

IBM System Storage N series



OnCommand Unified Manager Operations Manager Administration Guide for Use with Core Package 5.2

Contents

Preface	17
Supported features	17
Websites	17
Getting information, help, and service	17
Before you call	18
Using the documentation	18
Hardware service and support	18
Firmware updates	18
How to send your comments	19
Introduction to Operations Manager	20
What DataFabric Manager server does	20
What a license key is	20
Access to Operations Manager	20
Information to customize in the DataFabric Manager server	21
Administrator accounts on the DataFabric Manager server	21
Authentication methods on the management server	22
Authentication with native operating system	22
Authentication with LDAP	22
Understanding information relevant to 7-Mode and clustered environments	23
7-Mode: List of tasks and information	24
Clustered environment: List of tasks and information	38
What the discovery process is	49
Discovery by the DataFabric Manager server	49
What SNMP is	49
When to enable SNMP	49
SNMP versions to discover and monitor storage systems	50
What the Preferred SNMP Version option is	51
How the DataFabric Manager server chooses network credentials for discovery	51
Discovery process using SNMPv1 or SNMPv3	52
Monitoring process using SNMPv1	52

Monitoring process using SNMPv3	53
Setting SNMPv1 or SNMPv3 as the preferred version for storage system discovery	53
Setting SNMPv1 as the only SNMP version	53
Setting SNMPv1 or SNMPv3 to monitor a storage system	54
Modifying the network credentials and SNMP settings	54
Deleting SNMP settings for the network using the Operations Manager console	55
Addition of a storage system from an undiscovered network	55
Diagnosis of SNMP connectivity	55
What host discovery is	56
Ping methods in host discovery	56
What host-initiated discovery is	56
Discovery of vFiler units (7-Mode environments only)	57
Discovery of storage systems	57
Discovery of storage systems and networks	58
Methods of adding storage systems and networks	59
Guidelines for editing discovery options	59
Discovery of a cluster by Operations Manager (clustered environments only)	61
Adding clusters (clustered environments only)	62
Cluster monitoring tasks using Operations Manager (clustered environments only)	62
Unsupported tasks for cluster monitoring in Operations Manager (clustered environments only)	62
Tasks performed for the gateway SAN-attached storage management in Operations Manager (7-Mode environments only)	63
Limitations of the gateway SAN-attached storage management in Operations Manager (7-Mode environments only)	63
Tasks performed from the Storage Controller Details page for the gateway (7-Mode environments only)	64
Viewing configuration details of storage arrays connected to the gateway (7-Mode environments only)	64
Role-based access control in the DataFabric Manager server	65
What role-based access control does	65
Configuring vFiler unit access control (7-Mode environments only)	65
Logging in to the DataFabric Manager server	66

What default administrator accounts are	66
List of predefined roles in the DataFabric Manager server	67
Active Directory user group accounts	68
Adding administrative users	68
How roles relate to administrators	69
What predefined global roles are	69
What inheritance roles are	71
What capabilities are	72
Role precedence and inheritance	72
Creating roles	72
Modifying roles	72
What an RBAC resource is	73
Granting restricted access to RBAC resources	73
Access restriction on Windows Administrators group users	74
Access check for application administrators	75
How reports are viewed for administrators and roles	75
What global and group access control is	75
Management of administrator access	76
Prerequisites for managing administrator access (7-Mode environments only)	76
Limitations in managing administrator access	77
User access control for cluster management (clustered environments only)	77
Summary of the global group (7-Mode environments only)	77
Who local users are (7-Mode environments only)	77
What domain users are (7-Mode environments only)	83
What user groups are (7-Mode environments only)	85
What roles are (7-Mode environments only)	87
What jobs display (7-Mode environments only)	90
What groups and objects are	91
What group types are	91
What homogeneous groups are	91
What mixed-type groups are	92
What a Global group is	92
What hierarchical groups are	93
Creating groups	93

Creating groups from a report	94
What configuration resource groups are	94
Guidelines for creating groups	95
Guidelines for creating configuration resource groups	95
Guidelines for adding vFiler units to the appliance resource group (7-Mode environments only)	96
Editing group membership	96
What group threshold settings are	97
What group reports are	97
What summary reports are	97
What subgroup reports are	98
What cluster-related objects are (clustered environments only)	98
Creating a group of cluster objects (clustered environments only)	100
Storage monitoring and reporting	101
The DataFabric Manager server monitoring process	101
Cluster monitoring with Operations Manager (clustered environments only)	103
What the cluster-management LIF is (clustered environments only)	103
Information available on the Cluster Details page (clustered environments only)	103
Viewing the utilization of resources	104
Links to FilerView (7-Mode environments only)	105
Query intervals	105
Guidelines for changing monitoring intervals	106
What SNMP trap listener is	106
What SNMP trap events are	106
How SNMP trap reports are viewed	107
When SNMP traps cannot be received	107
SNMP trap listener configuration requirements	107
How SNMP trap listener is stopped	108
Configuration of SNMP trap global options	108
Information about the DataFabric Manager server MIB	108
What events are	109
Viewing events	109
Managing events	109
Operations on local configuration change events	109
List of events and severity types	110

Alarm configurations	134
Guidelines for configuring alarms	134
Creating alarms	134
Testing alarms	135
Comments in alarm notifications	136
Example of alarm notification in e-mail format	136
Example of alarm notification in script format	137
Example of alarm notification in trap format	137
Response to alarms	137
Deleting alarms	138
Working with user alerts	138
What user alerts are	138
Differences between alarms and user alerts	138
User alerts configurations	139
How you specify email addresses for alerts	139
Domains in user quota alerts	140
What the mailmap file is	140
Guidelines for editing the mailmap file	141
How the contents of the user alert are viewed	141
How the contents of the e-mail alert are changed	141
What the mailformat file is	141
Guidelines for editing the mailformat file	142
Introduction to DataFabric Manager server reports	143
Introduction to report options	143
Introduction to report catalogs	144
Custom reports you can create in Operations Manager	144
What performance reports do (7-Mode environments only)	148
Configuring custom reports	149
Deleting custom reports	150
Putting data into spreadsheet format	150
What scheduling report generation is	150
Methods to schedule a report	152
What Schedules reports are	154
What Saved reports are	156
Data export in the DataFabric Manager server	158
How to access DataFabric Manager server data	158

Where to find the database schema for the views	159
Data types for export	159
Files and formats for storing exported data	160
Format for exported the DataFabric Manager server data	160
Format for exported Performance Advisor data (7-Mode environments only)	161
Format for last updated timestamp	161
Security configurations	163
Types of certificates in the DataFabric Manager server	163
Self-signed certificates in the DataFabric Manager server	163
Trusted Certificate Authority (CA)-signed certificates in the DataFabric Manager server	163
Creating self-signed certificates in the DataFabric Manager server	164
Obtaining a trusted CA-signed certificate	164
Enabling HTTPS	165
Secure communications with the DataFabric Manager server	165
How clients communicate with the DataFabric Manager server	165
SecureAdmin for secure connection with the DataFabric Manager server clients	166
Requirements for security options	166
Guidelines to configure security options in Operations Manager	166
Managed host options	167
Where to find managed host options	167
Guidelines for changing managed host options	167
Comparison between global and storage system-specific managed host options	168
Limitations in managed host options	168
Changing password for storage systems in the DataFabric Manager server	169
Changing passwords on multiple storage systems	169
Issue with modification of passwords for storage systems	170
Using hosts.equiv to control authentication in the DataFabric Manager server (7-Mode environments only)	170
Editing HTTP and monitor services to run as different user	172
User quotas	173
Why you use quotas	173
Overview of the quota process	173

User quota management using Operations Manager	173
Prerequisites for managing user quotas using Operations Manager	174
Where to find user quota reports in Operations Manager	174
Monitor interval for user quotas in Operations Manager	175
Modification of user quotas in Operations Manager	175
Prerequisites to edit user quotas in Operations Manager	175
Editing user quotas using Operations Manager	176
Configuring user settings using Operations Manager	176
What user quota thresholds are	176
What the DataFabric Manager server user thresholds are	177
User quota thresholds	177
Ways to configure user quota thresholds in Operations Manager	177
Precedence of user quota thresholds in the DataFabric Manager server	177
Management of LUNs, Windows and UNIX hosts, and FCP targets ..	179
Management of SAN components	179
SAN and DataFabric Manager Host Agent software (7-Mode environments only)	180
List of tasks you can perform using DataFabric Manager Host Agent software (7-Mode environments only)	180
List of tasks performed to monitor targets and initiators	180
Prerequisites to manage targets and initiators	181
Prerequisites to manage SAN hosts	181
Reports for monitoring LUNs, FCP targets, and SAN hosts	182
Information available on the LUN Details page	183
Tasks performed from the LUN Details page	184
Information about the FCP Target Details page	184
Information provided in the Host Agent Details page (7-Mode environments only)	185
Tasks you can perform from the Host Agent Details page (7-Mode environments only)	185
How storage systems, SAN hosts, and LUNs are grouped	186
Granting access to storage systems, SAN hosts, and LUNs	186
Introduction to deleting and undeleting SAN components	187
Deleting a SAN component	187
How a deleted SAN component delete is restored	187
Where to configure monitoring intervals for SAN components	187

File system management	189
Access to storage-related reports	189
Storage capacity thresholds in Operations Manager	189
Modification of storage capacity thresholds settings	189
Changing storage capacity threshold settings for global group	190
Changing storage capacity threshold settings for an individual group	190
Changing storage capacity threshold settings for a specific aggregate, volume, or qtree	190
Management of aggregate capacity	191
Volume space guarantees and aggregate overcommitment	191
Available space on an aggregate	191
Considerations before modifying aggregate capacity thresholds	192
Aggregate capacity thresholds and their events	192
Management of volume capacity	196
Volume capacity thresholds and events	196
Normal events for a volume	201
Modification of the thresholds	201
Management of qtree capacity	201
Volume Snapshot copy thresholds and events	201
Qtree capacity thresholds and events	203
How Operations Manager monitors volumes and qtrees on a vFiler unit (7- Mode environments only)	204
How Operations Manager monitors qtree quotas	205
Where to find vFiler storage resource details (7-Mode environments only)	205
What clone volumes are	205
Identification of clones and clone parents	206
Why Snapshot copies are monitored	206
Snapshot copy monitoring requirements	206
Detection of Snapshot copy schedule conflicts	206
Dependencies of a Snapshot copy	207
What storage chargeback reports are	207
When is data collected for storage chargeback reports	208
Determine the current month's and the last month's values for storage chargeback report	208
Chargeback reports in various formats	208

The chargeback report options	209
Specifying storage chargeback options at the global or group level	210
The storage chargeback increment	210
Currency display format for storage chargeback	210
Specification of the annual charge rate for storage chargeback	211
Specification of the Day of the Month for Billing for storage chargeback ..	211
Formatted charge rate for storage chargeback	212
What happens when storage objects are deleted	212
Reports of deleted storage objects	212
Undeleting a storage object for monitoring	213
Storage system management	214
Management tasks you can perform from Operations Manager	214
Operations Manager components for managing your storage system	214
Storage system groups	215
Custom comment fields in Operations Manager	215
Consolidated data and reports for storage systems and vFiler units (7-Mode environments only)	216
Tasks performed by using the storage system and vFiler unit report pages (7-Mode environments only)	216
Where to find information about a specific storage system	217
Tasks performed from a Details page of Operations Manager	217
Editable options for storage system or vFiler unit settings	217
What the storage controller tools list is (7-Mode environments only)	218
What the cluster tools list is (clustered environments only)	219
What the Diagnose Connectivity tool does	219
The Refresh Monitoring Samples tool	219
The Run a Command tool	219
The Run Telnet tool	220
Console connection through Telnet	220
Managing HA pairs using the DataFabric Manager server (7-Mode environments only)	221
Viewing the status of an HA pair by using cluster console (7-Mode environments only)	221
What the takeover tool does (7-Mode environments only)	221
What the giveback tool does (7-Mode environments only)	222
DataFabric Manager server CLI to configure storage systems	223

Methods for remote configuration of storage systems (7-Mode environments only)	223
Prerequisites for running remote CLI commands from Operations Manager	223
Running commands on a specific storage system	224
Running commands on a group of storage systems from Operations Manager	224
Remote configuration of a cluster (clustered environments only)	224
Running commands on a specific cluster (clustered environments only) ...	225
Running commands on a specific node of a cluster (clustered environments only)	225
Storage system management using FilerView (7-Mode environments only)	226
What FilerView is (7-Mode environments only)	226
Configuring storage systems by using FilerView (7-Mode environments only)	226
Introduction to MultiStore and vFiler units (7-Mode environments only)	227
Monitoring of vFiler units with the DataFabric Manager server (7-Mode environments only)	227
Prerequisites for monitoring vFiler units (7-Mode environments only)	227
vFiler unit management tasks (7-Mode environments only)	228
Configuration of storage systems (7-Mode environments only)	230
Management of storage system configuration files (7-Mode environments only) .	230
Prerequisites to applying configuration files to storage systems and vFiler units (7-Mode environments only)	230
List of access roles to manage storage system configuration files (7-Mode environments only)	231
List of tasks for configuration management (7-Mode environments only)	231
What configuration files are (7-Mode environments only)	232
What a configuration plug-in is (7-Mode environments only)	232
Comparison of configurations (7-Mode environments only)	232
Verification of a successful configuration push (7-Mode environments only)	232
Configuring the DataFabric Manager server to monitor configuration changes (7-Mode environments only)	232
What a configuration resource group is (7-Mode environments only)	233

List of tasks for managing configuration groups (7-Mode environments only)	233
Guidelines for creating configuration groups (7-Mode environments only)	234
Creating configuration resource groups (7-Mode environments only)	234
Parent configuration resource groups (7-Mode environments only)	235
Configuring multiple storage systems or vFiler units (7-Mode environments only)	236
Maintenance and management	237
Accessing the DataFabric Manager server through the CLI	237
Where to find information about the DataFabric Manager server commands	237
Audit logging	238
Events audited in the DataFabric Manager server	238
Global options for audit log files and their values	239
Format of events in an audit log file	239
Permissions for accessing the audit log file	241
What the remote platform management interface is	241
RLM card monitoring in the DataFabric Manager server	241
Prerequisites for using the remote platform management interface	242
Scripts overview	242
Commands that can be used as part of the script	242
Package of the script content	243
What script plug-ins are	243
What the script plug-in directory is	243
What backup scripts do (7-Mode environments only)	244
Overview of the DataFabric Manager server database backup process	244
When to back up data	244
Where to back up data	245
Recommendations for disaster recovery	245
Backup storage and sizing	245
Limitations of Snapshot-based backups	245
Access requirements for backup operations	246
Changing the directory path for database backups	246
Starting database backup from Operations Manager	246
Scheduling database backups from Operations Manager	247
Specifying the backup retention count	247

Disabling database backup schedules	248
Listing database backups	248
Deleting database backups from Operations Manager	248
Displaying diagnostic information from Operations Manager	248
Exportability of a backup to a new location	249
What the restore process is	249
Restoring the database from the archive-based backup	249
Restoring the database from the Snapshot copy-based backup	250
Restoration of the database on different systems	250
Disaster recovery configurations (7-Mode environments only)	251
Disaster recovery using the N series Management Console data protection capability (7-Mode environments only)	251
Disaster recovery using SnapDrive (7-Mode environments only)	256
Troubleshooting in Operations Manager	257
AutoSupport in the DataFabric Manager server	257
Reasons for using AutoSupport	257
Types of AutoSupport messages in the DataFabric Manager server	257
Protection of private data by using AutoSupport	258
Configuring AutoSupport	258
DataFabric Manager server logs	259
Access to the DataFabric Manager server logs	259
Accessing the logs through the DataFabric Manager server CLI	259
Apache and Sybase log rotation in the DataFabric Manager server	260
Common DataFabric Manager server problems	260
Communication issues between the DataFabric Manager server and routers	260
E-mail alerts not working in the DataFabric Manager server	260
How discovery issues are resolved	260
Use of the Diagnose Connectivity tool for a managed storage system	261
Use of the Diagnose Connectivity tool for an unmanaged storage system	261
Where to find the Diagnose Connectivity tool in Operations Manager	261
Reasons why the DataFabric Manager server might not discover your network	262
Troubleshooting network discovery issues	262
Troubleshooting appliance discovery issues with Operations Manager	263
How configuration push errors are resolved	263

Issues related to SAN events	264
Offline LUNs	264
Snapshot copy of LUN not possible	264
High traffic in HBA Port	264
Custom reports do not display any data	265
Import and export of configuration files	266
How inconsistent configuration states are fixed	266
Data ONTAP issues impacting protection on vFiler units (7-Mode environments only)	266
Report fields and performance counters (7-Mode environments only)	268
Report fields and performance counters for filer catalogs (7-Mode environments only)	268
Report fields and performance counters for vFiler catalogs (7-Mode environments only)	270
Report fields and performance counters for volume catalogs (7-Mode environments only)	271
Report fields and performance counters for qtree catalogs (7-Mode environments only)	273
Report fields and performance counters for LUN catalogs (7-Mode environments only)	274
Report fields and performance counters for aggregate catalogs (7-Mode environments only)	274
Report fields and performance counters for disk catalogs (7-Mode environments only)	275
Protocols and port numbers used by the DataFabric Manager server	276
DataFabric Manager server communication	276
DataFabric Manager server access to storage systems	276
DataFabric Manager server access to host agents (7-Mode environments only)	277
DataFabric Manager server access to Open Systems SnapVault agents (7-Mode environments only)	277
SAN management	278
Discovery of SAN hosts by the DataFabric Manager server (7-Mode environments only)	278
SAN management using the DataFabric Manager server	279

Prerequisites for SAN management with the DataFabric Manager server .	279
List of tasks performed for SAN management	280
List of user interface locations to perform SAN management tasks	281
Reports for monitoring SANs	282
Location of SAN reports	282
The DataFabric Manager server managed SAN data in spreadsheet format	283
Where to find information for specific SAN components	283
Where to view LUN details of SAN components	283
Tasks performed from the LUN Details page for a SAN host	283
Information about FC target on a SAN host	284
Information about the FC switch on a SAN host	285
Access to the FC Switch Details page	285
Information about FC Switch on a SAN host	285
Tasks performed from the FC Switch Details page for a SAN host	285
Information about DataFabric Manager Host Agent software on a SAN host (7-Mode environments only)	286
Accessing the HBA Port Details page for a SAN host	286
Details about the HBA Port Details page	287
List of SAN management tasks	287
LUN management	287
Initiator group management	289
FC switch management	289
DataFabric Manager server options	289
DataFabric Manager server options for SAN management	289
Where to configure monitoring intervals for SAN components	291
Deleting and undeleting SAN components	291
Reasons for deleting and undeleting SAN components	291
Process of deleting SAN components	291
Process of undeleting SAN components	292
How SAN components are grouped	292
Restriction of SAN management access	292
Access control on groups of SAN components	292
Copyright information	293
Trademark information	294
Index	297

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 17).

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 17) 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 17).

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, ensure 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 17).

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 email 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

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 Operations Manager and the CLI from the IP address or Domain Name System (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; vFiler units, disks, aggregates, volumes, and qtrees on these storage systems; LUNs; and user quotas.

If the server is on Windows, Operations Manager is launched automatically and a welcome page appears.

You can use one of the following URLs to access Operations Manager:

- `http://server_ip_address:8080/dfm`
- `http://server_dnsname:8080`

Depending on your DNS setup, you might have to use the fully qualified name in this URL; for example, you should 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 69

Authentication methods on the management server

The management 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 management server to use Lightweight Directory Access Protocol (LDAP). If you configure LDAP, then the server uses it as the preferred method of authentication.

Despite the authentication method used, the server maintains its own database of user names and passwords for local users. (A local user might or might not be an administrator). For local users, the server does not use the native operating system for authentication; it performs authentication itself.

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 management server and configure the management server to communicate with your LDAP servers to retrieve information about LDAP users.

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

The following LDAP servers are compatible with the management server:

- Microsoft Active Directory
- OpenLDAP
- IBM Lotus LDAP
- Netscape LDAP Server

Understanding information relevant to 7-Mode and clustered environments

Because OnCommand Unified Manager supports both clustered Data ONTAP and 7-Mode environments, it is important to identify the information that is relevant to you.

The following labeling guidelines can help you to better understand the content that is provided in the OnCommand Unified Manager documentation:

- (7-Mode environments only)
Topic titles or information within a topic include this label when the topic or information is applicable to 7-Mode environments only. For example, the title *Adding storage systems (7-Mode environments only)* indicates that this task is relevant for 7-Mode environments only.
- (clustered environments only)
Topic titles or information within a topic include this label when the topic or information is applicable to the clustered environment only. For example, the title *Adding clusters (clustered environments only)* indicates that this task is relevant for the clustered environment only.
- No mode labels in titles
Topics that apply to both 7-Mode and clustered environments do not include any labels in the title. For example, the topic *Creating alarms for events* provides information that is relevant to both 7-Mode and clustered environments.
- No mode labels in topic information
Information within a topic that applies to both 7-Mode and clustered environments does not include any labels.

The documentation includes the following topics to help you navigate to the specific information that you want:

- Each section in the Help includes two topics that provide links to information in that section of the Help, which is relevant to each mode.
- Each product manual includes two topics that provide links to information throughout the manual, which is relevant to each mode.

7-Mode: List of tasks and information

OnCommand Unified Manager supports both 7-Mode and clustered environments; therefore, it is important to identify the information that relates specifically to your 7-Mode environment.

The following list provides links to all tasks and information related to Operations Manager administration in a 7-Mode environment:

Operations Manager

- [Introduction to Operations Manager](#) on page 20
- [What DataFabric Manager server does](#) on page 20
- [What a license key is](#) on page 20
- [Access to Operations Manager](#) on page 20
- [Information to customize in the DataFabric Manager server](#) on page 21
- [Administrator accounts on the DataFabric Manager server](#) on page 21
- [Authentication methods on the DataFabric Manager server](#) on page 22
- [Authentication with native operating system](#) on page 22
- [Authentication with LDAP](#) on page 22

Discovery

- [What the discovery process is](#) on page 49
- [Discovery by the DataFabric Manager server](#) on page 49
- [What SNMP is](#) on page 49
- [When to enable SNMP](#) on page 49
- [SNMP versions to discover and monitor storage systems](#) on page 50
- [What the Preferred SNMP Version option is](#) on page 51
- [SNMP version setup](#) on page 51
- [How DataFabric Manager chooses network credentials for discovery](#) on page 51
- [Discovery process using SNMPv1 or SNMPv3](#) on page 52
- [Monitoring process using SNMPv1](#) on page 52
- [Monitoring process using SNMPv3](#) on page 53
- [Setting SNMPv1 or SNMPv3 as the preferred version](#) on page 53
- [Setting SNMPv1 as the only SNMP version](#) on page 53
- [Setting SNMPv1 or SNMPv3 to monitor a storage system](#) on page 54
- [Modifying the network credentials and SNMP settings](#) on page 54
- [Deleting the SNMP settings for the network](#) on page 55
- [Addition of a storage system from an undiscovered network](#) on page 55
- [Diagnosis of SNMP connectivity](#) on page 55
- [What host discovery is](#) on page 56

- [*Ping methods in host discovery*](#) on page 56
- [*What host-initiated discovery is*](#) on page 56
- [*Discovery of vFiler units*](#) on page 57
- [*Discovery of storage systems*](#) on page 57
- [*Discovery of storage systems and networks*](#) on page 58
- [*Methods of adding storage systems and networks*](#) on page 59
- [*Guidelines for editing discovery options*](#) on page 59
- [*Tasks performed for the gateway SAN-attached storage management*](#) on page 63
- [*Limitations of the gateway SAN-attached storage management in Operations Manager*](#) on page 63
- [*Tasks performed from the Storage Controller Details page for the gateway*](#) on page 64
- [*Viewing configuration details of storage arrays connected to the gateway*](#) on page 64

Role-based access control (RBAC)

- [*Role-based access control in the DataFabric Manager server*](#) on page 65
- [*What role-based access control does*](#) on page 65
- [*Configuring vFiler unit access control*](#) on page 65
- [*Logging in to the DataFabric Manager server*](#) on page 66
- [*What default administrator accounts are*](#) on page 66
- [*List of predefined roles in the DataFabric Manager server*](#) on page 67
- [*Active Directory user group accounts*](#) on page 68
- [*Adding administrative users*](#) on page 68
- [*How roles relate to administrators*](#) on page 69
- [*What predefined global roles are*](#) on page 69
- [*What inheritance roles are*](#) on page 71
- [*What capabilities are*](#) on page 72
- [*Role precedence and inheritance*](#) on page 72
- [*Creating roles*](#) on page 72
- [*Modifying roles*](#) on page 72
- [*What an RBAC resource is*](#) on page 73
- [*Granting restricted access to RBAC resources*](#) on page 73
- [*Access restriction on Windows Administrators group users*](#) on page 74
- [*Access check for application administrators*](#) on page 75
- [*How reports are viewed for administrators and roles*](#) on page 75
- [*What a global and group access control is*](#) on page 75
- [*Management of administrator access*](#) on page 76
- [*Prerequisites for managing administrator access*](#) on page 76
- [*Limitations in managing administrator access*](#) on page 77
- [*Summary of the global group*](#) on page 77
- [*Viewing a specific summary page*](#) on page 77

- [Viewing users on the host](#) on page 77
- [Who local users are](#) on page 77
- [Viewing local users on the host](#) on page 78
- [Viewing local user settings on the host](#) on page 78
- [Adding local users to the host](#) on page 79
- [Editing local user settings on the host](#) on page 80
- [Users with Execute capability](#) on page 80
- [Pushing passwords to a local user](#) on page 80
- [Deleting local users from the host](#) on page 81
- [Pushing local users to hosts](#) on page 82
- [Monitoring changes in local user configuration](#) on page 82
- [Editing passwords](#) on page 82
- [What domain users are](#) on page 83
- [Viewing domain users on the host](#) on page 83
- [Adding domain users to the host](#) on page 83
- [Viewing domain user settings on the host](#) on page 83
- [Editing domain user settings on the host](#) on page 84
- [Removing domain users from all the user groups](#) on page 84
- [Pushing domain users to hosts](#) on page 84
- [Monitoring changes in domain user configuration](#) on page 85
- [What User groups are](#) on page 85
- [Viewing user groups on the host](#) on page 85
- [Adding user groups to a host](#) on page 85
- [Viewing the user group settings on a host](#) on page 86
- [Editing user group settings on a host](#) on page 86
- [Deleting user groups from a host](#) on page 87
- [Pushing user groups to hosts](#) on page 87
- [Monitoring changes in user group configuration](#) on page 87
- [What roles are](#) on page 87
- [Viewing roles on the host](#) on page 87
- [Adding roles to a host or user](#) on page 88
- [Viewing role settings on the host](#) on page 88
- [Editing role settings on the host](#) on page 89
- [Deleting roles from the host](#) on page 89
- [Pushing roles to the hosts](#) on page 89
- [Monitoring changes in role configuration](#) on page 90
- [What jobs display](#) on page 90
- [Pushing jobs](#) on page 90
- [Deleting push jobs](#) on page 90

Groups and objects

- [What groups and objects are](#) on page 91
- [What group types are](#) on page 91
- [What homogeneous groups are](#) on page 91
- [What mixed-type groups are](#) on page 92
- [What a Global group is](#) on page 92
- [What hierarchical groups are](#) on page 93
- [Creating groups](#) on page 93
- [Creating groups from a report](#) on page 94
- [What configuration resource groups are](#) on page 94
- [Guidelines for creating groups](#) on page 95
- [Guidelines for creating configuration resource groups](#) on page 95
- [Guidelines for adding vFiler units to the appliance resource group](#) on page 96
- [Editing group membership](#) on page 96
- [What group threshold settings are](#) on page 97
- [What group reports are](#) on page 97
- [What summary reports are](#) on page 97
- [What subgroup reports are](#) on page 98

Storage monitoring and reporting

- [Storage monitoring and reporting](#) on page 101
- [The DataFabric Manager server monitoring process](#) on page 101
- [Viewing utilization of resources](#) on page 104
- [Viewing the utilization of physical resources](#) on page 105
- [Links to FilerView](#) on page 105
- [Query intervals](#) on page 105
- [Guidelines for changing monitoring intervals](#) on page 106
- [What SNMP trap listener is](#) on page 106
- [What SNMP trap events are](#) on page 106
- [How SNMP trap reports are viewed](#) on page 107
- [When SNMP traps cannot be received](#) on page 107
- [SNMP trap listener configuration requirements](#) on page 107
- [How SNMP trap listener is stopped](#) on page 108
- [Configuration of SNMP trap global options](#) on page 108
- [Information about the DataFabric Manager server MIB](#) on page 108
- [What events are](#) on page 109
- [Viewing events](#) on page 109
- [Managing events](#) on page 109
- [Operations on local configuration change events](#) on page 109
- [List of events and severity types](#) on page 110

- [*Alarm configurations*](#) on page 134
- [*Guidelines for configuring alarms*](#) on page 134
- [*Creating alarms*](#) on page 134
- [*Testing alarms*](#) on page 135
- [*Comments in alarm notifications*](#) on page 136
- [*Example of alarm notification in e-mail format*](#) on page 136
- [*Example of alarm notification in script format*](#) on page 137
- [*Example of alarm notification in trap format*](#) on page 137
- [*Response to alarms*](#) on page 137
- [*Deleting alarms*](#) on page 138
- [*Working with user alerts*](#) on page 138
- [*What user alerts are*](#) on page 138
- [*Differences between alarms and user alerts*](#) on page 138
- [*User alerts configurations*](#) on page 139
- [*E-mail addresses for alerts*](#) on page 139
- [*Domains in user quota alerts*](#) on page 140
- [*What the mailmap file is*](#) on page 140
- [*Guidelines for editing the mailmap file*](#) on page 141
- [*How the contents of the user alert are viewed*](#) on page 141
- [*How the contents of the e-mail alert are changed*](#) on page 141
- [*What the mailformat file is*](#) on page 141
- [*Guidelines for editing the mailformat file*](#) on page 142
- [*Introduction to DataFabric Manager server reports*](#) on page 143
- [*Introduction to report options*](#) on page 143
- [*Introduction to report catalogs*](#) on page 144
- [*Reports you can create in Operations Manager*](#) on page 144
- [*What performance reports are*](#) on page 148
- [*Configuring custom reports*](#) on page 149
- [*Deleting custom reports*](#) on page 150
- [*Putting data into spreadsheet format*](#) on page 150
- [*What scheduling report generation is*](#) on page 150
- [*What Report Archival Directory is*](#) on page 151
- [*Additional capabilities for categories of reports*](#) on page 151
- [*What Report Schedules reports are*](#) on page 151
- [*Scheduling a report using the All submenu*](#) on page 151
- [*Scheduling a report using the Schedule submenu*](#) on page 152
- [*Methods to schedule a report*](#) on page 152
- [*Editing a report schedule*](#) on page 152
- [*Deleting a report schedule*](#) on page 153
- [*Enabling a report schedule*](#) on page 153

- [*Disabling a report schedule*](#) on page 153
- [*Running a report schedule*](#) on page 153
- [*Retrieving the list of enabled report schedules*](#) on page 154
- [*Retrieving the list of disabled report schedules*](#) on page 154
- [*Listing all the run results of a report schedule*](#) on page 154
- [*What Schedules reports are*](#) on page 154
- [*Listing all the schedules*](#) on page 154
- [*Adding a new schedule*](#) on page 155
- [*Editing a schedule*](#) on page 155
- [*Deleting a schedule*](#) on page 155
- [*What Saved reports are*](#) on page 156
- [*Listing the report outputs*](#) on page 156
- [*Listing the successful report outputs*](#) on page 156
- [*Listing the failed report outputs*](#) on page 156
- [*Viewing the output of report outputs from the status column*](#) on page 157
- [*Viewing the output of report outputs from the Output ID column*](#) on page 157
- [*Viewing the output details of a particular report output*](#) on page 157
- [*Data export in the DataFabric Manager server*](#) on page 158
- [*How to access DataFabric Manager server data*](#) on page 158
- [*Where to find the database schema for the views*](#) on page 159
- [*Data types for export*](#) on page 159
- [*Files and formats for storing exported data*](#) on page 160
- [*Format for exported DataFabric Manager data*](#) on page 160
- [*Format for exported Performance Advisor data*](#) on page 161
- [*Format for last updated timestamp*](#) on page 161

Security configurations

- [*Security configurations*](#) on page 163
- [*Types of certificates in the DataFabric Manager server*](#) on page 163
- [*Self-signed certificates in the DataFabric Manager server*](#) on page 163
- [*Trusted Certificate Authority \(CA\)-signed certificates in the DataFabric Manager server*](#) on page 163
- [*Creating self-signed certificates in DataFabric Manager*](#) on page 164
- [*Obtaining a trusted CA-signed certificate*](#) on page 164
- [*Enabling HTTPS*](#) on page 165
- [*Secure communications with the DataFabric Manager server*](#) on page 165
- [*How clients communicate with the DataFabric Manager server*](#) on page 165
- [*SecureAdmin for secure connection with the DataFabric Manager server clients*](#) on page 166
- [*Requirements for security options*](#) on page 166
- [*Guidelines to configure security options in Operations Manager*](#) on page 166

- [Managed host options](#) on page 167
- [Where to find managed host options](#) on page 167
- [Guidelines for changing managed host options](#) on page 167
- [Comparison between global and storage system-specific managed host options](#) on page 168
- [Limitations in managed host options](#) on page 168
- [Changing password for storage systems in the DataFabric Manager server](#) on page 169
- [Changing passwords on multiple storage systems](#) on page 169
- [Issue with modification of passwords for storage systems](#) on page 170
- [Using hosts.equiv to control authentication in the DataFabric Manager server](#) on page 170
- [Editing HTTP and monitor services to run as different user](#) on page 172

User quotas

- [User quotas](#) on page 173
- [Why you use quotas](#) on page 173
- [Overview of the quota process](#) on page 173
- [User quota management using Operations Manager](#) on page 173
- [Prerequisites for managing user quotas using Operations Manager](#) on page 174
- [Where to find user quota reports in Operations Manager](#) on page 174
- [Monitor interval for user quotas in Operations Manager](#) on page 175
- [Modification of user quotas in Operations Manager](#) on page 175
- [Prerequisites to edit user quotas in Operations Manager](#) on page 175
- [Editing user quotas using Operations Manager](#) on page 176
- [Configuring user settings using Operations Manager](#) on page 176
- [What user quota thresholds are](#) on page 176
- [What the DataFabric Manager server user thresholds are](#) on page 177
- [User quota thresholds](#) on page 177
- [Ways to configure user quota thresholds in Operations Manager](#) on page 177
- [Precedence of user quota thresholds in the DataFabric Manager server](#) on page 177

Management of LUNs, FCP Targets, and SAN hosts

- [Management of LUNs, Windows and UNIX hosts, and FCP targets](#) on page 179
- [Management of SAN components](#) on page 179
- [SAN and Host Agent software](#) on page 180
- [List of tasks you can perform using Host Agent software](#) on page 180
- [List of tasks performed to monitor targets and initiators](#) on page 180
- [Prerequisites to manage targets and initiators](#) on page 181
- [Prerequisites to manage SAN hosts](#) on page 181
- [Reports for monitoring LUNs, FCP targets, and SAN hosts](#) on page 182
- [Information available on the LUN Details page](#) on page 183
- [Tasks performed on the LUN Details page](#) on page 184

- [Information about the FCP Target Details page](#) on page 184
- [Information about the Host Agent Details page](#) on page 185
- [Tasks you can perform from the Host Agent Details page](#) on page 185
- [How storage systems, SAN hosts, and LUNs are grouped](#) on page 186
- [Granting access to storage systems, SAN hosts, and LUNs](#) on page 186
- [Introduction to deleting and undeleting SAN components](#) on page 187
- [Deleting a SAN component](#) on page 187
- [How a deleted SAN component delete is restored](#) on page 187
- [Where to configure monitoring intervals for SAN components](#) on page 187

File system management

- [File system management](#) on page 189
- [Access to storage-related reports](#) on page 189
- [Storage capacity thresholds in Operations Manager](#) on page 189
- [Modification of storage capacity thresholds settings](#) on page 189
- [Changing storage capacity threshold settings for global group](#) on page 190
- [Changing storage capacity threshold settings for an individual group](#) on page 190
- [Changing storage capacity threshold settings for a specific aggregate, volume, or qtree](#) on page 190
- [Management of aggregate capacity](#) on page 191
- [Volume space guarantees and aggregate overcommitment](#) on page 191
- [Available space on an aggregate](#) on page 191
- [Considerations before modifying aggregate capacity thresholds](#) on page 192
- [Aggregate capacity thresholds and their events](#) on page 192
- [Management of volume capacity](#) on page 196
- [Volume capacity thresholds and events](#) on page 196
- [Normal events for a volume](#) on page 201
- [Modification of the thresholds](#) on page 201
- [Management of qtree capacity](#) on page 201
- [Volume Snapshot copy thresholds and events](#) on page 201
- [Qtree capacity thresholds and events](#) on page 203
- [How Operations Manager monitors volumes and qtrees on a vFiler unit](#) on page 204
- [How Operations Manager monitors qtree quotas](#) on page 205
- [Where to find vFiler storage resource details](#) on page 205
- [What clone volumes are](#) on page 205
- [Identification of clones and clone parents](#) on page 206
- [Why Snapshot copies are monitored](#) on page 206
- [Snapshot copy monitoring requirements](#) on page 206
- [Detection of Snapshot copy schedule conflicts](#) on page 206
- [Dependencies of a Snapshot copy](#) on page 207

- [Thresholds on Snapshot copies](#) on page 207
- [Storage chargeback reports](#) on page 207
- [When is data collected for storage chargeback reports](#) on page 208
- [Determine the current month's and the last month's values for storage chargeback report](#) on page 208
- [Chargeback reports in various formats](#) on page 208
- [The chargeback report options](#) on page 209
- [Specifying storage chargeback options at the global or group level](#) on page 210
- [The storage chargeback increment](#) on page 210
- [Currency display format for storage chargeback](#) on page 210
- [Specification of the annual charge rate for storage chargeback](#) on page 211
- [Specification of the Day of the Month for Billing for storage chargeback](#) on page 211
- [Formatted charge rate for storage chargeback](#) on page 212
- [What happens when storage objects are deleted](#) on page 212
- [Reports of deleted storage objects](#) on page 212
- [Undeleting a storage object for monitoring](#) on page 213

Storage system management

- [Storage system management](#) on page 214
- [Management tasks you can perform from Operations Manager](#) on page 214
- [Operations Manager components for managing your storage system](#) on page 214
- [Storage system groups](#) on page 215
- [Custom comment fields in Operations Manager](#) on page 215
- [Consolidated data and reports for storage systems and vFiler units](#) on page 216
- [Tasks performed by using the storage systems and vFiler unit report pages](#) on page 216
- [Where to find information about a specific storage system](#) on page 217
- [Tasks performed from a Details page of Operations Manager](#) on page 217
- [Editable options for storage system or vFiler unit settings](#) on page 217
- [What the storage controller tools list is](#) on page 218
- [What the Diagnose Connectivity tool does](#) on page 219
- [The Refresh Monitoring Samples tool](#) on page 219
- [The Run a Command tool](#) on page 219
- [The Run Telnet tool](#) on page 220
- [Console connection through Telnet](#) on page 220
- [Managing HA pairs using the DataFabric Manager server](#) on page 221
- [Viewing the status of an HA pair by using cluster console](#) on page 221
- [What the takeover tool does](#) on page 221
- [What the giveback tool does](#) on page 222
- [DataFabric Manager server CLI to configure storage systems](#) on page 223
- [Methods for remote configuration of storage systems](#) on page 223

- [Prerequisites for running remote CLI commands from Operations Manager](#) on page 223
- [Running commands on a specific storage system](#) on page 224
- [Running commands on a group of storage systems from Operations Manager](#) on page 224
- [Storage system management using FilerView](#) on page 226
- [What FilerView is](#) on page 226
- [Configuring storage systems by using FilerView](#) on page 226
- [Introduction to MultiStore and vFiler units](#) on page 227
- [Monitoring of vFiler units with the DataFabric Manager server](#) on page 227
- [Prerequisites for monitoring vFiler units](#) on page 227
- [vFiler unit management tasks](#) on page 228

Configuration management of storage systems

- [Configuration of storage systems](#) on page 230
- [Management of storage system configuration files](#) on page 230
- [Prerequisites to applying configuration files to storage systems and vFiler units](#) on page 230
- [List of access roles to manage storage system configuration files](#) on page 231
- [List of tasks for configuration management](#) on page 231
- [What configuration files are](#) on page 232
- [What a configuration plug-in is](#) on page 232
- [Comparison of configurations](#) on page 232
- [Verification of a successful configuration push](#) on page 232
- [Configuring the DataFabric Manager server to monitor configuration changes](#) on page 232
- [What a configuration resource group is](#) on page 233
- [List of tasks for managing configuration groups](#) on page 233
- [Guidelines for creating configuration groups](#) on page 234
- [Creating configuration resource groups](#) on page 234
- [Parent configuration resource groups](#) on page 235
- [Parent group guidelines](#) on page 235
- [When to assign parent groups](#) on page 235
- [Properties of configuration files acquired from a parent](#) on page 235
- [Configuring multiple storage systems or vFiler units](#) on page 236

Maintenance and management

- [Maintenance and management](#) on page 237
- [Accessing the CLI](#) on page 237
- [Where to find information about the DataFabric Manager server commands](#) on page 237
- [Audit logging](#) on page 238
- [Events audited in DataFabric Manager](#) on page 238
- [Global options for audit log files and their values](#) on page 239
- [Format of events in an audit log file](#) on page 239

- [*Permissions for accessing the audit log file*](#) on page 241
- [*What remote platform management interface is*](#) on page 241
- [*RLM card monitoring in the DataFabric Manager server*](#) on page 241
- [*Prerequisites for using the remote platform management interface*](#) on page 242
- [*Scripts overview*](#) on page 242
- [*Commands that can be used as part of the script*](#) on page 242
- [*Package of the script content*](#) on page 243
- [*What script plug-ins are*](#) on page 243
- [*What the script plug-in directory is*](#) on page 243
- [*What backup scripts do*](#) on page 244
- [*Overview of the DataFabric Manager database backup process*](#) on page 244
- [*When to back up data*](#) on page 244
- [*Where to back up data*](#) on page 245
- [*Recommendations for disaster recovery*](#) on page 245
- [*Backup storage and sizing*](#) on page 245
- [*Limitation of Snapshot-based backups*](#) on page 245
- [*Access requirements for backup operations*](#) on page 246
- [*Changing the directory path for database backups*](#) on page 246
- [*Starting database backup from Operations Manager*](#) on page 246
- [*Scheduling database backups from Operations Manager*](#) on page 247
- [*Specifying the backup retention count*](#) on page 247
- [*Disabling database backup schedules*](#) on page 248
- [*Listing database backups*](#) on page 248
- [*Deleting database backups from Operations Manager*](#) on page 248
- [*Displaying diagnostic information from Operations Manager*](#) on page 248
- [*Exportability of a backup to a new location*](#) on page 249
- [*What the restore process is*](#) on page 249
- [*Restoring the database from the archive-based backup*](#) on page 249
- [*Restoring the database from the Snapshot copy-based backup*](#) on page 250
- [*Restoration of the database on different systems*](#) on page 250
- [*Disaster recovery configurations*](#) on page 251
- [*Disaster recovery using the Management Console data protection capability*](#) on page 251
- [*Limitation of disaster recovery support*](#) on page 251
- [*Prerequisites for disaster recovery support*](#) on page 251
- [*Setting up the DataFabric Manager server*](#) on page 252
- [*Recovering the DataFabric Manager server services*](#) on page 253
- [*Recovering the DataFabric Manager server services using the dfm datastore mirror connect command*](#) on page 254
- [*Failing back DataFabric Manager server services*](#) on page 255
- [*Disaster recovery using SnapDrive*](#) on page 256

Troubleshooting

- [Troubleshooting in Operations Manager](#) on page 257
- [AutoSupport in the DataFabric Manager server](#) on page 257
- [Reasons for using AutoSupport](#) on page 257
- [Types of AutoSupport messages in the DataFabric Manager server](#) on page 257
- [Protection of private data by using AutoSupport](#) on page 258
- [Configuring AutoSupport](#) on page 258
- [DataFabric Manager server logs](#) on page 259
- [Access to logs](#) on page 259
- [Accessing the logs through the DataFabric Manager server CLI](#) on page 259
- [Apache and Sybase log rotation in DataFabric Manager server](#) on page 260
- [Common DataFabric Manager server problems](#) on page 260
- [Communication issues between the DataFabric Manager server and routers](#) on page 260
- [E-mail alerts not working in the DataFabric Manager server](#) on page 260
- [How discovery issues are resolved](#) on page 260
- [Use of the Diagnose Connectivity tool for a managed storage system](#) on page 261
- [Use of the Diagnose Connectivity tool for an unmanaged storage system](#) on page 261
- [Where to find the Diagnose Connectivity tool in Operations Manager](#) on page 261
- [Reasons why the DataFabric Manager server might not discover your network](#) on page 262
- [Troubleshooting network discovery issues](#) on page 262
- [Troubleshooting appliance discovery issues with Operations Manager](#) on page 263
- [How configuration push errors are resolved](#) on page 263
- [Issues related to SAN events](#) on page 264
- [Offline LUNs](#) on page 264
- [Snapshot copy of LUN not possible](#) on page 264
- [High traffic in HBA Port](#) on page 264
- [Custom reports do not display any data](#) on page 265
- [Import and export of configuration files](#) on page 266
- [How inconsistent configuration states are fixed](#) on page 266
- [Data ONTAP issues impacting protection on vFiler units](#) on page 266

Report fields and performance counters

- [Report fields and performance counters](#) on page 268
- [Report fields and performance counters for filer catalogs](#) on page 268
- [Report fields and performance counters for vFiler catalogs](#) on page 270
- [Report fields and performance counters for volume catalogs](#) on page 271
- [Report fields and performance counters for qtree catalogs](#) on page 273
- [Report fields and performance counters for LUN catalogs](#) on page 274
- [Report fields and performance counters for aggregate catalogs](#) on page 274
- [Report fields and performance counters for disk catalogs](#) on page 275

Protocols and port numbers

- [Protocols and port numbers used by the DataFabric Manager server](#) on page 276
- [DataFabric Manager server communication](#) on page 276
- [DataFabric Manager server access to storage systems](#) on page 276
- [DataFabric Manager server access to host agents](#) on page 277
- [DataFabric Manager server access to Open Systems SnapVault agents](#) on page 277

SAN management

- [SAN management](#) on page 278
- [Discovery of SAN hosts by the DataFabric Manager server](#) on page 278
- [SAN management using the DataFabric Manager server](#) on page 279
- [Prerequisites for SAN management with the DataFabric Manager server](#) on page 279
- [List of tasks performed for SAN management](#) on page 280
- [List of user interface locations to perform SAN management tasks](#) on page 281
- [Reports for monitoring SANs](#) on page 282
- [Location of SAN reports](#) on page 282
- [DataFabric Manager server managed SAN data in spreadsheet format](#) on page 283
- [Where to find information for specific SAN components](#) on page 283
- [Where to view LUN details of SAN components](#) on page 283
- [Tasks performed from the LUN Details page for a SAN host](#) on page 283
- [Information about FC target on a SAN host](#) on page 284
- [Information about the FC switch on a SAN host](#) on page 285
- [Access to the FC Switch Details page](#) on page 285
- [Information about FC Switch on a SAN host](#) on page 285
- [Tasks performed from the FC Switch Details page for a SAN host](#) on page 285
- [Information about Host Agent software on a SAN host](#) on page 286
- [Accessing the HBA Port Details page access for a SAN host](#) on page 286
- [Details about the HBA Port Details page](#) on page 287
- [List of SAN management tasks](#) on page 287
- [LUN management](#) on page 287
- [Initiator group management](#) on page 289
- [FC switch management](#) on page 289
- [DataFabric Manager server options](#) on page 289
- [DataFabric Manager options for SAN management](#) on page 289
- [Where to configure monitoring intervals for SAN components](#) on page 291
- [Deleting and undeleting SAN components](#) on page 291
- [Reasons for deleting and undeleting SAN components](#) on page 291
- [Process of deleting SAN components](#) on page 291
- [Process of undeleting SAN components](#) on page 292
- [How SAN components are grouped](#) on page 292

- [*Restriction of SAN management access*](#) on page 292
- [*Access control on groups of SAN components*](#) on page 292

Clustered environment: List of tasks and information

Because Unified Manager supports both 7-Mode and clustered environments, it is important to identify the information that relates specifically to your clustered environment.

The following list provides links to all tasks and information related to Operations Manager administration in clustered environments:

Operations Manager

- [Introduction to Operations Manager](#) on page 20
- [What DataFabric Manager server does](#) on page 20
- [What a license key is](#) on page 20
- [Access to Operations Manager](#) on page 20
- [Information to customize in the DataFabric Manager server](#) on page 21
- [Administrator accounts on the DataFabric Manager server](#) on page 21
- [Authentication methods on the DataFabric Manager server](#) on page 22
- [Authentication with native operating system](#) on page 22
- [Authentication with LDAP](#) on page 22

Discovery

- [What the discovery process is](#) on page 49
- [Discovery by the DataFabric Manager server](#) on page 49
- [What SNMP is](#) on page 49
- [When to enable SNMP](#) on page 49
- [SNMP versions to discover and monitor storage systems](#) on page 50
- [What the Preferred SNMP Version option is](#) on page 51
- [SNMP version setup](#) on page 51
- [How DataFabric Manager chooses network credentials for discovery](#) on page 51
- [Discovery process using SNMPv1 or SNMPv3](#) on page 52
- [Monitoring process using SNMPv1](#) on page 52
- [Monitoring process using SNMPv3](#) on page 53
- [Setting SNMPv1 or SNMPv3 as the preferred version](#) on page 53
- [Setting SNMPv1 as the only SNMP version](#) on page 53
- [Setting SNMPv1 or SNMPv3 to monitor a storage system](#) on page 54
- [Modifying the network credentials and SNMP settings](#) on page 54
- [Deleting the SNMP settings for the network](#) on page 55
- [Addition of a storage system from an undiscovered network](#) on page 55

- [*Diagnosis of SNMP connectivity*](#) on page 55
- [*What host discovery is*](#) on page 56
- [*Ping methods in host discovery*](#) on page 56
- [*What host-initiated discovery is*](#) on page 56
- [*Discovery of storage systems*](#) on page 57
- [*Discovery of storage systems and networks*](#) on page 57
- [*Methods of adding storage systems and networks*](#) on page 59
- [*Guidelines for editing discovery options*](#) on page 59
- [*Discovery of a cluster by Operations Manager*](#) on page 61
- [*Adding clusters*](#) on page 62
- [*Cluster monitoring tasks using Operations Manager*](#) on page 62
- [*Unsupported tasks for cluster monitoring in Operations Manager*](#) on page 62

Role-based access control (RBAC)

- [*Role-based access control in the DataFabric Manager server*](#) on page 65
- [*What role-based access control does*](#) on page 65
- [*Logging in to the DataFabric Manager server*](#) on page 66
- [*What default administrator accounts are*](#) on page 66
- [*List of predefined roles in the DataFabric Manager server*](#) on page 67
- [*Active Directory user group accounts*](#) on page 68
- [*Adding administrative users*](#) on page 68
- [*How roles relate to administrators*](#) on page 69
- [*What predefined global roles are*](#) on page 69
- [*What inheritance roles are*](#) on page 71
- [*What capabilities are*](#) on page 72
- [*Role precedence and inheritance*](#) on page 72
- [*Creating roles*](#) on page 72
- [*Modifying roles*](#) on page 72
- [*What an RBAC resource is*](#) on page 73
- [*Granting restricted access to RBAC resources*](#) on page 73
- [*Access restriction on Windows Administrators group users*](#) on page 74
- [*Access check for application administrators*](#) on page 75
- [*How reports are viewed for administrators and roles*](#) on page 75
- [*What a global and group access control is*](#) on page 75
- [*Management of administrator access*](#) on page 76
- [*Prerequisites for managing administrator access*](#) on page 76
- [*Limitations in managing administrator access*](#) on page 77
- [*Controlled access to cluster management*](#) on page 77

Groups and objects

- [*What groups and objects are*](#) on page 91
- [*What group types are*](#) on page 91
- [*What homogeneous groups are*](#) on page 91
- [*What mixed-type groups are*](#) on page 92
- [*What a Global group is*](#) on page 92
- [*What hierarchical groups are*](#) on page 93
- [*Creating groups*](#) on page 93
- [*Creating groups from a report*](#) on page 94
- [*What configuration resource groups are*](#) on page 94
- [*Guidelines for creating groups*](#) on page 95
- [*Guidelines for creating configuration resource groups*](#) on page 95
- [*Editing group membership*](#) on page 96
- [*What group threshold settings are*](#) on page 97
- [*What group reports are*](#) on page 97
- [*What summary reports are*](#) on page 97
- [*What subgroup reports are*](#) on page 98
- [*What cluster-related objects are*](#) on page 98
- [*Creating a group of cluster objects*](#) on page 100

Storage monitoring and reporting

- [*Storage monitoring and reporting*](#) on page 101
- [*The DataFabric Manager server monitoring process*](#) on page 101
- [*Cluster monitoring with Operations Manager*](#) on page 103
- [*What the cluster-management LIF is*](#) on page 103
- [*Information available on the Cluster Details page*](#) on page 103
- [*Tasks performed from the Cluster Details page*](#) on page 104
- [*Viewing utilization of resources*](#) on page 104
- [*Viewing the utilization of logical resources*](#) on page 104
- [*Viewing the utilization of physical resources*](#) on page 105
- [*Query intervals*](#) on page 105
- [*Guidelines for changing monitoring intervals*](#) on page 106
- [*What SNMP trap listener is*](#) on page 106
- [*What SNMP trap events are*](#) on page 106
- [*How SNMP trap reports are viewed*](#) on page 107
- [*When SNMP traps cannot be received*](#) on page 107
- [*SNMP trap listener configuration requirements*](#) on page 107
- [*How SNMP trap listener is stopped*](#) on page 108
- [*Configuration of SNMP trap global options*](#) on page 108
- [*Information about the DataFabric Manager server MIB*](#) on page 108

- [*What events are*](#) on page 109
- [*Viewing events*](#) on page 109
- [*Managing events*](#) on page 109
- [*Operations on local configuration change events*](#) on page 109
- [*List of events and severity types*](#) on page 110
- [*Alarm configurations*](#) on page 134
- [*Guidelines for configuring alarms*](#) on page 134
- [*Creating alarms*](#) on page 134
- [*Testing alarms*](#) on page 135
- [*Comments in alarm notifications*](#) on page 136
- [*Example of alarm notification in e-mail format*](#) on page 136
- [*Example of alarm notification in script format*](#) on page 137
- [*Example of alarm notification in trap format*](#) on page 137
- [*Response to alarms*](#) on page 137
- [*Deleting alarms*](#) on page 138
- [*Working with user alerts*](#) on page 138
- [*What user alerts are*](#) on page 138
- [*Differences between alarms and user alerts*](#) on page 138
- [*User alerts configurations*](#) on page 139
- [*E-mail addresses for alerts*](#) on page 139
- [*Domains in user quota alerts*](#) on page 140
- [*What the mailmap file is*](#) on page 140
- [*Guidelines for editing the mailmap file*](#) on page 141
- [*How the contents of the user alert are viewed*](#) on page 141
- [*How the contents of the e-mail alert are changed*](#) on page 141
- [*What the mailformat file is*](#) on page 141
- [*Guidelines for editing the mailformat file*](#) on page 142
- [*Introduction to DataFabric Manager server reports*](#) on page 143
- [*Introduction to report options*](#) on page 143
- [*Introduction to report catalogs*](#) on page 144
- [*Reports you can create in Operations Manager*](#) on page 144
- [*Configuring custom reports*](#) on page 149
- [*Deleting custom reports*](#) on page 150
- [*Putting data into spreadsheet format*](#) on page 150
- [*What scheduling report generation is*](#) on page 150
- [*What Report Archival Directory is*](#) on page 151
- [*Additional capabilities for categories of reports*](#) on page 151
- [*What Report Schedules reports are*](#) on page 151
- [*Scheduling a report using the All submenu*](#) on page 151
- [*Scheduling a report using the Schedule submenu*](#) on page 152

- [*Methods to schedule a report*](#) on page 152
- [*Editing a report schedule*](#) on page 152
- [*Deleting a report schedule*](#) on page 153
- [*Enabling a report schedule*](#) on page 153
- [*Disabling a report schedule*](#) on page 153
- [*Running a report schedule*](#) on page 153
- [*Retrieving the list of enabled report schedules*](#) on page 154
- [*Retrieving the list of disabled report schedules*](#) on page 154
- [*Listing all the run results of a report schedule*](#) on page 154
- [*What Schedules reports are*](#) on page 154
- [*Listing all the schedules*](#) on page 154
- [*Adding a new schedule*](#) on page 155
- [*Editing a schedule*](#) on page 155
- [*Deleting a schedule*](#) on page 155
- [*What Saved reports are*](#) on page 156
- [*Listing the report outputs*](#) on page 156
- [*Listing the successful report outputs*](#) on page 156
- [*Listing the failed report outputs*](#) on page 156
- [*Viewing the output of report outputs from the status column*](#) on page 157
- [*Viewing the output of report outputs from the Output ID column*](#) on page 157
- [*Viewing the output details of a particular report output*](#) on page 157
- [*Data export in the DataFabric Manager server*](#) on page 158
- [*How to access DataFabric Manager server data*](#) on page 158
- [*Where to find the database schema for the views*](#) on page 159
- [*Two types of data for export*](#) on page 159
- [*Files and formats for storing exported data*](#) on page 160
- [*Format for exported DataFabric Manager data*](#) on page 160
- [*Format for last updated timestamp*](#) on page 161

Security configurations

- [*Security configurations*](#) on page 163
- [*Types of certificates in the DataFabric Manager server*](#) on page 163
- [*Self-signed certificates in the DataFabric Manager server*](#) on page 163
- [*Trusted Certificate Authority \(CA\)-signed certificates in the DataFabric Manager server*](#) on page 163
- [*Creating self-signed certificates in DataFabric Manager*](#) on page 164
- [*Obtaining a trusted CA-signed certificate*](#) on page 164
- [*Enabling HTTPS*](#) on page 165
- [*Secure communications with the DataFabric Manager server*](#) on page 165
- [*How clients communicate with the DataFabric Manager server*](#) on page 165

- [SecureAdmin for secure connection with the DataFabric Manager server clients](#) on page 166
- [Requirements for security options](#) on page 166
- [Guidelines to configure security options in Operations Manager](#) on page 166
- [Managed host options](#) on page 167
- [Where to find managed host options](#) on page 167
- [Guidelines for changing managed host options](#) on page 167
- [Comparison between global and storage system-specific managed host options](#) on page 168
- [Limitations in managed host options](#) on page 168
- [Changing password for storage systems in the DataFabric Manager server](#) on page 169
- [Changing passwords on multiple storage systems](#) on page 169
- [Issue with modification of passwords for storage systems](#) on page 170
- [Editing HTTP and monitor services to run as different user](#) on page 172

User quotas

- [User quotas](#) on page 173
- [Why you use quotas](#) on page 173
- [Overview of the quota process](#) on page 173
- [User quota management using Operations Manager](#) on page 173
- [Prerequisites for managing user quotas using Operations Manager](#) on page 174
- [Where to find user quota reports in Operations Manager](#) on page 174
- [Monitor interval for user quotas in Operations Manager](#) on page 175
- [Modification of user quotas in Operations Manager](#) on page 175
- [Prerequisites to edit user quotas in Operations Manager](#) on page 175
- [Editing user quotas using Operations Manager](#) on page 176
- [Configuring user settings using Operations Manager](#) on page 176
- [What user quota thresholds are](#) on page 176
- [What the DataFabric Manager server user thresholds are](#) on page 177
- [User quota thresholds](#) on page 177
- [Ways to configure user quota thresholds in Operations Manager](#) on page 177
- [Precedence of user quota thresholds in the DataFabric Manager server](#) on page 177

Management of LUNs, FCP Targets, and SAN hosts

- [Management of LUNs, Windows and UNIX hosts, and FCP targets](#) on page 179
- [Management of SAN components](#) on page 179
- [List of tasks performed to monitor targets and initiators](#) on page 180
- [Prerequisites to manage targets and initiators](#) on page 181
- [Prerequisites to manage SAN hosts](#) on page 181
- [Reports for monitoring LUNs, FCP targets, and SAN hosts](#) on page 182
- [Information available on the LUN Details page](#) on page 183
- [Tasks performed on the LUN Details page](#) on page 184

- [Information about the FCP Target Details page](#) on page 184
- [How storage systems, SAN hosts, and LUNs are grouped](#) on page 186
- [Granting access to storage systems, SAN hosts, and LUNs](#) on page 186
- [Introduction to deleting and undeleting SAN components](#) on page 187
- [Deleting a SAN component](#) on page 187
- [How a deleted SAN component delete is restored](#) on page 187
- [Where to configure monitoring intervals for SAN components](#) on page 187

File system management

- [File system management](#) on page 189
- [Access to storage-related reports](#) on page 189
- [Storage capacity thresholds in Operations Manager](#) on page 189
- [Modification of storage capacity thresholds settings](#) on page 189
- [Changing storage capacity threshold settings for global group](#) on page 190
- [Changing storage capacity threshold settings for an individual group](#) on page 190
- [Changing storage capacity threshold settings for a specific aggregate, volume, or qtree](#) on page 190
- [Management of aggregate capacity](#) on page 191
- [Volume space guarantees and aggregate overcommitment](#) on page 191
- [Available space on an aggregate](#) on page 191
- [Considerations before modifying aggregate capacity thresholds](#) on page 192
- [Aggregate capacity thresholds and their events](#) on page 192
- [Management of volume capacity](#) on page 196
- [Volume capacity thresholds and events](#) on page 196
- [Normal events for a volume](#) on page 201
- [Modification of the thresholds](#) on page 201
- [Management of qtree capacity](#) on page 201
- [Volume Snapshot copy thresholds and events](#) on page 201
- [Qtree capacity thresholds and events](#) on page 203
- [How Operations Manager monitors qtree quotas](#) on page 205
- [What clone volumes are](#) on page 205
- [Identification of clones and clone parents](#) on page 206
- [Why Snapshot copies are monitored](#) on page 206
- [Snapshot copy monitoring requirements](#) on page 206
- [Detection of Snapshot copy schedule conflicts](#) on page 206
- [Dependencies of a Snapshot copy](#) on page 207
- [Thresholds on Snapshot copies](#) on page 207
- [Storage chargeback reports](#) on page 207
- [When is data collected for storage chargeback reports](#) on page 208

- [*Determine the current month's and the last month's values for storage chargeback report*](#) on page 208
- [*Chargeback reports in various formats*](#) on page 208
- [*The chargeback report options*](#) on page 209
- [*Specifying storage chargeback options at the global or group level*](#) on page 210
- [*The storage chargeback increment*](#) on page 210
- [*Currency display format for storage chargeback*](#) on page 210
- [*Specification of the annual charge rate for storage chargeback*](#) on page 211
- [*Specification of the Day of the Month for Billing for storage chargeback*](#) on page 211
- [*Formatted charge rate for storage chargeback*](#) on page 212
- [*What happens when storage objects are deleted*](#) on page 212
- [*Reports of deleted storage objects*](#) on page 212
- [*Undeleting a storage object for monitoring*](#) on page 213

Storage system management

- [*Storage system management*](#) on page 214
- [*Management tasks you can perform from Operations Manager*](#) on page 214
- [*Operations Manager components for managing your storage system*](#) on page 214
- [*Storage system groups*](#) on page 215
- [*Custom comment fields in Operations Manager*](#) on page 215
- [*Where to find information about a specific storage system*](#) on page 217
- [*Tasks performed from a Details page of Operations Manager*](#) on page 217
- [*Editable options for storage system or vFiler unit settings*](#) on page 217
- [*What the cluster tools list is*](#) on page 219
- [*What the Diagnose Connectivity tool does*](#) on page 219
- [*The Refresh Monitoring Samples tool*](#) on page 219
- [*The Run a Command tool*](#) on page 219
- [*The Run Telnet tool*](#) on page 220
- [*Console connection through Telnet*](#) on page 220
- [*DataFabric Manager server CLI to configure storage systems*](#) on page 223
- [*Prerequisites for running remote CLI commands from Operations Manager*](#) on page 223
- [*Running commands on a specific storage system*](#) on page 224
- [*Running commands on a group of storage systems from Operations Manager*](#) on page 224
- [*Remote configuration of a cluster*](#) on page 224
- [*Running commands on a specific cluster*](#) on page 225
- [*Running commands on a specific node of a cluster*](#) on page 225

Maintenance and management

- [*Maintenance and management*](#) on page 237
- [*Accessing the CLI*](#) on page 237

- [*Where to find information about the DataFabric Manager server commands*](#) on page 237
- [*Audit logging*](#) on page 238
- [*Events audited in DataFabric Manager*](#) on page 238
- [*Global options for audit log files and their values*](#) on page 239
- [*Format of events in an audit log file*](#) on page 239
- [*Permissions for accessing the audit log file*](#) on page 241
- [*What remote platform management interface is*](#) on page 241
- [*RLM card monitoring in the DataFabric Manager server*](#) on page 241
- [*Prerequisites for using the remote platform management interface*](#) on page 242
- [*Scripts overview*](#) on page 242
- [*Commands that can be used as part of the script*](#) on page 242
- [*Package of the script content*](#) on page 243
- [*What script plug-ins are*](#) on page 243
- [*What the script plug-in directory is*](#) on page 243
- [*Overview of the DataFabric Manager database backup process*](#) on page 244
- [*When to back up data*](#) on page 244
- [*Where to back up data*](#) on page 245
- [*Recommendations for disaster recovery*](#) on page 245
- [*Backup storage and sizing*](#) on page 245
- [*Limitation of Snapshot-based backups*](#) on page 245
- [*Access requirements for backup operations*](#) on page 246
- [*Changing the directory path for database backups*](#) on page 246
- [*Starting database backup from Operations Manager*](#) on page 246
- [*Scheduling database backups from Operations Manager*](#) on page 247
- [*Specifying the backup retention count*](#) on page 247
- [*Disabling database backup schedules*](#) on page 248
- [*Listing database backups*](#) on page 248
- [*Deleting database backups from Operations Manager*](#) on page 248
- [*Displaying diagnostic information from Operations Manager*](#) on page 248
- [*Exportability of a backup to a new location*](#) on page 249
- [*What the restore process is*](#) on page 249
- [*Restoring the database from the archive-based backup*](#) on page 249
- [*Restoring the database from the Snapshot copy-based backup*](#) on page 250
- [*Restoration of the database on different systems*](#) on page 250

Troubleshooting

- [*Troubleshooting in Operations Manager*](#) on page 257
- [*AutoSupport in the DataFabric Manager server*](#) on page 257
- [*Reasons for using AutoSupport*](#) on page 257
- [*Types of AutoSupport messages in the DataFabric Manager server*](#) on page 257

- [*Protection of private data by using AutoSupport*](#) on page 258
- [*Configuring AutoSupport*](#) on page 258
- [*DataFabric Manager server logs*](#) on page 259
- [*Access to logs*](#) on page 259
- [*Accessing the logs through the DataFabric Manager server CLI*](#) on page 259
- [*Apache and Sybase log rotation in DataFabric Manager server*](#) on page 260
- [*Common DataFabric Manager server problems*](#) on page 260
- [*Communication issues between the DataFabric Manager server and routers*](#) on page 260
- [*E-mail alerts not working in the DataFabric Manager server*](#) on page 260
- [*How discovery issues are resolved*](#) on page 260
- [*Use of the Diagnose Connectivity tool for a managed storage system*](#) on page 261
- [*Use of the Diagnose Connectivity tool for an unmanaged storage system*](#) on page 261
- [*Where to find the Diagnose Connectivity tool in Operations Manager*](#) on page 261
- [*Reasons why the DataFabric Manager server might not discover your network*](#) on page 262
- [*Troubleshooting network discovery issues*](#) on page 262
- [*Troubleshooting appliance discovery issues with Operations Manager*](#) on page 263
- [*How configuration push errors are resolved*](#) on page 263
- [*Issues related to SAN events*](#) on page 264
- [*Offline FC Switch Port or Offline HBA Port*](#)
- [*Faulty FC Switch Port or HBA Port Error*](#)
- [*Offline LUNs*](#) on page 264
- [*Snapshot copy of LUN not possible*](#) on page 264
- [*High traffic in HBA Port*](#) on page 264
- [*Custom reports do not display any data*](#) on page 265
- [*Import and export of configuration files*](#) on page 266
- [*How inconsistent configuration states are fixed*](#) on page 266

Protocols and port numbers

- [*Protocols and port numbers used by the DataFabric Manager server*](#) on page 276
- [*DataFabric Manager server communication*](#) on page 276
- [*DataFabric Manager server access to storage systems*](#) on page 276

SAN management

- [*SAN management*](#) on page 278
- [*SAN management using the DataFabric Manager server*](#) on page 279
- [*Prerequisites for SAN management with the DataFabric Manager server*](#) on page 279
- [*List of tasks performed for SAN management*](#) on page 280
- [*List of user interface locations to perform SAN management tasks*](#) on page 281
- [*Reports for monitoring SANs*](#) on page 282
- [*Location of SAN reports*](#) on page 282

- [*DataFabric Manager server managed SAN data in spreadsheet format*](#) on page 283
- [*Where to find information for specific SAN components*](#) on page 283
- [*Where to view LUN details of SAN components*](#) on page 283
- [*Tasks performed from the LUN Details page for a SAN host*](#) on page 283
- [*Information about FC target on a SAN host*](#) on page 284
- [*Information about the FC switch on a SAN host*](#) on page 285
- [*Access to the FC Switch Details page*](#) on page 285
- [*Information about FC Switch on a SAN host*](#) on page 285
- [*Tasks performed from the FC Switch Details page for a SAN host*](#) on page 285
- [*Accessing the HBA Port Details page access for a SAN host*](#) on page 286
- [*Details about the HBA Port Details page*](#) on page 287
- [*List of SAN management tasks*](#) on page 287
- [*LUN management*](#) on page 287
- [*Initiator group management*](#) on page 289
- [*FC switch management*](#) on page 289
- [*DataFabric Manager server options*](#) on page 289
- [*DataFabric Manager options for SAN management*](#) on page 289
- [*Where to configure monitoring intervals for SAN components*](#) on page 291
- [*Deleting and undeleting SAN components*](#) on page 291
- [*Reasons for deleting and undeleting SAN components*](#) on page 291
- [*Process of deleting SAN components*](#) on page 291
- [*Process of undeleting SAN components*](#) on page 292
- [*How SAN components are grouped*](#) on page 292
- [*Restriction of SAN management access*](#) on page 292
- [*Access control on groups of SAN components*](#) on page 292

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 are not SNMP-enabled, you must enable SNMP for 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 has 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 that 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 require it only 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 59

[*Guidelines for editing discovery options*](#) on page 59

Related tasks

[*Modifying the network credentials and SNMP settings*](#) on page 54

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 54

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

You should be aware of 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 SNMPv3 monitoring.
The storage system credentials are not specified	The login and the password specified at the network level are used for SNMPv3 monitoring.
No credentials are provided at the network level	The login and the password specified at the global level are used for 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 authentication protocols: MD5 and SHA. You can configure SNMPv3 settings with either of the authentication protocols from the Network Credentials page or from the CLI.

Setting SNMPv1 or SNMPv3 as the preferred version for storage system discovery

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

Steps

1. Perform one of the following actions:
 - Select the **Network Credentials** submenu from the **Setup** menu.
 - 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 by using Operations Manager.

Steps

1. Choose one of the following options:
 - Click **Setup > Network Credentials**.
 - Click **Setup > Discovery**, and then 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 the values for the parameters as required.
4. Click **Update**.

Deleting SNMP settings for the network using the Operations Manager console

If you do not use a network, you can delete the SNMP settings for the network by using the Operations Manager console.

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.

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 diag hostname` 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 the 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 261

[Use of the Diagnose Connectivity tool for an unmanaged storage system](#) on page 261

[Where to find the Diagnose Connectivity tool in Operations Manager](#) on page 261

[Reasons why the DataFabric Manager server might not discover your network](#) on page 262

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 modifying the DataFabric Manager server details for host-initiated discovery, see the *DataFabric Manager Host Agent Installation and Administration Guide*.

Related information

The IBM N series support website: www.ibm.com/storage/support/nseries/

Discovery of vFiler units (7-Mode environments only)

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 169

[Changing passwords on multiple storage systems](#) on page 169

Discovery of storage systems

You must be aware of 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 the 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 located.

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>

3. 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 and 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

Before modifying the default values of the discovery options, you must decide which options you want to change and how that affects discovery.

Interval option

You can increase or decrease the default time interval between storage systems and network discovery attempts. After storage systems are discovered initially, you can determine the interval based on the number of networks and their size. Longer time intervals can lead to delays in discovery, but the discovery process is less likely to affect the network load.

The default interval value is 15 minutes.

Timeout option	<p>You can specify the time interval after which discovery is considered to have failed.</p> <p>Lengthening the interval avoids discovery queries on a local area network failing due to the long response times of a storage system.</p> <p>The default timeout value is 5 seconds.</p>
Host discovery option	<p>This option enables the discovery of storage systems, host agents and vFiler Units through SNMP.</p> <p>You can change the default value if any of the following situations exist:</p> <ul style="list-style-type: none">• 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. <p>Manually adding storage systems is faster than discovering storage systems in the following cases:</p> <ul style="list-style-type: none">• 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 option	<p>This option allows you to enable or disable discovery of host agents.</p> <p>You can change the default value if you want to disable the discovery of LUNs or storage area network (SAN) hosts and host agents.</p>
Network discovery	<p>This option enables the discovery of networks, including cluster networks.</p> <p>You can change the default value if you want the DataFabric Manager server to automatically discover storage systems on your entire network.</p> <p>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.</p>
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>

You can decrease the discovery limit if a smaller number of hops includes all the networks with storage systems that you want to discover. For example, you can 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.

The default discovery limit value is 15 hops.

Networks to discover

This option enables you to manually add or delete networks that the DataFabric Manager server scans for new storage systems.

You can change the default value if you want to add a network to the DataFabric Manager server that it cannot discover automatically or if you want to delete a network for which you no longer want storage systems to be discovered.

Network Credentials

This option enables you to specify, change, or delete an SNMP community that the DataFabric Manager server uses for a specific network or host.

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.

Discovery of a cluster by Operations Manager (clustered environments only)

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 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 clusters (clustered environments only)

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. If you want to change the default version, 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, enter the host name or IP address of the cluster that you want to add.
3. Click **Add**.

Result

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

Cluster monitoring tasks using Operations Manager (clustered environments only)

Using Operations Manager, you can perform various monitoring tasks on cluster nodes, including monitoring clusters and generating reports on the cluster and its objects.

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.
- Run commands remotely.

Unsupported tasks for cluster monitoring in Operations Manager (clustered environments only)

You cannot perform certain monitoring tasks related to a cluster due to lack of support for certain APIs, SNMP objects, and interfaces in clustered Data ONTAP systems.

The following features are not supported in Operations Manager:

- Management of the cluster administrator user profile and password
- Configuration management of clusters, controllers, and Vservers

- Configuration of the HA checker script
- Management of volume SnapMirror relationships, schedules, and jobs
- Receiving SNMP traps from a cluster
- SnapLock reports

Tasks performed for the gateway SAN-attached storage management in Operations Manager (7-Mode environments only)

Operations Manager enables you to discover 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 \(7-Mode environments only\)*](#) on page 64

Limitations of the gateway SAN-attached storage management in Operations Manager (7-Mode environments only)

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 (7-Mode environments only)

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 (7-Mode environments only)

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 have discovered 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, which provides the configuration details, such as the name of the storage array, the name of the gateway, array LUN count, the adapter used, and the name of the switch.

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 does

Role-based access control (RBAC) enables administrators to manage groups of users by defining roles. If you need to restrict access for specific functionality to selected administrators, you must set up administrator accounts for them. If you want to restrict the information that administrators can view and the operations they can perform, you must apply roles to the administrator accounts you create.

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

When you initiate an operation that requires specific privileges, the management server prompts you to log in. For example, to create administrator accounts, you must log in with Administrator account access.

Configuring vFiler unit access control (7-Mode environments only)

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

About this task

The following restrictions apply 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 that contains the volume.
- If a vFiler unit has a qtree assigned to it, the vFiler administrator cannot view details or reports for the volume that contains the qtree.

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.

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 superuser 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 <i>root</i> .
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

You must be aware of the roles assigned 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 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 available 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 that you are creating.
5. Type the email address for the administrator or administrator group.
6. Enter the pager number for the administrator or administrator group.
7. Click **Add**.

In Windows, when you add a user to the Administrators group, the user gets added as a local admin user.

After you finish

If you have not created a role for the user you created, you must create a role.

How roles relate to administrators

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

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

An operation must be specified for every role. You can assign multiple operation 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 view the description of an operation by using the `dfm role operation list` command:

```
dfm role operation list [-x] [operation-name]
```

Related concepts

[What global and group access control is](#) on page 75

[Guidelines for editing discovery options](#) on page 59

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 a set of predefined global roles that can be inherited to the user who is creating roles, as listed in the following table:

Role	Operations
Default	You cannot perform any operations.
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 the GlobalBackup, GlobalRead, and GlobalDataSet roles.
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, and create and delete alarms.
GlobalExecute	You can run commands on the storage system.
GlobalFailover	You can manage disaster recovery for datasets.
GlobalFullControl	<p>You can view and perform any operation on any object in the DataFabric Manager server database and configure administrator accounts.</p> <p>You cannot apply this role to accounts with group access control.</p>
GlobalMirror	You can create, delete, and update replication or failover policies.
GlobalPerfManagement	You can manage views (7-Mode environments only), event thresholds, and alarms, and view performance information in Performance Advisor.
GlobalProvisioning	You can provision primary dataset nodes and 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.
GlobalSDDataProtection	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	<p>You can list the Snapshot copies and the objects inside them.</p> <p>You can create, modify, and delete Snapshot copies.</p> <p>You can create clones of volumes, LUNs, and qtrees.</p> <p>You can restore volumes, LUNs, and qtrees from Snapshot copies.</p>
GlobalSDStorage	You can list, create, modify, and delete storage objects and their attributes.
GlobalWrite	You can view or write to the DataFabric Manager server database.

Note: Superusers are assigned the GlobalFullControl role in Operations Manager. For Linux, the superuser is the root user. For Windows, superusers 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. Click **Setup > Roles**.
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. Optional: To copy capabilities from an existing role, select that role from the **Inherit Capabilities** list and click >> to move the role to the list on the right.
5. Click **Add Role**.

Modifying roles

You can edit the roles created by using 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 the role inheritance by performing one of the following actions:
 - To disinherit a role, select the role from the list on the right, and click “<<” to remove it.
 - To inherit a role, select the role from the **Inherit Capabilities** list and click “>>”.
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 restriction on Windows Administrators group users

The DataFabric Manager server has the capability of restricting the access of Windows users who are part of the Administrators group.

When a user with administrator privileges installs the DataFabric Manager server, all Windows Administrators group users have GlobalFullControl access even when they are not associated with the DataFabric Manager server. The access for these Windows Administrators group users can be restricted using `dfm global option enableAdministratorRoot`.

The default value is `yes` where all the local administrators continue to have GlobalFullControl access and can manage the data of the user who installs the DataFabric Manager server. If the option is set to `no`, the Windows Administrators group users fall under RBAC mechanism. In such a situation, the Windows Administrators group users need to be provided with sufficient RBAC privileges to access the DataFabric Manager server. However, they can still run the following commands:

- `dfm plugin add`
- `dfm plugin delete`
- `dfm keystore rotatekey`
- `dfm service stop monitor`
- `dfm service start monitor`
- `dfm service list`
- `dfm datastore mirror setup`
- `dfm datastore mirror destroy`
- `dfm datastore mirror connect`
- `dfm autosupport send -e watal`
- `dfm autosupport view`
- `dfm backup create`
- `dfm backup restore`
- `dfm backup export`
- `dfm ssl server setup`
- `dfm service runas`

The user who installs the DataFabric Manager server gets GlobalFullControl access. If `dfm global option enableAdministratorRoot` is set to `no`, then the user who installs the DataFabric

Manager server can add the Windows Administrators group users as dfm users and assign RBAC privileges to them.

Access check for application administrators

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

The Core AccessCheck capability enables 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.

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 who is 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 (7-Mode environments only)

You must ensure that you meet the prerequisites for managing administrator access.

The prerequisites for managing administrator access on storage systems and vFiler units are as follows:

- You must be using Data ONTAP 7.1 or later.
- In the case of local users, you must have the minimum password age, maximum password age, and status fields, and they are available in Data ONTAP 7.1 and later.
- You must be able to reset passwords, and it 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.

User access control for cluster management (clustered environments only)

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 roles and capabilities on an object, the same capabilities are applicable to all the child objects. 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 (7-Mode environments only)

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

Viewing a specific summary page (7-Mode environments only)

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

Steps

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

Viewing users on the host (7-Mode environments only)

You can view the users on a host by 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 (7-Mode environments only)

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

Viewing local users on the host (7-Mode environments only)

You can view the local users on a host by 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 (7-Mode environments only)

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 are displayed.

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	User groups that the user belongs to.
Roles	Roles assigned to the local user.
Capabilities	Capabilities of the roles assigned to the local user who is part of the 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 the 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 has 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, after the maximum retries, 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 for 7-Mode*.

Related information

The IBM N series support website: www.ibm.com/storage/support/nseries

Adding local users to the host (7-Mode environments only)

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 following parameters:

Host Name	Name of the storage system or vFiler unit on which user is 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. Click **Local User** to add the user.

Editing local user settings on the host (7-Mode environments only)

You can edit the 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 following 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	User groups 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 (7-Mode environments only)

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

Users who do not have the Execute capability can modify the password by using the credentials that are provided.

Pushing passwords to a local user (7-Mode environments only)

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, click the password link in the Push column corresponding to the local user.

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

3. Specify the following 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	<p>Select the following from the respective list:</p> <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 the Password Management and Host User Management window.

Deleting local users from the host (7-Mode environments only)

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

Steps

- From any page, select **Management > Host Users > Local Users**.
- From the list of existing local users, select the local user that you want to delete.

3. Click **Delete Selected**.

Pushing local users to hosts (7-Mode environments only)

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 a host creates a job that is displayed in the Jobs tab of the Host User Management window.

Monitoring changes in the local user configuration (7-Mode environments only)

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

Steps

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

Editing passwords of local users (7-Mode environments only)

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 the 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 (7-Mode environments only)

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

Viewing domain users on a host (7-Mode environments only)

You can view the domain users on a host by using the Host Domain Users report.

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 a host (7-Mode environments only)

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 following parameters:

Host Name	Name of the storage system or vFiler unit.
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	User groups you want to be a member of. You can select one or more user groups from the list.

3. Click **Add Domain User**.

Viewing domain user settings on a host (7-Mode environments only)

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 are displayed:

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	User groups that the user belongs to.
Roles	Roles assigned to the domain user.
Capabilities	Capabilities of the roles assigned to the domain user as part of the user group.

Editing the domain user settings on a host (7-Mode environments only)

You can edit the settings of 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 the user group membership details.

Note: You cannot edit the host name and user name in the Edit Domain User section.

4. Click **Update**.

Removing domain users from all the user groups (7-Mode environments only)

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 (7-Mode environments only)

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 the domain user configuration (7-Mode environments only)

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

Steps

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

What user groups are (7-Mode environments only)

User groups are groups to which the users belong.

Viewing user groups on a host (7-Mode environments only)

You can view the user groups on a host by using the Host Usergroups report.

Steps

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

Adding user groups to a host (7-Mode environments only)

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 following parameters:

Host Name	Name of the storage system or vFiler unit.
Usergroup Name	Name of the user group.
Description	Description of the user group.
Select Roles	Capabilities of the roles. You can select one or more roles.

3. Click **Add Usergroup**.

Viewing the user group settings on a host (7-Mode environments only)

You can view the user group settings on 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 are displayed:

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 the user group settings on a host (7-Mode environments only)

You can edit the 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 following parameters:

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

Note: You cannot edit the host name in the Edit Usergroup section.

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

Deleting user groups from a host (7-Mode environments only)

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 user groups to hosts (7-Mode environments only)

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 to other storage systems or vFiler units.
3. Select the DataFabric Manager server group, storage system or vFiler unit to which you want to push the user group.
4. Select **OK**.
5. Click **Push**.

Monitoring changes in the user group configuration (7-Mode environments only)

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

Steps

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

What roles are (7-Mode environments only)

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 a role.

Viewing roles on hosts (7-Mode environments only)

You can view the role settings on 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 (7-Mode environments only)

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 following parameters:

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

3. Click the **Add Capabilities** link and select one or more capabilities that you want to add.
4. Click **OK**.
5. Click **Add Role**.

Viewing the role settings on a host (7-Mode environments only)

You can view the role settings on 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 are displayed:

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 the role settings on a host (7-Mode environments only)

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 following parameters:

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

Note: You cannot edit the host name and role name in the Edit Role section.

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

Deleting roles from a host (7-Mode environments only)

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 (7-Mode environments only)

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 that you want to push to other storage systems or vFiler units.
3. Select the DataFabric Manager server group or storage system to which you want to push the role.
4. Select **OK**.

5. Click **Push**.

Monitoring changes in the role configuration (7-Mode environments only)

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

Steps

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

What jobs display (7-Mode environments only)

Jobs display the status of the push jobs.

Viewing push jobs (7-Mode environments only)

You can view the status of the push jobs.

Step

1. Click **Management > Host Users > Jobs**.

Deleting push jobs (7-Mode environments only)

You can delete a push job.

Steps

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

What groups and objects are

A *group* is a collection of the DataFabric Manager server objects. Storage system elements that are 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*.

You can group objects based on characteristics such as the operating system of a storage system (Data ONTAP version), the storage systems at a location, or all the file systems that belong to a specific project or group in your organization.

You can add the following list of DataFabric Manager server objects to a resource group:

- Host (can include storage systems, Host Agents, virtual servers (clustered environments only), and vFiler units (7-Mode environments only))
- Volume
- Qtree
- Configuration
- LUN path
- Aggregate
- Disk
- Dataset (7-Mode environments only)
- Resource pool (7-Mode environments only)

You should consider the following information when 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. For example, they might be running 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
- 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 *OnCommand Unified Manager Guide to Common Provisioning and Data Protection Workflows for 7-Mode*.

Related information

The IBM N series support website: 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, their 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.
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 one of the following tabs: **Aggregate**, **File Systems**, or **LUN**.
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.
6. From the list of groups, select the parent group for the group you are creating.
You might have to expand the list to display the parent group that you want.
7. Click **Add**.

Result

The Current Groups list in the left-pane area is updated and displays the new group. You might have 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

Before creating a group, you should figure out how you want to configure the group, what the purpose of the group is, and how the groups affect your environment. For example, if you create too many groups, your environment could slow down.

You should consider the following guidelines:

- 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 copy to a group or move to a group in a group hierarchy.

However, a 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 the appliance resource group (7-Mode environments only)

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.
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. You do not see data about other object types in the parent group or the subgroups.

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.

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 duplicates, if any.

What cluster-related objects are (clustered environments only)

Operations Manager enables you to include cluster-related objects, such as nodes and Vservers, 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:

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

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 Vserver 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 Vserver or multiple Vservers, 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 Vserver 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 (clustered environments only)

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. 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

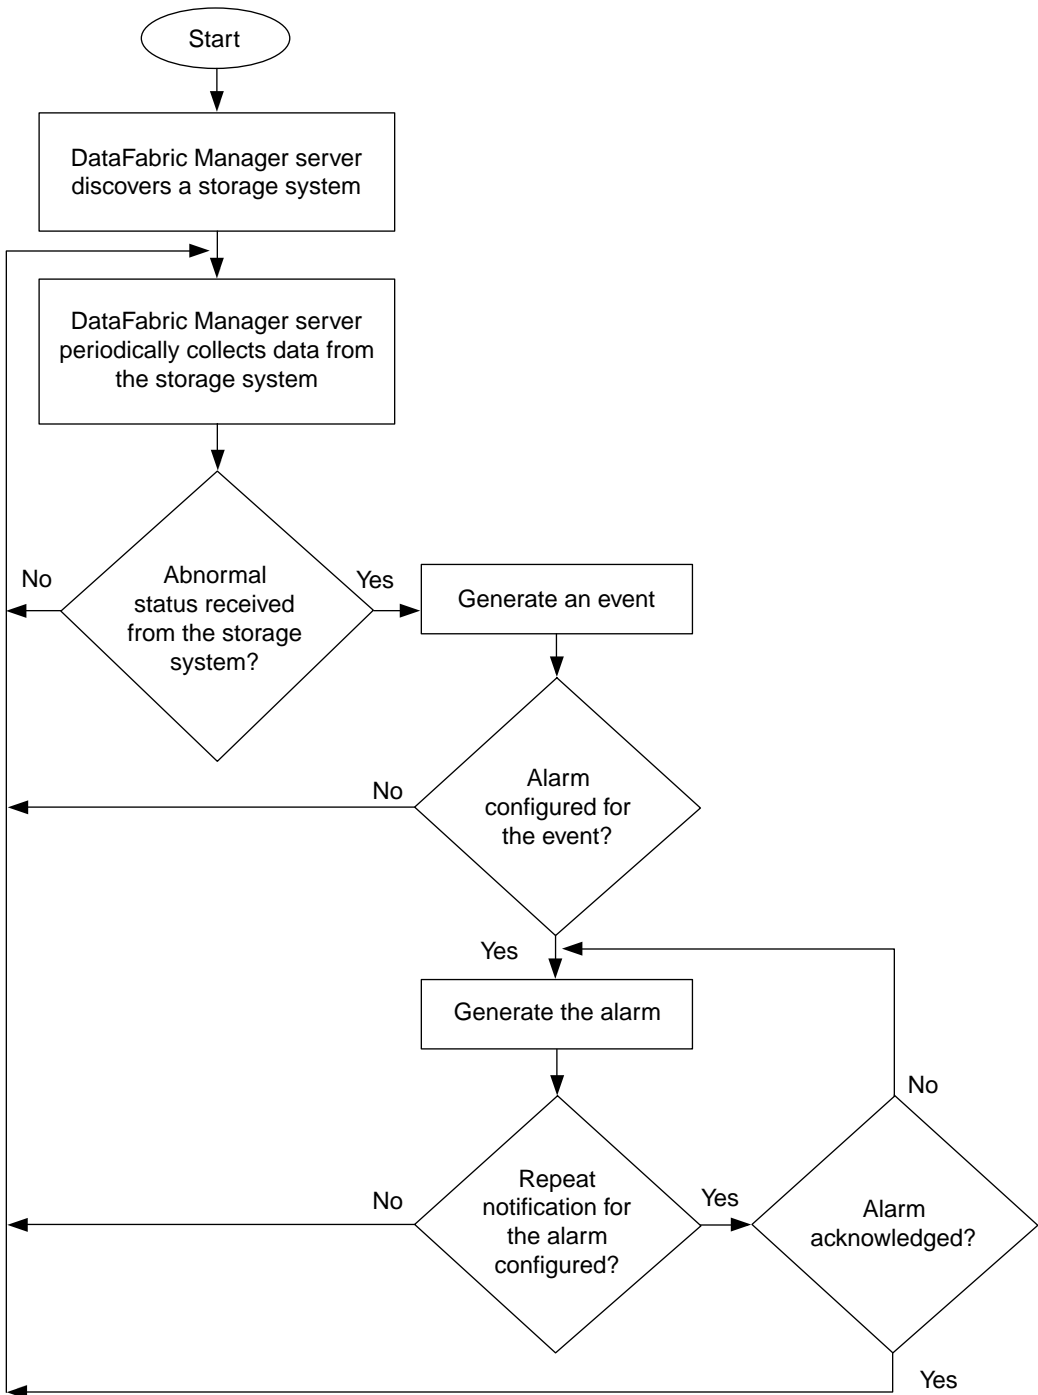
Monitoring and reporting functions in the DataFabric Manager server depend on event generation. You must configure the settings in Operations Manager to customize monitoring and to specify how and when you want to receive event notifications.

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 (clustered environments only)

You can use Operations Manager to discover and monitor storage systems in cluster environments and generate reports. Operations Manager uses SNMP or XML APIs for cluster monitoring.

What the cluster-management LIF is (clustered environments only)

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

The cluster-management LIF is associated with a cluster management server to provide a detailed view of the cluster. Operations Manager monitors clusters by using the appropriate SNMP version or APIs. You can gather information about the cluster resources, controller resources, and Vserver resources from the cluster-management LIF.

Information available on the Cluster Details page (clustered environments only)

The Cluster Details page for a cluster provides information, such as the cluster hierarchy, status of the cluster, number of logical interfaces, and ports. 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 Vservers
- 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/second
- CIFS Operations/second
- Network Traffic/second
- Logical Interface Traffic/second

Tasks performed from the Cluster Details page (clustered environments only)

You can perform various cluster management tasks from the Cluster Details page, such as viewing the resource utilization and gathering information about cluster objects. 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 objects through 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 the 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, as graphs, the information about the corresponding resources.
- Browsing the cluster and its objects
You can browse through the cluster and its objects from the "Cluster Hierarchy" section in the Cluster Details page. You can also browse to the corresponding report of a particular object. For example, you can expand the cluster, view the list of Vservers, and then click a specific Vserver 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 (clustered environments only)

By using Operations Manager, you can view, in graphs, the utilization of your logical resources, such as Vservers, logical interface traffic to the Vserver, 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 Vserver 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 Vserver's volume capacity used for 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 (7-Mode environments only)

In DataFabric Manager server 2.3 and later, UI pages that display information about some

DataFabric Manager server objects contain links to FilerView (indicated by the icon ). FilerView is a Web-based UI for managing 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 have to authenticate the storage system whose FilerView you are connecting to, by using one of the administrator user accounts on the storage system.

Note: Links to FilerView are not available for systems running Data ONTAP 8.1 or later operating in 7-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.

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 106

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 for 7-Mode*.

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 108

Related information

The IBM N series support website: 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. These events enable you to take action to prevent issues that can lead to poor performance and system unavailability. Events include an impact area, severity, and impact level.

Events are categorized by the type of impact area they encompass. Impact areas include availability, capacity, configuration, or protection. Events are also assigned a severity type and impact level that assist you in determining if immediate action is required.

You can configure alarms to send notification automatically when specific events or events of a specific severity occur.

Events are automatically logged and retained for a default of 180 days.

Event types are predetermined. You can manage event notification, but you cannot add or delete event types. However, you can modify the event severity by using the command-line interface.

Viewing events

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

Step

1. Click **Emergency, Critical, Error, Warning** on the Operations Manager main window to view reports of the respective event severity type.

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

You should be aware of all the events generated by the N series Management Console data protection and provisioning capabilities and Operations Manager and the associated event severity types.

You can use the links in the following table to move directly to the events for that object. Events are listed in alphabetical order by object type.

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

Event categories		
HA Configuration Controller on page 111 HA Configuration Interconnect on page 112 HA Configuration Partner on page 112 Agent on page 112 Aggregate on page 113 Alarm on page 113 CFO Interconnect on page 114 CFO Partner on page 114 CFO Settings on page 114 Cluster on page 114 Cluster Port on page 115 Comment Field on page 115 Configuration Changed on page 115 CPU on page 115 Data Protection on page 115 Database on page 116 Dataset on page 116 Dataset Backup on page 117 Dataset Conformance on page 117 Disks on page 118 Enclosures on page 118	Fibre Channel Switch Port on page 119 Fans on page 118 Filer Configuration on page 119 Global Status on page 119 HA State on page 119 HBA Port on page 120 Host on page 120 Host Agent on page 121 Inodes on page 121 Interface Status on page 121 Logical Interface on page 121 LUN on page 121 Management Station on page 122 Migration on page 122 NDMP on page 123 Network on page 123 Network Services on page 123 No Schedule Conflict on page 124 NVRAM Battery on page 124 OSSV on page 124 Performance Advisor on page 124 Power Supplies on page 125 Primary on page 125	Protection Policy on page 125 Protection Schedule on page 125 Provisioning Policy on page 125 Qtree on page 126 RPM on page 126 Resource Group on page 126 Resource Pool on page 126 SAN Host LUN on page 127 Script on page 127 SnapMirror on page 127 Snapshot(s) on page 128 SnapVault on page 129 SNMP Trap Listener on page 130 Space Management on page 130 Storage Services on page 130 Sync on page 131 Temperature on page 131 Unprotected Item on page 131 User on page 131 vFiler Unit on page 132 vFiler Unit Template on page 132 Vserver on page 132 Volume on page 133

HA configuration Controller

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

Event name	Severity
Takeover	Warning

HA configuration Interconnect

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

HA configuration Partner

Event name	Severity
Dead	Warning
May Be Down	Warning
OK	Normal

HA 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
Up	Normal

Aggregate

Event name	Severity
64-bit Upgrade	Information
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
Test	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

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
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

Event name	Severity
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
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

Event name	Severity
Initializing	Information
Nonconformant	Warning

Disks

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

Enclosures

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

HA State

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

Event name	Severity
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

Event name	Severity
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

Event name	Severity
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

Event name	Severity
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

Event name	Severity
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
Delete Aborted	Warning
Delete Completed	Information
Delete Failed	Error
Deleted	Information
Discovered	Information
Initialize Aborted	Warning

Event name	Severity
Initialize Completed	Normal
Initialize Failed	Error
Modified	Information
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

Event name	Severity
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
Space OK	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
Relationship Modified	Information
Replica Date OK	Normal
Replica Nearly Out of Date	Warning

Event name	Severity
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

Event name	Severity
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
E-mail Address OK	Normal
E-mail Address Rejected	Warning
Files Quota Almost Full	Warning
Files Quota Full	Error

Event name	Severity
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
Vserver Deleted	Information
Vserver Renamed	Information
Vserver Status Down	Error
Vserver Status Up	Normal

Volume

Event name	Severity
Almost Full	Warning
Automatically Deleted	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
Quota Almost Overcommitted	Warning
Quota Overcommitted	Error
Restricted	Normal
Snapshot Deleted	Normal
Space Normal	Normal

Event name	Severity
Space Reserve Depleted	Error
Space Reserve Nearly Depleted	Warning
Space Reserve 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 traphost, 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 109

Creating alarms

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

Steps

- Click **Control Center > Setup > Alarms**.
- 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: IP address of the source
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 alerts to all users who exceed their quota limits.

Alarms	User alerts
Alarms must 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 email address• A pager address• An SNMP traphost• A script that you write	User alerts are sent to the user who exceeds the user quota thresholds. User alerts are only in the form of an email message.

Alarms	User alerts
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.	<p>User alerts are sent only when the following user quota events occur:</p> <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.

How you specify email addresses for alerts

You can specify an email address on DataFabric Manager server to receive alerts. The email address that is used to send alerts depends on the DataFabric Manager server configuration.

The following list provides the different ways to specify an email address:

- Using the `dfm quota user` command
For more information about the `dfm quota user` command, see the DataFabric Manager server man pages.
- Using the `mailmap` file
If you have to specify email addresses of many users, 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 email addresses in one operation.

The following list details what occurs if you do not specify an email address:

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

Note: The DataFabric Manager server uses only that 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 email domain is not configured, 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 email addresses that contain the domain information, you must configure the domain in the DataFabric Manager server to ensure that email 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...
<i>DFM_EVENT_NAME</i>	Name of the event
<i>DFM_QUOTA_FILE_SYSTEM_NAME</i>	Name of the file system (volume or qtree) that caused the quota event
<i>DFM_QUOTA_FILE_SYSTEM_TYPE</i>	Type of file system (volume or qtree)
<i>DFM_QUOTA_PERCENT_USED</i>	Percentage of quota used
<i>DFM_QUOTA_USED</i>	Amount of disk space or number of files used
<i>DFM_QUOTA_LIMIT</i>	Total disk space or files quota
<i>DFM_QUOTA_TYPE</i>	Type of quota (disk space or files), depending on whether the disk space or files quota threshold was exceeded
<i>DFM_LINK_EVENT</i>	Hyperlink to the event
<i>DFM_QUOTA_USER_NAME</i>	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 command-line interface (CLI) or Operations Manager.

You can run reports and create custom reports from the CLI. The DataFabric Manager server also provides reports in Operations Manager interface, which enables you to perform the following operations:

- 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 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 (7-Mode environments only)
- Backup
- Disaster Recovery
- Data Protection Transfer
- Miscellaneous

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

Note: (7-Mode environments only) The report category Performance contains the performance characteristics of objects. 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.

Custom reports you can create in Operations Manager

You can use Operations Manager to create different types of reports, including reports on aggregates, array LUNs, datasets, controllers, clusters, datasets, events, ports, quotas, and backup, mirror, and other schedules.

Aggregates	The aggregate report displays information about the space utilization, capacity, performance characteristics (7-Mode environments only), and availability of 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 displays information about the LUN residing on third-party storage arrays that are attached to the gateway, including information such as the model, vendor, serial number of the LUN, and size. 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 displays information about array LUNs contained on the aggregates of the gateway, including information such as the

model, vendor, serial number of the LUN, and size. By default, you can view aggregate array LUNs reports from **Control Center > Home > Member Details > Aggregates > Report**.

Backup	(7-Mode environments only) The backup report displays information about the data transfer during a backup.
Clusters	(clustered environments only) The clusters report displays information about the clusters, such as the status, serial number of the cluster, number of Vservers and controllers associated with the cluster. You can view the reports from Control Center > Home > Member Details > Physical Systems > Report .
Controllers	The controllers report displays information about the cluster to which the controller belongs, including the model and serial number of the controller, and the system ID of the cluster. You can view the report from Control Center > Home > Member Details > Physical Systems > Report .
Dataset	(7-Mode environments only) The dataset report displays 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.
Disks	The disks report displays information about the disks in your storage system, such as the model, vendor, and size. You can view the performance characteristics (7-Mode environments only) 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 displays information about the event severity, user quotas, and SNMP traps. The information about all the 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 displays 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 displays FC switches that are deleted, that include user comments, 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 displays 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 displays information about all the 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 displays the status, and the used and available storage space 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.

Host Users

(7-Mode environments only) 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**

(7-Mode environments only) The host local users report displays 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**

(7-Mode environments only) The host domain users report displays 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**

(7-Mode environments only) The host usergroups report displays 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

(7-Mode environments only) The host roles report displays information about the existing roles on the host. By default, you can view the host roles reports from **Management > Host Roles > Report**.

**Interface
Groups**

The interface groups report displays information about all the defined cluster interface groups, ports, controllers, and active state of the group. You can view the interface groups report from **Control Center > Home > Member Details > Physical Systems > Report**.

**Logical
Interfaces**

The logical interfaces report displays information about the server, the current port that the logical interface uses, the status of the logical interface, the network address and mask, the type of role, and whether the interface is at the home port. You can view the logical interfaces report from **Control Center > Home > Member Details > Physical Systems > Report**.

LUN

The LUN report displays information and statistics about the LUNs and LUN initiator groups on the storage system, along with the performance characteristics (7-Mode environments only). By default, you can view LUN reports in the Member Details tab.

**History,
Performance
Report**

The performance events history report displays all the Performance Advisor events. By default, you can view the performance events history 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 displays information about the controllers that are connected, the type of role of the port, the 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 displays 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 displays information about the report outputs that are generated by the report schedules. By default, you can view report outputs reports from Reports > Schedule > Saved Reports .
Report Schedules	<p>The report schedules report displays information about the existing report schedules. By default, you can view report schedules reports from Reports > Schedule > Report Schedules.</p> <p>A report schedule is an association between a schedule and a report for the report generation to happen at that particular time.</p>
Resource Pools	The resource pools report displays 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	(7-Mode environments only) The SAN host report displays 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 displays 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 from Reports > Schedule > Schedules . The Schedules tab displays all the schedules. Schedules are separate entities that can be associated with reports.
Scripts	The scripts report displays 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 displays information about spare array LUNs of the gateway, such as the 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 .
storage systems	The storage systems report displays information about the capacity and operations of your storage systems, performance characteristics (7-Mode environments only), 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 displays information about the disk space usage and user quota thresholds collected from the monitored storage systems.
vFiler	(7-Mode environments only) The vFiler report displays the status, available protocols, storage usage, and performance characteristics (7-Mode environments only) 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	(clustered environments only) The Virtual Servers report displays 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	<p>The Volume report displays all the volumes with the following details, for the current month or for the past month:</p> <ul style="list-style-type: none"> • Name • Capacity • Available space • Snapshot capacity • Growth rates • Expendability • Chargeback by usage or allocation • Performance characteristics (7-Mode environments only) <p>By default, you can view the volume reports along with the file system reports in the Member Details tab.</p>

What performance reports do (7-Mode environments only)

Performance reports in Operations Manager provide information about the performance characteristics of an object.

You can perform the following operations with performance reports:

- You can view the performance characteristics of an object for the periods last, one day, one week, one month, three months, and one year.
 - You can view the performance counters related to various fields in catalogs.
 - You can use data consolidation, which is a statistical method to analyze data.
- Data consolidation is available only if you select the Performance option. You can view the average, minimum, maximum, or median value for the performance metrics over a period. This field is set to Average by default.

Related concepts

Report fields and performance counters (7-Mode environments only) on page 268

Configuring custom reports

You can configure custom reports in Operations Manager to customize the report to suit your needs. For example, you can configure a custom report with a relevant name, a comment to help you remember the use of this report, and select relevant fields for you report.

Steps

1. Select **Custom** from the **Reports** menu.
2. Enter a (short) name for the report, depending on how you want it displayed in the command-line interface (CLI).
3. Enter a (long) name for the report, depending on how you want it displayed in Operations Manager.
4. 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 have to expand the list to display the catalog you want to select.

8. In the **Choose From Available Fields** section, choose the fields you want to view:
 - To view fields related to the usage and configuration metrics of the object, click **Usage**.
 - (7-Mode environments only) To view fields related to performance metrics of the object, click **Performance**.
9. Select a field from **Choose From Available Fields**.
10. Enter a name for the field, depending on how 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. Specify the format of the field.

If you choose to not 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 that you want to include in the report.
14. Click **Move Up** or **Move Down** to reorder the fields.
15. (7-Mode environments only) If you selected **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.

18. Locate 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 apart from the Database Read capability.

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

Report category	Capabilities
Event Reports	Event Read capability
Policy Reports	Policy Read capability
Mirror Reports (7-Mode environments only)	Mirror 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 about storage system configuration and utilization. You can create customized reports to perform the following tasks:

- Forecast future capacity and bandwidth utilization requirements
- Present capacity and bandwidth utilization statistics
- Generate performance graphs
- Present 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 that is collected by the DataFabric Manager server
- Provides database access to the information that is 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, which eases the loading of data to user-specific database
- Allows you to schedule the export
- Allows you to specify the list of the counters to be exported
- Allows you to consolidate the sample values of the data export
- Allows you to customize the rate at which the performance counter data is exported (7-Mode environments only)

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 access the views that are defined within the embedded database of the DataFabric Manager server, you must 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 enables you to perform the following operations:

- Create a database user
- Delete a database user
- Enable database access to a database user
- Disable database access to a database user
Disable the database access denies the read permission on the DataFabric Manager server views for the user account.
- Change the password for the database user

All of these operations can be performed only through the command-line interface (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: To make an Open Database Connectivity (ODBC) connection to a remote database on the DataFabric Manager server, you should clear the **Shared memory** check box, and select the **TCP/IP** check box in the **Network** tab of the ODBC Data Source Administrator dialog box.

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.

Data types for export

You can export two types of data: the DataFabric Manager server data and Performance Advisor data.

- DataFabric Manager server data
The exporting of DataFabric Manager server data is controlled at the global level through the `dfmDataExportEnabled` option. By default, the value of the `dfmDataExportEnabled` global option is No.
- Performance Advisor data (7-Mode environments only)

The exporting of Performance Advisor data 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 that is performed, the counter data for the past seven days is exported. The sampling rate for the counter data export is customizable at the global level. By default, one sample is exported 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 *OnCommand Unified Manager Performance Advisor Administration Guide*.

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

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

The IBM N series support website: www.ibm.com/storage/support/nseries

Files and formats for storing exported data

The exported 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 (7-Mode environments only)

The data exported from Performance Advisor 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 counters in Performance Advisor.
- 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.

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	<p>This option enables you to set login protocols (RSH or SSH) that the DataFabric Manager server uses when connecting to the managed hosts.</p> <ul style="list-style-type: none"> Login connections. HA pair operations The <code>dfm run</code> command for running commands on the storage system <p>Change the default value if you want a secure connection for active/active configuration operations, running commands on the storage system.</p>
Administration Transport	<p>This option 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 option enables you to configure the administration port which, along with administration transport, monitors and manages storage systems.</p>

If you do not configure the port option at the storage system level, the default value for the corresponding protocol is used.

Note: If the value of the administration port is manually entered, when the value of the administration transport is changed for a host, you must ensure that you manually change the value of the administration port.

hosts.equiv option

This option enables users to authenticate storage systems when the user name and password are not provided.

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.

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 enables you to set passwords on all storage systems when you use the same authentication credentials on each system.

About this task

You should select the Global group to set the same passwords on all the systems at once.

Steps

1. Log in to Operations Manager.
2. Depending on the type of object for which you want to manage the passwords, select one of the following:

If you want to change the password for a...	Then...
Storage system	Click Management > Storage System > Passwords .
vFiler unit	Click 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 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 (7-Mode environments only)

You can control authentication of storage systems, vFiler units, and HA pair by using the `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 `hosts.equiv` file to authenticate the storage system or vFiler unit.

Steps

1. Edit the `/etc/hosts.equiv` file on the storage system to 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. 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. Depending on the operating system, perform the appropriate action:

If the operating system is...	Then...
Linux	<p>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>.</p> <p>By default, the HTTP service runs as <i>nobody</i> user on Linux.</p>
Windows	<p>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>.</p> <p>By default, the HTTP, server, scheduler, and monitor services run as <i>LocalSystem</i> user on Windows.</p>

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 have to provide the following entries:

```
DFM_HOST
```

```
DFM_HOST USER1
```

```
DFM_HOST nobody
```

On Windows operating system, you have 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 for 7-Mode*.

Related information

The IBM N series support website: 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.

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 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 is exceeded, the write operation fails, and a quota notification is sent. If any soft quota is 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 174

[Monitor interval for user quotas in Operations Manager](#) on page 175

[What user quota thresholds are](#) on page 176

Prerequisites for managing user quotas using Operations Manager

To monitor and manage user quotas by using Operations Manager, you must 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 quota privilege to view user quota reports and events so that you can configure user quotas for volumes and qtrees.
- 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 > Quotas > Report**.

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 by using the Edit Quota Settings page in Operations Manager to increase or decrease the limits associated with them.

Before you begin

You must have ensured that the storage system meets the prerequisites.

Steps

1. Click **Control Center > Home > Group Status > Quotas > Report > User Quotas, All**.
2. Click any quota-related fields for the required quota and modify the values.

Configuring user settings using Operations Manager

You can configure user settings, such as the email address of users and quota alerts, and set user quota threshold using Operations Manager.

Steps

1. Click **Control Center > Home > Group Status > Quotas > Report > User Quotas, All**.
2. Click the **Edit Settings** link in the lower left corner.
3. Edit settings such as email address of the user, quota alerts and user quota alerts for full threshold and nearly full threshold, resource tag, and so on.

You can leave the email address field blank if you want the DataFabric Manager server to use the default email 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

You can use Operations Manager 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 are installed in compliance with the specified SAN setup guidelines.

Note: IBM has announced the end of availability for the SAN license for the DataFabric Manager server. Existing customers can continue to license the SAN option with the DataFabric Manager server. DataFabric Manager server customers should check with their IBM sales representative about other IBM N series SAN management solutions.

For more information about setting up a SAN, see the *Data ONTAP SAN Administration Guide for 7-Mode*.

Related information

The IBM N series support website: 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 (7-Mode environments only)

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 you can perform using DataFabric Manager Host Agent software (7-Mode environments only)

After you install DataFabric Manager Host Agent software on a client host along with the DataFabric Manager server, you can perform various management tasks, such as monitoring system information for SAN hosts, and creating and managing LUNs.

You can perform the following 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.

Related information

The IBM N series support website: 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 182

Information available on the LUN Details page on page 183

Tasks performed from the LUN Details page on page 184

Information about the FCP Target Details page on page 184

Information provided in the Host Agent Details page (7-Mode environments only) on page 185

Tasks you can perform from the Host Agent Details page (7-Mode environments only) on page 185

Prerequisites to manage targets and initiators

Operations Manager does not report any data for your targets and initiators if you have not completed the specified hardware and software requirements for your SAN setup.

For more information about the supported hardware platforms, see the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 17). For more information about specific software requirements, see the *OnCommand Unified Manager Installation and Setup Guide*.

Related information

The IBM N series support website: www.ibm.com/storage/support/nseries

Prerequisites to manage SAN hosts

You must ensure that a proper network connection and software is 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 correct version of SnapDrive software installed.
For information about 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

The IBM N series support website: 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 182

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 provided in the Host Agent Details page (7-Mode environments only)

The Host Agent Details page contains information such as the status of a SAN host, events that occurred on a SAN host, and devices related to a SAN host.

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 running
- The operating system and the DataFabric Manager Host Agent version, and the 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 the SAN host
- The devices related to the SAN host, such as the storage systems accessible from the SAN host
- Graphs that provide 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*.

Tasks you can perform from the Host Agent Details page (7-Mode environments only)

The Host Tools list on the Host Agent Details page enables you to perform various tasks, such as editing settings, creating a LUN, and refreshing monitoring samples.

Edit Settings	Displays the Edit Host Agent Settings page, where you can configure the login, password, administration transport, and 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. You should specify a value on this page only if you want to change the global setting.
Create a LUN	Takes you to the LUN creation page, which enables you to create a LUN.
Diagnose Connectivity	Automates troubleshooting of connectivity issues.
Refresh Monitoring Samples	Obtains current monitoring samples from the SAN host.

Manage Host Agent	Enables you to edit settings for the Host Agent. You can edit details such as monitoring and management of API passwords, and 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 187

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 93

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 217

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 of aggregate overcommitment 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 required, 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 can 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 required to avoid write errors due to insufficient space.

For details about volume space reservations and aggregate overcommitment, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Related information

The IBM N series support website: 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 the management server. You can set alarms to monitor the capacity and committed space of an aggregate. You can also perform corrective actions based on the event generated.

You can configure alarms to send notification when 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 management 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 (%)

Description: Specifies the percentage at which an aggregate is full.

Note: You can reduce the number of Aggregate Full Threshold events generated by setting an Aggregate Full Threshold Interval. This causes the management server to generate an Aggregate Full event only if the condition persists for the specified time.

Default value: 90 percent

Event generated: Aggregate Full

Event severity: Error

Corrective Action: Perform one or more of the following actions:

- To free disk space, request your users to delete files that are no longer required from the volumes contained in the aggregate that generated the event.
- Add one or more disks to the aggregate that generated the event.

Note: After you add a disk to an aggregate, you cannot remove it without first destroying all flexible volumes in the aggregate to which the disk belongs. You must destroy the aggregate after all the flexible volumes are removed from the aggregate.

- 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. Therefore, it is 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 Snapshot reserve, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

Aggregate Nearly Full (%)

Description: Specifies the percentage at which an aggregate is nearly full.

Default value: 80 percent

The value for this threshold must be lower than the value for Aggregate Full Threshold for the management server to generate meaningful events.

Event generated: Aggregate Almost Full

Event severity: Warning

Corrective action: Perform one or more of the actions mentioned in the Aggregate Full section.

Aggregate Overcommitted (%)

Description: Specifies the percentage at which an aggregate is overcommitted.

Default value: 100 percent

Event generated: Aggregate Overcommitted

Event severity: Error

Corrective action: Perform one or more of the following actions:

- Create new free blocks in the aggregate by adding one or more disks to the aggregate that generated the event.

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 in the aggregate to which the disk belongs. You must destroy the aggregate after all the flexible volumes are destroyed.
- Temporarily free some already occupied blocks in the aggregate by taking unused flexible volumes offline.

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, it requires the space again.
- Permanently free some already occupied blocks in the aggregate by deleting unnecessary files.

Aggregate Nearly Overcommitted (%)

Description: Specifies the percentage at which an aggregate is nearly overcommitted.

Default value: 95 percent

The value for this threshold must be lower than the value for Aggregate Full Threshold for the management server to generate meaningful events.

Event generated: Aggregate Almost Overcommitted

Event severity: Warning

Corrective action: Perform one or more of the actions provided in the Aggregate Overcommitted section.

**Aggregate
Snapshot Reserve
Nearly Full
Threshold (%)**

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.

Default value: 80 percent

Event generated: Aggregate Snapshot Reserve Almost Full

Event severity: Warning

Corrective action: None

There is no way to prevent Snapshot copies from consuming disk space greater than the amount reserved for them. If you disable the aggregate Snapshot `autodelete` option, it is important to maintain a large enough reserve.

See the *Operations Manager Help* for instructions on how to identify the Snapshot copies you can delete. For more information about the Snapshot reserve, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

**Aggregate
Snapshot Reserve
Full Threshold
(%)**

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.

Default value: 90 percent

Event generated: Aggregate Snapshot Reserve Full

Event severity: Warning

Corrective action: None

There is no way to prevent Snapshot copies from consuming disk space greater than the amount reserved for them.

Note: A newly created traditional volume is tightly coupled 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 134

Related information

[The IBM N series support website: 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

The 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 threshold values:

Volume Full Threshold (%)	<p>Description: Specifies the percentage at which a volume is considered full.</p> <p>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, the DataFabric Manager server generates a Volume 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, the DataFabric Manager server generates the Volume Full events.• If the threshold interval is greater than the volume monitoring interval, the 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. <p>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.</p> <p>Default value: 90</p> <p>Event generated: Volume Full</p>
----------------------------------	--

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 for 7-Mode*.

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 the 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.

Volume Space Reserve Depleted Threshold (%)	<p>Default value: 80</p> <p>Event generated: Volume Space Reservation Nearly Depleted</p> <p>Event severity: Warning</p> <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>
Volume Quota Overcommitted Threshold (%)	<p>Default value: 90</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>

Perform one or more of the actions mentioned in Volume Quota Overcommitted.

**Volume Growth
Event Minimum
Change (%)**

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, the DataFabric Manager server generates a Volume Growth Abnormal event.

Default value: 1

Event generated: Volume Growth Abnormal

**Volume Snap
Reserve Full
Threshold (%)**

Description: Specifies the value (percentage) at which the space that is reserved for taking volume Snapshot copies is considered full.

Default value: 90

Event generated: Volume Snap Reserve Full

Event severity: Error

Corrective action: None

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 autodelete 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 *Operations Manager Help*.

**User Quota Full
Threshold (%)**

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 `/etc/quotas` file. If this limit is exceeded, the DataFabric Manager server generates a User Disk Space Quota Full event or a User Files Quota Full event.

Default value: 90

Event generated: User Quota Full

**User Quota Nearly
Full Threshold (%)**

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 `/etc/quotas` file. If this limit is exceeded, the DataFabric Manager server generates a User Disk Space Quota Almost Full event or a User Files Quota Almost Full event.

	Default value: 80
	Event generated: User Quota Almost Full
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 201

Related tasks

Creating alarms on page 134

Related information

The IBM N series support website: 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 196

Related references

[Guidelines for configuring alarms](#) on page 134

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.

The DataFabric Manager server provides 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 the 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, the 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, the DataFabric Manager server generates Qtree Full events. • If the threshold interval is greater than the volume monitoring interval, the 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 persists 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:

- Request users to delete files that are no longer required, 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:

- Request users to delete files that are no longer required, to free disk space.
- Increase the hard disk space quota for the qtree.

Related tasks

[Creating alarms](#) on page 134

Related information

The IBM N series support website: www.ibm.com/storage/support/nseries

How Operations Manager monitors volumes and qtrees on a vFiler unit (7-Mode environments only)

You can use Operations Manager to monitor volumes and qtrees in 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 the 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 (7-Mode environments only)

With Operations Manager, you can view the volumes and qtrees in 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 for 7-Mode*.

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

The IBM N series support website: 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

Snapshot copy monitoring and space management help you monitor and generate reports on Snapshot copies. You can determine how they influence your space management strategy.

By using the 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.

What storage chargeback reports are

You can create storage chargeback reports through Operations Manager to collect information about the amount of space used or allocated in specific storage objects (storage systems, volumes, qtrees) or a group of storage objects. Chargeback reports are useful if your organization bills other organizations, groups, or users in your company for the storage services they use.

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 systems, volumes, qtrees, or users. The billing reports contain information such as the 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. These reports can be generated in different formats, such as Perl, .txt, comma-separated values (.csv), Excel (.xls), and .xml files.

Storage chargeback reports provide an efficient way to track the space used and space that is allocated for generating bills based on your specifications.

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 chargeback report provides 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 use the `dfm report view -F perl report_name` command.

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
- 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 (clustered environments only)
- clusters-chargeback-last-month (clustered environments only)
- clusters-chargeback-allocation-this-month (clustered environments only)
- clusters-chargeback-allocation-last-month (clustered environments only)
- vservers-chargeback-this-month (clustered environments only)
- vservers-chargeback-last-month (clustered environments only)
- vservers-chargeback-allocation-this-month (clustered environments only)
- vservers-chargeback-allocation-last-month (clustered environments only)
- filers-chargeback-this-month (7-Mode environments only)
- filers-chargeback-last-month (7-Mode environments only)
- filers-chargeback-allocation-this-month (7-Mode environments only)
- filers-chargeback-allocation-last-month (7-Mode environments only)

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: $\text{Annual Rate} / 365 * \text{number of days in the billing period}$.
- 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: $\text{Annual Rate} / 12$.

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

In 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 the global group, the object is deleted only from that group; the 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

You can view a list of all the storage objects that are deleted from the DataFabric Manager server database various reports.

Following are the reports of deleted storage objects:

- Storage Systems, Deleted
- File Systems, Deleted
- Volumes, Deleted
- Qtrees, Deleted
- Aggregates, Deleted
- SAN Hosts, Deleted
- LUNs, Deleted
- Vservers, Deleted (clustered environments only)
- vFilers, Deleted (7-Mode environments only)

You can access these reports 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 you can perform from Operations Manager

You can use Operations Manager to manage your storage systems, including creating groups, viewing information about systems, events and quotas, and configuring alarms.

After the DataFabric Manager server is installed, it begins the process of discovering, monitoring, and gathering data about your supported storage systems. However, before you use the data to simplify your network administration tasks, you have to understand the different ways you can use Operations Manager to manage your storage systems.

You can use Operations Manager to perform the following tasks:

- Create groups.
- View the status of and obtain reports and information for a group of systems.
- View information about individual systems.
- Access the console of a storage system.
- View and respond to events.
- Configure alarms to 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.
- Enter values in custom comment fields.
- View user, qtree, and group quotas.
- Edit user quotas.
- View the HA pair status and perform takeover and giveback operations if the storage system is an active/active controller. (7-Mode environments only)
- Link to FilerView for a selected storage system or vFiler unit. (7-Mode environments only)

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 data and reports for storage systems and vFiler units (7-Mode environments only)

By using Operations Manager, you can view storage system and vFiler unit reports. You can view global and group information and view individual system data in detail from the Member Details report pages.

You can view storage system related information from **Control Center > Member Details > Physical Systems > Report list**.

You can view vFiler unit related information from **Control Center > Member Details > Virtual Systems > Report list**.

Tasks performed by using the storage system and vFiler unit report pages (7-Mode environments only)

You can view system data for all groups, generate spreadsheet reports, get information about storage systems, and open FilerView.

Similar to the other Operations Manager Control Center tab pages, the Appliances and vFiler reports enable you to view various 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
- Open 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 HA pair 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 by using Operations Manager.

You can use the Edit Storage Controller Settings page to specify or change storage system or vFiler unit settings. You can set global values for many settings by using the Options page. You do not have to modify storage system level 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 field specifies 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 for 7-Mode</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 that 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 two monitoring intervals. You can configure threshold intervals only for specific thresholds, as listed on the Options page.

Related information

The IBM N series support website: www.ibm.com/storage/support/nseries

What the storage controller tools list is (7-Mode environments only)

You can view the storage system or controller details by using storage controller tools in Operations Manager.

You can use storage controller tools to set up parameters that are required 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, and login and password.
- 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 in the lower left pane of Operations Manager.

What the cluster tools list is (clustered environments only)

You can view and modify the cluster or cluster controller details by using cluster tools in Operations Manager.

You can use cluster tools to set up parameters that are required to communicate with a cluster. You can perform the following tasks using cluster tools:

- Modify settings such as the primary IP address, monitoring options, and management options such as login and password.
- 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 on the lower left display of Operations Manager.

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

[*Methods for remote configuration of storage systems \(7-Mode environments only\)*](#) on page 223

[*DataFabric Manager server CLI to configure storage systems*](#) on page 223

[*Prerequisites for running remote CLI commands from Operations Manager*](#) on page 223

[*What the remote platform management interface is*](#) on page 241

Related tasks

[*Running commands on a specific storage system*](#) on page 224

[*Running commands on a group of storage systems from Operations Manager*](#) on page 224

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 HA pairs using the DataFabric Manager server (7-Mode environments only)

You can monitor and manage HA pairs from the cluster console of Operations Manager.

The cluster console enables you to view the status of an HA pair (controller and its partner) and perform takeover and giveback operations between the controllers.

For detailed information about HA pairs, see the *Data ONTAP High Availability and MetroCluster Configuration Guide for 7-Mode*.

Related information

The IBM N series support website: www.ibm.com/storage/support/nseries

Viewing the status of an HA pair by using cluster console (7-Mode environments only)

An authentication method must be set up for the DataFabric Manager server to authenticate the controller on which takeover and giveback operations should 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 HA pair.

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 HA pair.
4. Click **View Cluster Console** under **Storage Controller Tools**.

What the takeover tool does (7-Mode environments only)

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.

When 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 MetroCluster configurations only and is the equivalent of running the `cf forcetakeover -f -d` command. You can use this option if the partner cannot be recovered.

Note: The **Force a Takeover** and **Takeover After a Disaster** options are also available when the interconnect between the controller and its partner is down. It enables you to manually take over the partner.

After you select an option, the Status option on the Cluster Console page displays the status of the takeover operation. After 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 updated appropriately to indicate the HA pair operation each controller can perform.

What the giveback tool does (7-Mode environments only)

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 perform a giveback operation to its partner.

After you select **Giveback** for the controller, the Giveback page is displayed. You can select 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 the services and cancels 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.

After you have selected an option, the **Status** option on the Cluster Console page displays the status of the giveback operation. After the giveback 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 updated to indicate the HA pair operation each controller can 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*.

Methods for remote configuration of storage systems (7-Mode environments only)

The DataFabric Manager server provides three methods that you can use to remotely configure a storage system: accessing the storage system, accessing FilerView, and using the multiple storage system remote configuration feature.

When you monitor your storage system, you might have to alter the configuration settings on one or more storage systems. The DataFabric Manager server provides three methods that you can use to 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` or `ssh` to run a command on storage systems. Therefore, you must have enabled `rsh` access to your storage system 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 for 7-Mode*.

Note: The command operation uses only `ssh` to run remote CLI commands on clusters.

Related information

The IBM N series support website: www.ibm.com/storage/support/nseries

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 (clustered environments only)

You can remotely configure a cluster by using the DataFabric Manager server. As part of monitoring your clusters, you might have to run 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 (clustered environments only)

You can run commands on a specific cluster by using Operations Manager. However, you cannot execute RLM commands on clusters.

Before you begin

The user name and password must be set at the cluster level from the Edit Cluster Settings page or using the corresponding commands from the command-line interface (CLI).

Steps

1. Click **Control Center > Home > Member Details > Physical Systems > Report > Clusters, All**.
2. Click the specific cluster to go to the **Cluster Details** page.
3. Click the **Run a Command** link under **Cluster Tools**.
4. In the **Run Command** page, enter the command in the **Appliance Command** box.
5. Click **Run**.

Running commands on a specific node of a cluster (clustered environments only)

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

The credentials for the cluster and its node must be set.

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 (for example, 10 minutes) on the CLI of the DataFabric Manager server, to execute run command jobs that are time consuming and result in large outputs (for example, >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.
5. In the **Storage Controller Details** page, click the **Run a Command** link under **Storage Controller Tools**.

The Run Command page is displayed.

6. In the **Run Command** page, enter the command in the **Remote Platform Management Command** box.
7. Click **Run**.

Storage system management using FilerView (7-Mode environments only)

With Operations Manager, you can connect to a storage system by 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 launch FilerView, Operations Manager opens 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 (7-Mode environments only)

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 specific storage system or vFiler unit, you can 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 (7-Mode environments only)

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 (7-Mode environments only)

MultiStore is a software feature that enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage units on the network.

Each storage unit 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 the same way a storage system does.

You can create multiple vFiler units by 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 for 7-Mode*.

Related information

The IBM N series support website: www.ibm.com/storage/support/nseries

Monitoring of vFiler units with the DataFabric Manager server (7-Mode environments only)

You can monitor vFiler units by 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 (7-Mode environments only)

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, the 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**
The 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. The 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. The DataFabric Manager server collects details about vFiler unit SnapMirror relationships from the hosting storage system. The 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 (7-Mode environments only)

You can perform management tasks on a vFiler unit by using the DataFabric Manager server.

- Discover vFiler units on a hosting storage system.
- Group vFiler units for consolidated reporting.
- Monitor the health and general status of vFiler units.
- 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 administration.
- 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 hosting storage systems or only the vFiler units. By default, if you run a command on a vFiler unit, the settings of the hosting storage system is 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 (7-Mode environments only)

You can remotely configure multiple storage systems by 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 (7-Mode environments only)

You might prefer to centrally manage your storage system and vFiler configuration `/etc` files and registry options. Operations Manager enables you to create and manage configuration files that contain the 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 then push all of the configuration settings or some of the configuration settings to other storage systems, or groups of storage systems and vFiler units.

You can also ensure that the storage system and vFiler configuration conforms with the configuration that is pushed to it from Operations Manager.

Prerequisites to applying configuration files to storage systems and vFiler units (7-Mode environments only)

You must meet a set of requirements before applying configuration files to a group of storage systems and vFiler units.

You must meet the following requirements:

- 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 17).

List of access roles to manage storage system configuration files (7-Mode environments only)

You require specific access roles to perform management tasks with the storage system configuration files.

Task	Access role
Creating configuration files	Global Write
Deleting configuration files	Global Delete
Editing configuration files	Global Write
Exporting configuration files	Global Read
Importing configuration files	Global Write
Upgrading or reverting configuration file versions	Global Write

List of tasks for configuration management (7-Mode environments only)

You can perform a set of configuration management tasks by using the storage system configuration management feature.

You can perform the following tasks:

- 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 with 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 Data ONTAP version.
- 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 (7-Mode environments only) on page 233

What configuration files are (7-Mode environments only)

A configuration file is a set of configuration settings that storage systems in one or more groups can share.

Configuration files exist independently of groups and can be shared among groups. You can use Operations Manager to pull a configuration file from storage systems and save it.

What a configuration plug-in is (7-Mode environments only)

A configuration plug-in is an add-on library in a zip file that is required for the DataFabric Manager server to manage Data ONTAP. Configuration plug-ins enable you to upgrade or revert a configuration file that is stored in the DataFabric Manager server database to a different version.

For each Data ONTAP version, a Data ONTAP plug-in is provided.

Comparison of configurations (7-Mode environments only)

Operations Manager enables you to compare your configuration file settings with those of a template configuration. You can view the configuration comparison results in a report format, and use this report to identify the configuration settings that do not conform to those of the standard template.

You can also compare storage systems, vFiler units or groups of storage systems, vFiler units with a configuration file, and create jobs to obtain the comparison results. You can use Operations Manager to access the comparison job results.

Verification of a successful configuration push (7-Mode environments only)

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 (7-Mode environments only)

You can monitor configuration changes made to the storage systems by 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 triggers host discovery).

3. After the event is generated, select **Fix** in the **Event details** page to view the differences in the configuration setting.
4. 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 (7-Mode environments only)

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 a certain number of storage systems and have one or more files containing the required configuration settings. These configuration settings are listed in files called *configuration files*.

List of tasks for managing configuration groups (7-Mode environments only)

After you have added configuration files to a group, you can manage your configuration groups by performing a set of tasks, including removing an existing configuration file from a group's list, changing the order of files, specifying configuration overrides, and viewing Groups configuration summary.

You can perform the following 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 `dfm config push -R` command to specify a new retry limit.

- 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 job.

Guidelines for creating configuration groups (7-Mode environments only)

Before you create configuration groups, you must consider the versions of operating systems and Data ONTAP, and storage systems belonging to configuration resource groups.

- You can group storage systems running different operating system versions in the same configuration resource group.
- A storage system can belong to only one configuration resource group, but it can belong to other non-configuration resource groups.
- A storage system that is a member of a configuration resource group can belong to one or more groups.

Creating configuration resource groups (7-Mode environments only)

You can create configuration resource groups by creating an empty group and populating 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 the configuration files are 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 (7-Mode environments only)

You can specify a parent configuration resource group from which you 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 require to meet your deployment requirements. When you assign a parent group, only the parent group's configuration files are inherited, not the storage systems in the parent group.

Parent group guidelines (7-Mode environments only)

You should be aware of some guidelines about the inheritance from the parent files and the parent hierarchy before you assign a parent group.

- When you assign a parent, only the parent group's configuration files are inherited. The storage systems in the member group are not inherited.
- Parent groups can include 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: You must review the settings in a parent group so that they do not have unintended consequences on your storage systems.

When to assign parent groups (7-Mode environments only)

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 unnecessary settings before assigning a parent group.

You might not want to assign a parent group 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 (7-Mode environments only)

You must be aware of the properties of configuration files that are acquired from parent groups. For example, the configuration files are initially read-only.

When you include configuration files from another group, consider the following information:

- 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 (7-Mode environments only)

To configure storage systems or vFiler units, you must have enabled SNMP on the storage systems and the DataFabric Manager server must have 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 with 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 to the storage system or to the group.
8. Verify that the configuration changes are updated by reviewing the status of the push jobs.

Maintenance and management

You can configure and maintain the DataFabric Manager server through the CLI.

Accessing the DataFabric Manager server through the CLI

You can access the command-line interface (CLI) through 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 by performing the following action:

If you want to access the CLI on a... Then...	
Local Windows system	Use the command prompt window of the Windows system by clicking Start > Run.
Local Linux system	Use any shell prompt on your system.
Remote Windows or Linux system	<div><p>a. Start a Telnet connection from the remote host to the DataFabric Manager server that you want to access by entering the following command:</p><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, you should 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><p>b. Authenticate using your user name and password.</p></div>

Where to find information about the DataFabric Manager server commands

You can find information about the DataFabric Manager server commands either by accessing the help or by using the `dfm help` command.

- You can access the man pages through the table of contents of the *Operations Manager Help*. On a Linux system, you can access the man pages by running the `source/opt/IBMdfm/bin/vars.sh` command, and then the `man dfm` command.
- You can 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>

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-orid>
<admin-name-or-id>TOM-XP\dfmuser</admin-name-or-id></rbacadmin-
role-add>

Permissions for accessing the audit log file

Windows and UNIX users have separate permissions for accessing the `audit log` file.

The following users have both read and write permissions:

- Linux: -rw - - “root” users on Linux
- Windows users in the Administrators group

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.

Note: If the interpreter is installed after the DataFabric Manager server is installed, restart the DataFabric Manager services by using the following commands:

- `dfm service stop`
- `dfm service start`

The `dfm report -F Perl` command in the DataFabric Manager server generates reports for direct inclusion in Perl scripts. You can import 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 the DataFabric Manager server can execute scripts, you can use commands that are available as part of the DataFabric Manager server CLI to execute your script.

Your script can perform the following tasks:

- Use the `dfm run` command to execute commands on storage systems for which credentials are specified on the DataFabric Manager server.

- 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 have to package it as a .zip file.

The .zip file must contain the script, any data files that are required 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 (7-Mode environments only)

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 capability to run prebackup and postbackup scripts on specific primary directories, before and after data is backed up from those directories.

Note: The timeout value for backup scripts is 3600 seconds. This value cannot be modified.

For more information about 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 viewing the modifications of Performance Advisor are suspended during the backup process (7-Mode environments only).

Two types of backups are available:

- **Archive**
This 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 move an archive-based backup to a different system and restore it. This backup process is time-consuming.
- **Snapshot**
This 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.

Limitations 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 the 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 17).

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 select either archive-based backup or Snapshot-based backup.

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, and 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 by 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, you should use the `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 the restore operation 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 a 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 by using the `dfm backup restore` command.

Steps

1. Copy the backup file to 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.

The following message is displayed when the restore process is successful: `Completed restore`

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 the 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.

Result

The following message is displayed when the restore process is completed successfully: Completed restore.

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 (7-Mode environments only)

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 (7-Mode environments only)

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 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 (7-Mode environments only)

You must use SnapDrive to configure the DataFabric Manager server on Linux for disaster recovery.

Related information

[Technical Report 3655: Disaster Recovery Support for DataFabric Manager Data Using SnapDrive](#)

Prerequisites for disaster recovery support (7-Mode environments only)

You must ensure that you meet the prerequisites for disaster recovery support for the DataFabric Manager server data.

- You must be a Windows domain administrator.

- You must have the N series Management Console data protection capability license.
- You must be using 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 the Data ONTAP version for SnapDrive, see the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 17).
- 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` from the command-line interface.
- You should 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 that is used to administer SnapDrive must have full access to the Windows domain to which both the source and destination storage systems belong.

Related information

IBM N series interoperability matrix: www.ibm.com/systems/storage/network/interophome.html

Setting up the DataFabric Manager server for disaster recovery

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 performing 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. Configure the DataFabric Manager server for disaster recovery using the N series Management Console data protection capability features by performing 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 N series Management Console data protection capability Disaster Recovery license or 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 you must 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 the SnapMirror location information from the command output.

You require this information when you use 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 the DataFabric Manager server services.

The DataFabric Manager server services must be enabled only during disaster recovery, which you can do by using the `dfm datastore mirror connect` command.

Related tasks

[Scheduling database backups from Operations Manager](#) on page 247

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. Connect to the LUN using the Microsoft Management Console (MMC) **Connect Disk** wizard on the secondary storage system by completing 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 the 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 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

The IBM N series support website: 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 set of tasks to fail back the DataFabric Manager server services to the primary site.

Steps

1. Ensure that the DataFabric Manager server data at the source storage system is synchronized with the data in the destination storage system by completing the following steps:
 - a) Run the `dfdrm mirror list` command to find the 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 the 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) If the primary storage system is destroyed during a disaster, run the `dfdrm mirror initialize` command to create a relationship from the secondary storage system to the new primary storage system.
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, perform the following procedure:
 - a) Connect to the LUN using the **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. If the SnapMirror relationship is removed during the failback process, run the `snapmirror resync` command to resynchronize the data at the storage system level.
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 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 (7-Mode environments only)

If you do not have 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, *Disaster Recovery Support for DataFabric Manager Data Using SnapDrive*.

Related information

[*Technical Report 3655: 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 by using Operations Manager to monitor the health of storage systems.

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 the messages should contain.

Note: If you change the setting from "complete" to "minimal," any complete AutoSupport message that is not sent is cleared from the outbound 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 resend the AutoSupport notification, if previous attempts have failed.
7. Enter the time the system should wait before resending a failed AutoSupport notification.
8. (7-Mode environments only) Select **Include** to include the Performance Advisor AutoSupport data along with the DataFabric Manager server AutoSupport data.
9. Select **Include** to include 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 the DataFabric Manager server logs

You can access the DataFabric Manager server logs through Operations Manager or through the CLI.

You can access the DataFabric Manager server logs through the Diagnostics page. To access this page, you must 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 by using different directories, depending on whether you are using a Windows or a UNIX workstation:

- On a Windows workstation, you must enter `installation_directory\dfm\log`.
- On a UNIX workstation, you must 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 259

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 for the following reasons:

- 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 that are 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.

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.

Custom reports do not display any data

In certain scenarios, when you create a custom report and view it, the report lists only the field names and does not display any data.

Cause:

The options you select in the **Base Catalog** or **Related Catalogs** fields are not supported for the mode that is running on the system or object that is managed by OnCommand Unified Manager Core Package.

The following objects or items are supported in clustered environments only, and if you select these in the **Base Catalog** or **Related Catalogs** fields in a 7-Mode system, the custom report does not display any data:

- Cluster
- IfGrp
- Port
- LogicalIf
- Vserver

The following objects or items are supported in 7-Mode only, and if you select these in the **Base Catalog** or **Related Catalogs** fields in clustered environments, the custom report does not display any data:

- Dataset
- DPPolicy
- Host
- HostUsergroup
- HostUser
- HostLocalUser
- HostDomainUser
- PrimaryDirectory
- ResourcePool
- SANHost
- SnapmirrorRelationship
- StorageArray
- StorageArrayPort
- ArrayLUN
- vFiler

Corrective action:

Do not select objects or items for the **Base Catalog** and **Related Catalogs** fields that are not supported for the mode of the object or item you are currently managing.

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 \(7-Mode environments only\)*](#) on page 230

Data ONTAP issues impacting protection on vFiler units (7-Mode environments only)

You might encounter issues when creating SnapMirror and SnapVault relationships using vFiler units in Data ONTAP 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 uses 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 resynchronization 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.

Report fields and performance counters (7-Mode environments only)

Report fields and performance counters for filer catalogs (7-Mode environments only)

The DataFabric Manager server creates reports for filer catalogs.

The following table lists and describes the fields in the filer catalog and lists the 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 (7-Mode environments only)

The DataFabric Manager server creates reports for vFiler catalogs.

The following table lists and describes the fields of the vFiler catalog and lists the performance counters:

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

Field	Name/description	Performance counter
vFiler.MiscOps	vFiler Miscellaneous Ops/Sec	vfiler:vfiler_misc_ops
vFiler.NetThroughput	vFiler Network Throughput (KB/Sec)	vfiler:vfiler_nw_throughput
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 (7-Mode environments only)

The DataFabric Manager server creates reports for volume catalogs.

The following table lists and describes the fields of the volume catalog and lists the performance counters:

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

Field	Name/description	Performance counter
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
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

Field	Name/description	Performance counter
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
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 (7-Mode environments only)

The DataFabric Manager server creates reports for qtree catalogs.

The following table lists and describes the fields of the qtree catalog and lists the performance counters:

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

Field	Name/description	Performance counter
Qtree.PerfViolation Period	Qtree Perf Threshold Violation Period (Sec)	Not applicable

Report fields and performance counters for LUN catalogs (7-Mode environments only)

The DataFabric Manager server creates reports for LUN catalogs.

The following table lists and describes the fields of the LUN catalog and lists the performance counters:

Field	Name/description	Performance counter
LUN.TotalOps	LUN Total Ops/Sec	lun:total_ops
LUN.ReadOps	LUN Read Ops/Sec	lun:read_ops
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 (7-Mode environments only)

The DataFabric Manager server creates reports for aggregate catalogs.

The following table lists and describes the fields of the aggregate catalog and lists the performance counters:

Field	Name/description	Performance counter
Aggregate.TotalOps	Aggregate Total Ops/Sec	aggregate:total_transfers

Field	Name/description	Performance counter
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 (7-Mode environments only)

The DataFabric Manager server creates reports for disk catalogs.

The following table lists and describes the fields of the disk catalog and lists the performance counters:

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 used by the DataFabric Manager server

The DataFabric Manager server uses various networking protocols and port numbers to communicate with storage systems, host agents (7-Mode environments only), and Open Systems SnapVault agents (7-Mode environments only).

DataFabric Manager server communication

You might have to enable both HTTP and HTTPS transports if multiple administrators are accessing the workstation from different locations. Firewalls, non-trusted environments, or other circumstances might require a combination of HTTP and HTTPS transports.

Note: Reconfiguring these options through Operations Manager results in a message instructing you to restart the HTTP service from the CLI. You should use the `dfm service start http` command to restart the HTTP service.

DataFabric Manager server access to storage systems

The DataFabric Manager server uses a set of protocols and port numbers to access the storage systems.

Protocol	UDP port	TCP port
HTTP	N/A	80
HTTPS	N/A	443
RSH	N/A	514
SSH	N/A	22
Telnet	N/A	23
SNMP	161	N/A

DataFabric Manager server access to host agents (7-Mode environments only)

There are a set of protocols and port numbers used by the DataFabric Manager server to access your host agents. The DataFabric Manager server uses HTTP and HTTPS to access the host agents. HTTP uses the TCP port 4092, and HTTPS uses the TCP port 4093.

DataFabric Manager server access to Open Systems SnapVault agents (7-Mode environments only)

The DataFabric Manager server uses a set of protocols and port numbers to access the 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 SAN Administration Guide for 7-Mode*.

Note: DataFabric Manager server customers should check with their sales representative regarding other SAN management solutions.

Related information

The IBM N series support website: www.ibm.com/storage/support/nseries

Discovery of SAN hosts by the DataFabric Manager server (7-Mode environments only)

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 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 DataFabric Manager 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, you must 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

The IBM N series support website: 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, you should contact your sales representative for information about purchasing the license.

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 *Data ONTAP SAN Administration Guide for 7-Mode*.

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 Configuration Guide* (called *Fibre Channel and iSCSI Configuration Guide* in Data ONTAP 8.1 and earlier).

Note: For a list of supported Brocade switches, see the *SAN Configuration Guide* (called *Fibre Channel and iSCSI Configuration Guide* in Data ONTAP 8.1 and earlier) at the N series support website (accessed and navigated as described in [Websites](#) on page 17).

- 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 the correct version of the SnapDrive software installed on it, for LUN management by using the DataFabric Manager server.
To identify 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

The IBM N series support website: 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)	<p>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:</p> <ul style="list-style-type: none"> • 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 WWPN 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 (7-Mode environments only)

You can view the SAN host status, version details of the operating system and DataFabric Manager 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 when the last sample was collected and the configured polling interval for the SAN host
- Graphs that provide information about 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 WWPN 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.
- Connect to FilerView.

The LUN Details page provides a Manage LUNs with FilerView option in the Tools list, as shown in the previous example.

This option enables you to access FilerView, the Web-based UI of your storage system. You can perform the following LUN management functions from FilerView:

- Add or delete a LUN
- Modify configuration settings such as the size or status of a LUN
- Map a LUN to or unmap a LUN from initiator groups
- Create or delete initiator groups

The Tools list on the LUN Details page displays two options for FilerView. The `Invoke FilerView` option connects you to the main window of the UI of your storage system and the `Manage LUNs with FilerView` option connects you directly to the Manage LUNs window.

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.

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.
The default monitoring intervals are as follows:

- 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, you must 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 DataFabric 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 such a large amount of 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.

Copyright and trademark information

Copyright ©1994 - 2013 NetApp, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2013 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 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

7-Mode environments

- identifying information [23](#)
- Operations Manager topic list [24](#)

A

- access check
 - for application administrators [75](#)
- access check, for application administrator [65](#)
- access control
 - configuring for VFiler units [65](#)
 - on groups of SAN components [292](#)
 - precedence of, global and group [72](#)
- access roles
 - for managing configuration files [231](#)
- accessing
 - DataFabric Manager server through the CLI [237](#)
- Active Directory
 - user group accounts [68](#)
- active/active configuration
 - managing with DataFabric Manager server [221](#)
- adding
 - clusters [62](#)
- administrative users
 - adding [68](#)
- administrator access
 - managing [76](#)
 - requirements for managing [76](#)
- administrators
 - creating accounts [75](#)
 - default accounts [66](#)
 - everyone account [66](#)
 - types of access controls [75](#)
- Aggregate Details page [191](#), [192](#)
- aggregate full threshold [192](#)
- Aggregate Nearly Full threshold [192](#)
- aggregate nearly overcommitted threshold [192](#)
- aggregate overcommitted threshold [192](#)
- Aggregate Overcommitted threshold [192](#)
- aggregates
 - capacity graphs [192](#)
 - historical data [192](#)
 - name format [192](#)
 - relation to traditional volume [196](#)
- aggregates space availability [191](#)

- alarm notification
 - customize [136](#)
- alarms
 - acknowledging [137](#)
 - alerts and alarms, differences [138](#)
 - configuration guidelines [134](#)
 - creating [134](#)
 - e-mail messages [138](#), [141](#)
 - SNMP traps as [106](#)
 - testing [135](#)
 - trap format example [137](#)
- alerts
 - See user alerts [139](#)
 - using the mailmap file [139](#)
- annual charge rate, storage chargeback reports [211](#)
- Appliance Details page
 - tasks [217](#)
- appliances
 - console, connecting [220](#)
 - grouping [215](#)
- archive-based backup
 - restoring from file [249](#)
- archived reports
 - Report Archival Directory [151](#)
- audit log file
 - access permissions [241](#)
- authentication
 - LDAP [22](#)
 - methods [22](#)
 - supported servers [22](#)
- AutoSupport
 - configuring [258](#)

B

- backing up
 - access requirements [246](#)
 - deleting from Web-based user interface [248](#)
 - disabling schedules [248](#)
 - displaying information about [248](#)
 - starting from Web-based user interface [246](#)
 - storage and sizing [245](#)
- backup
 - directory location [245](#)
 - exporting to another location [249](#)
 - scheduling database [247](#)

- backup scripts
 - scripts
 - prebackup and postbackup [244](#)
 - what they do [244](#)
- backup types
 - archive backup [244](#)
 - Snapshot backup [244](#)
- billing cycle, storage chargeback reports [211](#)

C

- CA-signed certificate [163](#)
- capacity thresholds [201](#)
- catalogs
 - See reports [144](#)
- Certificate Authority (CA)-signed [163](#)
- Certificate Signing Request [163](#)
- certificate types [163](#)
- certificates
 - CA-signed [164](#)
 - definition of CA-signed [163](#)
 - obtaining [164](#)
 - role of, in secure connections [163](#)
 - self-signed, generating [164](#)
 - self-signed, security precautions for [163](#)
- certificates in DataFabric Manager server [163](#)
- chargeback reports
 - defined [207](#)
 - formats [208](#)
- CLI [237](#)
- Clone List report [206](#)
- clone volumes
 - Clone List report [206](#)
 - identifying [206](#)
 - parents, relation to [205](#)
- cluster details
 - cluster reports [103](#)
 - information available [103](#)
- cluster management
 - browse cluster components [104](#)
 - Cluster Details page [104](#)
 - cluster hierarchy information about cluster objects
 - CPU usage [104](#)
 - cluster objects
 - user access [77](#)
 - network traffic [104](#)
 - tasks you can perform [104](#)
 - user access control [77](#)
 - view physical system [104](#)
 - view utilization of physical resources [104](#)
 - volume capacity used [104](#)
- cluster monitoring [103](#)
- cluster objects
 - grouping [100](#)
- cluster tools
 - use of [219](#)
- cluster-management LIF
 - defined [103](#)
- cluster-related objects [98](#)
- clustered environment
 - Operations Manager topic list [38](#)
- clustered environments
 - identifying information [23](#)
- clusters
 - adding [62](#)
 - automatic discovery [61](#)
 - defined [98](#)
 - remote configuration [224](#)
 - running commands [225](#)
 - running commands in a node [225](#)
 - supported monitoring tasks [62](#)
 - unsupported monitoring tasks [62](#)
- commands
 - running them on a cluster [225](#)
- configuration
 - for disaster recovery [251](#)
- configuration details
 - storage arrays [64](#)
- configuration file properties
 - acquired from parent groups [235](#)
- configuration files
 - access roles for managing [231](#)
 - comparison with template configuration [232](#)
 - described [232](#)
 - manage [230](#)
 - prerequisites for storage systems and vFiler units [230](#)
 - verification of push operation [232](#)
- configuration groups
 - guidelines
 - for creating configuration groups [234](#)
 - guidelines for creating [234](#)
 - tasks to manage [233](#)
- configuration management
 - list of tasks for [231](#)
 - of storage systems and vFiler units [230](#)
- configuration plug-ins
 - defined [232](#)
- configuration report [64](#)
- configuration resource group

- what it is [233](#)
- configuration resource groups
 - assigning a parent group [235](#)
 - creating [234](#)
- configuring
 - AutoSupport [258](#)
 - configuring storage systems from FilerView [226](#)
 - custom reports [149](#)
 - multiple storage systems [236](#)
- connect storage system
 - Run Telnet tool [220](#)
- consolidated data
 - of storage systems and vFiler units [216](#)
- creating
 - configuration resource groups [234](#)
 - groups from a report [94](#)
- CSR [163](#)
- currency format, storage chargeback reports [210](#)
- custom comment fields
 - alarm notification by e-mail [136](#)
 - script format [137](#)
- custom reports
 - display issue [265](#)

D

- data export
 - types of [159](#)
- data protection capability
 - for disaster recovery [251](#)
- database
 - restoring from the Snapshot copy based backup [250](#)
- database backup
 - archive backup [244](#)
 - overview [244](#)
 - restoring [249](#)
 - scheduling [247](#)
 - Snapshot backup [244](#)
 - types of [244](#)
- DataFabric Manager
 - logging in to [66](#)
 - restoring backup file [249](#)
- DataFabric Manager Host Agent software
 - capabilities [179](#)
- DataFabric Manager server
 - access to host agents [277](#)
 - access to Open Systems SnapVault agents [277](#)
 - accessing through the CLI [237](#)
 - communication through HTTP and HTTPS [276](#)
 - controlling authentication through hosts.equiv [170](#)

- discovery by [49](#)
 - failing back the services [255](#)
 - groups [91](#)
 - information about commands [237](#)
 - monitoring configuration changes [232](#)
 - objects [91](#)
 - ports and protocols used [276](#)
 - protocols and port numbers [276, 277](#)
 - recovering the services [253](#)
 - restoring through the CLI [249](#)
 - setting up for disaster recovery [252](#)
- DataFabric Manager server options
 - global and local options [289](#)
 - objects [289](#)
- Day of the Month for Billing option [211](#)
- default role [69](#)
- deleted objects
 - what happens when they are deleted [212](#)
- deleting
 - SAN components [187](#)
 - SNMP settings [55](#)
- determine disk space [173](#)
- directory path for archive backups
 - changing [246](#)
- disaster recovery
 - configuration [251](#)
 - requirement for Linux [251](#)
 - requirements [251](#)
 - setting up the DataFabric Manager server [252](#)
 - using N series Management Console [251](#)
 - using SnapDrive [256](#)
- discovery
 - by the DataFabric Manager server [49](#)
 - DataFabric Manager Host Agent software [179](#)
 - of clusters [61](#)
 - of SAN hosts [278](#)
 - of storage systems [57](#)
 - SAN host [179](#)
 - storage arrays [63](#)
- discovery of
 - storage array ports [63](#)
- discovery options
 - guidelines for editing [59](#)
- Disk space hard limit option [175](#)
- Disk space soft limit option [175](#)
- documentation
 - identifying mode information [23](#)
- domain users
 - adding to a host [83](#)
 - defined [83](#)

- deleting from user groups [84](#)
- editing the settings of [84](#)
- pushing, to hosts [84](#)
- viewing settings of [83](#)
- viewing, on a host [83](#)

E

- editing
 - discovery options guidelines [59](#)
- email addresses
 - for user alerts, specifying [139](#)
- events
 - aggregate almost full [192](#)
 - aggregate almost overcommitted [192](#)
 - aggregate full [192](#)
 - aggregate overcommitted [192](#)
 - clearing configuration events [109](#)
 - definition of [109](#)
 - list of, and severity types [110](#)
 - list of, complete [110](#)
 - managing [109](#)
 - qtree full [203](#)
 - qtree nearly full [203](#)
 - user, notification of [138](#)
 - viewing [109](#)
 - volume almost full [196](#)
- Everyone, administrator account [66](#)
- exported data
 - files and formats of [160](#)

F

- FC switch
 - managing [289](#)
- FCP Target Details page [184](#)
- FCP targets [184](#)
- FCP topology [184](#)
- file formats
 - for exported data [161](#)
- FilerView
 - links to [105](#)
 - what it is [226](#)
- files
 - containing exported data [161](#)
- Files hard limit option [175](#)
- Files soft limit option [175](#)

G

- Gateway SAN-attached storage management [63](#)
- giveback tool
 - what it does [222](#)
- global access control
 - applying to administrator accounts [75](#)
 - description of [75](#)
 - precedence over group access control [69](#)
- global and local options
 - Host Agent Administration Transport [289](#)
 - SAN management [289](#)
- global group
 - summary of [77](#)
- global groups [92](#)
- global privileges
 - creating accounts [75](#)
- global roles [69](#)
- Global-only options
 - LUN monitoring interval [289](#)
 - managing SAN [289](#)
 - SAN device discovery [289](#)
- graphs
 - LUNs [183](#)
 - SAN hosts [185](#)
- group
 - adding members [96](#)
- group access control
 - applying [75](#)
 - description [75](#)
 - precedence over global access control [69](#)
- group access privileges
 - accounts [75](#)
- group membership [96](#)
- groups
 - changing thresholds [97](#)
 - cluster objects, creating [100](#)
 - creating, from a report [94](#)
 - defining roles [65](#)
 - definition of [91](#)
 - global [92](#)
 - guidelines for creating [95](#)
 - homogeneous [91](#)
 - managing [65](#)
 - reports (views) about [97](#)
- guidelines
 - creating groups [95](#)
 - editing discovery options [59](#)
 - for assigning parent groups [235](#)
 - for changing managed host options [167](#)

H

- HA configuration
 - viewing the status [221](#)
- hierarchical groups [93](#)
- homogeneous groups
 - types you can create [91](#)
- Host Agent
 - tasks you can perform [180](#)
- Host Agent Details page
 - tasks you can perform [185](#)
- host discovery [56](#)
- host-initiated discovery
 - defined [56](#)
- hosts.equiv file [170](#)
- hot backup mode
 - by using backup scripts [244](#)
- HTTPS
 - enabling [165](#)
 - requirements for [166](#)

I

- initiator group
 - connecting to FilerView [289](#)
 - managing [289](#)
 - viewing LUNs mapped to [183](#)
- interface groups
 - defined [98](#)

J

- jobs
 - deleting [90](#)
 - viewing the status of [90](#)
 - what they display [90](#)
- junctions
 - defined [98](#)

L

- LIFs
 - cluster management LIF, defined [98](#)
 - data LIF, defined [98](#)
 - node management LIF, defined [98](#)
- limitations
 - managed host options [168](#)
- local users
 - adding [79](#), [82](#)
 - defined [77](#)

- deleting [81](#)
 - editing passwords [82](#)
 - pushing [82](#)
 - pushing passwords [80](#)
 - viewing [78](#)
 - viewing settings [78](#)
- logical interfaces (LIFs)
 - See* LIFs
- logical resources
 - logical interface traffic [104](#)
 - utilization [104](#)
- logical storage
 - monitoring
 - qtree quotas [205](#)
 - monitoring qtree capacity threshold and events [203](#)
 - monitoring volume capacity threshold and events [196](#)
- logs
 - access to [259](#)
- LUN
 - connecting to FilerView [287](#)
 - managing [287](#)
- LUN Details page [183](#), [184](#)
- LUNs
 - deleting and undeleting [187](#)
 - destroying [184](#)
 - expanding [184](#)
 - graphs [183](#)
 - initiator groups mapped to [183](#)
 - reports in spreadsheet form [150](#)
 - status of [183](#)
 - stop monitoring [187](#)

M

- mailformat file [141](#)
- mailmap file
 - to specify email addresses [139](#)
- managed host options
 - guidelines for changing [167](#)
 - overview [167](#)
- Managed host options
 - limitations [168](#)
- methods
 - authentication [22](#)
- modification of passwords
 - issue [170](#)
- monitoring
 - clusters through Operations Manager [103](#)
 - flow chart of process [101](#)

- of clusters [62](#)
- process [101](#)
- query intervals [105](#)
- unsupported tasks for clusters [62](#)
- using SNMPv3 [53](#)
- monitoring options
 - guidelines for changing [106](#)
 - location to change [106](#)
- MultiStore
 - introduction to [227](#)

N

- namespaces
 - defined [98](#)
- NetApp Host Agent
 - SAN hosts [180](#)
- network credentials
 - modifying [54](#)
- network discovery
 - reasons for failure [262](#)
- nodes
 - running commands [225](#)

O

- objects
 - definition of [91](#)
- Open Systems SnapVault agents
 - protocols and ports used for accessing [277](#)
- Operations Manager
 - 7-Mode topic list [24](#)
 - access to [20](#)
 - clustered environment topic list [38](#)
 - deleting a backup from [248](#)
 - list of management tasks [214](#)
 - starting a backup from [246](#)
- options
 - Disk space hard limit [175](#)
 - Disk space soft limit [175](#)
 - Files hard limit [175](#)
 - Files soft limit [175](#)
 - for SAN management [289](#)
- order of applying privileges [72](#)
- overcommitment strategy [192](#)
- Overview [76](#)

P

- parent groups

- guidelines to assigning [235](#)
- when to assign [235](#)
- passwords
 - editing, of local users [82](#)
- performance reports
 - task you can perform with [148](#)
- Permanently stop monitoring [291](#)
- physical resources
 - controller [105](#)
 - CPU [105](#)
 - network traffic [105](#)
 - utilization [105](#)
- physical storage
 - aggregate capacity thresholds and events [192](#)
- ports
 - used by the DataFabric Manager server [276](#)
- ports and protocols
 - used by the DataFabric Manager server [276](#)
- predefined roles
 - list of [67](#)
- prerequisites
 - editing user quotas [175](#)
 - for managing user quotas [174](#)
 - to applying configuration files [230](#)
- primary roles
 - GlobalBackup [69](#)
 - GlobalDataProtection [69](#)
 - GlobalDataSet [69](#)
 - GlobalDelete role [69](#)
 - GlobalDistribution role [69](#)
 - GlobalEvent role [69](#)
 - GlobalFullControl role [69](#)
 - GlobalMirror role [69](#)
 - GlobalPerfManag [69](#)
 - GlobalQuota role [69](#)
 - GlobalRead role [69](#)
 - GlobalRestore role [69](#)
 - GlobalWrite role [69](#)
- protection issues
 - on vFiler units [266](#)
- protocols
 - used by the DataFabric Manager server [276](#)

Q

- qtrees
 - monitoring [204](#)
- quotas
 - editing, user [176](#)
 - how they work [173](#)

why you use [173](#)

R

RBAC

- access [73](#)
- for application administrators [74](#)
- how roles relate to administrators [69](#)
- what RBAC does [65](#)

RBAC resource [73](#)

remote configuration

- clusters [224](#)
- of storage systems [230](#)
- storage systems, methods of [223](#)

remote platform management

- RLM card management [241](#)
- system requirements [241](#)

Remote Platform Management

- configuring RLM card IP address [242](#)

report categories

- additional capabilities [151](#)

report generation

- scheduling [150](#)

report outputs

- list [156](#)

report schedule

- deleting [153](#)
- disabling [153](#)
- editing [152](#)
- listing results [154](#)
- retrieving [154](#)
- running [153](#)

Report Schedules report

- existing report schedules [151](#)

reports

- aggregate [144](#)
- aggregate array LUNs [144](#)
- aggregate catalog fields and performance counters [274](#)
- array LUNs [144](#)
- catalogs [144](#)
- clusters [144](#)
- configuring [149](#)
- controller [144](#)
- controllers [144](#)
- creating groups from [94](#)
- customizing options [143](#)
- dataset [144](#)
- deleting custom [150](#)
- disk catalog fields and performance counters [275](#)

disks [144](#)

events [144](#)

FC link [144](#)

FC switch [144](#)

FCP target [144](#)

file system [144](#)

filer catalog fields and performance counters [268](#)

group summary [144](#)

groups [97](#)

History performance reports [144](#)

Host domain users [144](#)

Host local users [144](#)

Host roles [144](#)

Host usergroups [144](#)

Host users [144](#)

interface groups [144](#)

logical interfaces [144](#)

LUN [144](#)

LUN catalog fields and performance counters [274](#)

of deleted storage objects [212](#)

performance events [144](#)

ports [144](#)

putting data in spreadsheets [150](#)

qtree catalog fields and performance counters [273](#)

quota [174](#)

Quotas [144](#)

Report outputs [144](#)

Report schedules [144](#)

resource pools [144](#)

SAN hosts [144](#)

schedules [144](#)

scripts [144](#)

spare array LUNs [144](#)

storage chargeback [207](#)

storage systems [144](#)

storage systems and vFiler units [216](#)

subgroups [98](#)

tasks performed from [216](#)

types of [143](#)

user quotas [144](#)

vFiler [144](#)

vFiler catalog fields and performance counters [270](#)

volume [144](#)

volume catalog fields and performance counters [271](#)

Vservers [144](#)

requirements

for disaster recovery [251](#)

for managing administrator access [76](#)

for managing SAN hosts [181](#)

for managing targets and initiators [181](#)

- for managing user quotas [174](#)
- resource groups
 - global [92](#)
- resource utilization [104](#)
- restoring
 - DataFabric Manager database [249](#)
 - the database from the Snapshot copy based backup [250](#)
 - the DataFabric Manager database [249](#)
- role-based access control
 - what it does [65](#)
- roles
 - adding to a host or user [88](#)
 - creating [72](#)
 - defined [87](#)
 - deleting from a host [89](#)
 - editing the settings of [89](#)
 - global [69](#)
 - global predefined [69](#)
 - group [71](#)
 - list of predefined [67](#)
 - modifying [72](#)
 - order applying [72](#)
 - pushing to hosts [89](#)
 - viewing settings for [88](#)
 - viewing, on hosts [87](#)
- run commands on storage systems
 - Run a Command tool [219](#)
- running commands
 - in a node of a cluster [225](#)
- running commands on storage systems [224](#)

S

- SAN components
 - access control [292](#)
 - configuring monitoring intervals [291](#)
 - deleting and undeleting [291](#)
 - deleting process [291](#)
 - grouping [292](#)
 - reasons, deleting, and undeleting [291](#)
 - undeleting process [292](#)
- SAN host
 - details [286](#)
- SAN hosts
 - administration transport [185](#)
 - discovery by DataFabric Manager server [278](#)
 - editing settings for [185](#)
 - host agents [185](#)
 - port for communication [185](#)

- requirements for managing [181](#)
- SAN management
 - FC switches [279](#)
 - Global and local options [289](#)
 - Global-only options [289](#)
 - GlobalSAN role [292](#)
 - license key [279](#)
 - list of tasks [287](#)
 - prerequisites [279](#)
 - SAN hosts, initiators [279](#)
 - tasks performed [280](#)
 - tasks performed, UI locations [281](#)
- SAN reports [282](#)
- saved reports
 - list of failed report outputs [156](#)
 - list of successful report outputs [156](#)
 - Viewing the output details of a particular report output [157](#)
 - viewing the output of a particular report [157](#)
- Saved reports
 - Viewing the output of a particular report output [157](#)
- schedule report
 - defined [154](#)
 - methods [152](#)
- schedules
 - adding [155](#)
 - deleting [155](#)
 - editing [155](#)
 - list [154](#)
- scheduling report [151](#), [152](#)
- scripts
 - overview [242](#)
- secure connections
 - components of [165](#)
 - enabling HTTPS for [165](#)
 - global options versus appliance options [168](#)
 - protocols used for [165](#)
 - requirements for, on workstation [166](#)
- Secure Shell [163](#)
- Secure Sockets Layer [165](#)
- self-signed certificates [163](#)
- server
 - accessing through the console [237](#)
- server authentication
 - through hosts.equiv [170](#)
- settings
 - editable options for systems [217](#)
- severity types
 - list of events and [110](#)
- Snapshot copies

- monitoring of [206](#)
- Snapshot copy
 - monitoring requirements [206](#)
- Snapshot-based backups
 - limitations [245](#)
- SNMP
 - setting v1 or v3 as the preferred version [53](#)
 - traps [106–108](#)
- SNMP connectivity
 - diagnosis [55](#)
- SNMP settings
 - deleting [55](#)
 - modifying [54](#)
- SNMPv3
 - monitoring process using [53](#)
- spreadsheets
 - LUN reports, putting into [150](#)
- SSH [165](#)
- SSL [165](#)
- stop monitoring a LUN [187](#)
- storage array ports
 - connected to the gateway, list of [64](#)
 - discovery [63](#)
- storage arrays
 - connected to the gateway, list of [64](#)
 - discovery [63](#)
- storage arrays of the gateway
 - configuration details [64](#)
- storage chargeback reports
 - defined [207](#)
- storage controller tools
 - use of [218](#)
- storage controllers
 - defined [98](#)
- storage objects
 - reports of deleted objects [212](#)
 - result of deleting [212](#)
- storage systems
 - accessing through FilerView [226](#)
 - changing passwords on multiple systems [169](#)
 - configuration management [230](#)
 - configuring from FilerView [226](#)
 - configuring multiple [236](#)
 - discovery of [57](#)
 - remote configuration [230](#)
 - remote configuration methods [223](#)
 - requirement for running remote commands [223](#)
- storage-related report [189](#)
- subgroup reports
 - what they are [98](#)

- summary page
 - viewing [77](#)
- Summary reports [97](#)
- Symbols
 - /etc/quotas file [175](#)
 - quotas file [173](#)

T

- takeover tool
 - what it does [221](#)
- targets, See FCP Targets [184](#)
- Temporarily stop monitoring [291](#)
- The Gateway SAN-attached storage management
 - limitations [63](#)
- thresholds
 - aggregate full [192](#)
 - aggregate full interval [192](#)
 - aggregate nearly full [192](#)
 - aggregate nearly overcommitted [192](#)
 - aggregate overcommitted [192](#)
 - changing group [97](#)
 - description of [176](#)
 - editing [189](#)
 - global and group levels [189](#)
 - overriding global settings [192](#)
 - precedence of [177](#)
 - user quota [176](#)
 - ways and locations to apply [177](#)
- traditional volumes
 - See volumes
- Trusted CA-signed certificates [163](#)

U

- undeleting (restoring)
 - objects [213](#)
 - SAN components [187](#)
- user access
 - cluster management [77](#)
- user accounts [66](#)
- user alert
 - e-mail message [138](#)
 - mailformat file [141](#)
- user alerts
 - alarms and alerts, differences [138](#)
 - contents of e-mail message [141](#)
 - Default Email Domain for Quota Alerts option [140](#)
 - defined [138](#)
- user capabilities [72](#)

- user groups
 - adding [85](#)
 - defined [85](#)
 - deleting [87](#)
 - deleting domain users from [84](#)
 - editing the settings [86](#)
 - pushing to hosts [87](#)
 - viewing the settings [86](#)
 - viewing, on a host [85](#)
- user quota
 - about editing [175](#)
 - mailformat file [141](#)
 - mailmap file [140](#)
- User Quota Monitoring Interval option [175](#)
- user quota thresholds
 - hard quota limit [177](#)
 - monitor [177](#)
 - settings [177](#)
 - soft quota limit [177](#)
- user quotas
 - alarms and alerts, differences [138](#)
 - alert See user alert [138](#)
 - editing [176](#), [227](#)
- user settings
 - configuring [176](#)
- users
 - adding, administrative [68](#)
 - viewing [77](#)
 - with Execute capability [80](#)
- using Operations Manager [103](#)
- utilization
 - logical resources [104](#)

V

- vFiler administrators
 - configuring [65](#)
- vFiler unit

- management tasks [228](#)
- vFiler units
 - access control, administrators [65](#)
 - adding to Resource groups [96](#)
 - configuration management [230](#)
 - discovery [57](#)
 - introduction to [227](#)
 - monitoring of [227](#)
 - monitoring volumes and qtrees in [204](#)
 - protection issues [266](#)
 - requirements for editing quotas [174](#)
 - resource details page [205](#)
- vFilers
 - discovery [57](#)
- view storage system details
 - Refresh Monitoring Samples tool [219](#)
- volume capacity thresholds
 - modify [201](#)
- volume full threshold [196](#)
- volume nearly full threshold [196](#)
- Volume Nearly No First Snapshot Threshold [201](#)
- Volume Snap Reserve Full Threshold [201](#)
- Volume Snapshot Count Threshold [201](#)
- Volume Snapshot No First Count Threshold [201](#)
- volume thresholds
 - Snapshot copy, events [201](#)
- Volume Too Old Snapshot Threshold [201](#)
- volumes
 - monitoring [204](#)
- Vserver [98](#)
- Vservers
 - defined [98](#)

W

- when to back up data [244](#)



NA 210-06255_A0, Printed in USA

GA32-1015-02

