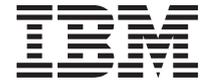


IBM System Storage N series



OnCommand Operations Manager Administration Guide For Use with Core Package 5.0

Contents

Introduction to Operations Manager	19
What DataFabric Manager server does	19
What a license key is	19
Access to Operations Manager	19
Information to customize in the DataFabric Manager server	20
Administrator accounts on the DataFabric Manager server	20
Authentication methods on the DataFabric Manager server	21
Authentication with native operating system	21
Authentication with LDAP	21
Preface	23
Supported features	23
Websites	23
Getting information, help, and service	23
Before you call	24
Using the documentation	24
Hardware service and support	24
Firmware updates	25
How to send your comments	25
What the discovery process is	27
Discovery by the DataFabric Manager server	27
What SNMP is	27
When to enable SNMP	27
SNMP versions to discover and monitor storage systems	28
What the Preferred SNMP Version option is	29
How the DataFabric Manager server chooses network credentials for discovery	29
Discovery process using SNMPv1 or SNMPv3	30
Monitoring process using SNMPv1	30
Monitoring process using SNMPv3	31
Setting SNMPv1 or SNMPv3 as the preferred version	31
Setting SNMPv1 as the only SNMP version	31
Setting SNMPv1 or SNMPv3 to monitor a storage system	32

Modifying the network credentials and SNMP settings	32
Deleting the SNMP settings for the network	33
Addition of a storage system from an undiscovered network	33
Diagnosis of SNMP connectivity	33
What host discovery is	34
Ping methods in host discovery	34
What host-initiated discovery is	34
Discovery of vFiler units	35
Discovery of storage systems	35
Discovery of storage systems and networks	36
Methods of adding storage systems and networks	37
Guidelines for editing discovery options	37
Discovery of a cluster by Operations Manager	39
Adding a cluster	40
Data ONTAP 8.0 cluster monitoring tasks using Operations Manager	40
Limitations of cluster monitoring in Operations Manager	41
Introduction to the gateway SAN-attached storage management	41
Limitations of the gateway SAN-attached storage management in Operations Manager	42
Tasks performed from the Storage Controller Details page for the gateway	42
Viewing configuration details of storage arrays connected to the gateway	42
Role-based access control in the DataFabric Manager server	45
What role-based access control is	45
Configuring vFiler unit access control	45
Logging in to the DataFabric Manager server	46
What default administrator accounts are	46
List of predefined roles in the DataFabric Manager server	47
Active Directory user group accounts	48
Adding administrative users	48
How roles relate to administrators	49
What predefined global roles are	49
What inheritance roles are	51
What capabilities are	52
Role precedence and inheritance	52

Creating roles	52
Modifying roles	52
What an RBAC resource is	53
Granting restricted access to RBAC resources	53
Access check for application administrators	54
How reports are viewed for administrators and roles	54
What global and group access control is	55
Management of administrator access	55
Prerequisites for managing administrator access	55
Limitations in managing administrator access	56
Controlled user access for cluster management	56
Summary of the global group	56
Who local users are	57
What domain users are	62
What User groups are	64
What roles are	67
What jobs display	69
Groups and objects	71
What group types are	71
What homogeneous groups are	72
What mixed-type groups are	72
What a Global group is	73
What hierarchical groups are	73
Creating groups	73
Creating groups from a report	74
What configuration resource groups are	74
Guidelines for creating groups	75
Guidelines for creating configuration resource groups	75
Guidelines for adding vFiler units to Appliance Resource group	76
Editing group membership	76
What group threshold settings are	77
What group reports are	77
What summary reports are	77
What subgroup reports are	78
What cluster-related objects are	78
Creating a group of cluster objects	80

Storage monitoring and reporting	81
The DataFabric Manager server monitoring process	81
Cluster monitoring with Operations Manager	83
What the cluster management logical interface is	83
Information available on the Cluster Details page	83
Viewing the utilization of resources	84
Links to FilerView	85
Query intervals	86
Guidelines for changing monitoring intervals	86
What SNMP trap listener is	86
What SNMP trap events are	87
How SNMP trap reports are viewed	87
When SNMP traps cannot be received	87
SNMP trap listener configuration requirements	88
How SNMP trap listener is stopped	88
Configuration of SNMP trap global options	88
Information about the DataFabric Manager server MIB	88
What events are	89
Viewing events	89
Managing events	90
Operations on local configuration change events	90
List of events and severity types	90
Alarm configurations	114
Guidelines for configuring alarms	114
Creating alarms	115
Testing alarms	115
Comments in alarm notifications	116
Example of alarm notification in e-mail format	116
Example of alarm notification in script format	117
Example of alarm notification in trap format	117
Response to alarms	118
Deleting alarms	118
Working with user alerts	118
What user alerts are	118
Differences between alarms and user alerts	119
User alerts configurations	119

E-mail addresses for alerts	120
Domains in user quota alerts	120
What the mailmap file is	121
Guidelines for editing the mailmap file	121
How the contents of the user alert are viewed	122
How the contents of the e-mail alert are changed	122
What the mailformat file is	122
Guidelines for editing the mailformat file	123
Introduction to DataFabric Manager server reports	123
Introduction to report options	124
Introduction to report catalogs	124
Reports in Operations Manager	125
What performance reports are	129
Configuring custom reports	129
Deleting custom reports	130
Putting data into spreadsheet format	131
What scheduling report generation is	131
Methods to schedule a report	133
What Schedules reports are	135
What Saved reports are	137
Data export in the DataFabric Manager server	139
How to access DataFabric Manager server data	139
Where to find the database schema for the views	140
Two types of data for export	140
Files and formats for storing exported data	141
Format for exported the DataFabric Manager server data	141
Format for exported Performance Advisor data	141
Format for last updated timestamp	142
Security configurations	143
Types of certificates in the DataFabric Manager server	143
Self-signed certificates in the DataFabric Manager server	143
Trusted Certificate Authority (CA)-signed certificates in the DataFabric Manager server	143
Creating self-signed certificates in the DataFabric Manager server	144
Obtaining a trusted CA-signed certificate	144
Enabling HTTPS	145

Secure communications with the DataFabric Manager server	145
How clients communicate with the DataFabric Manager server	146
SecureAdmin for secure connection with the DataFabric Manager server clients	146
Requirements for security options	146
Guidelines to configure security options in Operations Manager	147
Managed host options	147
Where to find managed host options	147
Guidelines for changing managed host options	147
Comparison between global and storage system-specific managed host options	148
Limitations in managed host options	149
Changing password for storage systems in the DataFabric Manager server	149
Changing passwords on multiple storage systems	150
Issue with modification of passwords for storage systems	150
Using hosts.equiv to control authentication in the DataFabric Manager server	150
Editing HTTP and monitor services to run as different user	152
What File Storage Resource Management does	153
How FSRM monitoring works	153
What capacity reports are	153
Difference between capacity reports and file system statistics	154
Prerequisites for FSRM	154
Setting up FSRM	155
DataFabric Manager Host Agent software overview	156
DataFabric Manager Host Agent communication	156
DataFabric Manager Host Agent software passwords for monitoring tasks	156
DataFabric Manager Host Agent software passwords for administration tasks	157
Managing host agents	157
Host Agent management tasks	157
Configuring host agent administration settings	159
Enabling administration access for one or more host agents	160
Enabling administration access globally for all host agents	160
What FSRM paths are	160
Adding CIFS credentials	161

Path management tasks	161
Adding SRM paths	162
Path names for CIFS	163
Conventions for specifying paths from the CLI	163
Viewing file-level details for a path	163
Viewing directory-level details for a path	164
Editing SRM paths	164
Deleting SRM paths	164
Automatically mapping SRM path	164
What path walks are	165
What File SRM reports are	165
Access restriction to file system data	166
Identification of oldest files in a storage network	166
FSRM prerequisites	167
Verifying administrative access for using FRSM	167
Verifying host agent communication	167
Creating a new group of hosts	168
Adding an FSRM path	168
Adding a path-walk schedule	168
Grouping the FSRM paths	169
Viewing a report that lists the oldest files	169
User quotas	171
Why you use quotas	171
Overview of the quota process	171
User quota management using Operations Manager	171
Prerequisites for managing user quotas using Operations Manager	172
Where to find user quota reports in Operations Manager	173
Monitor interval for user quotas in Operations Manager	173
Modification of user quotas in Operations Manager	173
Prerequisites to edit user quotas in Operations Manager	174
Editing user quotas using Operations Manager	174
Configuring user settings using Operations Manager	174
What user quota thresholds are	175
What the DataFabric Manager server user thresholds are	175
User quota thresholds	175
Ways to configure user quota thresholds in Operations Manager	175

Precedence of user quota thresholds in the DataFabric Manager server	176
Management of LUNs, Windows and UNIX hosts, and FCP targets ..	177
Management of SAN components	177
SAN and DataFabric Manager Host Agent software	178
List of tasks performed using DataFabric Manager Host Agent software	178
List of tasks performed to monitor targets and initiators	178
Prerequisites to manage targets and initiators	179
Prerequisites to manage SAN hosts	179
Reports for monitoring LUNs, FCP targets, and SAN hosts	180
Information available on the LUN Details page	181
Tasks performed from the LUN Details page	182
Information about the FCP Target Details page	182
Information about the Host Agent Details page	183
List of tasks performed from the Host Agent Details page	183
How storage systems, SAN hosts, and LUNs are grouped	184
Granting access to storage systems, SAN hosts, and LUNs	184
Introduction to deleting and undeleting SAN components	185
Deleting a SAN component	185
How a deleted SAN component delete is restored	185
Where to configure monitoring intervals for SAN components	185
File system management	187
Access to storage-related reports	187
Storage capacity thresholds in Operations Manager	187
Modification of storage capacity thresholds settings	187
Changing storage capacity threshold settings for global group	188
Changing storage capacity threshold settings for an individual group	188
Changing storage capacity threshold settings for a specific aggregate, volume, or qtree	188
Management of aggregate capacity	189
Volume space guarantees and aggregate overcommitment	189
Available space on an aggregate	189
Considerations before modifying aggregate capacity thresholds	190
Aggregate capacity thresholds and their events	190
Management of volume capacity	194
Volume capacity thresholds and events	194
Normal events for a volume	199

Modification of the thresholds	199
Management of qtree capacity	199
Volume Snapshot copy thresholds and events	199
Qtree capacity thresholds and events	201
How Operations Manager monitors volumes and qtrees on a vFiler unit	202
How Operations Manager monitors qtree quotas	203
Where to find vFiler storage resource details	203
What clone volumes are	203
Identification of clones and clone parents	204
Why Snapshot copies are monitored	204
Snapshot copy monitoring requirements	204
Detection of Snapshot copy schedule conflicts	204
Dependencies of a Snapshot copy	205
Storage chargeback reports	205
When is data collected for storage chargeback reports	206
Determine the current month's and the last month's values for storage chargeback report	206
Chargeback reports in various formats	206
The chargeback report options	207
Specifying storage chargeback options at the global or group level	208
The storage chargeback increment	208
Currency display format for storage chargeback	208
Specification of the annual charge rate for storage chargeback	209
Specification of the Day of the Month for Billing for storage chargeback .	209
Formatted charge rate for storage chargeback	210
What happens when storage objects are deleted	210
Reports of deleted storage objects	210
Undeleting a storage object for monitoring	211
Storage system management	213
Management tasks performed using Operations Manager	213
Operations Manager components for managing your storage system	213
Storage system groups	214
Custom comment fields in Operations Manager	214
Consolidated storage system and vFiler unit data and reports	215
Tasks performed by using the storage systems and vFiler unit report pages	215

Where to find information about a specific storage system	216
Tasks performed from a Details page of Operations Manager	216
Editable options for storage system or vFiler unit settings	216
What Storage Controller Tools list is	217
What Cluster Tools list is	218
What the Diagnose Connectivity tool does	218
The Refresh Monitoring Samples tool	218
The Run a Command tool	218
The Run Telnet tool	219
Console connection through Telnet	219
Managing active/active configuration with the DataFabric Manager server	220
Viewing the status of an active/active configuration by using cluster console	220
What the Takeover tool does	220
What the Giveback tool does	221
DataFabric Manager server CLI to configure storage systems	222
Remote configuration of a storage system	222
Prerequisites for running remote CLI commands from Operations Manager	222
Running commands on a specific storage system	223
Running commands on a group of storage systems from Operations Manager	223
Remote configuration of a cluster	223
Running commands on a specific cluster	224
Running commands on a specific node of a cluster	224
Storage system management using FilerView	225
What FilerView is	225
Configuring storage systems by using FilerView	225
Introduction to MultiStore and vFiler units	226
Why monitor vFiler units with the DataFabric Manager server	226
Prerequisites for monitoring vFiler units	226
vFiler unit management tasks	227
Configuration of storage systems	229
Management of storage system configuration files	229
Prerequisites to apply configuration files to storage systems and vFiler units	229

List of access roles to manage storage system configuration files	230
List of tasks for configuration management	230
What configuration files are	231
What a configuration plug-in is	231
Comparison of configurations	231
Verification of a successful configuration push	231
Configuring the DataFabric Manager server to monitor configuration changes	231
What a configuration resource group is	232
List of tasks for managing configuration groups	232
Considerations when creating configuration groups	233
Creating configuration resource groups	233
Parent configuration resource groups	233
Configuring multiple storage systems or vFiler units	235
Backup Manager	237
Backup management deployment scenario	237
System requirements for backup	238
What backup scripts do	238
What the Backup Manager discovery process is	238
Methods of storage system discovery	239
What SnapVault relationship discovery is	239
New directories for backup	239
Viewing directories that are not backed up	239
SnapVault services setup	240
Configuring the SnapVault license	240
Enabling NDMP backups	241
Management of SnapVault relationships	241
Adding secondary storage systems	241
Adding secondary volumes	242
Adding primary storage systems	242
Selecting primary directories or qtrees for backup	243
What backup schedules are	243
Best practices for creating backup relationships	243
Snapshot copies and retention copies	244
Requirements to create a backup schedule	244
Creating backup schedules	244

Local data protection with Snapshot copies	245
Snapshot copy schedule interaction	245
Management of discovered relationships	245
Enabling the DataFabric Manager server to manage discovered relationships	246
What lag thresholds are	246
Setting global thresholds	246
Setting local thresholds	247
Bandwidth limitation for backup transfers	247
Configuring backup bandwidth	247
List of CLI commands to configure SnapVault backup relationships	248
Primary directory format	250
Secondary volume format	251
Disaster Recovery Manager	253
Prerequisites for using Disaster Recovery Manager	253
Tasks performed by using Disaster Recovery Manager	254
What a policy is	254
What a replication policy does	254
What a failover policy does	255
Policy management tasks	256
Connection management	257
Connection management tasks	257
What the connection describes	258
What multipath connections are	258
Authentication of storage systems	259
Authentication of discovered and unmanaged storage systems	259
Addition of a storage system	259
Modification of NDMP credentials	259
Deletion of a storage system	260
Volume or qtree SnapMirror relationships	260
Decisions to make before adding a new SnapMirror relationship	260
Addition of a new SnapMirror relationship	261
Modification of an existing SnapMirror relationship	261
Modification of the source of a SnapMirror relationship	261
Reason to manually update a SnapMirror relationship	261
Termination of a SnapMirror transfer	262

SnapMirror relationship quiescence	262
View of quiesced SnapMirror relationships	262
Resumption of a SnapMirror relationship	262
Disruption of a SnapMirror relationship	262
View of a broken SnapMirror relationship	263
Resynchronization of a broken SnapMirror relationship	263
Deletion of a broken SnapMirror relationship	263
What lag thresholds for SnapMirror are	263
Where to change the lag thresholds	264
Lag thresholds you can change	264
Reasons for changing the lag thresholds	264
What the job status report is	264
Maintenance and management	265
Accessing the CLI	265
Where to find information about the DataFabric Manager server commands	265
Audit logging	266
Events audited in the DataFabric Manager server	266
Global options for audit log files and their values	267
Format of events in an audit log file	267
Permissions for accessing the audit log file	269
What the remote platform management interface is	269
RLM card monitoring in the DataFabric Manager server	269
Prerequisites for using the remote platform management interface	270
Scripts overview	270
Commands that can be used as part of the script	270
Package of the script content	271
What script plug-ins are	271
What the script plug-in directory is	271
What backup scripts do	271
Overview of the DataFabric Manager server database backup process	272
When to back up data	272
Where to back up data	272
Recommendations for disaster recovery	273
Backup storage and sizing	273
Limitation of Snapshot-based backups	273
Access requirements for backup operations	274

Changing the directory path for database backups	274
Starting database backup from Operations Manager	274
Scheduling database backups from Operations Manager	275
Specifying the backup retention count	275
Disabling database backup schedules	276
Listing database backups	276
Deleting database backups from Operations Manager	276
Displaying diagnostic information from Operations Manager	276
Exportability of a backup to a new location	277
What the restore process is	277
Restoring the database from the archive-based backup	277
Restoring the database from the Snapshot copy-based backup	277
Restoration of the database on different systems	278
Disaster recovery configurations	278
Disaster recovery using the N series Management Console data protection capability	279
Disaster recovery using SnapDrive	284
Troubleshooting in Operations Manager	285
AutoSupport in the DataFabric Manager server	285
Reasons for using AutoSupport	285
Types of AutoSupport messages in the DataFabric Manager server	285
Protection of private data by using AutoSupport	286
Configuring AutoSupport	286
DataFabric Manager server logs	287
Access to logs	287
Accessing the logs through the DataFabric Manager server CLI	287
Apache and Sybase log rotation in the DataFabric Manager server	287
Common DataFabric Manager server problems	287
Communication issues between the DataFabric Manager server and routers	288
E-mail alerts not working in the DataFabric Manager server	288
How discovery issues are resolved	288
Use of the Diagnose Connectivity tool for a managed storage system	289
Use of the Diagnose Connectivity tool for an unmanaged storage system ..	289
Where to find the Diagnose Connectivity tool in Operations Manager	289

Reasons why the DataFabric Manager server might not discover your network	289
Troubleshooting network discovery issues	290
Troubleshooting appliance discovery issues with Operations Manager	291
How configuration push errors are resolved	291
How File Storage Resource Manager (FSRM) issues are resolved	292
Issues related to SAN events	292
Offline LUNs	292
Snapshot copy of LUN not possible	293
High traffic in HBA Port	293
Import and export of configuration files	293
How inconsistent configuration states are fixed	294
Data ONTAP issues impacting protection on vFiler units	294
Appendix	297
Report fields and performance counters	297
Report fields and performance counters for filer catalogs	297
Report fields and performance counters for vFiler catalogs	299
Report fields and performance counters for volume catalogs	300
Report fields and performance counters for qtree catalogs	302
Report fields and performance counters for LUN catalogs	302
Report fields and performance counters for aggregate catalogs	303
Report fields and performance counters for disk catalogs	303
Protocols and port numbers	304
DataFabric Manager server communication	304
The DataFabric Manager server access to storage systems	304
The DataFabric Manager server access to host agents	305
The DataFabric Manager server access to Open Systems SnapVault agents	305
SAN management	305
Discovery of SAN hosts by the DataFabric Manager server	306
SAN management using the DataFabric Manager server	306
Reports for monitoring SANs	309
DataFabric Manager server options	316
DataFabric Manager server options for SAN management	316
How SAN components are grouped	319
Glossary	321

Copyright information	325
Trademark information	327
Index	329

Introduction to Operations Manager

Operations Manager is a Web-based UI of the DataFabric Manager server.

You can use Operations Manager for the following day-to-day activities on storage systems:

- Discover storage systems
- Monitor the device or the object health, the capacity utilization, and the performance characteristics of a storage system
- View or export reports
- Configure alerts and thresholds for event managements
- Group devices, vFiler units, host agents, volumes, qtrees, and LUNs
- Run Data ONTAP CLI commands simultaneously on multiple systems
- Configure role-based access control (RBAC)
- Manage host users, user groups, domain users, local users, and host roles

Note: DataFabric Manager server 3.8 and later supports not only IPv4, but also IPv6.

What DataFabric Manager server does

The DataFabric Manager server provides infrastructure services such as discovery, monitoring, role-based access control (RBAC), auditing, and logging for products in the Storage and Data suites.

You can script commands using the command-line interface (CLI) of the DataFabric Manager server software that runs on a separate server. The software does not run on the storage systems.

What a license key is

To use the DataFabric Manager server, you must enable the OnCommand Core Package license by using the license key. The license key is a character string that is supplied by IBM.

If you are installing the software for the first time, you enter the license key during installation. You can enter the license key in the Options window under Licensed Features. You must enable additional licenses to use other features, such as disaster recovery and backup.

Access to Operations Manager

You can access the Operations Manager and the CLI from the IP address or DNS name of the DataFabric Manager server.

After successfully installing the DataFabric Manager server software, the DataFabric Manager server starts discovering, monitoring, collecting, and saving information about objects in its database.

Objects are entities such as storage systems; the vFiler units, disks, aggregates, volumes, and qtrees on these storage systems; LUNs; and user quotas.

In the case of the server on Windows, Operations Manager launches automatically and a welcome page appears.

Use either of the following URLs to access Operations Manager:

`http://[server_ip_address]:8080/dfm`

`http://server_dnsname:8080`

Depending on your Domain Name System (DNS) setup, you might need to use the fully qualified name in the second URL; for example, use `tampa.florida.com` instead of `tampa`.

Information to customize in the DataFabric Manager server

You can use the DataFabric Manager server to configure storage system IP addresses or names, administrator access control, and alarms, set up SNMP communities and administrator accounts and create groups.

DataFabric Manager server 3.8 and later supports IPv6 along with IPv4. However, the following DataFabric Manager server features lack IPv6 support:

- LUN management
- Snapshot-based backups (because SnapDrive for Windows and SnapDrive for UNIX do not support IPv6 addressing)
- Disaster recovery
- High Availability (HA) over Veritas Cluster Servers (VCS)
- "hosts.equiv" file based authentication
- APIs over HTTPS do not work for storage systems managed using IPv6 addresses, when the option `httpd.admin.access` is set to a value other than `legacy`.
- Discovery of storage systems and host agents that exist on remote network
- Protocols such as RSH and SSH do not support IPv6 link local address to connect to storage systems and host agents.

Note: Link local address works with SNMP and ICMP only.

Administrator accounts on the DataFabric Manager server

You can use Operations Manager to set up administrator accounts on the DataFabric Manager server. You can grant capabilities such as read, write, delete, backup, restore, distribution, and full control to administrators.

The DataFabric Manager server software provides the following two different administrator accounts:

- Administrator—grants full access for the administrator who installed the software
- Everyone—allows users to have read-only access without logging in

Related concepts

How roles relate to administrators on page 49

Authentication methods on the DataFabric Manager server

The DataFabric Manager server uses the information available in the native operating system for authentication. The server does not maintain its own database of the administrator names and the passwords.

You can also configure the DataFabric Manager server to use Lightweight Directory Access Protocol (LDAP). If you configure LDAP, then the server uses it as the preferred method of authentication.

Authentication with native operating system

You do not need to configure any options to enable the DataFabric Manager server to use the native operating system for authentication.

Based on the native operating system, the DataFabric Manager server application supports the following authentication methods:

- For Windows: local and domain authentication
- For UNIX: local password files, and NIS or NIS+

Note: Ensure that the administrator name you are adding matches the user name specified in the native operating system.

Authentication with LDAP

You can enable LDAP authentication on the DataFabric Manager server and configure it to communicate with your LDAP servers in order to retrieve relevant data.

The DataFabric Manager server provides predefined templates for the most common LDAP server types. These templates provide predefined LDAP settings that make the DataFabric Manager server compatible with your LDAP server.

Preface

Supported features

IBM System Storage N series storage systems are driven by NetApp Data ONTAP software. Some features described in the product software documentation are neither offered nor supported by IBM. Please contact your local IBM representative or reseller for further details.

Information about supported features can also be found on the N series support website (accessed and navigated as described in [Websites](#) on page 23).

Websites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. The following web pages provide N series information:

- A listing of currently available N series products and features can be found at the following web page:
www.ibm.com/storage/nas/
- The IBM System Storage N series support website requires users to register in order to obtain access to N series support content on the web. To understand how the N series support web content is organized and navigated, and to access the N series support website, refer to the following publicly accessible web page:
www.ibm.com/storage/support/nseries/
This web page also provides links to AutoSupport information as well as other important N series product resources.
- IBM System Storage N series products attach to a variety of servers and operating systems. To determine the latest supported attachments, go to the IBM N series interoperability matrix at the following web page:
www.ibm.com/systems/storage/network/interophome.html
- For the latest N series hardware product documentation, including planning, installation and setup, and hardware monitoring, service and diagnostics, see the IBM N series Information Center at the following web page:
publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Getting information, help, and service

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains

information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your IBM N series product, and whom to call for service, if it is necessary.

Before you call

Before you call, make sure you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure they are connected.
- Check the power switches to make sure the system is turned on.
- Use the troubleshooting information in your system documentation and use the diagnostic tools that come with your system.
- Refer to the N series support website (accessed and navigated as described in [Websites](#) on page 23) for information on known problems and limitations.

Using the documentation

The latest versions of N series software documentation, including Data ONTAP and other software products, are available on the N series support website (accessed and navigated as described in [Websites](#) on page 23).

Current N series hardware product documentation is shipped with your hardware product in printed documents or as PDF files on a documentation CD. For the latest N series hardware product documentation PDFs, go to the N series support website.

Hardware documentation, including planning, installation and setup, and hardware monitoring, service, and diagnostics, is also provided in an IBM N series Information Center at the following web page:

publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Hardware service and support

You can receive hardware service through IBM Integrated Technology Services. Visit the following web page for support telephone numbers:

www.ibm.com/planetwide/

Firmware updates

IBM N series product firmware is embedded in Data ONTAP. As with all devices, it is recommended that you run the latest level of firmware. Any firmware updates are posted to the N series support website (accessed and navigated as described in [Websites](#) on page 23).

Note: If you do not see new firmware updates on the N series support website, you are running the latest level of firmware.

Verify that the latest level of firmware is installed on your machine before contacting IBM for technical support.

How to send your comments

Your feedback helps us to provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, please send them by e-mail to starpubs@us.ibm.com.

Be sure to include the following:

- Exact publication title
- Publication form number (for example, GC26-1234-02)
- Page, table, or illustration numbers
- A detailed description of any information that should be changed

What the discovery process is

The DataFabric Manager server discovers all the storage systems in your organization's network by default. You can add other networks to the discovery process or enable discovery on all the networks. Depending on your network setup, you can disable discovery entirely. You can disable auto-discovery if you do not want SNMP network walking.

Discovery by the DataFabric Manager server

The DataFabric Manager server depends on Simple Network Management Protocol (SNMP) to discover and periodically monitor storage systems.

If your storage systems is not SNMP-enabled, you must enable it before the server can discover them. You can enable SNMP on storage systems by using either FilerView or the Data ONTAP CLI.

If the routers, switches, or storage systems use SNMP communities other than “public,” you must specify the appropriate communities on the Edit Network Credentials page.

The server needs to locate and identify storage systems so that it can add them to its database. The server can monitor and manage only systems and networks that are in the database.

Automatic discovery is typically the primary process the server uses to discover storage systems and networks. In this process, the server and the systems (storage systems, vFiler units, and Host Agents) communicate automatically with each other.

Manual addition is secondary to the discovery process. You typically only need it for the storage systems and the networks that you add after the server discovers the infrastructure.

What SNMP is

Simple Network Management Protocol (SNMP) is an application-layer protocol that facilitates the exchange of management information between network devices.

SNMP is part of the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol suite. SNMP enables network administrators to manage network performance; find and solve network problems; and plan for network growth.

When to enable SNMP

You must enable SNMP on your storage systems before you install the DataFabric Manager server, if you want the DataFabric Manager server to discover the storage systems immediately.

SNMP is normally enabled on storage systems. You can verify this with `storage01> options snmp.enable`.

If SNMP is not enabled on your storage system, enable it with `Storage01> options snmp.enable on`. SNMPv1 uses community strings for authentication. Storage systems normally allow read-only access with the 'public' community.

You can also wait until after installing the software to enable SNMP on storage systems. However, this causes a delay in the server to discover the storage systems.

SNMP versions to discover and monitor storage systems

The DataFabric Manager server uses the SNMP protocol versions to discover and monitor the storage systems.

By default, the DataFabric Manager server uses SNMPv1, with public as the community string to discover the storage systems. To use a specific configuration on a network, you must add the networks required.

SNMPv1 is a widely used simple request/response protocol. SNMPv3 is an interoperable standards-based protocol with security and remote configuration capabilities.

SNMPv3 provides user-based security with separate authentication and authorization. It is a method to specify common credentials.

Note: SNMPv3 support is available only on storage systems running Data ONTAP 7.3 or later.

You can use SNMPv3 to discover and monitor storage systems if SNMPv1 is disabled.

Note: The user on the storage system whose credentials are specified in Operations Manager should have the login-snmp capability to be able to use SNMPv3.

The version specified in the Preferred SNMP Version option at the storage system level is used for monitoring the discovered storage system. If no version is specified at the storage system level, then either the network setting or the global setting is used. However, you can modify the SNMP version, if required.

Note: If the monitoring fails using the specified SNMP version, then the other SNMP version is not used for storage system monitoring.

Related concepts

[Methods of adding storage systems and networks](#) on page 37

Related tasks

[Modifying the network credentials and SNMP settings](#) on page 32

Related references

[Guidelines for editing discovery options](#) on page 37

What the Preferred SNMP Version option is

The Preferred SNMP Version option is a global or network-specific option that specifies the SNMP protocol version to be used first for discovery.

You can use Operations Manager to configure the option with values such as SNMPv1 or SNMPv3.

SNMP version setup

You must know of the settings that are used for the SNMP version preferred at the storage system level or network level.

If the SNMP version is...	Then...
Specified at the storage system level	The version preferred takes precedence over the network and global settings.
Not specified at the storage system level	Network setting is used.
Not specified at the network level	Global setting is used.

When the DataFabric Manager server is installed for the first time or updated, by default, the global and network setting uses SNMPv1 as the preferred version. However, you can configure the global and network setting to use SNMPv3 as the default version.

Related tasks

[Modifying the network credentials and SNMP settings](#) on page 32

How the DataFabric Manager server chooses network credentials for discovery

This table shows how the DataFabric Manager server chooses the network credentials for discovery.

If...	Then...
The discovery is running on a particular network and the network credentials are configured	The network credentials configured for that particular network are used for discovery.
No network exists	The network credentials configured as global settings are used for discovery.

Discovery process using SNMPv1 or SNMPv3

This table describes the discovery process for a storage system by using SNMPv1 or SNMPv3.

If...	Then...
The storage system is discovered using the preferred SNMP version (let us say, SNMPv1)	The discovered storage system is added with the preferred SNMP version as Global/Network Default. This implies that the network or global settings are used for monitoring.
The storage system is not discovered using SNMPv1	SNMPv3 is used for storage system discovery.
The discovery succeeds using SNMPv3	SNMPv3 is set as the preferred version for monitoring.

When all or most of the storage systems in a network are running only a particular SNMP version, then you are recommended to specify only that version as the preferred SNMP version for the network. This speeds up the discovery of storage systems running only a particular SNMP version.

You can prevent using a particular version of SNMP from being used for discovery. For example, if a particular version of SNMP is not in use in the network, then you can disable that SNMP version. This speeds up the discovery process.

Monitoring process using SNMPv1

This table shows how storage systems are monitored using SNMPv1.

If...	Then...
The Preferred SNMP Version option is set to SNMPv1, or the Preferred SNMP Version option is not set for the storage system, and the global or network setting is SNMPv1	The community string set at the network level is used for the SNMPv1 monitoring.
The community string is not specified at either global or network level	SNMPv1 is disabled and an event is generated to indicate the SNMP communication failure with the storage system.

Monitoring process using SNMPv3

This table shows how storage systems are monitored using SNMPv3.

If...	Then...
The Preferred SNMP Version option is set to SNMPv3, or the Preferred SNMP Version option is not set for the storage system, and the global or network setting is SNMPv3	The login and the password specified for the storage system are used for the SNMPv3 monitoring.
The storage system credentials are not specified	The login and the password specified at the network level are used for the SNMPv3 monitoring.
No credentials are provided at the network level	The login and the password specified at the global level are used for the SNMPv3 monitoring.
No credentials are provided at the global level	An event is generated to indicate the SNMP communication failure with the storage system.

The DataFabric Manager server supports SNMPv3 communication through auth protocols: MD5 and SHA. You can configure SNMPv3 settings with either of the auth protocols from the Network Credentials page or from the CLI.

Setting SNMPv1 or SNMPv3 as the preferred version

You can set SNMPv1 or SNMPv3 as the preferred version for storage system discovery on a specific network.

Steps

1. Select the Network Credentials submenu from the Setup menu.

Alternatively, select the Discovery submenu from the Setup menu and click the edit link corresponding to the Network Credentials option.

2. Provide values for each of the parameters requested.
3. Click **Add**.

Setting SNMPv1 as the only SNMP version

You can set SNMPv1 as the only SNMP version available to monitor all storage systems in a network.

Steps

1. Go to the **Network Credentials** page.

2. Click the **edit** link corresponding to the Edit field for the SNMPv3 enabled network.
3. In the Edit Network Credentials section, modify the value of the Preferred SNMP Version option to SNMPv1.
4. In the SNMPv3 Settings section, clear the Login and Password values.
5. Click **Update**.
6. If the storage system in the network has the Preferred SNMP Version option set to SNMPv3, then
 - a. Go to the **Edit Appliance Settings** page of the corresponding storage system.
 - b. Modify the value of the Preferred SNMP Version option to `Global/Network Default`.

Setting SNMPv1 or SNMPv3 to monitor a storage system

You can set SNMPv1 or SNMPv3 to monitor a storage system.

Steps

1. Click **Control Center > Member Details > Physical Systems > Storage Systems, All**.

A list of all the storage systems that are monitored by Operations Manager console is displayed.
2. Click the storage system for which you want to set the SNMP version.
3. From the **Edit Storage Controller Settings** page, click **Edit Settings** under Storage Controller in the left pane.
4. Select the preferred SNMP version option from the drop-down menu.
5. Click **Update**.

Modifying the network credentials and SNMP settings

You can modify the network credentials and SNMP settings using Operations Manager.

Steps

1. Select the **Network Credentials** submenu from the **Setup** menu.

Alternatively, select the **Discovery** submenu from the **Setup** menu and click the **edit** link corresponding to the `Network Credentials` option.
2. Click the **edit** link corresponding to the Edit field in the **Network Credentials** page.
3. Modify values for the parameters required.
4. Click **Update**.

Deleting the SNMP settings for the network

You can delete the SNMP settings for the network using Operations Manager.

Steps

1. Click **Setup > Network Credentials**.

The Network Credentials page displays the list of networks

2. Select the check box corresponding to the **Delete** field for the required network.
3. Click **Delete Selected**.

Addition of a storage system from an undiscovered network

You can add a single storage system to the DataFabric Manager server from an undiscovered network on which only SNMPv3 is enabled.

You can add the storage system by running the `dfm host add -N` command, with the appropriate values for the following storage system credentials:

- hostLogin
- hostPassword

In this case, the discovery is not enabled on the storage system's network.

Diagnosis of SNMP connectivity

You can diagnose the SNMP connectivity with a host by running the Diagnose Connectivity tool from Operations Manager.

You can access the Diagnose Connectivity tool from the Storage Controller Tools list located at the lower left of Operations Manager.

Alternatively, you can run the `dfm host diaghostname` command to diagnose the DataFabric Manager server's connectivity using SNMPv1 and SNMPv3 with a host. The credentials used for diagnosing connectivity using rsh and ssh are the host credentials. However, if host credentials are unspecified, then the network or global credentials are used for SNMPv3.

Related concepts

[Use of the Diagnose Connectivity tool for a managed storage system](#) on page 289

[Use of the Diagnose Connectivity tool for an unmanaged storage system](#) on page 289

[Where to find the Diagnose Connectivity tool in Operations Manager](#) on page 289

[Reasons why the DataFabric Manager server might not discover your network](#) on page 289

What host discovery is

The DataFabric Manager server automatically discovers storage systems and Host Agents that are in the same subnet as the server. When you install the DataFabric Manager server software, the Host Discovery option is enabled by default.

The discovery of networks, when enabled, is integrated with the discovery of storage systems and Host Agents. The discovery occurs at the same time.

Note: The DataFabric Manager server 3.8 supports discovery of IPv6 networks and hosts.

Ping methods in host discovery

The DataFabric Manager server uses SNMP queries for host discovery. You must enable SNMP on your storage systems and the routers for the DataFabric Manager server to monitor and manage systems.

By default, SNMP is enabled on storage systems.

Ping methods might include ICMP echo, HTTP, NDMP, or ICMP echo and SNMP. The latter ping method does not use HTTP to ping a host. Therefore, if a storage system (behind a transparent HTTP cache) is down and the HTTP cache responds, the server does not mistake the storage system to be running. The ICMP echo and SNMP ping method is the default for new installations.

Note: When you select ICMP echo and the SNMP ping method, the server uses ICMP echo first, and then SNMP, to determine if the storage system is running.

What host-initiated discovery is

Host-initiated discovery is based on the DNS SRV record, where the DataFabric Manager server details are maintained. Currently, host-initiated discovery is supported by DataFabric Manager Host Agent only.

Whenever a host initiates communication with the DataFabric Manager server, it does not identify the host until its network (IPv6) address details are added. You can add the host IPv6 network to the DataFabric Manager server by using the `dfm network add` command. When it receives this command, the host initiates a request to the DataFabric Manager server. After the DataFabric Manager server identifies the network, the host is added to the DataFabric Manager server host list.

For information about how you can modify the DataFabric Manager server details for host-initiated discovery, see the *DataFabric Manager Host Agent Installation and Administration Guide*.

Related information

DataFabric Manager Host Agent Installation and Administration Guide - www.ibm.com/storage/support/nseries/

Discovery of vFiler units

The OnCommand console DataFabric Manager server monitors the hosting storage systems to discover vFiler units. You must set authentication credentials for the hosting storage system to ensure that DataFabric Manager server discovers the vFiler units.

The server monitors the hosting storage system once every hour to discover new vFiler units that you configured on the storage system. The server deletes from the database the vFiler units that you destroyed on the storage system.

You can change the default monitoring interval from the Monitoring setup options, or by using the following CLI command:

```
dfm option set vFilerMonInterval=1hour
```

You can disable the vFiler discovery from the Discovery setup options, or by using the `dfm option set discovervfilers=no` CLI command.

When DataFabric Manager server discovers a vFiler unit, it does not add the network to which the vFiler unit belongs to its list of networks on which it runs host discovery. In addition, when you delete a network, the server continues to monitor the vFiler units in that network.

Related tasks

[Changing password for storage systems in the DataFabric Manager server](#) on page 149

[Changing passwords on multiple storage systems](#) on page 150

Discovery of storage systems

You must follow the process that the DataFabric Manager server uses to discover storage systems if the `Host Discovery` option is enabled and the `Network Discovery` option is disabled (the default value).

1. The DataFabric Manager server issues an SNMP GET request to all storage systems on the local network. The purpose of the request is to determine the identity of the storage systems. The local network is the network to which the DataFabric Manager server is attached.

2.

If...	Then...
-------	---------

The SNMP GET request is successful	<p>The DataFabric Manager server adds the discovered storage systems to its database and continues to step 3.</p> <p>If the storage system is a hosting storage system on which vFiler units are configured, the DataFabric Manager server also discovers those vFiler units.</p> <p>Note: vFiler units are discovered only after you set the credentials for the hosting storage system.</p>
------------------------------------	--

- The DataFabric Manager server repeats steps 1 through 2 until it has sent queries to all the networks in its database. The minimum interval for repeating the cycle is set by the Discovery Interval (the default is every 15 minutes) and the Discovery Timeout (the default is 2 seconds). The actual interval depends on the number of networks to scan and their size.

Note: The DataFabric Manager server repeats steps 1 to 2 to discover new storage systems. The minimum interval for repeating the discovery process is set by the Discovery Interval option.

Discovery of storage systems and networks

This table describes the process that the DataFabric Manager server uses to discover storage systems and networks if both the Host Discovery and Network Discovery options are enabled.

Stage	Description	
1.	The DataFabric Manager server issues an SNMP GET request to all hosts on the local network. The purpose of the request is to determine the system identity of the hosts. The local network is the network to which the DataFabric Manager server is attached.	
2.	If...	Then...
	The SNMP GET request is successful	<p>The DataFabric Manager server adds the discovered storage systems to its database and continues to Stage 4.</p> <p>If the storage system is a hosting storage system on which vFiler units are configured, the DataFabric Manager server also discovers those vFiler units.</p> <p>Note: vFiler units will be discovered only after you set the credentials for the hosting storage system.</p>

Stage	Description
3.	The DataFabric Manager server issues another SNMP GET request to routers that responded to the first SNMP request. The purpose of the request is to gather information about other networks to which these routers might be attached.
4.	When the DataFabric Manager server receives replies, if it finds networks that are not included in its database, it adds the new networks to its database.
5.	The DataFabric Manager server selects another network from its database and issues an SNMP GET request to all hosts on that network.
6.	The DataFabric Manager server repeats Stages 2 through 5 until it has sent SNMP queries to all the networks in its database. By default, the minimum interval for repeating the network discovery cycle is set at every 15 minutes.

Methods of adding storage systems and networks

You can apply a combination of methods to efficiently add storage systems and networks to the DataFabric Manager server.

- Keep the defaults for the discovery options of Host Discovery (Enabled), and Network Discovery (Disabled).
- Start the discovery process by manually adding one storage system from each network that has storage systems.
When you add a storage system, its network is added, too. Then other storage systems on the network are found automatically.
- After verifying that all of the storage systems have been added, you should disable host discovery to save network resources.
- After attaching a new storage system to your network, you must add hosts by using either Operations Manager or through the command line with the `dfm host add` command.
- If you set up a new network of storage systems, add one storage system so that its network and all other storage systems on it are found.

Guidelines for editing discovery options

You must follow a set of guidelines for changing the default values of the discovery options.

Interval This option specifies the period after which the DataFabric Manager server scans for new storage systems and networks.

You can change the default value if you want to increase the minimum time interval between system discovery attempts. This option affects the discovery interval only at the time of installation. After storage systems are discovered, the user should determine the interval based on the number of networks and their size.

If you choose a longer interval, there might be a delay in discovering new storage systems, but the discovery process is less likely to affect the network load.

The default is 15 minutes.

Timeout This option specifies the time interval after which the DataFabric Manager server considers a discovery query to have failed.

You can change the default value if you want to lengthen the time before considering a discovery to have failed (to avoid discovery queries on a local area network failing due to the long response times of a storage system).

The default is 5 seconds.

Host discovery This option enables the discovery of storage systems, host agents and vFiler Units through SNMP.

You can change the default value if any of the following situations exist:

- All storage systems that you expected the DataFabric Manager server to discover have been discovered and you do not want the DataFabric Manager server to continue scanning for new storage systems.
- You want to manually add storage systems to the DataFabric Manager server database.

Manually adding storage systems is faster than discovering storage systems in the following cases:

- You want the DataFabric Manager server to manage a small number of storage systems.
- You want to add a single new storage system to the DataFabric Manager server database.

Host agent discovery This option allows you to enable or disable discovery of host agents.

You can change the default value if you want to disable the discovery of LUNs or storage area network (SAN) hosts and host agents.

Network discovery This option enables the discovery of networks, including SAN and cluster networks.

You can change the default value if you want the DataFabric Manager server to automatically discover storage systems on your entire network.

Note: When the Network Discovery option is enabled, the list of networks on the Networks to Discover page can expand considerably as the DataFabric Manager server discovers additional networks attached to previously discovered networks.

Network Discovery Limit (in hops)	<p>This option sets the boundary of network discovery as a maximum number of hops (networks) from the DataFabric Manager server.</p> <p>You can change the default value if you want to increase this limit if the storage systems that you want the DataFabric Manager server to discover are connected to networks that are more than 15 hops (networks) away from the network to which the DataFabric Manager server is attached. The other method for discovering these storage systems is to add them manually.</p> <p>You can decrease the discovery limit if a smaller number of hops includes all the networks with storage systems you want to discover. For example, reduce the limit to six hops if there are no storage systems that must be discovered on networks beyond six hops. Reducing the limit prevents the DataFabric Manager server from using cycles to probe networks that contain no storage systems that you want to discover.</p> <p>The default is 15 hops.</p>
Networks to discover	<p>This option enables you to manually add or delete networks that the DataFabric Manager server scans for new storage systems.</p> <p>You can change the default value if you want to add a network to the DataFabric Manager server that it cannot discover automatically, or you want to delete a network for which you no longer want storage systems to be discovered.</p>
Network Credentials	<p>This option enables you to specify, change, or delete an SNMP community that the DataFabric Manager server uses for a specific network or host.</p> <p>You can change the default value if storage systems and routers that you want to include in the DataFabric Manager server do not use the default SNMP community.</p>

Discovery of a cluster by Operations Manager

The DataFabric Manager server automatically discovers a cluster that is in the same network as the DataFabric Manager server host. If the cluster is in a different network, you can discover it by specifying the appropriate SNMP version in the `Preferred SNMP Version` option in the Network Credentials page.

For a specified IP address and the SNMP version specified in the `Preferred SNMP Version` option, the following objects are queried:

- Name of the system (`sysName`)
- ID of the system object (`sysObjectId`)
- Cluster ID (`clusterIdentityUuid`)

If the query fails, you can use another version of SNMP to send queries. When the query succeeds, Operations Manager identifies the cluster based on the `sysObjectId`. For example, if the `sysObjectId` is `netappProducts.netappCluster`, the cluster is added to the network.

Note: For both SNMPv1 and SNMPv3, Operations Manager uses per-network configuration settings or the default network settings.

Adding a cluster

By using Operations Manager, you can add a cluster by specifying the IP address of the given cluster management logical interface.

About this task

Operations Manager uses the default SNMP version to discover a cluster. Otherwise, you can specify the appropriate SNMP version in the `Preferred SNMP Version` option in the Network Credentials page.

Steps

1. Click **Control Center > Home > Member Details > Physical Systems**.
2. In the **New storage system** field in the lower part of the page, enter the host name or IP address of the cluster you want to add.
3. Click **Add**.

Result

The cluster is added and displayed in the Clusters, All report.

Data ONTAP 8.0 cluster monitoring tasks using Operations Manager

Using Operations Manager, you can perform various monitoring tasks on cluster nodes running Data ONTAP 8.0 Cluster-Mode.

You can perform the following tasks on a cluster by using Operations Manager:

- Discover a cluster, either automatically or manually.
- Monitor the cluster.
- Configure the cluster to receive alerts.
- Generate various reports on the cluster and its components.
- Perform File SRM-related tasks.
- Execute commands remotely.

Limitations of cluster monitoring in Operations Manager

You cannot perform certain monitoring tasks related to a cluster due to the lack of support for certain APIs, SNMP objects, and interfaces in Data ONTAP 8.0 Cluster-Mode systems. However, the cluster monitoring feature is not applicable to systems running Data ONTAP 8.0 7-Mode.

The following features are not supported in Operations Manager:

- Management of the cluster administrator user profile and password.
- Configuration management of clusters, controllers, and virtual servers.
- Configuration of the high-availability checker script.
- Management of volume SnapMirror relationships, schedules, and jobs.
- Management of disks.
- Creation of SRM autopaths.
- Threshold monitoring of Performance Advisor.
- Monitoring of SAN and LUNs.
- Receiving SNMP traps from a cluster.
- SnapLock reports.
- Management of quotas.

Introduction to the gateway SAN-attached storage management

Operations Manager discovers storage arrays and storage array ports for the gateway through certain APIs. For discovery and monitoring of storage arrays and storage array ports, you must set the host login and password in the storage system.

You can perform the following tasks relating to SAN-attached storage of the gateway by using Operations Manager:

- Monitor storage arrays and storage array ports that are connected to the gateway.
- Monitor the storage load of array LUNs.
- Generate various reports on the gateway SAN-attached storage using back-end storage.

For more information about the gateway SAN-attached storage management reports, see the *Operations Manager Help*.

Related tasks

[Viewing configuration details of storage arrays connected to the gateway](#) on page 42

Limitations of the gateway SAN-attached storage management in Operations Manager

You cannot use Operations Manager to perform a high-level analysis of the average usage trend of a storage system's back-end storage, or the least used and most used array LUNs, aggregates, or adapters at a given point in time.

Operations Manager does not support the following features:

- Listing, adding, or deleting storage array and storage array port objects from the command-line interface
- Events for storage arrays and storage array ports
- RBAC security for storage arrays and storage array ports.
- A Details page for storage arrays and storage array ports

Tasks performed from the Storage Controller Details page for the gateway

You can view the list of storage arrays and the list of storage arrays ports connected to the gateway from the Storage Controller Details page. You can access this page by clicking the name of the gateway in the appropriate Storage System report.

You can perform the following tasks from the Storage Controller Details page for the gateway:

- View the list of storage arrays connected to the gateway.
You can view this list by clicking the number in the "Arrays connected to This Storage System" field.
- View the list of storage array ports connected to the gateway.
You can view this list by clicking the number in the "Array Ports connected to This Storage System" field.

Viewing configuration details of storage arrays connected to the gateway

You can view the configuration details of storage arrays connected to the gateway in the Storage Array Configuration report in Operations Manager.

Before you begin

Operations Manager must discover the storage arrays for which you want to view the configuration details.

Steps

1. Click **Control Center > Home > Member Details > Physical Systems**.
2. Select the Storage Array Configuration report from the Report drop-down menu.

Result

The Storage Array Configuration page is displayed. In this page, you can view configuration details, such as the name of the storage array, the name of the gateway, array LUN count, the adapter used, the name of the switch, and so on.

Role-based access control in the DataFabric Manager server

The DataFabric Manager server uses role-based access control (RBAC) for user login and role permissions.

What role-based access control is

The DataFabric Manager server uses role-based access control (RBAC) for user login and role permissions. If you have not changed the DataFabric Manager server's default settings for administrative user access, you do not need to log in to view information by using the DataFabric Manager server.

By default, you cannot view the DataFabric Manager server data. However, when you initiate an operation that requires specific privileges, the DataFabric Manager server prompts you to log in. For example, to create administrator accounts, you need to log in with Administrator account access.

RBAC allows administrators to manage groups of users by defining roles. If you need to restrict access to the database to specific administrators, you must set up administrator accounts for them. Additionally, if you want to restrict the information that these administrators can view and the operations they can perform, you must apply roles to the administrator accounts you create.

Configuring vFiler unit access control

An administrator who does not have any roles on a global level, but has enough roles on a group that contains only vFiler units is considered a vFiler administrator. The vFiler administrator does not have access to the host storage system information.

About this task

The following restrictions are applicable to vFiler units' administrators:

- If a vFiler unit has a volume assigned to it, the vFiler administrator cannot view details or reports for the aggregate in which the volume exists.
- If a vFiler unit has a qtree assigned to it, the vFiler administrator cannot view details or reports for the volume in which the qtree exists.

Note: The full name of a qtree contains a volume name (for example, 10.72.184.212:/hemzvol/hagar_root_backup_test) even though the vFiler unit does not contain the volume.

This procedure describes how to configure access control that allows an administrator to view and monitor vFiler units.

Steps

1. Create a group that contains vFiler units.
2. From the **Edit Group Membership** page, select vFiler units to add to the group.
3. From the **Roles** page, create a role for the vFiler administrator and assign it the following database operations: Delete, Read, and Write.
4. From the **Edit Administrator Settings** page, assign role to the vFiler administrator.

Logging in to the DataFabric Manager server

You can log in to the DataFabric Manager server by entering the administrator name and password on the Operations Manager interface.

Steps

1. From the **Control Center**, select **Log In**.
2. Type your administrator name and password.
3. Click **Log In**.

What default administrator accounts are

The DataFabric Manager server uses administrator accounts to manage access control and maintain security. When you install the DataFabric Manager server software, default administrator accounts are created: the “Administrator” and “Everyone” accounts. Administrator accounts have predefined roles assigned to them.

Administrator account The Administrator has super user privileges and can perform any operation in the DataFabric Manager server database and add other administrators. The Administrator account is given the same name as the name of the administrator who installed the software. Therefore, if you install the DataFabric Manager server on a Linux workstation, the administrator account is called root.

Everyone account After installing the DataFabric Manager server, you must log in as the Administrator and set up the Everyone account to grant view permission on this account. This is optional.

Note: Changes made will not be seen in the audit log.

List of predefined roles in the DataFabric Manager server

This table provides a list of roles to different administrator accounts in the DataFabric Manager server.

Administrator account	Roles
Administrator	<ul style="list-style-type: none"> • GlobalAlarm • GlobalBackup • GlobalConfigManagement • GlobalDataProtection • GlobalDataSet • GlobalDelete • GlobalEvent • GlobalExecute • GlobalFailover • GlobalFullControl • GlobalMirror • GlobalPerfManagement • GlobalProvisioning • GlobalQuota • GlobalRead • GlobalReport • GlobalResourceControl • GlobalRestore • GlobalSAN • GlobalSDConfig • GlobalSDDDataProtection • GlobalSDDDataProtectionAndRestore • GlobalSDFullControl • GlobalSDSnapshot • GlobalSDStorage • GlobalSRM • GlobalWrite
Everyone	<ul style="list-style-type: none"> • No roles

Active Directory user group accounts

The DataFabric Manager server recognizes two types of users namely Administrator and User, thereby allowing domain administrators the ability to define roles based on a company's organizational hierarchy.

To set up administrator accounts as a user group, use the following naming convention: <AD domain>\group_dfadmins.

In this example, all administrators who belong to group_dfadmins can log in to the DataFabric Manager server and inherit the roles specified for that group.

Adding administrative users

You can create administrator accounts from the Operations Manager console. Administrator accounts are either an individual administrator or a group of administrators.

Before you begin

The DataFabric Manager server user must be a local operating system user, or a domain user reachable by LDAP.

Steps

1. Log in to the Administrator account.
2. In the Operations Manager console, click **Setup > Administrative users**.
3. Type the name for the administrative user or domain name for the group of administrators.

When you add the user, they must be present locally.

4. If you have already created a role that you want to assign to this user or group of users, select the role in the left column of the displayed table and use the arrow button to move the role to the column on the right.

Roles in the column on the right are assigned to the user you are creating.

5. Type the e-mail address for the administrator or administrator group.
6. Enter the pager number for the administrator or administrator group.
7. Click **Add**.

After you finish

If you have not created a role for the user you created in this task, you must next create a role.

How roles relate to administrators

Role management allows the administrator who logs in with super-user access to restrict the use of certain DataFabric Manager server functions to other administrators.

The super-user can assign roles to administrators on an individual basis, by group, or globally (and for all objects in DataFabric Manager server).

An operation must be specified for every role. You can assign multiple operations levels if you want the administrator to have more control than a specific role provides. For example, if you want an administrator to perform both the backup and restore operations, you can create and assign to the administrator a single role that has both of these operations.

You can list the description of an operation by using the `dfm role operation list [-x] [<operation-name>]` command.

Related concepts

[What global and group access control is](#) on page 55

Related references

[Guidelines for editing discovery options](#) on page 37

What predefined global roles are

Administrators assigned with global roles can view information or configure settings for all groups in the DataFabric Manager server database, including the Global group.

The DataFabric Manager server provides the set of predefined global roles that can be inherited to the user creating roles as described in the following table:

Role	Operations
Default	None
GlobalAlarm	You can manage alarms. You can view, create, modify, or delete alarms.
GlobalBackup	You can create and manage backups.
GlobalConfigManagement	You can manage storage system configurations.
GlobalDataProtection	You can perform all the operations of GlobalBackup, GlobalRead, and GlobalDataSet.
GlobalDataSet	You can perform DataSet write and DataSet delete operations.

Role	Operations
GlobalDelete	You can delete information in the DataFabric Manager server database, including groups and members of a group, monitored objects, primary and secondary storage systems, and backup relationships, schedules, and retention policies.
GlobalEvent	You can view and acknowledge events in addition to create and delete alarms.
GlobalExecute	You can execute commands on storage system.
GlobalFailover	You can manage disaster recovery for datasets.
GlobalFullControl	You can view and perform any operation on any object in the DataFabric Manager server database and configure administrator accounts. You cannot apply this role to accounts with group access control.
GlobalMirror	You can create, destroy, and can update replication or failover policies.
GlobalPerfManagement	You can manage views, event thresholds, and alarms apart from viewing performance information in Performance Advisor.
GlobalProvisioning	You can provision primary dataset nodes and can attach resource pools to secondary or tertiary dataset nodes. You also have all the capabilities of the GlobalResourceControl, GlobalRead, and GlobalDataset roles for dataset nodes that are configured with provisioning policies.
GlobalQuota	You can view user quota reports and events.
GlobalRead	You can view the DataFabric Manager server database, backup configurations, events and alerts, and replication or failover policies.
GlobalReport	You can manage custom reports and report schedules.
GlobalResourceControl	You can add members to dataset nodes that are configured with provisioning policies.
GlobalRestore	You can perform restore operations from backups on secondary volumes.
GlobalSAN	You can create, expand, and destroy LUNs.
GlobalSDConfig	You can read, modify, and delete SnapDrive configuration.
GlobalSDDDataProtection	You can manage backups and datasets with SnapDrive.

Role	Operations
GlobalSDDataProtectionAndRestore	You can perform backup and restore operations with SnapDrive.
GlobalSDFullControl	You can perform operations specific to GlobalSDConfig, GlobalSDSnapshot, and GlobalSDStorage roles.
GlobalSDSnapshot	You can list the snapshots and the objects inside them. You can create, modify, and delete snapshots. You can create clones of volumes, luns, and qtrees. You can restore volumes, luns, and qtrees from snapshots.
GlobalSDStorage	You can list, create, modify and delete storage objects and their attributes.
GlobalSRM	You can view information collected by SRM path walks.
GlobalWrite	You can view or write to the DataFabric Manager server database.

Note: Super users are assigned the GlobalFullControl role in Operations Manager. For Linux, super user is the root user. For Windows, super-users belong to the administrators group.

What inheritance roles are

Administrators assigned with group roles can view or configure settings for the group to which they belong.

When you view roles for an administrator, the settings are those explicitly set for the administrator at the group level. For example, if administrators have the GlobalRead role, they implicitly have the Read role on all groups. Similarly, if administrators have the Read role on a parent group, they implicitly have the Read role on all the subgroups of that parent group.

Several other factors also affect the group role that is granted to an administrator:

- The capabilities granted to the administrator, "Everyone".
- The administrator's membership in Active Directory (AD) user groups that have been added to the DataFabric Manager server database.

Group roles are named similarly to the global roles that are defined in the previous table.

Note: Roles are carried forward prior to DataFabric Manager server 3.3.

What capabilities are

When creating roles, you must assign capabilities, a combination of operations and resources, to the role. You can view capabilities or edit them by modifying the operations that are associated with the resource. Resources can be groups of monitored objects, such as storage system and hosts.

Role precedence and inheritance

If an administrative user has both global and group roles on a group, the less restrictive (more powerful) of the two roles apply.

For example, if a user is assigned GlobalRead role and GlobalWrite role on a group, that user can view all groups. However, the user can change settings or run commands only on the storage systems of the specified group.

Role inheritance simplifies the task of assigning roles to administrators by letting you use defined roles. Specifying roles for a parent group implicitly grants those roles on its subgroups. You should grant roles conservatively at higher levels of the group hierarchy and specify additional roles as needed at the lower levels of the hierarchy.

Creating roles

You can create roles from the Setup menu in Operations Manager.

Steps

1. Select **Roles** from the **Setup** menu.
2. Click **Add Capabilities...** and, from the **Capabilities** window, select a resource from the resource tree.
3. Select the operations that you want to allow for the resource and click **OK**.
4. Optionally, to copy capabilities from an existing role, select that role from the **Inherit Capabilities** list and click \>> to move the role to the list at the right.
5. Click **Add Role**.

Modifying roles

You can edit the roles created from the Setup menu in Operations Manager.

Steps

1. Click **Setup > Roles**.
2. Find the role in the list of roles and click **edit**.
Note: You cannot modify the global administrative roles.
3. From the **Edit Role Settings** page, modify the basic settings of the role.

4. Click **Update**.
5. Modify role inheritance by doing one of the following:
 - To disinherit a role, select the role from the list at the right, and click “<<” to remove it.
 - To inherit a role, select the role from the “Inherit Capabilities” list and click “>>” to move the role to the Inherit Capabilities list.

Note: This step is optional.
6. Click **Update**.

What an RBAC resource is

An RBAC resource is an object on which an operation can be performed. In the DataFabric Manager server RBAC system, these resources include Aggregates, Controllers, Clusters, Volumes, Virtual servers, LUNs, Protection policies, Provisioning policies, vFiler templates, Hosts, and DataFabric Manager server Groups (except configuration groups).

A user with the Policy Write capability in the global scope can create schedules and throttles. A user with the Policy Write capability on the policy can modify data protection policies. Similarly, a user with the Policy Delete capability on a policy can delete that policy.

Note: On upgrading to DataFabric Manager server 3.6 or later, a user has the following capabilities:

- User with the Database Write capability in the global scope is assigned the Policy Write capability.
- User with the Database Delete capability is assigned the Policy Delete capability.

Granting restricted access to RBAC resources

You can grant restricted access to objects or resource groups in the DataFabric Manager server.

Steps

1. Create a user defined role by using the following command:

```
dfm role create role_name
```

Example

The following example shows you how to create a user role called EventRole using the CLI:

```
dfm role create EventRol
```

2. Add the following capabilities to the role created in Step 1:
 - Read capability on Global resource for events
 - Write capability on Global resource for events

Example

The following example shows you how to add the capabilities:

```
dfm role add EventRole DFM.Event.Read Global
```

```
dfm role add EventRole DFM.Event.Write Global
```

3. Assign the role created in Step 1 to user Everyone, using the following command:

```
dfm user role add Everyone EventRole
```

Note: You can also use Operations Manager GUI to perform Steps 1 through 3.

4. Open Operations Manager. Read and acknowledge events without logging in.
5. Ensure the user Everyone does not have the capability DFM.Database.Write.

Note: A user with the capability `DFM.Database.Write` can delete all events.

Access check for application administrators

The DataFabric Manager server introduces a new capability requirement to perform access check using RBAC.

The Core AccessCheck capability allows application administrators to check the capabilities of any arbitrary user. For example, if A wants to know the capability of B, A should have the capability to check B's capabilities.

However, any user is allowed to check their own capabilities, regardless of what they are.

When a user configures the client application, the Core AccessCheck capability has to be assigned to a role. Application administrators can check the access permissions of any user, only if they have the permission to do so.

A client application user configured on the DataFabric Manager server with this role allows the client application to check the access of all users.

Note: After upgrading to the DataFabric Manager server 3.6 or later, a user with the Database Read capability in the global scope is assigned the Core AccessCheck capability.

How reports are viewed for administrators and roles

You can view reports for administrators and roles from the CLI.

The following commands are used to generate reports for administrators and roles from the CLI:

- `dfm report role-admins`—Lists all administrators and the roles they are assigned, sorted by administrators.
- `dfm report admin-roles`—Lists all administrators and the roles they are assigned, sorted by role.

For information about how to use the CLI, see the DataFabric Manager server man pages for `dfm report` commands. The man pages specifically describe command organization and syntax.

What global and group access control is

Global access control authorizes an administrator to view and perform actions on any group in the DataFabric Manager server database. Group access control authorizes the administrator to view and perform actions only on the objects of the groups you specify.

However, the administrator cannot add objects to or remove objects from the groups.

You can apply global or group access control to administrator accounts.

You cannot directly create group access administrator accounts. You must first create a global administrator account and then grant access to specific groups. If you want an administrator to have access to specific groups only, create a global administrator account with no roles assigned.

Management of administrator access

You can manage administrator access on storage systems and vFiler units, to define and control the access to the resources, based on the role or job function of a user.

By managing administrator access on storage systems and vFiler units, you can complete the following tasks:

- Manage and control access on storage systems and vFiler units from the DataFabric Manager server.
- Monitor and manage user groups, local users, domain users and roles on storage systems and vFiler units.
- Create and modify identical local users, roles, and user groups on more than one storage system or vFiler unit.
- Edit user groups, local users, domain users, and roles on storage systems and vFiler units.
- Push user groups, local users, domain users, and roles from a storage system or vFiler unit to another storage system or vFiler unit.
- Modify passwords of local users on a storage system or vFiler unit.

Prerequisites for managing administrator access

There are a set of prerequisites that you must consider for managing administrator access.

Following are the prerequisites for managing administrator access on storage systems and vFiler units:

- You must be using Data ONTAP 7.1 or later.
- In the case of local users, the minimum password age, maximum password age, and status fields are available in Data ONTAP 7.1 and later.

- Resetting passwords is available only for storage systems running Data ONTAP 7.1 and later.
- To create or delete roles, user groups, or users on a host, you must have Core Control capability on the host.
- To modify roles, user groups, or users on a host, you must have Core Control and Database Read capabilities on the host.
- To list and view the details of roles, user groups, or users on a host, you must have Database Read capability on the host.
- To push roles, user groups, or users from host A to host B, you must have Database Read capability on host A and Core Control capability on host B.

Limitations in managing administrator access

Roles, user groups, and users without capabilities are not monitored, except for the user group Backup Operators and the users belonging to this user group.

Controlled user access for cluster management

By using Operations Manager, you can create roles and assign capabilities to control user access to selected cluster objects. You can identify users who have the required capabilities to access selected objects within a cluster. Only these users are provided access to manage the cluster objects.

When you provide a user with role capabilities on an object, the same capabilities are applicable to all the other objects contained in the parent object. For example, if you provide a user the `write` option on a cluster, the same `write` option is valid for the controllers, virtual servers, aggregates, and volumes contained in that cluster.

Summary of the global group

You can view the summary that is specific to the global group containing storage systems and vFiler units by selecting Host Users from the Management menu.

Viewing a specific summary page

You can view the summary page specific to a storage system or vFiler unit.

Steps

1. Click **Control Center > Home > Member Details > Physical Systems**.
2. Click desired the storage system or controller link.
3. From the left pane, under Storage Controller Tools, click **Host Users Summary**.

Viewing users on the host

You can view the users on a host using the Host Users report. The Host Users report displays information about the existing users on the host.

Steps

1. From any page, select **Management > Host Users**.
2. Select **Host Users, All** from the Report drop-down list.

Who local users are

Local users are the users created on storage systems and vFiler units.

Viewing local users on the host

You can view the local users on a host using the Host Local Users report. The Host Local Users report displays information about the existing local users on the host.

Steps

1. From any page, select **Management > Host Users**.
2. Select **Host Local Users, All** from the Report drop-down list.

Viewing local user settings on the host

You can view the local users on the storage systems or vFiler units.

Steps

1. From any page, select **Management > Host Users > Local Users**.
2. Click the **view** link corresponding to the local user.

Result

The following details of the selected local user appear:

Host Name	Name of the storage system or vFiler unit
User Name	Name of the local user
Description	Description of the local user
User Full-name	Full name of the local user
Usergroups	Usergroups that the user belongs to
Roles	Roles assigned to the local user

Capabilities	Capabilities of roles assigned to the local user part of user group
Minimum Password Age	Minimum number of days that a password must be used. The number of days should be less than or equal to maximum password age.
Maximum Password Age	Maximum number of days (0 to $2^{32}-1$) that a password can be used
Status	Displays the current status of the user account: <ul style="list-style-type: none"> • Enabled: The user account is enabled. • Disabled: The user account is disabled. • Expired: The user account is expired.

Note: Data ONTAP provides an option to set the maximum number of retries for the password, except for the root login. When the user fails to enter the correct password, even after the maximum retries, then the user account is disabled. The status of the user account is enabled only if the administrator resets the password for the user.

The user account expires if the user fails to change the password within the maximum password age.

For more information about maximum retries, see the *Data ONTAP System Administration Guide*.

Related information

Data ONTAP System Administration Guide - www.ibm.com/storage/support/nseries

Adding local users to the host

You can add a local user to a storage system or vFiler unit.

Steps

1. From any page, select **Management > Host Users > Local Users**.
2. Specify the parameters.

Host Name	Name of the storage system or vFiler unit which user is to be created
User Name	Name of the local user
Password	Password of the local user
Confirm Password	Confirm the password of the local user
User Full Name (optional)	Full name of the local user
Description (optional)	Description of the local user

Minimum Password Age (optional)	Minimum number of days that a password must be used
Maximum Password Age (optional)	Maximum number of days that a password must be used
Usergroup Membership	User groups you want the user to be a member of

3. Select one or more user groups from the list.
4. Add **Local User**.

Editing local user settings on the host

You can edit local user settings on a storage system or vFiler unit.

Steps

1. From any page, select **Management > Host Users > Local Users**.
2. Click the **user** link in the Edit column corresponding to the local user.
3. Edit the parameters.

User Full-name	Full name of the local user
Description	Description of the local user
Minimum Password Age (in days)	Minimum number of days that a password must be used
Maximum Password Age (in days)	Maximum number of days that a password must be used
Usergroup Membership	Usergroups you want to be a member of

Note: You cannot edit Host Name and User Name in the Edit Local User section.

4. Select one or more user groups from the list.
5. Click **Update**.

Users with Execute capability

The DataFabric Manager server users with the Execute capability can reset the password of a local user on storage system or vFiler unit using the credentials that are stored in the database.

Other users who do not have the Execute capability use the credentials that are provided, to modify the password.

Pushing passwords to a local user

You can push an identical password to a local user on multiple storage systems or vFiler units.

Steps

1. From any page, select **Management > Host Users > Local Users**.
2. From the List of Existing Local Users section, click the password link in the Push column corresponding to the local user.

If...	Then...
The local user is on the storage system	The Storage System Passwords page containing the section Modify Password on Storage Systems is displayed. Note: For more information about Storage System Passwords, see the <i>Operations Manager Help</i> .
The local user is on the vFiler	The vFiler Passwords page containing the section Modify Password on vFilers is displayed. Note: For more information about vFiler passwords, see the <i>Operations Manager Help</i> .

3. Specify the parameters.

User Name	Name of the local user
Old Password	Password of the local user
New Password	New password of the local user
Confirm New Password	Confirm the new password of the local user
Select groups and/or Storage systems	Select the following from the respective list: <ul style="list-style-type: none"> • Storage systems on which the local user exists • The DataFabric Manager server groups on which the local user exists
Apply to subgroups	Select the check box if the password change applies to the storage systems of the selected group and the subgroups of the selected group

4. Click **Update**.

Result

Note: Pushing an identical password creates a job that is displayed in the Jobs tab of Password Management and Host User Management window.

Deleting local users from the host

You can delete a local user from a storage system or vFiler unit.

Steps

1. From any page, select **Management > Host Users > Local Users**.
2. From the List of Existing Local Users section, select the local user that you want to delete.
3. Click **Delete Selected**.

Pushing local users to hosts

You can push a local user to a group of storage systems or vFiler units.

Steps

1. From any page, select **Management > Host Users > Local Users**.
2. Select the DataFabric Manager server group or storage system on which you want to push the local user.
3. Select **OK** in the Resources dialog box.
4. Click **Push**.

Result

Note: Pushing local users to host creates a job that is displayed in the Jobs tab of Host User Management window.

Monitoring changes in local user configuration

You can monitor changes in local user configuration on a storage system or vFiler unit.

Steps

1. Click **Setup > Alarms**.
2. Create a new alarm for the event Host User Modified.

Editing passwords

You can edit the password of a local user on a storage system or vFiler unit.

Steps

1. From any page, select **Management > Host Users > Local Users**.
2. Click the **password** link in the Edit column corresponding to the local user.

Note: You cannot edit Host Name and User Name in the Edit Password page.

3. Enter the old password.
4. Enter the new password.
5. Confirm the new password.
6. Click **Update**.

What domain users are

Domain users are the non-local users who belong to a Windows domain, and are authenticated by the domain.

Viewing domain users on the host

You can view the domain users on a host using the Host Domain Users report. The Host Domain Users report displays information about the existing domain users on the host.

Steps

1. From any page, select **Management > Host Users**.
2. Select **Host Domain Users, All** from the Report drop-down list.

Adding domain users to the host

You can add a domain user to a storage system or vFiler unit.

Steps

1. From any page, select **Management > Host Users > Domain Users**.
2. Specify the parameters.

Host Name	Name of the storage system or vFiler unit from the drop-down list
User Identifier (domain-name\username or SID)	Any one of the following: <ul style="list-style-type: none"> • Domain user name

- Security Identifier (SID) of the domain user

Usergroup Membership Usergroups you want to be a member of

3. Select one or more user groups from the list.
4. Click **Add Domain User**.

Viewing domain user settings on the host

You can view the domain user settings on the storage systems or vFiler units.

Steps

1. From any page, select **Management > Host Users > Domain Users**.
2. Click the **view** link corresponding to the domain user.

Result

The following details of the selected domain user appear:

Host Name	Name of the storage system or vFiler unit
User Name	Name of the domain user
SID	Security Identifier of the domain user
Usergroups	Usergroups that the user belongs to
Roles	Roles assigned to the domain user
Capabilities	Capabilities of roles assigned to the domain user as part of the user group

Editing domain user settings on the host

You can edit a domain user on a storage system or vFiler unit.

Steps

1. From any page, select **Management > Host Users > Domain Users**.
2. Click the **edit** link corresponding to the domain user.
3. Edit Usergroup Membership.

Usergroup Membership Usergroups you want to be a member of

Note: You cannot edit Host Name and User Name in the Edit Domain User section.

4. Click **Update**.

Removing domain users from all the user groups

You can remove a domain user from all the user groups.

Steps

1. From any page, select **Management > Host Users > Domain Users**.
2. Select the domain user that you want to remove.
3. Click **Remove From All Usergroups**.

Pushing domain users to hosts

You can push a domain user to a group of storage systems or vFiler units.

Steps

1. From any page, select **Management > Host Users > Domain Users**.
2. Click the **push** link corresponding to the domain user.
3. Select the DataFabric Manager server group, storage system or vFiler unit on which you want to push the domain user.
4. Select **OK**.
5. Click **Push**.

Monitoring changes in domain user configuration

You can monitor the changes in domain user configuration on a storage system or vFiler unit.

Steps

1. Click **Setup > Alarms**.
2. Create a new alarm for the event Host Domain User Modified.

What User groups are

User groups are groups to which the users belong.

Viewing user groups on the host

You can view the user groups on a host using the Host Usergroups report. The Host Usergroups report displays information about the existing user groups on the host.

Steps

1. From any page, select **Management > Host Users**.

2. Select **Host Usergroups, All** from the Report drop-down list.

Adding Usergroups to the host

You can add a user group to a storage system or vFiler unit.

Steps

1. From any page, select **Management > Host Users > Usergroups**.
2. Specify the parameters.

Host Name	Name of the storage system or vFiler unit from the drop-down list
Usergroup Name	Name of the user group
Description	Description of the user group
Select Roles	Capabilities of roles

3. Select one or more roles.
4. Click **Add Usergroup**.

Viewing Usergroup settings on the host

You can view the user group settings on the storage systems or vFiler units.

Steps

1. From any page, select **Management > Host Users > Usergroups**.
2. Click the view link corresponding to the user group.

Result

The following details of the selected user group appear:

Host Name	Name of the storage system or vFiler unit
Usergroup Name	Name of the user group
Description	Description of the user group
Roles	Roles assigned to the user group
Capabilities	Capabilities of the user group

Editing Usergroup settings on the host

You can edit user group settings on a storage system or vFiler unit.

Steps

1. From any page, select **Management > Host Users > Usergroups**.
2. Click the **edit** link corresponding to the user group that you want to edit.
3. Edit the parameters.

Usergroup Name	Name of the user group
Description	Description of the user group
Select Roles	Capabilities of roles

Note: You cannot edit Host Name in the Edit Usergroup section.

4. Select one or more roles.
5. Click **Update**.

Deleting Usergroups from the host

You can delete a user group from a storage system or vFiler unit.

Steps

1. From any page, select **Management > Host Users > Usergroups**.
2. Select the user group that you want to delete.
3. Click **Delete Selected**.

Pushing Usergroups to hosts

You can push identical user groups to a group of storage systems or vFiler units.

Steps

1. From any page, select **Management > Host Users > Usergroups**.
2. Click the **push** link of the user group that you want to push on other storage systems or vFiler units.
3. Select the DataFabric Manager server group, storage system or vFiler unit on which you want to push the user group.
4. Select **OK**.

5. Click **Push**.

Monitoring changes in Usergroup configuration

You can monitor the changes in user group configuration on a storage system or vFiler unit.

Steps

1. Click **Setup > Alarms**.
2. Create a new alarm for the event Host Usergroup Modified.

What roles are

A role is a set of capabilities that can be assigned to a group. You can predefine a role, or you can create or modify it.

Viewing roles on the host

You can view the role settings on the storage systems or vFiler units by using the Host Roles report.

Steps

1. From any page, click **Management > Host Users**.
2. Select **Host Roles, All** from the Report drop-down list.

Adding roles to a host or user

You can add a role to a storage system, vFiler unit, or user.

Steps

1. From any page, select **Management > Host Users > Roles**.
2. Specify the parameters.

Host Name Name of the storage system, vFiler unit, or user from the drop-down list

Role Name Name of the role

Description Description of the role

Capabilities Capabilities of the role

Click the Add Capabilities link.

3. Select one or more capabilities you want to add.
4. Click **OK**.
5. Click **Add Role**.

Viewing role settings on the host

You can view the roles on the storage systems or vFiler units.

Steps

1. From any page, select **Management > Host Users > Roles**.
2. Click the view link corresponding to the host role.

Result

The following details of the selected host role appear:

Host Name	Name of the storage system or vFiler unit
Role Name	Name of the role
Description	Description of the role
Capabilities	Capabilities of the role

Editing role settings on the host

You can edit a role on a storage system or vFiler unit.

Steps

1. From any page, select **Management > Host Users > Roles**.
2. Click the **edit** link corresponding to the host role that you want to edit.
3. Edit the parameters.

Description	Description of the role
Capabilities	Capabilities of the role
	Click the Edit link.

Note: You cannot edit Host Name and Role Name in the Edit Role section.

4. Select one or more capabilities you want to add.
5. Click **Ok**.
6. Click **Update**.

Deleting roles from the host

You can delete a role from a storage system or vFiler unit.

Steps

1. From any page, select **Management > Host Users > Roles**.
2. Select the host role that you want to delete.
3. Click **Delete Selected**.

Pushing roles to the hosts

You can push identical roles to a group of storage systems or vFiler units.

Steps

1. From any page, select **Management > Host Users > Roles**.
2. Click the **push** link of the host role you want to push on other storage systems or vFiler units.
3. Select the DataFabric Manager server group or storage system on which you want to push the role.
4. Select **OK**.
5. Click **Push**.

Monitoring changes in role configuration

You can monitor the changes in role configuration on a storage system or vFiler unit.

Steps

1. Click **Setup > Alarms**.
2. Create a new alarm for the event Host Role Modified.

What jobs display

Jobs display the status of the push jobs.

Pushing jobs

To view the status of the push jobs, select Management > Host Users >Jobs.

Deleting push jobs

You can delete a push job.

Steps

1. Select the push job that you want to delete.
2. Click **Delete**.

Groups and objects

A group is a collection of the DataFabric Manager server objects. You can group objects based on characteristics such as the operating system of a storage system (Data ONTAP version).

You can also group objects based on storage systems at a location, or all file systems that belong to a specific project or group in your organization.

Storage system elements monitored by the DataFabric Manager server, such as storage systems, aggregates, file systems (volumes and qtrees), and logical unit numbers (LUNs), are referred to as objects.

Following is a list of the DataFabric Manager server objects that can be added to a resource group:

- Host (can include storage systems, Host Agents, virtual servers, and vFiler units)
- Volume
- Qtree
- Configuration
- LUN Path
- Aggregate
- SRM Path
- Dataset
- Resource Pool
- Disk

Following is a set of considerations for creating groups:

- You can group similar or different objects in a group.
- An object can be a member of any number of groups.
- You can group a subset of group members to create a new group.
- You cannot create a group of groups.
- You can create any number of groups.
- You can copy a group or move a group in a group hierarchy.

What group types are

The DataFabric Manager server automatically determines the type of a group based on the objects it contains. If you place your cursor over an icon, to the left of a group name, on the left side of Operations Manager main window, you can quickly find out the type of objects the group contains.

What homogeneous groups are

You can group objects into sets of objects with common characteristics. They might, for example, have the same operating system or belong to a specific project or group in your organization.

You can create the following types of groups:

- Appliance Resource group () - contains storage systems, vFiler units, and host agents
- Aggregate Resource group () - contains aggregates only
- File System Resource group () - contains volumes and qtrees
- LUN Resource group () - contains LUNs only
- Configuration Resource group () - contains storage systems associated with one or more configuration files
- SRM path group - contains SRM paths only
- Dataset - is the data that is stored in a collection of primary storage containers, including all the copies of the data in those containers
- Resource pool - is a collection of storage objects from which other storage containers are allocated

Note: For more information about datasets and resource pools, see the *Provisioning Manager and Protection Manager Guide to Common Workflows for Administrators*.

Related information

[Provisioning Manager and Protection Manager Guide to Common Workflows for Administrators - www.ibm.com/storage/support/nseries](http://www.ibm.com/storage/support/nseries)

What mixed-type groups are

You can add objects of different types to the same group. Grouping objects from different homogeneous groups also constitutes a mixed-type group. For example, a mixed-type group can have a group of vFiler units and volumes. You can also group objects based on their geographical location, or by the client, they support.

Configuration resource groups can contain only storage systems, vFiler units, and configurations. Once created, you cannot add any objects to the configuration resource group in the DataFabric Manager server. The elements of a configuration resource group cannot be a part of any other homogeneous group. The DataFabric Manager server prevents you from adding other object types to configuration resource groups. If a group already contains objects other hosts, you cannot add a configuration file to the group.

What a Global group is

All objects monitored by DataFabric Manager server belong to the global group. By default, a group called global exists in the DataFabric Manager server database. All subgroups created also belong to the global group.

You cannot delete or rename the global group. When you delete an object, the DataFabric Manager server stops monitoring and reporting data for that object. Data collection and reporting is not resumed until the object is added back (recovered) to the database.

What hierarchical groups are

In addition to creating groups of objects, you can create subgroups within groups to establish a hierarchy of groups.

Hierarchical groups help you manage administrative privileges, because privileges granted to a parent group are implicitly granted to all its subgroups. Besides, following are the benefits of having hierarchical groups:

- You can determine the capacity of the group and the chargeback options.
- You can keep a record of trending, that is, the data growth rate of the group.
- You can select arguments for reports to be generated.

Creating groups

You can create a new group from the Edit Groups page. You can group objects based on storage systems at a location, or all file systems that belong to a specific project or group in your organization.

Before you begin

To create a group, you must be logged in as an administrator with a role having Database Write capability on the parent group. To create a group directly under the Global group, the administrator must have a role with Database Write capability on the Global group.

Steps

1. From the **Control Center**, click the **Edit Groups**.
2. In the **Group Name** field, type the name of the group you want to create.
See “Naming conventions” for groups.
3. From the list of groups, select the parent group for the group you are creating.

You might need to expand the list to display the parent group you want.

4. Click **Add.**

Result

The new group is created. The Current Groups list in the left-pane area is updated with the new group. You might need to expand the Current Groups list to display the new group.

Creating groups from a report

You can create a new group from a report in Operations Manager.

Steps

1. From the **Control Center**, click the **Member Details** tab.
2. Click **Aggregate**, **File Systems**, **LUN**, or tabs.
3. To the left of the list of objects in the main window, select the check boxes for the objects that you want to add to the group.
4. At the bottom left of the main window, click **Add to New Group**.
5. In the **Group Name** field, type the name of the group you want to create. See “Naming conventions” for groups.
6. From the list of groups, select the parent group for the group you are creating. You might need to expand the list to display the parent group you want.
7. Click **Add**.

The new group is created. The Current Groups list in the left-pane area is updated with the new group. You might need to expand the Current Groups list to display the new group.

What configuration resource groups are

A configuration resource group is a group of storage systems that share a set of common configuration settings. A configuration resource group allows you to designate groups of managed storage systems that can be remotely configured to share the same configuration settings.

A configuration resource group must contain some number of storage systems and have one or more files containing the desired configuration settings. These configuration settings are listed in files called configuration files. Configuration files exist independently of groups, and can be shared between groups. You can use Operations Manager to create configuration files and to specify the configuration settings that you want to include in them.

With Operations Manager, you can create and manage configuration files that contain configuration settings you want to apply to a storage system and vFiler unit or groups of storage systems and vFiler units. By using the storage system configuration management feature, you can pull configuration

settings from one storage system and vFiler unit and push the same or a partial set of settings to other storage systems or groups of storage systems and vFiler units. While pushing the configurations settings, you must ensure that storage system and vFiler unit configuration conforms to the configuration pushed to it from Operations Manager.

When you create configuration resource groups, consider the following:

- Storage systems running on different operating system versions can be grouped in the same configuration resource group.
- A storage system can belong to only one configuration resource group.
- A configuration resource group must have one or more configuration files associated with it. Otherwise, storage systems cannot be configured remotely.
- For configuration management, appropriate plug-ins must be associated.
- You cannot run any reports for a configuration resource group.

Besides specifying configuration settings by associating individual configuration files with a group of storage systems, you can also specify another configuration resource group from which to acquire configuration settings. Such a group is known as a parent group. For example, a previously created configuration resource group might already have most or all, of the settings you require.

Guidelines for creating groups

You should follow a set of guidelines when you create groups.

- You can group similar or mix-types objects in a group.
- An object can be a member of any number of groups.
- You can group a subset of group members to create a new group.
- You can create any number of groups.
- You can copy to a group or move to a group in a group hierarchy.
However, parent group cannot be moved to its child group.

Guidelines for creating configuration resource groups

You must use a set of guidelines when you create Configuration Resource groups.

Use the following guidelines when you create Configuration Resource groups:

- You can include storage systems, and vFiler units, with different model types and software versions.
- Storage system, and vFiler unit, can be a member of only one configuration resource group, but can still be a member of multiple groups for monitoring purposes.
- You cannot create a group of Configuration Resource groups.
- To apply settings to a Configuration Resource group, you must associate one or more configuration files with the group.

Guidelines for adding vFiler units to Appliance Resource group

You must consider a set of guidelines before adding vFiler units to a resource group.

- You can add vFiler units as members to an Appliance Resource group. The hosting Storage system and the storage resources (qtrees, volumes, and LUNs) assigned to the vFiler unit are also added as indirect members. When you remove a vFiler unit from a group, its related hosting Storage system and storage resources are also removed.
- If you add a hosting Storage system that is configured with vFiler units to a group, the vFiler units are also added as indirect members. When you remove a hosting storage system from a group, its vFiler units are also removed.
- If you add a storage resource assigned to a vFiler unit to a group, the vFiler unit is also added as an indirect member. If you remove the storage resources from a group, the vFiler unit is also removed.

Note: Indirect members are considered for determining the group status.

Editing group membership

In Operations Manager, you can add members to a group.

Steps

1. Go to the Groups area on the left side of Operations Manager and expand the list as needed to display the group to which you want to add members.
2. Click the name of the group to which you want to add members.
3. From the **Current Group** menu at the lower left of Operations Manager, click **Edit Membership**.
4. Select the object from the **Choose from All Available** list and click >> to move the object to the list at the right.

Operations Manager adds the selection to the group and updates the membership list displayed on the right side of the Edit Group Membership area.

What group threshold settings are

Group thresholds determine at what point you want the DataFabric Manager server to generate events regarding capacity problems with object groups. You can create an alarm for a group to send notification to designated recipients whenever a storage event occurs.

The thresholds you can change depend on the type of objects in a group. For example, you can change Appliance CPU Too Busy Threshold for only an Appliance Resource group. You can change Volume Full Threshold and Volume Nearly Full Threshold for only a File System Resource Group.

For a list of thresholds you can change for an object type, see the chapter where that type of object is the main topic of discussion in this guide.

Note: When you apply threshold changes to a group, the new threshold values are associated with the objects in the group. These new threshold values are not associated with the group. That is, if you add another object to a group, after applying a threshold change, the threshold value of the new object is not changed. The threshold value of the new object does not change if it is different from the current group. Additionally, if you apply threshold changes to an object that belongs to multiple groups, the threshold value is changed for this object across all groups.

For information about how to change the thresholds for a group of objects, see the *Operations Manager Help*.

What group reports are

Grouping objects enables you to view consolidated data reports, events, and status of objects.

For example, you can view the total storage capacity used, or events generated by all manufacturing storage systems, by creating a group of the storage systems, and using the Summary tab of Operations Manager.

What summary reports are

Summary reports are available for all groups, including the Global group.

- Status
- Group members
- Storage capacity used and available
- Events
- Storage chargeback information
- Monitored devices
- Physical space
- Storage system operating systems

- Storage system disks
- Capacity graphs

Note: You can view additional reports that focus on the objects in a group by clicking the name of the group and then clicking the appropriate Operations Manager tab.

What subgroup reports are

If you run a report on a group with subgroups, the data displayed includes data on applicable objects in the subgroups. For example, if you display the Aggregate Capacity Graph report on a parent group containing aggregates, you see data about the aggregates in the parent group. You can also see data about the aggregates in its subgroups. You do not see data about other object types in the parent group or the subgroups.

If you run a report on a mixed-type object group, Operations Manager runs the report on group members of the applicable type. For example, qtrees for the Qtree Growth report. Operations Manager combines the results, and then eliminates the duplicates, if any.

What cluster-related objects are

Operations Manager enables you to include cluster-related objects, such as controllers and virtual servers, in a group. This enables you to easily monitor cluster-related objects that belong to a particular group.

Note: This document discusses features that are associated with Data ONTAP 8.x Cluster-Mode, which is currently not offered by IBM.

The cluster-related objects are as follows:

- Virtual server** A single file-system namespace. A virtual server has separate network access and provides the same flexibility and control as a dedicated node. Each virtual server has its own user domain and security domain. It can span multiple physical nodes.
- A virtual server has a root volume that constitutes the top level of the namespace hierarchy; additional volumes are mounted to the root volume to extend the namespace. A virtual server is associated with one or more logical interfaces (LIFs) through which clients access the data on the storage server. Clients can access the virtual server from any node in the cluster, but only through the logical interfaces that are associated with the virtual server.
- Namespace** Every virtual server has a namespace associated with it. All the volumes associated with a virtual server are accessed under the virtual server's namespace. A namespace provides a context for the interpretation of the junctions that link together a collection of volumes.

Junction	A junction points from a directory in one volume to the root directory of another volume. Junctions are transparent to NFS and CIFS clients.
Logical interface	An IP address with associated characteristics, such as a home port, a list of ports to fail over to, a firewall policy, a routing group, and so on. Each logical interface is associated with a maximum of one virtual server to provide client access to it.
Cluster	A group of connected storage systems that share a global namespace and that you can manage as a single virtual server or multiple virtual servers, providing performance, reliability, and scalability benefits.
Storage controller	The component of a storage system that runs the Data ONTAP operating system and controls its disk subsystem. Storage controllers are also sometimes called controllers, storage appliances, appliances, storage engines, heads, CPU modules, or controller modules.
Ports	<p>A port represents a physical Ethernet connection. In a Data ONTAP cluster, ports are classified into the following three types:</p> <ul style="list-style-type: none">• Data ports Provide data access to NFS and CIFS clients.• Cluster ports Provide communication paths for cluster nodes.• Management ports Provide data access to Data ONTAP management utility.
Data LIF	A logical network interface mainly used for data transfers and operations. A data LIF is associated with a node or virtual server in a Data ONTAP cluster.
Node management LIF	A logical network interface mainly used for node management and maintenance operations. A node management LIF is associated with a node and does not fail over to a different node.
Cluster management LIF	A logical network interface used for cluster management operations. A cluster management LIF is associated with a cluster and can fail over to a different node.
Interface group	A single virtual network interface that is created by grouping together multiple physical interfaces.

Creating a group of cluster objects

By using Operations Manager, you can create a group of cluster objects for easier administration and access control. You can add objects such as clusters, controllers, aggregates, volumes, and virtual servers to a group.

Steps

1. Click **Control Center > Home > Member Details > Physical Systems > Report**.

Example

The navigational path to create a group of virtual servers is **Control Center > Home > Member Details > Virtual Systems > Report > Virtual Servers, All**.

2. Depending on the cluster objects you want to group, select the appropriate report from the Report drop-down list.
3. In the resulting report, select the cluster objects you want to include in a group.
4. From the buttons at the bottom of the page, click **Add To New Group**.
5. In the **Group Name** field, enter a name for the group.
6. Select an appropriate parent for your new group.
7. Click **Add**.

Storage monitoring and reporting

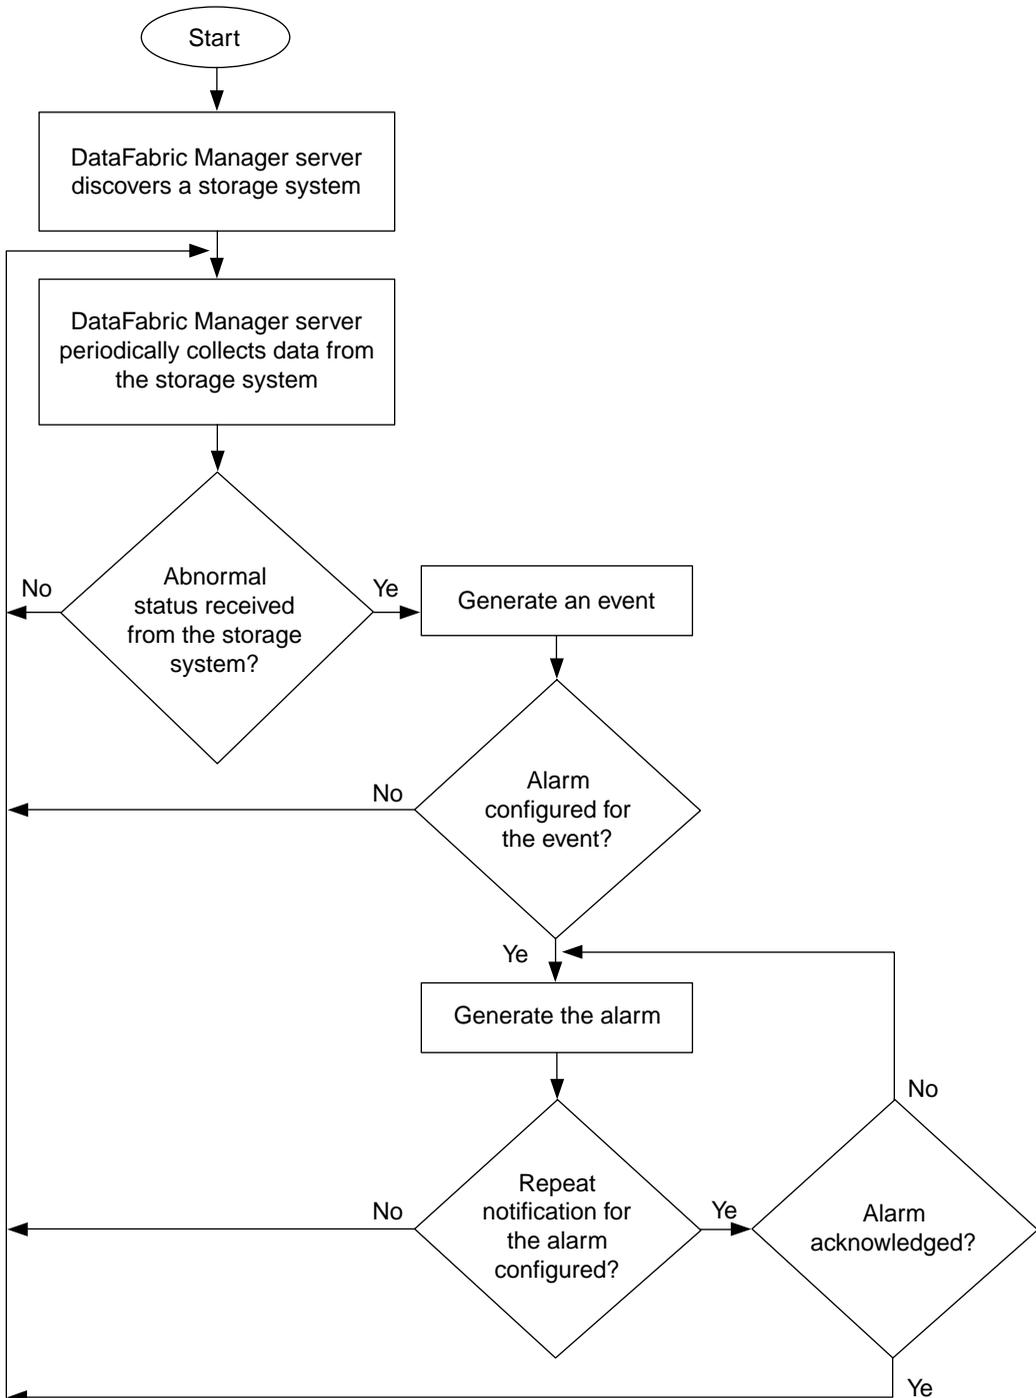
Monitoring and reporting functions in the DataFabric Manager server depend on event generation. You must configure the settings in Operations Manager to customize monitoring and to specify how and when you want to receive event notifications.

Operations Manager allows you to generate summary and detailed reports. Depending on which tab you select, Operations Manager returns the appropriate graph or selection of reports (for example, reports about storage systems, volumes, and disks).

The DataFabric Manager server monitoring process

The DataFabric Manager server discovers the storage systems supported on your network. The DataFabric Manager server periodically monitors data that it collects from the discovered storage systems, such as CPU usage, interface statistics, free disk space, qtree usage, and chassis environmental. The DataFabric Manager server generates events when it discovers a storage system, when the status is abnormal, or when a predefined threshold is breached. If configured to do so, the DataFabric Manager server sends a notification to a recipient when an event triggers an alarm.

The following flow chart illustrates the DataFabric Manager server monitoring process:



Cluster monitoring with Operations Manager

From the current release of the DataFabric Manager server, Operations Manager can discover, monitor, and generate reports of systems running Data ONTAP 8.0 Cluster-Mode. Operations Manager uses SNMP or XML APIs for cluster monitoring.

What the cluster management logical interface is

The cluster management logical interface is a virtual network interface that enables you to perform cluster management operations. This logical interface is associated with a cluster and can be failed over to a different node.

The cluster management logical interface is associated with a cluster management server to provide a detailed view of the cluster. Operations Manager monitors clusters running Data ONTAP 8.0 Cluster-Mode by using the appropriate SNMP version or APIs. You can gather information about the cluster resources, controller resources, and virtual server resources from the cluster management logical interface.

Information available on the Cluster Details page

The Cluster Details page for a cluster provides information such as the cluster hierarchy, status of the cluster, number of logical interfaces, ports, and so on. You can access the Clusters Details page by clicking the cluster name from any of the cluster reports.

The Cluster Details page displays the following cluster-related information:

- Status of the cluster
- Serial Number
- Uptime
- Primary IP address
- Number of controllers
- Number of virtual servers
- Number of ports
- Number of logical interfaces
- Contact and location
- Current events that have the cluster as the source
- Storage capacity
- Groups to which the cluster belongs
- List of the most recent polling sample date and the polling interval of all the events monitored for the cluster
- Graphs that display the following information:
 - Volume capacity used
 - Volume capacity used versus total capacity

- Aggregate capacity used
- Aggregate capacity used versus total capacity
- Aggregate space, usage versus committed
- CPU usage (in percentage)
- NFS Operations/sec
- NFS and CIFS Operations/sec
- CIFS Operations/sec
- Network Traffic/sec
- Logical Interface Traffic/sec

Tasks performed from the Cluster Details page

You can perform various cluster management tasks from the Cluster Details page. You can access the Clusters Details page by clicking the cluster name from any of the cluster reports.

You can perform the following tasks from the Cluster Details page:

- Exploring the physical view of the system
You can gather information about cluster components in the form of reports. For example, by clicking the number corresponding to the Controllers link, you can view the details of those controllers from the Controllers, All report.
- Viewing total utilization of physical resources
You can view the total utilization of physical resources, including CPU usage, network traffic at the cluster level, and volume capacity used. You can also view the graphical representation of the corresponding resources.
- Browsing the cluster and its components
You can browse through the cluster and its components from the "Cluster Hierarchy" section in the Cluster Details page. You can also browse to the corresponding report of a particular component. For example, you can expand the cluster, view the list of virtual servers, and then click a desired virtual server to view its details on the corresponding Virtual Server Details page.

Viewing the utilization of resources

You can view the graphical representation of the utilization of various physical and logical resources from the Details page in Operations Manager.

Viewing the utilization of logical resources

By using Operations Manager, you can view the graphical representation of the utilization of your logical resources such as virtual servers, logical interface traffic to the virtual server, and volumes.

You can also configure alarms to send notification whenever the utilization exceeds preset thresholds.

Steps

1. Click **Control Center > Home > Member Details > Virtual Systems > Report > Virtual Servers, All**.
2. Click the name of the virtual server for which you want to view the utilization of logical resources.
3. In the **Virtual Server Details** page, select the appropriate graph from the drop-down menu.

You can view the utilization graph on a daily, weekly, monthly, quarterly, or yearly basis. For example, to view the virtual server's volume capacity used for a period of one year, you can select Volume Capacity Used and click **1y**.

Viewing the utilization of physical resources

By using Operations Manager, you can view the graphical representation of the utilization of your physical resources such as CPU, network traffic to the controller, and so on.

Steps

1. Click **Control Center > Home > Member Details > Physical Systems > Report > Controllers, All**.
2. Click the name of the controller for which you want to view the utilization of physical resources.
3. In the **Storage Controller Details** page, select the appropriate graph from the drop-down menu.

You can view the utilization graph on a daily, weekly, monthly, quarterly, or yearly basis. For example, to view the controller's CPU usage for a period of one year, you can select CPU Usage (%) and click **1y**.

Links to FilerView

In the DataFabric Manager server 2.3 and later, UI pages displaying information about some DataFabric Manager server objects contain links, indicated by the icon to FilerView()¹, the Web-based UI for storage systems.

When you click the icon, you are connected to the FilerView location where you can view information about, and make changes to the object whose icon you clicked. Depending on your setup, you might need to authenticate the storage system whose FilerView you are connecting to, using one of the administrator user accounts on the storage system.

Note: Links to FilerView is not available for systems running Data ONTAP 8.0 Cluster-Mode.

Query intervals

The DataFabric Manager server uses periodic SNMP queries to collect data from the storage systems it discovers. The data is reported by the DataFabric Manager server in the form of tabular and graphical reports and event generation.

The time interval at which an SNMP query is sent depends on the data being collected. For example, although the DataFabric Manager server pings each storage system every minute to ensure that the storage system is reachable, the amount of free space on the disks of a storage system is collected every 30 minutes.

Guidelines for changing monitoring intervals

The SNMP query time intervals are specified by the global monitoring option that is located in the Monitoring Options section of the Options page. Although you should generally keep the default values, you might need to change some of the options to suit your environment. All the monitoring option values apply to all storage systems in all groups.

If you decrease the monitoring intervals, you receive more real-time data. However, the DataFabric Manager server queries the storage systems more frequently, thereby increasing the network traffic and the load on the DataFabric Manager server and the storage systems responding to the queries.

If you increase the monitoring interval, the network traffic and the storage system load are reduced. However, the reported data might not reflect the current status or condition of a storage system.

The DataFabric Manager server 3.1 and later includes an SNMP trap listener that improves the speed of event generation.

Related concepts

What SNMP trap listener is on page 86

What SNMP trap listener is

In addition to periodically sending out SNMP queries, DataFabric Manager server 3.1 and later include an SNMP trap listener as part of the server service. Event generation and alerting is faster than with SNMP queries because the proper monitoring mechanism is started immediately after the SNMP trap is received. In addition, monitoring is performed asynchronously, instead of waiting for the monitoring interval.

The SNMP trap listener listens for SNMP traps from monitored storage systems, if they have been manually configured to send traps to the DataFabric Manager server (over UDP port 162).

Note: The SNMP trap listener can receive SNMP traps only from storage systems that are supported on the DataFabric Manager server. Traps from other sources are dropped.

What SNMP trap events are

When the SNMP trap listener receives an SNMP trap, the DataFabric Manager server issues an Information event, but does not change the status of the host.

Instead, the corresponding monitor associated with the trap generates the proper event and continues to monitor the host to report status changes. The name associated with the SNMP trap Information event indicates the severity of the trap: for example, Error Trap. The trap severities are deduced from the last digit of the trap ID, as specified in the custom MIB. The SNMP traps received by the SNMP trap listener are specified in the custom MIB. For a complete list of traps and associated trap IDs, see the *Data ONTAP Network Management Guide*.

The following list describes the SNMP trap Information event types:

- Emergency Trap Received
- Alert Trap Received
- Critical Trap Received
- Error Trap Received
- Warning Trap Received 106 SNMP traps
- Notification Trap Received
- Information Trap Received

If the severity of a trap is unknown, the DataFabric Manager server drops the trap.

Related concepts

Information about the DataFabric Manager server MIB on page 88

Related information

Data ONTAP Network Management Guide - www.ibm.com/storage/support/nseries

How SNMP trap reports are viewed

You can use the Events tab to view reports about the SNMP traps that are received by the DataFabric Manager server.

The Events tab enables you to view a listing of all current SNMP traps or to sort them by severity. Each view provides information about each SNMP trap, for example, the name of the trap, the severity, and the condition that led to the error.

When SNMP traps cannot be received

The DataFabric Manager server cannot receive SNMP traps, if any of the following conditions exist:

- A system has not been configured to send traps to the DataFabric Manager server.
- The host is not a supported storage system.
- DataFabric Manager server version is before 3.1.

Additionally, the DataFabric Manager server cannot receive Debug traps.

SNMP trap listener configuration requirements

A set of configuration requirements must be met to enable reception of SNMP traps from managed storage systems.

On the DataFabric Manager server: No configuration is needed to start the SNMP trap listener on the DataFabric Manager server (the trap listener is automatically started after installation). The SNMP trap global options are also configured with default settings, although you might want to modify these settings.

On managed storage systems: You must manually add the DataFabric Manager server as a trap destination on all supported systems to be monitored. The traps must be sent to the DataFabric Manager server over UDP port 162.

Note: If another trap listener is listening on port 162, the startup of the built-in trap listener fails with an error and the Warning event is displayed in the Events page.

How SNMP trap listener is stopped

The SNMP trap listener is enabled by default. If you want to start or stop the SNMP trap listener, use the `snmpTrapListenerEnabled` CLI option.

Configuration of SNMP trap global options

You can configure the `SNMP trap global` options by accessing the `SNMP Trap Listener` options and the `Event` and `Alert` options on the Options page.

Configuration of the SNMP trap global options is not necessary at start-up. However, you might want to modify the global default settings.

The following global default settings can be modified:

- **Enable SNMP trap listener**
Use this option to enable or disable the SNMP trap listener.
- **SNMP Trap Listener Port**
Use this option to specify the UDP port on which the SNMP Manager Trap Listener receives traps. Supported storage systems can send SNMP traps only over UDP port 162.
- **SNMP Maximum Traps Received per window and SNMP Trap Window Size**
Use these two options to limit the number of SNMP traps that can be received by the trap listener within a specified period.

Information about the DataFabric Manager server MIB

The SNMP traps generated for the DataFabric Manager server events are specified in the DataFabric Manager server MIB.

The MIB at the following locations provides a complete list of the DataFabric Manager server SNMP traps and associated trap IDs:

- For Windows: `installation_directory\dfm\misc`
- For UNIX: `installation_directory/misc`

The DataFabric Manager server can send only the traps that are available in the MIB.

Note: The DataFabric Manager server can send information to an SNMP trap host only when an alarm for which the trap host is specified is generated. The DataFabric Manager server cannot serve as an SNMP agent; that is, you cannot query the DataFabric Manager server for information from an SNMP trap host.

What events are

Events are generated automatically when a predefined condition occurs or when an object crosses a threshold. *Event messages* inform you when specific events occur. All events are assigned a severity type and are automatically logged in the Events window.

You can configure alarms to send notification automatically when specific events or severity types occur. If an application is not configured to trigger an alarm when an event is generated, you can find out about the event by checking the Events window.

It is important that you take immediate corrective action for events with severity level Error, Critical, Emergency, or Worse. Ignoring such events can lead to poor performance and system unavailability.

Note: Event types are predetermined. Although you cannot add or delete event types, you can manage notification of events. However, you can modify the event severity type from the command-line interface.

Viewing events

You can view a list of all events that occurred and view detailed information about any event.

Step

1. View the events logged by Operations Manager in any of the following ways:
 - Click the **Events: Emergency, Critical, Error, Warning** link located at the top of the Operations Manager main window.
 - From the **Control Center** tab, click the **Events** tab located in the Group Summary page.
 - Select the Details pages for storage systems, SAN hosts, FC switches, HBA ports, and FCP targets. The Details pages provide lists of events related to the specific component.
 - From the **Backup Manager** tab or the **Disaster Recovery Manager** tab, click the **Events** tab.

Note: User quota threshold events can be viewed only with the User Quota Events report available through the Report drop-down list on the Events tab.

Managing events

If the DataFabric Manager server is not configured to trigger an alarm when an event is generated, you cannot find out about the event. However, to identify the event, you can check the events log on the DataFabric Manager server.

Steps

1. From an **Events** view, select the check box for the event that you want to acknowledge. You can select multiple events.
2. Click **Acknowledge Selected** to acknowledge the event that caused the alarm.
3. Find out the cause of the event and take corrective action.
4. Delete the event.

Operations on local configuration change events

After receiving the event, you have the choice of acknowledging, fixing, or deleting the event.

If you click Fix, a new window that shows the differences between the local configuration and the group settings is displayed.

From this window, you can accept, or, reject the local configuration changes made on the storage system. If you reject the configuration changes, the DataFabric Manager server undoes all the local changes. If you accept the configuration changes, the configuration settings listed are not modified during subsequent configuration pushes.

List of events and severity types

These tables list all of the events generated by the N series Management Console data protection and provisioning capabilities and Operations Manager and the associated event severity types. Events are listed in alphabetical order by object type.

Use the links in the following table to jump directly to the events for that object:

Note: Performance Advisor uses only the Normal and Error events.

Event categories		
Active/Active Configuration Controller on page 91	FC (Fibre Channel) Switch Port on page 99	Protection Policy on page 105
Active/Active Configuration Interconnect on page 92	Fans on page 99	Protection Schedule on page 105
Active/Active Configuration Partner on page 92	Filer Configuration on page 99	Provisioning Policy on page 106
Agent on page 92	Global Status on page 99	Qtree on page 106
Aggregate on page 93	HBA Port on page 100	Remote Platform Management (RPM) on page 106
Alarm on page 93	Host on page 100	Resource Group on page 106
CFO Interconnect on page 94	Host Agent on page 101	Resource Pool on page 107
CFO Partner on page 94	Inodes on page 101	SAN Host LUN Mapping on page 107
CFO Settings on page 94	Interface Status on page 101	Script on page 107
CFO This Storage System on page 94	Logical Interface on page 101	SnapMirror on page 107
Cluster on page 95	LUN on page 102	Snapshot(s) on page 109
Cluster port on page 95	Management Station on page 102	SnapVault on page 109
Comment Field on page 95	Migration on page 103	SNMP Trap Listener on page 110
Configuration Changed on page 95	NDMP on page 103	Space Management on page 110
CPU on page 96	Network on page 103	Storage Services on page 111
Data Protection on page 96	Network Services on page 104	Sync on page 111
Database on page 96	No Schedule Conflict on page 104	Temperature on page 111
Dataset on page 96	NVRAM Battery on page 104	Unprotected Item on page 111
Dataset Backup on page 98	OSSV (Open Systems SnapVault) on page 105	User on page 111
Dataset Conformance on page 98	Performance Advisor on page 105	vFiler Unit on page 112
Disks on page 98	Power Supplies on page 105	vFiler Unit Template on page 112
Enclosures on page 98	Primary on page 105	Vserver on page 112
		Volume on page 113

Active/Active Configuration Controller

Event name	Severity
Can Take Over	Normal
Cannot Takeover	Error

Event name	Severity
Dead	Critical
Takeover	Warning

Active/Active Configuration Interconnect

Event name	Severity
Down	Error
Not Present	Warning
Partial Failure	Error
Up	Normal

Active/Active Configuration Partner

Event name	Severity
Dead	Warning
May Be Down	Warning
OK	Normal

Active/Active Configuration Settings

Event name	Severity
Disabled	Normal
Enabled	Normal
Not Configured	Normal
Takeover Disabled	Normal
This Controller Dead	Warning

Agent

Event name	Severity
Down	Error
Login Failed	Warning
Login OK	Normal

Event name	Severity
Up	Normal

Aggregate

Event name	Severity
Almost Full	Warning
Almost Overcommitted	Warning
Deleted	Information
Discovered	Information
Failed	Error
Full	Error
Nearly Over Deduplicated	Warning
Not Over Deduplicated	Normal
Not Overcommitted	Normal
Offline	Error
Online	Normal
Overcommitted	Error
Over Deduplicated	Error
Restricted	Normal
Snapshot Reserve Almost Full	Warning
Snapshot Reserve Full	Warning
Snapshot Reserve OK	Normal
Space Normal	Normal

Alarm

Event name	Severity
Created	Information
Deleted	Information
Modified	Information

CFO Interconnect

Event name	Severity
Down	Error
Not Present	Warning
Partial Failure	Error
Up	Normal

CFO Partner

Event name	Severity
Dead	Warning
May Be Down	Warning
OK	Normal

CFO Settings

Event name	Severity
Disabled	Normal
Enabled	Normal
Not Configured	Normal
Takeover Disabled	Normal
This Node Dead	Warning

CFO This Storage System

Event name	Severity
Can Take Over	Normal
Cannot Take Over	Error
Dead	Critical
Takeover	Warning

Cluster

Event name	Severity
Cluster Discovered	Information
Cluster Reachable	Normal
Cluster Not Reachable	Critical
Cluster Renamed	Information
Cluster Node Added	Information
Cluster Node Removed	Information

Cluster port

Event name	Severity
Port Status Up	Normal
Port Status Down	Error
Port Status Undefined	Normal
Port Status Unknown	Normal
Port Role Changed	Information

Comment Field

Event name	Severity
Created	Information
Modified	Information
Destroyed	Information

Configuration Changed

Event name	Severity
Config Group	Information

CPU

Event name	Severity
Load Normal	Normal
Too Busy	Warning

Data Protection

Event name	Severity
Job Started	Information
Policy Created	Information
Policy Modified	Information
Schedule Created	Information
Schedule Modified	Information

Database

Event name	Severity
Backup Failed	Error
Backup Succeeded	Information
Restore Failed	Error
Restore Succeeded	Information

Dataset

Event name	Severity
Backup Aborted	Warning
Backup Completed	Normal
Backup Deleted	Information
Backup Failed	Error
Backup Prematurely Deleted	Warning
Created	Information
Deleted	Information

Event name	Severity
DR State Ready	Information
DR State Failover Over	Warning
DR State Failed Over	Information
DR State Failover Error	Error
DR Status Normal	Information
DR Status Warning	Warning
DR Status Error	Error
Initializing	Information
Job Failure	Warning
Member Clone Snapshot Discovered	Information
Member Clone Snapshot Status OK	Information
Member Dedupe Operation Failed	Error
Member Dedupe Operation Succeeded	Normal
Member Destroyed	Information
Member Destroy Operation Failed	Information
Member Resized	Information
Member Resize Operation Failed	Information
Modified	Information
Protected	Normal
Protection Failed	Error
Protection Lag Error	Error
Protection Lag Warning	Warning
Protection Suspended	Warning
Protection Uninitialized	Normal
Provisioning Failed	Error
Provisioning OK	Normal
Space Status: Normal	Normal

Event name	Severity
Space Status: Warning	Warning
Space Status: Error	Error
Write Guarantee Check - Member Resize Required	Warning
Write Guarantee Check - Member Size OK	Normal

Dataset Backup

Event name	Severity
Dataset Backup: Deleted	Information
Dataset Backup: Prematurely Deleted	Warning

Dataset Conformance

Event name	Severity
Conformant	Normal
Conforming	Information
Initializing	Information
Nonconformant	Warning

Disks

Event name	Severity
No Spares	Warning
None Failed	Normal
None Reconstructing	Normal
Some Failed	Error
Some Reconstructing	Warning
Spares Available	Normal

Enclosures

Event name	Severity
Active	Information

Event name	Severity
Disappeared	Warning
Failed	Error
Found	Normal
Inactive	Warning
OK	Normal

Fans

Event name	Severity
Many Failed	Error
Normal	Normal
One Failed	Error

FC (Fibre Channel) Switch Port

Event name	Severity
Faulty	Error
Offline	Warning
Online	Normal

Filer Configuration

Event name	Severity
Changed	Warning
OK	Normal
Push Error	Warning
Push OK	Normal

Global Status

Event name	Severity
Critical	Critical
Non Critical	Error

Event name	Severity
Non Recoverable	Emergency
OK	Normal
Other	Warning
Unknown	Warning

HBA Port

Event name	Severity
Offline	Warning
Online	Normal
Port Error	Error
Traffic High	Warning
Traffic OK	Normal

Host

Event name	Severity
Cluster Configuration Error	Error
Cluster Configuration OK	Normal
Cold Start	Information
Deleted	Information
Discovered	Information
Down	Critical
Identity Conflict	Warning
Identity OK	Normal
Login Failed	Warning
Login OK	Normal
Modified	Information
Name Changed	Information
SNMP Not Responding	Warning

Event name	Severity
SNMP OK	Normal
System ID Changed	Information
Up	Normal

Host Agent

Event name	Severity
Down	Error
Up	Normal
Host Agent: Login Failed	Warning

Inodes

Event name	Severity
Almost Full	Warning
Full	Error
Utilization Normal	Normal

Interface Status

Event name	Severity
Down	Error
Testing	Normal
Unknown	Normal
Up	Normal

Logical Interface

Event name	Severity
Logical Interface Status Up	Normal
Logical Interface Status Down	Error
Logical Interface Status Unknown	Normal
Logical Interface Migrated	Warning

LUN

Event name	Severity
Offline	Warning
Online	Normal
Snapshot Not Possible	Warning
Snapshot Possible	Normal

Management Station

Event name	Severity
Enough Free Space	Normal
File System File Size Limit Reached	Error
License Expired	Error
License Nearly Expired	Warning
License Not Expired	Normal
Load OK	Normal
Load Too High	Warning
Node Limit Nearly Reached	Warning
Node Limit OK	Normal
Node Limit Reached	Error
Not Enough Free Space	Error
Provisioning Manager Node Limit Nearly Reached	Warning
Provisioning Manager Node Limit Ok	Normal
Provisioning Manager Node Limit Reached	Error
Protection Manager Node Limit Nearly Reached	Warning
Protection Manager Node Limit Ok	Normal
Protection Manager Node Limit Reached	Error

Migration

Event name	Severity
Dataset Not Migrating	Normal
Dataset Migrating	Normal
Dataset Migrated With Errors	Warning
Dataset Migrated	Normal
Dataset Migrate Failed	Error
vFiler Unit Not Migrating	Normal
vFiler Unit Migrating	Normal
vFiler Unit Migrated With Errors	Warning
vFiler Unit Migrated	Normal
vFiler Unit Migrate Failed	Error

NDMP

Event name	Severity
Credentials Authentication Failed	Warning
Credentials Authentication Succeeded	Normal
Communication Initialization Failed	Warning
Communication Initialization Succeeded	Normal
Down	Warning
Up	Normal

Network

Event name	Severity
OK	Normal
Too Large	Warning

Network Services

Event name	Severity
CIFS Service - Up	Normal
CIFS Service - Down	Warning
NFS Service - Up	Normal
NFS Service - Down	Warning
iSCSI Service - Up	Normal
iSCSI Service - Down	Warning
FCP Service - Up	Normal
FCP Service - Down	Warning

No Schedule Conflict

Event name	Severity
Between Snapshot and SnapMirror Schedules	Normal
Between Snapshot and SnapVault Schedules	Normal

NVRAM Battery

Event name	Severity
Discharged	Error
Fully Charged	Normal
Low	Warning
Missing	Error
Normal	Normal
Old	Warning
Overcharged	Warning
Replace	Error
Unknown Status	Warning

OSSV (Open Systems SnapVault)

Event name	Severity
Host Discovered	Information

Performance Advisor

Event name	Severity
Enough Free Space	Normal
Not Enough Free Space	Error

Power Supplies

Event name	Severity
Many Failed	Error
Normal	Normal
One Failed	Error

Primary

Event name	Severity
Host Discovered	Information

Protection Policy

Event name	Severity
Created	Information
Deleted	Information
Modified	Information

Protection Schedule

Event name	Severity
Created	Information
Deleted	Information
Modified	Information

Provisioning Policy

Event name	Severity
Created	Information
Deleted	Information
Modified	Information

Qtree

Event name	Severity
Almost Full	Warning
Files Almost Full	Warning
Files Full	Error
Files Utilization Normal	Normal
Full	Error
Growth Rate Abnormal	Warning
Growth Rate OK	Information
Space Normal	Normal

Remote Platform Management (RPM)

Event name	Severity
Online	Normal
Unavailable	Critical

Resource Group

Event name	Severity
Created	Information
Deleted	Information
Modified	Information

Resource Pool

Event name	Severity
Created	Information
Deleted	Information
Modified	Information
Space Full	Error
Space Nearly Full	Warning
Space OK	Normal

SAN Host LUN Mapping

Event name	Severity
Changed	Warning

Script

Event name	Severity
Critical Event	Critical
Emergency Event	Emergency
Error Event	Error
Information Event	Information
Normal Event	Normal
Warning Event	Warning

SnapMirror

Event name	Severity
Abort Completed	Normal
Abort Failed	Error
Break Completed	Normal
Break Failed	Error
Date OK	Normal

Event name	Severity
Delete Aborted	Warning
Delete Completed	Information
Delete Failed	Error
Initialize Aborted	Warning
Initialize Completed	Normal
Initialize Failed	Error
Nearly Out of Date	Warning
Not Scheduled	Normal
Not Working	Error
Off	Normal
Out of Date	Error
Possible Problem	Warning
Quiesce Aborted	Warning
Quiesce Completed	Normal
Quiesce Failed	Error
Resume Completed	Normal
Resume Failed	Error
Resync Aborted	Warning
Resync Completed	Normal
Resync Failed	Error
Unknown State	Warning
Update Aborted	Warning
Update Completed	Normal
Update Failed	Error
Working	Normal

Snapshot(s)

Event name	Severity
Age Normal	Normal
Age Too Old	Warning
Count Normal	Normal
Count OK	Normal
Count Too Many	Error
Created	Normal
Failed	Error
Full	Warning
Schedule Conflicts with the SnapMirror Schedule	Warning
Schedule Conflicts with the SnapVault Schedule	Warning
Schedule Modified	Information
Scheduled Snapshots Disabled	Information
Scheduled Snapshots Enabled	Normal

SnapVault

Event name	Severity
Backup Aborted	Warning
Backup Completed	Information
Backup Failed	Error
Host Discovered	Information
Relationship Create Aborted	Warning
Relationship Create Completed	Information
Relationship Create Failed	Error
Relationship Delete Aborted	Warning
Relationship Delete Completed	Information
Relationship Delete Failed	Error
Relationship Discovered	Information

Event name	Severity
Relationship Modified	Information
Replica Date OK	Normal
Replica Nearly Out of Date	Warning
Replica Out of Date	Error
Restore Aborted	Warning
Restore Completed	Normal
Restore Failed	Error

SNMP Trap Listener

Event name	Severity
Alert Trap Received	Information
Critical Trap Received	Information
Emergency Trap Received	Information
Error Trap Received	Information
Information Trap Received	Information
Notification Trap Received	Information
Warning Trap Received	Information
Start Failed	Warning
Start OK	Information

Space Management

Event name	Severity
Space Management Job Started	Information
Space Management Job Succeeded	Information
Space Management Job Failed	Information

Storage Services

Event name	Severity
Storage Service Created	Information
Storage Service Modified	Information
Storage Service Destroyed	Information
Storage Service Dataset Provisioned	Information
Storage Service Dataset Attached	Information
Storage Service Dataset Detached	Information

Sync

Event name	Severity
SnapMirror In Sync	Information
SnapMirror Out of Sync	Warning

Temperature

Event name	Severity
Hot	Critical
Normal	Normal

Unprotected Item

Event name	Severity
Discovered	Information

User

Event name	Severity
Disk Space Quota Almost Full	Warning
Disk Space Quota Full	Error
Disk Space Quota OK	Normal
Disk Space Soft Limit Exceeded	Warning
Disk Space Soft Limit Not Exceeded	Normal

Event name	Severity
E-mail Address OK	Normal
E-mail Address Rejected	Warning
Files Quota Almost Full	Warning
Files Quota Full	Error
Files Quota Utilization Normal	Normal
Files Soft Limit Exceeded	Warning
Files Soft Limit Not Exceeded	Normal

vFiler Unit

Event name	Severity
Deleted	Information
Discovered	Information
Hosting Storage System Login Failed	Warning
IP Address Added	Information
IP Address Removed	Information
Renamed	Information
Storage Unit Added	Information
Storage Unit Removed	Information

vFiler Unit Template

Event name	Severity
Created	Information
Deleted	Information
Modified	Information

Vserver

Event name	Severity
Vserver Discovered	Information

Event name	Severity
Vserver Deleted	Information
Vserver Renamed	Information

Volume

Event name	Severity
Almost Full	Warning
Automatically Deleted	Information
Autosized	Information
Clone Deleted	Information
Clone Discovered	Information
Destroyed	Information
First Snapshot OK	Normal
Full	Error
Growth Rate Abnormal	Warning
Growth Rate OK	Normal
Maxdirsize Limit Nearly Reached	Information
Maxdirsize Limit Reached	Information
Nearly No Space for First Snapshot	Warning
Nearly Over Deduplicated	Warning
New Snapshot	Normal
Next Snapshot Not Possible	Warning
Next Snapshot Possible	Normal
No Space for First Snapshot	Warning
Not Over Deduplicated	Normal
Offline	Warning
Offline or Destroyed	Warning
Online	Normal
Over Deduplicated	Error

Event name	Severity
Quota Overcommitted	Error
Quota Almost Overcommitted	Warning
Restricted	Restricted
Snapshot Automatically Deleted	Information
Snapshot Deleted	Normal
Space Normal	Normal
Space Reserve Depleted	Error
Space Reservation Nearly Depleted	Error
Space Reservation OK	Normal

Alarm configurations

The DataFabric Manager server uses alarms to tell you when events occur. The DataFabric Manager server sends the alarm notification to one or more specified recipients: an e-mail address, a pager number, an SNMP trap host, or a script that you write.

You are responsible for which events cause alarms, whether the alarm repeats until it is acknowledged, and how many recipients an alarm has. Not all events are severe enough to require alarms, and not all alarms are important enough to require acknowledgment. Nevertheless, you should configure the DataFabric Manager server to repeat notification until an event is acknowledged, to avoid multiple responses to the same event.

The DataFabric Manager server does not automatically send alarms for the events. You must configure alarms for the events, you specify.

Guidelines for configuring alarms

When configuring alarms you must follow a set of guidelines.

- Alarms must be created for a group, either an individual group or the Global group.
If you want to set an alarm for a specific object, you must first create a group with that object as the only member. You can then configure an alarm for the newly created group.
- Alarms you create for a specific event are triggered when that event occurs.
- Alarms you create for a type of event are triggered when any event of that severity level occurs.
- Alarms can be configured based on events, event severity, or event class.

Related concepts

What events are on page 89

Creating alarms

You can create alarms from the Alarms page in Operations Manager.

Steps

1. Click **Control Center > Setup > Alarms**.
2. From the **Alarms** page, select the group that you want Operations Manager to monitor.
You might need to expand the list to display the one you want to select.
3. Specify what triggers the alarm: an event or the severity of event.
4. Specify the recipient of the alarm notification.
Note: If you want to specify more than one recipient or configure repeat notification, continue to Step 5.
5. Click **Add** to set the alarm. If you want to configure additional options, continue with Step 5.
6. Click **Advanced Version**.
7. Optionally, if you want to specify a class of events that should trigger this alarm, specify the event class.
You can use normal expressions.
8. Optionally, specify the recipients of the alarm notification.
Formats include administrator names, e-mail addresses, pager addresses, or an IP address of the system to receive SNMP traps (or port number to send the SNMP trap to).
9. Optionally, specify the period that Operations Manager sends alarm notifications.
10. Optionally, select **Yes** to resend the alarm notification until the event is acknowledged or **No** to notify the recipients only once.
11. Optionally, set the interval (in minutes) that Operations Manager waits before it tries to resend a notification.
12. Activate the alarm by selecting **No** in the Disable field.
13. Click **Add**.

Testing alarms

You can test the alarms from the Alarms page in Operations Manager.

Steps

1. Click **Control Center > Setup > Alarms**.
2. Find the alarm you want to test and click **Test**.

3. Read the message about the test alarm that the DataFabric Manager server generates, from the **Test Alarm** page and then click **Test**.

Result

The DataFabric Manager server generates an alarm and sends a notification to the recipient.

Comments in alarm notifications

By using the DataFabric Manager server, you can add details to alarm notifications, such as asset number, department code, location name, and support contact.

When you define a custom comment field, the DataFabric Manager server sends this information in the alarm notification to help you respond to the alarm.

Custom alarm notifications are sent by e-mail message or SNMP traps; you can also access them by executing scripts. Custom alarm notifications cannot be sent to pagers.

Example of alarm notification in e-mail format

You can configure alarms to send notifications in the form of e-mail messages.

This example shows alarm notification sent in the e-mail format:

```
From: DataFabric.Manager.on.lnx186-79@company.com  
[mailto:DataFabric.Manager.on.lnx186-79@company.com]  
Sent: Wednesday, July 14, 2010 3:20 PM  
To: xyz  
Subject: dfm: Normal event on deux-186-72:/  
os_f_ds_370_backup_118 (Volume Online)
```

```
A Normal event at 14 Jul 15:27 EDT on Volume  
os_f_ds_370_backup_118 on Active/Active Controller  
deux-186-72.lab.eng.btc.company.in:  
Volume Online.  
Volume online
```

```
Click below to see the source of this event.  
http://lnx186-79.lab.eng.btc.company.in:8080/
```

```
REMINDER: This alert will be repeated every 30 minutes, until  
the event is acknowledged.
```

```
*** Event details follow.***
```

```
General Information
```

```
-----  
DataFabric Manager Serial Number: 1-50-000001 Alarm  
Identifier: 4
```

```

Event Fields
-----
Event Identifier: 4081
Event Name: Volume Online
Event Description: Availability of the file system Event
Severity: Normal Event Timestamp: 14 Jul 15:27

Source of Event
-----
Source Identifier: 3044
Source Name: deux-186-72:/os_f_ds_370_backup_118
Source Type: Volume
Name of the host: deux-186-72.lab.eng.btc.company.in Type of
the host: Active/Active Controller Host identifier: 93

Event Arguments
-----
volName: deux-186-72:/os_f_ds_370_backup_118

```

Example of alarm notification in script format

Various event fields are passed to the script as environment variables. Characters other than [a-z][A-Z][0-9] are replaced with “_” (underscore).

This example shows environment variables passed to the sample script:

```

DFM_ALARM_ID="21"
DFM_EVENT_ID="312"
DFM_EVENT_NAME_GUI="Management Station: Node Limit Ok"
DFM_Source_ID="10.72.1.1/24"

```

Example of alarm notification in trap format

A new SNMP variable is added to all existing traps. The value of this variable is a string, which contains the name and values of all the defined comment fields. If this string is not empty, it is appended to the trap.

The format of this string is as follows:

Source	<i>IP address of</i>
<i>the source</i>	
Timestamp:	7
hours 25 minutes 12 seconds	
Enterprise:	.
1.3.6.1.4.1.789.3	
SNMP Version:	1
Specific:	11041

```
Generic:  
  enterpriseSpecific  
Variable buildings:  
Name:  
1.3.6.1.4.1.789.3.1.0  
Value:  
(octetString)1-50-000001
```

Response to alarms

When you receive an alarm, you should acknowledge the event and resolve the condition that triggered the alarm.

In addition, if the repeat notification feature is enabled and the alarm condition persists, you continue to receive notifications.

Deleting alarms

You can delete alarms from the Alarms page in Operations Manager.

Steps

1. Click **Control Center > Setup > Alarms**.
2. From the **Alarms** page, select the alarm for deletion.
3. Click **Delete Selected**.

Working with user alerts

The DataFabric Manager server can send an alert to you whenever it detects a condition or problem, in your systems that requires attention. You can configure the mail server so that the DataFabric Manager server can send alerts to specified recipients when an event occurs.

What user alerts are

By default, the DataFabric Manager server sends out user alerts (e-mail messages) to all users who exceed their quota limits.

Whenever a user event related to disk or file quotas occurs, the DataFabric Manager server sends an alert to the user who caused the event. The alert is in the form of an e-mail message that includes information about the file system (volume or qtree) on which the user exceeded the threshold for a quota.

You can disable the alerts for all users or for the users who have quotas.

Differences between alarms and user alerts

The DataFabric Manager server uses alarms to tell you when events occur. By default, the DataFabric Manager server sends out user alerts (e-mail messages) to all users who exceed their quota limits.

Alarms	User alerts
Alarms have to be configured for events before the DataFabric Manager server can send out notification to the specified recipients.	The DataFabric Manager server generates user alerts by default.
Alarms can be sent to one or more of the following recipients: <ul style="list-style-type: none"> • An e-mail address • A pager address • An SNMP traphost • A script that you write 	User alerts can be sent to the user who exceeds the user quota thresholds. User alerts can be only in the form of an e-mail message.
Alarms can be sent to only users listed as administrators on the Administrators page of Operations Manager.	User alerts are sent to any user with user quota information in the DataFabric Manager server database.
Alarms can be configured for any events with severity of Information or higher.	User alerts can be sent only when the following user quota events occur: <ul style="list-style-type: none"> • User Disk Space Quota Almost Full • User Disk Space Quota Full • User Files Quota Almost Full • User Files Quota Full

User alerts configurations

To receive user alerts, you must enable the `User Quota Alerts` option, configure e-mail addresses, and the e-mail domain. Optionally, you can configure the `mailmap` file.

If you want to enable or disable alerts to all users in the DataFabric Manager server database, use the `Enable User Quota Alerts` option in the Users section on the Options page.

You can enable, or disable alerts to only the users who have quotas configured on a specific file system, or a group of file systems. Use the `Enable User Quota Alerts` option on the Edit Volume Settings or Edit Qtree Settings page of that volume or qtree.

E-mail addresses for alerts

The e-mail address that the DataFabric Manager server uses to send alerts depends on the DataFabric Manager server configuration.

The following list specifies the checks the DataFabric Manager server makes before selecting an e-mail address for the user:

- There are three ways to specify an e-mail address:
 - Use the Edit User Settings page (**Quotas or File SRM > user name > Edit Settings**)
 - Use the `dfm quota user` command

Note: For more information about the `dfm quota user` command, see the DataFabric Manager server man pages.

- Use the `mailmap` file

If you need to specify e-mail addresses of many users, by using the Edit User Settings page for each user might not be convenient. Therefore, the DataFabric Manager server provides a `mailmap` file that enables you to specify many e-mail addresses in one operation.

- When you do not specify the e-mail address, the default e-mail domain is configured and appended to the user name, by the DataFabric Manager server. The resulting e-mail address is used to send the alert.

Note: The DataFabric Manager server uses only the part of the user name that is unique to the user. For example, if the user name is `company/joe`, Operations Manager uses `joe` as the user name.

- If a default e-mail domain is not configured, again Operations Manager uses the part of the user name that is unique to the user (without the domain information).

Note: If your SMTP server processes only e-mail addresses that contain the domain information, you must configure the domain in the DataFabric Manager server to ensure that e-mail messages are delivered to their intended recipients.

Domains in user quota alerts

You can use the `Default Email Domain for Quota Alerts` option, to specify the domain that Operations Manager appends to the user name when sending out a user quota alert.

The `Quota Alerts` option is in the User Options section on the Options page (**Setup menu > Options**).

If you specify a value for this option, it applies to all users except the ones who are listed with the e-mail domain information in the `mailmap` file.

What the mailmap file is

The `mailmap` file is a simple text file that contains a mapping between the user names and the e-mail addresses of these users.

The file is imported into the DataFabric Manager server database by using the `dfm mailmap import` command. Once the file has been imported, the information in the database is used to find the e-mail addresses of the users.

For more information about the `dfm mailmap import` command, see the DataFabric Manager server man pages.

Example of mailmap file

```
# Start mail map
USER windows_domain\joe joe@company.com
USER jane@nisl.company.com jane@company.com
USER chris@nisdomain1 chris
# End mail map
```

USER	A case-sensitive keyword that must appear at the beginning of each entry
user_name	The Windows or UNIX user name. If the name contains spaces, enclose the name in either double or single quotes.
e-mail address	E-mail address to which the quota alert is sent when the user crosses a user quota threshold

Guidelines for editing the mailmap file

You should follow a set of guidelines for editing the `mailmap` file.

- When specifying a user name for a UNIX user, you must specify the full NIS domain after the user name. For example, for UNIX user joe in NIS domain nisdomain1, specify `joe@nisdomain1` as the user name.
- The specified NIS domain must match the one configured on the storage system. If no domain information is configured on the storage system, you must also leave the domain in the `mailmap` file empty.
- Use one or more spaces or tabs to separate the fields in the file.
- Use the “#” character in the beginning of a line that is a comment.
- If blank lines exist in the file, they are ignored.

How the contents of the user alert are viewed

You can obtain the content of the current e-mail alert that the DataFabric Manager server sends by entering the following command at DataFabric Manager server console: `dfm quota mailformat export { path | - }`. Path is the location of the file containing e-mail alerts.

How the contents of the e-mail alert are changed

You can change the contents of the current e-mail alert by modifying the alerts in the mailformat file.

After you have changed the contents of the e-mail alert, you can import the modified file in to the DataFabric Manager server database by using the `dfm quota mailformat import` command.

For more information about the `dfm quota mailformat` command, see the DataFabric Manager server man pages.

What the mailformat file is

The `mailformat` file is a simple text file that enables you to customize the contents of the e-mail alert that is sent to the users.

This file must contain entries in the following format:

```
mail-headers
<empty line>
body
```

mail-headers The SMTP headers to be sent in the DATA section of the SMTP message

body Body of the e-mail

Any words that begin with `DFM_` are treated as the DataFabric Manager server variables and are replaced by their values. The following table lists the valid variables:

Variable	Variable is replaced with...
<code>DFM_EVENT_NAME</code>	Name of the event
<code>DFM_QUOTA_FILE_SYSTEM_NAME</code>	Name of the file system (volume or qtree) that caused the quota event
<code>DFM_QUOTA_FILE_SYSTEM_TYPE</code>	Type of file system (volume or qtree)
<code>DFM_QUOTA_PERCENT_USED</code>	Percentage of quota used
<code>DFM_QUOTA_USED</code>	Amount of disk space or number of files used
<code>DFM_QUOTA_LIMIT</code>	Total disk space or files quota
<code>DFM_QUOTA_TYPE</code>	Type of quota (disk space or files), depending on whether the disk space or files quota threshold was exceeded

Variable	Variable is replaced with...
<code>DFM_LINK_EVENT</code>	Hyperlink to the event
<code>DFM_QUOTA_USER_NAME</code>	Name of the user exceeding the quota threshold

Example of the mailformat file

```
From: IT Administrator
Subject: URGENT: Your DFM_QUOTA_TYPE quota on
DFM_QUOTA_FILE_SYSTEM_NAME
You (as user "DFM_QUOTA_USER_NAME") have used up
DFM_QUOTA_PERCENT_USED (DFM_QUOTA_USED out of DFM_QUOTA_LIMIT) of
available DFM_QUOTA_TYPE quota on DFM_QUOTA_FILE_SYSTEM_NAME.
Please delete files that you no longer need.
Event (For IT Use only): DFM_LINK_EVENT
-- IT Administrator
```

Guidelines for editing the mailformat file

You should follow a set of guidelines for editing the mailformat file.

- Ensure that the mailformat file conforms to SMTP protocol.
- Specify an empty line between the header and the body of the e-mail message.
- Use any of the headers recognized by the SMTP servers, such as “content type: text/html” to send an e-mail message with HTML formatted body.

Introduction to DataFabric Manager server reports

The DataFabric Manager server provides standard reports that you can view from the CLI or Operations Manager.

You can run reports and create custom reports from the CLI. However, the DataFabric Manager server provides reports in easy-to-use Operations Manager interface, in which you can do the following:

- View a report.
- Save a report in Excel format.
- Print a report.
- Create a report.
- Delete a custom report. You cannot delete a standard report.
- Use a custom report as a template to create another custom report.

By using the DataFabric Manager server 3.6 or later, you can search all the reports from **Reports menu** > **All**. All the reports are divided under the following categories:

- Recently Viewed

- Favorites
- Custom Reports
- Logical Objects
- Physical Objects
- Monitoring
- Performance
- Backup
- Disaster Recovery
- Data Protection Transfer
- Miscellaneous

For more information about these categories, see *Operations Manager Help*.

Note: The report category Performance contains the performance characteristics of objects. However, you can view the complete reports under their respective report categories.

Introduction to report options

The DataFabric Manager server provides standard reports that you can view from the CLI or Operations Manager.

The `dfm report` command specifies the report catalog object that you can modify to create the custom report. The custom report object has the following attributes that you can set:

- A short name (for CLI output)
- A long name (for GUI output)
- Field description
- The fields to display
- The report catalog it was created from

The custom report also has methods that let you create, delete, and view your custom reports. You can configure the report options in Operations Manager with respect to Name, Display tab, Catalogs, and Fields.

Introduction to report catalogs

The DataFabric Manager server provides report catalogs that you use to customize reports. You can set basic report properties from the CLI or Operations Manager.

Following are the basic report properties:

- A short report name (for CLI output)
- Long report name (for Operations Manager output)
- Field description
- The fields to display
- The report catalog it was created from

Every report that is generated by the DataFabric Manager server, including those you customize, is based on the catalogs.

For more information about how to use the CLI to configure and run reports, use the `dfm report help` command. The command specifically describes how to list a report catalog and its fields and command organization and syntax.

Reports in Operations Manager

There are different kinds of reports in Operations Manager.

- Aggregates** The aggregate report shows you space utilization, capacity, performance characteristics, and availability by the volumes on your aggregates. By default, you can view aggregate reports from **Control Center > Home > Member Details > Aggregates > Report**.
- Array LUNs** The array LUNs report shows you information about the LUN residing on the third-party storage arrays that are attached to the gateway. Information such as model, vendor, serial number of the LUN, and size is available in these reports. By default, you can view array LUNs reports from **Control Center > Home > Member Details > Physical Systems > Report**.
- Aggregate Array LUNs** The aggregate array LUNs report shows you information about array LUNs contained on the aggregates of the gateway. Information such as model, vendor, serial number of the LUN, and size is available in these reports. By default, you can view aggregate array LUNs reports from **Control Center > Home > Member Details > Aggregates > Report**.
- Backup** The backup report shows you information about the data transfer during a backup.
- Clusters** The Clusters report shows information about the clusters such as status, serial number of the cluster, number of virtual servers and controllers associated with cluster. You can view the Clusters report from **Control Center > Home > Member Details > Physical Systems > Report**.
- Controllers** The Controllers report shows information about the cluster to which the controller belongs, model and serial number of the controller, and the system ID of the cluster. You can view the Controllers report from **Control Center > Home > Member Details > Physical Systems > Report**.
- Dataset** The dataset report shows you information about the resource, protection, and the conformance status of the dataset. This report also displays information about the policy with which the dataset is associated.
- Dataset** The dataset report shows you information about the data transfer from individual mirror and backup relationships within a dataset.
- Disks** The disks report shows you information about the disks in your storage systems, such as model, vendor, and size. You can view the performance characteristics

and sort these reports by broken or spare disks, as well as by size. By default, you can view disks reports along with the controller reports in the Member Details tab.

- Events** The events report shows you information about event severity, user quotas, and SNMP traps. The information about all events, including deleted, unacknowledged, and acknowledged events, in the DataFabric Manager server database are available in these reports. By default, you can view event reports in the Group Status tab.
- FC Link** The FC link report shows you information about the logical and physical links of your FC switches and fabric interfaces. By default, you can view FC link reports along with the SAN reports in the Member Details tab.
- FC Switch** The FC switch report shows you FC switches that have been deleted, have user comments associated with them, or are not operating. By default, you can view FC switch reports along with the SAN reports in the Member Details tab.
- FCP Target** The FCP target report shows you information about the status, port state, and topology of the target. The FCP target also reports the name of the FC switch, the port to which the target connects and the HBA ports that the target can access. By default, you can view FCP target reports in the **Control Center > Home > Member Details > LUNs tab**.
- File System** The file system report shows you information about all file systems, and you can filter them into reports by volumes, Snapshot copies, space reservations, qtrees, and chargeback information. By default, you can view file system reports in the **Control Center > Home > Member Details > File systems tab**.
- Group Summary** The group summary report shows the status, storage space used, and storage space available for your groups. The FCP target report includes storage chargeback reports that are grouped by usage and allocation. By default, you can view group reports in the Group Status tab.
- Interface Groups** The Interface Groups report shows information about all the cluster interface groups defined, ports, controllers, active state of the group. You can view the Interface Groups report from **Control Center > Home > Member Details > Physical Systems > Report**.
- Host Users** The host users report shows you information about the existing users on the host. By default, you can view host users reports from **Management > Host Users > Report**.
- Host Local Users** The host local users report shows you information about the existing local users on the host. By default, you can view the host local users reports from **Management > Host Local Users > Report**.
- Host Domain Users** The host domain users report shows you information about the existing domain users on the host. By default, you can view the host domain users reports from **Management > Host Domain Users > Report**.

Host Usergroups	The host usergroups report shows you information about the existing user groups on the host. By default, you can view the host usergroups reports from Management > Host Usergroups > Report .
Host Roles	The host roles report shows you information about the existing roles on the host. By default, you can view the host roles reports from Management > Host Roles > Report .
Logical Interfaces	The Logical Interfaces report shows information about the server, the current port that the logical interface uses, status of the logical interface, the network address and mask, the type of role, and if the interface is at the home port or not. You can view the logical Interfaces report from Control Center > Home > Member Details > Physical Systems > Report .
LUN	The LUN report shows you information and statistics about the LUNs and LUN initiator groups on the storage systems, along with performance characteristics. By default, you can view LUN reports in the Member Details tab.
History Performance Report	The History, Performance Events report displays all the Performance Advisor events. By default, you can view the History, Performance Events reports from Group Status > Events > Report .
Mirror	The Mirror report displays information about data transfer in a mirrored relationship.
Performance Events	The Performance Events report displays all the current Performance Advisor events. By default, you can view the Performance Events reports from Control Center > Home > Group Status > Events > Report .
Ports	The Ports report shows information about the controllers that are connected, the type of role that the port portrays, status of the port, and sizes of data moving in and out of the port. You can view the Ports report from Control Center > Home > Member Details > Physical Systems > Report .
Quotas	The quota report shows you information about user quotas that you can use for chargeback reports. By default, you can view quota reports along with the group summary reports in the Group Status tab.
Report Outputs	The Report Outputs report shows you information about the report outputs that are generated by the report schedules. By default, you can view Report Outputs reports in the Reports > Schedule > Saved Reports .
Report Schedules	The Report Schedules report shows you information about the existing report schedules. By default, you can view Report Schedules reports in the Reports > Schedule > Report Schedules . A report schedule is an association between a schedule and a report for the report generation to happen at that particular time.

- Resource Pools** The Resource Pools report shows you information about the storage capacity that is available and the capacity that is used by all the aggregates in the resource pool. This report also displays the time zone and the status of the resource pool.
- SAN Host** The SAN host report shows you information about SAN hosts, including FCP traffic and LUN information, and the type and status of the SAN host. By default, you can view SAN host reports in the Member Details tab.
- Schedules** The Schedules report shows you information about the existing schedules and the names of the report schedules that are using a particular schedule. By default, you can view Schedules reports in the **Reports > Schedule > Schedules**. The Schedules tab displays all the schedules. Schedules are separate entities which can be associated with reports. The Manage Schedules link in the Reports-Add a Schedule and Reports-Edit a Schedule page points to this page.
- Scripts** The Scripts report shows you information about the script jobs and script schedules. By default, you can view the Scripts reports in the Member Details tab.
- Spare Array LUNs** The Spare Array LUNs report shows you information about spare array LUNs of the gateway, such as model, vendor, serial number of the LUN, and size. By default, you can view the Spare Array LUNs reports from **Control Center > Home > Member Details > Physical Systems > Report**.
- SRM** The SRM report shows you information about the SRM files, paths, directories, and host agents. The file space statistics which is reported by an SRM path walk differ from the "volume space used" statistics that is provided by the file system reports. By default, you can view SRM reports in the Group Status tab.
- storage systems** The storage systems report shows you information about the capacity and operations of your storage systems, performance characteristics, and the releases and protocols running on them. By default, you can view storage systems reports in the **Control Center > Home > Member Details > Physical Systems tab**.
- User Quotas** The User Quotas report shows you information about the disk space usage and user quota thresholds collected from the monitored storage systems. By default, you can view User Quotas reports along with the SRM reports in the Group Status tab.
- vFiler** The vFiler report shows you the status, available protocols, storage usage, and performance characteristics of vFiler units that you are monitoring with the DataFabric Manager server. By default, you can view the vFiler reports in the Member Details tab.
- Virtual Servers** The Virtual Servers report shows information about the associated cluster, the root volume on which the virtual server resides name of the service switch, NIS domain, and the status of the virtual server. You can view the Virtual Servers report from **Control Center > Home > Member Details > Virtual Systems > Report**.

Volume The Volume report shows you all volumes with the following details, for the current month or for the past month:

- Name
- Capacity
- Available space
- Snapshot capacity
- Growth rates
- Expendability
- Chargeback by usage or allocation

The reports also show the performance characteristics of volumes. By default, you can view the volume reports along with the file system reports in the Member Details tab.

What performance reports are

Performance reports in Operations Manager provide information regarding object's performance characteristics.

The following points describe the important features of performance reports:

- You can view the object's performance characteristics for the periods last, one day, one week, one month, three months, and one year.
- You can also view the performance counters related to various fields in catalogs.
- Data consolidation is available only if you select the option Performance. Data Consolidation is a statistical technique that helps you to analyze data. You can view the average, minimum, maximum, or median value for the performance metrics over a period; however, this field is set to Average by default.

Related concepts

[Report fields and performance counters](#) on page 297

Configuring custom reports

You can configure custom reports in Operations Manager.

Steps

1. Select **Custom** from the Reports menu.
2. Enter a (short) name for the report, as you want it to display in the CLI.
3. Optionally, enter a (long) name for the report, as you want it to display in Operations Manager.
4. Optionally, add comments to the report description.
5. Select the catalog from which the available report fields are based.

6. Select where you want the DataFabric Manager server to display this report in Operations Manager.
7. Select the related catalog from which you want to choose fields.
You might need to expand the list to display the catalog you want to select.
8. You can view two types of information fields in the “Choose From Available Fields” section:
 - To view fields related to the usage and configuration metrics of the object, click **Usage**.
 - To view fields related to performance metrics of the object, click **Performance**.
9. Select a field from "Choose From Available Fields."
10. Optionally, enter the name for the field, as you want it displayed on the report.
Make your field name abbreviated and as clear as possible. You must be able to view a field name in the reports and determine which field the information relates to.
11. Optionally, specify the format of the field.
If you choose not to format the field, the default format displayed is used.
12. Click **Add** to move the field to the Reported Fields list.
13. Repeat Steps 8 to 12 for each field you want to include in the report.
14. Optionally, click **Move Up** or **Move Down** to reorder the fields.
15. If you clicked Performance, select the required data consolidation method from the list.
16. Click **Create**.
17. To view this report, locate this report in the list at the lower part of the page and click the **display** tab name. Find the report from the Report drop-down list.

Deleting custom reports

You can delete a custom report you no longer need, in Operations Manager.

Steps

1. Select **Custom** from the Reports menu.
2. Find the report from the list of configured reports and select the report you want to delete.
3. Click **Delete**.

Putting data into spreadsheet format

You can put data from any of the about LUNs, SAN hosts, and FCP targets reports into spreadsheet format.

Before you begin

Reports about LUNs, SAN hosts, and FCP targets must be available on the LUNs page of Member Details tab.

Steps

1. Click **Member Details** on the **LUNs** page to view the reports.
2. You can view the data in report in a spreadsheet format, by clicking on the spreadsheet icon on the right side of the Report drop-down list.

Result

You can use the data in the spreadsheet to create your own charts and graphs or to analyze the data statistically.

What scheduling report generation is

Operations Manager allows you to schedule the generation of reports.

The report can include the following statistics:

- Volume capacity used
- CPU usage
- Storage system capacity
- Storage system up time

What Report Archival Directory is

Report Archival Directory is a repository where all the reports are archived.

You can modify the location of the destination directory by using the following CLI command: `dfm options set reportsArchiveDir=<destination dir>`.

When you modify the Report Archival Directory location, the DataFabric Manager server checks whether the directory is writable to archive the reports.

In the case of a Windows operating system, if the directory exists on the network, then the destination directory must be a UNC path. Besides, to save the reports, the scheduler service must run with an account that has write permissions on the directory. The server service must run with an account that has read and delete permissions on the directory to view and delete report output, respectively. The permissions for a service can be configured using Windows Service Configuration Manager.

Note: You require the Database Write capability on the Global group to modify the Report Archival Directory option.

Additional capabilities for categories of reports

You require report-specific read capability on the object in addition to the Database Read capability, for categories of reports such as SRM Reports, Event Reports, and so on.

The capabilities that you require for the categories of reports are as follows:

Report category	Capabilities
SRM Reports	SRM Read capability
Event Reports	Event Read capability
Mirror Reports	Mirror Read capability
Policy Reports	Policy Read capability
BackupManager Reports	BackupManager Read capability

What Report Schedules reports are

The Report Schedules report shows you information about the existing report schedules. A report schedule is an association between a report and a schedule for the report to be generated at a particular time.

By default, Report Schedules reports display in **Reports > Schedule > Report Schedules**.

Scheduling a report using the All submenu

You can schedule a report using the All submenu from the Reports menu in Operations Manager.

Steps

1. From any page, click **Reports > All** to display the **Report Categories** page.
By default, the Recently Viewed category appears.
2. Select a report of your choice.
3. Click **Show** to display the selected report.
4. Click the **Schedule This Report** icon, , located in the upper right corner of the page.
5. In the **Reports - Add a Schedule** page, specify the report schedule parameters.
For details about the report schedule parameters, see the *Operations Manager Help*.
6. Click **Add**.

Scheduling a report using the Schedule submenu

You can schedule a report using the Schedule submenu from the Reports menu.

Steps

1. From any page, click **Reports > Schedule** to display all the report schedules.
2. Click **Add New Report Schedule**.
3. In the **Reports - Add a Schedule** page, specify the report schedule parameters.
For details about the report schedule parameters, see the *Operations Manager Help*.
4. Click **Add**.

Methods to schedule a report

You can schedule a report in two possible methods from the Reports menu.

Following are the two methods with which you can schedule a report:

- Using the Schedule submenu from the Reports menu
- Using the All submenu from the Reports menu

Editing a report schedule

You can edit a report schedule using the Schedule submenu from the Reports menu.

Steps

1. From any page, click **Reports > Schedule** to display all the report schedules.
2. Click the report schedule that you want to edit.
Alternatively, click **Saved Reports** to list all the report outputs, and then click **Report Schedules** entry, that you want to edit.
3. In the Reports - Edit a Schedule page, edit the report schedule parameters.
4. Click **Update**.

Deleting a report schedule

You can delete a report schedule using the Schedule submenu from the Reports menu.

Steps

1. From any page, click **Reports > Schedule** to display all the report schedules.
2. Select the report schedule that you want to delete.

3. Click **Delete Selected**.

Enabling a report schedule

You can enable a report schedule using the Schedule submenu from the Reports menu.

Steps

1. From any page, click **Reports > Schedule** to display all the report schedules.
2. Select the report schedule that you want to enable.
3. Click **Enable Selected**.

Disabling a report schedule

You can disable a report schedule using the Schedule submenu from the Reports menu.

Steps

1. From any page, click **Reports > Schedule** to display all the report schedules.
2. Select the report schedule that you want to disable.
3. Click **Disable Selected**.

Running a report schedule

You can run a report schedule using the Schedule submenu from the Reports menu.

Steps

1. From any page, click **Reports > Schedule** to display all the report schedules.
2. Select the report schedule that you want to run.
3. Click **Run Selected**.

Retrieving the list of enabled report schedules

You can retrieve the list of enabled report schedules using the Schedule submenu from the Reports menu.

Steps

1. From any page, click **Reports > Schedule** to display all the report schedules.
2. Select the Report Schedules, Enabled entry from the Report drop-down list.

Retrieving the list of disabled report schedules

You can retrieve the list of disabled report schedules using the Schedule submenu from the Reports menu.

Steps

1. From any page, click **Reports > Schedule** to display all the report schedules.
2. Select the Report Schedules, Disabled entry from the Report drop-down list.

Listing all the run results of a report schedule

You can list all the run results of a report schedule using the Schedule submenu from the Reports menu.

Steps

1. From any page, click **Reports > Schedule** to display all the report schedules.
2. Click **Last Result Value** of a report schedule to display the run result for that particular report schedule.

What Schedules reports are

The Schedules report shows you information about the existing schedules and the names of the report schedules that are using a particular schedule.

By default, Schedules reports display in **Reports > Schedule > Schedules**. The Schedules tab displays all the schedules.

Schedules are separate entities that can be associated with reports.

Listing all the schedules

You can list all the schedules using the Schedule submenu from the Reports menu.

Steps

1. From any page, click **Reports > Schedule** to display all the report schedules.
2. Click **Schedules**.

Adding a new schedule

You can add a new schedule using the Schedule submenu from the Reports menu.

Steps

1. From any page, click **Reports > Schedule** to display all the report schedules.
2. Click the Schedules tab.
3. Click **Add New Schedule**.
4. In the Schedules - Add a Schedule page, specify the schedule parameters.
For details about the schedule parameters, see the *Operations Manager Help*.
5. Click **Add**.

Editing a schedule

You can edit a schedule using the Schedule submenu from the Reports menu.

Steps

1. From any page, click **Reports > Schedule** to display all the report schedules.
2. Click **Schedules**.
3. Click the schedule you want to edit.
4. In the Schedules - Edit a Schedule page, edit the schedule parameters.
5. Click **Update**.

Deleting a schedule

You can delete a schedule using the Schedule submenu from the Reports menu.

Steps

1. From any page, click **Reports > Schedule** to display all the report schedules.
2. Click **Schedules**.
3. Select the schedule that you want to delete.

Note: If the schedule is used by a report schedule, then the schedule cannot be selected for deletion.

4. Click **Delete Selected**.

What Saved reports are

The Saved reports display information about report outputs such as Status, Run Time, and the corresponding report schedule, which generated the report output.

By default, Saved reports display in **Reports > Schedule > Saved Reports**. The Saved Reports tab displays the list of all the report outputs that are generated by the report schedules.

Listing the report outputs

You can list the report outputs that are generated by all the report schedules using the Schedule submenu from the Reports menu.

Steps

1. From any page, click **Reports > Schedule** to display all the report schedules.
2. Click **Saved Reports** to display the list of report outputs.

Listing the successful report outputs

You can list the successful report outputs generated by all the report schedules using the Schedule submenu from the Reports menu.

Steps

1. From any page, click **Reports > Schedule** to display all the report schedules.
2. Click **Saved Reports**.
3. Select the Report Outputs, Successful entry from the Report drop-down list.

Listing the failed report outputs

You can list the failed report outputs that are generated by all the report schedules using the Schedule submenu from the Reports menu.

Steps

1. From any page, click **Reports > Schedule** to display all the report schedules.
2. Click **Saved Reports**.
3. Select the Report Outputs, Failed entry from the Report drop-down list.

Viewing the output of report outputs from the status column

There are two possible methods to view the output of a particular Report Output that is generated by a report schedule.

About this task

You can also view the output of a Report Output from the Output ID column in Operations Manager.

Steps

1. From any page, select **Schedule** from the Reports menu.
2. Click **Saved Reports**.
3. Click the link under the Status column corresponding to the report output.

Viewing the output of report outputs from the Output ID column

There are two possible methods to view the output of a particular Report Output, which is generated by a report schedule.

Steps

1. From any page, select **Schedule** from the Reports menu.
2. Click **Saved Reports**.
3. Click the Output ID column entry of the Report Output.
4. Click the Output link to view the output.

Viewing the output details of a particular report output

You can view the output details of a particular report output, which is generated by a report schedule in Operations Manager.

Steps

1. From any page, select **Schedule** from the Reports menu.
2. Click **Saved Reports**.
3. Click the Output ID entry of the report output.

Data export in the DataFabric Manager server

By using third-party tools, you can create customized reports from the data you export from the DataFabric Manager server and Performance Advisor.

Operations Manager reports are detailed reports of storage system configuration and utilization. Many administrators create customized reports to accomplish the following tasks:

- Forecasting future capacity and bandwidth utilization requirements
- Presenting capacity and bandwidth utilization statistics
- Generating performance graphs
- Presenting monthly Service Level Agreement (SLA) reports

Data export provides the following benefits:

- Saves effort in collecting up-to-date report data from different sources
- Provides database access to the historical data collected by the DataFabric Manager server
- Provides database access to the information provided by the custom report catalogs in the DataFabric Manager server
- Provides and validates the following interfaces to the exposed DataFabric Manager server views:
 - Open Database Connectivity (ODBC)
 - Java Database Connectivity (JDBC)
- Enables you to export the Performance Advisor and DataFabric Manager server data to text files, easing the loading of data to user-specific database
- Allows you to schedule the export
- Allows you to customize the rate at which the performance counter data is exported
- Allows you to specify the list of the counters to be exported
- Allows you to consolidate the sample values of the data export

How to access DataFabric Manager server data

By using third-party tools, you can create customized reports from the data you export from the DataFabric Manager server. By default, you cannot access the DataFabric Manager server views. To gain access to the views that are defined within the embedded database of the DataFabric Manager server, you need to first create a database user and then enable database access to this user.

You can access the DataFabric Manager server data through views, which are dynamic virtual tables collated from data in the database. These views are defined and exposed within the embedded database of the DataFabric Manager server.

Note: A database user is user-created and authenticated by the database server. Database users are not related to the DataFabric Manager server users.

Before you can create and give access to a database user, you must have the CoreControl capability. The CoreControl capability allows you to perform the following operations:

- Creating a database user
- Deleting a database user
- Enabling database access to a database user
- Disabling database access to a database user
Disabling the database access denies the read permission on the DataFabric Manager server views for the user account.
- Changing the password for the database user

All of these operations can be performed only through the CLI. For more information about the CLI commands, see the DataFabric Manager server manual (man) pages.

You can use a third-party reporting tool to connect to the DataFabric Manager server database for accessing views. Following are the connection parameters:

- Database name: `monitordb`
- User name: `<database user name>`
- Password: `<database user password>`
- Port: `2638`
- `dobroad`: `none`
- Links: `tcPIP`

Note: The .jar files required for iAnywhere and jConnect JDBC drivers are copied as part of the DataFabric Manager server installation. The new jar files are saved in the following directory path: `.../install/misc/dbconn`.

Where to find the database schema for the views

You can find the schema for the database views in the Operations Manager Help. Based on the database schema presented, you can choose the objects and the metrics you want to present.

Two types of data for export

Two types of data that you can export are the DataFabric Manager server data and Performance Advisor data.

- The DataFabric Manager server data
DataFabric Manager server data export is controlled at the global level through the `dfmDataExportEnabled` option. By default, the value of the `dfmDataExportEnabled` `global` option is `No`.
- Performance Advisor data
Performance Advisor data export is controlled at the global level and at the host level through the `perfDataExportEnabled` option. By default, the value of the `perfDataExportEnabled` option is `No`.

By default, in the first Performance Advisor data export run, the counter data for the last seven days is exported. The sampling rate for the counter data export is customizable at the global level. By default, one sample is exported at every 15 minutes.

You can consolidate the sample values of the data export only if the sampling interval is greater than the interval with which the counter data is collected. By default, the average method is used to consolidate the sample values. For more information about the samples and counters, see the *Performance Advisor Administration Guide*.

You can export the data, either on-demand or on-schedule, to text files using the CLI. For more information about the CLI commands, see the DataFabric Manager server manual (man) pages. To schedule the data export you must have the CoreControl capability.

Note: Database views are created within the DataFabric Manager server embedded database. This might increase the load on the database server if there are many accesses from the third-party tools to the exposed views.

Related information

Performance Advisor Administration Guide - www.ibm.com/storage/support/nseries

Files and formats for storing exported data

The exported the DataFabric Manager server and Performance Advisor data is stored in the export_<timestamp> directory located under the top-level directory specified by the dataExportDir global option. By default, the value of the dataExportDir global option is <DFM-install-dir>/dataExport.

Format for exported the DataFabric Manager server data

The exported DataFabric Manager server data is stored in files that are named after the views.

For example, you can store all the iGroup information in the iGroupView file, in the following format:

```
File name: iGroupView
Contents:
<iGroupId> <hostId> <type> <OSType>
... ..
... ..
```

The fields in each line of the file correspond to the columns in iGroupView.

Format for exported Performance Advisor data

The exported Performance Advisor data is stored in different files such as perfHosts, perfCounters, and perfObjInstances.

- perfHosts
This file contains information about the storage systems from which the counter data is collected.
- perfCounters

This file contains information about the various Performance Advisor counters.

- `perfObjInstances`
This file contains information about the performance object instances on storage systems for which the counter data is collected.
- `samples_<objType>_<hostId>`
This file contains the sample values that are collected at various timestamps for different counters and object instances.

The format of these files is as follows:

```
File name: perfHosts
Contents:
host-id host-name
... ..
File name: perfCounters
Contents:
counter-id counter-name description obj-type counter-unit
... ..
File name: perfObjInstances
Contents:
instance-id instance-name host-id obj-type obj-id
... ..
File name: samples_<objType>_<hostId>
Contents:
instance-id counter-id sample-time sample-value
... ..
```

Format for last updated timestamp

The last updated timestamp for both the DataFabric Manager server and Performance Advisor data export is stored in a configuration file named `export.conf` under the `dataExport` directory.

The entries in the `export.conf` file are in the following format:

```
Database Format Version: 1.0
Export Type: [Scheduled | On-demand]
Export Status: [Success | Failed | Canceled | Running]
Delimiter: [tab | comma]
Sampling Interval: <secs>
Consolidation Method: [average | min | max | last]
History: <secs>
DataFabric Manager Data Export Completion Timestamp: <timestamp>
Last PA data Export for following hosts at time <timestamp>
-----<host-name>-----
-----<host-name>-----
```

Security configurations

You can configure Secure Sockets Layer (SSL) in the DataFabric Manager server to monitor and manage storage systems over a secure connection by using Operations Manager.

Types of certificates in the DataFabric Manager server

The DataFabric Manager server uses the signed certificates for secure communication. Signed certificates provide your browser with a way to verify the identity of the storage system.

The DataFabric Manager server uses the following types of signed certificates:

- Self-signed certificates
- Trusted Certificate Authority (CA) signed certificates

Self-signed certificates in the DataFabric Manager server

You can generate self-signed certificates by using the DataFabric Manager server. You can set up the DataFabric Manager server as a Certificate Authority (CA), and generate self-signed certificates.

By issuing self-signed certificates, you can avoid the expense and delay of obtaining a certificate from an external trusted CA. Self-signed certificates are not signed by a mutually trusted authority for secure Web services.

When the DataFabric Manager server sends a self-signed certificate to a client browser, the browser has no way of verifying the identity of the DataFabric Manager server. As a result, the client browser displays a warning indicating to accept the certificate. After the browser accepts the certificate, the DataFabric Manager server is allowed to permanently import the certificate.

If you decide to issue self-signed certificates, you must safeguard access to, and communications with, the DataFabric Manager server and the file system that contains its Secure Socket Layer (SSL) - related private files.

Trusted Certificate Authority (CA) signed certificates in the DataFabric Manager server

You can generate trusted CA signed certificates using the DataFabric Manager server. When the DataFabric Manager server sends a trusted CA signed certificate to the client browser, the browser verifies the identity of the server.

You should obtain a trusted CA signed certificate by generating a Certificate Signing Request (CSR) in the DataFabric Manager server, and then submitting that request to a trusted authority for secure Web services. The DataFabric Manager server accepts certificates from Thawte, Verisign, and RSA.

Creating self-signed certificates in the DataFabric Manager server

You can generate self-signed certificate from the command-line interface (CLI) of the DataFabric Manager server. You can set up the DataFabric Manager server as a Certificate Authority (CA), and generate self-signed certificates.

Steps

1. Log into the DataFabric Manager server as the DataFabric Manager server administrator.
2. From the command-line interface, enter the following command:

```
dfm ssl server setup
```

3. Enter the following information when prompted:

- Key Size
- Certificate Duration
- Country Name
- State or Province
- Locality Name
- Organization Name
- Organizational Unit Name
- Common Name
- Email Address

Result

The DataFabric Manager server SSL server is now initialized with a self-signed certificate and the private key, `server.key` file is placed in the following DataFabric Manager server directory: `c:\Program Files\IBM\DataFabric Manager\DFM\conf\`.

Obtaining a trusted CA-signed certificate

You can obtain a certificate from a trusted Certificate Authority (CA) by running commands from the DataFabric Manager server command-line interface (CLI).

Steps

1. Enter the following command:

```
dfm ssl server req -o filename
```

The DataFabric Manager server creates a certificate signing request (CSR) file.

2. Submit the CSR to a CA for signing.
3. Import the signed certificate by entering the following command:

```
dfm ssl server import cert_filename
```

cert_filename is the certificate filename.

You can view the DataFabric Manager server private key and certificate files by running the command, `dfm ssl server show`.

Enabling HTTPS

You can use the `httpsEnabled` option by using the DataFabric Manager server CLI for the DataFabric Manager server to provide HTTPS services.

Before you begin

Ensure that you have set up the SSL server by using the `dfm ssl server setup` command.

Steps

1. Enter the following command:

```
dfm option set httpsEnabled=Yes
```

2. Change the HTTPS port by entering the following command:

```
dfm option set httpsPort=port_number
```

The default HTTPS port is 8443.

3. Stop the Web server by using the following command:

```
dfm service stop http
```

4. Start the Web server by using the command:

```
dfm service start http
```

This restarts the service using the trusted Certificate Authority (CA) certificate.

Secure communications with the DataFabric Manager server

Secure communications require a secure connection at both ends of each communications link. In the DataFabric Manager server, the two ends of a communication link consist of a secure server and secure managed host.

Clients, including browsers and managed storage systems, must use a secure connection to connect to the DataFabric Manager server. The DataFabric Manager server, in turn, uses a secure connection to connect to a storage system.

How clients communicate with the DataFabric Manager server

The DataFabric Manager server and the clients use a set of protocols such as HTTPS and Secure Shell (SSH) to communicate with each other.

The system on which the DataFabric Manager server is installed, the clients use the following combination of protocols running over SSL:

- Browsers use HTTPS to connect to a secure the DataFabric Manager server.
- The DataFabric Manager server connects to managed hosts by using SSH for operational purposes.
- The DataFabric Manager server connects to managed storage systems by using HTTPS for monitoring purposes.

SecureAdmin for secure connection with the DataFabric Manager server clients

To enable secure connection, you must have SecureAdmin installed on your storage systems.

SecureAdmin is an add-on software module that enables authenticated, command-based administrative sessions between an administrative user and storage systems over an intranet or the Internet.

This combination of SSL and SecureAdmin allows you to securely monitor and manage your storage systems in the DataFabric Manager server.

Related information

www.ibm.com/storage/support/nseries/

Requirements for security options

The security options in the DataFabric Manager server have the following requirements:

- If you disable HTTP and enable HTTPS, all browsers must connect to the DataFabric Manager server through HTTPS.
- If you want to enable secure connections from any browser, you must enable HTTPS transport on the DataFabric Manager server.
- You cannot disable both HTTP and HTTPS transports. The DataFabric Manager server does not allow that configuration.

To completely disable access to Operations Manager, stop the HTTP service at the CLI using the following command: `dfm service stop http`.

- You must select the default port for each transport type you have enabled. The ports must be different from each other.

Guidelines to configure security options in Operations Manager

You should configure security options in Operations Manager using secure protocol.

When configuring the DataFabric Manager server for SSL, you are responsible for safeguarding access to and communications with the DataFabric Manager server and the file system that contains the SSL-related private files. The DataFabric Manager server should not be accessed by not secure protocols, such as Telnet and RSH. Instead, use a secure, private network, or a secure protocol, such as SSH, to connect to the DataFabric Manager server.

Managed host options

You can configure the managed host options to ensure secure communication between the DataFabric Manager server and storage systems.

You can select conventional (HTTP) or secure (HTTPS) administration transport for API communication and conventional (RSH) or secure (SSH) login protocol for login connection.

Where to find managed host options

You can set managed host options by using both GUI and command-line interface.

The locations of managed host options are described in the following table:

Option type	GUI	Command-line interface
Global	Options page (Setup > Options)	<code>dfm option list</code> (to view) <code>dfm option set</code> (to set)
Appliance-specific	Edit Storage Controller Settings page (Controllers > controller name > Storage Controller Tools > Edit Settings)	<code>dfm host get</code> (to view) <code>dfm host set</code> (to set)

Guidelines for changing managed host options

You can change managed host options, such as the login protocol, transport protocol, port, and `hosts.equiv` option.

Login Protocol

This option enables you to set login protocols (RSH or SSH) that the DataFabric Manager server uses when connecting to the managed hosts.

- Login connections.
- Active/active configuration operations
- The `dfm run` command for running commands on the storage system

Change the default value if you want a secure connection for active/active configuration operations, running commands on the storage system.

Administration Transport	<p>This options enables you to select a conventional (HTTP) or secure (HTTPS) connection to monitor and manage storage systems through APIs (XML).</p> <p>Change the default value if you want a secure connection for monitoring and management.</p>
Administration Port	<p>This options enables you to configure the administration port which, along with administration transport, monitors and manages storage systems.</p> <p>If you do not configure the port option at the storage system-level, the default value for the corresponding protocol is used.</p>
hosts.equiv option	<p>This option enables users to authenticate storage systems when the user name and password are not provided.</p> <p>You must change the default value if you have selected the global default option and if you do not want to set authentication for a specific storage system.</p>

Note: If you do not set the transport and port options for a storage system, then the DataFabric Manager server uses SNMP to get storage system-specific transport and port options for communication. If SNMP fails, then the DataFabric Manager server uses the options set at the global level.

Comparison between global and storage system-specific managed host options

You can set managed host options globally, for all storage systems, or individually, for specific storage systems.

If you set storage system-specific options, the DataFabric Manager server retains information about the security settings for each managed storage system. It references this information when deciding whether to use one of the following options to connect to the storage system:

- HTTP or HTTPS
- RSH or SSH
- Login password
- hosts.equiv authentication

If a global setting conflicts with a storage system-specific setting, the storage system-specific setting takes precedence.

Note: You must use storage system-specific managed host options if you plan to use SecureAdmin on some storage systems and not on others.

Limitations in managed host options

You can enable managed host options, but you must accept the following known limitations:

- The DataFabric Manager server cannot connect to storage systems without SecureAdmin installed or to older storage systems that do not support SecureAdmin.
- On storage systems, SecureAdmin is not mutually exclusive with HTTP access. Transport behavior is configurable on the storage system with the `httpd.admin.access` option. The `http.admin.ssl.enable` option enables HTTPS access. For more information, see the documentation for your storage system.
- If you have storage systems running SecureAdmin 2.1.2R1 or earlier, HTTPS options do not work with self-signed certificates. You can work around this problem by using a trusted CA-signed certificate.
- If the `hosts.equiv` option and `login` are set, then the `hosts.equiv` option takes precedence.

Changing password for storage systems in the DataFabric Manager server

You can change the password for an individual storage system on the Edit Storage Controller Settings page using Operations Manager.

Steps

1. Go to the **Storage Controller Details** page for the storage system or hosting storage system (of the vFiler unit) and choose **Edit Settings** from the Storage Controller Tools (at the lower left of Operations Manager).

The Edit Storage Controller Settings page is displayed.

2. In the Login field, enter a user name that the DataFabric Manager server uses to authenticate the storage system or the vFiler unit.
3. In the Password field, enter a password that the DataFabric Manager server uses to authenticate the storage system or the vFiler unit.
4. Click **Update**.

Changing passwords on multiple storage systems

The DataFabric Manager server can set passwords on all storage systems when you use the same authentication credentials on each. Select the Global group to set the same passwords on all storage systems at once.

Steps

1. Log in to Operations Manager.
2. Depending on the type of storage system for which you want to manage passwords, select either of the following:

Type	Procedure
Storage system	Management > Storage System > Passwords
vFiler unit	Management > vFiler > Passwords

3. Enter the user name.
4. Enter the old password of the local user on the host.
5. Enter a new password for the storage system or groups of storage systems.
6. Reenter the new password exactly the same way in the Confirm New Password field.
7. Select the target storage system or target groups.
8. Click **Update**.

Issue with modification of passwords for storage systems

When modifying passwords for a large number of storage systems, you might get an error message if the length of your command input exceeds the specified limit.

This error occurs only when you are using the Operations Manager graphical user interface and not the CLI. If this error occurs, you can take either of the following corrective actions:

- Select fewer storage systems.
- Create a resource group and assign the selected storage systems to the group as members, then modify the password for the group.

Using `hosts.equiv` to control authentication in the DataFabric Manager server

You can control authentication of storage system, vFiler units, and active/active configuration using `hosts.equiv` file, without specifying the user name and password. If the storage system or vFiler

unit is configured with IPv6 addressing, then you cannot use the `host.equiv` file to authenticate the storage system or vFiler unit.

Steps

1. Edit the `/etc/hosts.equiv` file on the storage system and provide either the host name or the IP address of the system running the DataFabric Manager server, as an entry in the following format: `<host-name-or-ip-address>`.
2. You can edit the option on **Edit Appliance Settings** page in Operations Manager. Alternatively, provide the host name, or the IP address of the system running the DataFabric Manager server, and the user name of the user running the DataFabric Manager server CLI, in the following format: `<host-name-or-ip-address> <username>`.
- 3.

If the operating system is...	Then...
Linux	Provide the host name or the IP address of the system running the DataFabric Manager server, and the user name of the user running the HTTP service, in the following format: <code><host-name-or-ip-address> <username></code> .
Windows	Provide the host name, or the IP address of the system running the DataFabric Manager server, and the user name of the user running the HTTP, server, scheduler, and monitor services, in the following format: <code><host-name-or-ip-address> <username></code> .

Note:

By default, the HTTP service runs as `nobody` user on Linux.

By default, the HTTP, server, scheduler, and monitor services run as `LocalSystem` user on Windows.

If the DataFabric Manager server is running on a host named `DFM_HOST`, and `USER1` is running the `dfm` commands, then by default, on a Linux operating system, you need to provide the following entries:

```
DFM_HOST
```

```
DFM_HOST USER1
```

```
DFM_HOST nobody
```

On Windows operating system, you need to provide the following entries:

```
DFM_HOST
```

```
DFM_HOST USER1
```

```
DFM_HOST SYSTEM
```

For more information about configuring the `/etc/hosts.equiv` file, see the *Data ONTAP Storage Management Guide*.

Related information

[Data ONTAP Storage Management Guide-www.ibm.com/storage/support/nseries](http://www.ibm.com/storage/support/nseries)

Editing HTTP and monitor services to run as different user

You can configure HTTP and monitor services using Operations Manager.

Step

1.

Operating system	Command
Linux	<code>dfm service runas -u <user-name> http</code>
Windows	<code>dfm service runas -u <user-name> -p <password> [http] [monitor]</code>

Note: For security reasons the `<user-name>` cannot be “root” on Linux. On Windows hosts, `<user-name>` should belong to the administrator group.

What File Storage Resource Management does

File Storage Resource Manager (FSRM) enables you to gather file-system metadata and generate reports on different characteristics of that metadata.

The DataFabric Manager server interacts with the DataFabric Manager Host Agent residing on remote Windows, Solaris, or Linux workstations or servers (called hosts) to recursively examine the directory structures (paths) you have specified.

For example, if you suspect that certain file types are consuming excessive storage space on your storage systems, you can perform the following tasks:

1. Deploy one or more host agents.
2. Configure FSRM to walk a path.

The host agents might have a Data ONTAP LUN, volume, or qtree mounted. You can configure FSRM to generate reports periodically. These reports contain the following details:

- Files that are consuming the most space
- Files that are outdated or have been accessed recently
- Types of files (.doc, .gif, .mp3, and so on) on the file system

You can then decide how to most efficiently use your existing space.

Note: The File SRM tab in the Operations Manager includes other storage monitoring utilities: for example, chargeback and quota reporting.

How FSRM monitoring works

The DataFabric Manager server monitors directory paths that are visible to the host agent. Therefore, if you want to enable FSRM monitoring of storage systems, the remote host must mount a share using NFS or CIFS, or the host must use a LUN on the storage system.

Note: The DataFabric Manager server cannot obtain FSRM data for files that are located in volumes, which are not exported by CIFS or NFS. Host agents can also gather FSRM data about other file system paths that are not on a storage system: for example, local disk or third-party storage systems.

What capacity reports are

Capacity reports provide you with information about the file space statistics.

You can determine the following using the capacity reports:

- Total capacity of the storage system
- Amount of used space in the storage system
- Amount or percentage of free space available in the storage system

For example, you can determine the capacity of a volume by viewing the corresponding volume capacity report (**Reports > All > Volumes > Volume Capacity**).

Difference between capacity reports and file system statistics

The file space statistics that are reported by a path walk differ from the "volume space used" statistics reported by Operations Manager capacity reports.

This is due to the difficulty of determining how much space a file actually consumes on a volume. For example, most files consume slightly more space than the length of the file, depending on the block size. In addition, hard and soft links can cause a file to appear in more than one place and be counted twice. Therefore, do not use the results of a path walk to determine the amount of space used in a volume. Instead, use the DataFabric Manager server capacity reports.

Prerequisites for FSRM

The prerequisites for FSRM include File SRM license, connection to TCP/IP network, DataFabric Manager Host Agent software, and visible directory paths.

Prerequisite	Description
File SRM license	<p>You must have a valid File SRM license installed on your DataFabric Manager server. If you do not have an File SRM license, contact your sales representative.</p> <p>After you install the File SRM license, the Quotas tab in the Operations Manager is renamed "File SRM" and all the FSRM features become available.</p>
Connection to TCP/IP network	All FSRM hosts must be connected to a TCP/IP network that is either known to or discoverable by the DataFabric Manager server. The hosts must be connected to the network through an Ethernet port and must have a valid IP address.
DataFabric Manager Host Agent software	Each workstation from which FSRM data is collected must have DataFabric Manager Host Agent 2.0 (or later) installed. The 2.5 and later versions are recommended. For more information about the DataFabric Manager Host Agent, see the <i>DataFabric Manager Host Agent Installation and Administration Guide</i> .

Prerequisite	Description
Visible directory paths	All directory paths to be monitored must be visible to the host agent. For example, to enable FSRM monitoring, the host agent must mount a system share (volume or qtree), using NFS or CIFS, or the host agent must use a LUN on the system.

Related concepts

[How FSRM monitoring works](#) on page 153

Related information

[DataFabric Manager Host Agent Installation and Administration Guide - www.ibm.com/storage/support/nseries](http://www.ibm.com/storage/support/nseries)

Setting up FSRM

To set up and configure FSRM, you should perform a set of tasks such as identifying FSRM host agents, adding host agents, adding paths, and setting up path-walk schedules.

Steps

1. Identify FSRM host agents.

Note: To use host agents for FSRM, you must set your "login" to admin.

2. Add new host agents manually if they have not been discovered.
3. Set up host agent administration access on the hosts to be monitored. You can verify the host administration access by checking the **SRM Summary** page.
4. Add paths.
5. Set up path-walk schedules.

Related concepts

[How FSRM monitoring works](#) on page 153

Related tasks

[Configuring host agent administration settings](#) on page 159

[Enabling administration access for one or more host agents](#) on page 160

[Enabling administration access globally for all host agents](#) on page 160

[Adding SRM paths](#) on page 162

[Adding a path-walk schedule](#) on page 168

Related references

[Host Agent management tasks](#) on page 157

DataFabric Manager Host Agent software overview

DataFabric Manager Host Agent is a software that runs on Windows, Solaris, or Linux systems. It enables the DataFabric Manager server to collect SAN host bus adapter (HBA) information and remote file system metadata.

You can perform the following tasks by deploying DataFabric Manager Host Agent software on one or more hosts and enabling the File SRM license:

- Collect operating system version and host name information.
- Collect storage usage data at the file and directory levels.
- Identify and categorize various file-related information: for example, largest files, oldest files, files by owner, and files by type.
- Collect SAN host and SRM host data.

A workstation or server running the DataFabric Manager Host Agent software is called a host.

DataFabric Manager Host Agent communication

The DataFabric Manager server communicates with the DataFabric Manager Host Agent using HTTP or HTTPS.

You can specify the protocol to be used for communication in Operations Manager and when you install the DataFabric Manager Host Agent on your SAN host or SRM host agent. By default, both the DataFabric Manager Host Agent software and the DataFabric Manager server are configured to communicate with each other using HTTP.

DataFabric Manager Host Agent software passwords for monitoring tasks

Host agents have two user name and password pairs: one pair for monitoring and one pair for administration tasks. The default host agent user name and password allows monitoring only. You cannot perform FSRM functions.

Following are the default values for user name and password of the Host Agent software for monitoring tasks:

- User name=guest
- Password=public

Sessions, initiated by the DataFabric Manager server by using this user name and password are limited to basic monitoring operations. If you later decide to change the guest password on the host agent, you must also set the same user name and password in Operations Manager. You can specify the user name and password by using the `Host Agent Monitoring Password` option on the Options page.

DataFabric Manager Host Agent software passwords for administration tasks

Host agents have two user name and password pairs: one pair for monitoring and one pair for administration tasks. The administration user name and password allows read/write permission and is required for FSRM functions.

Following are the values for user name and password of the DataFabric Manager Host Agent for administration tasks:

- User name=admin
- Password=user-specified

You specify the password on the host agent's configuration UI ([http://name-of-agent: 4092/](http://name-of-agent:4092/)). This user name and password allows full access to the host agent. After setting the administration user name and password in the host agent, you must also set the same user name and password in Operations Manager on the Options page (**Setup menu > Options > Host Agent link**).

Note: This process of password change is applicable globally. To change passwords for one or more host agents, see "Enabling administration access for one or more host agents."

Related tasks

[Enabling administration access for one or more host agents](#) on page 160

Managing host agents

The DataFabric Manager server can discover host agents automatically; however, it does not use SNMP to poll for new host agents. Instead, a special-purpose agent called DataFabric Manager Host Agent is required for discovering, monitoring, and managing SAN and SRM hosts. DataFabric Manager Host Agent must be installed on each host agent that you want to monitor and manage with the DataFabric Manager server.

Host Agent management tasks

These are common host agent management tasks and the location of the Operations Manager user-interface page that enables you to complete them.

If you want to...	Go here...
Add a new host agent	Add Host Agents or Edit Host Agents page (Group Status tab > File SRM > Report drop-down list: SRM Summary > Add/Edit Hosts)

If you want to...	Go here...
Edit host agent settings (including passwords)	Edit Host Agent Settings page (Group Status tab > File SRM > Report drop-down list: SRM Summary > host agent name > Edit Settings in Tools list) OR Edit Host Agents page (Group Status tab > File SRM > Report drop-down list: SRM Summary > Add/Edit Hosts > appliance host name edit)
Edit host agent properties (for example, monitoring and management of API passwords, the HTTP port, and, HTTPS settings)	Edit Agent Settings page (Group Status tab > File SRM > Report drop-down list: SRM Summary > host agent name > Manage Host Agent in Host Tools list)
Configure host agent administration access	For one or more host agents: Edit Agent Logins page (Group Status tab > File SRM > Report drop-down list: SRM Summary > Edit Agent Logins) OR Create a global default using the Host Agent Options page (Setup menu > Options > Host Agent)
List of available host agents	If you have enabled a File SRM license on the workstation, the DataFabric Manager server automatically discovers all hosts that it can communicate with. Communication between the host agent and the DataFabric Manager server takes place over HTTP or HTTPS (port 4092 or port 4093, respectively). Host Agents, SRM view (Group Status tab > File SRM > Report drop-down list: SRM Summary > Host Agents, SRM)
Edit host agents	Add Host Agents or Edit Host Agents page (Group Status tab > File SRM > Report drop-down list: SRM Summary > Add/Edit Hosts link)
Review host agent global settings	Host Agent Options page (Setup menu > Options > Host Agent)
Obtain SRM host agent information	Host Agent Details page (Group Status tab > File SRM > Report drop-down list: SRM Summary > host agent name link)
Change the SRM host agent monitoring interval	Monitoring Options page (Setup menu > Options link > Monitoring Options link > Agent monitoring interval option)

If you want to...	Go here...
Modify DataFabric Manager Host Agent software settings	Edit Agent Settings page (Group Status tab > File SRM > Report drop-down list: SRM Summary > host agent name link > Manage Host Agent in Host Tools list)
Disable host agent discovery	Host Agent Discovery option on the Options page (Setup menu > Options link > Discovery > Host Agent Discovery field)
Delete a host agent	Add Host Agents or Edit Host Agents page (Group Status tab > File SRM > Report drop-down list: SRM Summary > Add/Edit Hosts link)

Configuring host agent administration settings

You can configure administration access such as monitoring only and management access, to the host agents.

Before you begin

You must enable administration access to your host agents before you can use the FSRM feature to gather statistical data.

About this task

Global options apply to all affected devices that do not have individual settings specified for them. For example, the Host Agent Login option applies to all host agents. The host agent access and communication options are globally set for all storage systems using the values specified in the Host Agent Options section on the Options page. Default values are initially supplied for these options. However, you should review and change the default values as necessary.

To enable administration access, the passwords set in Operations Manager must match those set for the DataFabric Manager Host Agent software.

Steps

1. Specify the DataFabric Manager server options:

Access type	DataFabric Manager server options
Monitoring only	Host Agent Login= <i>guest</i> Host Agent Monitoring Password
Management (required for FSRM)	Management Host Agent Login= <i>admin</i> Host Agent Management Password= <i>your-administration-password</i>

2. Specify the DataFabric Manager Host Agent options:

Access type	DataFabric Manager Host Agent options
Monitoring only	Monitoring API Password
Management (required for FSRM)	Monitoring API Password

Related concepts

[DataFabric Manager Host Agent software passwords for administration tasks](#) on page 157

Enabling administration access for one or more host agents

You can enable administration access for one or more host agents from the SRM Summary page.

Steps

1. From the **SRM Summary** page, click **Edit Agent Logins** in the Host Agents Total section.
2. Select the host agents for which you want to enable administration access.
3. Modify the fields, as needed, and then click **Update**.

Enabling administration access globally for all host agents

You can enable administration access globally for all host agents, from any Summary page.

Steps

1. From any **Summary** page, select **Options** from the Setup drop-down menu.
2. Select Host Agent from the **Edit Options list** (in the left pane).
3. Enter (or modify) the required information and then click **Update**.

This option changes all host agent login names and passwords, unless the host agent has a different login name or password already specified for it. For example, if an administrator has specified a password other than Global Default in the Password field of the Edit Host Agent Settings page, changing the global password option does not change the storage system password.

What FSRM paths are

File SRM paths define the location in the file system that is to be indexed for data.

Following are the properties of SRM paths:

- They must be defined for a specific host.
- They can be walked by managed host agents.
- They can be grouped like any other storage object.

- They can be mapped (linked) to volumes, qtrees, and LUNs.

Note: The FSRM path-walk feature can cause performance degradation. However, you can schedule your path walks to occur during low-use or non-business hours.

Adding CIFS credentials

To provide path names for CIFS, you must create a CIFS account.

Steps

1. Click **Setup > Options > Host Agent**.
2. In the **Host Agent Options** page, specify the CIFS account name in the **Host Agent CIFS Account** field.
3. In the **Host agent CIFS Password** field, type the password for the CIFS account.
4. Click **Update**.

Path management tasks

There are common path management tasks and the location of the Operations Manager user-interface page that enables you to complete them.

If you want to...	Go here..
Add paths	<p>To use the automapping feature: Create new SRM Path for this object link (Volume Details, Qtree Details, or LUN Details page > Create an SRM path)</p> <p>Note: The volume must be mounted on the managed Host Agent</p> <p>To manually add an SRM path: Add SRM Paths or Edit SRM Paths page (Group Status tab > File SRM > Report drop-down list: SRM Summary > Add/Edit Paths link)</p>
Create path-walk schedules	<p>Edit SRM Path Walk Schedules page (Group Status tab > File SRM > Report drop-down list: SRM Summary > Add/Edit Schedules link)</p>
Specify path-walk times	<p>SRM Path Walk Schedules Times page (Group Status tab > File SRM > Report drop-down list: SRM Summary > Add/Edit Schedules link > schedule name > Add Walk Times link)</p>
Manually start or stop an SRM path walk	<p>SRM Path Details page (SRM path name > Start)</p> <p>If an SRM path walk is in progress, the Start button changes to a Stop button.</p>

If you want to...	Go here..
Review SRM path details	SRM Path Walk Schedule Details page (schedule name)
Edit SRM paths	Add SRM Paths or Edit SRM Paths page (Group Status tab > File SRM > Report drop-down list: SRM Summary > Add/Edit Paths link) OR Edit SRM Path Settings page (SRM path name > Edit Settings in Host Tools list)
Review SRM path-walk schedule details	SRM Path Walk Schedule Details page (schedule name)
Edit SRM path-walk schedules	SRM Path Walk Schedule Times page (Group Status tab > File SRM > Report drop-down list: SRM Summary > Add/Edit Schedules link > schedule name > Add Walk Times link)
Delete SRM paths	Add SRM Paths or Edit SRM Paths page (Group Status tab > File SRM > Report drop-down list: SRM Summary > Add/Edit Paths link)
Delete SRM path-walk schedules	Edit SRM Path Walk Schedules page (Group Status tab > File SRM > Report drop-down list: SRM Summary > Add/Edit Schedules link > schedule name)
Delete SRM path-walk schedule times	SRM Path Walk Schedule Details page (schedule name)

Adding SRM paths

You can add SRM paths from the SRM Summary page.

Steps

1. From the **SRM Summary** page, click the **Add/Edit Paths** in the SRM Paths Total section.
2. From the SRM Host drop-down list in the Add a New SRM Path section, select the name of the host agent that you want to monitor.
3. Type a path name, select a schedule, and then click **Add SRM Path**.

Valid path entries

```
host:/u/earth/work
host:/usr/local/bin
host:/engineering/toi
```

```
host:C:\Program Files
```

```
For CIFS, you must specify the path as a UNC path, as follows: host:\\storage system
\share\dir
```

Path names for CIFS

For CIFS systems, always use Universal Naming Convention (UNC) path names.

In Windows operating systems, the UNC format is as follows:

```
\\servername\sharename\path\filename
```

The SRM feature does not convert mapped drives to UNC path names. For example, suppose that drive H: on the system host5 is mapped to the following path name:

```
\\abc\users\jones
```

The path entry `host5:H:\` fails because the FSRM feature cannot determine what drive H: is mapped to. The following path entry is correct:

```
host5:\\abc\users\jones
```

Conventions for specifying paths from the CLI

Unique conventions have to be followed for specifying paths from CLI in Windows and UNIX.

Windows requires that you use double quotation marks to enclose all strings that contain spaces. Following are the examples paths:

```
C:\dfm srm path add "inchon:C:\Program Files"
```

```
C:\dfm srm path add "oscar:/usr/home"
```

UNIX requires that you double all backslashes, unless the argument is enclosed in double quotation marks. This convention is also true for spaces in file names. For example:

```
$ dfm srm path add inchon:C:\\Program\ Files
```

```
$ dfm srm path add "inchon:C:\Program Files"
```

```
$ dfm srm path add oscar:/usr/local
```

Viewing file-level details for a path

You can view file-level details about an SRM path from the SRM Summary page.

Step

1. From the **SRM Summary** page, click a path name in the SRM Paths Total section.

Viewing directory-level details for a path

You can view directory-level details about an SRM path from the SRM Summary page.

Steps

1. From the **SRM Summary** page, click a path name in the SRM Paths Total section.
2. Click the **Browse Directories** link in the SRM Path Tools list (at the lower left of Operations Manager).
3. To view an expanded view of directory information that includes a listing of files by type and by user, click the **Extended Details** link (at the upper right corner of the File SRM tab window).

Editing SRM paths

You can edit SRM paths from the SRM Summary page.

Steps

1. From the **SRM Summary** page, click **Add/Edit Paths** in the SRM Paths Total section.
2. Select the SRM path you want to modify.
3. Modify the fields, as needed, and then click **Update**.

Deleting SRM paths

You can delete SRM paths from the SRM Summary page.

Steps

1. From the **SRM Summary** page, click **Add/Edit Paths** in the SRM Paths Total section.
2. Select the SRM paths you want to delete and then click **Delete**.

Automatically mapping SRM path

When you automatically map the SRM path, the initial path mapping will be correct. Subsequent changes on the host (running the DataFabric Manager Host Agent) or storage system can cause the path mapping to become invalid.

Note: Sometimes, you might not be able to do the following operations:

- When you cannot access SRM data, you cannot automatically map storage objects with an SRM path.
- If the SRM path is not on a storage device monitored by the DataFabric Manager server, you cannot associate an SRM path with a storage object.

Requirements for automatically mapping SRM path

You can automatically create a new path for an object using the “Create new SRM Path for this object” link on the details page for the object.

You must ensure the following, to create a new path for an object:

- The host agent is set up and properly configured.
- The host agent passwords match those set in the DataFabric Manager server.
- The host agent has access to the volume, qtree, or LUN:
 - If the host agent is a Windows host, you must ensure that the CIFS passwords match.
 - If the object is a LUN on a Windows host, SnapDrive must be installed and the LUN must be managed by SnapDrive.
 - If the host agent is a UNIX host, then the volume or qtree must be NFS mounted.
 - If the object is a LUN on a UNIX host, the LUN must be formatted and mounted directly into the file system (volume managers are not supported).
- The Host Agent Login and Management Password are set correctly.

You can also manually map SRM paths to volumes, qtrees, and LUNs.

What path walks are

A path walk is the process of recursively examining a directory path for file-level statistics. Path walks are scheduled using the DataFabric Manager server and executed by DataFabric Manager Host Agent software. The DataFabric Manager Host Agent scans all subdirectories of the specified directory path and gathers per-file and per-directory data.

SRM path-walk recommendations

SRM path walks can consume considerable resources on the SRM host agent and on the DataFabric Manager server. Therefore, schedule your SRM path walks to occur during off-peak hours. Also, do not schedule multiple, simultaneous SRM path walks on the same SRM host agent.

What File SRM reports are

Operations Manager provides three levels of file system statistics.

Following are the three levels of file system statistics provided by Operations Manager:

- Consolidated data gathered from all paths
- SRM path-specific data
 - This is a summary of the data for all directories in the specified path.
- Directory-level data
 - This contains the data for the specified directory only.

Note: If the host agent is installed on UNIX, File SRM feature tracks only the files having a file name extension that exactly matches the file type specification in Operations Manager. For example, files that end in .JPG will not match the .jpg file type if the host agent is on UNIX. Even

though they would match if the agent were on Windows. Running the host agent on Windows avoids this problem.

Viewing file system statistics

You can view the File SRM report for a group by clicking on the File SRM tab.

Steps

1. Click the group for which you want a File SRM report.
2. Click **File SRM** and select a report from the Report drop-down list. The following reports are available:
 - SRM Paths, All
 - SRM Directories, Largest
 - SRM Files, Largest
 - SRM Files, Least Recently Accessed
 - SRM Files, Least Recently Modified
 - SRM Files, Recently Modified
 - SRM File Types
 - SRM File Owners

Each FSRM report page displays statistics for the users who have storage space that are allocated on the objects (storage systems, aggregates, volumes, or qtrees) in your selected group.

You can list reports by using the CLI `dfm report list` command (without arguments) to display all available reports.

Access restriction to file system data

To restrict access to private or sensitive file system information, remove the GlobalSRM role from the access privileges in the Administrators page (Setup menu > Administrative Users link).

Related concepts

[How roles relate to administrators](#) on page 49

Identification of oldest files in a storage network

You can find the oldest files residing in its storage network using File SRM (FSRM) and archive the files to a NearStore system.

There is a list of high-level tasks to be performed to use FSRM, for identifying oldest files in a storage network.

- Check FSRM prerequisites
- Verify administrative access
- Verify host agent communication

- Create a new group
Group the host agents in a logical way. For example, group the engineering host agents together if you want to search for them separately.
- Add a FSRM path
- Add a schedule
- Group the FSRM paths
- View a report listing the oldest files

FSRM prerequisites

Before using the FSRM feature for the first time, you must verify that all prerequisites are met by referring to difference between, capacity reports and file system statistics.

Verifying administrative access for using FSRM

Before using the FSRM feature for the first time, you must verify that all prerequisites are met by referring to difference between, capacity reports and file system statistics. You must verify administrative access in Operations Manager for using FSRM.

Steps

1. Click **Setup > Options**.
2. Click **Host Agent** in the Edit Options section.
3. Change the Host Agent Login to `Admin`.
4. Verify that the Host Agent Management Password is set.
5. Click **Update** to apply the changes.
6. Click **Home** to return to the Control Center.
7. Click **File SRM** to return to the SRM Summary page.

Verifying host agent communication

To verify that the DataFabric Manager server can communicate with the host agents, the ABC administrator must complete a certain list of tasks.

Steps

1. Click **File SRM** then select **SRM Summary** from the Report drop-down list.
2. Check the list of host agents to view the status.
If the status is Unknown, the host agent login settings might not be properly configured.
3. If the status of one or more of the storage systems is Unknown, click **Edit Agent Logins**.
4. Select the host agents for engineering that the administrator wants to communicate with.

5. Edit the login or password information.
6. Click **Update**.
7. Click **File SRM** to return to the **SRM Summary** page.

Creating a new group of hosts

The ABC Company wants to find the oldest files in its engineering department first. To find these files, the ABC Company administrator groups together the host agents in the engineering domain. To create a new group of hosts, the administrator must complete a certain list of tasks.

Steps

1. From the **SRM Summary** page, select **Host Agents, SRM** from the Report drop-down list.
2. Select the host agent in the engineering domain.
3. From the buttons at the bottom of the page, click **Add To New Group**.
4. When prompted, enter the name `ENG` and Click **Add**.
The DataFabric Manager server refreshes.
5. Select **SRM Summary** from the Report drop-down list.

Adding an FSRM path

The administrator can add an FSRM path from the SRM Summary page.

Steps

1. From the **SRM Summary** page, **Add/Edit Paths**.
2. Select a host from the SRM Host drop-down list.
3. Enter the path to be searched in the Path field.
4. Click **Add A Schedule**.

Adding a path-walk schedule

To create the path-walk schedule, the administrator must complete a certain set of tasks.

Steps

1. Click the **Add/Edit Schedules** link in the **SRM Summary** page (File SRM tab).
2. In the Add a New Schedule section, enter a meaningful name for the schedule.
3. In the **Schedule Template** list, select a schedule or select None.
4. Click **Add**.

5. In the **SRM Path Walk Schedule Times**, select the days and times to start the SRM path walks.
6. Click **Update**.
7. Click **Home** to browse back to the main window.
8. Click **File SRM**.
9. **SRM Summary** page, click **Add/Edit Paths**.
10. In the Add a New SRM Path section, select a host agent to associate the new schedule with.
11. In the Schedule field, select the schedule name you created.
12. Click **Add SRM Path**.

Grouping the FSRM paths

To view consolidated data, the administrator must group the FSRM path. To do so, the administrator completes a set of tasks.

Steps

1. Click **File SRM**.
2. Select **SRM Paths, All** from the Report drop-down list.
3. Select the SRM path that the administrator wants to group.
4. From the buttons at the bottom of the page, click **New Group**.
5. When prompted, enter a name for the group.

The DataFabric Manager server adds the new group and refreshes.

Viewing a report that lists the oldest files

To view a report listing the oldest files in the SRM path, the ABC Company administrator must complete a set of tasks.

Steps

1. Click **Home**.
2. Click **File SRM**.
3. Select the engineering group in the Groups section at the left side of the tab window.
4. Select the report SRM Files, Least Recently Accessed from the Report drop-down list.
5. Review the data.

User quotas

You can use user quotas to limit the amount of disk space or the number of files that a user can use. Quotas provide a way to restrict or track the disk space, and number of files used by a user, group, or qtree. Quotas are applied to a specific volume or qtree.

Why you use quotas

You can use quotas to limit resource usage, to provide notification when resource usage reaches specific levels, or simply to track resource usage.

You specify a quota for the following reasons:

- To limit the amount of disk space or the number of files that can be used by a user or group, or that can be contained by a qtree
- To track the amount of disk space or the number of files used by a user, group, or qtree, without imposing a limit
- To warn users when their disk usage or file usage is high

Overview of the quota process

Quotas can be soft or hard. Soft quotas cause Data ONTAP to send a notification when specified thresholds are exceeded, and hard quotas prevent a write operation from succeeding when specified thresholds are exceeded.

When Data ONTAP receives a request to write to a volume, it checks to see whether quotas are activated for that volume. If so, Data ONTAP determines whether any quota for that volume (and, if the write is to a qtree, for that qtree) would be exceeded by performing the write operation. If any hard quota would be exceeded, the write operation fails, and a quota notification is sent. If any soft quota would be exceeded, the write operation succeeds, and a quota notification is sent.

User quota management using Operations Manager

You can view user quota summary reports, chargeback reports, user details, quota events, and so on.

You can perform the following user quota management tasks by using Operations Manager:

- View summary reports (across all storage systems) and per-quota reports with data, about files and disk space that is used by users, hard and soft quota limits, and the projected time when users exceed their quota limit
- View graphs of total growth and per-quota growth of each user
- View details about a user

- Obtain chargeback reports for users
- Edit user quotas when you edit user quotas through Operations Manager, the `/etc/quotas` file is updated on the storage system on which the quota is located.
- Configure and edit user quota thresholds for individual users, volumes, and qtrees. When you configure a user quota threshold for a volume or qtree, the settings apply to all user quotas on that volume or qtree.
- Notify users when they exceed the user quota thresholds configured in the DataFabric Manager server
- Configure the monitoring interval for user quota monitoring
- View and respond to user quota events
- Configure alarms to notify administrators of the user quota events

Related concepts

[Where to find user quota reports in Operations Manager](#) on page 173

[Monitor interval for user quotas in Operations Manager](#) on page 173

[What user quota thresholds are](#) on page 175

Prerequisites for managing user quotas using Operations Manager

To monitor and manage user quotas by using Operations Manager, ensure that your storage system meets certain prerequisites.

- Following are the prerequisites to monitor and edit user quotas assigned to vFiler units:
 - You should enable RSH or SSH access to the storage system and configure login and password credentials that are used to authenticate the DataFabric Manager server.
 - You must use the DataFabric Manager server to configure the root login name and root password of a storage system on which you want to monitor and manage user quotas.
 - You must configure and enable quotas for each volume for which you want to view the user quotas.
 - You must log in to Operations Manager as an administrator with the quota privilege to view user quota reports and events so that you can configure user quotas for volumes and qtrees.
- Additional requirements for editing quotas:
 - Directives, such as `QUOTA_TARGET_DOMAIN` and `QUOTA_PERFORM_USER_MAPPING` must not be present in the `/etc/quotas` file on the storage system.
 - The `/etc/quotas` file on the storage system must not contain any errors.

Where to find user quota reports in Operations Manager

You can view user quota reports in Operations Manager at **Control Center > Home > File SRM (or Quotas) > Report**.

If you have not installed a File SRM license, you cannot view the File SRM tab. However, you can access the reports from the Quotas tab. After you install the File SRM license, the Quotas tab is renamed as “File SRM.”

Monitor interval for user quotas in Operations Manager

You can use Operations Manager to view the monitoring interval at which the DataFabric Manager server is monitoring a user quota on a storage system.

The `User Quota Monitoring Interval` option on the Options page (**Setup > Options > Monitoring**) determines how frequently the DataFabric Manager server collects the user quota information from the monitored storage systems. By default, the user quota information is collected once every day; however, you can change this monitoring interval.

Note: The process of collecting the user quota information from storage systems is resource intensive. When you decrease the User Quota Monitoring Interval option to a low value, the DataFabric Manager server collects the information more frequently. However, decreasing the User Quota Monitoring Interval might negatively affect the performance of the storage systems and the DataFabric Manager server.

Modification of user quotas in Operations Manager

You can edit disk space threshold, disk space hard limit, disk space soft limit, and so on for a user quota in Operations Manager.

When you edit the options for a user quota, the `/etc/quotas` file on the storage system where the quota exists is appropriately updated.

- Disk space threshold
- Disk space hard limit
- Disk space soft limit
- Files hard limit
- Files soft limit

For more information about these fields, see the *Operations Manager Help*.

Prerequisites to edit user quotas in Operations Manager

If you want to edit user quota in Operations Manager, ensure that your storage system meets the prerequisites.

- You must configure the root login name and root password in the DataFabric Manager server for the storage system on which you want to monitor and manage user quotas.
- You must configure and enable quotas for each volume for which you want to view the user quotas.
- Operations Manager conducts vFiler quota editing by using the jobs. If a vFiler quota editing job fails, verify the `quota` file on the hosting storage system.

In addition, to protect the `quota` file against damage or loss, before starting a job, the DataFabric Manager server creates a backup file named `DFM (timestamp).bak`. If the job fails, you can recover data by renaming the `backup quota` file.

Editing user quotas using Operations Manager

You can edit user quotas using the Edit Quota Settings page in Operations Manager.

Before you begin

Ensure that the storage system meets the prerequisites before you edit user quotas in Operations Manager.

Steps

1. Click **Control Center > Home > Group Status > File SRM (or Quotas) > Report > User Quotas, All**.
2. Click any quota related fields for the required quota.

Configuring user settings using Operations Manager

You can configure user settings such as e-mail address of users and quota alerts, and set user quota threshold using Operations Manager.

Steps

1. Click **Control Center > Home > Group Status > File SRM (or Quotas) > Report > User Quotas, All**
2. Click the **Edit Settings** link in the lower left corner.
3. You can edit **E-mail Address** of the user, **Send Quota Alerts Now**, **User Quota Full Threshold (%)**, **User Quota Nearly Full Threshold (%)**, **Owner E-mail**, **Owner Name**, and **Resource Tag**.

You can leave the e-mail address field blank if you want the DataFabric Manager server to use the default e-mail address of the user.

4. Click **Update**.

What user quota thresholds are

User quota thresholds are the values that the DataFabric Manager server uses to evaluate whether the space consumption by a user is nearing, or has reached the limit that is set by the user's quota.

If these thresholds are crossed, the DataFabric Manager server generates user quota events.

By default, the DataFabric Manager server sends user alerts in the form of e-mail messages to the users who cause user quota events. Additionally, you can configure alarms that notify the specified recipients (DataFabric Manager server administrators, a pager address, or an SNMP trap host) of user quota events.

The DataFabric Manager server can also send a user alert when users exceed their soft quota limit; however, no thresholds are defined in the DataFabric Manager server for the soft quotas. The DataFabric Manager server uses the soft quota limits set in the `/etc/quotas` file of a storage system to determine whether a user has crossed the soft quota.

What the DataFabric Manager server user thresholds are

The DataFabric Manager server user quota thresholds are a percentage of the Data ONTAP hard limits (files and disk space) configured in the `/etc/quotas` file of a storage system.

The user quota threshold makes the user stay within the hard limit for the user quota. Therefore, the user quota thresholds are crossed even before users exceed their hard limits for user quotas.

User quota thresholds

You can set a user quota threshold to all the user quotas present in a volume or a qtree.

When you configure a user quota threshold for a volume or qtree, the settings apply to all user quotas on that volume or qtree.

The DataFabric Manager server uses the user quota thresholds to monitor the hard and soft quota limits configured in the `/etc/quotas` file of each storage system.

Ways to configure user quota thresholds in Operations Manager

You can configure user quota thresholds by applying the thresholds to all quotas of a specific user or on a specific file system or on a group of file systems using Operations Manager.

- Apply user quota thresholds to all quotas of a specific user
- Apply user quota thresholds to all quotas on a specific file system (volume or qtree) or a group of file systems

You can apply thresholds using the **Edit Quota Settings** links on the lower left pane of the Details page for a specific volume or qtree. You can access the Volume Details page by clicking

on a volume name at **Control Center > Home > Member Details > File Systems > Report > Volumes, All**. Similarly, for the Qtree Details page, clicking on the qtree name at **Control Center > Home > Member Details > File Systems > Report > Qtrees, All**

To apply settings to a group of file systems, select the group name from the Apply Settings To list on the quota settings page.

- Apply user quota thresholds to all quotas on all users on all file systems: that is, all user quotas in the DataFabric Manager server database

You can apply thresholds at **Setup > Options > Edit Options: Default Thresholds**.

Precedence of user quota thresholds in the DataFabric Manager server

The DataFabric Manager server prioritizes user quota threshold based on a specific user, a specific volume or a qtree, and, all users in the DataFabric Manager server.

The following list specifies the order in which user quota thresholds are applied:

1. User quota thresholds specified for a specific user
2. File systems (volumes and qtrees) user quota thresholds specified for a specific volume or qtree
3. Global user quota thresholds specified for all users in the DataFabric Manager server database

Management of LUNs, Windows and UNIX hosts, and FCP targets

Operations Manager is used to monitor and manage LUNs, Windows and UNIX hosts, and, FCP targets in your SANs. SANs on the DataFabric Manager server are storage networks that have been installed in compliance with the specified SAN setup guidelines. For more information about setting up a SAN, see the *Data ONTAP Block Access Management Guide for iSCSI and FC*.

Note: DataFabric Manager server customers should check with their sales representative regarding other SAN management solutions.

Related information

Data ONTAP Block Access Management Guide for iSCSI and FC - www.ibm.com/storage/support/nseries

Management of SAN components

To monitor and manage LUNs, FCP targets, and SAN hosts, the DataFabric Manager server must first discover them.

The DataFabric Manager server uses SNMP to discover storage systems, but SAN hosts must already have the DataFabric Manager Host Agent software installed and configured on them before the DataFabric Manager server can discover them.

After SAN components have been discovered, the DataFabric Manager server starts collecting pertinent data—for example, which LUNs exist on which storage systems. Data is collected periodically and reported through various Operations Manager reports. (The frequency of data collection depends on the values that are assigned to the DataFabric Manager server monitoring intervals.)

The DataFabric Manager server monitors LUNs, FCP targets, and SAN hosts for a number of predefined conditions and thresholds. For example, when the state of an HBA port changes to online or offline or when the traffic on an HBA port exceeds a specified threshold. If a predefined condition is met or a threshold is exceeded, the DataFabric Manager server generates and logs an event in its database. These events can be viewed through the details page of the affected object. Additionally, you can configure the DataFabric Manager server to send notification about such events (also known as alarms) to an e-mail address. You can also configure the DataFabric Manager server to send notifications to a pager, an SNMP trap host, or a script you write.

In addition to monitoring LUNs, FCP targets, and SAN hosts, you can use the DataFabric Manager server to manage these components. For example, you can create, delete, or expand a LUN.

SAN and DataFabric Manager Host Agent software

The DataFabric Manager server can automatically discover SAN hosts; however, it does not use SNMP to poll for new hosts.

DataFabric Manager Host Agent software discovers, monitors, and manages SANs on SAN hosts. You must install the DataFabric Manager Host Agent software on each SAN host that you want to monitor and manage with the DataFabric Manager server.

Note: To modify the global host agent monitoring interval for SAN hosts, you must change the SAN Host Monitoring Interval (**Setup > Options > Monitoring**).

List of tasks performed using DataFabric Manager Host Agent software

Once DataFabric Manager Host Agent software is installed on a client host along with the DataFabric Manager server, you can perform various management tasks.

After the DataFabric Manager Host Agent software is installed on a client host and you have installed the DataFabric Manager server with the Operations Manager license, you can perform a variety of tasks:

- Monitor basic system information for SAN hosts and related devices.
- Perform management functions such as creating, modifying, or expanding a LUN.
- View detailed LUN information.

Note: DataFabric Manager Host Agent is also used for File Storage Resource Management functions, with a File SRM license. For more information about the DataFabric Manager Host Agent software, see the *DataFabric Manager Host Agent Installation and Administration Guide*.

Related information

[DataFabric Manager Host Agent Installation and Administration Guide - www.ibm.com/storage/support/nseries](http://www.ibm.com/storage/support/nseries)

List of tasks performed to monitor targets and initiators

You can use Operations Manager to perform management tasks such as view reports; monitor, manage, and group LUNs; and respond to LUN and SAN host events.

Following is a list of tasks you can perform to monitor targets and initiators:

- View reports that provide information about LUNs, FCP targets, and SAN hosts.
- View details about a specific LUN, FCP target on a storage system, or SAN host.

- Group LUNs, storage systems in a SAN, or SAN hosts for efficient monitoring and management.
- Change the monitoring intervals for LUNs, and SAN hosts.
- View and respond to LUN and SAN host events.
- Configure the DataFabric Manager server to generate alarms to notify recipients of LUN and SAN host events.

Related concepts

Reports for monitoring LUNs, FCP targets, and SAN hosts on page 180

Information available on the LUN Details page on page 181

Tasks performed from the LUN Details page on page 182

Information about the FCP Target Details page on page 182

Information about the Host Agent Details page on page 183

List of tasks performed from the Host Agent Details page on page 183

Prerequisites to manage targets and initiators

Operations Manager does not report any data for your targets and initiators if you do not have the SAN set up according to the specified hardware and software requirements.

For more information about the supported hardware platforms, see the Interoperability Matrix. For more information about specific software requirements, see the *DataFabric Manager Installation and Upgrade Guide*.

Related information

DataFabric Manager Installation and Upgrade Guide - www.ibm.com/storage/support/nseries

Interoperability Matrix - www.ibm.com/servers/storage/nas/interophome.html

Prerequisites to manage SAN hosts

You must ensure a proper network connection and software installed on SAN hosts before you manage SAN hosts with the DataFabric Manager server.

- All SAN hosts to be managed by the DataFabric Manager server must be connected to a TCP/IP network either known to or discoverable by the DataFabric Manager server.
The SAN hosts must be connected to the network through an Ethernet port and must each have a valid IP address.
- Each SAN host must have the DataFabric Manager Host Agent software installed on it.
The DataFabric Manager Host Agent software is required for discovering, monitoring, and managing SAN hosts. For more information about the DataFabric Manager Host Agent software, see the *DataFabric Manager Host Agent Installation and Administration Guide*.
- For LUN management using the DataFabric Manager server, Windows SAN hosts must have the proper version of SnapDrive software installed.
To find out which SnapDrive version you received, see the DataFabric Manager server software CD.

Note: LUN management on UNIX SAN hosts using the DataFabric Manager server is not available. LUNs inherit access control settings from the storage system, volume, and qtree they are contained in. Therefore, to perform LUN operations on storage systems, you must have appropriate privileges set up on those storage systems.

Related information

DataFabric Manager Host Agent Installation and Administration Guide - www.ibm.com/storage/support/nseries

Reports for monitoring LUNs, FCP targets, and SAN hosts

Reports about LUNs, SAN hosts, and FCP targets that the DataFabric Manager server monitors, are available on the LUNs page of Member Details tab.

You can view reports by selecting from the Report drop-down list. If you want to view a report about a specific group, click the group name in the left pane of Operations Manager. You can view the following reports from the LUNs page:

- FCP Targets
- SAN Hosts, Comments
- SAN Hosts, All
- SAN Hosts, FCP
- SAN Hosts, iSCSI
- SAN Hosts, Deleted
- SAN Hosts Traffic, FCP
- SAN Host Cluster Groups
- SAN Host LUNs, All
- SAN Hosts LUNs, iSCSI
- SAN Hosts LUNs, FCP
- LUNs, All
- LUNs, Comments
- LUNs, Deleted
- LUNs, Unmapped
- LUN Statistics
- LUN Initiator Groups
- Initiator Groups

For more information about descriptions of each report field, see the *Operations Manager Help*.

Information available on the LUN Details page

The LUN Details page for a LUN consists of information such as status of the LUN, the LUN's storage system and so on.

You can access the Details page for a LUN by clicking the LUN path displayed in any of the reports. Following are the information about the LUN Details page:

- Status of the LUN
- Storage system on which the LUN exists
- Volume or qtree on which the LUN exists
- Initiator groups to which the LUN is mapped

You can access all LUN paths that are mapped to an initiator group by clicking the name of the initiator group.

Note: If a LUN is mapped to more than one initiator group, when you click an initiator group, the displayed page lists all the LUN paths that are mapped to the initiator group. Additionally, the report contains all other LUN mappings (LUN paths to initiator groups) that exist for those LUN paths.

- Size of the LUN
- Serial number of the LUN
- Description of the LUN
- Events associated with the LUN
- Groups to which the LUN belong
- Number of LUNs configured on the storage system on which the LUN exists and a link to a report displaying those LUNs
- Number of SAN hosts mapped to the LUN and a link to the report displaying those hosts
- Number of HBA ports that can access this LUN and a link to the report displaying those LUNs
- Time of the last sample collected and the configured polling interval for the LUN
- Graphs that display the following information:
 - LUN bytes per second—Displays the rate of bytes (bytes per second) read from and written to a LUN over time
 - LUN operations per second—Displays the rate of total protocol operations (operations per second) performed on the LUN over time

Related concepts

[Reports for monitoring LUNs, FCP targets, and SAN hosts](#) on page 180

Tasks performed from the LUN Details page

By using the **LUN Path Tools** links on the LUN Details page, you can perform various management tasks.

Following are the tasks you can perform on the LUN Details page:

Expand this LUN	Launches a wizard that helps you expand the LUN
Destroy this LUN	Launches a wizard that helps you destroy the LUN
Refresh Monitoring Samples	Obtains current monitoring samples from the storage system on which this LUN exists
Run a Command	Runs a Data ONTAP command on the storage system on which this LUN exists

Note: To manage a shared LUN on MSCS, perform the operation on the active controller. Otherwise, the operation fails.

You must have appropriate authentication to run commands on the storage system from the DataFabric Manager server.

Information about the FCP Target Details page

The FCP Target Details page contains information such as name of the storage system, operational status of the target, adapter speed, and so on.

You can access the FCP Target Details page for a target by clicking its port number in the FCP Targets report (**LUNs > View, FCP Targets**). The FCP Target Details page contains the following information:

- Name of the storage system on which target is installed
- Operational status of the target
- Adapter hardware and firmware versions
- Adapter speed
- FC topology of the target:
 - Fabric
 - Point-To-Point
 - Loop
 - Unknown
- Node name (WWNN) and port name (WWPN) of the target
- Number of other FCP targets on the storage system on which the target is installed (link to report)
- Time of the last sample collected and the configured polling interval for the FCP target

Information about the Host Agent Details page

The Host Agent Details page contains information such as status of SAN host, events that occurred on SAN host, devices related to SAN host, and so on.

You can access the Details page for a DataFabric Manager Host Agent by clicking its name in any of the SAN Host reports. The Details page for a Host Agent on a SAN host contains the following information:

- Status of the SAN host and the time since the host has been up
- The operating system and the DataFabric Manager Host Agent version, in addition to protocols and features running on the SAN host
- The MSCS configuration information about the SAN host, if any, such as the cluster name, cluster partner, and cluster groups to which the SAN host belongs
- The events that occurred on this SAN host
- The devices related to the SAN host, such as the storage systems accessible from the SAN host
- Graphs of information, such as the HBA port traffic per second or the HBA port frames for different time intervals

For more information about the SAN Host reports, see the *Operations Manager Help*.

List of tasks performed from the Host Agent Details page

The Host Tools list on the Host Agent Details page enables you to perform various tasks such as edit settings, create a LUN, refresh monitoring samples, and so on.

Edit Settings	Displays the Edit Host Agent Settings page, where you configure login, password, administration transport, port information for the SAN host, and the user name and password for CIFS access in Operations Manager. The login and password information is used to authenticate the DataFabric Manager server to the DataFabric Manager Host Agent software running on the SAN host. Specify a value on this page only if you want to change the global setting.
Create a LUN	Takes you to a LUN creation page that help you create a LUN.
Diagnose Connectivity	Automates connectivity troubleshooting.
Refresh Monitoring Samples	Obtains current monitoring samples from the SAN host.
Manage Host Agent	Allows you to edit settings for the Host Agent. You can edit details such as monitoring and management of API passwords, the HTTP and HTTPS ports. This enables remote upgrading and specifies a filewalk log path.

Related concepts

[Where to configure monitoring intervals for SAN components](#) on page 185

How storage systems, SAN hosts, and LUNs are grouped

You can group LUNs, storage systems, or SAN hosts for easier management and to apply access control.

When you create a group of storage systems or SAN hosts, the type of the created group is “Appliance resource group.” When you create a group of LUNs, the created group is “LUN resource group.”

Note: You cannot group HBA ports or FCP targets.

Related tasks

[Creating groups](#) on page 73

Granting access to storage systems, SAN hosts, and LUNs

You can allow an administrator access for managing all your SAN hosts and LUNs. The GlobalSAN role allows an administrator to create, expand, and destroy LUNs.

Before you begin

The GlobalSan role must be enabled for LUN management.

Step

1. To allow administrator access, go to the **Administrator** page and select **Setup menu > Administrative users**.

Option	Description
To create a new administrator	In the Administrators page, complete the Add a New Administrator option, and then select GlobalSAN from the Roles list.
To grant access to an existing administrator	In the Administrator page, from the List of administrators, click the Edit column of the administrator to be granted access, and then select GlobalSAN from the Roles list.

Introduction to deleting and undeleting SAN components

You can stop monitoring a SAN component (a LUN, a storage system, or a SAN host) with the DataFabric Manager server by deleting it from the Global group.

When you delete a SAN component from the Global group, the DataFabric Manager server stops collecting and reporting data about it. Data collection and reporting is not resumed until the component is again added by performing the undelete operation.

You cannot stop monitoring a specific FCP target or an HBA port, unless you first stop monitoring the storage system (for the FCP target) or the SAN host (for the HBA port) on which the target or the port exists.

Note: When you delete a SAN component from any group except Global, the component is deleted only from that group. The DataFabric Manager server does not stop collecting and reporting data about it. You must delete the SAN component from the Global group for the DataFabric Manager server to stop monitoring it altogether.

Deleting a SAN component

You can delete a SAN component from any of the reports related to that component.

Steps

1. Select the component you want to delete by clicking the check boxes in the left-most column of a report.
2. Click the **Delete Selected** button at the bottom of each report to delete the selected component.

How a deleted SAN component delete is restored

You can restore a deleted object, by selecting it and then clicking the Undelete button from the Deleted report.

All deleted objects are listed in their respective Deleted reports. For example, all deleted LUNs are listed in the LUNs Deleted report.

Where to configure monitoring intervals for SAN components

You can configure the global options on the Options page (Setup menu > Options) in Operations Manager.

To configure local options (for a specific object), you must be on the Edit Settings page of that specific object (**Details page > Tools list: Edit Settings**). For example, to access the LUNs Details

page, click **Member Details > LUNs > Report drop-down list: LUNS, All > LUN Path > LUN Path Tools: Edit Settings**.

Related concepts

Where to find information about a specific storage system on page 216

File system management

You can manage storage on storage system by using the data displayed by Operations Manager storage reports and the options you use to generate storage-related events and alarms.

Access to storage-related reports

You can view storage-related reports about storage objects that the DataFabric Manager server monitors.

The storage reports present information about a selected aspect of the storage object, such as chargeback, space availability, and status.

Note: The status specifies the current status of the selected storage object.

You can find storage-related reports on the tabs accessible from the Member Details tab: Physical Systems, Virtual Systems, File Systems, Aggregates, and LUNs tabs. Each tab has a Report drop-down list from which you can select the report you want to display.

For information about specific storage-related reports, see the *Operations Manager Help*.

Storage capacity thresholds in Operations Manager

Storage capacity threshold is the point at which the DataFabric Manager server generates events to report capacity problem.

Storage capacity thresholds determine at what point you want the DataFabric Manager server to generate events about capacity problems. You can configure alarms to send notification whenever a storage event occurs.

When the DataFabric Manager server is installed, the storage capacity thresholds for all aggregates, volumes, and qtrees are set to default values. You can change the settings as needed for an object, a group of objects, or the Global group.

Modification of storage capacity thresholds settings

You can change the storage capacity threshold settings globally, for a specific group, or for a specific aggregate, volume, or qtree. If you edit capacity threshold settings, the edited thresholds override the global thresholds.

Note: Changing storage capacity threshold settings for the Global group changes the default storage capacity settings for all groups and individual storage objects.

Changing storage capacity threshold settings for global group

Perform the following steps to change the storage capacity threshold settings for global group:

Steps

1. Select **Setup > Options**.
2. Click **Default Thresholds** in the left pane.
3. Edit the Default settings as needed.
4. Click **Update**.

Changing storage capacity threshold settings for an individual group

Perform the following steps to change the storage capacity threshold settings for an individual group:

Steps

1. Select **Control Center > Home > Member Details**.
2. Click **Aggregates**, to change aggregate options or, **File Systems** to change volume or qtree options.
3. Click the name of an aggregate, volume, or qtree.
4. Click **Edit Settings** under the **Tools** section in the left pane.
5. Edit the settings as needed.
6. Select the name of the group from **Apply Settings to** drop-down list.
7. Click **Update**.
8. Approve the change by clicking **OK** on the **verification** page.

Changing storage capacity threshold settings for a specific aggregate, volume, or qtree

Perform the following steps to change the storage capacity threshold settings for specific aggregate, volume, or qtree:

Steps

1. Click **Aggregates** to change Aggregate options or **File Systems** to change volume or qtree options.
2. Click the name of an aggregate, volume, or qtree.
3. Click **Edit Settings** under the **Tools** section in the left pane.
4. Edit the desired settings.

Note: To revert to the default settings, leave the fields empty.

5. Click **Update** and then click **OK**.

Management of aggregate capacity

You can manage aggregate capacity by gathering aggregate information and aggregate overcommitment information, by tracking aggregate space utilization, and by determining aggregate capacity threshold.

Volume space guarantees and aggregate overcommitment

You can use aggregate overcommitment to advertise more available space than the available space.

When managing storage resources, it is important to understand the role that aggregate overcommitment plays in space availability. To use aggregate overcommitment, you must create flexible volumes with a space guarantee of `none` or `file` so that the aggregate size does not limit the volume size. Each volume can be larger than its containing aggregate. You can use the storage space that the aggregate provides, as needed, by creating LUNs, or adding data to volumes.

By using aggregate overcommitment, the storage system can advertise more available storage than actually exists in the aggregate. With aggregate overcommitment, you could provide greater amounts of storage that you know would be used immediately. Alternatively, if you have several volumes that sometimes need to grow temporarily, the volumes can dynamically share the available space with each other.

Note: If you have overcommitted your aggregate, you must monitor its available space carefully and add storage as needed to avoid write errors due to insufficient space.

For details about volume space reservations, and aggregate overcommitment, see the *Data ONTAP Storage Management Guide*.

Related information

Data ONTAP Storage Management Guide - www.ibm.com/storage/support/nseries

Available space on an aggregate

With Operations Manager, you can determine the available space on an aggregate.

To help you determine the space availability on an aggregate, Operations Manager displays three values on the Aggregate Details page for each aggregate:

Aggregate size	Total size of the aggregate.
Capacity used	The total amount of disk space in use by volumes present in the aggregate.

Total committed capacity The total amount of committed disk space to volumes present in the aggregate. The total committed capacity can be larger than the total capacity by using aggregate overcommitment.

Considerations before modifying aggregate capacity thresholds

You must note the aggregate overcommitment point before changing aggregate capacity threshold.

Ensure that you take care of the following points when deciding whether to modify the aggregate capacity threshold:

Use of aggregate overcommitment strategy If you use the aggregate overcommitment strategy, you would want to increase the Aggregate Overcommitted threshold to more than 100. To determine how far beyond 100 to set the threshold, you must decide at what point the aggregate is too overcommitted. Ensure that you note the difference between your storage commitments and actual storage usage. Also, review the capacity graphs of historical data to get a sense of how the amount of storage used changes over time.

Set the Aggregate Full and Aggregate Nearly Full thresholds Set the Aggregate Full and Aggregate Nearly Full thresholds so that you have time to take corrective action if storage usage approaches capacity. Because the aggregate is overcommitted, you might want to set the Aggregate Full and Aggregate Nearly Full thresholds to values lower than the default. Lowering the thresholds generate an event well before completely filling the storage. Early notification gives you more time to take corrective action, such as installing more storage, before the storage space is full and write errors occur.

Non-usage of aggregate overcommitment If you do not use aggregate overcommitment as a storage-management strategy, you must leave the Aggregate Overcommitted and Nearly Overcommitted threshold values unchanged from their default.

Set the Aggregate Nearly Full threshold If an aggregate is routinely more than 80 percent full, set the Aggregate Nearly Full threshold to a value higher than the default.

Note: If you edit capacity thresholds for a particular aggregate, the edited thresholds override the global thresholds. You can edit thresholds for a particular aggregate, from the Aggregate Details page.

Aggregate capacity thresholds and their events

You can configure capacity thresholds for aggregates and events for these thresholds from DataFabric Manager server. You can set alarms to monitor the capacity and committed space of an aggregate. You can also take corrective actions based on the event generated.

You can configure alarms to send notification whenever an event related to the capacity of an aggregate occurs. For the Aggregate Full threshold, you can also configure an alarm to send notification only when the condition persists over a specified time.

By default, if you have configured an alarm to alert you to an event, the DataFabric Manager server issues the alarm only once per event. You can configure the alarm to repeat until you receive an acknowledgment.

Note: If you want to set an alarm for a specific aggregate, you must create a group with that aggregate as the only member.

You can set the following aggregate capacity thresholds:

Aggregate Full (%)	<p>Description: Specifies the percentage at which an aggregate is full.</p> <p>Note: To reduce the number of Aggregate Full Threshold events generated, you can set an Aggregate Full Threshold Interval. This causes DataFabric Manager server to generate an Aggregate Full event only if the condition persists for the specified time.</p> <p>Default value: 90 percent</p> <p>Event generated: Aggregate Full</p> <p>Event severity: Error</p> <p>Corrective Action</p> <p>Perform one or more of the following actions:</p> <ul style="list-style-type: none"> • To free disk space, ask your users to delete files that are no longer needed from volumes contained in the aggregate that generated the event. • You must add one or more disks to the aggregate that generated the event. <p>Note: After you add a disk to an aggregate, you cannot remove it without first destroying all flexible volumes present in the aggregate to which the disk belongs. You must destroy the aggregate after all the flexible volumes are removed from the aggregate.</p> <ul style="list-style-type: none"> • You must temporarily reduce the Snapshot reserve. By default, the reserve is 20 percent of disk space. If the reserve is not in use, reducing the reserve can free disk space, giving you more time to add a disk. There is no way to prevent Snapshot copies from consuming disk space greater than the amount reserved for them. It is, therefore, important to maintain a large enough reserve for Snapshot copies so that the active file system always has space available to create new files or modify existing ones. For more information about the Snapshot reserve, see the <i>Data ONTAP Data Protection Online Backup and Recovery Guide</i>.
Aggregate Nearly Full (%)	<p>Description: Specifies the percentage at which an aggregate is nearly full.</p> <p>Default value: 80 percent</p>

<p>Aggregate Overcommitted (%)</p>	<p>The value for this threshold must be lower than the value for Aggregate Full Threshold for DataFabric Manager server to generate meaningful events.</p> <p>Event generated: Aggregate Almost Full</p> <p>Event severity: Warning</p> <p>Corrective action</p> <p>Perform one or more of the actions mentioned in Aggregate Full.</p> <p>Description: Specifies the percentage at which an aggregate is overcommitted.</p> <p>Default value: 100 percent</p> <p>Event generated: Aggregate Overcommitted</p> <p>Event severity: Error</p> <p>Corrective action</p>
	<p>You should perform one or more of the following actions:</p> <ul style="list-style-type: none"> • You must create new free blocks in the aggregate by adding one or more disks to the aggregate that generated the event. <ul style="list-style-type: none"> Note: You must add disks with caution. After you add a disk to an aggregate, you cannot remove it without first destroying all flexible volumes present in the aggregate to which the disk belongs. You must destroy the aggregate after all the flexible volumes are destroyed. • You must temporarily free some already occupied blocks in the aggregate by taking unused flexible volumes offline. <ul style="list-style-type: none"> Note: When you take a flexible volume offline, it returns any space it uses to the aggregate. However, when you bring the flexible volume online again, it requires the space again. • Permanently free some already occupied blocks in the aggregate by deleting unnecessary files.
<p>Aggregate Nearly Overcommitted (%)</p>	<p>Description: Specifies the percentage at which an aggregate is nearly overcommitted.</p> <p>Default value: 95 percent</p> <p>The value for this threshold must be lower than the value for Aggregate Full Threshold for DataFabric Manager server to generate meaningful events.</p> <p>Event generated: Aggregate Almost Overcommitted</p> <p>Event severity: Warning</p> <p>Corrective action</p>

Aggregate Snapshot Reserve Nearly Full Threshold (%)	<p>Perform one or more of the actions provided in Aggregate Overcommitted.</p> <p>Description: Specifies the percentage of the Snapshot copy reserve on an aggregate that you can use before the system generates the Aggregate Snapshots Nearly Full event.</p> <p>Default value: 80 percent</p> <p>Event generated: Aggregate Snapshot Reserve Almost Full</p> <p>Event severity: Warning</p> <p>Corrective action: None</p> <p>There is no way to prevent Snapshot copies from consuming disk space greater than the amount reserved for them. If you disable the aggregate <code>Snapshot autodelete</code> option, it is important to maintain a large enough reserve.</p>
Aggregate Snapshot Reserve Full Threshold (%)	<p>See the <i>Operations Manager Help</i> for instructions on how to identify Snapshot copies you can delete. For more information about the Snapshot reserve, see the <i>Data ONTAP Data Protection Online Backup and Recovery Guide</i>.</p> <p>Description: Specifies the percentage of the Snapshot copy reserve on an aggregate that you can use before the system generates the Aggregate Snapshots Full event.</p> <p>Default value: 90 percent</p> <p>Event generated: Aggregate Snapshot Reserve Full</p> <p>Event severity: Warning</p> <p>Corrective action: None</p> <p>There is no way to prevent Snapshot copies from consuming disk space greater than the amount reserved for them.</p>

Note: A newly created traditional volume tightly couples with its containing aggregate so that the capacity of the aggregate determines the capacity of the new traditional volume. Therefore, you should synchronize the capacity thresholds of traditional volumes with the thresholds of their containing aggregates.

Related tasks

[Creating alarms](#) on page 115

Related information

[Data ONTAP Data Protection Online Backup and Recovery Guide - www.ibm.com/storage/support/nseries](http://www.ibm.com/storage/support/nseries)

Management of volume capacity

You can manage volume capacity by gathering volume information, by determining volume capacity threshold and events, by modifying volume capacity threshold, by setting volume Snapshot copy thresholds and events.

Volume capacity thresholds and events

DataFabric Manager server features thresholds to help you monitor the capacity of flexible and traditional volumes. You can configure alarms to send notification whenever an event related to the capacity of a volume occurs. You can also take corrective actions based on the event generated. For the Volume Full threshold, you can configure an alarm to send notification only when the condition persists over a specified period.

By default, if you have configured an alarm to alert you to an event, the DataFabric Manager server issues the alarm only once per event. You can configure the alarm to repeat until it is acknowledged.

Note: If you want to set an alarm for a specific volume, you must create a group with that volume as the only member.

You can set the following volume capacity thresholds:

Volume Full Threshold (%)

Description: Specifies the percentage at which a volume is considered full.

Note: To reduce the number of Volume Full Threshold events generated, you can set the Volume Full Threshold Interval to a non-zero value. By default, the Volume Full threshold Interval is set to zero. Volume Full Threshold Interval specifies the time during which the condition must persist before the event is triggered. Therefore, if the condition persists for the specified time, DataFabric Manager server generates a Volume Full event.

- If the threshold interval is 0 seconds or a value less than the volume monitoring interval, DataFabric Manager server generates the Volume Full events.
- If the threshold interval is greater than the volume monitoring interval, DataFabric Manager server waits for the specified threshold interval, which includes two or more monitoring intervals, and generates a Volume Full event only if the condition persisted throughout the threshold interval.

For example, if the monitoring cycle time is 60 seconds and the threshold interval is 90 seconds, the threshold event is generated only if the condition persists for two monitoring intervals.

Default value: 90

Event generated: Volume Full

Event severity: Error

Corrective action

Perform one or more of the following actions:

- Ask your users to delete files that are no longer needed, to free disk space.
- For flexible volumes containing enough aggregate space, you can increase the volume size.
- For traditional volumes containing aggregates with limited space, you can increase the size of the volume by adding one or more disks to the aggregate.

Note: Add disks with caution. After you add a disk to an aggregate, you cannot remove it without destroying the volume and its aggregate.

- For traditional volumes, temporarily reduce the Snapshot copy reserve. By default, the reserve is 20 percent of the disk space. If the reserve is not in use, reduce the reserve free disk space, giving you more time to add a disk. There is no way to prevent Snapshot copies from consuming disk space greater than the amount reserved for them. Therefore, it is important to maintain a large enough reserve for Snapshot copies. By maintaining the reserve for Snapshot copies, the active file system always has space available to create new files or modify existing ones. For more information about the Snapshot copy reserve, see the *Data ONTAP Data Protection Online Backup and Recovery Guide*.

Volume Nearly Full Threshold (%) Description: Specifies the percentage at which a volume is considered nearly full.

Default value: 80. The value for this threshold must be lower than the value for the Volume Full Threshold in order for DataFabric Manager server to generate meaningful events.

Event generated: Volume Almost Full

Event severity: Warning

Corrective action

Perform one or more of the actions mentioned in Volume Full.

Volume Space Reserve Nearly Depleted Threshold (%) Description: Specifies the percentage at which a volume is considered to have consumed most of its reserved blocks. This option applies to volumes with LUNs, Snapshot copies, no free blocks, and a fractional overwrite reserve of less than 100%. A volume that crosses this threshold is getting close to having write failures.

	<p>Default value: 80</p> <p>Event generated: Volume Space Reservation Nearly Depleted</p> <p>Event severity: Warning</p>
Volume Space Reserve Depleted Threshold (%)	<p>Description: Specifies the percentage at which a volume is considered to have consumed all its reserved blocks. This option applies to volumes with LUNs, Snapshot copies, no free blocks, and a fractional overwrite reserve of less than 100%. A volume that has crossed this threshold is getting dangerously close to having write failures.</p> <p>Default value: 80</p> <p>Event generated: Volume Space Reservation Depleted</p> <p>Event severity: Error</p> <p>When the status of a volume returns to normal after one of the preceding events, events with severity 'Normal' are generated. Normal events do not generate alarms or appear in default event lists, which display events of warning or worse severity.</p>
Volume Quota Overcommitted Threshold (%)	<p>Description: Specifies the percentage at which a volume is considered to have consumed the whole of the overcommitted space for that volume.</p> <p>Default value: 100</p> <p>Event generated: Volume Quota Overcommitted</p> <p>Event severity: Error</p> <p>Corrective action</p> <p>Perform one or more of the following actions:</p> <ul style="list-style-type: none">• Create new free blocks by increasing the size of the volume that generated the event.• Permanently free some of the occupied blocks in the volume by deleting unnecessary files.
Volume Quota Nearly Overcommitted Threshold (%)	<p>Description: Specifies the percentage at which a volume is considered to have consumed most of the overcommitted space for that volume.</p> <p>Default Value: 95</p> <p>Event generated: Volume Quota Almost Overcommitted</p> <p>Event Severity: Warning</p> <p>Corrective action:</p> <p>Perform one or more of the actions mentioned in Volume Quota Overcommitted.</p>

Volume Growth Event Minimum Change (%)	<p>Description: Specifies the minimum change in volume size (as a percentage of total volume size) that is acceptable. If the change in volume size is more than the specified value, and the growth is abnormal in relation to the volume-growth history, DataFabric Manager server generates a Volume Growth Abnormal event.</p> <p>Default value: 1</p> <p>Event generated: Volume Growth Abnormal</p>
Volume Snap Reserve Full Threshold (%)	<p>Description: Specifies the value (percentage) at which the space that is reserved for taking volume Snapshot copies is considered full.</p> <p>Default value: 90</p> <p>Event generated: Volume Snap Reserve Full</p> <p>Event severity: Error</p> <p>Corrective action: None</p> <p>There is no way to prevent Snapshot copies from consuming disk space greater than the amount reserved for them. If you disable the volume Snapshot <code>autodelete</code> option, it is important to maintain a large enough reserve. Disabling would ensure Snapshot copies that there is always space available to create new files or modify present ones. For instructions on how to identify Snapshot copies you can delete, see the <i>Operations Manager Help</i>.</p>
User Quota Full Threshold (%)	<p>Description: Specifies the value (percentage) at which a user is considered to have consumed all the allocated space (disk space or files used) as specified by the user quota. The user quota includes hard limit in the <code>/etc/quotas</code> file. If this limit is exceeded, DataFabric Manager server generates a User Disk Space Quota Full event or a User Files Quota Full event.</p> <p>Default value: 90</p> <p>Event generated: User Quota Full</p>
User Quota Nearly Full Threshold (%)	<p>Description: Specifies the value (percentage) at which a user is considered to have consumed most of the allocated space (disk space or files used) as specified by the user quota. The user quota includes hard limit in the <code>/etc/quotas</code> file. If this limit is exceeded, DataFabric Manager server generates a User Disk Space Quota Almost Full event or a User Files Quota Almost Full event.</p> <p>Default value: 80</p> <p>Event generated: User Quota Almost Full</p>

Volume No First Snapshot Threshold (%)	<p>Description: Specifies the value (percentage) at which a volume is considered to have consumed all the free space for its space reservation. This is the space that the volume needs when the first Snapshot copy is created.</p> <p>This option applies to volumes that contain space-reserved files, no Snapshot copies, a fraction of Snapshot copies overwrite reserve set to greater than 0, and where the sum of the space reservations for all LUNs in the volume is greater than the free space available to the volume.</p> <p>Default value: 90</p> <p>Event generated: Volume No First Snapshot</p>
Volume Nearly No First Snapshot Threshold (%)	<p>Description: Specifies the value (percentage) at which a volume is considered to have consumed most of the free space for its space reservation. This is the space that the volume needs when the first Snapshot copy is created.</p> <p>This option applies to volumes that contain space-reserved files, no Snapshot copies, a fractional overwrite reserve set to greater than 0, and where the sum of the space reservations for all LUNs in the volume is greater than the free space available to the volume.</p> <p>Default value: 80</p> <p>Event generated: Volume Almost No first Snapshot</p>

Note: When a traditional volume is created, it is tightly coupled with its containing aggregate so that its capacity is determined by the capacity of the aggregate. For this reason, you should synchronize the capacity thresholds of traditional volumes with the thresholds of their containing aggregates.

Related concepts

[Volume Snapshot copy thresholds and events](#) on page 199

Related tasks

[Creating alarms](#) on page 115

Related information

[Data ONTAP Data Protection Online Backup and Recovery Guide - www.ibm.com/storage/support/nseries](http://www.ibm.com/storage/support/nseries)

Normal events for a volume

Normal events do not generate alarms or appear in default event lists, which display events of Warning or worst severity.

To view normal events for a volume, do either of the following:

- Display the Volume Details page for the volume.
- Click the Events tab; then go to the Report drop-down list and select the History report.

Modification of the thresholds

You can set the thresholds to a value, higher or a lower, than the default.

You might want to set the thresholds to a value higher than the default if storage space is routinely more than 80 percent full. Leaving the Nearly Full Threshold at the default value might generate events that notify you that storage space is nearly full more often than you want.

You might want to set the thresholds to a value lower than the default. Lowering the threshold ensures that the DataFabric Manager server generates the event well before completely filling the storage. An early notification gives you more time to take corrective action before the storage space is full.

Management of qtree capacity

You can manage qtree capacity by gathering qtree information, tracking qtree capacity utilization, and determining qtree capacity threshold and events.

Volume Snapshot copy thresholds and events

You can set alarms whenever a Snapshot copy is taken on a flexible or a traditional volume.

The DataFabric Manager server features thresholds to help you monitor Snapshot copy usage for flexible and traditional volumes. You can configure alarms to send notification whenever a volume Snapshot copy event occurs.

By default, if you have configured an alarm to alert you to an event, the DataFabric Manager server issues the alarm only once per event. You can configure the alarm to repeat until it is acknowledged.

Note: If you want to set an alarm on a specific volume, you must create a group with that volume as the only member.

You can set the following volume Snapshot thresholds:

Volume Snap Reserve Full Threshold (%)	Description: Specifies the percentage at which the space that is reserved for taking volume Snapshot copies is considered full. Default value: 90 Event generated: Snapshot Full
---	--

Event severity: Warning

Corrective action:

1. Access the Volume Snapshot copy details report.
2. Select the Snapshot copies.
3. Click **Compute Reclaimable**.

**Volume Nearly
No First
Snapshot
Threshold (%)**

Description: Specifies the percentage at which a volume is considered to have consumed most of the free space for its space reservation. This is the space that the volume needs when the first Snapshot copy is created. This option applies to volumes that contain space-reserved files, no Snapshot copies, a fractional overwrite reserve set to greater than 0, and where the sum of the space reservations for all LUNs in the volume is greater than the free space available to the volume.

Default value: 80 percent

Event generated: Nearly No Space for First Snapshot

Event severity: Warning

**Volume No
First Snapshot
Threshold (%)**

Description: Specifies the percentage at which a volume is considered to have consumed all the free space for its space reservation. This is the space that the volume needs when the first Snapshot copy is created. This option applies to volumes that contain space-reserved files, no Snapshot copies, a fractional overwrite reserve set to greater than 0, and where the sum of the space reservations for all LUNs in the volume is greater than the free space available to the volume.

Default value: 90 percent

Event generated: No Space for First Snapshot

Event severity: Warning

**Volume
Snapshot
Count
Threshold**

Description: Specifies the number of Snapshot copies, which, if exceeded, is considered too many for the volume. A volume is allowed up to 255 Snapshot copies.

Default value: 250

Event generated: Too Many Snapshots

Event severity: Error

**Volume Too
Old Snapshot
Threshold**

Description: Specifies the age of a Snapshot copy, which, if exceeded, is considered too old for the volume. The Snapshot copy age can be specified in seconds, minutes, hours, days, or weeks.

Default value: 52 weeks

Event generated: Too Old Snapshot copies

Event severity: Warning

Related concepts

Volume capacity thresholds and events on page 194

Related references

Guidelines for configuring alarms on page 114

Qtree capacity thresholds and events

Operations Manager enables you to monitor qtree capacity and set alarms. You can also take corrective actions based on the event generated.

DataFabric Manager server features thresholds to help you monitor the capacity of qtrees. Quotas must be enabled on the storage systems. You can configure alarms to send notification whenever an event related to the capacity of a qtree occurs.

By default, if you have configured an alarm to alert you to an event, the DataFabric Manager server issues the alarm only once per event. You can configure the alarm to continue to alert you with events until it is acknowledged. For the Qtree Full threshold, you can also configure an alarm to send notification only when the condition persists over a specified period.

Note: If you want to set an alarm for a specific qtree, you must create a group with that qtree as the only member.

You can set the following qtree capacity thresholds:

Qtree Full (%)	<p>Description: Specifies the percentage at which a qtree is considered full.</p> <p>Note: To reduce the number of Qtree Full Threshold events generated, you can set a Qtree Full Threshold Interval to a non-zero value. By default, the Qtree Full threshold Interval is set to zero. The Qtree Full Threshold Interval specifies the time during which the condition must persist before the event is generated. If the condition persists for the specified amount of time, DataFabric Manager server generates a Qtree Full event.</p> <ul style="list-style-type: none"> • If the threshold interval is 0 seconds or a value less than the volume monitoring interval, DataFabric Manager server generates Qtree Full events. • If the threshold interval is greater than the volume monitoring interval, DataFabric Manager server waits for the specified threshold interval, which includes two or more monitoring intervals, and generates a Qtree Full event only if the condition persisted throughout the threshold interval.
-----------------------	--

For example, if the monitoring cycle time is 60 seconds and the threshold interval is 90 seconds, the threshold event is generated only if the condition persists for two monitoring intervals.

Default value: 90 percent

Event generated: Qtree Full

Event severity: Error

Corrective action

Perform one or more of the following actions:

- Ask users to delete files that are no longer needed, to free disk space.
- Increase the hard disk space quota for the qtree.

Qtree Nearly Full Threshold (%) Description: Specifies the percentage at which a qtree is considered nearly full.
Default value: 80 percent
Event severity: Warning

Corrective action

Perform one or more of the following actions:

- Ask users to delete files that are no longer needed, to free disk space.
- Increase the hard disk space quota for the qtree.

Related tasks

[Creating alarms](#) on page 115

Related information

[Data ONTAP Data Protection Online Backup and Recovery Guide - www.ibm.com/storage/support/nseries](http://www.ibm.com/storage/support/nseries)

How Operations Manager monitors volumes and qtrees on a vFiler unit

By using Operations Manager, you can monitor volumes and qtrees on a vFiler unit.

The DataFabric Manager server monitors storage resources (volumes and qtrees) that are configured on a vFiler unit.

During initial discovery, the DataFabric Manager server uses SNMP to discover the volumes and qtrees as a hosting storage system's objects. After it discovers a configured vFiler unit on the hosting storage system, the DataFabric Manager server assigns the resource objects to the vFiler unit.

The DataFabric Manager server maintains information in its database about volumes and qtrees that are removed or destroyed from a vFiler unit. As the volumes and qtrees are reassigned to other vFiler units, the DataFabric Manager server uses the stored information to update resource ownership.

How Operations Manager monitors qtree quotas

You can monitor qtree quotas by using Operations Manager. As the DataFabric Manager server monitors hosting storage systems for vFiler unit storage resources, it also provides information about qtree quotas.

Where to find vFiler storage resource details

With Operations Manager, you can view the volumes and qtrees on a vFiler unit in the vFiler Details page.

The vFiler Details page (**Member Details > vFilers**) provides you with a link to the volumes and qtrees assigned to a vFiler unit.

The Volume Details and Qtree Details pages provide you with details about the volumes and qtrees that are assigned to a vFiler unit.

What clone volumes are

Clone volumes are fully functional volumes that always exist in the same aggregate as their parent volumes.

Data ONTAP enables you to create writable copy of a volume, known as volume clone.

A clone is a point-in-time, writable copy of the parent volume. Changes made to the parent volume after the clone is created are not reflected in the clone. Clone volumes can themselves be cloned.

Clone volumes and their parent volumes share the same disk space for any data common to the clone and parent. This means that creating a clone is instantaneous and requires no additional disk space (until changes are made to the clone or parent). If you later decide you want to sever the connection between the parent and the clone, you can split the clone. This removes all restrictions on the parent volume and enables the space guarantee on the clone. For general information about clone volumes and clone parents, see the *Data ONTAP Storage Management Guide*.

The DataFabric Manager server helps you manage clone hierarchies by making it easier to view clone relationships between volumes. By using Operations Manager, you can view clone volume and the parent volume information.

Related information

Data ONTAP Storage Management Guide - www.ibm.com/storage/support/nseries

Identification of clones and clone parents

By using Operations Manager, you can view details of clones and their parents.

You can display the Volumes, Clone List report by selecting **Member Details > File Systems > Report drop-down list**. Cloned volumes have an entry in the Clone Parent column of the report, indicating the name of the parent. Volumes that have clone children have the names of those children, which are included in the Clones column. Each clone name links to the Volume Details page for the clone child.

Alternatively, the Volume Details page for a clone child includes the name of the clone parent, which is a link to the Volume Details page for the parent volumes.

If a volume is a clone parent, the Related Storage section of its Volume Details page includes a link to a list of its direct clone children.

Why Snapshot copies are monitored

The Snapshot copy monitoring and space management help you monitor and generate reports on Snapshot copies and how they influence your space management strategy.

By using DataFabric Manager Server, you can determine the following information about Snapshot copies:

- How much aggregate and volume space is used for Snapshot copies?
- Is there adequate space for the first Snapshot copy?
- Which Snapshot copies can be deleted?
- Which volumes have high Snapshot copy growth rates?
- Which volumes have Snapshot copy reserves that are nearing capacity?

See the *Operations Manager Help* for instructions.

Snapshot copy monitoring requirements

To use the Snapshot copy monitoring features, the DataFabric Manager server requires a valid login name and password for each system being monitored.

Detection of Snapshot copy schedule conflicts

By using Operations Manager, you can monitor conflicts between the Snapshot copy schedule and SnapMirror and SnapVault schedules.

When Snapshot copies are scheduled for a volume, the DataFabric Manager server monitors for conflicts between the Snapshot copy schedule and SnapMirror and SnapVault schedules on the same volume. Conflicts can cause scheduled Snapshot copies to fail. The Aggregate Details and Volume Details pages both feature a Protection area that indicates whether scheduled Snapshot copies and SnapMirror are enabled.

The DataFabric Manager server generates a schedule conflict event if a volume is configured with both Snapshot copy and SnapVault copy schedules. An event is also generated if a Snapshot copy is

scheduled at the same time as a SnapMirror transfer. The DataFabric Manager server generates these events only if the SnapVault or SnapMirror relationships are monitored on the volume.

Dependencies of a Snapshot copy

You can view the dependencies of a Snapshot copy, whether you can delete a Snapshot copy, and the steps to delete it using Operations Manager.

The Snapshots area of the Volume Details page displays information about, up to 10 of the most recent Snapshot copies for the volume. This includes the last time the Snapshot copies that were accessed and their dependencies, if any. This information helps you determine whether you can delete a Snapshot copy or, if the Snapshot copy has dependencies, what steps you need to take to delete the Snapshot copy.

To generate a page that lists dependent storage components for a Snapshot copy, and the steps you would need to take to delete the copy, click the hyperlinked text in the Dependency column for that Snapshot copy. The link is not available when the dependency is due to SnapMirror or SnapVault or to a FlexClone volume, that is offline.

Thresholds on Snapshot copies

You can set a threshold of a maximum number of Snapshot copies to determine when to delete a Snapshot copies.

You can avoid the problem of Snapshot failures due to inadequate space. The thresholds that address the number and the age of volume Snapshot copies let you know when you must delete Snapshot copies. You must set the Volume Snapshot Count threshold to define the maximum number of Snapshot copies for a volume.

You must set the Volume Snapshot Too Old threshold to indicate the maximum allowable age for Snapshot copies of the volume. The DataFabric Manager server generates events when it exceeds the thresholds.

Storage chargeback reports

You can create storage chargeback report using Operations Manager to collect the amount of space used or allocated in specific storage objects (storage system, volume, qtree) or a group of storage objects.

These reports can be generated in various formats, such as Perl, .txt, Comma-separated values (.csv), Excel (.xls), and .xml.

The storage chargeback feature of the DataFabric Manager server provides billing reports for the amount of space used or allocated in specific storage objects, or a group of storage objects. The storage objects include storage system, volume, qtree, or user. The billing reports contain information such as average use, length of billing cycle, rate per GB of space used, and charges based on the rate and use.

You can specify the day when the billing cycle starts, the rate, and the format for currency in the report.

Storage chargeback reports provide an efficient way to track space used and space that is allocated for generating bills based on your specifications. Chargeback reports are useful if your organization bills other organizations, groups, or users in your company for the storage services they use.

When is data collected for storage chargeback reports

You can collect data for chargeback reports for a specific period.

The data reported by the chargeback feature on a specific day is based on the last data sample that is collected before midnight GMT of the previous night. For example, if the last data sample before midnight on April 17 was collected at 11:45 p.m. GMT, the chargeback reports viewed on April 18 display details about average use, charges, and other data based on the sample collected on April 17 at 11:45 p.m.

Determine the current month's and the last month's values for storage chargeback report

You can calculate the current and previous month's chargeback report.

When you select a Chargeback, This Month or Last Month view, the data displayed pertains to the current or the last billing cycles, respectively. The Day of the Month for Billing option determines when the current month begins and the last month ends, as described in the following example:

Company A's DataFabric Manager server system is configured for the billing cycle to start on the fifth day of every month. If Chris (an administrator at Company A) views the Chargeback, this Month report on April 3, the report displays data for the period of March 5 through midnight (GMT) of April 2. If Chris views the Chargeback, Last Month report on April 3, the report displays data for the period of February 5 through March 4.

All chargeback reports contain Period Begin and Period End information that indicates when the billing cycle begins and ends for the displayed report.

Chargeback reports in various formats

You can generate chargeback report in various formats by running the `dfm report view -F format _report-name` command.

The DataFabric Manager server does not integrate with any specific billing program. However, the report does provide data in a spreadsheet form that you can use for other billing applications. The data is accessible through the spreadsheet icon on the right side of the Report drop-down list.

You can also generate chargeback data in other formats, such as Perl, .txt, Comma-separated values (.csv), Excel (.xls), and .xml by using the `dfm report` command. For example, to generate the chargeback reports in Perl format so that other billing applications can use the chargeback data, you should issue the following command: `dfm report view -F perl report_name`.

In this command, `report_name` is one of the following:

- groups-chargeback-this-month
- groups-chargeback-last-month
- groups-chargeback-allocation-this-month
- groups-chargeback-allocation-last-month
- filers-chargeback-this-month
- filers-chargeback-last-month
- filers-chargeback-allocation-this-month
- filers-chargeback-allocation-last-month
- volumes-chargeback-this-month
- volumes-chargeback-last-month
- volumes-chargeback-allocation-this-month
- volumes-chargeback-allocation-last-month
- qtrees-chargeback-this-month
- qtrees-chargeback-last-month
- qtrees-chargeback-allocation-this-month
- qtrees-chargeback-allocation-last-month
- clusters-chargeback-this-month
- clusters-chargeback-last-month
- clusters-chargeback-allocation-this-month
- clusters-chargeback-allocation-last-month
- vservers-chargeback-this-month
- vservers-chargeback-last-month
- vservers-chargeback-allocation-this-month
- vservers-chargeback-allocation-last-month

For more information about using the `dfm report` command, see the DataFabric Manager server man pages.

The chargeback report options

The chargeback report options enable you to specify the chargeback increment, the currency format, the chargeback rate, and the day when the billing cycle starts.

You can specify storage chargeback option at a global or a group level. The global level can be chargeback increment, the currency format, the chargeback rate, or specify an annual charge rate for objects in a specific group.

In addition to these global settings, you can specify an annual charge rate (per GB) for objects in a specific group. The annual charge rate specified for a group overrides the setting specified at the global level.

Specifying storage chargeback options at the global or group level

By using Operations Manager, you can set chargeback options at a global or group level.

Step

1. Following are steps to set chargeback options:

To apply changes to...	Go to...
All objects that the DataFabric Manager server manages	The Options page (Setup > Options link); then select Chargeback in the Edit Options section.
Objects in a specific group	The Edit Group Settings page (click Edit Groups in the left pane); then click the Edit column for the group for which you want to specify an annual charge rate.

The storage chargeback increment

The storage chargeback increment indicates how the charge rate is calculated.

You can specify storage chargeback increment using Operations Manager. You can specify this setting only at the global level. By default, the chargeback increment is Daily. The following values can be specified for this option:

- Daily** Charges are variable; they are adjusted based on the number of days in the billing period. The DataFabric Manager server calculates the charges as follows: Annual Rate / 365 * number of days in the billing period.
- Monthly** Charges are fixed; there is a flat rate for each billing period regardless of the number of days in the period. The DataFabric Manager server calculates the charges as follows: Annual Rate/12.

Currency display format for storage chargeback

You can specify currency formats to display in Operations Manager.

The Currency Format setting indicates the format to use for displaying currency amounts in Operations Manager.

By default, the format is \$ #,###.##, where # indicates a digit. If you need to specify any other format, use the following guidelines:

- You must specify four # characters before the decimal point.
A decimal point separates the integer part of a number from its fractional part. For example, in the number 5.67, the period (.) is the decimal point.
The symbol used as a decimal point depends on the type of currency. For example, a period (.) is used for US dollars and a comma (,) is used for Danish Kroner
- Although a decimal separator is optional in the currency format, if you use it, you must specify at least one # character after the decimal separator. For example, \$ #,###.# and JD #,###.###.

- You can optionally specify a thousands-separator.
A thousands-separator separates digits in numbers into groups of three. For example, the comma (,) is the thousands-separator in the number 567,890,123. The symbol used as a thousands-separator depends on the type of currency. For example, a comma (,) is used for US dollars and a period (.) is used for Danish Kroner.
- You can use any currency symbol, such as EUR or ¥, to suit your needs.
If the currency symbol you want to use is not part of the standard ASCII character set, use the code specified by the HTML Coded Character Set. For example, use ¥ for the Yen (¥) symbol.
- You can specify only one currency format per the DataFabric Manager server.
For example, if you specify \$ #, ###.## as your currency format for a specific installation, this format is used for all chargeback reports generated by that installation.

Specification of the annual charge rate for storage chargeback

You can set annual charge rate for storage chargeback at a global level or a specific group.

The Annual Charge Rate (per GB) setting indicates the amount to charge for storage space used per GB per year. You can specify this setting at the global level, in addition to specific groups.

By default, no rate is specified. You must specify a value for this option for the DataFabric Manager server to generate meaningful chargeback reports.

Specify this value in the x.y format, where x is the integer part of the number and y is the fractional part. For example, to specify an annual charge rate of \$150.55, enter 150.55.

Note: You must use a period (.) to indicate the fractional part of the number in the Annual Charge Rate box. Even if you are specifying a currency format that uses a comma (,) as the decimal separator. For example, to specify 150,55 Danish Kroner, enter 150.55.

Specification of the Day of the Month for Billing for storage chargeback

You can specify the day of the month from which the billing cycle begins.

The Day of the Month for Billing setting indicates the day of the month on which the billing cycle begins.

By default, this value is set to 1. The following values can be specified for this option:

1 through 28 These values specify the day of the month. For example, if you specify 15, it indicates the fifteenth day of the month.

-27 through 0 These values specify the number of days before the last day of the month. Therefore, 0 specifies the last day of the month.

For example, if you want to bill on the fifth day before the month ends every month, specify -4.

Formatted charge rate for storage chargeback

You can view the annual charge rate for storage chargeback in the specified format. For example, if the currency format is \$ #,###.## and the annual charge rate is 150.55, the Formatted Charge Rate option displays \$150.55.

The Formatted Charge Rate setting displays the annual charge rate value in the currency format. The value is automatically generated and displayed based on the currency format and the annual charge rate you specify. You cannot set or change this option.

What happens when storage objects are deleted

With Operations Manager, you can stop monitoring a storage object (aggregate, volume, or qtree) by deleting it from the Global group. When you delete an object, the DataFabric Manager server stops collecting and reporting data about it. Data collection and reporting is resumed only when the object is added back to the OnCommand console database.

Note: When you delete a storage object from any group other than Global, the object is deleted only from that group; DataFabric Manager server continues to collect and report data about it. You must delete the object from the Global group if you want the DataFabric Manager server to stop monitoring it.

Reports of deleted storage objects

All storage objects deleted from the DataFabric Manager server database are listed in various reports.

Following are the reports of deleted storage objects:

- Storage Systems, Deleted
- vFilers, Deleted
- File Systems, Deleted
- Volumes, Deleted
- Qtrees, Deleted
- Aggregates, Deleted
- SAN Hosts, Deleted
- LUNs, Deleted

Note: These reports are accessible from the Report drop-down list on the **Member Details** tab for each storage object (**Storage Systems**, **vFiler units**, **File Systems**, **Aggregates**, **SANs**, and **LUNs**).

Undeleting a storage object for monitoring

You can undelete a storage object using Operations Manager.

Steps

1. Select the check box next to each object you want to return to the database.
2. Click **Undelete**.

Storage system management

You can use Operations Manager to view the status of and report of groups, view configuration status, view and respond to events, configure alarms, and so on.

Management tasks performed using Operations Manager

You can use Operations Manager to manage your storage system.

As soon as the DataFabric Manager server is installed, it begins the process of discovering, monitoring, and gathering data about your supported storage systems. However, before you can use the data to simplify your network administration tasks, you need to understand the different ways you can use Operations Manager to manage your storage systems.

By using Operations Manager, you can do the following:

- Create groups.
- View the status of and obtain reports and information for a group of systems.
- View information for individual systems.
- Access the console of a storage system.
- View the active/active configuration status and perform takeover and giveback operations if the storage system is an active/active controller.
- View and respond to events.
- Configure alarms that send you notification if the DataFabric Manager server logs a specific type of event or severity of event.
- Edit the configuration settings of a storage system.
- Link to FilerView for a selected storage system or vFiler unit.
- Insert values into custom comment fields.
- View user, qtree, and group quotas.
- Edit user quotas.

Operations Manager components for managing your storage system

By using Operations Manager you can create or modify groups, view information about vFiler units, configure alarms, and so on.

You can perform the following tasks by using Operations Manager to manage your storage systems:

- Create new groups or modify, move, copy, or delete existing groups—**Control Center > Home > Edit Groups**.

- View information about all or a group of storage systems and access details of specific storage systems—**Control Center > Member Details > Physical Systems**.
- View information about vFiler units that are configured on hosting storage systems and access details of specific vFiler units —**Control Center > Member Details > Virtual Systems**.
- View and respond to system events—**Control Center > Group Status > Events**.
- Configure alarms for generated events, manage the DataFabric Manager server administrators, establish roles to manage the DataFabric Manager server access, and configure custom reports—**Control Center > Setup > Alarms**.
- Modify passwords for one or multiple storage systems, manage storage system configurations, and manage all scripts installed on the DataFabric Manager server—**Control Center > Management**.
- Compare storage system configurations and configuration files against a template and modify R global options—**Control Center > Management**.
- Configure host users and roles—**Control Center > Management > Host Users**.

Storage system groups

Using Operations Manager you can create and manage groups of storage systems.

Operations Manager is designed around the concept of groups. When a group is selected in the left pane of the Operations Manager main window, the pages change to display information relating to that group.

To display information about all your storage systems, select the Global group in the Groups pane on the left side of Operations Manager. The Global group is the default group containing the superset of all storage systems.

To display information about a specific group of systems, select the desired group name in the Groups pane on the left side of Operations Manager.

To manage your storage systems effectively, you should organize them into smaller groups so that you can view information only about objects in which you are interested. You can group your storage systems to meet your business needs, for example, by geographic location, operating system version, and storage system platform.

Custom comment fields in Operations Manager

You can create custom comment fields and associate them with specific storage systems, SAN hosts, FC switches, aggregates, volumes, qtrees, LUNs, groups, and quota users.

You can use the custom comment fields for any purpose. One example would be to associate a department code with a quota user, for use in chargeback for business accounting. You can use the Search function in Operations Manager to display every object that contains specific data in a custom comment field.

Using custom comment fields in Operations Manager has three aspects: creating the field, inserting data to create specific comments, and viewing comments.

Creating the comment field: You create the custom comment field using **Setup menu > Options > Custom Comment Fields**.

Inserting data: You insert data into the custom comment field in the Edit Settings page for the object you want to associate with the comment. For example, to associate a comment with a qtree, use the Edit Qtree Settings page for that qtree.

Viewing comments: You view custom comment data for multiple objects in the Comments report for the type of object, for example, the Qtrees Comments report.

You can also view the comments for a single object in its Details page—for example, the Qtree Details page for a specific qtree.

For detailed instructions on creating custom comment fields, see the *Operations Manager Help* for the Options page.

Consolidated storage system and vFiler unit data and reports

By using Operations Manager, you can view storage system and vFiler unit reports.

You can view storage system related information from **Control Center > Member Details > Physical Systems > Report list**.

You can view vFiler-related information from **Control Center > Member Details > Virtual Systems > Report list**.

You can view global and group information and select individual system data in detail using the Member Details report pages.

Tasks performed by using the storage systems and vFiler unit report pages

You can view system data for all groups, generate spreadsheet reports, get information about storage systems, and launch FilerView.

Similar to the other Operations Manager Control Center tab pages, the Appliances and vFiler reports enable you to view a wide variety of details in one place. You can perform the following tasks:

- View system data for all or a group of monitored systems
- Generate spreadsheet reports
- Obtain detailed information about a specific storage system
- Launch FilerView

Where to find information about a specific storage system

By using Operations Manager, you can view information about a specific storage system or a vFiler unit, and view or modify configurations of a storage system or a vFiler unit.

You can view the details by clicking the storage system or vFiler unit name on the Operations Manager reports.

The DataFabric Manager server regularly refreshes monitoring data for the entire group within which a storage system or vFiler unit resides, or you can click Refresh Group Monitors to manually refresh the data.

Tasks performed from a Details page of Operations Manager

You can view and modify storage system or vFiler unit configurations, view and check active/active configurations, view events related to the storage system or vFiler unit, and so on.

You can perform the following storage system management tasks from the Details page:

- View specific storage system or vFiler unit details.
- Edit the storage system or vFiler unit configuration using FilerView.
- View the active/active configuration status and perform takeover and giveback operations by using the cluster console (on active/active controllers only).
- Access the vFiler units that are hosted on a storage system.
- Check active/active controller configurations.
- Edit the storage system configuration using FilerView.
- Edit Remote LAN Module (RLM) port settings for the storage system.
- View events related to the storage system or vFiler unit.
- View graphing information specific to each type of storage system.

Editable options for storage system or vFiler unit settings

You can specify or change the storage system or vFiler unit settings using Operations Manager.

You can use the Edit Storage Controller Settings page to specify or change storage system or vFiler unit settings. Note, however, that you can set global values for many settings using the Options page. You do not need to modify storage system or vFiler unit-level settings unless they differ from your global values.

You can use the Edit Storage Controller Settings page to modify the following information:

IP address This is the IP address of the storage system that the DataFabric Manager server monitors.

You might want to change the storage system IP address if you want to use a different interface for administrative traffic.

Login and password	You should configure a login and password if you want to use Operations Manager to run a command on a system. Operations Manager uses this information to authenticate itself to the storage system on which the command is run. Configuration of login and password is mandatory.
Authentication	You can also set up authentication by using the <code>/etc/hosts.equiv</code> file on the storage system. For information about configuring the <code>/etc/hosts.equiv</code> file, see the <i>Data ONTAP Storage Management Guide</i> .
Threshold values	The threshold values indicate the level of activity that must be reached on the storage system before an event is triggered. By using these options, you can set specific storage system or group thresholds. For example, the Appliance CPU Too Busy threshold indicates the highest level of activity the CPU can reach before a CPU Too Busy event is triggered. Threshold values specified on this page supersede any global values specified on the Options page.
Threshold intervals	The threshold interval is the period of time during which a specific threshold condition must persist before an event is triggered. For example, if the monitoring cycle time is 60 seconds and the threshold interval is 90 seconds, the event is generated only if the condition persists for 2 monitoring intervals. You can configure threshold intervals only for specific thresholds, as listed on the Options page.

Related information

Data ONTAP Storage Management Guide - www.ibm.com/storage/support/nseries

What Storage Controller Tools list is

You can view storage system or controller details using Storage Controller Tools of Operations Manager.

You can use Storage Controller Tools to set up parameters needed to communicate with a storage system or a controller. You can perform the following tasks using Storage Controller Tools:

- Modify settings such as the primary IP address, remote platform management IP address, login and password, and so on.
- Diagnose the network connectivity of the storage system or the controller.
- Refresh the monitoring samples collected by Operations Manager.
- Run a command on the storage system or the controller.
- Connect to the device console.
- Gather information about the host (storage system or controller) users.

You can access Storage Controller Tools from the Details page for the storage system or hosting storage system (of a vFiler unit), or controller. The tools menu is located at the lower left display of Operations Manager.

Note: Storage Controller Tools is applicable to storage systems running Data ONTAP 8.0 7-Mode.

What Cluster Tools list is

You can view and modify cluster or cluster controller details using Cluster Tools of Operations Manager.

You can use Cluster Tools to set up parameters needed to communicate with a cluster. You can perform the following tasks using Cluster Tools:

- Modify settings such as the primary IP address, monitoring options, management options such as login and password, and so on.
- Diagnose the network connectivity of the cluster.
- Refresh the monitoring samples collected by Operations Manager.
- Run a command on the cluster.

You can access Cluster Tools from the Cluster Details page for the cluster. The tools menu is located at the lower left display of Operations Manager.

Note: Cluster Tools is applicable to storage systems running Data ONTAP 8.0 Cluster-Mode.

What the Diagnose Connectivity tool does

By using the Diagnose Connectivity tool, you can perform connectivity tests and review test outcome.

The Diagnose Connectivity tool queries the DataFabric Manager server database about a selected storage system, runs connectivity tests, and displays information and test outcomes. The sequence of steps depends on whether the storage system is managed or unmanaged. A managed storage system is one that is in the DataFabric Manager server database. An unmanaged storage system is one that is not in the DataFabric Manager server database.

The Refresh Monitoring Samples tool

You can view updated storage system details using Refresh Monitoring Samples in Operations Manager.

You can specify the frequency at which Operations Manager collects information by using the system information-monitoring interval.

The Run a Command tool

By using the Run a Command tool in Operations Manager, you can run commands on storage systems.

The Run a Command tool provides you with an interface to do the following:

- Run Data ONTAP commands on storage systems.
- Run any Remote LAN Module (RLM) command on the RLM card that is installed on a storage system.

Prerequisite

The DataFabric Manager server uses the following connection protocols for communication:

- Remote Shell (RSH) connection for running a command on a storage system
To establish an RSH connection and run a command on a storage system, the DataFabric Manager server must authenticate itself to the storage system. Therefore, you must enable RSH access to the storage system and configure login and password credentials that are used to authenticate Data ONTAP.
- Secure Socket Shell (SSH) connection for running a command on an RLM card, if the installed card provides a CLI.

Restrictions

The following restrictions exist:

- There are several Data ONTAP run commands that are available on storage systems, but are restricted in the DataFabric Manager server. For a list of restricted commands, see the *Operations Manager Help*.
- You cannot run a command on the Global group.

Related concepts

[Remote configuration of a storage system](#) on page 222

[DataFabric Manager server CLI to configure storage systems](#) on page 222

[Prerequisites for running remote CLI commands from Operations Manager](#) on page 222

[What the remote platform management interface is](#) on page 269

Related tasks

[Running commands on a specific storage system](#) on page 223

[Running commands on a group of storage systems from Operations Manager](#) on page 223

The Run Telnet tool

You can connect to the storage system using the Run Telnet tool in Operations Manager.

Console connection through Telnet

By using Operations Manager, you can connect the storage system console.

Use the Connect to Device Console tool to connect to the storage system console. The storage system must be connected to a terminal server for the DataFabric Manager server to connect to the storage system console.

Note: Before initiating the console connection, you must set the Console Terminal Server Address in the Edit Settings page for the storage system.

Managing active/active configuration with the DataFabric Manager server

You can monitor and manage active/active configuration with the cluster console of Operations Manager.

The cluster console enables you to view the status of an active/active configuration (controller and its partner) and perform takeover and giveback operations between the controllers.

For detailed information about active/active configurations, see the *Data ONTAP Storage Management Guide*.

Related information

Data ONTAP Storage Management Guide - www.ibm.com/storage/support/nseries

Viewing the status of an active/active configuration by using cluster console

An authentication method must be set up for the DataFabric Manager server to authenticate the controller on which takeover and giveback operations are to be performed. Login and password must be set for the storage system. You can access the cluster console from Operations Manager to view the status of an active/active configuration.

Steps

1. Click **Control Center > Home > Member Details > Physical Systems**.
2. From the Report drop-down list, select Active/Active Controllers, All.
3. Click the controller for which you want to view the status of the active/active configuration.
4. Click **View Cluster Console** under Storage Controller Tools.

What the Takeover tool does

You can use the Takeover tool from the Tools list to initiate a manual takeover of the controller's partner. The Takeover tool is available in the Tools list only when the controller whose Tools list you are viewing can take over its partner.

Once you select Takeover, the Takeover page is displayed. The Takeover page enables you to select the type of takeover you want the controller to perform. You can select from one of the following options:

Take Over Normally	This option is the equivalent of running the <code>cf takeover</code> command in which the controller takes over its partner in a normal manner. The controller allows
---------------------------	--

its partner to shut down its services before taking over. This option is used by default.

- Take Over Immediately** This option is the equivalent of running the `cf takeover -f` command in which the controller takes over its partner without allowing the partner to gracefully shut down its services.
- Force a Takeover** This option is the equivalent of running the `cf forcetakeover -f` command in which the controller takes over its partner even in cases when takeover of the partner is normally not allowed. Such a takeover might cause data loss.
- Takeover After a Disaster** This option is for MetroClusters only and is the equivalent of running the `cf forcetakeover -f -d` command. Use this option if the partner is unrecoverable.

Note: The `Force a Takeover` and `Takeover After a Disaster` options, are also available in circumstances when the interconnect between the controller, and its partner is down. It enables you to manually take over the partner.

Once you have made a selection, the `Status` option on the Cluster Console page displays the status of the takeover operation. Once the takeover operation is complete, the Cluster Console page displays the updated controller-icon colors. The Cluster Console page also displays the status of each controller. The Tools list of each controller is adjusted appropriately to indicate the active/active configuration operation each controller can now perform.

What the Giveback tool does

You can use the Giveback tool to initiate a giveback operation from a controller that has taken over its partner. The Giveback tool is available in the Tools list only when the controller whose Tools list you are viewing can give back to its partner.

Once you select Giveback for the controller, the Giveback page is displayed. You can select from one of the following giveback options:

- Give Back Normally** This option is the equivalent of the `cf giveback` command in which the controller performs a graceful shutdown of services and aborts CIFS operations. The controller also shuts down long-running jobs that are running on the controller on behalf of the taken over controller.
- Give Back Immediately** This option is the equivalent of the `cf giveback -f` command in which the controller does not gracefully shut down the services of the taken over controller.

Once you have selected an option, the `Status` option on the Cluster Console page displays the status of the giveback operation. Once the giveback operation is complete, the Cluster Console page displays the updated controller-icon colors. The Cluster Console page also displays status of each controller. The Tools list of each controller is adjusted appropriately to indicate the active/active configuration operation each controller can now perform.

DataFabric Manager server CLI to configure storage systems

The DataFabric Manager server enables you to run storage system commands such as `sysconfig`, `version`, and `install`, on a specific storage system or a group of storage systems.

You can run all commands, except for a few administrator commands. For a list of unsupported commands, see the *Operations Manager Help*.

Remote configuration of a storage system

You can remotely configure a storage system using the DataFabric Manager server.

As you monitor your storage systems, you might find that you need to alter the configuration settings on one or more storage systems. The DataFabric Manager server provides three methods by which you can remotely configure your storage systems:

- Accessing the storage system CLI
- Accessing FilerView
- Using the DataFabric Manager server multiple-storage system remote configuration feature

You can remotely configure the following DataFabric Manager server features:

- Host users management
- User quota management
- Password management
- Roles management

Prerequisites for running remote CLI commands from Operations Manager

Your storage systems must meet certain prerequisites to run remote CLI from Operations Manager

The Command operation uses `rsh` to run a command on storage systems. Therefore, you must have `rsh` access to your storage system enabled to run CLI commands from the DataFabric Manager server. By default, `rsh` access to a storage system is enabled.

For more information about enabling `rsh` on your storage system, see the *Data ONTAP System Administration Guide*.

Note: The command operation uses only `ssh` to run remote CLI commands on clusters running Data ONTAP 8.0 Cluster-Mode.

Running commands on a specific storage system

You can run a command for a specific storage system using Operations Manager.

Steps

1. Click **Control Center > Home > Member Details > Physical Systems**
2. Select Storage Systems, All report.
3. Click the storage system to go to the **Storage Controller Details** page for the storage system or hosting storage system (of a vFiler unit) that you want to run a command on.
4. Click **Run a Command** under Storage Controller Tools.
5. Enter the command in the Appliance Command box.
6. Click **Run**.

Running commands on a group of storage systems from Operations Manager

By using Operations Manager, you can run a command on a group of storage systems.

Steps

1. In the left pane of the Operations Manager window, select the group that you want to run a command on.
The **Group Summary** page is displayed.
2. Select **Run a Command** from the **Storage Controller Tools** menu.
The **Run Command** page is displayed.
3. Enter the command in the Appliance Command box.
4. Click **Run**.

Remote configuration of a cluster

You can remotely configure a cluster using the DataFabric Manager server. As you monitor your clusters, you might find that you need to execute commands to alter the configuration settings on one or more nodes in the cluster.

In Operations Manager, you can configure the credentials for the cluster by using either of the following methods:

- Accessing the CLI

- Accessing the Edit Storage Controller Settings or Edit Cluster Settings page (**Details > Tools List > Edit Settings**) from Operations Manager

Running commands on a specific cluster

You can run commands on a specific cluster (Data ONTAP 8.0 Cluster-Mode) using Operations Manager. However, you cannot execute RLM commands on clusters running Data ONTAP 8.0 Cluster-Mode.

Before you begin

You must set the user name and password at the cluster level from the Edit Cluster Settings page or using the corresponding CLI.

Steps

1. Click **Control Center > Home > Member Details > Physical Systems > Report > Clusters, All**.
2. Click the cluster to go to the **Cluster Details** page.
3. Click the **Run a Command** link under Cluster Tools.
The Run Command page is displayed.
4. Enter the command in the Appliance Command box.
5. Click **Run**.

Running commands on a specific node of a cluster

You can run only RLM commands for a specific node of a cluster by using Operations Manager. The DataFabric Manager server uses the Remote Platform Management IP Address (address of the RLM card) for running these RLM commands.

Before you begin

You must set the credentials for a cluster and its node.

Note: You must set the user name to naroot and set a password at the node level of the cluster from the Edit Storage Controller Settings page or using the corresponding CLI.

You can use the `dfm run cmd -t` command with the timeout set to a value (e.g. 10 mins) on the CLI of the DataFabric Manager server, to execute run command jobs that are time consuming and result in large outputs (e.g. >1500 lines).

Steps

1. Click **Control Center > Home > Member Details > Physical Systems > Report > Clusters, All**.

2. Click the cluster to go to the **Cluster Details** page.
3. In the **Cluster Details** page, click the number corresponding to "Controllers."
4. In the "Controllers, All" report, click the name of the controller.
The Storage Controller Details page is displayed.
5. Click the **Run a Command** link under Storage Controller Tools.
The Run Command page is displayed.
6. Enter the command in the Remote Platform Management Command box.
7. Click **Run**.

Storage system management using FilerView

With Operations Manager, you can connect to a storage system using FilerView.

In addition to providing access to the storage system, Operations Manager enables you to log in to the FilerView management UI of the storage system. When you invoke FilerView, Operations Manager spawns a new window. By using FilerView, you can edit the configuration settings of a storage system.

Note: You cannot remotely configure more than one storage system using this method.

What FilerView is

Operations Manager enables you to view information about storage systems and vFiler units from a Web-based UI called FilerView.

Pages displaying information about storage systems and vFiler units provide access to the Web-based UI, FilerView. You can access FilerView by clicking the icon next to the storage system or vFiler unit name in the details pages for events, storage systems, vFiler units, aggregates, LUNs, qtrees, and volumes.

To access FilerView for a selected storage system or vFiler unit, click the storage system icon next to the storage system or vFiler unit name in the respective details page.

Configuring storage systems by using FilerView

You can configure a storage system by using FilerView.

Steps

1. On the **Storage Controller Details** page, click the FilerView icon.
2. When prompted, provide your user name and the password.
3. Edit the settings.

Introduction to MultiStore and vFiler units

MultiStore is a software product that enables you to partition storage and network resources of a single storage system so that it appears as multiple storage units on the network.

Each “virtual filer” created as a result of the logical partitioning of the hosting storage system’s network and storage resources is called a vFiler unit. A vFiler unit, using the resources assigned, delivers storage system services to its clients as a storage system does.

You can create multiple vFiler units using MultiStore.

The storage resource assigned to a vFiler unit can be one or more qtrees or volumes. The network resource assigned can be one or more base IP addresses or IP aliases associated with network interfaces.

A vFiler unit can participate in a distinct IP address space called the IPspace. IP addresses defined for an IPspace are meaningful only within that space. A distinct routing system is maintained for each IPspace; no cross-IPspace traffic is routed.

For information about configuring and using vFiler units in your storage network, see the *Data ONTAP MultiStore Management Guide*.

Related information

[Data ONTAP MultiStore Management Guide - www.ibm.com/storage/support/nseries](http://www.ibm.com/storage/support/nseries)

Why monitor vFiler units with the DataFabric Manager server

You can monitor vFiler units using Operations Manager.

The DataFabric Manager server provides Storage Service Providers (SSPs) the same management interface for monitoring vFiler units and hosting storage systems. Hosting storage systems are physical storage systems on which a vFiler unit is configured.

Prerequisites for monitoring vFiler units

Before you enable monitoring of vFiler units, you must ensure that the hosting storage system is running a supported Data ONTAP release and it is part of the same routable network as the DataFabric Manager server. NDMP discovery must also be enabled.

You must meet the following requirements before monitoring vFiler units:

- Supported Data ONTAP release
The MultiStore monitoring feature supports hosting storage systems running Data ONTAP 7.1 or later.
Note: To run a command on a vFiler unit using a Secure Socket Shell (SSH) connection, you must ensure that the hosting storage system is running Data ONTAP 7.2 or later.
- Network connectivity

To monitor a vFiler unit, DataFabric Manager server and the hosting storage system must be part of the same routable network that is not separated by firewalls.

- **Hosting storage system discovery and monitoring**
You must first discover and monitor the hosting storage system before discovering and monitoring the vFiler units.
- **NDMP discovery**
DataFabric Manager server uses NDMP as the discovery method to manage SnapVault and SnapMirror relationships between vFiler units. To use NDMP discovery, you must first enable SNMP and HTTPS discovery.
- **Monitoring the default vFiler unit**
When you enable your core license, which includes MultiStore, Data ONTAP automatically creates a default vFiler unit on the hosting storage system unit called vfiler0. Operations Manager does not provide vfiler0 details.
- **Monitoring backup relationships**
For hosting storage systems that are backing up data to a secondary system, you must ensure that the secondary system is added to the vFiler group. DataFabric Manager server collects details about vFiler unit backup relationships from the hosting storage system. You can then view the backup relationships if the secondary storage system is assigned to the vFiler group, even though the primary system is not assigned to the same group.
- **Monitoring SnapMirror relationships**
For hosting storage systems that are mirroring data to a secondary system, you must ensure that the secondary system is added to the vFiler group. DataFabric Manager server collects details about vFiler unit SnapMirror relationships from the hosting storage system. DataFabric Manager server displays the relationships if the destination vFiler unit is assigned to the vFiler group, even though the source vFiler unit is not assigned to the same group.

vFiler unit management tasks

You can perform management tasks on a vFiler unit using the DataFabric Manager server.

- Discover vFiler units of a hosting storage system.
- Group vFiler units for consolidated reporting.
- Monitor vFiler health and general status.
- Obtain vFiler network and storage resource details.
- Obtain vFiler performance and usage reports.
- Control vFiler administrative access.
- Monitor and manage SnapVault relationships.
- Monitor and manage SnapMirror relationships.
- Manage user quota.
- Manage host admin.
- Manage your configuration.
- Run commands on vFiler units.

You can run a command on a vFiler unit as a root user with privileges set to host filers or only the vFiler units. By default, if you run a command on a vFiler unit, the settings of the host filer will be used for vFiler communication. However, if you want to use the vFiler unit settings for vFiler

communication, you must create a role in the DataFabric Manager server with privileges assigned to only the vFiler unit and run the command.

Configuration of storage systems

You can remotely configure multiple storage systems using Operations Manager.

By creating configuration resource groups and applying configuration settings to them, administrators can remotely configure multiple storage systems from the server on which the DataFabric Manager server is installed. Administrators can also manage CIFS data through configuration management.

Management of storage system configuration files

Many administrators prefer to centrally manage their storage system and vFiler configuration /etc files and registry options. With Operations Manager, you can create and manage configuration files that contain configuration settings you want to apply to a storage system and vFiler unit or groups of storage systems and vFiler units.

By using storage system configuration management, you can pull configuration settings from a storage system and vFiler unit.

You can now and push all of the configuration settings or a some of the configuration settings to other storage systems, or groups of storage systems and vFiler units.

You can also ensure that storage system and vFiler configuration conforms with the configuration that is pushed to it from Operations Manager.

Prerequisites to apply configuration files to storage systems and vFiler units

You must meet a set of requirements before applying configuration files to a group of storage systems and vFiler units.

- Ensure that you are assigned the Global Write and Global Delete access roles to add or delete a configuration file from a group.
- Set the login and password for the storage system or vFiler unit before you set up configuration groups.
- Obtain a Data ONTAP plug-in for each version of the configuration file that you use in the DataFabric Manager server. You must have write privileges for a group to push configurations to it. You can download the plug-ins from the N series support website (accessed and navigated as described in [Websites](#) on page 23).

Related information

The IBM N series support site -- www.ibm.com/storage/support/nseries/

List of access roles to manage storage system configuration files

You need specific access roles to perform management tasks with the storage system configuration files.

Task	Access role
Create configuration files	Global Write
Delete configuration files	Global Delete
Edit configuration files	Global Write
Export configuration files	Global Read
Import configuration files	Global Write
Upgrade or revert configuration file versions	Global Write

List of tasks for configuration management

You can complete a set of configuration management tasks by using the storage system configuration management feature.

- Pull a configuration file from a storage system or a vFiler unit
- View the contents of each configuration file.
- Edit the configuration file settings (`registry` options and `/etc` files).
- Copy or rename configuration files.
- Edit a configuration file to create a partial configuration file.
- Compare configuration files against a standard template.
- View the list of existing configuration files.
- Upgrade or revert file versions.
- Delete a configuration file.
- Import and export configuration files.
- Remove an existing configuration file from a group's configuration list.
- Change the order of files in the configuration list.
- Specify configuration overrides for a storage system or a vFiler unit assigned to a group.
- Exclude configuration settings from being pushed to a storage system or a vFiler unit.
- View Groups configuration summary for a version of Data ONTAP.
- Push configuration files to a storage system or a group of storage systems, or to vFiler units or a group of vFiler units.
- Delete push configuration jobs.
- View the status of push configuration jobs.

Related concepts

What a configuration resource group is on page 232

What configuration files are

A configuration file is a set of configuration settings that you want the storage systems in one or more groups to share.

Configuration files exist independently of groups and can be shared between groups. Use Operations Manager to pull a configuration file from storage systems and save it.

What a configuration plug-in is

A configuration plug-in is an add-on library in a zip file that is required by the DataFabric Manager server, to manage Data ONTAP. For each Data ONTAP version, a Data ONTAP plug-in is provided.

Configuration plug-ins provide the capability to upgrade or revert a configuration file that is stored in the DataFabric Manager server database to a different version.

Comparison of configurations

Operations Manager enables you to compare your configuration file settings against those of a template configuration.

You can also compare storage systems, vFiler units or groups of storage systems, vFiler units against a configuration file, and create jobs to obtain comparison results. Use Operations Manager to access the comparison job results.

You can view the configuration comparison results in a report format. Use this report to identify the configuration settings that do not conform to those of the standard template.

Verification of a successful configuration push

After you have initiated a configuration push, you can review the status of the push operation for each storage system or vFiler unit to which you pushed a configuration.

When you push a configuration to a storage system or vFiler unit, the DataFabric Manager server logs the push operation on the storage system, or vFiler unit. The DataFabric Manager server logs this operation as a message that contains information about the DataFabric Manager server station, and the administrator who started the push job.

Configuring the DataFabric Manager server to monitor configuration changes

You can monitor configuration changes made to the storage systems using the DataFabric Manager server.

Before you begin

The credentials for the storage system that you want to monitor must be set.

Steps

1. Pull the configuration setting from the storage system and save that as a new configuration file.
2. Modify the configuration file and push it to the storage system.

When changes are made to the configuration settings, the DataFabric Manager server generates a System Configuration: Changed Locally event after the periodic configuration conformance monitor is run (unless the user explicitly triggers host discovery).

3. After the event is generated, you can select **Fix** in the **Event details** page to see the differences in the configuration setting.
4. You can select **Accept** to save the new changes in the configuration file or click **Reject** to retain the previous settings.

Result

The DataFabric Manager server starts monitoring the configuration changes in the storage system.

What a configuration resource group is

A configuration resource group is a group of storage systems that share a set of common configuration settings.

You can designate groups of managed storage systems that can be remotely configured to share the same configuration settings.

A configuration resource group must contain some number of storage systems and have one or more files containing the desired configuration settings. These configuration settings are listed in files called configuration files.

List of tasks for managing configuration groups

After you have added configuration files to a group, you can manage your configuration groups by completing a set of tasks.

- Remove an existing configuration file from a group's configuration list.
- Change the order of files in the configuration list.
- Specify configuration overrides for a storage system or a vFiler unit assigned to a group.

Note: The DataFabric Manager server attempts to contact the storage system or vFiler unit, five times (default), to complete the configuration push job. You cannot reconfigure the number of retries from Operations Manager, but you can use the following CLI command to specify a new retry limit: `dfm config push -R`.

- Exclude configuration settings from being pushed to a storage system or a vFiler unit.
- View Groups configuration summary for any version of Data ONTAP.

- Push configuration files to a storage system or a group of storage systems, or to vFiler units or a group of vFiler units.
- Delete push configuration jobs.
- View the status of a push configuration jobs.

Considerations when creating configuration groups

Before you create configuration groups, you must consider versions of operating systems and Data ONTAP, and storage systems belonging to configuration resource groups.

- Storage systems running different operating system versions can be grouped in the same configuration resource group.
- A storage system can belong to only one configuration resource group, but can also belong to other non-configuration resource groups.
- A storage system that is a member of a configuration resource group can also belong to one or more groups.

Creating configuration resource groups

You can create configuration resource groups by creating an empty group and populate the group with storage systems.

Steps

1. Create an empty group.
2. From the Groups pane, select the group you want to edit.
3. From the Current Group pane, select **Edit Membership**.
4. Populate the group from the available members.
5. From the Current Group pane, select **Edit Storage System Configuration** to add one or more configuration files to the group.

Result

After configuration files have been associated with the group, the following icon is attached to the group name so that you can identify the group as a configuration resource group:



Parent configuration resource groups

You can specify a parent configuration resource group from which one can acquire configuration settings.

Assigning a parent group enables you to quickly set the majority of the configuration settings of a storage system. You can then add any other configuration files that you might need to meet your

deployment requirements. When you assign a parent group, you inherit only the parent group's configuration files, not the storage systems in the parent group.

Parent group considerations

You should consider some points about the inheritance from the parent files and the parent hierarchy before you assign a parent group.

- When you assign a parent, you inherit only the parent group's configuration files. You do not inherit the storage systems in the member group.
- Parent groups can have parents of their own. The configuration settings of all parents are added to the beginning of the child's configuration settings. There is no limit to the potential length of these parent "chains."

Note: Ensure to review the settings in a parent group so that they do not have unintended consequences on your storage systems.

When to assign parent groups

You should assign a parent group if you want to control all or most of the configuration settings of a storage system from Operations Manager.

Remember that when you assign a parent group, you inherit all configuration settings in the parent group. Therefore, you should carefully scan a parent's configuration for any undesirable settings before assigning a parent.

You would probably not want to assign a parent if you want to use only a few of a parent group's settings. For example, if an existing group contains most of the access control list (ACL) rules you require, you cannot assign the group as a parent. You also cannot add more ACL rules in another configuration file.

Properties of configuration files acquired from a parent

One of the properties of configuration files acquired from parent groups is that they are initially read-only.

When you include configuration files from another group, consider the following points:

- A configuration resource group can include configuration files from only one parent group.
- Configuration files acquired from parent groups are always read first. You cannot change the order in which the acquired files are read unless you re-order the configuration files from within the parent group.

Configuring multiple storage systems or vFiler units

To configure storage systems or vFiler units, you must have enabled SNMP on the storage systems and the DataFabric Manager server must have already discovered them.

Steps

1. Pull a configuration file from a storage system or a vFiler unit.
2. Click (**Management > Storage System or vFiler > Configuration Files > Edit Configuration File**).
3. Edit the file settings.
4. Click **Compare Configuration Files** to compare your storage system or vFiler configuration file against a standard template configuration.
5. Create a configuration resource group by adding a configuration file.
6. If necessary, click (**Edit Storage System Configuration or Edit vFiler Configuration > Edit Configuration Pushed for Appliance**) to specify configuration overrides for a specific storage system or vFiler unit, or exclude configuration settings from being pushed to the storage system or vFiler units.
7. Click **Edit Storage System Configuration** or **Edit vFiler Configuration** and push the configuration file or files out to the storage systems or to the group.
8. Verify that the configuration changes have taken effect by reviewing the status of the push jobs.

Backup Manager

You can manage disk-based backups for your storage systems using Backup Manager.

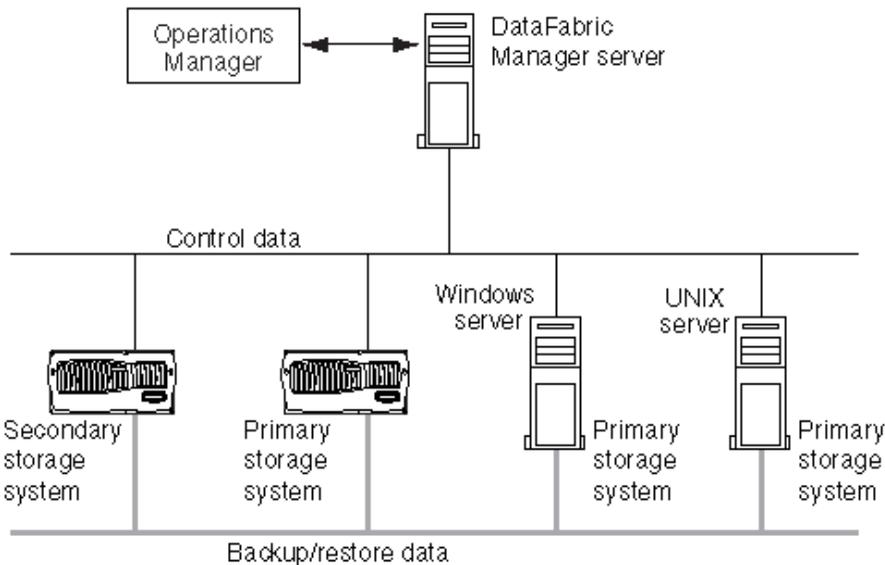
You can access it from the Backup tab in Operations Manager. Backup Manager provides tools for selecting data for backup, scheduling backup jobs, backing up data, and restoring data.

Note: Backup Manager does not support IPv6.

Backup management deployment scenario

The DataFabric Manager server uses the SnapVault technology of Data ONTAP to manage the backup and restore operations.

The following figure shows a sample configuration for backup management by using a DataFabric Manager server:



The configuration provides data protection between two storage systems and from a storage system to a UNIX or a Windows storage system.

System requirements for backup

You must meet a set of requirements for the DataFabric Manager server, primary storage system, and secondary storage system, for data backup.

The following table lists the requirements for each system:

System	Requirements
The DataFabric Manager server station	the N series Management Console data protection capability license
Primary storage system	<ul style="list-style-type: none"> • Data ONTAP 7.1 or later for storage systems and vFiler units • SnapVault primary license • SnapVault and NDMP enabled (configured using Data ONTAP commands or FilerView) • Open Systems SnapVault module for open systems platforms, such as UNIX, Linux, and Windows
Secondary storage system	<ul style="list-style-type: none"> • Data ONTAP 7.1 or later for storage systems • SnapVault secondary license • SnapVault and NDMP enabled (configured using Data ONTAP commands or FilerView) • Licenses for open systems platforms that are backing up data to secondary volumes on the secondary storage system

What backup scripts do

The prebackup and postbackup scripts help in bringing the databases into the hot backup mode before a backup is performed.

The DataFabric Manager server provides the ability to run prebackup and postbackup scripts on specific primary directories, before and after data has been backed up from those directories. For more information about the process of setting up such scripts to run on primary directories, see the DataFabric Manager server Backup man pages.

What the Backup Manager discovery process is

The Backup Manager performs three kinds of discovery process—storage system discovery, SnapVault relationship discovery, and New directories and qtree discoveries.

Methods of storage system discovery

Backup Manager provides two methods for discovering storage systems: SNMP and NDMP.

The DataFabric Manager server uses Simple Network Management Protocol (SNMP) to discover and monitor storage systems. The DataFabric Manager server uses Network Data Management Protocol (NDMP) primarily to communicate with primary and secondary storage systems.

When the DataFabric Manager server discovers a storage system, it adds the storage system to its database with the NDMP authentication credentials used for the connection. NDMP credentials are used to identify if the discovered storage system is a primary or a secondary storage system. If the DataFabric Manager server attempts to connect to an NDMP server, and the NDMP server rejects the authentication credentials, the DataFabric Manager server does not add the storage system to its database. Therefore, the DataFabric Manager server avoids spamming NDMP servers with “Login failed” errors.

When the DataFabric Manager server cannot authenticate a storage system with NDMP, it uses SNMP to discover primary and secondary storage systems. The DataFabric Manager server then adds the storage systems to its database without NDMP authentication credentials.

On authentication, Backup Manager communicates with the primary and secondary storage systems to perform backup and restore operations.

What SnapVault relationship discovery is

The DataFabric Manager server discovers and imports existing SnapVault relationships by using NDMP.

SnapVault relationship discovery is possible only if you have NDMP credentials for the primary and secondary storage systems.

New directories for backup

Backup administrators should know when new directories appear on primary storage systems so that they can schedule them for backup.

An Open Systems SnapVault monitor checks whether all Open Systems SnapVault hosts for which the DataFabric Manager server has valid NDMP authentication credentials are running. If the DataFabric Manager server discovers directories that are not backed up on Open Systems SnapVault hosts, it generates an Unprotected Item Discovered event.

Viewing directories that are not backed up

You can use either the CLI or the GUI to view the directories that are not backed up on Open Systems SnapVault hosts.

Steps

1. You can view the directories that are not backed up either by CLI or using the GUI.

Option	Description
By using the Command-Line Interface (CLI), execute the command...	<code>dfbm report primary-dirs-discovered</code>
By using Graphic User Interface (GUI), go to...	Directories not scheduled for backup view

- To disable the discovery of directories that are not backed up on Open Systems SnapVault hosts, execute the `dfbm primary dir ignore all` command at the CLI.

SnapVault services setup

Before you can use Backup Manager to back up your data using SnapVault relationships, you must prepare the primary storage system and the secondary storage system to use SnapVault.

The initial setup includes installing licenses, enabling SnapVault and NDMP services, and setting up the DataFabric Manager server access permissions.

To use SnapVault, you must have separate SnapVault licenses for both the primary and secondary storage systems—SnapVault primary license for primary storage system and SnapVault secondary license for secondary storage system. The Open Systems SnapVault agent does not require a license on the agent itself, but the secondary storage system requires Linux, UNIX, or Windows licenses.

Note: The setup procedure requires that you use the CLI of the storage system. If you want to use the Run a Command tool, you must first enable RSH on the storage system.

Configuring the SnapVault license

You must configure SnapVault license on your storage systems before you begin the backup operation.

Steps

- Enter `license add sv_primary_license`.
- Enter `license add sv_secondary_license`.
- Enter `options snapvault.enable on`.
- Enter `options snapvault.access host=snapvault_secondary` to name the secondary storage systems that you want to designate for backups.
- Enter `options snapvault.access host=snapvault_primary` to name the primary storage systems that you want to back up.

Enabling NDMP backups

You must enable the NDMP service and specify the DataFabric Manager server to enable NDMP backups on your storage systems.

Steps

1. Enter `ndmpd on` to enable NDMP service on each primary and secondary storage system.
2. Enter `options ndmpd.access host=dfm_server_host` to let the DataFabric Manager server perform backup and restore operations.

Management of SnapVault relationships

You can configure the DataFabric Manager server to manage SnapVault relationships.

The configuration involves tasks such as adding primary and secondary storage systems, adding secondary volumes, and selecting primary directories or qtrees for backup.

When the DataFabric Manager server discovers a primary or secondary storage system, Backup Manager lists the storage system in its backup reports. You can use either Operations Manager or the `dfbm` command to add a storage system.

Note: The DataFabric Manager server cannot discover Open systems platforms. Although you can add open systems platforms to the DataFabric Manager server database for backup management, you cannot manage these platforms with the DataFabric Manager server.

Adding secondary storage systems

You can add a secondary storage system to Backup Manager, from the Backup Summary page.

Before you begin

The SnapVault server feature must be licensed on the secondary storage system.

Steps

1. Click **Backup > Storage Systems**.
2. Click All Secondary Storage System from the **View** drop-down list.
3. In the **Secondary Storage Systems** page, enter the name (or IP address) of the secondary storage system.
4. Enter the NDMP user.
5. Obtain the NDMP password by entering the following command on the storage system:

```
ndmpd passwordusername
```

6. From the **Secondary Storage Systems** page, type the NDMP password.
7. Click **Add**.

Adding secondary volumes

You can add secondary volumes by selecting from a list of discovered volumes.

Before you begin

Ensure that you have added a secondary storage system to Backup Manager, so that the volumes of secondary storage system are automatically discovered by the DataFabric Manager server and are added to its database.

Steps

1. Click **Backup** > **Backup** and then the icon next to Secondary Volume.
2. From the **Secondary Volumes** page, select the secondary storage system.
3. Select a volume on the secondary storage system.

Note: If the DataFabric Manager server has not yet discovered a volume, you might need to click **Refresh**.

4. Click **Add**.

Adding primary storage systems

You can add a primary storage system to Backup Manager from the Primary Storage Systems page.

Steps

1. From the **Primary Storage Systems** page, enter the name (or IP address) of the primary storage system.
2. Enter the NDMP user name used to authenticate the primary storage system in the NDMP User field.
3. Obtain the NDMP password by entering the following command on the storage system:
ndmpd password username
4. If the primary storage system you are adding is an Open Platform system, you can configure a non-default value for the NDMP port on which the DataFabric Manager server will communicate with the system.
5. From the **Primary Storage Systems** page, enter the NDMP password.
6. Click **Add**.

Result

After you have added a primary storage system to Backup Manager, the DataFabric Manager server automatically discovers the primary directories, and adds them to its database. The DataFabric

Manager server also discovers primary qtrees, in the case of primary storage systems that support qtrees.

Selecting primary directories or qtrees for backup

You can schedule a primary directory or qtree for backup using the **Storage Systems** tab.

Before you begin

Before baseline transfers can start, you must add the primary directory to Backup Manager and configure its volumes or qtrees for backups to the secondary volume.

Steps

1. Click **Backup > Storage Systems**.
2. Select **Qtrees Not Scheduled For Backup** from the View drop-down list.
3. Select the qtree that you want to back up.
4. Click **Back Up**.

What backup schedules are

A backup schedule specifies how frequently data transfers are made from a primary directory or qtree to a secondary volume and how many Snapshot copies are retained on the secondary volume.

You must associate a backup schedule with a secondary volume before automatic backups can occur.

Note: Only one backup schedule can be associated with a secondary volume. Therefore, all backup relationships associated with a secondary volume must use the same backup schedule.

Best practices for creating backup relationships

Backups typically involve large amounts of data. Therefore, you might want to follow certain recommendations before creating backup relationships.

Following are the recommendations for creating backup relationships:

- Create backup relationships during off-peak hours so that any performance impact does not affect users.

Alternatively, you can specify a limit on the amount of bandwidth, a backup transfer can use.

Note: If a baseline backup transfer starts on a storage system when it is busy providing file services (NFS and CIFS), then the performance of file services is not impacted. These services are given higher priority than the backup. However, the backup takes longer to complete, because the storage system's resources are being consumed by services of higher priority.

- Avoid creating multiple backup relationships at the same time to avoid initiating multiple baseline transfers.

Snapshot copies and retention copies

If a scheduled or manual backup occurs, the DataFabric Manager server directs each primary storage system to create a Snapshot copy of its current data. Then, the DataFabric Manager server directs the secondary storage system to initiate a backup to its secondary volume, based on the Snapshot copy made on the primary storage system.

After a transfer to the secondary volume occurs, the DataFabric Manager server directs the secondary storage system to create a Snapshot copy of the entire secondary volume. This Snapshot copy is retained if a retention count other than zero is specified in the backup schedule. Otherwise, the copy is overwritten the next time a Snapshot copy of the secondary volume is created as a result of a new backup. If the number of Snapshot copies of the secondary volume being retained exceeds the number specified in the retention count, the oldest copy is purged.

Note: Unlike on supported primary storage systems, Snapshot copies of current data are not created on open systems platforms. Instead, entire changed files are transferred.

Requirements to create a backup schedule

Before you create a backup schedule, you must ensure that you have a name for the schedule, and have information about the weekly, nightly, and hourly schedules.

Creating backup schedules

To create a backup schedule, either use a template or customize an existing schedule. You can modify the schedule later and create a custom template.

Before you begin

When you create a backup schedule, you can identify it as the default backup schedule. Any secondary volumes subsequently added to Backup Manager are then automatically associated with this default backup schedule.

Steps

1. From the **Backup** page, click to open the **Backup Schedules** page.
2. Type a name for the schedule.
3. Create a schedule using one of the following methods:
 - Select a template.
 - Select **None** to create a schedule without using a template.
4. Click **Add** to add the schedule to the DataFabric Manager server database. The **Schedule Details** page is displayed.
5. Optionally, check the **Use as Default for New Secondary Volumes** check box to apply the schedule to all secondary volumes subsequently added to Backup Manager.

6. Optionally, enter the retention count for the hourly, weekly, and nightly schedule.
7. Optionally, click **Add a schedule** to open the **Edit Schedule** page and configure backup times for each hourly, weekly, and nightly schedule.
8. Click **Update**.

Local data protection with Snapshot copies

If you want to keep several Snapshot copies of your data on supported primary storage systems for local data protection, you must not rely on the Snapshot copies created for backup transfers with the DataFabric Manager server.

Use the Data ONTAP `snap sched` command or FilerView to provide local data protection on primary storage systems. However, you should turn off Snapshot copy scheduling configured with the `snapvault snap sched` command, to save resources.

You can also generate events to avoid conflict between Snapshot and SnapVault.

Snapshot copy schedule interaction

The backup schedules defined in the DataFabric Manager server do not affect the Snapshot copy schedules defined on the secondary storage systems.

Because, the backup schedules created in the DataFabric Manager server are independent of any Snapshot copy schedules that are defined on the secondary storage systems using the Data ONTAP `snapvault snap sched` and `snap sched` commands or FilerView.

Although all types of schedules can exist simultaneously, turn off all Snapshot copy creation and retention schedules configured with Data ONTAP commands on the secondary storage systems.

If you do not turn "off" the backup schedules that are defined for a relationship on a secondary storage system, backup transfers and retention of backups as defined by the commands continue to occur on the secondary storage system. Although such a situation does not lead to any data loss, it causes the primary and secondary storage systems to make unnecessary transfers, thus consuming resources on those storage systems. Hence, uses network bandwidth required for the backup transfers.

Management of discovered relationships

The DataFabric Manager server uses the following storage systems information to manage discovered relationships: system type, OS version, and NDMP credentials.

The DataFabric Manager server does not always have the basic information it needs to authenticate itself with discovered primary storage systems. Whenever you update the NDMP credentials of a storage system, the DataFabric Manager server contacts the storage system using NDMP to get this information. This also enables the DataFabric Manager server to verify the NDMP credentials.

Enabling the DataFabric Manager server to manage discovered relationships

You can enable the DataFabric Manager server to manage a discovered relationship by enabling NDMP, entering NDMP credentials, and associating a backup schedule.

Steps

1. Enable NDMP on the primary and secondary storage systems.
2. Enter the NDMP credentials for the primary and the secondary storage systems.
3. Associate a backup schedule with the secondary volume.

Note: Turn off all Snapshot copy schedules and policies defined for the imported backup relationship that were created using the Data ONTAP `snapvault snap sched` command.

What lag thresholds are

Lag thresholds are limits set on the time elapsed since the last successful backup. When those limits are exceeded, the DataFabric Manager server generates events of specific severity that indicates the acuteness of the event.

After you add a secondary volume to Backup Manager, the default values for lag thresholds are applied. However, you can change these lag thresholds, either for all volumes or for specific secondary volumes.

You can specify time in weeks, days, hours, minutes, or seconds. For example, to specify 20 hours and 15 minutes, enter 20.25 hours or 20:15:00 as the value for the threshold.

Setting global thresholds

The lag thresholds are applied by default to all secondary volumes in Backup Manager. To set a global option, you must complete a list of tasks.

Steps

1. Select **Setup > Options** and choose **Backup Default Thresholds** from the **Edit Options** menu at the left side of Operations Manager.
2. In the **SnapVault Replica Nearly Out of Date Threshold** field, enter the limit at which the backups on a secondary volume are considered nearly obsolete.
3. In the **SnapVault Replica Out of Date Threshold** field, enter the limit at which the backups on a secondary volume are considered obsolete.
4. Click **Update**.

Setting local thresholds

You must complete a list of tasks to change the lag thresholds for a specific secondary volume.

Steps

1. From any Backup Manager report, click the **secondary volume** name to access the **Secondary Volume Details** page.
2. In the **Lag WarningThreshold** field, specify the lag warning threshold limit in weeks, days, hours, minutes, or seconds.

Lag Warning threshold specifies the lag time after which the DataFabric Manager server generates the SnapVault Replica Nearly Out of Date event.

3. In the **Lag ErrorThreshold** field, specify the lag warning threshold limit in weeks, days, hours, minutes, or seconds.

Lag Error threshold specifies the lag time after which the DataFabric Manager server generates the SnapVault Replica Out of Date event.

4. Click **Update**.

Bandwidth limitation for backup transfers

You can specify a limit on the amount of bandwidth used when a backup transfer occurs. You specify this limit when you create a backup relationship or later when a need arises.

When you specify a limit for a backup relationship, the limit applies to all backup transfers—baseline and incremental—that occur for the relationship. However, if you do not want to apply a bandwidth limit permanently to a backup relationship, but still want to limit the amount of bandwidth that is used for the baseline transfer, you can apply the limit when you create a backup relationship. By doing this, the baseline transfer does not use more bandwidth than you specify. After a baseline transfer has occurred, you can remove the limit.

If you do not specify a limit for a backup relationship, the maximum available bandwidth for a transfer is used.

Note: The bandwidth limit applies only to the backup operations and not to the restore operations. For the restore operations, the maximum available bandwidth is always used.

Configuring backup bandwidth

You cannot specify a global bandwidth limit for backup relationships. You must specify a limit for each relationship individually.

Steps

1. Select the directory for which you want to configure a bandwidth limit by doing one of the following:

- For a new backup relationship, select the **Backup** tab to open the Backup page.
 - For an existing backup relationship, select a primary directory name from any view to open the Primary Directory Details page.
2. Enter the bandwidth limit that you want to impose on the backup transfers for this relationship.
 3. Click **Update**.

List of CLI commands to configure SnapVault backup relationships

To configure SnapVault backup relationships, you must execute a set of CLI commands.

CLI command	Description
dfbm backup list dfbm backup ls dfbm backup start	Initiates and browses backups for a secondary.
dfbm event list	Lists the backup events.
dfbm job abort dfbm job detail dfbm job list dfbm job purge	Manages backup jobs.
dfbm ndmp add dfbm ndmp delete dfbm ndmp modify dfbm ndmp list	Manages the list of user names and passwords used for Network Data Management Protocol (NDMP) discovery.
dfbm option list dfbm option set	Manages the global backup options that control the operation of Backup Manager.

CLI command	Description
<pre>dfbm primary dir add dfbm primary dir delete dfbm primary dir discovered dfbm primary dir ignore dfbm primary dir list dfbm primary dir modify dfbm primary dir relinquish dfbm primary dir unignore dfbm primary host add dfbm primary host delete dfbm primary host list dfbm primary host modify</pre>	<p>Manages primary directories that the DataFabric Manager server discovers.</p>
<pre>dfbm reports events dfbm reports events-error dfbm reports events-unack dfbm reports events-warning</pre>	<p>Runs reports on backup events.</p>
<pre>dfbm reports jobs dfbm reports jobs-1d dfbm reports jobs-30d dfbm reports jobs-7d dfbm reports jobs-aborted dfbm reports jobs-aborting dfbm reports jobs-completed dfbm reports jobs-failed dfbm reports jobs-running</pre>	<p>Runs reports on backup jobs.</p>
<pre>dfbm reports backups-by-primary dfbm reports backups-bysecondary</pre>	<p>Runs backup reports based on primary or secondary relationships.</p>
<pre>dfbm schedule add dfbm schedule create dfbm schedule delete dfbm schedule destroy dfbm schedule diag dfbm schedule modify</pre>	<p>Manages backup schedules.</p>

CLI command	Description
dfbm secondary host add	Manages hosts and volumes used as backup destinations.
dfbm secondary host delete	
dfbm secondary host list	
dfbm secondary host modify	
dfbm secondary volume add	
dfbm secondary volume delete	
dfbm secondary volume list	
dfbm secondary volume modify	

Primary directory format

If you want to add a primary directory to Backup Manager without browsing through Operations Manager, use the following format: `system_name:{drive_letter | volume_name}{path_name}`.

Note: The parameters are case-sensitive for UNIX systems, but not for Windows.

Supported storage systems

For a primary directory called `engineering/projects` in volume `vol1` of a storage system named `jupiter`, enter the following text:

```
jupiter:/vol1/engineering/projects
```

Windows system

For a primary directory called `engineering\projects` on the D drive of a Windows system named `mars`, enter the following text:

```
mars:D:\engineering\projects
```

Your capitalization could be different.

UNIX system

For a primary directory `/usr/local/share` on a UNIX system named `mercury`, enter the following text:

```
mercury:/usr/local/share
```

Secondary volume format

If you need to add a secondary volume to Backup Manager, use the following format:

`system_name:volume_name` For example, `pluto:/vol1`.

Disaster Recovery Manager

Disaster Recovery Manager is an application within the DataFabric Manager server that enables you to manage and monitor multiple SnapMirror relationships from a single interface.

Note: Disaster Recovery Manager does not support IPv6.

A SnapMirror relationship is the replication relationship between a source storage system or a vFiler unit and a destination storage system or a vFiler unit by using the SnapMirror feature. Disaster Recovery Manager provides a simple, Web-based method of monitoring and managing SnapMirror relationships between volumes and qtrees on your supported storage systems and vFiler units. You can view and manage all SnapMirror relationships through the Disaster Recovery tab of Operations Manager. You can also configure SnapMirror thresholds, so that Disaster Recovery Manager generates an event and notifies the designated recipients of the event.

For more information about SnapMirror, see the *Data ONTAP Data Protection Online Backup and Recovery Guide*.

Related information

[Data ONTAP Data Protection Online Backup and Recovery Guide - www.ibm.com/storage/support/nseries](http://www.ibm.com/storage/support/nseries)

Prerequisites for using Disaster Recovery Manager

Before you use Disaster Recovery Manager to monitor SnapMirror relationships, ensure that certain prerequisites are met.

- The Business Continuity Management license key must be installed.
 - Note:** The SnapMirror monitoring and management features are available only with the Business Continuity Management license. If you do not have this license, contact your sales representative.
- The SnapMirror destination storage systems must be running Data ONTAP 7.1.
 - Note:** Disaster Recovery Manager can discover and monitor only volume and qtree SnapMirror relationships in which the SnapMirror destinations have Data ONTAP 7.1 installed.
- The source and destination storage systems must have Data ONTAP 7.1 installed to perform any of the SnapMirror management tasks.
- The source and destination storage systems configured with vFiler units must be running Data ONTAP 7.1 to perform any of the SnapMirror relationship management and monitoring tasks.

Tasks performed by using Disaster Recovery Manager

You can manage policies, connections, and SnapMirror relationships; and authenticate storage systems using the Disaster Recovery Manager.

The Disaster Recovery Manager Home Page is the default page that appears when you click the **Disaster Recovery** tab. The page lists all the existing SnapMirror relationships.

What a policy is

A policy is a collection of configuration settings that you can apply to one or more SnapMirror relationships. The ability to apply a policy to more than one SnapMirror relationship makes a policy useful when managing many SnapMirror relationships.

There are two types of policies that you can create and apply to SnapMirror relationships:

- Replication
- Failover

What a replication policy does

A replication policy affects the way in which a source storage system replicates data to a destination storage system or a vFiler unit.

SnapMirror replication can occur asynchronously or synchronously; therefore, based on the type of replication, there are two policies:

- Asynchronous replication policy
- Synchronous replication policy

List of parameters for an asynchronous replication policy

You must set a list of parameters for an asynchronous replication policy.

Schedule	Specifies when an automatic update occurs. You can specify the schedule using Operations Manager or you can enter the schedule in the cron format. Note: For more information about scheduling using the cron format, see the <code>na_snapmirror.conf(5)</code> man page for Data ONTAP.
Maximum Transfer Speed	Specifies the maximum transfer speed, in kilobytes per second.
Restart	Specifies the restart mode that SnapMirror uses to continue an incremental transfer from a checkpoint if it is interrupted.

Lag Warning Threshold	Specifies the limit at which the SnapMirror destination contents are considered nearly obsolete. If this limit is exceeded, Disaster Recovery Manager generates a SnapMirror Nearly Out of Date event.
Lag Error Threshold	Specifies the limit at which the SnapMirror destination contents are considered obsolete. If this limit is exceeded, Disaster Recovery Manager generates a SnapMirror Out of Date event.
TCP Window Size	Specifies, in bytes, the amount of data that a source can send on a connection before it requires acknowledgment from the destination that the data was received.
Checksum	Specifies the use of a Cyclic Redundancy Check (CRC) checksum algorithm. Use the checksum option if the error rate of your network is high enough to cause an undetected error.

List of parameters for a synchronous replication policy

You must set a list of parameters for a synchronous replication policy.

Fully synchronous	Specifies full synchronicity between the source and the destination.
Semi-synchronous	Specifies the level of synchronicity between the source and the destination. The destination can be behind the source by 0 to 60 seconds or by 0 to 500 write operations.
Visibility interval	Specifies a time interval after which the transferred data becomes visible on the destination.
TCP Window Size	Specifies the amount of data that a source can send on a connection before it requires acknowledgment from the destination that the data was received.
Checksum	Specifies the use of a Cyclic Redundancy Check (CRC) checksum algorithm. Use the checksum option if the error rate of your network is high enough to cause an undetected error.

What a failover policy does

A failover policy affects the process used by a SnapMirror relationship to recover from a disaster.

Failover policies consist of a path to a user-installed script that would be called when a disaster occurs before and after the following events:

- SnapMirror break
- SnapMirror resynchronization

Policy management tasks

This table describes policy management tasks and provides the location of the user-interface page that enables you to complete the task.

If you want to...	Click...
Create a new replication policy	Disaster Recovery > Add a Mirror > Manage Replication Policies icon) OR Disaster Recovery > Mirrors-Mirrored > SnapMirror Relationship Source > Manage Replication Policies icon
Edit an existing replication policy	Disaster Recovery > Add a Mirror > Edit Replication Policy icon OR Disaster Recovery > Mirrors-Mirrored > SnapMirror Relationship Source > Policy name > Manage Replication Policies icon
Delete an existing replication policy Note: Remove the policy from SnapMirror relationships before deleting it; otherwise, Disaster Recovery Manager sends an error message.	Disaster Recovery > Add a Mirror > Manage Replication Policies icon OR Disaster Recovery > Mirrors-Mirrored > SnapMirror Relationship Source > Manage Replication Policies icon
Create a new failover policy	Disaster Recovery > Add a Mirror > Manage Failover Policies icon OR Disaster Recovery > Mirrors-Mirrored > SnapMirror Relationship Source > Manage Failover Policies icon
Edit an existing failover policy	Disaster Recovery > Add a Mirror > Edit Selected Failover Policy icon OR Disaster Recovery > Mirrors-Mirrored > SnapMirror Relationship Source > Policy name > Manage Failover Policies icon

If you want to...	Click...
Delete an existing failover policy Note: Remove the policy from SnapMirror relationships before deleting it; otherwise, Disaster Recovery Manager sends an error message.	Disaster Recovery tab > Add a Mirror link > Manage Failover Policies icon OR Disaster Recovery > Mirrors-Mirrored > SnapMirror Relationship Source > Manage Failover Policies icon

Connection management

You can specify one or two specific network paths between a source storage system or a vFiler unit, and a destination storage system or a vFiler unit using connection management.

The advantages of multiple paths between source and destination storage systems or vFiler units are as follows:

- Increased transfer bandwidth
- Networking failover capability

For more information, see the *Data ONTAP Data Protection Online Backup and Recovery Guide*.

Related information

Data ONTAP Data Protection Online Backup and Recovery Guide - www.ibm.com/storage/support/nseries

Connection management tasks

The table describes connection management tasks and provides the location of the user-interface page that enables you to complete the task.

If you want to...	Go here...
Create a connection	Connections page (Connections Disaster Recovery tab > Add a Mirror link > Manage Connections icon) OR (Disaster Recovery tab > Mirrors-Mirrored link > SnapMirror Relationship Source > Manage Connections icon)

If you want to...	Go here...
Edit a connection	Edit Connections page (Disaster Recovery tab > Add a Mirror link > Manage Connections icon > View drop-down list, connection name) OR (Disaster Recovery tab > Mirrors-Mirrored link > View drop-down list, connection name)
Delete a connection	Connections page (Disaster Recovery tab > Mirrors-Mirrored link > SnapMirror Relationship Source > Manage Connections icon)

What the connection describes

A connection specifies the parameters for one or two network paths between the source and the destination storage system or a vFiler unit.

The parameters that are specified by the connection are as follows:

Connection name Name of the connection.

Connection mode Defines the mode, that the paths use.

IP address pairs IP addresses of the source and destination storage systems and vFiler units that define the path that are used by the SnapMirror relationship. You must define one path, and you can define up to two paths.

What multipath connections are

Synchronous SnapMirror supports up to two paths for a particular SnapMirror relationship. The paths can be Ethernet, Fibre Channel, or a combination of Ethernet and Fibre Channel.

You can set the two paths to use one of the following two modes:

- **Multiplexing mode**—SnapMirror uses both paths at the same time, essentially load balancing the transfers. If one path fails, the transfers occur on the other path. After the failed path is repaired, the transfers resume using both paths.
- **Failover mode**—SnapMirror uses the first specified path as the desired path and uses the second specified path only if the first path fails.

Authentication of storage systems

The DataFabric Manager server uses the NDMP protocol to manage SnapMirror relationships, and therefore authentication of the storage systems is necessary.

Before you can perform any SnapMirror management task, the DataFabric Manager server must know the NDMP credentials for the source and destination storage systems.

Authentication of discovered and unmanaged storage systems

If you try to perform management tasks on a SnapMirror relationship in which a storage system is unauthenticated, Disaster Recovery Manager redirects you to the authentication page where you need to enter the NDMP credentials (user name and password).

Disaster Recovery Manager discovers storage systems enabled with SnapMirror on your network.

Disaster Recovery Manager might not identify the NDMP credentials for storage systems on your network. Storage systems for which the Disaster Recovery Manager fails to recognize the NDMP credentials are listed on the Authenticate Storage Systems page (**Disaster Recovery tab > Authenticate Storage systems link**).

Note: Disaster Recovery Manager can manage only storage systems running on Data ONTAP 7.1 or later.

Addition of a storage system

Before you manage the SnapMirror relationship using Disaster Recovery Manager, you must add the NDMP credentials.

When creating SnapMirror relationships by using storage systems that have not been SnapMirror enabled before, you must add the NDMP credentials. You can add the NDMP credentials on the Storage systems page when you create the new SnapMirror relationship (**Disaster Recovery tab > Add a Mirror link > Manage SnapMirror Hosts icon**).

Modification of NDMP credentials

You can edit the NDMP credentials of a storage system, for a selected SnapMirror relationship from the Edit a storage system page.

You might have to edit the NDMP credentials for selected SnapMirror relationships.

Change administrators and passwords as a security precaution

You can edit the NDMP credentials for a selected SnapMirror relationship from the Edit a storage system page (**Disaster Recovery tab > Mirrors-Mirrored link > SnapMirror Relationship source > Storage system**).

Note: The NDMP credentials are shared with the Backup Manager.

Deletion of a storage system

You can delete one or more storage systems from the managed storage system list if you do not use them in a SnapMirror relationship.

You must not delete a storage system if you are changing its NDMP password. Instead, you must change the NDMP password by editing the storage system.

Volume or qtree SnapMirror relationships

To create a new SnapMirror relationship, that is either a volume replication or a qtree replication, you must create the volume (restricted or not restricted) on the destination storage system.

If the new SnapMirror relationship is a volume replication, you must create the volume on the destination storage system and mark the volume as restricted before you can create the SnapMirror relationship. The FilerView link provides a shortcut to do this.

Note: Monitor must be running for you to be able to access the FilerView interface for a storage system.

If the new SnapMirror relationship is a qtree replication, ensure that the volume on the destination storage system where you want to replicate a qtree with SnapMirror is online and not restricted. Do not manually create a destination qtree.

On upgrading to the DataFabric Manager server 3.3, in a qtree SnapMirror relationship, you can select a qtree directly belonging to the vFiler unit by selecting the volume belonging to the storage system.

Decisions to make before adding a new SnapMirror relationship

You can add a new SnapMirror relationship by specifying the Relationship type, Source, Destination, Volume, Connection, and Policies.

Relationship type	If the new SnapMirror relationship is a volume replication, you must create the volume on the destination storage system and mark the volume as restricted before you can create the SnapMirror relationship. If the new SnapMirror relationship is a qtree replication, ensure that the volume on the destination storage system where you want to replicate a qtree with SnapMirror is online and not restricted.
Connection	Specifying a defined connection is optional. If you do not specify a defined connection, the default network route is used.
Policies	Replication and failover policies are optional. If you do not specify a replication policy, only a baseline transfer is performed. You can add a replication policy later to schedule when incremental transfers occur. If you do not specify a failover policy, no user scripts are called during a SnapMirror break or resynchronization.

Related concepts

[Connection management](#) on page 257

[What a policy is](#) on page 254

Addition of a new SnapMirror relationship

You can use the SnapMirror Relationships page to create a new SnapMirror relationship.

Adding a new SnapMirror relationship involves the following:

- By selecting the type of SnapMirror relationship, volume or qtree
- By selecting the source storage system and source volume or qtree

Note: If the SnapMirror relationship is a qtree replication, you select the volume and the qtree.

- Selecting the destination storage system and destination volume or qtree
- Selecting the connection
- Selecting the type of replication and failover policy

Modification of an existing SnapMirror relationship

You can edit an existing SnapMirror relationship from the Edit SnapMirror Relationship page.

To edit an existing SnapMirror relationship, use the Edit SnapMirror Relationship page.

Editing an existing SnapMirror relationship involves the following tasks:

- Creating or editing replication policies and failover policies
- Assigning a connection for a relationship

Related concepts

[Connection management](#) on page 257

[What a policy is](#) on page 254

Modification of the source of a SnapMirror relationship

You can change the source of a SnapMirror relationship from the Edit SnapMirror Relationship page.

To change the source of a SnapMirror relationship, use the Edit SnapMirror Relationship page (**Disaster Recovery tab > Home > SnapMirror relationship source**).

Reason to manually update a SnapMirror relationship

You can update an existing SnapMirror relationship from the Edit SnapMirror Relationship page.

Use the Edit SnapMirror Relationship page (**Disaster Recovery tab > Home > SnapMirror relationship source**) to update a SnapMirror relationship in between scheduled incremental updates.

By updating a SnapMirror relationship manually is useful because you might need to run an unscheduled update to prevent data loss. Data loss occurs due to scheduled or threatened power

outage or from a destination volume being taken offline for maintenance, repair, upgrade, or data migration.

Termination of a SnapMirror transfer

You can abort a SnapMirror transfer from the Edit SnapMirror Relationship page.

The **Abort** button on the Edit SnapMirror Relationship page is available only when SnapMirror transfers are in progress.

SnapMirror relationship quiescence

You can quiesce a SnapMirror relationship from the Edit SnapMirror Relationship page.

Use the Edit SnapMirror Relationship page (**Disaster Recovery tab > Home > SnapMirror relationship source**) to quiesce a SnapMirror relationship.

You can quiesce a SnapMirror relationship to block updates to the destination storage system after existing volume or qtree updates are complete. If a qtree is not in a stable state (is in transition), quiescing the SnapMirror relationship forces it into a stable state.

You can quiesce only volumes and qtrees that are online and that are SnapMirror destinations. You cannot quiesce a restricted or offline volume or a qtree in a restricted or offline volume.

View of quiesced SnapMirror relationships

You can use the Disaster Recovery Home page or the Quiesced SnapMirror Relationships page to view all quiesced SnapMirror relationships.

Resumption of a SnapMirror relationship

You can resume a SnapMirror relationship from the Edit SnapMirror Relationship page or Quiesced SnapMirror Relationships page.

Disruption of a SnapMirror relationship

You can break a quiesced SnapMirror relationship from the Edit SnapMirror Relationship page or the Quiesced SnapMirror Relationships page.

You can break SnapMirror relationships if you want to temporarily end a SnapMirror relationship between a source and a destination volume or qtree. When you break a relationship, the source from a destination volume or qtree releases, allowing the source to delete its base Snapshot copy for the SnapMirror relationship.

View of a broken SnapMirror relationship

You can view all the broken SnapMirror relationships from the Disaster Recovery Home page or the Broken SnapMirror Relationships page.

Resynchronization of a broken SnapMirror relationship

You can use the Edit SnapMirror Relationship page or the Broken SnapMirror Relationships page to perform resynchronization tasks related to SnapMirror relationships.

You can resynchronize a source and a destination volume or qtree in one of the following ways:

- You can resynchronize the SnapMirror relationship that you broke.
- You can reverse the functions of the source volume and the destination volume when you resynchronize.

Deletion of a broken SnapMirror relationship

You can delete a broken SnapMirror relationship from the Edit SnapMirror Relationship page or the Broken SnapMirror Relationships page.

Note: The only way to restore a deleted SnapMirror relationship is to initialize the relationship.

What lag thresholds for SnapMirror are

After Operations Manager is installed, the lag thresholds of all the SnapMirror source volumes and destination volumes are set to default values; however, you might want to change these thresholds.

You can change the threshold values of the following:

- All SnapMirror sources and destinations in the Disaster Recovery Manager database
- All SnapMirror sources and destinations of a specific group
- All mirrors of a SnapMirror source
- A specific SnapMirror source or destination

For more information on the thresholds and the default values, see *Operations Manager Help*.

Disaster Recovery Manager automatically generates SnapMirror events based on these thresholds. If you want to receive notifications in the form of e-mail messages, pager alerts, or SNMP traps when a SnapMirror event occurs, you can set up alarms.

Note: Disaster Recovery Manager does not generate events when lag thresholds of SnapMirror sources are crossed. Only lag thresholds of SnapMirror destinations are used for generating events.

Where to change the lag thresholds

You can change the lag thresholds from the Edit Policy page (Disaster Recovery tab > Home link > SnapMirror relationship Source > Manage Replication Policies icon > Policy name).

Lag thresholds you can change

You can change SnapMirror Lag Warning Threshold and SnapMirror Lag Error Threshold.

You can change the following lag thresholds:

- **SnapMirror Lag Warning Threshold**
This option specifies the limit at which the SnapMirror destination contents are considered nearly obsolete. If this limit is exceeded, Disaster Recovery Manager generates a SnapMirror Nearly Out of Date event.
You can specify the lag warning threshold limit in weeks, days, hours, minutes, or seconds. For example, to specify 20 hours and 15 minutes, enter 20.25 hours or 20:15:00 as the value for this option.
- **SnapMirror Lag Error Threshold**
This option specifies the limit at which the SnapMirror destination contents are considered obsolete. If this limit is exceeded, Disaster Recovery Manager generates a SnapMirror Out of Date event.
You can specify the lag error threshold limit in weeks, days, hours, minutes, or seconds. For example, to specify 20 hours and 15 minutes, enter 20.25 hours or 20:15:00 as the value for this option.

Reasons for changing the lag thresholds

You can change the default values of the lag thresholds to a lower value so that you are notified at an earlier time than the specified time.

If you use your SnapMirror destination to distribute data to remote sites, you must keep the latest data at the destination. After you have done this, you must set the SnapMirror schedule to transfer data from the source to the destination frequently. In such a case, Disaster Recovery Manager can generate an event sooner than the default lag time so that you can take corrective action sooner.

What the job status report is

The Job Status report displays the status of SnapMirror jobs along with information such as the time at which the jobs were started and their job IDs. The Job Status report is identical in appearance and function to the Job Status report for Backup Manager.

Maintenance and management

You can configure and maintain the DataFabric Manager server through the CLI.

Accessing the CLI

You can access the CLI using telnet or the console of the system on which the DataFabric Manager server is installed.

Step

1. Access the CLI on local or remote systems.

If you want to access the CLI on a...	Then...
Local Windows system	Use the command prompt window of the Windows workstation through Start > Run .
Local Linux system	Use any shell prompt on your workstation.
Remote Windows or Linux system	<ol style="list-style-type: none"> a. Make a Telnet connection from the remote host to the DataFabric Manager server you want to access by entering the following command: <pre>telnet hostname</pre> <p><i>hostname</i> is the host name or IP address of the system running the DataFabric Manager server.</p> <p>When connected to a terminal server, use the host name or IP address, and the port number of the terminal server, to access the console of the workstation:</p> <pre>telnet {term_server} : port_number</pre> b. Initiate authentication with your user name and password.

Where to find information about the DataFabric Manager server commands

There are two ways to find information about the DataFabric Manager server commands, by accessing the help and by using the `dfm help` command.

- Access the man pages through the table of contents of the *Operations Manager Help*. On a Linux system, you can access the man pages running the command `source /opt/IBMdfm/bin/vars.sh`, and then `man dfm`.
- Use the `dfm help` command, which has the following syntax: `dfm help command`.

Audit logging

Audit logging is the process of logging every activity performed by the DataFabric Manager server for a later review. The DataFabric Manager server logs all the activities in the `audit_log` file.

System administrators view the `audit_log` file for following reasons:

- Determine the recently changed configurations to understand why a problem is occurring.
- Determine when a specified change in the configuration of the system was made.
- Determine who made the specified change in the configuration of the system and ask them why the change was made.
- Identify attempts to subvert the security of the system.

The `audit_log` file resides in the default log directory of the DataFabric Manager server.

Events audited in the DataFabric Manager server

The DataFabric Manager server logs events and each of these events are recorded.

The events that can be audited in the DataFabric Manager server and the information that is recorded about each event are described here:

Authentication events	The DataFabric Manager server logs each authentication action that succeeded or failed. The user name associated with the authentication attempt is also recorded.
Authorization events	The DataFabric Manager server logs each authorization failure and the user name associated with it.
Command execution	The DataFabric Manager server logs the execution of each command. The complete command line (including options and arguments) is recorded in the <code>audit_log</code> file. The DataFabric Manager server also logs the name of the user who executed the command; the failure status of the command, if any; and the type of request: Web or CLI.
API calls	The DataFabric Manager server logs the invocation of any API by using the DataFabric Manager server service. The complete details of the API call and the authenticated user's name, on whose behalf, the API was invoked, are recorded in the <code>audit_log</code> file.
Scheduled actions	When the scheduler starts a job by invoking a CLI, the DataFabric Manager server logs the scheduled action and the user affiliated with it in the <code>audit_log</code> file.

In addition, a timestamp is recorded for each event. In the case of APIs, the IP address of the appliance from which the requests are received is logged. In the case of CLI requests, the IP address is always that of the DataFabric Manager server.

Global options for audit log files and their values

A global option `auditLogForever` is used to keep the audit log files forever. The valid values for this option are `yes` and `no` and the default value is `noauditLogEnabled`.

Note: When you set the `auditLogForever` global option to `yes`, the number of audit log files (each 3 MB in size) can grow excessively. You have to ensure that you have enough space on the DataFabric Manager server to keep the audit log files forever.

For audit logging, the `dfm option list` command requires global read capability, and the `dfm option set` command requires global write capability.

You must have Core Control Capability to modify `auditLogEnable` and `auditLogForever` global options. License features required: The `dfm option set` command requires an Operations Manager license.

Format of events in an audit log file

The format of events in the `audit log` file is as follows: `<timestamp> [<application-name>:<priority>]:<username>:<protocol> <label>:[ip-address]:<intent>:<message>`.

Example of events in the `audit log` file:

```
Apr 11 00:04:19 [dfm:NOTIC]:root:LOG:action:::Added new
administrator "testu1":
```

```
Apr 20 11:57:27 [dfm:NOTIC]:root:API:in:[10.72.1.2]:Listing database
```

```
Apr 23 14:56:40 [dfm:NOTIC]:root:WEB:in:[ABCD:EF01:2345:6789:ABCD:EF01:2345:6789]:dfm
report view
```

backups:

```
<ss><dfm-backup-directory>opt
```

```
<dfm-backup-directory><dfm-backup-verify-settings>
```

The `<application-name>` value denotes the application invoking the audit log facility. For example, it can be `dfm/dfbm/dfdrm` if a CLI or Operations Manager request is being audit logged from the `dfm/dfbm/dfdrm` executable.

In the case of APIs, `<application-name>` is the actual name of the application that called the API (for example, `dfm` and `sdu`). If the API was called by an external application other than the DataFabric Manager server and it does not pass the name of the application in the API input parameters, the field `::` is not printed in the log message.

The message priority field `<priority>` can have one of the following values:

- EMERG: unusable system
- ALERT: action required
- CRIT: critical conditions

- ERROR: error conditions
- WARN: warning conditions
- NOTIC: normal, but significant condition
- INFO: informational messages
- DEBUG: debug-level messages

The <username> field logs names of the users who invoked the CLIs and APIs.

The <protocol> field describes the source of the event being logged. The protocol label can have one of the following values:

- API: In the case of an API invocation
- CMD: In the case of CLI commands
- LOG: When an event is explicitly logged by the system
- WEB: In the case of an Operations Manager request

The <label> field describes the type of the event being logged. The message label can have one of the following values:

- IN: for input
- OUT: for output (for example, in the case of API calls)
- ERR: for error (for example, in the case of log in failures)
- ACTION: for actions initiated by the user

The <ip-address> value for APIs is the IP address of the system from which the API is invoked. In the case of the CLI, it is the IP address of the DataFabric Manager server. For requests coming through Operations Manager, it is the IP address of the workstation on which Operations Manager is installed.

The <intent> field describes the following information:

- If the protocol is API, it conveys the intention of the API.
- If the protocol is CMD, it is the actual command used.
- If the protocol is WEB, it is the URL of the Web page.
- If audit-log API is called, this field remains blank.

The <message> field content depends on the value of <protocol>, as follows:

- If <protocol> is API, the XML input to the API is logged, excluding the <IBM> element.
- If <protocol> is CMD, it contains output or an error message.
- If <protocol> is WEB, it is empty.

Following is an example:

```
July 04 22:11:59 [dfm:NOTIC]:IBM\tom:CMD:in:127.0.0.1:dfm host
password set -p ***** jameel:Started job:2
July 06 13:27:15 [dfm:NOTIC]:IBM\tom:WEB:in:127.0.0.1:dfm user
```

```
login username = tom password=*****:Logged in as<B>IBM\tom
<\B><BR>
```

```
July 06 14:42:55[dfm:NOTIC]:IBM\tom:API:in:127.0.0.1:Add a role
to a user:<rbac-admin-role-add> <role-name-or-id>4</role-name-or-id>
<admin-name-or-id>TOM-XP\dfmuser</admin-name-or-id></rbacadmin-
role-add>
```

Permissions for accessing the audit log file

For accessing the `audit log` file, Windows and UNIX users have separate permissions.

- Linux: `-rw - - root root` “root” users on Linux have both read and write permissions.
- Windows users in the Administrators group have both read and write permissions.

What the remote platform management interface is

The remote platform management interface enables you to remotely perform the routine system maintenance tasks such as resetting a system with backup firmware.

Following maintenance tasks can be performed using the interface on the DataFabric Manager server:

- Control system power, such as powering on or off the storage system
- Reset system firmware
- View system and event logs
- Obtain system sensor status

You can use the Run a Command tool in Operations Manager or use the `dfm run cmd -r` command on the CLI of the DataFabric Manager server to execute the remote platform management commands.

By using the interface, you can access the Remote LAN Module (RLM) cards on the N5000 series; N6000 series; and N7000 series storage systems.

RLM card monitoring in the DataFabric Manager server

The DataFabric Manager server monitors the RLM card installed on your storage system and obtains its status by issuing an ICMP ECHO request. When the card responds to the request, Operations Manager displays the card status as Online in the Appliance Details page.

By default, the DataFabric Manager server pings the card every 15 minutes. You can change this monitoring interval using the Operations Manager Tools menu.

Prerequisites for using the remote platform management interface

Before performing remote maintenance, by using the remote platform management interface, you must configure the RLM card IP address on the storage system.

After you have configured the card's IP address, you must set the following parameters in Edit Appliance Settings for the destination storage system:

Remote Platform Management IP address	This is the IP address of the RLM card on the storage system.
Appliance login	This is the login user name.
Appliance password	This is the login password.

For procedures to configure an RLM card IP address, see the *Operations Manager Help*.

Scripts overview

The installation, management, and execution of your scripts are supported by the DataFabric Manager server.

You begin creating a script by writing the script. You can use any scripting language, but keep in mind that your choice impacts your network administrators. The network administrator needs to install the interpreter you require on the system where the DataFabric Manager server is installed. It is recommended that you use a scripting language such as Perl. Perl is typically installed with the operating system on Linux and Windows workstations.

Besides, the `dfm report -F Perl` command in the DataFabric Manager server, generates reports for direct inclusion in Perl scripts. You have a way of importing information from the DataFabric Manager server reports into a Perl script. For more information, see the `dfm report` command man page.

See the *Operations Manager Help* for task procedures and option descriptions.

Commands that can be used as part of the script

Because scripts are executed by the DataFabric Manager server, your script is able to execute any of the commands available as part of the DataFabric Manager server CLI.

Your script can complete the following tasks:

- Use the `dfm run` command to execute commands on storage systems for which your DataFabric Manager server has been specified the credentials.
- Use the `dfm report` command to import information from Operations Manager reports into a Perl script.
- Use the `dfm event generate` command to enable your scripts to generate events.

Package of the script content

After you have written your script, you need to package it as a ZIP file.

The .zip file must contain the script, any data files that are needed by the script, and a file called `package.xml`.

For more information about `package.xml` and the XML schema information, see the *Operations Manager Help*.

What script plug-ins are

By using the script plug-in framework, you can perform several tasks in Operations Manager.

Scripts are installed using a .zip file that must contain the script, any data files that are needed by the script, and a file named `package.xml`.

The `package.xml` file contains information about the script, and might optionally contain definitions of new event types that your script generates.

After you have your .zip file ready, you can install it on the DataFabric Manager server to verify that your packaging functions correctly. After you verify the functionality, you can use, or distribute the script.

You can perform the following tasks using the script plug-in framework:

- Manage the scripts you add to the DataFabric Manager server
- Create and manage script jobs
- Create and manage schedules for the script jobs you have created
- Define new event classes during script installation or generate these events during script execution

For more information about creating scripts and the contents of the script .zip file, see the *Operations Manager Help*.

What the script plug-in directory is

By default, scripts are installed in the script-plugins subdirectory of the installation directory.

You can change the location of this directory by using the `scriptDir` option of the `dfm` option command. The option value must conform to the following requirements:

- The value must be an absolute Linux or Windows path.
- The value must include an existing directory on the DataFabric Manager server.

What backup scripts do

The prebackup and postbackup scripts help in bringing the databases into the hot backup mode before a backup is performed.

The DataFabric Manager server provides the ability to run prebackup and postbackup scripts on specific primary directories, before and after data has been backed up from those directories. For

more information about the process of setting up such scripts to run on primary directories, see the DataFabric Manager server Backup man pages.

Overview of the DataFabric Manager server database backup process

You can back up the DataFabric Manager server database, script plug-ins, and performance data without stopping any of the DataFabric Manager server services. However, data collection and view modifications of Performance Advisor are suspended during the backup process.

There are two types of backups:

- **Archive**
This backup process backs up your critical data in compressed form as a .zip file. The DataFabric Manager server data is automatically converted to an archive format and the DataFabric Manager server stores the backup in a local or remote directory. You can easily move an archive-based backup to a different system and restore it. This backup process is time-consuming.
- **Snapshot**
This backup process uses the Snapshot technology to back up the database. You can quicken the backup process through this approach. But you cannot transfer a Snapshot backup to a different system and restore it.

When to back up data

You must back up your data regularly before any upgrade operation, and before any maintenance on the system host on the DataFabric Manager server or an operating system upgrade.

You are requested to backup as part of the installation wizard steps during an upgrade.

You can backup data before:

- Upgrading DataFabric Manager server (or during the upgrade, as one of the installation wizard steps)
- Any maintenance on the system host on the DataFabric Manager server or Operating System upgrade

Where to back up data

By default, the DataFabric Manager server creates the archive backups in two directories.

DataFabric Manager server creates archives in the following directories:

- `/opt/IBMdfm/data` on a UNIX system
- `C:\Program Files\IBM\DataFabric Manager\DFM\data` on a Windows system

You can also specify a different target directory. The backup file has an added file name extension of .db, .gz, .zip, or .ndb, depending on the version of the DataFabric Manager server that you are running.

Note: The current version of the DataFabric Manager server uses .ndb format.

The Snapshot-based backups are volume Snapshot copies. Therefore, unlike in the archive backups, you do not have to specify a target directory in the Snapshot based backups.

Recommendations for disaster recovery

If you are using the archive-based backups, set the database backup location to a remote location. You can access the backup even if the DataFabric Manager server is not accessible.

Similar to archive-based backups, integration support of DataFabric Manager server with SnapVault or SnapMirror is available for Snapshot-based backups.

Backup storage and sizing

When you perform an archived backup, the DataFabric Manager server calculates the amount of disk space required to complete the backup successfully. You can view the backup information in the Database Backup Schedule Confirmation page.

In the case of Snapshot-based backups, the DataFabric Manager server does not calculate the amount of disk space required. Therefore, it is a good practice to provide enough disk space to hold the backups.

Limitation of Snapshot-based backups

To configure the DataFabric Manager server database for Snapshot-based backups on Windows and Linux, you must install the appropriate versions of SnapDrive software.

SnapDrive restrictions on Windows

To configure the DataFabric Manager server database for Snapshot-based backups on Windows, you must install SnapDrive 4.1 for Windows or later.

SnapDrive restrictions on Linux

To configure the DataFabric Manager server database for Snapshot-based backups on Linux, you must install SnapDrive 2.2.1 for UNIX or later. For information about supported versions, see the SnapDrive for UNIX Compatibility Matrix page on the N series support website (accessed and navigated as described in [Websites](#) on page 23).

Related information

The IBM N series support site -- www.ibm.com/storage/support/nseries/

Access requirements for backup operations

You must log in to the DataFabric Manager server with CoreControl capability in the Global scope to perform various backup operations.

You can perform the following backup operations:

- Create
- Delete
- Start
- Abort
- Set schedule
- Enable schedule
- Export
- Diagnose

To list the backups on a directory, and to get the status and schedules of backups, log in to Operations Manager with the GlobalRead role.

Changing the directory path for database backups

You can use Operations Manager to change the directory path for database backups if you want to back up to a different location.

Steps

1. Click **Setup > Options**.
2. Click **Database Backup under Edit Options**.
3. In the **Archive Backup Destination Directory** field, change the directory path if you want to back up to a different location.
4. Click **Update**.

Starting database backup from Operations Manager

You can manually start a database backup from Operations Manager.

Steps

1. Select **Setup > Database Backup**.
2. Select **Backup Type**.

You can choose between the archive-based and Snapshot-based backups.

3. Click **Back Up Now**.

The Database Backup Confirmation page is displayed if the Backup Type is Archive.

4. Click **Run Backup**.

After you finish

Note: To stop a backup that is in progress, click **Abort Backup**.

Scheduling database backups from Operations Manager

You can schedule a database backup to occur at a specific time on a recurring basis.

Before you begin

You must be authorized to perform all the steps of this task; your RBAC administrator can confirm your authorization in advance.

About this task

While scheduling database backup in the archive format, hourly backups and multiple backups in a day are not feasible. Because backups in the archive format take time.

Steps

1. Click **Setup > Database Backup**.
2. Select Backup Type.
You can choose between the archive-based and Snapshot-based backups.
3. Select **Enable Schedule** to activate the schedule.
4. Select the frequency to run the backup and enter the time to run it.
Entries are based on hourly, less frequent than hourly, daily, and weekly schedules.
5. Click **Update Schedule**.
6. Verify that the settings are correct, then click **Update Data**.

Specifying the backup retention count

You can specify the number of backups the DataFabric Manager server needs to retain.

Steps

1. Select **Setup > Options**.
2. Click **Database Backup**.
3. In the Database Backup Retention Count field, enter the count.
4. Click **Update**.

Disabling database backup schedules

You can temporarily disable the DataFabric Manager server database backup schedule.

Steps

1. Select **Setup > Database Backup**.
2. To disable the schedule, ensure that Enable Schedule check box is not selected.
3. Click **Update Schedule**.

Listing database backups

You can view information about the DataFabric Manager server database backups.

Steps

1. Select **Setup > Database Backup**.
2. Click **List Database Backups**.

Deleting database backups from Operations Manager

You can delete the DataFabric Manager server database backup using Operations Manager.

Steps

1. Select **Setup > Database Backup**.
2. Click **List Database Backups**.
3. Select **Delete Selected**.

Displaying diagnostic information from Operations Manager

The DataFabric Manager server might fail to detect the DataFabric Manager server setup on a LUN. You can investigate such backup-related issues using diagnostic information.

Steps

1. Click **Setup > Database Backup**.
2. Click **Update Schedule**.
Update schedule after you schedule a database backup.
3. Click **Diagnose**.

Exportability of a backup to a new location

You can export a Snapshot-based backup to a different location using the `dfm backup export <backup_name>` command. You can overwrite an existing backup at a new location by using the `dfm backup export -f <backup_name>` command.

To export the backup file to a new path, use `dfm backup export <backup_name> [target_filepath]` command. If `target_filepath` is not specified, the archive-based backup is created by default in the directory specified in the Archive Backup Destination Directory field using Operations Manager.

What the restore process is

You can restore the DataFabric Manager server database backups only through the CLI. You can restore all data by using the `dfm backup restore` command.

If, for some reason, the restore job fails, the DataFabric Manager server attempts to restore the database to its earlier state, and all temporary files that were created are deleted. If this attempt also fails, the CLI prompts you with a set of instructions to restore the original database.

Note: Do not run any `dfm` command during the DataFabric Manager server restore or upgrade operation. If some commands are run, they can interfere with the restore or upgrade operation by locking database tables and causing the operation to fail.

Restoring the database from the archive-based backup

You can restore the DataFabric Manager server database on the same system. You can specify the absolute path for the backup file using the `dfm backup restore` command.

Steps

1. Copy the backup file into the `databaseBackupDir` directory.
2. Type the following command at the command line: `dfm backup restore <backup_name>`

`backup_name` is the name of the file to which you saved your DataFabric Manager server database.

A “Completed restore” message is displayed when the restore process finishes successfully.

Restoring the database from the Snapshot copy-based backup

You can restore the DataFabric Manager server database on the same system.

Steps

1. Type the following command at the command line to display the names of backup copies of the database:

```
dfm backup list
```

2. Type the following command at the command line:

```
dfm backup restore <backup_name>
```

backup_name is the name of the backup copy in the database.

A “Completed restore” message is displayed when the restore process finishes successfully.

Restoration of the database on different systems

To restore the DataFabric Manager server database and other configuration from the archive-based backup on to another server, create a backup file (`dfm backup create`). Copy that backup file onto the new system, and then restore the backup on that system (`dfm backup restore`).

You can restore the database from a Snapshot-based backup on to another server. One way is to export the Snapshot-based backup in the archive format (`dfm backup export`), copy that backup file onto the new system, then restore the backup on that system (`dfm backup restore`).

Another way to restore the database from a Snapshot-based backup is to connect to the Snapshot copy using the SnapDrive commands, and run the `dfm datastore setup-n <target_dir_LUN>` command.

However, after you restore a database from one DataFabric Manager server to another, the local administrator account might be different on the new server. This restore operation would result in the local administrator account losing access to the restored DataFabric Manager server database. If this happens, you need to perform the following tasks:

- Log in to the new system as a user with GlobalFullControl.
- Add the local account administrator of the new system back into the DataFabric Manager server with GlobalFullControl capabilities.

Note: If there are no GlobalFullControl users that can access the new system, contact technical support for assistance. This can usually be avoided by ensuring that a domain user (who has permission to log in to both systems) exists in the DataFabric Manager server with GlobalFullControl role before migrating the database.

Disaster recovery configurations

You can configure the DataFabric Manager server for disaster recovery by using the N series Management Console data protection capability and SnapDrive.

Note: Disaster Recovery does not support IPv6 addressing.

Disaster recovery support enables you to recover the DataFabric Manager server services quickly on another site. Disaster recovery support prevents any data loss due to disasters that might result in a total site failure.

A disaster recovery plan typically involves deploying remote failover architecture. This remote failover architecture allows a secondary data center to take over critical operations when there is a disaster in the primary data center.

Disaster recovery using the N series Management Console data protection capability

If you have the N series Management Console data protection capability license, then you can use the N series Management Console data protection capability to configure the DataFabric Manager server for disaster recovery.

Snapshot-based backups are made at the primary site according to the backup schedule. By using the N series Management Console data protection capability and SnapMirror technology, the Snapshot-based backups are registered with the N series Management Console data protection capability. These Snapshot-based backups are then mirrored to the secondary site according to the the N series Management Console data protection capability Mirror policy.

If a catastrophic failure of the primary site occurs, the DataFabric Manager server services should be started on the secondary site using mirrored data by running the `dfm datastore mirror connect` command.

Limitation of disaster recovery support

You must use SnapDrive to configure the DataFabric Manager server on Linux for disaster recovery.

Related information

[*Disaster Recovery Support for DataFabric Manager Data Using SnapDrive*](#)

Prerequisites for disaster recovery support

Ensure that the prerequisites are met for disaster recovery support for the DataFabric Manager server data.

- You must be a Windows domain administrator.
- You must have a the N series Management Console data protection capability license.
- You must have SnapDrive for Windows 6.0 or later. The DataFabric Manager server is dependent on SnapDrive for Windows to provide disaster recovery support.
- You must be using the same version of Data ONTAP on both the source and destination storage systems. To ensure that you have the required Data ONTAP version for SnapDrive, see the IBM N series support site.
- You must have a SnapMirror license for both the source and destination storage systems.
- You must configure Snapshot-based backup for the DataFabric Manager server. To grant root access to the Windows domain account that is used by the SnapDrive service, you must configure the source and destination storage systems. You can configure the storage systems, by setting the `waf1.map_nt_admin_priv_to_root` option to "On" in the CLI.

- It is best to have a dedicated flexible volume for the DataFabric Manager server data, because Volume SnapMirror is used for mirroring the data.
- The DataFabric Manager server data must be stored in a LUN.
- The source and destination storage systems must be managed by the N series Management Console data protection capability.
- The Windows domain account that is used by the SnapDrive service must be a member of the local built-in group or local administrators group on both the source and destination storage systems.
- The Windows domain account used to administer SnapDrive must have full access to the Windows domain to which both the source and destination storage systems belong.

Related information

The IBM N series support site -- www.ibm.com/storage/support/nseries/

Setting up the DataFabric Manager server

You must perform a set of tasks to set up the DataFabric Manager server for disaster recovery.

Steps

1. Install or upgrade the DataFabric Manager server on the primary site by completing the following steps:
 - a. Install the DataFabric Manager server.
 - b. Install SnapDrive for Windows and configure it with the N series Management Console data protection capability.
 - c. Create an FC-based or iSCSI-based LUN on the storage system using SnapDrive.
 - d. Run the `dfm datastore setup` command to migrate the data to a directory in the FC-based or iSCSI-based LUN.
2. Configure a schedule for Snapshot-based backup.
3. Run the `dfm datastore mirror setup` command to create the application dataset.
4. To configure the DataFabric Manager server for disaster recovery using the the N series Management Console data protection capability features, complete the following steps:
 - a. Create a volume on the destination storage system having the same size and space configuration as the primary volume.

If you have either the the N series Management Console data protection capability Disaster Recovery license or the the N series Management Console provisioning capability license, secondary volume provisioning can take advantage of the policies provided by that license. If you do not have either of these licenses, then provision the secondary volume manually.

- b. Assign the provisioned secondary volume to the application dataset. For more information about how to use the N series Management Console data protection capability to assign resources to a dataset, see the *Management Console Help*.
 - c. Assign a schedule to the application dataset. For more information about how to use the N series Management Console data protection capability to assign schedules to a dataset, see the *Management Console Help*.
 - d. Ensure that there are no conformance issues.
5. Run the `dfm backup diag` command and note down the SnapMirror location information from the command output.

You would need this information while using the `dfm datastore mirror connect` command during the process of recovering the DataFabric Manager server.
 6. Install the DataFabric Manager server on the secondary site.
 7. Run the `dfm service disable` command on the secondary site to disable all DataFabric Manager server services.

The DataFabric Manager server services must be enabled only during disaster recovery, by using the `dfm datastore mirror connect` command.

Related tasks

[Scheduling database backups from Operations Manager](#) on page 275

Recovering the DataFabric Manager server services

You can recover the DataFabric Manager server services on the secondary site if a disaster occurs at the primary site.

Steps

1. To connect to the LUN using the Microsoft Management Console (MMC) Connect Disk wizard on the secondary storage system, complete the following steps:
 - a. Expand the `Storage` option in the left pane of the MMC.
 - b. If you are managing multiple instances of SnapDrive, double-click SnapDrive List. Otherwise, double-click SnapDrive.
 - c. Double-click the name of the SnapDrive host you want to manage.
 - d. Right-click Disks.
 - e. Select the `Connect Disk` option and follow the instructions on the Connect Disk wizard.

For more information about connecting virtual disks, see the *SnapDrive for Windows Installation and Administration Guide*.

2. Run the `dfm service enable` command to enable the services.

3. Run the `dfm datastore setup-n` command to configure the DataFabric Manager server to use the mirrored data.
4. Run the `dfm service start` command to start the DataFabric Manager server services.
5. Using the the N series Management Console data protection capability UI, change the dataset created for the DataFabric Manager server data from the current mode to suspended mode.
6. Run the `dfm options set` command to reset the `localHostName global` option on the secondary site, if the primary site is clustered using MSCS.
7. Run the `dfm service disable` command to disable the services on the primary site.
If the primary site is clustered using MSCS, offline the services before disabling them.

Related information

[SnapDrive for Windows Installation and Administration Guide - www.ibm.com/storage/support/nseries](http://www.ibm.com/storage/support/nseries)

Recovering the DataFabric Manager server services using the `dfm datastore mirror connect` command

You can use the `dfm datastore mirror connect` to recover the DataFabric Manager server services on the secondary site if a disaster occurs at the primary site.

Steps

1. Run the `dfm datastore mirror connect` command on the DataFabric Manager server at the secondary site to start the DataFabric Manager server services using mirrored data.

The `dfm datastore mirror connect` command performs the following operations:

- Breaks the mirror relationship between the source and destination DataFabric Manager server.
 - Connects to the mirrored volume or LUN using SnapDrive for Windows.
 - Enables the services using the `dfm service enable` command.
 - Configures the DataFabric Manager server to use the data from the mirrored location.
 - Starts the DataFabric Manager server services.
 - Puts the dataset created for the DataFabric Manager server data in suspended mode.
2. Run the `dfm options set` command to reset the `localHostName global` option on the secondary site, if the primary site is clustered using MSCS.
 3. Run the `dfm service disable` command to disable the services on the primary site.
If the primary site is clustered using MSCS, offline the services before disabling them.

Failing back DataFabric Manager server services

You must complete a list of tasks to fail back DataFabric Manager server services to the primary site.

Steps

1. Ensure that the DataFabric Manager server data at the source storage system is in synchronization with data at the destination storage system by completing the following steps:
 - a. Run the `dfdrm mirror list` command to find relationships between the source and destination storage systems.
 - b. Run the `dfdrm mirror resync -r` command to resynchronize the mirror relationships. This command reverses the mirror direction and starts updates.
 - c. If the SnapMirror relationship is removed during the process of recovering the DataFabric Manager server services, run the `snapmirror resync` command to resynchronize the data at the storage system level.
 - d. Run the `dfdrm mirror initialize` command to create a new relationship from the secondary storage system to the new primary storage system, if the primary storage system is destroyed during disaster.
2. Run the `dfm service disable` command to stop and disable the services at the secondary site.
3. To start the DataFabric Manager server services using the mirrored data on the primary site, use the `dfm datastore mirror connect` command.

Alternatively, perform the following steps:

- Run the `dfm datastore mirror connect` command at the CLI
 - Alternatively you perform the following procedure:
 - a. Connect to the LUN using MMC Connect Disk wizard on the primary storage system.
 - b. Run the `dfm service enable` command to enable the services.
 - c. Run the `dfm datastore setup -n` command to configure the DataFabric Manager server to use the mirrored data on the LUN.
 - d. Run the `dfm service start` command to start the DataFabric Manager server services.

Note: The `dfm datastore mirror connect` command does not support shared storage. Therefore, the command should not be used if the primary system is set up for cluster using MSCS.
4. Run the `dfdrm mirror resync -r` command to resynchronize the mirror relationships so that they are no longer reversed.
 5. Run the `snapmirror resync` command to resynchronize the data at the storage system level, if the SnapMirror relationship is removed during the failback process.

6. Run the `dfm host discover` command to discover the reversed relationships on the primary site, if they are not discovered already.
7. Run the `dfdrm mirror list` command to ensure that these relationships are discovered.
8. Run the `dfm datastore mirror destroy` command to destroy the application dataset created for the DataFabric Manager server data.
9. Run the `dfm datastore mirror setup` command to create a new application dataset for the DataFabric Manager server data.
10. Using the the N series Management Console data protection capability UI, import the SnapMirror relationship already established for the DataFabric Manager server data to the new application dataset.

For more information about how to use the N series Management Console data protection capability to import SnapMirror relationships, see the *Management Console Help*.

Disaster recovery using SnapDrive

If you do not have a the N series Management Console data protection capability license, then you can use SnapDrive to configure the DataFabric Manager server for disaster recovery.

For more information, you can view the technical report.

Related information

[*Disaster Recovery Support for DataFabric Manager Data Using SnapDrive*](#)

Troubleshooting in Operations Manager

Learn about the common issues with the DataFabric Manager server, how to troubleshoot those problems, and how to get technical assistance from your service provider.

You can contact technical support, if you cannot troubleshoot or resolve those problems.

AutoSupport in the DataFabric Manager server

You can use the AutoSupport feature of the DataFabric Manager server to send messages to technical support.

When you install or upgrade to DataFabric Manager server 3.3 or later, the scheduler service automatically enables AutoSupport after the first 24 hours of operation, if you have not disabled the feature. The DataFabric Manager server then starts to monitor the system's operations and logs a message that AutoSupport was enabled.

AutoSupport sends messages to technical support over secure HTTPS (by default), HTTP, or SMTP.

Note: AutoSupport in the DataFabric Manager server does not support IPv6 addressing.

Reasons for using AutoSupport

With the help of the AutoSupport feature in the DataFabric Manager server, you can detect potential problems and get quick help.

The AutoSupport feature sends messages to the technical support for the following reasons:

- Scripts have been programmed to automatically look for particular data in AutoSupport weekly reports that might indicate a potential problem.
- Technical support helps you to solve problems that AutoSupport detects.

A range of support might be provided, including automated parts replacement, e-mail contact, and contact by a technical support engineer, depending on the type of problem. For example, if a message is received that a disk failure occurred on your system, a replacement disk is automatically sent to you.

Types of AutoSupport messages in the DataFabric Manager server

AutoSupport tracks events and sends messages.

AutoSupport generates the following types of messages:

Event message A message that AutoSupport sends to recipients when an event tracked by AutoSupport occurs. The message contains information to help you diagnose and troubleshoot the problem.

Weekly report A weekly report is a general status message that AutoSupport sends automatically each week to recipients you have identified.

Note: If you are using a DataFabric Manager server demonstration license, the DataFabric Manager server does not send AutoSupport messages.

Protection of private data by using AutoSupport

You can ensure that private data, such as IP addresses and user names, are not included in the AutoSupport message.

Complete AutoSupport messages are required for normal technical support. Minimal AutoSupport messages omit sections and values that might be considered sensitive information, but greatly affect the level of support you can receive.

If you do not want to include private data, such as IP addresses, host names, and user names, set the `Autosupport Content` global option to `minimal`.

Configuring AutoSupport

You can configure AutoSupport using Operations Manager.

Steps

1. Click **Setup > Options > AutoSupport**.
2. From the **AutoSupport Settings** page, identify the administrator to be designated as the sender of the notification.
3. Specify the type of AutoSupport content that messages should contain.

Note: If this setting is changed from "complete" to "minimal," any complete AutoSupport message not sent is cleared from the outgoing message spool and notification of this is displayed on the console.
4. Select Yes to enable the inclusion of logs with the AutoSupport notification.
5. Specify the type of delivery --HTTP, HTTPS or SMTP-- for AutoSupport notification to technical support.
6. Enter the number of times the system should try to resend the AutoSupport notification before giving up, if previous attempts have failed.
7. Enter the time to wait before trying to resend a failed AutoSupport notification.
8. Select Include to include the Performance Advisor AutoSupport data along with the DataFabric Manager server AutoSupport data.
9. Select Include to include the the N series Management Console provisioning capability AutoSupport data along with the DataFabric Manager server AutoSupport data.
10. Click **Update**.

DataFabric Manager server logs

The DataFabric Manager server creates logs, which you can use to troubleshoot issues such as storage system discovery not working, CLI commands failing for unexpected reasons, and events not generating as expected.

Access to logs

You can access the DataFabric Manager server logs using the Operations Manager GUI or by using the CLI.

You can access the DataFabric Manager server logs through the Diagnostics page. To access this page, use the following URL:

http://mgmt_station:8080/dfm/diag

`mgmt_station` is the name or IP address of the workstation on which the DataFabric Manager server is installed.

You must scroll down to the Logs section to find the available DataFabric Manager server logs.

You can also access the DataFabric Manager server logs through the CLI, going to different directories depending on whether you are using a Windows or a UNIX workstation:

- On a Windows workstation, enter `installation_directory\dfm\log`.
- On a UNIX workstation, enter `installation_directory/log`.

Accessing the logs through the DataFabric Manager server CLI

You can access the DataFabric Manager server logs using the CLI.

On a Windows workstation, you can find the DataFabric Manager server logs in the following directory: `installation_directory\dfm\log`.

On a UNIX workstation, you can find the DataFabric Manager server logs in the following directory: `installation_directory/log`.

Apache and Sybase log rotation in the DataFabric Manager server

The DataFabric Manager server automatically rotates Apache and Sybase logs. Apache log files are rotated when they are 3,000 KB or larger. Sybase logs are rotated when they are 10,000 KB or larger.

Common DataFabric Manager server problems

You can resolve some of the common problems that you might encounter when using the DataFabric Manager server.

Communication issues between the DataFabric Manager server and routers

Communication between the DataFabric Manager server and routers fails due to problems such as a mismatch of SNMP community strings, SNMP disabled on the router, and so on.

The DataFabric Manager server relies on routers to discover networks other than the network to which it is attached. If the DataFabric Manager server fails to communicate with a router, it cannot discover other networks attached to that router. These are some typical reasons DataFabric Manager server fails to communicate with routers:

- The SNMP community strings do not match.
- SNMP is disabled on the router.
- The router is beyond the maximum number of hops set in the Network Discovery Limit option.

E-mail alerts not working in the DataFabric Manager server

If e-mail alerts do not work, you can verify the e-mail address in the `log` file.

If an alarm does not send an e-mail to the expected e-mail address, use the log files generated by the DataFabric Manager server to help troubleshoot the problem.

- Look in `alert.log` to see if the DataFabric Manager server attempted to send an e-mail to that address.
- Look in `dfmeventd.log` to see if errors were reported.

Related concepts

[DataFabric Manager server logs](#) on page 287

How discovery issues are resolved

You can use Diagnose Connectivity tool to troubleshoot discovery problems.

The DataFabric Manager server provides a Diagnose Connectivity tool that automates frequently used steps of the troubleshooting process for connectivity issues. Use this tool when you want to troubleshoot discovery problems.

This tool queries the DataFabric Manager server database about a selected storage system, runs connectivity tests, and displays information and test outcomes. The sequence of steps depends on whether the selected storage system is managed or unmanaged. A managed storage system is one that is in the DataFabric Manager server database. An unmanaged storage system is one that is not in the DataFabric Manager server database.

Use of the Diagnose Connectivity tool for a managed storage system

You can collect information about your managed storage system by using the Diagnose Connectivity tool.

The Diagnose Connectivity tool queries the database and displays a summary of information about the storage system:

- The name, DataFabric Manager server object ID, model, system ID, and OS version
- Whether the storage system is up, according to the DataFabric Manager server
- Results of the SNMP, ping, SSH, RLM, XML, RSH, and HTTP port tests.

Use of the Diagnose Connectivity tool for an unmanaged storage system

The Diagnose Connectivity tool runs the following tests:

- Determines if the storage system IP address falls within the range of networks discovered by the DataFabric Manager server
- Sends an SNMP GET request to determine if the DataFabric Manager server can use SNMP to communicate with the storage system
- If it is a supported storage system, shows sysName, sysObjectID, and productID
- Uses the ping utility to contact the IP address of the storage system and tests for connectivity
- If it is a supported storage system, it tests the following:
 - The RSH connection that uses the user name and password stored in the DataFabric Manager server for the storage system
 - The HTTP port

Note: For unmanaged storage systems, you must run the Diagnose Connectivity tool from the command line.

Where to find the Diagnose Connectivity tool in Operations Manager

You can use the Diagnose Connectivity tool for both managed storage systems, and clusters.

You can find the Diagnose Connectivity tool under Storage Controller Tools or Cluster Tools list in the left pane on a Details page. The Diagnose Connectivity tool is available only for managed storage systems. To diagnose connectivity on unmanaged storage systems, use one of the following methods:

- From the CLI, you should run the `dfm host diag` command.
- From Operations Manager, try to add the storage system. If the operation fails, an error message is displayed with a link labeled **Click here to troubleshoot**. You must click this link to run the Diagnose Connectivity tool.

Reasons why the DataFabric Manager server might not discover your network

The DataFabric Manager server might not discover your network if it fails to communicate with the router. The DataFabric Manager server might not discover your network if a SNMP community

string in the DataFabric Manager server does match the community string required by the router, switch, and storage system.

Three possible reasons why the DataFabric Manager server might not discover your network are as follows:

- The SNMP community string set in the DataFabric Manager server does not match the community string required by the router, switch, and storage system.
- The DataFabric Manager server fails to communicate with a router, and it cannot discover other networks attached to that router.
- You changed the IP address of the router for the network.

Troubleshooting network discovery issues

You can troubleshoot network discovery issues, such as the DataFabric Manager server not discovering specific networks, using Operations Manager.

Steps

1. Ensure that the `Network Discovery Enabled` option on the **Options** page is set to **Yes**.
2. Ensure that the router is within the maximum number of hops set in the `Network Discovery Limit` option.
3. From the command line, run the **Diagnose Connectivity** tool against the IP address of the router of the network to determine if the DataFabric Manager server can communicate with the router through SNMP.

If you changed the IP address for the router, you must change the primary IP address stored in the DataFabric Manager server on the Edit Settings page. You can also modify the primary IP address by entering the following CLI command:

```
dfm host set host-id hostPrimaryAddress=ip-address
```

4. Determine whether an SNMP community string other than the default (public) is required for the network device to which the undiscovered network is attached.

To set an SNMP community string in the DataFabric Manager server, click the **Options** link (in the Banner area), find **Discovery Options**, and then click the **edit** link beside **Network Credentials**. On the Network Credentials page, click the **edit** link at the right of the **SNMP Community** whose string you want to set.

5. If Steps 1 through 3 are not successful, add the network manually by using the Networks To Discover option on the **Options** page under **Discovery Options**.

Troubleshooting appliance discovery issues with Operations Manager

If the DataFabric Manager server does not discover a storage system, you can troubleshoot using Operations Manager.

Steps

1. Ensure that the `Host Discovery Enabled` option on the **Options** page is set to **Enabled**.
2. Click the **edit** link of the `Networks to Discover` option to check whether the network to which this appliance is attached has been discovered.

If the network to which this storage system is attached has not been discovered, follow the troubleshooting guidelines.

3. Determine whether an SNMP community string other than the default (public) is required for the network device to which the undiscovered network is attached.

To set an SNMP community string in the DataFabric Manager server, perform the following steps:

- a. Click the **Options** link (in the Banner area).
 - b. Find **discovery Options**.
 - c. Click the **edit** link beside **Network Credentials**.
 - d. On the **Network Credentials** page, click the **edit** link at the right of the SNMP Community whose string you want to set.
4. If Steps 1 through 3 are not successful, add the network manually by using the `Networks to Discover` option on the **Options** page under **Discovery Options**.

How configuration push errors are resolved

You can troubleshoot configuration push errors by analyzing logs created by the DataFabric Manager server.

The DataFabric Manager server logs the reason for the failure in the job's report, if one or more push jobs fail. You can access the job report by clicking the job ID in the Job ID column of the Status of Configuration Jobs option.

The following failure conditions are documented by the DataFabric Manager server:

- An attempt to push a configuration for a particular group was disallowed because the group already had a pending configuration push job. To fix the error, you must cancel the pending job and then re-push the new job.

- The storage system was offline. When a storage system is offline, the DataFabric Manager server continues to try to contact the storage system until you manually cancel the push.
- The DataFabric Manager server could not authenticate to the system. You can correct this problem by setting the correct login and password.
- The storage system downloaded the configuration, but did not update its configuration because it found the configuration to be invalid.

How File Storage Resource Manager (FSRM) issues are resolved

You can resolve FSRM issues by reviewing host agent credentials and network issues.

- If you cannot collect FSRM data, then review the settings of host agent administration and host agent password, and then check the login and password settings configured on the FSRM host agent and on the DataFabric Manager server.
- If a path walk fails, do the following:
 - Verify that the FSRM host agent is running.
 - Check for network problems between the FSRM host agent and the storage device. Network problems between the FSRM host agent and the storage device can cause the path walk to fail or return incomplete results.
- If a path walk takes a long time to complete: A path walk can take many hours to complete. You can monitor the status of the path walk that is in progress from the SRM Path Details page (**Control Center > Home > Group Status > File SRM > SRM path name**).

Issues related to SAN events

You can use Operations Manager to troubleshoot SAN events by resolving issues related to FC switch port, HBA port, and LUNs.

Offline LUNs

A LUN can be offline if it is taken over by an administrator or there is a conflict between the serial numbers of two LUNs.

Causes:

A LUN goes offline typically in one of two situations:

- An administrator, to perform maintenance or to apply changes to the LUN, such as to modify its size might have changed the status.
- The LUN has a serial number that conflicts with that of another LUN.

Corrective action:

- If an administrator did not change the LUN status, bring the LUN online from the storage system console.
- Check for a conflicting serial number and resolve the issue.

Snapshot copy of LUN not possible

If you cannot take a Snapshot copy of a LUN, you can expand the size of the volume.

Cause:

A Snapshot copy for the volume that contains the LUN cannot be taken because the amount of free space on this volume is less than the used and reserved space.

Corrective action:

Expand the size of the volume.

High traffic in HBA Port

High traffic in HBA Port occurs if the DataFabric Manager server threshold HBA Port Too Busy exceeds the permitted value.

Cause:

The DataFabric Manager server threshold HBA Port Too Busy has been exceeded.

Corrective action:

Determine the cause of high traffic. If traffic stays above the threshold for a long period, consider upgrading the infrastructure.

Import and export of configuration files

You can export your configuration files for archiving to the DataFabric Manager server or to your local computer. If you want to restore the configuration files, you can import the archived configuration files.

The DataFabric Manager server enables you to save all your configuration files as a text file. This helps you to revert to a working set of configuration files, if you need to undo future configuration changes. After creating the file, you should archive it in a different location for safekeeping. You can use the Export Configuration Files option on the Configurations page to save your configuration files to the DataFabric Manager server, or to your local computer.

If you need to restore the configuration files, you can import the archived file that you previously exported. You can also edit the configuration file with a text editor, to make any required changes, and then import it.

How inconsistent configuration states are fixed

To fix an inconsistent configuration state, re-push the configuration files to the required storage systems.

The DataFabric Manager server routinely monitors the storage systems in a configuration resource group to determine if their settings are inconsistent with the group's configuration files. To fix the inconsistent configuration state, re-push the configuration files to any storage systems that require updating.

Configuration modification events can also help you maintain configuration consistency by alerting you to configuration changes.

Related concepts

[Management of storage system configuration files](#) on page 229

Data ONTAP issues impacting protection on vFiler units

You might encounter errors while trying to create SnapMirror and SnapVault relationships by using vFiler units in DataFabric Manager server 7.2 and later.

Data ONTAP versions earlier than 7.2 do not support SnapMirror and SnapVault commands on vFiler units. To create SnapMirror and SnapVault relationships, you must use the hosting storage system. The DataFabric Manager server uses the hosting storage system to create SnapMirror and SnapVault relationships.

For Data ONTAP 7.2 and later, SnapMirror and SnapVault relationships can be created using vFiler units. However, the DataFabric Manager server continues using the hosting storage system to create, monitor, and manage these relationships. As a result, you might encounter the following issues:

- **Issue:**

If the `snapvault.access` and `snapmirror.access` options on the source storage system allow access only to the destination vFiler unit, relationship creation, scheduled backups, on-demand backups, SnapMirror updates, and SnapMirror resync from the DataFabric Manager server fails. The DataFabric Manager server displays the following error message: "Request denied by the source storage system. Check access permissions on the source."

Workaround:

To allow access to the destination hosting storage system, set the `snapmirror.access` and `snapvault.access` options on the source system.

- **Issue:**

If the `ndmpd.preferred_interfaces` option is not set on the source hosting storage system, the backups from the DataFabric Manager server might not use the correct network interface.

Workaround:

Set the `ndmpd.preferred_interfaces` option on the source hosting storage system.

- **Issue:**

The backups and SnapMirror updates from the DataFabric Manager server fail with the error message "Source unknown." This issue occurs when both of the following conditions occur:

- A relationship between two vFiler units is imported into the DataFabric Manager server by autodiscovery or added manually.
- The destination hosting storage system is not able to contact the IP address of the source vFiler unit.

Workaround:

Ensure that the host name or IP address of the source system that is used to create relationships can be reached from the destination hosting storage system.

Appendix

Report fields and performance counters

Report fields and performance counters for filer catalogs

The DataFabric Manager server creates reports for Filer Catalogs.

The table below lists and describes the various fields in Filer Catalogs and provides the corresponding performance counters.

Field	Name/description	Performance counter
Filer.TotalOpsperSec	Storage System Total Ops/Sec	system:total_ops
Filer.CIFSops	Storage System CIFS Ops/Sec	system:cifs_ops
Filer.NFSops	Storage System NFS Ops/Sec	system:nfs_ops
Filer.HTTPops	Storage System HTTP Ops/Sec	system:http_ops
Filer.iSCSIops	Storage System iSCSI Ops/Sec	system:iscsi_ops
Filer.FCPOps	Storage System FCP Ops/Sec	system:fcop_ops
Filer.NFSv3ReadOps	Storage System NFSv3 Read Ops/Sec	nfsv3:nfsv3_read_ops
Filer.NFSv3WriteOps	Storage System NFSv3 Write Ops/Sec	nfsv3:nfsv3_write_ops
Filer.NFSv4ReadOps	Storage System NFSv4 Read Ops/Sec	nfsv4:nfsv4_read_ops
Filer.NFSv4WriteOps	Storage System NFSv4 Write Ops/Sec	nfsv4:nfsv4_write_ops
Filer.NFSv3Avglatency	Storage System NFSv3 Avg Latency (millisec)	nfsv3:nfsv3_avg_op_latency
Filer.NFS4Avglatency	Storage System NFSv4 Avg Latency (millisec)	nfsv4:nfsv4_avg_op_latency
Filer.CPUBusy	Storage System CPU Busy (%)	system:cpu_busy
Filer.iSCSIReadOps	Storage System iSCSI Read Ops/Sec	iscsi:iscsi_read_ops

Field	Name/description	Performance counter
Filer.iSCSIWriteOps	Storage System iSCSI Write Operations	iscsi:iscsi_write_ops
Filer.CIFSLatency	Storage System CIFS Latency (millisec)	cifs:cifs_latency
Filer.NFSReadLatency	Storage System NFS Read Latency (millisec)	nfsv3:nfsv3_read_latency
Filer.NFSWriteLatency	Storage System NFS Write Latency (millisec)	nfsv3:nfsv3_write_latency
Filer.iSCSIRead Latency	Storage System iSCSI Read Latency (millisec)	iscsi:iscsi_read_latency
Filer.iSCSIWrite Latency	Storage System iSCSI Write Latency (millisec)	iscsi:iscsi_write_latency
Filer.FCPReadLatency	Storage System FCP Read Latency (millisec)	fcp:fcp_read_latency
Filer.FCPWriteLatency	Storage System FCP Write Latency (millisec)	fcp:fcp_write_latency
Filer.NASThroughput	Storage System NAS Throughput (KB/Sec)	system:nas_throughput
Filer.SANThroughput	Storage System SAN Throughput (KB/Sec)	system:san_throughput
Filer.DiskThroughput	Storage System Disk Throughput (KB/Sec)	system:disk_throughput
Filer.NetThroughput	Storage System Network Throughput (MB/Sec)	system:load_total_mbps
Filer.LoadInbound Mbps	Storage System Total Data Received (MB/Sec)	system:load_inbound_mbps
Filer.LoadOutbound Mbps	Storage System Total Data Sent (MB/Sec)	system:load_outbound_mbps
Filer.NetDataSent	Storage System Network Data Sent (KB/Sec)	system:net_data_sent
Filer.NetDataRecv	Storage System Network Data Receive (KB/Sec)	system:net_data_recv
Filer.LoadReadBytes Ratio	Storage System Ratio of disk data read and load outbound	system:load_read_bytes_ratio

Field	Name/description	Performance counter
Filer.LoadWriteBytes Ratio	Storage System Ratio of disk data write and load inbound	system:load_write_byte s_ratio
Filer.DiskDataRead	Storage System Disk Data Read (KB/Sec)	system:disk_data_read
Filer.DiskDataWritten	Storage System Disk Data Written (KB/Sec)	system:disk_data_written
Filer.FCPWriteData	Storage System FCP Write Data (B/Sec)	fcf:fcf_write_data
Filer.FCPReadData	Storage System FCP Read Data (B/Sec)	fcf:fcf_read_data
Filer.iSCSIWriteData	Storage System iSCSI Write Data (B/Sec)	iscsi:iscsi_write_data
Filer.iSCSIReadData	Storage System iSCSI Read Data (B/Sec)	iscsi:iscsi_read_data
Filer.ProcessorBusy	Storage System Processor Busy (%)	system:avg_processor_busy
Filer.NFSLatency	Storage System NFS Latency (millisec)	nfsv3:nfsv3_avg_op_latency
Filer.PerfViolation Count	Storage System Perf Threshold Violation Count	Not Applicable
Filer.PerfViolation Period	Storage System Perf Threshold Violation Period (Sec)	Not Applicable

Report fields and performance counters for vFiler catalogs

The DataFabric Manager server creates reports for vFiler Catalogs.

The following table lists and describes the fields of the vFiler Catalog:

Field	Name/description	Performance counter
vFiler.TotalOps	vFiler Total Ops/Sec	vfiler:vfiler_total_ops
vFiler.ReadOps	vFiler Read Ops/Sec	vfiler:vfiler_read_ops
vFiler.WriteOps	vFiler Write Ops/Sec	vfiler:vfiler_write_ops
vFiler.MiscOps	vFiler Miscellaneous Ops/Sec	vfiler:vfiler_misc_ops
vFiler.NetThroughput	vFiler Network Throughput (KB/Sec)	vfiler:vfiler_nw_throughput

Field	Name/description	Performance counter
vFiler.ReadBytes	vFiler Number of Bytes Read (KB/Sec)	vfiler:vfiler_read_bytes
vFiler.WriteBytes	vFiler Number of Bytes Write (KB/Sec)	vfiler:vfiler_write_bytes
vFiler.NetDataRecv	vFiler Network Data Received (KB/Sec)	vfiler:vfiler_net_data_recv
vFiler.NetDataSent	vFiler Network Data Sent (KB/Sec)	vfiler:vfiler_net_data_sent
vFiler.DataTransferred	vFiler Total Data Transferred (KB/Sec)	vfiler:vfiler_data_transferred
vFiler.PerfViolation Count	vFiler Perf Threshold Violation Count	Not Applicable
vFiler.PerfViolation Period	vFiler Perf Threshold Violation Period (Sec)	Not Applicable

Report fields and performance counters for volume catalogs

The DataFabric Manager server creates reports for Volume Catalogs.

The following table lists and describes the fields of the Volume Catalog:

Field	Name/description	Performance counter
Volume.TotalOps	Volume Total Ops/Sec	volume:total_ops
Volume.CIFSops	Volume CIFS Ops/Sec	volume:cifs_ops
Volume.NFSops	Volume NFS Ops/Sec	volume:nfs_ops
Volume.SANops	Volume SAN Ops/Sec	volume:total_san_ops
Volume.SANReadOps	Volume SAN Read Ops/Sec	volume:san_read_ops
Volume.SANWriteOps	Volume SAN Write Ops/Sec	volume:san_write_ops
Volume.SANOtherOps	Volume SAN Other Ops/Sec	volume:san_other_ops
Volume.ReadOps	Volume Read Ops/Sec	volume:read_ops
Volume.WriteOps	Volume Write Ops/Sec	volume:write_ops
Volume.OtherOps	Volume Other Ops/Sec	volume:other_ops
Volume.NFSReadOps	Volume NFS Read Ops/Sec	volume:nfs_read_ops
Volume.NFSWriteOps	Volume NFS Write Ops/Sec	volume:nfs_write_ops

Field	Name/description	Performance counter
Volume.NFSOtherOps	Volume NFS Other Ops/Sec	volume:nfs_other_ops
Volume.CIFSReadOps	Volume CIFS Read Ops/Sec	volume:cifs_read_ops
Volume.CIFSWriteOps	Volume CIFS Write Ops/Sec	volume:cifs_write_ops
Volume.CFSOtherOps	Volume CIFS Other Ops/Sec	volume:cifs_other_ops
Volume.FlexCache ReadOps	Volume FlexCache Read Ops/Sec	volume:flexcache_read_ops
Volume.FlexCache WriteOps	Volume FlexCache Write Ops/Sec	volume:flexcache_write_ops
Volume.FlexCache OtherOps	Volume FlexCache Other Ops/Sec	volume:flexcache_other_ops
Volume.Latency	Volume Latency (millisec)	volume:avg_latency
Volume.CIFSLatency	Volume CIFS Latency (millisec)	volume:cifs_latency
Volume.NFSLatency	Volume NFS Latency (millisec)	volume:nfs_latency
Volume.SANLatency	Volume SAN Latency (millisec)	volume:san_latency
Volume.ReadLatency	Volume Read Latency (millisec)	volume:read_latency
Volume.WriteLatency	Volume Write Latency (millisec)	volume:write_latency
Volume.OtherLatency	Volume Other Latency (millisec)	volume:other_latency
Volume.CIFSRead Latency	Volume CIFS Read Latency (millisec)	volume:cifs_read_latency
Volume.CIFSWrite Latency	Volume CIFS Write Latency (millisec)	volume:cifs_write_latency
Volume.CFSOther Latency	Volume CIFS Other Latency	volume:cifs_other_latency
Volume.SANRead Latency	Volume SAN Read Latency (millisec)	volume:san_read_latency
Volume.SANWrite Latency	Volume SAN Write Latency (millisec)	volume:san_write_latency

Field	Name/description	Performance counter
Volume.SANOther Latency	Volume SAN Other Latency (millisec)	volume:san_other_latency
Volume.NFSRead Latency	Volume NFS Read Latency (millisec)	volume:nfs_read_latency
Volume.NFSWrite Latency	Volume NFS Write Latency	volume:nfs_write_latency
Volume.NFSOther Latency	Volume NFS Other Latency (millisec)	volume:nfs_other_latency
Volume.Data Throughput	Volume Throughput (KB/Sec)	volume:throughput
Volume.PerfViolation Count	Volume Perf Threshold Violation Count	Not Applicable
Volume.PerfViolation Period	Volume Perf Threshold Violation Period (Sec)	Not Applicable

Report fields and performance counters for qtree catalogs

The DataFabric Manager server creates reports for Qtree Catalogs.

The following table lists and describes the fields of the Qtree Catalog:

Field	Name/description	Performance counter
Qtree.CIFSops	Qtree CIFS Ops/Sec	qtree:cifs_ops
Qtree.NFSops	Qtree NFS Ops/Sec	qtree:nfs_ops
Qtree.InternalOps	Qtree Internal Ops/Sec	qtree:internal_ops
Qtree.PerfViolation Count	Qtree Perf Threshold Violation Count	Not Applicable
Qtree.PerfViolation Period	Qtree Perf Threshold Violation Period (Sec)	Not Applicable

Report fields and performance counters for LUN catalogs

The DataFabric Manager server creates reports for LUN Catalogs.

The following table lists and describes the fields of the LUN Catalog:

Field	Name/description	Performance counter
LUN.TotalOps	LUN Total Ops/Sec	lun:total_ops
LUN.ReadOps	LUN Read Ops/Sec	lun:read_ops

Field	Name/description	Performance counter
LUN.WriteOps	LUN Write Ops/Sec	lun:write_ops
LUN.OtherOps	LUN Other Ops/Sec	lun:other_ops
LUN.Latency	LUN Latency (millisec)	lun:avg_latency
LUN.Throughput	LUN Throughput (KB/Sec)	lun:throughput
LUN.ReadData	LUN Read Data (KB/Sec)	lun:read_data
LUN.WriteData	LUN Write Data (KB/Sec)	lun:write_data
LUN.PerfViolation Count	LUN Perf Threshold Violation Count	Not Applicable
LUN.PerfViolation Period	LUN Perf Threshold Violation Period (Sec)	Not Applicable

Report fields and performance counters for aggregate catalogs

The DataFabric Manager server creates reports for Aggregate Catalogs.

The following table lists and describes the fields of the Aggregate Catalog:

Field	Name/description	Performance counter
Aggregate.TotalOps	Aggregate Total Ops/Sec	aggregate:total_transfers
Aggregate.UserReads	Aggregate User Reads (per_sec)	aggregate:user_reads
Aggregate.UserWrites	Aggregate User Writes (per_sec)	aggregate:user_writes
Aggregate.CPReads	Aggregate Reads done during CP (per_sec)	aggregate:cp_reads
Aggregate.PerfViolation Count	Aggregate Perf Threshold Violation Count	Not Applicable
Aggregate.PerfViolation Period	Aggregate Perf Threshold Violation Period (Sec)	Not Applicable

Report fields and performance counters for disk catalogs

The DataFabric Manager server creates reports for Disk Catalogs.

The following table lists and describes the fields of the Disk Catalog:

Field	Name/description	Performance counter
Disk.ReadOps	Disk User Read Ops/Sec	disk:user_reads
Disk.WriteOps	Disk User Write Ops/Sec	disk:user_writes
Disk.CPReads	Disk Reads initiated for CP processing (per_sec)	disk:cp_reads
Disk.ReadLatency	Disk Read Latency (millisec)	disk:user_read_latency
Disk.WriteLatency	Disk Write Latency (millisec)	disk:user_write_latency
Disk.CPReadLatency	Disk CP Read Latency (millisec)	disk:cp_read_latency
Disk.Throughput	Disk Throughput (blocks/Sec)	disk:throughput
Disk.Utilization	Disk Utilization (%)	disk:disk_busy
Disk.PerfThreshold Violations	Disk Perf Threshold Violations Count	Not Applicable
Disk.PerfViolation Period	Disk Perf Threshold Violation Period (Sec)	Not Applicable

Protocols and port numbers

The DataFabric Manager server uses various networking protocols and port numbers to communicate with storage systems, host agents (for SRM and SAN management), and Open Systems SnapVault agents.

DataFabric Manager server communication

You might need to enable both HTTP and HTTPS transports if multiple administrators are accessing the workstation from different locations. The presence of firewalls, non-trusted environments, or other circumstances might require a combination of HTTP and HTTPS transport.

Note: Reconfiguring these options from Operations Manager results in a message instructing you to restart the HTTP service from the CLI. Use the following command to restart the HTTP service:
`dfm service start http.`

The DataFabric Manager server access to storage systems

There are a set of protocols and port numbers used by the DataFabric Manager server, to access your storage systems.

Protocol	UDP port	TCP port
HTTP		80

Protocol	UDP port	TCP port
HTTPS		443
RSH		514
SSH		22
Telnet		23
SNMP	161	

The DataFabric Manager server access to host agents

There are a set of protocols and port numbers used by the DataFabric Manager server to access your host agents.

Protocol	UDP port	TCP port
HTTP		4092
HTTPS		4093

The DataFabric Manager server access to Open Systems SnapVault agents

There are a set of protocols and port numbers used by the DataFabric Manager server to access your Open Systems SnapVault agents.

Protocol	UDP port	TCP port
HTTP		10000
SNMP	161	

SAN management

You can use the DataFabric Manager server to monitor and manage components—such as logical unit numbers (LUNs), Fibre Channel (FC) switches, and Windows and UNIX SAN hosts—of your storage area networks (SANs).

The SANs are storage networks that have been installed in compliance with the "SAN setup guidelines" by IBM. For information about setting up a SAN, see the *Data ONTAP Block Access Management Guide for iSCSI and FC*.

Note: DataFabric Manager server customers should check with their sales representative regarding other SAN management solutions.

Related information

[Data ONTAP Block Access Management Guide for iSCSI and FC - www.ibm.com/storage/support/nseries](http://www.ibm.com/storage/support/nseries)

Discovery of SAN hosts by the DataFabric Manager server

The DataFabric Manager server discovers SAN hosts with the DataFabric Manager Host Agent software installed on each SAN host.

The DataFabric Manager server can automatically discover SAN hosts; however, it does not use SNMP to poll for new SAN hosts. Instead, the special-purpose DataFabric Manager Host Agent software discovers, monitors, and manages SANs on SAN hosts. You must install the DataFabric Manager Host Agent software on each SAN host that you want to monitor and manage with the DataFabric Manager server.

The DataFabric Manager server communicates with the DataFabric Manager Host Agent software using HTTP or HTTPS. If both the DataFabric Manager server and the DataFabric Manager Host Agent software are configured to use HTTP, port 4092 is used for communication; however, if HTTPS is configured on both, port 4093 is used.

You can specify the protocol to use for communication in the DataFabric Manager server and when you install the DataFabric Manager Host Agent software on your SAN host. By default, both the Host Agent and the DataFabric Manager server are configured to use HTTP.

Note: If you choose to use HTTPS for communication between the DataFabric Manager server and a SAN host, ensure that both the DataFabric Manager server and the DataFabric Manager Host Agent software on the SAN host are configured to use HTTPS. If the Host Agent is configured to use HTTP and the DataFabric Manager server is configured to use HTTPS, communication between the SAN host and the DataFabric Manager server does not occur. Conversely, if the DataFabric Manager Host Agent software is configured to use HTTPS and the DataFabric Manager server is configured to use HTTP, communication between the two occurs, but HTTP is used for communication.

For more information about the DataFabric Manager Host Agent software, see the *DataFabric Manager Host Agent Installation and Administration Guide*.

Related information

[DataFabric Manager Host Agent Installation and Administration Guide - www.ibm.com/storage/support/nseries](http://www.ibm.com/storage/support/nseries)

SAN management using the DataFabric Manager server

By using the DataFabric Manager server, you can perform tasks such as viewing reports about LUNs, targets, and HBA ports. You can also group LUNs, storage systems in a SAN, FC switches, or SAN hosts, and so on.

Prerequisites for SAN management with the DataFabric Manager server

Different prerequisites apply for SAN management with the DataFabric Manager server.

For the DataFabric Manager server

You must have the SAN Management license key installed on your DataFabric Manager server.

All SAN monitoring and management features including LUN and FC switch monitoring are available only if you have the SAN Management license key installed.

If you do not have this license, contact your sales representative to find out how you can purchase one.

For IBM N series storage systems (targets)

The DataFabric Manager server does not report any data for your SAN if you do not have it set up according to the guidelines specified by IBM.

SAN deployments are supported on specific hardware platforms running Data ONTAP. For information about the supported hardware platforms, see the *SAN Configuration Guide*.

For FC switches

- To enable discovery of FC switches, the following settings must be enabled:
 - discoverEnabled (available from the CLI only)
 - Host Discovery (**Setup > Options > Edit Options: Discovery**)
 - SAN Device Discovery (**Setup > Options > Edit Options: Discovery**)
- The DataFabric Manager server can discover and monitor only FC switches, specifically Brocade Silkstorm switches, configured in a SAN setup as specified in the SAN Setup Overview for FCP guide.

Note: For a list of supported Brocade switches, see the SAN Configuration Guide at www.ibm.com/storage/support/nseries/.

- All FC switches to be managed by the DataFabric Manager server must be connected to a TCP/IP network either known to or discoverable by the DataFabric Manager server. The FC switches must be connected to the network through an Ethernet port and must have a valid IP address.

For SAN hosts (initiators)

- All SAN hosts to be managed by the DataFabric Manager server must be connected to a TCP/IP network either known to or discoverable by the DataFabric Manager server. The SAN hosts must be connected to the network through an Ethernet port and must each have a valid IP address.
- Each SAN host must have the DataFabric Manager Host Agent software installed on it. The DataFabric Manager Host Agent software is required for discovering, monitoring, and managing

SAN hosts. For more information about the Host Agent software, see the *DataFabric Manager Host Agent Installation and Administration Guide*.

- The Windows SAN hosts must have proper version of the SnapDrive software installed on it, for LUN management by using the DataFabric Manager server. To find out which SnapDrive version you must have installed, see the DataFabric Manager server software download pages at www.ibm.com/storage/support/nseries/.

Note: LUN management on UNIX SAN hosts by using the DataFabric Manager server is not currently available.

Related information

DataFabric Manager Host Agent Installation and Administration Guide - www.ibm.com/storage/support/nseries/

List of tasks performed for SAN management

You can perform different tasks for SAN management by using the DataFabric Manager server.

- View reports that provide information about all LUNs, targets, FC switches, SAN hosts, and HBA ports in a SAN.
- View details about a specific LUN, a target on a storage system, an FC switch, a SAN host, and an HBA port.
- Perform management functions such as configuring an FC switch, and creating, modifying, or expanding a LUN.
- Group LUNs, storage systems in a SAN, FC switches, or SAN hosts for efficient monitoring and management.
- Change the monitoring intervals for LUNs, FC switches, and SAN hosts.
- View SAN events and logs containing information related to LUN management and respond to SAN events.
- Configure the DataFabric Manager server to generate alarms to notify recipients of the SAN events.

List of user interface locations to perform SAN management tasks

You can perform SAN management tasks in the Operations Manager Control Center UI.

UI locations to perform SAN management tasks:

SANs tab (Control Center > Home > Member Details)

To view reports about all or a group of SAN components (LUNs, FC targets on a storage system, FC switches, and SAN hosts). You also use this tab to perform the following tasks:

- Access the details about a specific SAN component.
- Perform LUN and FC switch management functions.
- Create groups of LUNs, FC switches, and SAN hosts.

- Configure settings for SAN hosts such as the administration port, on which the DataFabric Manager server should poll the SAN host. You can also configure the type of protocol (HTTP or HTTPS) the DataFabric Manager server should use to communicate with SAN hosts.

Options link (Setup > Options) To enable and disable the discovery of SAN components and change the monitoring intervals for FC switches, LUNs, and SAN hosts in the DataFabric Manager server database

Events tab (Control Center > Home > Group Status > Events) To view and respond to SAN events

Alarms link (Control Center > Home > Group Status > Alarms) To configure alarms for SAN events

Reports for monitoring SANs

You can view various reports about the SAN components that the DataFabric Manager server manages, from the SANs tab.

Location of SAN reports

Reports about the SAN components that the DataFabric Manager server manages are available on the SANs tab. You can view reports by selecting reports in the Report drop-down list. If you want to view a report about a specific group, you specify the group by clicking the group name in the left pane of the DataFabric Manager server window.

You can view the following reports from the SANs page:

- LUNs, Comments
- Fibre Channel Switches, All
- Fibre Channel Switches, Deleted
- Fibre Channel Switches, Comments
- Fibre Channel Switches, Compact
- Fibre Channel Switches, Down
- Fibre Channel Switch Environmentals
- Fibre Channel Switch Locations
- Fibre Channel Switch Firmware
- Fibre Channel Switches, Up
- Fibre Channel Switches, Uptime
- Fibre Channel Switch Ports

- Fibre Channel Links, Physical
- Fibre Channel Links, Logical
- FCP Targets
- HBA Ports, All
- HBA Ports, FCP
- HBA Ports, iSCSI
- SAN Hosts, Comments
- SAN Hosts, All
- SAN Hosts, FCP
- SAN Hosts, iSCSI
- SAN Hosts, Deleted
- SAN Hosts Traffic, FCP
- SAN Host Cluster Groups
- SAN Host LUNs, All
- SAN Host LUNs, iSCSI SAN Host LUNs, FCP
- LUNs, All
- LUNs, Comments
- LUNs, Deleted
- LUNs, Unmapped
- LUN Statistics
- LUN Initiator Groups
- Initiator Groups

The DataFabric Manager server managed SAN data in spreadsheet format

You can put data from reports for SAN components that the DataFabric Manager server manages into a spreadsheet format.

When you view a report, you can bring up the data in spreadsheet format by clicking the spreadsheet icon on the right side of the Report drop-down list.

You can use the data in the spreadsheet to create your own charts and graphs or to analyze the data statistically.

Where to find information for specific SAN components

You can view information about SAN components from the Details page of a SAN component.

The Details page of a SAN component provides information specific to that component. For example, the FC Switch Details page provides the following details:

- Firmware version
- Uptime
- Status of each port on the switch

The LUN Details page provides the following details:

- The status and size of a LUN
- Events associated with the LUN
- All groups to which the LUN belongs

In addition, you can obtain graphs of information about the SAN components, access the management tools for these components, and view events associated with these components.

Where to view LUN details of SAN components

You can view LUN details of SAN components using the DataFabric Manager server.

You can access the LUN Details page for a LUN by clicking the path of the LUN in any of the reports.

Tasks performed from the LUN Details page for a SAN host

From the LUN Details page, you can diagnose connectivity, expand a LUN, destroy LUN, edit settings, and so on.

The Tools list on the LUN Details page enables you to select the following tasks:

- Diagnose Connectivity
- Expand this LUN—Launches a wizard that helps you to expand the LUN.
- Destroy this LUN—Launches a wizard that helps you to destroy the LUN.
- Edit Settings
- Refresh Monitoring Samples—Obtains current monitoring samples from the storage system on which this LUN exists
- FilerView—Launches FilerView, the Web-based UI of the storage system on which the LUN exists.
- Manage LUNs with FilerView—Launches FilerView and displays a page where you can manage the LUN.
- Run a Command—Runs a Data ONTAP command on the storage system on which this LUN exists.

Note: You must setup the appropriate authentication to run commands on the storage system.

Information about FC target on a SAN host

You can view FCP target details of a SAN host from the FCP Target Details page.

The FCP Target Details page contains the following information:

- Name of the FC switch and the port to which the target connects
- Status of the target
- Name of the storage system on which target is installed
- Port state of the target. Port state can be one of the following:
 - Startup

- Uninitialized
- Initializing Firmware
- Link Not Connected
- Waiting For Link Up
- Online
- Link Disconnected
- Resetting
- Offline
- Offline By User System
- Unknown
- Specifics about the target such as hardware version, firmware version, and speed of the target
- FC topology of the target. Topology can be one of the following:
 - Fabric
 - Point-To-Point
 - Loop
 - Unknown
- WWNN and WWPNN of the target
- Other FCP targets on the storage system on which the target is installed
- Time of the last sample collected and the configured polling interval for the FCP target

Information about the FC switch on a SAN host

You can view FCP switch details of a SAN host and the status of the switch. You also view events associated with the switch, and the number of devices connected with the switch.

Access to the FC Switch Details page

You can access FC Switch Details page from SAN monitoring reports.

You can access the FC Switch Details page for a switch by clicking the name of the switch in any of the reports for SAN monitoring.

Information about FC Switch on a SAN host

You can view FC Switch details of a SAN host from the FC Switch Details page.

- Status of the switch
- Firmware version installed on the switch
- The length of time that the switch has been up since the last reboot
- Contact information for the switch such as the administrator name and location of the switch
- Events associated with the FC switch
- FC switch port status-- A graphical layout of the switch ports, with the color of the switch port that indicates the status of the port:
 - Green—indicates that the port is online and working normally.

- Yellow—indicates that the port is connected to a GBIC, but is not synchronized (No Sync).
- Red—indicates that the port is offline or not working normally.
- Black—indicates that there is no GBIC connected.
- Number of devices connected to the switch and a link to a report that lists those devices
- The DataFabric Manager server groups to which the FC switch belongs
- Time when the last data sample was collected and the configured polling interval for the switch
- Graph of FC traffic per second on the switch.
You can view the traffic over a period of one day, one week, one month, one quarter (three months), or one year.

Tasks performed from the FC Switch Details page for a SAN host

You can edit the FC Switch settings, refresh monitoring samples, run FabricSwitch, and run Telnet using the FC Switch Details page.

You can use the Tools list on the FC Switch Details page to select the tasks to perform for the switch whose Details page you are on. The tasks are as follows:

- **Edit Settings**—Displays the Edit FC Switch Settings page where you configure the login and password information in the DataFabric Manager server for the switch. The DataFabric Manager server requires the login and password information to connect to a switch using the Telnet program.
- **Refresh Monitoring Samples**—Obtains current monitoring samples from the FC switch.
- **Invoke FabricWatch**—Connects you to FabricWatch, the Web-based UI, of the Brocade switch. You might want to connect to FabricWatch to manage and configure the switch.
- **Run Telnet**—Connect to the CLI of the switch using the Telnet program.
The DataFabric Manager server requires the login and password information to connect to a switch using the Telnet program.

Information about DataFabric Manager Host Agent software on a SAN host

You can view the SAN host status, version details of the operating system and Host Agent, number of HBAs, and so on in the Host Agent Details page.

The Details page for a SAN host contains the following information:

- Status of the SAN host and the time since the host has been up
- The operating system and DataFabric Manager Host Agent software version running on the SAN host
- The SAN protocols available on the host
- The MSCS configuration information about the SAN host, if any, such as the cluster name, cluster partner, and cluster groups to which the SAN host belongs
- The events that occurred on this SAN host
- The number of HBAs and HBA ports on the SAN host
- The devices related to the SAN host such as the FC switches connected to it and the storage systems accessible from it

- The number of LUNs mapped to the SAN host and a link to the list of those LUNs
- The number of initiator groups that contain the SAN host and a link to the list of those initiator groups
- Time of the last sample collected and the configured polling interval for the SAN host
- Graphs of information such as the HBA port traffic per second or the HBA port frames per second over a period of one day, one week, one month, one quarter (three months), or one year

Accessing the HBA Port Details page for a SAN host

You can access the HBA Port Details page for a SAN host from the SAN host reports.

Steps

1. Click **Member Details > SANs > Report drop-down > HBA Ports, All**.
2. Click the name of the HBA port.

The HBA Port Details page is displayed.

Details about the HBA Port Details page

You can use the HBA Port Details page to view the status of HBA port, view the HBA protocol, view events that occurred in the HBA port, and so on.

The Details page for an HBA port contains the following information:

- Status and state of the HBA port
- Name of the SAN host on which the HBA is installed
- The protocol available on the HBA
- Specifics about the HBA such as the name, serial number, model, hardware version, driver version, and firmware version
- WWNN of the HBA and WWPNN of the HBA port
- FC switch port to which the HBA port connects
- The events that occurred on this HBA port
- The number of HBA ports on the HBA and a link to the list of those ports
- The number of HBA ports on the SAN host on which the HBA port exists and a link to the list of those ports
- The number of storage systems accessible to the HBA port and a link to the list of those storage systems
- The number of LUNs mapped to the HBA port and a link to the list of those LUNs
- The number of initiator groups that contain the HBA port and a link to the list of those initiator groups
- Time of the last sample collected and the configured polling interval for the HBA port
- Graphs of information such as the HBA port traffic per second or the HBA port frames per second over a period of one day, one week, one month, one quarter (three months), or one year

List of SAN management tasks

You can perform various SAN management tasks by using the DataFabric Manager server.

- For LUNs: Create, expand, and destroy LUNs; map a LUN to or unmap a LUN from initiator groups.
- For initiator groups: Create or delete initiator groups.
- For FC switches: Configure and view the current configuration of a switch.

LUN management

You can manage a LUN in two ways by using the DataFabric Manager server.

- Use the LUN management options available in the DataFabric Manager server:
 - The Host Agent Details page provides a `Create a LUN` option in the Tools list to create LUNs.
When you select the `Create a LUN` option, a wizard is launched that takes you through the process of creating a LUN.
 - The LUN Details page provides the two LUN management options in the Tools list: `Expand this LUN` and `Destroy this LUN`.
These LUN management options launch wizards specific to their function. The wizards take you through the process of expanding or destroying a LUN.

By using a wizard available in the Tools list on the Host Agent Details page and LUN Details page, you can create, expand, and destroy a LUN. Before you run the wizard, ensure the following:

- The SAN host management options are appropriately set on the Options page or the Edit Host Agent Settings page.
- To manage a shared LUN on an MSCS, perform the operation on the active node of the cluster. Otherwise, the operation fails.

Initiator group management

You can manage initiator groups by connecting to FilerView, the Web-based management interface, of your storage system. The LUN Details page provides a Manage LUNs with FilerView option in the Tools list. This option enables you to connect to FilerView.

FC switch management

The FC Switch Details page provides the `Invoke FabricWatch` option in the Tools list. You can use this option to connect to FabricWatch, the Web-based management tool for the Brocade SilkWorm switches.

DataFabric Manager server options

The DataFabric Manager server uses the values of options in its database to determine whether to automatically discover new objects or not. The DataFabric Manager server determines, how frequently to monitor objects, and what threshold value to use to generate an event.

Note: The DataFabric Manager server object is an entity that is monitored or managed by the DataFabric Manager server. Examples of the DataFabric Manager server objects are storage systems, LUNs, FC switches, and user quotas.

When the DataFabric Manager server is installed, these options are assigned default values; however, you can change these values. The options can be changed globally—to apply to all objects in the DataFabric Manager server database, or locally—to apply to a specific object or a group of objects in the DataFabric Manager server database. Some options can be set globally, but not locally.

When both global and local options are specified for an object, the local options override the global options.

DataFabric Manager server options for SAN management

There are two DataFabric Manager server options for SAN management, namely, Global-only option and Global and local options.

Global-only options:

- SAN Device Discovery
This option specifies whether to enable or disable the automatic discovery of SAN components (LUNs, FC switches, SAN hosts).
By default, this option is enabled.
- LUN monitoring interval, Fibre Channel monitoring interval, and SAN Host monitoring interval
The monitoring intervals determine how frequently the DataFabric Manager server collects information about an object.
Following are the default monitoring intervals:
 - For LUNs, 30 minutes
 - For Fibre Channel, 5 minutes

- For SAN Host, 5 minutes

Global and local options:

- **Host Agent Login**
This option specifies the user name that is used to authenticate to the DataFabric Manager Host Agent software, for SAN monitoring and management.
By default, SAN monitoring is enabled; therefore, the user name guest is used.
If you want to enable SAN management in addition to monitoring, you must select the user name admin.
- **Host Agent Monitoring Password**
This option specifies the password that is used for the user name guest to authenticate to the Host Agent software for SAN monitoring.
By default, public is used as the password; however, you can change it. If you change the password in the DataFabric Manager server, ensure that you change the password to the same value in the Host Agent software running on the SAN hosts. Otherwise, the DataFabric Manager server cannot communicate with the SAN host.
- **Host Agent Management Password**
This option specifies the password that is used for the user name admin to authenticate to the DataFabric Manager Host Agent software for SAN monitoring and management.
There is no default value for the management password. Therefore, you must specify a value for this option before you can use the LUN management features through the DataFabric Manager server. The password you specify for this option must match the password specified in the Host Agent software running on the SAN hosts. Otherwise, the DataFabric Manager server cannot communicate with the SAN host.
- **Host Agent Administration Transport**
This option specifies the protocol, HTTP or HTTPS, used to connect to the DataFabric Manager Host Agent software.
By default, this option is set to HTTP.
- **Host Agent Administration Port**
This option specifies the port that is used to connect to the Host AgentDataFabric Manager Host Agent software.
By default, 4092 is used for HTTP and 4093 for HTTPS.
- **HBA Port Too Busy Threshold**
This threshold specifies the value, as a percentage, at which an HBA port has so much incoming and outgoing traffic that its optimal performance is hindered.
If this threshold is crossed, the DataFabric Manager server generates an HBA Port Traffic High event. By default, this threshold is set to 90 for all HBA ports.

Where to configure monitoring intervals for SAN components

You can configure the global options on the Options page.

To configure options locally (for a specific object), you must be on the Edit Settings page of that specific object (**Details page > Tools list > Edit Settings**).

Deleting and undeleting SAN components

You can stop monitoring a SAN component (a LUN, an FC switch, a storage system, or a SAN host) with the DataFabric Manager server by deleting it from the Global group.

When you delete a SAN component, the DataFabric Manager server stops collecting and reporting data about it. Data collection and reporting is not resumed until the component is added back (using the Undelete button) to the database.

You cannot stop monitoring a specific FC target or an HBA port unless you stop monitoring the storage system or the SAN host on which the target or the port exists.

Note: When you delete a SAN component from any group except Global, the component is deleted only from that group; the DataFabric Manager server does not stop collecting and reporting data about it. You must delete the SAN component from the Global group for the DataFabric Manager server to stop monitoring it.

Reasons for deleting and undeleting SAN components

You might want to delete a SAN component if you want to stop monitoring the SAN component either temporarily or permanently. You might want to undelete a SAN component if you want to resume the monitoring of the component that you previously deleted.

Temporarily stop monitoring SAN component

You might want to delete a SAN component if you want to perform maintenance on the component and do not want the DataFabric Manager server to generate events and alarms (if configured) during the maintenance process.

Permanently stop monitoring the SAN component

You might want to delete a component if it exists on a non-mission critical network and does not need to be monitored, but it has been discovered by the DataFabric Manager server. A non-mission critical network could be a laboratory network.

Process of deleting SAN components

You can delete a SAN component from any of the reports related to that component.

First, you select the components you want to delete in the left-most column of a report. Then, you click **Delete** at the bottom of each report to delete the selected components.

Process of undeleting SAN components

All deleted objects are listed in their respective Deleted reports. You can delete a SAN component from these reports.

All deleted FC switches are listed in the FC Switches, Deleted report. All deleted LUNs are listed in the LUNs, Deleted report. All deleted SAN hosts are listed in the SAN hosts, Deleted report.

You can undelete an object by selecting it from its Deleted report and clicking **Undelete** at the bottom of the report.

How SAN components are grouped

You can group SAN components—LUNs, storage systems, SAN hosts, or FC switches—to manage them easily and to apply access control.

Storage systems, SAN hosts, and FC switches are considered storage systems for the purpose of creating groups. Therefore, when you create a group of storage systems, SAN hosts, or FC switches, the type of the created group is “Appliance resource group.” In addition, you can add storage systems, SAN hosts, and FC switches to an Appliance resource group.

However, a group containing LUNs cannot contain any other objects—storage systems, SAN hosts, or FC switches. Therefore, when you create a group of LUNs, the created group is “LUN resource group.”

Note: You cannot group HBA ports or FCP targets.

Restriction of SAN management access

You can allow an administrator to manage your SAN hosts and devices, by selecting the **GlobalSAN** role on the Administrators page (**Setup menu > Administrative Users**). The GlobalSAN role allows an administrator to create, expand, and destroy LUNs. This role can be global or per-group.

Access control on groups of SAN components

Just as with other DataFabric Manager server groups, you can apply access control on groups of SAN components to restrict the scope of tasks an administrator can perform.

You can apply access control to a group (group-based access control) by assigning a role to an administrator account for that group. The role defines the administrator’s authority to view and perform operations on the components in that group.

Note: By default, the DataFabric Manager server is configured to provide global Read access across all groups, including Global, to user Everyone.

Glossary

- associated view** A performance view available for an object.
An associated view can be a canned view or a custom view. Canned views are associated to an object based on the object type. Custom views can be associated to an object by the user.
- cluster** A cluster refers to a group of connected nodes (storage systems) that share a global namespace and that you can manage as a single virtual server or multiple virtual servers, providing performance, reliability, and scalability benefits.
- counter** A statistical measurement of activity on a storage system or storage subsystem that is provided by Data ONTAP.
- controller** Controller also known as storage controller refers to the component of a storage system that runs the Data ONTAP operating system and controls its disk subsystem. Storage controllers are also sometimes called storage appliances, appliances, storage engines, heads, CPU modules, or controller modules.
- custom views** User-defined views that includes metrics from specific instances.
- dataset** A collection of storage sets along with configuration information associated with data. The storage sets associated with a dataset include a primary storage set used to export data to clients, and the set of replicas and archives that exist on other storage sets. Datasets represent exportable user data.
- hierarchical groups** The customizable grouping of storage systems and vFiler units whose performance can be monitored and displayed through Performance Advisor. All systems that are monitored by the DataFabric Manager server and whose performance can be monitored by Performance Advisor are displayed in the Groups browser panel of Performance Advisor.
- historical data** Data that is archived by the performance-monitoring server on the DataFabric Manager server. All the data that is included in the Performance Advisor default views is also archived as historical data. The historical data is accessible to any Performance Advisor that can connect to the workstation. Historical data is collected on an ongoing basis, independent of whether a client has the associated performance view open or not. Historical data can be used for diagnosing past performance problems or for short-term trend analysis.
- hosting storage system** The physical storage system on which one or more vFiler units are configured. Some counters that Performance Advisor tracks apply to both storage systems and vFiler units. Other counters (for example, CPU usage) apply only to storage systems and the associated host of a vFiler unit. The hosting storage system is also referred to as the “vFiler host.”

logical objects	Object types that represent storage containers such as volume, qtree, LUNs, vFiler units, and datasets are known as Logical Objects.
logical hierarchy	The hierarchy that displays the logical objects and instances.
lower threshold	The type of threshold set for an event generation when the counter value falls and remains below the lower threshold value for longer than the Threshold Interval specified.
managed object	Any object that has an identity and a name in the DataFabric Manager server object table. A managed object is an object that is contained within a DataFabric Manager server group. Volumes, aggregates, qtrees, and LUNs are examples of managed objects.
Management Console	The client platform for the Java-based applications. Management Console runs on a Windows or Linux workstation, separate from the system on which the DataFabric Manager server is installed.
object	Typically there is an object associated with each hardware or software subsystem within Data ONTAP. Examples of hardware objects are processor, disk, NVRAM, and networking card objects. FCP, iSCSI, CIFS, and NFS are the examples of software protocol objects. WAFL, RAID and target are the examples of internal objects specific to Data ONTAP. Virtual objects like the system object capture key statistics across all the other objects in one single place. Examples of system objects are avg_processor_busy, nfs_ops, cifs_ops, and net_data_recv.
Operations Manager	Operations Manager is the Web-based user interface of the DataFabric Manager server, from which you can monitor and manage multiple storage systems, and clusters on systems running Data ONTAP 8.0 Cluster-Mode. Operations Manager is used for day-to-day monitoring, alerting, and reporting on storage infrastructure.
Performance Advisor	The Performance Advisor component installed on the Management Console platform enables you to monitor the performance of storage systems and vFiler units. The user interface of Performance Advisor contains only performance-monitoring information. This Performance Advisor interface is distinct from Operations Manager, which contains other DataFabric Manager server information.
performance-monitoring server	The Performance Advisor component that is enabled on the DataFabric Manager server to collect and archive performance data at regular intervals from the monitored storage systems and vFiler units. Performance Advisor gathers sets of data that the performance-monitoring server has collected to generate its graphical charts and views.
performance view	A collection of one or more counters accessible through Performance Advisor for a group of storage systems or subsystems through one or more charts.

	Performance Advisor displays canned views that are built in to the system, and custom views.
physical objects	Object types that represent the physical resources in a storage system such as disk, aggregates, memory, network interfaces, resource pools, and RAID groups are known as Physical Objects.
physical hierarchy	The hierarchy that displays the physical objects and instances.
real-time data	Data that is passed through Performance Advisor for display but is not stored. Real-time data is collected only as long as the performance view window displaying it is open on Performance Advisor. Real-time data is accessible only to the client by which it is retrieved. Real-time data is suitable for diagnosing immediate performance issues.
resource pool	A managed object in the DataFabric Manager server, containing storage provisioning resources like storage systems, aggregates, and spare disks.
retention period	The amount of previously collected data for a counter that is accessible to the user. You can set the retention period based on units of a minute, hour, day, week, month, or year.
sample rate	The rate at which data is collected for a counter. You can set the sample rate based on units of a minute.
storage set	Containers that are used for delegation, replication, and in some cases, substorage provisioning. The only container of merit in a storage set is a volume (flexible or traditional). A storage set contains a group of volumes whereas a volume should be in at most one storage set.
storage system	An appliance that is attached to a computer network and is used for data storage. IBM N series storage systems and NearStore systems are examples of storage systems.
templates and canned templates	A template is a view definition that applies to object instances known to Performance Advisor. A set of pre-configured templates in Performance Advisor is known as Canned Template.
threshold interval	The amount of time in seconds for which an event generation is suppressed before Performance Advisor decides that a counter has crossed a specified threshold and an event needs to be generated. The same interval is also be used to generate a normal event.
unmanaged object	Objects apart from the managed objects belong to the class of unmanaged objects. An unmanaged object does not have a unique identity in the DataFabric Manager server table.

- upper threshold** The type of threshold set for an event generation when the counter value exceeds and remains above the higher threshold value for longer than the Threshold Interval specified.
- vFiler unit** One or more virtual storage systems that can be configured on a single physical storage system licensed for the MultiStore feature. DataFabric Manager server 3.4 and later enables monitoring and management of vFiler units.
- view** A collection of related panels represented together and displayed by the Performance Advisor client.

Copyright and trademark information

Copyright ©1994 - 2011 NetApp, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2011 IBM Corporation. All rights reserved.

US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

No part of this document covered by copyright may be reproduced in any form or by any means— graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

References in this documentation to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's or NetApp's intellectual property rights may be used instead of the IBM or NetApp product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM and NetApp, are the user's responsibility.

No part of this document covered by copyright may be reproduced in any form or by any means— graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT

(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S.A. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, FAServer, FastStak, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service

Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe on any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, N.Y. 10504-1785
U.S.A.

For additional information, visit the web at:
<http://www.ibm.com/ibm/licensing/contact/>

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM web sites are provided for convenience only and do not in any manner serve as an endorsement of those web sites. The materials at those web sites are not part of the materials for this IBM product and use of those web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Index

- ## A
- access check, for application administrator 45
 - access control
 - configuring
 - vFiler units 45
 - on groups of SAN components 319
 - precedence of, global and group 52
 - access roles
 - manage configuration files 230
 - accessing CLI 265
 - accessing DataFabric Manager server
 - CLI 265
 - Active Directory
 - user group accounts 48
 - active/active configuration
 - cluster console, accessing 220
 - managing with DataFabric Manager server 220
 - add
 - administration access for host agents 160
 - administration passwords, host agents 156, 157
 - automapping SRM paths 164, 165
 - SRM paths 160
 - adding clusters 40
 - adding primary, secondary storage systems 241
 - Administration Port option 147
 - Administration Transport option 147
 - Administrator access
 - managing 55
 - administrators
 - creating accounts 55
 - default accounts
 - everyone account 46
 - types of access controls 55
 - Aggregate Details page 189, 190
 - aggregate full threshold 190
 - Aggregate Nearly Full threshold 190
 - aggregate nearly overcommitted threshold 190
 - aggregate overcommitted threshold 190
 - Aggregate Overcommitted threshold 190
 - aggregates
 - capacity graphs 190
 - chargeback (billing). See storage chargeback 205
 - historical data 190
 - name format 190
 - relation to traditional volume 194
 - aggregates space availability 189
 - alarm notification
 - customize 116
 - alarms
 - acknowledging 118
 - alerts and alarms, differences 119
 - configuration guidelines 114
 - creating 115
 - e-mail messages 118, 122
 - SNMP traps as 86
 - testing 115
 - alerts
 - See user alerts 119
 - annual charge rate, storage chargeback reports 209
 - Appliance Details page
 - tasks 216
 - appliance management. See appliances, managing 213
 - appliances
 - commands, running 222
 - configuration changes 222
 - console, connecting 219
 - grouping 214
 - managing, administrative access 216
 - archive backup 272
 - archived reports
 - Report Archival Directory 131
 - assign parent groups 234
 - authenticate controller 220
 - authentication
 - LDAP 21
 - methods 21
 - requirements
 - NDMP 245
- ## B
- backing up
 - access requirements 274
 - deleting from Web-based user interface 276
 - directory location 272
 - disabling schedules 276
 - displaying information about 276
 - restoring from backup file 277
 - starting from Web-based user interface 274
 - storage and sizing 273
 - backup
 - requirements 244

- retention copies 244
- Snapshot copies 244
- backup management
 - hot mode for databases 238, 271
 - scripts, pre and postbackup 238, 271
- Backup Manager
 - adding primary storage systems 242
 - adding secondary storage systems 241
 - adding secondary volumes 242
 - authentication 239
 - discovering 239
 - discovery process 238
 - initial setup 240
 - NDMP
 - SNMP 239
 - primary directory
 - format 250
 - secondary volume
 - format 251
 - select primary directories, qtrees 243
 - SnapVault relationships 240
- backup relationships
 - bandwidth limitation 247
 - baseline transfer 247
 - configuring bandwidth
 - limit 247
 - create 243
- backup schedules
 - creating 244
 - Snapshot copies 243
- BackupManager reports 132
- billing cycle, storage chargeback reports 209
- billing reports. See storage chargeback 205
- Business Continuance license 238
- Business Continuance Management license 253

C

- CA-signed certificate 143
- capacity 83
- capacity reports 153
- capacity thresholds 199
- catalogs
 - See reports 124
- Certificate Authority (CA)-signed 143
- Certificate Signing Request 143
- certificate types 143
- certificates
 - CA-signed 144
 - definition of CA-signed 143

- obtaining 144
- role of, in secure connections 143
- self-signed, generating 144
- self-signed, security precautions for 143
- certificates in DataFabric Manager server 143
- CIFS, SRM requirements 163
- CLI 265
- CLI, accessing 265
- Clone List report 204
- clone volumes
 - Clone List report 204
 - identifying 204
 - parents, relation to 203
- cluster components
 - aggregates 80
 - controllers 80
 - grouping 80
 - virtual servers 80
 - volumes 80
- cluster details
 - cluster objects 83
 - cluster reports 83
 - information available 83
- cluster discovery
 - automatic discovery 39
- cluster management
 - browse cluster components 84
 - capabilities 56
 - Cluster Details page 84
 - cluster hierarchy information about cluster objects
 - CPU usage 84
 - cluster objects 56
 - controlled user access 56
 - gather cluster information 83
 - logical interface 83
 - login interface 83
 - network traffic 84
 - tasks performed 84
 - view physical system 84
 - view utilization of physical resources 84
 - volume capacity used 84
- cluster mode 40
- cluster monitoring
 - features supported 41
 - limitations 41
 - using APIs 83
 - using SNMP 83
- Cluster Tools 218
- cluster-related objects 78
- clusters 78

- Clusters 125
 - configuration details of
 - storage arrays 42
 - configuration files
 - acquired from parent groups 234
 - compare configurations 231
 - manage 229
 - prerequisites, for storage systems and vFiler Units 229
 - properties 234
 - pull 231
 - tasks to manage 230
 - template configuration 231
 - verify successful push 231
 - configuration group
 - remove configuration file 232
 - tasks to manage 232
 - configuration management
 - storage system 229
 - configuration plug-ins
 - Data ONTAP plug-in 231
 - configuration report 42
 - configuration resource groups
 - assign parent group 233
 - configuration settings 232
 - configuration using DataFabric Manager server
 - data protection 237
 - configure alarms 84
 - configuring multiple storage systems 235
 - connect storage system
 - Run Telnet tool 219
 - considerations
 - assigning parent group
 - inherit, configuration files of parent group 234
 - console access to DataFabric Manager server 265
 - Controllers 125
 - CPU usage 83
 - Create a LUN wizard 183
 - CSR 143
 - currency format, storage chargeback reports 208
 - custom comment fields
 - alarm notification by e-mail 116
 - script format 117
 - trap format 117
- D**
- database backup
 - archive backup 272
 - process 272
 - Restore process 277
 - scheduling 275
 - Snapshot backup 272
 - types 272, 275
 - database scripts, pre and postbackup 238, 271
 - DataFabric Manager
 - groups
 - objects 71
 - logging in to 46
 - restoring backup file 277
 - DataFabric Manager Host Agent software
 - administration access 159
 - capabilities 177
 - communication ports 156
 - overview 156
 - passwords 156, 157
 - DataFabric Manager server
 - access to host agents 305
 - access to Open Systems SnapVault agents 305
 - access to storage systems 304
 - communication 304
 - dfm backup command (CLI) 277
 - HTTP, HTTPS 304
 - protocols and port numbers 304, 305
 - DataFabric Manager server authentication 220
 - DataFabric Manager server options
 - global and local options 316
 - objects 316
 - Day of the Month for Billing option 209
 - default backup schedule 244
 - default role 49
 - delete
 - SRM paths 164
 - deleted objects
 - what happens when storage objects are deleted 210
 - deleting
 - SAN components 185
 - determine disk space 171
 - directories
 - not backed up
 - viewing 239
 - directory path for archive backups
 - changing 274
 - Disaster Recovery Management
 - policy management 256
 - Disaster Recovery Manager
 - connection management 257
 - defined 253
 - managing SnapMirror relationships 253
 - storage system authentication 259

- tasks, general 253
- discoveries
 - new, qtrees 239
 - Open Systems SnapVault monitor 239
- discovery
 - DataFabric Manager Host Agent software 177
 - SAN host 177
- discovery of
 - storage array ports 41
 - storage arrays 41
- Disk space hard limit option 173
- Disk space soft limit option 173
- Domain users
 - Pushing 64
 - Viewing settings 63

E

- edit
 - SRM paths 164
- Edit Quota Settings page 174
- Event Reports 132
- events
 - aggregate almost full 190
 - aggregate almost overcommitted 190
 - aggregate full 190
 - aggregate overcommitted 190
 - clearing configuration events 90
 - definition of 89
 - list of, complete 90–96, 98–107, 109–113
 - managing 90
 - qtree full 201
 - qtree nearly full 201
 - user, notification of 119
 - viewing 89
 - volume almost full 194
- Everyone, administrator account 46
- example
 - assigning parent groups 234

F

- failover mode, of multiple paths 258
- FC switch
 - managing 316
- FCP Target Details page 182
- FCP targets 182
- FCP topology 182
- File Storage Resource Management (FSRM)
 - See FSRM
- FilerView
 - configuring storage systems 225
 - links 85

- Files hard limit option 173
- Files soft limit option 173
- FSRM
 - monitoring requirements 153
 - what FSRM does 153

G

- Gateway SAN-attached storage management 41
- giveback 221
- GlobalDataProtection 49
- global access control
 - applying to administrator accounts 55
 - description of 55
 - precedence over group access control 49
- global and local options
 - HBA Port Too Busy Threshold 316, 317
 - Host Agent Administration Port 316, 317
 - Host Agent Administration Transport 316, 317
 - Host Agent login 316, 317
 - Host Agent Management Password 316, 317
 - Host Agent Monitoring Password 316, 317
 - SAN management 316, 317
- Global Delete 230
- global groups 73
- global privileges
 - creating accounts 55
- Global Read 230
- Global SAN role 49
- global thresholds
 - setting 246
- Global Write 230
- Global-only options
 - LUN monitoring interval 316, 317
 - managing SAN 316, 317
 - SAN device discovery 316, 317
- global, group information
 - location 215
- GlobalBackup role 49
- GlobalDataSet 49
- GlobalDelete role 49
- GlobalDistribution role 49
- GlobalEvent role 49
- GlobalFullControl role 49
- GlobalMirror role 49
- GlobalPerfManag 49
- GlobalQuota role 49
- GlobalRead role 49
- GlobalRestore role 49
- GlobalSRM role 49
- GlobalWrite role 49

- graphs
 - LUNs 181
- Graphs
 - SAN hosts 183
- group
 - adding members 76
- group access control
 - applying 55
 - description 55
 - precedence over global access control 49
- group access privileges
 - accounts 55
- group membership 76
- groups
 - changing thresholds 77
 - definition of 71
 - global 73
 - guidelines for creating 75
 - managing 75
 - reports (views) about 77

H

- hierarchical groups 73
- host agent
 - administration access 159
 - passwords 159
- Host Agent Details page 183
- host agents
 - types of 156
- host discovery 34
- hosts.equiv 150
- hosts.equiv file 150
- hosts.equiv option 147
- hot backup mode 238, 271
- HTTPS
 - enabling 145
 - requirements for 146

I

- initiator group
 - connecting to FilerView 316
 - managing 316
 - viewing LUNs mapped to 181
- installing
 - licenses 240
- interface groups 78
- Interface groups 125

J

- jobs 69, 70

- junctions 78

L

- lag thresholds
 - SnapMirror 263
- LDAP
 - authentication 21
- license key
 - Business Continuanance Management 253
 - SRM 154
- licenses, permissions 240
- LIFs
 - cluster management LIF 78
 - data LIF 78
 - node management LIF 78
- limitations
 - managed host options 149
- list of storage array ports connected to the gateway 42
- list of storage arrays connected to the gateway 42
- local threshold
 - setting 247
- local users 57
- Local users
 - Adding 58, 61
 - Deleting 61
 - Editing passwords 62
 - Pushing 61
 - Pushing passwords 60
 - Viewing settings 57
- Logical interfaces 125
- logical interfaces (LIFs)
 - See LIFs
- logical resources
 - logical interface traffic 84
 - utilization 84
 - virtual servers 84
- logical storage
 - monitoring
 - qtree capacity threshold and events 201
 - qtree quotas 203
 - volume capacity threshold and events 194
- login and password
 - SAN hosts 183
- Login Protocol option 147
- LUN
 - connecting to FilerView 315
 - managing 315
- LUN Details page 181, 182
- LUNs

- creating 183
- deleting and undeleting 185
- destroying 182
- expanding 182
- graphs 181
- initiator groups mapped to 181
- reports in spreadsheet form 131
- status of 181
- stop monitoring 185

M

- mailformat file 122, 123
- mailmap file 120
- manage discovered relationships
 - enable DataFabric Manager server 246
- managed host options
 - guidelines for changing 147
 - overview 147
- Managed host options
 - limitations 149
- managing administrator access
 - domain users 62
- Mirror Reports 132
- modification of passwords
 - issue 150
- modify
 - cluster controller details 218
 - cluster details 218
- monitoring
 - flow chart of process 81
 - process 81
 - query intervals 86
- monitoring of clusters
 - cluster discovery 40
 - configuration 40
 - Data ONTAP 8.0 40
 - tasks performed 40
- monitoring options
 - guidelines for changing 86
 - location to change 86
- monitoring passwords, host agents 156, 157
- multiple paths
 - failover mode 258
 - multiplexing mode 258
- MultiStore 226

N

- namespaces 78
- NDMP service
 - enabling 241

- network settings 39
- node mode 40

O

- Open Systems SnapVault hosts 239
- Operations Manager
 - deleting a backup from 276
 - starting a backup from 274
- options
 - Disk space hard limit 173
 - Disk space soft limit 173
 - Files hard limit 173
 - Files soft limit 173
- order of applying privileges 52
- overcommitment strategy 190
- Overview 55

P

- parent group 80
- passwords for host agents 156, 157
- paths
 - adding paths 162
 - CIFS requirements 163
 - CLI quoting conventions 163
 - UNC requirements, 163
 - valid path formats 162
- paths, managing
 - automapping 164, 165
 - deleting 164
 - editing 164
 - quick reference 161
 - viewing details, directory-level 164
- Permanently stop monitoring 318
- physical resources
 - controller 85
 - CPU 85
 - network traffic 85
 - utilization 85
- physical storage
 - monitoring
 - aggregate capacity thresholds and events 190
- Policy Reports 132
- Ports 125
- potential length, parent chains 234
- predefined roles 47
- preferred SNMP version option 40
- prerequisites
 - editing user quotas 174

- managing user quotas 172
- managing, SnapMirror relationships 253
- SRM components 154
- system backup 238
- preset thresholds 84
- push jobs 235

Q

- quick reference for tasks 157
- quiesced SnapMirror relationships 262
- quotas
 - process 171
 - why you use 171
- Quotas tab renamed to SRM 154

R

RBAC

- access 53
- how roles relate to administrators 49

RBAC resource 53

remote configuration

- accessing CLI 223
- cluster 223
- Operations Manager UI 223

remote platform management

- RLM card management 269
- system requirements 269

Remote Platform Management

- configuring RLM card IP address 270

Remote Platform Management IP Address 224

report categories

- capabilities 132

report generation

- scheduling 131

report outputs

- list 137

report schedule

- deleting 133
- disabling 134
- editing 133
- listing results 135
- retrieving 134
- running 134

Report Schedules report

- existing report schedules 132

reports

- aggregate 125
- aggregate array LUNs 125

- array LUNs 125
- catalogs 124
- configuring 129
- controller 125
- customizing options 124
- dataset 125
- deleting custom 130
- disks 125
- events 125
- FC link 125
- FC switch 125
- FCP target 125
- file system 125
- group summary 125
- groups 77
- History performance reports 125
- Host domain users 125
- Host local users 125
- Host roles 125
- Host usergroups 125
- Host users 125
- LUN 125
- performance events 125
- putting data in spreadsheets 131
- Quotas 125
- Report outputs 125
- Report schedules 125
- resource pools 125
- SAN hosts 125
- schedules 125
- scripts 125
- spare array LUNs 125
- SRM 125
- storage systems 125
- types 123
- user quotas 125
- vFiler 125
- viewing, tasks performed 215
- volume 125

requirements

- SRM 154
- SRM license requirements 172

resource groups

- global 73

resource utilization 84

restoring DataFabric Manager database 277

restricting access to SRM data 166

roles

- create 52
- global 49

- GlobalWrite
 - GlobalDelete 229
 - group 51
 - modify 52
 - order applying 52
 - Roles
 - configuring 67
 - run commands
 - on a cluster 224
 - run commands on storage systems
 - Run a Command tool 218, 219
 - running commands
 - specific cluster node 224
 - specific node of a cluster 224
 - using RLM card 224
 - running commands on storage systems 223
- S**
- SAN components
 - access control 319
 - configuring monitoring intervals 318
 - deleting and undeleting 318
 - deleting process 318
 - grouping 319
 - reasons, deleting, and undeleting 318
 - undeleting process 319
 - SAN hosts
 - administration transport 183
 - editing settings for 183
 - host agents 183
 - port for communication 183
 - SAN management
 - FC switches 307
 - Global and local options 316, 317
 - Global-only options 316, 317
 - GlobalSAN role 319
 - license key 307
 - list of tasks 315
 - prerequisites 307
 - SAN hosts, initiators 307
 - tasks performed 308
 - tasks performed, UI locations 308
 - SAN reports 309
 - saved reports
 - list of failed report outputs 137
 - list of successful report outputs 137
 - viewing the output of a particular report 138
 - Saved reports
 - Viewing the output details of a particular report output 138
 - Viewing the output of a particular report output 138
 - scenario example, identify oldest files 166
 - schedule report
 - defined 135
 - methods 133
 - schedules
 - adding 136
 - deleting 136
 - editing 136
 - list 135
 - scheduling report 132, 133
 - scripts, pre and postbackup 238, 271
 - secure connections
 - components of 145
 - enabling HTTPS for 145
 - global options versus appliance options 148
 - protocols used for 146
 - requirements for, on workstation 146
 - Secure Shell 143
 - Secure Sockets Layer 146
 - self-signed certificates 143
 - server
 - console access 265
 - set up, steps for FSRM 155
 - SnapMirror
 - Business Continuance Management license key 253
 - lag thresholds 263, 264
 - license requirement for managing 253
 - relationships, managing 253
 - synchronous, modes 258
 - SnapMirror Lag Error threshold 264
 - SnapMirror Lag Warning threshold 264
 - SnapMirror relationships
 - resynchronize 263
 - Snapshot backup 272
 - Snapshot copies
 - local data protection 245
 - schedules, interaction 245
 - Snapshot copy
 - monitoring requirements 204
 - Snapshot-based backups
 - limitations 273
 - SnapVault 240
 - SnapVault backup relationships
 - CLI commands 248
 - configure 248
 - SnapVault relationship

- NDMP 239
- SnapVault relationships
 - configuration of DataFabric Manager server 241
- SNMP
 - traps 86–88
 - SNMP version 39
- spreadsheets
 - LUN reports, putting into 131
- SRM
 - comparison to capacity reports 154
 - license keys 154
 - password, setting 160
 - passwords, setting 160
- SRM hosts
 - administration access
 - configuration settings 159
 - described 156
 - enabling administration access 160
 - passwords, types of 156, 157
 - quick reference of management tasks 157
- SRM Reports 132
- SRM, File
 - See FSRM
- SSH 146
- SSL 146
- stop monitoring a LUN 185
- storage array ports
 - discover 41
- storage arrays
 - discover 41
- storage arrays of the gateway
 - configuration details 42
- Storage Controller Tools 217
- storage controllers 78
- storage objects
 - deleted objects 210
- storage systems
 - remote configuration 229
- storage-related report 187
- Subgroup reports 78
- Summary reports 77
- Symbols
 - /etc/quotas file 173, 174
 - quotas file 171
- synchronous SnapMirror
 - failover mode 258
 - multiplexing mode 258

T

- takeover, tool 220

- targets, See FCP Targets 182
- templates
 - customizing reports 123
- Temporarily stop monitoring 318
- The Gateway SAN-attached storage management
 - limitations 42
- thresholds
 - aggregate full 190
 - aggregate full interval 190
 - aggregate nearly full 190
 - aggregate nearly overcommitted 190
 - aggregate overcommitted 190
 - changing group 77
 - description of 175
 - editing 187
 - global and group levels 187
 - overriding global settings 190
 - precedence of 176
 - SnapMirror 263
 - user quota 175
 - ways and locations to apply 175
- traditional volumes
 - See volumes
- Trusted CA-signed certificates 143

U

- undeleting (restoring)
 - objects 211
 - SAN components 185
- user accounts 46
- user alert
 - e-mail message 118
 - mailformat file 122, 123
 - mailmap file 120
- user alerts
 - alarms and alerts, differences 119
 - contents of e-mail message 122
 - Default Email Domain for Quota Alerts option 120
 - defined 118
- user capabilities 52
- user quota
 - about editing 173
 - mailformat file 122
 - mailmap file 120
- User Quota Monitoring Interval option 173
- user quota thresholds
 - hard quota limit 175
 - monitor 175
 - settings 175

- soft quota limit 175
- user quotas
 - alarms and alerts, differences 119
 - alert See user alert 118
 - editing 226
- Usergroups
 - Adding 65
 - Deleting 66
 - Editing settings 66
 - Pushing 66
 - Viewing settings 65
- using FSRM
 - administrative access 167

V

- verify administrative access 167
- vFiler administrator
 - configuring 45
- vFiler unit
 - management tasks 227
- vFiler units
 - access control
 - administrators 45
 - adding to Resource groups 76
 - discovery 35
- vFiler Units
 - editing quotas 172

- vFilers
 - discovery 35
- view storage system details
 - Refresh Monitoring Samples tool 218
- viewing data 165
- virtual servers 78
- Virtual servers 125
- volume capacity thresholds
 - modify 199
- volume full threshold 194
- volume nearly full threshold 194
- Volume Nearly No First Snapshot Threshold 199
- Volume Snap Reserve Full Threshold 199
- Volume Snapshot Count Threshold 199
- Volume Snapshot No First Count Threshold 199
- volume thresholds
 - Snapshot copy, events 199
- Volume Too Old Snapshot Threshold 199
- volumes
 - chargeback (billing). See storage chargeback 205

W

- when to back up data 272
- workstation
 - command line access (CLI) 265
 - Web-based access (GUI) 265



NA 210-05177_A0, Printed in USA

GA32-1015-00

