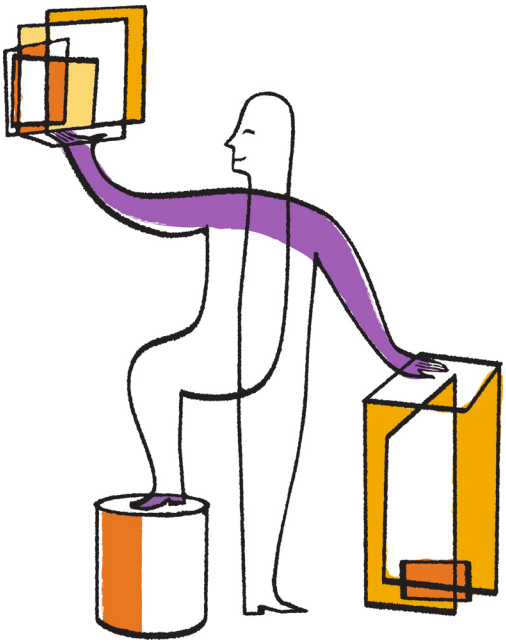
Release Candidate Documentation - Contents Subject To Change

OnCommand® System Manager 2.2

Help For 7-Mode

For Use with Data ONTAP®



NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.

Telephone: +1(408) 822-6000
Fax: +1(408) 822-4501
Support telephone: +1 (888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-07544_A0
April 2013

Contents

Welcome to OnCommand System Manager Help 10
System Manager 11
Understanding System Manager 11
Storage resource management 11
Storage system discovery 12
Credential caching 12
System logging 13
Window layout customization 13
Icons used in the application interface 14
Support for troubleshooting issues in System Manager 15
How log file rotation works 15
Creating a support bundle 16
Uploading a support bundle 16
Supportability Dashboard 16
What the network configuration checker is 17
Synchronization of active and persistent values 18
What the /etc/rc file format is 18
Configuring System Manager 19
Adding storage systems 19
Removing storage systems 20
Discovering storage systems 20
Saving your storage system credentials 21
Configuring system logging 21
Viewing System Manager application information 21
Configuring the SNMP timeout value 21
Configuring a default web browser 22
Verifying network configuration for storage systems 22
Editing the /etc/rc file 23
Window descriptions 23
Home tab 23
Dashboard window 25
Monitoring storage systems using the dashboard 26

Storage window 27
Data ONTAP storage architecture overview 27
Storage units for managing disk space 28
Where to find additional 7-Mode information 29
Configuring storage systems 31
Creating an NFS datastore for VMware 31
Viewing storage system details 32
Storage 33
Volumes 33
Understanding volumes 33
Configuring volumes 49
Managing volumes 55
Monitoring volumes 61
Window descriptions 62
Shares 66
Configuring shares 66
Managing shares 68
Window descriptions 68
Exports 69
Configuring exports 69
Managing exports 71
Window descriptions 72
LUNs 73
Understanding LUNs 73
Configuring LUNs 78
Managing LUNs 81
Monitoring LUNs 84
Window descriptions 85
Array LUNs 86
Understanding array LUNs 86
Configuring array LUNs 88
Managing array LUNs 89
Window descriptions 90
Quotas 91
Understanding quotas 91
Configuring quotas 98

Managing quotas 99
Monitoring quotas 100
Window descriptions 101
Qtrees 102
Understanding qtrees 102
Configuring qtrees 105
Managing qtrees 106
Monitoring qtrees 107
Window descriptions 107
Aggregates 108
Understanding aggregates 108
Configuring aggregates 127
Managing aggregates 130
Monitoring aggregates 136
Window descriptions 136
Disks 138
Understanding disks 138
Configuring disks 142
Managing disks 143
Monitoring disks 144
Window descriptions 144
vFiler Units 146
Understanding vFiler units 146
What vFiler units are 146
The default vFiler unit 146
What an IPspace is 146
Configuring vFiler units 147
Creating vFiler units 147
Deleting vFiler units 147
Managing vFiler units 148
Editing vFiler units 148
Starting or stopping vFiler units 149
Window descriptions 149
vFiler units window 149
Data Protection 151
SnapMirror 151

Understanding SnapMirror technology 151
Configuring SnapMirror relationships 153
Managing SnapMirror relationships 156
Window descriptions 161
SnapVault 163
Understanding SnapVault relationships 163
Configuring SnapVault relationships 165
Managing SnapVault relationships 167
Window descriptions 173
Configuration 177
Local Users and Groups > Users 177
Understanding local users 177
Configuring local users 178
Managing local users 179
Window descriptions 181
Local Users and Groups > Groups 182
Configuring local groups 182
Managing local groups 184
Window descriptions 184
Network > DNS 185
Understanding DNS 185
Configuring DNS 186
Managing DNS 187
Window descriptions 188
Network > Network Interfaces 188
Understanding interfaces 188
Configuring interfaces 197
Managing interfaces 199
Window descriptions 201
Network > Network Files 203
Understanding network file configuration 203
Configuring network files 204
Managing network files 205
Window descriptions 206
Network > NIS 207
Understanding NIS 207

Configuring NIS 209
Managing NIS 210
Window descriptions 211
Protocols > CIFS 211
Understanding CIFS 211
Configuring CIFS 213
Managing CIFS 215
Monitoring CIFS 221
Window descriptions 221
Protocols > NFS 222
Understanding NFS 222
Managing NFS 223
Window descriptions 224
Protocols > iSCSI 224
Understanding iSCSI 224
Configuring iSCSI 228
Managing iSCSI 230
Monitoring iSCSI 233
Window descriptions 233
Protocols > FC/FCoE 235
Understanding FC/FCoE 235
Configuring FC/FCoE 236
Managing FC/FCoE 236
Window descriptions 237
Security > Password/RSH 238
Understanding password/RSH 238
Configuring password/RSH 238
Window descriptions 240
Security > SSH/SSL 241
Understanding SSH and SSL 241
Managing SSH and SSL 245
Window descriptions 247
System Tools > AutoSupport 248
Understanding AutoSupport 248
Configuring AutoSupport 250
Managing AutoSupport 251

Window descriptions 252
System Tools > DateTime 253
Understanding date and time management 253
Configuring date and time settings 253
Window descriptions 254
System Tools > Licenses 255
Understanding licenses 255
Managing licenses 256
Window descriptions 258
System Tools > SNMP 259
Understanding SNMP 259
Configuring SNMP 260
Managing SNMP 260
Window descriptions 261
System Tools > NDMP 261
Understanding NDMP 261
Configuring NDMP service 261
Managing NDMP service 262
Window description 263
System Tools > Halt/Reboot 263
Halting storage systems 263
Rebooting storage systems 264
Window descriptions 264
Flash Cache 265
Understanding Flash Cache 265
Managing Flash Cache 266
Window descriptions 266
Diagnostics 268
CIFS 268
Understanding CIFS diagnostics 268
Monitoring CIFS diagnostics 268
Window descriptions 269
Session 270
Viewing sessions 270
Window descriptions 270
System Health 271

Understanding system health 271
Monitoring the health of your system 272
Window descriptions 274
Flash Pool Statistics 275
Window descriptions 275
Logs > Syslog 275
Understanding Syslog messages 275
Managing Syslog messages 277
Monitoring Syslog messages 278
Window descriptions 278
Logs > Audit Log 279
Understanding audit log 279
Managing audit log 280
Window descriptions 281
Logs > SnapMirror Log 281
Format of SnapMirror log files 281
Window description 283
HA Configuration 284
Understanding HA configuration 284
HA configuration 284
What an HA pair is 284
How the nodes in an HA pair provide redundancy 284
How HA pairs support nondisruptive operations and fault tolerance 285
What happens during takeover 286
Managing HA configuration 286
Enabling or disabling HA configuration 286
Initiating a takeover 287
Performing a giveback operation 287
Halting a storage system 288
Window descriptions 288
HA Configuration window 288
Copyright information 290
Trademark information 291
How to send your comments 292
Index 293

Welcome to OnCommand System Manager Help

The Help includes information about how to configure, manage, and monitor Data ONTAP 7.3.x (starting from 7.3.7) and Data ONTAP 8.0.x and 8.1.x operating in 7-Mode storage systems and storage objects by using OnCommand System Manager (abbreviated to System Manager). The table of contents, search, index, and favorites in the Help system help you find the relevant information required to achieve your goals.

The structure of the Help is similar to what you see on the UI. Help is also available from each window and its respective tabs. You can learn about a specific window parameter by clicking the Help icon ().

System Manager

Understanding System Manager

System Manager is a graphical management interface that enables you to manage storage systems and storage objects (such as disks, volumes, and aggregates) and perform common management tasks related to storage systems from a Web browser.

You can use System Manager to manage storage systems running the following versions of Data ONTAP:

- Data ONTAP 7.3.x (starting from 7.3.7)
- Data ONTAP 8.0.4 and 8.0.5 operating in 7-Mode
- Data ONTAP 8.1.2
- Data ONTAP 8.2

Note: In the Data ONTAP 8.x operating in 7-Mode product name, the term *7-Mode* signifies that the 8.x release operates in the same way as the prior Data ONTAP 7.1, 7.2, and 7.3 release families.

You can also use System Manager to manage V-Series systems.

System Manager enables you to perform many common tasks such as the following:

- Configure and manage storage objects, such as disks, aggregates, volumes, qtrees, and quotas.
- Configure protocols, such as CIFS and NFS, and provision file sharing.
- Configure protocols such as FC and iSCSI for block access.
- Verify and configure network configuration settings in the storage systems.
- Create and manage vFiler units.
- Set up and manage SnapMirror relationships and SnapVault relationships.
- Manage HA configurations and perform takeover and giveback operations.

Note: System Manager replaces FilerView as the tool to manage storage systems running Data ONTAP 8.1 or later.

Related tasks

Discovering storage systems on page 20

Storage resource management

You can use System Manager to manage the resources of your storage system.

Some of the important management tasks related to storage resource that you can perform in System Manager are as follows:

- Manage volumes and disks
- Increase data availability through Snapshot copies
- Back up and recover data
- Create aggregates, LUNs, and qtrees
- Manage shares, exports, and CIFS sessions
- Manage network interfaces
- Check the dashboard for performance of the storage objects and status
- Monitor system health

Access to your favorite topics

You can quickly access a particular subject that you often look up by bookmarking topics in the **Favorites** tab of the Help system.

Storage system discovery

The Discover Storage Systems dialog box lists all the storage systems discovered by System Manager. You can use this dialog box to discover storage systems or high-availability pair on a network subnet and add them to the list of managed systems.

The SNMP protocols SNMPv1 and SNMPv2c must be enabled, and an SNMP community must be configured to discover storage systems in System Manager. However, SNMPv1 and SNMPv2c protocols are less secure when compared to SNMPv3 protocol. In storage systems running Data ONTAP 8.1 or a later version in the Data ONTAP 8.1 family, SNMPv1 and SNMPv2c protocols are enabled by default.

When you add one of the systems in high-availability pair, the partner system is automatically added to the list of managed systems. You can type the IP address in any of the following formats: A.B.C.D, A.B.C, A.B.C.*, or A.B.C.D/24.

Credential caching

You can cache (save) your storage system server login and password information for future System Manager sessions.

By default, credential caching in System Manager is turned on. You have to supply your user name and password the first time you log in to a storage system. If both nodes of an HA pair have the same credentials, you have to supply the credentials only once.

After you enable the credential caching option, all storage system credentials are encrypted and saved to the user settings file. When you update storage system credential information, the user settings file is updated and saved. If System Manager shuts down unexpectedly, the saved credentials are available the next time you start System Manager. If you clear the credential caching option, all of the encrypted credentials are immediately erased from the user settings file.

System logging

System logging is an essential tool for application troubleshooting. It is helpful to enable system logging so that if there is a problem with an application, the problem can be located. You can enable System Manager logging at runtime without modifying the application binary.

Log output can be so voluminous that it quickly becomes overwhelming. System Manager enables you to refine the logging output by selecting which type of log statements are output. By default, system logging is set to INFO. You can choose one of the following log levels:

- OFF
- FATAL
- ERROR
- WARN
- INFO
- DEBUG
- TRACE

These levels function hierarchically. The log level set to OFF indicates no logging of messages. TRACE level logging includes all logs ranging from DEBUG to FATAL.

Window layout customization

System Manager enables you to customize the window layout. By customizing the windows, you can control which data is viewable or how it is displayed.

Sorting	You can click the column headings to sort the column entries in ascending order and display the sort arrows (and). You can then use the sort arrows to specify the order in which entries are displayed.
Filtering	You can use the filter icon ( ) to display only those entries that match the conditions provided. You can then use the character filter (?) or string filter (*) to narrow your search. You can apply filters to one or more columns. **Note:** If an entry in the column contains "?" or "*", to use the character filter or string filter, you must enclose "?" or "*" in square brackets.
Hiding or redisplaying the columns	You can click the column display icon ( ) to select the columns you want to display.
Customizing the layout	You can drag the bottom of the list of objects area up or down to resize the main areas of the window. You can also display or hide the list of related objects and list of views panels. You can drag the vertical dividers to resize the width of the columns or other areas of the window.

Icons used in the application interface

You can view the icons in the interface to get quick information about systems and operations, and to open other windows such as the application Help.

Icons that do not have labels in the interface are defined in the following tables.

Home tab icons

You might see the following icons in the **Home** tab.

Icon	Name	Description
	Individual system	The system type is an individual storage system.
	High availability pair	The system type is a high-availability pair.
	Unknown system	The system type is unknown or unavailable.

Dashboard window icons

You might see the following icons when viewing the dashboard for a selected storage system or HA pair.

Icon	Name	Description
	Help button	Opens a help window with information about that dashboard pane.
	Warning	There are minor issues, but none that require immediate attention.
	Error	Problems that might eventually result in downtime and therefore require attention.
	Critical	The storage system is not serving data or cannot be contacted. Immediate attention is required.
	Link arrow button	If this is displayed next to a line item in a dashboard pane, clicking it links to another page where you can get more information about the line item or make changes to the line item.

Support for troubleshooting issues in System Manager

If you encounter any issues when using the System Manager application, you can create a support bundle that includes your system configuration data and log files. You can send this bundle to technical support to help troubleshoot the issues.

The bundle contains the following data:

- System configuration details such as the version of the application, the name of the operating system hosting the application, and the browser used to launch the application
- The application configuration information, including the name, IP address, status, type, model, and ID of the storage systems that is being currently managed by the user who is logged in and is using System Manager
- Log files created by System Manager application

  These files record the errors that occur in the application during the course of managing the storage systems.

Note: Sensitive information such as storage system credentials is *not* collected as part of the bundle.

How log file rotation works

You must be aware of how log files are rotated in System Manager in order to save space, to easily maintain data, and to secure data.

You can use the log files generated by System Manager to debug issues because these files record the operations performed by System Manager. Because the number of log files increase with the operations performed, they are *rotated* to save system space and to ease file management.

The System Manager log files are rotated when its size crosses 10 MB. After the first rotation of the log files, the log files are renamed `System Manager.log.1` and subsequent logs are written to `SystemManager.log`. When the size of `SystemManager.log` exceeds 10 MB, the log files are renamed `SystemManager.log.1` and the older `SystemManager.log.1` files are renamed `SystemManager.log.2`. Similarly, the log files are rotated five times and at any time there is a maximum of six log files, namely, `SystemManager.log`, `SystemManager.log.1`, `SystemManager.log.2`, `SystemManager.log.3`, `SystemManager.log.4` and `SystemManager.log.5`. When all these log files are available in the system and when the size of `SystemManager.log` exceeds 10 MB, the rotation is repeated as described, and the existing `SystemManager.log.5` is removed, and `SystemManager.log.4` is renamed `SystemManager.log.5`.

The log files are stored as `SystemManager.log` in the following locations:

- `C:\users\<user>\NetApp\SystemManager` if System Manager is running on Windows Vista, Windows 7, or Windows Server 2008 operating system.
- `C:\Documents and Settings\<user>\NetApp\SystemManager` if System Manager is running on versions older than Windows Vista operating system.

- `<user_home>/NetApp/SystemManager` if System Manager is running on Linux operating system.

Creating a support bundle

You can create a support bundle from System Manager and send it to technical support to analyze and resolve issues with System Manager.

Steps

1. In the System Manager application window, click **Help > Support Bundle**.
2. Click **Click here to generate the file** to create the support bundle.

Result

The `SystemManagerSupportBundle.zip` file is downloaded to the specified location.

Uploading a support bundle

After you generate the support bundle, you must upload the support bundle to the NetApp Support Site to send this bundle to technical support to help troubleshoot the issues.

Before you begin

You must have generated a support bundle.

Steps

1. Open a support case to obtain a case number in one of the following ways:
   - Contact NetApp Support: +1 (888) 463-8277.
   - Log in to the NetApp Support Site.
2. Go to the NetApp File Upload Utility site and enter information when prompted.
3. Enter the case number obtained in Step 2.
4. Select the file type as **Non-Core** from the list.
5. Upload the support bundle.

Related information

NetApp File Upload Utility: support.netapp.com/upload

Supportability Dashboard

You can use the Supportability Dashboard to access product documentation and AutoSupport tools, download software, and visit sites such as the Community and NetApp University for additional information.

The Supportability Dashboard contains the following sources of information:

Community

Provides access to online collaborative resources on a range of NetApp products.

NetApp Support Site

Provides access to technical assistance and troubleshooting tools.

NetApp University

Provides course material to learn about NetApp products.

Downloads

Provides access to NetApp firmware and software that you can download.

Documentation

Provides access to NetApp product documentation.

My AutoSupport

Provides access to AutoSupport tools and processes.

What the network configuration checker is

The network configuration checker tool verifies the network configuration settings in the `/etc/rc` files with the active configuration settings in the storage systems managed by System Manager.

Note: You can use the network configuration checker tool to verify configuration mismatches only on storage systems running Data ONTAP 7.3.3 or later.

Before the configuration settings of a storage system are changed, System Manager compares the active and persistent configuration values, and the command sequence in the `/etc/rc` file.

System Manager creates a backup of the `/etc/rc` files and `/etc/hosts` files before making any changes, which enables you to restore the configuration settings from these backup files. The backed-up information is overwritten every time a networking operation is performed from System Manager. You can restore corrupt configuration settings in the `/etc/rc` file from the `rc.sysmgr.bak` file, and in the `/etc/hosts` file from the `hosts.sysmgr.bak` file.

The tool verifies the following network configuration settings:

- VLAN (VLAN tags)
- Interface group (type, policy, links)
- Interface configuration (IP address, aliases, netmask, MTU, Windows server, trusted interface, flow control)
- Route (destination, metric, next-hop)

Note: System Manager does not manage script-based `/etc/rc` files, because the files cause mismatches between the active and persistent settings. The following are examples of script-based `/etc/rc` files: `source -v/tc/myhostname` and `source -v/etc/myifconfigs`.

Synchronization of active and persistent values

You must synchronize the active and persistent values in the `/etc/rc` file to prevent the loss of configuration changes made to the storage system on reboot.

Running network commands from the command-line interface results in mismatches between the active and persistent values in the `/etc/rc` files.

Although System Manager detects such configuration mismatches, you cannot use System Manager to edit the `/etc/rc` file to update these values. You must use the command-line interface to eliminate the configuration mismatches by updating the persistent values in the `/etc/rc` file.

For example, consider an `/etc/rc` file that consists of the `hostname myhost`, `ifconfig e0a 1.2.3.4 netmask 255.255.255.0`, and `savecore` commands when the storage system is running. If the `ifconfig e0b 5.6.7.8 netmask 255.255.255.0` command is run only from the command-line interface without updating the `/etc/rc` file, then on storage system reboot, the e0b interface is not configured because the persistent value is not updated in the `/etc/rc` file.

What the /etc/rc file format is

The `/etc/rc` files consist of a sequence of commands that are executed in a particular order which determines the format of these files.

The sequence of commands in the `/etc/rc` files is as follows:

1. `hostname`
2. `vif or interface group`
3. `vlan`
4. `ifconfig`
5. `vfiler`
6. `route`
7. `routed`
8. `options`
9. `savecore`

If the order of these commands is modified in the `/etc/rc` files, System Manager displays a failure message.

Note: For storage system running Data ONTAP 8.0 or later, you must change the vif commands in the `/etc/rc` file to interface group commands because vif commands are deprecated from Data ONTAP 8.0 onwards. If you do not change the vif commands, System Manager reports an error in

the sequence of commands preventing changes to the configuration settings until the command sequence is corrected.

Configuring System Manager

Adding storage systems

Before you use System Manager to manage your storage systems, you have to add them to System Manager. You can also add storage systems that are in a high-availability (HA) configuration.

Before you begin

- Your storage systems must be running a supported version of Data ONTAP.
- SSL must be enabled on the storage system.

About this task

If you are adding one of the storage systems from an HA pair, the partner node is automatically added to the list of managed systems. If a high-availability partner node is down, you can add the working storage node.

Steps

1. From the **Home** tab, click **Add**.
2. Type the fully qualified DNS host name or the IPv4 address of the storage system.

   You can specify the IPv6 address of the storage system if you are adding a system that is running a supported version of Data ONTAP operating in 7-Mode.
3. Click the **More** arrow.
4. Select the method for discovering and adding the storage systems:
   - SNMP
     You must specify the SNMP community and SNMP version.
   - Credentials
     You must specify the user name and password.
5. Click **Add**.

Removing storage systems

You can remove one or more storage systems from the list of managed systems in System Manager. You have to select and remove only one of the storage systems in an HA configuration to remove both systems.

Step

1. From the **Home** tab, select one or more storage systems from the list of managed systems and click **Remove**.

Discovering storage systems

You can use the Discover Storage Systems dialog box to discover storage systems, clusters, or storage systems in an HA pair on a network subnet and add them to the list of managed systems.

Before you begin

- Your storage systems must be running a supported version of Data ONTAP.

About this task

If you are adding one of the storage controllers from an HA pair, the partner system is automatically added to the list of managed systems.

You can only discover storage systems and clusters running SNMPv1 and SNMPv2.

Steps

1. From the **Home** tab, click **Discover**.
2. In the **Discover Storage Systems** dialog box, type the subnet IP address and click **Discover**.
3. Select one or more storage systems from the list of discovered systems and click **Add Selected Systems**.
4. Verify that the storage system or the HA pair that you added is included in the list of managed systems in the System Manager application window.

Related concepts

Understanding System Manager on page 11

Saving your storage system credentials

You can save or cache your storage system user name and password for future System Manager sessions.

Steps

1. In the System Manager application window, click **Tools > Options**.
2. Select **Enable password caching** and click **Save and Close**.

Configuring system logging

You can enable logging for your system and select the level of detail recorded.

Steps

1. In the System Manager application window, click **Tools > Options**.
2. In the **Options** dialog box, select **TRACE** log level.
3. Click **Save and Close**.

Viewing System Manager application information

You can use the Help menu on the menu bar to view information about System Manager.

Steps

1. In the System Manager application window, click **Help > About NetApp OnCommand System Manager**.
2. Click **Configuration**.

Configuring the SNMP timeout value

You can configure the amount of time System Manager waits for a storage system to respond to an SNMP request. You can increase the SNMP timeout value if there is more latency in your network. By default, it is set to two seconds.

Steps

1. In the System Manager application window, click **Tools > Options**.
2. Set the SNMP timeout value, in seconds.
3. Click **Save and Close**.

Configuring a default web browser

You can configure a default web browser from which you can view and manage your storage systems using System Manager.

Steps

1. In the System Manager application window, click **Tools > Options**.
2. In the **Options** dialog box, set the default browser in **Browser Path**.
3. Click **Save and Close**.

Verifying network configuration for storage systems

You can use Network Configuration Checker to compare the network configuration settings in the storage system with the settings in the `/etc/rc` file and to identify any mismatches.

Before you begin

The user name and password for the storage system must be provided.

Network Configuration Checker does not verify storage systems for which the user name and password are not provided. If the authentication fails, the storage systems are highlighted.

About this task

Note: You must not run Network Configuration Checker on a node when the node is taken over by its partner.

You can only view the configuration settings for storage systems running Data ONTAP versions earlier than 7.3.3.

Steps

1. In the System Manager application window, click **Tools > Network Configuration Checker**.
2. In **Network Configuration Checker**, click **Check Mismatch** to verify any mismatches in the network configuration settings.

   The following statuses might be displayed:

   - Mismatch Found
   - No Mismatch
   - Error

     This status indicates that an error has occurred while attempting to read the `/etc/rc` file.
3. If any mismatches are found, click the status link displayed for more information.
4. Click **Close**.

Editing the /etc/rc file

You edit the storage system's boot configuration file, the `/etc/rc` file, to modify the commands that the system runs at boot time.

About this task

The storage system's boot configuration file is named `rc` and is in the `/etc` directory of its default volume (the default is `/vol/vol0/etc/rc`).

Steps

1. Make a backup copy of the `/etc/rc` file.
2. Edit the `/etc/rc` file.

   Note: Do not add CIFS commands to `/etc/rc`. Doing so can cause problems when the storage system boots if CIFS is not fully initialized or the commands cause deadlocks between the `/etc/rc` file and CIFS.

3. Ensure that entries in the `/etc/rc` file are listed in the following order:

```
hostname system_name
ifgrp commands
vlan commands
ifconfig commands
vfiler commands
route commands
[any other commands]
```

4. Save the edited file.
5. Reboot the storage system to test the new configuration.

   If the new configuration does not work as you want, repeat Step 2 through Step 4.

Window descriptions

Home tab

The **Home** tab enables you to view the storage systems that you are managing. You can discover and add storage systems from this tab.

- *Command buttons* on page 24
- *Systems list* on page 24

Command buttons

Login	Opens the management window for a selected storage system, which enables you to manage storage objects, vFiler units, and mirror relationships. You can also configure users, groups, network settings, protocols, system security, and system tools.
Discover	Opens the Discover Storage Systems dialog box, which enables you to discover storage systems with preferred SNMP options and add storage systems to the list of managed systems.
Add	Opens the Add a System dialog box, which enables you to add storage systems.
Remove	Removes one or more selected storage systems from the list of managed systems.
Refresh	Updates the information in the window.

Systems list

The systems list displays the list of managed storage systems and the address, status, type, operating system version, model, and ID of each system.

Storage system name	Specifies the storage system name.
Address	Specifies the IP address of the storage system.
Status	Specifies the current status of the storage system.
Type	Specifies the type of storage system as an HA pair or a stand-alone storage system.
Version	Specifies the version number of the operating system.
Model	Specifies the storage system model.
System ID	Specifies the ID of the storage system.

Dashboard window

The dashboard contains multiple panels that provide cumulative at-a-glance information about your system and its performance. You can use the Dashboard window to view information about space and CPU utilization, the status of storage objects, notifications, system properties, network throughput, and protocol operations.

System tab

The System tab includes the following panels:

Storage Capacity	Displays the storage capacity of the node such as the used space, available space in aggregates, spare disks, and unowned disks.
Notifications/ Reminders	Displays any notifications or reminders about issues in the storage system and pending configuration settings. Notifications or reminders about the HA status and configuration errors, disk failures, insufficient spare disks, license mismatches in the HA pair, SSL, and DNS are displayed.
Aggregates	Displays the total number of aggregates and the number of offline aggregates, if any.
Volumes	Displays a graphical view of the space utilization by the volumes.
Properties	Displays storage system attributes such as the model, system ID, Data ONTAP version, the duration for which the system has been running, and SnapLock ComplianceClock time.
Disks	Displays the number of disks available in the storage system along with the number of spare disks, failed disks, and unowned disks. A link is provided to the Disks window.

Performance tab

The Performance tab includes the following panels:

CPU Utilization	Displays a graphical view of the CPU utilization of the storage systems.
I/O Throughput	Displays a graphical view of the network throughput and disk throughput.
Protocol Ops	Displays the operations per second associated with the CIFS, NFSv3, FC/FCoE, and iSCSI protocols.
Protocol Latency	Displays the latency (in milliseconds) associated with the CIFS, NFSv3, FC/ FCoE, and iSCSI protocols.

Note: Some charts in the System Manager dashboard might be displayed with a dark grey background when viewed in Internet Explorer 9.0.

Monitoring storage systems using the dashboard

The dashboard enables you to monitor the health and performance of storage systems. You can also identify hardware problems and storage configuration issues by using the dashboard.

Before you begin

Adobe Flash Player 11.0 or later must be installed on your host system.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click the topmost entry, which is the name of the storage system.
3. View the details in the dashboard panels.

Storage window

If you have not configured your storage system, the Frequent Tasks window enables you to access the Storage Configuration wizard. If you have already configured your storage system, you can click the other links such as "Create Volume" to manage the resources of your storage system.

Frequent Tasks

Storage Configuration Wizard	Launches the Storage Configuration wizard, which enables you to configure your storage system or a high-availability configuration.
Create Aggregate	Launches the Create Aggregate wizard, which enables you to create aggregates.
Create Volume	Opens the Create Volume dialog box, which enables you to create volumes.
Create LUN	Launches the Create LUN wizard, which enables you to create LUNs.
Create Qtree	Open the Create Qtree dialog box, which enables you to create qtrees.
Create Export	Opens the Create Export dialog box, which enables you to create NFS exports.
Provision Storage for VMware	Starts the Create NFS Datastore for VMware wizard, which enables you to create an NFS datastore for VMware.
Create SnapMirror Relationship	Launches the SnapMirror Relationship Create wizard, which enables you to create a SnapMirror relationship from a source volume or a qtree.

Note: The Frequent Tasks window displays only the **Storage Configuration Wizard** link if you have not configured your storage system.

Data ONTAP storage architecture overview

Storage architecture refers to how Data ONTAP provides data storage resources to host or client systems and applications. Data ONTAP distinguishes between the physical layer of data storage resources and the logical layer.

- The physical layer includes disks, array LUNs, virtual disks, RAID groups, plexes, and aggregates.

  Note: A *disk* is the basic unit of storage for storage systems that use Data ONTAP to access native disk shelves. An *array LUN* is the basic unit of storage that a third-party storage array

provides to a storage system that runs Data ONTAP. A virtual disk is the basic unit of storage for a storage system that runs Data ONTAP-v.

- The logical layer includes the file systems—volumes, qtrees, logical unit numbers (LUNs)—and the directories and files that store data.

  Note: LUNs are storage target devices in iSCSI and FC networks.

Aggregates provide storage to volumes. Aggregates can be composed of either disks or array LUNs, but not both. Data ONTAP organizes the disks or array LUNs in an aggregate into one or more RAID groups. RAID groups are then collected into one or two plexes, depending on whether RAID-level mirroring (SyncMirror) is in use. Aggregates can have two formats: 32-bit and 64-bit. An aggregate's format affects its maximum size.

Volumes are data containers. Clients can access the data in volumes through the access protocols supported by Data ONTAP. These protocols include Network File System (NFS), Common Internet File System (CIFS), HyperText Transfer Protocol (HTTP), Web-based Distributed Authoring and Versioning (WebDAV), Fibre Channel (FC), and Internet SCSI (iSCSI).

You can partition volumes and control resource usage using qtrees. You can create LUNs for use in a SAN environment, using the FC or iSCSI access protocols. Volumes, qtrees, and LUNs contain directories and files.

Note: V-Series systems also support native disk shelves.

Storage units for managing disk space

To properly provision storage, it is important to define and distinguish between the different units of storage.

The following list defines the various storage units:

Plexes

A collection of one or more Redundant Array of Independent Disks (RAID) groups that together provide the storage for one or more Write Anywhere File Layout (WAFL) file system aggregates or traditional volumes.

Data ONTAP uses plexes as the unit of RAID-level mirroring when the SyncMirror software is enabled.

Aggregates

The physical layer of storage that consists of the disks within the RAID groups and the plexes that contain the RAID groups.

It is a collection of one or two plexes, depending on whether you want to take advantage of RAID-level mirroring. If the aggregate is unmirrored, it contains a single plex. Aggregates provide the underlying physical storage for traditional and FlexVol volumes.

Traditional or flexible volumes	A traditional volume is directly tied to the underlying aggregate and its properties. When you create a traditional volume, Data ONTAP creates the underlying aggregate based on the properties you assign with the `vol create` command, such as the disks assigned to the RAID group and RAID-level protection. A FlexVol volume is a volume that is loosely coupled to its containing aggregate. A FlexVol volume can share its containing aggregate with other FlexVol volumes. Thus, a single aggregate can be the shared source of all the storage used by all the FlexVol volumes contained by that aggregate. You can use either traditional or FlexVol volumes to organize and manage system and user data. A volume can hold qtrees and LUNs. After you set up the underlying aggregate, you can create, clone, or resize FlexVol volumes without regard to the underlying physical storage. You do not have to manipulate the aggregate frequently.
Qtrees	A qtree is a subdirectory of the root directory of a volume. You can use qtrees to subdivide a volume in order to group LUNs.
LUNs	A logical unit of storage that represents all or part of an underlying physical disk. You can create LUNs in the root of a volume (traditional or flexible) or in the root of a qtree. **Note:** You should not create LUNs in the root volume because it is used by Data ONTAP for system administration. The default root volume is /vol/vol0.

For detailed information about storage units, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Related information

Data ONTAP documentation on the NetApp Support Site-support.netapp.com

Where to find additional 7-Mode information

System Manager Help provides basic Data ONTAP operating in 7-Mode conceptual information to help you perform tasks using System Manager. For in-depth conceptual information to help you configure, monitor, and manage storage objects and storage systems, you can see the Data ONTAP documentation available on the NetApp Support Site.

You might find the following Data ONTAP documentation useful:

Data ONTAP Storage Management Guide for 7-Mode	Describes how to configure, operate, and manage the storage resources for storage systems running Data ONTAP operating in 7-Mode, using disks, RAID groups,

	aggregates, volumes, FlexClone volumes, files and LUNs, FlexCache volumes, deduplication, compression, qtrees, and quotas.
Data ONTAP System Administration Guide for 7-Mode	Describes general system administration for storage systems that run Data ONTAP software.
Data ONTAP High Availability and MetroCluster Configuration Guide for 7-Mode	Describes how to install and manage high-availability configurations.
Data ONTAP MultiStore Management Guide for 7-Mode	Describes how to administer vFiler units (virtual storage systems) with the MultiStore software available by license with Data ONTAP operating in 7-Mode.
Data ONTAP Network Management Guide for 7-Mode	Describes how to configure and manage networks associated with storage systems running Data ONTAP operating in 7-Mode.
Data ONTAP Storage Efficiency Management Guide for 7-Mode	Describes the features and functionalities that help to significantly improve storage utilization.
Data ONTAP SAN Administration Guide for 7-Mode	Describes how to configure and manage the iSCSI and FC protocols for SAN environments.
Data ONTAP File Access and Protocols Management Guide for 7-Mode	Describes how to manage file access on storage systems with Data ONTAP operating in 7-Mode for NFS, CIFS, HTTP, FTP, and WebDAV protocols.
Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode	Describes how to back up and recover data using Data ONTAP operating in 7-Mode online backup and recovery features.
Data ONTAP Archive and Compliance Management Guide for 7-Mode	Describes how to archive and protect data for compliance purposes.

Related information

Documentation: By Product Library: support.netapp.com/documentation/productsatoz/index.html

Configuring storage systems

You can use the Storage Configuration wizard to configure your storage system or an HA configuration. You must separately configure each storage system when you configure an HA configuration.

Before you begin

Your storage systems must be running one of the following versions of Data ONTAP:

- Data ONTAP 7.3.x (starting from 7.3.7)
- Data ONTAP 8.0 or later in the 8.0 release family operating in 7-Mode
- Data ONTAP 8.1 or later in the 8.1 release family operating in 7-Mode

Note: In the Data ONTAP 8.x operating in 7-Mode product name, the term *7-Mode* signifies that the 8.x release has the same features and functionality found in the prior Data ONTAP 7.1, 7.2, and 7.3 release families.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage**.
3. Click the **Storage Configuration** wizard.
4. Type or select information as prompted by the wizard.
5. Confirm the details and click **Finish** to complete the wizard.

Creating an NFS datastore for VMware

You can use the Create NFS Datastore for VMware wizard to create an NFS datastore for VMware. You can create a volume for the NFS datastore and specify the ESX servers that can access the NFS datastore.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage**.
3. Click **Provision Storage for VMware**.
4. Type or select information as prompted by the wizard.
5. Confirm the details and click **Finish** to complete the wizard.

Viewing storage system details

You can use the **Home** tab to view the details of a storage system, such as the name, IP address, status of a storage system, and version of Data ONTAP that the storage system is running.

Steps

1. From the **Home** tab, select the storage system that you want to view information about from the displayed list of managed systems.
2. Review the details.

Storage

Volumes

Understanding volumes

What volumes are

Volumes are data containers that enable you to partition and manage your data. Understanding the types of volumes and their associated capabilities enables you to design your storage architecture for maximum storage efficiency and ease of administration.

Volumes are the highest-level logical storage object. Unlike aggregates, which are composed of physical storage resources, volumes are completely logical objects.

System Manager supports two types of volumes, traditional and flexible. However, you can create only flexible volumes (FlexVol volumes) by using System Manager.

Understanding the root volume

The storage system's root volume contains special directories and configuration files that help you administer the storage system. Understanding the facts about the root volume helps you manage it.

The following facts apply to the root volume:

- How the root volume is installed and whether you need to create it yourself depend on the storage system:
  - For FAS systems and V-Series systems ordered with disk shelves, the root volume is a FlexVol volume that is installed at the factory.
  - For a V-Series system that does not have a disk shelf, you install the root volume on an array LUN.

    For more information about setting up a V-Series system, see the *Data ONTAP Software Setup Guide for 7-Mode*.
  - For systems running virtual storage, the Data ONTAP-v installation process creates a single aggregate by using all currently defined virtual disks and creates the root FlexVol volume in that aggregate.

    For more information about system setup, see the Installation and Administration Guide that came with your Data ONTAP-v system.
- The default name for the root volume is `/vol/vol0`.

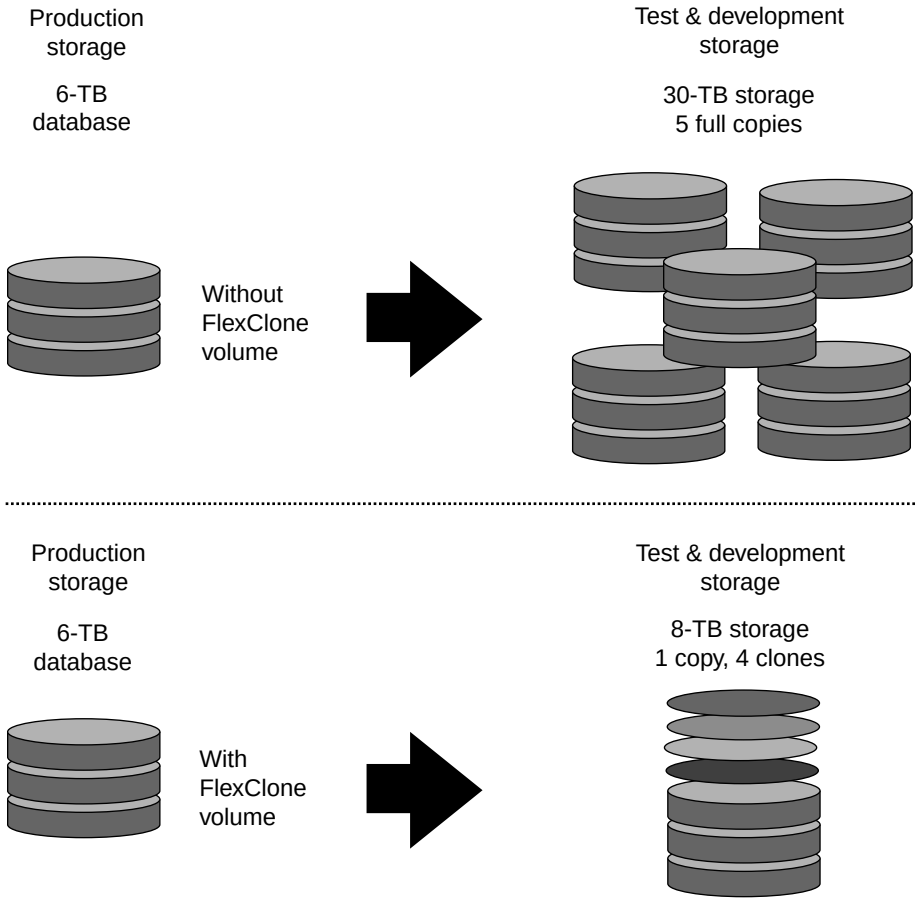
  You can designate a different volume to be the new root volume. Starting in Data ONTAP 8.0.1, you can designate a 64-bit volume to be the new root volume.
- The root volume's fractional reserve must be 100%.
- A VM-aligned volume is not supported as a root volume.

How FlexClone volumes save space

Understanding how FlexClone volumes save space enables you to maximize your storage efficiency. FlexClone volumes provide writeable volume copies that use only the space which is required to hold new data.

FlexClone volumes can be created instantaneously without interrupting access to the parent FlexVol volume. A FlexClone volume is initialized with a Snapshot copy and updated continually when data is written to the volume.

The following figure illustrates the space savings of test and development storage without and with FlexClone volumes.



For more information about FlexClone volumes, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Related information

Documentation on the NetApp Support Site: support.netapp.com

How FlexClone volumes work

FlexClone volumes can be managed similarly to regular FlexVol volumes, with a few important differences. For instance, the changes made to the parent FlexVol volume after the FlexClone volume is created are not reflected in the FlexClone volume.

The following list outlines important facts about FlexClone volumes:

- A FlexClone volume is a point-in-time, writable copy of the parent FlexVol volume.
- You must install the license for the FlexClone feature before you can create FlexClone volumes.
- A FlexClone volume is a fully functional FlexVol volume similar to its parent.
- A FlexClone volume is always created in the same aggregate as its parent.
- A traditional volume cannot be used as the parent of a FlexClone volume.
- Because a FlexClone volume and its parent share the same disk space for common data, creating a FlexClone volume is instantaneous and requires no additional disk space (until changes are made to the FlexClone volume or its parent).
- A FlexClone volume is created with the same volume guarantee as its parent.

  The volume guarantee setting is enforced for the new FlexClone volume only if there is enough space in the containing aggregate.
- A FlexClone volume is created with the same space reservation and fractional reserve settings as its parent.
- A FlexClone volume is created with the same Snapshot schedule as its parent.
- The common Snapshot copy shared between a FlexClone volume and its parent volume cannot be deleted while the FlexClone volume exists.
- While a FlexClone volume exists, some operations on its parent are not allowed, such as deleting the parent volume.
- You can sever the connection between the parent volume and the FlexClone volume.

  This is called *splitting* the FlexClone volume. Splitting removes all restrictions on the parent volume and causes the FlexClone to use its own additional disk space rather than sharing space with its parent.

  Attention: Splitting a FlexClone volume from its parent volume deletes all existing Snapshot copies of the FlexClone volume, and disables the creation of new Snapshot copies while the splitting operation is in progress.

  If you want to retain the Snapshot copies of the FlexClone volume, you can move the FlexClone volume to a different aggregate by using the `vol move` command. During the volume move operation, you can also create new Snapshot copies, if required. For more information about the volume move operation, see the *Data ONTAP SAN Administration Guide for 7-Mode*.

- Quotas applied to the parent volume are *not* automatically applied to the FlexClone volume.

- The clone of a SnapLock volume is also a SnapLock volume, and inherits the expiry date of the parent volume. This date cannot be changed, and the volume cannot be destroyed before the expiry date. For more information about SnapLock volumes, see the *Data ONTAP Archive and Compliance Management Guide for 7-Mode*.
- When a FlexClone volume is created, any LUNs present in the parent volume are present in the FlexClone volume but are unmapped and offline.

Note: For more detailed information about FlexClone volumes, refer to the *Data ONTAP Storage Management Guide for 7-Mode*.

Related information

support.netapp.com

How splitting a FlexClone volume from its parent works

Splitting a FlexClone volume from its parent removes any space optimizations that are currently used by the FlexClone volume. After the split, both the FlexClone volume and the parent volume require the full space allocation determined by their volume guarantees. The FlexClone volume becomes a normal FlexVol volume.

You must be aware of the following considerations related to clone-splitting operations:

- When you split a FlexClone volume from its parent, all existing Snapshot copies of the FlexClone volume are deleted.

  Note: If you want to retain the Snapshot copies of the FlexClone volume, you can move the FlexClone volume to a different aggregate by using the `vol move` command. During the volume move operation, you can also create new Snapshot copies, if required. For details of the `vol move` command, see the *Data ONTAP SAN Administration Guide for 7-Mode*.
- No new Snapshot copies can be created of the FlexClone volume for the duration of the split operation.
- Because the clone-splitting operation is a copy operation that might take considerable time to carry out, Data ONTAP provides the `vol clone split stop` and `vol clone split status` commands to stop or check the status of a clone-splitting operation.
- The clone-splitting operation proceeds in the background and does not interfere with data access to either the parent or the clone volume.
- The FlexClone volume must be online when you start the split operation.
- The parent volume must be online for the split operation to succeed.
- If you take the FlexClone volume offline while splitting is in progress, the operation is suspended; when you bring the FlexClone volume back online, the splitting operation resumes.
- If the FlexClone volume has a DP or LS mirror, it cannot be split from its parent volume.
- After a FlexClone volume and its parent volume have been split, they cannot be rejoined.

Considerations for creating a FlexClone volume from a SnapMirror source or destination volume

You can create a FlexClone volume from the source or destination volume in an existing volume SnapMirror relationship. However, doing so could prevent future SnapMirror replication operations from completing successfully.

Replication might not work because when you create the FlexClone volume, you might lock a Snapshot copy that is used by SnapMirror. If this happens, SnapMirror stops replicating to the destination volume until the FlexClone volume is destroyed or is split from its parent. You have two options for addressing this issue:

- If you require the FlexClone volume on a temporary basis and can accommodate a temporary stoppage of the SnapMirror replication, you can create the FlexClone volume and either delete it or split it from its parent when possible.

  The SnapMirror replication continues normally when the FlexClone volume is deleted or is split from its parent.
- If a temporary stoppage of the SnapMirror replication is not acceptable, you can create a Snapshot copy in the SnapMirror source volume, and then use that Snapshot copy to create the FlexClone volume. (If you are creating the FlexClone volume from the destination volume, you must wait until that Snapshot copy replicates to the SnapMirror destination volume.)

  This method of creating a Snapshot copy in the SnapMirror source volume allows you to create the clone without locking a Snapshot copy that is in use by SnapMirror.

FlexClone volumes and LUNs

You can clone FlexVol volumes that contain LUNs and LUN clones.

Note: LUNs in this context refer to the LUNs that Data ONTAP serves to clients, not to the array LUNs used for storage on a storage array.

When you create a FlexClone volume, LUNs in the parent volume are present in the FlexClone volume but they are not mapped and they are offline. To bring the LUNs in the FlexClone volume online, you need to map them to igroups. When the LUNs in the parent volume are backed by Snapshot copies, the FlexClone volume also inherits the Snapshot copies.

If the parent volume contains LUN clones (LUNs created by using the `lun clone` command), the FlexClone volume inherits the LUN clones and their base Snapshot copies. In this case, the LUN clone's base Snapshot copy in the parent volume shares blocks with the base Snapshot copy in the FlexClone volume. You cannot delete the LUN clone's base Snapshot copy in the parent volume while the base Snapshot copy in the FlexClone volume still exists.

If the parent volume contains FlexClone files or FlexClone LUNs (LUNs created by using the `clone start` command), the FlexClone volume also contains FlexClone files and FlexClone LUNs, which share storage with the FlexClone files and FlexClone LUNs in the parent volume.

How FlexVol volumes work

FlexVol volumes allow you to manage the logical layer of the file system independently of the physical layer of storage. Multiple FlexVol volumes can exist within a single separate, physically defined aggregate structure of disks and RAID groups. FlexVol volumes contained by the same aggregate share the physical storage resources, RAID configuration, and plex structure of that aggregate.

FlexVol volumes represent a significant administrative improvement over traditional volumes. Using multiple FlexVol volumes enables you to do the following:

- Perform administrative and maintenance tasks (for example, backup and restore) on individual FlexVol volumes rather than on a single, large file system.
- Set services (for example, Snapshot copy schedules) differently for individual FlexVol volumes.
- Minimize interruptions in data availability by taking individual FlexVol volumes offline to perform administrative tasks on them while the other FlexVol volumes remain online.
- Save time by backing up and restoring individual FlexVol volumes instead of all the file systems an aggregate contains.

Options for resizing volumes

You can use the Volume Resize wizard to change your volume size, adjust the Snapshot reserve, delete Snapshot copies, and dynamically see the results of your changes.

The Volume Resize wizard displays a bar graph that displays the current space allocations within the volume, including the amount of used and free space. When you make changes to the size or Snapshot reserve of the volume, this graph is updated dynamically to reflect the changes.

You can also use the **Calculate space** button to determine the amount of space that is freed by deleting selected Snapshot copies.

You can use the Volume Resize wizard to make the following changes to your volume:

Change the volume size	You can change the total volume size to increase or decrease storage space.
Adjust Snapshot reserve	You can adjust the amount of space reserved for Snapshot copies to increase or decrease storage space.
Delete Snapshot copies	You can delete Snapshot copies to reclaim volume space. **Note:** Snapshot copies that are being used or that have dependencies cannot be deleted.
Autogrow	You can specify the limit to which the volume can be grown automatically, if required.

What a Snapshot copy is

A Snapshot copy is a frozen, read-only image of a traditional volume, a flexible volume, or an aggregate that captures the state of the file system at a point in time. Snapshot copies are your first line of defense to backup and restore data.

When Snapshot copies are created, Data ONTAP maintains a configurable Snapshot copy schedule that creates and deletes Snapshot copies automatically for each volume. You can also create and delete Snapshot copies manually.

You can store up to 255 Snapshot copies at one time on each volume.

You can specify the percentage of disk space that Snapshot copies can occupy. The default space reserved for Snapshot copies is zero percent for SAN and VMware volumes. For NAS volumes, it is five percent on storage systems running Data ONTAP 8.1.

What the Snapshot copy reserve is

The Snapshot copy reserve sets a specific percent of the disk space for Snapshot copies. By default, the Snapshot copy reserve is 20 percent of the disk space. However, for a FlexVol volume, the Snapshot copy reserve is set to 5 percent by default. The active file system cannot consume the Snapshot copy reserve space, but the Snapshot copy reserve, if exhausted, can use space in the active file system.

Managing the Snapshot copy reserve involves the following tasks:

- Ensuring that enough disk space is allocated for Snapshot copies so that they do not consume active file system space
- Keeping disk space consumed by Snapshot copies below the Snapshot copy reserve
- Ensuring that the Snapshot copy reserve is not so large that it wastes space that could be used by the active file system

FlexClone volumes and shared Snapshot copies

When volume guarantees are in effect, a new FlexClone volume uses the Snapshot copy it shares with its parent to minimize its space requirements. If you delete the shared Snapshot copy, you might increase the space requirements of the FlexClone volume.

For example, suppose that you have a 100-MB FlexVol volume that has a volume guarantee of `volume`, with 70 MB used and 30 MB free, and you use that FlexVol volume as a parent volume for a new FlexClone volume. The new FlexClone volume has an initial volume guarantee of `volume`, but it does not require a full 100 MB of space from the aggregate, as it would if you had copied the volume. Instead, the aggregate needs to allocate only 30 MB (100 MB – 70 MB) of free space to the clone.

Now, suppose that you delete the shared Snapshot copy from the FlexClone volume. The FlexClone volume can no longer optimize its space requirements, and the full 100 MB is required from the containing aggregate.

Note: If you are prevented from deleting a Snapshot copy from a FlexClone volume due to "insufficient space in the aggregate" it is because deleting that Snapshot copy requires the allocation of more space than the aggregate currently has available. You can either increase the size of the aggregate, or change the volume guarantee of the FlexClone volume.

How volume guarantees work with FlexVol volumes

Volume guarantees (sometimes called *space guarantees*) determine how space for a volume is allocated from its containing aggregate--whether the space is preallocated for the entire volume or for only the reserved files or LUNs in the volume, or whether space for user data is not preallocated.

The guarantee is an attribute of the volume.

Volume guarantee types can be `volume` (the default type), `file`, or `none`

- A guarantee type of `volume` allocates space in the aggregate for the volume when you create the volume, regardless of whether that space is used for data yet.
  This approach to space management is called *thick provisioning*. The allocated space cannot be provided to or allocated for any other volume in that aggregate.
  When you use thick provisioning, all of the space specified for the volume is allocated from the aggregate at volume creation time. The volume cannot run out of space before the amount of data it contains (including Snapshot copies) reaches the size of the volume. However, if your volumes are not very full, this comes at the cost of reduced storage utilization.
- A guarantee type of `file` allocates space for the volume in its containing aggregate so that any reserved LUN or file in the volume can be completely rewritten, even if its blocks are being retained on disk by a Snapshot copy.
  However, writes to any file in the volume that is not reserved could run out of space.
  Before configuring your volumes with a guarantee of `file`, you should refer to Technical Report 3965.
- A guarantee of `none` allocates space from the aggregate only as it is needed by the volume.
  This approach to space management is called *thin provisioning*. The amount of space consumed by volumes with this guarantee type grows as data is added instead of being determined by the initial volume size, which might leave space unused if the volume data does not grow to that size.
  The maximum size of a volume with a guarantee of `none` is not limited by the amount of free space in its aggregate.
  Writes to LUNs or files (including space-reserved files) contained by that volume could fail if the containing aggregate does not have enough available space to accommodate the write. If you configure your volumes with a volume guarantee of `none`, you should refer to Technical Report 3965 for information about how doing so can affect storage availability.

When space in the aggregate is allocated for the guarantee for an existing volume, that space is no longer considered free in the aggregate. Operations that consume free space in the aggregate, such as creation of aggregate Snapshot copies or creation of new volumes in the containing aggregate, can occur only if there is enough available free space in that aggregate; these operations are prevented from using space already allocated to another volume.

When the free space in an aggregate is exhausted, only writes to volumes or files in that aggregate with preallocated space are guaranteed to succeed.

Guarantees are honored only for online volumes. If you take a volume offline, any allocated but unused space for that volume becomes available for other volumes in that aggregate. When you try to bring that volume back online, if there is insufficient available space in the aggregate to fulfill its guarantee, it will remain offline. You must force the volume online, at which point the volume's guarantee will be disabled.

What kind of space management to use for FlexVol volumes

The type of space management you should use for FlexVol volumes depends on many factors, including your tolerance for out-of-space errors, whether you plan to overcommit your aggregates, and your rate of data overwrite.

The following table can help you determine which space management capabilities best suit your requirements.

Note: *LUNs* in this context refer to the LUNs that Data ONTAP serves to clients, not to the array LUNs used for storage on a storage array.

If...	Then use...	Typical usage	Notes
You want management simplicity	FlexVol volumes with a guarantee of `volume`	NAS file systems	This is the easiest option to administer. As long as you have sufficient free space in the volume, writes to any file in this volume always succeed.

If...	Then use...	Typical usage	Notes
You need even more effective storage usage than file space reservation provides You use automatic space preservation and actively monitor available space on your aggregate and can take corrective action when needed You have space savings from deduplication and compression and want to use the free space made available Snapshot copies are short-lived Your rate of data overwrite is relatively predictable and low	FlexVol volumes with all of the following characteristics: • Guarantee of `volume` • Reservations enabled for LUNs and files that require writes to succeed • Fractional reserve < 100%	LUNs (with active space monitoring) Databases (with active space monitoring)	With fractional reserve <100%, it is possible to use up all available space, even with reservations on. Before enabling this option, be sure either that you can accept failed writes or that you have correctly calculated and anticipated storage and Snapshot copy usage.
You want to use thin provisioning You actively monitor available space on your aggregate and can take corrective action when needed You want to share free space at the aggregate level to increase overall storage utilization	FlexVol volumes with a guarantee of `none`	Storage providers who need to provide storage that they know will not be used immediately Storage providers who need to allow available space to be shared dynamically between volumes	With an overcommitted aggregate, writes can fail due to insufficient space.

FlexClone volumes and space guarantees

A FlexClone volume inherits its initial space guarantee from its parent volume. For example, if you create a FlexClone volume from a parent volume with a space guarantee of `volume`, then the

FlexClone volume's initial space guarantee will be `volume` also. You can change the FlexClone volume's space guarantee.

For example, suppose that you have a 100-MB FlexVol volume with a space guarantee of `volume`, with 70 MB used and 30 MB free, and you use that FlexVol volume as a parent volume for a new FlexClone volume. The new FlexClone volume has an initial space guarantee of `volume`, but it does not require a full 100 MB of space from the aggregate, as it would if you had copied the volume. Instead, the aggregate needs to allocate only 30 MB (100 MB minus 70 MB) of free space to the clone.

If you have multiple clones with the same parent volume and a space guarantee of `volume`, they all share the same shared parent space with each other, so the space savings are even greater.

Note: The shared space depends on the existence of the shared Snapshot copy (the base Snapshot copy that was used to create the FlexClone volume). If you delete this shared Snapshot copy, you lose the space savings provided by the FlexClone volume.

Thin provisioning for greater efficiencies using FlexVol volumes

With thin provisioning, when you create volumes and LUNs for different purposes in a given aggregate, you do not actually allocate any space for those volumes in advance. The space is allocated as data is written to the volumes.

The unused aggregate space is available to other thin provisioned volumes and LUNs. By allowing as-needed provisioning and space reclamation, thin provisioning can improve storage utilization and decrease storage costs.

A FlexVol volume can share its containing aggregate with other FlexVol volumes. Therefore, a single aggregate is the shared source of all the storage used by the FlexVol volumes it contains. Flexible volumes are no longer bound by the limitations of the disks on which they reside. A FlexVol volume can be sized based on how much data you want to store in it, rather than on the size of your disk. This flexibility enables you to maximize the performance and capacity utilization of the storage systems. Because FlexVol volumes can access all available physical storage in the system, improvements in storage utilization are possible.

Example

A 500-GB volume is allocated with only 100 GB of actual data; the remaining 400 GB allocated has no data stored in it. This unused capacity is assigned to a business application, even though the application might not need all 400 GB until later. The allocated but unused 400 GB of excess capacity is temporarily wasted.

With thin provisioning, the storage administrator provisions 500 GB to the business application but uses only 100 GB for the data. The difference is that with thin provisioning, the unused 400 GB is still available to other applications. This approach allows the application to grow transparently, and the physical storage is fully allocated only when the application needs it. The rest of the storage remains in the free pool to be used as needed.

How Data ONTAP can automatically provide more space for full FlexVol volumes

Data ONTAP uses two methods for automatically making more space available for a FlexVol volume when that volume is nearly full: allowing the volume size to increase, and deleting Snapshot copies.

Data ONTAP can automatically provide more free space for the volume by using one of the following methods:

- Increase the size of the volume when it is nearly full (`try_first` option set to `volume_grow`).
  This method is useful if the volume's containing aggregate has enough space to support a larger volume. You can configure Data ONTAP to increase the size in increments and set a maximum size for the volume.
- Delete Snapshot copies when the volume is nearly full (`try_first` option set to `snap_delete`).
  For example, you can configure Data ONTAP to automatically delete Snapshot copies that are not linked to Snapshot copies in cloned volumes or LUNs, or you can define which Snapshot copies you want Data ONTAP to delete first—your oldest or newest Snapshot copies. You can also determine when Data ONTAP should begin deleting Snapshot copies—for example, when the volume is nearly full or when the volume's Snapshot reserve is nearly full.
  For more information about deleting Snapshot copies automatically, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

How security styles affect access to your data

Every qtree and volume has a security style setting—NTFS, UNIX, or mixed. The setting determines whether files use Windows NT or UNIX (NFS) security. How you set up security styles depends on what protocols are licensed on your storage system.

Although security styles can be applied to volumes, they are not shown as a volume attribute, and they are managed for both volumes and qtrees using the `qtree` command. The security style for a volume applies only to files and directories in that volume that are not contained in any qtree. The volume security style does not affect the security style for any qtrees in that volume.

The following table describes the three security styles and the effects of changing them.

Security style	Description	Effect of changing to this style
NTFS	For CIFS clients, security is handled using Windows NTFS ACLs. For NFS clients, the NFS UID (user ID) is mapped to a Windows SID (security identifier) and its associated groups. Those mapped credentials are used to determine file access, based on the NTFS ACL. **Note:** To use NTFS security, the storage system must be licensed for CIFS. You cannot use an NFS client to change file or directory permissions on qtrees with the NTFS security style.	If the change is from a mixed qtree, Windows NT permissions determine file access for a file that had Windows NT permissions. Otherwise, UNIX-style (NFS) permission bits determine file access for files created before the change. **Note:** If the change is from a CIFS storage system to a multiprotocol storage system, and the `/etc` directory is a qtree, its security style changes to NTFS.
UNIX	Files and directories have UNIX permissions.	The storage system disregards any Windows NT permissions established previously and uses the UNIX permissions exclusively.
Mixed	Both NTFS and UNIX security are allowed: a file or directory can have either Windows NT permissions or UNIX permissions. The default security style of a file is the style most recently used to set permissions on that file.	If NTFS permissions on a file are changed, the storage system recomputes UNIX permissions on that file. If UNIX permissions or ownership on a file are changed, the storage system deletes any NTFS permissions on that file.

Note: When you create an NTFS qtree or change a qtree to NTFS, every Windows user is given full access to the qtree, by default. You must change the permissions if you want to restrict access to the qtree for some users. If you do not set NTFS file security on a file, UNIX permissions are enforced.

For more information about file access and permissions, see the *Data ONTAP File Access and Protocols Management Guide for 7-Mode*.

Using deduplication to increase storage efficiency

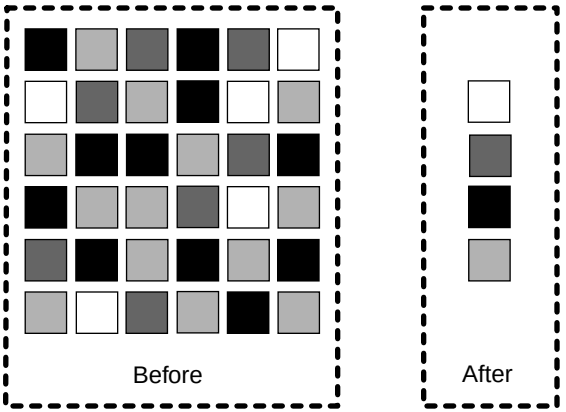
Deduplication is a Data ONTAP feature that reduces the amount of physical storage space required by eliminating duplicate data blocks within a FlexVol volume.

Deduplication works at the block level on an active file system, and uses the WAFL block-sharing mechanism. Each block of data has a digital signature that is compared with all other signatures in a

data volume. If an exact block match exists, a byte-by-byte comparison is done for all the bytes in the block, and the duplicate block is discarded and its disk space is reclaimed.

You can configure deduplication operations to run automatically or according to a schedule. You can deduplicate new and existing data, or only new data. You cannot enable deduplication on the root volume.

Deduplication removes data redundancies, as shown in the following illustration:



For more information about deduplication, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Related information

Documentation on the NetApp Support Site: support.netapp.com

Storage efficiency

Storage efficiency enables you to store the maximum amount of data for the lowest cost and accommodates rapid data growth while consuming less space. You can use technologies such as RAID-DP, FlexVol, Snapshot copies, deduplication, SnapMirror, and FlexClone to increase storage utilization and decrease storage costs. When used together, these technologies help to achieve increased performance.

- High-density disk drives, such as serial advanced technology attachment (SATA) drives mitigated with RAID-DP technology, provide increased efficiency and read performance.
- RAID-DP is a double-parity RAID6 implementation that protects against dual disk drive failures.
- Thin provisioning enables you to maintain a common unallocated storage space that is readily available to other applications as needed. It is based on the FlexVol technology.
- Snapshot copies are a point-in-time, read-only view of a data volume, which consumes minimal storage space. Two Snapshot copies created in sequence differ only by the blocks added or changed in the time interval between the two. This block incremental behavior limits the associated consumption of storage capacity.

- Deduplication saves storage space by eliminating redundant data blocks within a FlexVol volume.
- SnapMirror technology is a flexible solution for replicating data over local area, wide area, and Fibre Channel networks. It can serve as a critical component in implementing enterprise data protection strategies. You can replicate your data to one or more storage systems to minimize downtime costs in case of a production site failure. You can also use SnapMirror technology to centralize the backup of data to disks from multiple data centers.
- FlexClone technology copies data volumes, files, and LUNs as instant virtual copies. A FlexClone volume, file, or LUN is a writable point-in-time image of the FlexVol volume or another FlexClone volume, file, or LUN. This technology enables you to use space efficiently, storing only data that changes between the parent and the clone.
- The unified architecture integrates multiprotocol support to enable both file-based and block-based storage on a single platform. With V-Series systems, you can virtualize your entire storage infrastructure under one interface, and you can apply all the preceding efficiencies to your non-NetApp systems.

Guidelines for using deduplication

You must remember certain guidelines about system resources and free space when using deduplication.
The guidelines are as follows:

- If you have a performance sensitive solution, carefully consider the performance impact of deduplication and measure the impact in a test setup before deploying deduplication.
- Use the Data ONTAP version 7.3. Deduplication requires at a minimum ONTAP 7.2.5.1, but ONTAP 7.3 is recommended.
- Deduplication is a background process that consumes system resources while it is running. If the data does not change very often in a FlexVol volume, it is best to run deduplication less frequently. Multiple concurrent deduplication operations running on a storage system lead to a higher consumption of system resources.
- You must ensure that sufficient free space exists for deduplication metadata in the volumes and aggregates.
- For releases earlier than Data ONTAP 8.1, you cannot increase the size of a volume that contains deduplicated data beyond the maximum supported size limit, either manually or by using the `autogrow` option.
- For releases earlier than Data ONTAP 8.1, you cannot enable deduplication on a volume if it is larger than the maximum volume size. However, you can enable deduplication on a volume after reducing its size within the supported size limits.
- If deduplication is used on the source volume, use deduplication on the destination volume.
- Use automatic mode when possible so deduplication runs only when significant additional data has been written to each flexible volume.
- Run deduplication before creating a Snapshot copy to obtain maximum savings.
- Set the Snapshot reserve for greater than 0 if Snapshot copies are used.

Space savings with data compression

Data compression, an optional feature of Data ONTAP, enables you to reduce the physical capacity required to store data on storage systems by compressing data blocks within a FlexVol volume. You use data compression on primary, secondary, and archive storage tiers.

You can use data compression to store more data in less space, thereby reducing the time and bandwidth required to replicate data during volume SnapMirror transfers.

You can run data compression on regular files, virtual local disks, and LUNs. However, file system internal files, NT streams, and volume metadata are not compressed.

After you enable data compression in a FlexVol volume, all subsequent writes to the volume are compressed. However, existing data remains uncompressed. You can use the data compression scanner to compress the existing data.

Data compression is a licensed feature. You need to work with your NetApp sales team or NetApp partner sales team to request a NetApp Data compression license.

What SnapLock volumes are

SnapLock volumes are of two types—SnapLock Compliance volume and SnapLock Enterprise volume.

The SnapLock Compliance volume provides WORM protection for files and also restricts the storage administrator's ability to perform any operations that might modify or erase retained WORM records. SnapLock volumes use volume ComplianceClock to enforce the retention periods. Use SnapLock Compliance in strictly regulated environments that require information to be retained for a specified period of time, such as those governed by SEC Rule 17a-4.

The SnapLock Enterprise volume provides WORM protection for files with a trusted model of operation to manage the systems. SnapLock Enterprise allows the administrator to destroy SnapLock Enterprise volumes before all locked files on the volume reach their expiry date.

You cannot use a SnapLock volume as a regular volume for data storage. In most cases, SnapLock volumes behave identically to regular volumes, but there are some specific and critical differences as far as functionality and administration are concerned that make the SnapLock volume unsuitable for use as regular volumes. Specific examples include the following:

- Renaming directories on SnapLock volumes are not allowed.
- Transition of the file attribute from writable to read-only commits a file to the WORM state.
- Administrative interfaces are restricted (drastically for SnapLock Compliance volumes).

What retention period is

A retention period is the time period after which Data ONTAP permits the deletion of a write once, read many (WORM) file on a SnapLock volume. It is the duration for which a file is retained in WORM state.

Regulatory environments require that records be retained for a long period. Every record committed to the WORM state on a SnapLock volume can have an individual retention period associated with it. Data ONTAP enforces retention of these records until the retention period ends. After the retention period is over, the records can be deleted but not modified. Data ONTAP does not automatically delete any record. All records must be deleted using an application or manually.

The retention period is calculated by using the volume ComplianceClock. You can extend the retention period of an existing WORM file to infinite, however, you cannot shorten the retention period.

Types of SnapLock volumes

SnapLock volumes are of two types—SnapLock Compliance volumes and SnapLock Enterprise volumes.

The SnapLock Compliance volume provides WORM protection for files and restricts the storage administrator's capability to perform any operations that might modify or erase retained WORM records. SnapLock volumes use volume ComplianceClock to enforce the retention periods.

The SnapLock Enterprise volume provides WORM protection for files with a trusted model of operation to manage the systems. The administrator can destroy SnapLock Enterprise volumes before all the locked files on the volume reach their expiry date.

Configuring volumes

Creating FlexVol volumes

You can create a FlexVol volume for your data by using the Create Volume dialog box. You should always create a volume for your data rather than storing data in the root volume. You cannot create traditional volumes through System Manager.

Before you begin

The storage system must contain a non-root aggregate.

About this task

You cannot enable data compression on a volume if you are using Data ONTAP-v storage.

Steps

1. From the **Home** tab, double-click the appropriate storage system.

2. In the navigation pane, click **Storage > Volumes**.
3. Click **Create**.
4. If you want to change the default name, specify a new name.
5. Select the containing aggregate for the volume.

   If the containing aggregate is enabled for SnapLock compliance, the volume created has mandatory SnapLock protection.
6. Select the type of storage for which you are creating this volume.
7. Specify the size of the volume and the percentage of the total volume size that you want to reserve for Snapshot copies.
8. If you want to enable thin provisioning for the volume, select **Thin Provisioned**.

   When thin provisioning is enabled, space is allocated to the volume from the aggregate only when data is written to it.
9. If you want to enable deduplication, compression, or both on this volume, make the necessary changes in the **Storage Efficiency** tab.

   You cannot enable compression on a 32-bit volume.

   System Manager uses the default deduplication schedule. If the specified volume size exceeds the limit required for running deduplication, the volume is created and deduplication is not enabled.
10. Click **Create**.
11. Verify that the volume you created is included in the list of volumes in the **Volume** window.

    The volume is created with UNIX style security and UNIX 700 "read write execute" permissions for the Owner.

Related references

Volumes window on page 62

Creating FlexClone volumes

You can create a FlexClone volume when you need a writable, point-in-time copy of an existing flexible volume. You might want to create a copy of a FlexVol volume for testing or to provide access to the volume for additional users, without giving them access to the production data.

Before you begin

- The FlexClone license must be installed on the storage system.
- The volume that you want to clone must be online, and a non-root volume.

About this task

You can create a FlexClone volume from a SnapLock Enterprise volume, but not from a SnapLock Compliance volume.

The base Snapshot copy that is used to create a FlexClone volume of a SnapMirror destination is marked as busy and cannot be deleted. If a FlexClone volume is created from a Snapshot copy that is not the most recent Snapshot copy, and that Snapshot copy no longer exists on the source volume, all SnapMirror updates to the destination volume fail.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Volumes**.
3. Select the volume from the volume list.
4. Click **Clone** > **Create** > **Volume**.
5. In the **Create FlexClone Volume** dialog box, type the name of the FlexClone volume you want to create.
6. If you want to enable thin provisioning for the new FlexClone volume, select **Thin Provisioned**.

   By default, this setting is the same as that of the parent volume.
7. Create a new Snapshot copy or select an existing Snapshot copy that you want to use as the base Snapshot copy for creating the new FlexClone volume.
8. Click **Clone**.

Related references

Volumes window on page 62

Creating FlexClone files

You can create a FlexClone file, which is a writable copy of a parent file. You can use these copies to test applications.

Before you begin

- The file that is cloned must be part of the active file system.
- The FlexClone license must be installed on the storage system.

About this task

Note: You can create a FlexClone file of a parent file that is within a volume by accessing the parent file from the volume it resides in and not the parent volume.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Volumes**.
3. From the **Clone** menu, click **Create** > **File**.

4. Select the file that you want to clone and specify a name for the FlexClone file.
5. Click **Clone**.

Result

The FlexClone file is created in the same volume as the parent file.

Related references

Volumes window on page 62

Deleting volumes

You can delete a FlexVol volume when you no longer require the data it contains or if you have copied the data it contains to another location. When you delete a volume, all the data in the volume is destroyed and you cannot recover this data.

Before you begin

- If the FlexVol volume is cloned, the FlexClone volumes must be either split from the parent volume or be destroyed.
- The volume must be unmounted and in the offline state.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Volumes**
3. Select the volumes that you want to delete and click **Delete**.
4. Select the confirmation check box and click **Delete**.

Related references

Volumes window on page 62

Setting the Snapshot copy reserve

You can reserve space (in percentage) for Snapshot copies in a FlexVol volume. Setting the Snapshot copy reserve ensures that enough disk space is allocated for Snapshot copies so that they do not consume active file system space.

About this task

The default space reserved for Snapshot copies is zero percent for SAN and VMware volumes. For NAS volumes, the space reserved is 20 percent on storage systems running Data ONTAP versions earlier than 8.1 and five percent on storage systems running Data ONTAP 8.1.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Volumes**.
3. Select the volume for which you want to set the Snapshot copy reserve.
4. Click **Snapshot Copies** > **Configure**.
5. Type or select the percentage of the volume space that you want to reserve for Snapshot copies and click **OK**.

Related references

Volumes window on page 62

Creating Snapshot copies

You might want to create a Snapshot copy of a volume outside a specified schedule to capture the state of the file system at a specific point in time.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Volumes**.
3. Select the volume on which you want to create the Snapshot copy.
4. Click **Snapshot Copies** > **Create**.
5. In the **Create Snapshot Copy** dialog box, if you want to change the default name, specify a new name for the Snapshot copy.

   The default name of a Snapshot copy consists of the volume name and the timestamp.
6. Click **Create**.
7. Verify that the Snapshot copy you created is included in the list of Snapshot copies in the **Snapshot Copies** tab.

Related references

Volumes window on page 62

Deleting Snapshot copies

You can delete a Snapshot copy to conserve space or free disk space, or you can delete the Snapshot copy if it is no longer required. If you want to delete a Snapshot copy that is busy or locked, then you have to release the Snapshot copy from the application that is using it.

About this task

- You cannot delete the base Snapshot copy in a parent volume if a FlexClone volume is using that Snapshot copy.
  The base Snapshot copy is the Snapshot copy that is used to create the FlexClone volume, and always displays the status "busy" and Application Dependency as "busy,vclone" in the parent volume.
- You cannot delete locked a Snapshot copy that is used by a SnapMirror relationship.
  The Snapshot copy is locked and is required for the next update.

For more information about deleting busy Snapshot copies, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode* for your version of Data ONTAP.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Volumes**.
3. Select the appropriate volume that contains the Snapshot copy you want to delete.
4. Click **Snapshot Copies** in the lower pane of the **Volumes** window.
5. In the lower window pane, select the Snapshot copy that you want to delete.
6. Click **Delete**.
7. Select the confirmation check box and click **Delete**.

Related references

Volumes window on page 62

Managing volumes

Editing the volume properties

You can modify volume properties, such as the volume name, security style, fractional reserve, and space guarantee settings. You can also modify storage efficiency settings (deduplication schedule and compression) and space reclamation settings.

About this task

System Manager enables you to set the fractional reserve to either zero percent or 100 percent. Data compression is not supported on 32-bit volumes.

Data compression is not supported on Data ONTAP-v storage.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Volumes**.
3. Select the volume whose properties you want to edit and click **Edit**.
4. Click the appropriate tab to display the properties or settings you want to change.
5. Make the necessary changes.

   You cannot modify the name of a SnapLock Compliance volume.
6. Click **Save and Close** to save your changes and close the dialog box.

Related references

Volumes window on page 62

Changing the status of a volume

You can change the status of a FlexVol volume when you want to take the volume offline, bring it back online, or restrict access to the volume. However, you cannot take a root volume offline.

Before you begin

- If you want a volume to be the target of a volume copy or a SnapMirror replication operation, the volume must be in restricted state.

About this task

You can take a volume offline to perform maintenance on the volume, move it, or destroy it. When a volume is offline, it is unavailable for read or write access by clients.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage > Volumes**.
3. Select the volume for which you want to modify the status.
4. From the **Status** menu, click the volume status that you want.
5. In the confirmation dialog box, click the button for the volume status that you want.

Related references

Volumes window on page 62

Configuring deduplication on a volume

If you have not configured deduplication when you created the volume, you can do so later from the Edit dialog box. Deduplication saves storage space by eliminating redundant data blocks within the volume.

Before you begin

- Deduplication license must be enabled on the storage system.
- The volume must be online.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage > Volumes**.
3. Select the volume for which you want to configure deduplication.
4. Click **Edit**, and then click **Storage Efficiency**.
5. Select **Enable Storage Efficiency**.
6. Select one of the following schedules:
   - On-demand
   - Automated
     Deduplication is automatically run when 20 percent new data is written to the volume.
   - Scheduled
7. If you selected the **Scheduled** option, set the schedule by specifying the days on which you want deduplication to run and the number of times and frequency at which deduplication is run.
8. Click **Save and Close** to save your changes.

Related references

Volumes window on page 62

Changing the deduplication schedule

You can change the deduplication schedule by choosing to run deduplication manually, automatically, or on a schedule that you specify.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage > Volumes**.
3. Select the volume for which you want to modify the deduplication schedule.
4. Click **Edit**, and click **Storage Efficiency**.
5. Change the deduplication schedule as required.
6. Click **Save and Close** to save your changes.

Related references

Volumes window on page 62

Running deduplication operations

You can run deduplication immediately after creating a FlexVol volume or schedule deduplication to run at a specified time.

Before you begin

- The deduplication license must be enabled on the storage system.
- Deduplication must be enabled on the volume.
- The volume must be online and mounted.

About this task

Deduplication is a background process that consumes system resources during the operation; therefore, it might impact other operations that are in progress. You must cancel deduplication before you can perform any other operation.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage > Volumes**.
3. Select the volume for which you want to run deduplication.
4. Click **Storage Efficiency**.
5. In the **Storage Efficiency** dialog box, if you are running deduplication on the volume for the first time, run deduplication on the entire volume data by selecting **Scan Entire Volume**.

6. Click **Start**.
7. Check the status of the deduplication operation in the Storage Efficiency tab of the Volumes window.

Related references

Volumes window on page 62

Splitting a FlexClone volume from its parent volume

If you want the FlexClone volume to have its own disk space, rather than using that of its parent, you can split it from its parent. After the split, the FlexClone volume becomes a normal flexible volume.

Before you begin

The FlexClone volume must be online.

About this task

The clone-splitting operation deletes all the existing Snapshot copies of the clone. The Snapshot copies that are required for SnapMirror updates are also deleted. Therefore, any further SnapMirror updates might fail.

You can pause the clone-splitting operation, if you have to perform any other operation on the volume. You can resume the process after the operation is complete.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage > Volumes**.
3. Select the FlexClone volume that you want to split from its parent volume.
4. Click **Clone > Split**.
5. Confirm the clone split operation and click **Start Split** in the confirmation dialog box.

Related references

Volumes window on page 62

Resizing volumes

When your volume reaches nearly full capacity, you can increase the size of the volume, delete some Snapshot copies, or adjust the Snapshot reserve. You can use the Volume Resize wizard to provide more free space.

About this task

For a volume that is configured to grow automatically, you can modify the limit to which the volume can grow automatically, based on the increased size of the volume.

You cannot resize traditional volumes.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Volumes**.
3. Select the volume that you want to resize.
4. Click **Resize**.
5. Type or select information as prompted by the wizard.
6. Confirm the details and click **Finish** to complete the wizard.
7. Verify the changes you made to the available space and total space of the volume in the **Volumes** window.

Related references

Volumes window on page 62

Restoring a volume from a Snapshot copy

You can restore a volume to a state recorded in a previously created Snapshot copy to retrieve lost information. When you restore a Snapshot copy, the restore operation overwrites the existing volume configuration. Any changes made to the data in the volume after the Snapshot copy was made are lost.

Before you begin

- The SnapRestore license must be installed on your system.
- If the FlexVol volume you want to restore contains a LUN, the LUN must be unmounted or unmapped.
- There must be enough available space for the restored volume.
- Users accessing the volume must be notified that you are going to revert a volume, and that the data from the selected Snapshot copy replaces the current data in the volume.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Volumes**.
3. Select the volume that you want to restore from a Snapshot copy.
4. Click **Snapshot Copies** > **Restore**.
5. Select the appropriate Snapshot copy and click **Restore**.
6. Select the confirmation check box and click **Restore**.

Related references

Volumes window on page 62

Scheduling automatic Snapshot copies

You can set up a schedule for making automatic Snapshot copies of a FlexVol volume. You can specify the time and frequency of making the copies and specify the number of Snapshot copies that are saved.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Volumes**.
3. Click **Snapshot Copies** > **Configure**.
4. Select **Enable scheduled snapshots**.
5. Type or select the maximum number of Snapshot copies associated with each schedule.

   You can retain a maximum of 255 Snapshot copies.
6. Select one or more hours at which you want a Snapshot copy made.

   If you do not specify the time, Snapshot copies are created every hour.
7. Click **OK** to save your changes and start your Snapshot copy schedule.

Related references

Volumes window on page 62

Renaming Snapshot copies

You can rename a Snapshot copy to help you organize and manage your Snapshot copies.

Steps

1. From the **Home** tab, double-click the appropriate storage system.

2. In the navigation pane, click **Storage** > **Volumes**.
3. Select the appropriate volume that contains the Snapshot copy that you want to rename.
4. Click **Snapshot Copies** in the lower pane of the **Volumes** window.
5. In the lower window pane, select the Snapshot copy that you want to rename.
6. Click **Rename**.
7. Specify the new name and click **Rename**.
8. Verify the Snapshot copy name in the **Snapshot copies** tab of the **Volumes** window.

Related references

Volumes window on page 62

Hiding the Snapshot copy directory

You can hide the Snapshot copy directory (.snapshot) so that it is not visible when you view your volume directories. By default, the `.snapshot` directory is visible.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Volumes**.
3. Select the volume for which you want hide the Snapshot copy directory.
4. Click **Snapshot Copies** > **Configure**.
5. Ensure that **Make snapshot directory (.snapshot) visible** is not selected, and then click **Ok**.

Related references

Volumes window on page 62

Monitoring volumes

Viewing FlexClone volumes hierarchy

You can view the hierarchy of FlexClone volumes and their parent volumes by using the **View Hierarchy** option from the **Clone** menu.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Volumes**.
3. Select the volume from the volume list.

4. Click **Clone > View Hierarchy**.

   Volumes that have at least one child FlexClone volume are displayed. The FlexClone volumes are displayed as children of their respective parent volumes.

Related references

Volumes window on page 62

Viewing the Snapshot copies list

You can view a list of all the saved Snapshot copies for a selected volume from the **Snapshot Copies** tab in the lower pane of the Volumes window. You can also rename, restore, or delete the Snapshot Copy.

Before you begin

The volume must be online.

About this task

You can view Snapshot copies for only one volume at a time.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage > Volumes**.
3. In the upper pane of the **Volumes** window, select the volume for which you want to view the Snapshot copies.
4. In the lower pane, click **Snapshot Copies**.

   You can view the list of available Snapshot copies for the selected volume.

Window descriptions

Volumes window

You can use the Volumes window to manage your volumes and display information about them.

- *Command buttons* on page 62
- *Volume list* on page 63
- *Details area* on page 64

Command buttons

Create	Opens the Create Volume dialog box, which enables you to add a new volume.

Edit	Opens the Edit Volume dialog box, which enables you to modify a selected volume.
Delete	Deletes the selected volume or volumes.
Clone	Provides a list of clone options, including the following: • **Create** Creates a clone of the selected volume or a clone of a file from the selected volume. • **Split** Splits the clone from the parent volume. • **View Hierarchy** Display information about the clone hierarchy.
Status	Changes the status of the selected volume or volumes to one of the following statuses: • Online • Offline • Restrict
Snapshot Copies	Provides a list of Snapshot options, including the following: • **Create** Displays the Create Snapshot dialog box, which you can use to create a new Snapshot copy of the selected volume. • **Configure** Configures the Snapshot settings. • **Restore** Restores a Snapshot copy of the selected volume.
Resize	Opens the Volume Resize wizard, which enables you to change the volume size. This option is available only for FlexVol volumes.
Storage Efficiency	Opens the Storage Efficiency dialog box, which you can use to manually start deduplication or to abort a running deduplication operation. This button is displayed only if deduplication is enabled on the storage system.
Refresh	Updates the information in the window.

Volume list

The volume list displays the name of and storage information about each volume.

Name	Displays the name of the volume.
Aggregate	Displays the name of the aggregate.

Status	Displays the status of the volume.
Thin Provisioned	Displays whether space guarantee is set for the selected volume. Valid values for online volumes are "Yes" and "No."
Type	Displays the type of volume: traditional or flexible.
Root volume	Displays whether the volume is a root volume.
% Used	Displays the amount of space (in percentage) that is used in the volume.
Available Space	Displays the available space in the volume.
Total Space	Displays the total space in the volume. This includes space that is reserved for Snapshot copies.
Storage Efficiency	Displays whether deduplication is enabled or disabled for the selected volume.
Clone	Displays whether the volume is a FlexClone volume.
SnapLock Type	Displays whether the volume is a SnapLock Compliance volume or a SnapLock Enterprise volume.

Details area

The area below the volume list contains four tabs that display detailed information about the selected volume.

Details tab

Displays general information about the selected volume, such as the maximum and current file count on the volume.

Space Allocation tab

Displays the allocation of space in the volume.

- Bar graph
  Displays, in graphical format, details about the volume space.
- Volume
  Displays the total data space of the volume and the space reserved for Snapshot copies.
- Available
  Displays the amount of space that is available in the volume for data and for Snapshot copies, and the total space available in the volume.
- Used
  Displays the amount of space in the volume that is used for data and for Snapshot copies, and the total volume space that is used.

The **Space Allocation** tab displays different components, depending on whether the volume is configured for NAS or SAN. For a NAS volume, the space tab displays the following information:

- Used data space

- Available data space
- Used Snapshot reserve space
- Available Snapshot reserve space (this is applicable only if the snap reserve is greater than zero)

For a SAN volume, the space tab displays the following information:

- Space used by data in LUNs
- Available space
- Space used by Snapshot copies

Snapshot Copies tab

Displays, in tabular format, the Snapshot copies of the selected volume. This tab contains the following command buttons:

- **Create**
  Opens the Create Snapshot Copy dialog box, which enables you to create a new Snapshot copy of the selected volume.
- **Rename**
  Opens the Rename Snapshot Copy dialog box, which enables you to rename a selected Snapshot copy.
- **Delete**
  Deletes the selected Snapshot copy.
- **Restore**
  Restores the Snapshot copy.
- **Refresh**
  Updates the information in the window.

Storage Efficiency tab

The **Storage Efficiency** tab displays information in four panes.

- Bar graph
  Displays, in graphical format, the volume space used by data and Snapshot copies. You can view the space used details before and after applying storage efficiency savings.
- Details
  Displays information about deduplication properties, including whether deduplication is enabled on the volume, the deduplication status, and the current schedule. The space savings due to compression and deduplication applied on the data on the volume is also available.
- Last run details
  Provides details about the last-run deduplication operation on the volume.
- Graph legend
  Explains the symbols that are displayed on the graph.

Related tasks

Creating FlexVol volumes on page 49
Creating FlexClone volumes on page 50
Creating FlexClone files on page 51
Deleting volumes on page 52
Setting the Snapshot copy reserve on page 52
Deleting Snapshot copies on page 54
Creating Snapshot copies on page 53
Editing the volume properties on page 55
Changing the status of a volume on page 55
Configuring deduplication on a volume on page 56
Changing the deduplication schedule on page 57
Running deduplication operations on page 57
Splitting a FlexClone volume from its parent volume on page 58
Resizing volumes on page 59
Restoring a volume from a Snapshot copy on page 59
Scheduling automatic Snapshot copies on page 60
Renaming Snapshot copies on page 60
Hiding the Snapshot copy directory on page 61
Viewing FlexClone volumes hierarchy on page 61

Shares

Configuring shares

Creating a CIFS share

You can create a share that enables you to specify a folder, qtree, or volume that CIFS users can access.

Before you begin

You must have installed the CIFS license before you set up and start CIFS.

About this task

When you reconfigure CIFS on storage systems running Data ONTAP 8.x operating in 7-Mode from the CIFS Setup wizard, all existing user-created CIFS shares are deleted. However, the default CIFS shares that are created by Data ONTAP are not deleted, but their access permissions are reset to the default values. For more information, see the customer support bulletin CSB-1207-02.

When you reconfigure CIFS on storage systems running Data ONTAP 8.1.2 operating in 7-Mode, an error message is displayed stating that the CIFS shares are deleted. You can ignore this message.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage > Shares**.
3. Click **Create**.
4. Click **Browse** and select the folder, qtree, or volume that should be shared.
5. Specify a name for the new CIFS share.
6. Provide a description for the share and click **Create**.

Result

The share is created with the access permissions set to Full Control for Everyone in the group.

Related tasks

Setting up CIFS on page 213

Related references

Shares window on page 68

Related information

Customer support bulletin: support.netapp.com/info/communications/index.html

Stopping share access

Stopping share access stops the sharing of a folder, qtree, or volume. You can stop share access in the Shares window.

Before you begin

You must have the CIFS license.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage > Shares**.
3. From the list of shares, select the share that you want to stop sharing and click **Stop Sharing**.
4. Select the confirmation check box and click **Stop**.
5. Verify that the share is no longer listed in the **Shares** window.

Related references

Shares window on page 68

Managing shares

Editing share settings

You can modify the settings of a share, such as the number of users allowed for the share, the symbolic link settings, and the virus scan options. You can also modify share permissions by specifying the group or users who can access the share and the type of access to the share.

Before you begin

You must have the CIFS license.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Shares**..
3. Select the share you want to modify from the share list and click **Edit**.
4. Modify the share settings as required.
5. Click **Save and Close** to save your changes and close the dialog box.
6. Verify the changes you made to the selected share.

Related references

Shares window on page 68

Window descriptions

Shares window

You can use the Shares window to manage your shares and display information about them.

- *Command buttons* on page 68
- *Shares list* on page 69
- *Details area* on page 69

Command buttons

Create	Opens the Create Share dialog box, which enables you to create a share.
Edit	Opens the settings dialog box, which enables you to modify the properties of a selected share.

Stop Sharing	Stops the selected object from being shared.
Refresh	Updates the information in the window.

Shares list

The shares list displays the name and path of each share.

Share Name	Specifies the name of the share.
Path	Specifies the complete path name of an existing folder, qtree, or volume that is shared. Path separators can be backward or forward slashes, although Data ONTAP displays them as forward slashes.
Comment	Specifies the description for the share.

Details area

The area below the shares list displays the share properties and the access rights for each share.

Properties	Displays the share properties, such as the name of the share, the caching settings, and the virus scan, and volume states.
Share access control	Displays the access rights of the domain users and local users for the share.

Related tasks

Creating a CIFS share on page 66
Stopping share access on page 67
Editing share settings on page 68

Exports

Configuring exports

Creating NFS exports

You can create an NFS export to make file system paths on your storage system available for mounting by NFS clients. NFS clients can mount resources only after the resources have been exported and made available for mounting.

Before you begin

- The NFS license must be enabled on the storage system.
- You must have the following information:

- File system path to be exported
- Access privileges of the NFS clients (read-only, read-write, or root)
- Security types that an NFS client must support to access the file system path
- Anonymous access settings

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Exports**.
3. Click **Create**.
4. Click **Browse** and select the path to the volume, directory, or file to be exported.
5. In the **Export path** field, specify the path for accessing the exported path from a host.
6. Click **Add** in the Host Permissions section.
7. In the **Add Export Rule** dialog box, specify the required settings and click **Add**.
8. Click **Create**.

Related references

Exports window on page 72

Deleting NFS exports

You can delete one or more NFS exports in the **Exports** window and make file system paths on your storage system unavailable for mounting by NFS clients.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Exports**.
3. Select one or more exports that you want to delete from the exports list and click **Delete**.
4. Select the confirmation check box and click **Delete**.

Related references

Exports window on page 72

Managing exports

Adding export rules

The export rule specifies client permissions, security type, and anonymous access settings. You can use the **Add Export Rule** dialog box to add an export rule.

Before you begin

You must have the following information:

- Access privileges of NFS clients (read-only, read-write, or root)
- Security types that an NFS client must support to access the file system path
- Anonymous access settings

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Exports**.
3. Select the export for which you want to add an export rule.
4. Click **Add** in the **Client Permissions for Export** area.
5. In the **Add Export Rule** dialog box, specify the security type that an NFS client can use to access the file system path and the NFS clients and their access privileges.
6. Select the anonymous access settings.
7. Click **Add**.

Related references

Exports window on page 72

Editing NFS export rules

You can use the Edit Export Rule dialog box to change the security type, NFS clients and their access privileges, and anonymous access settings for an NFS export.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Exports**.
3. Select the NFS export that you want to edit.

   The client permissions details are displayed in the lower pane.
4. Select the security type in the **Client Permissions For Export** area, and click **Edit**.

5. In the **Edit Export Rule** dialog box, make the changes as required and click **Modify**.

Related references

Exports window on page 72

Window descriptions

Exports window

You can use the Exports window to manage NFS exports and display information about them.

Command buttons

Create	Opens the Create Export dialog box, which you can use to create an NFS export.
Delete	Deletes the selected exports.
Refresh	Updates the information in the window.

Exports list

The exports list displays the name and mount point of each NFS export.

Details area

The area below the exports list displays the security type that the NFS client must support to access the selected NFS export. You can add, edit, or delete an export rule for a selected NFS export. You can view the client permission details by either security type or NFS clients.

Related tasks

Creating NFS exports on page 69
Deleting NFS exports on page 70
Adding export rules on page 71
Editing NFS export rules on page 71

LUNs

Understanding LUNs

Guidelines for working with FlexVol volumes that contain LUNs

When you work with FlexVol volumes that contain LUNs, you must change the default settings for Snapshot copies. You can also optimize the LUN layout to simplify administration.

Snapshot copies are required for many optional features, such as the SnapMirror feature, SyncMirror feature, dump and restore, and ndmpcopy.

When you create a volume, Data ONTAP automatically performs the following:

- Reserves 5 percent of the space for Snapshot copies
- Schedules Snapshot copies

Because the internal scheduling mechanism for taking Snapshot copies within Data ONTAP does not ensure that the data within a LUN is in a consistent state, you should change these Snapshot copy settings by performing the following tasks:

- Turn off the automatic Snapshot copy schedule.
- Delete all existing Snapshot copies.
- Set the percentage of space reserved for Snapshot copies to zero.

You should use the following guidelines to create volumes that contain LUNs:

- Do not create any LUNs in the system's root volume.
  Data ONTAP uses this volume to administer the storage system. The default root volume is /vol/vol0.
- You should use a SAN volume to contain the LUN.
- You should ensure that no other files or directories exist in the volume that contains the LUN.
  If this is not possible and you are storing LUNs and files in the same volume, you should use a separate qtree to contain the LUNs.
- If multiple hosts share the same volume, you should create a qtree on the volume to store all the LUNs for the same host.
  This is a best practice that simplifies LUN administration and tracking.
- To simplify management, you should use naming conventions for LUNs and volumes that reflect their ownership or the way that they are used.

See the *Data ONTAP Storage Management Guide for 7-Mode* for more information.

Related information

Documentation: By Product Library: support.netapp.com/documentation/productsatoz/index.html

LUN size and type

When you create a LUN, you must specify the LUN size and the type for your host operating system.

The LUN Multiprotocol Type, or operating system type, determines the layout of data on the LUN, the geometry used to access that data, and the minimum and maximum sizes of the LUN. After the LUN is created, you cannot modify the LUN host operating system type.

Guidelines for using LUN multiprotocol type

The LUN multiprotocol type, or operating system type, specifies the operating system of the host accessing the LUN. It also determines the layout of data on the LUN, the geometry used to access that data, and the minimum and maximum size of the LUN.

The following table describes the LUN multiprotocol type values and the guidelines for using each type:

LUN multiprotocol type	When to use
AIX	If your host operating system is AIX.
HP-UX	If your host operating system is HP-UX.
Hyper-V	If you are using Windows Server 2008 Hyper-V and your LUNs contain virtual hard disks (VHDs). **Note:** For raw LUNs, you can use the type of child operating system as the LUN multiprotocol type.
Linux	If your host operating system is Linux.
NetWare	Your host operating system is NetWare.
OpenVMS	If your host operating system is OpenVMS.
Solaris	If your host operating system is Solaris and you are not using Solaris EFI labels.
Solaris EFI	If you are using Solaris EFI labels. **Note:** Using any other LUN multiprotocol type with Solaris EFI labels might result in LUN misalignment problems. For more information, see the Solaris Host Utilities documentation and release notes.
VMware	If you are using ESX Server and your LUNs will be configured with VMFS. **Note:** If you configure the LUNs with RDM, you can use the guest operating system as the LUN multiprotocol type.

LUN multiprotocol type	When to use
Windows	If your host operating system is Windows 2000 Server, Windows XP, or Windows Server 2003 using the MBR partitioning method.
Windows GPT	If you want to use the GPT partitioning method and your host is capable of using it. Windows Server 2003, Service Pack 1 and later are capable of using the GPT partitioning method, and all 64-bit versions of Windows support it.
Windows 2008	If your host operating system is Windows Server 2008; both MBR and GPT partitioning methods are supported.
Xen	If you are using Xen and your LUNs will be configured with Linux LVM with Dom0. **Note:** For raw LUNs, you can use the type of guest operating system as the LUN multiprotocol type.

For information about supported hosts, see the Interoperability Matrix.

Related information

NetApp Interoperability Matrix: support.netapp.com/NOW/products/interoperability

LUN clones

LUN clones are writable, space-efficient clones of parent LUNs. Creating LUN clones is highly space-efficient and time-efficient because the cloning operation does not involve physically copying any data. Clones aid in space storage utilization of the physical aggregate space.

You can clone a complete LUN without the need of a backing Snapshot copy in a SAN environment. The cloning operation is instantaneous and clients that are accessing the parent LUN do not experience any disruption or outage. Clients can perform all normal LUN operations on both parent entities and clone entities. Clients have immediate read-write access to both the parent and cloned LUN.

Clones share the data blocks of their parent LUNs and occupy negligible storage space until clients write new data either to the parent LUN, or to the clone. By default, the LUN clone inherits the space-reserved attribute of the parent LUN. For example, if the parent LUN is thinly provisioned, the LUN clone is also thinly provisioned.

Note: When you clone a LUN, you must ensure that the volume has enough space to contain the LUN clone.

Resizing a LUN

You can resize a LUN to be bigger or smaller than its original size. When you resize a LUN, you have to perform the steps on the host side that are recommended for the host type and the application that is using the LUN.

Initiator hosts

Initiator hosts can access the LUNs mapped to them. When you map a LUN on a storage system to the igroup, you grant all the initiators in that group access to that LUN. If a host is not a member of an igroup that is mapped to a LUN, that host does not have access to the LUN.

Guidelines for mapping LUNs to igroups

There are several important guidelines that you must follow when mapping LUNs to an igroup.

- You can map two different LUNs with the same LUN ID to two different igroups without having a conflict, provided that the igroups do not share any initiators or only one of the LUNs is online at a given time.
- You should ensure that the LUNs are online before mapping them to an igroup. You should not map LUNs that are in the offline state.
- You can map a LUN only once to an igroup.
- You can map a LUN only once to a specific initiator through the igroup.
- You can add a single initiator to multiple igroups, but the initiator can be mapped to a LUN only once. You cannot map a LUN to multiple igroups that contain the same initiator.
- You cannot use the same LUN ID for two LUNs mapped to the same igroup.

VMware RDM

When you perform raw device mapping (RDM) on VMware, the operating system type of the LUN must be the operating system type of the guest operating system.

What igroups are

Initiator groups (igroups) are tables of FC protocol host WWPNs or iSCSI host node names. You can define igroups and map them to LUNs to control which initiators have access to LUNs.

Typically, you want all of the host's HBAs or software initiators to have access to a LUN. If you are using multipathing software or have clustered hosts, each HBA or software initiator of each clustered host needs redundant paths to the same LUN.

You can create igroups that specify which initiators have access to the LUNs either before or after you create LUNs, but you must create igroups before you can map a LUN to an igroup.

Initiator groups can have multiple initiators, and multiple igroups can have the same initiator. However, you cannot map a LUN to multiple igroups that have the same initiator.

Note: An initiator cannot be a member of igroups of differing ostypes. Also, a given igroup can be used for FC protocol or iSCSI, but not both.

Required information for creating igroups

There are a number of attributes required when creating igroups, including the name of the igroup, type of igroup, ostype, iSCSI node name for iSCSI igroups, and WWPN for FCP igroups.

igroup name

The igroup name is a case-sensitive name that must satisfy several requirements.

The igroup name:

- Contains 1 to 96 characters. Spaces are not allowed.
- Can contain the letters A through Z, a through z, numbers 0 through 9, hyphen ("-"), underscore ("_"), colon (":"), and period (".").
- Must start with a letter or number.

The name you assign to an igroup is independent of the name of the host that is used by the host operating system, host files, or Domain Name Service (DNS). If you name an igroup aix1, for example, it is not mapped to the actual IP host name (DNS name) of the host.

Note: You might find it useful to provide meaningful names for igroups, ones that describe the hosts that can access the LUNs mapped to them.

igroup type

The igroup type can be mixed type, iSCSI, or FC/FCoE.

igroup ostype

The ostype indicates the type of host operating system used by all of the initiators in the igroup. All initiators in an igroup must be of the same ostype. The ostypes of initiators are `solaris`, `windows`, `hpux`, `aix`, `netware`, `xen`, `hyper_v`, `vmware`, and `linux`.

You must select an ostype for the igroup.

What ALUA is

Data ONTAP 7.2 added support for the asymmetric logical unit access (ALUA) features of SCSI, also known as SCSI Target Port Groups or Target Port Group Support.

ALUA is an industry standard protocol for identifying optimized paths between a storage system and a host. ALUA enables the initiator to query the target about path attributes, such as primary path and secondary path. It also allows the target to communicate events back to the initiator. It is beneficial because multipathing software can be developed to support any array; proprietary SCSI commands are no longer required.

Note: You cannot enable ALUA on iSCSI igroups.

Attention: You must ensure that your host supports ALUA before enabling it. Enabling ALUA for a host that does not support it can cause host failures during cluster failover.

Related information

Interoperability Matrix: support.netapp.com/NOW/products/interoperability

Configuring LUNs

Creating LUNs

You can create LUNs for an existing aggregate, volume, or qtree when there is available free space. You can create a LUN in an existing volume or create a new FlexVol volume for the LUN.

About this task

If you specify the LUN ID, System Manager checks the validity of the LUN ID before adding it. If you do not specify a LUN ID, Data ONTAP automatically assigns one.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **LUNs**.
3. In the **LUN Management** tab, click **Create**.
4. Type or select information as prompted by the wizard.
5. Confirm the details and click **Finish** to complete the wizard.

Related references

LUNs window on page 85

Deleting LUNs

You can delete LUNs and return the space used by the LUNs to their containing aggregates or volumes.

Before you begin

- The LUN must be offline.
- The LUN must be unmapped from all initiator hosts.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **LUNs**.

3. In the **LUN Management** tab, select one or more LUNs that you want to delete and click **Delete**.
4. Select the confirmation check box and click **Delete**.

Related references

LUNs window on page 85

Creating initiator groups

You can use the Create Initiator Group dialog box to create an initiator group. Initiator groups enable you to control host access to specific LUNs.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage > LUNs**.
3. Click **Initiator Groups**, and click **Create**.
4. In the **General** tab, specify the initiator group name, operating system, and supported protocol for the group.
5. Optional: Enable ALUA by selecting the check box.

   This check box is enabled if you select FC protocol for the initiator group.
6. In the **Initiators** tab, add the initiators.

   While adding initiators, ensure that the initiators and port sets are on the same subnet.
7. Click **Create**.

Related references

LUNs window on page 85

Deleting initiator groups

You can use the **Initiator Groups** tab to delete initiator groups.

Before you begin

All the LUNs mapped to the initiator group must be manually unmapped.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage > LUNs**.
3. Click **Initiator Groups**.

4. Select one or more initiator groups that you want to delete and click **Delete**.
5. Click **Delete**.
6. Verify that the initiator groups you deleted are no longer displayed in the Initiator Groups tab.

Related references

LUNs window on page 85

Adding initiators

You can use the Edit Initiator Group dialog box to add initiators to an initiator group. An initiator is provided access to a LUN when the initiator group that it belongs to is mapped to that LUN.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage > LUNs**.
3. Click **Initiators Groups**.
4. Select the initiator group to which you want to add initiators and click **Edit**.
5. In the **Edit Initiator Group** dialog box, click **Initiators**.
6. Click **Add**.
7. Specify the initiator name and click **OK**.
8. Click **Save and Close**.

Related references

LUNs window on page 85

Deleting initiators from an initiator group

You can use the **Initiator Groups** tab to delete an initiator.

Before you begin

All the LUNs mapped to the initiator group that contains the initiators must be manually unmapped.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage > LUNs**.
3. Click **Initiator Groups** and click **Edit**.
4. In the **Initiators** tab, select one or more initiators that you want to delete and click **Delete**.

Related references

LUNs window on page 85

Managing LUNs

Editing LUNs

You can use the LUN properties dialog box to change the name, description, size, space reservation setting, or the mapped initiator hosts of a LUN.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage > LUNs**.
3. Click **LUN Management**.
4. Select the LUN that you want to edit from the list of LUNs, and click **Edit**.
5. Make the changes as required.
6. Click **Save and Close** to save your changes and close the dialog box.

Related references

LUNs window on page 85

Editing initiator groups

You can use the Edit Initiator Group dialog box to change the name of an existing initiator group and its operating system. You can add initiators to or remove initiators from the initiator group. You can also enable or disable ALUA for an FC initiator group.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage > LUNs**.
3. Click **Initiator Groups**, and then click **Edit**.
4. Click the appropriate tab to display the properties or settings that you want to change.
5. Make the necessary changes.
6. Click **Save and Close** to save your changes and close the dialog box.
7. Use the **Initiator Groups** tab to verify the changes that you made to the selected initiator group.

Related references

LUNs window on page 85

Editing initiators

You can use the Edit Initiator Group dialog box to change the name, operating system type the of an existing initiator in an initiator group.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **LUNs**.
3. Click **Initiators Groups**.
4. Select the initiator group to which the initiator belongs and click **Edit**.
5. In the **Edit Initiator Group** dialog box, click **Initiators**.
6. Select the initiator that you want to edit and click **Edit**.
7. Change the name and click **OK**.
8. Click **Save and Close**.

Related references

LUNs window on page 85

Bringing LUNs online

You can use the **LUN Management** tab to bring selected LUNs online and make them available to the host.

Before you begin

Any host application accessing the LUN must be quiesced or synchronized.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **LUNs**.
3. In the **LUN Management** tab, select one or more LUNs that you want to bring online.
4. Click **Status** > **Online**.
5. Click **Online**.

Related references

LUNs window on page 85

Taking LUNs offline

You can use the **LUN Management** tab to take selected LUNs offline and make them unavailable for block protocol access.

Before you begin

Any host application accessing the LUN must be quiesced or synchronized.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **LUNs**.
3. In the **LUN Management** tab, select one or more LUNs that you want to take offline.
4. Click **Status** > **Offline**.
5. Click **Offline**.

Related references

LUNs window on page 85

Cloning LUNs

LUN clones enable you to create multiple readable and writable copies of a LUN. You might want to create a temporary copy of a LUN for testing or to make a copy of your data available to additional users, without providing them access to the production data.

Before you begin

- The FlexClone license must be installed on the storage system.
- When a LUN is thinly provisioned, the volume that contains the LUN must have enough space to accommodate changes to the clone.

Note: A thickly provisioned LUN clone requires as much space as the thickly provisioned parent LUN. If your storage system is running a version of Data ONTAP earlier than 8.1, the LUN clone operation might fail because of lack of space, and you might see the following error message: `Vdisk internal error`.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **LUNs**.
3. Select the LUN that you want to clone and click **Clone**.
4. If you want to change the default name, specify a new name.

5. Click **Clone**.
6. Verify that the LUN clone you created is listed in the **LUNs** window.

Related references

LUNs window on page 85

Monitoring LUNs

Viewing LUN information

You can use the **LUN Management** tab to view details about a LUN, such as its name, status, size, and type.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **LUNs**.
3. In the **LUN Management** tab, select the LUN that you want to view information about from the displayed list of LUNs.
4. Review the LUN details in the **LUNs** window.

Viewing initiator groups

You can use the **Initiator Groups** tab to view all the initiator groups and the initiators mapped to these initiator groups, and the LUNs and LUN ID mapping to the initiator groups.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **LUNs**.
3. Click **Initiator Groups** and review the initiator groups that are listed in the upper pane.
4. Select an initiator group to view the initiators that belong to it, which are listed in the **Initiators** tab in the lower pane.
5. Select an initiator group to view the LUNs mapped to it, which are listed in the **Mapped LUNs** in the lower pane.

Window descriptions

LUNs window

You can use the LUNs window to create and manage LUNs, and to display information about LUNs. You can also add, edit, or delete initiator groups and initiator IDs.

- *Tabs* on page 85
- *Command buttons* on page 85
- *LUNs list* on page 85
- *Details area* on page 86

Tabs

LUN Management	Enables you to create, clone, edit settings of, or delete LUNs.
Initiator Groups	Enables you to add, edit settings of, or delete initiator groups and initiator IDs.

Command buttons

Create	Opens the Create LUN wizard, which enables you to create LUNs.
Clone	Opens the Clone LUN dialog box, which enables you to clone the selected LUNs.
Edit	Opens the Edit LUN dialog box, which enables you to edit settings of the selected LUN.
Delete	Deletes the selected LUN.
Status	Specifies the status of the LUN.
Refresh	Updates the information in the window.

LUNs list

The LUNs list displays the name of and storage information about each LUN.

Name	Specifies the name of the LUN.
Container Path	Specifies the name of the file system (volume or qtree), which contains the LUN.
Thin Provisioned	Specifies whether thin provisioning is enabled.
Available Size	Specifies the space available in the LUN.
Total Size	Specifies total space in the LUN.
%Used	Specifies the total space (in percentage) that is used.
Type	Specifies the LUN type.

Status	Specifies the status of the LUN.

Details area

The area below the LUNs list displays LUN properties such as the LUN serial number and LUN description. You can view the initiator groups and initiator details associated with the selected LUN by clicking the corresponding tabs in the interface.

Related tasks

Creating LUNs on page 78
Deleting LUNs on page 78
Creating initiator groups on page 79
Deleting initiator groups on page 79
Adding initiators on page 80
Deleting initiators from an initiator group on page 80
Editing LUNs on page 81
Editing initiator groups on page 81
Editing initiators on page 82
Bringing LUNs online on page 82
Taking LUNs offline on page 83
Cloning LUNs on page 83

Array LUNs

Understanding array LUNs

About disks and array LUNs

Disks provide the basic unit of storage for storage systems that use Data ONTAP to access native disk shelves. Array LUNs are the basic unit of storage that a third-party storage array provides to a storage system that runs Data ONTAP.

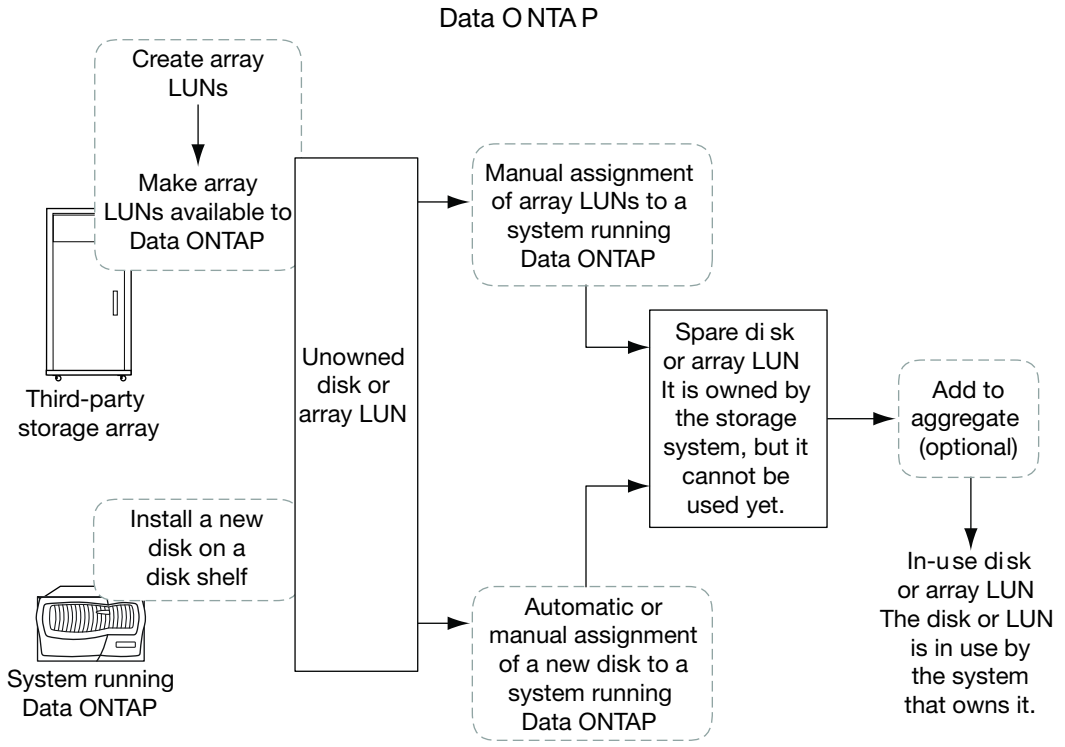
Data ONTAP enables you to assign ownership to your disks and array LUNs, and to add them to an aggregate. Data ONTAP also provides a number of ways to manage your disks, including removing them, replacing them, and sanitizing them. Because array LUNs are provided by the third-party storage array, you use the third-party storage array for all other management tasks for array LUNs.

You can create an aggregate using either disks or array LUNs. Once you have created the aggregate, you manage it using Data ONTAP in exactly the same way whether it was created from disks or array LUNs.

How disks and array LUNs become available for use

When you add a disk or array LUN to a system running Data ONTAP, the disk or array LUN goes through several stages before it can be used by Data ONTAP to store data or parity information.

The process for making a disk available for use differs slightly from the process for making an array LUN available for use. Both processes are shown in the following diagram.



The process for disks includes the following actions:

1. The administrator physically installs the disk into a disk shelf.
   Data ONTAP can see the disk but the disk is still unowned.
2. If the system is configured to support disk autoassignment, Data ONTAP assigns ownership for the disk; otherwise, the administrator must assign ownership of the disk manually.
   The disk is now a spare disk.
3. The administrator or Data ONTAP adds the disk to an aggregate.
   The disk is now in use by that aggregate. It could contain data or parity information.

The process for array LUNs includes the following actions:

1. The storage array administrator creates the array LUN and makes it available to Data ONTAP.
   Data ONTAP can see the array LUN but the array LUN is still unowned.

2. The Data ONTAP administrator assigns ownership for the array LUN to a V-Series system.
   The array LUN is now a spare array LUN.
3. The Data ONTAP administrator adds the array LUN to an aggregate.
   The array LUN is now in use by that aggregate and is used to contain data.

Rules for mixing array LUNs in an aggregate

Data ONTAP does not support mixing different types of storage in the same aggregate because it causes performance degradation.

There are restrictions on the types of array LUNs that you can mix in the same aggregate, which you must observe when you add array LUNs to an aggregate. Data ONTAP does not prevent you from mixing different types of array LUNs, but it does prevent you from mixing native disks and array LUNs in the same aggregate.

For array LUN aggregates, you cannot mix the following in the same aggregate:

- Array LUNs from storage arrays from different vendors
- Array LUNs from storage arrays from the same vendor but from different storage array families

  Note: Storage arrays in the same family share the same characteristics---for example, the same performance characteristics. See the *V-Series Implementation Guide for Third-Party Storage* for information about how Data ONTAP defines family members for the vendor.
- Array LUNs from FC and SATA drives
  You can deploy FC and SATA drives behind the same V-Series system. However, you cannot mix array LUNs from SATA disks and FC disks in the same aggregate, even if they are from the same series and the same vendor. Before setting up this type of configuration, consult your authorized reseller to plan the best implementation for your environment.

Configuring array LUNs

Creating an aggregate from spare array LUNs

You can use the Create Aggregate dialog box to create a new aggregate from selected spare array LUNs.

Before you begin

The ownership of an array LUN must be changed to spare, making the array LUN available for use.

About this task

Note: For aggregates for third-party storage, you cannot have array LUNs from storage arrays from different vendors in the same aggregate.

Steps

1. From the **Home** tab, double-click the appropriate storage system.

2. In the navigation pane, click **Storage > Array LUNs**.
3. Select one or more spare array LUNs and click **Create Aggregate**.
4. Specify a name for the aggregate, and then click **Create**.

Related references

Array LUNs window on page 90

Managing array LUNs

Adding array LUNs to an aggregate

You can use the Add Disks to Aggregate dialog box to add spare array LUNs to an existing aggregate.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage > Array LUNs**.
3. Select one or more spare array LUNs that you want to add to the aggregate and click **Add to Aggregate**.
4. Select the aggregate to which you want to add the spare array LUNs and click **Add**.

Related references

Array LUNs window on page 90

Assigning array LUNs

You can use the Make Spare dialog box to assign spare array LUNs to storage systems.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage > Array LUNs**.
3. Select the array LUNs that you want to assign to the storage system and click **Make Spare**.

Result

The array LUN is now assigned to the storage system.

Window descriptions

Array LUNs window

The Array LUNs window enables you to assign ownership to your array LUNs, and to add them to an aggregate.

The Array LUNs link in the left navigation pane appears only for V-Series systems.

- *Command buttons* on page 90
- *Array LUN list* on page 90
- *Details area* on page 90

Command buttons

Create Aggregate	Opens the Create Aggregate dialog box, which enables you to create a new aggregate using spare array LUNs. **Note:** This button is enabled only if there is at least one spare array LUN.
Add to Aggregate	Opens the Add Disks to Aggregate dialog box, which enables you to add spare array LUNs to an existing aggregate. **Note:** This button is enabled only if there is at least one spare array LUN.
Refresh	Updates the information in the window.

Array LUN list

The array LUN list displays information such as the name, state, and vendor for each array LUN.

Name	Specifies the name of the array LUN.
State	Specifies the state of the array LUN.
Model	Specifies the V-Series system model.
Vendor	Specifies the V-Series system vendor.
Used Space	Specifies the space used by an array LUN.
Size	Specifies the size of the array LUN.
Container	Specifies the aggregate or traditional volume to which this array LUN belongs.

Details area

The area below the array LUNs list displays detailed information about the selected array LUN.

Related tasks

Creating an aggregate from spare array LUNs on page 88
Adding array LUNs to an aggregate on page 89

Quotas

Understanding quotas

About quotas

Quotas provide a way to restrict or track the disk space and number of files used by a user, group, or qtree. You specify quotas using the `/etc/quotas` file. 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

Types of quotas

Quotas can be classified on the basis of targets they are applied to.

The following are the types of quotas based on the targets they are applied to:

User quota	The target is a user. The user can be represented by a UNIX user name/UID, a Windows SID, a file or directory whose UID matches the user Windows user name in pre-Windows 2000 format, or a file or directory with an ACL owned by the user's SID. You can apply it to a volume or a qtree.
Group quota	The target is a group. The group is represented by a UNIX group name, a GID, or a file or directory whose GID matches the group. Data ONTAP does not apply group quotas based on a Windows ID. You can apply it to a volume or a qtree.

Qtree quota	The target is a qtree, specified by the path name to the qtree. You can determine the size of the target qtree.
Default quota	Automatically applies a quota limit to a large set of quota targets without creating separate quotas for each target. Default quotas can be applied to all three types of quota target (users, groups, and qtrees). A quota target with an asterisk mark (*) denotes a default quota. The quota type is determined by the value of the type field.

How you specify UNIX users for quotas

You can specify a UNIX user for a quota using one of three formats: the user name, the UID, or a file or directory owned by the user.

To specify a UNIX user for a quota, you can use one of the following formats:

- The user name, as defined in the `/etc/passwd` file or the NIS password map, such as jsmith.

  Note: You cannot use a UNIX user name to specify a quota if that name includes a backslash (\) or an @ sign. This is because Data ONTAP treats names containing these characters as Windows names.

- The UID, such as 20.
- The path of a file or directory owned by that user, so that the file's UID matches the user.

  Note:

  If you specify a file or directory name, you should choose a file or directory that will last as long as the user account remains on the system.

  Specifying a file or directory name for the UID does not cause Data ONTAP to apply a quota to that file or directory.

How you specify Windows users for quotas

You can specify a Windows user for a quota using one of three formats: the Windows name in pre-Windows 2000 format, the SID, or a file or directory owned by the SID of the user.

To specify a Windows user for a quota, you can use one of the following formats:

- The Windows name in pre-Windows 2000 format.
- The security ID (SID), as displayed by Windows in text form, such as S-1-5-32-544.
- The name of a file or directory that has an ACL owned by that user's SID.

  Note:

  If you specify a file or directory name, you should choose a file or directory that will last as long as the user account remains on the system.

  For Data ONTAP to obtain the SID from the ACL, the ACL must be valid.

If the file or directory exists in a UNIX-style qtree, or if the storage system uses UNIX mode for user authentication, Data ONTAP applies the user quota to the user whose *UID*, not SID, matches that of the file or directory.

Specifying a file or directory name to identify a user for a quota does not cause Data ONTAP to apply a quota to that file or directory.

How you specify a user name in pre-Windows 2000 format

The pre-Windows 2000 format, for example `engineering\john_smith`, is used by the quotas file for specifying Windows users.

Keep in mind the following rules when creating pre-Windows 2000 format user names:

- The user name must not exceed 20 characters.
- The NetBIOS form of the domain name must be used.

How you specify a Windows domain using the QUOTA_TARGET_DOMAIN directive

Using the QUOTA_TARGET_DOMAIN directive in the quotas file enables you to specify the domain name only once for a group of Windows users.

The QUOTA_TARGET_DOMAIN directive takes an optional argument. This string, followed by a backslash (\), is prefixed to the name specified in the quota entry. Data ONTAP stops adding the domain name when it reaches the end of the quotas file or another QUOTA_TARGET_DOMAIN directive.

Example

The following example illustrates the use of the QUOTA_TARGET_DOMAIN directive:

```
QUOTA_TARGET_DOMAIN corp
roberts    user@/vol/vol2         900M      30K
smith      user@/vol/vol2         900M      30K
QUOTA_TARGET_DOMAIN engineering
daly       user@/vol/vol2         900M      30K
thomas     user@/vol/vol2         900M      30K
QUOTA_TARGET_DOMAIN
stevens    user@/vol/vol2         900M      30K
```

The string corp\ is added as a prefix to the user names of the first two entries. The string engineering\ is added as a prefix to the user names of the third and fourth entries. The last entry is unaffected by the QUOTA_TARGET_DOMAIN entry because the entry contains no argument.

The following entries produce the same effects:

```
corp\roberts       user@/vol/vol2         900M      30K
corp\smith         user@/vol/vol2         900M      30K
engineering\daly   user@/vol/vol2         900M      30K
```

```
engineering\thomas user@/vol/vol2          900M      30K
stevens            user@/vol/vol2          900M      30K
```

Quota limits

You can apply a disk space limit or limit the number of files for each quota type. If you do not specify a limit for a quota, none is applied. The maximum quota limit is 16383 GB or 16 TB - 1 on systems running versions earlier than Data ONTAP 8.0. On systems running Data ONTAP 8.0 7-Mode, the maximum quota limit is 1073741823 GB.

Disk space soft limit	Disk space limit applied to soft quotas.
Disk space hard limit	Disk space limit applied to hard quotas.
Threshold limit	Disk space limit applied to threshold quotas.
Files soft limit	The maximum number of files on a soft quota.
Files hard limit	The maximum number of files on a hard quota.

Quota management

System Manager includes several features that help you to create, edit, or delete quotas. You can create a user, group, or tree quota and you can specify both disk and file level quota limits. All quotas are established on a per-volume basis.

After creating a quota, you can perform the following tasks:

- Enable and disable quotas
- Resize quotas

How default quotas work

You can use default quotas to apply a quota to all instances of a given quota type. For example, a default user quota affects all users on the system for the specified volume. In addition, default quotas enable you to modify your quotas easily.

You can use default quotas to automatically apply a limit to a large set of quota targets without having to create separate quotas for each target. For example, if you want to limit most users to 10 GB of disk space, you can specify a default user quota of 10 GB of disk space instead of creating a quota for each user. If you have specific users for whom you want to apply a different limit, you can create explicit quotas for those users. (Explicit quotas—quotas with a specific target or list of targets —override default quotas.)

In addition, default quotas enable you to use resizing rather than reinitialization when you want quota changes to take effect. For example, if you add an explicit user quota to a volume that already has a default user quota, you can activate the new quota by resizing.

Default quotas can be applied to all three types of quota target (users, groups, and qtrees).

Default quotas do not necessarily have specified limits; a default quota can be a tracking quota.

Default user quota example

The following quotas file uses a default user quota to apply a 50-MB limit on each user for vol1:

```
#Quota target type           disk  files  thold  sdisk  sfile
#-----------  ----           ----  -----  -----  -----  -----
*             user@/vol/vol1  50M
```

If any user on the system enters a command that would cause that user's data to take up more than 50 MB in vol1 (for example, writing to a file from an editor), the command fails.

How quotas work with qtrees

You can create quotas with a qtree as their target; these quotas are called *tree quotas*. You can also create user and group quotas for a specific qtree. In addition, quotas for a volume are sometimes inherited by the qtrees contained by that volume.

How tree quotas work

You can create a quota with a qtree as its target to limit how large the target qtree can become. These quotas are also called *tree quotas*.

When you apply a quota to a qtree, the result is similar to a disk partition, except that you can change the qtree's maximum size at any time by changing the quota. When applying a tree quota, Data ONTAP limits the disk space and number of files in the qtree, regardless of their owners. No users, including root and members of the BUILTIN\Administrators group, can write to the qtree if the write operation causes the tree quota to be exceeded.

Note: The size of the quota does not guarantee any specific amount of available space. The size of the quota can be larger than the amount of free space available to the qtree.

How user and group quotas work with qtrees

Tree quotas limit the overall size of the qtree. To prevent individual users or groups from consuming the entire qtree, you specify a user or group quota for that qtree.

Example user quota in a qtree

Suppose you have the following quotas file:

```
#Quota target type           disk files thold sdisk sfile
#-----------  ----           ---- ----- ----- ----- -----
*             user@/vol/vol1  50M   -     45M
jsmith        user@/vol/vol1  80M   -     75M
```

It comes to your attention that a certain user, kjones, is taking up too much space in a critical qtree, qt1, which resides in vol2. You can restrict this user's space by adding the following line to the quotas file:

```
kjones        user@/vol/vol2/qt1  20M  -   15M
```

How qtree changes affect quotas

When you delete, rename, or change the security style of a qtree, the quotas applied by Data ONTAP might change, depending on the current quotas being applied.

How renaming a qtree affects quotas

When you rename a qtree, its ID does not change. As a result, all quotas applicable to the qtree continue to be applicable, without reinitializing quotas. However, before you reinitialize quotas, you must update the quota with the new qtree name to ensure that the quota continues to be applied for that qtree.

How deleting a qtree affects tree quotas

When you delete a qtree, all quotas applicable to that qtree, whether they are explicit or derived, are no longer applied by Data ONTAP.

If you create a new qtree with the same name as the one you deleted, the quotas previously applied to the deleted qtree are not applied automatically to the new qtree until you reinitialize quotas. If a default tree quota exists, Data ONTAP creates new derived quotas for the new qtree.

If you don't create a new qtree with the same name as the one you deleted, you can delete the quotas that applied to that qtree to avoid getting errors when you reinitialize quotas.

How changing the security style of a qtree affects user quotas

ACLs apply in qtrees using NTFS or mixed security style, but not in qtrees using UNIX security style. Therefore, changing the security style of a qtree might affect how quotas are calculated. You should always reinitialize quotas after you change the security style of a qtree.

If you change a qtree's security style from NTFS or mixed to UNIX, any ACLs on files in that qtree are ignored as a result, and file usage is charged against UNIX user IDs.

If you change a qtree's security style from UNIX to either mixed or NTFS, previously hidden ACLs become visible, any ACLs that were ignored become effective again, and the NFS user information is ignored.

Note: If no ACL existed before, the NFS information continues to be used in the quota calculation.

Attention: To make sure that quota usages for both UNIX and Windows users are properly calculated after you change the security style of a qtree, always reinitialize quotas for the volume containing that qtree.

Example

Suppose NTFS security is in effect on qtree A, and an ACL gives Windows user corp\joe ownership of a 5-MB file. User corp\joe is charged with 5 MB of disk space usage for qtree A.

Now you change the security style of qtree A from NTFS to UNIX. After quotas are reinitialized, Windows user corp\joe is no longer charged for this file; instead, the UNIX user corresponding to the UID of the file is charged for the file. The UID could be a UNIX user mapped to corp\joe or the root user.

Note: Only UNIX group quotas apply to qtrees. Changing the security style of a qtree, therefore, does not affect the group quotas.

How quotas work with users and groups

When you specify a user or group as the target of a quota, the limits imposed by that quota are applied to that user or group. However, some special groups and users are handled differently. There are different ways to specify IDs for users, depending on your environment.

When a full quota reinitialization is required

Although resizing quotas is faster, you must do a full quota reinitialization if you make certain or extensive changes to your quotas.

A full quota reinitialization is necessary in the following circumstances:

- You create a quota for a target that has not previously had a quota.
- You change user mapping in the `usermap.cfg` file and you use the QUOTA_PERFORM_USER_MAPPING entry in the quotas file.
- You change the security style of a qtree from UNIX to either mixed or NTFS.
- You change the security style of a qtree from mixed or NTFS to UNIX.
- You remove users from a quota target with multiple users, or add users to a target that already has multiple users.
- You make extensive changes to your quotas.

Example quotas changes that require initialization

Suppose you have a volume that contains three qtrees and the only quotas in the volume are three tree quotas. You decided to make the following changes:

- Add a new qtree and create a new tree quota for it.
- Add a default user quota for the volume.

Both of these changes require a full quota initialization. Resizing would not make the quotas effective.

Configuring quotas

Creating quotas

Quotas enable you to restrict or track the disk space and number of files used by a user, group, or qtree. You can use the Add Quota wizard to create a quota and apply it to a specific volume or qtree.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Quotas**.
3. In the **User Defined Quotas** tab, click **Create**.
4. Type or select information as prompted by the wizard.
5. Confirm the details and click **Finish** to complete the wizard.

After you finish

You can use the local user name or RID to create user quotas. If you create the user quota or group quota using the username or group name, then the `/etc/passwd file` and `/etc/group file` must be updated respectively.

Related references

Quotas window on page 101

Deleting quotas

You can delete one or more quotas as your users and their storage requirements and limitations change.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Quotas**.
3. Select one or more quotas that you want to delete and click **Delete**.
4. Select the confirmation check box and click **Delete**.

Related references

Quotas window on page 101

Managing quotas

Editing quota limits

You can use the Edit Limits dialog box to edit the disk space threshold, the hard and soft limits on the amount of disk space that the quota target can use, and the hard and soft limits on the number of files that the quota target can own.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Quotas**.
3. Select the quota that you want to edit and click **Edit Limits**.
4. Edit the quota settings as required.
5. Click **Save and Close** to save your changes and close the dialog box.
6. Verify the changes that you made to the selected quota in the **User Defined Quotas** tab.

Related references

Quotas window on page 101

Activating or deactivating quotas

You can activate or deactivate quotas on one or more selected volumes on your storage system, as your users and their storage requirements and limitations change.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Quotas**.
3. In the **Quota Status on Volumes** tab, select one or more volumes for which you want to activate or deactivate quotas.
4. Click either **Activate** or **Deactivate**, as required.
5. If you are deactivating a quota, select the confirmation check box and click **OK**.
6. Check the Status column to verify the quota status on the volumes.

Related references

Quotas window on page 101

Resizing quotas

You can use the Resize Quota dialog box to adjust the currently active quotas in the specified volume so that they reflect the changes that you have made to a quota.

Before you begin

Quotas must be enabled for the volumes for which you want to resize quotas.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Quotas**.
3. In the **Quota Status on Volumes** tab, select one or more volumes for which you want to resize the quotas.
4. Click **Resize**.

Related references

Quotas window on page 101

Monitoring quotas

Viewing quota information

You can use the Quotas window to view quota details such as the volume and the qtrees to which the quota is applied, the type of quota, the user or group to which the quota is applied, and the space and file usage.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Quotas**.
3. Perform the appropriate action:

If...	Then...
You want to view details of all the quotas that you created	Click **User Defined Quotas**.
You want to view the details of the quotas that are currently active	Click **Quota Reports**.

4. Select the quota that you want to view information about from the displayed list of quotas.
5. Review the quota details.

Window descriptions

Quotas window

You can use the Quotas window to create, display, and manage information about quotas.

- *Tabs* on page 101
- *Command buttons* on page 101
- *Quotas list* on page 101
- *Details area* on page 102

Tabs

User Defined Quotas	You can use the **User Defined Quotas** tab to view details of the quotas that you create and to create, edit, or delete quotas.
Quota Report	You can use the Quota Report tab to view the space and file usage and to edit the space and file limits of quotas that are active.
Quota Status on Volumes	You can use the Quota Status on Volumes tab to view the status of a quota and to turn quotas on or off and to resize quotas.

Command buttons

Create	Launches the Create Quota wizard, which enables you to create quotas.
Edit Limits	Opens the Edit Limits dialog box, which enables you to edit settings of the selected quota.
Delete	Deletes the selected quota from the quotas list.
Refresh	Updates the information in the window.

Quotas list

The quotas list displays the name and storage information for each quota.

Volume	Specifies the volume to which the quota is applied.
Qtree	Specifies the qtree associated with the quota. "All Qtrees" indicates that this quota is associated with all qtrees.
Type	Specifies the quota type: user, or group, or tree.
User/Group	Specifies a user or a group associated with the quota. "All Users" indicates that the quota is associated with all users. "All groups" indicates that the quota is associated with all groups.

Details area

The area below the quotas list displays the quota details such as the quota error, space usage and limits, and file usage and limits.

Related tasks

Creating quotas on page 98
Deleting quotas on page 98
Editing quota limits on page 99
Activating or deactivating quotas on page 99
Resizing quotas on page 100

Qtrees

Understanding qtrees

What a qtree is

A qtree is a logically defined file system that can exist as a special subdirectory of the root directory within either a traditional volume or a FlexVol volume. You can create up to 4995 qtrees per volume. There is no maximum limit for the storage system as a whole. You can create qtrees for managing and partitioning your data within the volume.

In general, qtrees are similar to volumes. However, they have the following key differences:

- Snapshot copies can be enabled or disabled for individual volumes but not for individual qtrees.
- Qtrees do not support space reservations or space guarantees.

There are no restrictions on how much disk space can be used by the qtree or how many files can exist in the qtree.

Qtree options

You must specify the following when creating a qtree: a name for the qtree and the volume in which the qtree resides. By default, the security style of a qtree is the same as that for the root directory of the volume. By default, oplocks are enabled for each qtree. If you disable oplocks for the entire storage system, oplocks are not set even if you enable oplocks on a per-qtree basis.

Related concepts

Qtree name restrictions on page 103
About the CIFS oplocks setting on page 104
Security styles on page 104

When to use qtrees

Qtrees enable you to partition your data without incurring the overhead associated with a volume. You might create qtrees to organize your data, or to manage one or more of the following factors: quotas, backup strategy, security style, and CIFS oplocks setting.

The following list describes examples of qtree usage strategies:

- Quotas
  You can limit the size of the data used by a particular project, by placing all of that project's files into a qtree and applying a tree quota to the qtree.
- Backups
  You can use qtrees to keep your backups more modular, to add flexibility to backup schedules, or to limit the size of each backup to one tape.
- Security style
  If you have a project that needs to use NTFS-style security, because the members of the project use Windows files and applications, you can group the data for that project in a qtree and set its security style to NTFS, without requiring that other projects also use the same security style.
- CIFS oplocks settings
  If you have a project using a database that requires CIFS oplocks to be off, you can set CIFS oplocks to `off` for that project's qtree, while allowing other projects to retain CIFS oplocks.

Qtree name restrictions

Qtree names can be no more than 64 characters in length. In addition, using some special characters in qtree names, such as commas and spaces, can cause problems with other Data ONTAP capabilities, and should be avoided.

The following characters should be avoided in qtree names:

- Space
  Spaces in qtree names can prevent SnapMirror updates from working correctly.
- Comma
  Commas in qtree names can prevent quotas from working correctly for that qtree, unless the name is enclosed in double quotation marks.

Related concepts

Qtree options on page 102

Security styles

Storage systems running Data ONTAP operating system supports different types of security styles for a storage object. By default, the security style of a qtree is the same as that for the root directory of the volume.

UNIX The user's UID and GID, and the UNIX-style permission bits of the file or directory determine user access. The storage system uses the same method for determining access for both NFS and CIFS requests.

If you change the security style of a qtree or a volume from NTFS to UNIX, the storage system disregards the Windows NT permissions that were established when the qtree or volume used the NTFS security style.

NTFS For CIFS requests, Windows NT permissions determine user access. For NFS requests, the storage system generates and stores a set of UNIX-style permission bits that are at least as restrictive as the Windows NT permissions.

The storage system grants NFS access only if the UNIX-style permission bits allow the user access.

If you change the security style of a qtree or a volume from UNIX to NTFS, files created before the change do not have Windows NT permissions. For these files, the storage system uses only the UNIX-style permission bits to determine access.

Mixed Some files in the qtree or volume have the UNIX security style and some have the NTFS security style. A file's security style depends on whether the permission was last set from CIFS or NFS.

For example, if a file currently uses the UNIX security style and a CIFS user sends a set-ACL request to the file, the file's security style is changed to NTFS. If a file currently uses the NTFS security style and an NFS user sends a set-permission request to the file, the file's security style is changed to UNIX.

Related concepts

Qtree options on page 102

About the CIFS oplocks setting

Usually, you should leave CIFS oplocks (opportunistic locks) on for all volumes and qtrees. This is the default setting. However, you might turn CIFS oplocks off under certain circumstances.

CIFS oplocks enable the redirector on a CIFS client in certain file-sharing scenarios to perform client-side caching of read-ahead, write-behind, and lock information. A client can then work with a file (read or write it) without regularly reminding the server that it needs access to the file. This improves performance by reducing network traffic.

You might turn CIFS oplocks off on a volume or a qtree under either of the following circumstances:

- You are using a database application whose documentation recommends that CIFS oplocks be turned off.
- You are handling critical data and cannot afford even the slightest data loss.

Otherwise, you can leave CIFS oplocks on.

For more information about CIFS oplocks, see the CIFS section of the *Data ONTAP File Access and Protocols Management Guide for 7-Mode*.

Related concepts

Qtree options on page 102

Configuring qtrees

Creating qtrees

Qtrees enable you to manage and partition your data within the volume. You can use the Create Qtree dialog box to add a new qtree to a volume on your storage system.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Qtrees**.
3. Click **Create**.
4. In the **Details** tab, type a name for the qtree.
5. Select the volume to which you want to add this qtree.

   The Volume browse list includes only volumes that are online.
6. If you want to disable oplocks for the qtree, clear **Enable Oplocks for files and directories in this Qtree**.

   By default, oplocks are enabled for each qtree.
7. If you want to change the default inherited security style, select a new one.

   The default security style of the qtree is the security style of the volume that contains the qtree.
8. If you want to restrict disk space usage, click the **Quotas** tab.
   a) If you want to apply quota on the qtree, click **Qtree quota** and specify the disk space limit.
   b) If you want to apply quota for all the users on the qtree, click **User quota** and specify the disk space limit.
9. Click **Create**.
10. Verify that the new qtree you created is included in the list of qtrees in the **Qtrees** window.

Related references

Qtrees window on page 107

Deleting qtrees

You can delete a qtree and reclaim the disk space it uses within a volume. When you delete a qtree, all quotas applicable to that qtree are no longer applied by Data ONTAP.

Before you begin

- The qtree status must be normal.
- The qtree must not contain any LUN.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Qtrees**.
3. Select one or more qtrees that you want to delete and click **Delete**.
4. Select the confirmation check box and click **Delete**.
5. Verify that the qtree you deleted is no longer included in the list of qtrees in the **Qtrees** window.

Related references

Qtrees window on page 107

Managing qtrees

Editing qtrees

You can change the security style of the qtree or to enable or disable opportunistic locks (oplocks) on a qtree.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Qtrees**.
3. Select the qtree that you want to edit and click **Edit**.
4. In the **Edit Qtree** dialog box, edit the settings as required.
5. Click **OK**.
6. Verify the changes you made to the selected qtree in the **Qtrees** window.

Related references

Qtrees window on page 107

Monitoring qtrees

Viewing qtree information

You can use the Qtrees window to view the volume that contains the qtree; the name, security style, and status of the qtree; and the oplocks status.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Qtrees**.
3. Select the qtree that you want to view information about from the displayed list of qtrees.
4. Review the qtree details in the **Qtrees** window.

Window descriptions

Qtrees window

You can use the Qtrees window to create, display, and manage information about qtrees.

- *Command buttons* on page 107
- *Qtree list* on page 107

Command buttons

Create Opens the Create Qtree dialog box, which enables you to create a new qtree.

Edit Opens the Edit Qtree dialog box, which enables you to change the security style and to enable or disable oplocks (opportunistic locks) on a qtree.

Delete Deletes the selected qtree.

Note: This button is disabled unless the selected qtree has a name and the qtree status is normal.

Refresh Updates the information in the window.

Qtree list

The qtree list displays the volume in which the qtree resides and the qtree name.

Name Specifies the name of the qtree.

Volume	Specifies the name of the volume in which the qtree resides.
Security Style	Specifies the security style of the qtree.
Status	Specifies the current status of the qtree.
Oplocks	Specifies whether the oplocks setting is enabled or disabled for the qtree.

Related tasks

Creating qtrees on page 105
Deleting qtrees on page 106
Editing qtrees on page 106

Aggregates

Understanding aggregates

Aggregate management

System Manager includes several features that help you to create, edit, or delete aggregates.

When you create an aggregate, you must provide the following information:

- A name for the aggregate
- RAID type (double parity or RAID4), which specifies the level of RAID protection that you want to provide for this aggregate

  Note: RAID0 is used only for array LUNs and VSA systems.
- Disks to include in the aggregate
- Type of aggregate (for example, SnapLock, SyncMirror, and Hybrid)

How you use aggregates to provide storage to your volumes

To support the differing security, backup, performance, and data sharing needs of your users, you group the physical data storage resources on your storage system into one or more aggregates. You can design and configure your aggregates to provide the appropriate level of performance and redundancy for your storage requirements.

Each aggregate has its own RAID configuration, plex structure, and set of assigned disks or array LUNs. The aggregate provides storage, based on its configuration, to its associated FlexVol volumes.

Aggregates have the following characteristics:

- They can be composed of disks or array LUNs.
- They can be mirrored or unmirrored.
- They can be 64-bit or 32-bit format.

- If they are composed of disks, they can be single-tier (composed of only HDDs or only SSDs), or they can be Flash Pools, which include both of those storage types in two separate tiers.

For information about best practices for working with aggregates, see *Technical Report 3437: Storage Subsystem Resiliency Guide*.

Related information

TR 3437: Storage Subsystem Resiliency Guide

Introduction to 64-bit and 32-bit aggregate formats

Aggregates are either 64-bit or 32-bit format. 64-bit aggregates have much larger size limits than 32-bit aggregates. 64-bit and 32-bit aggregates can coexist on the same storage system.

32-bit aggregates have a maximum size of 16 TB; 64-bit aggregates' maximum size depends on the storage system model. For the maximum 64-bit aggregate size of your storage system model, see the *Hardware Universe* (formerly the *System Configuration Guide*) at *support.netapp.com/knowledge/docs/hardware/NetApp/syscfg/index.shtml*.

By default, newly created aggregates are 32-bit for storage systems running Data ONTAP versions earlier than 8.1, and 64-bit for storage systems running Data ONTAP 8.1 or later.

You can expand 32-bit aggregates to 64-bit aggregates by increasing their size beyond 16 TB. 64-bit aggregates, including aggregates that were previously expanded, cannot be converted to 32-bit aggregates.

About using thin provisioning with FlexVol volumes

Using thin provisioning, you can appear to provide more storage than is actually available from a given aggregate, as long as not all of that storage is currently being used. Thin provisioning is also called *aggregate overcommitment*.

The storage provided by the aggregate is used up only as reserved LUNs are created or data is appended to files in the volumes.

Note: The aggregate must provide enough free space to hold the metadata for each FlexVol volume it contains. The space required for a FlexVol volume's metadata is approximately 0.5 percent of the volume's configured size.

When the aggregate is overcommitted, it is possible for writes (hole writes or overwrites) to LUNs or files in volumes contained by that aggregate to fail if there is not sufficient free space available to accommodate the write.

You can configure a thinly-provisioned volume to automatically secure more space from its aggregate when it needs to. However, if you have overcommitted your aggregate, you must monitor your available space and add storage to the aggregate as needed to avoid write errors due to insufficient space.

For more information about thin provisioning, see Technical Reports 3563 and 3483.

Related information

TR-3563: NetApp Thin Provisioning
TR 3483: Thin Provisioning in a NetApp SAN or IP SAN Enterprise Environment

How Data ONTAP RAID groups work

A RAID group consists of one or more data disks or array LUNs, across which client data is striped and stored, and up to two parity disks, depending on the RAID level of the aggregate that contains the RAID group.

RAID-DP uses two parity disks to ensure data recoverability even if two disks within the RAID group fail.

RAID4 uses one parity disk to ensure data recoverability if one disk within the RAID group fails.

RAID0 does not use any parity disks; it does not provide data recoverability if any disks within the RAID group fail.

How Data ONTAP uses RAID to protect your data and data availability

Understanding how RAID protects your data and data availability can help you administer your storage systems more effectively.

For native storage, Data ONTAP uses RAID-DP (double-parity) or RAID Level 4 (RAID4) protection to ensure data integrity within a group of disks even if one or two of those disks fail. Parity disks provide redundancy for the data stored in the data disks. If a disk fails (or, for RAID-DP, up to two disks), the RAID subsystem can use the parity disks to reconstruct the data in the drive that failed.

For third-party storage, Data ONTAP stripes data across the array LUNs using RAID0. The storage arrays, not Data ONTAP, provide the RAID protection for the array LUNs that they make available to Data ONTAP.

RAID types

RAID-DP provides double-parity disk protection. RAID4 provides single-parity disk protection against single-disk failure within a RAID group.

With RAID4, if there is a second disk failure before data can be reconstructed from the data on the first failed disk, there is data loss. To avoid data loss when two disks fail, you can select RAID-DP. RAID-DP provides two parity disks to protect you from data loss when two disk failures occur in the same RAID group before the first failed disk can be reconstructed.

For array LUNs, Data ONTAP uses RAID0 RAID groups to determine where to allocate data to the LUNs on the storage array. The RAID0 RAID groups are not used for RAID data protection. The storage arrays provide the RAID data protection.

Understanding RAID disk types

Data ONTAP classifies disks as one of four types for RAID: data, hot spare, parity, or dParity. The RAID disk type is determined by how RAID is using a disk; it is different from the Data ONTAP disk type.

Data disk	Holds data stored on behalf of clients within RAID groups (and any data generated about the state of the storage system as a result of a malfunction).
Spare disk	Does not hold usable data, but is available to be added to a RAID group in an aggregate. Any functioning disk that is not assigned to an aggregate but is assigned to a system functions as a hot spare disk.
Parity disk	Stores row parity information that is used for data reconstruction when a single disk drive fails within the RAID group.
dParity disk	Stores diagonal parity information that is used for data reconstruction when two disk drives fail within the RAID group, if RAID-DP is enabled.

RAID protection levels for disks

Data ONTAP supports two levels of RAID protection for aggregates composed of disks in native disk shelves: RAID-DP and RAID4. RAID-DP is the default RAID level for new aggregates.

For more information about choosing RAID protection levels, see *Technical Report 3437: Storage Subsystem Resiliency Guide*.

Related information

TR 3437: Storage Subsystem Resiliency Guide

What RAID-DP protection is

If an aggregate is configured for RAID-DP protection, Data ONTAP reconstructs the data from one or two failed disks within a RAID group and transfers that reconstructed data to one or two spare disks as necessary.

RAID-DP provides double-parity disk protection when the following conditions occur:

- There is a single-disk or double-disk failure within a RAID group.
- There are media errors on a block when Data ONTAP is attempting to reconstruct a failed disk.

The minimum number of disks in a RAID-DP group is three: at least one data disk, one regular parity disk, and one double-parity (or dParity) disk.

If there is a data-disk or parity-disk failure in a RAID-DP group, Data ONTAP replaces the failed disk in the RAID group with a spare disk and uses the parity data to reconstruct the data of the failed disk on the replacement disk. If there is a double-disk failure, Data ONTAP replaces the failed disks in the RAID group with two spare disks and uses the double-parity data to reconstruct the data of the failed disks on the replacement disks.

RAID-DP is the default RAID type for all aggregates.

What RAID4 protection is

RAID4 provides single-parity disk protection against single-disk failure within a RAID group. If an aggregate is configured for RAID4 protection, Data ONTAP reconstructs the data from a single failed disk within a RAID group and transfers that reconstructed data to a spare disk.

The minimum number of disks in a RAID4 group is two: at least one data disk and one parity disk.

If there is a single data or parity disk failure in a RAID4 group, Data ONTAP replaces the failed disk in the RAID group with a spare disk and uses the parity data to reconstruct the failed disk's data on the replacement disk. If no spare disks are available, Data ONTAP goes into degraded mode and alerts you of this condition.

Attention: With RAID4, if there is a second disk failure before data can be reconstructed from the data on the first failed disk, there will be data loss. To avoid data loss when two disks fail, you can select RAID-DP. This provides two parity disks to protect you from data loss when two disk failures occur in the same RAID group before the first failed disk can be reconstructed.

Note: Nondisruptive upgrade is not supported for aggregates configured for RAID4. For more information about nondisruptive upgrade, see the *Data ONTAP Upgrade and Revert/Downgrade Guide for 7-Mode*.

How RAID groups are named

Within each aggregate, RAID groups are named rg0, rg1, rg2, and so on in order of their creation. You cannot specify the names of RAID groups.

About RAID group size

A RAID group has a maximum number of disks or array LUNs that it can contain. This is called its maximum size, or its size. A RAID group can be left partially full, with fewer than its maximum number of disks or array LUNs, but storage system performance is optimized when all RAID groups are full.

Considerations for sizing RAID groups for disks

Configuring an optimum RAID group size for an aggregate made up of disks requires a trade-off of factors. You must decide which factor—speed of recovery, assurance against data loss, or maximizing data storage space—is most important for the aggregate that you are configuring.

You change the size of RAID groups on a per-aggregate basis. You cannot change the size of an individual RAID group.

HDDs

You should follow these guidelines when sizing your RAID groups for HDD disks:

- All RAID groups in an aggregate should have the same number of disks.

If this is impossible, any RAID group with fewer disks should have only one less disk than the largest RAID group.

- The recommended range of RAID group size is between 12 and 20.
  The reliability of SAS and FC disks can support a RAID group size of up to 28 if needed.
- If you can satisfy the first two guidelines with multiple RAID group sizes, you should choose the larger size.

SSDs

You should follow these guidelines when sizing your RAID groups for SSD disks:

- All RAID groups in an aggregate should have the same number of disks.
  If this is impossible, any RAID group with fewer disks should have only one less disk than the largest RAID group.
- The recommended range of RAID group size is between 20 and 28.

Flash Pools

For Flash Pools, the SSD tier must have the same RAID group size as the HDD tier. You should use the HDD guidelines to determine the RAID group size for the entire Flash Pool.

Considerations for Data ONTAP RAID groups for array LUNs

Setting up Data ONTAP RAID groups for array LUNs requires planning and coordination with the storage array administrator so that the administrator makes the number and size of array LUNs you need available to Data ONTAP.

For array LUNs, Data ONTAP uses RAID0 RAID groups to determine where to allocate data to the LUNs on the storage array. The RAID0 RAID groups are not used for RAID data protection. The storage arrays provide the RAID data protection.

Note: Data ONTAP RAID groups are similar in concept to what storage array vendors call RAID groups, parity groups, disk groups, Parity RAID groups, and other terms.

Follow these steps when planning your Data ONTAP RAID groups for array LUNs:

1. Plan the size of the aggregate that best meets your data needs.
2. Plan the number and size of RAID groups that you need for the size of the aggregate.
   Follow these guidelines:
   - RAID groups in the same aggregate should be the same size with the same number of LUNs in each RAID group. For example, you should create four RAID groups of 8 LUNs each, not three RAID groups of 8 LUNs and one RAID group of 6 LUNs.
   - Use the default RAID group size for array LUNs, if possible. The default RAID group size is adequate for most organizations.

     Note: The default RAID group size is different for array LUNs and disks.
3. Plan the size of the LUNs that you need in your RAID groups.

- To avoid a performance penalty, all array LUNs in a particular RAID group should be the same size.
- The LUNs should be the same size in all RAID groups in the aggregate.

4. Ask the storage array administrator to create the number of LUNs of the size you need for the aggregate.
   The LUNs should be optimized for performance, according to the instructions in the storage array vendor documentation.
5. Create all the RAID groups in the aggregate at the same time.

   Note: Do not mix array LUNs from storage arrays with different characteristics in the same Data ONTAP RAID group.

   Note: If you create a new RAID group for an existing aggregate, be sure that the new RAID group is the same size as the other RAID groups in the aggregate, and that the array LUNs are the same size as the LUNs in the other RAID groups in the aggregate.

How Data ONTAP works with hot spare disks

A hot spare disk is a disk that is assigned to a storage system but is not in use by a RAID group. It does not yet hold data but is ready for use. If a disk failure occurs within a RAID group, Data ONTAP automatically assigns hot spare disks to RAID groups to replace the failed disks.

How many hot spares you should have

Having insufficient spares increases the risk of a disk failure with no available spare, resulting in a degraded RAID group. The number of hot spares you should have depends on the Data ONTAP disk type.

MSATA disks, or disks in a multi-disk carrier, should have four hot spares during steady state operation, and you should never allow the number of MSATA hot spares to dip below two.

For RAID groups composed of SSDs, you should have at least one spare disk.

For all other Data ONTAP disk types, you should have at least one matching or appropriate hot spare available for each kind of disk installed in your storage system. However, having two available hot spares for all disks provides the best protection against disk failure. Having at least two available hot spares provides the following benefits:

- When you have two or more hot spares for a data disk, Data ONTAP can put that disk into the maintenance center if needed.
  Data ONTAP uses the maintenance center to test suspect disks and take offline any disk that shows problems.
- Having two hot spares means that when a disk fails, you still have a spare available if another disk fails before you replace the first failed disk.

A single spare disk can serve as a hot spare for multiple RAID groups.

What disks can be used as hot spares

A disk must conform to certain criteria to be used as a hot spare for a particular data disk.

For a disk to be used as a hot spare for another disk, it must conform to the following criteria:

- It must be either an exact match for the disk it is replacing or an appropriate alternative.
- If SyncMirror is in use, the spare must be in the same pool as the disk it is replacing.
- The spare must be owned by the same system as the disk it is replacing.

What a matching spare is

A matching hot spare exactly matches several characteristics of a designated data disk. Understanding what a matching spare is, and how Data ONTAP selects spares, enables you to optimize your spares allocation for your environment.

A matching spare is a disk that exactly matches a data disk for all of the following criteria:

- Effective Data ONTAP disk type
  The effective disk type can be affected by the value of the `raid.disktype.enable` option, which affects which disk types are considered to be equivalent.
- Size
- Speed (RPM)
- Checksum type (BCS or AZCS)

What an appropriate hot spare is

If a disk fails and no hot spare disk that exactly matches the failed disk is available, Data ONTAP uses the best available spare. Understanding how Data ONTAP chooses an appropriate spare when there is no matching spare enables you to optimize your spare allocation for your environment.

Data ONTAP picks a non-matching hot spare based on the following criteria:

- If the available hot spares are not the correct size, Data ONTAP uses one that is the next size up, if there is one.
  The replacement disk is downsized to match the size of the disk it is replacing; the extra capacity is not available.
- If the available hot spares are not the correct speed, Data ONTAP uses one that is a different speed.
  Using drives with different speeds within the same aggregate is not optimal. Replacing a disk with a slower disk can cause performance degradation, and replacing a disk with a faster disk is not cost-effective.
- If the failed disk is part of a mirrored aggregate and there are no hot spares available in the correct pool, Data ONTAP uses a spare from the other pool.
  Using drives from the wrong pool is not optimal because you no longer have fault isolation for your SyncMirror configuration.

If no spare exists with an equivalent disk type or checksum type, the RAID group that contains the failed disk goes into degraded mode; Data ONTAP does not combine effective disk types or checksum types within a RAID group.

RAID protection for third-party storage

Third-party storage arrays provide the RAID protection for the array LUNs that they make available to V-Series systems, not Data ONTAP.

Data ONTAP uses RAID 0 (striping) for array LUNs. Data ONTAP supports a variety of RAID types on the storage arrays, except RAID 0 because RAID 0 does not provide storage protection.

When creating *RAID groups* on storage arrays, you need to follow the best practices of the storage array vendor to ensure that there is an adequate level of protection on the storage array so that disk failure does not result in loss of data or loss of access to data.

Note: A *RAID group* on a storage array is the arrangement of disks that together form the defined RAID level. Each RAID group supports only one RAID type. The number of disks that you select for a RAID group determines the RAID type that a particular RAID group supports. Different storage array vendors use different terms to describe this entity—RAID groups, parity groups, disk groups, Parity RAID groups, and other terms.

V-Series systems support native disk shelves as well as third-party storage. Data ONTAP supports RAID4 and RAID-DP on the native disk shelves connected to a V-Series system but does not support RAID4 and RAID-DP with array LUNs.

See the *V-Series Implementation Guide for Third-Party Storage* to determine whether there are specific requirements or limitations about RAID types for your storage array.

What happens when you add larger disks to an aggregate

What Data ONTAP does when you add disks to an aggregate that are larger than the existing disks depends on the RAID level (RAID4 or RAID-DP) of the aggregate.

- When an aggregate configured for RAID4 protection is created, Data ONTAP assigns the role of parity disk to the largest disk in each RAID group.

  When an existing RAID4 group is assigned an additional disk that is larger than the group's existing parity disk, then Data ONTAP reassigns the new disk as parity disk for that RAID group.
- When an aggregate configured for RAID-DP protection is created, Data ONTAP assigns the role of dParity disk and regular parity disk to the largest and second largest disk in the RAID group.

  When an existing RAID-DP group is assigned an additional disk that is larger than the group's existing dParity disk, then Data ONTAP reassigns the new disk as the regular parity disk for that RAID group and restricts its capacity to be the same size as the existing dParity disk. Note that Data ONTAP does *not* replace the existing dParity disk, even if the new disk is larger than the dParity disk.

  Note: Because the smallest parity disk limits the effective size of disks added to a RAID-DP group, you can maximize available disk space by ensuring that the regular parity disk is as large as the dParity disk.

Note: If needed, you can replace a capacity-restricted disk with a more suitable (smaller) disk later, to avoid wasting disk space. However, replacing a disk already in use in an aggregate with a larger disk does not result in any additional usable disk space; the new disk is capacity-restricted to be the same size as the smaller disk it replaced.

Maximum number of RAID groups

Data ONTAP supports up to 400 RAID groups per storage system or HA pair.

When configuring your aggregates, keep in mind that each aggregate requires at least one RAID group and that the total of all RAID groups in a storage system cannot exceed 400.

How Flash Pools work

The Flash Pool technology enables you to add one or more RAID groups composed of SSDs to an aggregate that consists of RAID groups of HDDs.

The SSDs provide a high-performance cache for the active data set of the data volumes provisioned on the Flash Pool, which can increase the performance of the volumes associated with the aggregate without requiring the aggregate to be composed entirely of SSDs.

The SSD cache does not contribute to the size of the aggregate as calculated against the maximum aggregate size. For example, even if an aggregate is at the maximum aggregate size, you can add an SSD RAID group to it.

The HDD RAID groups in a Flash Pool behave the same as HDD RAID groups in a standard aggregate, following the same rules for mixing disk types, sizes, speeds, and checksums.

The checksum type, RAID type, and RAID group size values can be configured for the SSD cache RAID groups and HDD RAID groups independently.

There is a platform-dependent maximum size for the SSD cache. For information about this limit for your platform, see the *Hardware Universe*.

Rules for mixing HDD types in aggregates

You can mix disks from different loops or stacks within the same aggregate. Depending on the value of the `raid.mix.hdd.disktype` RAID options, you can mix certain types of HDDs within the same aggregate, but some disk type combinations are more desirable than others.

When the appropriate `raid.mix.hdd.disktype` option is set to `off`, single-tier aggregates and the HDD tier of Flash Pools can be composed of only one Data ONTAP disk type. This setting ensures that your aggregates are homogeneous, and requires that you provide sufficient spare disks for every disk type in use in your system.

The default value for the `raid.mix.hdd.disktype.performance` option is `off`, to prevent mixing SAS and FC-AL disks.

The default value for the `raid.mix.hdd.disktype.capacity` option is `on`. For this setting, the SATA, BSAS, FSAS, and ATA disk types are considered to be equivalent for the purposes of creating and adding to aggregates, and spare management.

To maximize aggregate performance and for easier storage administration, you should avoid mixing FC-AL-connected and SAS-connected disks in the same aggregate. This is because of the performance mismatch between FC-AL-connected disk shelves and SAS-connected disk shelves. When you mix these connection architectures in the same aggregate, the performance of the aggregate is limited by the presence of the FC-AL-connected disk shelves, even though some of the data is being served from the higher-performing SAS-connected disk shelves.

MSATA disks cannot be mixed with any other disk type in the same aggregate.

Disks using Storage Encryption have a Data ONTAP disk type of SAS. However, they cannot be mixed with any other disk type, including SAS disks that are not using Storage Encryption. If any disks on a storage system use Storage Encryption, all of the disks on the storage system (and its high-availability partner node) must use Storage Encryption.

Note: If you set a `raid.mix.hdd.disktype` option to `off` for a system that already contains aggregates with more than one type of HDD, those aggregates continue to function normally and accept both types of HDDs. However, no other aggregates composed of the specified disk type will accept mixed HDD types as long as that option is set to `off`.

For information about best practices for working with different types of disks, see *Technical Report 3437: Storage Best Practices and Resiliency Guide*.

Related information

TR 3437: Storage Best Practices and Resiliency Guide

How you can use effective Data ONTAP disk type for mixing HDDs

Starting with Data ONTAP 8.1, certain Data ONTAP disk types are considered equivalent for the purposes of creating and adding to aggregates, and managing spares. Data ONTAP assigns an effective disk type for each disk type. You can mix HDDs with the same effective disk type.

When the `raid.disktype.enable` option is set to `off`, you can mix certain types of HDDs within the same aggregate. When the `raid.disktype.enable` option is set to `on`, the effective disk type is same as the Data ONTAP disk type. Aggregates can be created using only one disk type. The default value for the `raid.disktype.enable` option is `off`.

Starting with Data ONTAP 8.2, the option `raid.mix.hdd.disktype.capacity` must be set to `on` to mix disks of type SATA, BSAS, FSAS, and ATA. The option `raid.mix.hdd.disktype.performance` must be set to `on` to mix disks of type FCAL and SAS.

The following table shows how the disk types map to the effective disk type:

Data ONTAP disk type	Effective disk type
FCAL	SAS

Data ONTAP disk type	Effective disk type
SAS	SAS
ATA	SATA
SATA	SATA
BSAS	SATA
FCAL and SAS	SAS
ATA and SATA	SATA
ATA, SATA, and BSAS	SATA

Requirements for using Flash Pools

Flash Pools have some configuration requirements that you should be aware of before planning to use them in your storage architecture.

A Flash Pool can use either RAID-DP or RAID4 protection (but not both in the same aggregate).

Flash Pools can be created from mirrored aggregates; however, the SSD configuration must be kept the same for both plexes.

Flash Pools cannot be used on all platforms. For a list of the platforms that support Flash Pools, see the *Hardware Universe*.

Flash Pools cannot be used in the following configurations:

- 32-bit aggregates
- Aggregates that use third-party storage
- Aggregates that use the ZCS checksum type
- SnapLock aggregates
- A storage system that uses Storage Encryption

You can use Flash Pools and the Flash Cache module (WAFL external cache) in the same system. However, data stored in a Flash Pool is not cached in the Flash Cache module. Flash Cache is reserved for data stored in aggregates composed of only HDDs. For more information about Flash Cache and WAFL external cache, see the *Data ONTAP System Administration Guide for 7-Mode*.

You must either disable the automatic creation of aggregate Snapshot copies or enable the automatic deletion of aggregate Snapshot copies for a Flash Pool. These operations must be done by using the Data ONTAP Command Line Interface. For information about automatic aggregate Snapshot copy creation and deletion, see the *Data ONTAP System Administration Guide for 7-Mode*.

If you create a Flash Pool using an aggregate that was created using Data ONTAP 7.1 or earlier, the volumes associated with that Flash Pool will not support write caching.

When you cannot use aggregates composed of SSDs

Aggregates composed of SSDs have some restrictions on when they can be used.

You cannot use aggregates composed of SSDs with the following configurations or technologies:

- SnapLock
- Storage Encryption

What SyncMirror is

SyncMirror is an optional feature of Data ONTAP. It is used to mirror data to two separate aggregates. It allows for real-time mirroring of data to matching aggregates physically connected to the same storage system.

SyncMirror provides for synchronous mirroring of data, implemented at the RAID level. You can use SyncMirror to create aggregates that consist of two copies of the same WAFL file system. The two copies, known as plexes, are simultaneously updated. Therefore, the copies are always identical. The two plexes are directly connected to the same system.

The following provides information about the activities of SyncMirror:

- SyncMirror can be used to mirror aggregates and traditional volumes. (A traditional volume is essentially an aggregate with a single volume that spans the entire aggregate.)
- SyncMirror cannot be used to mirror FlexVol volumes. However, FlexVol volumes can be mirrored as part of an aggregate.
- SyncMirror is different from synchronous SnapMirror.

For more information about aggregates and volumes, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Related information

Data ONTAP Information Library page: support.netapp.com//documentation/productsatoz/index.html

Advantages of using SyncMirror

A SyncMirror aggregate has two plexes. This setup provides a high level of data availability because the two plexes are physically separated.

For a system using disks, the two plexes are on different shelves connected to the system with separate cables and adapters. Each plex has its own collection of spare disks. For a system using third-party storage, the plexes are on separate sets of array LUNs, either on one storage array or on separate storage arrays.

Note: You cannot set up SyncMirror with disks in one plex and array LUNs in the other plex.

Physical separation of the plexes protects against data loss if one of the shelves or the storage array becomes unavailable. The unaffected plex continues to serve data while you fix the cause of the failure. Once fixed, the two plexes can be resynchronized.

Another advantage of mirrored plexes is faster rebuild time.

In contrast, if an aggregate using SnapMirror for replication becomes unavailable, you can use one of the following options to access the data on the SnapMirror destination (secondary).

- The SnapMirror destination cannot automatically take over the file serving functions. However, you can manually set the SnapMirror destination to allow read-write access to the data.
- You can restore the data from the SnapMirror destination to the primary (source) storage system.

An aggregate mirrored using SyncMirror requires twice as much storage as an unmirrored aggregate. Each of the two plexes requires an independent set of disks or array LUNs. For example, you need 2,880 GB of disk space to mirror a 1,440-GB aggregate—1,440 GB for each plex of the mirrored aggregate.

Protection provided by RAID and SyncMirror

Combining RAID and SyncMirror provides protection against more types of disk failures than using RAID alone.

You can use RAID in combination with the SyncMirror functionality, which also offers protection against data loss due to disk or other hardware component failure. SyncMirror protects against data loss by maintaining two copies of the data contained in the aggregate, one in each plex. Any data loss due to disk failure in one plex is repaired by the undamaged data in the other plex.

For more information about SyncMirror, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

The following tables show the differences between using RAID alone and using RAID with SyncMirror:

Table 1: RAID-DP and SyncMirror

Criteria	RAID-DP alone	RAID-DP with SyncMirror
Failures protected against	• Single-disk failure • Double-disk failure within a single RAID group • Multiple-disk failures, as long as no more than two disks within a single RAID group fail	• All failures protected against by RAID-DP alone • Any combination of failures protected against by RAID-DP alone in one plex, concurrent with an unlimited number of failures in the other plex • Storage subsystem failures (HBA, cables, shelf), as long as only one plex is affected
Failures *not* protected against	• Three or more concurrent disk failures within a single RAID group • Storage subsystem failures (HBA, cables, shelf) that lead to three or more concurrent disk failures within a single RAID group	• Three or more concurrent disk failures in a single RAID group on both plexes
Required disk resources per RAID group	*n* data disks + 2 parity disks	2 x (*n* data disks + 2 parity disks)
Performance cost	Almost none	Low mirroring overhead; can improve performance
Additional cost and complexity	None	SyncMirror license and configuration

Table 2: RAID4 and SyncMirror

Criteria	RAID4 alone	RAID4 with SyncMirror
Failures protected against	• Single-disk failure • Multiple-disk failures, as long as no more than one disk within a single RAID group fails	• All failures protected against by RAID4 alone • Any combination of failures protected against by RAID4 alone in one plex, concurrent with an unlimited number of failures in the other plex • Storage subsystem failures (HBA, cables, shelf), as long as only one plex is affected
Failures *not* protected against	• Two or more concurrent disk failures within a single RAID group • Storage subsystem failures (HBA, cables, shelf) that lead to two or more concurrent disk failures within a single RAID group	• Two or more concurrent disk failures in a single RAID group on both plexes
Required disk resources per RAID group	*n* data disks + 1 parity disk	2 x (*n* data disks + 1 parity disk)
Performance cost	None	Low mirroring overhead; can improve performance
Additional cost and complexity	None	SyncMirror license and configuration

Table 3: RAID0 and SyncMirror

Criteria	RAID0 alone	RAID0 with SyncMirror
Failures protected against	No protection against any failures RAID protection is provided by the RAID implemented on the third-party storage array.	Any combination of array LUN, connectivity, or hardware failures, as long as only one plex is affected

Criteria	RAID0 alone	RAID0 with SyncMirror
Failures *not* protected against	No protection against any failures RAID protection is provided by the RAID implemented on the storage array.	Any concurrent failures that affect both plexes
Required array LUN resources per RAID group	No extra array LUNs required other than *n* data array LUNs	2 x *n* data array LUNs
Performance cost	None	Low mirroring overhead; can improve performance
Additional cost and complexity	None	SyncMirror license and configuration

What mirrored aggregates are

A mirrored aggregate is a single WAFL storage file system with two physically separated and synchronously up-to-date copies on disks or array LUNs. These copies are called plexes. Data ONTAP typically names the first plex `plex0` and the second plex `plex1`.

Each plex is a physical copy of the same WAFL file system, and consists of one or more RAID groups. As SyncMirror duplicates complete WAFL file systems, you cannot use the SyncMirror feature with a FlexVol volume—only aggregates (including all contained FlexVol volumes) are supported.

How mirrored aggregates work

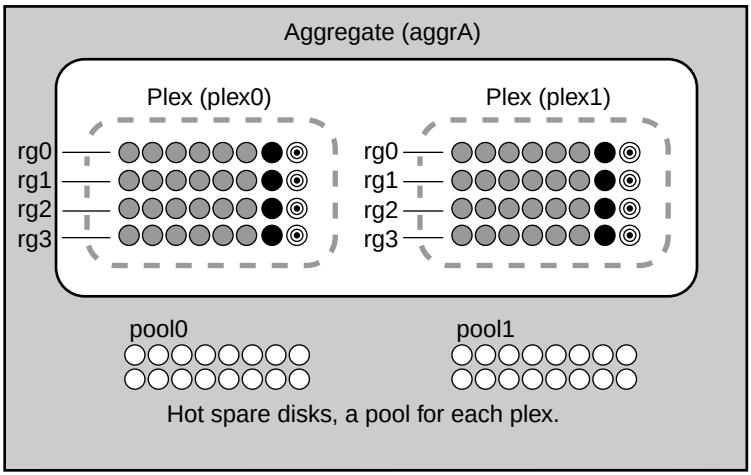
Mirrored aggregates have two *plexes* (copies of their data), which use the SyncMirror functionality to duplicate the data to provide redundancy.

When SyncMirror is enabled, all the disks or array LUNs are divided into two pools, and a copy of the plex is created. The plexes are physically separated (each plex has its own RAID groups and its own pool), and the plexes are updated simultaneously. This provides added protection against data loss if more disks fail than the RAID level of the aggregate protects against or there is a loss of connectivity, because the unaffected plex continues to serve data while you fix the cause of the failure. After the plex that had a problem is fixed, you can resynchronize the two plexes and reestablish the mirror relationship.

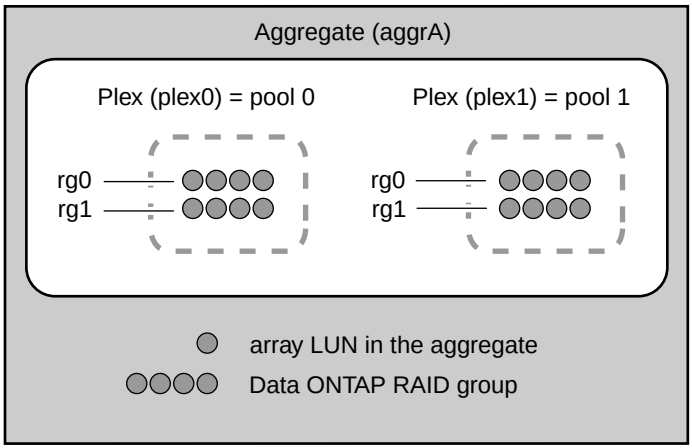
Note: Before an aggregate can be enabled for mirroring, the storage system must have the syncmirror_local license installed and enabled, and the storage configuration must support RAID-level mirroring.

In the following diagram of a storage system using disks, SyncMirror is enabled and implemented, so Data ONTAP copies plex0 and automatically names the copy plex1. Plex0 and plex1 contain copies of one or more file systems. In this diagram, 32 disks were available prior to the SyncMirror

relationship being initiated. After initiating SyncMirror, the spare disks are allocated to pool0 or pool1.



The following diagram shows a storage system using array LUNs with SyncMirror enabled and implemented.



Considerations for using mirrored aggregates

If you want to use mirrored aggregates, you can either create a new aggregate with two mirrored plexes, or add a plex to an existing aggregate.

Note: A mirrored aggregate can have only two plexes.

The rules for the selection of disks or array LUNs, for using as mirrored aggregates, are as follows:

- Disks or array LUNs selected for each plex must be in different pools.

- The same number of disks or array LUNs must be in both the plexes.
- Disks are selected first on the basis of equivalent bytes per sector (bps) size, then on the basis of the size of the disk.
- If there is no equivalent-sized disk, Data ONTAP uses a larger-capacity disk, and limits the size to make it identically sized.
- Data ONTAP names the plexes of the mirrored aggregate.

Note: When creating an aggregate, Data ONTAP selects disks from the plex which has the most available disks. You can override this selection policy by specifying the disks to use.

How disks are assigned to plexes

You need to understand how Data ONTAP assigns disks to plexes in order to configure your disk shelves and host adapters.
When a mirrored aggregate is created, Data ONTAP uses spare disks from a collection of disks, to create two disk pools, pool0 and pool1.

When assigning a disk to a pool, Data ONTAP determines the shelf for the disk and ensures that the disks in pool0 are from different shelves than the disks in pool1. So, before enabling SyncMirror, you should ensure that the disks are installed in at least two shelves and the shelves are connected to the system with separate cables and adapters. Disk pools must be physically separate to ensure high availability of the mirrored aggregate.

Disks from pool0 are used to create `plex0` while disks from pool1 are used to create `plex1`.

Plexes local to the host node in an HA pair must be connected to the disk pool named pool0. pool0 consists of the storage attached to host adapters in slots 3 through 7.

Note: Pool rules for MetroCluster configurations that use switches are different.

For more information about V-Series system slot assignments, see the *Hardware Universe*.

Rules for adding disks to a mirrored aggregate

You need to follow certain rules regarding the distribution and size of disks when adding disks to a mirrored aggregate.

- The number of disks must be even, and the disks must be equally divided between the two plexes.
- The disks for each plex must come from different disk pools.
- The disks that you add must have equivalent bytes per sector (bps) sizes.

When adding new disks to a RAID group, the utilization of the new disks depends on the RAID level used. If the storage capacity of the new disks is more than the disks already in the RAID group, the larger-capacity disks might be downsized to suit the RAID group.

- RAID-DP: Larger-capacity disks are downsized to size of parity disks.
- RAID-4: Larger-capacity disks can replace the parity disks.

The states of a plex

A plex can either be in an online state or in an offline state. In the online state, the plex is available for read or write access and the contents of the plex are current. In an offline state, the plex is not accessible for read or write.
An online plex can be in the following states.

- Active—The plex is available for use.
- Adding disks or array LUNs—Data ONTAP is adding disks or array LUNs to the RAID group or groups of the plex.
- Empty—The plex is part of an aggregate that is being created and Data ONTAP needs to zero out one or more of the disks or array LUNs targeted to the aggregate before adding the disks to the plex.
- Failed—One or more of the RAID groups in the plex failed.
- Inactive—The plex is not available for use.
- Normal—All RAID groups in the plex are functional.
- Out-of-date—The plex contents are out of date and the other plex of the aggregate has failed.
- Resyncing—The plex contents are being resynchronized with the contents of the other plex of the aggregate.

Configuring aggregates

Creating aggregates

You can create an aggregate or a Flash Pool to provide storage for one or more FlexVol volumes.

Before you begin

- You must have enough spare disks to create an aggregate.
- You must be aware of platform-specific and workload-specific best practices for the Flash Pool SSD tier size and configuration.
  For more information, see *Technical Report 4070: NetApp Flash Pool Design and Implementation Guide*.
- If you want to create mirrored aggregates, the SyncMirror license must be enabled.
- If you want to create a SnapLock aggregate, the SnapLock compliance license, or the SnapLock enterprise license, or both must be installed on the storage system.

About this task

- You cannot combine disks of different sizes when creating an aggregate or a Flash Pool using System Manager.
  You can initially create an aggregate with disks of the same size and add disks of a different size later.
- You cannot combine disks with different checksum types when creating an aggregate or a Flash Pool using System Manager.

You can initially create an aggregate with a single checksum type and add storage of a different checksum type later.

- You cannot combine disks of different sizes when creating an aggregate or a Flash Pool using System Manager even if there are enough spare disks of different sizes.

Using the Create Aggregate wizard, you can perform the following:

- Create a SnapLock aggregate if the SnapLock license is enabled.
- Enable a Flash Pool.

  Note: You cannot downgrade or revert the Data ONTAP version on your storage system after the Flash Pool is enabled.
- Specify the disk group to include.

  The minimum hot spare rule is applied to the disk group having the largest disk size.
- Specify the capacity disk size and the number of capacity disks to include.

  For storage systems running Data ONTAP 8.1 or a later version in the 8.1 release family, the minimum number of disks for RAID-DP is 3 and for RAID4 is 2. For storage systems running Data ONTAP 8.2, the minimum number of disks for RAID-DP is 5 and for RAID4 is 3.
- Specify the cache disk size and the number of cache disks to include.

  The cache disk size and the number of cache disks that you can include are based on the selected disk group.
- Specify the type of disks to include.

  All existing HDDs and SSDs that you want to include in the aggregate must be of the same size.
- Specify the RAID type for the RAID groups on the aggregate.
- Specify the RAID group size.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Aggregates**.
3. Click **Create**.
4. Type or select information as prompted by the wizard.
5. Confirm the details and click **Finish** to complete the wizard.

Result

The aggregate is created with the specified configuration and added to the list of aggregates in the Aggregates window.

Related references

Aggregates window on page 136

Related information

TR 4070: NetApp Flash Pool Design and Implementation Guide

Mirroring an aggregate

You can mirror aggregates to provide a high level of data availability. A mirrored aggregate consists of two plexes and has two copies of its data. You can use the Aggregate window to mirror an aggregate.

Before you begin

- SyncMirror license must be enabled on the storage system.
- The storage system must have disks in both the plexes.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Aggregates**.
3. Select the aggregate that you want to mirror and click **Mirror**.
4. Click **Mirror**.

Related references

Aggregates window on page 136

Deleting aggregates

You can delete aggregates when you no longer require the data in the aggregates. However, you cannot delete the root aggregate because it contains the root volume, which contains the system configuration information.

Before you begin

- All the FlexVol volumes contained by the aggregate must be deleted.
- The aggregate must be offline.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Aggregates**.
3. Select one or more aggregates that you want to delete and click **Delete**.
4. Select the confirmation check box and click **Delete**.
5. Verify that the deleted aggregates are no longer displayed in the Aggregates window.

Related references

Aggregates window on page 136

Managing aggregates

Editing aggregate settings

You can use the Edit Aggregate dialog box to change the aggregate name, RAID type, RAID group size, and to add capacity disks and cache disks to the aggregate. You can also convert the aggregate to a Flash Pool. However, you cannot modify the properties of a SnapLock Compliance aggregate.

Before you begin

- You must have sufficient spare disks to add additional disks to an aggregate.

About this task

When you add disks to an aggregate on storage systems running Data ONTAP 7.3.7, new disks are added only to the most recently created RAID group. When the existing RAID groups become full after the disks are added, new RAID groups are created and disks are added to the new RAID groups. The previously created RAID groups remain at their current size unless you add the disks to them.

You cannot edit a Flash Pool that has mixed RAID types for storage systems running Data ONTAP 8.2.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Aggregates**.
3. Select the aggregate that you want to edit and click **Edit**.
4. In the **Edit Aggregate** dialog box, perform the appropriate steps:

If you want to...	Do this...
Increase the storage	Specify the number of disks to add and click **Save and Close**. The spare disks from the default disk group are added to all the RAID groups of the aggregate.

If you want to...	Do this...
Specify which disks to use when increasing storage	**a.** Click **Advanced**, and select a disk group from the Advanced Disk Selection window. The minimum hot spare rule is applied to the disk group having the largest disk size. The disk count displayed in the Advanced Disk Selection window does not include the hot spare. **b.** Specify the RAID group to which you want to add the disks. You can add disks to all RAID groups or a specific RAID group, or create a new RAID group and add the disks. • For HDDs, you cannot add disks to all RAID groups if all the parity disks of the aggregate are not equal in size (within a ten percent tolerance range). You should create a new RAID group to add the disks, or specify an existing RAID group. • If you select a disk group with a disk size that is smaller than the selected RAID group parity disk's size, a warning message indicating a suboptimal configuration is displayed. • If you select a disk group with a disk size that is larger than the selected RAID group parity disk's size, a warning message indicating a downsized configuration is displayed. **c.** Specify the number of disks to add. **d.** Click **Save and Close**.

5. In the **Edit Aggregate** dialog box, perform the following steps if you want to modify the RAID type or RAID group size:
   a) Click **Change**.
   b) In the **RAID Details** window, specify the required details.
   c) Click **Save and Close**.

Result

The selected aggregate is updated with the new information. The updated information is also displayed in the **Details** tab of the Aggregates window.

Related references

Aggregates window on page 136

Taking a plex offline

A plex can either be in an online state or in an offline state. When a plex is offline, it is not available for read or write access. You can use the Aggregate window to take a plex offline.

Before you begin

The plex must be part of a mirrored aggregate and both plexes must be online.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Aggregates**.
3. Select the appropriate mirrored aggregate and click **Plexes** in the lower pane.
4. Select the plex you want to take offline and click **Offline**.
5. Select the confirmation check box and click **Offline**.

Related references

Aggregates window on page 136

Bringing a plex online

A plex can either be in an online state or in an offline state. In the online state, the plex is available for read or write access and the contents of the plex are current. You can use the Aggregate window to bring a plex online.

Before you begin

The plex must be part of a mirrored aggregate.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Aggregates**.
3. Select the appropriate mirrored aggregate and click **Plexes** in the lower pane.
4. Select a plex you want to bring online and click **Online**.
5. Select the confirmation check box and click **Online**.

Related references

Aggregates window on page 136

Destroying a plex

You can destroy a plex if you want to stop mirroring the aggregate, or if there is a problem with the plex. You can use the Aggregate window to destroy or remove a plex from a mirrored aggregate.

Before you begin

The plex must be offline.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Aggregates**.
3. Select the appropriate mirrored aggregate and click **Plexes** in the lower pane.
4. Select the plex that you want to destroy and click **Destroy**.
5. Click **Destroy** in the confirmation window.

Result

Destroying a plex results in an unmirrored aggregate, because the aggregate now has only one plex.

Related references

Aggregates window on page 136

Splitting a mirrored aggregate

You can use the **Aggregate** management tab to split a mirrored aggregate. You might split a mirrored aggregate to move a mirrored aggregate to another location. Splitting a mirrored aggregate removes the relationship between its two plexes and creates two independent unmirrored aggregates. After splitting, both the aggregates are online.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Aggregates**.
3. Select the appropriate mirrored aggregate and click **Plexes** in the lower pane.
4. Select a plex you want to split, and click **Split**.
5. Click **Split** in the confirmation window.
6. If you want to change the default name for the newly created aggregate, specify the new name.
7. Click **Split**.

Related references

Aggregates window on page 136

Upgrading to a 64-bit aggregate

For storage systems running Data ONTAP 8.1 and later, System Manager enables you to upgrade an existing 32-bit aggregate to a 64-bit aggregate by adding disks to increase their size beyond 16 TB.

Before you begin

If you want to add disks to the aggregate:

- All the existing disks in the aggregate must be of the same size.
- You must have sufficient homogeneous spare disks of the same size as the existing disks in the aggregate.

About this task

You cannot upgrade a SyncMirror aggregate to 64 bit.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Aggregates**.
3. Select the 32-bit aggregate whose size you want to increase and click **Add Disks**.
4. Select one or more homogeneous disks that are of the same size as the existing disks in the aggregate and click **Add**.
5. Select the confirmation check box and click **Add**.

   The **Upgrade to 64 bit aggregate Wizard** is displayed.
6. Type or select information as prompted by the wizard.
7. Confirm the details and click **Finish** to complete the wizard.

Related references

Aggregates window on page 136

Changing the state of an aggregate

An aggregate may be online, restricted, or offline. You can use the Aggregate window to take an aggregate offline, bring it back online, or restrict access to the aggregate. An aggregate cannot be restricted or taken offline if it has FlexVol volumes or mounted volumes.

About this task

When an aggregate is online, read and write access to volumes hosted on this aggregate is allowed. When an aggregate is offline, no read or write access is allowed. You can put the aggregate into a restricted state if you want the aggregate to be the target of an aggregate copy or SnapMirror replication operation.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Aggregates**.
3. Select the aggregate for which you want to modify the state.
4. From the **Status** menu, click the aggregate state you want.
5. In the confirmation dialog box, click **Offline** or **Restrict**, as appropriate.

Related references

Aggregates window on page 136

Converting an aggregate to a Flash Pool

You can convert a non-root aggregate that is composed of HDDs to a Flash Pool by adding one or more RAID groups composed of SSDs. The SSD tier functions as a high-performance cache to the working data set, increasing the performance of the aggregate without using SSDs for the entire aggregate.

Before you begin

- You must have identified a valid 64-bit non-root aggregate composed of HDDs to convert it to a Flash Pool.
- The aggregate must not be a SnapLock aggregate.
- The aggregate must not contain any array LUNs.
- You must have determined the SSDs you plan to add, and these SSDs must be owned by the node on which you are creating the Flash Pool.
- All the SSDs in the spare pool must be of the same size.

About this task

You should be aware of platform-specific and workload-specific best practices for Flash Pool SSD tier size and configuration. For more information, see *Technical Report 4070: NetApp Flash Pool Design and Implementation Guide*.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Aggregates**.
3. Select the aggregate that you want to convert to a Flash Pool and click **Edit**.
4. Select the option for enabling Flash Pool.
5. Specify the cache disk size and the number of SSDs that you want to add to create a Flash Pool.

The default value is the minimum number of cache disks that is required to create a RAID group. The default number of cache disks is 3 for RAID-DP and 2 for RAID4.

6. Click **Save and Close**.
7. Verify the changes you made to the selected aggregate in the **Details** tab in the **Aggregates** window.

Related information

TR 4070: NetApp Flash Pool Design and Implementation Guide

Monitoring aggregates

Viewing aggregate information

You can use the Aggregates window to view the name and status of and the space information about an aggregate.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage > Aggregates**.
3. Select the aggregate that you want to view information about from the displayed list of aggregates.
4. Review the aggregate details in the **Aggregates** window.

Window descriptions

Aggregates window

You can use the Aggregates window to create, display, and manage information about aggregates.

- *Command buttons* on page 136
- *Aggregate list* on page 137
- *Details area* on page 137

Command buttons

Create	Starts the Create Aggregate wizard, which enables you to create an aggregate.
Edit	Opens the Edit Aggregate dialog box, which enables you to change the name of an aggregate or the level of RAID protection you want to provide for this aggregate.
Delete	Deletes the selected aggregate. **Note:** This button is disabled for the root aggregate.

Status Displays the status of the selected aggregate. The status can be one of the following:

- Online
  Read and write access to volumes contained in this aggregate is allowed.
- Offline
  Some operations, such as parity reconstruction, are allowed, but data access is not allowed.
- Restrict
  No read or write access is allowed.

Mirror Opens a dialog box, which enables you to create a mirrored aggregate.

Refresh Updates the information in the window.

Aggregate list

The aggregate list displays the name and the space usage information for each aggregate.

Name	Displays the name of the aggregate.
Used (%)	Displays the percentage of space used in the aggregate.
Available Space	Displays the available space in the aggregate.
Used Space	Displays the amount of space that is used for data in the aggregate.
Total Space	Displays the total space of the aggregate.
Volume Count	Displays the number of volumes associated with the aggregate.
Disk Count	Displays the number of disks used to create the aggregate.
Status	Displays the current status of the aggregate.
Flash Pool	Displays the total cache size of a Flash Pool. A value of -NA- indicates that the aggregate is not a Flash Pool.
SnapLock	Displays the type of SnapLock attribute: compliance or enterprise. If this field is blank, it indicates that the SnapLock attribute was not set on the aggregate.

Details area

The area below the aggregate list displays detailed information about the selected aggregate.

Details tab	Displays detailed information about the selected aggregate.
Volumes tab	Displays details about the total number of volumes present on the aggregate, total aggregate space, and the space committed to by the aggregate. Details about the available space, total space, and the percentage of space utilization of each volume on the selected aggregate are also displayed.

Disk Layout tab	Displays disk layout information, such as the status, disk type, RAID type, checksum, RPM, and RAID group for the selected aggregate. The disk port associated with disk primary path and the disk name with the disk secondary path, for a multipath configuration are also displayed.

Related tasks

Creating aggregates on page 127
Mirroring an aggregate on page 129
Deleting aggregates on page 129
Editing aggregate settings on page 130
Taking a plex offline on page 131
Bringing a plex online on page 132
Destroying a plex on page 132
Splitting a mirrored aggregate on page 133
Upgrading to a 64-bit aggregate on page 134
Changing the state of an aggregate on page 134

Disks

Understanding disks

Disk management

System Manager includes several features that help you to create an aggregate from selected disks and add spare disks to an existing aggregate.

You can select the individual disks you want to use to create an aggregate by scrolling through the list of available disks in the Create Aggregate dialog box. You must select at least two disks (one data disk and one parity disk) for RAID-4 and at least three disks (one data disk, a regular parity disk, and a double-parity disk) for RAID-DP.

Three kinds of disks are available for the storage system's file system:

Data	Holds data stored on behalf of clients and data generated about the state of the storage system as a result of a malfunction.
Hot spare	Does not hold usable data, but is available for addition to an aggregate. You can also add a hot spare disk to an aggregate by adding the disk to a traditional volume contained by the aggregate.
Parity	Stores data reconstruction information.

How Data ONTAP works with hot spare disks

A hot spare disk is a disk that is assigned to a storage system but is not in use by a RAID group. It does not yet hold data but is ready for use. If a disk failure occurs within a RAID group, Data ONTAP automatically assigns hot spare disks to RAID groups to replace the failed disks.

What happens when you add storage to an aggregate

By default, Data ONTAP adds new disks or array LUNs to the most recently created RAID group until it reaches its maximum size. Then Data ONTAP creates a new RAID group. Alternatively, you can specify a RAID group you want to add storage to.

When you create an aggregate or add storage to an aggregate, Data ONTAP creates new RAID groups as each RAID group is filled with its maximum number of disks or array LUNs. The last RAID group formed might contain fewer disks or array LUNs than the maximum RAID group size for the aggregate. In that case, any storage added to the aggregate is added to the last RAID group until the specified RAID group size is reached.

If you increase the RAID group size for an aggregate, new disks or array LUNs are added only to the most recently created RAID group; the previously created RAID groups remain at their current size unless you explicitly add storage to them.

Note: You are advised to keep your RAID groups homogeneous when possible. If needed, you can replace a mismatched disk with a more suitable disk later.

How disk checksum types affect aggregate and spare management

There are two checksum types available for disks used by Data ONTAP: BCS (block) and AZCS (zoned). Understanding how the checksum types differ and how they impact storage management enables you to manage your storage more effectively.

Both checksum types provide the same resiliency capabilities; BCS optimizes data access speed and capacity for disks that use 520 byte sectors. AZCS provides enhanced storage utilization and capacity for disks that use 512 byte sectors (usually SATA disks, which emphasize capacity).

Aggregates have a checksum type, which is determined by the checksum type of the disks that compose the aggregate. The following configuration rules apply to aggregates, disks, and checksums:

- Checksum types cannot be combined within RAID groups.
  This means that you must consider checksum type when you provide hot spare disks.
- When you add storage to an aggregate, if it has a different checksum type than the storage in the RAID group to which it would normally be added, Data ONTAP creates a new RAID group.
- An aggregate can have RAID groups of both checksum types.
  These aggregates have a checksum type of `mixed`.
- For mirrored aggregates, both plexes must have the same checksum type.
- Disks of a different checksum type cannot be used to replace a failed disk.
- You cannot change the checksum type of a disk.

Checksum type by Data ONTAP disk type

You should know the Data ONTAP disk type and checksum type of all of the disks you manage, because these disk characteristics impact where and when the disks can be used.

The following table shows the checksum type by Data ONTAP disk type:

Data ONTAP disk type	Checksum type
SAS or FC-AL	BCS
SATA/BSAS/FSAS/ATA	BCS
SSD	BCS
MSATA	AZCS

Spare requirements for multi-disk carrier disks

Maintaining the proper number of spares for disks in multi-disk carriers is critical for optimizing storage redundancy and minimizing the amount of time Data ONTAP must spend copying disks to achieve an optimal disk layout.

You must maintain a minimum of two hot spares for multi-disk carrier disks at all times. To support the use of the Maintenance Center, and to avoid issues caused by multiple concurrent disk failures, you should maintain at least four hot spares for steady state operation, and replace failed disks promptly.

If two disks fail at the same time with only two available hot spares, Data ONTAP might not be able to swap the contents of both the failed disk and its carrier mate to the spare disks. This scenario is called a *stalemate*. If this happens, you are notified through EMS messages and AutoSupport messages. When the replacement carriers become available, you must follow the instructions provided by the EMS messages or contact technical support to recover from the stalemate.

Shelf configuration requirements for multi-disk carrier disk shelves

You can combine multi-disk carrier disk shelves with single-disk carrier disk shelves (standard disk shelves) on the same storage system. However, you cannot combine the two disk shelf types in the same stack.

Aggregate requirements for disks from multi-disk carrier disk shelves

Aggregates composed of disks from multi-disk carrier disk shelves must conform to some configuration requirements.

The following configuration requirements apply to aggregates composed of disks from multi-disk carrier disk shelves:

- The RAID type must be RAID-DP.
- The format must be 64-bit.

- All HDDs in the aggregate must be the same Data ONTAP disk type.
  The aggregate can be a Flash Pool.
- If the aggregate is mirrored, both plexes must have the same Data ONTAP disk type (or types, in the case of a Flash Pool).
- The aggregate cannot be a traditional volume.

Considerations for using disks from a multi-disk carrier disk shelf in an aggregate

Observing the requirements and best practices for using disks from a multi-disk carrier disk shelf in an aggregate enables you to maximize storage redundancy and minimize the impact of disk failures.

Disks in multi-disk carriers always have the Data ONTAP disk type of MSATA. MSATA disks cannot be mixed with HDDs from a single-carrier disk shelf in the same aggregate.

The following disk layout requirements apply when you are creating or increasing the size of an aggregate composed of MSATA disks:

- Data ONTAP prevents you from putting two disks in the same carrier into the same RAID group.
- Do not put two disks in the same carrier into different pools, even if the shelf is supplying disks to both pools.
- Do not assign disks in the same carrier to different nodes.

How Data ONTAP avoids RAID impact when a multi-disk carrier must be removed

Data ONTAP takes extra steps to ensure that both disks in a carrier can be replaced without impacting any RAID group. Understanding this process helps you know what to expect when a disk from a multi-disk carrier disk shelf fails.

A multi-disk carrier disk shelf, such as the DS4486, has double the disk density of other SAS disk shelves. It accomplishes this by housing two disks per disk carrier. When two disks share the same disk carrier, they must be removed and inserted together. This means that when one of the disks in a carrier needs to be replaced, the other disk in the carrier must also be replaced, even if it was not experiencing any issues.

Removing two data or parity disks from an aggregate at the same time is undesirable, because it could leave two RAID groups degraded, or one RAID group double-degraded. To avoid this situation, Data ONTAP initiates a disk evacuation operation for the carrier mate of the failed disk, as well as the usual reconstruction to replace the failed disk. The disk evacuation operation copies the contents of the carrier mate to a disk in a different carrier so the data on that disk remains available when you remove the carrier. During the evacuation operation, the status for the disk being evacuated shows as `evacuating`.

In addition, Data ONTAP tries to create an optimal layout that avoids having two carrier mates in the same RAID group. Depending on how the other disks are laid out, achieving the optimal layout can require as many as three consecutive disk evacuation operations. Depending on the size of the disks and the storage system load, each disk evacuation operation could take several hours, so the entire swapping process could take an entire day or more.

If insufficient spares are available to support the swapping operation, Data ONTAP issues a warning and waits to perform the swap until you provide enough spares.

How to determine when it is safe to remove a multi-disk carrier

Removing a multi-disk carrier before it is safe to do so can result in one or more RAID groups becoming degraded, or possibly even a storage disruption. System Manager enables you to determine when it is safe to remove a multi-disk carrier.

When a multi-disk carrier has to be replaced, the following events must have occurred before you can remove the carrier safely:

- An AutoSupport message must have been logged indicating that the carrier is ready to be removed.
- An EMS message must have been logged indicating that the carrier is ready to be removed.
- The state of both disks in the carrier must be displayed as `broken` in the Disks window.

  You must remove the disks only after the carrier mate of a failed disk is evacuated. You can click Details to view the disk evacuation status in the Properties tab of the Disks window.
- The fault LED (amber) on the carrier must be lit continuously indicating that it is ready for removal.
- The activity LED (green) must be turned off indicating there is no disk activity.
- The shelf digital display only shows the shelf ID number.

Attention: You cannot reuse the carrier mate of a failed disk. When you remove a multi-disk carrier that contains a failed disk, you must replace with a new carrier.

For more information about how to determine when it is safe to remove a faulted disk carrier, see the hardware guide for your disk shelf model on the NetApp Support Site.

Related information

NetApp Support Site: support.netapp.com

Configuring disks

Creating an aggregate from spare disks

You can use the Create Aggregate dialog box to create an aggregate from selected spare disks and provide disk space to one or more FlexVol volumes.

Before you begin

- Depending on the RAID type, an appropriate number of compatible spare disks must be available.
- For a SnapLock aggregate, the SnapLock Compliance license, or the SnapLock Enterprise license, or both must be installed on the storage system.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Disks**.
3. Select the appropriate number of compatible spare disks and click **Create Aggregate**.

   Depending on the RAID type, you must retain a minimum of three spare disks for RAID-DP and two spare disks for RAID4.
4. Specify a name for the aggregate.
5. Select the RAID type.

   The options that are enabled depend on the number of disks that are selected.

   RAID0 option is available only if the storage system is a V-Series system or a Data ONTAP-v storage.
6. Click **Create**.
7. Verify that the aggregate you created is displayed in the **Aggregates** window.

Related references

Disks window on page 144

Managing disks

Adding disks to an aggregate

You can use the Add Disks to Aggregate dialog box to add spare disks to an existing aggregate to increase its size and provide more storage space to its contained FlexVol volumes.

About this task

You can add the following disks to an aggregate:

- Disks of the same effective disk type that are contained in an aggregate
- SSD disks, if the aggregate already contains other SSD disks

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Disks**.
3. Select one or more spare disks that you want to add to the aggregate and click **Add to Aggregate**.
4. Select the aggregate to which you want to add the spare disks and click **Add**.
5. Verify that the Aggregate column displays the aggregate name to which you added the disk.

Related references

Disks window on page 144

Monitoring disks

Viewing disk information

You can use the Disks window to view the name, size, and container of a disk.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Storage** > **Disks**.
3. Select the disk that you want to view information about from the displayed list of disks.
4. Review the disk details.

Window descriptions

Disks window

You can use the Disks window to manage the spare disks in your storage system and to create new aggregates or increase the size of an existing aggregate using these disks.

- *Command buttons* on page 144
- *Disk list* on page 144
- *Details area* on page 145

Command buttons

Create Aggregate	Opens the Create Aggregate dialog box, which enables you to create a new aggregate using spare disks. **Note:** This button is enabled only if the user selects at least two spare disks.
Add to Aggregate	Opens the Add Disks to Aggregate dialog box, which enables you to add spare disks to an existing aggregate. **Note:** This button is enabled only if the user selects at least one spare disk.
Refresh	Updates the information in the window.

Disk list

Name	Displays the name of the disk.
State	Displays the state of the disk.

Type	Displays the type of the disk.
Firmware Version	Displays the firmware version of the disk.
RPM	Displays the speed of the disk drive.
Effective Size	Displays the usable space available on the disk.
Physical Space	Displays the total physical space of the disk.
Aggregate	Displays the aggregate to which this disk belongs.
Shelf	Displays the shelf in which the physical disks are located.
Bay	Displays the bay within the shelf for the physical disk.
Pool	Displays the name of the pool to which the selected disk is assigned. A value of -NA- indicates that SyncMirror is not licensed.
Checksum	Displays the type of the checksum.

Details area

The area below the disk list displays detailed information about the selected disk, including information about the containing aggregate or volume (if applicable). The RAID state is zeroing for a spare disk that is in the process of being zeroed out.

Related tasks

Creating an aggregate from spare disks on page 142
Adding disks to an aggregate on page 143

vFiler Units

Understanding vFiler units

What vFiler units are

A *vFiler unit* is a partition of a storage system and the associated network resources. Each vFiler partition appears to the user as a separate storage system on the network and functions as a storage system.

Access to vFiler units can be restricted so that an administrator can manage and view files only on an assigned vFiler unit, not on other vFiler units that reside on the same storage system. In addition, there is no data flow between vFiler units. When using vFiler units, you can be sure that no sensitive information is exposed to other administrators or users who store data on the same storage system.

To use vFiler units you must have the MultiStore software licensed on the storage system that is hosting the vFiler units.

The default vFiler unit

When you enable MultiStore, Data ONTAP automatically creates a default vFiler unit on the hosting storage system that is named vfiler0. The vfiler0 unit owns all the resources of the storage system.

When you create vFiler units and assign resources to them, the resources are assigned from vfiler0. Therefore, vfiler0 owns all resources that are not owned by nondefault vFiler units.

The default vFiler unit exists as long as MultiStore is enabled. On a storage system with MultiStore enabled, you cannot rename or destroy vfiler0.

All information provided about the vFiler units is applicable to vfiler0, unless noted otherwise.

What an IPspace is

An IPspace defines a distinct IP address space in which vFiler units can participate. IP addresses defined for an IPspace are applicable only within that IPspace. A distinct routing table is maintained for each IPspace. No cross-IPspace traffic routing happens.

Note: IPspaces support IPv4 and IPv6 addresses on their routing domains.

Each IPspace has a unique loopback interface assigned to it. The loopback traffic of each IPspace is completely isolated from the loopback traffic on other IPspaces.

Configuring vFiler units

Creating vFiler units

You can partition the storage and network resources of a single storage system so that it appears as multiple storage systems called vFiler units. You can use the Create vFiler unit wizard to create vFiler units.

Before you begin

- The MultiStore license must be installed on the storage system.
- You need the following information:
  - Networking details
    IP address space in which the vFiler unit can participate, the IP address of the vFiler unit, and the interface to which the IP address is bound.
    The DNS and NIS domain name and server details for the vFiler unit.
  - Protocols
    Protocols allowed on the vFiler unit.
  - Administration details
    Administrator host name or IP address and password of the vFiler unit's root user.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **vFiler units**, and then click **Create**.
3. Type or select information as requested by the wizard.
4. Confirm the details and click **Finish** to complete the wizard.

Related references

vFiler units window on page 149

Deleting vFiler units

You can delete or destroy a vFiler unit and return storage resources back to the hosting storage system. On a storage system with the MultiStore license enabled, you cannot destroy vfiler0.

Before you begin

- LUNs that are mapped to the vFiler unit's storage must be unmapped.
- The vFiler unit must be stopped.

- If there are multiple vFiler units in an IPspace, routes used by other vFiler units must not be associated with the vFiler unit that you want to delete. Otherwise, deleting the vFiler unit makes the other vFiler units in the IPspace inaccessible.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **vFiler units**.
3. Select the vFiler unit that you want to delete and click **Delete**.
4. Select the confirmation check box and click **Delete**.

Related references

vFiler units window on page 149

Managing vFiler units

Editing vFiler units

You can edit the settings for a vFiler unit, such as the protocols allowed and additional paths associated with the vFiler unit.

About this task

You cannot change the settings of the default vFiler unit (vfiler0).

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **vFiler units**.
3. Select the vFiler unit that you want to edit and click **Edit**.
4. In the **Edit vFiler settings** dialog box, modify the required settings.
5. Click **Save and Close** to save your changes and close the dialog box.
6. Use the **vFiler units** window to verify the changes that you made to the selected vFiler unit.

Related references

vFiler units window on page 149

Starting or stopping vFiler units

You can start a vFiler unit that is in the stopped state. After a vFiler unit is started it can receive packets of data from clients. You can stop a vFiler unit to troubleshoot or destroy a vFiler unit. You can use the vFiler units window to start or stop a vFiler unit.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **vFiler units**.
3. Select the appropriate vFiler unit and click either **Start** or **Stop**, as required.

Related references

vFiler units window on page 149

Window descriptions

vFiler units window

You can use the vFiler units window to create, display, and manage information about the vFiler units.

- *Command buttons* on page 149
- *vFiler units list* on page 149
- *Details area* on page 150

Command buttons

Create	Opens the Create vFiler unit wizard to create vFiler units and attach storage objects to it.
Edit	Opens the Edit vFiler unit settings dialog box to edit settings of selected vFiler units.
Delete	Deletes the selected vFiler units.
Start	Starts the selected vFiler units to keep it in a running state so that the vFiler unit can receive packets from clients. For example, if iSCSI is licensed on the storage system, starting a vFiler unit starts iSCSI packet processing for that vFiler unit.
Stop	Stops the selected vFiler units from receiving packets from clients.
Refresh	Updates the information in the window.

vFiler units list

Name	Specifies the name of a vFiler unit.

Status	Specifies whether the vFiler unit is running.
IPspace	Specifies the IPspace used.
Allowed Protocols	Specifies the protocols that clients can use to access the vFiler units.
RSH	Specifies whether the RSH protocol is enabled. You can execute RSH commands for a vFiler unit if the RSH protocol is enabled.

Details area

The area below the vFiler units list displays detailed information about the selected vFiler unit.

Details tab	Displays the details of selected vFiler units such as root path (complete path to an existing volume or a qtree), and DNS and NIS domain names. Additionally, it also displays details about the DNS & NIS servers and administrative host.
Storage tab	Displays the storage objects managed by selected vFiler unit.
Network tab	Displays the vFiler unit's network details including IP address, netmask, and interface used.

Related tasks

Creating vFiler units on page 147
Deleting vFiler units on page 147
Editing vFiler units on page 148
Starting or stopping vFiler units on page 149

Data Protection

SnapMirror

Understanding SnapMirror technology

Using SnapMirror for Data protection

You can use SnapMirror, an optional feature of Data ONTAP, to replicate data from specified source volumes or qtrees to specified destination volumes or qtrees, respectively.

You require a separate license to use SnapMirror.

After the data is replicated to the destination storage system, you can access the data on the destination to perform the following actions:

- Provide users immediate access to mirrored data if the source goes down.
- Restore the data to the source to recover from disaster, data corruption (qtrees only), or user error.
- Archive the data to tape.
- Balance resource loads.
- Back up or distribute the data to remote sites.

System Manager cannot manage SnapMirror relationships that are configured using SnapMirror connections, vFiler units, or preferred interfaces. System Manager uses the storage system name that is specified in the SnapMirror relationship to query the storage system. The host resolution fails because the connection name, vFiler unit name, or the preferred interface name is not same as the storage system name. You must add both source and destination systems to System Manager.

How SnapMirror works

SnapMirror replicates data from a source volume or qtree to a partner destination volume or qtree, respectively, by using Snapshot copies. Before using SnapMirror to copy data, you need to establish a relationship between the source and the destination.

The SnapMirror feature performs the following operations:

1. Creates a Snapshot copy of the data on the source volume.
2. Copies it to the destination, which can be a read-only volume or qtree.
3. Updates the destination to reflect incremental changes on the source, as per the schedule you specify.

The result of this process is an online, read-only volume or qtree that contains the same data as the source at the time of the most recent update.

Each of the following replication methods consists of a pair of operations, one operation each at the source storage system and the destination storage system:

- Volume SnapMirror replication
- Qtree SnapMirror replication

If a storage system is the source for one replication and the destination for another replication, it uses two replication operations. Similarly, if a storage system is the source as well as the destination for the same replication, it uses two replication operations.

Applications of SnapMirror

SnapMirror is used to replicate data. Its qualities make SnapMirror useful in several scenarios, including disaster recovery, data backup, and data restoration.

You can copy or use the data stored on a SnapMirror destination. The additional advantages of SnapMirror make it useful in data retrieval situations such as those described in the following table:

Situation	How to use SnapMirror
Disaster recovery: You want to provide immediate access to data after a disaster has made a qtree, volume, or system unavailable.	You can make the destination writable so clients can use the same data that was on the source volume the last time data was copied.
Disaster recovery testing: You want to test the recovery of data and restoration of services in the event of a disaster.	You can use FlexClone technology on the SnapMirror destination, and test for disaster recovery, without stopping or pausing other replication operations.
Data restoration: You want to restore lost data on a qtree or volume source from its mirrored qtree or volume SnapMirror partner.	You can temporarily reverse the roles for the source and destination qtrees or volumes and copy the mirrored information back to its source.
Application testing: You want to use an application on a database, but you want to test it on a copy of the database in case the application damages the data.	You can make a copy of the database to be used in the application testing to ensure that the data on the source cannot be lost.

Deployment of SnapMirror

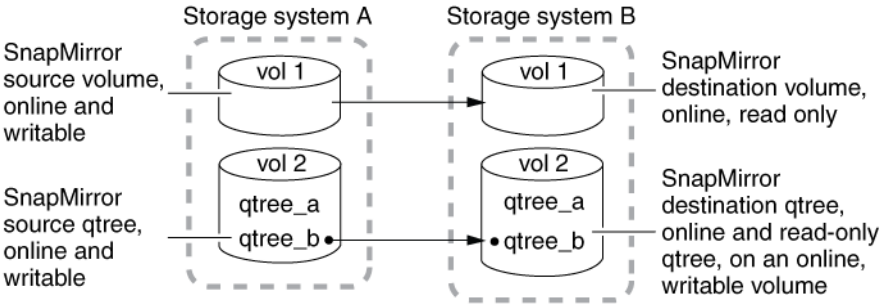
A basic deployment of SnapMirror consists of source volumes and qtrees, and destination volumes and qtrees.

Source volumes or qtrees: In a SnapMirror configuration, source volumes and qtrees are the data objects that need to be replicated. Normally, users of storage can access and write to source volumes and qtrees.

Destination volumes or qtrees: In a SnapMirror configuration, destination volumes and qtrees are data objects to which the source volumes and qtrees are replicated. The destination volumes and qtrees are read-only, and usually placed on a separate system than the source. The destination volumes and qtrees can be accessed by users in case the source becomes unavailable. The administrator can use SnapMirror commands to make the replicated data at the destination accessible and writable.

Note: Destination volumes have to be writable when using qtree SnapMirror for replication.

The following illustration depicts a basic SnapMirror deployment:



Configuring SnapMirror relationships

Adding remote access

You can use the Remote Access dialog box to provide the destination system access to the source volume or qtree in order to mirror a volume or qtree from the source system to the remote destination system.

Before you begin

The `snapmirror.access` option must be set to `legacy`.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection** > **SnapMirror**, and then click **Remote Access**.
3. In the **Remote Access** dialog box, click **Add**.
4. Type the IP address or host name of the remote system.
5. Browse to select the source volume or qtree to be accessed by the remote system, click **Select**, and then click **OK**.

   You can allow access by the destination system to all the volumes on the source system.
6. Click **OK**.

Related references

SnapMirror window on page 161

Creating SnapMirror relationships

You can use mirroring technology to replicate data from a source volume or qtree to a destination volume or qtree at regular intervals or on demand.

Before you begin

- The SnapMirror license must be enabled on both the source and the destination storage systems.
- For SnapMirror volume replication, the capacity of the destination volume must be greater than or equal to the capacity of the source volume.
- The SnapMirror destination volume must not be the root volume of a storage system.
- The destination system must be running a Data ONTAP version from the same release family or later than that of the source system.
- Both source and destination systems must be managed by System Manager.

  Note: During disaster recovery, when the source system is down, the Quiesce and Break operations can be performed by System Manager from the destination system.
- The destination storage system must have access to the source storage system.

About this task

The storage system can either be the source system or the destination system for the new SnapMirror relationship that you create. You can create a volume SnapMirror relationship by using a FlexClone volume or its parent as the source volume. However, you cannot create a volume SnapMirror relationship by using either a FlexClone volume or its parent as the destination volume.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapMirror**.
3. Click **Create**.
4. Type or select information as requested by the wizard.
5. Confirm the details and click **Finish** to complete the wizard.
6. Verify that the SnapMirror relationship you created is included in the list of SnapMirror relationships in the **SnapMirror** window.

   If you choose not to initialize the SnapMirror relationship during creation, then it is not displayed in the SnapMirror window of the source system. You have to initialize it from the SnapMirror window of the destination system.

   Note: System Manager does not record the storage system's fully qualified domain name (FQDN) in the `snapmirror.conf` file.

Related references

SnapMirror window on page 161

Deleting SnapMirror relationships

You can delete a SnapMirror relationship and permanently end the relationship between the source and destination volumes or qtrees. Deletion of a SnapMirror relationship allows the source to delete the Snapshot copies associated with that destination.

Before you begin

The SnapMirror relationship between the source and destination must be broken.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapMirror**.
3. Select the SnapMirror relationship that you want to delete and click **Delete**.
4. Select the confirmation check box and click **Delete**.

Related references

SnapMirror window on page 161

Deleting remote access

Remote access allows a SnapMirror destination to copy from the SnapMirror source. You can delete remote access provided to a SnapMirror destination from the Remote Access dialog box.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapMirror**, and then click **Remote Access**.
3. In the **Remote Access** dialog box, select the volume or qtree that is accessed by a remote system and click **Delete**.
4. Select the confirmation check box and click **Remove**.

Related references

SnapMirror window on page 161

Managing SnapMirror relationships

Editing SnapMirror relationship properties

You can use the Edit SnapMirror Relationship dialog box to edit the schedule for data transfer and the data transfer rate for an asynchronous SnapMirror relationship.

About this task

You can use System Manager to edit a SnapMirror relationship. However, you should edit the SnapMirror relationship by updating the `/etc/snapmirror.conf` file in the following scenarios:

- If the SnapMirror relationship is a synchronous or semi-synchronous SnapMirror relationship.
- If any option other than the data transfer rate is specified.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapMirror**.
3. Select the SnapMirror relationship that you want to edit and click **Edit**.
4. In the **Edit SnapMirror Relationship** dialog box, modify the properties as required.
5. Click **Save and Close** to save your changes and close the dialog box.

Related references

SnapMirror window on page 161

Initializing SnapMirror destinations

When you start a SnapMirror relationship for the first time, you have to initialize the relationship. Initializing a relationship consists of a complete baseline transfer of data from a source volume or qtree to the destination. You can use the SnapMirror window to initialize a SnapMirror relationship.

Before you begin

For a volume SnapMirror relationship, the destination volume must be in a restricted state.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapMirror**.
3. Select the SnapMirror relationship that you want to initialize.
4. Click **Operations > Initialize**.
5. Click **Initialize**.

Related references

SnapMirror window on page 161

Updating SnapMirror relationships

You can use the SnapMirror window to initiate an unscheduled SnapMirror update of the destination. You might have to perform a manual update to prevent data loss due to an upcoming power outage, scheduled maintenance, or data migration.

Before you begin

The SnapMirror relationship must be in snapmirrored state.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapMirror**.
3. Select the SnapMirror relationship that you want to update.
4. Click **Operations > Update**.
5. Click **Update**.

Result

The destination volume or qtree is now updated with data from the source volume or qtree.

Related references

SnapMirror window on page 161

Quiescing SnapMirror destinations

A SnapMirror destination is quiesced to stabilize the destination before taking a Snapshot copy. This operation enables active SnapMirror transfers to finish and disables future transfers for the mirroring relationship. You can use the SnapMirror window to quiesce a SnapMirror relationship.

About this task

You can quiesce only SnapMirror relationships that are in snapmirrored state.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapMirror**.
3. Select the SnapMirror relationship that you want to quiesce.
4. Click **Operations > Quiesce**.

5. Select the confirmation check box and click **Quiesce**.

Related references

SnapMirror window on page 161

Resuming SnapMirror relationships

You can use the SnapMirror window to resume a quiesced SnapMirror relationship for FlexVol volumes. When you resume the relationship, normal data transfer to the SnapMirror destination is resumed and all SnapMirror activities are restarted.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapMirror**.
3. Select the SnapMirror relationship that you want to resume.
4. Click **Operations > Resume**.
5. Click **Resume**.

Result

Data transfer to the SnapMirror destination resumes for the selected SnapMirror relationship.

Related references

SnapMirror window on page 161

Breaking SnapMirror relationships

If the SnapMirror source becomes unavailable or if you want to make the SnapMirror destination a read/write volume or qtree, you can break the SnapMirror relationship by using the SnapMirror window.

Before you begin

The SnapMirror destination must be quiesced.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapMirror**.
3. Select the SnapMirror relationship that you want to break.
4. Click **Operations > Break**.
5. Select the confirmation check box and click **Break**.

Related references

SnapMirror window on page 161

Resynchronizing SnapMirror relationships

You can use the SnapMirror window to reestablish a SnapMirror relationship that was broken. You can perform a resynchronization operation to recover from a disaster that disabled the source volume or qtree.

About this task

When you perform a resynchronization operation, the contents on the SnapMirror destination are overwritten by the contents on the source. The resynchronization operation can cause loss of data written to the destination volume after the base Snapshot copy was created.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapMirror**.
3. Select the SnapMirror relationship that you want to resynchronize.
4. Click **Operations > Resync**.
5. Select the confirmation check box and click **Resync**.

Related references

SnapMirror window on page 161

Reverse resynchronizing SnapMirror relationships

You can use the SnapMirror window to reestablish a SnapMirror relationship that was broken. In a reverse resynchronization operation, you reverse the functions of the source and destination, and the source volume or qtree is converted to a copy of the original destination volume or qtree.

About this task

When you perform reverse resynchronization, the contents on the SnapMirror source are overwritten by the contents on the destination. This operation can cause data loss on the source.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapMirror**.
3. Select the SnapMirror relationship that you want to reverse resynchronize.
4. Click **Operations > Reverse Resync**.

5. Select the confirmation check box and click **Reverse Resync**.

Related references

SnapMirror window on page 161

Aborting a SnapMirror transfer

You can abort a volume or qtree replication operation before the data transfer is complete. You can abort a scheduled update, a manual update, or an initial SnapMirror data transfer.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapMirror**.
3. Select the SnapMirror relationship.
4. Click **Operations > Abort**.
5. Select the confirmation check box and click **Abort**.

Related references

SnapMirror window on page 161

Editing remote access

You can edit the remote access provided to a remote destination system from the Remote Access dialog box. You can provide access to another volume or qtree of the source system.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapMirror**, and then click **Remote Access**.
3. In the **Remote Access** dialog box, select the remote destination system whose remote access you want to edit and click **Edit**.
4. Select the volume or qtree to be accessed by the remote system and click **OK**.
5. Click **OK**.

Related references

SnapMirror window on page 161

Window descriptions

SnapMirror window

You can use the SnapMirror window to create, display, and manage SnapMirror relationships.

- *Command buttons* on page 161
- *SnapMirror relationships list* on page 161
- *Details area* on page 162

Command buttons

Create	Opens the SnapMirror Relationship Create wizard, which enables you to create a SnapMirror relationship from a source volume or a qtree.
Edit	Opens the Edit SnapMirror Relationship dialog box, which enables you to edit the schedule and the data transfer rate of a SnapMirror relationship.
Delete	Deletes the SnapMirror relationship.
Operations	Displays the operations that can be performed on a SnapMirror relationship.
Remote Access	Opens the Remote Access dialog box, which enables you to manage the access to source volumes or qtrees from remote destination systems.
Refresh	Updates the information in the window.

SnapMirror relationships list

Source	Specifies the volume or qtree from which data is mirrored in a SnapMirror relationship.
Destination	Specifies the volume or qtree to which data is mirrored in a SnapMirror relationship.
SnapMirror Type	Specifies the type of the SnapMirror relationship.
State	Specifies the state of the SnapMirror relationship as source, or snapmirrored, or broken-off.
Status	Specifies the SnapMirror relationship status as idle or transferring.
Transfer Status	Specifies the status of the data transfer.
Lag Time	Specifies the difference between the current time and the timestamp of the Snapshot copy that was successfully transferred last to the destination storage system.

Details area

The details area includes details about the SnapMirror relationship, such as data transfer rates, the status, and the schedule of the relationship.

Related tasks

Adding remote access on page 153
Creating SnapMirror relationships on page 154
Deleting SnapMirror relationships on page 155
Deleting remote access on page 155
Editing SnapMirror relationship properties on page 156
Initializing SnapMirror destinations on page 156
Updating SnapMirror relationships on page 157
Quiescing SnapMirror destinations on page 157
Resuming SnapMirror relationships on page 158
Breaking SnapMirror relationships on page 158
Resynchronizing SnapMirror relationships on page 159
Reverse resynchronizing SnapMirror relationships on page 159
Aborting a SnapMirror transfer on page 160
Editing remote access on page 160

What lag time is

Lag time is the amount of time by which the destination system lags behind the source system.

The lag time is the difference between the current time and the timestamp of the Snapshot copy that was last successfully transferred to the destination system. The lag time will always be at least as much as the duration of the last successful transfer, unless the clocks on the source and destination systems are not synchronized. The lag time can be negative if the time zone of the destination system is behind the time zone of the source system.

SnapVault

Understanding SnapVault relationships

What SnapVault is

SnapVault is a disk-based storage backup feature of Data ONTAP. SnapVault enables data stored on multiple systems to be backed up to a central, secondary system quickly and efficiently as read-only Snapshot copies.

In the event of data loss or corruption on a system, backed-up data can be restored from the SnapVault secondary system with less downtime and uncertainty than is associated with conventional tape backup and restore operations.

The following terms are used to describe the SnapVault feature:

- Primary system—a system whose data is to be backed up
- Secondary system—a system to which data is backed up
- Primary system qtree—a qtree on a primary system whose data is backed up to a secondary qtree on a secondary system
- Secondary system qtree—a qtree on a secondary system to which data from a primary qtree on a primary system is backed up
- SnapVault relationship—the backup relationship between a qtree on a primary system and its corresponding secondary system qtree
- SnapVault Snapshot copy—the backup images that SnapVault creates at intervals on its primary and secondary systems

  SnapVault Snapshot copies capture the state of primary qtree data on each primary system. This data is transferred to secondary qtrees on the SnapVault secondary system. The secondary system creates and maintains versions of Snapshot copies of the combined data for long-term storage and possible restore operations.
- SnapVault Snapshot basename—a name that you assign to a set of SnapVault Snapshot copies. As incremental Snapshot copies for a set are taken and stored on both the primary and secondary systems, the system appends a number (0, 1, 2, 3, and so on) to the basenames to track the most recent and earlier Snapshot updates.
- SnapVault baseline transfer—an initial complete backup of a primary storage qtree to a corresponding qtree on the secondary system
- SnapVault incremental transfer—a follow-up backup to the secondary system that contains only the changes to the primary storage data between the current and last transfer actions

Data protection using SnapVault

SnapVault enables you to protect data on primary systems by maintaining a number of read-only versions of that data on a SnapVault secondary system and the SnapVault primary system.

The data that is backed up and restored through SnapVault depends on the primary system. On systems running Data ONTAP, a qtree is the basic unit of SnapVault backup and restore. SnapVault backs up specified qtrees on the primary system to associated qtrees on the SnapVault secondary system. If necessary, data is restored from the secondary qtrees to their associated primary qtrees. You can back up the qtrees from multiple primary systems to the associated qtrees on a single SnapVault secondary volume.

What basic SnapVault deployment is

The basic SnapVault backup system deployment consists of a primary system and a secondary system.

Primary storage systems

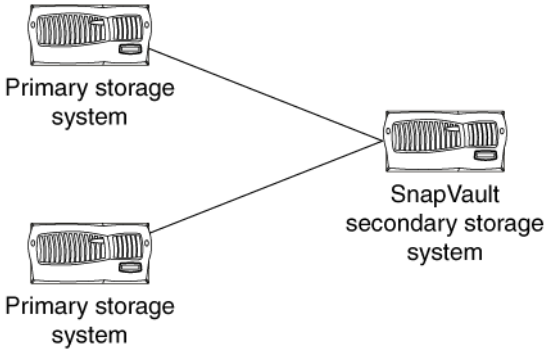
Primary systems are the platforms that run Data ONTAP to be backed up.

- On primary systems, SnapVault backs up primary qtree data, non-qtree data, and entire volumes, to qtree locations on the SnapVault secondary systems.

Secondary storage system

The SnapVault secondary system is the central disk-based unit that receives and stores backup data from the system as Snapshot copies. Any system can be configured as a SnapVault secondary system; however, the recommended hardware platform is a NearStore system.

The following figure shows a basic SnapVault deployment:



Configuring SnapVault relationships

Configuring remote access for all storage systems

You can use the Remote Access dialog box to configure storage systems to grant or deny data access to all remote storage systems in a SnapVault relationship.

About this task

You must configure remote access on both the source system and destination system to establish a successful SnapVault relationship.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapVault**.
3. Click the **Relationships** tab, and then select **Remote Access**.
4. In the **Remote Access** dialog box, select the appropriate option to grant or deny access to all remote storage systems.
5. Click **Ok**.

Configuring remote access for a storage system

You can use the Remote Access dialog box to configure storage systems to grant or deny data access to specific remote systems, in addition to managing data access between the primary and secondary systems in a SnapVault relationship.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapVault**.
3. Click the **Relationships** tab, and then select **Remote Access**.
4. In the **Remote Access** dialog box, select the option for data access to specific remote systems and perform the appropriate action:

If you want to...	Do the following...
Provide data access to specific remote systems	**a.** Click **Add**. **b.** Select the remote system from the drop-down list, and click **Save**.

If you want to...	Do the following...
Deny data access to specific remote systems	Select the remote system that you want to delete, and then click **Delete**.

5. Click **Ok**.

Creating SnapVault relationships

You can use the SnapVault Relationship Create Wizard to back up data from a qtree on a primary system to its corresponding secondary system qtree.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapVault**.
3. In the **Relationships** tab, click **Create**.
4. Type or select information as prompted by the wizard.
5. Confirm the details and click **Finish** to complete the wizard.

Result

The SnapVault relationship is created.

Creating SnapVault schedules to manage data backup

You can create a SnapVault schedule on a volume for regular backup of data by using the Create Schedule dialog box. You can also create a transfer schedule on a destination volume to transfer data from a source volume or qtree to a destination qtree, and create a backup schedule on a source volume or qtree to create Snapshot copies.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapVault**.
3. In the **Schedules** tab, click **Create**.

   The Create Schedule dialog box is displayed.
4. In the **General** tab, specify the following schedule types in addition to the schedule details:

If you want to...	Specify the following schedule type...
Transfer data from a source volume or qtree to a destination qtree	**Transfer**
Create a backup schedule on a volume	**Backup**

5. In the **Snapshot Retention** tab, specify details such as the maximum Snapshot copies to retain, attempts to create Snapshot copies, warning count, whether old Snapshot copies should be preserved, and the SnapLock retention period.
6. Click **Create**.

Result

The SnapVault schedule is displayed in the SnapVault window.

Deleting SnapVault schedules

You can delete a SnapVault schedule on a volume that is no longer necessary by using the **Schedules** tab.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapVault**.
3. Click the **Schedules** tab.
4. Select the SnapVault schedule that you want to delete from the list of SnapVault schedules.
5. Click **Delete**.
6. In the **Delete Schedule** dialog box, click **Delete** to confirm.

Result

The selected SnapVault schedule is deleted.

Managing SnapVault relationships

Starting SnapVault relationships

You can initiate the baseline data transfer in a SnapVault relationship from a source qtree to a destination qtree by using the Relationships tab.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapVault**.
3. In the **Relationships** tab, select the SnapVault relationship in which you want to start or reinitialize the data transfer.
4. Click **Operations > Start**.
5. In the **Start** dialog box, click **Start**.

Result

The data transfer is initiated from the source volume or qtree to the destination qtree for the selected SnapVault relationship.

Updating SnapVault relationships

You can use the Relationships tab to update the destination qtree with data from the latest Snapshot copies, or older Snapshot copies from the source volume or qtree, in a SnapVault relationship.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapVault**.
3. Click the **Relationships** tab.
4. Select the SnapVault relationship that you want to update.
5. Click **Operations > Update**.
6. In the **Update** dialog box, specify the required details, and then click **Update**.
7. In the prompt that is displayed, click **Update**.

Result

The destination qtree is updated with the latest data from the source volume or qtree.

Deleting SnapVault relationships

You can use the Relationships tab to stop data transfers from source qtrees to destination qtrees, delete destination qtrees, and delete the SnapVault relationship.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapVault**.
3. Click the **Relationships** tab.
4. Select the SnapVault relationship in which you want to stop the data transfer.
5. Click **Operations > Stop**.
6. In the **Stop** dialog box, select the confirmation check box, and then click **Stop**.

Result

The data transfer to the selected destination qtree is stopped, and it is deleted along with the SnapVault relationship.

Releasing SnapVault relationships

You can use the Relationships tab to release a destination qtree from a SnapVault relationship if you no longer want to back up data to the qtree.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapVault**.
3. Click the **Relationships** tab.
4. Select the SnapVault relationship that you want to release.
5. Click **Operations > Release**.
6. In the **Release** dialog box, select the confirmation check box, and then click **Release**.

Result

The destination qtree is released from the selected SnapVault relationship.

Resynchronizing SnapVault relationships

You can use the Relationships tab to restore a broken SnapVault relationship and restart data backups to the destination qtree from the restored qtree.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapVault**.
3. Click the **Relationships** tab.
4. Select the SnapVault relationship for which you want to restart data backup.
5. Click **Operations > Resync**.
6. In the **Resync** dialog box, select the confirmation check box, and then click **Resync**.

Result

The data backup is restarted for the selected SnapVault relationship.

Restoring data in SnapVault relationships

You can use the Relationships tab to restore data from a destination qtree to an existing qtree or a new qtree (source) in a SnapVault relationship. You can also restore data from a specific Snapshot copy or the latest Snapshot copy from the destination volume.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapVault**.
3. Click the **Relationships** tab.
4. Select the SnapVault relationship that you want to restore.
5. Click **Operations > Restore**.
6. In the **Restore** dialog box, specify the required details, and then click **Restore**.
7. In the prompt that is displayed, select the confirmation check box, and then click **Restore**.

Result

The data from the destination qtree is restored to the source qtree in the selected SnapVault relationship.

Aborting SnapVault relationships

You can stop a data transfer that is in progress in a SnapVault relationship by using the Relationships tab.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapVault**.
3. In the **Relationships** tab, select the SnapVault relationship in which you want to stop active data transfer.
4. Click **Operations > Abort**.
5. In the **Abort** dialog box, select the confirmation check box, and then click **Abort**.

Result

The ongoing data transfer in the selected SnapVault relationship is stopped.

Deleting check points in SnapVault relationships

You can stop a data transfer that is in progress in a SnapVault relationship as well as delete the check points by using the Relationships tab.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapVault**.
3. In the **Relationships** tab, select the SnapVault relationship in which you want to stop active data transfer and delete the check points.
4. Click **Operations > Hard Abort**.
5. In the **Hard Abort** dialog box, select the confirmation check box, and then click **Hard Abort**.

Result

The check points are deleted and active data transfer in the selected SnapVault relationship is stopped.

Editing SnapVault relationship properties

You can use the Edit SnapVault Relationship dialog box to change the data transfer rate, number of attempts, and source path of a SnapVault relationship.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapVault**.
3. Click the **Relationships** tab.
4. Select the SnapVault relationship that you want to edit from the list of SnapVault relationships and click **Edit**.
5. In the **Edit SnapVault Relationship** dialog box, specify the required values, and then click **Save and Close**.

Result

The properties of the SnapVault relationship are modified.

Viewing SnapVault schedules

You can use the Schedules tab to view the list of SnapVault schedules, in a tree or list format, to sort the list or filter the list.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapVault**.
3. Click the **Schedules** tab.
4. Depending on how you want to view the schedules, perform the appropriate action:

If you want to view the schedules in the...	Then...
Tree format	Click **View as: Tree**.
List format	Click **View as: List**.

Result

The SnapVault schedules are displayed in the chosen format.

Editing SnapVault schedules

You can edit a SnapVault schedule on a volume from the Edit Schedule dialog box, to modify configuration settings such as type of schedule and the schedule time line.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Data Protection > SnapVault**.
3. Click the **Schedules** tab.
4. Select the SnapVault schedule that you want to edit from the list of SnapVault schedules and click **Edit**.

   The Edit Schedule dialog box is displayed.
5. In the **General** tab, modify the schedule details and the schedule type.
6. In the **Snapshot Retention** tab, modify the maximum Snapshot copies to retain, attempts to create Snapshot copies, warning count, whether old snapshot copies should be preserved, and the SnapLock retention period as available.
7. Click **Save and Close** to save your changes and close the dialog box.

Result

The properties of the SnapVault schedule are modified.

Window descriptions

SnapVault Window

You can use the SnapVault window to create, display, and manage SnapVault relationships.

- *Relationship tab* on page 173
- *Schedules tab* on page 174

Relationships tab

Command buttons

You can use the following command buttons from the Relationships tab:

- **Create**
  Opens the SnapVault Relationship Create wizard, which enables you to create a SnapVault relationship between a source qtree and a destination qtree.
- **Edit**
  Opens the Edit SnapVault Relationship dialog box, which enables you to edit the data transfer rate, attempts, and source path of a SnapVault relationship.
- **Remote Access**
  Opens the Remote Access dialog box, which enables you to manage access to source volumes or qtrees from remote destination systems, and access to a destination qtree from a remote source system.
- **Refresh**
  Updates the information in the window.

Operations

Displays the operations that can be performed on a SnapVault relationship.

Relationships list

You can view the following information from the Relationships tab:

- Source
  Specifies the source volume or qtree from which data is backed up in a SnapVault relationship.
- Destination
  Specifies the destination qtree to which data is backed up in a SnapVault relationship.
- State
  Specifies the SnapVault state such as snapvaulted, unknown, broken-off, or uninitialized.
- Status
  Specifies the SnapVault relationship status such as idle or transferring.
- Transfer Status

Specifies the status of the data transfer.

- Lag Time
  Specifies the difference between the current time and the timestamp of the Snapshot copy that was last successfully transferred to the destination storage system. It indicates the time difference between the data that is currently on the source system and the latest data stored on the destination system.
- Type
  Specifies the type of SnapVault relationship.
- Base Snapshot copy
  Specifies the name of the base Snapshot copy.

Relationships areas

You can view the following information from the Relationships tab:

- Details area
  The details area includes details about the SnapVault relationship, such as the data transfer rates, status, size, and schedule of the relationship.
- Associated Schedules area
  The associated schedules area displays the SnapVault schedules on the specified destination volume or source volume.
- Snapshot copies area
  The Snapshot copies area displays the Snapshot copies that are created by the SnapVault relationship on the specified destination volume or source volume.

Schedules tab

Command buttons

You can use the following command buttons from the Schedules tab:

- **Create**
  Opens the Create Schedule dialog box, which enables you to create a SnapVault schedule.
- **Edit**
  Opens the Edit Schedule dialog box, which enables you to edit the SnapVault schedule details and the schedule type.
- **Delete**
  Deletes the SnapVault schedule.
- **Create Snapshot copy**
  Creates a new Snapshot copy on the specified volume, based on the defined SnapVault schedule.
- **View as**
  Displays the SnapVault schedules for the volumes in a tree or list format.
- **Refresh**
  Updates the information in the window.

Schedules list You can view the following information from the Schedules tab:

- Volume
  Specifies the volume for which the SnapVault schedule is displayed.
- Schedule
  Specifies the schedule for a volume.
- Schedule Status
  Specifies the schedule status as idle, active, saving, or aborting.
- Schedule Details
  Specifies the schedule details such as day and time.
- Maximum Snapshot copies to Retain
  Specifies the maximum Snapshot copies that must be retained by the schedule.
- Relationships Count
  Specifies the number of SnapVault relationships that the volume is in.

Schedules areas You can view the following information from the Schedules tab:

- Details area
  The details area includes details about the SnapVault schedule, such as the number of Snapshot copies, schedule type, and role of the volume (whether it is a source or a destination).
- Associated Relationships area
  The associated relationships area displays the SnapVault relationships that are associated with the volume, state, and status of the SnapVault relationship.
- Snapshot copies area
  The Snapshot copies area displays the Snapshot copies that are created by SnapVault on the specified schedule.

What retention count is

Retention count specifies the number of Snapshot copies that are retained by the SnapVault schedule.

What remote access is

Remote access enables you to control data access between the primary and secondary storage systems in a SnapVault relationship. For example, a secondary storage system requires access to data in a primary system for data backup, while a primary system requires access to data in a secondary system for data restore.

Types of SnapVault schedules

Understanding the types of SnapVault schedules helps you select the appropriate schedule to transfer data from a source volume or qtree to a destination qtree.

The types of SnapVault schedules are as follows:

- Transfer schedule

  Transfer schedules are created on the destination volume and specify the schedule at which data is transferred from the source volume or qtree to the destination qtree. You can create these transfer schedules on a destination volume that has a SnapVault license.
- Backup schedule

  Backup schedules are created on the source volume and destination volume, and specify the schedule at which a new Snapshot copy is created on the source volume. The backup schedule creates a Snapshot copy of the primary data that is transferred to the destination volume. After data transfer is completed, the backup schedule on the destination volume creates a backup of the transferred data.

Types of SnapVault relationships

You must be aware of the types of SnapVault relationships in order to establish a relationship between a source volume or qtree and a destination qtree, and back up data from the source to the destination to prevent data loss.

The types of SnapVault relationships are as follows:

- Volume SnapVault relationship

  In a volume SnapVault relationship, all the data in the qtrees from a source volume is backed up to a destination qtree as folders. You cannot restore these qtrees stored as folders because each qtree must be restored separately; therefore, restoring the qtrees is not recommended.
- Qtree SnapVault relationship

  In a qtree SnapVault relationship, all the data from a source qtree is backed up to a destination qtree.
- Non-qtree SnapVault relationship

  In a non-qtree SnapVault relationship, all the data in a source volume is backed up to a destination qtree. You cannot back up data in a qtree within the source volume.

Configuration

Local Users and Groups > Users

Understanding local users

What local users and groups are

You can use local users and groups to secure and manage user accounts and groups stored locally on a storage system.

A user is an account that is authenticated on a storage system. Users can be placed into storage system groups to grant them capabilities on the storage system. When your system is first installed and CIFS is configured in Workgroup mode, a user named "administrator" is automatically created. This user login can be used to access shares with a blank password. You should change the password for this built-in account to increase security on your system.

A group is a collection of users that can be granted one or more roles. Groups can be predefined, created, or modified. When CIFS is enabled, groups act as Windows groups.

You can use local users and groups to limit the ability of users to perform certain actions by assigning them rights and permissions. A right authorizes a user to perform certain actions on a computer, such as backing up files and folders or shutting down a computer. A permission is a rule associated with an object (usually a file, folder, or printer), and it regulates which users have access to the object. You cannot use local users and groups to view local user and group accounts after a member server is promoted to a domain controller.

When you should create local user accounts

There are several reasons for creating local user accounts on your storage system.

You should create one or more local user accounts if your system configuration meets the following criteria:

- If, during setup, you configured the storage system to be a member of a Windows workgroup. In this case, the storage system must use the information in local user accounts to authenticate users.
- If your storage system is a member of a domain:
  - Local user accounts enable the storage system to authenticate users who try to connect to the storage system from an untrusted domain.
  - Local users can access the storage system when the domain controller is down or when network problems prevent your storage system from contacting the domain controller.
    For example, you can define a BUILTIN\Administrator account that you can use to access the storage system even when the storage system fails to contact the domain controller.

Note: If, during setup, you configured your storage system to use UNIX mode for authenticating users, the storage system always authenticates users using the UNIX password database.

Configuring local users

Creating local users

You can create a local user and assign that user to one or more predefined groups, giving that user the roles and capabilities associated with those groups. You can have a maximum of 96 administrative users on a storage system.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Click **Configuration** > **Local Users and Groups** > **Users**.
3. Click **Create**.
4. In the **New User** dialog box, type the login name for the new user.

   User names are not case-sensitive.
5. Type the full name of the user and a description that helps you identify this new user.
6. Type the password that the user uses to connect to the server, and then confirm the password.
7. Select the group type that best suits the access level the user requires, and then click **Add**.

Related references

Users window on page 181

Deleting local users

You can delete a local user to remove that user's access to the system.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Local Users and Groups** > **Users**.
3. Select the local user that you want to delete.
4. Click **Delete**.
5. Select the confirmation check box and click **Delete**.

Related references

Users window on page 181

Managing local users

Editing the password duration for a local user

You can modify the duration that a local user password is effective. Setting a shorter duration increases the security of the system access.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Local Users and Groups** > **Users**.
3. Select the local user whose password duration you want to modify and click **Edit**.
4. In the **General** tab, type the minimum number of days that the user must have the password before they can change it.

   This value is by default set to zero.
5. Type the maximum number of days that the user can use the password before they have to change it.
6. Click **Save and Close** to save your changes and close the dialog box.

Related references

Users window on page 181

Editing a local user's full name and description

You can modify the local user's full name and description to help you to better identify a local user. You cannot modify the user name.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Local Users and Groups** > **Users**.
3. Select the local user whose full user name and description you want to modify.
4. Click **Edit**.
5. In the **General** tab, type the new full name and description.
6. Click **Save and Close** to save your changes and close the dialog box.

Related references

Users window on page 181

Assigning a local user to a group

You can assign a user to one or more predefined groups, and provide that user the roles and capabilities associated with those groups.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Local Users and Groups > Users**.
3. Select the local user that you want to assign to a predefined group.
4. Click **Edit**.
5. Click **Member Of**.
6. Click **Add**.
7. Select the group that corresponds with the access level you want to assign to the user and click **Add**.
8. Click **OK** to save your changes.

Related references

Users window on page 181

Changing the local user's password

You can use the Set Password dialog box to change the password for a local user.

About this task

You should know the current password of the local user if you do not have the necessary permissions to reset the password.

You should not use certain special characters, such as the less than symbol (<), greater than symbol (>), ampersand (&), or forward slash (/) in the new password.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Local Users and Groups > Users**.
3. Select the local user whose password you want to change and click **Set Password**.
4. In the **Set Password** dialog box, type the current password and the new password, confirm the new password, and then click **Modify**.

Related references

Users window on page 181

Resetting the local user's password

You can use the Reset Password dialog box to change or reset the password of another user. By default, only root and members of the Administrators group have this capability.

Before you begin

You must have the necessary permissions to perform this task.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Local Users and Groups > Users**.
3. Select the local user for which you want to reset the password and click **Set Password**.
4. In the **Reset Password** dialog box, type the new password, confirm the new password, and then click **Reset**.

Related references

Users window on page 181

Window descriptions

Users window

You can use the Users window to create and modify user accounts that enable local users to access your storage system.

Command buttons

Create	Opens the New User dialog box, which enables you to create new users.
Edit	Opens the user Properties dialog box, which enables you to edit properties of the selected user.
Delete	Deletes the selected local user account.
Set Password	Displays the Reset Password dialog box, which enables you to set the password for the selected user.
Refresh	Updates the information in the window.

Users list

Name	Specifies the login name of the local user.
Full name	Specifies the full name of the local user.
Description	Provides a description of the local user account.

Related tasks

Creating local users on page 178
Deleting local users on page 178
Editing the password duration for a local user on page 179
Editing a local user's full name and description on page 179
Assigning a local user to a group on page 180
Changing the local user's password on page 180
Resetting the local user's password on page 181

Local Users and Groups > Groups

Configuring local groups

Creating user groups

You can create a group and give that group the capabilities associated with a predefined role.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Local Users and Groups** > **Groups**.
3. Click **Create**.
4. In the **Create Group** dialog box, type the name and description for your new group.
5. Select the appropriate role for your new group, and then click **Add**.

Related references

Groups window on page 184

Assigning local users to a user group

You can assign one or more users to a group, giving those users the roles and capabilities associated with the group.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Local Users and Groups** > **Groups**.
3. Select the group to which you want to add one or more users.
4. Click **Edit**.
5. In the **General** tab, click **Add**.
6. In the **Local/Domain Users** dialog box, select the user that you want to add to the group and click **Add**.

   If you have configured CIFS with Active Directory domain authentication, you can add a domain user using the following format:

   `<domain>\<user>`
7. Repeat Step 5 through Step 6 to add multiple users to the group.
8. Click **Save and Close** to save your changes and close the dialog box.

Related references

Groups window on page 184

Deleting user groups

You can delete a user group, when you no longer need it. You cannot delete a default group.

Before you begin

All users must be removed from the group.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Local Users and Groups** > **Groups**.
3. Select the group that you want to delete.
4. Click **Delete**.
5. Select the confirmation check box and click **Delete**.

Related references

Groups window on page 184

Managing local groups

Editing user group properties

You can modify the description of a group to make it easier to identify the group. You can add local users to a group or remove local users from a group. You can also edit the roles of the group.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Local Users and Groups** > **Groups**.
3. Select the group that you want to modify.
4. Click **Edit**.
5. In the **General** tab, modify the description of the group as required.
6. Add or remove users from the group.
7. In the **Roles** tab, add or remove roles.
8. Click **Save and Close** to save your changes and close the dialog box.

Related references

Groups window on page 184

Window descriptions

Groups window

You can use the Groups window to create a local group, assign local users to the group, modify the group description, and remove a group.

Command buttons

Create	Displays the Create Group dialog box, which enables you to create new groups.
Edit	Displays the user groups Properties dialog box, which enables you to edit properties of the selected group.
Delete	Deletes the selected group.
Refresh	Updates the information in the window.

Groups list

Name	The name of the group.
Description	The description of the group.

Related tasks

Creating user groups on page 182
Assigning local users to a user group on page 183
Deleting user groups on page 183
Editing user group properties on page 184

Network > DNS

Understanding DNS

How to configure DNS to maintain host information

You can maintain host information centrally using DNS. With DNS, you do not have to update the `/etc/hosts` file every time you add a new host to the network.

If you have several storage systems on your network, maintaining host information centrally saves you from updating the `/etc/hosts` file on each storage system every time you add or delete a host.

If you configure DNS later, you must take the following actions:

- Specify DNS name servers.
- Specify the DNS domain name of your storage system.
- Enable DNS on your storage system.

If you want to use primarily DNS for host-name resolution, you should specify it ahead of other methods in the hosts section of the `/etc/nsswitch.conf` file.

Correct host-name resolution depends on correctly configuring the DNS server. If you experience problems with host-name resolution or data availability, check the DNS server in addition to local networking.

How to use dynamic DNS to update host information

You can use dynamic DNS updates to prevent errors and save time when sending new or changed DNS information to the primary master DNS server for your storage system's zone. Dynamic DNS allows your storage system to automatically send information to the DNS servers as soon as the information changes on the system.

Without dynamic DNS updates, you must manually add DNS information (DNS name and IP address) to the identified DNS servers when a new system is brought online or when existing DNS

information changes. This process is slow and error-prone. During disaster recovery, manual configuration can result in a long downtime.

For example, if you want to change the IP address on interface e0 of *storagesystem1*, you can simply configure e0 with the new IP address. The storage system *storagesystem1* automatically sends its updated information to the primary master DNS server.

Note: Data ONTAP supports a maximum of 64 Dynamic Domain Name Server (DDNS) aliases.

Configuring DNS

Enabling or disabling DNS

You can use the Edit DNS Settings dialog box to enable or disable DNS on a storage system. DNS is disabled by default.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Network > DNS**.
3. Click **Edit**.
4. Either select or clear **Enable DNS**, as appropriate.
5. Click **Save and Close**.

Related references

DNS window on page 188

Adding or editing the DNS domain name

You can maintain host information centrally using DNS. You can use the Edit DNS Settings dialog box to add or modify the DNS domain name of your storage system.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Network > DNS**.
3. Click **Edit**.
4. In the DNS domain name and DNS search domains boxes, type or modify the DNS domain name and the DNS search domain name.
5. Click **Save and Close**.

Related references

DNS window on page 188

Managing DNS

Enabling or disabling dynamic DNS

You can use the Edit DNS Settings dialog box to enable or disable dynamic DNS on your storage system. Dynamic DNS is disabled by default.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Network** > **DNS**.
3. Click **Edit**, and then click **Dynamic DNS**.
4. Either select or clear **Enable dynamic DNS**, as appropriate.
5. Click **Save and Close**.

Related references

DNS window on page 188

Setting dynamic DNS updates

You can use the **Dynamic DNS** tab to specify the DNS time-to-live (TTL) value for every DNS update sent from your storage system. The TTL value defines the time for which a DNS entry is valid on the DNS server. By default, the TTL value is set to 24 hours.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Network** > **DNS**.
3. Click **Edit**.
4. In the **Dynamic DNS** tab, select the TTL values for dynamic DNS updates.
5. Click **Save and Close**.

Related references

DNS window on page 188

Window descriptions

DNS window

The DNS window enables you to view the current DNS settings for your system.

Command buttons

Edit Opens the Edit DNS Settings dialog box, which you can use to either enable or disable dynamic DNS or to add DNS domain names.

Refresh Updates the information in the window.

Related tasks

Enabling or disabling DNS on page 186
Adding or editing the DNS domain name on page 186
Enabling or disabling dynamic DNS on page 187
Setting dynamic DNS updates on page 187

Network > Network Interfaces

Understanding interfaces

Interface links for a virtual interface

The list includes only an Ethernet interface or a virtual interface (vif). Enabled interfaces, interfaces that are included in another vif, an existing VLAN interface, and a physical VLAN are not included in the list.

Related concepts

Network interface naming on page 188
Guidelines for configuring interface groups on page 190

Network interface naming

Network interface names are based on whether the interface is a physical or virtual network interface. Physical interfaces are assigned names based on the slot number of the adapter. Interface group names are user specified. VLANs are named by combining the interface name and VLAN ID.

Physical interfaces are automatically assigned names based on the slot where the network adapter is installed. Because physical interfaces are Ethernet interfaces, they are identified by a name consisting of "e," the slot number of the adapter, and the port on the adapter (if multi-port adapter). A multiport adapter has letters or numbers imprinted next to its ports.

- e<*slot_number*> if the adapter or slot has only one port
- e<*slot_number*><*port_letter*> if the adapter or slot has multiple ports

Interface group names are user specified. An interface group's name should meet the following criteria:

- It must begin with a letter.
- It must not contain any spaces.
- It must not contain more than 15 characters.
- It must not already be in use by another interface or interface group.

VLAN interface names are in the following format:

- <*physical_interface_name*>-<*vlan_ID*>
- <*ifgrp_name*>-<*vlan_ID*>

The following table lists interface types, interface name formats, and example of names that use these identifiers.

Interface type	Interface name format	Examples of names
Physical interface on a single-port adapter or slot	e<*slot_number*>	e0 e1
Physical interface on a multiple-port adapter or slot	e<*slot_number*><*port_letter*>	e0a e0b e0c e0d e1a e1b
Interface group	Any user-specified string that meets certain criteria	web_ifgrp ifgrp1
VLAN	<*physical_interface_name*>-<*vlan-ID*> or <*ifgrp_name*>-<*vlan_ID*>	e8-2 ifgrp1-3

Host names

When you run the `setup` command on a storage system for the first time, Data ONTAP creates a host name for each installed interface by appending the interface name to the host name of the storage system.

Note: The interface host names are not advertised by DDNS, but are available in the `/etc/hosts` file.

The following table shows examples of host names appended with the interface names.

Interface type	Host name
Single-port Ethernet interface in slot 0	toaster-e0
Quad-port Ethernet interface in slot 1	toaster-e1a toaster-e1b toaster-e1c toaster-e1d

Related concepts

Interface links for a virtual interface on page 188

Guidelines for configuring interface groups

Before creating and configuring interface groups, you should follow certain guidelines about the type, MTU size, speed, and media of the underlying interfaces.

The following guidelines apply when you create and configure interface groups on your storage system:

- The network interfaces that are part of an interface group should be on the same network adapter.
- You can configure a maximum of eight network interfaces in a single interface group.
- You cannot include a VLAN interface in an interface group.
- The interfaces that form an interface group must have the same Maximum Transmission Unit (MTU) size.

  If you attempt to create or add to an interface group and the member interfaces have different MTU sizes, Data ONTAP automatically modifies the MTU size to be the same. To ensure that the desired MTU size is configured, you can use the `ifconfig` command to configure the MTU size of the interface group after it is created. You need to configure the MTU size only if you are enabling jumbo frames on the interfaces.
- When an interface on a TOE NIC is in an interface group, the TOE functionality is disabled on all TOE NICs.
- You can include any interface, except the e0M management interface that is present on some storage systems.
- You should not mix interfaces of different speeds or media in the same multimode interface group.
- You should set the same flow control settings for all the underlying physical network interfaces that constitute an interface group.
- You should set the flow control settings of all the network interfaces to `none`.

Some switches might not support multimode link aggregation of ports configured for jumbo frames. For more information, see your switch vendor's documentation.

Related concepts

Interface links for a virtual interface on page 188

Network interface configuration

Configuring network interfaces involves assigning IP addresses, setting network parameters and hardware-dependent values, specifying network interfaces, and viewing your storage system's network configuration.

When you configure network interfaces, you can do any or all of the following:

- Assign an IP address to a network interface.
- Set parameters such as network mask, broadcast address, and prefix length.

  Note: If IPv6 is enabled on your storage system, you can set only the prefix length. IPv6 does not have a network mask and does not support broadcast addresses.
- Set hardware-dependent values such as media type, MTU size, and flow control.
- Specify whether the interface should be attached to a network with firewall security protection.
- Specify whether the network interface must be registered with Windows Internet Name Services (WINS), if CIFS is running and at least one WINS server has been configured.
- Specify the IP address of an interface or specify the interface name on an HA pair partner for takeover mode.

  Note: When using IPv6 in an HA pair, you can specify only the partner interface name (and not the IP address) on the HA pair for takeover mode.
- View the current configuration of a specific interface or all interfaces that exist on your storage system.

Network interfaces on your storage system

Your storage system supports physical network interfaces, such as Ethernet and Gigabit Ethernet interfaces, and virtual network interfaces, such as interface group and virtual local area network (VLAN). Each of these network interface types has its own naming convention.

Your storage system supports the following types of physical network interfaces:

- 10/100/1000 Ethernet
- Gigabit Ethernet (GbE)
- 10 Gigabit Ethernet

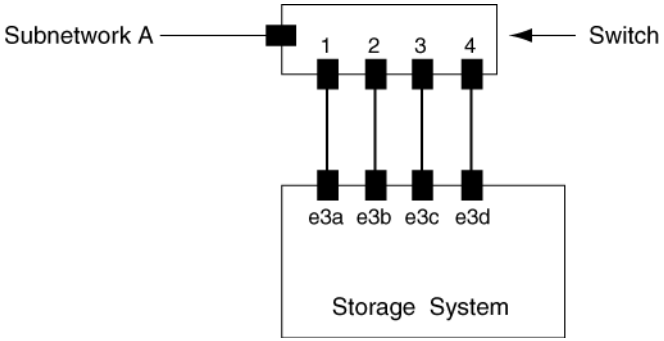
In addition, some storage system models have a physical network interface named e0M. It is a low-bandwidth interface of 100 Mbps and is used only for Data ONTAP management activities, such as running a Telnet, SSH, or RSH session.

How interface groups work in Data ONTAP

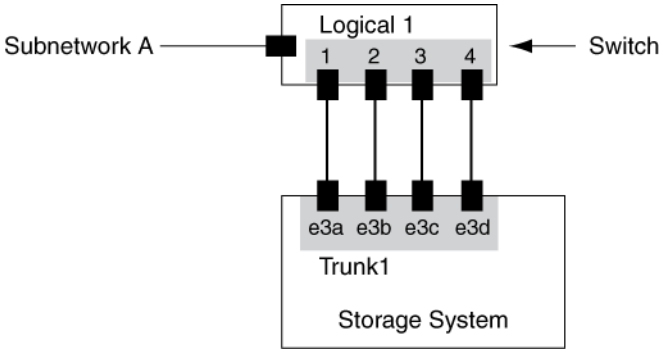
An interface group is a feature in Data ONTAP that implements link aggregation on your storage system. Interface groups provide a mechanism to group together multiple network interfaces (links)

into one logical interface (aggregate). After an interface group is created, it is indistinguishable from a physical network interface.

The following figure shows four separate network interfaces, e3a, e3b, e3c, and e3d, before they are grouped into an interface group.



The following figure shows the four network interfaces grouped into a single interface group called Trunk1.



Different vendors refer to interface groups by the following terms:

- Virtual aggregations
- Link aggregations
- Trunks
- EtherChannel

Interface groups provide several advantages over individual network interfaces:

- Higher throughput
  Multiple interfaces work as one interface.
- Fault tolerance
  If one interface in an interface group goes down, your storage system stays connected to the network by using the other interfaces.
- No single point of failure

If the physical interfaces in an interface group are connected to multiple switches and a switch goes down, your storage system stays connected to the network through the other switches.

Types of interface groups

You can create three different types of interface groups on your storage system: single-mode, static multimode, and dynamic multimode interface groups.

Each interface group provides different levels of fault tolerance. Multimode interface groups provide methods for load balancing network traffic.

Starting with Data ONTAP 7.3.1, IPv6 supports both single-mode and multimode interface groups.

Load balancing in multimode interface groups

You can ensure that all interfaces of a multimode interface group are equally utilized for outgoing traffic by using the IP address, MAC address, round-robin, or port based load-balancing methods to distribute network traffic equally over the network ports of a multimode interface group.

The load-balancing method for a multimode interface group can be specified only when the interface group is created.

IP address and MAC address load balancing

IP address and MAC address load balancing are the methods for equalizing traffic on multimode interface groups.

These load-balancing methods use a fast hashing algorithm on the source and destination addresses (IP address and MAC address). If the result of the hashing algorithm maps to an interface that is not in the UP link-state, the next active interface is used.

Note: Do not select the MAC address load-balancing method when creating interface groups on a storage system that connects directly to a router. In such a setup, for every outgoing IP frame, the destination MAC address is the MAC address of the router. As a result, only one interface of the interface group is used.

IP address load balancing works in the same way for both IPv4 and IPv6 addresses.

Standards and characteristics of Ethernet frames

Frame size and Maximum Transmission Unit (MTU) size are the two important characteristics of an Ethernet frame. The standard Ethernet (IEEE 802.3) frame size is 1,518 bytes. The MTU size specifies the maximum number of bytes of data that can be encapsulated in an Ethernet frame.

The frame size of a standard Ethernet frame (defined by RFC 894) is the sum of the Ethernet header (14 bytes), the payload (IP packet, usually 1,500 bytes), and the Frame Check Sequence (FCS) field (4 bytes). You can change the default frame size on Gigabit Ethernet network interfaces.

The MTU size specifies the maximum payload that can be encapsulated in an Ethernet frame. For example, the MTU size of a standard Ethernet frame is 1,500 bytes; this is the default for storage systems. However, a jumbo frame, with an MTU size of 9,000 bytes, can also be configured.

Flow control

Flow control enables you to manage the flow of frames between two directly connected link-partners. Flow control can reduce or eliminate dropped packets due to overrun.

To achieve flow control, you can specify a flow control option that causes packets called Pause frames to be used as needed. For example, link-partner A sends a Pause On frame to link-partner B when its receive buffers are nearly full. Link-partner B suspends transmission until it receives a Pause Off frame from link-partner A or a specified timeout threshold is reached.

Flow control options

You can use the flow control option to view and configure flow control settings. If you do not specify a flow control option when configuring a network interface, the configured flow control setting defaults to full.

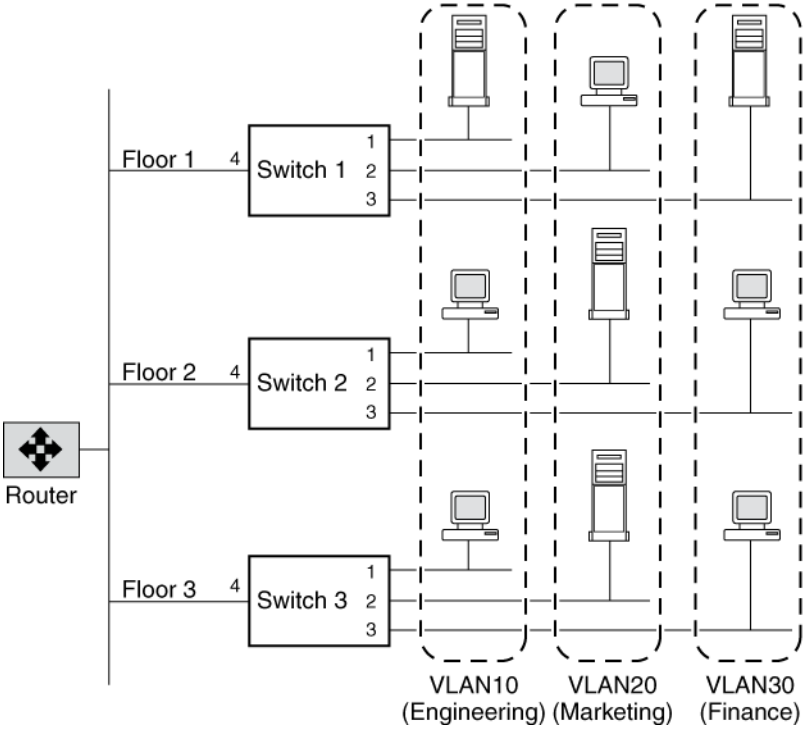
The following table describes the values you can specify for the flow control option.

Flow control value	Description
none	No flow control
receive	Able to receive flow control frames
send	Able to send flow control frames
full	Able to send and receive flow control frames

How VLANs work

Traffic from multiple VLANs can traverse a link that interconnects two switches by using VLAN tagging. A VLAN tag is a unique identifier that indicates the VLAN to which a frame belongs. A VLAN tag is included in the header of every frame sent by an end-station on a VLAN.

On receiving a tagged frame, the switch inspects the frame header and, based on the VLAN tag, identifies the VLAN. The switch then forwards the frame to the destination in the identified VLAN. If the destination MAC address is unknown, the switch limits the flooding of the frame to ports that belong to the identified VLAN.



For example, in this figure, if a member of VLAN 10 on Floor 1 sends a frame for a member of VLAN 10 on Floor 2, Switch 1 inspects the frame header for the VLAN tag (to determine the VLAN) and the destination MAC address. The destination MAC address is not known to Switch 1. Therefore, the switch forwards the frame to all other ports that belong to VLAN 10, that is, port 4 of Switch 2 and Switch 3. Similarly, Switch 2 and Switch 3 inspect the frame header. If the destination MAC address on VLAN 10 is known to either switch, that switch forwards the frame to the destination. The end-station on Floor 2 then receives the frame.

Advantages of VLANs

VLANs provide a number of advantages, such as ease of administration, confinement of broadcast domains, reduced network traffic, and enforcement of security policies.

VLANs provide the following advantages:

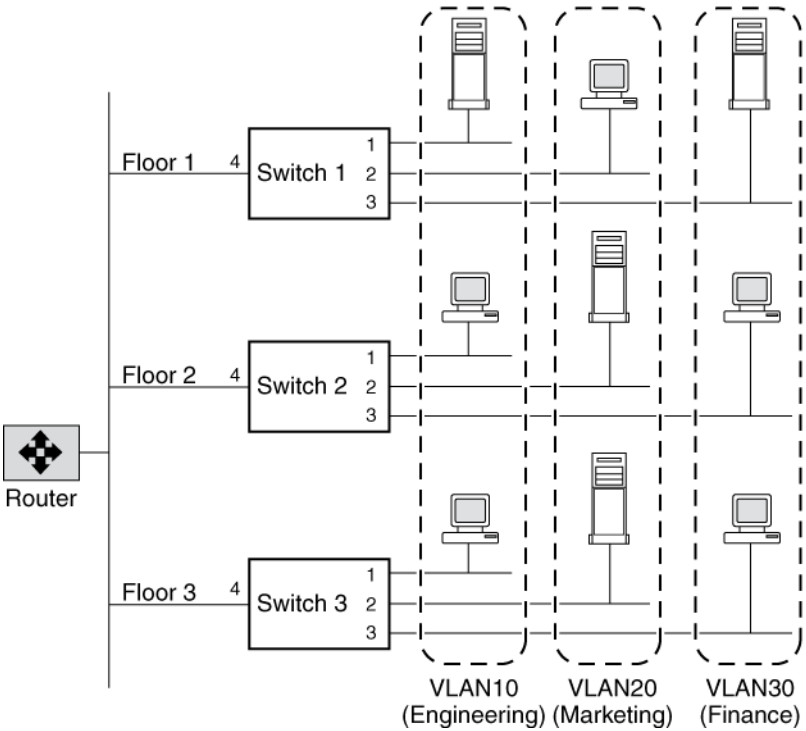
- VLANs enable logical grouping of end-stations that are physically dispersed on a network. When users on a VLAN move to a new physical location but continue to perform the same job function, the end-stations of those users do not need to be reconfigured. Similarly, if users change their job function, they need not physically move: changing the VLAN membership of the end-stations to that of the new team makes the users' end-stations local to the resources of the new team.
- VLANs reduce the need to have routers deployed on a network to contain broadcast traffic. Flooding of a packet is limited to the switch ports that belong to a VLAN.

- Confinement of broadcast domains on a network significantly reduces traffic.
  By confining the broadcast domains, end-stations on a VLAN are prevented from listening to or receiving broadcasts not intended for them. Moreover, if a router is not connected between the VLANs, the end-stations of a VLAN cannot communicate with the end-stations of the other VLANs.

VLAN tags

A VLAN tag is a unique identifier that indicates the VLAN to which a frame belongs. Generally, a VLAN tag is included in the header of every frame sent by an end-station on a VLAN.

On receiving a tagged frame, the switch inspects the frame header and, based on the VLAN tag, identifies the VLAN. The switch then forwards the frame to the destination in the identified VLAN. If the destination MAC address is unknown, the switch limits the flooding of the frame to ports that belong to the identified VLAN.



For example, in this figure, port 4 on Switch 1, Switch 2, and Switch 3 allows traffic from VLANs 10, 20, and 30. If a member of VLAN 10 on Floor 1 sends a frame for a member of VLAN 10 on Floor 2, Switch 1 inspects the frame header for the VLAN tag (to determine the VLAN) and the destination MAC address. The destination MAC address is not known to Switch 1. Therefore, the switch forwards the frame to all other ports that belong to VLAN 10, that is, port 4 of Switch 2 and Switch 3. Similarly, Switch 2 and Switch 3 inspect the frame header. If the destination MAC address

on VLAN 10 is known to either switch, that switch forwards the frame to the destination. The end-station on Floor 2 then receives the frame.

Configuring interfaces

Adding interface aliases

You can use the Add Alias dialog box to add an alias, which is an alternate IP address for an interface, when you change the IP address of an interface to a new address. You can use the alias to continue accepting packets to the old IP address. You cannot add an alias to a physical VLAN.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Network > Network Interfaces**.
3. Click **Edit**.
4. In the **Advanced** tab, click **Add**.
5. Type the IP address and a subnet mask of the alias.
6. Click **Save**, and then click **Save and Close**.

Related references

Network Interfaces window on page 201

Creating virtual interfaces

You can use the Create VIF wizard to create a virtual interface (vif), which enables you to implement link aggregation on your storage system. You can group together multiple network interfaces into one logical interface.

Before you begin

The status of the physical interface must be down.

About this task

You can only view the configuration settings for storage systems running Data ONTAP versions earlier than 7.3.3.

Note: You cannot add or remove trunks from existing interface groups.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Network > Network Interfaces**.

3. Click **Create VIF**.
4. Type or select information as prompted by the wizard.
5. Verify that the vif you created is included in the list of interfaces in the **Network Interfaces** window.

Related references

Network Interfaces window on page 201

Creating VLAN interfaces

You can create a VLAN for ease of administration, confinement of broadcast domains, reduced network traffic, and enforcement of security policies. You cannot add an interface alias to a physical VLAN, but you can add an alias to VLAN interfaces.

Before you begin

The status of the physical interface and virtual interface must be down.

To create a VLAN from a virtual interface, you must ensure that the virtual interface name does not exceed 10 characters. Otherwise, some VLAN tags might not be visible and you might not be able to create a VLAN.

About this task

You can only view the configuration settings for storage systems running Data ONTAP versions earlier than 7.3.3.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Network > Network Interfaces**.
3. Click **Create VLAN**.
4. Select a physical interface for the VLAN from the drop-down list.

   The drop-down list includes only an Ethernet interface or an interface group. Enabled interfaces, interfaces that are included in another interface group, an existing VLAN interface, or a physical VLAN are not included in the list.
5. Type a VLAN tag in the VLAN tag box and click **Add**.

   Note: You cannot add duplicate VLAN tags.
6. Click **Create**.
7. Verify that the VLAN you created is included in the list of VLANs in the **Network Interfaces** window.

Related references

Network Interfaces window on page 201

Managing interfaces

Editing interface aliases

You can use the Edit Alias dialog box to modify an interface alias. You can change the alias IP address and the subnet mask.

About this task

If you enable IPv6 from the command-line interface when the System Manager session is on, System Manager does not detect the change in IPv6 status. Therefore, you must refresh the System Manager session to enable IPv6.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Network** > **Network Interfaces**.
3. Click **Edit**.
4. In the **Advanced** tab, select the alias IP address that you want to modify and click **Edit**.
5. Change the IP address or the subnet mask of the alias.
6. Click **Save**, and then click **Save and Close**.

Related references

Network Interfaces window on page 201

Editing virtual interfaces

You can use the Edit Network Interface dialog box to modify interface parameters, such as the IP address, network mask, and MTU size.

About this task

You can only view the configuration settings for storage systems running Data ONTAP versions earlier than 7.3.3.

Note: If the Network Configuration Checker generates a false alarm due to some misconfiguration, then you cannot modify the configuration settings as the editing capability is disabled.

If you enable IPv6 from the command-line interface when the System Manager application is running, System Manager does not detect the change in IPv6 status. Therefore, you must refresh the System Manager session to enable IPv6.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Network > Network Interfaces**.
3. Select the virtual interface that you want to modify from the network interface list and click **Edit**.
4. Click the appropriate tab to display the properties or settings that you want to change.
5. Make the necessary changes.
6. Click **Save and Close**.
7. Verify the changes that you made to the selected virtual interface in the **Network Interfaces** window.

Related references

Network Interfaces window on page 201

Editing network interfaces

You can use the Edit Network Interface dialog box to change network interface parameters, such as the IP address, network mask and MTU size. You can specify the interface name on an HA pair partner for takeover mode, and add, edit, or remove an interface alias.

About this task

You can only view the configuration settings for storage systems running Data ONTAP versions earlier than 7.3.3.

If you enable IPv6 from the command-line interface when the System Manager application is running, System Manager does not detect the change in IPv6 status. Therefore, you must refresh the System Manager session to enable IPv6.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Network > Network Interfaces**.
3. Select the network interface that you want to modify from the interface list and click **Edit**.
4. Click the appropriate tab to display the properties or settings you want to change.
5. Make the necessary changes.
6. Click **Save and Close**.
7. Verify the changes you made to the selected interface in the **Network Interfaces** window.

Related references

Network Interfaces window on page 201

Enabling or disabling network interfaces

You can enable or disable a network interface from the Network Interfaces window.

About this task

You can only view the configuration settings for storage systems running Data ONTAP versions earlier than 7.3.3.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Network > Network Interfaces**.
3. Select the network interface that you want to enable or disable.
4. From the **Status** menu, click either **Enable** or **Disable**, as appropriate.
5. If you are disabling the network interface, click **OK**.

Related references

Network Interfaces window on page 201

Window descriptions

Network Interfaces window

You can use the Network Interfaces window to view a list of network interfaces available in your storage system.

- *Command buttons* on page 201
- *Interface list* on page 202
- *Details area* on page 202

Command buttons

You can only view the configuration settings for storage systems running Data ONTAP versions earlier than 7.3.3.

Create VIF	Opens the Create VIF wizard, which enables you to create virtual interfaces for storage systems running Data ONTAP 7.3.3 or later. **Note:** You cannot add a virtual interface if there are no available interfaces.
Create VLAN	Opens the Create VLAN Interface dialog box, which enables you to add a new VLAN interface for storage systems running Data ONTAP 7.3.3 or later.

Note: You cannot add a VLAN interface if there are no available interfaces.

Edit Opens the Edit Network Interface dialog box, which enables you to modify network interfaces for storage systems running Data ONTAP 7.3.3 or later.

Note: You cannot edit the settings of physical VLANs, trunked interfaces, and interfaces that are used to manage System Manager.

Status Updates the status of the selected network interface for storage systems running Data ONTAP 7.3.3 or later. The interface status can be one of the following:

- Enable
  Enables the selected network interface.
- Disable
  Disables the selected network interface. You cannot disable a physical VLAN or an interface that is a part of the vif.

Note: You cannot modify the status of physical VLANs, trunked interfaces, and interfaces used to manage System Manager.

Refresh Updates the information in the window.

Interface list

The interface list displays the name, type, IP address, and the status of each interface.

Name	Specifies the name of the interface.
Type	Specifies the type of the interface.
IP Address	Specifies the IP address of the storage system.
Status	Specifies the current status of the interface.

Details area

The area below the interface list displays detailed information about the selected interface.

General tab	Displays configuration details for the selected interface.
Alias tab	Displays details on the alias for a selected interface.

Related tasks

Adding interface aliases on page 197
Creating virtual interfaces on page 197
Creating VLAN interfaces on page 198
Editing interface aliases on page 199

Editing virtual interfaces on page 199
Editing network interfaces on page 200
Enabling or disabling network interfaces on page 201

Network > Network Files

Understanding network file configuration

How to maintain host-name information

Data ONTAP relies on correct resolution of host names to provide basic connectivity for storage systems on the network. If you are unable to access the storage system data or establish sessions, there might be problems with host-name resolution on your storage system or on a name server.

Host-name information can be maintained in one or all of the following ways in Data ONTAP:

- In the `/etc/hosts` file on your storage system's default volume
- On a Domain Name System (DNS) server
- On a Network Information Service (NIS) server

If you use more than one of the resources for host-name resolution, the order in which they are used is determined by the `/etc/nsswitch.conf` file.

How the /etc/hosts file works

Data ONTAP uses the `/etc/hosts` file to resolve host names to IP addresses. You need to keep the `/etc/hosts` file up-to-date. Changes to the `/etc/hosts` file take effect immediately.

When Data ONTAP is first installed, the `/etc/hosts` file is automatically created with default entries for the following interfaces:

- Local host
- All interfaces on your storage system

The `/etc/hosts` file resolves the host names for the storage system on which it is configured. This file cannot be used by other systems for name resolution.

For more information about file formats, see the na_hosts(5) man page.

You can add IP address and host name entries in the `/etc/hosts` file in the following two ways:

- Locally—You can add entries by using the command-line interface.
- Remotely—If the file has many entries and you have access to an NIS makefile master, you can use the NIS makefile master to create the `/etc/hosts` file. This method prevents errors that might be caused by editing the file manually.

Hard limits for the /etc/hosts file

You need to be aware of the hard limits on the line size and number of aliases when you edit the `/etc/hosts` file.

The hard limits are as follows:

- Maximum line size is 1022 characters.
  The line size limit includes the end of line character. You can enter up to 1021 characters per line.
- Maximum number of aliases is 34.

Note: There is no limit on file size.

Configuring network files

Adding hosts

You can use the Add Host dialog box to add the IP address of a host or host name entries in the`/etc/hosts` file. Data ONTAP uses this file on the storage system's default volume, NIS, and DNS to resolve host names.

About this task

The `/etc/hosts` file contains information about the known hosts on the network. Each internet IP address is associated with the official host name and any host name aliases.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Network** > **Network Files**.
3. In the **Etc Hosts** tab, click **Add**.
4. Specify properties such as the IP address, the host name, and the aliases of the local host you want to add.
5. Click **OK**.
6. Verify that the local host information that you added is included in the list of host configurations in the **Etc Hosts** tab.

Related references

Network Files window on page 206

Deleting hosts

You can use the **Delete Host** dialog box delete a host name entry in the `/etc/hosts` file.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Network** > **Network Files**.
3. In the **Etc Hosts** tab, select a local host and click **Delete**.
4. Select the confirmation check box and click **Delete**.

Related references

Network Files window on page 206

Managing network files

Editing hosts

You can use the Edit Host dialog box to change the IP address or host name entries in the `/etc/hosts` file.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Network** > **Network Files**.
3. In the **Etc Hosts** tab, select a host configuration from the list and click **Edit**.
4. Change any of the fields for this host and click **OK**.
5. Use the **Etc Hosts** tab to verify the changes you made to the selected host.

Related references

Network Files window on page 206

Editing configuration files

You can edit the configuration files, such as `/etc/hosts.equiv`, `/etc/nsswitch.conf`, and `/etc/netgroup`, from the Network Files window.

About this task

Data ONTAP uses the `/etc/hosts` file to resolve host names to IP addresses. If you use a Domain Name System (DNS) server or a Network Information Service (NIS) server, the order in which they are used is determined by the `/etc/nsswitch.conf` file. The `/etc/netgroup` file defines

network-wide groups used for permission checking when fielding requests for remote mounts, remote logins, and remote shells. For remote mounts, the information in the netgroup file is used to classify machines. For remote logins and remote shells, the file is used to classify users.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Network** > **Network Files**.
3. In the **Others** tab, click the configuration file that you want to modify.
4. Make the necessary changes and click **OK**.

Related references

Network Files window on page 206

Window descriptions

Network Files window

You can use the Network Files window to manage network configuration files and add, edit, or remove local host information.

- *Tabs* on page 206
- *Command buttons* on page 206
- *Network files list* on page 207

Tabs

/etc/hosts	You can use the **/etc/hosts** tab to manage network configuration files and add, edit, or remove local host information.
Others	You can use the **Others** tab to edit other configuration files.

Command buttons

Add	Opens the Add Host dialog box, which enables you to add IP address or host name entries in the `/etc/hosts` file.
Edit	Opens the Edit Host dialog box, which enables you to change the IP address or host name entries in the `/etc/hosts` file.
Delete	Deletes the selected local host information.
Refresh	Updates the information in the window.

Network files list

Address	Specifies the IP address of the local host.
Host Name	Specifies the name of the local host.
Aliases	Specifies the aliases of the local host.

Related tasks

Adding hosts on page 204
Deleting hosts on page 205
Editing hosts on page 205
Editing configuration files on page 205

Network > NIS

Understanding NIS

How to use NIS to maintain host information

NIS enables you to centrally maintain host information. In addition, NIS enables you to maintain user information.

NIS provides the following methods for resolving the storage system's host name:

- Using the `/etc/hosts` file on the NIS server
  You can download the `/etc/hosts` file on the NIS server to your storage system's default volume for local host-name lookup.
- Using a hosts map that is maintained as a database on the NIS server
  The storage system uses the hosts map to query during a host lookup request across the network.
- Using the ipnodes map that is maintained as a database on the NIS server
  The ipnodes map is used for host lookup when IPv6 is enabled on your storage system.

  Note: The ipnodes database is supported only on Solaris NIS servers. To resolve a host name to an address, your storage system (with IPv6 enabled) first looks in the ipnodes database. If the IP address is not present in the ipnodes database, the application looks in the hosts database. However, if IPv6 is not enabled, then your storage system looks only in the hosts database and does not refer the ipnodes database.

How using NIS slaves can improve performance

Host-name resolution by using a hosts map can have a performance impact because each query for the hosts map is sent across the network to the NIS server. You can improve the performance of your storage system by downloading the maps and listening for updates from the NIS master server.

The NIS slave improves performance by establishing contact with an NIS master server and performing the following two tasks:

- Downloading the maps from the NIS master server
  You can download the maps from the NIS master server to the NIS slave by running the `yppush` command from the NIS server. You can also download the maps by disabling and then enabling the NIS slave from your storage system. After the maps are downloaded, they are stored in the `/etc/yp/nis_domain_name` directory. The NIS slave then services all the NIS requests from your storage system by using these maps. The NIS slave checks the NIS master every 45 minutes for any changes to the maps. If there are changes, they are downloaded.
- Listening for updates from the NIS master
  When the maps on the NIS master are changed, the NIS master administrator can optionally notify all slaves. Therefore, in addition to periodically checking for updates from the NIS master, the NIS slave also listens for updates from the master.

You cannot configure the NIS slave during the setup procedure. To configure the NIS slave after the setup procedure is complete, you need to enable NIS slave by setting `options nis.slave.enable` to `on`.

Note: The NIS slave does not respond to remote NIS client requests and therefore cannot be used by other NIS clients for name lookups.

Guidelines for using NIS slaves

When using an NIS slave, you should follow certain guidelines, such as the available space in the storage system, conditions for enabling DNS, and supported configurations.

The following guidelines apply when using the NIS slave:

- The root volume of your storage system must have sufficient space to download maps for the NIS slave. Typically, the space required in the root volume is same as the size of the maps on the NIS server.
- If the root volume does not have enough space to download maps, the following occurs:
  - An error message is displayed informing you that the space on the disk is not sufficient to download or update the maps from the NIS master.
  - If the maps cannot be downloaded, the NIS slave is disabled. Your storage system switches to using hosts map on the NIS server for name resolution.
  - If the maps cannot be updated, your storage system continues to use the old maps.
- If the NIS master server was started with the `-d` option or if the `hosts.byname` and `hosts.byaddr` maps are generated with the `-b` option, your storage system must have DNS

enabled, DNS servers must be configured, and the hosts entry in the `/etc/nsswitch.conf` file must contain DNS as an option to use for host name lookup.

If you have your NIS server configured to perform host name lookups using DNS, or if you use DNS to resolve names that cannot be first resolved using the `hosts.by*` maps, using the NIS slave causes those lookups to fail. This is because when the NIS slave is used, all lookups are performed locally using the downloaded maps. However, if you configure DNS on your storage system, the lookups succeed.

You can use the NIS slave for the following:

- Interface groups and VLAN interfaces
- vFiler units
- HA pairs

Note: In an HA pair, you should ensure that the `nis.servers` options value is the same on both nodes and that the `/etc/hosts` file on both nodes can resolve the name of the NIS master server.

Things to consider when binding NIS servers to storage systems

There are certain guidelines that you must follow before binding a NIS server to your storage system.

Keep the following in mind before performing the binding procedure:

- Using the NIS broadcast feature can incur security risks.
- You can specify NIS servers by IP address or host name. If host names are used, ensure that each host name and its IP address are listed in the `/etc/hosts` file of your storage system. Otherwise, the binding with the host name fails.
- You can only specify IPv4 addresses or server names that resolve to IPv4 addresses using the `/etc/hosts` file on your storage system.

Configuring NIS

Enabling or disabling NIS

NIS enables you to centrally maintain host and user information. You can use the Edit NIS Settings dialog box to enable or disable NIS on your storage system. NIS is disabled by default.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Network > NIS**.
3. Click **Edit**.
4. Either select or clear **Enable NIS**, as appropriate.
5. Click **Save and Close**.

Related references

NIS window on page 211

Adding or editing the NIS domain name

You can maintain host information centrally using NIS. You can use the Edit NIS Settings dialog box to add or modify the NIS domain name of your storage system.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Network** > **NIS**.
3. Click **Edit**.
4. In the NIS domain name field, type or change the NIS domain name.
5. Click **Save and Close**.

Related references

NIS window on page 211

Managing NIS

Enabling or disabling an NIS slave

You can enable an NIS slave on your storage system to reduce traffic over your network. You can use the Edit NIS Settings dialog box to enable or disable an NIS slave on your storage system. The NIS slave is disabled by default.

About this task

If you enable and then later disable the NIS slave, the storage system reverts to the original configuration, where it contacts an NIS server to resolve host names.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Network** > **NIS**.
3. Click **Edit**.
4. In the **Advanced** tab, either select or clear **Enable NIS slave**, as appropriate.
5. Schedule the caching of NIS group information by performing the appropriate action:

If you want to...	Then...
Immediately update the NIS group information	Click **Now**.
Update the NIS group information at regular intervals	Click **Every** and specify the time interval.

6. Click **Save and Close**.

Related references

NIS window on page 211

Window descriptions

NIS window

The NIS window enables you to view the current NIS settings for your storage system.

Command buttons

Edit Opens the Edit NIS Settings dialog box, which you can use to enable or disable NIS, add or modify the NIS domain name, and enable or disable the NIS slave.

Refresh Updates the information in the window.

Related tasks

Enabling or disabling NIS on page 209
Adding or editing the NIS domain name on page 210
Enabling or disabling an NIS slave on page 210

Protocols > CIFS

Understanding CIFS

About CIFS and SMB

Data ONTAP supports all of the most common file protocols, including the CIFS protocol to enable file sharing from host storage systems. When your system is first installed and CIFS is configured in Workgroup mode, a login named "administrator" is automatically created. You can use this login to access shares with a blank password.

The CIFS protocol is used to share files. CIFS is the method of transport for Windows Shares. CIFS is an extension of the Server Message Block (SMB) protocol, which is a file-sharing protocol used on Windows and UNIX systems. SMB runs over several different types of networks, including TCP/IP. For most purposes, SMB is superseded by CIFS.

CIFS license

Your storage system requires a software license to enable CIFS service. This license is installed on the storage system at the factory per your order; therefore, you do not typically need to enter license codes when you initially configure your system. If CIFS license is not installed on the storage system, then System Manager does not list 'cifs' in the Licenses window (**Configuration > System Tools > Licenses**).

You need to enter license codes only if any of the following conditions applies:

- You purchased a storage system with a software version earlier than Data ONTAP 4.0 and you are upgrading it.
- You want to enable CIFS, which was not previously licensed for your storage system.
- You reinstalled your file system on an existing system that was not shipped with it installed.

In these cases, you are provided with the appropriate license codes when the software upgrade kit is shipped to you or when you are given instructions for obtaining the software upgrade over the Internet.

What CIFS auditing does

System Manager enables you to use CIFS auditing to monitor reads and writes of a specified file on the storage system by a specified user.

You can use System Manager to set up auditing of the following events:

- Logon and logoff events
- File access events
- Account management

The file on the storage system must be in a mixed or NTFS volume or qtree. You cannot audit events on a file in a UNIX volume or qtree.

You can specify the logging of successes, failures, or both, for any type of event.

What an event log is

You can use the event log to see the file access information gathered by CIFS auditing. The log is in Windows NT format and can be viewed by the Event Viewer.

By default, the event log is `/etc/log/adtlog.evt`. You can specify another file as the event log and an alternative maximum file size.

You cannot update the event log when it is being viewed by a client. To prevent losing event information that is gathered when the event log is open, System Manager does not write to the event log as event information is being collected. Instead, it updates the event log when you manually save the log from System Manager.

About home directories on the storage system

Data ONTAP maps home directory names to user names, searches for home directories that you specify, and treats home directories slightly differently than regular shares

Data ONTAP offers the share to the user with a matching name. The user name for matching can be a Windows user name, a domain name followed by a Windows user name, or a UNIX user name. Home directory names are not case-sensitive.

When Data ONTAP tries to locate the directories named after the users, it searches only the paths that you specify. These paths are called home directory paths. They can exist in different volumes.

The following differences exist between a home directory and other shares:

- You cannot change the share-level ACL and the comment for a home directory.
- The cifs shares command does not display the home directories.
- The format of specifying the home directory using the Universal Naming Convention (UNC) is sometimes different from that for specifying other shares.

If you specify `/vol/vol1/enghome` and `/vol/vol2/mktghome` as the home directory paths, Data ONTAP searches these paths to locate user home directories. If you create a directory for jdoe in the `/vol/vol1/enghome` path and a directory for jsmith in the `/vol/vol2/mktghome` path, both users are offered a home directory. The home directory for jdoe corresponds to the `/vol/vol1/enghome/jdoe` directory, and the home directory for jsmith corresponds to the `/vol/vol2/mktghome/jsmith` directory.

Configuring CIFS

Setting up CIFS

You can set up CIFS from the CIFS Set Up wizard. If the CIFS service is already running, completing the CIFS Set Up wizard stops and restarts the CIFS service.

Before you begin

- The CIFS license must be installed on your storage system.
- While configuring CIFS in the Active Directory domain, you must ensure the following requirements are met.
  - DNS must be enabled and configured correctly.
  - The storage system must be able to talk to the domain controller using the fully qualified domain name (FQDN).
  - The time differences (clock skew) between the storage system time and the domain time must not be more than the skew time that is configured in Data ONTAP.

Steps

1. From the **Home** tab, double-click the appropriate storage system.

2. In the navigation pane, click **Configuration** > **Protocols** > **CIFS**.
3. In the **Configuration** tab, click Setup.
4. Type or select information as prompted by the wizard.
5. Confirm the details and click **Finish** to complete the wizard.

Related tasks

Creating a CIFS share on page 66

Related references

CIFS window on page 221

Configuring CIFS and NFS auditing

You can configure CIFS and NFS auditing on your storage system to troubleshoot access problems, check for suspicious activity on a system, or investigate a security breach.

Before you begin

- The file or directory to be audited must be in a mixed or NTFS volume or qtree.
- Access to individual files and directories must be activated according to Windows documentation.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Protocols** > **CIFS**.
3. In the **Auditing** area, click **Edit**.
4. In the **Edit Auditing** dialog box, select the appropriate check boxes to enable CIFS and NFS auditing.
5. If you are configuring NFS auditing, click **Browse** and select the appropriate NFS audit filter file.
6. Specify the general settings for the audit log file.
7. Select the check boxes corresponding to the types of events you want to audit.
8. Click **Save and Close** to save your changes and close the dialog box.

Related references

CIFS window on page 221

Managing CIFS

Editing the general properties for CIFS

You can modify the general properties for CIFS, such as the server description, idle timeout for a CIFS session, Snapshot access mode, and maximum concurrent operations.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Protocols** > **CIFS**.
3. In the **Configuration** tab, click **Options**.
4. In the **CIFS Options** dialog box, click **General** and make the necessary changes.
5. Click **Save and Close** to save your changes and close the dialog box.

Related references

CIFS window on page 221

Editing the networking properties for CIFS

You can modify the CIFS networking options and add or remove WINS servers and NetBIOS aliases. You can also enable or disable NetBIOS over TCP.

Before you begin

- If you are adding a WINS server, the WINS server name or IP address must be available.
- If you are adding a NetBIOS alias, the NetBIOS alias name must be available.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Protocols** > **CIFS**.
3. In the **Configuration** tab, click **Options**.
4. In the **CIFS Options** dialog box, click **Networking** and make the necessary changes.
5. Click **Save and Close** to save your changes and close the dialog box.

Related references

CIFS window on page 221

Editing the access security properties for CIFS

You can set the restriction level for your CIFS session and enable or disable SMB signing.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Protocols > CIFS**.
3. In the **Configuration** tab, click **Options**.
4. In the **CIFS Options** dialog box, click **Access Security** and make the necessary changes.
5. Click **Save and Close** to save your changes and close the dialog box.

Related references

CIFS window on page 221

Adding home directory paths

You can specify one or more paths that can be used by the storage system to resolve the location of users' CIFS home directories. You can add a home directory path by using the Edit Home Directories dialog box.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Protocols > CIFS**.
3. Click **Configuration**.
4. In the **Home Directories** area, click **Edit**.
5. In the **Edit Home Directories** dialog box, specify the naming style that is used for home directories.
6. Specify the paths used by the storage system to search for users' CIFS home directories.
7. Click **Add**, and click **Save and Close**.

Related references

CIFS window on page 221

Deleting home directory paths

You can delete a home directory path when you do not want the storage system to use the path to resolve the location of users' CIFS home directories. You can delete a home directory path by using the Edit Home Directories dialog box.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Protocols** > **CIFS**.
3. Click **Configuration**.
4. In the Home Directories area, click **Edit**.
5. In the **Edit Home Directories** dialog box, select the home directory path that you want to delete and click **Delete**.
6. Click **Save and Close** to save your changes and close the dialog box.

Related references

CIFS window on page 221

Stopping and restarting CIFS

You can stop and then optionally restart the CIFS service from the CIFS window. When you stop CIFS, all the sessions connected to the service are stopped and all the shared folders on the host storage system are unavailable.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Protocols** > **CIFS**.
3. In the **Configuration** tab, click **Stop** to stop the CIFS service.
4. If you want to restart CIFS, click **Start**.

Related references

CIFS window on page 221

Saving your audit log

You can save your audit log either to the default location or to a different location.

Before you begin

CIFS auditing must be enabled.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Protocols > CIFS**.
3. In the **Auditing** area, click **Edit**.
4. If you want to save the audit log file in a different location, enter the new location, or click **Browse** and select the path.
5. Click **Save and Close** to save your changes and close the dialog box.

Related references

CIFS window on page 221

Clearing your audit log

You can clear your audit log if you want the audit information to restart from a certain point.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Protocols > CIFS**.
3. In the **Auditing** area, click **Clear Log**.
4. Click **Clear** on the confirmation prompt.

Related references

CIFS window on page 221

Enabling or disabling audit events

You can enable or disable your audit event as required.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Protocols > CIFS**.
3. Click **Configuration**.
4. In the Auditing area, click **Edit**.
5. In the **Edit Auditing** dialog box, either select or clear the type of auditing check box, as required.
6. Click **Save and Close** to save your changes and close the dialog box.

Related references

CIFS window on page 221

Resetting CIFS domain controllers

You have to reset the CIFS connection to domain controllers for the specified domain. Failure to reset the domain controller information can cause a connection failure.

About this task

You have to update the discovery information of the storage system's available domain controller after you add or delete a domain from the list of preferred domain controllers. You can update the storage system's available domain controller discovery information in Data ONTAP through the command-line interface (CLI).

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Protocols** > **CIFS**.
3. In the **Domain** tab, click **Reset**.

Related references

CIFS window on page 221

Enabling a trace on a CIFS domain controller

You can enable a trace to log all the domain controller discovery and connection activities on the storage system. The trace logs can be used to diagnose domain controller connection problems on the storage system.

About this task

All the domain controller address discovery and connection activities on the storage system are logged to syslog. This information, by default, is logged in the `/etc/messages` file and the console.

Note: Enabling a trace on a CIFS domain controller might impact system performance.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Protocols** > **CIFS**.
3. If the appropriate tab is not displayed, click **Domain**, and then click **Edit**.
4. Select the option for enabling a trace log and click **OK**.

Related references

CIFS window on page 221

Scheduling the frequency of password changes

You can schedule the domain password to be changed once a week to improve the security of the storage system.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Under the selected host storage system, click **Configuration > Protocols > CIFS**.
3. In the **Domain** tab, click **Edit**.
4. Select the check box to schedule a weekly password change and click **OK**.

Result

The password change occurs at approximately 1:00 a.m. on Sundays.

Related references

CIFS window on page 221

Translating user or group names to security identifiers

You can use the CIFS window to translate a Windows NT user or group name to its corresponding textual Windows NT security identifier (SID), or a textual NT SID to its corresponding Windows NT user or group name.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Protocols > CIFS**.
3. In the **Configuration** tab, click **Look up** in the **CIFS** area.
4. Enter the user name, group name or SID, and click **Look up**.
5. Click **Close**.

Related references

CIFS window on page 221

Monitoring CIFS

Viewing CIFS domain information

You can view information about the domain controllers and LDAP servers that the storage system is connected to.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Click **Domain**.
3. Review the information about the connected domain controllers and connected servers.

Window descriptions

CIFS window

You can use the CIFS window to manage your CIFS sessions and domain controllers.

- *Tabs* on page 221
- *Command buttons* on page 221
- *Details area* on page 222

Tabs

Configuration tab	Enables you to manage your CIFS sessions.
Domain tab	Enables you to view, test, and reset your CIFS domain controllers.

Command buttons

Setup	Launches the CIFS Setup wizard, which enables you to set up CIFS.
Start	Starts the CIFS session.
Stop	Stops the CIFS session. **Note:** Stopping the CIFS session causes all shared sessions on your storage system to become unavailable.
Options	Displays the CIFS Options dialog box, which enables you to modify the CIFS properties.
Refresh	Updates the information in the window.
Lookup	Opens CIFS Look Up SID/Name dialog box, which enables you to find the user name and group names you want to add to the CIFS session.

Edit	Opens a dialog box, which enables you to modify CIFS auditing or home directory information for CIFS.
Clear log	Deletes all of the information in the log file.

Details area

CIFS	Specifies the CIFS session details such as the authentication type.
Auditing	Specifies the status of CIFS and NFS auditing, and the location of the log file.
Home directories	Specifies home directory paths and the style that determines how you want PC user names to be mapped to home directory entries.

Related tasks

Setting up CIFS on page 213
Configuring CIFS and NFS auditing on page 214
Editing the general properties for CIFS on page 215
Editing the networking properties for CIFS on page 215
Editing the access security properties for CIFS on page 216
Adding home directory paths on page 216
Deleting home directory paths on page 217
Stopping and restarting CIFS on page 217
Saving your audit log on page 217
Clearing your audit log on page 218
Enabling or disabling audit events on page 218
Resetting CIFS domain controllers on page 219
Enabling a trace on a CIFS domain controller on page 219
Scheduling the frequency of password changes on page 220
Translating user or group names to security identifiers on page 220

Protocols > NFS

Understanding NFS

NFS concepts

NFS clients can access your storage system using the NFS protocol provided Data ONTAP can properly authenticate the user.

When an NFS client connects to the Vserver, Data ONTAP obtains the UNIX credentials for the user by checking different name services, depending on the name services configuration of the Vserver.

The options are local UNIX accounts, NIS domains, and LDAP domains. You must configure at least one of them so Data ONTAP can successfully authorize the user. You can specify multiple name services and the order in which they are searched.

In a pure NFS environment with UNIX volume security styles, this configuration is sufficient to authenticate a user connecting from an NFS client and provide the proper file access.

If you are using mixed or NTFS volume security styles, Data ONTAP must obtain a CIFS user name for the UNIX user for authentication with a Windows domain controller. This can happen either by mapping individual users using local UNIX accounts or LDAP domains, or by using a default CIFS user instead. You can specify for the Vserver which name services are searched in which order, or specify a default CIFS user.

Managing NFS

Editing NFS settings

You can edit the NFS settings, such as enabling or disabling NFSv3 and NFSv4, enabling or disabling read and write delegations for NFSv4 clients, and enabling NFSv4 ACLs.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Protocols > NFS**.
3. Click **Edit**.
4. Make the necessary changes.
5. Click **Save and Close** to save your changes and close the dialog box.

Related references

NFS window on page 224

Enabling or disabling the NFS service

You can enable or disable the NFS service from the NFS window.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Protocols > NFS**.
3. Click either **Enable** or **Disable**, as required.

Related references

NFS window on page 224

Window descriptions

NFS window

You can use the NFS window to display and configure your NFS settings.

Command buttons

Enable	Enables the NFS service.
Disable	Disables the NFS service.
Edit	Opens the Edit NFS Settings dialog box, which enables you to edit NFS settings.
Refresh	Updates the information in the window.

Related tasks

Editing NFS settings on page 223
Enabling or disabling the NFS service on page 223

Protocols > iSCSI

Understanding iSCSI

What iSCSI is

The iSCSI protocol is a licensed service on the storage system that enables you to transfer block data to hosts using the SCSI protocol over TCP/IP. The iSCSI protocol standard is defined by RFC 3720.

In an iSCSI network, storage systems are targets that have storage target devices, which are referred to as LUNs (logical units). A host with an iSCSI host bus adapter (HBA), or running iSCSI initiator software, uses the iSCSI protocol to access LUNs on a storage system. The iSCSI protocol is implemented over the storage system's standard Ethernet interfaces using a software driver.

The connection between the initiator and target uses a standard TCP/IP network. No special network configuration is needed to support iSCSI traffic. The network can be a dedicated TCP/IP network, or it can be your regular public network. The storage system listens for iSCSI connections on TCP port 3260.

Related information

RFC 3270 - www.ietf.org/rfc/rfc3270.txt
RFC 3720 - www.ietf.org/rfc/rfc3720.txt

What iSCSI nodes are

In an iSCSI network, there are two types of nodes: targets and initiators. Targets are storage systems, and initiators are hosts. Switches, routers, and ports are TCP/IP devices only, and are not iSCSI nodes.

How iSCSI works with HA pairs

HA pairs provide high availability because one system in the HA pair can take over if its partner fails. During failover, the working system assumes the IP addresses of the failed partner and can continue to support iSCSI LUNs.

The two systems in the HA pair should have identical networking hardware with equivalent network configurations. The target portal group tags associated with each networking interface must be the same on both systems in the configuration. This ensures that the hosts see the same IP addresses and target portal group tags whether connected to the original storage system or connected to the partner during failover.

Target portal group management

A target portal group is a set of one or more storage system network interfaces that can be used for an iSCSI session between an initiator and a target. A target portal group is identified by a name and a numeric tag. If you want to have multiple connections per session across more than one interface for performance and reliability reasons, then you must use target portal groups.

Note: If you are using MultiStore, you can also configure non-default vFiler units for target portal group management based on IP address.

For iSCSI sessions that use multiple connections, all of the connections must use interfaces in the same target portal group. Each interface belongs to one and only one target portal group. Interfaces can be physical interfaces or logical interfaces (VLANs and interface groups).

Prior to Data ONTAP 7.1, each interface was automatically assigned to its own target portal group when the interface was added. The target portal group tag was assigned based on the interface location and could not be modified. This works fine for single-connection sessions.

You can explicitly create target portal groups and assign tag values. If you want to increase performance and reliability by using multi-connections per session across more than one interface, you must create one or more target portal groups.

Because a session can use interfaces in only one target portal group, you might want to put all of your interfaces in one large group. However, some initiators are also limited to one session with a given target portal group. To support multipath I/O (MPIO), you need to have one session per path, and therefore more than one target portal group.

When a new network interface is added to the storage system, that interface is automatically assigned to its own target portal group.

Initiator security

You can select from the following authentication methods:

- none
  There is no authentication for the initiator.
- deny
  The initiator is denied access when it attempts to authenticate to the storage system.
- CHAP
  The initiator logs in using a Challenge Handshake Authentication Protocol (CHAP) user name and password. You can specify a CHAP password or generate a random password.
- default
  The initiator uses the default security settings. The initial setting for default initiator security is none.

In CHAP authentication, the storage system sends the initiator a challenge value. The initiator responds with a value calculated using a one-way hash function. The storage system then checks the response against its own version of the value calculated using the same one-way hash function. If the values match, the authentication is successful.

How iSCSI communication sessions work

During an iSCSI session, the initiator and the target communicate over their standard Ethernet interfaces, unless the host has an iSCSI HBA or a CNA.

The storage system appears as a single iSCSI target node with one iSCSI node name. For storage systems with a MultiStore license enabled, each vFiler unit is a target with a different iSCSI node name.

On the storage system, the interface can be an Ethernet port, interface group, UTA, or a virtual LAN (VLAN) interface.

Each interface on the target belongs to its own portal group by default. This enables an initiator port to conduct simultaneous iSCSI sessions on the target, with one session for each portal group. The storage system supports up to 1,024 simultaneous sessions, depending on its memory capacity. To determine whether your host's initiator software or HBA can have multiple sessions with one storage system, see your host OS or initiator documentation.

You can change the assignment of target portals to portal groups as needed to support multi-connection sessions, multiple sessions, and multipath I/O.

Each session has an Initiator Session ID (ISID), a number that is determined by the initiator.

How iSCSI authentication works

During the initial stage of an iSCSI session, the initiator sends a login request to the storage system to begin an iSCSI session. The storage system permits or denies the login request according to one of the available authentication methods.

The authentication methods are as follows:

- Challenge Handshake Authentication Protocol (CHAP)—The initiator logs in using a CHAP user name and password.

  You can specify a CHAP password or generate a random password. There are two types of CHAP user names and passwords:

  - Inbound—The storage system authenticates the initiator.

    Inbound settings are required if you are using CHAP authentication.
  - Inbound—The storage system authenticates the initiator.

    Inbound settings are required if you are using CHAP authentication without RADIUS.
  - Outbound—This is an optional setting to enable the initiator to authenticate the storage system.

    You can use outbound settings only if you defined an inbound user name and password on the storage system.

    RADIUS can be used in conjunction with CHAP for initiator authentication. With this method, the initiator logs in using a CHAP user name and password, but authentication is managed from a centralized RADIUS server rather than locally on the storage system.
- deny—The initiator is denied access to the storage system.
- none—The storage system does not require authentication for the initiator.

You can define a list of initiators and their authentication methods. You can also define a default authentication method that applies to initiators that are not on this list.

The default iSCSI authentication method is none, which means any initiator not in the authentication list can log in to the storage system without authentication. However, you can change the default method to deny or CHAP.

If you use iSCSI with vFiler units, the CHAP authentication settings are configured separately for each vFiler unit. Each vFiler unit has its own default authentication mode and list of initiators and passwords.

To configure CHAP settings for vFiler units, you must use the command line.

For information about managing vFiler units, see the sections on iSCSI service on vFiler units in the *Data ONTAP 7-Mode MultiStore Management Guide*.

Related information

Data ONTAP documentation on NetApp support site -- support.netapp.com

What CHAP authentication is

The Challenge Handshake Authentication Protocol (CHAP) enables authenticated communication between iSCSI initiators and targets. When you use CHAP authentication, you define CHAP user names and passwords on both the initiator and the storage system.

During the initial stage of an iSCSI session, the initiator sends a login request to the storage system to begin the session. The login request includes the initiator's CHAP user name and CHAP algorithm. The storage system responds with a CHAP challenge. The initiator provides a CHAP response. The storage system verifies the response and authenticates the initiator. The CHAP password is used to compute the response.

Configuring iSCSI

Creating iSCSI aliases

An iSCSI alias is a user-friendly identifier that you assign to an iSCSI target device (in this case, the storage system) to make it easier to identify the target device in user interfaces. You can use the Edit iSCSI Service Configurations dialog box to create an iSCSI alias.

About this task

An iSCSI alias is a string of 1 to 128 printable characters, and must not include spaces.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Protocols** > **iSCSI**.
3. In the **Service** tab, click **Edit**.
4. Type an iSCSI alias in the **Target Alias** field and click **OK**.

Related references

iSCSI window on page 233

Creating target portal groups

If you want to use multi-connection iSCSI sessions to improve performance and reliability, then you must use target portal groups to define the interfaces available for each iSCSI session.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Protocols** > **iSCSI**.
3. In the **Target Portal Group** tab, click **Create**.

4. Type the name of the target portal group and select the numeric tag for the portal group.

   If you leave the tag field blank, the system assigns the next available tag value.

5. Select the interfaces to include in the target portal group and click **Create**.

Related references

iSCSI window on page 233

Deleting target portal groups

You can delete one or more user-defined target portal groups. Deleting a target portal group removes the group from the storage system. Interfaces that belonged to the group are returned to their individual default target portal groups. You cannot delete system-defined portal groups.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Protocols > iSCSI**.
3. In the **Target Portal Group** tab, select the target portal group that you want to delete and click **Delete**.
4. Select the confirmation check box and click **Delete**.

Related references

iSCSI window on page 233

Enabling or disabling the iSCSI service on storage system interfaces

You can control which network interfaces are used for iSCSI communication by enabling or disabling the interfaces. When the iSCSI service is enabled, iSCSI connections and requests are accepted over those network interfaces that are enabled for iSCSI, but not over disabled interfaces.

Before you begin

You must terminate any outstanding iSCSI connections and sessions currently using the interface. By default, the iSCSI service is enabled on all Ethernet interfaces after you enable the iSCSI license.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Network > NIS**.
3. In the **iSCSI Interfaces** area, select the interface on which you want to enable or disable the iSCSI service.
4. Click **Enable** or **Disable**, as required.

Related references

iSCSI window on page 233

Adding the security method for iSCSI initiators

You can use the Add Initiator Security dialog box to add an initiator and specify the security method that is used to authenticate the initiator.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Protocols** > **iSCSI**.
3. In the **Initiator Security** tab, click **Add** in the **Initiator Security** area.
4. Specify the initiator name and the security method to authenticate the initiator.

   For CHAP authentication, you must provide the user name and password, and confirm your password for inbound settings. For outbound settings, this login information is optional.
5. Click **OK**.

Related references

iSCSI window on page 233

Managing iSCSI

Editing default security settings

You can use the Edit Default Security dialog box to edit the default security settings for iSCSI initiators that are connected to the storage system.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Protocols** > **iSCSI**.
3. In the **Default Security** box on the **Initiator Security** tab, click **Edit**.
4. Change the security type.

   For CHAP authentication, you must provide the user name and password, and confirm your password for inbound settings. For outbound settings, this login information is optional.
5. Click **OK**.

Related references

iSCSI window on page 233

Editing initiator security

The security style configured for an initiator specifies how the authentication is done for that initiator during the iSCSI connection login phase. You can change the security for selected iSCSI initiators by changing the authentication method.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Protocols > iSCSI**.
3. In the **Initiator Security** tab, select one or more initiators from the initiator list and click **Edit** in the **Initiator Security** area.
4. Change the security type.

   For CHAP authentication, you must provide the user name and password and confirm your password for inbound settings. For outbound settings, this is optional.
5. Click **OK**.
6. Verify the changes you made in the **Initiator Security** tab.

Related references

iSCSI window on page 233

Changing the default iSCSI initiator authentication method

You can change the default iSCSI authentication method, which is the authentication method that is used for any initiator that is not configured with a specific authentication method.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Protocols > iSCSI**.
3. In the **Initiator Security** tab, click **Edit** in the **Default Security** area.
4. Change the security type.

   For CHAP authentication, you must provide the user name and password and confirm your password for inbound settings. For outbound settings, this is optional.
5. Click **OK**.

Related references

iSCSI window on page 233

Setting the default security for iSCSI initiators

You can remove the authentication settings for an initiator and use the default security method to authenticate the initiator.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Protocols > iSCSI**.
3. In the **Initiator Security** tab, select the initiator whose security setting you want change.
4. Click **Set Default** in the **Initiator Security** area, and then click **Set Default** in the confirmation box.

Related references

iSCSI window on page 233

Editing a target portal group

You can edit a user-defined target portal group by adding interfaces or removing interfaces from it. When you add interfaces, the specified interfaces are removed from their current groups and added to the group. When you remove interfaces, the specified interfaces are removed from the group and returned to their individual default target portal groups.

About this task

You cannot edit system-defined target portal groups.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Protocols > iSCSI**.
3. In the **Target Portal Group** tab, select the portal group that you want to edit and click **Edit**.
4. Select the interfaces that you want to add to or remove from the portal group and click **Save**.

Related references

iSCSI window on page 233

Starting or stopping the iSCSI service

You can start or stop the iSCSI service on your storage system.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Protocols** > **iSCSI**.
3. Click either **Start** or **Stop**, as required.

Related references

iSCSI window on page 233

Monitoring iSCSI

Viewing initiator security information

You can use the Initiator Security tab to view the default authentication information and all the initiator-specific authentication information.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Protocols** > **iSCSI**.
3. In the **Initiator Security** tab, review the details.

Window descriptions

iSCSI window

You can use the iSCSI window to start or stop iSCSI service, change a storage system iSCSI node name, and create or change the iSCSI alias of a storage system. You can also add or change the initiator security setting for an iSCSI initiator that is connected to your storage system.

Tabs

Service	You can use the **Service** tab to start or stop iSCSI service, change a storage system iSCSI node name, and create or change the iSCSI alias of a storage system.
Initiator Security	You can use the **Initiator Security** tab to add or change the initiator security setting for an iSCSI initiator that is connected to your storage system.

Target Portal Group	You can use the Target Portal Group tab to manage a group of one or more storage system network interfaces that can be used for an iSCSI session between an initiator and a target.

Command buttons

Edit	Opens Edit iSCSI Service Configurations dialog box, which enables you to change iSCSI node name and iSCSI alias of the storage system.
Start	Starts the iSCSI service.
Stop	Stops the iSCSI service.
Refresh	Updates the information in the window.

Details area

The details area displays information about the status of iSCSI service, iSCSI target node name, and iSCSI target alias. You can use this area to enable or disable the iSCSI service on a network interface.

Related tasks

Creating iSCSI aliases on page 228
Creating target portal groups on page 228
Deleting target portal groups on page 229
Enabling or disabling the iSCSI service on storage system interfaces on page 229
Adding the security method for iSCSI initiators on page 230
Editing default security settings on page 230
Editing initiator security on page 231
Changing the default iSCSI initiator authentication method on page 231
Setting the default security for iSCSI initiators on page 232
Editing a target portal group on page 232
Starting or stopping the iSCSI service on page 233

Protocols > FC/FCoE

Understanding FC/FCoE

What FC is

FC is a licensed service on the storage system that enables you to export LUNs and transfer block data to hosts using the SCSI protocol over a Fibre Channel fabric.

What FC nodes are

In an FC network, nodes include targets, initiators, and switches.

Targets are storage systems, and initiators are hosts. Nodes register with the Fabric Name Server when they are connected to an FC switch.

How FC target nodes connect to the network

Storage systems and hosts have adapters, so they can be directly connected to each other or to FC switches with optical cables. For switch or storage system management, they might be connected to each other or to TCP/IP switches with Ethernet cables.

When a node is connected to the FC SAN, it registers each of its ports with the switch's Fabric Name Server service, using a unique identifier.

The FCoE protocol

Fibre Channel over Ethernet (FCoE) is a new model for connecting hosts to storage systems. Like the traditional FC protocol, FCoE maintains existing FC management and controls, but it uses a 10-gigabit Ethernet network as the hardware transport.

Setting up an FCoE connection requires one or more supported converged network adapters (CNAs) in the host, connected to a supported data center bridging (DCB) Ethernet switch. The CNA is a consolidation point and effectively serves as both an HBA and an Ethernet adapter.

In general, you can configure and use FCoE connections the same way you use traditional FC connections.

Configuring FC/FCoE

Starting or stopping the FC or FCoE service

The FC service enables you to manage FC target adapters for use with LUNs. You have to start the FC service to bring the adapters online and allow access to the LUNs on the storage system. You can stop the FC service to take the FC adapters offline and prevent access to the LUNs.

Before you begin

- The FC license must be installed.
- An FC adapter must be present in the target storage system.

About this task

If your storage system is running Data ONTAP versions earlier than 7.3.2, the left pane displays FCP as the Fibre Channel protocol, and if your storage system is running Data ONTAP 7.3.2 or later, FC/FCoE is displayed.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Protocols**, and then click the Fibre Channel protocol.
3. Click either **Start** or **Stop**, as appropriate.
4. If you are stopping the FC or FCoE service, click **Stop**.

Related references

FC/FCoE window on page 237

Managing FC/FCoE

Changing an FC or FCoE node name

If you replace a storage system chassis and reuse it in the same Fibre Channel SAN, the node name of the replaced storage system in certain cases might be duplicated. You can change the node name of the storage system by using the Edit Node Name dialog box.

About this task

If your storage system is running Data ONTAP versions earlier than 7.3.2, the left pane displays FCP as the Fibre Channel protocol, and if your storage system is running Data ONTAP 7.3.2 or later, FC/FCoE is displayed.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Protocols**, and then click the Fibre Channel protocol.
3. Click **Edit**.
4. Type the new name and click **OK**.

Related references

FC/FCoE window on page 237

Window descriptions

FC/FCoE window

You can use the FC/FCoE window to start or stop the FC service. If your storage system is running Data ONTAP versions earlier than 7.3.2, the left pane displays FCP as the Fibre Channel protocol, and if your storage system is running Data ONTAP 7.3.2 or later, FC/FCoE is displayed.

Command buttons

Edit	Opens the Edit Node Name dialog box, which enables you to change the FC or FCoE node name.
Start	Starts the FC/FCoE service.
Stop	Stops the FC/FCoE service.
Refresh	Updates the information in the window.

FC/FCoE details

The details area displays information about the status of FC/FCoE service, the node name, and the FC/FCoE adapters.

Related tasks

Starting or stopping the FC or FCoE service on page 236
Changing an FC or FCoE node name on page 236

Security > Password/RSH

Understanding password/RSH

When to configure RSH

You can use a remote shell (RSH) to run a command on a remote host.

You can use the RSH security feature to specify a host name or IP address from which to execute a command.

What trusted hosts are

You can use the trusted host feature to limit the hosts from which you can access your storage system. Access is typically made through a telnet connection or a Web browser.

The default value for this trusted host is "All" which means that you can connect to your storage system through any host via a telnet or HTTP connection. To restrict host access, you must specify the IP address of the host machine or machines that you want to specify as trusted.

Configuring password/RSH

Changing the system password

You can change the storage system password for increased security. The system password is also the password for the root user account.

Before you begin

The current system password must be available.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Security** > **Password/RSH**.
3. Click **Change Password.**
4. Type your current password in the appropriate field.

   If you have the capability to change the password of other users, you do not have to enter the current password.
5. Type your new password in the appropriate fields.
6. Click **Change**.

Related references

Password/RSH window on page 240

Adding or deleting RSH host names

You can control which hosts can access the storage system through a Remote Shell session for administrative purposes. You can restrict Secure Shell access to the storage system by specifying the host name and user ID.

Before you begin

The following information must be available:

- Host name or IP address
- User ID

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Security > Password/RSH**.
3. Click **Edit**.
4. In the **Edit Security Settings** dialog box, click **RSH settings**.
5. Choose the appropriate action:
   - To add an RSH host name and user ID, type the host name or IP address and the user ID in the appropriate fields and click **Add**.
     You can repeat this step to add more host names and user IDs.
   - To delete an RSH host name and user ID, select the name or IP address that you want to delete and click **Delete**.
6. Click **OK** to save your changes.

Related references

Password/RSH window on page 240

Managing trusted hosts

You can specify the hosts that are allowed to access a storage system. These hosts are considered trusted hosts of that storage system. You can also specify that all hosts are trusted or that none of the hosts are trusted. Setting trusted hosts to None prevents access to the hosts from System Manager.

Before you begin

The name or IP address of host that you want to specify as trusted must be available.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Security**.
3. Click **Password/RSH**.
4. Click **Edit**.
5. In the **Edit Security Settings** dialog box, click the **Trusted hosts** tab.
6. Perform the appropriate action:

If...	Then...
You want to specify that all hosts are trusted hosts on your system and you want to allow access for all hosts	Click **Any host**.
You want to specify that no host is trusted on your system and you want to disable access for all hosts	Click **None**.
You want to specify that some hosts are trusted on your system and you want to restrict access to specific hosts	**a.** Click **Selected hosts**. **b.** Add the host name or IP addresses of the hosts.

7. Click **OK** to save your changes.

Related references

Password/RSH window on page 240

Window descriptions

Password/RSH window

The Password/RSH window enables you to view trusted hosts and RSH settings for your system. You can use the window command buttons to change your system password and modify your trusted hosts and RSH settings.

- *Command buttons* on page 240
- *Trusted hosts lists* on page 241
- *RSH settings* on page 241

Command buttons

Edit	Opens the Edit Security Settings dialog box, which enables you to add and delete trusted hosts and change your RSH settings.
Change password	Opens the Reset Password dialog box, which enables you to change your system password.

Refresh	Updates the information in the window.

Trusted hosts lists

Host name/IP address	Displays the host name or IP address for hosts that are designated as trusted.

RSH settings

Host name/IP address	Displays the host name or IP address for RSH host.
User ID	Displays the user ID that is required to establish the RSH session with the host.

Related tasks

Changing the system password on page 238
Adding or deleting RSH host names on page 239
Managing trusted hosts on page 239

Security > SSH/SSL

Understanding SSH and SSL

SSL certificates

SSL uses a certificate to provide a secure connection between the storage system and a Web browser. An SSL certificate enables encryption of sensitive information during online transactions. Each SSL certificate contains unique, authenticated information about the certificate owner. A Certificate Authority verifies the identity of the certificate owner when it is issued.

Secure protocols and storage system access

Using secure protocols improves the security of your storage system by making it very difficult for someone to intercept a storage system administrator's password over the network, because the password and all administrative communication are encrypted.

If your storage system does not have secure protocols enabled, you can set up SecureAdmin, which provides a secure communication channel between a client and the storage system by using one or both of the following protocols—SSH and SSL.

Note: SecureAdmin is set up automatically on storage systems shipped with Data ONTAP 8.0 or later.

- Secure Shell (SSH) protocol

SSH provides a secure remote shell and interactive network session.

- Secure Sockets Layer (SSL) protocol
  SSL provides secure web access for Data ONTAP APIs.

Understanding the SSH protocol

The Secure Shell (SSH) protocol performs public-key encryption using a host key and a server key. SSH improves security by providing a means for the storage system to authenticate the client and by generating a session key that encrypts data sent between the client and storage system.

The SSH server version running on Data ONTAP is Data ONTAP SSH version 1.0, which is equivalent to OpenSSH server version 3.4p1. For information about the Common Vulnerabilities and Exposures (CVE) fixes implemented in Data ONTAP, see the Suspected Security Vulnerabilities page on the NetApp Support Site.

Data ONTAP supports the SSH 1.x protocol and the SSH 2.0 protocol.

Data ONTAP supports the following SSH clients:

- OpenSSH client version 4.4p1 on UNIX platforms
- SSH Communications Security client (SSH Tectia client) version 6.0.0 on Windows platforms
- Vandyke SecureCRT version 6.0.1 on Windows platforms
- PuTTY version 0.6.0 on Windows platforms
- F-Secure SSH client version 7.0.0 on UNIX platforms

SSH uses three keys to improve security:

- Host key
  SSH uses the host key to encrypt and decrypt the session key. You determine the size of the host key, and Data ONTAP generates the host key when you configure SecureAdmin.

  Note: SecureAdmin is set up automatically on storage systems shipped with Data ONTAP 8.0 or later.

- Server key
  SSH uses the server key to encrypt and decrypt the session key. You determine the size of the server key when you configure SecureAdmin. If SSH is enabled, Data ONTAP generates the server key when any of the following events occur:

  - You start SecureAdmin
  - An hour elapses
  - The storage system reboots

- Session key
  SSH uses the session key to encrypt data sent between the client and storage system. The session key is created by the client. To use the session key, the client encrypts the session key using the host and server keys and sends the encrypted session key to the storage system, where it is decrypted using the host and server keys. After the session key is decrypted, the client and storage system can exchange encrypted data.

The following table shows how Data ONTAP creates a secure session between the storage system and client.

Stage	What the client does	What the storage system does
1	The client sends an SSH request to the storage system.	The storage system receives the SSH request from the client.
2		The storage system sends the public portion of the host key, and the server key if SSH 1.x is used, to the client.
3	The client stores the public portion of the host key for future host authentication.	
4	The client generates a random session key.	
5	The client encrypts the session key by using the public portion of the host key, and the server key if SSH 1.x is used, and sends it to the storage system.	
6		The storage system decrypts the session key using the private portions of the host key, and the server key if SSH 1.x is used.
7	The storage system and the client exchange information that they encrypt and decrypt using the session key.	

Note: Some characters, such as question mark (?), period (.), asterisk (*), and caret (^), can have special meaning for the command interpreter running on the client. The client command interpreter might replace the character with an environment-specific value prior to passing it to the SSH program. To prevent a replacement, you can use an escape sequence before the character (`ssh` *`ip_address`* `\?`) or enclose the character in quotes (`ssh` *`ip_address`* `'?'`).

Data ONTAP supports password authentication and public-key-based authentication. It does not support the use of a `.rhosts` file or the use of a `.rhosts` file with RSA host authentication.

Data ONTAP supports the following encryption algorithms:

- RSA/DSA 1024 bit
- 3DES in CBC mode
- HMAC-SHA1
- HMAC-MD5

Related information

Suspected Security Vulnerabilities page: support.netapp.com/NOW/knowledge/docs/olio/scanner_results

The SSL protocol

The Secure Sockets Layer (SSL) protocol improves security by providing a digital certificate that authenticates storage systems and allows encrypted data to pass between the system and a browser. SSL is built into all major browsers. Therefore, installing a digital certificate on the storage system enables the SSL capabilities between system and browser.

Data ONTAP supports SSLv2, SSLv3, and Transport Layer Security version 1.0 (TLSv1.0). You should use TLSv1.0 or SSLv3 because it offers better security protections than previous SSL versions.

As a precautionary measure due to security vulnerability CVE-2009-3555, the SSL renegotiation feature is disabled in Data ONTAP.

How to manage SSL

SSL uses a certificate to provide a secure connection between the storage system and a Web browser. If your storage system does not have SSL enabled, you can set up SecureAdmin to enable SSL and allow administrative requests over HTTPS to succeed.

SecureAdmin is set up automatically on storage systems shipped with Data ONTAP 8.0 or later. For these systems, Secure protocols (including SSH, SSL, and HTTPS) are enabled by default, and nonsecure protocols (including RSH, Telnet, FTP, and HTTP) are disabled by default.

Two types of certificates are used—self-signed certificate and certificate-authority-signed certificate.

- Self-signed certificate
  A certificate generated by Data ONTAP. Self-signed certificates can be used as is, but they are less secure than certificate-authority signed certificates, because the browser has no way of verifying the signer of the certificate. This means the system could be spoofed by an unauthorized server.
- Certificate authority (CA) signed certificate
  A CA-signed certificate is a self-signed certificate that is sent to a certificate authority to be signed. The advantage of a certificate-authority-signed certificate is that it verifies to the browser that the system is the system to which the client intended to connect.
  To enhance security, starting with Data ONTAP 8.0.2, Data ONTAP uses the SHA256 message-digest algorithm to generate digital certificates (including CSRs and root certificates) on the storage system.

Public-key-based authentication

Setting up key-based authentication requires an RSA key pair (a private and public key) in addition to the host and server keys. Public-key-based authentication differs between the two versions of SSH; SSH 1.x uses an RSA key pair and SSH 2.0 uses a DSA key pair in addition to an RSA key pair.

For both versions of SSH, you must generate the key pairs and copy the public key to the storage system.

Managing SSH and SSL

Enabling or disabling SSH

You can use SSH for authentication and secure communication between a client and the storage system. You can use the Edit SSH Settings dialog box to enable or disable the SSH protocol on your storage system.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Security**.
3. Click **SSH/SSL**.
4. In the **SSH Settings** area, click **Edit SSH**.
5. Either select or clear the check box for the SSH protocol version that you want to use.
6. Click **OK**.

Related references

SSH/SSL window on page 247

Generating SSH keys

You can use the Generate SSH Keys dialog box to generate a host key and a server key that are required for a secure connection between a client and your storage system.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **Security** > **SSH/SSL**.
3. In the **SSH Settings** area, click **SSH Setup**.

   Note: While setting up SSH, existing SSH settings are overwritten.
4. Select the check box to disable SSH.

   This check box is visible only if either or both versions of SSH are enabled.
5. Click **Setup** in the confirmation window.

Related references

SSH/SSL window on page 247

Editing SSH settings

You can enable or disable the SSH service for SSH 1.x clients and SSH 2.0 clients. You can specify the SSH idle sessions timeout to close an SSH connection if the connection is idle for a period of time.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Security > SSH & SSL**.
3. In the **SSH Settings** area, click **Edit SSH**.
4. Modify the settings as required and click **OK**.

Related references

SSH/SSL window on page 247

Enabling or disabling SSL

You can use SSL for secure communication between a client and the storage system. Enabling SSL allows administrative requests over HTTPS to succeed. Disabling SSL disallows all administrative requests over HTTPS.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Security**.
3. Click **SSH/SSL**.
4. In the **SSL area**, click either **Enable SSL** or **Disable SSL**, as appropriate.
5. Select the confirmation check box and click **Disable SSL**.

Related references

SSH/SSL window on page 247

Generating an SSL certificate

You can use the Generate SSL Certificate dialog box to generate a self-signed SSL certificate.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Security**.

3. Click **SSH/SSL**.
4. Click **SSL Certificate > Generate SSL Certificate**.
5. Type the required information in each field and click **Generate**.

Related references

SSH/SSL window on page 247

Installing an SSL certificate

You can use the Install SSL Certificate dialog box to browse to a CA signed certificate, or paste the contents of an SSL certificate file.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Security**.
3. Click **SSH/SSL**.
4. Click **SSL Certificate > Install SSL certificate**.
5. Copy and paste the signed certificate into the text box and click **Install**.

Related references

SSH/SSL window on page 247

Window descriptions

SSH/SSL window

You can use the SSH/SSL window to configure the security of your storage system. You can also modify the Secure Shell (SSH) protocol settings or enable and disable the Secure Sockets Layer (SSL) protocol.

Command buttons

Edit SSH	Opens the Edit SSH Settings dialog box, which enables you to change your storage system's SSH settings.
Setup SSH keys	Generates the host key and the server key.
Enable/Disable SSL	Enables or disables SSL.
SSL Certificate	Allows you to generate, view, or install an SSL certificate. Select one of the following:

- Generate SSL Certificate
  Opens the Generate SSL Certificate dialog box, which enables you to specify information required to generate a self-signed SSL certificate.
- View CSR
  Opens the View SSL Certificate dialog box, which enables you to view a read-only Certificate Signing Request.
- Install CA signed certificate
  Opens the Install SSL Certificate dialog box, which enables you to install an SSL certificate on the SSL server.

Refresh Updates the information in the window.

Related tasks

Enabling or disabling SSH on page 245
Generating SSH keys on page 245
Editing SSH settings on page 246
Enabling or disabling SSL on page 246
Generating an SSL certificate on page 246
Installing an SSL certificate on page 247

System Tools > AutoSupport

Understanding AutoSupport

Overview of the AutoSupport feature

The AutoSupport feature monitors the storage system's operations and sends automatic messages to technical support to alert for potential system problems. If necessary, technical support contacts you at the email address that you specify to help resolve a potential system problem.

AutoSupport is enabled by default when you configure your storage system for the first time. AutoSupport begins sending messages to technical support 24 hours after AutoSupport is enabled. You can reduce the 24-hour period by upgrading or reverting the system, modifying the AutoSupport configuration, or changing the time of the system to be outside of the 24-hour period.

Some scenarios when AutoSupport messages are generated are when the storage system reboots or when events occur on the storage system that require corrective action from the system administrator technical support or by using the command-line interface.

AutoSupport messages can be sent by SMTP, HTTP, or HTTPS (Hypertext Transfer Protocol over Secure Sockets Layer). HTTPS is the default. If an AutoSupport message cannot be sent successfully, an SNMP trap is generated.

For more information about AutoSupport, see the NetApp Support Site.

Related information

NetApp Support Site: support.netapp.com

AutoSupport transport protocols

AutoSupport supports HTTPS, HTTP, and SMTP as the transport protocols for delivering AutoSupport messages to NetApp technical support. All of these protocols run on IPv4 or IPv6 based on the address family the name resolves to. If you enable AutoSupport messages to your internal support organization, those messages are sent by SMTP.

Protocol availability varies with the destination of the AutoSupport messages:

- If you enable AutoSupport to send messages to NetApp technical support, you can use any of the following transport protocols:

Protocol and port	Description
HTTPS on port 443	This is the default protocol. You should use this whenever possible. The certificate from the remote server is validated against the root certificate, unless you disable validation. The delivery uses an HTTP PUT request. With PUT, if the request fails during transmission, the request restarts where it left off. If the server receiving the request does not support PUT, the delivery uses an HTTP POST request.
HTTP on port 80	This protocol is preferred over SMTP. The delivery uses an HTTP PUT request. With PUT, if the request fails during transmission, the request restarts where it left off. If the server receiving the request does not support PUT, the delivery uses an HTTP POST request.
SMTP on port 25	You should use this protocol only if the network connection does not allow HTTPS or HTTP, because SMTP can introduce limitations on message length and line length.

- If you configure AutoSupport with specific e-mail addresses for your internal support organization, those messages are always sent by SMTP.

For example, if you use the recommended protocol to send messages to NetApp technical support and you also want to send messages your internal support organization, your messages would be transported via both HTTPS and SMTP, respectively.

The protocols require the following additional configuration:

- If you use HTTP or HTTPS to send AutoSupport messages to NetApp technical support and you have a proxy, you must identify the proxy's URL. If the proxy uses a port other than the default

port, which is 3128, you can specify the proxy's port. You can also specify a username and password for proxy authentication.

- If you use SMTP to send AutoSupport messages either to your internal support organization or to NetApp technical support, you must have an external mail server. The storage system does not function as a mail server—it requires an external mail server at your site to send mail. The mail server must be a host that listens on the SMTP port (25), and it must be configured to send and receive 8-bit Multipurpose Internet Mail Extensions (MIME) encoding. Example mail hosts include a UNIX host running an SMTP server such as the sendmail program and a Windows NT server running the Microsoft Exchange server. You can have one or more mail hosts.

No matter what transport protocol you use, you can use IPv4 or IPv6 addresses based on the address family that the name resolves to.

AutoSupport severity types

AutoSupport messages have severity types that help you understand the purpose of each message—for example, to draw immediate attention to a critical problem, or only to provide information.

Messages have one of the following severities:

- Critical: critical conditions
- Error: error conditions
- Warning: warning conditions
- Notice: normal but significant condition
- Info: informational message
- Debug: debug-level messages

If your internal support organization receives AutoSupport messages via email, the severity appears in the subject line of the email message.

Configuring AutoSupport

Setting up AutoSupport

You can use the Edit AutoSupport Settings dialog box to specify an email address from which email notifications are sent and add multiple email host names.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > System Tools > AutoSupport**.
3. Click **Edit**.
4. In the **E-mail Recipient** tab, type the email address from which email notifications are sent, specify the email recipients and the message content for each email recipient, and add the mail hosts.

   Note: You can add up to five email addresses of the host names.

5. In the **Others** tab, select a transport protocol for delivering the email messages from the drop-down list and specify the HTTP or HTTPS proxy for HTTP.
6. Click **OK**.
7. Verify the configuration you have set for AutoSupport.

Related references

AutoSupport window on page 252

Managing AutoSupport

Enabling or disabling AutoSupport

You can enable or disable AutoSupport on your storage system. AutoSupport is enabled by default.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > System Tools > AutoSupport**.
3. Either select **Enable** or **Disable**, as required.
4. Click **OK**.
5. Verify that the AutoSupport status correctly displays the change you made.

Related references

AutoSupport window on page 252

Adding AutoSupport email recipients

You can use the **E-mail recipient** tab to add email addresses of recipients of AutoSupport notifications.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > System Tools > AutoSupport**.
3. Click **Edit**.
4. In the **E-mail recipient** tab, type the address of the email recipient, specify whether the recipient receives a full message or a short message, and click **Add**.
5. Click **OK**.
6. Verify that the details you specified are displayed in the **AutoSupport** window.

Related references

AutoSupport window on page 252

Testing AutoSupport

You can use the AutoSupport Test dialog box to test the AutoSupport configuration.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > System Tools > AutoSupport**.
3. Click **Test**.
4. In the AutoSupport subject box, enter the text *Test AutoSupport* or any text that notifies the recipients that you are testing AutoSupport.
5. In the **AutoSupport Test** dialog box, click **Test**.

Result

An email message with the subject "Test AutoSupport" or the text that you typed in the AutoSupport subject box is sent to the specified recipients.

Related references

AutoSupport window on page 252

Window descriptions

AutoSupport window

The AutoSupport window enables you to view the current AutoSupport settings for your system. You can also change your system's AutoSupport settings.

Command buttons

Enable	Enables AutoSupport notification.
Disable	Disables AutoSupport notification.
Edit	Opens the Edit AutoSupport Settings dialog box, which enables you to specify an email address from which email notifications are sent and to add multiple email addresses of the host names.
Test	Opens the AutoSupport Test dialog box, which enables you to generate an AutoSupport test message.
Refresh	Updates the information in the window.

Details area

The details area displays AutoSupport setting information such as the status of AutoSupport, the transport protocol used, and the name of the proxy server.

Related tasks

Setting up AutoSupport on page 250
Enabling or disabling AutoSupport on page 251
Adding AutoSupport email recipients on page 251
Testing AutoSupport on page 252

System Tools > DateTime

Understanding date and time management

Guidelines for setting system date and time

Keeping the system date and time correct is important to ensure that the storage system can service requests correctly.

To automatically keep your storage system time synchronized, you need the name of at least one time server. For best results, supply the name of more than one time server if one becomes unavailable.

There are two protocols you can use for time synchronization: SNTP and rdate. SNTP (Simple Network Time Protocol) is more accurate; therefore, it is the preferred protocol.

If you cannot access an SNTP server, you can use rdate. Many UNIX servers can function as an rdate server; work with your system administrator to set up or identify an rdate server in your environment.

Configuring date and time settings

Setting the date, time, and time zone for storage systems

You can use the Edit DateTime dialog box to manually set the date, time, and time zone for your storage system. However, for an HA configuration, you cannot modify the date, time, and time zone settings for the failed node or the partner node after a takeover occurs.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **System Tools** > **DateTime**.
3. Click **Edit**.
4. Select the time zone.

5. If you want to manually specify the date and time for your storage system, select **Manual**, and specify the date and time.
6. If you want to use a time daemon to set the date and time, select **Automatic**.
   a) Select either **SNTP** or **RDate** as the time protocol.

      Note: Starting with Data ONTAP 8.0, Network Time Protocol (NTP) is the only supported protocol for time synchronization.

   b) Specify up to five time servers to synchronize the time.

      Note: The **Up**, **Down**, and **Delete** buttons are unavailable if you delete all the time servers from the list.
7. Click **OK**.
8. Verify the changes you made to the date and time settings.

Related references

DateTime window on page 254

Window descriptions

DateTime window

The DateTime window enables you to view the current date and time settings for your storage system. You can also change your system's date and time settings.

Command buttons

Edit	Opens the Edit DateTime dialog box, which enables you to manually set the date, time, and time zone for your storage system.
Refresh	Updates the information in the window.

Details area

The details area displays information about the date, time, and time zone for your storage system.

Related tasks

Setting the date, time, and time zone for storage systems on page 253

System Tools > Licenses

Understanding licenses

License requirements

System Manager is an unlicensed application and is free to download, install, and use. However, you require storage system software licenses to enable certain services and features on your storage system, such as NFS.

Depending on the platform model, some features require license keys. A license key enables you to unlock and use a single product or multiple products. License keys are provided on a per-system basis and must be added on each system for features to work correctly. Some features do not require individual license keys; they are provided free of cost or along with other features when you install a license key for a software pack.

You can find license keys for your initial or add-on software orders at the NetApp Support Site under **My Support > Software Licenses**. For instance, you can search with a system's serial number to find all license keys for the system, and you can search with a sales order number to find license keys for all systems on the order. If you cannot locate your license keys from the Software Licenses page, you should contact your sales or support representative.

Preinstalled software licenses

Many of the software licenses that are required for your storage system services and features should be installed on the storage system at the factory. Therefore, you should not have to enter the license code during initial setup of the storage system except for some special circumstances described below.

CIFS	The storage system requires a software license to enable CIFS service. The license is installed on the storage system at the factory per your order; therefore, the initial setup of your storage system does not involve entering license codes.
Fibre Channel Protocol (FCP)	Fibre Channel Protocol (FCP) is a service that enables you to manage Fibre Channel target adapters for use with LUNs. The storage system requires a software license to enable the FCP service. You are provided with the appropriate license codes when your storage system or software is shipped from the factory or when you are provided instructions for obtaining the software over the Internet.
HTTP	The HTTP software license is required to enable HTTP service.
NFS	The storage system requires a software license to enable NFS services. The license is installed on the storage system at the factory per your order; therefore, you should not have to enter the license code for this software.

SnapRestore	SnapRestore enables you to revert a volume or file quickly to the state it was in when a particular Snapshot copy was created. The storage system requires a license to enable the SnapRestore service.
UNIX	The storage system requires a UNIX software license to enable NFS services.
Windows Shares (CIFS)	The storage system requires a software license to enable CIFS service. The license is installed on the storage system at the factory per your order; therefore, you should not have to enter the license code for this software.

Software licenses that must be installed

The following software license must be installed to support the iSCSI feature.

Note: For high availability configurations, you must install licenses on both the systems.

iSCSI The iSCSI service enables you to manage adapters that support the iSCSI protocol on your storage system. The storage system requires a software license to enable the iSCSI service. You are provided with the appropriate license codes when your storage system or software is shipped from the factory or when you are given instructions for obtaining the software over the Internet.

Related information

NetApp Support Site: support.netapp.com

Managing licenses

Adding licenses

If your storage system software was installed at the factory, System Manager automatically adds the software to its list of licenses. If the software was not installed at the factory or if you want to add additional software licenses, you can add the software license through the Add Licenses dialog box.

Before you begin

- The software license code for the specific Data ONTAP service must be available.
- In an HA configuration, both nodes must have the software license installed for the same functionality, such as SyncMirror, CIFS, or NFS.

  If a takeover occurs, the takeover node can provide only the functionality for the licenses installed on it. If the takeover node does not have a license that was used by the partner node to serve data, your HA configuration loses the functionality after a takeover.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > System Tools > Licenses**.

3. In the **Licenses** window, click **Add**.
4. In the **Add Licenses** dialog box, enter the software license key and click **Add**.

   You can add multiple licenses by entering the software license keys, separated by commas.
5. Click **Close**.

Result

The software license is added to your storage system. The newly added license is also included in the list of licenses in the Licenses window.

Related references

Licenses window on page 258

Deleting licenses

You can use the Licenses window to delete any software license installed on a storage system.

Before you begin

The software license you want to delete must not be used by any service or feature.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > System Tools > Licenses**.
3. In the **Licenses** window, select the software license that you want to delete, and then click **Delete**.
4. Select the confirmation check box and click **Delete**.

Result

The software license is deleted from your storage system. The deleted license is also removed from the list of licenses in the Licenses window.

Related references

Licenses window on page 258

Enabling or disabling licenses

For systems running Data ONTAP 8.1 or later operating in 7-Mode, you can enable or disable certain licensed services, making them available or unavailable for the storage system.

Steps

1. From the **Home** tab, double-click the appropriate storage system.

2. In the navigation pane, click **Configuration > System Tools > Licenses**.
3. Select the licensed service that you want to enable or disable.
4. Click either **Enable** or **Disable**, as required.

Related references

Licenses window on page 258

Window descriptions

Licenses window

Your storage system arrives from the factory with pre-installed software. If you want to add or remove a software license after you receive the storage system, you can use the Licenses window.

Command buttons

Add	Opens the Add License window, which enables you to add new software licenses.
Delete	Deletes the software license that you select in the software license list.
Refresh	Updates the information in the window.
Enable	Enables the software license that you select in the software license list.
Disable	Disables the software license that you select in the software license list.
Package	Displays the name of the license package.
State	Displays the state of the software license, whether it is enabled, disabled, or not applicable.
Description	Displays information about the license package—for example, the type of license.

Related tasks

Adding licenses on page 256
Deleting licenses on page 257
Enabling or disabling licenses on page 257

System Tools > SNMP

Understanding SNMP

What the SNMP agent does

The storage system includes an SNMP agent that responds to queries and sends traps to network management stations.

The SNMP agent on the storage system has read-only privileges—that is, it cannot be used to take corrective action in response to a trap.

Note: Starting with Data ONTAP 7.3.1, the SNMP agent supports IPv6 transport.

How to configure the SNMP agent

You need to configure the SNMP agent on your storage system to set SNMP values and parameters.

To configure the SNMP agent on your storage system, you need to perform the following tasks:

- Verify that SNMP is enabled.

  Note: SNMP is enabled by default in Data ONTAP.
- If you are running SNMPv3, configure SNMPv3 for read-only access.
- Enable traps.
  Although SNMP is enabled by default, traps are disabled by default.
- Specify host names of one or more network management stations.
  Traps can only be sent when at least one SNMP management station is specified as a traphost. Trap notifications can be sent to a maximum of eight network management stations.

  Note: The SNMP agent can send traps over IPv6 transport to the traphosts whose IPv6 address is configured on the storage system. You can specify traphosts by their IPv6 addresses, but not by their host names.

You can perform the following tasks after configuring SNMP:

- Provide courtesy information about storage system location and contact personnel.
- Specify SNMP communities.
  Community strings function as group names to establish trust between SNMP managers and clients. Data ONTAP supports only read-only communities.

  Note: No more than eight communities are allowed.

Note: Storage systems in an HA configuration can have different SNMP configurations.

Configuring SNMP

Setting SNMP information

You can use the Edit SNMP Settings dialog box to update information about the storage system location and contact personnel, and to specify SNMP communities.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > System Tools > SNMP**.
3. Click **Edit**.
4. In the **General** tab, specify the storage system contact personnel and location, and SNMP communities.
5. Click **OK**.
6. Verify the changes you made to the SNMP settings.

Related references

SNMP window on page 261

Managing SNMP

Enabling or disabling SNMP traps

SNMP traps enable you to monitor the health and state of various components of the storage system. You can use the **Trap hosts** tab to enable or disable SNMP traps on your storage system. Although SNMP is enabled by default, traps are disabled by default.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > System Tools > SNMP**.
3. Click **Edit**.
4. In the **Trap hosts** tab, either select or clear **Enable traps**.
5. If you enable SNMP traps, add the host name or IP address of the hosts to which the traps are sent.
6. Click **OK**.

Related references

SNMP window on page 261

Window descriptions

SNMP window

The SNMP window enables you to view the current SNMP settings for your system. You can also change your system's SNMP settings.

Command buttons

Edit	Opens the Edit SNMP Settings dialog box, which enables you to specify SNMP communities and enable or disable traps.
Refresh	Updates the information in the window.

Details

The details area displays information about the status of SNMP and traps for your storage system.

Related tasks

Setting SNMP information on page 260
Enabling or disabling SNMP traps on page 260

System Tools > NDMP

Understanding NDMP

NDMP management

The Network Data Management Protocol (NDMP) is a standardized protocol for controlling backup, recovery, and other types of data transfer between primary and secondary storage devices, such as storage systems and tape libraries.

By enabling NDMP protocol support on a storage system, you enable that storage system to carry out communications with NDMP-enabled commercial network-attached backup applications (also called *Data Management Applications* or *DMAs*), data servers, and tape servers participating in backup or recovery operations. NDMP also provides low-level control of tape drives and medium changers.

Configuring NDMP service

Enabling or disabling the NDMP service

Enabling the NDMP service on your storage system allows NDMP-compliant data protection applications to communicate with the storage system. After you disable the NDMP service, the

storage system continues processing all requests on already established sessions, but rejects new sessions.

Before you begin

The storage system must be running Data ONTAP 8.1 or later operating in 7-Mode.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > System Tools > NDMP**.
3. Click either **Enable** or **Disable**, as required.
4. If you are disabling the NDMP service, select the confirmation check box and click **Disable**.

Related references

NDMP window on page 263

Managing NDMP service

Stopping NDMP sessions

You can stop an NDMP session if the session is not responding. The specified session stops processing its current requests and moves to an inactive state. This allows hung sessions to be cleared without requiring a reboot.

Before you begin

The storage system must be running Data ONTAP 8.1 operating in 7-Mode.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > System Tools > NDMP**.
3. Select the NDMP session that you want to stop and click **Terminate Session**.

Related references

NDMP window on page 263

Window description

NDMP window

You can use the NDMP window to enable the NDMP service and to view the active NDMP sessions for your system.

Command buttons

Enable	Enables NDMP service.
Terminate Session	Terminates NDMP sessions.
Refresh	Updates the information in the window.

Related tasks

Enabling or disabling the NDMP service on page 261
Stopping NDMP sessions on page 262

System Tools > Halt/Reboot

Halting storage systems

You can use the **Halt and Reboot** window to halt or shut down a storage system. You may shut down a storage system to perform maintenance on it.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **System Tools**.
3. Click **Halt and Reboot**.
4. Perform the appropriate action:

If...	Then...
You want to allow clients to terminate connections and perform a clean shutdown of the storage system after an interval of time	Select **Wait for clients** and specify the time.
You want the storage system to perform a core dump, without flushing cached data, before halting	Select **Dump core**.

5. Click **Halt**.
6. Select the check box in the confirmation window and click **Halt**.

Related references

Halt/Reboot window on page 264

Rebooting storage systems

Rebooting a storage system is commonly performed to allow modified configuration files to take effect or to run a newly installed version of Data ONTAP. You can use the **Halt and Reboot** window to reboot a storage system. Rebooting stops and then restarts the storage system.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration** > **System Tools**.
3. Click **Halt and Reboot**.
4. Perform the appropriate action:

If...	Then...
You want to allow clients to terminate connections gracefully and specify the time before rebooting	Select **Wait for clients** and specify the time.
You want the storage system to perform a core dump before rebooting	Select **Dump core**.

5. Click **Reboot**.
6. Select the check box in the confirmation window and click **Reboot**.

Related references

Halt/Reboot window on page 264

Window descriptions

Halt/Reboot window

You can use the Halt/Reboot window to halt or reboot a storage system.

Command buttons

Halt Halts a storage system. You can halt a storage system to perform maintenance on it.

Reboot Reboots a storage system. You can reboot a storage system to allow modified configuration files to take effect or to run a newly installed version of Data ONTAP.

Related tasks

Halting storage systems on page 263
Rebooting storage systems on page 264

Flash Cache

Understanding Flash Cache

Increasing WAFL cache memory

You can increase Write Anywhere File Layout (WAFL) cache memory in a system that has a caching module installed (Performance Acceleration Module (PAM), Flash Cache module, or Flash Cache 2 module). To increase the WAFL cache memory, you use the WAFL external cache, a software component of Data ONTAP.

WAFL external cache provides extra WAFL cache memory to improve the performance of the storage system by reducing the number of disk reads. You can control how user data blocks are cached by changing the mode of operation for a caching module. You can keep the default mode (normal user data blocks) or you can choose metadata mode or low-priority blocks mode.

To change the mode, you need to use the Data ONTAP CLI by running the `options flexscale` commands. For more information, see the *Data ONTAP System Administration Guide for 7-Mode*.

You should verify that the WAFL external cache functionality is enabled after you install a caching module.

Note: WAFL external cache does not require a separate license if your system is running Data ONTAP 8.1 or later.

Note: Not all systems have a caching module installed. Therefore, not all systems can utilize the WAFL external cache functionality.

WAFL external cache does not cache data that is stored in a RAID group composed of SSDs.

If you use WAFL external cache on storage systems with a high-availability configuration, you must ensure that the WAFL external cache options are the same on both nodes. Otherwise, a takeover can result in lower performance due to the lack of WAFL external cache on the remaining node.

Besides the Data ONTAP options that you can use to manage WAFL external cache, a diagnostic command is available for sanitizing a caching module. For more information, see the *Diagnostics Guide*.

Managing Flash Cache

Enabling or disabling Flash Cache

You can enable or disable the WAFL external cache functionality for a storage system that has a PAM or Flash Cache module installed.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Configuration > Flash Cache**.
3. Click **Enable** or **Disable**, as required.

Window descriptions

Flash Cache window

You can use the Flash Cache window to enable or disable Flash Cache for a storage system that has a Flash Cache module installed. You can also view the read workload statistics.

Command buttons

Enable/Disable	Enables or disables Flash Cache.

Flash Cache Read Workload

Displays a graph specifying the rate of read workload served by the disks and the Flash Cache module, thereby indicating the performance of the storage system.

Details area

Displays information about the system read latency (in seconds), the caching mode that specifies the caching configuration, the state of Flash Cache (enabled or disabled), and the size of the Flash Cache (in GB). If there are multiple Flash Cache cards, the total cache size from all the cards is displayed.

How Flash Cache works

Using Flash Caches improves the performance of a storage system.

You can configure Flash Cache and disks based on the workload requirements of a storage system. By determining the read workload (number of read operations) served by Flash Cache and disks, you can analyze the performance of the storage system.

Flash Cache does not contain any data during storage system boot or when control is returned to the storage system after a takeover event. Therefore, disks serve all the data read requests of the storage system.

The Flash Cache module is slowly populated with data when data read requests are served. Because the data read requests served by Flash Cache are faster than those served by the disks, the performance of the storage system improves.

Data read requests served by the Flash Cache module replace the data read requests served by the disks and, therefore, the performance improvement in the storage system is directly related to the disk reads that are replaced. To understand the impact of Flash Cache on storage system performance, you must view the read workload graph when the Flash Cache contains data.

Diagnostics

CIFS

Understanding CIFS diagnostics

CIFS diagnostics

You can view current CIFS activities and statistics for a selected storage system in the **Diagnostics CIFS** window.

CIFS client monitoring

If you enable per-client monitoring, the application can display client-based CIFS activities. The output can be sorted by client name, operations per second, read operations, read size per second, suspicious events per second, write operations, and write size per second.

Note: Enabling CIFS client monitoring might impact system performance.

CIFS statistics

If you click the CIFS Diagnostics window **Statistics** button, the application displays a copy of the current counts and percentages of all CIFS operations and a number of internal statistics that might be used when diagnosing performance and other problems.

If the per-client flag is on, you can query a user or a host CIFS statistic. If more than one match is found, the application lists all the matched users or host names and the sum of their statistics. You can reset all CIFS operation counters, including per-client counters to zero.

Note: Enabling CIFS statistics queries might impact system performance.

Monitoring CIFS diagnostics

Monitoring CIFS diagnostics

You can view current CIFS activities and statistics for a selected storage system. You can sort the output by client name, operations per second, read operations, read size per second, suspicious events per second, write operations, and write size per second.

Before you begin

CIFS must be licensed and enabled on the storage system.

About this task

You can view the CIFS statistics if you are using Internet Explorer as your browser. However, if you are using Firefox as your browser, you have to view the CIFS statistics from the CLI.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Diagnostics > CIFS**.

   CIFS monitoring is started for the selected storage system in the CIFS Diagnostics window. Monitoring continues until you select a different storage system.
3. Click **Statistics** and view detailed CIFS statistics.
4. If you want to enable CIFS statistics per client, click **Edit**, select **Enable CIFS statistics per client**, and then click **OK**.

   Note: The per-client statistics feature is turned off by default. This feature tracks counts and percentages for non-blocking and blocking CIFS operations. Because of the quantity of information, this feature might affect system performance.

Related references

CIFS diagnostics window on page 269

Window descriptions

CIFS diagnostics window

You can use the CIFS diagnostics window to view current information about CIFS activities.

Command buttons

Statistics	Opens the CIFS Statistics dialog box for the selected storage system.
Edit	Opens the Edit Diagnostic dialog box.
Refresh	Updates the information in the window.

CIFS diagnostics list

User information	Displays the client IP address or host name.
Operations/sec	Displays the CIFS operations per second for the client.
Read Operations (KB/sec)	Displays the total number of read operations for the client.
Read size/sec	Displays the rate for read operations per second.

Suspicious events/sec	Displays the number of suspicious events per second.
Write operations (KB/sec)	Displays the number of write operations per second.
Write size/sec	Displays the rate for write operations per second.

Related tasks

Monitoring CIFS diagnostics on page 268

Session

Viewing sessions

You can monitor all of the CIFS sessions activity on your storage system and view session information in the Sessions window. You can view the volumes accessed and names of shares and files opened by connected users.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Diagnostics > Session**.

Related references

Session window on page 270

Window descriptions

Session window

You can use the Session window to view detailed information on your system's CIFS sessions.

Command buttons

Refresh	Updates the information in the window.

Session list

The top table in the Session window displays a list of all current sessions on your system.

User	Specifies the name of the user for the session.
Computer	Specifies the name of the user's computer.
IP address	Specifies the IP address for the user's computer.

# Open shares	Specifies the number of open shares.
# Open directories	Specifies the number of open directories.
# Open files	Specifies the number of open files.

Accessed volume list

This pane provides a list of volumes accessed by the selected user.

Related tasks

Viewing sessions on page 270

System Health

Understanding system health

How you can respond to system health alerts

When a system health alert occurs, you can learn more about it, acknowledge it, repair the underlying condition, and prevent it from occurring again.

When a health monitor raises an alert, you can respond in any of the following ways:

- Get information about the alert, which includes the affected resource, alert severity, probable cause, possible effect, and corrective actions.
- Get detailed information about the alert, such as the time when the alert was raised and whether anyone else has acknowledged the alert already.
- Get health-related information about the state of the affected resource or subsystem, such as a specific shelf or disk.
- Acknowledge the alert to indicate that someone is working on the problem, and identify yourself as the "Acknowledger."
- Resolve the problem by taking the corrective actions provided in the alert, such as fixing cabling to resolve a connectivity problem.
- Delete the alert, if the system didn't automatically clear it.
- Suppress an alert to prevent the system from notifying you about the same alert again, and identify yourself as the "Suppressor."

  Suppressing is useful when you understand a problem. After you suppress an alert, it can still occur but the subsystem health remains OK even when the alert occurs.

What health monitors are available

In addition to the overall System Health Monitor, there currently is one individual health monitor called Node Connectivity, which is for the Storage subsystem.

Health monitor name (identifier)	Subsystem name (identifier)	Purpose
Node Connectivity (node-connect)	Storage (SAS-connect)	Monitor shelves, disks, and adapters at the node level to ensure that they have appropriate pathing and connections.
System	n/a	Aggregate other health monitors

Monitoring the health of your system

You can proactively manage your system by monitoring a single, integrated health status.

If the status is degraded, you can view details about the problem, including the probable cause and recommended recovery actions. After you resolve the problem, the system health status automatically returns to OK.

The system health status reflects multiple separate health monitors. A degraded status in an individual health monitor causes a degraded status for the overall system health.

Currently, there are two health monitors: an overall System Health Monitor and a Node Connectivity health monitor for the Storage subsystem.

Acknowledging system health alerts

For storage systems running Data ONTAP 8.1, you can acknowledge and respond to system health alerts for SAS connectivity from System Manager. You can use the information displayed to take the recommended action and correct the problem reported by the alert.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Diagnostics > System Health**.
3. Click the arrow icon next to the name of subsystem.
4. Select the alert that you want to acknowledge and click **Acknowledge**.
5. Type your name and click **Acknowledge**.

Related references

System Health window on page 274

Suppressing system health alerts

You can suppress system health alerts that do not require any intervention from you.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Diagnostics > System Health**.
3. Click the arrow icon next to the name of subsystem.
4. Select the alert that you want to suppress and click **Suppress**.
5. Type your name and click **Suppress**.

Related references

System Health window on page 274

Deleting system health alerts

You can delete system health alerts that you have already responded to.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Diagnostics > System Health**.
3. Click the arrow icon next to the name of subsystem.
4. Select the alert that you want to delete and click **Delete**.
5. Click **OK**.

Related references

System Health window on page 274

Window descriptions

System Health window

You can use the System Health window to learn more about system health alerts. You can also acknowledge, delete, and suppress alerts from the window.

Command buttons

Acknowledge	Enables you to acknowledge the selected alert to indicate that the problem is being addressed and identifies the person who clicks the button as the "Acknowledger."
Suppress	Enables you to suppress the selected alert to prevent the system from notifying you about the same alert again and identifies you as the "Suppressor."
Delete	Deletes the selected alert.
Refresh	Updates the information in the window.

Alerts list

SubSystem (No. of Alerts)	Displays the name of the subsystem for which the alert is generated. Only SAS connection subsystems are supported.
Alert ID	Displays the alert ID.
Node	Displays the name of the node for which the alert is generated.
Severity	Displays the severity of the alert as Unknown, Other, Information, Degraded, Minor, Major, Critical, or Fatal.
Resource	Displays the resource that generated the alert, such as a specific shelf or disk.
Time	Displays the time when the alert was generated.

Details area

The details area displays detailed information about the alert, such as the time when the alert was generated and whether the alert has been acknowledged. The area also includes information about the probable cause and possible effect of the condition generated by the alert, and the recommended actions to correct the problem reported by the alert.

Related tasks

Acknowledging system health alerts on page 272
Suppressing system health alerts on page 273
Deleting system health alerts on page 273

Flash Pool Statistics

Window descriptions

Flash Pool Statistics window

You can view the real-time SSD tier read and write workloads for a selected Flash Pool.

Displaying Statistics for Flash Pool	From the list of Flash Pools, select the Flash Pool whose statistics you want to view.

SSD Cache Read Workload

Displays a graphical view of the total read requests that are sent to the Flash Pool in comparison with the read operations that are performed by the SSD tier.

SSD Cache Write Workload

Displays a graphical view of the total write requests that are sent to the Flash Pool in comparison with the write operations that are performed by the SSD tier.

Logs > Syslog

Understanding Syslog messages

What Syslog messages are

You can monitor the status and operation of managed storage systems by using the Event Management System (EMS) output in Syslog.

Events are generated automatically when a predefined condition occurs or when an object crosses a threshold. When an event occurs, status alert messages might be generated as a result of the event.

EMS is a subsystem in the Data ONTAP kernel where event indications are posted, and from which notification services, such as Syslog, monitor for individual event types. EMS collects event data from various parts of the Data ONTAP kernel and provides a set of filtering and event forwarding mechanisms.

The syslog.conf configuration file

Message logging is done by a `syslogd` daemon. By default, all system messages (except those with debug-level severity) are sent to the console and logged in the `/etc/messages` file.

The `/etc/syslog.conf` configuration file on the storage system's root volume is the configuration file for the syslogd daemon and it determines how system messages are logged.

Syslog messaging configuration options

You can configure which types of messages to log for a storage system, based upon your combinations of facility and severity level. The facility is the part of the system that is generating the message. For example, defining message type `kern.err`, invokes logging of all error level events from the kernel.

You can combine the following facilities with the available Syslog severity levels:

Facility	Definition
kern	Messages generated by the storage system kernel.
daemon	System daemons, such as the rshd daemon or the routing daemon.
auth	Authentication system messages, such as those logged for Telnet sessions.
cron	The storage system's internal cron facility.
local7	The storage system's audit logging facility. All messages coming from the audit logging facility are logged at level debug.
*	An asterisk acts as a wildcard and designates all facilities (except local7). For example, use `*.err` to see all messages with severity level err from all facilities (except local7).

Syslog message severity levels

The Syslog messages use a different scheme of severity levels than the System Manager monitoring. This is because the Syslog messages are based on EMS messages. The following table defines the possible Syslog message severity levels and shows how they relate to EMS severity levels.

Syslog severity	EMS severity	Description
*	Not applicable	An asterisk acts as a wildcard and designates all severity levels. For example, use `kern.*` to see all severity level messages generated by the kernel.
emerg	EMERGENCY	A panic condition that causes a disruption of normal service.
alert	ALERT	A condition that you should correct immediately, such as a failed disk.
crit	CRITICAL	A critical condition, such as a disk error.
err	ERROR	An error condition, such as a bad configuration file.

Syslog severity	EMS severity	Description
warning	WARNING	A condition that might become an error if not corrected.
notice	NOTICE	A condition that is not an error, but that might require special attention.
info	INFORMATION	Information, such as the hourly uptime message.
debug	DEBUG	Information used for diagnostic purposes.

Message logging locations

You can configure where a particular message type is logged. You can log messages in the following locations:

- A console (`dev/console`)
- A file (`/etc/messages`)
- A remote system (`@adminhost`)

Managing Syslog messages

Editing Syslog messaging configuration

You can use the Configure Syslog dialog box to edit an existing messaging configuration and specify how system messages are logged. By default, all system messages (except those with debug-level severity) are sent to the console and logged in the `/etc/messages` file.

About this task

The `/etc/syslog.conf` configuration file on the root volume of the storage system determines how system messages are logged.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Diagnostics > Logs > Syslog**.
3. Click **Edit**.
4. Click **Basic** and select the severity of system messages and specify where the messages are sent.
5. Click **Advanced** to directly modify the contents of the `/etc/syslog.conf` file.

   Note: If you click the **Basic** button after making changes to your messaging configuration, the contents of the advanced section are erased and replaced with the basic configuration.

6. Click **OK**.

Related references

Syslog window on page 278

Monitoring Syslog messages

Monitoring status using Syslog messages

You can monitor the status and operation of managed storage systems using the Syslog output.

Before you begin

The Syslog filters, the EMS events that you want notification of, and the locations for the output must be configured.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Diagnostics** > **Logs** > **Syslog**.
3. Sort the events in the table in the upper-right pane by clicking the column headings for severity level, name, date and time, or text.
4. Select one of the entries in the table to display EMS details for the event in the **EMS Details** pane.

   Details for the event are displayed, including the EMS source of the event, if it is available. If no EMS message is associated with the event, N/A (not applicable) is displayed.

Related references

Syslog window on page 278

Window descriptions

Syslog window

You can use the Syslog window to view Syslog messages.

- *Command buttons* on page 278
- *Syslog message list* on page 279
- *Details area* on page 279

Command buttons

Edit Opens the Configure Syslog dialog box, which enables you to change your messaging configuration.

Refresh Updates the information in the window.

Syslog message list

Severity	Sorts the message list by message severity level.
Event	Sorts the message list by the source EMS event for messages.
Date/Time	Sorts the list by the date and time of the event for messages.
Message	Sorts the list by the message text.

You can use the navigation toolbar at the bottom of the list to navigate to different records of the list. However, if you are managing storage systems running Data ONTAP 7.2.*x*, you can navigate only to the next page. Also, the page count might not be displayed.

Details area

The area below the Syslog message list displays details of the selected message, including a pointer to Syslog Translator.

Related tasks

Editing Syslog messaging configuration on page 277
Monitoring status using Syslog messages on page 278

Logs > Audit Log

Understanding audit log

Understanding audit logging

An audit log is a record of commands executed at the console, through a Telnet shell or an SSH shell, or by using the `rsh` command. All the commands executed in a source file script are also recorded in the audit log. Administrative HTTP operations are logged. All login attempts to access the storage system, with success or failure, are also audit-logged.

In addition, changes made to configuration and registry files are audited. Read-only APIs by default are not audited but you can enable auditing with the `auditlog.readonly_api.enable` option.

By default, Data ONTAP is configured to save an audit log. The audit log data is stored in the `/etc/log` directory in a file called `auditlog`.

For configuration changes, the audit log shows the following information:

- What configuration files were accessed
- When the configuration files were accessed
- What has been changed in the configuration files

For commands executed through the console, a Telnet shell, an SSH shell, or by using the `rsh` command, the audit log shows the following information:

- What commands were executed
- Who executed the commands
- When the commands were executed

The maximum size of the audit-log file is specified by the `auditlog.max_file_size` option. The maximum size of an audit entry in the audit-log file is 511 characters. An audit entry is truncated to 511 characters if it exceeds the size limit.

Every Saturday at midnight, the `/etc/log/auditlog` file is copied to `/etc/log/auditlog.0`, `/etc/log/auditlog.0` is copied to `/etc/log/auditlog.1`, and so on. This also occurs if the audit-log file reaches the maximum size specified by `auditlog.max_file_size`.

The system saves audit-log files for six weeks, unless any audit-log file reaches the maximum size, in which case the oldest audit-log file is discarded.

You can access the audit-log files using your NFS or CIFS client, or using HTTP.

Note: You can also configure auditing specific to your file access protocol. For more information, see the *Data ONTAP File Access and Protocols Management Guide for 7-Mode*.

For information about forwarding audit logs to a remote syslog log host, see the na_auditlog(5) man page.

Managing audit log

Enabling or disabling audit logging

You can record commands that are executed at the console in an audit log. The audit log enables system administrators track user actions and monitor system activity. By default, Data ONTAP is configured to save an audit log. You can enable or disable audit logging in the Audit Log window.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. In the navigation pane, click **Diagnostics > Logs > Audit Log**.
3. Click either **Enable** or **Disable**, as appropriate.

Related references

Audit Log window on page 281

Window descriptions

Audit Log window

You can use the Audit Log window to track user actions and monitor system activity.

Command buttons

Enable/Disable	Enables or disables audit logging.
Refresh	Updates the information in the window.

Audit log list

Type	Displays the message type.
Source	Displays the message source.
User name	Displays the names of the users who invoked the CLIs and APIs.
IP	Displays the IP address of the host where the user performed the action.
Date time	Displays the date and time of the action.
Application	Displays the application invoking the audit log facility.
Priority	Displays the priority of the message.

Details area

The details area displays information about the audit log such as the message and the priority of message.

Related tasks

Enabling or disabling audit logging on page 280

Logs > SnapMirror Log

Format of SnapMirror log files

Understanding the format of SnapMirror log files can help you better handle issues related to SnapMirror transfers.
The log file is in the following format:

```
type timestamp source_system:source_path dest_system:dest_path event_info
```

type can be one of the following: `src`, `dst`, `log`, `cmd`. *type* specifies whether the record is for the source side (`src`) or destination side (`dst`) of the transfer. Certain events apply to only one side. The

type `log` indicates a record about the logging system itself, for example, `Start_Logging` and `End_Logging`. The type `cmd` indicates a record of user commands, for example, `Release_command` and `Resync_command`.

`timestamp` is expressed in `ctime` format, for example:
`Fri Jul 27 20:41:09 GMT.`

`event_info` includes the following event names:
`Request (IP address | transfer type) Start Restart (@ num KB) End (num KB done) Abort (error_msg) Defer (reason) Rollback_start Rollback_end Rollback_failed Start_Logging End_Logging Wait_tape New_tape Snapmirror_on Snapmirror_off Quiesce_start Quiesce_end Quiesce_failed Resume_command Break_command Release_command Abort_command Resync_command Migrate_command`

The `Request` event on the source side includes the IP address of the system that made the transfer request; the `Request` event on the destination side includes the type of transfer. At the end of each successful transfer, the `End` event also reports the total size of the transfer in KB. Error messages are included with the `Abort` and `Defer` events.

Example

The following is an example of a log file from the source side:

```
log Fri Jul 27 20:00:01 GMT - - Start_Logging
cmd Fri Jul 27 20:00:20 GMT - - Snapmirror_on
src Fri Jul 27 20:41:09 GMT system1:vol1 system2:vol1 Request (10.56.17.133)
src Fri Jul 27 20:41:32 GMT system1:vol1 system2:vol1 Abort (Destination not allowed)
src Fri Jul 27 20:45:31 GMT system1:vol0 system1:vol1 Request (10.56.17.132)
src Fri Jul 27 20:45:35 GMT system1:vol0 system1:vol1 Start
src Fri Jul 27 20:51:40 GMT system1:vol0 system1:vol1 End (26200 KB)
src Fri Jul 27 22:41:09 GMT system1:/vol/vol1/qtA system2:/vol/vol1/qtB Request
(10.56.17.133)
src Fri Jul 27 22:41:12 GMT system1:/vol/vol1/qtA system2:/vol/vol1/qtB Start
src Fri Jul 27 22:41:13 GMT system1:/vol/vol1/qtA system2:/vol/vol1/qtB Abort (Non-
unicode directory found in source qtree.)
src Fri Jul 27 22:45:53 GMT system1:/vol/vol1/qtb system2:/vol/vol1/qsmb Request
(10.56.17.133)
src Fri Jul 27 22:45:56 GMT system1:/vol/vol1/qtb system2:/vol/vol1/qsmb Start
src Fri Jul 27 22:45:59 GMT system1:/vol/vol1/qtb system2:/vol/vol1/qsmb End (3800 KB)
cmd Fri Jul 27 22:50:29 GMT system1:/vol/vol1/qtb system2:/vol/vol1/qsmb
Release_command
```

Example

The following is an example of a log file from the destination side:

```
dst Fri Jul 27 22:50:18 GMT system1:vol0 system1:vol1 Request (Initialization)
dst Fri Jul 27 22:50:20 GMT system1:vol0 system1:vol1 Abort (Destination is not
restricted)
dst Fri Jul 27 22:57:17 GMT system1:/vol/vol1/qtA system2:/vol/vol1/qtB Request
(Initialize)
dst Fri Jul 27 22:57:24 GMT system1:/vol/vol1/qtA system2:/vol/vol1/qtB Start
dst Fri Jul 27 22:57:36 GMT system1:/vol/vol1/qtA system2:/vol/vol1/qtB End (55670 KB)
dst Fri Jul 27 23:10:03 GMT system1:/vol/vol1/qtA system2:/vol/vol1/qtB Request
(Scheduled)
dst Fri Jul 27 23:10:07 GMT system1:/vol/vol1/qtA system2:/vol/vol1/qtB Start
dst Fri Jul 27 23:10:18 GMT system1:/vol/vol1/qtA system2:/vol/vol1/qtB End (12900 KB)
cmd Sat Jul 28 00:05:29 GMT - system2:/vol/vol1/qtB Quiesce_start
```

```
cmd Sat Jul 28 00:05:29 GMT - system2:/vol/vol1/qtB Quiesce_end
cmd Sat Jul 28 00:05:40 GMT - system2:/vol/vol1/qtB Break_command
cmd Sat Jul 28 00:41:05 GMT system1:/vol/vol1/qtA system2:/vol/vol1/qtB Resync_command
log Sat Jul 28 00:41:10 GMT - - End_Logging
```

Example

The following is an example of a log file from a retrieve (from tape) request:

```
dst Fri Jun 22 03:07:34 GMT filer_1:rst0l filer_1:bigtwo Request (retrieve)
dst Fri Jun 22 03:07:34 GMT filer_1:rst0l filer_1:bigtwo Start
dst Fri Jun 22 05:03:45 GMT filer_1:rst0l filer_1:bigtwo Wait_tape
dst Fri Jun 22 15:16:44 GMT filer_1:rst0l filer_1:bigtwo New_tape
dst Fri Jun 22 17:13:24 GMT filer_1:rst0l filer_1:bigtwo Wait_tape
dst Fri Jun 22 17:56:43 GMT filer_1:rst0l filer_1:bigtwo New_tape
dst Fri Jun 22 18:10:37 GMT filer_1:rst0l filer_1:bigtwo End (98602256 KB)
```

Window description

SnapMirror Log window

You can use the SnapMirror Log window to load the SnapMirror log file and view the contents of the log. The log files record the SnapMirror data transfer history. The details help you to verify that the transfers are occurring as planned, and check how long the transfers take and how well the system setup works.

Command buttons

Load Loads the selected SnapMirror log file. The latest logs are stored in the file named `snapmirror`. The older logs are named `snapmirror.0` and `snapmirror.1`.

Note: There might be one or more SnapMirror log files.

Refresh Updates the information in the window.

SnapMirror log list

Source Displays the volume or qtree from which data is mirrored in a SnapMirror relationship.

Destination Displays the volume or qtree to which data is mirrored in a SnapMirror relationship.

Date time Displays the date and time of the SnapMirror operation.

Action Displays the name of the event.

Message Displays the message related to the event.

HA Configuration

Understanding HA configuration

HA configuration

System Manager includes several features that enable you to keep operating a storage system even if its partner system in an HA configuration stops functioning. Takeover is the process in which a node takes over the storage of its partner. Giveback is the process in which the storage is returned to the partner.

When one storage system in an HA configuration undergoes a system failure and cannot reboot, the partner system in the HA configuration can take over the failed system's functions and serve network clients the data on the failed system's disks. This is known as a takeover. Additionally, you can issue a manual takeover at any time, to force a takeover: for instance, to allow scheduled maintenance to be performed on a storage system.

After the failed partner is running normally again, you issue a giveback, which returns the identity from the emulated storage system to the failed system, resulting in a return to normal operation.

What an HA pair is

An HA pair is two storage systems (nodes) whose controllers are connected to each other either directly or, in the case of a fabric-attached MetroCluster, through switches and FC-VI interconnect adapters. In this configuration, one node can take over its partner's storage to provide continued data service if the partner goes down.

You can configure the HA pair so that each node in the pair shares access to a common set of storage, subnets, and tape drives, or each node can own its own distinct set of storage.

The controllers are connected to each other through an HA interconnect. This allows one node to serve data that resides on the disks of its failed partner node. Each node continually monitors its partner, mirroring the data for each other’s nonvolatile memory (NVRAM or NVMEM). The interconnect is internal and requires no external cabling if both controllers are in the same chassis.

Takeover is the process in which a node takes over the storage of its partner. *Giveback* is the process in which that storage is returned to the partner. Both processes can be initiated manually or configured for automatic initiation.

How the nodes in an HA pair provide redundancy

To configure and manage nodes in an HA pair, you should be familiar with how the nodes in the HA pair provide redundancy.

- The controllers in the HA pair are connected to each other either through an HA interconnect consisting of adapters and cable, or, in systems with two controllers in the same chassis, through an internal interconnect. The nodes use the interconnect to do the following tasks:

    - Continually check whether the other node is functioning
    - Mirror log data for each other's NVRAM or NVMEM
    - Synchronize each other's time
- They use two or more disk shelf loops, or third-party storage, in which the following conditions apply:
    - Each node manages its own disks or array LUNs.
    - Each node in takeover mode manages its partner's disks or array LUNs. For third-party storage, the partner node takes over read/write access to the array LUNs owned by the failed node until the failed node becomes available again.

      Note: Disk ownership is established by Data ONTAP or the administrator, rather than by which disk shelf the disk is attached to.

      For more information about disk ownership, see the *Data ONTAP 7-Mode Storage Management Guide*.

- They own their spare disks, spare array LUNs, or both and do not share them with the other node.
- They each have mailbox disks or array LUNs on the root volume that do the following tasks:
    - Maintain consistency between the pair
    - Continually check whether the other node is running or whether it has performed a takeover
    - Store configuration information that is not specific to any particular node
- They can reside on the same Windows domain or on different domains.

How HA pairs support nondisruptive operations and fault tolerance

- Fault tolerance
  When one node fails or becomes impaired and a takeover occurs, the partner node continues to serve the failed node's data.
- Nondisruptive software upgrades or hardware maintenance
  When you halt one node and a takeover occurs (automatically, unless you specify otherwise), the partner node continues to serve data for the halted node while you upgrade or perform maintenance on the node you halted.

The HA pair supplies nondisruptive operation and fault tolerance due to the following aspects of their configuration:

- The controllers in the HA pair are connected to each other either through an HA interconnect consisting of adapters and cable, or, in systems with two controllers in the same chassis, through an internal interconnect.
  The nodes use the interconnect to perform the following tasks:
    - Continually check whether the other node is functioning
    - Mirror log data for each other's NVRAM or NVMEM
    - Synchronize each other's time
- They use two or more disk shelf loops, or third-party storage, in which the following conditions apply:

    - Each node manages its own disks or array LUNs.
    - In case of takeover, the surviving node provides read/write access to the partner's disks or array LUNs, until the failed node becomes available again.

      Note: Disk ownership is established by Data ONTAP or the administrator, rather than by which disk shelf the disk is attached to.
- They own their spare disks, spare array LUNs, or both, and do not share them with the other node.
- They each have mailbox disks or array LUNs on the root volume that do the following tasks:
    - Maintain consistency between the pair
    - Continually check whether the other node is running or whether it has performed a takeover
    - Store configuration information

What happens during takeover

When a takeover occurs, the unimpaired partner node takes over the functions and disk drives of the failed node by creating an emulated storage system.

The emulated system performs the following tasks:

- Assumes the identity of the failed node
- Accesses the failed node’s disks, array LUNs, or both and serves its data to clients

The partner node maintains its own identity and its own primary functions, but also handles the added functionality of the failed node through the emulated node.

Note: When a takeover occurs, existing CIFS sessions are terminated. A graceful shutdown of the CIFS sessions is not possible, and some data loss could occur for CIFS users.

Managing HA configuration

Enabling or disabling HA configuration

You can enable a partner node to take over the storage of its failover partner if the partner fails. You can use the HA Configuration window to enable or disable HA configuration. HA configuration is enabled by default.

Steps

1. From the **Home** tab, double-click the appropriate HA configuration.
2. In the navigation pane, click **HA Configuration**.
3. Click either **Enable HA** or **Disable HA**, as appropriate.
4. Select the confirmation check box and click either **Enable** or **Disable**, as appropriate.

Related references

HA Configuration window on page 288

Initiating a takeover

You can use the Takeover Operation dialog box to start the takeover of a storage system that you want to disable so that you can perform repairs or software upgrades.

About this task

You can perform a normal or a forced takeover. In a normal takeover, the HA configuration is checked for the following on both nodes:

- Cluster failover status
- License mismatch
- Date and time settings
- Network interfaces

When you initiate a forced takeover, the HA configuration checks are skipped.

Steps

1. From the Home tab, double-click the appropriate HA configuration.
2. Click **HA Configuration**.
3. Click **Takeover** and select the appropriate storage system from the list.
4. Specify the takeover options and click **Takeover**.
5. Verify that the takeover was successfully completed in the **HA Configuration** window.

Related references

HA Configuration window on page 288

Performing a giveback operation

You can use the Giveback Operation dialog box to issue a giveback, returning the identity of the partner from the emulated storage system to the partner. You can perform a normal giveback, which is a giveback in which you terminate processes on the partner node, or a forced giveback.

Steps

1. From the **Home** tab, double-click the appropriate HA configuration.
2. Click **HA Configuration**.
3. Click **Giveback**.
4. Select a giveback option and click **Giveback**.

Note: If there are open files or a core dump is in progress, you can select the **Force giveback** option. If not, you can select the **Normal** option.

Related references

HA Configuration window on page 288

Halting a storage system

You can halt one of the storage systems in an HA configuration without a takeover by the partner system. You may halt a storage system when you have to perform maintenance on both the storage system and its disks and want to avoid an attempt by the partner node to write to those disks.

Steps

1. From the Home tab, double-click the appropriate HA configuration.
2. Click **HA Configuration**.
3. From the **Halt system** menu, select the appropriate partner storage system.
4. Select the confirmation check box and click **Halt**.
5. Verify that the storage system is halted in the **HA Configuration** window.

Related references

HA Configuration window on page 288

Window descriptions

HA Configuration window

You can use the HA Configuration window to enable and disable HA configuration, complete a takeover, or to issue a giveback.

Command Buttons

Enable/Disable HA	Opens a dialog box to disable or enable high availability.
Takeover	Opens the Takeover Operation dialog box.
Giveback	Opens the Giveback Operation dialog box, which enables you to issue a giveback and return the identity of the partner from the emulated storage system to the partner.
Halt system	Halts the selected storage system.
Refresh	Updates the information in the window.

Related tasks

Enabling or disabling HA configuration on page 286
Initiating a takeover on page 287
Performing a giveback operation on page 287
Halting a storage system on page 288

Copyright information

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

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. 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

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, SnapIntegrator, 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.

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 *www.ibm.com/legal/copytrade.shtml*.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

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.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to *doccomments@netapp.com*. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

/etc/hosts file
  about *203*
  hard limits *204*
  host-name resolution *203*
/etc/nsswitch.conf file *203*
/etc/rc file
  format of *18*

32-bit aggregates
  format explained *109*
  upgrading to 64-bit size *134*
64-bit aggregates
  format explained *109*
  increasing from 32-bit size *134*
7-Mode
  additional information *29*

A

aborting
  a SnapMirror transfer *160*
  See also stopping
access
  secure protocols and storage system *241*
access properties
  CIFS, editing *216*
access to remote data
  using SnapMirror *152*
activating and deactivating
  quotas *99*
active values
  synchronization of *18*
active/active systems
  about *284*
  halting *288*
  halting storage systems *288*
adding
  array LUNs to an aggregate *89*
  AutoSupport email recipients *251*
  disks to aggregate *143*
  DNS domain name *186*
  hosts *204*
  initiators *80*
  licenses *256*
  network interfaces, aliases *197*
  remote access *153*
  RSH sessions *239*
  storage systems *19*
  VLAN interfaces *198*
aggregate overcommitment
  about *109*
aggregates
  64-bit, 32-bit formats explained *109*
  about *108*
  adding array LUNs to *89*
  adding cache disks *130*
  adding capacity disks *130*
  adding disks *143*
  adding disks to *139*
  adding smaller disks to *116*
  characteristics of *108*
  compatible disk types *118*
  composed of SSDs, when you cannot use *120*
  configuration requirements for multi-disk carrier
  disk shelves *140*
  considerations for using disks from multi-disk
  carriers in *141*
  converting to Flash Pools *135*
  creating *127*
  creating from spare array LUNs *88*
  creating from spare disks *142*
  defined *28*
  deleting *129*
  editing *130*
  format explained *109*
  how Flash Pools work *117*
  how you use *108*
  mirrored, defined *124*
  mirroring *129*
  RAID type *110*
  rules for mixing array LUNs in *88*
  rules for mixing HDD types in *117*
  upgrading to 64-bit *134*
  viewing information *136*
Aggregates window *136*
alias
  adding, network *197*
  editing, network *199*
aliases
  creating an iSCSI *228*

ALUA
defined *77*
architecture
overview of Data ONTAP storage *27*
array LUNs
adding to an aggregate *89*
assigning *89*
creating aggregate from spare *88*
Array LUNs window *90*
assigning
array LUNs *89*
ASUP
See AutoSupport
audit events
disabling for CIFS *218*
enabling for CIFS *218*
audit log
clearing *218*
disabling *280*
enabling *280*
saving CIFS log *217*
audit logging
introduction *279*
logs
audit, introduction *279*
auditing
CIFS
about *212*
configuring CIFS and NFS *214*
authentication
changing default initiator method of *231*
iSCSI *227*
public-key-based overview *244*
with SSL *244*
AutoSupport
about *248*
email recipients, adding *251*
enabling or disabling *251*
mail host support for *249*
requirements for *249*
setting up *250*
severity types *250*
technical support and *248*
testing the configuration *252*
transport protocol *249*
AutoSupport window *252*
AZCS type checksums
effect on aggregate management *139*
effect on spare management *139*

B

backup
creating schedules for *166*
types of SnapVault relationships for *176*
backups
using SnapMirror *152*
BCS type checksums
effect on aggregate management *139*
effect on spare management *139*
bringing LUNs online *82*
browser
configuring default web *22*
browsers, improving security through *244*

C

carriers
determining when to remove multi-disk *142*
how Data ONTAP avoids RAID impact when removing multi-disk *141*
spare requirements for multi-disk *140*
changing
aggregate state *134*
default iSCSI initiator security method *231*
CHAP
defined *228*
iSCSI authentication *227*
using with vFiler units *227*
characteristics
aggregate *108*
check points
deleting them in a SnapVault relationship *171*
checksum types
by Data ONTAP disk type *140*
described *139*
effect on aggregate and spare management *139*
CIFS
about *211*
adding NetBIOS aliases *215*
adding the home directory *216*
adding WINS servers *215*
auditing
about *212*
clearing the audit log *218*
configuring CIFS and auditing *214*
deleting NetBIOS aliases *215*
deleting the home directory *217*
deleting WINS servers *215*
diagnostics *268*

diagnostics, monitoring *268*
disabling audit events *218*
editing access security properties *216*
editing general properties *215*
editing idle timeout *215*
editing network properties *215*
editing opportunistic locks (oplocks) *215*
editing protocol properties *216*
editing server description *215*
enabling a trace on domain controllers *219*
enabling audit events *218*
event log
about *212*
license *212*
resetting domain controllers *219*
restarting *217*
saving the audit log *217*
scheduling domain password changes *220*
setting up *213*
stopping *217*
translating names to SIDs
group name *220*
group names *220*
translating to SID *220*
viewing domain information *221*
CIFS (Common Internet File System) *211*
CIFS Diagnostics window *269*
CIFS sessions terminated on takeover *286*
CIFS shares
creating *66*
CIFS window *221*
clones
creating, of LUNs *83*
command sequence
ifgrp *18*
See also interface group
in the /etc/rc file *18*
Common Internet File System (CIFS) *211*
configuration files
editing *205*
configuration settings
verifying, network *22*
configuring
remote access *153*
storage systems *31*
converting
aggregates to Flash Pools *135*
creating
aggregates *127*
an aggregate from spare disks *142*
CIFS shares *66*
Flash Pools *127*
FlexVol volumes *49*
initiator groups *79*
local users *178*
LUN clones *83*
LUNs *78*
qtrees *105*
quotas *98*
SnapMirror relationships *154*
Snapshot copies *53*
SnapVault backup and transfer schedules *166*
SnapVault relationships *166*
target portal groups *228*
user groups *182*
creating aggregates
from spare array LUNs *88*
credential caching, about *12*
credentials
saving *21*
customization
window layout *13*
customizing SSH settings *246*

D

dashboard
monitoring systems *26*
Dashboard window *25*
data
restoring *170*
data access
on the destination *151*
data compression *48*
Data ONTAP
additional 7-Mode information *29*
Data ONTAP-v systems
root volume, introduction to *33*
data protection
using SnapMirror *151*
using SnapVault *164*
date
setting
date and time *253*
setting the date *253*
time
setting the time *253*
DateTime window *254*
deduplication
changing schedule *57*

configuring *56*
FlexVol volume
maximum size *47*
maximum size with deduplication *47*
increasing storage efficiency using *45*
maximum volume size *47*
running on volumes *57*
default browser
configuring *22*
default initiator security
editing *230*
default quota *91*
default quotas
how they work *94*
deleting
aggregates *129*
FlexVol volumes *52*
hosts *205*
initiator groups *79*
initiators from an initiator group *80*
licenses *257*
LUNs *78*
qtrees *106*
quotas *98*
remote access *155*
RSH sessions *239*
SnapMirror relationships *155*
Snapshot copies *54*
SnapVault relationships *168*
SnapVault schedules *167*
target portal groups *229*
deleting a vFiler unit *147*
diagnostics
CIFS *268*
CIFS monitoring *268*
disabling
AutoSupport *251*
dynamic DNS *187*
NDMP service *261*
network interfaces *201*
NFS service *223*
SSL *246*
disaster recovery
using SnapMirror *152*
discovering
storage systems *20*
disk
types for RAID *111*
disk ownership
application to array LUNs *87*
application to disks *87*
disk
ownership
about *87*
disk shelves
aggregate configuration requirements for multi-disk carrier *140*
configuration requirements for multi-disk carrier *140*
disk space hard limit *94*
disk space soft limit *94*
disks
adding disks to aggregate *143*
adding smaller to aggregate *116*
adding them to aggregates *130*
assigned to plexes *126*
considerations for using from multi-disk carriers *141*
creating an aggregate from spare disks *142*
evacuation process, about *141*
how available for Data ONTAP use *87*
managing *138*
matching spares defined *115*
minimum required hot spare *114*
rules for mixing HDD types in aggregates *117*
spare requirements for multi-disk carrier *140*
spare, appropriate *115*
using effective disk types *118*
viewing
disk information *144*
viewing information *144*
Disks window *144*
distinct IP address space *146*
DNS
about *185*
adding domain name *186*
dynamic updates *185*
enabling *186*
enabling or disabling dynamic DNS *187*
host-name resolution *185*, *203*
setting dynamic DNS updates *187*
DNS window *188*
domain account
scheduling password changes *220*
domain controllers
enabling a trace *219*
resetting *219*
viewing information *221*
domain information
viewing *221*
Domain Name System (DNS) *185*
dynamic DNS

about *185*
Dynamic Host Configuration Protocol (DHCP) *185*

E

editing
aggregates *130*
data transfer rate *156*
default security settings *230*
DNS domain name *186*
domain name *186*
FlexVol volume properties *55*
hosts *205*
initiator groups *81*
initiator name *82*
initiator security for iSCSI *231*
LUNs *81*
network aliases *199*
NFS settings *223*
qtrees *106*
quotas *99*
remote access *160*
SnapMirror schedule *156*
SnapVault relationship *171*
SnapVault schedules *172*
target portal groups *232*
vFiler units *148*
Editing share
general settings *68*
options *68*
permissions *68*
effective disk types
grouping disks *118*
emails
adding recipients, AutoSupport *251*
enabling
AutoSupport *251*
DNS *186*
dynamic DNS *187*
NDMP service *261*
network interfaces *201*
NFS service *223*
NIS *209*
NIS slave *210*
SSH *245*
SSL *246*
encryption
with SSL *244*
Ethernet *224*
evacuation process for disks, about *141*
event log
about *212*

F

FAS systems
root volumes, introduction to *33*
fault tolerance *285*
FC
ALUA *77*
FC/FCoE window *237*
FCoE
converged network adapters *235*
data center bridging *235*
Ethernet switch *235*
traditional FC *235*
FCP
changing node name *236*
defined *235*
node connection *235*
nodes defined *235*
starting and stopping *236*
files
creating FlexClone *51*
files hard limit *94*
files soft limit *94*
Flash Cache
disabling *266*
enabling *266*
how it improves performance *266*
read workload *266*
Flash Cache family of modules *265*
Flash Cache window *266*
Flash Pool Statistics window *275*
Flash Pools
converting aggregates to *135*
creating *127*
how they work *117*
requirements for using *119*
statistics window *275*
FlexClone files
creating *51*
FlexClone volumes
about *35*
creating *50*
how they save space *34*
shared Snapshot copies and *39*
space guarantees and *42*
flexible volumes
described *28*

FlexVol volumes
about *38*
automatically adding space for *44*
creating *49*
creating Snapshot copies *53*
deleting *52*
editing properties *55*
how volume guarantees work with *40*
renaming Snapshot copies *60*
resizing *59*
setting reserve for Snapshot copies *52*
space management for *41*
thick provisioning for *40*
thin provisioning for *40*
try_first volume option *44*
flow control
about *194*
forced takeover *287*
formats
64-bit, 32-bit aggregates explained *109*
frame
about *193*
characteristics *193*
flow *194*
frame size *193*
jumbo frame *193*
MTU size *193*
Pause Off *194*
Pause On *194*
free space
automatically increasing *44*

G

generating
SSH keys *245*
SSL certificate *246*
giveback
performing, for HA pairs *287*
group quota *91*
groups
about *177*
assigning a local user *180*, *183*
creating target portal *228*
creating user *182*
deleting *183*
deleting target portal *229*
editing description *184*
editing target portal *232*
Groups window *184*
guidelines
for creating LUNs *73*
LUN mapping *76*
LUN type *74*

H

HA configuration
disabling *286*
enabling *286*
HA configuration window *288*
HA configurations
benefits of *285*
characteristics of *284*
definition of *284*
HA interconnect *285*
HA pairs
and iSCSI *225*
performing giveback *287*
HBA *76*, *224*
HDDs
rules for mixing types in aggregates *117*
Help
about the *10*
home directories
defined *213*
home directory
adding for CIFS *216*
deleting for CIFS *217*
Home tab *23*
host
naming *188*
host name
about *188*
resolution, with /etc/hosts file *203*
resolution, with DNS *185*
resolution, with NIS *207*
host-name resolution
about *203*
using /etc/hosts file *203*
using DNS *185*
using NIS *207*
hosts
managing trusted *239*
hot spares
appropriate *115*
defined *114*, *139*
matching, defined *115*
minimum needed *114*
what disks can be used as *115*

See also spares
hybrid aggregates
See Flash Pools

I

icons, definitions *14*
increasing
aggregate size to 64 bit *134*
increasing cache memory *265*
initiating takeover *287*
initiator groups
adding initiators *80*
creating *79*
defined *76*
deleting *79*
deleting initiators *80*
editing *81*
editing initiators *82*
name rules *77*
naming *77*
ostype of *77*
requirements for creation *77*
type *77*
viewing *84*
initiator security
viewing iSCSI *233*
initiators
adding *80*
adding security for iSCSI *230*
changing the name *82*
deleting from an initiator group *80*
setting default security for iSCSI *232*
installing SSL certificate *247*
interface group
about *191*
dynamic multimode *193*
load balancing *193*
load balancing, IP address based *193*
load balancing, MAC address based *193*
manage *190*
naming *188*
single-mode *193*
static multimode *193*
types *193*
interfaces
enabling or disabling iSCSI service *229*
IP address
configuration *191*
iSCSI
ALUA *77*
changing default initiator security *231*
creating aliases *228*
disabling on interface *229*
editing initiator security *231*
enabling on interface *229*
explained *224*
how communication sessions work *226*
initiator security
setting default *232*
initiator security, viewing *233*
nodes defined *225*
security *227*
target portal groups defined *225*
using with HA pairs *225*
iSCSI initiators
adding security *230*
iSCSI service
starting *233*
stopping *233*
iSCSI window *233*

K

keys
public-based, authentication overview *244*

L

lag time *162*
license
FC *235*
licenses
adding *256*
CIFS *212*
deleting *257*
disabling *257*
enabling *257*
requirements *255*
Licenses window *258*
load balancing
IP address based *193*
MAC address based *193*
multimode interface groups *193*
round-robin *193*
using SnapMirror *152*
local groups
assigning a local user *183*
creating *182*
deleting *183*

editing description *184*
local user accounts
when to create *177*
local users
about *177*
assigning to a group *180*
creating *178*
deleting *178*
editing description *179*
editing full name *179*
password, changing *180*
password, editing duration *179*
password, resetting *181*
local users and groups
about *177*
localhost *203*
log files
how they are rotated *15*
LUN
host operating system type *74*
multiprotocol type *74*
LUN clones
creating *83*
LUN creation
host operating system type *74*
LUNs
bringing online *82*
creating *78*
creating clones *83*
deleting *78*
editing *81*
guidelines for creating *73*
initiator hosts *76*
mapping guidelines *76*
resizing *76*
size and type *74*
taking offline *83*
viewing information about *84*
LUNs (array)
Data ONTAP RAID groups with *113*
how available for Data ONTAP use *87*
rules for mixing in an aggregate *88*
LUNs window *85*

M

mail host support for AutoSupport *249*
mailbox disks *284*
mailbox disks in the HA pair *285*
mirroring
for data protection *151*
mirroring, NVMEM or NVRAM log *285*
modifying
local user's password *180*
NIS domain name *210*
quotas *99*
remote access *160*
SnapVault schedules *172*
monitoring system status
using Syslog messages *278*
MPIO *76*
multi-disk carrier
spare requirements for *140*
multi-disk carrier disk shelves
aggregate configuration requirements for *140*
multi-disk carrier shelves
configuration requirements for *140*
multi-disk carriers
considerations for using disks from *141*
determining when to remove *142*
how Data ONTAP handles when removing *141*
multimode interface groups
load balancing, IP address based *193*
load balancing, MAC address based *193*
multiprotocol type
guidelines *74*

N

name restrictions
qtree *103*
name rules
igroups *77*
NDMP
about *261*
stopping a session *262*
NDMP service
disabling *261*
enabling *261*
NDMP window *263*
network configuration
verification tool *17*
verifying settings *22*
network configuration checker
defined *17*
network files
adding *204*
hosts
deleting *205*
editing *205*

Network Files window *206*
Network Information Service (NIS) *207*
network interface
    10 Gigabit Ethernet *191*
    10/100/1000 Ethernet *191*
    100 Mbps *191*
    100BT *191*
    configuration *191*
    Gigabit Ethernet *191*
    naming *188*
    types *191*
network interfaces
    adding aliases *197*
    creating
        virtual interface *197*
    creating VLANs *198*
    disabling *201*
    editing
        network interfaces *200*
        virtual interface *199*
    enabling *201*
    flow controlflow control
        about *194*
    interface alias, adding *197*
    interface alias, editing *199*
    links *188*
    virtual interfaces
        creating *197*
        editing *199*
    virtual interfaces, creating *197*
    virtual interfaces, editing *199*
Network Interfaces window *201*
NFS
    concepts *222*
    creating exports *69*
    disabling audit events *218*
    editing export rules *71*
    editing the settings *223*
    enabling audit events *218*
    exports
        adding an export rule *71*
        deleting *70*
    saving audit log *217*
NFS datastore
    creating for VMware *31*
NFS service
    disabling *223*
    enabling *223*
NIS
    about *207*
    adding domain name *210*
    administrative commands
        yppush *208*
    binding master *209*
    considerations *209*
    enabling *209*
    enabling NIS slave *210*
    host-name resolution *203*, *207*
    hosts map *207*
    ipnodes map *207*
    IPv6 support *207*
    master *208*
    modifying domain name *210*
    slave *208*
NIS (Network Information Service) *207*
NIS and /etc/hosts file *209*
NIS slave
    about *208*
    guidelines *208*
    improve performance *208*
NIS window *211*
nodes
    FCP *235*
    iSCSI *225*
nondisruptive operations *285*
normal takeover *287*
NVMEM log mirroring *285*
NVRAM log mirroring *285*

O

online Help
    about *10*
options
    nis.server *208*

P

PAM
    See Flash Cache
PAM (Performance Acceleration Module) *265*
parent FlexVol volumes
    splitting FlexClone volumes from *36*
password
    changing *238*
    scheduling changes for domain accounts *220*
password duration
    editing for local users *179*
Password/RSH window *240*
passwords

local users, changing *180*
local users, resetting *181*
paths
ALUA *77*
pause frame *194*
Performance Acceleration Module *265*
performance improvements, in storage systems
WAFL external cache *265*
persistent values
synchronization of *18*
plex
defined *28*, *124*
plexes
bring online *132*
destroying *132*
mirroring *129*
splitting *133*
take offline *131*
portal groups
creating target *228*
deleting target *229*
editing target *232*
pre-Windows 2000 format
rules for specifying user names in *93*
protocol properties
CIFS, editing *216*
protocols
introduction to SSH *242*
secure, and storage system access *241*
public-key-based
authentication overview *244*

Q

qtree quota *91*
qtrees
about *102*
creating *105*
defined *28*
deleting *106*
deletion, quotas and *96*
editing *106*
name restrictions *103*
options *102*
releasing from a SnapVault relationship *169*
renaming, quotas and *96*
security style *104*
viewing information *107*
when to use *103*
Qtrees window *107*
quotas
activating and deactivating *99*
creating *98*
default *91*
deleting *98*
editing *99*
group *91*
hard *91*
how they work with qtrees *95*
managing *94*
qtree *91*
qtree deletion, and *96*
qtree rename and *96*
reinitialization, when required *97*
resizing *100*
security style changes and *96*
soft *91*
threshold *91*
tree *95*
UNIX users and *92*
user *91*
user and group, working with qtrees *95*
viewing information about *100*
why you use *91*
Windows users and *92*
Quotas window *101*

R

RAID
avoiding impact to when replacing multi-disk carriers *141*
protection with SyncMirror and *121*
RAID 0
how Data ONTAP uses for array LUNs *116*
RAID disk types *111*
RAID groups
adding disks to *139*
definition *110*
maximum number allowed *117*
naming convention *112*
size *112*
sizing considerations for disks *112*
with array LUNs, considerations *113*
RAID types
editing *130*
RAID-DP *111*
RAID-level mirroring
described *28*
RAID4

described *112*
raw device mapping *76*
RDM *76*
read workload statistics
Flash Cache *266*
reinitializing quotas *97*
remote access
adding *153*
configuring in storage systems *165*
deleting *155*
editing *160*
what it is *175*
removing
multi-disk carriers, determining when it is safe *142*
removing storage systems *20*
requirements
Flash Pool use *119*
licenses *255*
resizing
FlexVol volumes *59*
quotas *100*
resizing volumes
options for *38*
restarting
SnapMirror relationship *158*
restoring
data *170*
restrictions
qtree name *103*
resynchronizing
SnapVault relationships *169*
retention count
SnapVault *175*
retention period
about retention period *49*
root volumes
introduction to *33*
rotation
of log files *15*
RSH
about *238*
RSH sessions
adding *239*
deleting *239*
rules
for mixing HDD types in aggregates *117*

S

schedule deduplication
changing *57*
schedule types
SnapVault *175*
secure protocols
and storage system access *241*
Secure Shell
See SSH protocol
Secure Sockets Layer
See SSL
SecureAdmin
improving security with SSL *244*
securing styles
changing, quotas and *96*
security
editing the default settings *230*
setting iSCSI initiator default *232*
viewing iSCSI initiator *233*
security styles
affect on data access *44*
setting date and time
guidelines *253*
setting up
AutoSupport *250*
CIFS *213*
severity
AutoSupport *250*
shares
creating, CIFS *66*
Shares
disabling *67*
Shares window *68*
shelves
aggregate configuration requirements for multi-disk carrier *140*
configuration requirements for multi-disk carrier *140*
sizing
RAID groups for disks considerations *112*
SnapLock
Compliance volumes *49*
Enterprise volumes *49*
SnapLock Compliance volume *48*
SnapLock Enterprise volume *48*
SnapMirror
deployment *152*
format of log files *281*
log file examples *281*
qtree replication *151*
uses *152*
volume replication *151*
SnapMirror Log window *283*

SnapMirror relationship
resuming *158*
SnapMirror relationships
breaking *158*
creating *154*
deleting *155*
initializing *156*
managing through System Manager *151*
properties
editing *156*
quiescing *157*
resynchronizing *159*
reverse resynchronizing *159*
updating *157*
SnapMirror transfers
aborting *160*
SnapMirror volumes
FlexClone volumes considerations for *37*
SnapMirror window *161*
Snapshot copies
automatic scheduling *60*
creating *53*
deleting *54*
directory, making invisible *61*
renaming *60*
restoring a volume from *59*
scheduling *60*
setting reserve *52*
understanding *39*
viewing list of *62*
SnapVault
basic deployment *164*
creating transfer schedules *166*
for data protection, using *164*
remote access *175*
remote access, granting or denying *165*
retention count *175*
schedule types *175*
what is *163*
SnapVault deployment
introduction *164*
SnapVault relationship
releasing *169*
restoring data *170*
SnapVault relationships
adding specific remote systems *165*
creating *166*
deleting check points *171*
deleting specific remote systems *165*
editing *171*
reinitializing *167*
resynchronizing *169*
starting *167*
stopping *168*, *170*
types of *176*
updating *168*
SnapVault schedules
deleting *167*
editing *172*
viewing *172*
SnapVault window *173*
SNMP
agent *259*
agent, configure *259*
enabling SNMP traps *260*
for storage system discovery *12*
setting information *260*
SNMP window *261*
software efficiency
FlexVol volumes *43*
space
increasing for full FlexVol volumes *44*
space efficiency
configuring deduplication *56*
space guarantees
See volume guarantees
space management
what kind to use *41*
spare disks
appropriate *115*
defined *114*, *139*
matching, defined *115*
what disks can be used as *115*
spare disks in the HA pair *285*
spares
minimum needed *114*
requirements for multi-disk carriers *140*
splitting
FlexClone volumes from parent volumes *36*
Splitting FlexClone volumes *58*
SSDs
aggregates composed of, when you cannot use *120*
how used in Flash Pools *117*
SSH
customizing settings *246*
enabling *245*
generating keys *245*
SSH protocol
introduction to *242*
SSH/SSL window *247*

SSL
certificate
generating *246*
installing *247*
certificates *241*
enabling or disabling *246*
how to manage *244*
SSL (Secure Sockets Layer) protocol
authentication with *244*
improving security with *244*
starting
iSCSI service *233*
SnapVault relationships *167*
states of an aggregate *134*
stopping
an NDMP session *262*
iSCSI service *233*
SnapVault relationships *168*, *170*
storage
rules for mixing array LUNs in an aggregate *88*
storage architecture
overview of Data ONTAP *27*
storage efficiency
data compression *48*
how FlexClone volumes help achieve *34*
Storage node *27*
storage system
access and secure protocols *241*
storage system credentials, saving *21*
storage systems
adding *19*
configuring *31*
discovering *20*
discovery of *12*
editing boot configuration file in *23*
halting *263*
monitoring *26*
rebooting *264*
removing *20*
resource management
about *11*
viewing information *32*
storage units
types *28*
support bundle
creating *16*
for troubleshooting *15*
uploading *16*
support for AutoSupport, mail host *249*
supportability dashboard *16*
SyncMirror
advantages *120*
aggregate *124*
description *120*
mirrored aggregates, create *125*
plexes *28*
protection with RAID and *121*
Syslog messages
monitoring status using *278*
understanding *275*
Syslog messaging configuration
editing *277*
Syslog window *278*
System Health window *274*
system logging
about *13*
configuring *21*
log levels *13*
System Manager
about *11*
supported Data ONTAP versions *11*
tasks you can perform from *11*
system password, changing *238*

T

takeover
CIFS sessions and *286*
what happens during *286*
taking LUNs offline *83*
tape backup
using NDMP protocol *261*
tape backup and recovery
using NDMP *261*
target portal groups
about *225*
creating *228*
deleting *229*
editing *232*
technical support
creating a support bundle for *16*
terminating
an NDMP session *262*
terminology
RAID groups
on a storage array *116*
testing
AutoSupport configuration *252*
thick provisioning
for FlexVol volumes *40*

thin provisioning
about *109*
for FlexVol volumes *40*
using FlexVol volumes *43*
threshold soft limit *94*
traditional volumes
described *28*
transfers
aborting, SnapMirror *160*
Transport Layer Security (TLS) protocol *244*
tree quotas *95*
troubleshooting
support bundle for *15*
trusted hosts
about *238*
managing *239*
try_first volume option *44*

U

UNIX users, specifying for quotas *92*
updating
SnapMirror relationships *157*
SnapVault relationships *168*
upgrading
aggregates from 32 bit to 64 bit *134*
uploading
support bundle *16*
user groups
creating *182*
user names
rules for specifying in pre-Windows 2000 format *93*
translating to SID *220*
user quota *91*
users
changing password of local *180*
resetting password of local *181*
Users window *181*

V

V-Series systems
root volumes, introduction to *33*
version
viewing information about *21*
vFiler unit
default *146*
starting *149*
vFiler units
authentication using CHAP *227*
creating *147*
defined *146*
editing *148*
viewing
aggregate information *136*
initiator groups *84*
iSCSI initiator security *233*
LUN information *84*
qtree information *107*
quota information *100*
SnapVault schedules *172*
version information *21*
Viewing FlexClone hierarchy *61*
viewing storage system information *32*
VLAN
naming *188*
tags *196*
VLAN interfaces
creating *198*
VLANs
advantages of *195*
tagging *194*
VMware
creating NFS datastore *31*
volume guarantees
effect on maximum FlexVol volume size *40*
how they work with FlexVol volumes *40*
volume status
changing *55*
Volume window *62*
volumes
automatically adding space for *44*
changing the status *55*
configuring deduplication *56*
deduplication
changing schedule *57*
defined *33*
FlexClone
creating *50*
FlexVol volumes *38*
how FlexClone type saves space *34*
how you use aggregates to provide storage to *108*
resizing options *38*
restoring from Snapshot copies *59*
running deduplication *57*
scheduling Snapshot copies *60*
Snapshot copies
making directory invisible *61*
understanding *39*
viewing list of Snapshot copies *62*

W

WAFL (Write Anywhere File Layout) *265*
WAFL external cache
    about *265*
web browser
    configuring *22*
what is
    SnapVault *163*
window layout
    customization *13*
Windows users, specifying for quotas *92*
Write Anywhere File Layout (WAFL) *265*