Storage Systems*

Keep in mind the following rules for host security groups on IBM storage systems.

- 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.
- 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 453.
- The management server can read the names of host security groups created by the native tool.
- There can be only one initiator per host security group.
- A volume can be assigned to more than one initiator.
- You can select any number of ports from *one* to *all* when creating the host security group.
- You can create the host security group 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 that all ports are in the host security group. If no ports are selected, the default is used.

### *Host Security Groups on Sun 6920 and 6940 Storage Systems*

Keep in mind the following rules for host security groups on 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 host security group.
- You must create the host security group with one initiator and at least one LUN.
- When the management server creates the host security group, the name of the host security group is the initiator WWN.
- The management server can read the names of host security groups created by the native tool.
- Creation of a new host security group with an initiator that already exists in another host security group adds the volumes to the previously created host security group.
- The management server displays the ports connected to the initiator you selected for creating an host security group. 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.

### *Host Security Groups on Xiotech Storage Systems*

Keep in mind the following rules for host security groups on Xiotech storage systems.

- The initiators used in a host security group must be registered with the array.
- Any given host security group 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 contains the following topics:

- “Accessing Information About Host Security Groups” on page 446
- “Creating Host Security Groups” on page 448
- “Editing a Host Security Group” on page 450
- “Deleting a Host Security Group” on page 453

## 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 Fiber 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. For more information, see “General Provisioning Issues” on page 454.
- 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 for the storage system for which you want to access host security group information.
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.

The following information is available on this page:

- Port (If applicable) - Displays 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 **+Filter**, select the port that you would like to view, and click **Apply**. Only the information for the selected port is displayed in the table below.

---

- Name - Displays 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 has not been discovered.
- Volumes - Displays 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 (If applicable) - Displays optional settings on the port that describe how the host accesses the port. Multiple options exist. For more information, see you documentation for HDS storage system.

If the table contains more than 10 entries, the following buttons for navigating through the table are enabled:

-  - Move to the first page.
-  - Move back one page.
-  - Move forward one page.
-  - Move 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 448
- “Editing a Host Security Group” on page 450
- “Deleting a Host Security Group” on page 453
- “Setting the Host Mode for IBM Storage Systems” on page 453


## Creating Host Security Groups

Host Security Groups define which initiators (HBA ports) have access to specified storage volumes. They are associated with a Fiber 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. For more information “Rules for Creating Host Security Groups” on page 441.
- Each type of storage system treats host security groups differently. For more information, see “General Provisioning Issues” on page 454.

To create a host security group you assign a name to the host security group, assign a host mode (if applicable to your system), and then add initiators and volumes to the group:

1. Click **Provisioning** (  ).
2. In the right pane, click the **Storage Systems** tab.
3. Click the **Provision** button for 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. enter a unique name for the host security group in the Name box.

---

**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

LSI 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. (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 host security groups differently. For more information, see “Rules for Creating Host Security Groups” on page 441.


---

3. *HDS storage system*: Click the **Options** button to the right of the Host Mode box, select a host mode resembling the port settings for your environment, and then click **OK**. If your host mode is not listed, enter it in the Host Mode box.
4. *HDS storage systems*: If your storage system supports a second host mode, enter the second host mode in the Second Host Mode box.  
  
A second host mode is an optional setting on the port that describes 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. For more information, see “Setting the Host Mode for IBM Storage Systems” on page 453.
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:
  - enter 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 finished with adding initiators, click **Close**.
5. To remove an initiator from the host security group, click the **Delete** () button. To remove multiple HBA initiators from the list, select the HBA ports you want to remove and then click **Remove Selected**.
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 box at the top of the window.

---

**Note** – For LSI storage systems LUN numbers cannot be duplicated, and that the management server can use an existing LUN number if the access mode for the created LUN is “No Access.”

---


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 **Close**.
5. To remove a volume from the host security group, click the **Delete** () button. To remove multiple volumes from the list, select the volumes you want to remove and then click **Remove Selected**.
6. Click **Finish**.

## Editing a Host Security Group

To edit a host security group:

1. Click **Provisioning** () .
2. In the right pane, click the **Storage Systems** tab.

3. Click the **Provision** button for the storage system in which you want to access information about host security groups.
4. Click **Step 3 Host Security Group**.
5. Click the **Edit** () button for 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 host security groups differently. For more information, see “Rules for Creating Host Security Groups” on page 441.

---

3. *HDS storage system*: Click the **Options** button to the right of the Host Mode box. Select a host mode resembling the port settings for your environment and then click **OK**. If your host mode is not listed, enter it in the Host Mode box.
4. *HDS storage systems*: If your storage system supports a second host mode, enter the second host mode in the Second Host Mode box.

---

**Note** – A second host mode is an optional settings on the port that describes 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. For more information, see “Setting the Host Mode for IBM Storage Systems” on page 453.
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 **Add**, and then do one of the following:
    - enter 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.

Then click **Add**. Click **Close** to exit the window.

- **Delete an initiator** - Click the **Delete** () button for 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

---

**Caution** – 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 **Add**, 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 box at the top of the window.

Then click **Add**. Click **Close** to exit the window.

---


**Note** – For LSI storage systems, LUN numbers cannot be duplicated. The management server can use an existing LUN number if the access mode for the created LUN is “No Access.”

---

---

**Caution** – Removing an HBA can cause hosts that are using it to lose access to their storage. This may result in the loss of data.

---

- **Remove a volume** - Click **Delete** () button for 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 a Host Security Group

---

**Caution** – 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. Click the **Provision** button for the storage system in which you want to access information about host security groups.
4. Click **Step 3 Host Security Group**.
5. Click the **Delete** (  ) button for 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, but can be set by assigning the host mode to the `smi.ProvisioningIbmEss.hostConnectionProfile` property.

To set the host mode for an IBM storage system:

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 `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, right-click the selected text, and 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 box. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the box and select **Paste**.

6. Assign a host mode to  
`smi.ProvisioningIbmEss.hostConnectionProfile` in the Custom Properties box. For example:

```
smi.ProvisioningIbmEss.hostConnectionProfile=AIX
```

where AIX is the host mode.

7. Click **Save**.

## General Provisioning Issues

This section contains the following topics:

- “Provisioning Can Make a Device Inaccessible” on page 454
- “Provisioning Does Not Make an Operating System Aware of a Device” on page 454

### Provisioning Can Make a Device Inaccessible

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. 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 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 contains the following topics:

- “Issues Specific to CLARiiON Storage Systems” on page 455
- “Issues Specific to EMC Symmetrix Storage Systems” on page 456
- “Issues Specific to HDS Storage Systems” on page 457
- “Issues Specific to HP Storage Systems” on page 461

- “Issues Specific to LSI Storage Systems” on page 462

## Issues Specific to CLARiiON Storage Systems

The following issues are specific to CLARiiON® storage systems:

- “Making the Management Server a Privileged User for CLARiiON” on page 455
- “EMC Navisphere CLI Is Required” on page 455

### *Making 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

For the management server on Windows, use Navisphere Manager to 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

For the management server on UNIX systems, add the following to the privilege user section:

```
SYSTEM@root
```

When you discover the CLARiiON storage system, provide its IP address and the user name and password to be 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 456
- “Process Has an Exclusive Lock Message” on page 456
- “Some EMC Volumes, Their LUNs and LUN Maskings Are Hidden” on page 457

### *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.

### *Process Has an Exclusive Lock Message*

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. 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 storage system.



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 whether someone is performing an operation that requires a lock, such as LUN mapping, LUN masking, or Get Details. This applies even if one of the processes is being used by a third-party product, such as for LUN masking. If this is the case, wait until the process is complete. Remove the lock manually only if you are certain that no other processes are happening on the storage system. To learn how to remove the lock, refer to the documentation for the Symmetrix storage system.

If a provisioning failure causes 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 similar to 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, and 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 contains the following topics:

- “About Provisioning on HDS Storage Systems” on page 458
- “Host Security Group Names on HDS Storage Systems” on page 458
- “Unable to Provision When HDS CruiseControl Is Enabled” on page 459
- “Increasing the Wait Time for the Management Server” on page 459
- “Initiator Ports Cannot Be Used for Provisioning” on page 461
- “Mapping Issue on HDS 9900V Storage Systems” on page 461

## *About Provisioning on HDS Storage Systems*

---

**Caution** – 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 already-created volumes.

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 from the Free LDEVs list when you create a volume, the LDEV is removed 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 requested 4-GB volume is larger than one LDEV, two of the Free LDEVs will be used for the 4-GB volume.

## *Host Security Group Names on HDS Storage Systems*

In releases prior 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 host security group 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 host security group 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 host security group and its value is stored on the device itself. The management server uses 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 host security group 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. Note that if it is run after an upgrade, the management server will need to be resynchronized so that the database is updated with the renamed host security groups. 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 UNIX systems, run `HsgNameUpgrader.sh`. The Host Security Group Name Upgrader tool opens.
3. Enter a host name and credentials in the Server section.
4. Click **Load**. 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 **Process Select Array** if you want to process a single array, or click **Process All Arrays** 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 active. After two retries, the management server stops attempting to contact HiCommand Device Manager.

For example, assume you initiated a provisioning command to add the host security group. The management server waits 20 minutes for a response. If it receives no response during that 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 host security groups with the same configuration of LUNs and 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 second command tries to delete the same host security group, but the target host security group is deleted when the first command is completed, and 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. To change the provisioning timeout property:

1. Select **Configuration > Product Health**, and 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. Repeat step 1 to return to the Advanced page.
5. Paste the copied text into the Custom Properties box. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the box and select **Paste**.
6. Make any necessary changes in the Custom Properties box. Change the value of the `cimom.provider.hds.ProvisioningTimeout` property. (Note that 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. For example:  
`cimom.provider.hds.ProvisioningTimeout=3600000`
7. Click **Save**.

### *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.

### *Mapping Issue on HDS 9900V Storage Systems*

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**

This section contains the following topics:

- “Cannot Always Delete Selected Volume on MSA” on page 461

### *Cannot Always Delete Selected Volume on MSA*

MSA volumes must be deleted in the reverse order of their creation. For example, if you have six volumes, and you want to delete the second one you created, you must delete the volumes one at a time, starting with the volume created sixth and continuing with the fifth, fourth, third, and then the second. Attempting to delete other volumes will return a generic error code 4.

### *Selective Storage Presentation must be enabled on MSA*

Selective Storage Presentation (SSP) for the array must be enabled for provisioning to work.

## Issues Specific to LSI Storage Systems

This section contains the following topics:

- “Creating and Deleting Storage Pools” on page 462
- “Creating and Deleting Storage Volumes” on page 462

### *Creating and Deleting Storage Pools*

For LSI, 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 LSI 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 LSI 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 Table 11-6, “Volume Usage,” on page 462

---

**Note** – No volume-to-LUN masking is done by default.

---

**TABLE 11-6** 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

---

**Caution** – 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 it.

---

## Path Provisioning

---

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 contains the following topics:

- “About Path Provisioning” on page 463
- “How Path Provisioning Works” on page 465
- “About the User Interface” on page 468
- “How to Use Path Provisioning” on page 466
- “About the User Interface” on page 468
- “Default System Action Templates” on page 469
- “Creating a System Action Template” on page 502
- “Modifying a System Action Template” on page 504
- “Adding a Host” on page 504
- “Creating a Host Security Group” on page 505
- “Scheduling Provisioning Jobs” on page 506
- “Executing Provisioning Jobs” on page 508
- “Monitoring Provisioning Jobs” on page 509
- “Deleting Multiple Jobs” on page 509
- “Naming Conventions” on page 509
- “Using Multipathing with Path Provisioning” on page 510
- “Customizing Path Provisioning” on page 511
- “Assigning a Template to a Role” on page 516

---

## About Path Provisioning


Path Provisioning lets you schedule provisioning tasks to take place 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., and when you come in later that morning, review 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 **Refresh**. The **Refresh** button updates the following:

- The Path Provisioning screen with the following changes:
  - Changes made in Provisioning. For example, if you use the wizards in Provisioning to create a host security group, when you access Path Provisioning, your changes are not shown until you refresh.
  - Changes from executed jobs. After a job is executed in Path Provisioning, the Path Provisioning screen is not updated until you click **Refresh** or 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
  - Protection 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 the  icon next to them, indicating they are unreachable. You can still select these storage ports 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. Therefore, if the management server cannot detect a switch, as with a DAS connection, the user interface assumes the storage ports are unreachable.
- A port 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 HBAs displayed may not have a connection to the selected storage system. This is done to provide flexibility. For example, you can select a disconnected HBA for a job you want to take place when the HBA is connected to the storage system.

**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, or zone aliases.
- 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, 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.
- Only manageable fabrics will be displayed in the Path Provisioning. If no provisioning can be done on the fabric (any vendor) means it will not be displayed in Path Provisioning, but will be displayed in Provisioning wizard.
- 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.

#### For Brocade switches:

While configuring Path Provisioning for a Brocade switch, you might occasionally see the following message in the Message Console tab:

```
Failed Path Provisioning Job, <job information>, cause:
CIM_ERR_FAILED
Fabric Session is Locked: SessionState = 2
Try again later
```

Click **Refresh** to see the final status. This message generally means that the Brocade fabric is busy and the management server was unable to complete the operation. Try the operation again at a later time.

---

## 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).
- 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).
- A set of tasks of the same type (such as mapping multiple LUNs).
- Multiple LUN masking tasks or a job consisting of several tasks of different types (such as, Map LUN, zone required ports, mask LUN) that go together as a combination).
- 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. A job with all its required details (such as, parameters to invoke Provider Service methods) is stored in a job queue. At the scheduled time, the scheduler 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 job is displayed in the State column of the Provision Job section located in the lower pane of the screen. A job can have one of the following statuses:

- **Created** - The job has been created, but it will not be executed. The job cannot be viewed by others and is deleted when the Web browser is closed. See “Scheduling Provisioning Jobs” on page 506 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.
- **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. Provisioning allows you to either select a default template from the System Action combo-box or create a system action template consisting of the provisioning tasks best suited to your system, as described in “Creating a System Action Template” on page 502.

The default provisioning templates are listed below:

- **Volume Creation + LUN Security + Zone Operation** - Lets you create a meta volume, map a volume to a Fibre Channel port and host HBAs (HSG), and then create a zone. See “Volume Creation, LUN Security, and Zone Operation” on page 469 for more information.
- **Meta Volume Creation** - Lets you create a meta volume. See “Creating a Meta Volume” on page 476 for more information.
- **LUN Security** - Lets you map a meta volume to Fibre Channel port and host HBAs (HSG). See the topic “LUN Security” on page 478 for more information.
- **Zone Operation** - Lets you perform a zone operation. See “Zone Operation” on page 482 for more information.
- **Volume Creation + LUN Security** - Lets you create a meta volume and then map the meta volume to a Fibre Channel port. See “Volume Creation and LUN Security” on page 487 for more information.
- **LUN Security and Zone Operation** - Lets you create a host security group with the host HBA WWN along with zoning operations. See “LUN Security and Zone Operation” on page 492.
- **Volume Assignment** - Lets you assign a volume to existing host security groups. See “Volume Assignment” on page 497.

Complete the steps in the various panes. A step that is not required for the action is disabled after all data has been loaded.

Schedule the task as described in the topic, “Scheduling Provisioning Jobs” on page 506.


---

**Note** – You can control which templates users can access. See “Assigning a Template to a Role” on page 516 for more information.

---

The Table 12-1, “Overview for Path Provisioning,” on page 468 provides an overview of the steps required for Path Provisioning.

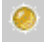

**TABLE 12-1** Overview for Path Provisioning

Step	Description	Where to Find Additional Information
1	Select a system action.	See the topics: <ul style="list-style-type: none"> <li>• “Volume Creation, LUN Security, and Zone Operation” on page 469</li> <li>• “Creating a Meta Volume” on page 476</li> <li>• “LUN Security” on page 478</li> <li>• “Zone Operation” on page 482</li> <li>• “Volume Creation and LUN Security” on page 487</li> <li>• “LUN Security and Zone Operation” on page 492</li> <li>• “Volume Assignment” on page 497</li> </ul>
2	Complete the steps.	
3	<i>Optional:</i> Schedule a provisioning job that you want to take place at a later time.	See “Scheduling Provisioning Jobs” on page 506.
4	Execute the job. The job does not run until you click the <b>Execute Job</b>  button is clicked. Once this button is clicked, the job is saved in the management server database.	See “Executing Provisioning Jobs” on page 508.

## About the User Interface

The Path Provisioning feature provides the following on its toolbar:

TABLE 12-2 Feature Toolbar

Button	Description
	Change Observer button - Monitors changes in the database status on the server. When changes are detected, the button turns gold. Click the gold button and a pop-up window displays the elements that have changed on the server. When no changes are detected, the button is greyed out.
	Reloads the Change Observer button to display the latest changes to elements on the server.
Configure Templates	Lets you create and configure a System Action Template. For more information, see “Creating a System Action Template” on page 502 and “Modifying a System Action Template” on page 504.
Assign Templates	Lets you assign a template to a role. For more information, see “Assigning a Template to a Role” on page 516.

## Default System Action Templates

This section describes the contents of the default system action templates and the options included in each. It also provides instructions for using each of the options.


This section contains the following topics:

- “Volume Creation, LUN Security, and Zone Operation” on page 469
- “Creating a Meta Volume” on page 476
- “LUN Security” on page 478
- “Zone Operation” on page 482
- “Volume Creation and LUN Security” on page 487
- “LUN Security and Zone Operation” on page 492
- “Volume Assignment” on page 497
- “Providing a LUN Number” on page 502
- “Adding a Host” on page 504


## Volume Creation, LUN Security, and Zone Operation

You can use Path Provisioning to create meta volumes. Map that meta volume to a Fibre Channel port and host HBA, and then 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, clear the options you do not want in your next job. For example, assume you created a volume, and now you want to create a new volume on the same host. Clear the Volume pane by clicking the  button.
- If you want to clear all the steps, except for the Step 1 (storage systems) action, select another option from the System Action combo-box.
- (*HDS storage systems only*) Before you can create a volume, you must delete some unmapped LDEVs using the standard provisioning tool.

To delete LDEVs:

- a. Click **Provisioning** (  ).
- b. Click the Storage Systems tab, then the **Provision** button for the storage system.
- c. Click **Step 2 Volume**.
- d. Select the desired number of LDEVs for the LUSE volume, and then click **Delete Selected Volumes**.
- e. 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 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 the storage system on which you want to create the metavolume.

Keep in mind the following:

- The S column 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 below the pane.

The selected storage system's name is displayed below the Storage System pane. The Host pane is populated. Notice in the figure below that some hosts have a red X over their icon. This means the host is not accessible.

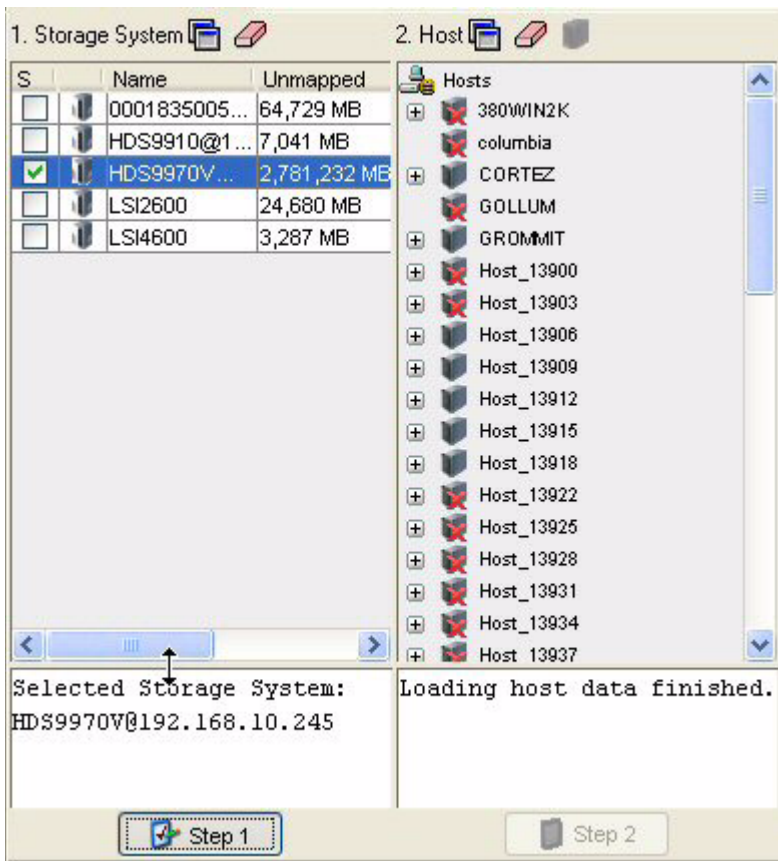


FIGURE 12-1 Selecting a Storage System

## Step 2 - Select a Host

Keep in mind the following:


- If you select host and storage ports that belong to an existing zone alias and you have the **Display Zone Aliases** option selected in Customize Zone Options dialog box, the existing zone alias is automatically selected and highlighted.


- If you select hosts and storage ports that are not contained in an existing zone alias, the new hosts and storage ports are added into the existing zone alias after the provisioning job finishes successfully.
1. Wait for all data to be loaded. All data has been loaded when you see the following messages:
    - Host data Loaded.
    - Loading volume data finished.
    - Loading HSG data finished.
    - Loading zone data finished.

The Step 2 button is disabled until data has been loaded

2. Select a host that is accessible.


Keep in mind the following:

- To create a provisioning job for multipathing, click the  button above the Host pane, and then select the option for multipathing. When you select this option, you must select the same host for both paths. See “Host Customize Dialog Box” on page 512 for more information.
- To automatically create a zone if a zone does not meet a preset criteria:

- a. Click the  button above the Zone pane.
- b. Select the option **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**.
- c. Set the criteria. See “Customize Zone Options Dialog Box” on page 513 for more information about setting the criteria.
- d. 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” on page 509 for more information about the naming requirements for a zone. After you enter a zone name, the new zone is displayed in the Zone pane, but it will not actually be created until the job runs.

- To configure zoning manually, click the  button above the Zone pane, and then deselect the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**. In the zone pane, select a zone or create one manually.

3. Click **Step 2**.

Information about the selected port, such as its WWN, 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. To select multiple volumes in Windows, hold down 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** – Lists available free extents that can be used to create a meta volume or LUSE. 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. For more information, see


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 438 for more information.

---

**Caution** – 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:

- To narrow the type of volumes displayed in the Volumes pane by using the Customize Volume Options dialog box () , located above the Volumes pane. See “Host Customize Dialog Box” on page 512 for more information.

---

**Note** – The Customize Volume Options dialog box is not available for the HP EVA.

---

- 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 502 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 505 for information on how to create a host security group. See “General Provisioning Issues” on page 454 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 a zone has not been selected or created yet, most likely the option **Automatically Configure Zoning** is not selected in the Customize Zone Options dialog box (). The management server assumes you want to select a pre-existing zone or create one manually.






- **To reuse a zone** - Select a zone in the Zone pane and then click **Step 5**, and expand a fabric node to view its zones.
- **To create a zone** - Select a fabric in the zone pane, click the  button, and then enter a name for the zone. For more information, see “Creating a Zone” on page 486.

TABLE 12-3 Zone Icons

Icon	Description
	<ul style="list-style-type: none"><li>• Above Zone pane - Used to create zones.</li><li>• In the Zone pane - Represents a zone.</li></ul>
	Zone Alias
	Port
	The fabric cannot be reached

---


**Caution** – *McDATA switches only*: Path Provisioning looks for the names of the active zone set and of the active zones and verifies that all of their saved counterparts are matched 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 12-2** Selecting a Fabric


2. Click the  button located above the Zone pane.
3. Enter a zone name in the dialog box. For naming conventions, see “Naming Conventions” on page 509.
4. Click **OK**.  
The new zone is added to the Zone pane.
5. Click **Create Job**.  
The job is added to the bottom pane.
6. Take one of the following actions:
  - 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 506.

# Creating a Meta Volume


Keep in mind the following when creating meta volumes:

- If you have options still selected from a previous job, clear the options you do not want in your next job. For example, assume you previously created a volume. To create a new volume on the same host, clear the Volume pane by clicking the  button.
- To clear the action taken in all Steps except Step 1, 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:

- a. Click **Provisioning** (  ).
- b. Click the Storage Systems tab, then the **Provision** button for the storage system.
- c. Click **Step 2 Volume**.
- d. Select the desired number of LDEVs for the LUSE volume, and then click **Delete Selected Volumes**.
- e. Take note of the array group from which you deleted the LDEVs. You need this information to create the LUSE volume.

To use Path Provisioning to create meta volumes:

1. Click **Provisioning** (  ).
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. Select **Meta Volume Creation** from the System Action combo-box.
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 column heading in the Storage Systems pane means that only a single selection is allowed.

---

6. Click the **Step 1** button below the pane.

The selected storage system's name 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 438 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 Customize Volume Options dialog box () located above the Volumes pane. See “Host Customize Dialog Box” on page 512.

---

**Note** – The Customize Volume Options dialog box is not available for the HP EVA.

---


- 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 **Create Job**.


The job is added to the bottom pane.

10. Take one of the following actions:


- 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 506.

## LUN Security

Keep in mind the following:

- If you have options still selected from a previous job, 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. 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.

To use Path Provisioning to designate subsystem LUN security:

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 column heading in the Storage Systems pane means that only a single selection is allowed.

---

3. Click the **Step 1** button below the pane.

The selected storage system’s name is displayed below the Storage System pane. The Host pane is populated. Notice in “Selecting a Storage System” on page 479 that some hosts have a red X over their icon. This means the host is not accessible.




FIGURE 12-3 Selecting a Storage System

## Step 2 - Select a Host

Keep in mind the following:


- If you select host and storage ports that belong to an existing zone alias and you have the **Display Zone Aliases** option selected in Customize Zone Options dialog box, the existing zone alias is automatically selected and highlighted.
  - If you select hosts and storage ports that are not contained in an existing zone alias, the new hosts and storage ports are added into the existing zone alias after the provisioning job finishes successfully.
1. Wait for all data to be loaded. All data has been loaded when you see the following messages:
    - Host data Loaded.
    - Loading volume data finished.
    - Loading HSG data finished.


The Step 2 button is disabled until data has been loaded

2. Take one of the following actions:

- Select a host that is accessible.
- Add a host that is not currently connected to the network by clicking the  button. See “Adding a Host” on page 480.


Keep in mind the following:

- To create a provisioning job for multipathing, click the  button above the Host pane, and then select the option for multipathing. When you select this option, you must select the same host for both paths. See “Host Customize Dialog Box” on page 512 for more information.
- To automatically create a zone if a zone does not meet a preset criteria:

- a. Click the  button above the Zone pane.
- b. Select the option **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**.
- c. Set the criteria. See “Customize Zone Options Dialog Box” on page 513 for more information about setting the criteria.
- d. 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” on page 509 for more information about the naming requirements for a zone. After you enter a zone name, the new zone is displayed in the Zone pane, but it will not actually be created until the job runs.

- To configure zoning manually, click the  button above the Zone pane, and then deselect the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**. In the zone pane, select a zone or create one manually.

3. Click the **Step 2** button.

Information about the selected port, such as its WWN, 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.



---

**Caution** – Make sure the added host is physically connected to the network before the scheduled job runs.

---

1. Click the  button.
2. Enter a name for the host in the Host Name box.
3. Enter a port name of the host in the Port WWN box.
4. Click the **Add** button.
5. Repeat Steps 2 and 3 for multiple ports.
6. When you are done with your changes, click **OK**.  
The host is added to the list of hosts.
7. 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 438 for more information.

---


**Caution** – 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 502 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 Customize Volume Options dialog box () located above the Volumes pane. See “Host Customize Dialog Box” on page 512.

---

**Note** – The Customize Volume Options dialog box is not available for the HP EVA.

---

- 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 505 for information on how to create a host security group. See “General Provisioning Issues” on page 454 for information on how your storage system handles host security groups.

2. Click **Step 4**.

3. Click **Create Job**.


The job is added to the bottom pane.

4. Take one of the following actions:


- 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 506.

## Zone Operation

Keep in mind the following:

- If you have options still selected from a previous job, 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. 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 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 column heading in the Storage Systems pane means that only a single selection is allowed.

---

3. Click the **Step 1** button below the pane.

The selected storage system's name 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.

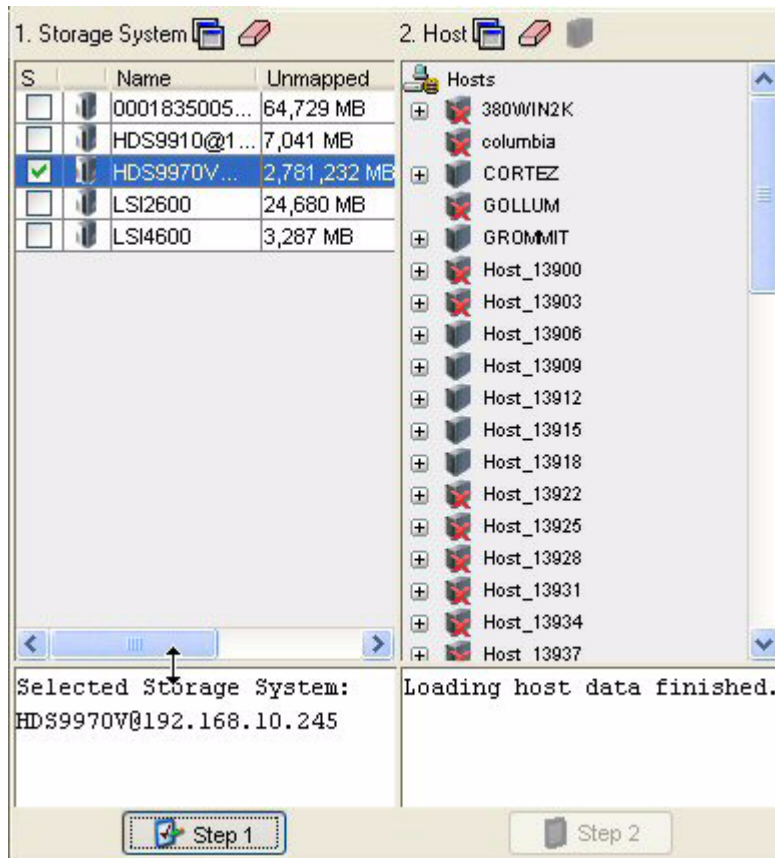


FIGURE 12-4 Selecting a Storage System

## Step 2 - Select a Host

Keep in mind the following:

- If you select host and storage ports that belong to an existing zone alias and you have the **Display Zone Aliases** option selected in Customize Zone Options dialog box, the existing zone alias is automatically selected and highlighted.
- If you select hosts and storage ports that are not contained in an existing zone alias, the new hosts and storage ports are added into the existing zone alias after the provisioning job finishes successfully.


1. Wait for all data to be loaded. When all data has been loaded, the following messages are displayed.:


- Host data Loaded.
- Loading HSG data finished.
- Loading zone data finished.

The Step 2 button is disabled until data has been loaded.

2. Select a host that is accessible.

Keep in mind the following:

- To create a provisioning job for multipathing, click the  button above the Host pane, and then select the option for multipathing. When you select this option, you must select the same host for both paths. See “Host Customize Dialog Box” on page 512 for more information.
- To automatically create a zone if a zone does not meet a preset criteria:

a. Click the  button above the Zone pane.


b. Select the option **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**.

c. Set the criteria. See “Customize Zone Options Dialog Box” on page 513 for more information about setting the criteria.

d. 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” on page 509 for more information about the naming requirements for a zone. After you enter a zone name, the new zone is displayed in the Zone pane, but it will not actually be created until the job runs.

- To configure zoning manually, click the  button above the Zone pane, and then deselect the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**. In the zone pane, select a zone or create one manually.

3. Click **Step 2**.


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 505 for information on how to create a host security group. See “General Provisioning Issues” on page 454 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 a zone has not been selected or created yet, most likely the option **Automatically Configure Zoning** is not selected in the Customize Zone Options dialog box (). The management server assumes you want to select a pre-existing zone or create one manually.






- **To reuse a zone** - Select a zone in the Zone pane and then click **Step 5**, and expand a fabric node to view its zones.
- **To create a zone** - Select a fabric in the zone pane, click the  button, and then enter a name for the zone. For more information, see “Creating a Zone” on page 486.

TABLE 12-4 Zone Icons

Icon	Description
	<ul style="list-style-type: none"><li>• Above Zone pane - Used to create zones.</li><li>• In the Zone pane - Represents a zone.</li></ul>
	Zone Alias
	Port
	The fabric cannot be reached

---


**Caution** – (McDATA switches only) Path Provisioning looks for the names of the active zone set and of the active zones and verifies that all of their saved counterparts are matched 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 12-5** Selecting a Fabric

2. Click the  button located above the Zone pane.
3. Enter a zone name in the dialog box. For naming conventions, see “Naming Conventions” on page 509.
4. Click **OK**.  
The new zone is added to the Zone pane.
5. Click **Create Job**.  
The job is added to the bottom pane.
6. Take one of the following actions:
  - 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 506.

## Volume Creation and LUN Security


Keep in mind the following:

- If you have options still selected from a previous job, 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. 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:

- a. Click **Provisioning** (  ).
- b. Click the Storage Systems tab, then the **Provision** button for the storage system.
- c. Click **Step 2 Volume**.
- d. Select the desired number of LDEVs for the LUSE volume, and then click **Delete Selected Volumes**.
- e. 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 column 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 below the pane.  
The selected storage system's name is displayed below the Storage System pane.



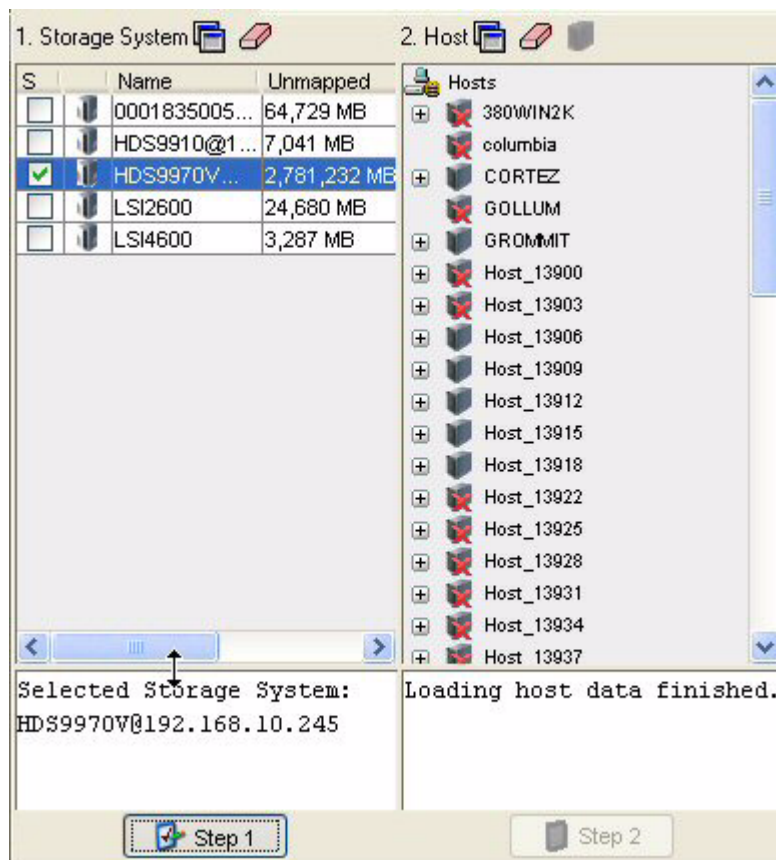


FIGURE 12-6 Selecting a Storage System

## Step 2 - Select a Volume

To select a volume:

- 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 438 for more information.

---

**Caution** – 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 Customize Volume Options dialog box () located above the Volumes pane. See “Host Customize Dialog Box” on page 512.

---

**Note** – The Customize Volume Options dialog box is not available for the HP EVA.

---

- If the LUN has already been selected and Step 4 is clicked, skip this step or click the  button.

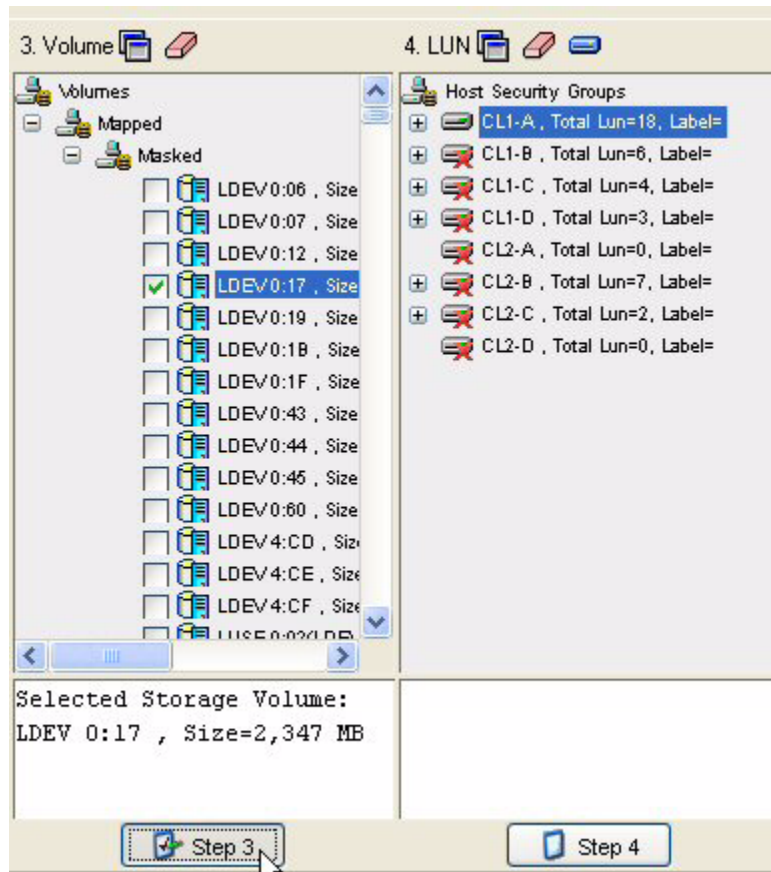


FIGURE 12-7 Selecting a Volume

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 502 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 505 for information on how to create a host security group. See “General Provisioning Issues” on page 454 for information on how your storage system handles host security groups.
2. Click **Step 4**.

3. Click **Create Job**.

The job is added to the bottom pane.


4. Take one of the following actions:

- 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 506.

## LUN Security and Zone Operation


Keep in mind the following:

- If you have options still selected from a previous job, 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. 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 a storage system on which you want to create the metavolume.

---

**Note** – The S column heading in the Storage Systems pane means that only a single selection is allowed.

---

3. Click the **Step 1** button below the pane.

The selected storage system's name is displayed below the Storage System pane. The Host pane is populated. Notice in the following pane that some hosts have a red X over their icon. This means the host is not accessible.

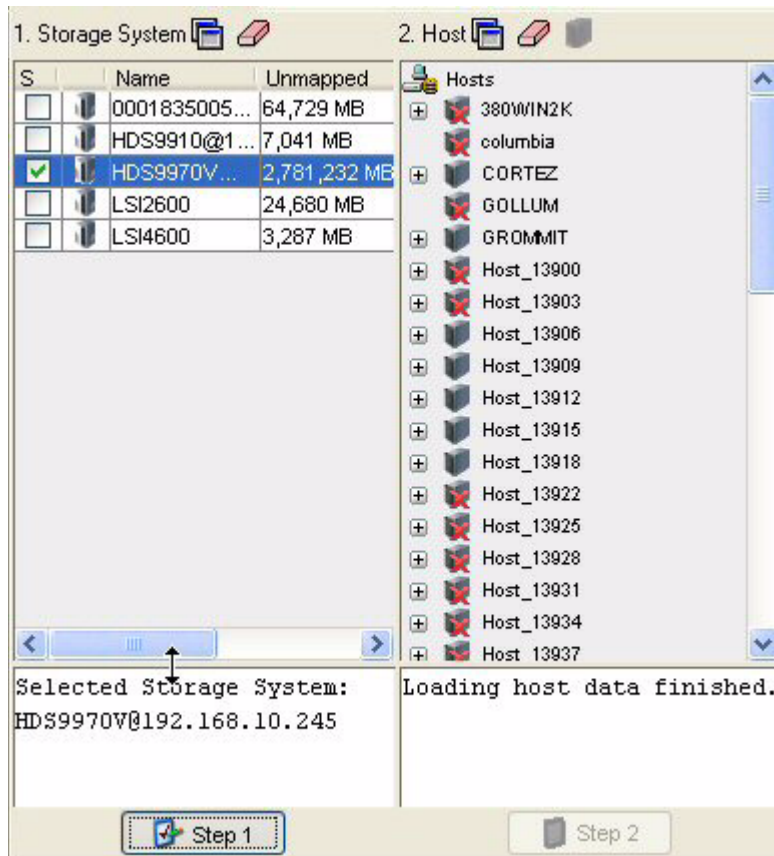


FIGURE 12-8 Selecting a Storage System

## Step 2 - Select a Host

Keep in mind the following:


- If you select host and storage ports that belong to an existing zone alias and you have the **Display Zone Aliases** option selected in Customize Zone Options dialog box, the existing zone alias is automatically selected and highlighted.
- If you select hosts and storage ports that are not contained in an existing zone alias, the new hosts and storage ports are added into the existing zone alias after the provisioning job finishes successfully.

1. Wait for all data to be loaded. All data has been loaded when you see the following messages:


- Host data Loaded.
- Loading volume data finished.
- Loading HSG data finished.


The Step 2 button is disabled until data has been loaded

2. Take one of the following actions:

- Select a host that is accessible.
- Add a host that is not currently connected to the network by clicking the  button. See “Adding a Host” on page 480.

Keep in mind the following:

- To create a provisioning job for multipathing, click the  button above the Host pane, and then select the option for multipathing. When you select this option, you must select the same host for both paths. See “Host Customize Dialog Box” on page 512 for more information.
- To automatically create a zone if a zone does not meet a preset criteria:

a. Click the  button above the Zone pane.


b. Select the option **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**.

c. Set the criteria. See “Customize Zone Options Dialog Box” on page 513 for more information about setting the criteria.

d. 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” on page 509 for more information about the naming requirements for a zone. After you enter a zone name, the new zone is displayed in the Zone pane, but it will not actually be created until the job runs.

- To configure zoning manually, click the  button above the Zone pane, and then deselect the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**. In the zone pane, select a zone or create one manually.

3. Click **Step 2**.

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.


---

**Caution** – Make sure the added host is physically connected to the network before the scheduled job runs.

---


1. Click the  button.
2. Enter the name for the host in the Host Name box.
3. Enter the port name of the host in the Port WWN box.
4. Click **Add**.
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 505 for information on how to create a host security group. See “General Provisioning Issues” on page 454 for information on how your storage system handles host security groups.
2. Click **Step 4**.
3. Click **Create Job**.  
The job is added to the bottom pane.
4. Take one of the following actions:
  - 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 506.


## 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 a zone has not been selected or created yet, most likely the option **Automatically Configure Zoning** is not selected in the Customize Zone Options dialog box (). The management server assumes you want to select a pre-existing zone or create one manually.

- **To reuse a zone** - Select a zone in the Zone pane and then click **Step 5**, and expand a fabric node to view its zones.
- **To create a zone** - Select a fabric in the zone pane, click the  button, and then enter a name for the zone. For more information, see “Creating a Zone” on page 496

**TABLE 12-5** Zone Icons

Icon	Description
	<ul style="list-style-type: none"><li>• Above Zone pane - Used to create zones.</li><li>• In the Zone pane - Represents a zone.</li></ul>
	Zone Alias
	Port
	The fabric cannot be reached

---

**Caution** – (McDATA switches only) Path Provisioning looks for the names of the active zone set and of the active zones and verifies that all of their saved counterparts are matched 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 12-9** Selecting a Fabric

2. Click the  button located above the Zone pane.
3. Enter a zone name in the dialog box. For naming conventions, see “Naming Conventions” on page 509.
4. Click **OK**.  
The new zone is added to the Zone pane.
5. Click **Create Job**.  
The job is added to the bottom pane.
6. Take one of the following actions:
  - 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 506.

## Volume Assignment


Keep in mind the following:

- If you have options still selected from a previous job, 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. 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:

- a. Click **Provisioning** (  ).
- b. Click the Storage Systems tab, then the **Provision** button for the storage system.
- c. Click **Step 2 Volume**.
- d. Select the desired number of LDEVs for the LUSE volume, and then click **Delete Selected Volumes**.
- e. 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 a storage system on which you want to create the metavolume.

---

**Note** – The S column 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 below the pane.  
The selected storage system's name is displayed below the Storage System pane.

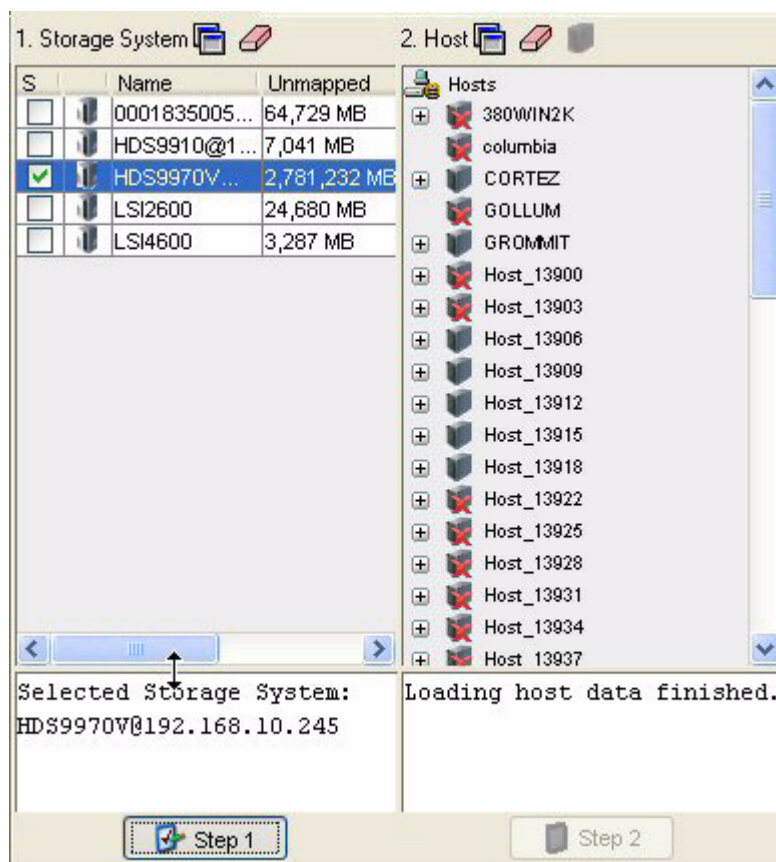


FIGURE 12-10 Selecting a Storage System

## 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 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 438 for more information.

---

**Caution** – 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 Customize Volume Options dialog box () located above the Volumes pane. See “Host Customize Dialog Box” on page 512.

---

**Note** – The Customize Volume Options dialog box is not available for the HP EVA.

---

- If the LUN has already been selected and Step 4 is clicked, skip this step or click the  button.




FIGURE 12-11 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 502 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 505 for information on how to create a host security group. See “General Provisioning Issues” on page 454 for information on how your storage system handles host security groups.
2. Click **Step 4**.

3. Click **Create Job**.

The job is added to the bottom pane.

4. Take one of the following actions:

- 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 506.

## Providing a LUN Number

Keep in mind the following when providing a LUN number:

- 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.
- *LSI 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 figure below. Click **OK** when you are done. This dialog box is displayed if your storage system requires you to provide a LUN number.

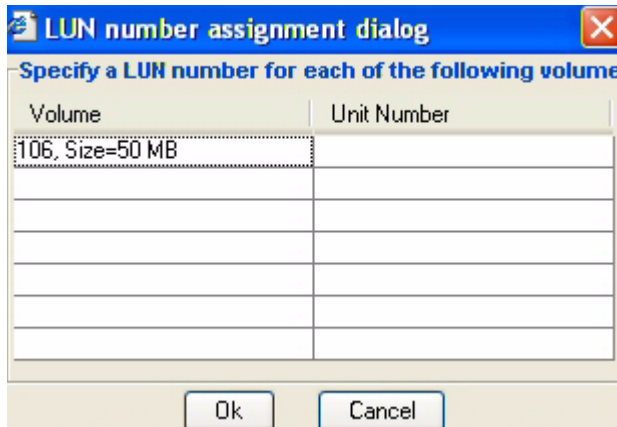



FIGURE 12-12 Specifying a LUN Number

---

## 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 dialog box.
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 box, and then click **Apply**.
6. Select a master template on which you want to base your new template. For more information about the default master templates, see “Default System Action Templates” on page 469.
7. Click the following tabs and select the options you want for your template from each tab.

---

**Note** – Not all tabs are available for all default templates.

---

- **Storage Options** - See “Storage System Customize Dialog Box” on page 511 for more information for these options.
  - **Host Options** - See “Host Customize Dialog Box” on page 512 for more information about these options.
  - **Volume Options** - See “Customize Volume Options Dialog Box” on page 512 for more information about these options.
  - **Host Security Options** - See “Customize HSG Options” on page 513 for more information about these options.
  - **Zone Options** - See “Customize Zone Options Dialog Box” on page 513 for more information about these options.
8. When you have selected all of the options you want in your template, take one of the following actions:
    - Click **Apply** to apply your changes and keep the Provisioning Template Configuration dialog box open.
    - Click **OK** to apply your changes and leave the Provisioning Template Configuration dialog box.
    - Click **Cancel** to leave the Provisioning Template dialog box 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** 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, take one of the following actions:
  - Click **Apply** if you want to apply your changes and keep the Provisioning Template Configuration dialog box open.
  - Click **OK** if you want to apply your changes and leave the Provisioning Template Configuration dialog box.
  - Click **Cancel** if you want to leave the Provisioning Template dialog box without saving your changes. You will lose your changes unless you had previously clicked **Apply**.

---


# Adding a Host

---

**Caution** – 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. This feature is only available when you select **LUN Security** from the System Action menu.

1. Click the  button.
2. Enter a name for the host in the Host Name box.
3. Enter the port WWN of the host in the Port WWN box.



4. Click **Add**.
5. Repeat Steps 3 and 4 for multiple ports.
6. When you are done with your changes, click **OK**.  
The host is added to the list of hosts.
7. Physically connect the host to the network before the job runs.

---

## Creating a Host Security Group

Keep in mind the following:


- Each storage system handles host security groups differently. See “General Provisioning Issues” on page 454.
- *For LSI storage systems:* When the **Volume Creation and LUN Security** option is selected on the System Action menu, you cannot add a host security group to an LSI storage system or to any storage system that LSI resells under a different brand.
- *For IBM storage systems:* You cannot assign the host mode in the user interface. You must modify an internal property. See “Setting the Host Mode for IBM Storage Systems” on page 453 for more information.

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 12-13 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, enter a unique name.

---

**Note** – For Symmetrix storage systems, you are not asked for the name of the host security group.

---

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 only*:
  - a. Select the host mode for the host security group.
  - b. Provide a second host mode if applicable.
  - c. Click **Create Host Security Group**.
9. 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 before you can schedule it. See “About Path Provisioning” on page 463 for more information.
- You cannot delete a job once it has started.
- You can deselect a job by clicking its check mark or by clicking the Clear Selection  button in the Provision Jobs pane.

To determine the status of a job, look 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 506 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 has failed.
- **Ended** - The job has finished.


To schedule a provisioning job:



1. Click the **Create Job** button in the lower pane.


The job is assigned 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 box () is selected for multipathing. If the Host Customize dialog box is set for multipathing, select a 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 Box” on page 512.

2. Schedule the job by selecting the job and then clicking the  button in the Provision Jobs pane.
  - a. *Optional:* In the Time box, change the time displayed. The management server automatically displays a time five minutes from when you clicked the . Enter the time in 24-hour format with hours and minutes separated by a colon. For example, enter 23:15 for 11:15 p.m.
  - b. Select a date when you want the job to start.
  - c. Click **Set**.
 

If you click **Set** 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 will be executed at the scheduled time.


To execute a scheduled job, select the job in the Provision Jobs pane. Detailed information appears in the Job Console pane, located in the lower-right corner of the window. The Message Console tab, located also in the lower-right corner, provides information such as whether 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 management server database.
- 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 or according to the order of the jobs appearing in the Provision Jobs table.
- The management server can detect when the requested volume and/or zone has already been created. For example, assume you created a job to create a volume, and the next few jobs use this volume. Before creating the volume, these jobs will first determine whether the volume has already been created.

---



## Executing Provisioning Jobs

If you want to save and execute a job, you must click the **Execute Job** () button. When you click that 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 in 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 immediately or at a pre-determined time in the future:

- To start a job now, click the **Execute Job** () button before the job is scheduled.
- To start a job at a pre-determined time in the future, schedule the job first, and then click the **Execute Job** () button. The job is executed according to the time set on the management server.

To view the latest information in Path Provisioning, click **Refresh**. For details about what the **Refresh** button updates, see “About Path Provisioning” on page 463.

---


# 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. Use this feature when you are not sure if a job has ended.


---

# Deleting Multiple Jobs

---

**Caution** – 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 **Delete** () button until all selected jobs have been removed, as shown in the figure below.

---

To deselect a job, click the check marks or click the **Clear Selection** () button in the Provision Jobs pane.

---

# Naming Conventions

Observer the following naming conventions:

**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

To use provisioning with multipathing, set the Host Customize dialog box to the multipathing option, and 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** - Create a meta volume, map a volume to a Fibre Channel port and host HBAs (HSG), and then create a zone. See “Volume Creation, LUN Security, and Zone Operation” on page 469 for more information.
  - **LUN Security** - Map meta volume to Fibre Channel port and host HBAs (HSG). See the topic “LUN Security” on page 478 for more information.
  - **Zone Operation** - Create a new zone. See “Zone Operation” on page 482 for more information.
2. Select a storage system, and then click the **Step 1** button.
3. Click the  button above the Host pane.
4. Select the following option:
 

Multipath: Select more than one port within a single server.
5. Select a port on a host, and then click **Step 2**.
6. Select a volume, host security group, and/or zone, as described in the following topics:
  - “LUN Security” on page 478
  - “Zone Operation” on page 482
7. Select the second path from the Path combo box.

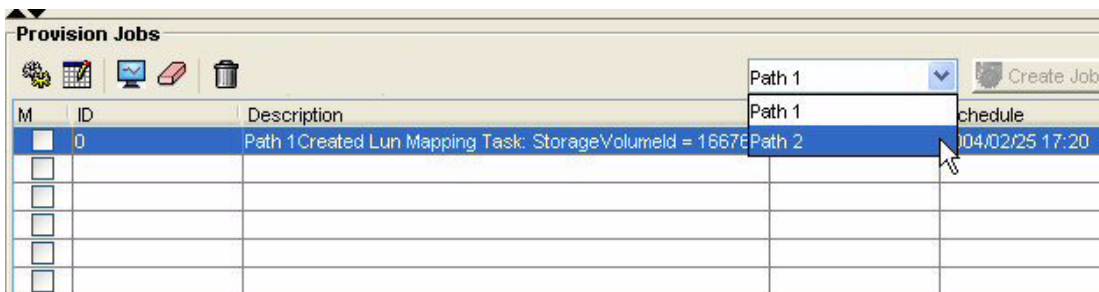


FIGURE 12-14 Selecting the Second Path

8. Repeat Step 6.

You do not need to select a system action or storage system, and you must select the same host as the one used for the first path.

9. Set the schedule for the job as described in the topic, “Scheduling Provisioning Jobs” on page 506.

To narrow the types of volumes displayed in the Volume pan, set the Customize Volume Options dialog box. The Customize Volume Options dialog box is not available for the HP EVA.

## Customizing Path Provisioning

This section contains the following topics:

- “Storage System Customize Dialog Box” on page 511
- “Host Customize Dialog Box” on page 512
- “Customize Volume Options Dialog Box” on page 512
- “Customize HSG Options” on page 513
- “Customize Zone Options Dialog Box” on page 513

## Storage System Customize Dialog Box

The Storage System Customize dialog box lets you specify the type of storage systems displayed in the Storage System pane. To specify the type of storage systems displayed in the Storage System pane:

1. 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. This menu is blank if no storage systems have been assigned to storage tiers. Storage systems are assigned to storage tiers in Chargeback. See “Adding Asset Information” on page 713 and “Adding General Information” on page 714 for information about assigning storage systems to a storage tier.
2. Click **OK** when you have finished making your selections.  
The Storage System pane is updated.

## Host Customize Dialog Box

The Host Customize dialog box lets you use multipathing with Path Provisioning. To set multipathing:

1. 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 510 for more information on how to use multipathing with Path Provisioning.
2. Click **OK** when you have finished making your selections.

## Customize Volume Options Dialog Box

---

**Note** – The Customize Volume Options dialog box is not available for the HP EVA.

---

1. If you want to view all volumes, select the **Show All Volumes** option and then, if desired, select the **Show All Volumes** option. You then have the option to select the **Hide mapped volumes** option.
2. Select one of the following options for metavolumes:
  - **Concatenating** - Only concatenating metavolumes are displayed.
  - **Striped** - Only striped metavolumes are displayed. (Applies only to EMC storage systems)



3. Click **OK**.


The Volume pane is updated.

## Customize HSG Options

1. To specify how LUN are mapped, select one or more of the following options in the Customize HSG Options dialog box:
  - **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 than 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 than an unassigned port. The management server looks for any storage system port that is zoned to the selected HBA. If the management server cannot find a storage system port zoned to the selected HBA, it selects a port with the fewest LUNs.
  - **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. Two 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 have made all your selections.

## Customize Zone Options Dialog Box

When the option Automatically create new zone if no existing zone containing HBA and storage system ports is detected is selected in the Customize Zone Options

dialog box () , 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. To set the criteria for automatically configuring zones, select 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 to see if an identical zone exists. An identical zone contains only the 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 provide a name for the 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.

---

**Caution** – (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 WWNs, 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 as long as the host and storage system are able to access each other through the network.


To create zones through switch port zoning, select the **Use Switch Port Zoning Mode** option, and then select ports on a switch to create your zone. For example, if ports 1 and 2 on a switch are designated as a zone, whatever is connected to either of these ports will be a part of the zone.

### Display Zone Alias Option

Select the Display Zone Alias option to prevent you from creating duplicate zone aliases on the same port. If you have this option selected when path provisioning detects a Fabric support zone alias provisioning feature, all zone aliases in this Fabric will be also listed on the Path Provisioning Zone panel. If the selected ports are part of any existing zone alias provisioning feature, such zone aliases will be selected and highlighted to prevent you from creating duplicate zone aliases on the same port.

## Automatically Configure Zoning

To automatically configure zoning:

1. Click the  button above the Zone pane before you select a host.
2. Select **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**.
3. Select one of the options discussed in the previous bulleted list, and click **Apply**.
4. Select a storage system, and click **Step 1**.
5. Select a host, and click **Step 2**.

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.
6. If you are asked for a zone name, enter the zone name and then click **OK**.

See “Naming Conventions” on page 509 for restrictions on naming zones.

The new zone is displayed in the Zone pane, but it will not be created until 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, and then enter a zone name. See the following topic for more information.

TABLE 12-6 Zone Icons


Icon	Description
	Its location on the page determines its use: <ul style="list-style-type: none"><li>• <b>Displayed above Zone pane</b> - Button for creating zones.</li><li>• <b>In Zone pane</b> - Icon for a zone.</li></ul>
	Zone Alias

**TABLE 12-6** Zone Icons (Continued)


Icon	Description
	Port
	The fabric cannot be reached

To create a zone:

1. Select a fabric in the zone pane.



**FIGURE 12-15** Selecting a Fabric

2. Click the  button located above the Zone pane.
3. Enter a zone name in the dialog box. For naming conventions, see "Naming Conventions" on page 509.
4. Click **OK**.

The new zone is added to the Zone pane.

---

## Assigning a Template 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 with administrator privileges can access the Volume Assignment template.


---

**Caution** – 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, and click the name of the role to see which privileges are assigned to your role.

---

To assign templates to a role:

1. Click **Provisioning** (  ).
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. Click **Assign Templates**.
4. Open the **Select Role** menu and select the role you want to have access to the Path Provisioning template.

For a role to appear in the **Select Role** 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** menu.
5. In the Provisioning Templates pane, select one of the following options for access to the templates for the users belonging to the role you selected in step 4:
    - Select the templates from the Assigned Provisioning Templates panel you want the users to be able to access and click **Add**.
    - Click **Add All >>** to make all templates available to the users assigned to the selected role.
    - Select the templates you do not want the users to be able to access and click **<Remove**.
    - Click **<< Remove All** to make all of the templates unavailable to the users assigned to the selected role.

The templates the users belonging to the selected role will be able to access appear in the Assigned Provisioning Templates panel.

6. When you are done, take one of the following actions:
  - Click **Apply** to apply your changes and keep the Provisioning Template Configuration dialog box open.
  - Click **OK** to apply your changes and leave the Provisioning Template Configuration dialog box.
  - Click **Cancel** to leave the Provisioning Template dialog box without saving your changes. You will lose your changes unless you had previously clicked **Apply**.



## Running Reports

---

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 contains the following topics:

- “About Reporter” on page 519
- “Accessing Reporter” on page 523
- “Viewing Reports” on page 523
- “Viewing Report Collectors for an Element” on page 525
- “Refreshing a Report” on page 526
- “Changing the Formatting of a Report” on page 527
- “Opening a Report in a New Window” on page 527
- “Maximizing the Screen Space for a Report” on page 527
- “Filtering Data in Global Reports” on page 529
- “Sending a Report by E-mail” on page 530
- “Managing E-mail Schedules for Reports” on page 531
- “Creating Custom Reports” on page 536

---

## 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 selected report appears in the right pane.

---

**Caution** – If the message `The report does not contain any data` is displayed, verify that the collector for the report is running. You may also want to verify that a collector is running if you believe the report is not displaying the latest


information. Since data collectors gather the information for reports, if the collector is not running, its reports are not updated. See “Managing Collectors for Reports” on page 252.

---

## Available Reports

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.

The management server also provides some reports with pie charts. If the report is available as a pie chart, the  icon appears next to the report name.

The software provides reports for the following:

- **Global** - These reports provide data gathered from multiple management servers. For example, if 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 260 for more information.
- **Asset Management** - These reports provide information regarding 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 705.
- **System** - These reports are enterprise-wide. 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 bus adapters (HBAs) the management server detects
- **Host** - Data about the hosts the management server monitors, such as reports on host storage allocation and total host utilization. The management server treats host clusters as managed elements, and the capacity calculator intelligently avoids double counting of the capacity from individual nodes at the cluster level.
- **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. ISL Trunking information is also available for supported Brocade switches in related reports that include *SAN Components Not Zoned*, *Individual Switch Report*, *Switch Port I/O Performance*, *Port Utilization by Connection Type*, and *Switch Port Traffic*.
- **Protection Explorer** - Data about backups, such as reports about the status of the daily backup, backup volume, and media availability. To learn more about backups, see “About Protection Explorer” on page 753.
- **Applications** - These reports provide information about an application, such as Oracle or Microsoft Exchange.
- **Hosts** - These reports provide information about a host. The management server treats host clusters as managed elements, and the capacity calculator intelligently avoids double counting of the capacity from individual nodes at the cluster level.
- **NAS** - These reports provide information about NAS storage devices.
- **Storage Systems** - These reports provide information about a selected storage system.
- **Switches** - These reports provide information about a selected switch.
- **Tape Libraries** - These reports provide information about a selected tape library.
- **Recent** - This report lists the last 10 reports viewed. This option is not displayed when you first access Reporter.

## Troubleshooting Reporter

- While the report cache is being refreshed, reports generally contain either no data or incomplete/incorrect data. See “Refreshing a Report” on page 526.
- Some reports display data trends for several days in the future. For example, if 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. Keep in mind that the data trends are just assumptions and should not be treated as fact.

- The Application Dependency report for UNIX-based applications that use raw volumes does not include the Unix host on which the application resides.
- Some collectors and reports can take minutes, an hour, or up to a day to generate the first bit of data. For example, Remote and Local Host Utilization History reports do not show data for the first day.
- In File Server SRM reports, file extensions may sometimes be truncated in reports. For example myfilename.txt may appear as myfilename.tx with a missing t in the report.
- 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
- The following reports generally presents the volumes and the associated mapped and unmapped capacities. The units are always post-RAID. Primordial pools (whose units are always raw) are not listed in any of the utilization reports. For arrays supporting flexible pools, the detail for each of the possible and supported RAID types are displayed.
  - Storage System Reports:
    - Utilization
  - System Reports:
    - Storage Allocation Utilization by Organization
    - Storage Allocation VS Utilization by Organization
    - Storage System Capacity
    - Storage System Utilization
  - Global Reports:
    - Global Storage System Utilization Details
    - Global Storage System Utilization Summary
    - Global Storage System Utilization Details by Vendor
    - Global Storage System Utilization Summary by Vendor
- 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 those 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. For example, assume that you belong to two organizations: OnlyHosts and OnlySwitches. If you set your organization filter to display the elements in OnlySwitches and not in OnlyHosts, your reports display only the elements in OnlySwitches. If you send an Assets by Department report by e-mail, the report displays only information about the elements that are currently allowed through your organization filter. You can set the organization filter when you create an e-mail schedule. See “Sending a Report by E-mail” on page 530.


The following reports display all information regardless of a user's organizations:

- Event Reports
- Application Reports
- Hosts Reports\*
- Storage System Reports\*
- Switch Reports\*

\*These reports 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. If you are having difficulty viewing a report or you believe the report is not displaying the latest information, you might want to verify that a collector is running. See “Viewing Report Collectors for an Element” on page 525 for more information.

In general, while reports are being populated they may contain no data or incomplete/incorrect data.

---

**Caution** – The elements you see in the report are based on your organizations and which organizations are selected in the organization filter displayed at the top of the page.

---

To view a report:


1. Access Reporter as described in “Accessing Reporter” on page 523.
2. In the middle pane, expand the tree and click the type of report you want.

The report is displayed in the right pane. A list of the available filters is displayed at the top of the pane. To hide the list of available filters, click 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 525

---

3. To change the format of the report, select one of the following from the **Format** menu and then click **Run Report**:
  - **HTML (Default)** - Displays the report in a Web page by default.
  - **PDF** - 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 menu, and then click **Run Report**.
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 boxes, by clicking the  icon, and then click **Run Report**.
6. To send a report by e-mail, use the Scheduled Deliveries tab. See “Sending a Report by E-mail” on page 530 for more information.

# Report Parameters

Many reports offer an additional set of parameters for filtering information. If additional parameters are available, they will be displayed 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 **Run Report**. 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 **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 box. If you decide to change the start time, make sure the date is in the  
yyyy-mm-dd hh:mm based on the 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, enter the number of minutes in the Interval box.

---

**Caution** – 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 **Start**.
5. To stop a collector, click **Stop**.
6. To view a report, click its link. See “Viewing Reports” on page 523 for more information.

---

## Refreshing a Report

To view the latest information in a report:

1. Select **Configuration > Reports**
2. Click the **Report Cache** tab.
3. Click **Refresh Now**.

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. To change the schedule for refreshing a report, see “Refreshing the Report Cache” on page 259.

If you are still seeing old information after you click **Refresh Now**, verify that your database is being updated within the appropriate time frame for your organization. Some of the information contained in the database depends on whether the collectors are running and when they last ran. To 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 226.

---

## Changing the Formatting of a Report

Reports are displayed in HTML format by default. To change the formatting of the report, select one of the following options from the **Format** menu in Reporter, and then click **Run Report**:

- **HTML (Default)** - Displays the report in a Web page by default.
- **PDF** - 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** - Displays the report in Microsoft Excel, providing you have a copy of Microsoft Excel already installed.
- **XML** - 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. While viewing a report, select the **Open in new window** option.
2. Click **Run Report**.

A new window opens, and it displays just the report.

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 panes, as described below:

- To hide the middle pane, click the border between the tree listing the reports and the main pane and drag the border to the left side of the page, as shown in the following figure:



**FIGURE 13-1** Hiding the Middle Pane






3. To display the middle pane, drag the border to the right.
4. 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
	Protection Explorer
	System Explorer
	Capacity Explorer



**TABLE 13-1** Buttons Displayed When the Left Pane Is Closed (*Continued*)

Button	Provides Access to...
	Performance Explorer
	Event Manager
	Provisioning
	Policy Manager
	Business Tools
	Chargeback
	Reporter

5. To open the left pane, click the section between the arrows in the upper-left border.

## Filtering Data in Global Reports

**Caution** – 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. If your license allows you to 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 see only 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 that you have global reporting set up correctly. See “Setting Up Global Reporter” on page 260.

To filter data global reports:

1. Access Reporter as described in “Accessing Reporter” on page 523.
2. Expand the Global node.

3. Select a global report.
4. In the right pane, click **Filter Data**.
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 is applied to all global reports.

---

## Sending a Report by E-mail

You can e-mail an attached report in PDF, XML or Microsoft Excel format. 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 531.

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 221.
- 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, which occurs every six hours by default. Since the reports display the data that is in the cache, e-mails sent too long after a refresh will show old data. See “Scheduling a Report Cache Refresh” on page 260.

To send a report by e-mail:

1. Access Reporter as described in “Accessing Reporter” on page 523.
2. Expand the tree in the middle pane, and click the report you want to send by e-mail.
3. Click **E-mail this Report**.

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 221.

4. In the top box, enter 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. (Protection Explorer reports only) Select the period of time you want displayed in the report by entering a start date and end date in the appropriate boxes.
6. From the **Format** 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.
7. *Optional*: Modify the subject and the message.
8. Click **OK**.

The report is sent.

---

## Managing E-mail Schedules for Reports

This section contains the following topics:


- “Adding an E-mail Schedule for a Report” on page 531
- “Editing an E-mail Schedule for a Report” on page 534
- “Deleting an E-mail Schedule for a Report” on page 535
- “Viewing E-mail Schedules for a Report” on page 535

### 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.

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 221.
- Schedule your reports to be e-mailed soon after a report cache refresh, which occurs every six hours by default. Since the reports display the data that is in the cache, e-mails sent too long after a refresh will show old data. For example, if you add an e-mail schedule to send a report daily at 7 a.m., and schedule your report cache to refresh 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 260.

- 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, select **Configuration > Reports**, and then click the **Scheduled Deliveries** tab at the top of the screen.
- The elements in the report you send are based on the organizations selected from the Organization Filters tab. For example, assume you belong to two organizations: OnlyHosts and OnlySwitches. If you only select OnlySwitches 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 remains 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 523.
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 **Add E-mail Schedule**.
5. Verify that the Properties tab is displayed.
6. In the top box, enter 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 box, enter a subject for the e-mail messages you plan to send.

---

**Note** – Provide the name of the report in the Subject box so users can distinguish the message from others.

---

8. In the Message box, enter 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 send 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** 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 that the user has an understanding of XML.
10. (Any report having a date parameter or Protection Explorer reports) From the **Include Data for the Last** menu, select the period of time you want displayed in the report.
11. In the Time to Run box, enter the time you want to send the report in the 24-hour format. For example, if you want a report sent at 2:15 p.m., you would enter 14:15 in the Time to Run box.
12. 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** - Enter how frequently (in days) you want the report to be sent. For example, if you enter 15, the report is sent every 15 days.
  - **Weekly** - If you select Weekly, use the **Frequency** menu to select the day of the week on which you want the report sent.
  - **Monthly** - If you select 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, and then select **First** or **Last** from the menu.
    - To send the report on a specified day during the month, select the second option, and then enter the day on which you want the report sent. If you enter a day that is not in the month, (for example 30 for February) the report is sent on the last day of the month.
13. 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.
14. Select the organizations containing the elements you want used in the report. If all organizations are already 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

15. Click **Finish**.

The schedule is created.


## Editing an E-mail Schedule for a Report

---

**Caution** – Only the e-mail schedules created by the current user are listed. To view the e-mail schedules for all reports, select **Configuration > Reports**, and 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 523.
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. For details about the options available for e-mail schedules, see “Adding an E-mail Schedule for a Report” on page 531.


## Deleting an E-mail Schedule for a Report

---

**Caution** – Only the e-mail schedules created by the current user are listed. To view the e-mail schedules for all reports, select **Configuration > Reports**, and then click the **Scheduled Deliveries** tab.

---

To delete an e-mail schedule:

1. Access Reporter as described in “Accessing Reporter” on page 523.
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**.
4. Click the **Delete** () button corresponding to the e-mail schedule you want to remove. The report is deleted.

## Viewing E-mail Schedules for a Report

---

**Caution** – Only the e-mail schedules created by the current user are listed. To view the e-mail schedules for all reports, select **Configuration > Reports**, and 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 523.
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 plus a brief summary of what it is about.
Format	The format of the report sent: <ul style="list-style-type: none"><li>• PDF</li><li>• Microsoft EXCEL</li><li>• XML</li></ul>
Edit	Click the <b>Edit</b> (  ) button to edit a schedule of the report. See “Adding an E-mail Schedule for a Report” on page 531 and “Editing an E-mail Schedule for a Report” on page 534 for information about the options displayed in this window.
Delete	Click the <b>Edit</b> (  ) button to remove the report schedule.

---

## Creating Custom Reports

This section contains the following topics:

- “About Creating Custom Reports” on page 537
- “Configuring Report Designer to Work with the Management Server” on page 538
- “Designing Custom Reports” on page 540
- “Integrating Custom Reports” on page 552
- “Detailed Schema Information” on page 554
- “Views from Previous Releases” on page 598
- “Implementing Custom Reports on Sun Solaris” on page 604



# About Creating Custom Reports

**Caution** – 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.

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 reports as you create them. Once you are satisfied with the customized reports, you can merge them onto the management server so they are accessible from the management server console.

The following figure shows how Report Designer fits into the reporting architecture. Use Report Designer to create the reports, and then deploy the reports on the management server.

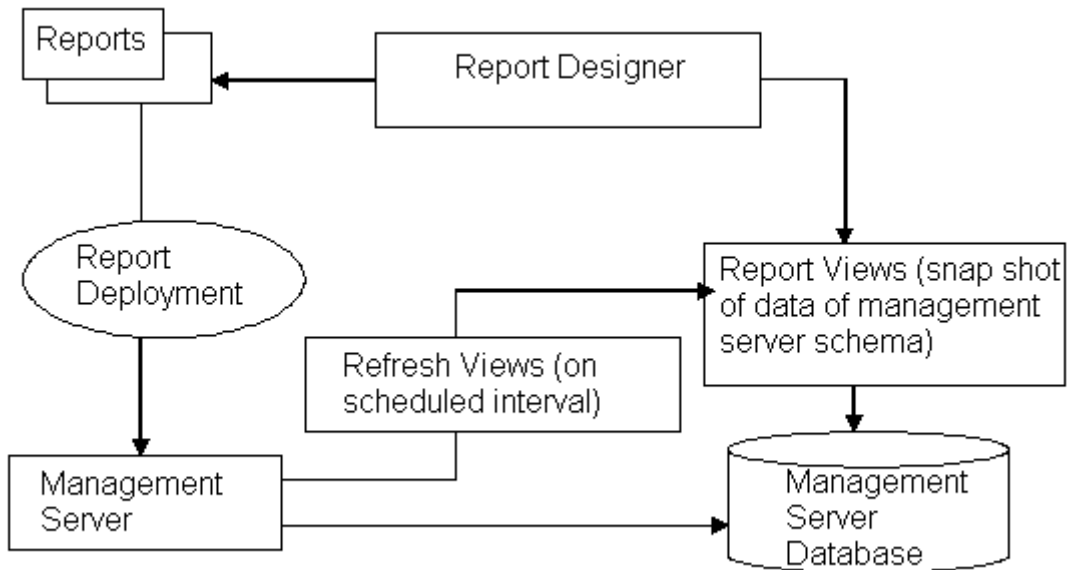


FIGURE 13-4 Report Architecture

## About Materialized Views

When you create reports, you can use pre-existing schema materialized views. A materialized view is a snapshot of data from the database, created from a query. Report Designer refers to these materialized views as “tables.” A materialized view is refreshed based on its collector's schedule for obtaining the latest data. The default refresh time is every six hours. The current and deprecated materialized views are provided at the end of this chapter.

## Configuring Report Designer to Work with the Management Server

---

**Caution** – 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 box, enter a name for the catalog, such as `custom.cat`.
4. Save the catalog in a directory created 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 finish creating the catalog, the Catalog Browser window is displayed. (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 **Open**.

6. In the Catalog Browser window, expand the Default node in the tree, and then right-click **Connection** and select **New Connection**.

7. Verify that the OracleOraHome92TNSListener service is running on the management server.

---

**Note** – 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. Enter the following command for the JDBC driver:

```
oracle.jdbc.driver.OracleDriver
```

- d. Enter the following command in the JDBC URL box:

```
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. Enter the following in the User Name box: `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. Enter the following in the Password box: `appiq`

9. Click **OK**.

Report Designer searches for JDBC driver.

If Report Designer cannot find the JDBC drivers, you may have entered incorrect path information.

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 sure which tables you may be using, you may want to select them all.

11. Select the tables and then click **Add**.

12. When you have finished adding tables, click **Done**.

The tables populate the Catalog Browser window.

13. Click **File > Save Catalog** to save the catalog.

## Designing Custom Reports

---

**Caution** – This section assumes you have already installed and configured Report Designer and integrated it with the management server.

---

The procedures in this section 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.

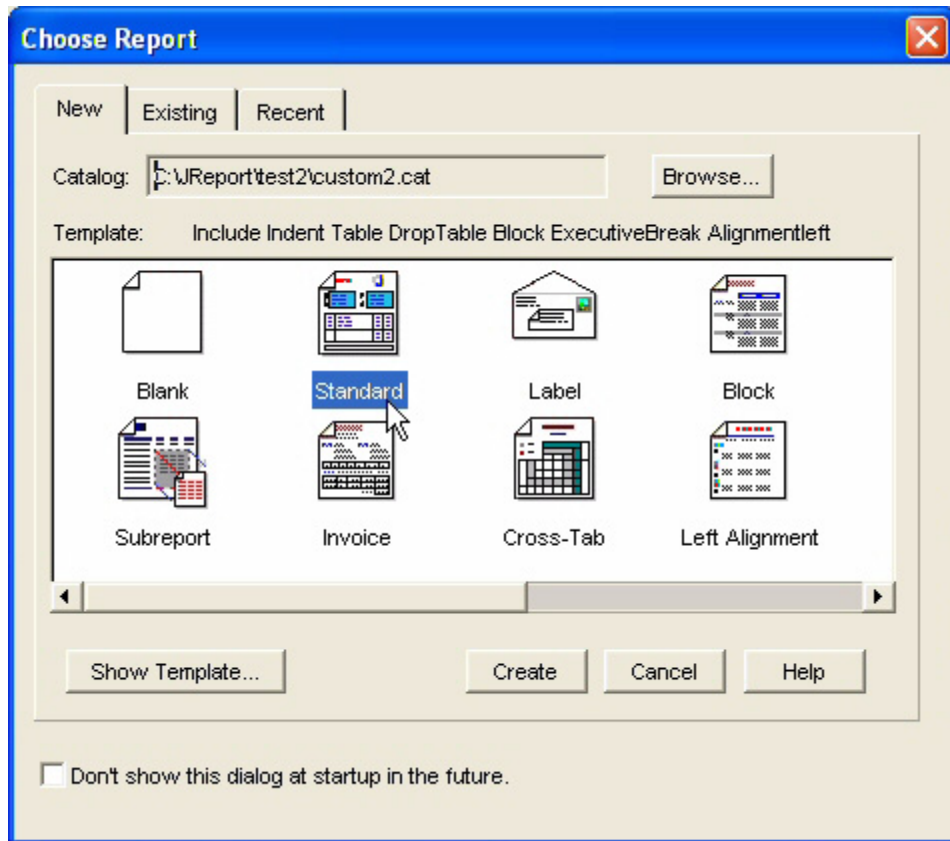
---

## Creating Standard Reports

The Report Form Creation Wizard is used to create Standard reports. The procedure uses Microsoft Windows and Report Designer 6.0.

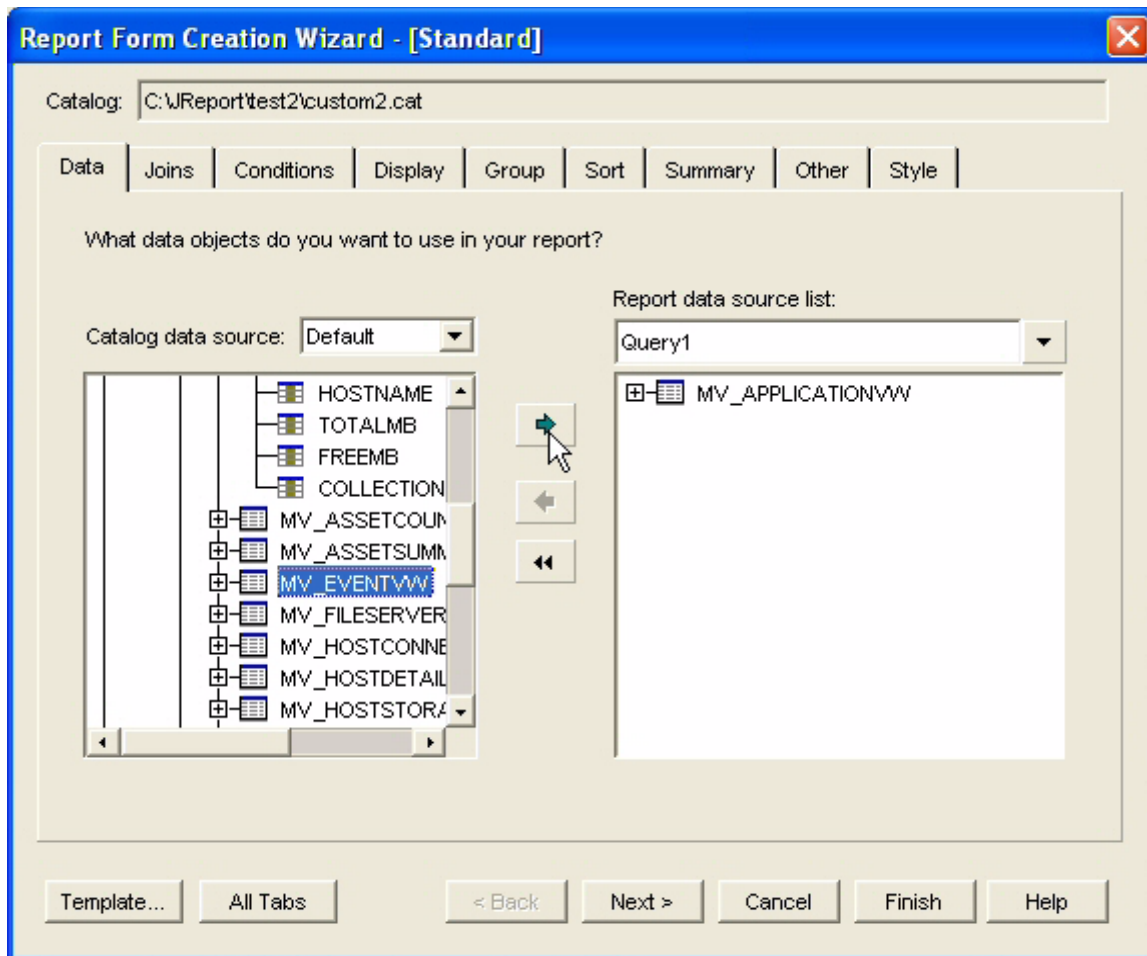
To create standard reports using the Report Form Creation Wizard:

1. Open Report Designer.
2. If the Choose Report screen is not displayed, select **File > New** in Report Designer.
3. Click the **Standard Report** icon and then click **Create** to start the Report Form Creation Wizard.



**FIGURE 13-5** 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. 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 that you do not want all data displayed in the report. To find a definition of the listings in a table, see “Detailed Schema Information” on page 554. When you are done, click **Next**.



**FIGURE 13-6** 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 view (table) 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 criterion is truncated in a table, you can expand the size of the table by clicking a table border and dragging it to the appropriate position.

For definitions of the search criteria, see "Detailed Schema Information" on page 554.

Refer to the online help for Report Designer for more information. When you finish selecting and linking materialized views, click **Next**.

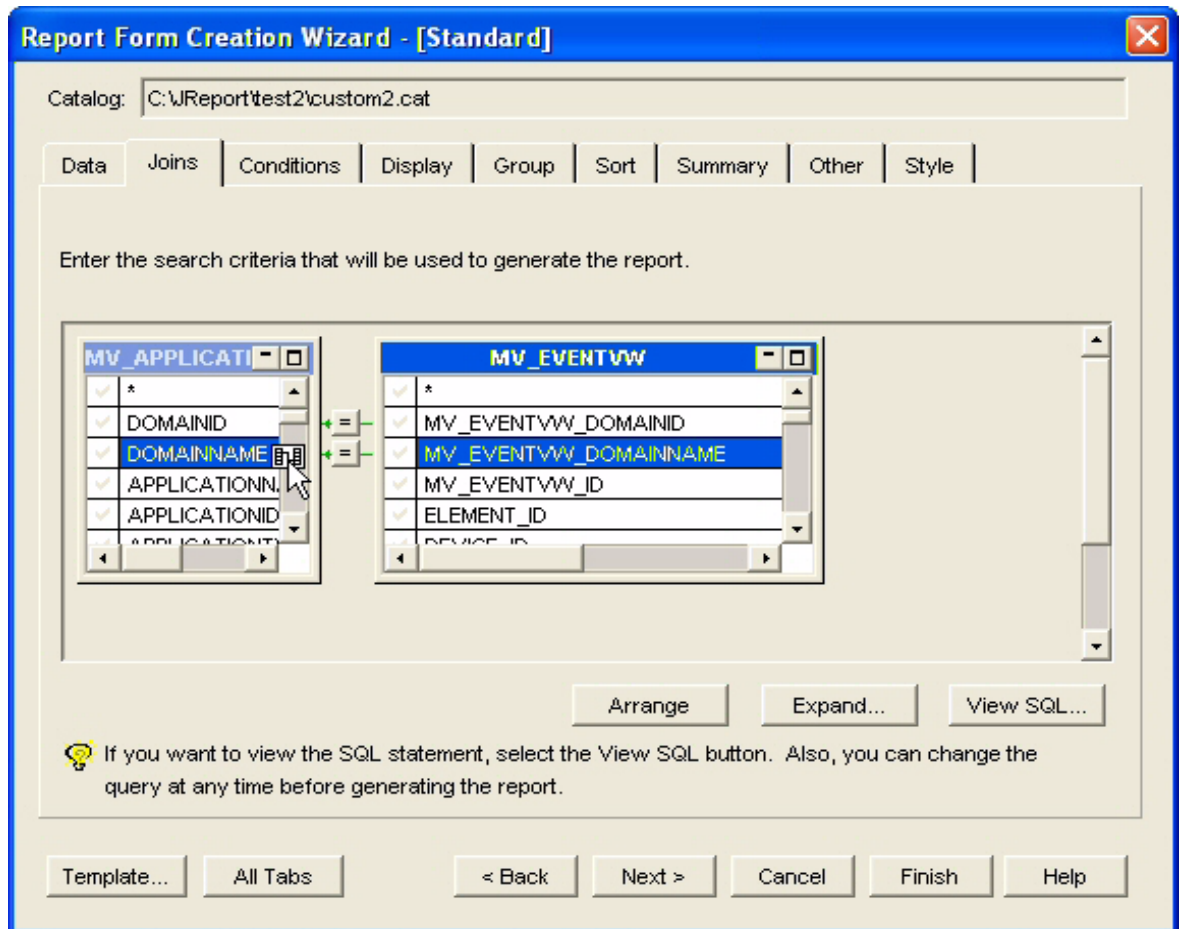


FIGURE 13-7 Linking Common Data in Tables for a Standard Report

7. Enter search criteria that will be used to generate the report. For example, if you want the report to display information only about Oracle applications, enter a search criterion that tells Report Designer to display data from Oracle applications. When you are done, click **Next**.

Report Form Creation Wizard - [Standard]

Catalog: C:\Report\test2\custom2.cat

Data | Joins | Conditions | Display | Group | Sort | Summary | Other | Style

Enter the search criteria that will be used to generate the report.

Expression: MV\_APPLICATION\VV.APPLI ... Operator: = Expression: ORACLE More: End

☒ AND ☐ QBE ☐ Select Distinct ☐ Ignore Predicate If Param Is Null View SQL...

If you want to view the SQL statement, select the View SQL button. Also, you can change the query at any time before generating the report.

Template... All Tabs < Back Next > Cancel Finish Help

**FIGURE 13-8** Creating Search Criteria for Standard Reports

8. To select the data you want displayed in the report, click the data source in the Data Source pane, and then click the arrow pointing right. The report field associated with the selected data 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. Similarly, 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 554.



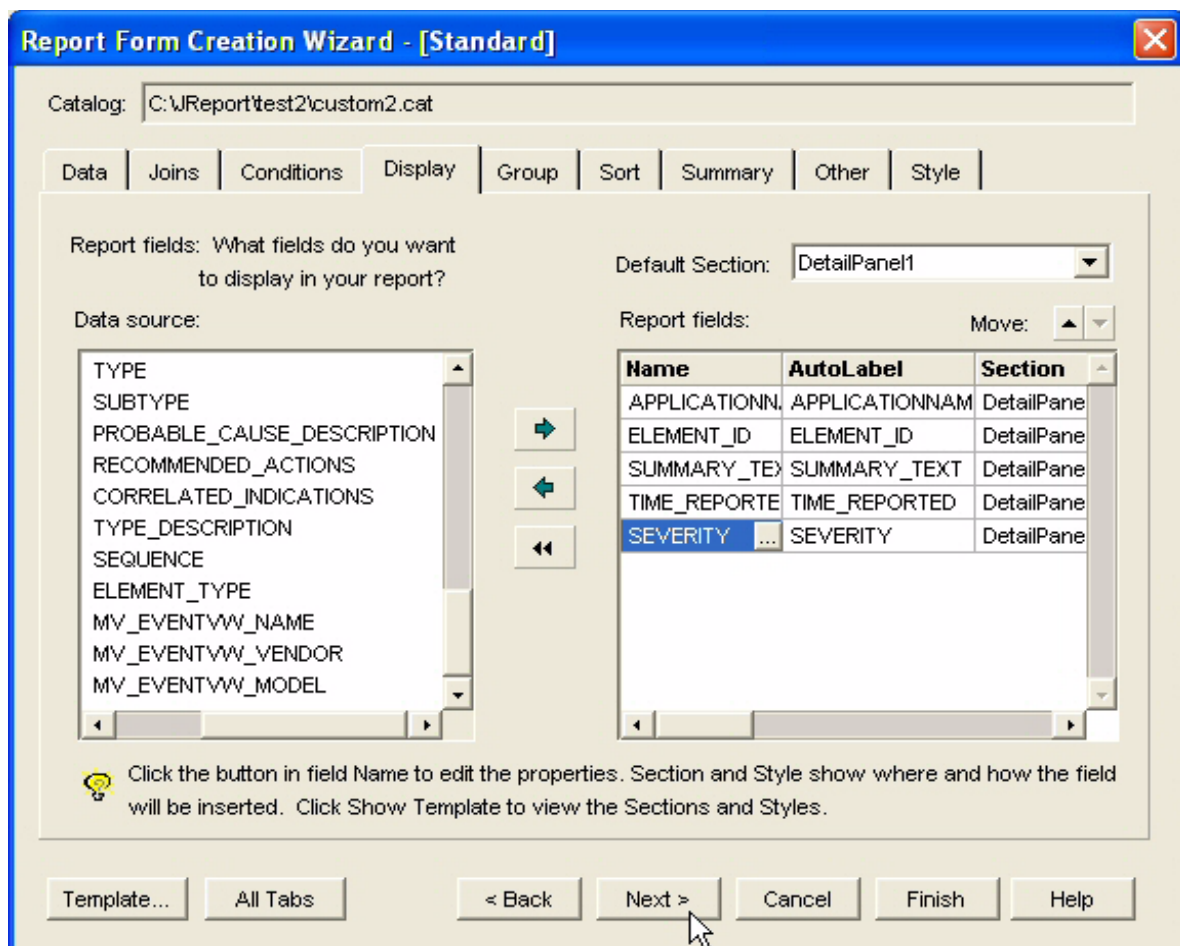


FIGURE 13-9 Deciding Which Data Should Appear in the Report

9. Select the fields in the left pane in the order you want them sorted in your report, and 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. If you then select APPLICATIONID, applications will first be sorted by application name and then by their application ID.

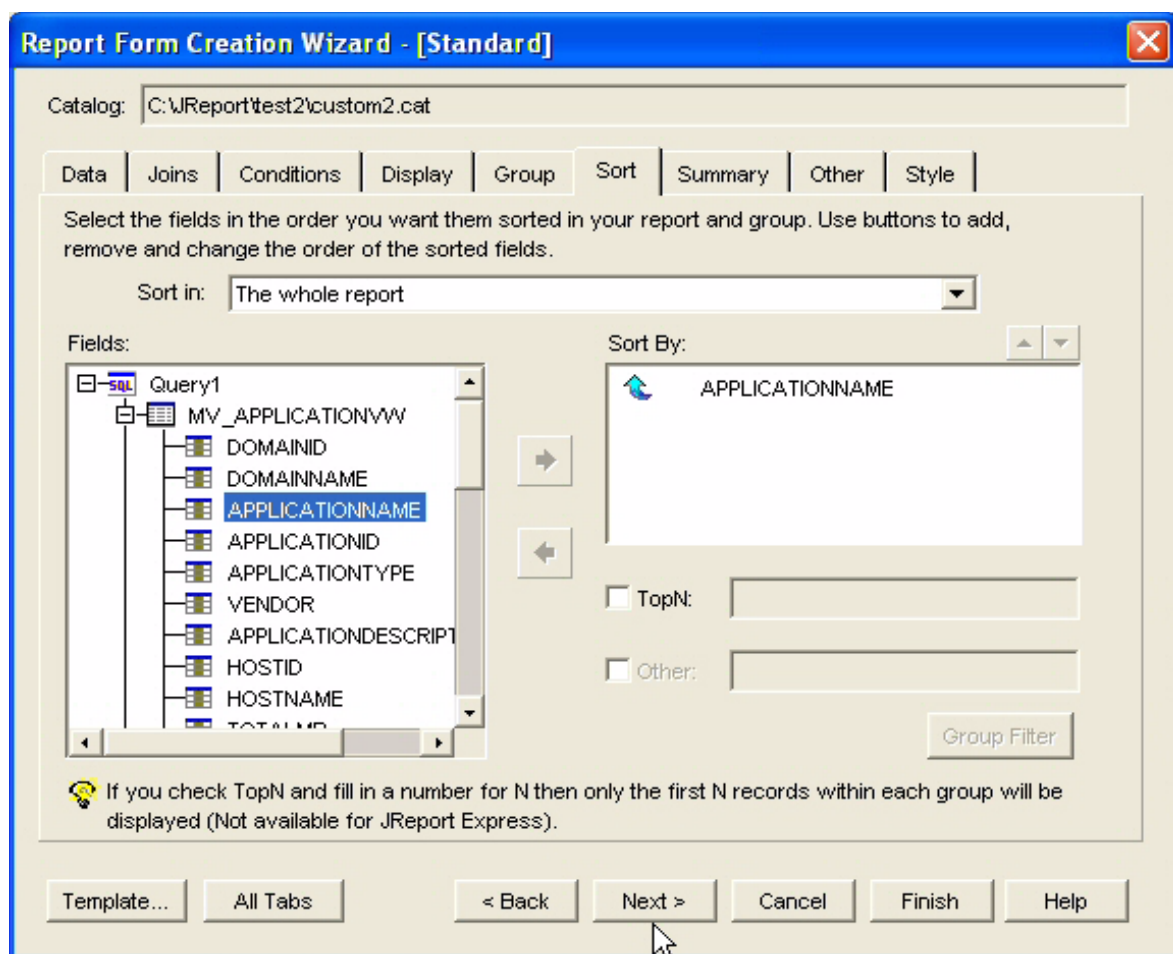


FIGURE 13-10 Sorting Information in the Report

10. Use the Style tab to determine the layout of the report. When you are done, click **Finish**.

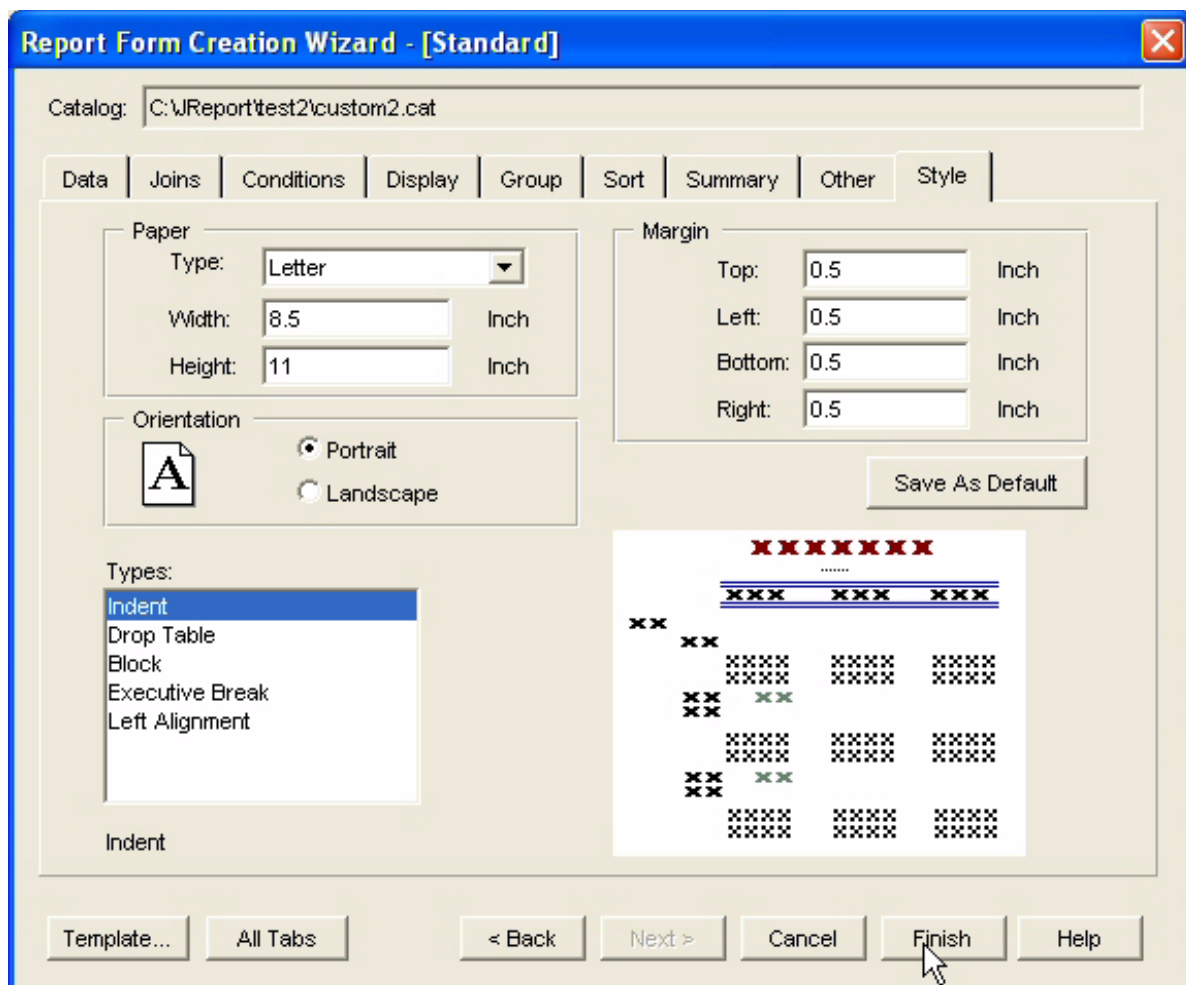


FIGURE 13-11 Selecting the Layout of the Report

The report template is displayed. You will not see any data reported, only placeholders, as shown in the following figure.

<i>Report Title</i>				
APPLICATIONNAME	ELEMENT_ID	SUMMARY_TEXT	TIME_REPORTED	SEVERITY
XXXXXXXXXX				
XXXXXXXXXX	#####	XXXXXXXXXX	MM/dd/yy	#####

FIGURE 13-12 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 537 as a guideline, and then access Reporter on the management server and verify that 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 that 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 that the connection information for the catalog is correct. Open the catalog and expand the Connection node. Verify that the IP address/DNS name listed is correct.
 If you are still having problems, verify that the network from your computer to the management server is stable. Try accessing the management server console from your computer as a test.
13. 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 design the report.

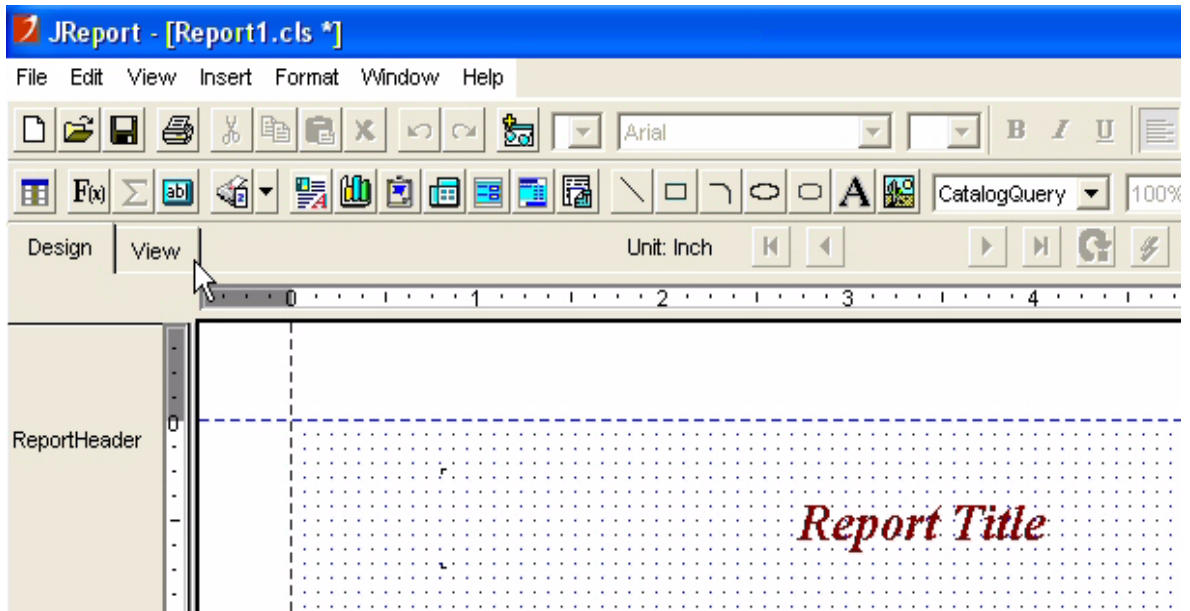


FIGURE 13-13 Result of Clicking the View Tab

## Managing Custom Reports (Importing and Deleting)

You can manage custom reports using the graphical user interface. This interface enables you to import and delete custom reports and is accessed from the **Reports** tab under the **Configuration** choice. The interface eliminates the need to manually deploy custom reports at customer locations.

To use this feature, click **Configuration** on the main screen. Then, click **Reports**. From the **Reports** choices, click **Manage Custom Reports**.

The following screen displays.

Scheduled Deliveries Data Collection Data Aging Report Cleanup Report Cache Global Reporter **Manage Custom Reports**

To import and view a custom report:

1. Package the report definition, report template and database into a zip file.
2. Use the Browse button below to specify the name/path of the zip file.
3. Use the Import button below to import the zip file.
4. View the report from the "System" link of the Reporter
5. For more information Please Click [Help](#)

Import Reports

Zip file name

Imported Reports

Customer/Report Name	Report File(s)	Delete

**FIGURE 13-14** Manage Custom Reports Screen

The screen lists these instructions to import and view a custom report:

1. Package the report definition, report template, and database into a zip file.
2. Use the **Browse** button in the display to specify the name/path of the zip file.
3. Use the **Import** button in the display to import the zip file.
4. View the report from the "system" link of the Reporter.

The instructions also remind you to refer to **Help** for additional information.

After you package the necessary files into a .zip file, use the **Browse** button to navigate to the desired .zip file. When ready to upload the selected file, click the **Import** button.

After importing, a screen display similar to the following shows the imported report files. The imported reports can then be viewed from the "system" link of the Reporter.

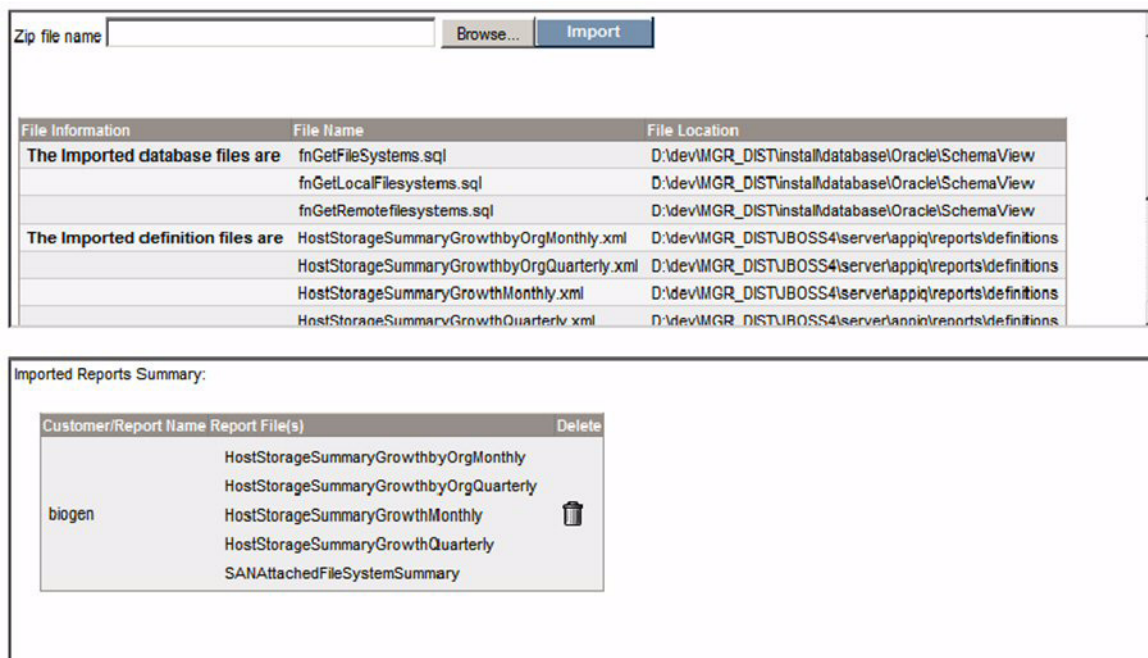


FIGURE 13-15 Screen Displays Custom Reports Available

## Importing Custom Reports Error Messages

If you attempt to load a file type other than .zip, the system will give an error message saying that it is not a valid file type.

If you try to import a report that already exists, the system gives an error message saying the report already exists. In this situation, you must delete the report that already exists, then import the new report.

## Deleting Custom Reports

If you want to delete imported custom reports by clicking **Delete**, the system displays a warning message. If you confirm that you want to delete by clicking **Yes**, the system deletes the reports and supporting files. You can confirm that the files are deleted by navigating to the **Reporter** tab and clicking the link.

# Integrating Custom Reports

When you are satisfied with your custom report, you must integrate it with the management server so that other users can access it. To make the report accessible, you must first deploy the custom report to the management server, and then you must integrate the report so that it is accessible from Reporter. Custom reports appear with their own icon in the reports tree.

---

**Caution** – 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, if you have two reports that use a different catalog, you would create two subdirectories under reports\custom, one for each catalog. 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 deploy and 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, and 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\customTreeNodes if it does not 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. Populate the tree in Reporter with your custom reports.

- a. Create the directory %JBOSS4\_DIST%\server\appiq\reports\definitions\custom if it does not already exist.
- b. Create an XML file in that directory for your report, for example:  
CustomHostReport.xml
- c. Open the new XML file, and 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 must 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 following directory: %JBOSS4\_DIST%\server\appiq\reports\definitions\custom

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 following directory:

%JBOSS4\_DIST%\server\appiq\reports\definitions\custom


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.

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

%JBOSS4\_DIST%\server\appiq\reports\customTreeNodes. Do not use 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. 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. You must use the new views for any new report development. Verify that your existing reports will work correctly against these new views. Some of the views have changed and may not work in existing reports. See “Views from Previous Releases” on page 598 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 570.
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 570.
MVC_OPTIONALTABLEVW	Provides information about the names and values of the optional value. See Table 13-25, “MVC_OPTIONALTABLEVW,” on page 572.
MVC_ASSETSUMMARYVW	Provides summary information about assets.
MVCA_BU_MASTERSERVERSUMMARY	Provides summary information about the backup manager. See Table 13-46, “MVCA_BU_MASTERSERVERSUMMARY,” on page 582.
MVCA_BU_MEDIASERVERSUMMARY	Provides summary information about media managers. See Table 13-47, “MVCA_BU_MEDIASERVERSUMMARY,” on page 583.

**TABLE 13-3** Description of the Report Views (*Continued*)

Materialized View (Tables)	Description
MVCA_BU_CLIENTSUMMARY	Provides summary information about clients. See Table 13-48, "MVCA_BU_CLIENTSUMMARY," on page 583.
MVCA_BU_MEDIASUMMARY	Provides summary information about the media. See Table 13-49, "MVCA_BU_MEDIASUMMARY," on page 583.
MVCA_BU_JOBSSUMMARY	Provides summary information about jobs. See Table 13-50, "MVCA_BU_JOBSSUMMARY," on page 585.
MVCA_BU_LIBRARYSUMMARY	Provides summary information about libraries. See Table 13-51, "MVCA_BU_LIBRARYSUMMARY," on page 585.
MVC_HOSTSUMMARYVW	Provides summary information about hosts. See Table 13-4, "MVC_HOSTSUMMARYVW," on page 560.
MVC_APPLICATIONSUMMARYVW	Provides summary information about applications. See Table 13-33, "MVC_APPLICATIONSUMMARYVW," on page 577.
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 578.
MVCA_DBAPPCAPACITYVW	Provides capacity information for a supported database application. See Table 13-35, "MVCA_DBAPPCAPACITYVW," on page 578.
MVCA_EXCHANGEAPPCAPACITYVW	Provides capacity information for Microsoft Exchange. See Table 13-36, "MVCA_EXCHAPPCAPACITYVW," on page 579.
MVCA_VIRTUALAPPCAPACITYVW	Provides capacity information for a virtual application. See Table 13-37, "MVCA_VIRTUALAPPCAPACITYVW," on page 579.
MVC_EVENTSVW	Provides event information, such as description and time reported. See Table 13-18, "MVC_EVENTSVW," on page 568.
MVCA_FSRM_VOLUMESUMMARYVW	Provides summary information about File Server SRM volumes. Table 13-38, "MVCA_FSRM_VOLUMESUMMARYVW," on page 579.

**TABLE 13-3** Description of the Report Views (*Continued*)

Materialized View (Tables)	Description
MVCA_FSRM_AGESUMMARYVW	Provides summary information about ages in File Server SRM. See Table 13-39, "MVCA_FSRM_AGESUMMARYVW," on page 580.
MVCA_FSRM_EXTDETAILSUMMARYVW	Provides summary information about extent details. See Table 13-40, "MVCA_FSRM_EXTDETAILSUMMARYVW," on page 580.
MVCA_FSRM_DIRDETAILSUMMARYVW	Provides information about directories in File Server SRM. See Table 13-41, "MVCA_FSRM_DIRDETAILSUMMARYVW," on page 580.
MVCA_FSRM_USERSUMMARYVW	Provides information about users in File Server SRM. See Table 13-42, "MVCA_FSRM_USERSUMMARYVW," on page 581.
MVCA_FSRM_TOPNFILES	
MVCA_FSRM_AGEDFILEDETAILS	Provides information about the age properties of files in File Server SRM. See Table 13-44, "MVCA_FSRM_AGEDFILEDETAILS," on page 581.
MVCA_FSRM_LARGEDIRINFO	
MVC_STORAGEPOOLSUMMARYVW	Provides summary information about storage pools. See Table 13-9, "MVC_STORAGEPOOLSUMMARYVW," on page 562.
MVC_STORAGESYSTEMSUMMARYVW	Provides summary information about a storage system. See Table 13-8, "MVC_STORAGESYSTEMSUMMARYVW," on page 562.
MVC_STORAGEVOLUMESUMMARYVW	Provides summary information about a storage volume. See Table 13-10, "MVC_STORAGEVOLUMESUMMARYVW," on page 563.
MVC_HOSTDISKDRIVEVW	Provides information about host disk drives. See Table 13-7, "MVC_HOSTDISKDRIVEVW," on page 562.
MVC_HOSTVOLUMESUMMARYVW	Provides summary information about host volumes. See Table 13-6, "MVC_HOSTVOLUMESUMMARYVW (logical volumes)," on page 561.

**TABLE 13-3** Description of the Report Views (*Continued*)

Materialized View (Tables)	Description
MVC_DISKEXTENTSUMMARYVW	Provides summary information about disk extents. See Table 13-26, "MVC_DISKEXTENTSUMMARYVW," on page 572.
MVC_STORAGEVOLUMEPORTS	Provides information about storage volume ports. See Table 13-27, "MVC_STORAGEVOLUMEPORTS," on page 573.
MVC_VOLUMEDISKDRIVEVW	Provides information about volume disk drives. See Table 13-28, "MVC_VOLUMEDISKDRIVEVW," on page 573.
MVC_STORAGEPROCESSORSUMMARYVW	Provides information about storage processors. Table 13-29, "MVC_STORAGEPROCESSORSUMMARYVW," on page 574.
MVC_DISKDRIVESUMMARYVW	Provides summary information about disk drives. See Table 13-30, "MVC_DISKDRIVESUMMARYVW," on page 574.
MVC_DISK_EXTENTVW	Provides information about disk extents. See Table 13-31, "MVC_DISK_EXTENTVW," on page 575.
MVC_SWITCHSUMMARYVW	Provides summary information about switches. See Table 13-11, "MVC_SWITCHSUMMARYVW," on page 564.
MVC_PORTSUMMARYVW	Provides summary information about ports. See Table 13-12, "MVC_PORTSUMMARYVW," on page 565.
MVC_ZONESUMMARYVW	Provides summary information about zones. See Table 13-13, "MVC_ZONESUMMARYVW," on page 566.
MVC_ZONEVW	Provides information about zones. See Table 13-14, "MVC_ZONEVW," on page 567.
MVC_PATHVW	Provides path information. See Table 13-15, "MVC_PATHVW," on page 567.
MVC_SUBPATHVW	Provides subpath information. See Table 13-16, "MVC_SUBPATHVW," on page 567.
MVC_MULTIPATHVW	Provides multipath information. See Table 13-17, "MVC_MULTIPATHVW," on page 568.
MVC_SWITCHCONFIGVW	Provides switch configuration information. See Table 13-24, "MVC_SWITCHCONFIGVW," on page 572.

**TABLE 13-3** Description of the Report Views (*Continued*)

Materialized View (Tables)	Description
MVC_HOSTCAPACITYVW	Provides host capacity information. See Table 13-21, "MVC_HOSTCAPACITYVW," on page 571.
MVC_STORAGESYSTEMCONFIGVW	Provides storage system configuration information. See Table 13-22, "MVC_STORAGESYSTEMCONFIGVW," on page 571.
MVC_STORAGEPOOLCONFIGVW	Provides information about storage pool configurations. See Table 13-23, "MVC_STORAGEPOOLCONFIGVW," on page 571.
MVC_DISCOVERYDETAILSVW	Provides information about quarantined elements. See Table 13-54, "MVC_DISCOVERYDETAILSVW," on page 586.
MVC_HOSTRELATIONVW	Provides information about cluster hosts and related host members. See Table 13-55, "MVC_HOSTRELATIONVW," on page 587.
MVC_APPLICATIONRELATIONVW	Provides information about cluster applications and related application members. Table 13-56, "MVC_APPLICATIONRELATIONVW," on page 587.
MVC_STORGETIERDETAILVW	Provides information about Chargeback storage tiers. See Table 13-57, "MVC_STORAGETIERDETAILVW," on page 588.
MVCA_BU_OPTIONALTABLEVW	Provides information about Protection Explorer optional values. See Table 13-58, "MVCA_BU_OPTIONALTABLEVW," on page 588.
MVCA_BU_DRIVESTATVW	Provides information about Protection Explorer drive stats. See Table 13-59, "MVCA_BU_DRIVESTATVW," on page 589.
MVCA_EXCHMAILBOXDETAILVW	Provides Microsoft Exchange mailbox details. See Table 13-60, "MVCA_EXCHMAILBOXDETAILVW," on page 589.
MVCA_EXCHPUBLICFOLDERDETAILVW	Provides Microsoft Exchange public folder details. See Table 13-61, "MVCA_EXCHPUBLICFOLDERDETAILVW," on page 590.

**TABLE 13-3** Description of the Report Views (*Continued*)

Materialized View (Tables)	Description
MVCA_EXCHANGESTORESUMMARYVW	Provides Microsoft Exchange store details. See Table 13-62, "MVCA_EXCHANGESTORESUMMARYVW," on page 590.
MVCA_EXCHSTORGROUPSUMMARYVW	Provides Microsoft Exchange storage group details. See Table 13-63, "MVCA_EXCHSTORGROUPSUMMARYVW," on page 591.
MVCA_FSRM_FILEREPORTDATAVW	Provides information about File Server SRM files. See Table 13-64, "MVCA_FSRM_FILEREPORTDATAVW," on page 591.
MVCA_FSRM_DIRREPORTDATAVW	Provides information about File Server SRM directories. See Table 13-65, "MVCA_FSRM_DIRREPORTDATAVW," on page 591.
MVCA_FSRM_REPORTRULEVW	Provides information about File Server SRM rules. See Table 13-66, "MVCA_FSRM_REPORTRULEVW," on page 592.
MVCS_HOSTMEMORYSTATSVW	Provides host memory performance statistics. See Table 13-67, "MVCS_HOSTMEMORYSTATSVW," on page 592.
MVCS_HOSTCPUSTATSVW	Provides host CPU performance statistics. See Table 13-68, "MVCS_HOSTCPUSTATSVW," on page 593.
MVCS_EVACTRLSTATSVW	Provides EVA controller statistics. See Table 13-69, "MVCS_EVACTRLSTATSVW," on page 593.
MVCS_EVADISKSTATSVW	Provides EVA disk statistics. See Table 13-70, "MVCS_EVADISKSTATSVW," on page 594.
MVCS_EVAHOSTFCPORTSTATSVW	Provides EVA FC port statistics. See Table 13-71, "MVCS_EVAHOSTFCPORTSTATSVW," on page 595.
MVCS_EVASPAGVOLUMESTATSVW	Provides EVA storage pool statistics. See Table 13-72, "MVCS_EVASPAGVOLUMESTATSVW," on page 596.
MVCS_EVASSAGVOLUMESTATSVW	Provides EVA storage system AG statistics. See Table 13-72, "MVCS_EVASPAGVOLUMESTATSVW," on page 596.

**TABLE 13-3** Description of the Report Views (*Continued*)

Materialized View (Tables)	Description
MVCS_EVASTORAGESYSTEMSTATSVW	Provides EVA storage system statistics. See Table 13-73, "MVCS_EVASTORAGESYSTEMSTATSVW," on page 597.
MVCS_EVAVOLUMESTATSVW	Provides EVA volume statistics. See Table 13-74, "MVCS_EVAVOLUMESTATSVW," on page 597.
MVIEWCORE_STATUS	Provides information about the core views. The core views are the views starting with mvc, mvca, and mvcs. See Table 13-52, "MVIEWCORE_STATUS," on page 586.
MVIEW_STATUS	Shows the status of the materialized views for reports. Table 13-53, "MVIEW_STATUS," on page 586.

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
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
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\_STORAGEPOOLSUMMARYVW**

Column Name	Type	Description
StoragePoolID	NUMBER(38)	Storage Pool ID
StoragePoolName	VARCHAR2(256)	Pool Name

**TABLE 13-9 MVC\_STORAGEPOOLSUMMARYVW (Continued)**

Column Name	Type	Description
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
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)

**TABLE 13-10 MVC\_STORAGEVOLUMESUMMARYVW (Continued)**

Column Name	Type	Description
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
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)	

**TABLE 13-11 MVC\_SWITCHSUMMARYVW (Continued)**

Column Name	Type	Description
ZoningInstalled	NUMBER(1)	
MaxModuleNumber	NUMBER(38)	
CurrentZoningEnforcement	VARCHAR2(254)	
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)	
Physical_State	VARCHAR2(254)	
Port_ID	NUMBER(38)	
Port_Speed	NUMBER(38)	

**TABLE 13-12** MVC\_PORTSUMMARYVW *(Continued)*

Column Name	Type	Description
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)	AppIQ 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
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
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

**TABLE 13-16 MVC\_SUBPATHVW (Continued)**

Column Name	Type	Description
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)	



**TABLE 13-18** MVC\_EVENTSVW (Continued)

Column Name	Type	Description
Element_ID	NU MBE R(38)	
Device_ID	NU MBE R(38)	
Summary_Text	VAR CHA R2(4 000)	
Time_Reported	Date	
Severity	NU MBE R(38)	
Cleared	NU MBE R(1)	
Source	VAR CHA R2(2 54)	
Type	NU MBE R(38)	
SubType	NU MBE R(38)	
Probable_Cause_Description	VAR CHA R2(4 000)	
Recommended_Actions	VAR CHA R2(2 54)	
Correlated_Indications	VAR CHA R2(2 54)	

**TABLE 13-18** MVC\_EVENTSVW *(Continued)*

Column Name	Type	Description
Type_Description	VAR CHA R2(2 54)	
Element_type	NU MBE R(38)	
Name	VAR CHA R2(2 56)	
Vendor	VAR CHA R2(2 54)	
Model	VAR CHA R2(2 54)	

**TABLE 13-19** MVC\_ORGANIZATIONVW

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\_ORGRELATIONVW

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\_HOSTCAPACITYVW**

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
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

**TABLE 13-23 MVC\_STORAGEPOOLCONFIGVW** (Continued)

Column Name	Type	Description
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
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
DiskExtentName	VARCHAR2(256)	Name of the extent
DomainID	NUMBER(38)	Domain ID
DiskExtentDescription	VARCHAR2(1024)	Description of the extent
Status	NUMBER(38)	Operational status

**TABLE 13-26 MVC\_DISKEXTENTSUMMARYVW** *(Continued)*

Column Name	Type	Description
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 (1 of 2)**

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

**TABLE 13-30 MVC\_DISKDRIVESUMMARYVW** (Continued) (2 of 2)

Column Name	Type	Description
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

**TABLE 13-32 MVC\_ASSETSUMMARY (Continued)**

Column Name	Type	Description
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
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



**TABLE 13-32 MVC\_ASSETSUMMARY (Continued)**

Column Name	Type	Description
DATEPURCHASED	DATE	Asset Date purchased
COST	NUMBER(36,2)	Asset cost
VALUE	NUMBER(36,2)	
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
ASSETFIXCOSTTAXPERDEPT ERYEAR	NUMBER	Asset Fixed cost tax per department per Year
STORAGFIXCOSTTAXPERDEPT PERYEAR	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
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)	
TotalMB	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)	

**TABLE 13-38** MVCA\_FSRM\_VOLUMESUMMARYVW (Continued)

Column Name	Type	Description
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\_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 13-44** MVCA\_FSRM\_AGEDFILEDETAILS

Column Name	Type	Description
VolumeID	NUMBER(38)	

**TABLE 13-44** MVCA\_FSRM\_AGEDFILEDETAILS (Continued)

Column Name	Type	Description
FileName	Varchar2(254)	
FileSize	NUMBER(38)	
FileAge	NUMBER(38)	
Timestamp	Timestamp(6)	
DomainID	NUMBER(38)	

**TABLE 13-45** MVCA\_FSRM\_LARGEDIRINFO

Column Name	Type	Description
VolumeID	NUMBER(38)	
DirName	Varchar2(256)	
DirSize	NUMBER(38)	
TotalFiles	NUMBER(38)	
TotalDirs	NUMBER(38)	
Timestamp	Timestamp(6)	
DomainID	NUMBER(38)	

**TABLE 13-46** MVCA\_BU\_MASTERSERVERSUMMARY

Column Name	Type	Description
MasterServerID	NUMBER(38)	
MasterServerName	Varchar2(256)	
HostID	NUMBER(38)	
Vendor	Varchar2(254)	
Description	Varchar2(1024)	
Status	NUMBER	
ProductName	Varchar2(254)	
LicenseKey	Varchar2(256)	
LicenseFeatures	Varchar2(256)	
DomainID	NUMBER(38)	

**TABLE 13-47 MVCA\_BU\_MEDIASERVERSUMMARY**

Column Name	Type	Description
MediaServerID	NUMBER	
MediaServerName	Varchar2(256)	
MasterServerID	NUMBER	
HostID	NUMBER(38)	
Vendor	Varchar2(254)	
Description	Varchar2(1024)	
Status	NUMBER	
ProductName	Varchar2(254)	
LicenseKey	Varchar2(256)	
LicenseFeatures	Varchar2(256)	
DomainID	NUMBER(38)	

**TABLE 13-48 MVCA\_BU\_CLIENTSUMMARY**

Column Name	Type	Description
ClientID	NUMBER(38)	
ClientName	Varchar2(256)	
MasterServerID	NUMBER	
HostID	NUMBER(38)	
Vendor	Varchar2(254)	
Description	Varchar2(1024)	
Status	NUMBER	
ProductName	Varchar2(254)	
DomainID	NUMBER(38)	

**TABLE 13-49 MVCA\_BU\_MEDIASUMMARY**

Column Name	Type	Description
MediaID	NUMBER(38)	
MediaName	Varchar2(256)	
TapeLibraryID	NUMBER	
PoolID	NUMBER	

**TABLE 13-49** MVCA\_BU\_MEDIASUMMARY (Continued)

Column Name	Type	Description
MasterServerID	NUMBER(38)	
TLMediaID	Varchar2(32)	
Type	Varchar2(32)	
Barcode	Varchar2(32)	
MediaPoolName	Varchar2(256)	
RobotType	Varchar2(32)	
RobotNumber	NUMBER	
RobotSlot	NUMBER	
RobotHost	Varchar2(128)	
VolumeGroup	Varchar2(64)	
Created	Date	
Assigned	Date	
LastMounted	Date	
FisrtMounted	Date	
ExpirationDate	Date	
NumberOfMounths	NUMBER	
MaxMountsAllocated	NUMBER	
Density	Varchar2(64)	
TimeAllocated	Date	
LastWritten	Date	
Expir	Varchar2(32)	
LastRead	Date	
Mbytes	NUMBER	
NImages	NUMBER	
VImages	NUMBER	
RL	Varchar2(64)	
TotalRestores	NUMBER	
MediaStatus	Varchar2(16)	
Vendor	Varchar2(254)	
Description	Varchar2(1024)	
DomainID	NUMBER(38)	



**TABLE 13-50 MVCA\_BU\_JOBSUMMARY**

Column Name	Type	Description
JobID	NUMBER(38)	
TemplateID	NUMBER	
TemplateName	Varchar2(64)	
MasterServerID	NUMBER	
ClientID	NUMBER	
BUJobID	NUMBER	
JobState	Varchar2(16)	
JobStatus	Varchar2(16)	
ScheduleName	Varchar2(32)	
StorageUnit	Varchar2(64)	
BUTargetServer	Varchar2(128)	
FilesLastWritten	NUMBER	
StartTime	Date	
EndTime	Date	
Description	Varchar2(256)	
Time	NUMBER	
RetentionPeriod	Varchar2(16)	
Compression	Varchar2(16)	
Priority	Varchar2(16)	
KBLastWritten	NUMBER	
FileListCount	NUMBER	

**TABLE 13-51 MVCA\_BU\_LIBRARYSUMMARY**

Column Name	Type	Description
TapeLibraryID	NUMBER	
TapeLibraryName	Varchar2(256)	
Vendor	Varchar2(254)	
Description	Varchar2(1024)	
MediaServerID	NUMBER	
MasterServerID	NUMBER	
Type	Varchar2(64)	

**TABLE 13-51** MVCA\_BU\_LIBRARYSUMMARY (Continued)

Column Name	Type	Description
RobotType	Varchar2(64)	
RobotNumber	NUMBER	
TotalNoOfSlots	NUMBER	
TotalSlotsInUse	NUMBER	
TotalNumberOfDrives	NUMBER	
RobotDevicePath	Varchar2(128)	
DomainID	NUMBER(38)	

**TABLE 13-52** MVIEWCORE\_STATUS

Name	Type
MVIEWNAME	NOT NULL VARCHAR2(30)
LAST_REFRESH_TIME	DATE
TOTALREFRESHTIME	VARCHAR2(32)
STATUS	VARCHAR2(10)

**TABLE 13-53** MVIEW\_STATUS

Name	Type
MVIEWNAME	NOT NULL VARCHAR2(30)
LAST_REFRESH_TIME	DATE
TOTALREFRESHTIME	VARCHAR2(32)
STATUS	VARCHAR2(10)

**TABLE 13-54** MVC\_DISCOVERYDETAILSVW

Name	Description
ElementID	ID of the quarientend element

**TABLE 13-54 MVC\_DISCOVERYDETAILSVW**

Name	Description
ElementName	Name of the quarientied element
Address	IP address of the element
ElementType	Type of element
DiscoveryGroup	Name of the discovery group
Enabled	Enabled quarientied or not
Status	Status of the element
ParentAddressID	Parent address
Mapping_version	

**TABLE 13-55 MVC\_HOSTRELATIONVW**

Name	Description
DomainID	Domain ID
ID	ID of the cluster
ClusterHostID	ID of the cluster host
ClusterHostname	Cluster Host name
MemberHostID	Member host ID
MemberRelation	Member relation
State	State of the cluster
CluterHostModelType	Model type
CluterHostVendor	Vendor of the cluster
Description	Description
Status	Status of the cluster
ObjectType	Object type

**TABLE 13-56 MVC\_APPLICATIONRELATIONVW**

Name	Description
ApplicationClusterID	ID of the cluster application

**TABLE 13-56 MVC\_APPLICATIONRELATIONVW**

<b>Name</b>	<b>Description</b>
ApplicationClusterName	Name of the cluster application
AppClusterMemberID	ID of the application cluster member
AppClusterMemberName	Name of the application cluster member

**TABLE 13-57 MVC\_STORAGETIERDETAILVW**

<b>Name</b>	<b>Description</b>
Memberid	ID for the storage members
Domainid	Domain of the member
Member	Name of the member
SSID	Storage system ID of the member
Storagesystem	Name of the storage system
TierID	Storage tier id
TierName	Name of the storage tier
CostPerGB	cost per GB of the tier
TotalCapacity	Total capacity of the member
Storagetype	Type of storage

**TABLE 13-58 MVCA\_BU\_OPTIONALTABLEVW**

<b>Name</b>	<b>Description</b>
Basetableid	ID of the basetable
Basetablename	Name of the basetable
Optionalname	Optional name
Optionalvalue	Optional value

**TABLE 13-59** MVCA\_BU\_DRIVESTATVW

Name	Description
Driveid	ID of the drive
Collectiontime	Timestamp of the collection
ID	
Robotnumber	Robot number
Serialnumber	Serial number
Status	Status of the drive
Devicetime	
APPLICATION	Application name
Mediainuse	Media usage

**TABLE 13-60** MVCA\_EXCHMAILBOXDETAILVW

Name	Description
UserMailboxID	Mailbox ID
UserMailBoxName	Mailbox name
ServerName	Name of Exchange server
StoreID	Store ID of the mailbox
MailboxMessageSizeBytes	Messages size
UserMailBoxSizebytes	Mailbox size
Legacy_dn	
EmailAddress	Email address
NTUserName	Name of NT user
LastLogonTime	Timestamp of last logon
Storage_limit_info	Storage limit
CountofNormalMessages	Count of normal messages
CountofAssociatedMessages	Count of attachment messages
Applicationid	ID of the exchange application

**TABLE 13-61** MVCA\_EXCHPUBLICFOLDERDETAILVW

Name	Description
PublicFolderID	Public folder ID
PublicFolderName	Public folder name
EmailAddress	Email address of the user
Legacy_dn	
Server_name	Name of the server
Applicationid	ID of the exchange application
FolderPath	Path of the folder
StoreID	ID of the store
CountofContacts	Count of contacts in the mailbox
CountofMessages	Count of messages
Associated_content_count	Associated messages count
MessageSizeinFolderbytes	Message size in bytes
FolderSizeBytes	Public folder size
CreationTime	Timestamp of creation
LastAccessTime	Timestamp of last access
LastModifiedTime	Timestamp of the last modification of the file
CountofOwners	Count of owner of folder

**TABLE 13-62** MVCA\_EXCHANGESTORESUMMARYVW

Name	Description
StoreID	ID of the store
StorageGroupID	Storage group ID
ActiveDirectoryName	Directory name
StoreName	Name of the store
Filepath	File path of the store
Isonline	Online status
Private_store	

**TABLE 13-63** MVCA\_EXCHSTORGROUPSUMMARYVW

Name	Description
StorageGroupID	ID of the storage group
StorageGroupName	Name of the storage group
ActiveDirectoryName	Name of active directory
ApplicationID	Application ID

**TABLE 13-64** MVCA\_FSRM\_FILEREPORTDATAVW

Name	Description
Volumeid	ID of FSRM volume
Volumename	Name of volume
Ruleid	Rule ID
Rulename	Name of the rule created
Fullpath	Path of the file
Filename	Name of the file
Totalsize	Total size of file
Accesstime	Timestamp of access time
Createtime	Timestamp of creation time
Modifytime	Timestamp of modified time
Attributes	
Owner	Owner name of file

**TABLE 13-65** MVCA\_FSRM\_DIRREPORTDATAVW

Name	Description
Volumeid	ID of FSRM volume
Volumename	Name of volume
Ruleid	Rule ID
Rulename	Name of the rule created

**TABLE 13-65** MVCA\_FSRM\_DIRREPORTDATAVW

Name	Description
Fullpath	Path of the file
Filename	Name of the file
Totalsize	Total size of file
Accesstime	Timestamp of access time
Createtime	Timestamp of creation time
Modifytime	Timestamp of modified time
Attributes	
Owner	Owner name of file

**TABLE 13-66** MVCA\_FSRM\_REPORTRULEVW

Name	Description
Ruleid	ID of the FSRM rule created
Rule_name	Name of the rule created
Description	Description
Collection_type	Type of collection
Condition_sq	

**TABLE 13-67** MVCS\_HOSTMEMORYSTATSVW

Name	Description
Hostid	Hosti ID
CAPACITYTYPE	Type of capacity (Raw, Daily, Weekly, Monthly)
Timestamp	Collection time
Percentphysicalused	Percentage of physical memory used
Freephysicalmemory	Percentage of physical memory free
Percentvirtualused	Percentage of virtual memory used



**TABLE 13-67** MVCS\_HOSTMEMORYSTATSVW

Name	Description
Freevirtualmemory	Percentage of virtual memory free

**TABLE 13-68** MVCS\_HOSTCPUSTATSVW

Name	Description
Hostid	Hosti ID
CAPACITYTYPE	Type of capacity (Raw,Daily,Weekly,Monthly)
Processorid	ID of the processor
Name	Name of the processor
Timestamp	Collection time
Pctprocesstime	

**TABLE 13-69** MVCS\_EVACTRLSTATSVW

Name	Description
ID	
COLLECTIONTIME	
STATSTYPE	
DEVICETIME	
DURATION	
AVGREADLATENCY	
AVGREADSIZE	
AVGWRELATENCY	
AVGWRELATENCY	
CPUPERCENT	
DATAAXFERPERCENT	
DELTAREADIOS	
DELTAREADLATENCY	
DELTAWRITEIOS	
DELTAWRELATENCY	
PCTREADIOS	

**TABLE 13-69** MVCS\_EVACTRLSTATSVW

Name	Description
PCTWRITEIOS	
READDATARATE	
READRATE	
TOTALDATARATE	
TOTALIORATE	
WRITEDATARATE	
WRITERATE	

**TABLE 13-70** MVCS\_EVADISKSTATSVW

Name	Description
COLLECTIONTIME	
ID	
STATSTYPE	
DEVICETIME	
DURATION	
AVGDRIVELATENCY	
AVGQUEUEDEPTH	
AVGREADLATENCY	
AVGREADSIZE	
AVGWritelatency	
AVGWritesize	
DELTADrivelatency	
DELTAREADIOS	
DELTAREADLATENCY	
DELTATOTALIOS	
DELTAWRITEIOS	
DELTAWritelatency	
PCTREADIOS	
PCTWRITEIOS	
READDATARATE	
READRATE	
TOTALDATARATE	
TOTALIORATE	

**TABLE 13-70 MVCS\_EVADISKSTATSVW**

Name	Description
WRITEDATARATE	
WRITERATE	

**TABLE 13-71 MVCS\_EVAHOSTFCPORTSTATSVW**

Name	Description
ID	
COLLECTIONTIME	
STATSTYPE	
DEVICETIME	
DURATION	
AVGQUEUEDEPTH	
AVGREADLATENCY	
AVGWRELATENCY	
BADCRCERR	
DELTAREADIOS	
DELTAREADLATENCY	
DELTAWRITEIOS	
DELTAWRELATENCY	
DISCARDFRAMES	
LINKFAILURE	
LOSSOFSIGNAL	
LOSSOFSYNCH	
PCTREADIOS	
PCTWRITEIOS	
PROTOCOLERROR	
READDATARATE	
READRATE	
RECEIVEEOFA	
TOTALDATARATE	
TOTALIORATE	
WRITEDATARATE	
WRITERATE	

**TABLE 13-72** MVCS\_EVASPAGVOLUMESTATSVW

Name	Description
ID	
COLLECTIONTIME	
STATSTYPE	
DEVICETIME	
DURATION	
AVGREADHITLATENCY	
AVGREADMISSLATENCY	
AVGREADSIZE	
AVGWritelatency	
AVGWritesize	
DELTAREADHITIOS	
DELTAREADHITLATENCY	
DELTAREADMISSIOS	
DELTAREADMISSLATENCY	
DELTAWRITEIOS	
DELTAWRITELATENCY	
FLUSHDATARATE	
FLUSHRATE	
MIRRORDATARATE	
PCTREADIOS	
PCTWRITEIOS	
PREFETCHDATARATE	
READDATARATE	
READHITDATARATE	
READHITRATE	
READMISSDATARATE	
READMISSRATE	
READRATE	
TOTALDATARATE	
TOTALIORATE	
WRITEDATARATE	
WRITERATE	

**TABLE 13-73 MVCS\_EVASTORAGESYSTEMSTATSVW**

Name	Description
ID	
COLLECTIONTIME	
STATSTYPE	
DEVICETIME	
DURATION	
TOTALDATARATE	
TOTALIORATE	

**TABLE 13-74 MVCS\_EVAVOLUMEESTATSVW**

ID	
COLLECTIONTIME	
STATSTYPE	
DEVICETIME	
DURATION	
AVGREADHITLATENCY	
AVGREADMISSLATENCY	
AVGREADSIZE	
AVGWRELATENCY	
AVGWWRITEIZE	
DELTAREADHITIOS	
DELTAREADHITLATENCY	
DELTAREADMISSIOS	
DELTAREADMISSLATENCY	
DELTAWRITEIOS	
DELTAWRELATENCY	
FLUSHDATARATE	
FLUSHRATE	
MIRRORDATARATE	
PCTREADIOS	
PCTWRITEIOS	

**TABLE 13-74** MVCS\_EVAVOLUMESTATSVW

PREFETCHDATARATE	
READDATARATE	
READHITDATARATE	
READHITRATE	
READMISSDATARATE	
READMISSRATE	
READRATE	
TOTALDATARATE	
TOTALIORATE	
WRITEDATARATE	
WRITERATE	

## 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 lists the views from earlier releases and the corresponding new views. You must use the new views for any new report development. Verify that your existing reports will work correctly against these new views. Some of the views have changed and may not work in existing reports.

**TABLE 13-75** Views from Previous Releases

Legacy View	Alternate Core Views
MV_STORAGESYSTEMPORTUTILVW	MVC_STORAGESYSTEMCONFIGVW MVC_STORAGESYSTEMSUMMARYVW
MV_SSFRONTENDVW	MVC_LUNMAPPINGVW MVC_PORTSUMMARYVW MVC_STORAGEVOLUMEPORTS MVC_STORAGEVOLUMESUMMARYVW MVC_PORTSUMMARYVW MVC_STORGAEPOOLSUMMARYVW MVC_STORAGEPROCESSORSUMMARYVW MVC_STORAGESYSTEMSUMMARYVW

**TABLE 13-75** Views from Previous Releases (*Continued*)

Legacy View	Alternate Core Views
MV_SSBACKENDDETAILVW	MVC_CARDSUMMARYVW MVC_DISKEXTENTSUMMARYVW MVC_DISK_EXTENTVW MVC_DISKDRIVESUMMARYVW MVC_STORAGESYSTEMSUMMARYVW
MV_SSLOGICALDETAILVW	MVC_STORGAEPOOLSUMMARYVW MVC_STORAGEVOLUMESUMMARYVW MVC_VOLUMEDISKDRIVEVW MVC_DISKEXTENTSUMMARYVW
MV_SSAVAILABLEVOLUMEVW	MVC_STORAGESYSTEMSUMMARYVW MVC_STORAGEVOLUMESUMMARYVW MVC_STORGAEPOOLSUMMARYVW MVC_STORAGEVOLUMEPORTS MVC_STORAGEPROCESSORSUMMARYVW
MV_TEMPMAAPPEDVOLSUMMARYVW	MVC_STORAGEVOLUMESUMMARYVW MVC_STORAGESYSTEMSUMMARYVW MVC_STORAGEPROCESSORSUMMARYVW MVC_STORAGEVOLUMEPORTS MVC_PROTOCOLCONTROLLERVW MVC_LUNMAPPINGVW
MV_LUNSPERFASUMMARYVW	MVC_STORAGESYSTEMCONFIGVW MVC_STORAGESYSTEMSUMMARYVW MVC_STORAGEPROCESSORSUMMARYVW MVC_PORTSUMMARYVW MVC_PORTCONTROLLERMAPVW MVC_PROTOCOLCONTROLLERVW MVC_STORAGEVOLUMESUMMARYVW MVC_STORAGEVOLUMEPORTS MVC_LUNMAPPINGVW MVC_CARDSUMMARYVW, MV_HOSTSUM
MV_FABRICADAPCAPLUNVW	MVC_STORAGESYSTEMSUMMARYVW MVC_STORAGEPROCESSORSUMMARYVW MVC_PORTSUMMARYVW MVC_PORTCONTROLLERMAPVW MVC_PROTOCOLCONTROLLERVW MVC_STORAGEVOLUMEPORTS MVC_STORAGEVOLUMESUMMARYVW MVC_STORGAEPOOLSUMMARYVW
MV_TEMPSTORSYSTEMSUMMARYVW	MVC_STORAGESYSTEMSUMMARYVW MVC_STORAGEVOLUMESUMMARYVW MVC_STORAGESYSTEMCONFIGVW MVC_STORAGEPROCESSORSUMMARYVW MVC_STORAGEPOOLCONFIGVW MVC_STORGAEPOOLSUMMARYVW

**TABLE 13-75** Views from Previous Releases (*Continued*)

Legacy View	Alternate Core Views
MV_TEMPSTORPOOLSUMMARYVW	MVC_STORAGESYSTEMSUMMARYVW MVC_STORAGEPROCESSORSUMMARYVW MVC_STORAGEPOOLCONFIGVW MVC_STORGAEPOOLSUMMARYVW
MV_TEMPFRONTENDVW	MVC_STORAGESYSTEMSUMMARYVW MVC_STORAGEVOLUMESUMMARYVW MVC_STORAGEPROCESSORSUMMARYVW MVC_STORAGEPOOLCONFIGVW MVC_POOLSUMMARYVW
MV_LUNMAPPINGVW	MVC_LUNMAPPINGVW
MV_HOSTSUMMARYVW	MVC_HOSTSUMMARYVW MVC_HBASUMMARYVW MVC_CARDSUMMARYVW MVC_PORTSUMMARYVW
MV_TEMPHOSTLOGICALVW	MVC_HOSTDISKDRIVEVW MVC_SUBPATHVW MVC_PATHVW MVC_HOSTVOLUMESUMMARYVW MVC_HOSTSUMMARYVW MVC_HBASUMMARYVW MVC_CARDSUMMARYVW MVC_PORTSUMMARYVW
MV_TEMPHOSTCARDVW	MVC_PORTSUMMARYVW MVC_HBAPORTTARGETS MVC_PORTSUMMARYVW MVC_CARDSUMMARYVW MVC_PORTSUMMARYVW
MV_HOSTSTORAGESUMMARYVW	MVC_HOSTSUMMARYVW MVC_HOSTVOLUMESUMMARYVW MVC_HOSTCAPACITYVW MVC_OPTIONALTABLEVW
MV_HOSTSTORAGEBYOSVW	MVC_HOSTSUMMARYVW MVC_HOSTCAPACITYVW
MV_HOSTSTORAGEALLOCATIONVW	MVC_HOSTSUMMARYVW MVC_HOSTCAPACITYVW MVC_SUBPATHVW MVC_STORAGEVOLUMESUMMARYVW MVC_STORAGEVOLUMEPORTS MVC_STORGAEPOOLSUMMARYVW MVC_APPLICATIONSUMMARYVW MVC_PORTSUMMARYVW MVC_SWITCHSUMMARYVW



**TABLE 13-75** Views from Previous Releases (*Continued*)

Legacy View	Alternate Core Views
MV_HOSTCONNECTIVITYVW	MVC_PATHVW MVC_HOSTSUMMARYVW MVC_SUBPATHVW MVC_STORAGEVOLUMESUMMARYVW MVC_STORGAEPOOLSUMMARYVW MVC_CARDSUMMARYVW MVC_PORTSUMMARYVW MVC_STORAGESYSTEMSUMMARYVW
MV_HOSTVXVMVW	MVC_DISKEXTENTSUMMARYVW MVC_HOSTSUMMARYVW MVC_DISK_EXTENTVW MVC_DISKDRIVESUMMARYVW MVC_OPTIONALTABLEVW, MVC_PATHVW MVC_SUBPATHVW MVC_HOSTDISKDRIVEVW MVC_HOSTVOLUMESUMMARYVW
MV_HOSTVMVW	MVC_HOSTDISKDRIVEVW MVC_PATHVW MVC_HOSTVOLUMESUMMARYVW MVC_HOSTSUMMARYVW MVC_HOSTDISKDRIVEVW
MV_HOSTLOGICALVOLUMEVW	MVC_HOSTCAPACITYVW
MV_HOSTFSVOLUMEVW	MVX_HOSTSUMMARYVW MVC_HOSTCAPACITYV MVC_PATHVW MVC_SUBPATHVW MVC_HOSTVOLUMESUMMARYVW MVC_STORGAEPOOLISUMMARYVW MVC_STORAGEVOLUMESUMMARYVW
MV_HOSTRAWVOLUMEVW	MVC_PATHVW MVC_HOSTSUMMARYVW MVC_SUBPATHVW MVC_STORAGEVOLUMESUMMARYVW MVC_STORGAEPOOLSUMMARYVW MVC_DISKEXTENTSUMMARYVW
MV_HOSTUNUSEDVOLUMEVW	MVC_PATHVW MVC_HOSTSUMMARYVW MVC_SUBPATHVW MVC_STORAGEVOLUMESUMMARYVW MVC_STORGAEPOOLSUMMARYVW
MV_HOSTDISKDRIVEVW	MVC_HOSTSUMMARYVW MVC_DISKDRIVESUMMARYVW MVC_DISK_EXTENTVW MVC_DISKEXTENTSUMMARYVW

**TABLE 13-75** Views from Previous Releases (*Continued*)

Legacy View	Alternate Core Views
MV_HOSTSSDEPENDENCYVW	MVC_HOSTSUMMARYVW MVC_SUBPATHVW MVC_STORAGEVOLUMESUMMARYVW MVC_STORAGESYSTEMSUMMARYVW MVC_STORAGEPROCESSORSUMMARYVW MVC_PORTSUMMARYVW MVC_PATHVW MVC_HOSTVOLUMESUMMARYVW
MV_HOSTAPPDEPENDENCYVW	MVC_APPLICATIONSUMMARYVW MVC_HOSTVOLUMESUMMARYVW MVC_STORAGEVOLUMESUMMARYVW MVC_SUBPATHVW MVC_PATHVW MVC_PORTSUMMARYVW
MV_HOSTSPERFAVW	MVC_HOSTSUMMARYVW MVC_CARDSUMMARYVW MVC_PORTSUMMARYVW MVC_LUNMAPPINGVW MVC_STORAGEVOLUMESUMMARYVW MVC_PROTOCOLCONTROLLERVW MVC_PORTCONTROLLERMAPVW MVC_PORTSUMMARYVW MVC_STORAGEPROCESSORSUMMARYVW MVC_STORAGESYSTEMSUMMARYVW
MV_HOSTDISKPARTITIONVW	MVC_SUBPATHVW MVC_VOLUMEDISKDRIVEVW MVC_HOSTDISKDRIVEVW MVC_HOSTVOLUMESUMMARYVW
MV_TEMP SWITCHBYTESINTERVALVW	
MV_TEMP SWITCHCONNECTEDVW	MVC_PORTSUMMARYVW MVC_SWITCHSUMMARYVW
MV_TEMP CONNECTEDHOSTVW	MVC_PORTSUMMARYVW MVC_CARDSUMMARYVW MVC_HOSTSUMMARYVW MVC_SWITCHSUMMARYVW
MV_TEMP CONNECTEDSTORAGEVW	MVC_PORTSUMMARYVW MVC_STORAGEPROCESSORSUMMARYVW MVC_STORAGESYSTEMSUMMARYVW MVC_SWITCHSUMMARYVW MVC_HOSTSUMMARYVW
MV_TEMPZONEVW	MVC_ZONEVW MVC_ZONESUMMARY

**TABLE 13-75** Views from Previous Releases (*Continued*)

Legacy View	Alternate Core Views
MV_SWITCHDETAILVW	MVC_SWITCHSUMMARYVW MVC_ZONESUMMARY MVC_ZONEVW MVC_PORTSUMMARYVW MVC_CARDSUMMARYVW MVC_HOSTSUMMARYVW
MV_AVAILABLEPORTVW	MVC_SWITCHSUMMARYVW
MV_TOTALPORTSVW	MVC_PORTSUMMARYVW
MV_SANZONEPORTWWNVW	MVC_PORTSUMMARYVW MVC_SWITCHSUMMARYVW MVC_CARDSUMMARYVW MVC_PORTSUMMARYVW MVC_ZONEVW
MV_SANCOMNOTLOGINVW	MVC_PORTSUMMARYVW MVC_ZONESUMMARY MVC_ZONEVW MVC_SWITCHSUMMARYVW MVC_CARDSUMMARYVW
MV_ZONEDETAILSVW	MVC_SWITCHSUMMARYVW MVC_ZONESUMMARY MVC_ZONEVW
MV_EVENTVW	MVC_EVENTVW
MV_ORGANIZATIONVW	MVC_ORGANIZATIONVW
MV_ORGRELATIONVW	MVC_ORGRELATIONVW
MV_APPLICATIONVW	MVC_APPLICATIONSUMMARYVW
MV_DBAPPCHARGEBACKVW	MVC_APPLICATIONSUMMARYVW MVCA_DBAPPINSTCAPACITYVW MVCA_DBAPPPHYCAPACITYVW MVCA_EXCHAPPCAPACITYVW MVCA_VIRTUALAPPCAPACITYVW
MV_APPDEPENDENCYVW	MVC_APPLICATIONSUMMARYVW MVC_HOSTVOLUMESUMMARYVW MVC_SUBPATHVW MVC_PATHVW MVC_PORTSUMMARYVW
MV_ASSETSUMMARYVW	MVC_ASSETSUMMARYVW MVC_OPTIONALTABLEVW
MV_ASSETCOUNTVW	MVC_APPLICATIONSUMMARYVW MVC_HOSTSUMMARYVW MVC_STORAGESYSTEMSUMMARYVW MVC_SWITCHSUMMARYVW MVC_TAPELIBRARYSUMMARYVW

# 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.

## Event Management

---

This chapter contains the following topics:


- “About Event Manager” on page 605
- “Viewing Event Details” on page 611
- “Clearing Events” on page 612
- “Configuring the Clearing of Events” on page 613
- “Configuring the Deletion of Events” on page 614
- “Deleting Events” on page 615
- “Sorting Events” on page 615
- “Adding Journal Entries” on page 616
- “Changing the CLARiiON Event Polling Interval” on page 616
- “Brocade Events” on page 617
- “Filtering Events” on page 619

---

## 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 when a device connected to a Brocade switch has gone offline. Event Manager provides the following information about the events:

- **ID**- The identification number assigned to the event
- **Severity** - Identifies the severity level.
- **Time** - The time the event was recorded.
- **Element** - The source of the event. An element can be a switch, host, application, fabric or anything else on the network.
- **Summary Text** - A brief explanation of the event. When you click the text, the details of the event are displayed.
- **Event Type** - Specifies whether the source of this event is an application, a host, etc.
- **Count** - The total count of similar events.

- **Cleared** - Specifies whether an event is cleared. See “Clearing Events” on page 612.
- **Delete** - Click the  icon to remove an event.

Event Manager also provides several buttons at the top of its screen:

- **Delete Selected** - Deletes all selected events.
- **Delete All** - Deletes all events lists.
- **Clear Selected** - Marks the selected events as cleared.
- **Clear All** - Marks all events as cleared.
- **Un-clear Selected** - Removes the clear status from events that are selected.
- **Un-clear All** - Removes the clear status from all events.

### Event Manager Summary Metrics

The top of the Event Manager pane provides a subtotal of the severity rating of the events and the number of events from the different element types. If you do not see the Event Manager summary metrics, you may need to expand the Summary node located at the top of the page, as shown in the following figure.

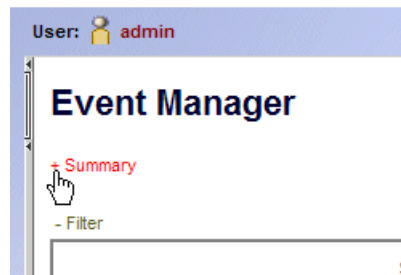


FIGURE 14-1 Accessing summary information

The Event Manager summary metrics are displayed after the Summary node is expanded.

You may need to click the **Refresh** button occasionally for the latest Event Manager summary metrics.


The definition for each severity level varies according to the type of element; however, there are some generic definitions:

**TABLE 14-1** Severity Levels

Severity Level	Description
Unknown	The event does not fall into the other categories; further information can not be obtained from it.
Informational	Provides informational data. For example, for a Brocade switch, 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.
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 offline.

## Accessing Event Manager

To access Event Manager, do one of the following:

- Click **Event Manager** () in the left pane.
- To view events from a specific element, take one of the following actions:
  - 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.

---

**Note** – You cannot access the **Element Type** filter using this method.

---








- Right-click the element in Capacity Explorer, Performance Explorer or System Explorer and then select the **Show Events** option in the menu.
- Select a discovered host in Protection Explorer and then click **Events** in the lower-right corner.

To change your user preferences for Event Manager, by selecting **Configuration > Events**. See “Controlling the Display of Cleared and Deleted Events” on page 241.

# Event Manager Icons

The following icons are displayed in Event Manager.

**TABLE 14-2** Icons in Event Manager

Icon	Description
	Event is marked cleared. See the topic, “Clearing Events” on page 612 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.

## Events Supported

Event Manager does not support events from all discovered elements. See the following table for more information about which elements Event Manager supports.

**TABLE 14-3** Supported Hardware for Events

Hardware	Events Supported?	Additional Information
Brocade switches	Y	Not all events that show up in the webtool appear in Event Manager.
Cisco switches SNMP	Y	Need to configure the switch or proxy to send traps to the management server.
CNT switches SMI	N	
McDATA switches SWAPI to EFCM	Y	Not all events that show up in EFCM appear in Event Manager. Also applies to EMC Connectrix switches.
McDATA switches SNMP through proxy	Y	Need to configure the switch or proxy to send traps to the management server. Also applies to EMC Connectrix switches.
McDATA switches SNMP to switches	Y	Need to configure the switch or proxy to send traps to the management server. Also applies to EMC Connectrix switches.



**TABLE 14-3** Supported Hardware for Events (*Continued*)

Hardware	Events Supported?	Additional Information
QLogic SNMP switches	Y	Need to configure the switch or proxy to send traps to the management server. Also applies to Sun StorEdge switches.
CLARiiON storage systems	Y	
LSI and Sun 6130 storage systems	Y	
HDS storage systems	Y	
HP EVA	Y	To receive events from HP EVA storage systems, you must add the management server address as an SNMP trap host in Command View. See "Obtaining SNMP Traps using Command View EVA" on page 71.
HP XP with Command View XP	N	
HP XP with Command View XP AE	Y	
HP XP with XP Provider	Y	
HP MSA storage systems	N	
IBM ESS storage systems	N	
Sun 35xx storage systems	N	
Sun 6920 and 6940 storage systems	N	
Symmetrix storage systems	Y	
Xiotech storage systems	N	
HP NAS Filers	N	
NetApp Filers	Y	To receive events from NetApp filers, you must add the management server address as an SNMP trap host on the NetApp filer.
Sun NAS Devices	N	
Tape Libraries	N	

### Viewing Events from the Management Server

By default the management server displays events from all of the elements, regardless of the user's organization. However, it does not display its own events. To view events from the management server:

- Select the **All** or **<Product Name of the Management Server>** option from the Show Element Type menu. See "" on page 629 for more information about the menu options.

- Click the **Customize** button next to the **Show Element Type** 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**, and then click **Apply Filters**.

### **Avoiding Excessive Notification**

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 kept 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.

### **Issues with Sun StorEdge or QLogic Switches**

- 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.

### **Issues with NetApp Filers**

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.

### **Issues with CNT InVsn Enterprise Manager**

Event Manager displays events it receives from CNT InVsn Enterprise Manager. As of version 9.5, 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 that Event Manager not display events from additional firmware versions.
- Events from Brocade Fabric Watch are not supported when the Brocade switch has been discovered through the SMI-S provider.

### **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 offline/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 44 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, enable SNMP trap forwarding only 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, and 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.

### **Issue with CLARiiON Storage Systems**

When you manage a CLARiiON storage system, extraneous events appear in Event Manager such as `CRU Bound`, `CRU Enabled`, and various `Success` messages. These do not indicate problems and can safely be ignored.

---

## Viewing Event Details

To access event details:

1. Access Event Manager as described in “Accessing Event Manager” on page 607.
2. In Event Manager click the event summary, as shown in the following figure.

Delete Selected	Delete All	Clear Selected	Clear All	Un-Clear Selected	Un
-----------------	------------	----------------	-----------	-------------------	----

Page  of 42 Showing Element(s) 1-25 out of 1050 Total (2 Selected)

<input type="checkbox"/>	ID	Severity	Time	Element	Summary Text	Event Type
<input type="checkbox"/>	10421		2007-07-25 06:08		User admin logged in	applq_event
<input type="checkbox"/>	10420		2007-07-25 05:49		2007-07-25 08:48: A single device came online	cim_alert
<input type="checkbox"/>	10419		2007-07-25 05:49		2007-07-25 08:48: A single device came online	cim_alert
<input type="checkbox"/>	10418		2007-07-25 05:49		2007-07-25 08:48: A single device came online	cim_alert

**FIGURE 14-2** Accessing Event Details

The event details are displayed.

The Event Details pane provides information on one or more of the following items:

- **Serial Number** - The number assigned to the event.
- **Status** - Indicates whether the event has been cleared.
- **Component** - If the event came from a component of an element, the component is listed.
- **Element** - The source of the event. An element can be a switch, host, application, fabric, or anything else on the network. If this box is blank, the event did not come from an element.
- **Severity** - The severity level, which can be one of the following:
  - **Clear**
  - **Unknown**
  - **Informational**
  - **Warning**
  - **Minor**
  - **Major**
  - **Critical**
- **Time Reported** - The time and date the event was reported to the management server.
- **Type** - A brief label of the event.
- **Summary Text** - An explanation of the event.
- **Probable Cause** - An explanation of a probable cause.
- **Recommended Actions** - Provides recommendations.
- **Journal Entries** - Use this box to enter additional information and then click **Add Journal Entry**. This box is limited to 4,000 characters. See “Adding Journal Entries” on page 616.

---

**Note** – Events listed in Event Manager may not be attributed to the correct source until Get Details has completed.

---

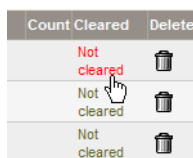
---

## Clearing Events

If you have already reviewed an event, you might want to mark it as “cleared” so you can keep track of which events you have already reviewed.

To clear events, do one of the following:

- To clear an event - Click the **Not cleared** text for the event. Notice that the text turns red, as shown in the following figure:





Count Cleared	Delete
Not cleared	
Not cleared	
Not cleared	

FIGURE 14-3 Clearing an event

- To clear several events - Select the events and then click **Clear Selected**.
- To clear all events - Click **Clear All**.

When events are cleared, a clear icon (○) appears in the Cleared column for the event.

---

## 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 mark 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 14-4** 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 before clearing an event, do the following:

1. Select **Configuration > Events** to access the Events page.
2. Do one of the following:
  - If you never want an event of a specified severity level marked as cleared, select the **Never** option next to the severity level in the Automatic Clear Delay column.
  - 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 enter the number in the adjacent box in the Automatic Clear Delay column:

**Seconds**

**Minutes**

**Hours**

**Days**

**Weeks**

3. Click **Save Changes**.

---

## Configuring the Deletion of Events

The management server automatically deletes events after two weeks by default. For each severity level you can specify 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. Select **Configuration > Events** to access the Events page.
  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 for deleting an event, select one of the following units of measurement from the combo box and enter the number in the adjacent box:
      - Seconds**
      - Minutes**
      - Hours**
      - Days**
      - Weeks**
- For example, if you want events that are a week old deleted, select **Weeks** in the combo box in the Automatic Delete Delay column, and then enter 1.
3. Click **Save Changes**.

---

## Deleting Events

To delete an event, do one of the following:

- To delete one element at a time - Click the trash can icon next to the event.
- To delete specific elements - Select the events you want to delete, and then click the **Delete Selected** button at the top of the screen.
- To delete all elements - Click the **Delete All** button at the top of the screen.

---

## 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 605.
2. In the Event Manager table, click the column title corresponding to the attribute you want to sort on:

- **ID** - The identification number assigned to the event
  - **Severity** - Shows the severity level.
  - **Time** - Displays the time and date the management server was aware of the event. The time and date are displayed in the following format: YYYY-DD-MM HH:MM
  - **Element** - The source of the event. An element can be a switch, host, application, fabric, or anything else on the network.
  - **Summary Text** - A brief explanation of the event. When you click the text, the details of the event are displayed.
  - **Event Type** - Specifies whether the source of this event is an application, a host, etc.
  - **Count** - The total count of similar events.
  - **Cleared** - Indicates whether an event has been cleared.
- 

## 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 605.
2. To access the Event Details page for an event, click the text for the event in the **Summary Text** column.
3. In the Journal Entries box, type the entry for the event. This box is limited to 4,000 characters.
4. Click **Add Journal Entry**.

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 too many “information” messages from the CLARiiON storage system.

---

**Caution** – 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, if 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 such a long time.

---

To change the polling interval:

1. Select **Configuration > Product Health**, and 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, right-click the selected text, and then select **Copy**.
4. Repeat step 1 to return to the Advanced page.
5. Paste the copied text into the Custom Properties box. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the box and select **Paste**.
6. Make your changes in the Custom Properties box. Remove the hash (#) symbol in front of the property to make sure the property is not commented out.
7. Change the value assigned to the `cimom.ClariionEventPollInterval` property. Note that 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**.

**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 contains the following topics:

- “Brocade Switch Events” on page 617
- “Supported Brocade Events” on page 618

### Brocade Switch Events

When a Brocade switch generates an event, it assigns a code instead of an event severity level to the event. 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 14-5** 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.
3	Critical	EV_CONNECTED_OBJE CT_OFFLINE	A device connected to a switch has gone offline.
4	Major	EV_CONNECTED_OBJE CT_ONLINE	A device connected to a switch has come online.

**TABLE 14-5** Brocade Switch Events (*Continued*)

*Code	Event Severity Level	*Name	Description
5	Info	EV_RSCN	For those RSCN events not covered by EV_OBJ_XXX codes listed above. Examples: 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 term does not appear in the event description, but is provided for clarity.

## 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 Access library events:** Events about proxy switch health.
- **Zoning events:** Events about zoning, such as zone-related state change notification.

---

## Filtering Events

This section contains the following topics:

- “Setting up a filter” on page 619
- “Selecting a custom time period” on page 622
- “Resetting a filter” on page 624
- “Setting up advanced filtering” on page 625
- “Clearing Advance Filtering Options” on page 629

# Setting up a filter

The management server provides several types of event filters which lets you 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:

- **Time period**
- **Severity level**
- **Element type**
- **Summary text**
- **Element name**
- Cleared status
- Specific element

To set up a filter:

1. To access the filter for Event Manager, click the Filter heading at the top of the page, as shown in the following figure:




FIGURE 14-4 Accessing the filter feature

The filtering feature is displayed, as shown in the following figure:




FIGURE 14-5 The Filter feature in Event Manager

2. Select a time from the **Time Period** combo box. You can also select a customized time as described in "Selecting a custom time period" on page 622.




FIGURE 14-6 Selecting a time period

3. Select which events of a severity type you want displayed. Use the Control and Shift keys to select multiple severities. The following options are provided:
  - **All** - Events of all severities are displayed.
  - **Unknown** - Only events of the severity type, unknown, are displayed.
  - **Informational** - Only events of the severity type, informational, are displayed.
  - **Warning** - Only events of the severity type, warning, are displayed.
  - **Minor** - Only events of the severity type, minor, are displayed.
  - **Major** - Only events of the severity type, major, are displayed.
  - **Critical** - Only events of the severity type, critical, are displayed.
4. Select which events of an element type you want displayed. Use the Control and Shift keys to select multiple severities. The following options are provided:
  - All
  - Application
  - Host
  - Switch
  - Storage System
  - Tape Library
  - Fabric
  - Other
  - <product name>
5. To set the filter by summary text, enter the text you want to be used for the filter in the **Summary Text Contains** field.
 

You may want to use this option when you are seeing several events that span over several elements or severity levels.

For example, let's assume you are wondering if someone else is logging into the management server as admin. You can find how often the admin user has logged into the management server over the last few days by entering the following text in the **Summary Text Contains** field:

User admin logged in
6. To set the filter by element name enter text in the **Element Name Contains** field.

This feature can be helpful if you are interested in events from elements that have similar names.

For example, let's assume you have a naming convention for hosts, where all hosts that belong to the engineering group begin with engineering. Event Manager could display only events from those hosts by entering engineering in the **Element Name Contains** field.

7. To set a filter by cleared status, select one or more of the following by using the Control and Shift keys.
  - **All** - All events regardless of their clear status
  - **All But Clear** - All events except for those marked clear
  - **Clear** - Only events marked clear
8. If you want events to be merged together, select the option, **Collapse all events with same severity, type and element**.

This feature is very useful when numerous events of the same severity, type and element are being picked up by the management server. If you select the option, **Collapse all events with same severity, type and element**, the management server displays just unique events with a total count in front of them, rather than list all of the events individually.
9. If you want Event Manager to filter elements automatically at a set interval, select the **Automatically filter every** option and select one of the following:
  - 30 seconds
  - 1 minute
  - 5 minutes
  - 10 minutes
10. When you are done setting your options, click the **Filter** button.

---

**Note** – You can update the events displayed by clicking the **Filter** button.

---

## Selecting a custom time period


---

**Note** – You may want to select a custom time period for the filter when troubleshooting an issue.

---


To select a custom time period:

1. Select the **Custom Time Period** option.



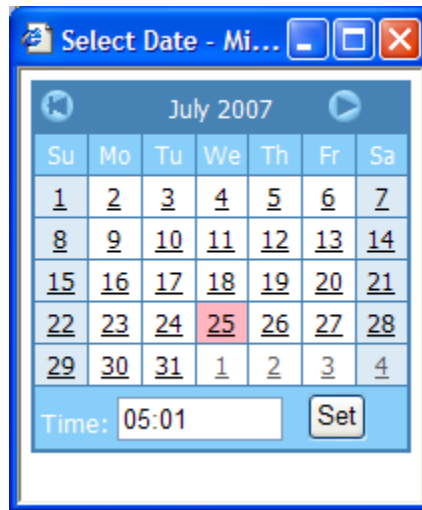
**FIGURE 14-7** Custom Time Period

2. Event Manager only displays events that occur after the start date and time and before the end date and time:
  - a. Click the Calendar icon to the right of the **Start Date** field.



**FIGURE 14-8** Selecting a start date for filtering

- b. In the Calendar, select the start date.



**FIGURE 14-9** Selecting a date

- c. In the **Time** field, enter the start time.

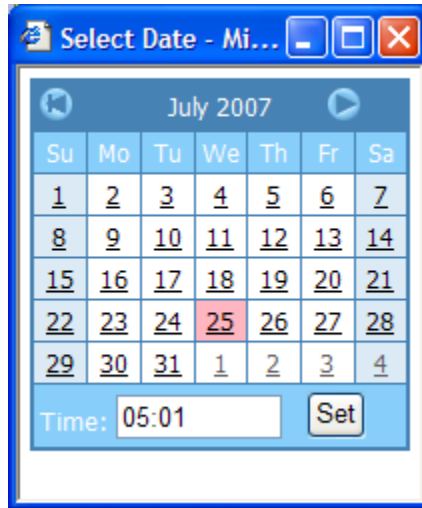
The time is based on a 24-hour clock. For example, if you want Event Manager to display events occurring after 8 p.m. on the specified date, you would enter 20:00.

- d. When you are done setting the start time and date, click **Set**.

3. To set the end date and time for the filter. Event Manager only displays events that occur after the start date and time and before the end date and time:

- a. Click the Calendar icon to the right of the **End Date** field.
- b. In the Calendar, select the end date.





**FIGURE 14-10** Selecting a date

- c. In the **Time** field, enter the end time.

The time is based on a 24-hour clock. For example, if you want Event Manager to display events occurring before 8 p.m. on the specified date, you would enter 20:00.

- d. When you are done setting the end time and date, click **Set**.

## Resetting a filter

To revert to the default settings for the filter:

1. To access the filter for Event Manager, click the Filter heading at the top of the page, as shown in the following figure:

## Event Manager

+ Summary

+ Filter



Delete Selected

Delete All

**FIGURE 14-11** Accessing the filter feature

The filtering feature is displayed, as shown in the following figure:

The screenshot shows the 'Filter' panel in the Event Manager interface. It includes a 'Custom Time Period' checkbox, a 'Time Period' dropdown set to 'Last one day', and a 'Severities' list with options: All, Unknown, Informational, Warning, Minor, and Major. The 'Element Types' list includes: All, Application, Host, Switch, Storage System, and Tape Library. There are input fields for 'Summary Text Contains' and 'Element Name Contains'. A 'Cleared' dropdown is set to 'All'. A checkbox for 'Collapse all events with same severity, type and element' is present. At the bottom right, there are 'Filter' and 'Reset' buttons, and an 'Automatically filter every:' section with a checkbox and a '30 seconds' dropdown. A '+ Advanced Options' link is at the bottom left.

**FIGURE 14-12** The Filter feature in Event Manager

2. Click the **Reset** button.

The filter is reset.

## Setting up advanced filtering

Event Manager has an advanced filtering feature that lets you provide detailed information for filtering.

To set up advanced filtering:

1. To access the filter for Event Manager, click the **Filter** heading at the top of the page, as shown in the following figure:

## Event Manager

+ Summary

+ Filter



Delete Selected

Delete All

**FIGURE 14-13** Accessing the filter feature

The filtering feature is displayed, as shown in the following figure:

The screenshot shows the 'Filter' dialog box in the Event Manager. It includes several sections: 'Custom Time Period' with a checkbox and a 'Time Period' dropdown set to 'Last one day'; 'Severities' and 'Element Types' dropdown menus, both currently set to 'All'; 'Summary Text Contains' and 'Element Name Contains' text input fields; a 'Cleared' dropdown set to 'All'; a checkbox for 'Collapse all events with same severity, type and element'; a checkbox for 'Automatically filter every:' with a '30 seconds' dropdown; and 'Filter' and 'Reset' buttons. A '+ Advanced Options' link is visible at the bottom left.

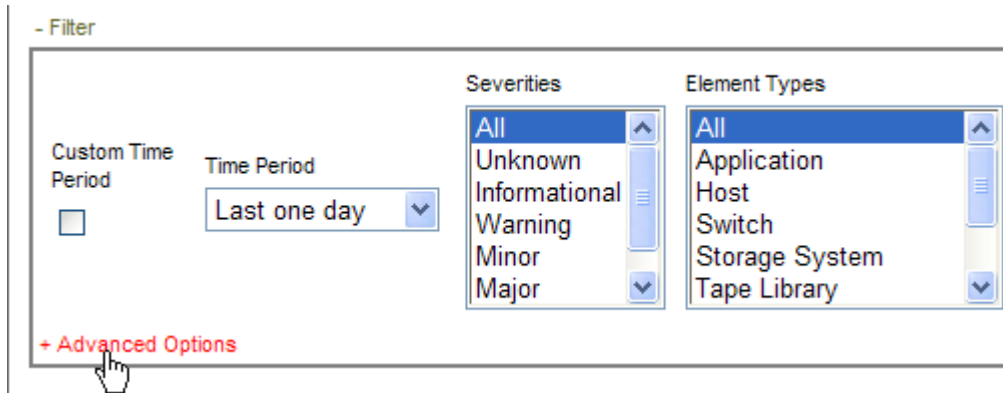
**FIGURE 14-14** The Filter feature in Event Manager

2. Expand the **Advanced Options** heading, as shown in the following figure:

---


**Caution** – Once the **Advanced Options** heading has been expanded, the **Element Name Contains** field becomes inactive. Any values entered in the **Element Name Contains** field are ignored as long as the **Advanced Options** heading is expanded.

---



**FIGURE 14-15** Accessing advanced options for filtering

3. Click the **Add** button as shown in the following figure:



**FIGURE 14-16** Clicking the Add button

4. Select the element type you want to add. Then, click **Next**.
5. Select one or more elements you want to add to the filter. Then, click **Next**.

Only the elements of the type you specified in the previous window are displayed.

For example, if you selected hosts in the previous window, only hosts are displayed in this screen.

If the element you are looking for does not appear on the first page of the table, use the navigation tools at the top of the table to page through the list of elements.

The hosts you select are listed in the navigation filter.

For example, in the following figure Host\_13081 and Host\_10380 were selected for advanced filtering, so they are listed under the Advanced Filtering heading, as shown in the following figure:

## Event Manager

+ Summary

- Filter

Time Period  

Last one day

Severities  

All  
Informational  
Minor  
Major  
Warning  
Critical

Element Types  

All  
Host  
Switch  
Storage System  
Fabric  
Application

Summary Text Contains

Element Name Contains

- Advanced Options  

NameDelete

Host\_13081 (discovered)  
Host\_10380 (discovered)

Add

Clear Data

Delete Selected

Delete All

Clear Selected

Clear All

FIGURE 14-17 Listing of elements

- Verify that **Advanced Options** heading is expanded as shown in the following figure. **Advanced Options** must be expanded for advanced filtering to work.




FIGURE 14-18 Advance filtering options

7. Click the **Filter** button.

Event Manager displays the elements specified under **Advanced Options**.

## Clearing Advance Filtering Options

You can clear the filtering set for advanced options, by clicking the **Clear Data** button under the Advanced Options heading.

## Viewing Performance Data

---

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 contains the following topics:

- “About Performance Explorer and Array Performance Pack” on page 631
- “General Considerations for Performance Explorer” on page 639
- “Accessing Performance Explorer” on page 640
- “Creating Performance Charts” on page 640
- “The Toolbars in Performance Explorer” on page 641
- “Comparing the Performance of Different Elements” on page 644
- “Viewing Summary Charts” on page 645
- “Viewing Trending Information for Performance” on page 645
- “Removing Performance Data from a Graph” on page 646
- “Setting a Custom Period” on page 647
- “Monitoring Options” on page 649
- “Monitoring with Direct Attached Storage” on page 658
- “Supported Host Configurations for Monitoring” on page 659
- “Sudden Dips Displayed in Certain Charts in Performance Explorer” on page 661
- “Values Continue to Increase in Charts for Aggregated Drives and Aggregate Volumes” on page 662

---

# About Performance Explorer and Array Performance Pack

Performance Explorer provides a graphical representation of the performance history of a managed element, such as bytes transmitted per second for a switch. From this performance information, you can also generate charts and customize reports.

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.

The licensed optional Array Performance Pack provides additional HP EVA array performance monitoring and reporting capability for Performance Explorer. The Array Performance Pack enhances both the Performance Explorer and the Reporter capabilities.

## Array Performance Pack Requirements

The following paragraphs describe important requirements and considerations for licensing and using the Array Performance Pack:

- Licensing Requirements and Setup
- Software Requirements
- EVAPerf Requirements
- Specifying Data Collectors
- EVA Array Discovery

### Licensing Requirements and Setup

Your Array Performance Pack license determines how many EVA arrays you can select for performance monitoring. Additional licenses can be purchased if you need to monitor more EVA arrays than your current license allows.

As part of the license setup, you must specify which EVA arrays are to be monitored. You can make changes to your selection at a later time to accommodate new monitoring and reporting needs. Array Performance Pack license setup is discussed in the chapter, *Managing Licenses*.

Setup for the Array Performance Pack also requires you to specify the data collectors used to implement your specific reporting needs. This is discussed in the section, *Specifying Data Collectors*, later in this chapter.



## Software Requirements

Refer to the Support Matrix for detailed supported software requirements for the EVA Array Performance pack. In summary, they include:

**Table 1: Performance Pack Configurations**

	CV EVA 6.0.2	CV EVA 7.0	CV EVA 7.0.1
VCS 3.028	X	X	X
VCS 3.1xx	X	X	X
VCS 4.0xx	X	X	X
XCS 5.1xx	X	Unsupported	X
XCS 6.0xx	X	X	X
XCS 6.1xx	X	X	X

- Management server software Build 6.0.2 or later.
- Command View EVA version 7.x highly recommended (minimum Build: 6.0.2). Command View EVA v7.0 is not supported with EVA XCS code version v5.1xx.
- A maximum of 8 EVAs per CV EVA server are supported.
- VCS 4.004 is not supported for EVA 5000 disk arrays.
- Management server performance provider does not support EVA VCS code earlier than 4.x.

Command View version v7.01, or later, is highly recommended to take best advantage of the enhancements. The required Command View EVA versions support the new built-in provider, which replaces the SMI-S provider used previously for EVA array discovery.

For Command View versions prior to v7.0.1, there is a maximum of 8 EVA arrays supported per Command View EVA server.

The maximum number of vdisks supported per EVA array is 512.

---

**Note** – You must discover your EVA arrays for the Array Performance Pack to work.

---

## EVAPerf Data Collector Requirements

The management server requires the EVAPerf Data Collector (EVA Perf) service to be running on each Command View EVA server being monitored. This service is installed by default with the Command View EVA suite, but sets the service to

“enabled” and “manual startup”. It does not automatically start the service. Therefore, you may want to consider changing the service startup parameters to “automatic” and “restart on failure”.

Although EVAPerf version 6.0.2 or newer is supported, it is strongly recommended that you run EVAPerf version 7.0.1, or later, to take advantage of fixes for service crashes and instability.

RPC over Port 860 must be “enabled” in any firewalls between the management server and the Command View EVA server(s).

## Specifying Data Collectors

After applying Array Performance Pack licenses to the target EVAs, you must select which corresponding EVA-specific “performance data collectors” to enable in the Data Collectors for Performance page. Each EVA Storage System-specific row in the list of performance data collectors represents a distinct set of performance metrics for the corresponding EVA array. When you start a collector, you will be prompted for a sample interval and start time for the collection of the corresponding metrics. This can be edited after a collector is started by selecting the icon for the collector. Be aware the type, quantity, and frequency of data collectors started can affect system performance. For example, EVA physical disk metrics collection is I/O intensive, so should be avoided unless needed.

To choose collectors, click **Configuration** at the management software home screen, then click **Performance**. When the Data Collection screen displays, click the **Data Collection** tab.

A screen similar to the following displays.

Data Collection

Data Aging

### Performance Data Collectors

You can start, stop and edit data collectors for elements listed in the table below. If a data collector is stopped, performance charts will not have up-to-date data. In some instances, performance charts may display no data. The default changes to collectors are application wide.

[Default Collector Settings](#) for newly-discovered devices

+ Filter

Edit Selected...

Start Selected

Stop Selected

Page 1 of 8

Showing 31-40 out of 76 Total (1 Selected)

Display 10 rows

Select All Pages

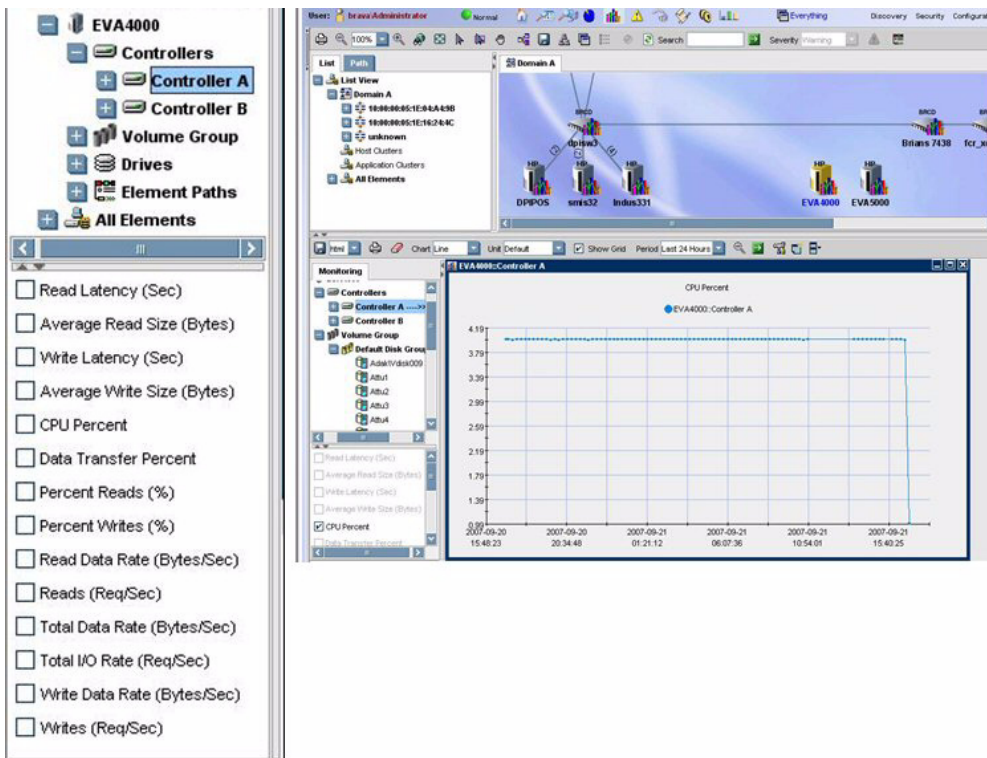
Unselect All Pages

Element	Element Type	Collector Type	Enabled	Interval (Minutes)	Next Scheduled Run	Edit	Action
<input type="checkbox"/> EVA5000	Storage System	HPEVA Controller Observer + Statistics	<input checked="" type="checkbox"/>	30	2007-11-05 17:40		<div>Stop</div>
<input type="checkbox"/> EVA5000	Storage System	HPEVA Storage System Aggregated Volume Observer + Statistics	<input checked="" type="checkbox"/>	30	2007-11-05 17:42		<div>Stop</div>
<input type="checkbox"/> EVA5000	Storage System	HPEVA Physical Disk Observer + Statistics	<input checked="" type="checkbox"/>	30	2007-11-05 17:40		<div>Stop</div>
<input type="checkbox"/> EVA5000	Storage System	HPEVA Storage System Observer + Statistics	<input checked="" type="checkbox"/>	30	2007-11-05 17:42		<div>Stop</div>
<input type="checkbox"/> EVA5000	Storage System	HPEVA Host Fibre Channel Port Observer + Statistics	<input checked="" type="checkbox"/>	30	2007-11-05 17:41		<div>Stop</div>
<input type="checkbox"/> EVA5000	Storage System	HPEVA Volume Observer + Statistics	<input checked="" type="checkbox"/>	30	2007-11-05 17:41		<div>Stop</div>
<input type="checkbox"/> EVA5000	Storage System	HPEVA Storage Pool Aggregated Volume Observer + Statistics	<input checked="" type="checkbox"/>	30	2007-11-05 17:42		<div>Stop</div>
<input type="checkbox"/> EVA9K_1	Storage System	HPEVA Controller Observer + Statistics	<input checked="" type="checkbox"/>	30	2007-11-05 17:43		<div>Stop</div>
<input type="checkbox"/> EVA9K_1	Storage System	HPEVA Volume Observer + Statistics	<input checked="" type="checkbox"/>	30	2007-11-05 17:45		<div>Stop</div>
<input type="checkbox"/> EVA9K_1	Storage System	HPEVA Physical Disk Observer + Statistics	<input checked="" type="checkbox"/>	30	2007-11-05 17:45		<div>Stop</div>

FIGURE 15-1 Data Collector Selection

Select the desired collectors from those listed in your display; then, click **Start Selected** for multiple collectors or click **Action** for single collectors. You might prefer to start and run the collectors only for specific needs, rather than running all of them all the time; otherwise it may affect overall system response. To edit the start time and interval for multiple collectors, click **Spread Start Time** to stagger the starting times for multiple collectors. This minimizes the impact on system performance.

When licensed appropriately, you can review the collected data in Performance Explorer. You can expand the device sub-elements (controllers, volume group, etc.) in the navigation tree, as shown in the following representative screen.



**FIGURE 15-2** Expanded View of Available Metrics in Performance Manager

The list of sub-elements and metrics vary depending upon the device type. The selections determine the performance data retrieved and available for analysis.

For example, the following representative screen display shows information about data rates associated with the highlighted EVA array host port.

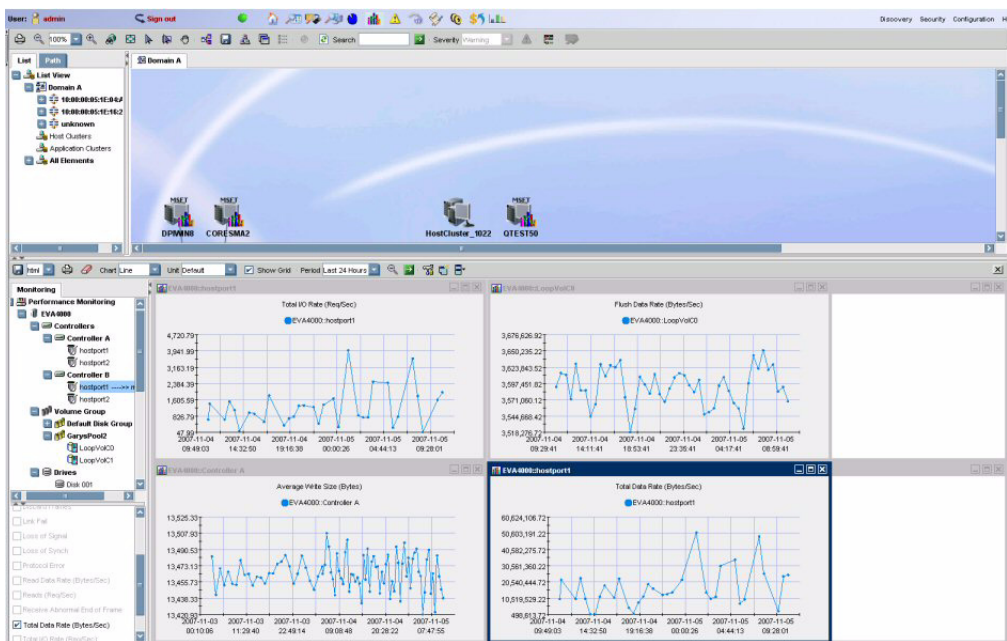


FIGURE 15-3 EVA Array Host Port Data Rates

## EVA Array Discovery

The following paragraphs describe important considerations in the EVA array discovery process to support the EVA Array Performance Pack.

### *Discovery Considerations Using Command View EVA*

If you performed EVA array discovery previously using the SMI-S provider (as with earlier Command View EVA versions), Array Performance Pack requires discovery (or rediscovery) of EVA Arrays using the supported Command View EVA version, which utilizes the built-in provider. (By default, Management Server Build 6.x discovers EVAs using the built-in provider.)

You should have run Get Details on an EVA before the Performance Monitoring tree will be populated with the corresponding EVA Elements used for corresponding metric selection.

Any configuration changes to the EVA array (i.e., addition or deletion of mapped volumes) requires a Get Details in order for the changes to be reflected in the Performance Monitoring tree, and in the actual performance data collected. Also, changes in the number of elements (i.e., disks or volumes) can impact the aggregated performance metric data.

If you make changes to your selection of EVA arrays later, you must ensure the selected new EVA arrays have been discovered using the required Command View EVA version. For more information regarding discovery procedures, refer to the management server installation guide.

Command View EVA supports a standby configuration whereby an array can be managed by multiple Command View EVA stations simultaneously, with only one of them actively managing a specific array at a time. Discovery will include both the EVA array being actively managed and the array not being actively managed (Standby Command View station). In such cases, there are limitations to the information that can be gathered, based upon whether both the managed and the unmanaged (standby) arrays are discovered by the management server.

If only the unmanaged array is discovered, only the top level array information can be collected.

If both the managed and the unmanaged (standby) arrays are discovered using the required version of Command View EVA with the new built-in provider, both the managed and unmanaged arrays must remain in the same discovery group. They can be moved to a different discovery group, but both must remain together in the same discovery group.

## EVA Metrics and Considerations

Be aware of the following considerations relating to the collecting metrics and general usage considerations:

- EVA performance data cannot be collected for EVA arrays that have controller passwords set via the Operator Control Panel (OCP).
- Only EVA volumes that are mapped to a host have associated performance data.
- The management server will not be able to gather host port metrics if the EVA controller has two switches attached to it with the same DomainID and the same switch port number is used in each switch.

## EVAPerf Considerations

Be aware of these EVAPerf Data Collector service considerations:

- EVAPerf Data Collector service can sometimes crash, especially in larger environments. This is less likely with EVAPerf 7.0.1. Restarting the service will enable performance data collection to succeed. Consider changing the service startup parameters to “automatic” and “restart on failure”.
- If the total number of vdisks configured on the EVA exceeds 512 using EVAPerf versions prior to v7.0.1, some mapped vdisks may not have any metrics associated with them.
- Occasionally, a physical disk might not have reporting because the EVAPerf Data Collector returns an unknown identifier for it.
- Vdisk and physical disk metric data viewed in the EVAPerf CLI on the Command View EVA server might appear inconsistent with the management server’s data. This is due to an EVAPerf CLI situation in which only point-in-time samples are returned for sample intervals greater than 30 seconds (this is for all VCS versions).
- Socket connections between the management server and the EVAPerf Data Collector service on port 860 are not immediately closed after stopping the EVA performance collectors in the management server, or when un-licensing EVA arrays for performance data collection in the management server. However, the connection will eventually time out and be closed after the default idle time for the connection expires. In these situations, although the connections are still open, no I/O occurs over the connection.
- The minimum collection interval that can be set for the EVA performance data collectors is 1 minute. The collection interval for real-time metrics is 20 seconds and cannot be changed. Frequent collection of large amounts of performance data, combined with regular I/O, may impact Command View EVA responsiveness. However, this does not impact the EVA’s data I/O. The repeat collection interval is two hours and is not configurable.
- Only Vdisks that mapped to a host have performance metrics associated with them. Only grouped physical disks have performance metrics associated with them.
- The default idle time-out is 2 hours and is not configurable.

---

## General Considerations for Performance Explorer

Keep in mind the following about Performance Explorer:

- If you see the message “There is not enough data to produce a chart [chart\_title] at this time,” lessen the frequency option or select **All** in the **Edit Chart Property** dialog box.
- Verify the performance collector for that element is enabled (**Configuration > Performance**). See “Managing Performance Collection” on page 270 for more information about enabling performance collectors.

- 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 “The Toolbars in Performance Explorer” on page 641.
- The aggregated volume and aggregated drive filters are no longer shown for LSI 5884 storage systems or any of such a system’s controllers. These filters were available in builds earlier than 4.0 of the management server.
- If one or more scheduled data points on a graph seem to be missing, it may be that an error occurred collecting data at that time. Check the CIMOM log for errors for the time frame covering sample interval of the missing data point(s).


---

**Caution** – 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. Charts that were active in Performance Explorer when Get Details was started are not restarted.

---

---

## Accessing Performance Explorer

To access Performance Explorer, click Performance Explorer (  ).


---

## Creating Performance Charts

To create a performance chart for an element:

1. Access Performance Explorer as described in “Accessing Performance Explorer” on page 640.
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 component, such as a port on a switch.
4. In the pane under the tree, select a monitoring option.  
See “Monitoring Options” on page 649 for more information.






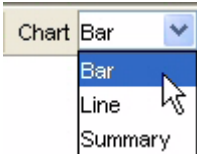

5. Use the Chart and Unit combo box to modify the chart. When you are done with your selections, click the  button in the lower pane. To learn more about these features, see “The Toolbars in Performance Explorer” on page 641
6. To monitor more than one element in a chart, see “Comparing the Performance of Different Elements” on page 644.

# 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 “The Toolbar in System Explorer” on page 312 for information about the toolbar in the upper pane.

The toolbar in the lower pane provides the following information:

TABLE 15-1    Toolbar in Lower Pane of Performance Explorer

Icon	Description
	Saves the graph in three formats: HTML, XLS, CSV.
	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: <ul style="list-style-type: none"><li>• <b>Bar</b> - Displays each data point as a bar. The data for the different elements is displayed side by side.</li><li>• <b>Line</b> - 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.</li><li>• <b>Summary</b> - 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 645.</li></ul>

**TABLE 15-1** Toolbar in Lower Pane of Performance Explorer (Continued)

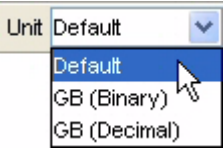

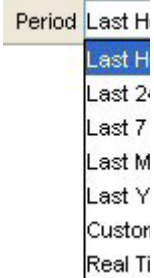








Icon	Description
	<p>Lets you determine the unit of measurement in the graph. Select one of the following options, and then click the  button:</p> <ul style="list-style-type: none"> <li>• <b>Default</b> - Displays the data in its default unit, which is usually megabytes.</li> <li>• <b>GB (Binary)</b> - Displays the data in binary gigabytes. (The computer handles the data in binary format. A gigabyte is equal to 1,073,741,824 bytes.)</li> <li>• <b>GB (Decimal)</b> - Displays the data in gigabytes. (A gigabyte is equal to 1,000,000,000 bytes.)</li> </ul>
	<p>Lets you format the graph to provide data within the time period specified. Select the option to the left of the <b>Period</b> combo box. Select one of the following options, and then click the  button:</p> <ul style="list-style-type: none"> <li>• <b>Last Hour</b> - Information collected in the last hour is reported. If you select Last Hour, the only frequency available is All.</li> <li>• <b>Last 24 Hours</b> - Information collected in the last 24 hours is reported.</li> <li>• <b>Last 7 Days</b> - Information collected in the last seven days is reported.</li> <li>• <b>Last Month</b> - Information collected in the last month is reported.</li> <li>• <b>Last Year</b> - Information collected in the last year is reported.</li> <li>• <b>Custom...</b> - Lets you created custom settings for the graph.</li> <li>• <b>Real Time</b> - Displays the data as it is currently being gathered by the management server. If you select Real Time, the only option for frequency is All and unit is Default.</li> </ul> <p>To set a custom period, select <b>Custom</b>. See “Setting a Custom Period” on page 647 for more information.</p>
	<p>Zoom-out button allows you to return to the original data set.</p>
	<p>Applies the current period, frequency and trending information to the chart.</p>
	<p>This icon appears in the tool bar of the chart. It lets you filter out the additional data series if the chart contains more than one series of data.</p>

TABLE 15-1 Toolbar in Lower Pane of Performance Explorer (Continued)

Icon	Description
	<p>Lets you modify the performance data displayed in the graph and change graph settings. When you select it, the following options are displayed:</p> <p><b>Performance data in the graph</b></p> <ul style="list-style-type: none"> <li>• <b>Data and Delete note pad</b> - Displays elements on the chart (including statistics for the element) and allows you to delete some or all of them.</li> </ul> <p><b>Specify graph settings</b></p> <p>Chart:</p> <ul style="list-style-type: none"> <li>• <b>Type</b> - Allows you to select from the line, bar, or summary chart</li> <li>• <b>Unit</b> - Scales the y-axis by Gigabyte. You can select either decimal or binary value.</li> <li>• <b>Line size</b> - Affects the size of the line for line charts.</li> <li>• <b>Frequency</b> - Allows you to chart information at specified intervals.</li> <li>• <b>Threshold Visible</b> - Mandatory field to set a threshold or graph an existing threshold.</li> <li>• <b>Threshold Value</b> - Actual value of the threshold identified by a red line displayed on line and bar charts.</li> <li>• <b>Show Grid</b> - Check this field to display the Chart Grid.</li> </ul> <p>Period:</p> <ul style="list-style-type: none"> <li>• <b>Period bullet</b> - Lets you format the graph to provide data within the time period specified</li> <li>• <b>Custom</b> - Allows you to customize a time period value by selecting start and end times from a calendar pop-up.</li> <li>• <b>Trend</b> - Specifies the number of trending data points you want to see. Trending information only applies to line and bar charts with a single statistic. The time period of the trend is determined by the frequency setting.</li> </ul>
	<p>Creates a chart window.</p>
 	<p>Lets you arrange the chart windows as follows:</p> <ul style="list-style-type: none"> <li>• <b>Tiled</b> - Displays the windows within their own tile.</li> <li>• <b>Cascade</b> - Displays the windows on top of each other, with the active window on top.</li> <li>• <b>Minimize All</b> - Minimizes all windows.</li> <li>• <b>Restore All</b> - Opens all windows that have been minimized.</li> </ul>

---

# 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 640.

2. Click the element in the topology whose performance you want to see.

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 under the **Application Path** node.

4. Select the graph monitoring option 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 > Hosts > HostA > Element Paths > Application Path**.

8. Select the same type of component you selected previously for Cortez.

9. Select the same type of graph monitoring option 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 summarizes the values for a single statistic. Multiple statistics can be shown with multiple lines.

---

**Note** – You cannot see more than one element at a time on a chart that has multiple elements.


---

The line has the following attributes:

- The vertical length of the line indicates the minimum and maximum value of the statistic within the selected data time frame.
- The green marker marks the median.
- The blue marker marks the average value.

Summary charts are currently not supported for real time performance display.

To view a summary chart:

1. Access Performance Explorer as described in “Accessing Performance Explorer” on page 640.
2. Create a chart as described in “Creating Performance Charts” on page 640.
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 charts. 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. Keep in mind that 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 640.
2. Click the element whose performance you want to see.
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 Graph Editing Dialog window, enter a number in the Trend box.

The number corresponds to the number of frequency intervals for which the trending information will be provided. For example, if you enter 5 in the Trend box, the chart provides trending information for five frequency intervals, such as five weeks if weeks was selected from the Frequency menu. If the frequency is set to **All**, the trend interval is an hour.

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 menu and you only have two days of data, a chart is not displayed, regardless of the value in the Trend box. 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 640.
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 Graph Editing Dialog window, click the **Delete** () 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, assume 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.

To set a custom period:

1. Access Performance Explorer as described in “Accessing Performance Explorer” on page 640.
2. Click the element whose performance you want to see.
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 component, such as a port on a switch.
4. In the lower-left pane on the Monitoring tab, click a performance monitoring option.  
The performance monitoring options vary according to the type of element.
5. Select **Custom** from the Period combo box in the lower pane.

The Performance Graph Editing Dialog window appears when you select the Custom option and there are no previous custom settings.

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 box.
8. Enter the time in the time box. 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 **Set**.

The start time and date are displayed in the Start box.

11. To set the end date, click the calendar  button to the right of the End box. Repeat steps 8 through 10.

The ending time and date are displayed in the End box.

12. Click **OK**.

13. Click the  button.



# Monitoring Options

Table 15-2, “About the Monitoring Options,” on page 649 describes some of the monitoring options.

**TABLE 15-2** 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)	LSI storage systems	The average input/out size (bytes/sec)
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.

**TABLE 15-2** About the Monitoring Options (*Continued*)

Available Monitoring Options	Available to Which Elements	Description
Buffer Hit Ratio	Oracle	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	LSI storage systems	Bytes transferred on a drive.
Bytes Transmitted (MB/Sec)	<ul style="list-style-type: none"><li>• Storage systems</li><li>• Host (port for HBA card)</li><li>• NAS filer (IP port)</li><li>• Switch Port</li></ul>	Number of bytes transmitted by the port per second
Bytes Received (MB/Sec)	<ul style="list-style-type: none"><li>• Storage systems</li><li>• Host (port for HBA card)</li><li>• NAS filer (IP port)</li><li>• Switch port</li></ul>	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 (not the host CPU) are using 100 percent of the Sybase CPU. The host CPU could be pegged at 60 percent, while the Sybase CPU is pegged at 100 percent.
CRC Errors (Errors/Sec)	<ul style="list-style-type: none"><li>• Storage systems</li><li>• Host (port for HBA card)</li><li>• Switch port</li></ul>	The number of cyclic redundancy check (CRC) errors per second.

**TABLE 15-2** About the Monitoring Options (*Continued*)

Available Monitoring Options	Available to Which Elements	Description
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.
Dictionary Hit Ratio	Oracle	The ratio you use to determine if the shared pool is large enough to store dictionary cache data adequately. If the ratio is less than 95 percent, investigate further. You may need to increase <code>shared_pool_size</code> .
Disk Read (KB/second) Not available on HP-UX hosts because HP-UX hosts do not return read/write data separately.	Disk drives	The speed at which the disk is read. To receive this data from a 64-bit AIX host, the <code>bos.perf.libperfstat</code> file must be installed on the host. See the Prerequisites section in the chapter about 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 <code>bos.perf.libperfstat</code> file must be installed on the host. See the Prerequisites section in the chapter about 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 <code>bos.perf.libperfstat</code> file must be installed on the host. See the Prerequisites section in the chapter about installing the CIM extension on AIX in the installation guide.

**TABLE 15-2** About the Monitoring Options (*Continued*)

Available Monitoring Options	Available to Which Elements	Description
Disk Write (KB/second) Not available on HP-UX hosts because HP-UX hosts do not return read/write data separately.	Disk drives	The speed to which the disk is written. To receive this data from a 64-bit AIX host, the <code>bos.perf.libperfstat</code> file must be installed on the host. See the Prerequisites section in the chapter about installing the CIM extension on AIX in the installation guide.
Exchange Services	Microsoft Exchange	Services Microsoft Exchange depends on to operate. A red square indicates the service is stopped; a green circle indicates the service is 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 I/O percentage against Total I/O. 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 (not supported on Microsoft Exchange 2007)	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 <code>bos.perf.libperfstat</code> 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.
Free Virtual Memory (KB)	Host	The amount of free virtual memory on the host. To receive this data from an AIX host, the <code>bos.perf.libperfstat</code> file must be installed on the host.

**TABLE 15-2** About the Monitoring Options (*Continued*)

Available Monitoring Options	Available to Which Elements	Description
In Memory Sort Ratio	Oracle	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 <code>sort_area_size</code> or <code>pga_aggregate_target</code> (for Oracle 9i if the <code>workarea_size_policy</code> is <code>auto</code> .)
Inode Cache Hits Count per Second	NAS filers (system)	The number of <b>inode</b> cache hits per second.
Inode Cache Misses Count per Second	NAS filers (system)	The number of <b>inode</b> cache misses per second.
Invalid CRC Errors (errors/second)	Storage systems	The speed at which invalid cyclic redundancy check (CRC) errors are found.
Library Cache Hit Ratio	Oracle	The number of times that 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 <code>shared_pool_size</code> .
Link Failures (failures/second)	<ul style="list-style-type: none"> <li>Storage systems</li> <li>Host (port for HBA card)</li> <li>Switch port</li> </ul>	The number of link failures per second.
Mail Box Count	Microsoft Exchange Stores	The number of mail boxes for a store.
Memory Usage Percentage (%)	Sybase	Percentage of the Sybase memory (not the host memory) being used.
Messages Awaiting Directory Lookup Queue Size	Microsoft Exchange (not supported on Microsoft Exchange 2007)	The number of messages waiting for the recipient to be resolved in the global catalog, so that the message can be sent.
Messages to be Routed Queue Size	Microsoft Exchange (not supported on Microsoft Exchange 2007)	The number of messages waiting to be routed to the recipient.

**TABLE 15-2** About the Monitoring Options (*Continued*)

<b>Available Monitoring Options</b>	<b>Available to Which Elements</b>	<b>Description</b>
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 cache 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.
Packets Transmitted (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 (not supported on Microsoft Exchange 2007)	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.

**TABLE 15-2** About the Monitoring Options (*Continued*)

<b>Available Monitoring Options</b>	<b>Available to Which Elements</b>	<b>Description</b>
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 about 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.
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 about installing the CIM extension on AIX in the installation guide.
Percent Read (%)	LSI 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)	LSI 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.

**TABLE 15-2** About the Monitoring Options (*Continued*)

Available Monitoring Options	Available to Which Elements	Description
Recovered Errors	LSI Storage Systems	The number of recovered errors on the drive.
Redo Buffer Allocation Retries Ratio	Oracle	The number of times DSP had to wait to write to the log buffer. If the ratio is more than 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 the log file. If the ratio is more than .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 <b>inode files</b> on a NAS file system.
Retried Requests	LSI storage systems	The number of retried requests for a drive.
Send Queue Size	Microsoft Exchange Stores (not supported on Microsoft Exchange 2007)	The size of the send queue.
SMTP Local Delivery Queue Size	Microsoft Exchange (not supported on Microsoft Exchange 2007)	The number of messages in the SMTP Local Delivery Queue, the queue hold mail box of local recipients.
SMTP Server Queues Summary	Microsoft Exchange (not supported on Microsoft Exchange 2007)	The total number of messages in the following queues: <ul style="list-style-type: none"> <li>• Final Destination Currently Unreachable</li> <li>• Messages Awaiting Directory Lookup</li> <li>• Messages to Be Routed</li> <li>• Pre-submission</li> <li>• SMTP Local Delivery</li> </ul>
Storage Group Size	Microsoft Exchange storage groups	The size of the storage group in megabytes.



**TABLE 15-2** About the Monitoring Options (*Continued*)

Available Monitoring Options	Available to Which Elements	Description
Store size (MB)	Microsoft Exchange Stores	The size of the log in megabytes.
System Event Time Waited (ms)	Oracle	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	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	LSI storage systems	The number of timeouts on the drive.
Total Bandwidth (Bytes/Sec)	LSI storage systems	The total bandwidth.
Total Inodes	NAS filers (file systems)	The total number of <b>inode files</b> for a NAS file system.
Total IOs	LSI storage systems	The total input/output.
Total IO Rate (IOs/Sec)	LSI storage systems	The total input/output rate.
Transaction Log Size (MB)	Microsoft Exchange Transaction Logs	The size of the store in megabytes.
Unrecovered Errors	LSI storage systems	The number of unrecovered errors on a drive.
Used Inodes	NAS filers (file systems)	The number of unused <b>inode files</b> 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 <code>bos.perf.libperfstat</code> file must be installed on the host.
Write IO Rate (Writes/Sec)	LSI storage systems	The input/output of the write rate.

**TABLE 15-2** About the Monitoring Options (*Continued*)

Available Monitoring Options	Available to Which Elements	Description
Write Operations	EMC storage systems	Write operations (bytes/second).

---

## Managing Late Data or Errors

If you are performing real time data collection, and the element is not returning the information in time, you are shown a message in red resembling the following:

Data is late or an error occurred...

The software cannot obtain the information in a timely manner because of 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.
- The specific element is already being monitored in real time on another chart.

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 disk drives on a server, you might want to try collecting information for one disk drive.

Performance Explorer will continue to attempt to retrieve data from the element until the chart is closed.

---

## Monitoring with Direct Attached Storage

A port on a storage system that is 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

Table 15-3, “Host Monitoring Support,” on page 660 describes which host configurations the management server can monitor. The management server supports configurations that have a Y next to them. Unsupported 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 following table and one of them cannot be monitored, monitoring is not supported for any 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, neither of those devices can be monitored, because Solaris 9 Sun Foundation Suite Leadville with MPXIO is not supported, as shown with the following formula:

A monitorable configuration (Y) + an 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 15-3** Host Monitoring Support

Host Configuration	Monitoring Supported?
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
AIX 5.2, 5.3 with AMS	N
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
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
Redhat 2.1 HDLM	N
Redhat 3.0	Y
Redhat 3.0 HDLM	N
Redhat 4.0	Y
Redhat 4.0 Itanium	Y
SGI ProPack 3.0	Y
SGI ProPack 3.0 XVM	Y
SGI ProPack 3.0 CXFS	Y (only on node sending I/O)

**TABLE 15-3** Host Monitoring Support (*Continued*)

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
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

## 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  $2^{31}$  (2147483648) integers. The counters revert back to zero after reaching  $2^{31}$  (2147483648) integers, and 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, and 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 contains  $2^{31}$  (2147483648) integers. Instead, the values continue to rise smoothly.

The following charts for aggregate drives are impacted: ReadIOs, WriteIOs, TotalIOs, Bytes Transferred, Unrecovered Errors, Recovered Errors, Timeouts, Retried Requests. Charts for aggregated drives are also provided for arrays, controllers, and volume groups.

The following charts for aggregate volumes are impacted: Bytes Read, Bytes Read Large, Bytes Written, Bytes Written Large, ReadIOs Large, WriteIOs Large, ReadIOs, WriteIOs, TotalIOs, Bytes Transferred, ReadHitIOs. Charts for aggregated drives are also provided for arrays, controllers, and volume groups.

## Finding an Element's Storage Capacity

---

**Caution** – 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 contains the following topics:

- “About Capacity Explorer” on page 663
- “Accessing Capacity Explorer” on page 666
- “The Toolbars in Capacity Explorer” on page 666
- “Finding the Capacity of an Element” on page 668
- “Obtaining Utilization Reports” on page 675
- “Printing Elements in Capacity Explorer” on page 676
- “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

**Caution** – 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 with the exception of Windows 2000 hosts. See “How Capacity Differs in Chargeback and Capacity Explorer” on page 731 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 807.
- 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 with the exception of Windows 2000 hosts.
- 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.

## **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 is highlighted in the topology.

## **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.

## **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**




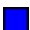

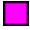
## ■ Port Utilization by Fabric

See “Obtaining Utilization Reports” on page 675 for more information.

### 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. The following table lists the colors:

**TABLE 16-1** Color Coding for Capacity Explorer

Color	Description
	Description varies according to element type: <ul style="list-style-type: none"><li>• <b>For storage systems</b> — The space is mapped.</li><li>• <b>For all elements except storage systems</b> — The space is used.</li></ul>
	Unallocated
	Unused raw (storage systems only)
	Unmapped (storage systems only)

The colors indicate that the element displayed in the following figure is about 75 percent available. The rest of it is being used.



**FIGURE 16-1** 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 668.

### Capacity Chart Tab

The Capacity Chart tab lets you 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** (  ).

## 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 “The Toolbar in System Explorer” on page 312 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:

---

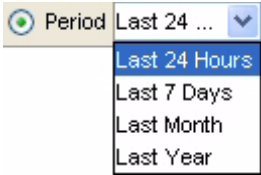



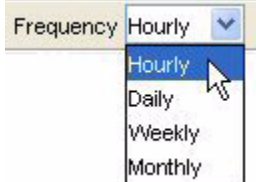
**Note** – Capacity Explorer provides additional tabs for NetApp NAS devices. The toolbars are the same.

---

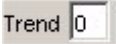



**TABLE 16-2** Toolbar in Lower Pane of Capacity Explorer

Icon	Description
	<p>Lets you switch between the following:</p> <ul style="list-style-type: none"><li>• <b>Bar</b> — Displays each data point as a bar. The data for the different elements is displayed side by side.</li><li>• <b>Line</b> — 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.</li></ul>

**TABLE 16-2** Toolbar in Lower Pane of Capacity Explorer (*Continued*)

Icon	Description
	<p>Lets you format the graph to provide data within the time period specified. Select the option to the left of the <b>Period</b> combo box. Select one of the following from the menu. To update the chart, click the  button.</p> <ul style="list-style-type: none"> <li>• <b>Last 24 Hours</b> — Information collected in the last 24 hours is reported. This option is only available to hosts and applications.</li> <li>• <b>Last 7 Days</b> — Information collected in the last seven days is reported.</li> <li>• <b>Last Month</b> — Information collected in the last month is reported.</li> <li>• <b>Last Year</b> — Information collected in the last year is reported.</li> </ul>
	<p>: the graph to provide starting and ending time specified.</p>
	<p><i>Applications and hosts only:</i> Lets you change the display frequency. The options are the following:</p> <ul style="list-style-type: none"> <li>• <b>Hourly</b> — The information is displayed in hourly increments.</li> <li>• <b>Daily</b> — The information is displayed in daily increments.</li> <li>• <b>Weekly</b> — The information is displayed in weekly increments.</li> <li>• <b>Monthly</b> — The information is displayed in monthly increments.</li> </ul>

**TABLE 16-2** Toolbar in Lower Pane of Capacity Explorer (*Continued*)

Icon	Description
	<p>Lets you set trending information.</p> <p>When a switch or storage system is selected, the frequency box is set to hourly.</p> <p>See “Viewing Trending Information for Storage Capacity” on page 678.</p> <p>Keep in mind the following:</p> <ul style="list-style-type: none"><li>• An element's performance can drastically change in the future. Keep in mind that the data trends are just assumptions and should not be treated as fact.</li><li>• Trending requires at least two sets of data gathered within the frequency specified.</li></ul>
	<p>Applies period, frequency and trending information.</p>
	<p>Lets you filter out the additional data series if the chart contains more than one series of data.</p>
	<p>Lets you print a graph.</p>

---

## Finding the Capacity of an Element

---

**Note** – Capacity Explorer rounds the data it displays. As a result, the totals you add for a property may be different from the data displayed in the Summary column. For example, if you add the total capacity from each data pool and compare that total to the number for Total Capacity displayed under the Summary heading, you will most likely find that the results differ.

---

To find the capacity of an element:

1. Access Capacity Explorer as described in “Accessing Capacity Explorer” on page 666.

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.

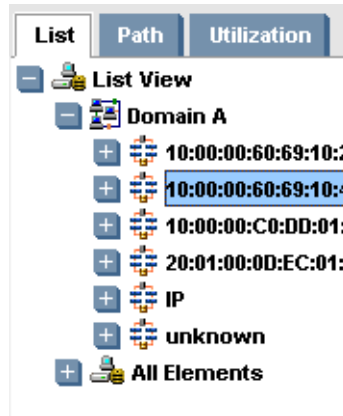


FIGURE 16-2 Viewing the Capacity of Elements in a Fabric

See “About Capacity Explorer” on page 663 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 group (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 box 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, or your license does not allow you to discover a particular storage system. See the support matrix 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:

---

**Note** – Aggregate tabs only apply to NetApps hosts.

---

- **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 an SNMP trap is generated.
- **Quota Used**
- **Percentage Used**

Capacity Explorer displays the following options under the Capacity Summary and Capacity columns on the Snapshot Data tab. Definitions are provided here.

---

**Note** – Snapshot tabs only apply to NetApps hosts.

---

- Total Volume Space — Sum of space on all volumes
- Total Space Reserved — Sum of the space reserved for snapshots across all volumes
- Total Space Used — Sum of space used by snapshots on all volumes

---

**Note** – **Total Space Used** is the sum of the largest total space used on each volume.

---

- Cumulative% of Total Vol — Percentage of space used for snapshots from the total volume space across all volumes.
- Cumulative% of Used Vol — Percentage of space used for snapshots from the total used space across all volumes.

Capacity Explorer displays the following information about each volume snapshot:

- Snapshot Name
- % Reserved
- % of Total Vol
- % of Used Vol
- Cumulative% of Total Vol
- Cumulative% of Used Vol
- Total Space Used

## Capacity Information for Storage Systems

Capacity Explorer provides three tabs for storage systems:

- **Raw Capacity** - See “Viewing the raw capacity of a storage system” on page 672.
- **Post RAID** - See “Viewing post-RAID information” on page 673.
- **Capacity Chart** - See “Viewing Capacity Charts” on page 676.

### Viewing the raw capacity of a storage system

The Raw Capacity tab displays the raw capacity in megabytes for the ports.

- **Total Raw** - The sum of used raw and unused raw capacity.
- **Used Raw** - Raw disk capacity that is consumed by RAID groups or other such disk groups on the array. Disks that have been configured for use in provisioning volumes, regardless of whether volumes have been allocated from those disk groups. For example, on CLARiiON storage systems these are disks that are in RAID groups.



- **Unused Raw** - Raw disk capacity not currently configured into any kind of RAID or disk group. Disks that have not been configured for use in provisioning volumes. For example, on the EVA or MSA, these are not part of any volume group. For enterprise arrays that are preconfigured at installation, Unused Raw will typically be zero.
  - **Percentage Used** - The percentage of raw disk capacity used.
- 

**Caution** – The raw capacity values come directly from the SMI instrumentation of storage arrays, where raw capacity is modeled as primordial storage pools.

---

The Raw Capacity tab provides information about the raw capacity used on a storage system. A graphic displaying what which percentage of the storage system has used raw capacity and unused raw capacity.

## Viewing post-RAID information

You can view post-RAID information from the Post RAID tab. The following information is displayed on the tab.

displays the following information under the **Capacity Summary** and **Capacity** columns:

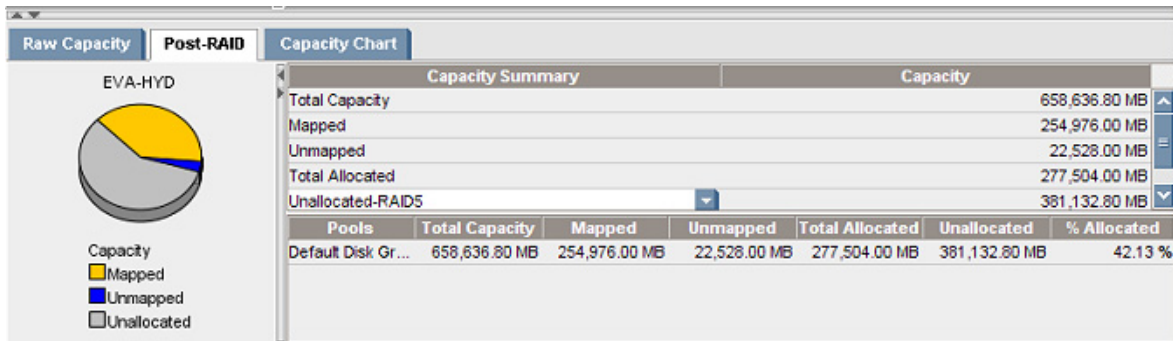
- **Total Capacity** - The sum of mapped, unmapped and unallocated capacities.
  - **Mapped** - Sum of mapped volume capacities allocated from storage pools. For a volume to be 'mapped' it must be have a logical mapping to at least one host initiator.
  - **Unmapped** - Sum of unmapped volume capacities allocated from storage pools. An unmapped volume is storage committed as a single volume but not visible or potentially visible to any initiator.
  - **Total Allocated** - Sum of capacity in storage pools not available for creating volumes.
  - **Unallocated** - Sum of capacity in storage pools available for creating volumes.
- 

**Note** – Primordial pools are not available for creating volumes, so they do not contribute to this total.

---

- **Unused Raw** - Raw disk capacity not currently configured into any kind of RAID or disk group. Disks that have not been configured for use in provisioning volumes. For example, on the EVA, these are not part of any volume group. For enterprise arrays that are preconfigured at installation, Unused Raw will typically be zero.
- **Percentage Allocated**

**Caution** – For arrays that permit RAID choice when creating volumes (for exaple the EVA), the concept of flexible RAID Pools is introduced. In this case, the amount of unallocated space is dependent on the RAID level chosen for new volumes. Therefore, the values for Unallocated and Total Capacity displayed may be modified dynamically by choosing a RAID level for unallocated space.



**FIGURE 16-3** Post-RAID tab

The following table describes how the properties are calculated for HDS array groups.

**TABLE 16-3** 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	The sum of the 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	The sum of the 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 do not have LUNs but are not on the management server's "free" list
Total Allocated	sum of volume.size for all volumes	All LDEVs that are not on the management server's "free" list

**TABLE 16-3** 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
Unallocated	pool.totalRemainingSpace	LDEVs on the management server's "free" list plus the total free space

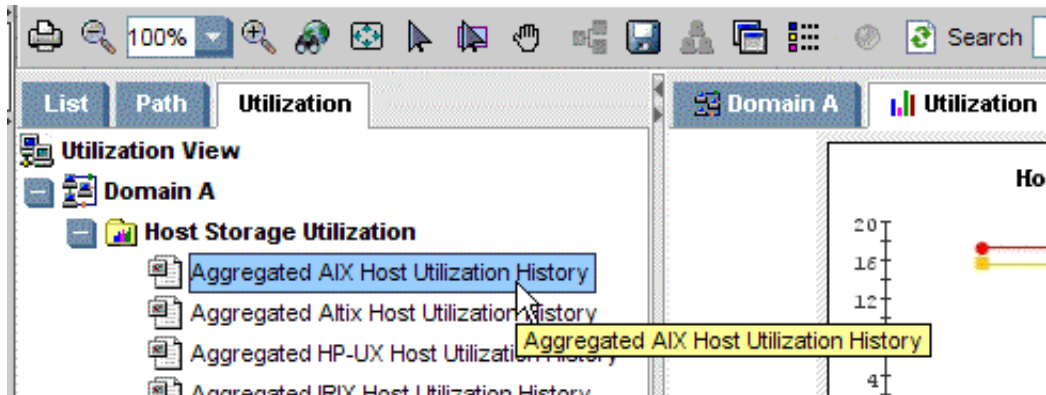
## Obtaining Utilization Reports

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 666.
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 16-4** 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 666.
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. For more information, see “Configuring Print Settings” on page 222.

---

# Viewing Capacity Charts

Capacity Explorer provides a graphical representation of the capacity history of an element, such as port summary information for switches.

The following lists the types of capacity charts available:

- Volume
- Aggregate
- Quota
- Snapshot

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 that the performance collector for the element is enabled (**Configuration > Performance > Data Collection**). See “Managing Performance Collection” on page 270 for more information about enabling performance collectors.
- An hourly roll-up of capacity chart data occurs daily at 2 a.m. Capacity chart data is not available until after Get Details is performed and the hourly roll-up occurs.

- Switches and storage systems display data from the last time Get Details was performed.
- If you see the message “There is not enough data to produce a chart [chart\_title] at this time,” lessen the frequency option, so that the amount of time listed in the **Frequency** 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 666.
2. Click the element whose capacity history you want to see.
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** – To change the orientation of the chart, hold down the mouse button when you click the chart, and continue holding it while you move the mouse.

---


5. To change the reporting period, do one of the following, and 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 menu.
    - Last 24 hours** — This option is not available to switches and storage systems.
    - Last 7 Days**
    - Last Month**
    - Last Year**
  - **Display data within a starting and ending time** — You can format the graph to provide data within the starting and ending time specified:
    - a. Select the option to the left of the Start box.
    - b. Click the  icon.
    - c. In the Time box, enter the time you want the graph to start, using the 24-hour format.
    - d. Select a date and click **Set**.

- e. Repeat the steps for the End box.



---

**Caution** – If you change the date in the box 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** menu.
  - **Hourly**
  - **Daily**
  - **Weekly**
  - **Monthly**
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, enter an integer in the Trend box.

The integer corresponds to the number of frequency intervals for which the trending information will be provided. For example, if you enter 5 in the Trend box, the chart provides trending information for five frequency intervals, such as five weeks if weeks was selected from the Frequency menu.

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. Keep in mind that 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 666.
2. Click the element whose performance history you want to see.
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, enter a number in the Trend box.

The number corresponds to the number of frequency intervals for which the trending information will be provided. For example, if you enter 5 in the Trend box, the chart provides trending information for five frequency intervals, such as five weeks if weeks was selected from the Frequency 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** menu, and you have only two days of data, a chart is not displayed, regardless of the value in the Trend box. 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 that 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 as follows:

```
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 run 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 does this happen? When you run the `df -k` command, the computer runs 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 calculates the total capacity by using the equation discussed previously (used capacity + available capacity) and displaying 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.



## Managing Policies

---

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 contains the following topics:

- “About Policy Manager” on page 681
- “Accessing Policy Manager” on page 682
- “Creating Policies” on page 683
- “Modifying Policies” on page 692
- “Viewing Policies” on page 697
- “Deactivating a Policy” on page 697
- “Deleting Policies” on page 698
- “Providing E-mail Notification for a Policy” on page 698
- “Providing Event Generation for a Policy” on page 699
- “Providing a Custom Command for a Policy” on page 700

---

## 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 events occurs:

- A new element is discovered
- Successful provisioning occurred
- An event occurs on one or more specified elements

---

**Caution** – Collectors gather information for reports, monitoring and File Server SRM. Policies are not triggered when a collector is running. If you configured a policy to be triggered when a host has a critical event and the collector is running

when a critical event occurs, the policy is not triggered. You can modify a collector's schedule, by going to **Configuration > Reports, Configuration > Performance** or **Configuration > File SRM**.

---

Policy Manager allows you to create the following types of policies:

- **Utilization policies** - Monitor 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.

- **Protection policies** - Monitor the backups on your network. For example, polices can be created to notify you if a backup failed on a backup client.
- **Infrastructure policies** - Monitor the following:
  - the discovery of a new element
  - successful provisioning
  - the occurrence of an event on one or more specified elements

---


## Accessing Policy Manager

This section describes the various techniques for accessing Policy Manager. Although there is only one way to access policies for discovery, provisioning, and events; there are multiple ways to access policies. This 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 policies 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 policies for backup elements** - Select one of the following elements in Protection Explorer, and then click **Policies** in the lower-right corner:
  - Backup Clients
  - Backup Library
  - Master backup server
  - Master backup media

To access all policies in Policy Manager:

1. Click **Policy Manager** (  ).
2. Click a policy in the Policy Manager tree.

---

## Creating Policies

This section contains the following topics:

- “Actions Available for When a Policy Condition is Fulfilled” on page 683
- “Creating a Utilization or Backup Policy” on page 684
- “Creating Policies for Discovery” on page 688
- “Creating Policies for Provisioning” on page 689
- “Creating Policies for Events” on page 690
- “Testing a Utilization Policy” on page 692

## Actions Available for When a Policy Condition is Fulfilled

When you create or modify a policy, you must select an action to occur when the policy condition is fulfilled. More than one action can be assigned to a policy. The following actions are available:

- **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, and the event appears in Event Manager. After you select a severity level, click **OK**. For a list of the severity levels you can select from, see “Severity Levels” on page 684.
- **Execute a Custom Command** - Policy Manager executes a custom command on the management server when the condition is fulfilled. Enter a command that will execute the script in the box, and click **OK**.

The software assumes you are in the %JBSS4\_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`, where `POLICY_NAME` provides the policy name and `POLICY_DESCRIPTION` provides a description of the policy. See “Software Environment Variables for Scripting” on page 362 for more information.

Prefix the command with “start” if the custom command triggers a user interface component, for example, Microsoft Internet Explorer or a command prompt window. For example, if you want the custom command to open a command prompt window and list the contents in the directory, you would prefix the command as follows:

```
start dir
```

## Severity Levels

When you create or modify a policy, you can select the severity level for which you want Policy Manager to generate an event. The following severity levels are available:

---

**Caution** – 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.

## Creating a Utilization or Backup Policy

You can create a utilization or backup 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 221.
- Policies that are triggered for virtual applications are also triggered for file servers.

To create a utilization or Backup 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 **Add**.
4. Select a policy, and click **Next**. See Table 17-1, “Policy Templates,” on page 686 for information about the policy types available.

---

**Note** – This step is not applicable for backup libraries.

---

5. In the Name box on the Policy Properties tab, enter a name for the policy or keep the default.
6. In the Description box, modify the description for the policy or keep the default.
7. In the Re-arm Period box, specify the amount of time (in minutes) after the policy executes before it can execute again. The re-arm period is useful for limiting the number of times the same actions will execute.

---

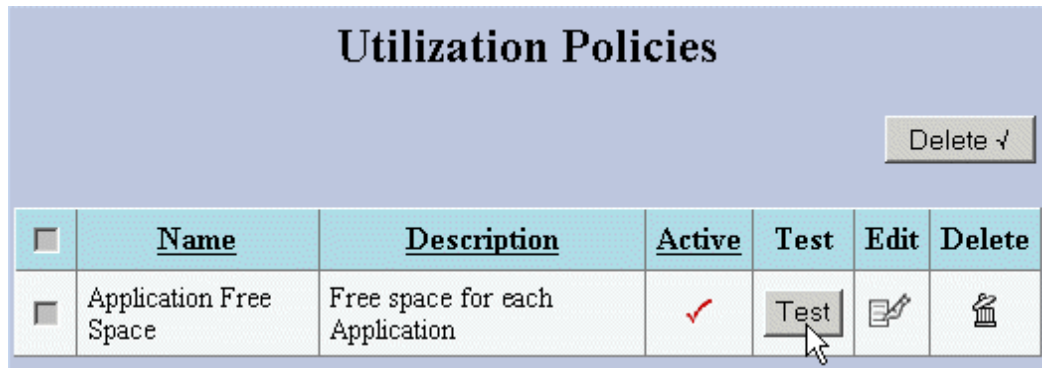
**Caution** – Specify shorter periods for important applications.

---

8. If you are setting a policy for a backup client, go to Step 11.
9. In the Condition menu, specify a comparison operator.
10. (Available for only backup libraries). To the left of the comparison operator, select the media pool you want to monitor.
11. To the right of the Comparison Operator menu, take the following action, depending on the box displayed:
  - Enter a percentage.
  - Enter an amount in gigabytes (GB).
  - Type the number of available media that will trigger an alert. For example, if you want to be alerted when the number of available media for a storage pool is less than two, you would set the conditional to less than (<). You would then enter 2 in the Media box.
12. For trending policies, enter the number of days in the Historic period box (min=7; max=180) and Projection period box (min=1; max=180).
13. Select an action to occur when the policy condition is fulfilled. For more information, see “Actions Available for When a Policy Condition is Fulfilled” on page 683.
14. Click **Finish**.

15. 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 17-1** Testing a Newly Created Utilization Policy

## About the Policy Templates

The following table provides descriptions of the policy templates.

**TABLE 17-1** Policy Templates

Policy Type	Description	Applies to
Backup failure	A backup session for the selected backup client failed.	Backup Clients
Backup partially complete	A backup session for the selected backup client finished with partial success.	Backup Clients
Backup successful	A backup session for the selected backup client finished successfully.	Backup Clients
Backup warning	A backup session for the selected backup client finished with a warning.	Backup Clients

**TABLE 17-1** Policy Templates (*Continued*)

Policy Type	Description	Applies to
Free Space	<p>The amount of free space on one of the following:</p> <ul style="list-style-type: none"> <li>• A host</li> <li>• A Microsoft Exchange instance</li> <li>• A database instance, such as Microsoft SQL Server, Sybase, Caché, or Oracle.</li> <li>• A file server</li> <li>• A virtual application</li> </ul>	<ul style="list-style-type: none"> <li>• Hosts</li> <li>• Backup clients</li> <li>• Microsoft Exchange</li> <li>• Microsoft SQL Server</li> <li>• Oracle</li> <li>• Sybase</li> <li>• Caché</li> <li>• File servers</li> <li>• Virtual applications</li> </ul>
Free Space Trending	<p>A forecast of amount of free space on one of the following:</p> <ul style="list-style-type: none"> <li>• A host</li> <li>• A database instance, such as Microsoft SQL Server, Sybase, Caché, or Oracle.</li> <li>• A file servers</li> <li>• A virtual applications</li> </ul>	<ul style="list-style-type: none"> <li>• Hosts</li> <li>• Microsoft SQL Server</li> <li>• Oracle</li> <li>• Sybase</li> <li>• Caché</li> <li>• File servers</li> <li>• Virtual applications</li> </ul>
Percent Used	<p>Description varies according to element type:</p> <ul style="list-style-type: none"> <li>• <b>Hosts</b> - Percent of storage used for a host.</li> <li>• <b>Microsoft Exchange instances</b> - Percent of storage used for a Microsoft Exchange instance.</li> <li>• <b>Database instances</b> - Percent of storage used for a database instance, such as Microsoft SQL Server, Sybase, Caché, or Oracle.</li> <li>• <b>Switches</b> - Percent of ports used for a switch.</li> <li>• <b>Storage System</b> - Percent of ports used for a storage system</li> <li>• <b>File Servers</b> - Percent of disk space used.</li> <li>• <b>File Server User</b> - Percent of storage used for a file server user.</li> <li>• <b>Virtual Application</b> - Percent of storage used for a virtual application.</li> </ul>	<ul style="list-style-type: none"> <li>• Hosts</li> <li>• Backup clients</li> <li>• Microsoft Exchange</li> <li>• Microsoft SQL Server</li> <li>• Oracle</li> <li>• Sybase</li> <li>• Caché</li> <li>• Storage systems</li> <li>• Switches</li> <li>• File servers</li> <li>• File server users</li> <li>• Virtual applications</li> </ul>

**TABLE 17-1** Policy Templates (*Continued*)

Policy Type	Description	Applies to
Trending for Percent Used	<p>Forecast of percent used. Description varies according to element type:</p> <ul style="list-style-type: none"> <li>• <b>Host</b> - Percent of storage used for a host.</li> <li>• <b>Database instance</b> - Percent of storage used for a database instance, such as Microsoft SQL Server, Sybase, Caché, or Oracle.</li> <li>• <b>Switch</b> - Percent of ports used for a switch. <ul style="list-style-type: none"> <li>■ <b>File Server</b> - Percent of disk space used.</li> <li>■ <b>Virtual Application</b> - Percent of storage used for a virtual application.</li> </ul> </li> </ul>	<ul style="list-style-type: none"> <li>• Hosts</li> <li>• Microsoft SQL Server</li> <li>• Oracle</li> <li>• Sybase</li> <li>• Caché</li> <li>• Switches</li> <li>• File servers</li> <li>• Virtual applications</li> </ul>
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
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.	<ul style="list-style-type: none"> <li>• Switches</li> <li>• Storage systems</li> </ul>
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.

---

**Caution** – 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 221.

---

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 Infrastructure Policies node, and then select **New Element Discovery**.
3. Click the **Add** button in the right pane.
4. In the Name box, enter a name for the policy.
5. In the Description box, enter a description for the policy.
6. Select one or more element types.

When a condition is fulfilled on a selected 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 act when the event is cleared. If you do not select "Fire when event is cleared," the policy is triggered when the event is received.
8. Select one of the following options from the Summary Text list to establish how Policy Manager will scan the summary text and respond:
  - **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. No additional text is required.
  - **Contains** - If the event's summary text contains the text you specify here, Policy Manager sends an e-mail, generates an event, or runs a custom command. Enter the desired text in the box that appears to the right of the Summary Text menu.
  - **Matches Regular Expression** - If the event's summary text matches the the expression you specify here, Policy Manager sends an e-mail, generates an event, or runs a custom command. Enter the desired expression in the box that appears to the right of the Summary Text menu.

9. Select an action to occur when the policy condition is fulfilled. For more information, see “Actions Available for When a Policy Condition is Fulfilled” on page 683.
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.

---

**Caution** – 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 221.

---

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, and then click **Provisioning**.
3. Click **Add**.
4. In the Name box, enter a name for the policy.
5. In the Description box, enter 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 act when the event is cleared. If you do not select "Fire when event is cleared," the policy is triggered when the event is received.
8. Select an action to occur when the policy condition is fulfilled. For more information, see “Actions Available for When a Policy Condition is Fulfilled” on page 683.
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 a specific type of event occurs 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 221.

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 Infrastructure Policies node, and then click **Events**.
3. Click **Add**.
4. In the Name box, enter a name for the policy.
5. In the Description box, enter a description for the policy.
6. In the Re-arm Period box, specify the amount of time (in minutes) after the policy executes before it can execute again. The re-arm period is useful for limiting the number of times the same actions will execute.

---

**Caution** – 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 act when the event is cleared. If you do not select “Fire when event is cleared,” the policy is triggered when the event is received.
9. In the Severity menu, specify a comparison operator.
10. To the right of the Severity menu, select one of the following severity levels:

---

**Caution** – 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.
11. Select one of the following options from the Summary Text list to establish how Policy Manager will scan the summary text and respond::
- **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 text you specify here, Policy Manager sends an e-mail, generates an event, or runs a custom command. Enter the desired text in the box that appears to the right of the Summary Text menu.
  - **Matches Regular Expression** - If the event's summary text matches the the expression you specify here, Policy Manager sends an e-mail, generates an event, or runs a custom command. Enter the desired expression in the box that appears to the right of the Summary Text menu.
12. Select an action to occur when the policy condition is fulfilled. For more information, see “Actions Available for When a Policy Condition is Fulfilled” on page 683.
13. Click **OK**.

## Testing a Utilization Policy

After you create or modify a utilization policy, test it to verify that 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 again, 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.

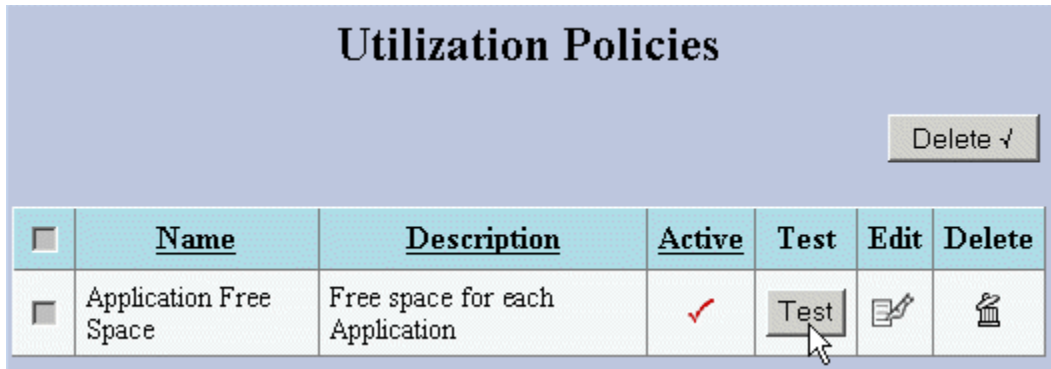


FIGURE 17-2 Testing a Utilization Policy

## Modifying Policies

This section contains the following topics:

- “Modifying Utilization and Backup Policies” on page 693
- “Modifying Discovery Policies” on page 694
- “Modifying Provisioning Policies” on page 695
- “Modifying Policies for Events” on page 695

## Modifying Utilization and Backup Policies

To modify a utilization or backup policy:

1. Access Policy Manager as described in the topic, “Accessing Policy Manager” on page 682.
2. Click the **Edit** (📝) button corresponding to the policy you want to modify.
3. In the Name box, change the name for the policy.
4. In the Description box, change the description for the policy.
5. Select or deselect the **Active** option to activate or de-activate the policy.
6. In the Re-arm Period box, specify the amount of time (in minutes) after the policy executes before it can execute again. The re-arm period is useful for limiting the number of times the same actions will execute.

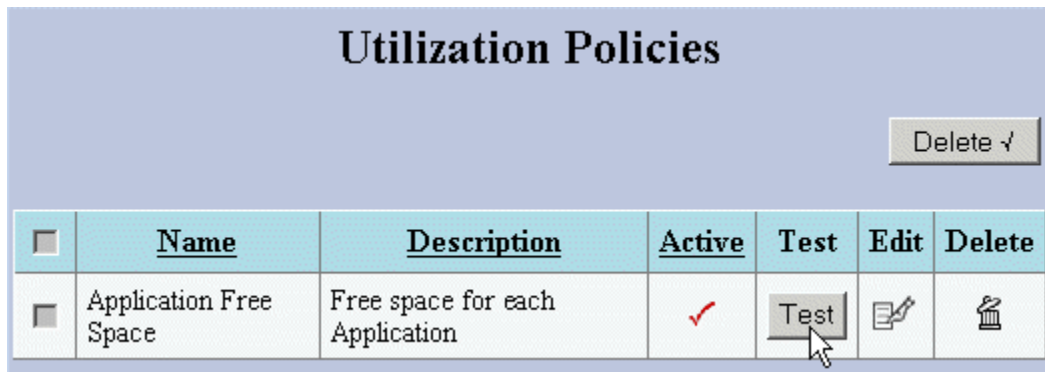
---

**Caution** – Specify shorter periods for important applications.

---

7. *Skip this step for backup clients:* In the Condition menu, change the conditions of the policy.
8. For trending policies, change the number of days in the Historic period box (min=7; max=180) or Projection period box (min=1; max=180).
9. Select an action to occur when the policy condition is fulfilled. For more information, see “Actions Available for When a Policy Condition is Fulfilled” on page 683.
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 17-3** 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, and then click **New Element Discovery**.
3. Click the **Edit** () button corresponding to the policy you want to modify.

4. In the Name box, change the name for the policy.
5. In the Description box, 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 act when the event is cleared. If you do not select "Fire when event is cleared," the policy is triggered when the event is received.
8. Select an action to occur when the policy condition is fulfilled. For more information, see "Actions Available for When a Policy Condition is Fulfilled" on page 683.
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 221.

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, and then click **Provisioning**.
3. Click the **Edit** () button corresponding to the policy you want to modify.
4. In the Name box, change the name for the policy.
5. In the Description box, 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 act when the event is cleared. If you do not select "Fire when event is cleared," the policy is triggered when the event is received.
8. Select an action to occur when the policy condition is fulfilled. For more information, see "Actions Available for When a Policy Condition is Fulfilled" on page 683.
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, and then click **Events**.
3. Click the **Edit** () button corresponding to the policy you want to modify.
4. In the Name box, change the name for the policy.
5. In the Description box, enter the description for the policy.
6. In the Re-arm Period box, specify the amount of time (in minutes) after the policy executes before it can execute again. The re-arm period is useful for limiting the number of times the same actions will execute.

---

**Caution** – 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 act when the event is cleared. If you do not select "Fire when event is cleared," the policy is triggered when the event is received.
9. In the Severity menu, specify a comparison operator.
10. To the right of the Severity menu, you can change the severity level. For a list of the severity levels you can select from, see "Severity Levels" on page 684.



11. To change how Policy Manager scans the summary text, select one of the following from the Summary Text 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 text you specify here, Policy Manager sends an e-mail, generates an event, or runs a custom command. Enter the desired text in the box that appears to the right of the Summary Text menu.
  - **Matches Regular Expression** - If the event's summary text matches the the expression you specify here, Policy Manager sends an e-mail, generates an event, or runs a custom command. Enter the desired expression in the box that appears to the right of the Summary Text menu.
12. Select an action to occur when the policy condition is fulfilled. For more information, see “Actions Available for When a Policy Condition is Fulfilled” on page 683.
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 17-2** 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 <b>Edit</b> (  ) button to edit a policy.
Delete	Click the <b>Delete</b> (  ) 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 **Edit** () button corresponding to the policy you want to modify.
3. Deselect the **Active** option.
4. Click **Finish**.

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. Take one of the following actions:
  - **Delete a policy** - If you want to delete just one policy, click the **Delete** () button corresponding to the policy you want to delete.
  - **Delete several policies at once** - If you want to delete several policies at one time, 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, and 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 events occur:

- A new element is discovered
- Provisioning is successful
- An event occurs 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 221 for more information.
- The following instructions assume you have already created a policy. If you have not yet created a policy, see the following topics:
  - “Creating a Utilization or Backup Policy” on page 684
  - “Creating Policies for Discovery” on page 688
  - “Creating Policies for Provisioning” on page 689
  - “Creating Policies for Events” on page 690

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 **Edit** () button corresponding to the policy you want to modify.
3. Click **Send E-mail**.

Policy Manager sends an e-mail when the condition is fulfilled. The software verifies that the address is entered in the correct format.

4. Enter a comma-separated list of e-mail addresses, and then click **OK**.
5. Click **Finish**.

---

# 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

- Provisioning is successful
- An event occurs 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 or Backup Policy” on page 684
- “Creating Policies for Discovery” on page 688
- “Creating Policies for Provisioning” on page 689
- “Creating Policies for Events” on page 690

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 **Edit** () button corresponding to the policy you want to modify.
3. Click **Generate Event**.  
Select the severity level for which you want Policy Manager to generate an event. For a list of the severity levels you can select from, see “Severity Levels” on page 684.
4. After you select an event level, click **OK**.
5. Click **Finish**.

---

## 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
- Provisioning is successful
- An event occurs 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 or Backup Policy” on page 684
- “Creating Policies for Discovery” on page 688
- “Creating Policies for Provisioning” on page 689
- “Creating Policies for Events” on page 690

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 **Edit** () button corresponding to the policy you want to modify.
3. Click **Execute Custom Command**.
4. Enter a command that will execute the script in the box, and click **OK**.

The software assumes you are in the %JBASS4\_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, where POLICY\_NAME provides the policy name and POLICY\_DESCRIPTION provides a description of the policy. See “Software Environment Variables for Scripting” on page 362 for more information.

Prefix the command with “start” if the custom command triggers a user interface component, for example, Microsoft Internet Explorer or a command prompt window. For example, if you want the custom command to open a command prompt window and list the contents in the directory, you would prefix the command as follows:

```
start dir
```

5. Click **Finish**.

Policy Manager executes a remote script on the management server when the condition is fulfilled.



# Chargeback

---

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**). Chargeback helps you manage departmental ownership, track cost, and assemble business reports, thus making inquiries, such as audits and inventory reviews, easier.

This chapter contains the following topics:

- “About Chargeback” on page 703
- “Setting Up Chargeback” on page 705
- “Accessing Chargeback” on page 705
- “Creating an Asset Record” on page 706
- “Changing the Status of an Element” on page 707
- “Saving Chargeback Information” on page 708
- “Viewing Assets” on page 708
- “Defining Storage Tiers This section contains the following topics:” on page 709
- “Adding Asset Information” on page 712
- “Managing Departments” on page 717
- “Setting the Infrastructure Cost” on page 718
- “Setting Up Asset and Storage Based Chargeback” on page 719
- “Viewing Chargeback” on page 737
- “Chargeback Reports” on page 741
- “Filtering Assets” on page 749

---

## About Chargeback

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 286 for more information.

First set up your chargeback as described in the topic, “Setting Up Chargeback” on page 705. When you are done with adding your chargeback information, you can view chargeback as follows:

- **By element** - Displays chargeback for a single element. See “Viewing Chargeback by Element” on page 738 for more information.
- **By department** - Displays chargeback for a department. See “Viewing Chargeback by Department” on page 739 for more information.
- **By owner** - Displays chargeback for an owner. See “Viewing Chargeback by Owner” on page 740 for more information.

Chargeback also provides the following reports. See “Viewing Chargeback Reports” on page 741 for more information.

- **Array Based Assets** - Displays the following 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 information 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 the following storage-based chargeback information 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 742 and “Adding an E-mail Schedule for a Chargeback Report” on page 743 for more information.

Chargeback helps you track 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 707 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 749.


# Setting Up Chargeback

To be able to view chargeback, you must first complete the steps in the following table.

TABLE 18-1    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 that the management server does not detect. Just create an asset record for each element.	“Creating an Asset Record” on page 706
2	Set the status of the asset.	“Changing the Status of an Element” on page 707
3	<i>Optional:</i> Add asset information for asset management.	“Adding Asset Information” on page 712
4	Add departments.	“Adding Departments” on page 717
5	Set up Chargeback.	<ul style="list-style-type: none"><li>• “Setting Up Asset-Based Chargeback” on page 720</li><li>• “Setting Up Storage-Based Chargeback” on page 726</li></ul>

# Accessing Chargeback

To access Chargeback, click **Chargeback** ().

---

# Creating an Asset Record

---


**Caution** – 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. Create an asset record for the application or hardware the management server does not monitor, and then follow the steps for setting up Chargeback, as described in the topic, “Setting Up Chargeback” on page 705.

After you create a record, the element for which you created the record 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.

---

**Caution** – You can easily remove an element's record by clicking the **Delete**  button. When you remove an element's record, the management server no longer monitors that element. See “Deleting Elements from the Product” on page 97 for more information.

---

To create a record:

1. Access Chargeback, as described in “Accessing Chargeback” on page 705.
2. Click **New**.
3. In the **Add Asset Record** window, enter 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.

---

**Caution** – Once you are done with changing the status of your elements, save your settings. See “Saving Chargeback Information” on page 708.

---

To change the status of an element:

1. Access Chargeback, as described in “Accessing Chargeback” on page 705.

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 In Use to 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 **Set Status**.
- c. From the Asset Status menu, select the new status for the elements you selected.
- d. Click **OK**.



**FIGURE 18-1** 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 705.

The following is displayed in the right pane:

- **Name** - The name of the element.
- **Status** - The status 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 707 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 705 for more information.
- **Vendor/Model** - The vendor and/or model of an element.
- **Type** - The type of element, such as an application. Table 18-2, “Element Type Icons,” on page 709 shows the icons for the element types.
- **Serial Number** - The serial number of the element.
- **Description** - Description of the element.

TABLE 18-2 Element Type Icons

Graphic	Element Type
	Application
	Host
	Switch
	Storage System

2. To remove an asset record, click the **Delete** () button corresponding to the record you want to remove.

## Defining Storage Tiers This section contains the following topics:

- “Creating a New Storage Tier” on page 710
- “Adding Elements to a Storage Tier” on page 710
- “Removing Elements from a Storage Tier” on page 711
- “Editing a Storage Tier” on page 712
- “Deleting a Storage Tier” on page 712

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 a maximum of 64).

Follow these steps to create a new storage tier:

1. Access Chargeback, as described in “Accessing Chargeback” on page 705.
2. Click the **Ownership** tab in the right pane.
3. Click the **Chargeback Method** menu, and select **Storage-based**.
4. Click **Set Storage Tier Cost**. The **Storage Tiers** window is displayed.
5. Click **Create New Storage Tier**.
6. Enter the necessary information in the Name, Monthly Cost/GB, and Description boxes.
7. Click **OK**. 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 705.
2. Click the **Ownership** tab in the right pane.
3. Click the **Chargeback Method** menu and select **Storage-based**.

4. Click **Set Storage Tier Cost**. The **Storage Tiers** window is displayed.
5. Click **Add/Remove Elements** for the storage tier to which you want to add elements. The **Add or Remove Storage Elements from Tier** window is displayed.
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** 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 **Filter**.
9. If you would like to see a list of the ports associated with a particular volume, select the storage volumes you want to add to the storage tier. Click the **+Ports** link in the **Ports** column.
10. Click the **Storage Pools** tab and select a storage system from the **Showing Pools for Storage System** menu.
11. If you would like to see a list of the volumes associated with a particular pool, select the storage pools you want to add to the storage tier. Click the **+Volumes** link in the **Volumes** column.
12. When you are finished selecting the storage systems, volumes, and pools you want to add to the storage tier, click **Add Selected Elements to Storage Tier**. The selected elements are added to the **Assets for Tier** section.
13. Click **OK**. 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 710.
2. In the **Assets for Tier** section, select the elements you would like to remove from the storage tier.
3. Click **Remove Selected Elements from Storage Tier**. The selected elements are removed from the storage tier.
4. Click **OK**. 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 710.
2. Click the **Edit** () button for the storage tier you want to edit.
3. Update the information for the storage tier.
4. Click **OK**. 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 710.
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 **OK**. You will be returned to the **Storage Tiers** window.

---

## Adding Asset Information

This section contains the following topics:

- “Adding Asset Information” on page 713
- “Adding General Information” on page 714
- “Adding Staff Information” on page 715
- “Adding Geographic Information” on page 716
- “Adding Licensing and Warranty Information” on page 716
- “Adding Custom Information” on page 716



# Adding Asset Information


Chargeback provides a handy way for you to keep track of your asset information for an element. In addition to warranty and licensing information, you can also store 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 find not only the contact information for that switch, but also the location of that switch.

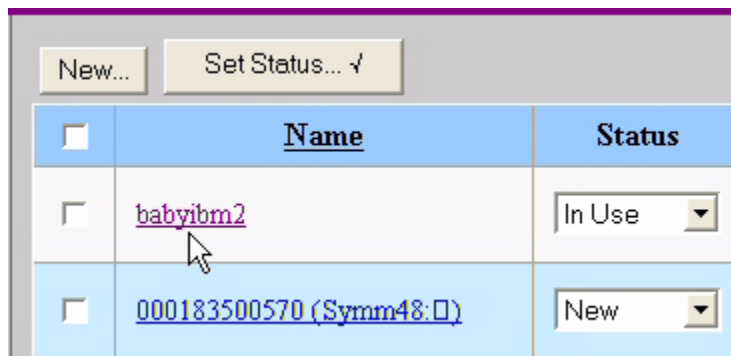
---

**Caution** – 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 286 for more information.

---

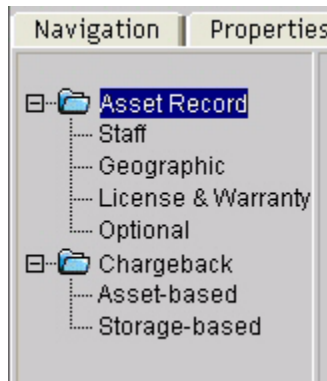
You can view and add asset information by doing the following:

1. Access Chargeback, as described in “Accessing Chargeback” on page 705.
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 18-2** Accessing an Element's Asset Information

3. To access the different types of asset information, click the Asset Record node or one of 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 18-3** 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 741 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 703 and “Setting Up Chargeback” on page 705.

---

**Note** – This page enforces the maximum number of characters you can enter in a box. When you can no longer add additional characters to a box, you have reached the maximum number of characters that can be entered for that box.

---

- **Custom Name** - A name you assign to the element. See “Assigning a Custom Name” on page 378 for more information.
- **Date Created** - The date the element was discovered.

- **Date Last Modified** - The 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 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 company that supplied the element.
- **Model** - The model of the element.
- **Serial Number** - The 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 This section contains the following topics:” on page 709 for more information.

## Adding Staff Information

This page provides contact information about the element.

---

**Caution** – 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 of the locations of all your elements, 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 street 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 box 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 box, 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 by following these steps:

1. In the Name box, assign a name for the box, for example, Backup Contact. The name cannot be more than 50 characters.
2. In the Value box, provide the information for the box, for example, Joe Smith. Do not enter more than 255 characters. For example, Joe Smith.
3. Repeat steps 1 and 2 for each custom property you want to add.

4. When you are done, click **Save Changes**.

---

## Managing Departments

This section contains the following topics:

- “Adding Departments” on page 717
- “Editing a Department” on page 717
- “Removing a Department from Chargeback” on page 718

### 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 705.
2. Click the **Departments** tab above the table.
3. Click **New**.
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 705.
2. Click the **Departments** tab above the table.
3. Click the **Edit**  button corresponding to the department you want to edit.


4. In the Edit Department window, you can edit all boxes except the department number.
5. Click **OK**.

## Removing a Department from Chargeback

Over time, some 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 assigned only to the department that is removed, it no longer has an owner. However, 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 owned by the Server department, but only by 50 percent.

To remove a department from the list:

1. Access Chargeback, as described in “Accessing Chargeback” on page 705.
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. Asset and storage-based infrastructure costs that you specify are calculated as part of your monthly chargeback results. See “About Asset-based and Storage-based Infrastructure Cost” on page 719 for more information. Follow these steps to set up infrastructure costs on the management server:

1. Access **Chargeback**, as described in “Accessing Chargeback” on page 705.
2. Click the **Ownership** tab in the right pane.
3. Click the **Chargeback Method** menu, and select either **Asset-based** or **Storage-based**.

4. Click **Set Infrastructure Cost**.
5. Enter the monthly infrastructure charge for calculating either asset-based or storage-based infrastructure costs, whichever you chose in step 3.
6. Click **OK**.

## About Asset-based and Storage-based Infrastructure Cost

Asset-based and storage-based infrastructure cost are optional chargeback calculations that you can specify on the Ownership page. Setting an asset-based or storage-based infrastructure cost calculates a monthly infrastructure charge, which is identical for each department and is applied once each month on top of the department's total ownership cost. Modifying the infrastructure charge impacts the asset-based or storage-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. Storage-based infrastructure cost 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 asset-based or 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 contains the following topics:

- "Setting Up Asset-Based Chargeback" on page 720
- "Setting Up Storage-Based Chargeback" on page 726
- "Editing Percentage of Ownership" on page 729
- "Removing Department Ownership of an Element" on page 730
- "How Capacity Differs in Chargeback and Capacity Explorer" on page 731
- "How a Depreciation Method Is Calculated" on page 732

# 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, you must perform these tasks:


1. Specify Financial information, as described in “Step 1 - Specify Financial information” on page 721.
2. Assign a Departmental Ownership Percentage, as described “Step 2 - Assign Departmental Ownership Percentage” on page 722.
3. Review the asset-based chargeback result, as described in “Step 3 - Review Asset-based Chargeback Result” on page 723.

---

**Caution** – You must have already added your departments, as described in the topic, “Editing a Department” on page 717.

---

To provide information for asset-based chargeback:

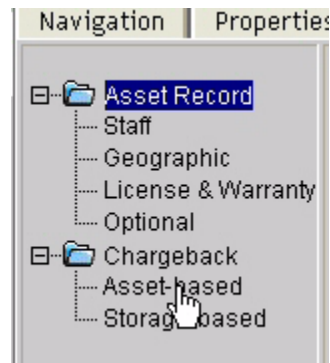
1. Access Chargeback, as described in “Accessing Chargeback” on page 705.
2. Do one of the following:
  - Click the  icon corresponding to the element.
  - Click the link for an element in the right pane, as shown in the following figure.



<input type="button" value="New..."/> <input type="button" value="Set Status... ✓"/>		
<input type="checkbox"/>	<u>N</u> ame	S
<input type="checkbox"/>	<a href="#">babyibm2</a>	In Use ▼
<input type="checkbox"/>	<a href="#">000183500570 (Symm48:□)</a>	New ▼

**FIGURE 18-4** Accessing an Element's Asset Information

- Click the **Asset-based** node under the Chargeback node, as shown in the following figure.



**FIGURE 18-5** 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 that 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.
- *Required:* **Date Purchased** - To update 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 through 3000.
- **Reseller** - The company that directly sold you the element.
- *Required:* **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.
- *Required:* **Depreciation Period (months)** - The time period in which you plan to keep the element.
- *Required:* **Depreciation Method** - How the depreciation is calculated. Select one of the following:

**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 732.

**Fixed Declining Balance** - This 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 733.

**Double Declining Balance** - This 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 735.

- **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 boxes:

Date Purchased

Original Cost (\$)

Depreciation Salvage Value (\$)

Depreciation Period

Depreciation Method

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 **Add Ownership**.
3. Select a department from the **Department** menu.
4. Enter the percentage of ownership in the Ownership % box. If you do not see a department listed, add it to the list as described in the topic, “Adding Departments” on page 717. 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 steps 1 through 5 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. For details, see “About Asset-based and Storage-based Infrastructure Cost” on page 719.
  - a. Click **Set Infrastructure Cost**.
  - b. Enter the monthly infrastructure charge.
  - c. Click **OK**.
8. When you are done with assigning the element to a department, click **Save Changes**.

## Step 3 - Review Asset-based Chargeback Result

---

**Caution** – 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.

---

**Caution –** 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 **cortez1**

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 18-6 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 perform these tasks:

- Assign Departmental Ownership Percentage, as described in “Step 1 - Assign Departmental Ownership Percentage” on page 727.
- Review Storage Tier Cost, as described in “Step 2 - Review Storage Tier Cost” on page 728.
- Review Storage Dependency and Cost, as described in “Step 3 - Review Storage Dependency and Cost” on page 729.
- Review storage-based chargeback result, as described in “Step 4 - Review Storage-Based Chargeback Result” on page 729.

Keep in mind the following:

- You must have already added your departments, as described in the topic, “Adding Departments” on page 717.
- 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 that your organization and roles allow you access. See “About Security for the Management Server” on page 175 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, “How Capacity Differs in Chargeback and Capacity Explorer” on page 731 for more information.

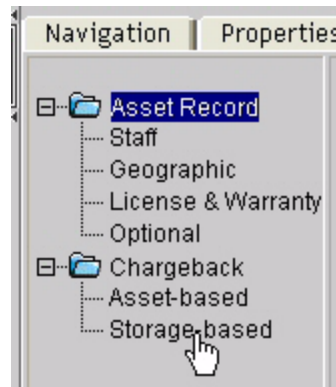
To provide information for storage-based Chargeback:

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.

<input type="button" value="New..."/> <input type="button" value="Set Status... ✓"/>		
<input type="checkbox"/>	<u>Name</u>	Status
<input type="checkbox"/>	<a href="#">babyibm2</a>	In Use ▼
<input type="checkbox"/>	<a href="#">000183500570 (Symm48:□)</a>	New ▼

**FIGURE 18-7** Accessing an Element's Asset Information

- Click the **Storage-based** node under the Chargeback node, as shown in the following figure.



**FIGURE 18-8** 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 **Add Ownership**.
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 717. 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 % box.
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 details, see “About Asset-based and Storage-based Infrastructure Cost” on page 719.
  - a. Click **Set Infrastructure Cost**.
  - b. Type the monthly infrastructure charge.
  - c. Click **OK**.
8. When you are done with assigning the element to a department, click **Save Changes**.

## 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.



Click the **Storage Tiers** link to review your storage tier costs. See “Defining Storage Tiers This section contains the following topics:” on page 709 for more information about setting up storage tiers.

## Step 3 - Review Storage Dependency and Cost

---

**Caution** – 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 **Storage/volume dependency and cost details** link to view the details.

The Storage Dependency for Application table is displayed if both 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 that 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 assigned tiers are contributing to the total cost.

## Step 4 - Review Storage-Based Chargeback Result

---

**Caution** – 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 that 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 % box, enter 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, for example, when 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 it from the list of departments as described in the topic, "Removing a Department from Chargeback" on page 718.

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 that 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.

# How Capacity Differs in Chargeback and Capacity Explorer

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. Notice that chargeback capacity is 0.89 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 18-9 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. When you 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 18-10 Capacity in Capacity Explorer

# How a Depreciation Method Is Calculated

This section contains the following topics:

- “Calculating Straight Line Depreciation” on page 732
- “Calculating Fixed Declining Balance” on page 733
- “Calculating Double Declining Balance” on page 735

## 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 rolled back to today.

*Example:* Assume the purchase date of an element is January 15, 2003. The management server adjusts the purchase 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 the purchase date and the period ending date. This determines how many month's 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 through 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 depreciate 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 the end of last month, the management server simply subtracts 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 at 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 purchase 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 rolled 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 needs 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) 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 depreciates 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 knows 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 the 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:

$$\text{Asset value of the month } (\$2395) - \text{Depreciation for the month } (\$100.59) = \$2294.41$$

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 rolled back to today.

Example: Assume the purchase date of an element is January 15, 2003. The management server adjusts the purchase 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 uses 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) in the following formula to find the rate of depreciation 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 knows 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

---


**Note** – If you see empty values, make sure chargeback has been set up, as described in “Setting Up Asset-Based Chargeback” on page 720 and “Setting Up Storage-Based Chargeback” on page 726.

---

This section contains the following topics:

- “Viewing Chargeback by Element” on page 738
- “Viewing Chargeback by Department” on page 739
- “Viewing Chargeback by Owner” on page 740

## Viewing Chargeback by Element


You can view chargeback for an element by clicking the  icon next to the element listed on the Asset tab.

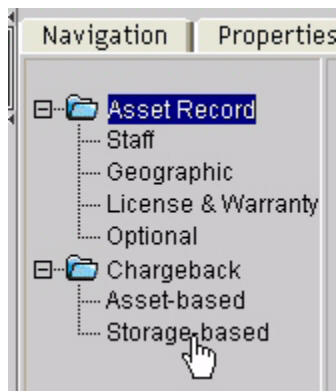
---

**Caution** – 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 705.
2. Click the  icon next corresponding to the element for which 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 18-11** Accessing Storage-Based Chargeback

Chargeback information for the element is displayed.

# 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.

---

**Caution** – 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 705.
2. Click the **Departments** tab in the right pane.
3. Click the  icon next corresponding to the element for which 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 specified amount of storage.
- **Total Storage-based Chargeback** - How much it costs the department for an application to use a specified 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:

- Asset-Based Chargeback lists the asset, depreciation, ownership percentage, and ownership cost.

- Storage-Based Chargeback lists the application, storage allotted, storage used, 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 to quickly identify 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 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.

---

**Note** – You can sort elements according to a column heading. Click a column heading in the table to sort the data. The arrow next to a column heading indicates 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 Name** - 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.
- **Application or Asset Name** - The name of the associated element or application.
- **Vendor** - The company that supplied the element.
- **Serial Number** - The serial number of the element.
- **Ownership Percentage** - The percentage of the element that the department owns.

- *Storage-based only:* **Storage (GB)** - The amount in gigabytes that the application uses. This value includes mounted and unmounted storage. Any volumes the application can access are included in the storage calculations.
- *Storage-based only:* **Storage Cost** - How much it costs to run the storage that the application uses.
- *Storage-based only:* **Partition Size** - The partition size used.
- **Ownership Cost** - How much it costs the department to use the asset.

---

## Chargeback Reports

This section contains the following topics:

- “Viewing Chargeback Reports” on page 741
- “E-mailing a Chargeback Report” on page 742
- “Managing E-mail Schedules for Chargeback Reports” on page 743

## Viewing Chargeback Reports

You can access chargeback reports in either 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 705.
- 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 259 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, and then expand the **Reporter > Chargeback** nodes in the tree in the middle pane.
2. Click one of the default reports:
  - **Array-Based Chargeback** - Displays the following 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 information 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 the following storage-based chargeback information 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 741 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, and then click **Run Report**.

## 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 743.

---

**Caution** – 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 221.

---

To e-mail a report:

1. Do one of the following:
  - To access the reports through Reporter, click the **Reporter** button, and 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 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 box, enter the recipient's e-mail address.

The software verifies that the address entered has a correct form. To send multiple addresses, use a comma (,) to separate addresses, for example:

john.example@domainname.com, jerry.example@domainname.com

5. From the **Format** 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 that the user has an understanding of XML.
6. *Optional:* Modify the subject and message.
7. Click **OK**.

The report is sent.

## Managing E-mail Schedules for Chargeback Reports

This section contains the following topics:

- “Adding an E-mail Schedule for a Chargeback Report” on page 743
- “Editing an E-mail Schedule for a Chargeback Report” on page 745
- “Deleting E-mail Schedules for a Chargeback Report” on page 747
- “Viewing E-mail Schedules for a Chargeback Report” on page 747
- “Viewing the History of an E-mail Chargeback Schedule” on page 748

### 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 221.
- 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, and 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 705.

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, and 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 741 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 **Add E-mail Schedule**.
6. In the To box, enter the recipient's e-mail address.

The software verifies that the address entered has a correct form. To send multiple addresses, use a comma (,) to separate addresses, for example:

john.example@domainname.com, jerry.example@domainname.com
7. In the Subject box, enter a subject for the e-mail messages you plan to send.

---

**Note** – Provide the name of the report in the subject box so users can distinguish this message from others.

---

8. In the Message box, enter 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** 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 that the user has an understanding of XML.
10. In the Time to Run box, enter 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 enter 14:15.



11. Select an option to determine how frequently you want to send the report. See “Setting the Frequency at which Reports are Sent” on page 745.
12. Click **OK**.

The schedule is created.

### *Setting the Frequency at which Reports are Sent*

When adding or editing a report, there are several options for how frequently the report is sent. The following options are available:

- **Daily** - If you select daily, select how frequently you want the management server to send the report.

**Everyday** - The report is sent every day.

**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 a person’s replacement while the person is away on vacation.

- **Weekly** - If you selected weekly, use the **Frequency** 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, and then select **First** or **Last** from the menu.

To send the report on a specified day during the month, select the second option, and then enter the day on which you want the report sent. Keep in mind that the number of days in a month varies. So if you enter 30 in this box, users will not receive a report in February. Also, if you enter 29 in this box, users do not receive the report in February during non-leap years.

## Editing an E-mail Schedule for a Chargeback Report

---


**Caution** – 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, and 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 705.

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, and 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 741," 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 box, change the recipient's e-mail address.

The software verifies that the address entered has a correct form. To send multiple addresses, use a comma (,) to separate addresses, for example:

john.example@domainname.com, jerry.example@domainname.com
7. In the Subject box, change the subject of the e-mail.
8. In the Message box, change the message describing the report.
9. From the **Format** 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 that the user has an understanding of XML.
10. In the Time to Run box, enter 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 enter 14:15.
11. Select an option to determine how frequently you want to send the report. See "Setting the Frequency at which Reports are Sent" on page 745.
12. Click **OK**.

## Deleting E-mail Schedules for a Chargeback Report


---

**Caution** – Only the e-mail schedules created by the current user are listed. To view the e-mail schedules for all reports, select **Configuration > Reports** in the upper-right corner of the screen, and 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 705.
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, and 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 741 for more information about each of the reports.
4. Click the **Scheduled Deliveries** tab.
5. Click the **Delete** () button corresponding to the e-mail schedule you want to remove.

## Viewing E-mail Schedules for a Chargeback Report

---

**Caution** – Only the e-mail schedules created by the current user are listed. To view the e-mail schedules for all reports, select **Configuration > Reports** in the upper-right corner of the screen, and 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 705.
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, and 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, as described in Table 18-3, “Viewing E-mail Schedules for a Chargeback Report,” on page 748.

**TABLE 18-3** 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: <ul style="list-style-type: none"> <li>• PDF</li> <li>• Microsoft EXCEL</li> <li>• XML</li> </ul>
Last Delivered	The time the last report was sent to the recipient.
History	Click the <b>View</b> button to display the times and dates when the report was sent. You can also delete a historical entry by clicking the <b>Delete</b>  button for the corresponding entry.
Edit	Click the <b>Edit</b>  button to edit a schedule of the report. See the topic, “Adding an E-mail Schedule for a Chargeback Report” on page 743 for information about the options displayed in this window.
Delete	Click the <b>Delete</b>  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 **Delete**  button for the corresponding entry.

To view the history of an e-mail schedule:

1. Access Chargeback, as described in “Accessing Chargeback” on page 705.
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, and 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 **View**.  
You are shown when the report was sent.
  6. To remove a historical entry, click the **Delete** () button.
- 

## Filtering Assets

This section contains the following topics:

- “About Filtering Assets” on page 749
- “Selecting an Element Type for Chargeback” on page 750
- “Customizing the Element Type Filter for Chargeback” on page 750
- “Filtering Assets by Status” on page 751
- “Customizing the Asset Status Filter for Chargeback” on page 751
- “Hiding Filters in Chargeback” on page 752

## 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.

As another example, you could use the filters to find out which elements are missing or repaired by doing 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, click **Apply Filters**.

## 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** menu in Chargeback. When you are asked if you want to apply your changes, if you want to apply them now, click **Yes**. If you want to apply them at a later time, click **No**, then click **Apply Filters** when you are ready for your changes to take effect. Chargeback displays only the elements you specified in your filter.

**TABLE 18-4** 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** 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** 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, if you want to apply them now, click **Yes**. If you want to apply them at a later time, click **No**, then click **Apply Filters** when you are ready for your changes to take effect. Chargeback displays only the elements you specified in your 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 an option from the **Show Element Type** 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, if you want to apply them now, click **Yes**. If you want to apply them at a later time, click **No**, then click **Apply Filters** when you are ready for your changes to take effect. Chargeback displays only the elements 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** menu in Chargeback.

For example, you can specify you want only assets that are missing displayed in Chargeback.

To select more than one asset for filtering:

1. Click the **Custom** button next to the **Show Status** 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, if you want to apply them now, click **Yes**. If you want to apply them at a later time, click **No**, then click **Apply Filters** when you are ready for your changes to take effect. Chargeback displays only the elements you specified in your 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 **-Filters** link in the upper-left corner of Chargeback.

To display the filters, click the **+Filters** link in the upper-left corner of Chargeback.



## Managing Backups

---

Depending on your license, the Protection Explorer feature may not be available. See the List of Features to determine if you have access to Protection Explorer. The List of Features is accessible from the Documentation Center (**Help > Documentation Center**).

This chapter contains the following topics:

- “About Protection Explorer” on page 753
- “Viewing Sessions that are Running” on page 757
- “Determining if the Last Scheduled Backup was Successful” on page 758
- “About the User Interface” on page 761
- “About the Summary Backup Charts” on page 772
- “About the Tabs in the Topology Lower Pane” on page 773
- “Modifying Summary Backup Charts” on page 777
- “Viewing Charts for a Backup Manager Host” on page 778
- “Printing Summary Charts To print a summary chart:” on page 779

---

## About Protection Explorer

The Protection Explorer feature allows you to:

- Monitor the overall status of the backup process
- Visualize the backup configuration and recoverability of a file, directory, volume, or server
- View the status of the physical infrastructure supporting the backup process, backup application, backup server, network, tape library, and media
- Provide information on reasons for backup failures and advisory information for configuring new backup schedules

Protection Explorer monitors the backup applications running on discovered hosts. To determine which backup applications are supported, see the support matrix, which is accessible from the Documentation Center.

The management server is able to detect the presence of the following after you obtain backup details:

- **Backup Application** — A backup application, such as NetBackup or HP Data Protector, serving as the master in a backup hierarchy. A backup application is responsible for managing other media managers.
- **Backup Manager Host** — A managed host that is running the backup application. The IP address of the backup manager host must be specified in Step 1 of discovery before the backup application can be discovered.
- **Media Manager Application** — A backup application functioning as a server to control the media in a backup hierarchy. A media manager application can be responsible for managing different types of hardware, such as tapes and drives.
- **Media Manager Host** — A host that has the backup application running as the media manager application. A media manager application and its host can be discovered through the backup manager host (similar to the way that hosts are detected through a switch). If a media manager host is discovered through a backup manager host, the media manager host is considered to be “unmanaged,” meaning that the management server has discovered it, but cannot obtain additional information about the element. If the IP address of the media manager host is specified in Step 1 of discovery, the media manager host will be considered “managed.”
- **Media** — Any device that is used to store backup data, such as tape.
- **Media Pool** — A logical grouping of the backup media
- **Sessions** — Scheduled and executed backup sessions
- **Tape Library** — A device hosting a collection of tapes
- **Robot** — An automated device inside the tape library; responsible for manipulating the tapes
- **Backup Client** — A host that is being backed up by a backup application. A backup client can be managed as a non-generic element if its IP address appears in the discovery list. Otherwise, backup clients that are identified through the backup application are considered generic elements.

---

**Note** – In the Protection Explorer feature, only a single click on an element in the topology is required to obtain more information about the element.

---

## Requirements for Using Protection Explorer

---

**Caution** – The CIM extension supports only one backup solution on a host. For this reason, Protection Explorer does not support both NetBackup Master Server and HP Data Protector Cell Manager on the same host. If NetBackup Master Server and Data

Protector Cell Manager are installed on the same host, by default only Data Protector Cell Manager is discovered. NetBackup Master Server is ignored by the CIM extension.

---

Before you can use the Protection Explorer feature, you must take the following actions:

1. Install a CIM extension on the host running the backup application. See the installation guide for information about installing CIM extensions.
2. Discover the host running the backup application, selecting the **Include backup details** option. See “Step 1 — Discovering Your Hosts and Backup Manager Hosts” on page 105. It is recommended that you also select the **Include infrastructure detail** option, so you can also monitor and manage the host itself.

---

**Caution** – Make sure you have at least 500 MB available if you are using the host as a backup manager host in a large environment (for example, 300 clients, 25,000 sessions, and 500,000 images).

---


3. Schedule backup collection for your backup manager hosts as described in “Scheduling Backup Collection for Backup Managers” on page 245

## Determining if You Have Enough Media to Run a Backup

If you are performing many and/or large backups, you should make sure you have enough media available for the backup. Protection Explorer provides several methods for determining if you have enough media:

- **Media tab** — Provides information about the discovered media, including its usage count.
- **Media Pool tab** — Provides information about media in the pools, such as whether it is Available, Allocated, Frozen, or Suspended.
- **Media Summary reports** — Provides information about all discovered media over a defined time period.

To use these methods:

1. Access Protection Explorer by clicking **Protection Explorer** ()
2. Click the **Topology** tab on the right side of the window.
3. Expand the Backup Applications node in the left pane, and then select a backup manager host.

4. If necessary, expand the lower pane so you can view the Media and Media Pool tabs.

5. Click the **Media** tab in the lower pane.

The following information is displayed:

- **Media ID** — The identification number of the media
- **Media Pool** — The name of the media pool
- **Usage Count** — How often the media has been used
- **Barcode** — The barcode associated with the media
- **Retention** — How long the media is retained
- **State** — Whether the media is Full, Available, or Active

To learn more about a specific media, select its row. Additional information is displayed in the lower-right pane.

6. To learn more about the media pool that contains the media, click the **Media Pool** tab. The following information is displayed:

- **Media Pool** — The name of the media pool
- **Backup Manager** — The name of the backup manager host to which the media pool belongs
- **Media Manager** — The name of the media manager to which the media pool belongs
- **Library** — The name of the library to which the media pool belongs
- **Available** — The media is available for backup.
- **Allocated** — The media is currently either actively being used or has a valid backup on it.
- **Frozen** — The media will never become available for backup again, but it is still available for restores.
- **Suspended** — The media will not be used again until all backups written to it expire. It is still available for installations however.

7. To view summary media information for selected media in a report format:

- a. Access Reporter by clicking **Reporter** ()
- b. Click **Reporter > Protection Explorer** to expand the tree.
- c. Click the Media Summary report.

Protection Explorer displays the total media, designating each as either Available, Allocated, Frozen, or Suspended

You can also set up a policy that will notify you when the number of available media for a storage pool is running low. For example, you could create a policy that sends you an e-mail when the number of available media for a storage pool is less than two. See “Creating a Utilization or Backup Policy” on page 684.

# Accessing Protection Explorer

To access Protection Explorer, click **Protection Explorer** ()

---


## Viewing Sessions that are Running

---

**Note** – For information about configuring session monitoring, see “Session Monitoring” on page 246.

---

To view sessions that are running:

1. Access Protection Explorer by clicking **Protection Explorer** ()
2. Click the **Topology** tab on the right side of the window.
3. Right-click a backup manager application, and then select **Show Running Sessions**. The Running Sessions page is displayed.
4. If desired, select filter criteria, and then click **Apply Filter**. The table is updated to display only the sessions that meet the filter criteria you entered.

The following information about each session is displayed:

- **Job ID** — The identifier assigned to the session
- **Backup Manager** — The name of the backup manager
- **Media Manager** — The name of the media manager
- **Clients** — The names of the clients
- **Backup Policies** — The names of the backup policies
- **Schedules** — The names of the schedules
- **Session Status** — The session status: Success or Failure
- **Session State** — The session state: Done, Queued, or Active
- **Media Used** — The type of media used for the backup
- **Start Time** — The starting time and date of the session
- **End Time** — The end time and date of the session
- **Size** — The size of the session in kilobytes (KB)
- **Files** — The number of files.

---

**Note** – After removing Caché and installing Cim extensions, the first session monitoring collection will collect all of the sessions for the past week. One hour after a fresh installation of Caché, these sessions will be removed from the Running Sessions page.

---

---

**Note** – The Running Sessions page is automatically refreshed every two minutes. To manually refresh the page, click **Getting Latest Sessions**.

---

---


## Determining if the Last Scheduled Backup was Successful

Protection Explorer provides several tools to help you determine if the last scheduled backup was successful. The quickest way to do this is from the Summary tab in Protection Explorer. See the following topics to learn more:

- “Viewing the Summary Backup Charts” on page 758
- “Viewing Backup Results for a Backup Manager Host” on page 759
- “Viewing Backup Results for a Client” on page 759
- “Viewing Backup Information for a Client” on page 760
- “Viewing Backup Reports” on page 760

## Viewing the Summary Backup Charts


To access summary information about last night’s backup from the backup charts:

1. Access Protection Explorer by clicking **Protection Explorer** ().
2. Click the **Summary** tab on the right side of the window.

By default, the Backup SLA Performance chart is displayed in the upper-left pane on the Summary page. This chart includes the overall results of the backups made in the last 14 days. It tells you if the overall backup was successful, partially successful, or failed. See “About the Summary Backup Charts” on page 772 and “Modifying Summary Backup Charts” on page 777 for more information about this chart and about modifying it.

# Viewing Backup Results for a Backup Manager Host

To quickly view the results of the backup sessions that a backup application has performed:


1. Access Protection Explorer by clicking **Protection Explorer** ()
2. Click the **Topology** tab on the right side of the window.
3. In the left pane, expand the Backup Applications node, and then select the backup manager host.
4. If necessary, expand the lower pane so you can see the tabs.
5. Click the Sessions tab.

Protection Explorer displays the following information for each session the backup application performed:

- **Session ID** — The identifier assigned to the session
  - **Client** — The name of the client
  - **Backup Policy** — The name of the backup policy
  - **Schedule** — The name of the schedule
  - **Status** — The status of the backup: Success, Partial, or Failure
  - **Start Time** — The starting time and date of the backup
  - **End Time** — The end time and date of the backup
  - **Duration** — The amount of time in seconds it took for the backup to be displayed
  - **Size** — The size of the backup in gigabytes (GB)
  - **Files** — The number of files backed up
6. To learn more about a session, select the session's row in the table and see the Session Detail, Policy Detail, and Schedule Detail tabs in the lower-right pane for more information.

# Viewing Backup Results for a Client

To view the results of the last backup for a client:


1. Access Protection Explorer by clicking **Protection Explorer** ()
2. Click the **Topology** tab on the right side of the window.
3. In the left pane, expand the Clients node, and then select the client.

4. In the topology pane, check the color of the check mark above the icon for the client. If the icon is:
  - **Green** — The last backup on the client was successful.
  - **Yellow** — The backup on the client was partially successful.
  - **Red** — The backup on the client failed.
5. To view detailed backup information for a client, see “Viewing Backup Information for a Client” on page 760.

## Viewing Backup Information for a Client

Protection Explorer tracks of backup information for a client for the last 30 days. The Protection tab gives an at-a-glance view of the backup coverage for a selected element. For example, you can select the period for the coverage and review the policy, schedule, and results for the backups executed for that period.

To obtain detailed information about the backup sessions for a client:

1. Access Protection Explorer by clicking **Protection Explorer** ().
2. Click the **Topology** tab on the right side of the window.
3. Select the client from the topology pane or from the tree in the left pane.
4. If necessary, expand the lower pane so you can view the Properties and Protection tabs.
5. Click the **Protection** tab to view the following information for each backup policy:
  - **Backup Policy** — The name of the backup policy
  - **Schedule** — The name of the schedule
  - **Date** — The date, end time, and status
6. To learn more about a policy or schedule, select a cell in the table. Additional information is provided on the Policy Detail and Schedule Detail tabs in the lower-right pane.

## Viewing Backup Reports


One way to quickly inform others of the results of last night’s backups is to send them the reports that provide information about the backup. These reports can be sent manually or you can configure Reporter to e-mail these reports automatically to a list of user accounts. The following reports are provided:

- Backup Volume by Clients



- Backup Volume by Policy
- Backup Volume Overview
- Daily Backup Job
- Daily Failed or Partial Backup Job
- Daily Restore Job
- Hosts at Risk
- Image Data Volume by Clients
- Image Data Volume by Policy
- Library Utilization
- Media Availability
- Media Properties
- Media Summary
- Media Used
- Most Common Job Failure Reasons
- Most Frequently Failing Hosts
- Restore SLA Summary
- SLA Summary

To access the backup reports:

1. Access Reporter by clicking **Reporter** (.
2. Expand the tree in Reporter.
3. Expand the Protection Explorer node.
4. Click one of the reports in the Reporter tree.
5. To e-mail the report, click **E-mail Report**.

---

**Note** – Before you can e-mail a report, you must make the management server aware of the e-mail server. See “Setting Up E-mail Notification” on page 221.

---

To learn more about scheduling the e-mailing of reports, see “Adding an E-mail Schedule for a Report” on page 531.

## About the User Interface

Protection Explorer gives has an easy-to-use interface that provides the following options:

- **Toolbar** — Provides buttons and menus to help you modify the topology and charts in Protection Explorer. See “About the Toolbars in Protection Explorer” on page 764.


- **Summary and Topology tabs** — Provide information about individual elements. The following tabs are available:
  - **Summary tab** — The summary charts for backup elements are displayed. See “About the Summary Backup Charts” on page 772 and “Modifying Summary Backup Charts” on page 777.
  - **Topology tab** — The topology of the backup elements is displayed.
- **Lower pane on the Topology tab** — The lower pane on the Topology tab is displayed when you select a discovered element. The tabs displayed are determined by the backup element you selected. See “About the Tabs in the Topology Lower Pane” on page 773.
- **Access to the Navigation, Events, Collectors, Policies, and Chargeback** — When you click an element on the Topology tab, the following links in the lower-right corner are enabled if that feature is supported for the selected element:
  - **Navigation** — Displays the navigation information for an element, such as which storage systems are connected to the element. See “About the Navigation Tab” on page 370.
  - **Events** — Displays the events for the element. See “About the Events Tab” on page 392.
  - **Collectors** — Provides links to data collectors and reports about the element. See “About the Collectors Tab” on page 395.
  - **Policies** — Allows you to set up policies for the element. See “About the Policies Tab” on page 396.
  - **Chargeback** — Allows you to provide chargeback information about the element. See “Asset Attributes of an Element” on page 393.

You can also select an element and use the right-click menu options to obtain additional information. See “Right-Click Menu Options on the Topology Tab” on page 768.








## About the Topology Icons in Protection Explorer

Table 19-1, “Topology Icons in Protection Explorer,” on page 762 describes the icons that appear in the topology in Protection Explorer.

**TABLE 19-1** Topology Icons in Protection Explorer

Icon	Description
	When a client computers is shown with a green check mark, the backup on the computer was successful.

**TABLE 19-1** Topology Icons in Protection Explorer (*Continued*)

Icon	Description
	When a client computers is shown with a yellow question mark, the backup on the computer was partial.
	When a client computers is shown with a red X, the backup on the computer failed.
	Host
	Master backup server (media). This image is an example of a master backup server for NetBackup.
	Backup server (media), such as NetBackup
	Tape Library
	Tape Drive

# About the Toolbars in Protection Explorer







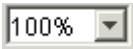



Protection Explorer has two toolbars.

- The main toolbar that appears at the top of its screen.
- The toolbar for charts that appears in the middle of its screen.














## About the Main Toolbar

Table 19-2, “The Toolbar in Protection Explorer,” on page 764 provides a brief description of the buttons and menus on the main toolbar in Protection Explorer. This toolbar is available at the top of the Protection Explorer screen.

**TABLE 19-2** The Toolbar in Protection Explorer

Button	Description
	Saves the current topology or summary page, so that when you return to Protection Explorer, the saved layout or summary is restored. This option can be especially useful if you want to keep the new location of elements you have moved. When you click the button, you are asked if you want the layout to apply to all users. <ul style="list-style-type: none"><li>• <b>Yes</b> — All users who log into the management server can view the topology or summary you created. Only users with system configuration capability can save their layout for all other users</li><li>• <b>No</b> — No other users can view the topology or summary you saved.</li></ul>
	Prints the topology or the summary.
	Allows you to view information from past backups. For example, if you want to view a backup from last March, just click the Calendar icon, then select the date.
	Allows you to modify the summary charts. Enabled when the Summary tab is active.
	Magnifies the view. Enabled when the Topology tab is active.
	Decreases the magnification. Enabled when the Topology tab is active.
	Allows you to set the magnification to a percentage of the default magnification. Enabled when the Topology tab is active.
	Opens a smaller pane, which provides a global view of the topology. This allows you to position the main view to a certain section of the topology. For more information, see “Using the Global View” on page 337. Enabled when the Topology tab is active.
	Fits the topology to the window, so you can view the entire topology. Enabled when the Topology tab is active.
	Allows you to move an element in the topology. See “Arranging Elements in the Topology” on page 335. Enabled when the Topology tab is active.



**TABLE 19-2** The Toolbar in Protection Explorer (*Continued*)

Button	Description
	Allows you to move the entire topology at once. Click the <b>Pan</b> (  ) button, click any place in the topology, and then drag the mouse to a new location. Enabled when the Topology tab is active.
	Opens a new window, containing the topology. This feature allows you to view different domains of the topology at one time. Enabled when the Topology tab is active. See “About the New Window Option” on page 387 for more information.
	Allows you to change the topology layout. See “Changing the Topology Settings” on page 766. Enabled when the Topology tab is active.
	Restores the topology layout to the last saved version. Enabled when the Topology tab is active.
Find <input data-bbox="264 618 422 661" type="text"/>	Allows you to find an element by name or by Worldwide Name (WWN) in the topology. Enabled when the Topology tab is active. To find an element, enter the name or part of the name in the Find box, and then click the <b>Find Next</b> (  ) button. The management server highlights the elements that match in the topology and in the tree. If the management server has found multiple elements matching your search criteria, click the <b>Find Next</b> (  ) button to find the next element that matches your search criteria. To view the previous element that matches the search criteria, click the <b>Find Previous</b> (  ) button. The <b>Find Previous</b> (  ) button is disabled when only one element meets your search criteria.
	Exports the topology to an XML file that can be viewed in Microsoft Visio. See “Exporting the Topology to Visio” on page 767 for more information. Enabled when the Topology tab is active.
	Displays links between shared libraries. Additional connections between media servers and tape libraries, and media servers and disk drives are displayed. If the additional links between shared libraries are currently displayed, clicking the Show MultiPath button a second time will hide the links.
	Change Observer button - Monitors changes in the database status on the server. When changes are detected, the button turns amber. Click on the amber button and a pop-up window displays the elements that have changed on the server. When no changes are detected, the button is greyed out
	Reloads the Change Observer button to display the latest changes to elements on the server.


## About the Toolbar for Charts

The toolbar options described in Table 19-3, “Toolbar for Charts,” on page 766 are only available when you are viewing a chart on the Topology tab. To view a chart on the Topology tab, click an element in the topology.

**TABLE 19-3** Toolbar for Charts

Option	Description
	Converts the data in the chart to a list in a separate browser.
	Click to print a chart. See “Printing Summary Charts To print a summary chart:” on page 779 for more information.
<input type="checkbox"/> Invert Chart	Click to switch the X and Y axes in a chart.
Title <input type="text" value="Backup Volume"/>	To change the chart displayed, select another chart from the Title menu. See “About the Summary Backup Charts” on page 772 for more information.
Period <div> <div>Last 7 days</div> <div>Last 7 days</div> <div>Last 14 days</div> <div>Last 30 days</div> </div>	To change the period displayed in the chart, select a period from the Period menu.
Average Service Level <input type="text" value="85%"/>	In Service Level Agreement charts, Protection Explorer provides a green line that serves as a baseline. Use the Average Service Level menu to change the location of this baseline. The default baseline value is 95% of usage.

## Changing the Topology Settings

The **Display Layout Settings Dialog** () button allows you to modify the following properties of the topology in Protection Explorer:

- **Direction** — Horizontal or Vertical. The direction of the topology is Horizontal by default, with multiple elements of the same type displayed in a row. If you select Vertical, multiple elements of the same type are displayed in a column.
- **Alignment** — Left, Right, or Center. The default alignment of the topology is Center. You can change the alignment of the topology to be left- or right-justified. For example, if you select the Left, the backup clients are aligned along the left side of the topology window.



- **Horizontal Spacing** — The number of spaces in pixels between elements in a row
- **Vertical Spacing** — The number of spaces in pixels between elements in a column

---

**Note** – To restore the layout to the default settings, click **Defaults**.

---

To change the layout settings:

1. Access Protection Explorer by clicking **Protection Explorer** ()
2. Click the **Display Layout Settings Dialog** () button.
3. Select one of the following directions:
  - Horizontal
  - Vertical
4. Select one of the following alignments:
  - Left
  - Right
  - Center
5. To change the horizontal spacing, enter a new number in the Horizontal Spacing box.
6. To change the vertical spacing, enter a new number in the Vertical Spacing box.
7. Click **OK**.

---


**Note** – You may need to use the scroll buttons to see the rearranged topology.

---

## 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 **Protection Explorer** ()
2. Click **Export to Visio**.
3. Name the file, and then select the directory in which you want the file to be saved.

4. Click **Save**. The XML file is saved to the directory that you selected. For information about configuring Visio and viewing the exported file, see “Viewing the Topology in Microsoft Visio” on page 340.

## Right-Click Menu Options on the Topology Tab

When you right-click an element on the Topology tab or in the Backup Applications tree, a list of options is displayed. The options displayed in the menu depend on the type of element you selected.

Table 19-4, “Right-Click Menu Options on the Topology Tab,” on page 768 provides an explanation of the menu options displayed when an element is right-clicked on the Topology tab or in the Backup Applications tree.

**TABLE 19-4** Right-Click Menu Options on the Topology Tab

Right-Click Menu Option	Description
Go To Navigation Details	Directs you to the Navigation page. If the element is labeled Discovered, you are shown the Properties page. An element is labeled unmanaged when the management server has become aware of it, but cannot obtain additional information about it. See “About the Navigation Tab” on page 370.
Show Events	Events for the selected element are displayed. See “About the Events Tab” on page 392.
Show Policies	The backup policies for the selected element are displayed. See “About the Policies Tab” on page 396.
Show Collectors	The report collectors for the selected element are displayed. See “About the Collectors Tab” on page 395.
Show Chargeback	The chargebacks for the selected element are displayed. See “About the Monitoring Tab” on page 396.



**TABLE 19-4** Right-Click Menu Options on the Topology Tab (*Continued*)

Right-Click Menu Option	Description
Update Element Data	<p>The management server gathers new and changed details from the element and then redraws the topology with the updated information.</p> <p><b>Important:</b></p> <ul style="list-style-type: none"> <li>• Do not update element data during Get Topology or Get Details. To determine if the management server is getting the topology or all element details, look at the label near the Status button.</li> <li>• 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 then select <b>Update Element Data</b>, the deleted LUNs still appear in the user interface. You must perform Get Details for the deleted LUNs to be removed from the user interface.</li> </ul> <p>For more information, see “Get Details” on page 91.</p>
External Tools	<p>Provides several ways to access an element:</p> <ul style="list-style-type: none"> <li>• <b>Telnet</b> — Allows you to access a host or a switch through the telnet utility. The Telnet feature is only accessible to Web browsers on Microsoft Windows operating systems.</li> <li>• <b>Browse</b> — Allows you to access the main Web page for a host or a switch.</li> <li>• <b>Set Up External Tools</b> — Allows you to add URLs for accessing the management tools for the storage system. In some instances, the management tool for the storage system is directly accessible from this menu (for example, HiCommand for HDS storage systems and Command View for HP XP storage systems).</li> </ul> <p>See “Using External Tools” on page 369</p>
Add Virtual Application	<p>Allows you to add an unsupported application so you can monitor it. For example, you might want to add a virtual application so you can monitor software that was created uniquely for your company.</p> <p>See “Creating a Virtual Application” on page 390.</p>
Reachable Storage	<p>Provides information about the storage accessible from a selected host, such as:</p> <ul style="list-style-type: none"> <li>• Free volumes on current storage systems</li> <li>• LUNs mapped to a host, but not mounted with a file system</li> <li>• Free volumes on other storage systems in host fabrics</li> <li>• Free volumes on all other storage systems</li> </ul>
Custom Commands	<p>Allows you to run a custom command on an element; for example to start an executable or a script. See “Setting Up Custom Commands” on page 358.</p>

**TABLE 19-4** Right-Click Menu Options on the Topology Tab (*Continued*)

Right-Click Menu Option	Description
Go to System Explorer	Provides a topology that allows you to view how the devices in your network are connected. See “About System Explorer” on page 307.
Show MultiPath/Remove MultiPath	Displays links between shared libraries. Additional links between media servers and tape libraries, and media servers and disk drives are displayed. If the additional links between shared libraries are currently displayed, the menu option becomes Remove MultiPath. Selecting Remove MultiPath hides the links between shared libraries.

The charts in Protection Explorer provide a wealth of information about your backups. You can obtain detailed information about a data point displayed in a chart by right-clicking the data point. For example, assume you are looking at a Service Level Agreement (SLA) chart on the Summary tab and you want to obtain more information about a backup performed yesterday. You could right-click the bar **yesterday’s date**, and then select **Show Details** to display the Sessions tab showing the additional details of that data point, such as the backup status of each client, in addition to the start and end time of the backup on that client.

**TABLE 19-5** Right-Click Menu Options on the Summary Tab

Right-Click Menu Option	Description
Summary Page Settings	The settings for the charts provided on the Summary tab are displayed. See “About the Summary Backup Charts” on page 772 and “Modifying Summary Backup Charts” on page 777.
Go To Topology	A graphical representation of the path of an element is displayed. This also includes multipathing. See “Viewing Element Topology” on page 379.
Show Details	The additional information about the data point you right-clicked is displayed. See Table 19-6, “To Obtain Additional Information from a Chart on the Summary Tab,” on page 771.

Table 19-6, “To Obtain Additional Information from a Chart on the Summary Tab,” on page 771 explains what is displayed when you click **Show Details** on the Summary tab’s right-click menu option.

---

**Note** – When additional information is not available for a data point, Show Details is disabled.

---

**TABLE 19-6** To Obtain Additional Information from a Chart on the Summary Tab

When you right click a bar and select Show Details in the following chart...	The Sessions tab displays the following information:
Service Level Agreement	Clients that were backed up, whether successfully or not. The failures are displayed first. To obtain details about a session, select the session in the Session tab and then expand the View the Details pane on the far right. The status you right-clicked is highlighted in the Sessions tab.
Backup Volume	Clients in the backup, sorted by size.
Windows Utilization Chart	Clients sorted by duration of the backup. The time span represented by the bar you clicked is highlighted in the Sessions tab. For example, assume a bar in the Windows utilization tab shows a duration of seven hours. To determine which sessions were running during that time, right-click the bar, and then select <b>Show Details</b> .
Largest/Longest Sessions	Clients in the backup, sorted by duration of the session.

You can also obtain additional information from some of the charts that are displayed on the bottom pane of the Topology tab. See Table 19-7, “Show Details for Tabs on the Lower Pane of the Topology Tab,” on page 771.

**TABLE 19-7** Show Details for Tabs on the Lower Pane of the Topology Tab

Right-click...	Select Show Details to view...
Any data point on the Charts tab	The Sessions tab for the data point you right-clicked.
Any element on the Servers tab	The Sessions tab showing the sessions for the element you right-clicked.
Any element on the Resources tab	The Media tab for the element you right-clicked.
Any element on the Media Pools tab	The Media tab for the elements contained in the media pool you right-clicked.

---

# About the Summary Backup Charts

Protection Explorer displays six summary backup charts on the Summary tab by default and offers many other charts as well. To learn how to display the various charts and/or modify which charts display by default, see “Modifying Summary Backup Charts” on page 777.

**TABLE 19-8** Protection Explorer Summary Charts

View	Description
Servers	Displays the servers Protection Explorer monitors with the following information for each server: <ul style="list-style-type: none"><li>• Volume</li><li>• Sessions</li><li>• Failed</li><li>• Partial</li><li>• Successful</li></ul>
Resources	Displays the resources Protection Explorer monitors, showing the following information for each server: <ul style="list-style-type: none"><li>• Media Pools</li><li>• Available media</li><li>• Allocated</li><li>• Frozen</li><li>• Suspended</li></ul> <p>In the Available Media, Allocated, Frozen, and Suspended columns, the first number shows the number of available online media, and the second number shows the number of available offline media. Note that this information is available only to the backup manager host.</p>
Service Level Agreements (SLAs)	Displays the performance of backup SLAs, showing the percentage of the following types of sessions for each SLA: <ul style="list-style-type: none"><li>• Successful sessions</li><li>• Partial sessions</li><li>• Failed sessions</li></ul>
Backup Volume	Displays the backup volume of all backup applications in gigabytes (GB). This chart can also display the backup volume of a backup manager host.

**TABLE 19-8** Protection Explorer Summary Charts (*Continued*)

View	Description
Window Utilization	Displays the number of hours it takes for all backup sessions on a server to run. Keep in mind this time may seem extended if you have overlapping sessions. For example, bsession1 starts at 11 p.m. on Monday. While it is running, bsession2 starts. At 2 a.m. on Tuesday, bsession1 stops, but bsession2 continues to run until 9 a.m. on Tuesday. The Windows Utilization report shows the backup sessions running for 10 hours - from the beginning of the bsession1 to the end of bsession2.
Backup Manager Hosts with Most Executed Sessions	Displays the five backup manager hosts with the most executed sessions. Only successful sessions are counted.
Most Unsuccessful Backup Manager Hosts	Displays the five backup manager hosts with the most unsuccessful sessions. Unsuccessful sessions include failed and partially completed sessions.
Servers with Most Available Media	Displays the five backup manager hosts with the largest number of media in the Available state and for each of the displayed hosts, the chart shows media that is Allocated, Frozen, or Suspended.
Servers with Fewest Available Media	Displays the five backup manager hosts with the lowest number of media in the Available state. The chart displays the following states: <ul style="list-style-type: none"><li>• Allocated</li><li>• Frozen</li><li>• Suspended</li><li>• Available</li></ul>
Five Largest Sessions	Displays the five largest sessions in gigabytes (GB).
Five Longest Sessions	Displays the five longest sessions in seconds.

## About the Tabs in the Topology Lower Pane

The lower pane on the Topology tab is displayed when you select a discovered backup element. Different tabs are displayed according to the element type you selected. To learn more about these tabs, see Table 19-9, “Tabs in the Lower Pane of Protection Explorer Topology,” on page 773.

**TABLE 19-9** Tabs in the Lower Pane of Protection Explorer Topology

Tab	Element Type	Description
Properties	All elements	Provides property information for an element, including information about whether the element supports backup.
Protection	Clients	Provides information about the last time the client was backed up. See “Viewing Charts for a Backup Manager Host” on page 778 for more information.
Charts	<ul style="list-style-type: none"> <li>• Backup Manager Hosts</li> <li>• Media Managers</li> <li>• Tape Libraries</li> </ul>	Shows a chart for the selected element.
Servers	<ul style="list-style-type: none"> <li>• Backup Manager Hosts</li> <li>• Media Managers</li> </ul>	<p>Displays the servers Protection Explorer monitors with the following information for each server:</p> <ul style="list-style-type: none"> <li>• <b>Volume</b> — The size of the volume backed up, in kilobytes</li> <li>• <b>Sessions</b> — The number of backup sessions that have run in the specified time</li> <li>• <b>Failed</b> — The number of failed sessions during the specified time</li> <li>• <b>Partial</b> — The number of partial sessions during the specified time</li> <li>• <b>Successful</b> — The number of successful sessions within the specified time</li> </ul>
Resources	<ul style="list-style-type: none"> <li>• Backup Manager Hosts</li> <li>• Media Managers</li> <li>• Tape Libraries</li> </ul>	<p>Displays the resources Protection Explorer monitors with the following for each server:</p> <ul style="list-style-type: none"> <li>• <b>Media Pools</b> — The number of media pools that the backup manager host can access</li> <li>• <b>Available Media</b></li> <li>• <b>Allocated</b></li> <li>• <b>Frozen</b></li> <li>• <b>Suspended</b></li> </ul> <p>In the Available Media, Allocated, Frozen, and Suspended columns, the first number shows the number of available online media; the second number shows the number of available offline media. Note that this information is available only to the backup manager host.</p>

**TABLE 19-9** Tabs in the Lower Pane of Protection Explorer Topology (*Continued*)

Tab	Element Type	Description
Sessions	<ul style="list-style-type: none"> <li>• Backup Manager Hosts</li> <li>• Media Managers</li> </ul>	<p>Displays the following information for the sessions assigned to a backup server:</p> <ul style="list-style-type: none"> <li>• <b>Session ID</b> — The identifier for the session</li> <li>• <b>Client</b> — The DNS name of the computer on which the session is taking place</li> <li>• <b>Backup Policy</b> — The name of the backup policy</li> <li>• <b>Schedule</b> — The name of the schedule for the session</li> <li>• <b>Status</b> — The status of the session</li> <li>• <b>Start Time</b> — The time the session started</li> <li>• <b>End Time</b> — The time when the session ended</li> <li>• <b>Duration</b> — The amount of time in seconds the session ran</li> <li>• <b>Size</b> — The size of the session</li> <li>• <b>Files</b> — The number of files that were backed up</li> </ul>
Media	<ul style="list-style-type: none"> <li>• Backup Manager Hosts</li> <li>• Media Managers</li> <li>• Tape Libraries</li> </ul>	<p>Displays the following information for the media attached to a backup server or tape library:</p> <ul style="list-style-type: none"> <li>• <b>Media ID</b> — The identifier for the media</li> <li>• <b>Media Pool</b> — The media pool to which the media belongs</li> <li>• <b>Usage Count</b> — How often the media is used</li> <li>• <b>Retention</b> — How long the media is retained</li> <li>• <b>State</b> — Whether the media is Full, Available, or Active</li> </ul>
Media Pool	<ul style="list-style-type: none"> <li>• Backup Manager Hosts</li> <li>• Media Managers</li> <li>• Tape Libraries</li> </ul>	<p>Displays the following information for the media pools containing the selected element:</p> <ul style="list-style-type: none"> <li>• <b>Media Pool</b> — The media pool to which the media belongs</li> <li>• <b>Backup Manager</b> — The name of the backup manager host in the media pool</li> <li>• <b>Library</b> — The name of the library in the media pool</li> <li>• <b>Available Media</b> — The number of available media</li> <li>• <b>Allocated</b> — The number of allocated media</li> <li>• <b>Frozen</b> — The number of frozen media.</li> <li>• <b>Suspended</b> — The number of suspended media</li> </ul>

**TABLE 19-9** Tabs in the Lower Pane of Protection Explorer Topology (*Continued*)

Tab	Element Type	Description
Drive Utilization	<ul style="list-style-type: none"><li>• Tape Libraries</li><li>• Drives</li></ul>	<p>Displays the following information for the drives in a tape library:</p> <ul style="list-style-type: none"><li>• <b>Library</b> — The name of the tape library that contains the drive</li><li>• <b>Drive</b> — The name of the drive</li><li>• <b>Media ID</b> — The media identifier</li><li>• <b>Status</b> — The running status of the drive</li></ul> <p>For information about configuring drive monitoring, see “Drive Monitoring” on page 247.</p>

---

## Sorting Information in the Lower Pane

You can sort the information displayed on the tabs in the lower pane by clicking the heading of a column. You can also sort more than one column at a time. The sorting feature for multiple columns can be extremely useful. For example, if you have several clients with failed backups, you would take the following steps to sort the table to show the clients according to their status:

1. Click the **Status** heading in the session column to sort the sessions according to status.
2. Press the **CTRL** key, and then click the **Client** heading.

The clients are sorted first according to their status and second according to their client name. You can now easily view all clients with failed sessions in alphabetical order.

You can sort as many columns as you want on a tab. Notice that the arrow indicates an ascending or descending sort order. The arrow also decreases in size for each additional column that is sorted. The largest arrow corresponds to the column that is sorted first; the second largest arrow corresponds to the second sort, and so on.



# Modifying Summary Backup Charts

You can modify Protection Explorer to display charts other than the default. To learn more about the data and options available in the Summary Backup charts, see Table 19-8, “Protection Explorer Summary Charts,” on page 772.

The Summary Settings page shown below displays a grid that lists the charts available on the Summary tab in Protection Explorer. To change a chart, select a grid on the this page and change the settings as described in the steps in this section.

The top four grids on the Summary Settings page correspond to the top four charts on the Summary tab, and the lower four grids on the Summary Settings page correspond to the lower Summary tab. There are only two tables on the bottom half of the Summary tab because each of the tables spans two grids.

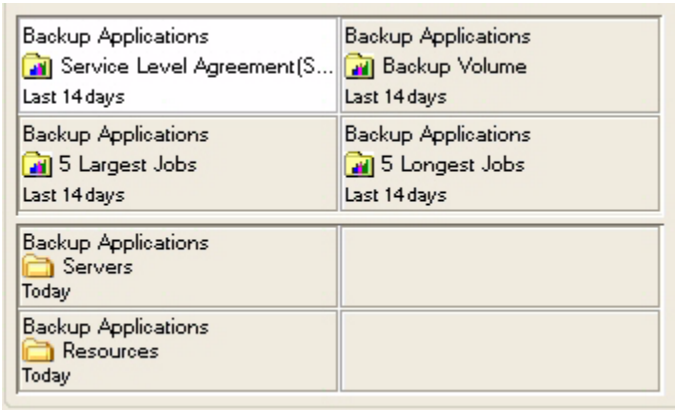


FIGURE 19-1 Summary Settings Page for Protection Explorer Charts

You can also modify the Summary tab instantly by clicking one of the buttons displayed in the Summary Settings page. These buttons are described in Table 19-10, “Buttons on the Summary Settings Page,” on page 777.

TABLE 19-10 Buttons on the Summary Settings Page

Button	Description
Clear All	Clears the settings for all the charts.
Clear	Clears the settings for the selected chart.
Revert	Returns the Summary Settings page to the previous setting.
Defaults	Returns the Summary Settings page to the default setting.

To modify a chart displayed on the Summary tab in Protection Explorer:

1. Access Protection Explorer as described in “Accessing Protection Explorer” on page 757.

2. Click the  icon.

3. To change the title for the summary page, enter a new title in the Title box.

4. Select the grid in which you want the chart to appear on the screen.

5. Select one of the following options from the Backup Element menu:

- **Backup Applications** — The chart includes the results from all backup applications.
- **A specific backup element** — The chart includes the results from only the backup application you selected.

6. Select the type of chart you want from the View menu.

To learn more about the available charts, see Table 19-8, “Protection Explorer Summary Charts,” on page 772.

7. Select a period for coverage from the Period menu.

8. Click **OK**.

The Summary page is updated with your changes and the chart accessed from the selected grid will reflect your changes.

---

## Viewing Charts for a Backup Manager Host

To quickly view charts for a backup manager host:

---

**Note** – To learn more about the charts in Protection Explorer, see Table 19-8, “Protection Explorer Summary Charts,” on page 772.

---

1. Access Protection Explorer as described in “Accessing Protection Explorer” on page 757.

2. Select the backup manager host on the Topology tab.


3. Click the **Charts** tab in the lower pane.

4. Select a chart from the Title menu.
5. Select a period of coverage for the chart.

See Table 19-3, “Toolbar for Charts,” on page 766 for information about the toolbar for charts.

## Printing Summary Charts

To print a summary chart:

1. Access a backup summary chart by clicking an element on the Topology tab.
2. Scroll to the bottom of the screen.
3. Click the **Print** (  ) button.
4. Click **Landscape** at the top of the new window if you want the picture to be printed in landscape format. To revert to portrait format, click **Portrait**.
5. To change the magnification of the image on the printed page, select the desired percentage.
6. Click **Print** when you are ready to print the chart.



## Business Tools

---

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 781
  - “Using Business Tools in Remote AD/LDAP Authentication Mode” on page 782
  - “Using the HBA Replacement Automator” on page 783
  - “Setting up Risk Analysis” on page 784
  - “Global Change Management” on page 785
- 

## About the Business Tools

---

**Caution** – 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 784 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 784 for more information.
- **Global Change Management** - It lets you record the properties and connections of managed elements. See “Global Change Management” on page 785.
- **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 783 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 Business Tools in Remote AD/LDAP Authentication Mode

To use Business Tools in remote AD/LDAP authentication mode:

1. Log in to where you have installed the Management Server as an administrator or a person with administrator privileges.
2. Run the CLI tool from a command prompt. Enter the following:

```
%CLI_DIR%\cli\bin
```

At the command prompt type, enter:

```
appiconfig
```

3. At the username prompt, enter the AD/LDAP username you have designated in your login-handler.xml file as the primary admin user.
4. Specify the AD/LDAP password for this user
5. Specify the server as localhost.
6. Accept the defaults for remaining prompts. The following message is displayed:  
CLI configuration is set; ready to issue commands
7. Log in to the Management Server using the specified AD username and password.
8. Launch Business Tools.

---

**Note** – If you change the password on the AD server for the user you have specified, you will have to execute these steps again and provide the new password.

---

---

## 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.

---

**Caution** – 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.

---

## Setting up Risk Analysis

The risk analysis tools flag HBAs and switches that are a risk in your environment. 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
```

---

**Caution** – 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 flags 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. The following sample output displays the analyzed elements and the properties that have changed:

```
-----
---
Global Configuration Change Manager

(c) Copyright 2002-2006 Hewlett-Packard Development Company, L.P.
All rights reserved.
```

Select a reference configuration, or save current configuration:

[1] config

[2] config1

[0] SAVE CURRENT CONFIGURATION

Enter selection number: Switch:1001:clbrocade3

Switch:1002:clbrocade2

Switch:1003:clbrocade1

Host:1000:COLO-WINHOST2

Host:1708:Host\_1708

Host:1714:Host\_1714

Host:1717:Host\_1717

Host:1724:Host\_1724

Host:1729:Host\_1729

Host:1732:Host\_1732

Host:1735:Host\_1735

Host:1739:Host\_1739

Host:1743:Host\_1743

Host:1747:Host\_1747

Host:1750:Host\_1750

Host:1753:Host\_1753

StorageSystem:1004:2107.75ABNY2

StorageSystem:1005:2107.75ABNY1

StorageSystem:1711:Emulex LP8000 FV3.93A0 Dv5-2.30a2 COLO-WINUTIL

StorageSystem:1721:IBM 2105F20 .134

computing differences...

Switch clbrocade2...

Host \_1735...

Host \_1714...

Host \_1747...

Host \_1753...

StorageSystem 2107.75ABNY2...

Host \_1743...

Host \_1729...

```

StorageSystem IBM      2105F20      .134...
StorageSystem 2107.75ABNY1...
Host _1724...
Switch clbrocade1...
Host _1750...
Host _1739...
Switch clbrocade3...
Host _1732...
Host COLO-WINHOST2...
Host _1717...
StorageSystem Emulex LP8000 FV3.93A0 Dv5-2.30a2 COLO-WINUTIL...
Host _1708...

```

```

-----
---
```

#### CHANGED PROPERTIES:

```
Switch clbrocade3
```

```
<contained element=Switch clbrocade3////>
```

```
PortType: was G, now F
```

```
State: was Enabled but Offline, now Enabled
```

```
Status: was Other, now OK
```

```
WWNofConnectedPort: was , now 10000000c95164d1
```

```
Switch clbrocade2
```

```
<contained element=Switch clbrocade2////>
```

```
PortType: was G, now F
```

```
State: was Enabled but Offline, now Enabled
```

```
Status: was Other, now OK
```

```
WWNofConnectedPort: was , now 10000000c9516579
```

```
Host COLO-WINHOST2
```

```
NEW LogicalDisk K:
```

```
NEW HostTargetMapping_1026
```

```
NEW HostTargetMapping_1045
```

```
NEW HostTargetMapping_1040
```

```
NEW DiskPartition Disk #1, Partition #0
```

```
NEW HostTargetMapping_1031
```

```
NEW LogicalDisk J:
```

```

NEW DiskDrive 6005076303ffc640000000000000110c:c0t0d1p4
NEW DiskDrive 6005076303ffc640000000000000110f:c0t0d4p3
NEW HostTargetMapping_1027
NEW DiskDrive 6005076303ffc640000000000000110d:c0t0d2p3
NEW DiskPartition Disk #5, Partition #1
NEW HostTargetMapping_1032
NEW DiskPartition Disk #6, Partition #0
NEW DiskDrive 6005076303ffc640000000000000110c:c0t0d1p3
NEW HostTargetMapping_1044
<contained element=Host COLO-WINHOST2////>
    PortID: was 0, now 198144
    PortSpeed: was 0 Gb/s, now 2 Gb/s
    State: was LinkDown, now Online
    WWNofConnectedPort: was , now 200600051e34924b
NEW DiskDrive 6005076303ffc640000000000000110e:c0t1d3p4
NEW HostTargetMapping_1035
NEW DiskPartition Disk #6, Partition #2
NEW DiskDrive 6005076303ffc640000000000000110f:c0t1d4p4
NEW DiskDrive 6005076303ffc640000000000000110b:c0t0d0p3
NEW DiskDrive 6005076303ffc6400000000000001110:c0t0d5p3
NEW DiskDrive 6005076303ffc640000000000000110b:c0t1d0p3
NEW LogicalDisk I:
NEW LogicalDisk G:
NEW HostTargetMapping_1030
NEW LogicalDisk H:
NEW HostTargetMapping_1028
NEW LogicalDisk E:
NEW HostTargetMapping_1041
NEW HostTargetMapping_1047
NEW DiskPartition Disk #6, Partition #1
NEW LogicalDisk N:
NEW DiskDrive 6005076303ffc6400000000000001110:c0t1d5p4
NEW HostTargetMapping_1039
NEW DiskDrive 6005076303ffc6400000000000001110:c0t1d5p3
NEW DiskPartition Disk #6, Partition #3
NEW DiskDrive 6005076303ffc640000000000000110b:c0t0d0p4

```

```

NEW LogicalDisk L:
NEW HostTargetMapping_1046
NEW HostTargetMapping_1036
NEW DiskPartition Disk #4, Partition #0
NEW DiskPartition Disk #5, Partition #0
NEW DiskDrive 6005076303ffc640000000000000110f:c0t1d4p3
NEW DiskPartition Disk #3, Partition #0
NEW DiskDrive 6005076303ffc640000000000000110e:c0t0d3p4
NEW DiskPartition Disk #2, Partition #1
NEW HostTargetMapping_1043
NEW DiskDrive 6005076303ffc640000000000000110d:c0t0d2p4
NEW HostTargetMapping_1038
NEW DiskDrive 6005076303ffc640000000000000110e:c0t1d3p3
NEW LogicalDisk Q:
NEW HostTargetMapping_1024
NEW DiskDrive 6005076303ffc640000000000000110e:c0t0d3p3
NEW LogicalDisk O:
NEW DiskDrive 6005076303ffc640000000000000110f:c0t0d4p4
<contained element=Host COLO-WINHOST2////>
    PortID: was 0, now 133120
    PortSpeed: was 0 Gb/s, now 2 Gb/s
    State: was LinkDown, now Online
    WWNofConnectedPort: was , now 200800051e34c539
NEW LogicalDisk P:
NEW HostTargetMapping_1034
NEW DiskDrive 6005076303ffc640000000000000110c:c0t1d1p3
NEW DiskDrive 6005076303ffc640000000000000110d:c0t1d2p4
NEW LogicalDisk M:
NEW HostTargetMapping_1025
NEW HostTargetMapping_1042
NEW DiskPartition Disk #2, Partition #0
NEW HostTargetMapping_1037
NEW DiskDrive 6005076303ffc640000000000000110b:c0t1d0p4
NEW DiskDrive 6005076303ffc6400000000000001110:c0t0d5p4
NEW DiskDrive 6005076303ffc640000000000000110c:c0t1d1p4
NEW HostTargetMapping_1029

```

```
NEW DiskDrive 6005076303ffc640000000000000110d:c0t1d2p3
NEW HostTargetMapping_1033
NEW DiskPartition Disk #2, Partition #2
```

```
-----
---
```

```
NO PHYSICAL TOPOLOGY CHANGES DETECTED
```

```
=====
```

```
process on the server ended
```





# Troubleshooting

---

This chapter contains the following topics:

- “Troubleshooting Installation/Upgrade” on page 793
- “Configuring the Java Console” on page 798
- ““Data is late or an error occurred” Message” on page 799
- “appstorm.<timestamp>.log Filled with Connection Exceptions” on page 799
- “Receiving HTTP ERROR: 503 When Accessing the Management Server” on page 800
- “Permanently Changing the Port a CIM Extension Uses (UNIX Only)” on page 802
- “Configuring UNIX CIM Extensions to Run Behind Firewalls” on page 804
- “Volume Names from Ambiguous Automounts Are Not Displayed” on page 807
- “Solaris Management Server Suddenly Restarts” on page 808
- “Installing the Software Security Certificate” on page 808
- “Troubleshooting Discovery and Get Details” on page 811
- “Troubleshooting Topology Issues” on page 823
- “Troubleshooting the Java Plug-in” on page 838
- “Troubleshooting Provisioning” on page 839
- “Troubleshooting Hardware” on page 840

---

## Troubleshooting Installation/Upgrade

This section provides help with troubleshooting installations and upgrades.

- “If Your Installation or Upgrade Failed, Capture the Logs” on page 794
- “Checking Installation Log Files” on page 795
- ““The environment variable ‘perl5lib’ is set.” Message” on page 795
- ““SEVERE: OUI-10029...” Message” on page 796
- “Brocade API Switches Displaying Stale Data” on page 796
- “Troubleshooting the Oracle Database (Windows)” on page 796

# If Your Installation or Upgrade Failed, Capture the Logs

(Windows management servers only) You can quickly gather system information and log files for troubleshooting by running the `srmCapture.cmd` program in `<installation directory>/tools`.

---

**Caution** – The `srmCapture.cmd` program requires that `zip.exe` is in the same folder as `srmCapture.cmd`. If you are missing `zip.exe`, you can find it in the `tools` directory of the management server CD.

---

The following information is gathered by `srmCapture.cmd`:

- List of environment variables, look for file `srmListEnvVar.txt`.
- Results from running `ipconfig /all`, look for file `srmListIpconfigAll.txt`.
- Results from running `netstat -noab`, look for file `srmListNetstatNoab.txt`.
- Results from running `netstat -rte`, look for file `srmListNetstatRte.txt`.
- Results from running `netsh diag show test`, look for file `srmListNetshDiagShowTest.txt`.
- Install wizard log files (all files are found in `%systemdrive%\srmInstallLogs`).
- `srmwiz.ini`.
- Oracle export log file.
- File SRM log files.
- File SRM configuration files.
- Oracle log files
- Zero G registry content

If you see a message resembling the following, “Current location, `d:\Tools`, is not writable,” the current working subdirectory is not writable. The `srmCapture.cmd` program will go through the following directories in order until it finds one that is writeable:

1. `%temp%`
2. `%tmp%`
3. `%systemdrive%`

## Checking Installation Log Files

The following log files are generated by the installer and can be found on the management server in the following directories:


- **C:\srmlInstallLogs** includes these log files:
  - **srmlInstall.log** — This is the master log file of the installation wizard session. It provides information for troubleshooting installation of the management server and related components.
  - **srmlInstallOracle10g.log** — Log file that provides information about the Oracle 10g database installation.
  - **srmlInstallSrm.log** — Log file that provides information about the management server installation.
  - **srmOracle<monthyear>Patch.log** — Log file that provides information about the installation of the specified Oracle patch.

Where <monthyear> is the date of release of the specified Oracle patch.

See the Troubleshooting chapter in the installation guide for more information about installations and upgrades.

## “The environment variable ‘perl5lib’ is set.” Message

(Windows only) If the perl5lib environment variable is set, the installation/upgrade fails with the following message:



**FIGURE 21-1** Perl5lib environment variable message

This variable may have been set by another application. The environment variable may have also been set if your upgrade of Oracle was suddenly stopped, for example, as a result of a power outage. You must remove the perl5lib environment

variable before you can run the installation/upgrade again. For information about removing environment variables, refer to the documentation for the Windows operating system.

## “SEVERE: OUI-10029...” Message

The installation wizard lets you specify an installation location for Oracle 10g. If you specify a location that is being used by another program or if you specify the Oracle DVD drive, Oracle displays the following message:

```
SEVERE: OUI-10029: You have specified a non-empty directory to install
this product. It is recommended to specify either an empty or a non-
existent directory. You may, however, choose to ignore this message
if the directory contains Operating System generated files or
subdirectories like lost+found
```

If you see this message, contact customer support. Engineering has found this message to indicate the installation of your Oracle database may have failed.

## Brocade API Switches Displaying Stale Data

All Brocade API switches are placed in quarantine after you upgrade to Build 6.0. This means previous data is preserved but you can no longer update the data using Get Details. Therefore, data such as topology, zoning information will be stale until you migrate to Brocade SMI-A. See “Discovering Brocade Switches” on page 35.

## Troubleshooting the Oracle Database (Windows)

This section provides Oracle troubleshooting help:

- “Use Only the Installation Wizard (or Unix Scripts) to Install/Upgrade Oracle” on page 796
- “Existing Oracle Database Is Detected” on page 798

### Use Only the Installation Wizard (or Unix Scripts) to Install/Upgrade Oracle

With this release of the product, the Oracle database is automatically installed using the new Installation Wizard (or Unix scripts) developed to install the management server along with the Oracle database used by the management server. Installing Oracle separately is no longer recommended.

---

**Caution** – Do not install the Oracle database separately, the management server Installation Wizard (or Unix scripts) automatically configures the Oracle database for use with the management server. If you install the Oracle database separately, the database will not meet the configuration settings required by the management server.

---

## Cancelling an Installation or Upgrade Before Completion

If you cancel the installation of the management server after the Oracle database is installed, you must use the Oracle scripts to remove the Windows Registry entries and other Oracle changes and files that were partially installed or future installations of the management server will fail. See “Uninstalling Oracle Using the Oracle Scripts” on page 797.

## Uninstalling Oracle Using the Oracle Scripts

With this release of the management server, the Oracle database is automatically installed by the Windows installation wizard installer and the Linux and Solaris installation scripts along with various Oracle files and some Windows Registry changes. If you cancel the installation, you must use the Oracle removal scripts included on the CD-ROM set for the management server to completely remove all of the Oracle files. If the Linux and Solaris installation scripts or the Windows installer wizard detects any Oracle files during a re-installation, the installation will fail. If you need to uninstall Oracle for any reason, you must follow these steps:

### Windows:

1. Put the Oracle DVD in the DVD drive of the management server for Windows.
2. Locate the following scripts on the DVD:
  - `removeOracle9i.vbs`
  - `removeOracle10g.vbs`
3. Open a Command window and enter a script name as follows:

```
cscript d:\removeOracle10g.vbs  
removeOracle10g.vbs
```

The script runs in the command window and removes the Oracle files from the server.

### Linux:

1. Put the Oracle DVD in the DVD drive of the management server for Linux.

2. Log on to the management server as root.
3. Run either of the scripts by entering the following at the command line:

```
<ORACLEDVD>/uninstallOracle9i.sh
```

```
<ORACLEDVD>/UninstallDatabase.sh
```

The remaining Oracle files are uninstalled from the management server for Linux.

### **Solaris:**

1. Put the Oracle DVD in the DVD drive of the management server for Solaris.
2. Log on to the management server as root.
3. Run either of the scripts by entering the following at the command line:

```
<ORACLEDVD>/uninstallOracle9i.sh
```

```
<ORACLEDVD>/uninstallOracle10g.sh
```

The remaining Oracle files are uninstalled from the management server for Solaris.

## Re-installing the Management Server

See “Cancelling an Installation or Upgrade Before Completion” on page 797.

## Existing Oracle Database Is Detected

If the Windows installation wizard installer (or the Unix installation scripts) detects an existing Oracle database, the following message is displayed: Existing Oracle Database is Detected. See “Uninstalling Oracle Using the Oracle Scripts” on page 797.

---

# Configuring the Java Console

It is recommended you configure your Java Console as follows for optimal performance. Please refer to the documentation for your Java Console for more information on how to make these changes.

To increase:

- The Memory, add -Xmx128m to the Java console
- The heap size, add -Xms128m to the Java console

---

## “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.

---

## appstorm.<timestamp>.log Filled with Connection Exceptions

When an Oracle redo log becomes corrupt, the management server is unable to connect to the database. Whenever this occurs, the management server writes to the `appstorm.<timestamp>.log` file.

To correct this problem, stop the management server and Oracle, and 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_appstorm.<timestamp>.log` file, which can be found in one of the following locations:
  - **Windows:** `\oracle\admin\APPIQ\bdump`.
  - **Unix systems:** `$ORACLE_BASE/admin/APPIQ/bdump`

You can verify if the redo log listed in the `alert_appstorm.<timestamp>.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;
```

The following sections describe how to start the database for the management server.



## Windows

In the Services window, make sure the OracleOraHome92TNSListener service has started and is set to automatic. See 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.

## Unix systems

To verify the Oracle service has started, enter the following at the command prompt:

```
# ps -ef | grep ora
```

If the service has started, output resembling the following is displayed:

```
/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
```

---

**Caution** – 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. To save space, the text has been shortened:

```
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
```

---

**Caution** – 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 is 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 box in the Add Address for Discovery page (**Discovery > Setup > Add Address**), enter 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 add it again. 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 802.
- 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 21-1** Troubleshooting Firewalls

<b>Configur- ation</b>	<b>Manual Start Parameters for CIM Extension</b>	<b>If Mentioned in cim.extension.parameters</b>	<b>Step 1 Discovery and RMI Registry Port</b>
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.	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
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 192.168.250.10:9012 Communication Port: 1234, 5678, 9012

**TABLE 21-1** Troubleshooting Firewalls (*Continued*)

<b>Configur- ation</b>	<b>Manual Start Parameters for CIM Extension</b>	<b>If Mentioned in cim.extension.parameters</b>	<b>Step 1 Discovery and RMI Registry Port</b>
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
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	172.16.10.10:1234OR 172.16.20.20:5678OR 172.16.30.30:9012 Communication Port: 1234, 5678, 9012

**TABLE 21-1** Troubleshooting Firewalls (*Continued*)

Configur- ation	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, cimom.Dcxws.agency.firstwai t=200000 cimom.Dcxws.agency.timeou t=200000	Any IP that is reachable Communication Port: 4673
No DNS, never resolve.		jboss.properties cimom.Dcxws.agency.firstwai t=200000 cimom.Dcxws.agency.timeou t=200000	Any IP that is reachable Communication Port: 4673
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 192.168.250.10:9012. 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 or 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-separated 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-separated 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.



---

**Caution** – 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.

---

This section contains the following topics:

- “Installing the Certificate by Using Microsoft Internet Explorer 6.0” on page 808
- “Changing the Security Certificate to Match the Name of the Server” on page 809

## Installing the Certificate by Using Microsoft Internet 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**.
3. When you are told there is a problem with the site's security certificate, click **View Certificate**.
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**.

# 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` program to delete the original certificate, and then use the `generateAppiqKeystore` program to create a new certificate and to copy the new certificate to the management server.

## Windows

To change the certificate on Windows:

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.

## Sun Solaris and Linux

To change the certificate on Sun Solaris and Linux:

1. Go to the [Install\_Dir] directory and run the following command:

```
eval `./usersvars.sh`
```

---

**Caution** – The quotes in the example must be entered as left single quotes as shown.

---

2. Go to the following directory:

```
[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 Discovery and Get Details

This section contains the following topics:

- “Troubleshooting Mode” on page 812
- “Unable to discover Emulex host bus adapters” on page 813
- “CIMOM Service Not Starting After Trying to Discover Sybase or SQL Server Applications” on page 813
- “Configuring E-mail Notification for Get Details” on page 814
- “Increasing the Time-out Period and Number of Retries for Switches in Progress” on page 815
- ““Connection to the Database Server Failed” Error” on page 816
- “Using the Test Button to Troubleshoot Discovery” on page 817
- “DCOM Unable to Communicate with Computer” on page 819
- “Duplicate Listings/Logs for Brocade Switches in Same Fabric” on page 819
- “Element Logs Authentication Errors During Discovery” on page 821
- “EMC Device Masking Database Does Not Appear in Topology (AIX Only)” on page 821
- “Management Server Does Not Discover Another Management Server's Database” on page 821
- “Microsoft Exchange Drive Shown as a Local Drive” on page 821
- “Unable to Discover Microsoft Exchange Servers” on page 822
- “Nonexistent Oracle Instance Is Displayed” on page 822
- “Requirements for Discovering Oracle” on page 822
- “Do Not Run Overlapping Discovery Schedules” on page 822
- ““This storage system uses unsupported firmware. ManagementClassName: class\_name” Message” on page 823
- “Troubleshooting Topology Issues” on page 823
- “Incorrect Topology Sometimes Displayed for CNT Switches” on page 828
- “Unable to Find Elements on the Network” on page 829
- “Unable to See Path Information” on page 829
- “Device Locking Mechanism for Brocade Element Manager Query/Reconfiguration” on page 829
- “A Discovered Sun StorEdge A5000 JBOD Does Not Display Its WWN Properly” on page 829
- “Unable to Monitor McDATA Switches” on page 830
- “Unable to Detect a Host Bus Adapter” on page 831
- “Navigation Tab Displays Removed Drives as Disk Drives” on page 831
- “Unable to Obtain Information from a CLARiiON Storage System” on page 831
- “Discovery Fails Too Slowly for a Nonexistent IP Address” on page 832
- ““CIM\_ERR\_FAILED” Message” on page 833
- “CIM\_ERR\_FAILED When Trying to Activate a Zone Set Using McDATA SWAPI” on page 835

- “Communicating with HiCommand Device Manager Over SSL” on page 835
- “Unable to Discover a UNIX Host Because of DNS or Routing Issues” on page 836
- “ERROR replicating APPIQ\_EVAStorageVolume during Get Details for an EVA array” on page 837

## 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.

When Troubleshooting Mode is enabled during Get Details, the following additional information is provided to assist in the identification of configuration issues:

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

## Unable to discover Emulex host bus adapters

The Emulex driver does not contain the required library that is required by the management server. You must install Emulex HBAnywhere software so that the management server can discover hosts configured with HBAnywhere and hbatest can detect the Emulex host bus adapter.

## CIMOM Service Not Starting After Trying to Discover Sybase or SQL Server Applications

If your management server is running on Linux, you will not be able to discover Sybase or SQL Server applications. If you already added a Sybase or SQL Server entry to be managed in the Discovery setup page and performed a Get All Element Details operation, entries for the Sybase or SQL server will be added to the oracle listener configuration file. On the next system reboot, or on the next restart of the Oracle service, the Oracle listener will error out, and the CIMOM service will not start.

To correct the issue:

1. Edit `ORA_HOME/network/admin/listener.ora` and remove the `SID_DESC` text blocks containing the `PROGRAM=hsodbc` string.

where `ORA_HOME` is the Oracle home

For example: `. /opt/oracle/product/9.2.0.4`

If you have a `SID_DESC` block similar to the text block below, remove this entire block.

```
(SID_DESC =  
(SID_NAME = SQLSERVERSID)  
(ORACLE_HOME = /opt/oracle/product/9.2.0.4)  
(PROGRAM = hsodbc)
```

2. Restart Oracle with the following command:  
`/etc/init.d/dbora restart`
3. Restart the appstormanager service.
4. After the service has started, delete any Sybase or SQL entries from the Application tab in the discovery setup page. This is necessary to prevent them from being re-added to the `listener.ora` on further discoveries.

## 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. See the User Guide for more information.

2. Add or edit the e-mail address for the Admin account.

The status reports for Get Details are sent as follows:

- `gaedemail` property is empty - The e-mail is sent to users whose roles have System Configuration selected.
- `gaedemail` property is populated - The e-mail is sent only to users whose e-mail is assigned to the `gaedemail` property.

3. If you want additional users to receive the status reports for Get Details, do the following:

- a. Select **Configuration > Product Health**, and then click **Advanced** in the **Disk Space** tree.

- b. Click **Show Default Properties** at the bottom of the page.

- c. Copy the `gaedemail` property.

- d. Return to the Advanced page.

- e. Paste the copied text into the Custom Properties box.

- f. Assign the e-mail accounts you want to receive the report to the `gaedemail` property. For example, if you want `user1@mycompany.com` and `user2@mycompany.com` to receive these status reports, modify the `gaedemail` property in the Custom Properties box as follows:

```
gaedemail=user1@mycompany.com;user2@mycompany.com
```

---

**Note** – Make sure the hash (#) symbol is removed from the `gaedmail` property.

---

- g. When you are done, click **Save**.

## 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 switch 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. If it does not receive a response from the switch a second time, the management server says it cannot contact the switch.

To change the time-out period and number of retries for switches, modify the properties as described in the following steps:

1. Access the management server.
2. Select **Configuration > Product Health**, and 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 21-2, "Time-out Properties," on page 815.
5. Return to the Advanced page.
6. Paste the copied text into the Custom Properties box.
7. 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 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 21-2** 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>

9. To modify the number of retries, repeat steps 4 through 6 by copying and pasting the property specified in the table below. 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 21-3** Retry Properties

Switch	Property
McDATA/Connectrix discovered through SNMP	<code>cimom.McData.Snmp.Retries</code>
Cisco	<code>cimom.Cisco.Snmp.Retries</code>
Other switches discovered through SNMP: <ul style="list-style-type: none"> <li>• Sun StorEdge</li> <li>• QLogic</li> </ul>	<code>cimom.snmp.switch.retries</code>

10. When you are done, click **Save**.

## “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 these items, 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, and 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 Win32 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
HDL 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/Logs for Brocade Switches in Same Fabric

### Duplicate listings: Targets tab

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
```

### Duplicate Logs

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
```

## Duplicate entries for the same element on the Get Details page

If an element is discovered through two different protocols, it may be listed twice on the Get Details page.

If you want to change the protocol used to discover an element that has already been discovered, delete the element before attempting to rediscover it. See “Deleting Elements from the Product” on page 97.

For some elements, duplicate entries may result if a second protocol is available. For example, you could choose to discover an element through a supported API, but if the element supports SMI-S, and the SMI-S provider is also available, the element could be discovered again. In this example, you could fix the issue by disabling the SMI-S provider.

## 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.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. See 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, see 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 the 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 rescheduled 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 `appstorm.<timestamp>.log` file in the following directory:

```
[Install_Dir]\jbossandjetty\server\appiq\logs
```

## "This storage system uses unsupported firmware. ManagementClassName: class\_name" Message

The following message is displayed when an LSI storage system is discovered, and is running unsupported firmware:

```
This storage system uses unsupported firmware. ManagementClassName:
class_name
```

Where `class_name` is the management class name for the unsupported array.

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 contains the following topics:

- “About the Topology” on page 824
- “Undiscovered Hosts Display as Storage Systems” on page 827
- “Solaris Machines Appear to Have Extra QLogic HBAs” on page 828
- “No Stitching for Brocade Switches with Firmware 3.2.0” on page 828
- “Link Between a Brocade Switch and a Host Disappears from the Topology” on page 828
- “Incorrect Topology Sometimes Displayed for CNT Switches” on page 828
- “Unable to Find Elements on the Network” on page 829
- “Unable to See Path Information” on page 829
- “Device Locking Mechanism for Brocade Element Manager Query/Reconfiguration” on page 829
- “A Discovered Sun StorEdge A5000 JBOD Does Not Display Its WWN Properly” on page 829
- “Sun 6920 Storage Systems: “ReplicatorSQLException: Database create error” During Get Details” on page 830
- “Mirrored Volumes Cannot Be Provisioned on Sun 6920 Storage Systems” on page 830
- “Unable to Monitor McDATA Switches” on page 830
- “Unable to Detect a Host Bus Adapter” on page 831
- “Navigation Tab Displays Removed Drives as Disk Drives” on page 831
- “Unable to Obtain Information from a CLARiiON Storage System” on page 831
- “Discovery Fails Too Slowly for a Nonexistent IP Address” on page 832
- ““CIM\_ERR\_FAILED” Message” on page 833
- “Communicating with HiCommand Device Manager Over SSL” on page 835
- “Unable to Discover a UNIX Host Because of DNS or Routing Issues” on page 836

### About the Topology

The software determines the topology by looking at the following:


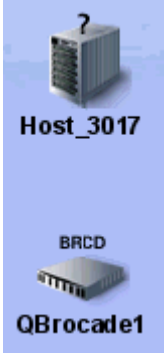
- **Fibre Channel switch** - The Fibre Channel switch contains a list of all elements within the fabric. The software obtains a detailed listing of all elements connected to the switch fabric.



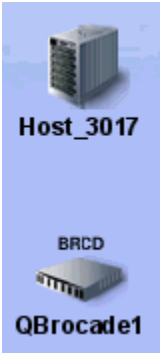

- **A host containing a Host Bus Adapter (HBA)** - All Fibre Channel host adapters look for available elements attached to the HBA. This information is gathered by CIM extensions and sent to the management server.
- **A proxy connected to the SAN** - Include a proxy that has a direct connection or a SAN connection to the management server. An example of a proxy is the EMC Solutions Enabler or Hitachi HiCommand Device Manager. LSI 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 21-4, “Troubleshooting Discovery and Get Details,” on page 825 provides details about how to correct problems that might occur during discovery and data collection.

**TABLE 21-4** Troubleshooting Discovery and Get Details

Scenario	Description	What to do
 <p>The host appears discovered and it is connected to the switch.</p>	<p>The software is aware of the host, but it cannot obtain additional information about it.</p>	<p>Verify that a CIM extension is installed on the host.</p> <p>Try discovering the element again, and then run Get Details.</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 828.</p>	<p>Verify that the host is on and the network cables are connected to it.</p> <p>Try discovering the element again, and then run Get Details.</p>

**TABLE 21-4** Troubleshooting Discovery and Get Details (*Continued*)

Scenario	Description	What to do
 <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 828.</p>	<p>Try getting the topology again:</p> <ol style="list-style-type: none"> <li>1. Click the <b>Discovery</b> menu, and then click the <b>Topology</b> tab.</li> <li>2. Verify the element is selected and click <b>Get Topology</b>.</li> </ol>
 <p>The element appears discovered, but a connected switch does not appear.</p>	<p>The switch has not been discovered.</p>	<p>Try discovering the switch again.</p> <ol style="list-style-type: none"> <li>1. Click the <b>Discovery</b> menu.</li> <li>2. Click the <b>Setup</b> tab and the <b>Add Address</b> button on the IP Addresses tab.</li> <li>3. Enter the IP address or DNS Name of the switch, and then enter its user name and password. Click <b>OK</b>.</li> <li>4. Verify the element is selected.</li> <li>5. Click <b>Start Discovery</b>.</li> <li>6. After discovery has completed, click the <b>Topology</b> tab.</li> <li>7. Verify the element is selected and click <b>Get Topology</b>.</li> </ol>
<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 WinMgmt.exe process is not running. This process starts WMI.*</li> <li>• In the Windows Event Log, DCOM error messages are shown.</li> </ul>	<p>An invalid user account was entered</p>	<p>Enter a valid user account that has administrative privileges so it can start WMI.</p>

\*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 the management server.

---

**Caution** – 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 103 for more information.

---

## Undiscovered Hosts Display as Storage Systems

On rare occasions, the management server displays undiscovered hosts as storage systems in System Explorer. To resolve this issue, provide the host’s world wide name (WWN) as described in the following steps:

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. Select **Configuration > Product Health**, and then click **Advanced** in the **Disk Space** tree.
3. Click **Show Default Properties** at the bottom of the page.
4. Copy the following property:  
`#hostPortWWNs=`
5. Return to the Advanced page.
6. Paste the copied text into the Custom Properties box.
7. Uncomment the `hostPortWWNs` property by removing the hash mark (#) in front of `hostPortWWNs`.
8. Enter the host’s WWN in hexadecimal format. Multiple WWNs can be entered as a comma-separated list. For example:  
`hostPortWWNs=00-01-C9,00-01-C8`
9. Click **Save**.

## 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.
- 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 446.

## 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.

## Sun 6920 Storage Systems: “ReplicatorSQLException: Database create error” During Get Details

While performing a Get Details, the Sun 6920 provider will return the error “ReplicatorSQLException: Database create error” under certain circumstances. This error appears in the management server logs but can be safely ignored. Sun Microsystems is aware of this issue.

## Mirrored Volumes Cannot Be Provisioned on Sun 6920 Storage Systems

Mirrored volumes are not represented properly by the management server. You cannot use the management server to provision mirrored volumes on Sun 6920 storage system.

## 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 “Discovering Switches, Storage Systems, NAS Devices, and Tape Libraries” on page 21.

---

**Caution** – 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, if 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_XXXXX` after getting element details. The `cfgadmn` command makes the software realize the drive has been removed. See 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 because the service processor is 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 IOs per second.

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 or three minutes and 45 seconds on Unix systems. 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 the `cimom.CimXmlClientHttpConnectTimeout` property to more than 20 seconds on Windows or three minutes and 45 seconds on Unix systems, 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. Select **Configuration > Product Health**, and 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.
5. Return to the Advanced page.
6. Paste the copied text into the Custom Properties box.
7. Make your changes in the Custom Properties box. 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 milliseconds 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**.



## “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 tries 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 Get Details operation the management server performs 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. Select **Configuration > Product Health**, and 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`.
4. Return to the Advanced page.
5. Paste the copied text into the Custom Properties box.
6. Make your changes in the Custom Properties box by changing the value of `cimom.mcData.swapIThrottle`. For example, the default is 200 ms. To change the value to 800 ms, 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. Verify if you can re-establish communication with EFCM by following the steps in “Re-establishing Communication with EFCM” on page 834. 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 to EFCM fails, there is a network problem that must be resolved. Once network connectivity is restored, click 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 indicate that it still cannot communicate with EFCM, wait approximately three minutes for the lost SWAPI connection to time out, and then click the **Test** button again. If this 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. Note that 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, see the support matrix to determine if the EFCM and switch versions are all supported. Contact technical support for further information.

## 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.

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, run Get Details for this element to update the zoning information. See “Get Details” on page 91 for more 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 that you want to communicate through a secure connection (SSL) and another that 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 following steps. 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. Select **Configuration > Product Health**, and 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.
4. Return to the Advanced page.
5. Paste the copied text into the Custom Properties box.
6. Make your changes in the Custom Properties box. 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.

## 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 to 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**, and then click **Advanced** in the **Disk Space** tree.

2. Paste the following text into the Custom Properties box.

```
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 200,000 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 starts to question whether its partner is still alive. If one entity (management server or extension) does not receive a message from the other during the interval set by the `timeout` property, it sends an "are you there" message. If that message is not acknowledged during the interval set by the `timeout` property, the entity concludes that the connection is no longer functioning. The CIM extension stops attempting to make a connection. When this occurs on the side of the management server, the management server attempts to re-connect (and continues the attempt until the host becomes available). The default value is 1,000 ms. You are modifying it to wait 200,000 ms or 3.33 minutes.

3. Click **Save**.

## ERROR replicating APPIQ\_EVAStorageVolume during Get Details for an EVA array

Errors similar to `ERROR replicating APPIQ_EVAStorageVolume` might occur when an EVA-specific data cache is updated during a Get Details operation. For example, when Data Protector creates a snapshot, a new virtual disk is automatically created on the EVA array, and the EVA database used by the management server is updated to reflect this change.

If the EVA database is changed during a Get Details operation, small replication errors may be seen as a result. The array information will be updated with the correct information next time Get Details runs.

## Recalculating the Topology

When Recalculating the topology or running Get Details, other tasks using the management server can be delayed because the management server must recalculate the topology, which is a resource intensive operation. Recalculation occurs after a Get Details when provisioning is done, and when you choose to recalculate the topology manually.

During the recalculation period, you may not be able to log into the application. If you are already logged into the application, navigation may not be possible until the topology recalculation is complete.

---

## Troubleshooting the Java Plug-in

This section contains the following topics:

- “Java Applet Has Data from a Different Build of Management Server Software” on page 838
- “OutOfMemoryException Messages” on page 839
- “Improving Reload Performance in System Explorer” on page 839

### Java Applet Has Data from a Different Build 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 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 contains the following topics:

- “Cannot Access a Resource Owned by Another Controller” on page 839
- “Error -56” on page 840
- ““Can't delete this zone” Message” on page 840
- “Changes in EFC Manager Requiring Get Details” on page 840

### Cannot Access a Resource Owned by Another Controller

If you receive a message about not being able to access a resource owned by another controller, it is because you tried to access a controller that has not been discovered. You should discover all controllers on the LSI storage system.

For example, assume you discovered only one of the controllers on an LSI 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 “Discovering Switches, Storage Systems, NAS Devices, and Tape Libraries” 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 or problems. To solve the problem, make sure the switch is physically connected to the network, and 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 see the following message when you try to delete a zone, move the zone to an inactive zone set, and 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 contains the following topics:

- “About Swapping Host Bus Adapters” on page 841
- ““Fork Function Failed” Message on AIX Hosts” on page 841
- “Known Driver Issues” on page 841
- “Known Device Issues” on page 841
- ““mailbox command 17 failure status FFF7” Message” on page 845
- ““Process Has an Exclusive Lock” Message” on page 845



## 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 are 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 Table 21-5, “Known Device Issues,” on page 842 provides a description of the known device issues. You can find the latest information about device issues in the release notes.

**TABLE 21-5** 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 occurs, you 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 /opt/APPQcime/tools/start 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>
Hosts running SGI IRIX version 6.5.22 or 6.5.24	NA	<p>If a host is 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.</p>

**TABLE 21-5** Known Device Issues (*Continued*)

Device	Software	Description
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 <code>/folder</code>, 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 <code>/folder</code> 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. Once you discover the host with the switches and storage, it reports its drives as being external. It reports 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.”
Solaris host	HDLM	<p>If you do a Get Details for the host by itself, on the bindings page, the controller number begins with <code>c-1</code>, for example, <code>c-1t0d58</code>. Perform Get Details on the host with storage and switches. The controller numbers are displayed correctly.</p>

**TABLE 21-5** Known Device Issues (*Continued*)

Device	Software	Description
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 fiber 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 1 in the SCSI bus column of the Disk Drives page.</p>
Any host	NA	<p>The Unmounted Volume box 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, or if your license does not allow you to discover a particular storage system. See the support matrix to determine which storage systems you can discover. The List of Features is accessible from the Documentation Center (<b>Help &gt; Documentation Center</b>).</p>

**TABLE 21-5** Known Device Issues (*Continued*)

Device	Software	Description
IBM Storage Systems	Subsystem Device Driver (SDD) or MPIO (multipath I/O)	If you discover 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 HBAAPI 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.

```
SYMAPI routine SymDevMaskSessionStart failed with error code 188: The operation failed because another process has an exclusive lock on the local Symmetrix.
```

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.

After the management server has detected the lock on the Symmetrix storage system, it tries to access the storage system for 15 minutes and then 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 before you remove the lock manually. Be sure that no other processes are occurring on the storage system. To learn how to remove the lock, see 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.
```

# Glossary

---

---

## 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.
- allocated** When media is referred to as allocated in Protection Explorer, the media is currently in use, either actively being used or it has a valid backup on it.
- available** When media is referred to as available in Protection Explorer, it is available for backup.

---

## 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.

---

## 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.

---

## F

**frozen** When media is referred to as frozen in Protection Explorer, the media will never become available again, but it is still available for restores.



## **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.

---

## **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, SQL Server, Caché, 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. LSI 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.
- suspended** When media is referred to as suspended in Protection Explorer, the media will not be used again until all backups written to it expire. It is still available for installations.
- 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, 59

## A

about

- Access tab, 317
- asset attributes, 393
- Business Tools, 781
- buttons in System Explorer, 312
- Capacity Explorer, 663
- Chargeback reports, 741
- custom commands, 358
- custom reports, 537
- CXFS, 397
- Event Manager, 605
- Events tab, 392
- filtering events, 620
- groups, 351
- Home page, 5
- List tab, 315
- monitoring options, 649
- Navigation screen, 370
- Path tab, 324
- Performance Explorer, 640
- Policies tab, 396
- Policy Manager, 681
- Protection Explorer toolbar, 764
- provisioning, 399
- Provisioning tab, 391
- remote console, 366
- Reporter, 519
- Reports tab, 395
- security, 175

- storage tiers, 705
- summary backup charts, 772
- System Explorer, 307
- Topology Screen, 379
- virtual applications, 852
- zone aliases, 853
- zone sets, 853

access point, 847

Access tab, 311, 317

accessing

- asset attributes, 708
- Database Admin Utility, 289
- domain controller, 113
- e-mail schedules, 535
- host security groups, 446
- logs, 235
- management server, 11
- Navigation screen, 375
- Policy Manager, 682
- properties, 375
- Protection Explorer, 757
- storage pools, 430
- System Explorer, 310
- trending information, 678
- volume information, 432
- zone sets, 412
- zoning information, 409

accessing information

- zone aliases, 406

account

- password, 184

Accounters, 781

accounts

- users, 182
- activating
  - a zone set, 417
  - zone sets, 412
- active
  - zone sets, 412
- Active Client Logons, 649
- Active Directory, 822
- active zone set, 847
- adding
  - asset information, 713
  - asset record, 706
  - custom command, 359
  - custom information, 716
  - departments, 717
  - discovery policies, 689
  - discovery schedule, 226
  - domain controller, 113, 156
  - elements, 192, 195
  - e-mail schedule, 531
  - event policies, 691
  - general information, 714
  - geographic information, 716
  - host security groups, 448
  - hosts, 480, 495, 504
  - IP address, 29
  - IP range, 27
  - journal entries, 616
  - licensing information, 716
  - new elements, 102
  - organizations, 192
  - provisioning policies, 690
  - roles, 190
  - staff information, 715
  - storage pool, 429
  - storage volume, 434
  - switches, 55
  - TNS Listener Port, 155
  - user accounts, 182
  - utilization policies, 684
  - virtual application, 326
  - virtual applications, 390
  - warranty information, 716
  - zone alias members, 407
  - zone members, 411
  - zone sets, 412, 413
  - zones, 410
- Advisors, 781
- aging statistics, 257
- AIX, 822
- API data
  - Brocade switches, 796
- API proxy, 851
- appiq.log, 233, 238
- APPIQ\_OWNER account, 113
- APPIQ\_USER, 156
- application
  - virtual, 326
- Application Administrator role, 175
- applications
  - business costs, 346
  - discovering, 113
  - virtual, 390
- archive destination, 299
- archive mode
  - changing, 296
- arranging
  - elements, 335
- asset attributes
  - elements, 393
  - viewing, 708
- asset based
  - reports, 741
- asset information
  - adding, 713
- Asset Management, 781
- asset record
  - creating, 706
- asset-based chargeback
  - setting up, 720
- assigning
  - business costs, 346
  - costs, 705
  - Custom Name, 378
  - e-mail schedules, 531
- authentication errors
  - SNMP, 821
- automatically
  - clearing events, 241
  - deleting events, 241
- Automators, 781
- available pools, 430
- Average Blocks Requested, 649



Average Delivery Time, 649

## B

- backing up
  - database, 287, 289
- Backup Manager
  - collector status, 247
- backup results, 758
- backup servers
  - viewing collectors, 244
- backups
  - managing, 753
- bar chart, 641
- binding
  - persistent, 849
- block size
  - cache, 462
- Bridge Agent, 44
- Brocade
  - switches, 618
  - tracing, 238
- Brocade Rapid program, 88
- Brocade switches, 88
  - API data, 796
  - discovering, 35
  - stale data, 796
- browsing
  - hosts, 369
  - switches, 369
- Buffer Hit Ratio, 649
- building
  - topology, 89
- business cost
  - setting, 326
- business costs
  - assigning, 346
- Business Tools, 781
  - risk Analysis, 784
- buttons
  - Protection Explorer, 764
  - remote console, 366, 368
  - right, 668
  - System Explorer, 312
- Bytes Received, 649
- Bytes Transmitted, 649

## C

- cache
  - refreshing reports, 259
- cache block size
  - changing, 462
- cache read-ahead multiplier, 427
- cache settings
  - modifying, 427, 462
- calculating
  - fixed declining balance, 733
  - straight line depreciation, 732
- calculations for
  - double declining balance, 735
- capacity charts, 676
- Capacity Explorer
  - about, 663
  - capacity charts, 676
  - Capacity tab, 663
  - finding capacity, 668
  - global view, 666
  - HDS, 668
  - legend, 663
  - magnifying, 666
  - Path tab, 663
  - printing, 666, 676
  - toolbar, 666
  - trending, 678
  - utilization reports, 675
  - Utilization tab, 663
- Capacity Manager
  - finding capacity, 668
  - HDS, 668
- Capacity tab, 663
- categorizations
  - storage, 705
- certificate
  - installing, 14
- changing
  - archive destination, 299
  - archive mode, 296
  - cache block size, 462
  - cache settings, 427
  - custom commands, 361
  - database passwords, 291
  - date, 277
  - discovery policies, 694
  - domain controller, 113, 156

- e-mail address, 186
- e-mail schedules, 256
- event policies, 696
- external tools, 369
- fabric name, 326, 355
- frequency, 641
- full name, 186
- license, 217
- logging, 238
- login name, 186
- number of retries, 58, 815
- Oracle Listener Password, 300
- organizations, 195
- password, 90, 156, 185, 186
- phone number, 186
- properties, 232
- provisioning policies, 695
- roles, 190
- SNMP trap listener, 57
- status, 707
- summary backup charts, 777
- time, 277
- time-out period, 58, 815
- TNS Listener Port, 155
- tools, 369
- topology, 666, 766
- user account, 184
- user name, 90
- user preferences, 187
- user profile, 186
- utilization policies, 693
- zone set, 414

changing ports

- zone alias, 408

Chargeback, 781

- about, 703
- asset information, 713
- asset record, 706
- asset-based, 703, 720
- custom information, 716
- custom properties, 716
- customizing filters, 750
- department views, 739
- double declining balance, 735
- e-mail schedule, 743
- e-mailing reports, 742
- filtering assets, 751
- fixed declining balance, 733
- general information, 714
- geographic information, 716
- hiding filters, 752
- infrastructure cost, 718
- licensing information, 716
- reports, 741
- saving, 708
- schedules editing, 745
- selecting element type, 750
- setting up, 705, 720, 726
- staff information, 715
- storage subsystem reports, 741
- storage tiers, 709
- storage-based, 703, 726
- straight line depreciation, 732
- viewing by element, 738
- viewing by enterprise, 740
- viewing report schedules, 747
- warranty information, 716

chart

- bar, 641
- clearing, 641
- filtering data, 641
- line, 641
- printing, 641
- saving, 641
- summary, 641
- time period, 641

charts

- custom periods, 647
- removing data, 646
- summary, 645
- thresholds, 646
- viewing, 645

checking

- database, 291

child organizations, 175

choosing

- fabrics, 348

CIM

- log levels, 238

CIM Extension

- port, 803

CIM Extensions, 1

CIM Object Manager, 848

- tracing, 238

cimom.CimXmlClientHttpConnectTimeout, 832

cimom.ClariionEventPollInterval, 617

- cimom.emc.skipRefresh, 62
- cimom.hds.exclude, 66
- cimom.log, 233
- cimom.provider.hds.hiCommandTimeout, 457
- cimom.symmetrix.exclude, 61
- CIO role, 175
- CISCO switches
  - topology, 311
  - VSAN, 311
- CLARiiON, 422, 455
  - polling interval, 617
- CLARiiON storage systems
  - no data, 831
- cleanup
  - reports, 258
- clearing
  - elements, 31
  - events, 241, 242, 392, 613
  - graph, 641
- clearing events, 241
- closing
  - Topology windows, 337
- CNT
  - switches, 40
- cold backup, 298
- collector status
  - Backup Manager, 247
  - Protection Explorer, 247
- collectors
  - editing schedule, 246
  - for reports, 395
  - managing, 270
  - Protection Explorer, 244
  - reports, 252, 277
  - scheduling, 245
  - setting time, 277
  - starting, 252, 525
  - stopping, 252
- Command View EVA
  - SNMP traps, 71
- command window
  - log file, 238
- commands
  - custom, 359, 361, 362
  - remote, 326
- Common Information Model (CIM), 847
- comparing
  - element performance, 644
- components, 3
- configuring
  - business cost, 326
  - Chargeback, 705, 720, 726
  - e-mail notification, 221
  - Java Console, 798
- console
  - buttons, 366
  - remote, 366
- controller
  - removing, 113
- conventions
  - naming zones, 509
- copying
  - zone sets, 416
- copying text, 368
- costs
  - assigning, 705
  - business, 346
- CR, 368
- CRC Errors, 649
- creating
  - asset record, 706
  - custom commands, 359
  - custom name, 378
  - custom reports, 540
  - departments, 717
  - discovery policies, 689, 690
  - e-mail schedule, 531
  - event policies, 691
  - global reports, 260
  - host security groups, 448
  - journal entries, 616
  - LUSE volumes, 457
  - new password, 186
  - organizations, 192
  - roles, 190
  - storage pool, 429
  - storage volume, 434
  - topology, 21
  - user accounts, 182
  - utilization policies, 684
  - virtual application, 326
  - virtual applications, 390
  - volumes, 455

- zone alias members, 407
  - zone in fabric, 410
  - zone sets, 412, 413
- creating a protection policy, 684
- CRLF, 368
- Current View combo box, 666
- custom
  - periods, 647
- custom commands
  - about, 358
  - adding, 359
  - deleting, 362
  - editing, 361
  - setting up, 326
  - stopping, 359
- custom name
  - truncated, 350
- custom properties, 716
- custom reports, 537
  - integrating, 552
- customer information
  - adding, 716
- Customer Support
  - generating a support database, 303
- customer support, 235
- customized time
  - filtering, 622
- Customizing, 232
- customizing
  - properties, 232
- cut off
  - custom name, 350
- CXFS, 397

## D

- data
  - backing up, 289
  - outdated (Brocade switches), 796
  - performance, 640
  - trending, 678
- data aging statistics, 257
- database
  - AIX, 822
  - backing up, 287, 289
  - changing passwords, 291
  - checking status, 291
  - cold backup, 296, 298
  - Database Admin Utility, 293
  - export, 292
  - reinitializing, 294, 302
  - restarting, 295
  - RMAN restore, 297
- Database Admin Utility
  - accessing, 289
  - archive destination, 299
  - archive mode, 296
  - changing passwords, 291
  - cold backup, 298
  - database passwords, 291
  - exporting database, 292
  - refreshing, 290, 300
  - reinitializing database, 294
  - resetting tablespace, 295
  - restarting database, 295
  - restoring cold backup, 296
  - RMAN restore, 297
  - temp tablespace, 295
  - viewing logs, 300
  - viewing/downloading logs, 299
- Database Admin Utility log file
  - resetting logs, 300
- database connection failed
  - error, 817
- date
  - setting, 277
- DCOM
  - unable to communicate, 820
- deactivating
  - policies, 698
  - zone sets, 412
- definitions
  - ports, 373
- deleting
  - custom commands, 362
  - departments, 718
  - domain controller, 113
  - element, 326
  - elements, 31, 97, 344
  - e-mail schedules, 535, 747
  - events, 243, 392, 614, 615
  - fabrics, 326, 355
  - host security groups, 453
  - multiple jobs, 509
  - organizations, 196

- performance data, 646
- policies, 698
- roles, 191
- storage pool, 431
- storage pools, 430
- switches, 55
- TNS Listener Port, 155
- user accounts, 186
- volumes, 438, 455
- zone aliases, 409
- zone sets, 412, 415, 840
- zones, 412, 840
- deleting events, 241
- department ownership, 718
  - removing, 730
- departments
  - adding, 717
  - deleting, 718
  - editing, 717
  - removing, 718
- designing
  - custom reports, 540
- detailed
  - tracing, 238
- detailed schema, 554
- details
  - event, 611
  - obtaining, 91
- detecting
  - IP range, 27
  - McDATA switches, 55
  - switches, 55
- device
  - about, 848
  - finding, 454
  - inaccessible, 454
- device issues, 841
- devices
  - deleting, 97, 344
  - updating, 343
- different
  - Java applet, 838
- different version
  - Java applet, 838
- direct attached storage, 658
- directory lookup, 649
- disabling
  - discovery schedule, 227
  - schedule, 227
- discovered
  - hosts, 334, 351
  - storage systems, 353
- discovered address
  - modifying, 90
- Discovered Elements
  - deleting elements, 344
- discovered elements
  - deleting elements, 97
- discovering
  - applications, 113
  - Brocade switches, 35, 88
  - CNT switches, 40
  - DNS Name, 29
  - EMC Solutions Enabler, 60
  - event details, 611
  - HDS storage systems, 65
  - HDS systems, 66
  - HP XP storage systems, 70, 74
  - IBM storage systems, 77
  - IP address, 29
  - McDATA switches, 44
  - Microsoft Exchange, 113, 138, 822
  - NetApp filers, 85
  - new elements, 102
  - Oracle, 113, 115
  - Oracle clusters, 115
  - passwords, 24
  - policies, 694
  - SQL servers, 123
  - storage system, 21
  - storage systems, 63, 79
  - Sun StorEdge storage systems, 79, 80, 81
  - Sun StorEdge switches, 42
  - switches, 21, 35
  - Sybase, 113, 134
  - Symmetrix systems, 61
  - troubleshooting, 822, 823, 824, 829, 845
  - user names, 24
- discovery
  - authentication errors, 821
  - Emulex host bus adapters, 813
  - quarantine, 100
  - time-out, 832
  - troubleshooting, 817
  - Windows proxy, 162

- discovery groups, 90
- discovery policies
  - creating, 689
- discovery requirements
  - Oracle, 823
- discovery schedule
  - adding, 226
  - disabling, 227
  - editing, 228
  - removing, 228
- discovery settings
  - importing, 31
  - saving, 33
- disk drive, 232, 831
- Disk Read, 649
- disk space monitoring
  - results of, 231
- Disk Utilization, 649
- Disk Write, 649
- displaying
  - deleted Oracle instances, 822
  - element home page, 326
  - element impact, 344
  - element properties, 375
  - log file, 238
  - port details, 326
- DNS, 822
- Domain Administrator role, 175
- domain controller
  - access, 156
  - accessing, 113, 156
  - removing, 113
- domain controller access, 113, 156
- double declining
  - calculating, 735
- dragging
  - elements, 666
- drivers
  - fixing, 841
- drives
  - Microsoft Exchange, 822
  - uninitialized, 831
- duplication
  - zone sets, 418
  - zones, 418

## E

- edit
  - event policies, 696
- editing
  - cache settings, 427
  - custom commands, 361
  - date, 277
  - departments, 717
  - discovery policies, 694
  - discovery schedule, 228
  - e-mail address, 186
  - e-mail schedule, 534
  - e-mail schedules, 256, 745
  - fabric name, 326, 355
  - frequency, 641
  - full name, 186
  - host security groups, 450
  - logging, 238
  - login name, 186
  - organizations, 195, 196
  - ownership percentage, 730
  - password, 185, 186
  - phone number, 186
  - Protection collectors, 246
  - provisioning policies, 695
  - roles, 190
  - status, 707
  - time, 277
  - user account, 184
  - user preferences, 187
  - user profile, 186
  - utilization policies, 693
  - zone set, 414
  - zone sets, 412
- editing ports
  - zone alias, 408
- EFC Manager, 44, 840
- element
  - about, 848
  - asset attributes, 393
  - deleting, 326
  - impact, 344
  - performance, 640
  - refresh, 326
- element created in Chargeback
  - about, 848
- element details
  - obtaining, 91

- element home page
  - viewing, 326
- element performance
  - comparing, 644
- element properties
  - viewing, 375
- element topology, 326
- element type
  - selecting, 630
- elements
  - adding, 192, 195
  - arranging, 335
  - capacity, 668
  - deleting, 31, 97, 344
  - dragging, 666
  - managing, 195
  - modifying, 90
  - organization, 195
  - removing, 196
  - topology, 89, 335
  - unable to find, 824, 829
  - unmanaged, 849
  - updating, 343
- e-mail
  - notification, 699
  - reports, 530
  - schedules, 256, 530
- e-mail address
  - changing, 186
- e-mail notification
  - configuring, 221
- e-mail schedule
  - adding, 531
  - editing, 534
  - organizations, 531
  - viewing, 747
- e-mail schedules, 536
  - deleting, 535, 747
  - editing, 745
  - viewing, 535
- e-mailing
  - chargeback reports, 742
  - logs, 235
  - reports, 743
- EMC CLARiiON, 63
- EMC Navisphere Agent, 455
- EMC Solutions Enabler, 60
- EMC Symmetrix, 422, 456
- Empty Chart message, 676
- Emulex host bus adapters, 813
- entries
  - journal, 616
- error
  - database connection failed, 817
  - error -56, 840
- Error 503, 800
- error message
  - exclusive lock, 845
- errors
  - authentication, 821
- event
  - details, 611
  - icons, 349
- event filter
  - customized time, 622
- Event Manager
  - about, 605
  - clearing, 613
  - clearing events, 241
  - deleting events, 241
  - events, 613, 615, 618
  - journal entries, 616
  - SNMP traps, 605
  - sorting, 615
  - switches, 618
- event policies
  - creating, 691
- event severity
  - setting, 699
- event types
  - all, 630
  - all but management server, 630
  - applications, 630
  - fabrics, 630
  - hosts, 630
  - management servers, 630
  - other, 630
  - storage systems, 630
  - switches, 630
- events
  - clearing, 241, 242, 392, 613
  - deleting, 243, 392, 614, 615
  - filtering, 620
  - filters, 392

- managing, 620
- policies, 696
- polling, 617
- removing, 613
- severity, 392
- severity levels, 607
- sorting, 392, 615
- switches, 618
- unclearing, 392
- viewing, 392

Events tab

- about, 392

exceptions, 839

excluding

- HDS systems, 66
- switches, 53
- Symmetrix systems, 61

exclusive lock

- error message, 845

export

- database, 292

external tools, 369

- setting up, 369

## F

fabric

- name, 355
- properties, 377

fabrics

- changing name, 326
- filtering, 348, 666
- removing, 326, 355
- viewing, 377

File Server SRM, 849

file systems

- about, 397

files

- all logs, 235

filtering

- assets, 751
- customized time, 622
- events, 620
- fabrics, 348, 666
- global reports, 529
- organizations, 197
- volumes, 433

filtering data, 641

filters

- Chargeback, 750
- hiding, 752
- host dependency, 512
- LUN dependency, 513
- storage system dependency, 511
- volume dependency, 512
- zone dependency, 513

Final Destination, 649

finding

- applications, 113
- capacity, 668
- devices, 454
- element impact, 344
- event details, 611
- hosts, 113
- IP address, 29
- IP range, 27
- new elements, 102
- performance data, 640
- storage systems, 21
- switches, 21
- trending information, 678

first time

- users, 2

fixed declining

- balance, 733

fixing

- drivers, 841

formatting

- reports, 527

forwarding

- SNMP traps, 605

Free Physical Memory, 649

frequency

- changing, 641

full name

- changing, 186

## G

general information

- adding, 714

Generating a support database

- support database
- customer support, 303

generic

- hosts, 334



- geographic information
  - adding, 716
- Get Details
  - email notification, 814
- getting
  - element details, 91
- getting details, 91
  - applications, 113
  - hosts, 113
- Global Reporter server
  - about, 849
- global reporting view, 849
- global reports
  - filtering, 529
  - setting up, 260
  - tnsnames.ora file, 266
- global view, 666
  - using, 337, 387
- graph
  - bar, 641
  - clearing, 641
  - filtering data, 641
  - line, 641
  - printing, 641
  - saving, 641
  - time period, 641
- graphs
  - saving, 640
- grey screen, 310
- grouping
  - hosts, 351
  - storage systems, 353
- groups
  - about, 351

## H

- hard zone, 849
- HBA Risk Analysis, 784
- HBAs
  - replacing, 783
  - swapping, 841
- HDS, 422
  - capacity definitions, 668
  - timeout, 457
- HDS 9200, 457
- HDS 9900, 457

- HDS 9900V, 457
- HDS storage systems
  - discovering, 65
- HdsSkipRefresh, 68
- Help button
  - remote console, 368
- Help Desk role, 175
- hiCommandTimeout, 457
- hiding
  - filters, 752
- hiding tabs, 347
- hierarchy
  - organizations, 175
- host
  - not in topology, 824, 829
- host bus adapter
  - unable to detect, 831
- host bust adapters
  - replacing, 783
- host dependency
  - host, 512
- host persistent binding, 849
- host security groups
  - accessing, 446
  - creating, 448
  - deleting, 453
  - editing, 450
- hosts
  - adding, 480, 495, 504
  - browsing, 369
  - discovered, 334
  - discovering, 113
  - generic, 334
  - grouping, 351
  - removing, 31
  - telnet to, 369
  - ungrouping, 352
- hot-swapped
  - drives, 831
- HP XP storage systems, 70, 74
- HTTP Error 503, 800
- HTTPS, 14

## I

- IBM storage systems
  - discovering, 77

- icons
  - event, 349
- impact
  - element, 344
  - showing, 326
- importing
  - database, 293
  - discovery settings, 31
- importing database, 293
- inaccessible
  - device, 454, 839
- increasing
  - Java heap size, 798
  - memory, 839
- information
  - obtaining element, 91
  - performance, 640
- infrastructure cost
  - setting, 718
- initiator WWN, 849
- In-memory Sort Ratio, 649
- inode file
  - about, 849
- installing
  - Java plug-in, 12
  - security certificate, 14
- integrating
  - custom reports, 552
- internal
  - drives, 831
- intervals
  - events polling of, 617
- Invalid CRC Errors, 649
- IP range, 27
- issues
  - devices, 841

## J

- Java applet
  - different version, 838
- Java Console
  - increading heap size
  - increasing
    - Java memory, 798
  - increasing memory, 798
- Java plug-in, 839

- installing, 12
- jboss.properties
  - modifying, 232
- jobs
  - provisioning, 509
- journal entries, 616

## K

- keeping active
  - remote console, 366
- killing
  - commands, 359
  - custom commands, 359

## L

- layout
  - arranging, 335
- levels
  - event, 618
  - severity, 607
- LF, 368
- Library Cache Hit Ratio, 649
- license
  - modifying, 217
- licensing information
  - adding, 716
- line chart, 641
- Link Failures, 649
- List tab, 311, 315
- local drives, 822
- locating
  - storage systems, 21
  - switches, 21
- log file
  - displaying, 238
- log messages
  - viewing, 58
- logging
  - changing, 238
  - viewing, 240
- logging off, 19
- login name
  - modifying, 186
- logs
  - accessing, 235
  - e-mailing, 235

- refreshing, 300
- resetting, 300
- saving, 235
- viewing, 300
- viewing/downloading, 299
- LUN numbering
  - requirements, 502
- LUN Security
  - Path Provisioning, 478
- LUSE volumes
  - cannot create, 457

**M**

- magnifying
  - Capacity Explorer, 666
  - topology, 666
- MALs, 850
- Managed Access Points
  - about, 850
- managed access points, 850
- managed application license, 850
- managed object, 850
- management server
  - about, 1
  - accessing, 11
  - assigning a name, 378
  - backup, 289
  - components, 3
  - security, 175
  - stopping, 18
  - variables, 362
- managing
  - backups, 753
  - collectors, 252
  - elements, 192, 195, 196
  - events, 620
  - performance collectors, 270
  - switches, 54
- mapped, 850
- MAPs, 850
  - about, 850
- materialized view, 850
  - about, 850
- materialized views
  - previous, 598
- maximizing
  - screen space, 527

- maximum thresholds, 646
- McDATA switches, 830
  - adding, 55
  - discovering, 44
- members
  - zone sets, 412
- memory
  - increasing, 839
- messages
  - data is late, 799
  - OutOfMemoryException, 839
- messages waiting, 649
- meta device, 850
- meta volume, 850
- metadata client, 397
- metadata server, 397
- Microsoft Exchange
  - Adding domain controllers, 138
  - deleting domain controllers, 140
  - discovering, 113, 138, 822
  - drive M, 822
  - failover clusters, 140
- Microsoft Exchange Services, 649
- minimum
  - thresholds, 646
- missing element, 850
- mixed mode authentication, 124
- modifying
  - cache block size, 462
  - cache settings, 427, 462
  - Chargeback, 705
  - custom command, 361
  - database passwords, 291
  - date, 277
  - departments, 717
  - discovered address, 90
  - discovery IP address, 30
  - discovery policies, 694
  - DNS name for discovery, 30
  - domain controller, 113, 156
  - elements, 90
  - e-mail address, 186
  - e-mail schedule, 534
  - e-mail schedules, 256
  - event policies, 696
  - external tools, 369
  - fabric name, 326, 355

- filters, 750, 751
- frequency, 641
- full name, 186
- jboss.properties, 232
- license, 217
- logging, 238
- login name, 186
- Oracle Listener Password, 300
- organizations, 195
- ownership percentage, 730
- password, 90, 156, 185, 186
- phone number, 186
- properties, 232
- provisioning policies, 695
- roles, 190
- schedules, 255
- SNMP trap listener, 57
- status, 707
- summary backup charts, 777
- time, 277
- TNS Listener Port, 155
- topology, 666, 766
- user account, 184
- user name, 90
- user preferences, 187
- user profile, 186
- utilization policies, 693
- zone members, 411
- zone set, 414
- zone sets, 412
- modifying ports
  - zone alias, 408
- monitoring
  - provisioning jobs, 509
- monitoring options
  - about, 649
- moving
  - elements, 335, 666
  - groups, 335
  - topology, 666
- multipathing, 379, 850

## N

- name
  - fabric, 355
  - storage pool, 430
  - storage tiers, 705
  - zone sets, 412

- naming
  - storage tier, 709
- naming conventions
  - for zones, 509
- naming organizations, 175
- navigating elements, 370
- navigation details, 326
- Navigation screen
  - about, 370
  - accessing, 375
- Navigation tab
  - duplication, 418
- Navisphere, 369
- Navisphere Agent, 455
- NetApp filers
  - discovering, 85
- netcnfg, 60
- nethost, 60
- new elements
  - adding, 102
- new password, 186
- new window, 527
- New Window option, 387
- no data
  - CLARiON storage systems, 831
- nonexistent IP addresses, 832
- nonexistent Oracle instances, 822
- notification
  - e-mail, 221, 699
- number of retries
  - changing, 58, 815

## O

- obtaining
  - performance data, 640
  - security certificate, 14
  - topology information, 89
  - utilization reports, 675
- old reports, 259
- opening
  - new window, 387, 666
- opening reports
  - new window, 527
- Oracle
  - deleted instances, 822

- discovering, 113, 115
- discovery requirements, 823
- Oracle Listener Password, 300
- Oracle TNS Listener Port, 155
- organizations
  - about, 175
  - adding, 192
  - deleting, 196
  - editing, 195, 196
  - elements, 192, 195, 196
  - e-mailing reports, 531
  - filtering, 197
  - properties, 189
  - users, 189
  - viewing, 194
- organizing
  - topology, 351
- outdated reports, 259
- OutOfMemoryException, 839
- overview
  - SAN zoning, 401, 403
- owners
  - adding, 717
  - editing, 717
  - removing, 718
- ownership
  - percentage, 730
  - removing, 718

**P**

- Pan button, 666
- parent organizations, 175
- Parse CPU, 649
- password
  - changing, 90, 156, 184, 185, 186
- path information
  - unable to find, 829
- Path Provisioning
  - adding hosts, 480, 495, 504
  - deleting jobs, 509
  - executing, 508
  - filters, 513
  - LUN Security, 478
  - monitoring jobs, 509
  - storage system dependency filter, 511
  - Volume Creation, 476
  - Volume Creation and LUN Security, 487

- Volume Creation, LUN Security, and Zone
  - Operation, 469, 502
  - volume dependency filter, 512
  - zone dependency filter, 513
  - Zone Operation, 482
- Path tab, 311, 324, 663
- paths, 379
- percentage
  - ownership, 730
- performance
  - finding, 640
  - statistics, 273
- performance collectors
  - managing, 270
  - starting, 272
  - stopping, 273
- Performance Explorer
  - about, 640
  - accessing, 640
  - chart format, 641
  - charts, 640
  - clearing graph, 641
  - comparing, 644
  - custom periods, 647
  - filtering data, 641
  - frequency, 641
  - Java plug-in, 12
  - multiple elements, 641
  - printing, 641
  - removing data, 646
  - saving graph, 641
  - summary charts, 645
  - thresholds, 646
  - time period, 641
  - trending, 640
- periods
  - custom, 647
- persisitent binding, 849
- phone number
  - editing, 186
- Physical Memory Used, 649
- planning organizations, 175
- policies
  - creating, 684
  - deactivating, 698
  - deleting, 698
  - discovery, 689, 694

- event, 691, 696
- provisioning, 690, 695
- testing, 684, 692
- utilization, 684, 693
- viewing, 682, 697
- Policies tab, 396
- Policy Manager
  - about, 681
  - accessing, 682
  - creating policies, 681
  - deactivating policies, 681, 698
  - deleting policies, 681, 698
  - modifying policies, 681
  - viewing policies, 681, 697
- polling
  - events, 617
- pool
  - information, 430
  - storage, 429
- pools
  - accessing, 430
  - deleting, 430
- port
  - CIM Extension, 803
- port details
  - showing, 326
- ports
  - definitions, 373
  - viewing, 343
- predicting, 640
  - performance, 640
  - storage capacity, 678
- Pre-submission Queue Size, 649
- previous materialized views, 598
- printing
  - backup summary charts, 779
  - Capacity Explorer, 666, 676
  - graphs, 641
  - System Explorer, 338
  - topology, 388, 666
- privileges
  - roles, 175
- problems
  - drivers, 841
- process
  - exclusive lock, 845
- Processor Utilization, 649
- profile
  - user, 186
- properties
  - customizing, 232
  - fabric, 377
  - organizations, 189
  - roles, 188
  - viewing, 375
- Protection Explorer
  - accessing, 757
  - backup results, 758
  - buttons, 764
  - charts, 773
  - collector status, 247
  - collectors, 244
  - editing collectors, 246
  - jobs, 773
  - lower pane tabs, 773
  - master server charts, 778
  - media, 773
  - printing, 779
  - Properties tab, 773
  - Protection, 773
  - resources, 773
  - scheduling collectors, 245
  - servers, 773
  - summary backup charts, 758, 772, 777
  - topology, 766
- provider, 851
- provisioning
  - about, 399
  - adding hosts, 480, 495, 504
  - copying zone sets, 416
  - LUN numbering, 502
  - monitoring jobs, 509
  - policies, 695
  - troubleshooting, 454, 455, 456, 839, 840, 845
- provisioning jobs
  - scheduling, 506
- provisioning policies
  - creating, 690
- Provisioning tab, 391
- proxy server, 851

## Q

- quarantine
  - adding elements, 100

- clearing elements, 100
- queue size, 649

## R

- RAID level, 430
- Rapid program, 88
- read caching, 427
- Read Operations, 649
- Read Requests, 649
- recalculating topology, 5
- Receive Queue Size, 649
- receiving
  - SNMP traps, 605
- refresh
  - element, 326
- refresh view, 526
- refreshing
  - Database Admin Utility, 290
  - element properties, 375
  - logs, 300
  - report cache, 259
  - reports, 526
  - Symmetrix systems, 62
- reinitializing
  - database, 294, 302
  - warnings, 302
- remote
  - commands, 326
  - console, 366
- remote console
  - buttons, 368
  - copying text, 368
  - keeping active, 366
  - menu options, 368
- remote drives, 822
- remote scripts
  - running, 700
- remote sites
  - filtering, 529
- removing
  - custom commands, 362
  - department ownership, 730
  - departments, 718
  - discovery schedule, 228
  - domain controller, 113
  - element, 326

- elements, 31, 97, 195, 196, 344
- e-mail schedules, 535, 747
- events, 243, 392, 613, 614, 615
- fabrics, 326, 355
- host security groups, 453
- multiple jobs, 509
- organizations, 196
- ownership, 718
- performance data, 646
- policies, 698
- roles, 191
- schedules, 255
- storage pool, 431
- storage pools, 430
- switches, 55
- TNS Listener Port, 155
- user accounts, 186
- volumes, 438
- zone aliases, 409
- zone members, 411
- zone sets, 415, 840
- zones, 840

- removing ports
  - zone alias, 408
- renaming
  - storage tiers, 705
- replacing
  - fabric name, 355
  - HBAs, 783

Reporter

- about, 519
- custom, 537

reports

- asset based, 741
- Chargeback, 743
- cleanup, 258
- collectors, 252
- custom, 537
- data aging statistics, 257
- deleting e-mail schedules, 535
- designing, 540
- detailed schema, 554
- e-mail, 530
- e-mail schedule, 536
- e-mail schedules, 255, 256, 531, 535
- e-mailing, 742
- formatting, 527
- global, 260, 529

- HTML, 527
- integrating, 552
- Microsoft Excel, 527
- new window, 527
- organizations, 531
- PDF, 527
- refreshing, 526
- refreshing cache, 259
- screen space, 527
- sending, 530
- starting collectors, 525
- storage based, 741
- storage subsystem, 741
- utilization, 675
- XML, 527
- Reports tab
  - about, 395
- Requests Serviced, 649
- resetting
  - tablespace, 295
- restarting
  - database, 295
- restoring
  - cold backup, 296
  - database, 291
  - RMAN, 297
- right button, 668
- risk analysis
  - setting up, 784
- RMAN
  - destination, 299
- RMAN backup
  - viewing results, 288
- RMAN backups, 287
- RMAN restore, 297
- roles
  - about, 175
  - adding, 190
  - Application Administrator, 175
  - CIO, 175
  - deleting, 191
  - Domain Administrator, 175
  - editing, 190
  - Element Control privilege, 175
  - Full Control privilege, 175
  - Help Desk, 175
  - privileges, 175
  - properties, 188
  - Server Administrator, 175
  - Storage Administrator, 175
  - users, 188
  - View privilege, 175
- running
  - cold backup, 298
  - remote scripts, 700
- Rx Bytes, 649
- S**
  - SAN, 851
  - SAN Zoning Overview, 401
  - saving
    - chargeback information, 708
    - discovery settings, 33
    - graph, 641
    - graphs, 640
    - logs, 235
    - settings to a file, 33
    - topology, 312, 666, 764
  - scanning
    - IP range, 27
  - schedules
    - deleting, 535
    - disabling, 227
    - editing, 745
    - e-mail, 256, 530, 535
    - modifying, 255
    - removing, 255
  - scheduling
    - database backup, 287
    - discovery, 226, 227, 228
    - provisioning jobs, 506
    - report cleanup, 258
    - reports, 743
  - screen space
    - maximizing, 527
  - scripting
    - variables, 362
  - scripts
    - remote, 700
    - stopping, 359
  - searching
    - topology, 666
  - security
    - Management server, 175



- roles, 190
- security certificate
  - installing, 14
- selecting
  - element type, 630
  - fabrics, 348
  - severity levels, 607
- Send Queue Size, 649
- sending
  - logs, 235
  - reports, 530
- server
  - accessing, 11
  - backup, 289
- Server Administrator role, 175
- services
  - stopping, 18
- sets
  - zone, 412
- setting
  - business cost, 326
  - date, 277
  - discovery passwords, 24
  - discovery user name, 24
  - event severity, 699
  - infrastructure cost, 718
  - report collectors, 395
  - storage cost, 709
  - storage tier, 709
  - thresholds, 646
  - time, 277
  - time period, 641
- setting up
  - Chargeback, 705, 720, 726
  - e-mail notification, 221
  - external tools, 369
  - global reports, 260
  - remote commands, 326
- settings
  - cache, 462
- severity
  - events, 392
- Severity combo box, 666
- severity levels, 349, 618
  - all, 607
  - all but clear, 607
  - clear, 607
  - critical, 607
  - informational, 607
  - major, 607
  - minor, 607
  - unknown, 607
  - warning, 607
- showing
  - element impact, 344
  - impact, 326
  - port details, 326
- signing out, 19
- size
  - cache block, 462
  - storage pool, 430
- SMTP, 221, 851
- SNIA specification, 851
- SNMP
  - authentication errors, 821
  - trap forwarding, 605
- SNMP trap destinations, 219
- SNMP trap listener
  - changing, 57
- soft zone, 851
- sorting
  - events, 392, 615
- specifying
  - custom periods, 647
- SQL Server
  - authentication modes, 124
- SQL servers
  - discovering, 123
- staff information
  - adding, 715
- Standard XML Connector, 1
- starting
  - collectors, 252, 525
  - performance collectors, 272
- Statistics, 232
  - data aging, 257
  - performance, 273
- status
  - changing, 707
  - protection collectors, 247
- stop button
  - remote console, 368
- stopping

- collectors, 252
- performance collectors, 273
- SAN details, 93
- scripts, 359
- services, 18
- storage
  - assigning costs, 705
- Storage Administrator role, 175
- storage area network, 851
- storage based
  - reports, 741
- storage capacity
  - predicting, 678
- storage cost
  - setting, 709
- storage elements
  - viewing, 332
- storage pool, 851
  - creating, 429
  - deleting, 431
- storage pools
  - accessing, 430
  - deleting, 430
- storage subsystem based
  - reports, 741
- storage system proxy, 851
- storage systems, 79, 232
  - discovering, 21, 79
  - grouping, 353
  - polling interval, 617
  - removing, 31
  - ungrouping, 354
- storage tier
  - setting, 709
- storage tiers, 705
- storage-based
  - Chargeback, 703
- storage-based chargeback
  - setting up, 726
- Store Group Size, 649
- Store Size, 649
- straight line depreciation
  - calculating, 732
- suggested topics, 2
- summary backup charts, 758
  - about, 772

- Sun StorEdge
  - SNMP trap listener, 57
- Sun StorEdge A5000, 829
- Sun StorEdge storage systems, 79, 80, 81
- Sun StorEdge switches, 42
- swapped
  - drives, 831
- swapping HBAs, 841
- switch
  - events, 618
- switch port zoning, 851
- Switch Risk Analysis, 784
- switches
  - adding, 55
  - browsing, 369
  - CISCO, 311
  - discovering, 21, 35
  - duplication, 418
  - excluding, 53
  - managing, 54
  - McDATA, 44, 55, 830
  - name truncated, 350
  - number of retries, 58, 815
  - port status, 373
  - removing, 31, 55
  - telnet to, 369
  - time-out period, 58, 815
  - unable to monitor, 830
- Sybase
  - discovering, 113, 134
- System Event Times, 649
- System Explorer
  - about, 307
  - accessing, 310
  - buttons, 312
  - can't access, 839
  - Cisco switches, 311
  - deleting elements, 97, 344
  - Discovered hosts, 334
  - element impact, 344
  - event icons, 349
  - global views, 337, 387
  - groups, 351
  - hiding tabs, 347
  - Java plug-in, 12
  - menu options, 326
  - printing, 338

- scripting variables, 362
  - storage elements, 332
  - Topology windows, 337
  - update, 343
  - user interface, 311
  - viewing ports, 343
  - virtual applications, 390
- T**
- tablespace
    - resetting, 295
  - tabs
    - Access, 317
    - hiding, 347
    - List, 315
    - Path, 324
  - telnet to
    - hosts, 369
    - switches, 369
  - testing
    - policies, 684, 692
  - text copying, 368
  - thresholds, 646
  - tiers
    - storage, 705
  - time
    - setting, 277
  - timeout
    - HDS, 457
  - time-out period
    - changing, 58
  - TNS Listener Port
    - changing, 155
  - tnsnames.ora file, 266
  - toolbar
    - Capacity Explorer, 666
    - Protection Explorer, 764
    - System Explorer, 312
  - tools
    - external, 369
  - topology
    - AIX, 822
    - building, 89
    - changing, 666
    - Cisco switches, 311
    - filtering, 666
    - host not appearing, 824, 829
    - moving, 666
    - new window option, 666
    - organizing, 351
    - printing, 388, 666
    - saving, 666, 764
    - updating, 5
    - VSANs, 311
  - topology issues, 824
  - Topology tab
    - about, 379
  - Total Cost/Month for Each Departments, 741
  - tracing
    - detailed, 238
  - trap forwarding, 605
  - trap generation, 219
  - traps
    - SNMP, 605
  - trending
    - Capacity Explorer, 678
    - reports, 519
  - triangle
    - icons, 349
  - troubleshooting
    - CLARiiON storage systems, 455
    - discovery, 817
    - discovery and getting element details, 817, 820, 822, 823, 824, 829, 845
    - EMC storage systems, 456
    - EMC Symmetrix storage systems, 456
    - HDS storage systems, 457
    - JavaScript blockers, 310
    - LSI storage systems, 462
    - Microsoft Exchange, 822
    - popups, 310
    - provisioning, 454, 455, 456, 839, 840, 845
  - truncated
    - custom name, 350
  - Tx Bytes, 649
  - types of
    - monitoring, 649
- U**
- unable to
    - discover, 817
  - Unable to access resource, 839
  - unable to detect
    - host bus adapter, 831

- unable to find
  - elements, 829
  - path information, 829
- unable to retrieve data, 841
- unclearing
  - events, 392
- ungrouping
  - hosts, 352
  - storage systems, 354
- uninitialized
  - drives, 831
- updating
  - element properties, 375
  - elements, 343
- updating topology, 5
- uring, 798
- used pools, 430
- user accounts
  - creating, 182
  - deleting, 186
- user audit log, 237
- user interface
  - System Explorer, 311
- user name
  - changing, 90
- user profile
  - modifying, 186
- users
  - about, 175
  - adding, 182
  - first time, 2
  - organizations, 189
  - roles, 188, 190
- using
  - global view, 337, 387
  - journal entries, 616
- utilization
  - policies, 692, 693
- utilization policies
  - creating, 684
- Utilization tab, 663

## V

- variables
  - scripting, 362
- viewing

- all logs, 235
- asset attributes, 708
- capacity, 668
- capacity charts, 676
- chargeback by element, 738
- charts, 645
- data aging statistics, 273
- database status, 291
- department views, 739
- element asset attributes, 393
- element home page, 326
- element impact, 344
- element properties, 375
- e-mail schedules, 535, 747
- event details, 611
- events, 392
- fabrics, 348, 377
- log file, 238
- log messages, 58
- logging, 240
- logs, 300
- organization properties, 189
- organizations, 194
- performance data, 640
- policies, 682, 697
- ports, 343
- report refresh, 526
- RMAN backup results, 288
- schedule history, 748
- scheduled e-mail deliveries, 255
- security certificate, 14
- storage elements, 332
- summary backup charts, 758
- topology, 21
- trending information, 678
- utilization reports, 675
- volume information, 432
- zone aliases, 406
- viewing/downloading
  - logs, 299
- views
  - global, 337, 387
- virtual application
  - about, 852
  - adding, 326, 390
  - creating, 390
- Virtual Memory Used, 649
- volume

- meta, 850
- volumes
  - creating, 455
  - deleting, 438, 455
  - filtering, 433
  - unable to create, 457
  - viewing information, 432
- VSAN, 311

## W

- warnings
  - reinitializing, 302
- warranty information
  - adding, 716
- WBEM, 852
  - about, 852
- window
  - opening reports, 527
- Windows Management Instrumentation, 852
- Windows proxy
  - discovery, 162
- WinMgmt.exe, 824
- WMI, 852
  - about, 852
- Worldwide Name, 853
- WorldWide Name zoning, 851
- Worldwide Name zoning, 852
  - about, 853
- write caching, 427
- Write Operations, 649
- WWN, 853

## X

- XFS, 397
- Xiotech storage systems, 82

## Z

- zone
  - hard, 853
  - soft, 853
- zone alias
  - about, 853
  - creating, 407
  - deleting, 409
  - modifying ports, 408
- zone hard

- about, 853
- zone member, 853
- zone members
  - adding, 411
  - removing, 411
- zone set
  - creating, 413
- zone sets
  - accessing, 412
  - activating, 417
  - copying, 416
  - creating, 412, 413
  - deleting, 415, 840
  - duplication, 418
  - editing, 412
  - members, 412
  - modifying, 414
  - removing, 412
- zones
  - creating, 410
  - deleting, 412, 840
  - duplication, 418
  - naming conventions, 509
  - removing, 412
  - viewing aliases, 406
- Zoning Library
  - duplication, 418

