

OnCommand™ System Manager 2.0

Help For Use With Data ONTAP® Cluster-Mode

NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089 USA
Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 4-NETAPP
Documentation comments: doccomments@netapp.com
Information Web: <http://www.netapp.com>
Part number: 215-06034_A0
September 2011

Contents

Welcome to System Manager Help	9
System Manager	11
Understanding System Manager	11
Storage resource management	11
Storage system discovery	12
Credential caching	12
System logging	12
Window layout customization	13
Access to your favorite topics	13
Icons used in the application interface	13
Support for troubleshooting issues in System Manager	14
Creating a support bundle	15
Configuring System Manager	15
Adding storage systems	15
Removing storage systems	16
Discovering storage systems	16
Saving your storage system credentials	16
Configuring system logging	17
Viewing System Manager application information	17
Configuring the SNMP timeout value	17
Window descriptions	17
Home tab	17
Cluster	19
What a cluster is	19
What a node in the cluster is	19
Storage units for managing disk space	20
Where to find additional Data ONTAP Cluster-Mode information	21
Dashboard window	21
Viewing storage system details	22
Storage > Aggregates	23
Understanding aggregates	23
Configuring aggregates	26

Managing aggregates	28
Monitoring aggregates	30
Window descriptions	31
Storage > Disks	32
Understanding disks	32
Configuring disks	41
Managing disks	42
Monitoring disks	42
Window descriptions	43
Storage > Array LUNs	44
Understanding array LUNs	44
Configuring array LUNs	46
Managing array LUNs	47
Window descriptions	47
Configuration > Security > Users	48
Understanding users	48
Managing users	49
Window descriptions	50
Configuration > Security > Roles	52
Understanding roles	52
Managing roles	52
Window descriptions	53
Configuration > Services > LDAP	54
Understanding LDAP	54
Managing LDAP	54
Window descriptions	56
Configuration > Services > Kerberos	57
Understanding Kerberos	57
Managing Kerberos	57
Window descriptions	58
Configuration > System Tools > DateTime	59
Managing the system time (cluster administrators only)	59
Window descriptions	60
Configuration > System Tools > Licenses	61
Managing feature licenses (cluster administrators only)	61
Window descriptions	63

Configuration > System Tools > SNMP	64
Understanding SNMP	64
Managing SNMP on the cluster (cluster administrators only)	64
Window descriptions	66
Configuration > Snapshot Policies	66
Understanding Snapshot policies	66
Managing Snapshot policies	66
Window descriptions	68
Configuration > Schedules	68
Understanding schedules	69
Managing schedules	69
Window descriptions	70
Configuration > Peers	71
Understanding cluster peers	71
Managing cluster peers	74
Window descriptions	75
Diagnostics > Job	76
Understanding jobs	76
Window descriptions	77
Diagnostics > System Health	77
Understanding system health	77
Monitoring the health of your system	79
Window descriptions	80
Diagnostics > Logs > Events	81
Window descriptions	81
Diagnostics > Logs > SnapMirror Log	82
SnapMirror files	82
Window description	83
Vserver	85
What a Vserver is	85
Vserver management	85
Types of Vservers	86
Why you use Vservers	86
Managing Vservers	87
Window descriptions	90
Storage > Volumes	90

Understanding volumes	90
Configuring volumes	105
Managing volumes	110
Monitoring volumes	117
Window descriptions	118
Storage > Namespace	122
Managing namespace	122
Window descriptions	123
Storage > Shares	124
Understanding shares	124
Configuring shares	125
Managing shares	127
Window descriptions	127
Storage > LUNs	128
Understanding LUNs	129
Configuring LUNs	134
Managing LUNs	138
Monitoring LUNs	142
Window descriptions	142
Storage > Qtrees	144
Understanding qtrees	144
Configuring qtrees	147
Managing qtrees	148
Monitoring qtrees	149
Window descriptions	149
Storage > Quotas	150
Understanding quotas	150
Configuring quotas	156
Managing quotas	157
Monitoring quotas	159
Window descriptions	160
Policies > Export Policies	161
Understanding export policies	161
Managing export policies	162
Window descriptions	164
Policies > Efficiency Policies	165

Understanding efficiency policies	165
Managing efficiency policies	165
Window descriptions	167
SnapMirror	167
Understanding SnapMirror technology	167
Configuring SnapMirror relationships	168
Managing SnapMirror relationships	169
Window descriptions	174
Configuration > Network Interfaces	175
Understanding interfaces	175
Configuring interfaces	176
Managing interfaces	176
Window descriptions	177
Configuration > Protocols > CIFS	178
Understanding CIFS	178
Configuring CIFS	179
Managing CIFS	180
Monitoring CIFS	184
Window descriptions	184
Configuration > Protocols > NFS	185
Understanding NFS	185
Managing NFS	185
Window descriptions	186
Configuration > Protocols > iSCSI	186
Understanding iSCSI	186
Configuring iSCSI	189
Managing iSCSI	192
Monitoring iSCSI	195
Window descriptions	195
Configuration > Protocols > FC/FCoE	196
Understanding FC/FCoE	196
Configuring FC/FCoE	197
Managing FC/FCoE	198
Window descriptions	198
Configuration > Security > Users	199
Managing users	199

Window descriptions	200
Configuration > Security > Roles	202
Understanding roles	202
Managing roles	203
Window descriptions	204
Configuration > Services > NIS	204
Managing NIS domains	205
Window descriptions	205
Configuration > Services > LDAP	206
Managing LDAP	206
Configuration > Services > Kerberos	207
Managing Kerberos	207
Configuration > Services > DNS	207
Configuring DNS	207
Window descriptions	208
Configuration > Local Users and Groups > UNIX	209
Window descriptions	209
Configuration > Local Users and Groups > Name Mapping	209
Understanding name mapping	209
Window descriptions	211
Storage Controllers	213
Configuration > Ports/Adapters	213
Window descriptions	213
Configuration > NDMP	214
Window descriptions	214
Configuration > AutoSupport	214
Window descriptions	214
HA Configuration	215
Window descriptions	215
Copyright information	217
Trademark information	219
How to send your comments	221
Index	223

Welcome to System Manager Help

The Help includes information about how to configure, manage, and monitor Data ONTAP 8.1 Cluster-Mode storage systems and storage objects by using System Manager. The table of contents, search, index, and favorites in the Help system help you find the relevant information needed to achieve your goals.

The structure of the Help is similar to what you see on the GUI. Help is also available from each window and its respective tabs. To learn about a specific window parameter, you can click the Help icon ().

System Manager

Understanding System Manager

System Manager enables you to manage storage systems and storage objects, such as disks, volumes, and aggregates. System Manager is a Web-based graphical management interface to manage common functions related to storage systems from a Web browser.

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

- Data ONTAP 7.x (starting from 7.2.3)
- Data ONTAP 8.x 7-Mode
- Data ONTAP 8.1 Cluster-Mode

Note: In the Data ONTAP 8.x 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.

You can also 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.
- Set up and manage SnapMirror relationships.
- Perform cluster management, storage node management, and Vserver management operations in a cluster environment.
- Create and configure Vservers, manage storage objects associated with a Vserver, and manage Vserver services.
- Monitor and manage your storage systems.

Related tasks

[Discovering storage systems](#) on page 16

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

Storage system discovery

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

When you add one of the systems in a 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 appear.

Filtering

You can use the filter icon () to display only those entries that match the conditions provided. You can use the character filter (?) or string filter (*) to narrow your search. You can apply filters to one or more columns. The column heading is highlighted if a filter is applied.

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 choose to display or hide the "list of related objects" and "list of views" panels. You can drag vertical dividers to resize the width of columns or other areas of the window.

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.

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 below.

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.

Icon	Name	Description
	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 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 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 now and is using System Manager.
- Log files created by System Manager application – These files records the error 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.

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. Create the support bundle.

Configuring System Manager

Adding storage systems

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

Before you begin

- All the nodes in the cluster must be running Data ONTAP 8.1.

Steps

1. From the **Home** tab, click **Add**.
2. Type the fully qualified DNS host name, or the IPv4 address of the storage system.
3. Click the **More** arrow.
4. Select the method for discovering and adding the storage system:

- **SNMP**

You must specify the SNMP community and SNMP version.

Note: If the storage system is running a Data ONTAP release in the Data ONTAP 7.2 release family, use SNMP version1.

- **Credentials**

You must specify the user name and password.

5. Click **Add**.

Result

The storage system that you added is included in the list of managed systems.

Removing storage systems

You can remove one or more storage systems from the list of managed systems in System Manager.

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 on a network subnet and add them to the list of managed systems.

Before you begin

- All the nodes in the cluster must be running Data ONTAP 8.1.

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 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 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**.

Window descriptions

Home tab

The Home tab enables you to view managed storage systems. You can use the window command buttons to discover and add storage systems.

- [Command buttons](#) on page 18
- [Systems list](#) on page 18

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 a cluster.
Version	Specifies the version number of the operating system.
Model	Specifies the storage system model.
System ID	Specifies the ID of the storage system.

Cluster

What a cluster is

You can group pairs of nodes together to form a scalable cluster. Creating a cluster enables up to 24 individual nodes (or twelve pairs of nodes) to pool their resources and distribute work across the cluster, while presenting administrators with a single entity to manage. Clustering also enables continuous service to end users if individual nodes go offline or are taken offline.

Each node in the cluster can access the same volumes as any other node in the cluster. The total file-system namespace, which comprises all of the volumes and their resultant paths, is global across the cluster.

When a new node is added to a cluster, there is no need to update clients to point to the new node. The existence of the new nodes is transparent to the clients.

A key concept in the cluster is *quorum*, which is a simple majority of nodes that participate in the cluster. Each node in the cluster participates in a voting protocol that elects one node *master*; each remaining node is a *secondary*. The master is responsible for synchronizing information across the cluster. When quorum is formed, it is maintained by continual voting; if the master goes offline, a new master is elected by the nodes that remain online. Quorum must be maintained for the cluster to function fully. Because there is the possibility of a tie in a cluster that has an even number of nodes, one node can be configured to hold *epsilon*, which is an extra fractional voting weight. For instance, in a four-node cluster, quorum can be maintained if two nodes cast the same vote and one of those nodes holds epsilon. The system administrator can set epsilon on one node at most. Epsilon can be unassigned and reassigned as needed.

If you have a two-node cluster, you must configure cluster high availability (HA). For more information, see the *Data ONTAP Cluster-Mode High Availability Configuration Guide*.

The nodes in a cluster communicate over a dedicated, physically secure Ethernet network. The cluster logical interfaces (LIFs) on each node in the cluster must be on the same subnet. For information about network management for cluster and nodes, see the *Data ONTAP Cluster-Mode Network Management Guide*.

For information about setting up a cluster or joining a node to the cluster, see the *Data ONTAP Cluster-Mode Setup Guide*.

What a node in the cluster is

A *node* is a controller that runs the Data ONTAP Cluster-Mode software. In a cluster, it is connected to other nodes over a cluster network.

A node is also connected to the disk shelves that provide physical storage for the Data ONTAP Cluster-Mode system or to third-party storage arrays that provide array LUNs for Data ONTAP use.

Services and components that are controlled by the node, not by the cluster, can be managed by using the `system node` commands.

A *node Vserver* represents a node in the cluster. The cluster setup process automatically creates a node Vserver for each node in the cluster. If you run the `vserver show` command, the output includes node Vservers in the list.

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.
- 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 FlexVol volumes.
- Flexible volumes** 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 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 flexible volume or in the root of a qtree.

For detailed information about storage units, see the *Data ONTAP Cluster-Mode Logical Storage Management Guide* and *Data ONTAP Cluster-Mode Physical Storage Management Guide*.

Related information

[Data ONTAP documentation on the NetApp Support Site-support.netapp.com](http://Data-ONTAP-documentation-on-the-NetApp-Support-Site-support.netapp.com)

Where to find additional Data ONTAP Cluster-Mode information

System Manager Help provides basic Data ONTAP Cluster-Mode conceptual information that is relevant to the tasks that you can perform using System Manager. For more in-depth conceptual information that you might require for configuration, monitoring, and management of storage objects and storage systems, you can see the Data ONTAP documentation available on the NetApp Support Site.

You can access the following related Data ONTAP Cluster-Mode documentation from the NetApp Support Site:

Physical Storage Management Guide	Describes how to configure, operate, and manage the physical storage resources for storage systems running Data ONTAP Cluster-Mode. It covers disks, RAID groups, plexes, and aggregates.
Logical Storage Management Guide	Describes how to configure, operate, and manage the storage resources for storage systems running Data ONTAP Cluster-Mode. It covers volumes, FlexClone volumes, files and LUNs, deduplication, compression, qtrees, and quotas.
System Administration Guide	Describes general system administration for storage systems that run Data ONTAP software.
High Availability Configuration Guide	Describes installation and management for high availability configurations.
Vserver Administrator Capabilities Overview	Describes the capabilities of a Vserver administrator to administer Vservers (virtual storage systems) in the Data ONTAP Cluster-Mode.
Network Management Guide	Describes how to configure and manage networks associated with storage systems running Data ONTAP Cluster-Mode and how to manage file access with NFS and CIFS protocols.
Block Access Management Guide for iSCSI and FC	Describes how to configure and manage the iSCSI and FC protocols for SAN environments.
File Access and Protocols Management Guide	Describes how to configure and manage networks and how to manage file access with NFS and CIFS protocols.
Data Protection Management Guide	Describes how to back up and recover data using Data ONTAP Cluster-Mode online backup and recovery features.

Related information

Data ONTAP Cluster-Mode Information Library page - http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_cmode_index.shtml

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 various aspects of your

storage management environment, such as space and CPU utilization, the status of storage objects, notifications, system properties, network throughput, and protocol operations.

The available panels for storage systems running Data ONTAP Cluster-Mode are as follows:

- **Properties** - Displays cluster attributes like cluster management IP address, Data ONTAP version, number of nodes and Vservers in the cluster, and number of stopped Vservers, if any.
- **System Health** - Displays the overall health of the cluster and the health of each subsystem in the cluster. A link is provided to the System Health Monitor.
- **Alarms** - Displays alarms for the following:
 - Average client latency (performance of NFS and CIFS operations)
 - CPU utilization
 - Data network health (ratio of number of packet errors to total number of packets)
 - Data port utilization
 - Storage space utilization
- **Storage** - Displays the total number of aggregates and the number of offline aggregates, if any.

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 8.0 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.

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 > Aggregates

Understanding aggregates

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. These aggregates provide storage to the volume or volumes that they contain.

Each aggregate has its own RAID configuration, plex structure, and set of assigned disks or array LUNs. You create an aggregate to hold one or more FlexVol volumes—the logical file systems that share the physical storage resources, RAID configuration, and plex structure of that common containing aggregate.

Aggregates can be mirrored or unmirrored. An unmirrored aggregate has only one plex; mirrored aggregates have two plexes. For Data ONTAP 8 Cluster-Mode, only unmirrored aggregates are supported at this time.

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

Related information

[*TR 3437: Storage Best Practices and Resiliency Guide*](#)

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.

For native storage, Data ONTAP uses RAID-DP or RAID4 groups to provide parity protection. For third-party storage, Data ONTAP uses RAID0 groups to optimize performance and storage utilization. The storage arrays provide the parity protection for third-party storage.

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.

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.

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.

Understanding 64-bit and 32-bit aggregates

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 or cluster. They interoperate for some Data ONTAP features but not all.

32-bit aggregates have a maximum size of 16 TB; 64-bit aggregates have a maximum size of up to 100 TB, depending on the storage system model. For the maximum 64-bit aggregate size of your storage system model, see the *System Configuration Guide*.

You decide the format of an aggregate when you create it. By default, newly created aggregates are 64-bit.

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

Related information

System Configuration Guide: now.netapp.com/NOW/knowledge/docs/hardware/NetApp/syscfg/

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.

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.

How moving a volume works

You can move a volume from one aggregate or node to another within the same Vserver for capacity utilization, improved performance, and to satisfy service-level agreements.

Moving a volume does not disrupt client access when the system is moving the volume. Moving a volume occurs in multiple phases. A new volume is made on the destination aggregate. The data from the original volume is copied to the new volume. During this time, the original volume is intact and available for clients to access. At the end of the move process, client access is temporarily blocked. During this time the system performs a final replication from the source volume to the destination volume, swaps the identities of the source and destination volumes, and the destination

becomes the source. After completing the move, the system routes client traffic to the new source volume and resumes client access.

The move is not disruptive to client access because the time in which client access is blocked ends before clients notice a disruption and timeout. Client access is blocked for 45 seconds by default. If the volume move operation cannot finish in the time that access is denied, the system aborts this final phase of the volume move operation and allows client access. The system will run the final phase of the volume move operation until the volume move completes or until the default of three attempts is reached. If the volume move operation fails after the third attempt, the process goes into a cutover deferred state and waits for you to initiate the final phase.

You can change the amount of time client access is blocked or the number of times the final phase of the volume move operation, known as cutover attempts, are run if the defaults are not adequate. Also you can determine what the system does if the volume move operation cannot be completed during the time client access is blocked. See the volume move start man page for details about moving a volume without disrupting client access.

Requirements for performing a volume move

Before you move a volume nondisruptively, you must be aware of the type of volumes you can move and the operations that might conflict with the volume move. The volume move does not start if the volume has unsupported settings or if there are conflicting operations.

- Your FAS system or V-Series system must be running Data ONTAP 8.0 or later.
- The volume must be online.
- You cannot move the following types of volumes:
 - A root volume
 - A FlexCache volume
 - A compressed 64-bit volume to a 32-bit volume
- The source volume must be consistent.
- The following conflicting operations must not be running:
 - SnapRestore of the source volume or the containing aggregate
 - WAFLIron operation on the source or the destination aggregate
 - Revert operation on the storage system

Note: FlexClone volumes of the source volume are not moved along with the source volume.

Related information

[Data ONTAP documentation on the NetApp Support Site-support.netapp.com](http://support.netapp.com)

Configuring aggregates

Creating aggregates

You can use the Create Aggregate wizard to create an aggregate to provide storage to one or more FlexVol volumes. You must specify the RAID type, RAID group size, and disk type.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays either the cluster hierarchy or the storage controllers hierarchy.
3. In the navigation pane, click **Storage > Aggregates**.
4. Click **Create**.
5. Type or select information as requested by the wizard.
6. Confirm the details and click **Finish** to complete the wizard.

Related references

[Aggregates window](#) on page 31

Deleting aggregates

You can delete aggregates when you no longer need the data in that aggregate. However, you cannot delete the root aggregate because it contains the root volume, which contains the system configuration information.

Before you begin

- The aggregate must be offline.
- All of the FlexVol volumes contained by the aggregate must be deleted.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays either the cluster hierarchy or the storage controllers hierarchy.
3. In the navigation pane, click **Storage > Aggregates**.
4. Select one or more aggregates that you want to delete and click **Delete**.
5. Select the confirmation check box and click **Delete**.

Related references

[Aggregates window](#) on page 31

Managing aggregates

Adding disks to an aggregate

You can add hot spare disks to an aggregate to replace disks or to increase its size so that it can provide more storage space to its contained volumes. You can use the **Add Disks to Aggregate** dialog box to add hot spare disks to an existing aggregate.

Before you begin

An appropriate number of hot spare disks of the matching disk type, RPM, and checksum type must be available.

About this task

You cannot add disks to a mirrored aggregate if one of the plexes is offline.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays either the cluster hierarchy or the storage controllers hierarchy.
3. In the navigation pane, click **Storage > Aggregates**.
4. Select an aggregate to which you want to add disks and click **Add Disks**.
5. Click **Select Disks** and specify the number of matching spare disks that you want to add to the aggregate and click **OK**.
6. In the **Advanced** area, specify the RAID group to which you want to add the disks.
You can also modify the RAID group size, if required.
7. Click **Add**.
8. Verify that the disks are added to the aggregate.

Related references

[Aggregates window](#) on page 31

Editing aggregate settings

You can use the Edit Aggregate dialog box to change the aggregate name and RAID type. You can edit the RAID type to change the level of protection.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays either the cluster hierarchy or the storage controllers hierarchy.
3. In the navigation pane, click **Storage > Aggregates**.
4. Select the aggregate that you want to edit and click **Edit**.
5. In the **Edit Aggregate** dialog box, make the necessary changes.
6. Click **Save and Close**.
7. Verify the changes you made to the selected aggregate.

Related references

[Aggregates window](#) on page 31

Moving volumes

You can move a FlexVol volume to a different aggregate or node within the same Vserver.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays either the cluster hierarchy or the storage controllers hierarchy.
3. In the navigation pane, click **Storage > Aggregates**.
4. Select an aggregate and click **Volume Move**.
5. Type or select information as requested by the wizard.
6. Confirm the details and click **Finish** to complete the wizard.

Upgrading to a 64-bit aggregate

For storage systems running Data ONTAP 8.1 and later, you can upgrade 32-bit aggregates to 64-bit aggregates by increasing their size beyond 16 TB. System Manager enables you to add disks to a 32-bit aggregate and upgrade it when the aggregate size crosses 16 TB.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays either the cluster hierarchy or the storage controllers hierarchy.
3. In the navigation pane, click **Storage > Aggregates**.
4. Select the 32-bit aggregate whose size you want to increase and click **Add Disks**.
5. Select one or more disks and click **Add**.
6. Select the confirmation check box and click **Add**.

This opens the **Upgrade to 64 bit aggregate Wizard** for upgrading the 32-bit aggregate.

7. Type or select information as requested by the wizard.
8. Confirm the details and click **Finish** to complete the wizard.

Related references

[Aggregates window](#) on page 31

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. Ensure that the left navigation pane displays either the cluster hierarchy or the storage controllers hierarchy.
3. In the navigation pane, click **Storage > Aggregates**.
4. Select the aggregate that you want to view information about from the displayed list of aggregates.
5. 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 31
- [Aggregate list](#) on page 31
- [Details area](#) on page 32

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.
Add Disks	Opens the Add Disks to Aggregate dialog box, which enables you to add a disk to an existing 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.
Volume Move	Starts the Volume Move wizard, which enables you to move a volume without disrupting client access.
Refresh	Updates the information in the window.

Aggregate list

The aggregate list displays the name and the space usage information for each aggregate.

Name	Specifies the name of the aggregate.
Node	Specifies the name of the node, which contains the aggregate. This is available only at the cluster level.

Used (%)	Specifies the percentage of space used in the aggregate.
Available Space	Specifies the available space in the aggregate.
Used Space	Specifies the amount of space that is used for data in the aggregate.
Total Space	Specifies the total space of the aggregate.
Volume Count	Specifies the number of volumes associated with the aggregate.
Disk Count	Specifies the number of disks used to create the aggregate.
Status	Specifies the current status of 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.
Disk Layout tab	Displays disk layout information, such as the disk type, RAID type, and RAID group for the selected aggregate.
Volumes tab	Displays information about the volumes associated with the select aggregate, including the list of associated volumes and the available space in each volume.

Related tasks

[Creating aggregates](#) on page 27

[Deleting aggregates](#) on page 27

[Editing aggregate settings](#) on page 29

[Adding disks to an aggregate](#) on page 28

[Upgrading to a 64-bit aggregate](#) on page 30

Storage > Disks

Understanding disks

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.

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 data reconstruction information within RAID groups.

dParity disk Stores double-parity information within RAID groups, if RAID-DP is enabled.

How Data ONTAP reports disk types

Data ONTAP associates a type with every disk. This type is used to determine whether a disk can be used as a matching spare or added to an aggregate. Data ONTAP reports some disk types differently than the industry standards.

When this guide refers to a disk type, it is the type used by Data ONTAP unless otherwise specified.

The following table shows Data ONTAP disk types and how they map to industry standard disk types.

Note: Data ONTAP disk types are different from RAID disk types.

Data ONTAP disk type	Disk connection type	Industry standard disk type	Description
ATA	FC-AL	SATA	
BSAS	SAS	SATA	Bridged SAS–SATA disks with added hardware to enable them to be plugged into a SAS shelf.
FCAL	FC-AL	FC	
LUN	FC (point-to-point or switched)	LUN	A logical storage device backed by third-party storage and used by Data ONTAP as a disk. In this document, these LUNs are referred to as <i>array LUNs</i> to distinguish them from the LUNs that Data ONTAP serves to clients.
SAS	SAS	SAS	

Data ONTAP disk type	Disk connection type	Industry standard disk type	Description
SATA	SAS	SATA	Available only as internal disks for the 20xx systems.
SSD	SAS	SSD	Solid-state disks

Disks using Storage Encryption have a Data ONTAP disk type of SAS. However, they cannot be intermixed 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 HA partner node) must use Storage Encryption.

For a specific configuration, the disk types supported depend on the storage system model, the disk shelf type, and the I/O modules installed in the system. For more information about the types of disks supported by your configuration, see the *System Configuration Guide*.

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*](#)

[System Configuration Guide: now.netapp.com/NOW/knowledge/docs/hardware/NetApp/syscfg/](http://now.netapp.com/NOW/knowledge/docs/hardware/NetApp/syscfg/System Configuration Guide)

Available disk capacity by disk size

To maintain compatibility across different brands of disks, Data ONTAP rounds down ("right-sizes") the amount of space available for user data.

Because of right-sizing, informational commands such as `sysconfig` show a lower number for available space than the disk's rated capacity (you use rated capacity if you specify disk size when creating an aggregate). The available disk space is rounded down as shown in the following table.

Note: For this table, GB = 1,000 MB.

The capacity numbers in this table do not take into account the 10 percent of disk space that Data ONTAP reserves for its own use.

Disk size	Right-sized capacity	Available blocks
FC disks		
144 GB	136 GB	278,528,000
300 GB	272 GB	557,056,000
450 GB	418 GB	856,064,000

Disk size	Right-sized capacity	Available blocks
600 GB	560 GB	1,146,880,000
ATA disks—FC connected		
250 GB	211 GB	432,901,760
320 GB	274 GB	561,971,200
500 GB	423 GB	866,531,584
750 GB	635 GB	1,301,618,176
1 TB	847 GB	1,735,794,176
2 TB	1,695 GB	3,472,315,904
SAS disks—SAS connected		
300 GB	272 GB	557,056,000
450 GB	418 GB	856,064,000
600 GB	560 GB	1,146,880,000
SATA disks—SAS connected		
500 GB	423 GB	866,531,584
750 GB	635 GB	1,301,618,176
1 TB	847 GB	1,735,794,176
2 TB	1,695 GB	3,472,315,904
3 TB	2,538 GB	5,198,943,744
SSD disks—SAS connected		
100 GB	84 GB	173,208,064

Disk speeds supported by Data ONTAP

For hard disk drives, which use rotating media, speed is measured in revolutions per minute (RPM). Faster disks provide more disk input/output operations per second (IOPS) and faster response time.

It is best to use disks of the same speed in an aggregate.

Data ONTAP supports the following rotational speeds for disks:

- FCAL disks (FC-AL connected)
 - 10K RPM
 - 15K RPM

- ATA disks (FC-AL connected)
 - 5.4K RPM
 - 7.2K RPM
- SAS disks (SAS-connected)
 - 15K RPM
- SATA disks (SAS-connected)
 - 7.2K RPM
- BSAS disks (SAS-connected)
 - 7.2K RPM

Solid-state disks, or SSDs, are flash memory-based devices and therefore the concept of rotational speed does not apply to them. SSDs provide more IOPS and faster response times than rotating media.

For more information about supported disk speeds, see the *System Configuration Guide*.

Related information

TR 3285: 15,000 RPM Fibre Channel Disk Drives: A Best-Practice Guide for Optimizing System Performance

System Configuration Guide: now.netapp.com/NOW/knowledge/docs/hardware/NetApp/syscfg/

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.

RAID protection levels for disks

Data ONTAP supports two levels of RAID protection for disks in native disk shelves, RAID-DP and RAID4. RAID-DP can protect against double-disk failures or failures during reconstruction. RAID4 can protect against single-disk failures. You assign RAID level on a per-aggregate basis.

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

Related information

TR 3437: Storage Best Practices and Resiliency Guide

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.

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 it does not provide storage protection. The storage arrays provide the data protection, not Data ONTAP.

When creating "RAID groups" on storage arrays, 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.

About degraded mode

When a disk fails, Data ONTAP can continue to serve data, but it must reconstruct the data from the failed disk using RAID parity. When this happens, the affected RAID group is said to be in *degraded mode*. The performance of a storage system with one or more RAID groups in degraded mode is decreased.

A RAID group goes into degraded mode in the following scenarios:

- A single disk fails in a RAID4 group.
After the failed disk is reconstructed to a spare, the RAID group returns to normal mode.
- One or two disks fail in a RAID-DP group.
If two disks have failed in a RAID-DP group, the RAID group goes into *double-degraded mode*.
- A disk in a RAID4 group is taken offline by Data ONTAP.
After the offline disk is brought back online, the RAID group returns to normal mode.

Note: If another disk fails in a RAID-DP group in double-degraded mode or a RAID4 group in degraded mode, data loss could occur (unless the data is mirrored). For this reason, always minimize the amount of time a RAID group is in degraded mode by ensuring that appropriate hot spares are available.

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.

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 possible.

Note: The replacement disk is downsized to match the size of the disk it is replacing; the extra capacity is not available.

- If the hot spares are not the correct speed, Data ONTAP uses one that is a different speed.

Note: 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 with a faster disk is not a cost-effective solution.

- If SyncMirror is in use and the hot spares are not in the correct pool, Data ONTAP uses a spare from the other pool.

Note: Using drives from the wrong pool is not optimal because you no longer have fault isolation for your SyncMirror configuration. Warning messages go to the logs and console to alert you to this issue.

- The hot spare must be of the same disk type (FC, SAS, and so on) as the failed disk, or of a type that is considered to be equivalent.
- If the hot spares are ZCS disks, they can be added only to zoned checksum aggregates. ZCS disks cannot be added to block checksum aggregates.

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.

How many hot spares you should have

At a minimum, 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 for all disks provides the following benefits:

- At least two hot spares must be available in order to put a disk into the maintenance center.

- 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.

Note: One disk can be the hot spare for multiple disks.

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.

In most cases, the default RAID group size is the best size for your RAID groups. However, you can change the maximum size of your RAID groups.

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

Configuring an optimum RAID group size for an aggregate requires making trade-offs among several different factors. Adding more data disks to a RAID group increases the striping of data across those disks, which typically improves I/O performance. Additionally, a smaller percentage of disks is used for parity rather than data. However, with more disks in a RAID group, there is a greater risk that one of the disks might fail.

Note: With RAID-DP, you can use larger RAID groups because they offer more protection. A RAID-DP group is more reliable than a RAID4 group that is half its size, even though a RAID-DP group has twice as many disks. Thus, the RAID-DP group provides better reliability with the same parity overhead.

Large RAID group configurations offer the following advantages:

- More data drives available. An aggregate configured into a few large RAID groups requires fewer drives reserved for parity than that same aggregate configured into many small RAID groups.
- Small improvement in storage system performance. Write operations are generally faster with larger RAID groups than with smaller RAID groups.

Small RAID group configurations offer the following advantages:

- Shorter disk reconstruction times. In case of disk failure within a small RAID group, data reconstruction time is usually shorter than it would be within a large RAID group.
- Decreased risk of data loss due to multiple disk failures. The probability of data loss through double-disk failure within a RAID4 group or through triple-disk failure within a RAID-DP group is lower within a small RAID group than within a large RAID group.

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.

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.

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 spare disks must be available.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays either the cluster hierarchy or the storage controllers hierarchy.
3. In the navigation pane, click **Storage > Disks**.
4. Select the appropriate number of spare disks and click **Create Aggregate**.

You must retain at least one spare disk as hot spare on the system.

5. Specify a name for the aggregate.
6. Select the RAID type.

The options that are enabled depend on the number of disks that are selected and whether the storage system is a V-Series system or a VSA system.

7. Click **Create**.

Related references

[Disks window](#) on page 43

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 so that it can provide more storage space to its contained volumes.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays either the cluster hierarchy or the storage controllers hierarchy.
3. In the navigation pane, click **Storage > Disks**.
4. Select one or more spare disks that you want to add to the aggregate and click **Add to Aggregate**.
5. Select the aggregate to which you want to add the spare disks and click **Add**.

Related references

[Disks window](#) on page 43

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. Ensure that the left navigation pane displays either the cluster hierarchy or the storage controllers hierarchy.

3. In the navigation pane, click **Storage > Disks**.
4. Select the disk that you want to view information about from the displayed list of disks.
5. 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 43
- [Disk list](#) on page 43
- [Details area](#) on page 44

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	Specifies the name of the disk.
State	Specifies the state of the disk.
Node Name	Specifies the name of the node that contains the aggregate. This is available only at the cluster level.
Type	Specifies the type of the disk.
Firmware Version	Specifies the firmware version of the disk.
RPM	Specifies the speed of the disk drive.
Effective Size	Specifies the usable space available on the disk.
Physical Space	Specifies the total physical space of the disk.
Aggregate	Specifies the aggregate to which this disk belongs.

Shelf	Specifies the shelf in which the physical disks are located.
Bay	Specifies the bay within the shelf for the physical disk.
Pool	Specifies the name of the pool to which the selected disk is assigned.

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 41

[Adding disks to an aggregate](#) on page 42

Storage > Array LUNs

Understanding array LUNs

About disks and array LUNs

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 third-party storage arrays provide to storage systems that run 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. After 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.

Note: Data ONTAP prevents you from mixing native disks and array LUNs in the same aggregate.

For aggregates for third-party storage, you cannot mix the following storage types 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 your vendor for information about how Data ONTAP defines family members for the vendor.

- Array LUNs from storage arrays with 4-Gb HBAs and array LUNs from storage arrays with 2-Gb HBAs
- Array LUNs from Fibre Channel and SATA drives

You can deploy Fibre Channel and SATA drives behind the same V-Series system. However, you cannot mix array LUNs from SATA disks and Fibre Channel 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.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.
3. In the navigation pane, click **Storage > Array LUNs**.
4. Select one or more spare array LUNs and click **Create Aggregate**.

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

5. Specify a name for the aggregate and click **Create**.

Related references

[Array LUNs window](#) on page 47

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. Ensure that the left navigation pane displays the cluster hierarchy.
3. In the navigation pane, click **Storage > Array LUNs**.
4. Select one or more spare array LUNs that you want to add to the aggregate and click **Add to Aggregate**.
5. Select the aggregate to which you want to add the spare array LUNs and click **Add**.

Related references

[Array LUNs window](#) on page 47

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 48
- [Array LUN list](#) on page 48
- [Details area](#) on page 48

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.
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.
Node name	Specifies the name of the node to which this array LUN belongs.
Array name	Specifies the name of the array.

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 46

[Adding array LUNs to an aggregate](#) on page 47

Configuration > Security > Users

Understanding users

User accounts (cluster administrators only)

You can create, modify, lock, unlock, or delete a cluster user account, reset a user's password, or display information about all user accounts.

You can manage cluster user accounts in the following ways:

- Creating a login method for a user by specifying the user's account name, the access method, the authentication method, and, optionally, the access-control role that the user is assigned
- Displaying users' login information, such as the account name, allowed access method, authentication method, access-control role, and account status
- Modifying the access-control role that is associated with a user's login method

Note: It is best to use a single role for all the access and authentication methods of a user account.

- Deleting a user's login method, such as the access method or the authentication method
- Changing the password for a user account
- Locking a user account to prevent the user from accessing the system
- Unlocking a previously locked user account to enable the user to access the system again

Roles

You can use an access-control role to control the level of access a user has to the system. In addition to using the predefined roles, you can create new access-control roles, modify them, delete them, or specify account restrictions for the users of a role.

Managing users

Adding a cluster user account

You can add a cluster user account and specify a login user method to access the storage system.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.
3. Click **Configuration > Security > Users**.
4. Click **Add**.
5. Type the user name for the new user.
6. Type the password that the user uses to connect to the storage system, then confirm the password.
7. Add one or more user login methods, and then click **Add**.

Editing a cluster user account

You can edit a cluster user account by modifying the user login methods to access the storage system.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.
3. Click **Configuration > Security > Users**.
4. Select the user account that you want to modify, and then click **Edit**.
5. Modify one or more user login methods, and then click **Modify**.

Changing passwords for cluster user accounts

You can use System Manager to reset the password for a cluster user account.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.
3. Click **Configuration > Security > Users**.
4. Select the user account for which you want to modify the password and click **Change Password**.
5. In the **Reset Password** dialog box, type the new password, confirm the new password, and then click **Change**.

Locking or unlocking cluster user accounts

You can either lock or unlock cluster user accounts, as required.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.
3. Click **Configuration > Security > Users**.
4. Select the user account whose account status you want to modify and click either **Lock** or **Unlock**, as appropriate.

Window descriptions

Users window

You can use the Users window to manage user accounts, reset a user's password, or display information about all user accounts.

Command buttons

Add	Opens the Add User dialog box, which enables you to add user accounts.
Edit	Opens the Modify User dialog box, which enables you to modify user login methods.
	Note: It is best to use a single role for all access and authentication methods of a user account.
Change Password	Opens the Reset Password dialog box, which enables you to reset the user password.
Lock	Locks the user account.
Refresh	Updates the information in the window.

Users list

User	Specifies the name of the user account.
Account Locked	Specifies whether the user account is locked.

Details area

The details area specifies the user login methods used by the selected user to access the storage system.

Application	Specifies the access method that a user can use to access the storage system. The supported access methods include the following: • System console (console) • HTTP(S) (http) • Data ONTAP API (ontapi) • Service Processor or RLM (service-processor) • SSH (ssh)
Authentication	Displays the default supported authentication method, which is "password."
Role	Specifies the role of a selected user.

Configuration > Security > Roles

Understanding roles

Roles and permissions

The cluster administrator can restrict a user's access to only a specified set of commands by creating a restricted access-control role and assigning it to a user.

You can manage access-control roles in the following ways:

- Creating an access-control role and specifying the command or command directory that the role's users can access.
- Controlling the level of access the role has for the command or command directory and specifying a query that applies to the command or command directory.
- Modifying an access-control role's access to a command or command directory.
- Displaying information about access-control roles, such as the role name, the command, or command directory that a role can access, the access level, and the query.
- Deleting an access-control role.
- Restricting a user's access to only a specified set of commands.
- Displaying Data ONTAP APIs and their corresponding CLI commands.

Predefined roles for cluster administrators

The predefined roles that Data ONTAP provides include `admin`, `readonly`, and `none`.

The following table describes the Data ONTAP predefined roles and their level of access.

Role name	Level of access the role has for the specified command directory	Default granted capabilities
<code>admin</code>	<code>all</code>	Grants all possible capabilities.
<code>readonly</code>	<code>readonly</code>	Grants read-only capabilities.
<code>none</code>	<code>none</code>	Grants no capabilities.

You can create additional roles by using the `security login role create` command.

Managing roles

Editing roles (cluster administrator only)

You can modify an access-control role's access to a command or command directory and restrict a user's access to only a specified set of commands.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.
3. Click **Configuration > Security > Roles**.
4. Select the role that you want to modify, and then click **Edit**.
5. Modify role attributes, and then click **Modify**.

Adding roles (cluster administrators only)

You can add an access-control role and specify the command or command directory that the role's users can access. You can also control the level of access the role has to the command or command directory and specify a query that applies to the command or command directory.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.
3. Click **Configuration > Security > Roles**.
4. Click **Add**.
5. Type the role name and add the role attributes.
6. Click **Add**.

Window descriptions

Roles window

You can use the Roles window (**Configuration > Security > Roles**) to manage roles for user accounts.

Command buttons

- | | |
|-------------|--|
| Add | Opens the Add Role dialog box, which enables you to create an access-control role and specify the command or command directory that the role's users can access. |
| Edit | Opens the Edit Role dialog box, which enables you to add or modify role attributes. |

Refresh Updates the information in the window.

Roles list

The roles list provides a list of roles that are available to be assigned to users.

Details area

The details area displays the role attributes, such as the command or command directory that the selected role can access, the access level, and the query.

Configuration > Services > LDAP

Understanding LDAP

Using LDAP services

Data ONTAP supports LDAP for user authentication, file access authorization, user lookup and mapping services between NFS and CIFS. If the Vserver is set up to use LDAP as a name service using the `-ns-switch ldap` option or for name mapping using the `-nm-switch ldap` option, you should create an LDAP configuration for it.

About this task

An LDAP server enables you to centrally maintain user information.

If you store your user database on an LDAP server, you can configure your Vserver to look up user information in the LDAP database. For example, on your LDAP server, you can store logins and passwords for administrative users of the console and the rsh, telnet, http, https, and ssh protocols, making it possible for you to centrally manage them.

Requirements for Active Directory LDAP servers

You need several things to use Active Directory as your LDAP server.

You need these things to use Active Directory as your LDAP server:

- A valid CIFS license
- Your storage system joined to an Active Directory domain
- A two-way trust relationship established between your storage system's domain and your LDAP server's domain, if they are different

Managing LDAP

Adding an LDAP client configuration

You can add LDAP client configuration to use the LDAP services. You must first set up an LDAP client to use LDAP services.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.
3. Click **Configuration > Services > LDAP**.
4. Click **Add**.
5. Type the name of the LDAP client.
6. Add either the Active Directory domain or the Active Directory server.
7. Click **Binding** and provide the authentication details.
8. Click **Save and Close**.

Editing an LDAP client configuration

You can modify an LDAP client configuration, which enables you to retrieve user information.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.
3. Click **Configuration > Services > LDAP**.
4. Select the LDAP client.
5. Click **Edit**.
6. Modify the LDAP client configuration.
7. Click **Save and Close**.

Deleting an LDAP client configuration

You can delete an LDAP client configuration when you do not want a Vserver to be associated with it.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.

3. Click **Configuration > Services > LDAP**.
4. Select the LDAP client.
5. Click **Delete**.
6. Select **OK to delete the selected LDAP client(s)** and click **Delete**.

Window descriptions

LDAP window

You can use the LDAP window to create LDAP clients for user authentication, file access authorization, user lookup, and mapping services between NFS and CIFS.

Command buttons

- Add** Opens the Create LDAP Client dialog box, which enables you to create and configure LDAP clients at the cluster level. This command button is not available at the Vserver level.
- Edit** Opens the Edit LDAP Client dialog box, which enables you to edit LDAP client configurations at the cluster level. You can use this command button to edit active LDAP clients at the Vserver level.
- Delete** Opens the Delete LDAP Client(s) dialog box, which enables you to delete LDAP client configurations at the cluster level. You can use this command button to delete an active LDAP client at the Vserver level.
- Refresh** Updates the information in the window.

LDAP client list

Specifies the details of LDAP clients in a tabular format. This list is available only at the cluster level.

- | | |
|---|---|
| LDAP Client Config | Specifies the name of the LDAP client configuration that you specified. |
| Active Directory Domain | Specifies the Active Directory domain for each LDAP client configuration. |
| Active Directory Servers | Specifies the Active Directory server for each LDAP client configuration. |
| Preferred Active Directory Servers | Specifies the preferred Active Directory server for each LDAP client configuration. |

Details

The details area displays the LDAP client details such as the Active Directory domain, the Preferred Active Directory Server, and the LDAP client name associated with the Vserver. This area is available only at the Vserver level.

Configuration > Services > Kerberos

Understanding Kerberos

NFS and Kerberos

Kerberos can be used to provide authentication between Vservers and NFS clients.

Kerberos is a network authentication protocol developed at the Massachusetts Institute of Technology. Kerberos uses encryption keys to provide authentication for client/server applications. Authentication provides verification of a user's or process's identity to a server. In the Data ONTAP environment, Kerberos is used to provide authentication between Vservers and NFS and CIFS clients. For more information about Kerberos, see web.mit.edu/kerberos/www/.

Related information

web.mit.edu/kerberos/www/

About Kerberos authentication

With Kerberos authentication, upon connection to your storage system, the client negotiates the highest possible security level. However, if the client is unable to use Kerberos authentication, Microsoft NTLM or NTLM V2 is used to authenticate with the CIFS server.

Managing Kerberos

Creating Kerberos realm configurations

You can use System Manager to create a Kerberos realm configuration, which enables Vservers to use Kerberos security services for NFS.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.
3. Click **Configuration > Services > Kerberos**.
4. Click **Create**.
5. Type or select information as requested by the wizard.

6. Confirm the details and click **Finish** to complete the wizard.

Editing Kerberos realm configurations

You can use System Manager to modify a Kerberos realm configuration at the cluster level.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.
3. Click **Configuration > Services > Kerberos**.
4. Select the Kerberos realm configuration that you want to modify and click **Edit**.
5. Type or select information as requested by the wizard.
6. Confirm the details and click **Finish** to complete the wizard.

Deleting Kerberos realm configurations

You can use System Manager to delete a Kerberos realm configuration.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.
3. Click **Configuration > Services > Kerberos**.
4. Select one or more Kerberos realm configurations that you want to delete and click **Delete**.
5. Select **OK to delete the selected Kerberos realm(s)**.
6. Click **Delete**.

Window descriptions

Kerberos window

You can use the Kerberos window to provide authentication between Vservers and NFS clients.

Command buttons

- | | |
|---------------|---|
| Create | Opens the Kerberos Realm Create wizard, which enables you to configure a Kerberos realm to retrieve user information. This command button is available at the cluster level only. |
| Edit | Opens the Kerberos Realm Edit wizard, which enables you to edit a Kerberos realm configuration based on the need for Vserver authentication and authorization. |

Delete Opens the Delete Kerberos Realm(s) dialog box, which enables you to delete Kerberos realm configurations. This command button is available at the cluster level only.

Refresh Updates the information in the window.

Kerberos configuration list

Specifies the details of Kerberos configuration list in a tabular format.

Configuration Name	Specifies the name of the Kerberos realm configuration.
Realm	Specifies the name of the Kerberos realm.
KDC Vendor	Specifies the name of the Kerberos Distribution Center (KDC) vendor.
KDC IP Address	Specifies the KDC IP address used by the configuration.
Interface Name	Specifies the logical interface associated with the Kerberos configuration for a Vserver.
Service Principal Name	Specifies Service Principal Name (SPN) that matches the Kerberos configuration.
Kerberos Status	Specifies whether Kerberos is enabled.

Details area

The details area displays information such as the KDC port and the administrative-server port number of the selected Kerberos realm configuration. This is available at the cluster level only.

Configuration > System Tools > DateTime

Managing the system time (cluster administrators only)

Problems can occur when the storage system clock is inaccurate. You can manually set the time zone, date, and time on a node to reflect the location of your preference. However, it is best to keep your system time synchronized automatically by using the Network Time Protocol (NTP).

Data ONTAP 8.0 and 8.0.1 use NTP v4 by default. To address situations where certain time servers support only NTP v3, starting with Data ONTAP 8.0.2, the NTP version to be used for communicating with a newly configured NTP server defaults to v3 instead of v4. The NTP Daemon continues to use the highest supported version (v4 in this case) to communicate with the time servers that were configured prior to Data ONTAP 8.0.2.

You can manage the system time with NTP in the following ways:

- Enabling or disabling the use of NTP for the cluster
By default, NTP is disabled.
- Displaying whether NTP is currently enabled for the cluster

- Associating a node with an NTP server
- Specifying the preferred NTP server and version for a node
- Displaying information about NTP servers that are associated with a node or the cluster
- Dissociating a node from an NTP server

You use the `system services ntp config` and `system services ntp server` commands to manage the NTP service. For more information, see the man pages. You can get a list of public NTP time servers from the NTP Public Services page.

If you do not use NTP, you can manually manage the system time in the following ways:

- Setting or modifying the time zone, date, and time on a node
- Displaying the current time zone, date, and time settings of a node or the cluster

You use the `system node date` commands to manage the system time manually. For more information, see the man pages.

Related information

[NTP Public Services - http://support.ntp.org](http://support.ntp.org)

Setting the time zone

You can use the Edit DateTime dialog box to set the time zone for your storage system. However, for an HA configuration, you cannot modify the 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. Ensure that the left navigation pane displays the cluster hierarchy.
3. In the navigation pane, click **Configuration > System Tools > DateTime**.
4. Click **Edit**.
5. Select the time zone.
6. If you want to use the NTP service, select **Enable NTP service** and add the time server.
7. Click **OK**.
8. Verify the changes you made to the date and time settings.

Related references

[DateTime window](#) on page 61

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 time zone](#) on page 60

Configuration > System Tools > Licenses

Managing feature licenses (cluster administrators only)

Your sales or support representative can provide you with license keys that you use to enable certain features on your storage system.

License keys are provided on a per-cluster basis, based on the number of nodes in the cluster. License keys have to be added on each node in the cluster for features to work correctly. If you add new nodes to your cluster, your sales or support representative will provide you with licenses that are valid for the new total number of nodes in the cluster.

Depending on which platform you are using, some features require license keys. Features and products may not require license keys or can be enabled when you add certain other keys for a software bundle.

You can add, delete, and display information about feature licenses by using the `system license` commands.

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, you can add the software license from the **Add License** dialog box.

Before you begin

The software license code for the specific Data ONTAP service must be available.

About this task

The disk sanitization license can be added only from the command-line interface.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.
3. In the navigation pane, click **Configuration > System Tools > Licenses**.
4. Click **Add**.
5. In the **Add License** dialog box, enter the software license code, and click **Add**.

If you are adding a license to an HA configuration, System Manager checks the HA configurations. If the service is not licensed on the partner node, System Manager prompts you to add the license on the partner node.

Result

Your software license is added immediately.

Related references

[Licenses window](#) on page 63

Deleting licenses

You can use the Licenses window to delete an expired software license. However, you cannot delete the disk sanitization license after it is installed on your system.

Before you begin

You must have checked the list of required licenses to ensure that the software license you want to delete is not used by other services or features.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.
3. In the navigation pane, click **Configuration > System Tools > Licenses** to display the **Licenses** window.
4. Select the software license that you want to delete, then click **Delete**.
5. Select the confirmation check box and click **Delete**.

Related references

[Licenses window](#) on page 63

Window descriptions

Licenses window

Your storage system arrives from the factory with pre-installed software. If you want to add or remove a license after you receive the storage system, you can use the Licenses window to add or delete software licenses.

Command buttons

Add Opens the Add License dialog box, which enables you to add new software licenses.

Delete Deletes the software license that you select in the software license list.

Enable Enables the software license that you select in the software license list.

Note: This option is only available in Data ONTAP 8.1 7-Mode.

Disable Disables the software license that you select in the software license list.

Note: This option is only available in Data ONTAP 8.1 7-Mode.

Refresh Updates the information in the window.

Software license list

This list provides the following information about each license installed on your storage system:

Name Displays the name of the software license.

Key Displays the license keys, which enables certain features on your storage system.

State Displays the state of the software license whether it is enabled or disabled.

Limit Displays the limit on the number of nodes for the selected license.

Expires On The expiration date for the software license, if applicable.

Related tasks

[Adding licenses](#) on page 62

[Deleting licenses](#) on page 62

Configuration > System Tools > SNMP

Understanding SNMP

Options when configuring SNMP

Select the desired options when configuring SNMP.

SNMP	Enable or disable the appliance SNMP agent.
Location	Enter the location of the SNMP agent.
Contact	Enter the contact person for the SNMP agent.
Host	Enter the destination host IP address.
Community	Enter the SNMP community name to send with the trap.
SNMPv1/v2c	Enable or disable SNMPv1/v2c.
Read Community	Configure a single read-only community string for SNMPv1 and SNMPv2.
SNMPv3	Enable or disable SNMPv3.
Username and password	Enter a username and password to use with SNMPv3.

Managing SNMP on the cluster (cluster administrators only)

Enabling SNMP provides a mechanism for monitoring your cluster to avoid issues before they occur and to respond to issues when they occur. Managing SNMP involves enabling SNMP, configuring SNMP users, and configuring SNMP traphosts for specific events.

If you enable SNMP in Data ONTAP, SNMP managers can query your storage system's SNMP agent for information. The SNMP agent gathers information and forwards it to the SNMP managers. The SNMP agent also generates trap notifications whenever specific events occur. The SNMP agent on the storage system has read-only privileges—that is, it cannot be used for any set operations or for taking a corrective action in response to a trap. In Data ONTAP Cluster-Mode, SNMP is enabled cluster-wide.

For diagnostic and other network management services, Data ONTAP provides an SNMP agent compatible with SNMP versions v1, v2c, and v3. SNMPv3 offers advanced security by using passphrases and encryption.

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. Ensure that the left navigation pane displays the cluster hierarchy.
3. In the navigation pane, click **Configuration > System Tools > SNMP**.
4. Click **Edit**.
5. In the **General** tab, specify the storage system contact personnel and location, and SNMP communities.
6. Click **OK**.
7. Verify the changes you made to the SNMP settings.

Related references

[SNMP window](#) on page 66

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. Ensure that the left navigation pane displays the cluster hierarchy.
3. In the navigation pane, click **Configuration > System Tools > SNMP**.
4. Click **Edit**.
5. In the **Trap hosts** tab, either select or clear **Enable traps**.
6. If you enable SNMP traps, add the host name or IP address of the hosts to which the traps are sent.
7. Click **OK**.

Related references

[SNMP window](#) on page 66

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 65

[Enabling or disabling SNMP traps](#) on page 65

Configuration > Snapshot Policies

Understanding Snapshot policies

About Snapshot policies

When applied to a volume, a Snapshot policy specifies a schedule or schedules according to which Snapshot copies are made and the maximum number of Snapshot copies that each schedule can create. A Snapshot policy can include up to five schedules.

Managing Snapshot policies

Adding Snapshot policies

You can add a Snapshot policy to specify the frequency and the maximum number of automatically created Snapshot copies.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.
3. Click **Configuration > Snapshot Policies**.
4. Click **Add**.
5. Type the policy name and add one or more Snapshot schedules.
6. Click **Add**.

Editing Snapshot policies

You can make changes to a previously created Snapshot policy if you want to modify its schedule or the maximum number of Snapshot copies that are created.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.
3. Click **Configuration > Snapshot Policies**.
4. Select the Snapshot policy and click **Edit**.
5. In the **Edit Snapshot Policy** dialog box, make the necessary changes.
6. Click **OK**.
7. Verify the changes you made to the selected Snapshot policy.

Deleting Snapshot policies

You can use System Manager to delete Snapshot policies. If you delete a Snapshot policy that is being used by one or more volumes, Snapshot copies of the volume or volumes are no longer created according to the deleted policy.

Before you begin

You must have dissociated the Snapshot policy from each volume that uses it.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.
3. Click **Configuration > Snapshot Policies**.
4. Select the Snapshot policy and click **Delete**.
5. Select the confirmation check box and click **Delete**.

Window descriptions

Snapshot Policies window

You can use the Snapshot Policies window to manage Snapshot policy tasks, such as adding, editing, and deleting Snapshot policies.

Command buttons

- Add** Opens the Add Snapshot Policy dialog box, which enables you to add schedules and specify the maximum number of Snapshot copies to be retained in a policy.
- Edit** Opens the Edit Snapshot Policy dialog box, which enables you to modify the frequency at which Snapshot copies have to be made and the maximum number of Snapshot copies to be retained.
- Delete** Opens the Delete dialog box, which enables you to delete the selected Snapshot policy.
- View as** Enables you to view the Snapshot policies either as a list or as a tree.
- Status** Opens the menu, which you can use to either enable or disable the selected Snapshot policy.
- Refresh** Updates the information in the window.

Snapshot policy list

- | | |
|---|--|
| Policy/Schedule Name | Specifies the name of the Snapshot policy and the schedules in the policy. |
| Status | Specifies the status of the Snapshot policy. |
| Maximum Snapshots to be retained | Specifies the maximum number of Snapshot copies to be retained. |

Configuration > Schedules

Understanding schedules

Schedules

You can configure many tasks (for instance, volume Snapshot copies and mirror replications) to run on specified schedules. Schedules that are run at specified schedules are known as *cron* schedules because of their similarity to UNIX cron schedules. Schedules that are run at intervals are known as *interval* schedules.

You can manage schedules in the following ways:

- Creating a cron schedule or an interval schedule
- Displaying information about all the schedules
- Modifying a cron schedule or an interval schedule
- Deleting a cron schedule or an interval schedule

You cannot delete a schedule that is currently in use by a running job.

The cluster administrator can perform all the schedule management tasks.

Managing schedules

Creating schedules

You can create schedules to run a job at a specific time or at regular periods.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.
3. Click **Configuration > Schedules**.
4. Click **Create**.
5. Type the schedule name and choose the schedule.
6. Click **Create**.

Editing schedules

You can make changes to a previously created schedule if it does not meet your requirements.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.

3. Click **Configuration > Schedules**.
4. Select the schedule that you want to modify and click **Edit**.
5. Modify the schedule details and click **OK**.

Deleting schedules

You can use System Manager to delete the schedules.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.
3. Click **Configuration > Schedules**.
4. Select the schedule that you want to delete and click **Delete**.
5. Modify the schedule details and click **OK**.
6. Select the confirmation check box and click **Delete**.

Window descriptions

Schedules window

You can use the Schedules window to manage scheduled tasks, such as creating, displaying information about, modifying, and deleting schedules.

Command buttons

- | | |
|----------------|--|
| Create | Opens the Create Schedule dialog box, which enables you to create time-based and interval schedules. |
| Edit | Opens the Edit Schedule dialog box, which enables you to edit the selected schedules. |
| Delete | Opens the Delete Schedule dialog box, which enables you to delete the selected schedules. |
| Refresh | Updates the information in the window. |

Schedules list

- | | |
|-------------|---|
| Name | Specifies the name of the schedule. |
| Type | Specifies the type of the schedule--time-based or interval-based. |

Details area

The details area displays information about when a selected schedule is run.

Configuration > Peers

Understanding cluster peers

What a cluster peer is

The cluster peer feature allows two clusters to coordinate and share resources between them.

Connecting one cluster to another cluster in a peer relationship

You connect clusters together in a cluster peer relationship to share information and to provide access to operations on the peer cluster.

About this task

Connecting clusters together requires network ports, network interfaces configured with the intercluster role, and creating the cluster peer relationship.


Cluster peer network topologies

You want to connect clusters together in such a way that the clusters in a peer relationship will always be able to communicate with each other.

The best network topology when you have multiple intercluster LIFs connecting clusters in a peering relationship is full mesh connectivity. Full mesh connectivity means that all of the intercluster LIFs of a cluster can communicate with all of the intercluster LIFs on all of the clusters to which you want it to communicate.

For example, Cluster D has the following LIF connections to Clusters S1, S2, and S3:

- D2a is connected to S1a and S1b, S2a and S2b, and S3a and S3b
- D2b is connected to S1a and S1b, S2a and S2b, and S3a and S3b




With this topology, there are no missing routes from Cluster D to Clusters S1, S2, and S3.

If, for some reason, you cannot configure full mesh connectivity between all of the LIFs of one cluster to other clusters, you can configure full mesh connections between some LIFs of one cluster to other clusters. This provides the full connectivity between clusters, but can result in an initial and temporary slowing in performance and in the systems issuing EMS warnings. The slowed performance can occur because Data ONTAP might need to define the route from one cluster to another before data can transfer.

For example, Cluster D has the following LIF connections to Clusters S1, S2, and S3:

- D1a is connected to S1a and S1b
- D1b is connected to S1a and S1b
- D2a is connected to S2a and S2b
- D2b is connected to S2a and S2b
- D3a is connected to S3a and S3b
- D3b is connected to S3a and S3b




With this topology, routes in which LIFs are directly connected to each other are defined. If there is not a direct connection, for example, when data must go between D1a to S3b, a small amount of time is taken for Data ONTAP to determine a path between Cluster D and Cluster S3.

If the network topology used is not full mesh connectivity, the peer relationships will fail because not all of the possible routes between intercluster LIFs exist.

For example, Cluster D has the following LIF connections to Cluster S1, S2, and S3:

- D2a to S1a, S2a, and S3a
- D2b to S1b, S2b, and S3b



Managing cluster peers

Editing cluster peer relationships

You can edit a cluster peer relationship to enable new nodes to participate in the intercluster communication.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.
3. Click **Configuration > Peers**.
4. Select the cluster peer and click **Edit**.
5. In the dialog box to edit the cluster peer relationship, make the necessary changes.
6. Click **OK**.
7. Verify the changes you made to the selected cluster peer.

Creating cluster peer relationships

You can create a cluster peer relationship to connect clusters together so that the clusters in a peer relationship can communicate with each other.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.
3. Click **Configuration > Peers**.
4. Click **Create**.
5. Type or select information as requested by the wizard.
6. Confirm the details and click **Finish** to complete the wizard.

Deleting cluster peer relationships

You can delete a cluster peer relationship from the Peer window.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the cluster hierarchy.
3. Click **Configuration > Peers**.
4. Select the cluster peer that you want to delete, and then click **Delete**.
5. Select the confirmation check box and click **Delete**.

Window descriptions

Peers window

You can use the Peers window to manage peer relationships, which enable you to move data from one cluster to another.

Command buttons

- Create** Starts the Cluster Peer Create wizard, which enables you to establish a relationship with a remote cluster.
- Edit** Opens a dialog box, which enables you to edit settings related to a cluster peer.
- Delete** Opens the Delete Cluster Peer Relationship dialog box, which enables you to delete the selected peer cluster relationship.

Refresh Updates the information in the window.

Peer cluster list

Peer Cluster Specifies the name of the peer cluster in the relationship.

Availability Specifies whether peer cluster is available for communication.

Details area

The details area displays detailed information about the selected peer cluster relationship.

Diagnostics > Job

Understanding jobs

Jobs

Jobs are asynchronous task and typically long-running volume operations, such as copying, moving, or mirroring data. Jobs are placed in a job queue and are run when resources are available. The cluster administrator can perform all the tasks related to job management.

A job can be one of the following categories:

- A *server-affiliated* job is placed in queue by the management framework to be run in a specific node.
- A *cluster-affiliated* job is placed in queue by the management framework to be run in any node in the cluster.
- A *private* job is specific to a node and does not use the replicated database (RDB) or any other cluster mechanism.

You require the advanced privilege level or higher to run the commands to manage private jobs.

You can manage jobs in the following ways:

- Displaying job information, including the following:
 - Jobs on a per-node basis
 - Cluster-affiliated jobs
 - Completed jobs
 - Job history
- Monitoring a job's progress
- Displaying information about the initialization state for job managers.

You can determine the outcome of a completed job by checking the event log.

Window descriptions

Job window

You can use the Job window to manage job tasks such as displaying job information and monitoring a job's progress.

Command button

Refresh Updates the information in the window.

Tabs

Current Jobs This tab displays information about the job tasks that are in progress.

Job History This tab displays information about the completed job tasks.

Job list

Job ID	Displays ID of the job.
Start Time	Displays the start time of the job.
Job Name	Displays the name of the job.
Controller Name	Displays the name of the node.
State	Displays the state of the job.
Job Description	Displays the description of the job.
Progress	Displays the state of the job.
Schedule Name	Displays the name of the schedule
Job Log ID	Displays the progress of the job with respect to the schedule.
Completion Message	Displays the completion message for the job.
End Time	Displays the end time of the job.

Diagnostics > System Health

Understanding system health

What health monitors are available

In addition to the overall System Health Monitor, there currently are two individual health monitors called System Connectivity and Node Connectivity, which are both for the Storage subsystem.

Health monitor name (identifier)	Subsystem name (identifier)	Purpose
System Connectivity (system-connect)	Storage (SAS-connect)	Monitor shelves at the cluster level to ensure that all shelves always have appropriate pathing to two HA clustered nodes
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

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.

Monitoring the health of your system

You can proactively manage your system by monitoring a single, integrated health status for the entire cluster. If the status is Degraded, you can drill down to individual health monitors to view details about the problem, including the probable cause and recommended recovery actions. After you resolve the problem, the system status automatically returns to OK.

Currently, there are three health monitors: an overall System Health Monitor and two individual health monitors 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. Ensure that the left navigation pane displays the cluster hierarchy.
3. In the navigation pane, click **Diagnostics > System Health**.
4. Click the arrow icon next to the name of subsystem.
5. Select the alert that you want to acknowledge and click **Acknowledge**.
6. Type your name and click **Acknowledge**.

Related references

[System Health window](#) on page 80

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. Ensure that the left navigation pane displays the cluster hierarchy.
3. In the navigation pane, click **Diagnostics > System Health**.
4. Click the arrow icon next to the name of subsystem.
5. Select the alert that you want to suppress and click **Suppress**.
6. Type your name and click **Suppress**.

Related references

[System Health window](#) on page 80

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. Ensure that the left navigation pane displays the cluster hierarchy.
3. In the navigation pane, click **Diagnostics > System Health**.
4. Click the arrow icon next to the name of subsystem.
5. Select the alert that you want to delete and click **Delete**.
6. Click **OK**.

Related references

[System Health window](#) on page 80

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 the alerts from the window.

Command buttons

Acknowledge	Acknowledges the selected alert to indicate that the problem is being addressed and identifies the person who clicks the button as the "Acknowledger."
Suppress	Suppresses 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 with the number of alerts.
Alert ID	Displays the alert ID
Controller	Displays the name of the affected storage controller in the cluster.

Severity	Displays the severity of the alert.
Resource	Displays the resource that is affected such as a specific shelf or disk.
Time	Displays the time when the alert was raised.

Details area

The details area displays detailed information about the alert, such as the time when the alert was raised and whether anyone else has acknowledged the alert already. The details also include information about the probable cause, possible effect, and corrective actions.

Related tasks

[Acknowledging system health alerts](#) on page 79

[Suppressing system health alerts](#) on page 79

[Deleting system health alerts](#) on page 80

Diagnostics > Logs > Events

Window descriptions

Events window

You can use the Events window to view the event log and configure event notification.

Command buttons

Config Opens the Events Configuration dialog box, which enables you to configure event notification.

Refresh Updates the information in the window.

Events list

Time Displays the time at which the event occurred.

Node Displays the node on which the event occurred.

Severity Displays the severity of the event. Severity levels and their explanations are as follows:

- EMERGENCY—The system is unusable
- ALERT—Action must be taken immediately
- CRITICAL—Critical condition
- ERROR—Error condition
- WARNING—Warning condition

- NOTICE—Normal but significant condition
- INFORMATIONAL—Information message
- DEBUG—Debugging message

Source Displays the source of the event.

Event Displays the detail description of the event.

Details area

The area below the events list displays the event details including event description, message name, and sequence number of the selected event.

Diagnostics > Logs > SnapMirror Log

SnapMirror files

SnapMirror uses configuration files, log files, and other files for its operation.

The following table lists the files used by SnapMirror, along with their corresponding functions:

File	Function
/etc/snapmirror.conf	Enables you to specify SnapMirror source and destination relationships, along with the following settings: • SnapMirror update schedules for a relationship • Type of relationship; single path, multipath, or failover • Other options for a given SnapMirror relationship
/etc/snapmirror.allow	Enables you to specify the SnapMirror destinations that are allowed to copy from the system. Note: You can also use the options <code>snapmirror.access</code> command to specify the allowed destinations. However, if both the options <code>snapmirror.access</code> command and the <code>/etc/snapmirror.allow</code> file are used, options <code>snapmirror.access</code> takes precedence. This can affect the initialization of SnapMirror relationships.
/etc/log/snapmirror.x	Records the SnapMirror data transfer history. Note: There might be one or more SnapMirror log files. The latest logs are stored in the file named <code>snapmirror</code>. The older logs are named <code>snapmirror.0</code> and <code>snapmirror.1</code>.

File	Function
/etc/hosts	SnapMirror uses the entries in this file to resolve host names.

For more information about SnapMirror files, see the following man pages:

- `na_snapmirror.conf`
- `na_snapmirror.allow`
- `na_hosts`

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 from which data is mirrored in a SnapMirror relationship.
Destination	Displays the volume to which data is mirrored in a SnapMirror relationship.
Date time	Displays the date and time of the SnapMirror operation.
Transfer Type	Displays the transfer type associated with the SnapMirror operation.
Action	Displays the name of the event.
Message	Displays the message related to the event.
Job Number	Displays the job number associated with the SnapMirror operation.
Log level	Displays the log level.
Log source	Displays the source of the log.

Vserver

What a Vserver is

A Vserver is a secure virtual storage server, which supports multiple protocols and unified storage. A Vserver contains data volumes and one or more LIFs through which it serves data to the clients.

A Vserver securely isolates the shared virtualized data storage and network, and appears as a single dedicated server to the clients. Each Vserver has a separate administrator authentication domain and can be managed independently by a Vserver administrator.

A Vserver represents a single file system and a unique global namespace. A global namespace enables the NAS clients to access data without specifying the physical location of the data. Global namespace also enables the cluster and Vserver administrators to manage distributed data storage as a single file system.

The volumes of each Vserver are related through junctions and are mounted on junction paths. The file system of each volume appears to be mounted at the junctions. The root volume of a Vserver is at the top level of the namespace hierarchy; additional volumes are mounted to the Vserver's root volume to extend the global namespace.

A Vserver's data volumes can contain files and LUNs. Vserver provides file-level data access by using NFS and CIFS protocols for the NAS clients, and block-level data access by using iSCSI and FCP protocols for SAN hosts.

Vserver management

You must configure the services and protocols for a Vserver to serve data. To delegate the Vserver administration to a Vserver administrator, you must set up a password and unlock the vsadmin user account.

The cluster administrator can perform all Vserver management tasks. The cluster administrator can create, configure, modify, rename, or delete a Vserver, display information about the Vserver, or start or stop the Vserver.

The Vserver administrator can perform only the following Vserver administrative tasks:

- Displaying information about the Vserver
- Configuring data access protocols
- Configuring services, such as NIS and DNS
- Managing volumes and Snapshot copies
- Viewing jobs
- Monitoring network configuration

For more information about Vserver administrator capabilities, see the *Data ONTAP Cluster-Mode Vserver Administrator Capabilities Overview Guide*.

Types of Vservers

A cluster consists of three types of Vservers, which help in managing the cluster and its resources, and the data access to the clients and applications.

A cluster contains the following types of Vservers:

- Admin Vserver
- Node Vserver
- Data Vserver

The cluster setup process automatically creates the admin Vserver for the cluster and the node Vserver for each node in the cluster. The admin Vserver represents the cluster, and node Vservers represents the individual nodes of the cluster.

The Data Vserver represents the data serving Vservers. After the cluster setup, a cluster administrator must create cluster or data Vservers and add volumes to these Vservers to facilitate data access from the cluster. A cluster must have at least one data Vserver to serve data to its clients.

Note: Data Vserver is also known as Cluster Vserver.

Why you use Vservers

Vservers provide data access to clients without regard to physical storage or heterogeneous structure, similar to any storage system. Vservers provide nondisruptive operation, scalability, security and supports unified storage.

A Vserver has the following benefits:

- Nondisruptive operation
Vservers are resilient, can operate continuously and nondisruptively for as long as they are needed. Vservers help cluster to operate continuously across software and hardware upgrades, addition and removal of nodes, and all administrative operations.
- Scalability
Vservers support on demand data throughput and the storage requirements.
- Security
A Vserver appears as a single independent server, thus allowing multiple Vservers to co-exist while ensuring no data flow between them.
- Unified Storage
Vservers can serve data concurrently through multiple data access protocols. A Vserver provides file-level data access by using NAS protocols, such as CIFS and NFS, and block-level data access by using SAN protocols, such as iSCSI and FCP. A Vserver can serve data to SAN and NAS clients independently at the same time.
- Delegation of management
A Vserver can have its own user and administration authentication. Vserver administrators can manage the Vservers that they are authorized to access. However, Vserver administrators have less privileges than cluster administrators.

Managing Vservers

Creating Vservers

You can create a Vserver to set up client access to data.

Before you begin

- The cluster must have at least one non-root aggregate.
- You must have synchronized the time across the cluster by configuring and enabling NTP to prevent CIFS creation and authentication failures.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the cluster.
4. Click **Create**.
5. Type or select information as requested by the wizard.
6. Confirm the details and click **Finish** to complete the wizard.

Result

The Vserver you created is started automatically. By default, the vsadmin user account is created and is in the locked state. The vsadmin role is assigned to the default vsadmin user account.

Editing Vserver settings

You can use System Manager to edit Vserver settings.

About this task

You can modify the values of the following properties of a Vserver:

- Name Service Switch
- Name Mapping Switch
- List of protocols to serve data

Note: The CIFS protocol that is configured on a Vserver continues to serve data even when you disallow it on that Vserver.

- List of the aggregates available to create volumes

You can specify the aggregate names for the Vserver. A Vserver administrator can view the list of available aggregates to create volumes on the Vserver.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the appropriate cluster.
4. Select the Vserver that you want to edit and click **Edit**.
5. Make the necessary changes and click **Save and Close**.

Deleting Vservers

You can delete Vservers that you no longer require from the storage system configuration.

Before you begin

You must have completed the following tasks:

1. Disabled the Snapshot copies, and DP and LS mirrors for all the volumes.
2. Deleted all the igroups that belong to the Vserver manually.
3. Deleted all the volumes in the Vserver.
4. Deleted the CIFS server.
5. Deleted any customized user accounts and roles that are associated with the Vserver.
6. Stopped the Vserver.

About this task

- If you delete a Vserver, the CIFS server and any shares defined on the CIFS server are deleted automatically. Any Windows-to-UNIX or UNIX-to-Windows name mappings defined for the Vserver are also deleted because they require both the CIFS and NFS servers.
- If you want to move data from one Vserver to another Vserver before you delete the first Vserver, you can use SnapMirror commands. For more information about SnapMirror, see the *Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide*.
- If you delete a Vserver that is configured to use Kerberos, or modify a Vserver to use a different Service Principal Name (SPN), the Vserver's original service principal is not automatically deleted or disabled from the Kerberos realm. You must manually delete or disable the principal. You must have the Kerberos realm administrator's user name and password to delete or disable the principal.
- When you delete a Vserver, the associated sis policies are also deleted.

Steps

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

2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the appropriate cluster.
4. Select the Vserver that you want to delete and click **Delete**.
5. Select the confirmation check box and click **Delete**.

After you finish

Reenable the Snapshot copies, and DP and LS mirrors for all the volumes.

Starting Vservers

You can use System Manager to provide data access from a Vserver by starting the Vserver.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the appropriate cluster.
4. Select the Vserver that you want to start and click **Start**.

Result

The Vserver starts serving data to clients.

Stopping Vservers

You can stop a Vserver if you want to troubleshoot any issue with the Vserver, delete the Vserver, or stop data access from the Vserver. You cannot stop a Vserver during storage failover (SFO).

Before you begin

All the clients connected to the Vserver must be disconnected.

Attention: If any clients are connected to a Vserver when you stop it, data loss might occur.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the appropriate cluster.
4. Select the Vserver that you want to stop and click **Stop**.

Result

The Vserver stops serving data to clients. A Vserver administrator cannot log in to the Vserver.

Window descriptions**Vserver window**

You can use the Vserver window to manage your Vservers and display information about them.

Command buttons

- Create** Opens the Create Vserver wizard, which enables you to create a new Vserver.
- Edit** Opens the Edit Vserver dialog box, which enables you to modify the properties of a selected Vserver.
- Delete** Deletes the selected Vserver or Vservers.
- Start** Starts the selected Vserver
- Stop** Stops the selected Vserver.
- Manage** Manages the selected Vserver.
- Refresh** Updates the information in the window.

Vserver list

The Vserver list displays the name of and allowed protocols on each Vserver.

- Name** Specifies the name of the Vserver.
- Status** Specifies the status of the Vserver.
- Allowed Protocols** Specifies the allowed protocols, such as CIFS and NFS, on each Vserver.

Details area

The area below the Vserver list displays detailed information about the selected Vserver.

Storage > Volumes**Understanding volumes**

What volumes are

Volumes are file systems that hold user data that is accessible using one or more of the access protocols supported by Data ONTAP. System Manager supports two types of volumes: traditional and flexible. However, you can create only flexible volumes (FlexVols) using System Manager.

Understanding node root volumes

A node's root volume contains special directories and configuration files for that node.

The root volume is installed at the factory on FAS systems and on V-Series systems ordered with disk shelves.

The directory name for a node's root volume is `/mroot`, which is accessible only through the systemshell and with guidance from technical support. The root volume is reserved for system files, log files, and core files. You should not store user data in the root volume. Also, you must not change the preconfigured size for the root volume or modify the content of the root directory, unless technical support instructs you to do so. If you need to modify system configurations, you use Data ONTAP commands to do so instead of editing configuration files directly in the root directory.

You cannot set the root volume's fractional reserve to any value other than 100%.

The root aggregate contains the root volume. Starting with Data ONTAP 8.1, new systems are shipped with the root volume in a dedicated, 64-bit root aggregate that contains at least three disks. The root aggregate should contain only the root volume. You must not include or create data volumes in the root aggregate.

By default, the storage system is set up to use a hard disk drive (HDD) aggregate for the root aggregate. When no HDDs are available, the system is set up to use a solid-state drive (SSD) aggregate for the root aggregate.

You should contact technical support if you need to designate a different volume to be the new root volume or move the root volume to another aggregate.

About FlexClone volumes

FlexClone volumes provide writeable volume copies that use only as much storage as necessary to hold new data. They can be created instantaneously without interrupting access to the parent FlexVol volume. Additionally, they use the space very efficiently, storing only data that changes between the parent and the clone.

A FlexClone volume is initialized with a Snapshot copy and updated continuously thereafter.

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

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




Figure 1: Space savings of test and development storage without and with FlexClone volumes

Related information

Storage Management Guide — http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml

How FlexClone volumes work

FlexClone volumes can be managed similarly to regular FlexVol volumes, with a few key differences.

The following list outlines some key facts about FlexClone volumes:

- A FlexClone volume is a point-in-time, writable copy of the parent volume. Changes made to the parent volume after the FlexClone volume is created are not reflected in the FlexClone volume.

- FlexClone volumes are fully functional volumes; you manage them just as you do the parent volume.
 - A FlexClone volume is always created in the same aggregate as its parent.
 - A FlexClone volume is always created in the same Vserver as its parent.
 - A FlexClone volume can itself be cloned to create another FlexClone volume.
 - A FlexClone volume and its parent share the same disk space for any common data.
This means that 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 space guarantee as its parent.
The space 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.
 - 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.
- Quotas applied to the parent volume are *not* automatically applied to the FlexClone volume.
 - 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 7-Mode Storage Management Guide*.

Related information

Data ONTAP documentation on NOW - http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml

How splitting a FlexClone volume from its parent works

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

It is helpful to know the following about clone-splitting operations:

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

About creating a FlexClone volume from volumes currently in a SnapMirror relationship

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 your need for the FlexClone volume is temporary, and you can accept the temporary cessation of SnapMirror replication, you can create the FlexClone volume and either delete it or split it from its parent when possible. At that time, the SnapMirror replication continues normally.
- If you cannot accept the temporary cessation of SnapMirror replication, 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 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 FlexClone files or FlexClone LUNs (LUN created by using the `volume file clone create` command), those FlexClone files and FlexClone LUNs are present

in the FlexClone volume and they share storage with the FlexClone files and FlexClone LUNs in the parent volume.

How flexible volumes work

Flexible volumes (also called 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.

What a Snapshot copy is

A Snapshot copy is a read-only image of a traditional volume, a FlexVol 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 back up and restore data.

For information about traditional volumes, FlexVol volumes, or aggregates, see the *Data ONTAP 7-Mode Storage Management Guide*.

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 setting is 20 percent of the total (both used and unused) space on the disk. However, the default setting for a root FlexVol volume is 5 percent of the total disk space.

Snapshot configuration

You can configure Snapshot copies by setting a schedule to an existing Snapshot policy. In Data ONTAP Cluster-Mode, you can have a maximum of 128 Snapshot copies of a regular FlexVol parent volume. You can change the maximum number of Snapshot copies for a Snapshot policy's schedule.

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 Snapshot copy reserve can be used only by Snapshot copies, not by the active file system.

If the active file system runs out of disk space, any disk space still remaining in the Snapshot copy reserve is not available for use by the active file system.

Note: Although the active file system cannot consume disk space reserved for Snapshot copies, Snapshot copies can exceed the Snapshot copy reserve and consume disk space normally available to the active file system.

Managing the Snapshot copy reserve involves the following tasks:

- Ensuring that enough disk space is set aside 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

Data compression and deduplication


Beginning with Data ONTAP 8.0.1, data compression is supported with deduplication.

When both data compression and deduplication are enabled on a FlexVol volume, the data is first compressed and then deduplicated. Therefore, deduplication can further increase the space savings by removing duplicate blocks in the FlexVol volume.

Though data compression and deduplication can be enabled on a FlexVol volume, the savings might not be the sum of the savings when each is run individually on a data set. The combined savings can yield higher savings than running deduplication or data compression individually.

You can achieve better savings when you run the data compression scanner before deduplication. This is because data compression scanner cannot run on data that is locked by deduplication, but deduplication can run on compressed data.

The following illustration shows how data is first compressed and then deduplicated:



When you run deduplication on a FlexVol volume that contains uncompressed data, it scans all the uncompressed blocks in the FlexVol volume and creates a digital fingerprint for each of the blocks.

Note: If a FlexVol volume has compressed data, but the compression option is disabled on that volume, then you might lose the space savings when you run the `sis undo` command.

FlexClone volumes and shared Snapshot copies

When space 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 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 – 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 space guarantee of the FlexClone volume.

What volume guarantees are

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, for only the reserved files or LUNs in the volume, or whether no space at all is preallocated for the volume.

The volume guarantee is an attribute of the volume. It is persistent across storage system reboots, takeovers, and givebacks. Volume guarantee values can be volume (the default value), file, or none.

- A volume guarantee 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. This is the default value.

A volume guarantee of volume provides the easiest management strategy-- the volume cannot run out of space before the amount of data it contains reaches the size of the volume. However, this manageability comes at the cost of increased storage requirements.

- A volume guarantee of file reserves space for the volume from its containing aggregate so that any file in the volume with space reservation enabled 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 do not have space reservation enabled could run out of space.

- A FlexVol volume that has a volume guarantee of none reserves no space for user data; space is allocated from the aggregate on an as-needed basis, and 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. This approach to space management is called *thin provisioning*.

When space in the aggregate is reserved for space 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 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 committed to another volume.

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

Note: Volume guarantees are honored only for online volumes. If you take a volume offline, any committed but unused space for that volume becomes available for other volumes in that aggregate. When you bring that volume back online, if there is not sufficient available space in the aggregate to fulfill its space guarantees, you must use the force option, and the volume's space guarantees are disabled. When a volume's space guarantee is disabled, this is reflected in the output of the `volume show` command.

What kind of space management to use

The type of space management you should use 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 space guarantee of volume OR Traditional volumes	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 will always succeed.

If...	Then use...	Typical usage	Notes
Writes to certain files must always succeed You want to overcommit your aggregate	FlexVol volumes with a space guarantee of file OR Traditional volume AND space reservation enabled for files that require writes to succeed	LUNS Databases	This option enables you to guarantee writes to specific files.
You need even more effective storage usage than file space reservation provides You actively monitor available space on your volume and can take corrective action when needed Snapshot copies are short-lived Your rate of data overwrite is relatively predictable and low	FlexVol volumes with a space guarantee of volume OR Traditional volume AND Space reservation on for files that require writes to succeed AND 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 space 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 overcommit your aggregate You actively monitor available space on your aggregate and can take corrective action when needed	FlexVol volumes with a space guarantee of none	Storage providers who need to provide storage that they know will not immediately be used Storage providers who need to allow available space to be dynamically shared 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 for different purposes in a given aggregate, you do not actually allocate any space for those volumes in advance. The space is allocated only when the application host needs it.

The unused aggregate space is available for the thinly provisioned volumes to expand or for creating new volumes. 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 is a pool of storage that 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, dramatic 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 500 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 truly needs it. The rest of the storage remains in the free pool to be used as needed. Storage administrators can set thresholds, so they are alerted when more disks need to be added to the pool.

The following illustration compares thin provisioning with traditional provisioning:




Figure 2: Thin provisioning versus traditional provisioning

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

Data ONTAP can automatically make more free space available for a FlexVol volume when that volume is nearly full. You can choose to make the space available by first allowing the volume size to increase, or by first 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. (`space-mgmt-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. (`space-mgmt-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 Cluster-Mode Data Protection Guide*.

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 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 <code>/etc</code> 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 Cluster-Mode File Access and Protocols Management Guide*.

Using deduplication to increase storage efficiency

Deduplication is an optional feature of Data ONTAP that can improve physical storage space 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, 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, on any FlexVol volume.

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




Figure 3: How deduplication removes data redundancies

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.
- 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.
- 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.

Options for resizing volumes

You can use the Volume Resize wizard to change your flexible 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 shows the current space allocations within the volume, including how much space is used and how much is free. 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 how much space 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.

Configuring volumes

Creating volumes

You can create a new flexible volume (FlexVol volume) by using the Create Volume dialog box. Using multiple FlexVol volumes enables you to perform administrative tasks on some of your data, while the rest remains available to users. You cannot create traditional volumes by using System Manager.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Volumes**.
4. Click **Create**.
5. If you want to change the default name, specify a new name.
6. Select the containing aggregate for the volume.

7. Select the type of storage for which you are creating this volume.

Select **Data Protection** if you are creating a SnapMirror destination volume. You will have read-only access to this volume.

8. Specify the size of the volume and the percentage of the total volume size that you want to reserve for Snapshot copies.

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.

9. If you want to enable thin provisioning for the volume, select **Thin Provisioned**.

With thin provisioning, space is allocated to the volume from the aggregate only when data is written to it.

10. If you want to enable deduplication, compression, or both on this volume, make the necessary changes in the **Storage Efficiency** tab.

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.

11. Click **Create**.

Related references

[Volumes window](#) on page 118

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 flexible 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 and flexible volume.

About this task

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. Ensure that the left navigation pane displays the Vservers hierarchy.

3. In the navigation pane, select the Vserver and click **Storage > Volumes**.
4. Select the volume from the volume list.
5. Click **Clone > Create > Volume**.
6. In the **Create FlexClone Volume** dialog box, type the name of the FlexClone volume you want to create.
7. 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.
8. 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.
9. Click **Clone**.

Related references

[Volumes window](#) on page 118

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.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Volumes**.
4. From the **Clone** menu, click **Create > File**.
5. Select the file that you want to clone and specify a name for the FlexClone file.
6. Click **Clone**.

Result

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

Related references

[Volumes window](#) on page 118

Deleting volumes

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

Before you begin

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

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Volumes**.
4. Select one or more volumes that you want to delete and click **Delete**.
5. Select the confirmation check box and click **Delete**.

Related references

[Volumes window](#) on page 118

Setting the Snapshot copy reserve

You can reserve space (in percentage) in a volume for Snapshot copies.

About this task

The default space reserved for Snapshot copies is zero percent for SAN and VMware volumes.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Volumes**.
4. Select the volume for which you want to set the Snapshot copy reserve.
5. Click **Snapshot Copies > Configure**.
6. 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 118

Deleting Snapshot copies

You can delete a Snapshot copy to conserve space or free disk space, or if the Snapshot copy 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 displays the status "busy, vclone" in the parent volume.
- You cannot delete locked Snapshot copies because the SnapMirror or SnapVault relationship requires these copies for the next update.

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

Steps

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

Related references

[Volumes window](#) on page 118

Creating Snapshot copies

You can create a Snapshot copy of a volume 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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Volumes**.
4. Select the volume on which you want to create the Snapshot copy.
5. Click **Snapshot Copies > Create**.
6. In the **Create Snapshot Copy** dialog box, if you want to change the default name, specify a new name for the Snapshot copy.

Valid characters are ASCII characters, numerals, hyphens (-), underscores (_), periods (.) and the plus symbol.

The default name of a Snapshot copy consists of the volume name and the timestamp.

7. Click **Create**.

Related references

[Volumes window](#) on page 118

Managing volumes

Editing volume properties

You can use the **Edit Volume** dialog box to modify volume properties, such as the volume name, security style for the volume, fractional reserve, and space guarantee settings. You can also modify storage efficiency settings (deduplication schedule, enable compression) and space reclamation settings (volume autogrow and Snapshot copies autodelete).

About this task

System Manager enables you to set the fractional reserve to either zero percent or 100 percent.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Volumes**.

4. Select the volume whose properties you want to modify and click **Edit**.
5. Click the appropriate tab to display the properties or settings you want to change.
6. Make the necessary changes.
7. Click **Save and Close** to save your changes and close the dialog box.

Related references

[Volumes window](#) on page 118

Changing your volume status

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

About this task

If you want a volume to be the target of a volume copy or SnapMirror replication operation, you have to put the volume into restricted state. 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.

For NAS volumes, you need to unmount the volume before you can take it offline.

Steps

1. Ensure that the left navigation pane displays the Vservers hierarchy.
2. In the navigation pane, select the Vserver and click **Storage > Volumes**.
3. In the navigation pane, click **Storage > Volumes**.
4. Select the volume for which you want to modify the status.
5. From the **Status** menu, click the volume status you want.
6. In the confirmation dialog box, click the button for the volume status that you want.

Related references

[Volumes window](#) on page 118

Configuring deduplication on a volume

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

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Volumes**.
4. Select the volume for which you want to configure deduplication.
5. Click **Edit**, and then click **Storage Efficiency**.
6. Select **Enable storage efficiency**.
7. Select one of the following schedules:
 - On-demand
 - Automated
Deduplication is automatically run when 20 percent new data is written to the volume.
 - Scheduled
8. If you choose the **Scheduled** option, select the storage efficiency policy.
9. Click **Save and Close** to save your changes.

Related references

[Volumes window](#) on page 118

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Volumes**.
4. Select the volume for which you want to modify the deduplication schedule.
5. Click **Edit**, and click **Storage Efficiency**.

6. Change the deduplication schedule as required.
7. Click **Save and Close** to save your changes.

Related references

[Volumes window](#) on page 118

Running deduplication

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

Before you begin

- Deduplication must be enabled on the volume.

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 need to abort deduplication before you can perform any other operation.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Volumes**.
4. Select the volume for which you want to run deduplication.
5. Click **Storage Efficiency**.
6. 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**.
7. Click **Start**.

Related references

[Volumes window](#) on page 118

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 need 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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Volumes**.
4. Select the FlexClone volume that you want to split from its parent volume.
5. Click **Clone > Split**.
6. Confirm the clone split operation and click **Start Split** in the confirmation dialog box.

Related references

[Volumes window](#) on page 118

Resizing FlexVol volumes

When your FlexVol volume becomes nearly full, 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. You can also decrease the size of FlexVol volumes.

About this task

For a volume that is set to grow automatically, you can modify the limit to which the volume can grow automatically, taking into account the increased size of the volume.

Steps

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

2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Volumes**.
4. Select the volume that you want to resize.
5. Click **Resize**.
6. Type or select information as requested by the wizard.
7. Confirm the details and click **Finish** to complete the wizard.

Related references

[Volumes window](#) on page 118

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. Keep in mind that restoring a Snapshot copy 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

- SnapRestore license must be installed on your system.
- If the 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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Volumes**.
4. Select the volume that you want to restore from a Snapshot copy.
5. Click **Snapshot Copies > Restore**.
6. Select the appropriate Snapshot copy and click **Restore**.
7. Select the confirmation check box and click **Restore**.

Related references

[Volumes window](#) on page 118

Scheduling automatic Snapshot copies

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

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Volumes**.
4. Select the volume from the volume list.
5. Click **Snapshot Copies > Configure**.
6. Select **Enable scheduled snapshots**.
7. Select a Snapshot policy.
Only policy based Snapshot policies are available.
8. Click **OK** to save your changes and start your Snapshot copy schedule.

Related references

[Volumes window](#) on page 118

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

Valid characters are ASCII characters, numerals, hyphens (-), underscores (_), periods (.) and the plus symbol.

Related references

[Volumes window](#) on page 118

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Volumes**.
4. Select the volume for which you want hide the Snapshot copy directory.
5. Click **Snapshot Copies > Configure**.
6. Ensure that **Make snapshot directory (.snapshot) visible** is not selected, and then click **OK**.

Related references

[Volumes window](#) on page 118

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Volumes**.
4. Select the volume from the volume list.
5. 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 118

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 select the Snapshot copy you want to rename, restore, or delete.

Before you begin

The volume must be online.

About this task

You can view Snapshot copies for only one volume at a time. If you select more than one volume, the lower pane displays the following message: Select a single item from the table to view the item details.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Volumes**.
4. In the upper pane of the **Volumes** window, select the volume for which you want to see Snapshot copies.
5. In the lower pane, click **Snapshot Copies**.

Window descriptions

Volumes window

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

- [Command buttons](#) on page 118
- [Volume list](#) on page 119
- [Details area](#) on page 120

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.
Move	Opens the Move Volume dialog box, which you can use to move volumes from one aggregate or node to another within the same virtual server.
Refresh	Updates the information in the window.

Volume list

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

Name	Specifies the name of the volume.
Aggregate	Specifies the name of the aggregate.
Status	Specifies the status of the volume.

Thin Provisioned	Specifies whether space guarantee is set for the selected volume. Valid values for online volumes are "Yes" and "No."
Type	Specifies the type of volume--rw for read-write, ls for load sharing, or dp for data protection.
Root volume	Specifies whether the volume is a root volume.
% Used	Specifies the amount of space (in percentage) that is used in the volume.
Available Space	Specifies the available space in the volume.
Total Space	Specifies the total space in the volume. This includes space that is reserved for Snapshot copies.
Storage Efficiency	Specifies whether deduplication is enabled or disabled for the selected volume.
Clone	Specifies whether the volume is a FlexClone 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 capacity 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.
- Details
Displays information about deduplication properties, including whether deduplication is enabled on the volume, the deduplication status, and the current schedule.
- 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 volumes](#) on page 105

[Creating FlexClone volumes](#) on page 106

[Creating FlexClone files](#) on page 107

[Deleting volumes](#) on page 108
[Setting the Snapshot copy reserve](#) on page 108
[Deleting Snapshot copies](#) on page 109
[Creating Snapshot copies](#) on page 110
[Editing volume properties](#) on page 110
[Changing your volume status](#) on page 111
[Configuring deduplication on a volume](#) on page 112
[Changing the deduplication schedule](#) on page 112
[Running deduplication](#) on page 113
[Splitting a FlexClone volume from its parent volume](#) on page 114
[Resizing FlexVol volumes](#) on page 114
[Restoring a volume from a Snapshot copy](#) on page 115
[Scheduling automatic Snapshot copies](#) on page 116
[Renaming Snapshot copies](#) on page 116
[Hiding the Snapshot copy directory](#) on page 117
[Viewing FlexClone volumes hierarchy](#) on page 117

Storage > Namespace

Managing namespace

Mounting volumes

You can use the Namespace window in System Manager to mount volumes to a junction in Vserver namespace.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Namespace**.
4. Click **Mount** and select the volume to be mounted.
5. If you want to change the default junction name, specify a new name.
6. Click **Browse** and select a junction path to mount the volume.
7. Click **OK** and click **Mount**.

Unmounting volumes

You can use Namespace window in System Manager to unmount volumes from a junction in Vserver's namespace.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Namespace**.
4. Select the volume that needs to be unmounted and click **Unmount**.
5. Select the confirmation check box and click **Unmount**.

Changing export policies

You can change the export policy associated with a volume to redefine the client access to data.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Namespace**.
4. Select the volume and click **Change Export Policy**.
5. Select the export policy and click **Change**.

Window descriptions

Namespace window

You can use the Namespace window to manage the NAS namespace of a Vserver.

Command buttons

Mount	Opens the Mount Volume dialog box, which enables you to mount a volume to a junction in a Vserver's namespace.
Unmount	Opens the Unmount Volume dialog box, which enables you to unmount a volume from its parent volume.
Change Export Policy	Opens the Change Export Policy dialog box, which enables you to change the existing export policy associated with a volume.
View as:	Enables you to the namespace list either as a list or as a tree.

Refresh Updates the information in the window.

Namespace list

Path Specifies the junction path of the mounted volume.

Volume Specifies the name of the volume mounted on the path.

Export Policy Specifies the export policy of the mounted volume.

Security Style Specifies the security style for the volume. Possible values include UNIX (for UNIX mode bits), NTFS (for CIFS ACLs), and Mixed (for mixed NFS and CIFS permissions).

Storage > Shares

Understanding shares

CIFS home directory concepts

The Data ONTAP CIFS home directory feature enables you to configure a share that maps to different directories based on the user that connects to it and a set of variables. Instead of having to create separate shares for each user, you can configure a single share with a few home directory parameters to define a user's relationship between an entry point (the share) and their home directory (a directory on the Vserver).

There are four variables that determine how a user is mapped to a directory:

Share name This is the name of the share that you create that the user connects to. It can be static (for example: home), dynamic (for example: %w), or a combination of the two. You must set the homedirectory property for this share.

The share name can use the following dynamic names:

- %w (the user's Windows user name)
- %d (the user's Windows domain name)

Share path This is the relative path, defined by the share and therefore associated with one of the share names, that is appended to each search path to generate the user's entire home directory path from the root of the Vserver. It can be static (for example: home), dynamic (for example: %w), or a combination of the two.

Search path This is the absolute path from the root of a Vserver that you specify to tell Data ONTAP where to search for home directories. You specify one or more search paths by using the `vserver cifs home-directory search-path add` command. If

you specify multiple search paths, Data ONTAP tries them in the order specified until it finds a valid path.

Directory This is the user's actual home directory that you create for the user. It is usually the user's name. You must create it in one of the directories defined by the search paths.

As an example, consider the following setup:

- User: John Smith
- User domain: acme
- User name: jsmith
- Vserver name: vs1
- Home directory share name #1: home - share path: %w
- Home directory share name #2: %w - share path: %d/%w
- Search path #1: /vol/aggr0home/home
- Search path #2: /vol/aggr1home/home
- Search path #3: /vol/aggr2home/home
- Home directory: /vol/aggr1home/home/jsmith

Scenario 1: The user connects to \\vs1\home. This matches the first home directory share name and generates the relative path jsmith. Data ONTAP now searches for a directory named jsmith by checking each search path in order:

- /vol/aggr0home/home/jsmith does not exist. Moving on to search path #2.
- /vol/aggr1home/home/jsmith does exist. Therefore search path #3 is not checked. The user is now connected to his home directory.

Scenario 2: The user connects to \\vs1\jsmith. This matches the second home directory share name and generates the relative path acme/jsmith. Data ONTAP now searches for a directory named acme/jsmith by checking each search path in order:

- /vol/aggr0home/home/acme/jsmith does not exist. Moving on to search path #2.
- /vol/aggr1home/home/acme/jsmith does not exist. Moving on to search path #3.
- /vol/aggr2home/home/acme/jsmith does not exist. The home directory does not exist, therefore the connection fails.

Configuring shares

Creating shares

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

Before you begin

CIFS must be set up and started.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Shares**.
4. Click **Create Share**.
5. Click **Browse** and select the folder, qtree, or volume to be shared.
6. If you want to change the default name, specify a new name.
7. Provide a description for the share and click **Create**.

Result

The share is created with the access permissions set to full control for the Everyone group.

Related tasks

[Setting up CIFS](#) on page 180

Related references

[Shares window](#) on page 127

Stopping share access

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

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Shares**.
4. From the list of shares, select the share that you want to stop sharing and click **Stop Sharing**.
5. Select the confirmation check box and click **Stop**.

Related references

[Shares window](#) on page 127

Creating home directory shares

System Manager enables you to create a home directory share and manage home directory search paths.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Shares**.
4. Click **Create Home Directory** and provide the pattern information that determines how a user is mapped to a directory.
5. Click **Create**.

Managing shares

Editing share settings

You can use System Manager to modify the settings of a share.

Steps

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

Related references

[Shares window](#) on page 127

Window descriptions

Shares window

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

- [Command buttons](#) on page 128
- [Shares list](#) on page 128

- [Details area](#) on page 128

Command buttons

Create Share	Opens the Create Share dialog box, which enables you to create a share.
Create Home Directory	Opens the Create Home Directory Share dialog box, which enables you to create a new home directory 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.
Home Directory	Specifies the name of the home directory 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 shares](#) on page 125

[Stopping share access](#) on page 126

[Editing share settings](#) on page 127

Storage > LUNs

Understanding LUNs

Guidelines for working with volumes that contain LUNs

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

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

When you create a volume, Data ONTAP automatically:

- Reserves 20 percent of the space for Snapshot copies
- Schedules Snapshot copies

Because the internal scheduling mechanism for taking Snapshot copies within Data ONTAP has no means of ensuring that the data within a LUN is in a consistent state, it is recommended that you 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.

It is strongly recommended that you use the following guidelines to create traditional or FlexVol 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.
- Use a SAN volume to contain the LUN.
- Ensure that no other files or directories exist in a volume that contains a LUN.
If this is not possible and you are storing LUNs and files in the same volume, use a separate qtree to contain the LUNs.
- If multiple hosts share the same volume, create a qtree on the volume to store all LUNs for the same host. This is a recommended best practice that simplifies LUN administration and tracking.
- Ensure that the volume option `create_ucose` is set to on.
- To simplify management, 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 more information.

Related information

Data ONTAP documentation on NOW - http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml

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.

LUN multiprotocol type guidelines

The LUN Multiprotocol Type, or operating system type, specifies the OS 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.

Note: Not all Data ONTAP versions support all LUN multiprotocol types. You should consult the Interoperability Matrix to get the most up-to-date information.

The following table describes the guidelines for using each LUN Multiprotocol Type:

LUN Multiprotocol Type	When to use
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 your Solaris Host Utilities documentation and release notes.
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.
hpux	If your host operating system is HP-UX.
aix	If your host operating system is AIX.
linux	If your host operating system is Linux.
netware	Your host operating system is Netware.

LUN Multiprotocol Type	When to use
openvms	If your host operating system is OpenVMS.
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.
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.
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.

Note: If you are using SnapDrive for Windows, the LUN Multiprotocol Type is automatically set.

Related information

[Interoperability Matrix--support.netapp.com/NOW/products/interoperability](http://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 or a specific initiator.
- 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.
- You cannot map a LUN to both FC and iSCSI igroups if ALUA is enabled on one of the igroups.

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 ostyles.

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 ostyles of initiators are `solaris`, `windows`, `hpux`, `aix`, `netware`, `xen`, `hyper_v`, `vmware`, and `linux`.


You must select an ostype for the igroup.

How igroups and port sets limit access to LUNs


You can limit access to your LUNs through igroups and port sets. Until you map LUNs to an igroup, you cannot access to your LUNs. You use port sets to limit which LIFs an initiator can access.

A port set is a collection of LIFs. If you do not bind those igroups to a port set, any initiator within an igroup can access any mapped LUN through any LIF.

The way an initiator connects to a LUN is through an igroup. In the following example, `initiator1` can access `LUN1` through `LIF1` and `LIF2`. Without a port set, `initiator1` can access `LUN1` over any LIF.



You can limit access to a LUN by binding igroups to port sets. In the following example, initiator1 can access LUN1 through LIF1. However, initiator1 cannot access LUN1 through LIF2 because LIF2 is not in portset1.



Configuring LUNs

Creating LUNs

You can use the Create LUN wizard to 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 flexible 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.

You can create a LUN of size up to 16 TB.

Steps

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

Related references

[LUNs window](#) on page 143

Deleting LUNs

You can delete LUNs and return the space that they had used to their containing aggregates or volumes. You can also delete the LUN-containing volume if it has no other LUNs, files, or directories.

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > LUNs**.
4. In the **LUN Management** tab, select one or more LUNs that you want to delete and click **Delete**.
5. Select the confirmation check box and click **Delete**.

Related references

[LUNs window](#) on page 143

Creating initiator groups

You can use the **Create Initiator Group** dialog box to create a new initiator group. Initiator groups enable you to control host access to specific LUNs. You use port sets to limit which LIFs an initiator can access.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > LUNs**.
4. Click **Initiator Groups**, and click **Create**.
5. In the **General** tab, specify the initiator group name, operating system, port set, and supported protocol for the group.
6. In the **Initiators** tab, add the initiators.

7. Click **Create**.

Related references

[LUNs window](#) on page 143

Deleting initiator groups

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

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > LUNs**.
4. Click **Initiator Groups**.
5. Select one or more initiator groups that you want to delete and click **Delete**.
6. Click **Delete**.

Related references

[LUNs window](#) on page 143

Adding initiators

You can use the Edit Initiator Group dialog box to add initiators to an initiator group. An initiator gets 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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > LUNs**.
4. Click **Initiators Groups**.
5. Select the initiator group to which you want to add initiators and click **Edit**.
6. In the **Edit Initiator Group** dialog box, click **Initiators**.
7. Click **Add**.

8. Specify the initiator name and click **OK**.
9. Click **Save and Close**.

Related references

[LUNs window](#) on page 143

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > LUNs**.
4. Click **Initiator Groups**, and click **Edit**.
5. In the **Initiators** tab, select one or more initiators that you want to delete and click **Delete**.

Related references

[LUNs window](#) on page 143

Creating port sets

You can use the Create Portset dialog box to create a new port set.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > LUNs**.
4. Click **Portsets**, then click **Create**.
5. In the **Create Portset** dialog box, select the type of protocol.
6. Choose the network interface that you want to associate with the port set.
7. Click **Create**.

Deleting port sets

You can use the Portsets tab to delete a port set.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > LUNs**.
4. Click **Portsets**, then select a port set.
5. Click **Delete**.
6. Confirm the deletion by clicking **Delete**.

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > LUNs**.
4. Click **LUN Management**.
5. Select the LUN that you want to edit from the list of LUNs and click **Edit**.
6. Make the changes as required.
7. Click **Save and Close** to save your changes and close the dialog box.

Related references

[LUNs window](#) on page 143

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 change the port set associated with the initiator group.

Steps

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

Related references

[LUNs window](#) on page 143

Editing initiators

You can use the **Edit Initiator Group** dialog box to change the name of an existing initiator in an initiator group.

Steps

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

Related references

[LUNs window](#) on page 143

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > LUNs**.
4. In the **LUN Management** tab, select one or more LUNs that you want to bring online.
5. Click **Status > Online**.

Related references

[LUNs window](#) on page 143

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > LUNs**.
4. In the **LUN Management** tab, select one or more LUNs that you want to take offline.
5. Click **Status > Offline**.
6. Click **Offline**.

Related references

[LUNs window](#) on page 143

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 giving 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 thick provisioned LUN clone requires as much space as the thick 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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > LUNs**.
4. Select the LUN that you want to clone and click **Clone**.
5. If you want to change the default name, specify a new one.
6. Click **Clone**.

Related references

[LUNs window](#) on page 143

Editing port sets

You can use the Portsets tab to edit settings related to port sets.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > LUNs**.
4. Click **Portsets**, and select the port set you want to edit.

5. Click **Edit**.
6. Make the necessary changes.
7. Click **Save and Close**.

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > LUNs**.
4. In the **LUN Management** tab, select the LUN that you want to view information about from the displayed list of LUNs.
5. Review the LUN details in the **LUNs** window.

Viewing initiator groups

You can use the Initiator Groups tab to view all the initiator groups, initiators, and LUNs that are mapped to an initiator group.

Steps

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

You can view the LUNs that are mapped to the initiator group in the **Mapped LUNs** tab.

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 143
- [Command buttons](#) on page 143
- [LUNs list](#) on page 143
- [Details area](#) on page 144

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.

Portsets Enables you to create, edit settings of, or delete port sets.

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 134

[Deleting LUNs](#) on page 135

[Creating initiator groups](#) on page 135

[Deleting initiator groups](#) on page 136

[Adding initiators](#) on page 136

[Deleting initiators from an initiator group](#) on page 137

[Editing LUNs](#) on page 138

[Editing initiator groups](#) on page 139

[Editing initiators](#) on page 139

[Bringing LUNs online](#) on page 140

[Taking LUNs offline](#) on page 140

[Cloning LUNs](#) on page 141

Storage > 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 flexible volume. You can create up to 4995 qtrees per volume. There is no maximum for the storage system as a whole. You can easily 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-qtrees basis.

Related concepts

[Qtree name restrictions](#) on page 145

[About the CIFS oplocks setting](#) on page 146

[Security styles](#) on page 146

When you use qtrees

Qtrees allow 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, 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.
- 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

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 144

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 144

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 Cluster-Mode File Access and Protocol Management Guide*.

Related concepts

[Qtree options](#) on page 144

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Qtrees**.
4. Click **Create**.
5. In the **Details** tab, type a name for the qtree.
6. Select the volume to which you want to add this qtree.

The Volume browse list includes only volumes that are online.

7. 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.

8. 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.

9. 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.

10. Click **Create**.

Related references

[Qtrees window](#) on page 149

Deleting qtrees

You can delete a qtree and 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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Qtrees**.
4. Select one or more qtrees that you want to delete and click **Delete**.
5. Select the confirmation check box and click **Delete**.

Related references

[Qtrees window](#) on page 149

Managing qtrees

Editing qtrees

You can use the Edit Qtree dialog box to 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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Qtrees**.
4. Select the qtree that you want to edit and click **Edit**.
5. In the **Edit Qtree** dialog box, edit the settings as required.
6. Click **OK**.
7. Verify the changes you made to the selected qtree in the **Qtrees** window.

Related references

[Qtrees window](#) on page 149

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Qtrees**.
4. Select the qtree that you want to view information about from the displayed list of qtrees.
5. 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 149
- [Qtree list](#) on page 149

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 147

[Deleting qtrees](#) on page 148

[Editing qtrees](#) on page 148

Storage > 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. Quotas are applied to a specific volume or qtree.

Why you use quotas

You can use quotas to limit resource usage, to provide notification when resource usage reaches specific levels, or simply to track resource usage.

You specify a quota for the following reasons:

- To limit the amount of disk space or the number of files that can be used by a user or group, or that can be contained by a qtree
- To track the amount of disk space or the number of files used by a user, group, or qtree, without imposing a limit
- To warn users when their disk usage or file usage is high

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, 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 prepended 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.

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.

A quota is indicated by a target that is either an empty string ("") or an asterisk (*), depending on the context:

- When you create a quota using the `volume quota policy rule create` command, setting the `-target` parameter to an empty string ("") creates a default quota.
- In the output of the `volume quota policy rule show` command, a default quota appears with an empty string ("") as the Target.
- In the output of the `volume quota report` command, a default quota appears with an asterisk (*) as the ID and Quota Specifier .

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.

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 using Data ONTAP, the quota rules for that qtree are automatically updated. If you rename a qtree using your CIFS or NFS client, you must update any quota rules for that qtree.

Note: If you rename a qtree using your CIFS or NFS client and do not update quota rules for that qtree with the new name before you reinitialize quotas, quotas will not be applied to the qtree and explicit quotas for the qtree—including tree quotas and user or group quotas for the qtree—might be converted into derived quotas.

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.

Whether the quota rules persist depends on where you delete the qtree:

- If you delete a qtree using Data ONTAP, the quota rules for that qtree are automatically deleted, including tree quota rules and any user and group quota rules configured for that qtree.
- If you delete a qtree using your CIFS or NFS client, you must delete any quota rules for that qtree to avoid getting errors when you reinitialize quotas. If you create a new qtree with the same name as the one you deleted, the existing quota rules are not applied to the new qtree until 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 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 provide a way 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 new quota and apply it to a specific volume or qtree.

Steps

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

Related references

[Quotas window](#) on page 160

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Quotas**.
4. Select one or more quotas that you want to delete and click **Delete**.
5. Select the confirmation check box and click **Delete**.

Related references

[Quotas window](#) on page 160

Managing quotas

Editing quota limits

You can use the **Edit Limits** dialog box to modify 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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Quotas**.
4. Select the quota that you want to edit and click **Edit Limits**.
5. Modify the quota settings as required.
6. Click **Save and Close** to save your changes and close the dialog box.
7. Verify the changes that you made to the selected quota in the **User Defined Quotas** tab.

Related references

[Quotas window](#) on page 160

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Quotas**.
4. In the **Quota Status on Volumes** tab, select one or more volumes for which you want to activate or deactivate quotas.
5. Click either **Activate** or **Deactivate**, as appropriate.
6. If you are deactivating a quota, select the confirmation check box and click **OK**.

Related references

[Quotas window](#) on page 160

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Quotas**.
4. In the **Quota Status on Volumes** tab, select one or more volumes for which you want to resize the quotas.
5. Click **Resize**.

Related references

[Quotas window](#) on page 160

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Storage > Quotas**.
4. 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 .

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

6. 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 160
- [Command buttons](#) on page 160
- [Quotas list](#) on page 160
- [Details area](#) on page 161

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	Opens the Create Quota wizard.
Edit Limits	Opens the Edit Limits dialog box.
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 157

[Deleting quotas](#) on page 157

[Editing quota limits](#) on page 158

[Activating or deactivating quotas](#) on page 158

[Resizing quotas](#) on page 159

Policies > Export Policies

Understanding export policies

Export policy and rule concepts

Export policies enable you to restrict access to volumes to clients that match specific IP addresses and specific authentication types. Clients cannot access data on a Vserver until you create an export policy and export rules.

Each volume is associated with exactly one export policy. Each export policy is identified by a unique name and a unique numeric ID. A Data ONTAP cluster can contain up to 1,024 export policies. Each Vserver has at least one export policy called `default`, which contains no rules. This export policy cannot be deleted, although it can be renamed or modified. Each volume on a Vserver by default is associated with the default export policy.

Export policies consist of individual *export rules*. An export policy can contain a large number of rules (approximately 4,000). Each rule specifies access permissions to volumes for one or more clients. The clients can be specified by hostname, IP address, or netgroup.

Rules are processed in the order in which they appear in the export policy. The rule order is dictated by the rule index number.

The rule also specifies the authentication types that are required for both read-only and read-write operations. To have any access to a volume, matching clients must authenticate with the authentication type specified by the read-only rule. To have write access to the volume, matching clients must authenticate with the authentication type specified by the read-write rule. If a client makes an access request that is not permitted by the applicable export policy, the request fails with a permission-denied message. If a client IP address does not match any rule in the volume's export policy, then access is denied. If an export policy is empty, then all accesses are implicitly denied.

Export rules can use host entries from a netgroup.

You can modify an export policy dynamically on a running Data ONTAP system.

Managing export policies

Creating export policies

You can create an export policy so that certain clients can access specific volumes.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Policies > Export Policies**.
4. Click **Create Policy** and specify a policy name.
5. If you want to create a new export policy by copying the rules from an existing export policy, select the Vserver and the export policy.
6. Click **Add** to add rules to the policy, and click **Create**.

Deleting export policy rules

System Manager enables you to delete export policy rules.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Policies > Export Policies**.
4. Select the export policy rule that you want to delete and click **Delete Rule**.
5. Click **Delete**.

Deleting export policies

System Manager enables you to delete export policies.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Policies > Export Policies**.
4. Select the export policy that you want to delete and click **Delete Policy**.
5. Select the confirmation check box and click **Delete**.

Modifying export policy rules

System Manager enables you to modify export policy rules.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Policies > Export Policies**.
4. Select the export policy rule and click **Modify Rule**.
5. Make the changes as required.
6. Click **OK**.

Adding rules to export policies

You can add rules to an export policy, which enables you to define client access to data.

Before you begin

You must have created an export policy to add the export rules to.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Policies > Export Policies**.
4. Select the export policy and click **Add Rule**.
5. Specify the client that requires access to the data.
6. If you want to modify the rule index number, select the appropriate rule index number.
7. Select the access protocol.
8. Select one or more security types and access rules.
9. Click **OK**.

Renaming export policies

System Manager enables you to rename export policies.

Steps

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

2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Policies > Export Policies**.
4. Select the export policy that you want to rename and click **Rename Policy**.
5. Specify a new policy name and click **Modify**.

Window descriptions

Export Policies window

You can use the Export Policies window to create, display, and manage information about export policies.

Command buttons

Create Policy	Opens the Create Export Policy dialog box, which enables you to create an export policy and add export rules. You can also copy the export rules from another Vserver.
Rename Policy	Opens the Rename Policy dialog box, which enables you to rename the selected export policy.
Delete Policy	Opens the Delete Export Policy dialog box, which enables you to delete the selected export policy.
Add Rule	Opens the Create Export Rule dialog box, which enables you to add an export rule to the selected export policy.
Modify Rule	Opens the Modify Export Rule dialog box, which enables you to modify the attributes of the selected export rule.
Delete Rule	Opens the Delete Export Rule dialog box, which enables you to delete the selected export rule.
Refresh	Updates the information in the window.

Efficiency policies list

Policy	Specifies the name of the new export policy.
Client Specification	Specifies the client to which the rule applies. You can specify the match in any of the following formats: • As a host name; for instance, host1 • As an IPv4 address; for instance, 10.1.12.24 • As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0.

- As a netgroup, with the netgroup name preceded by the @ character; for instance, @netgroup
- As a domain name preceded by the "." symbol; for instance, .example.com

Note: You must not enter an IP address range, such as 10.1.12.10-10.1.12.70. Entries in this format are interpreted as a text string and treated as a host name. You can enter the IPv4 address 0.0.0.0/0 to provide access to all hosts.

Access Protocols Specifies the access protocol.

You can specify multiple access protocols for an export rule. If you do not specify an access protocol, the default value of "any" is used.

Read-Only Rule Specifies one or more security types for read-only access.

Read-Write Rule Specifies one or more security types for read-write access.

Superuser Specifies the security type or types for superuser access.

Policies > Efficiency Policies

Understanding efficiency policies

Efficiency policy

An efficiency policy is a job schedule for deduplication operation. A job schedule is a time-based cron (For example: every Tuesday at midnight).

Interval based schedule (For example: every 15 hours) cannot be used to create an efficiency policy. There are no default or built-in efficiency policies. You need to create an efficiency policy before enabling storage efficiency on a FlexVol volume.

Managing efficiency policies

Adding efficiency policies

You can use System Manager to add efficiency policies to run the deduplication operation on a FlexVol volume for a specific time and to specify the job schedule.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Policies > Efficiency Policies**.

4. Click **Add** and specify the policy name.
5. Select the schedule. You can specify the maximum runtime duration of the efficiency policy, if required.

System Manager allows efficiency policies only with time-based schedules.

6. Click **Add**.

Editing efficiency policies

System Manager enables you to modify the attributes of an efficiency policy.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Policies > Efficiency Policies**.
4. Click **Edit** and make the changes as required.
5. Click **Edit**.

Deleting efficiency policies

System Manager enables you to delete an efficiency policy that is no longer required.

Before you begin

The efficiency policy must be disabled.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Policies > Efficiency Policies**.
4. Select the efficiency policy that you want to delete and click **Delete**.
5. Select the confirmation check box and click **Delete**.

Enabling or disabling efficiency policies

You can either enable or disable an efficiency policy, as required.

Steps

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

2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Policies > Efficiency Policies**.
4. Select one or more efficiency policies which you want to enable or disable.
5. Click either **Enable** or **Disable**, as appropriate.
6. If you are disabling an efficiency policy, select the confirmation check box and click **OK**.

Window descriptions

Efficiency Policies window

You can use the Efficiency Policies window to create, display, and manage information about efficiency policies.

Command buttons

- Add** Opens the Add Efficiency Policy dialog box, which enables you to run a deduplication operation on a FlexVol volume for a specified duration.
- Edit** Opens the Edit Efficiency Policy dialog box, which enables you to modify the schedule and maximum runtime of the deduplication operation on a FlexVol volume.
- Delete** Opens the Delete Efficiency Policy dialog box, which enables you to delete the selected efficiency policy.
- Status** Open the drop-down menu, which provides options to enable or disable the selected efficiency policy.
- Refresh** Updates the information in the window.

Efficiency policies list

- | | |
|------------------------|--|
| Policy | Specifies the name of an efficiency policy. |
| Status | Specifies the status of an efficiency policy. |
| Schedule | Specifies the name of the schedule for the deduplication operation. |
| Maximum Runtime | Specifies the maximum runtime duration of an efficiency policy. If this value is not specified, the efficiency policy will run till the operation is complete. |

SnapMirror

Understanding SnapMirror technology

How SnapMirrors work

You create a SnapMirror relationship between a source volume and a destination volume. The source volume is a read-write volume that clients can access and modify. The destination volume is a read-only volume that exports a Snapshot copy to clients for read-only access.

Snapshot copies are used to update destination volumes. Snapshot copies are transferred from the source volume to the destination volume using an automated schedule or manually; therefore, mirrors are updated asynchronously. You use the set of `snapmirror` commands to create and manage SnapMirror relationships.

Uses for data protection mirrors

You can create data protection mirrors to solve a number of storage administration tasks, such as backing up data for archiving, recovering data when disasters occur, and distributing data to various sites.

Configuring SnapMirror relationships

Creating SnapMirror relationships

You can use mirroring technology to replicate data from a source volume to a destination volume, at regular intervals or on demand. You can use the **SnapMirror Relationship Create** wizard to create a mirroring relationship between a source volume and a destination volume.

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 cannot be the root volume of a storage 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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **SnapMirror**.

4. Click **Create**.
5. Type or select information as requested by the wizard.
6. Confirm the details and click **Finish** to complete the wizard.
7. Verify that the SnapMirror relationship you created is included in the list of SnapMirror relationships in the **SnapMirror** window.

If the SnapMirror relationship is not initialized 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.

Related references

[SnapMirror window](#) on page 174

Deleting SnapMirror relationships

You can delete a SnapMirror relationship and permanently end a SnapMirror relationship between a source and destination pair of volumes. 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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **SnapMirror**.
4. Select the SnapMirror relationship that you want to delete and click **Delete**.
5. Select the confirmation check box and click **Delete**.

Related references

[SnapMirror window](#) on page 174

Managing SnapMirror relationships

Editing SnapMirror relationship properties

You can use the **Edit SnapMirror Relationship** dialog box to edit the schedule for data transfer for an asynchronous SnapMirror relationship.

About this task

You can use System Manager to edit a SnapMirror relationship.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **SnapMirror**.
4. Select the SnapMirror relationship that you want to edit and click **Edit**.
5. In the **Edit SnapMirror Relationship** dialog box, modify the properties as required.
6. Click **Save and Close** to save your changes and close the dialog box.

Related references

[SnapMirror window](#) on page 174

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 to the destination. You can use the SnapMirror window to initialize a SnapMirror relationship.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **SnapMirror**.
4. Select the SnapMirror relationship that you want to initialize.
5. Click **Operations > Initialize**.
6. Click **Initialize**.

Related references

[SnapMirror window](#) on page 174

Updating SnapMirror relationships

You can use the SnapMirror window to initiate an unscheduled SnapMirror update of the destination. You may 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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **SnapMirror**.
4. Select the SnapMirror relationship that you want to update.
5. Click **Operations > Update**.
6. Click **Update**.

Related references

[SnapMirror window](#) on page 174

Quiescing SnapMirror destinations

A SnapMirror destination is quiesced to stabilize the destination before taking a Snapshot copy. It allows active SnapMirror transfers to finish and temporarily prevents new data transfers. 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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **SnapMirror**.
4. Select the SnapMirror relationship that you want to quiesce.
5. Click **Operations > Quiesce**.
6. Select the confirmation check box and click **Quiesce**.

Related references

[SnapMirror window](#) on page 174

Resuming SnapMirror relationships

You can use the SnapMirror window to resume a quiesced SnapMirror relationship. This allows normal data transfer to the SnapMirror destination and restart all SnapMirror activities.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **SnapMirror**.
4. Select the SnapMirror relationship that you want to resume.
5. Click **Operations > Resume**.

Related references

[SnapMirror window](#) on page 174

Breaking SnapMirror relationships

If the SnapMirror source becomes unavailable or if you wish to use the SnapMirror destination for writing and reading, you can break the SnapMirror relationship. You can use the SnapMirror window to break a SnapMirror relationship and make the destination volume writable. It is best to quiesce the SnapMirror relationship before breaking it.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **SnapMirror**.
4. Select the SnapMirror relationship that you want to break.
5. Click **Operations > Break**.
6. Select the confirmation check box and click **Break**.

Related references

[SnapMirror window](#) on page 174

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.

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **SnapMirror**.
4. Select the SnapMirror relationship that you want to resynchronize.
5. Click **Operations > Resync**.
6. Select the confirmation check box and click **Resync**.

Related references

[SnapMirror window](#) on page 174

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 is converted to a copy of the original destination volume.

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **SnapMirror**.
4. Select the SnapMirror relationship that you want to reverse resynchronize.
5. Click **Operations > Reverse Resync**.

6. Select the confirmation check box and click **Reverse Resync**.

Related references

[SnapMirror window](#) on page 174

Aborting a SnapMirror transfer

You can abort a volume 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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **SnapMirror**.
4. Select the SnapMirror relationship.
5. Click **Operations > Abort**.
6. Select the confirmation check box and click **Abort**.

Related references

[SnapMirror window](#) on page 174

Window descriptions

SnapMirror window

You can use the SnapMirror window to create, display, and manage SnapMirror relationships.

- [Command buttons](#) on page 174
- [SnapMirror relationships list](#) on page 175
- [Details area](#) on page 175

Command buttons

Create	Opens the SnapMirror Relationship Create wizard to create a SnapMirror relationship from a source volume.
Edit	Opens the Edit SnapMirror Relationship dialog box to edit schedule of a SnapMirror relationship.
Delete	Deletes the SnapMirror relationship.
Operations	Displays the operations possible on a SnapMirror relationship.

Refresh Updates the information in the window.

SnapMirror relationships list

Source Specifies the volume from which data is mirrored in a SnapMirror relationship.

Destination Specifies the volume to which data is mirrored in a SnapMirror relationship.

Health Specifies the health of a SnapMirror relationship.

State Specifies the mirror state as "source", or "snapmirrored", or "broken-off."

Status Specifies the relationship status as "idle" or "transferring".

Details area

The details area include SnapMirror relationship details such as data transfer rates, status, and schedule of the relationship.

Related tasks

[Creating SnapMirror relationships](#) on page 168

[Deleting SnapMirror relationships](#) on page 169

[Editing SnapMirror relationship properties](#) on page 170

[Initializing SnapMirror destinations](#) on page 170

[Updating SnapMirror relationships](#) on page 171

[Quiescing SnapMirror destinations](#) on page 171

[Resuming SnapMirror relationships](#) on page 172

[Breaking SnapMirror relationships](#) on page 172

[Resynchronizing SnapMirror relationships](#) on page 173

[Reverse resynchronizing SnapMirror relationships](#) on page 173

[Aborting a SnapMirror transfer](#) on page 174

Configuration > Network Interfaces

Understanding interfaces

Ports and adapters

Ports are grouped under nodes and the nodes are displayed based on the selected protocol category. For example, if the data is served using the FC protocol, then only the nodes with FCP adapters are displayed. The hosted interface count helps you in choosing a port which is less loaded.

Configuring interfaces

Creating network interfaces

You can create a network interface to access the data and manage Vservers.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Network Interfaces**.
4. Click **Create**.
5. Type or select information as requested by the wizard.
6. Confirm the details and click **Finish** to complete the wizard.

Deleting network interfaces

You can delete a network interface by using System Manager.

Before you begin

The status of network interface must be disabled.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Network Interfaces**.
4. Select the interface and click **Delete**.
5. Select the confirmation check box and click **Delete**.

Managing interfaces

Editing network interfaces

You can modify the properties of a network interface by using System Manager.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Network Interfaces**.
4. Select the interface and click **Edit**.
5. In the dialog box to edit interfaces, make the necessary changes.
6. Click **Save and Close**.
7. Verify the changes you made to the selected network interface.

Window descriptions

Network Interfaces

You can use the Network Interfaces window to view a list of network interfaces and create, edit, or delete interfaces from your storage system.

Command buttons

Create	Opens the Network Interface Create wizard, which enables you to create network interfaces to serve data and manage Vservers.
Edit	Opens a dialog box, which enables you to modify the attributes of a network interface.
Delete	Deletes the selected network interface.
Status	Updates the status of the selected network interface.
Send to Home	Sets the current hosted port to home port. This command button is enabled only when the selected interface is hosted on a non-home port and when the home port is available.
Refresh	Updates the information in the window.

Interface list

Interface Name	Specifies the name of a network interface.
Data Protocol Access	Specifies the protocol used to access data.

Management Access	Specifies whether the interface is used to manage Vserver.
IP Address/WWPN	Specifies the IP address or WWPN of the interface.
Current Port	Specifies the name of node and port where the interface is hosted.
Status	Specifies the status of the interface

Details area

The area below the interface list displays detailed information about the selected interface.

Configuration > Protocols > CIFS

Understanding CIFS

CIFS concepts

CIFS clients can access files on a Vserver using the CIFS protocol provided Data ONTAP can properly authenticate the user.

When a CIFS client connects to a Vserver, Data ONTAP authenticates the user with a Windows domain controller. Data ONTAP uses two methods to obtain the domain controllers to use for authentication:

- It queries DNS servers in the domain that the Vserver is configured to use for domain controller information.
- It queries a list of preferred domain controllers you can optionally specify.

Next, Data ONTAP must obtain UNIX credentials for the user. It does this by using mapping rules on the Vserver or a LDAP server, or by using a default UNIX user instead. You can specify for a Vserver which mapping services are searched in which order, or specify a default UNIX user.

Data ONTAP then checks different name services for UNIX credentials for the user, depending on the name services configuration of a 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.

Supported CIFS clients and domain controllers

Before you can use CIFS with your Vserver, you need to know which CIFS clients and domain controllers Data ONTAP supports.

For the latest information on which CIFS clients and domain controllers Data ONTAP supports, see the CIFS (Windows File Services) Compatibility Matrix and the NetApp Interoperability Matrix Tool on support.netapp.com/NOW/products/interoperability.

How CIFS clients can access UNIX symbolic links

You must understand certain concepts about how Data ONTAP enables you to manage symbolic links. This is important to provide access to CIFS users connecting to the storage system.

A symbolic link is a file created in a UNIX environment that contains a reference to another file or directory. If a client accesses a symbolic link, it is redirected to the target file or directory that the symbolic link refers to.

Data ONTAP provides CIFS clients the ability to follow UNIX symbolic links on the storage system. This feature is optional and you can configure it on a per-share basis with one of the following three options:

- Enabled with read/write access
- Enabled with read-only access
- Disabled by hiding symbolic links from CIFS clients

There are two types of symbolic links:

Relative A relative symbolic link contains a reference to the file or directory relative to its parent directory. Therefore the path of the file it is referring to should not begin with a slash (/). A relative symbolic link always refers to a file or directory within the same file system. If you enable symbolic links on a share, relative symbolic links work without further configuration.

Absolute An absolute symbolic link contains a reference to a file or directory in the form of an absolute path. Therefore the path of the file it is referring to should begin with a slash (/). It is treated as an absolute path location of the file from the root of the file system. An absolute symbolic link can refer to a file or directory within or outside of the file system of the symbolic link. If the target is not in the same local file system, the symbolic link is called a widelink. If you enable symbolic links on a share, absolute symbolic links do not work right away. You must first create a mapping between the UNIX path of the symbolic link to the destination CIFS path. When creating absolute symbolic link mappings, you specify whether it is a local or widelink. If you create an absolute symbolic link to a file or directory outside of the local share but set the locality to local, Data ONTAP disallows access to the target.

Note that if a client attempts to delete a local symbolic link (absolute or relative), only the symbolic link is deleted, not the target file or directory. However, if a client attempts to delete a widelink, it might delete the actual target file or directory that the widelink refers to. Data ONTAP does not have control over this since the client can explicitly open the target file or directory outside the storage system and delete it.

Configuring CIFS

Setting up CIFS

You can use System Manager to set up CIFS.

Before you begin

- The CIFS license must be installed on your storage system.
- While configuring CIFS in Active Directory domain:
 - 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 drift between the storage system time and the domain time must not be more than the drift time configured in Data ONTAP.

Note: System Manager currently supports organizational units (OU) for active directory domains only on systems running Data ONTAP 8.1 Cluster-Mode. If you want to include your system in an OU other than the default "Computers", you must use the `CIFS setup` command.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols > CIFS**.
4. In the **Configuration** tab, click **Set up**.
5. Specify the NetBIOS name and the Active Directory domain details as requested by the **CIFS Server Set Up** dialog box.
6. Click **Setup**.

Related tasks

[Creating shares](#) on page 125

Related references

[CIFS window](#) on page 184

Managing CIFS

Editing CIFS general properties

You can modify the general properties for CIFS, such as default UNIX and Windows user.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols > CIFS**.
4. In the **Configuration** tab, click **Options**.
5. In the **Options** dialog box, make the necessary changes and click **OK**.

Related references

[CIFS window](#) on page 184

Adding home directory paths

You can specify one or more paths to be used by the storage system to resolve the location of user' CIFS home directories.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols > CIFS**.
4. Click **Configuration**.
5. In the **Home Directories** area, click **Edit**.
6. In the **Edit Home Directory Search Paths** dialog box, specify the paths used by the storage system to search for users' CIFS home directories.
7. Click **Add**, then click **Save and Close**.

Related references

[CIFS window](#) on page 184

Deleting home directory paths

You can delete a home directory path when you do not want the storage system to use it to resolve the location of user' CIFS home directories.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols > CIFS**.
4. Click **Configuration**.
5. In the Home Directories area, click **Edit**.
6. In the **Edit Home Directory Search Paths** dialog box, select the home directory path that you want to delete and click **Delete**.
7. Click **Save and Close** to save your changes and close the dialog box.

Related references

[CIFS window](#) on page 184

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 storage system's available domain controller discovery information 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 CLI.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols > CIFS**.
4. In the **Domain** tab, click **Reset**.

Related references

[CIFS window](#) on page 184

Updating CIFS group policy configuration

You need to update the group policy after the policy configuring is changed using the CLI. You can use the CIFS window to update the group policy.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols > CIFS**.
4. Click **Domain**.
5. In the Group Policy area, select the group policy configuration that you want to update and click **Update**.

Enabling or disabling CIFS group policy

You can enable or disable CIFS group policy configuration from the CIFS window.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols > CIFS**.
4. Click **Domain**.
5. In the Group Policy area, select the group policy configuration that you want to enable or disable and click either **Enable** or **Disable**, as appropriate.

Reloading CIFS group policy

You need to reload a CIFS group policy if the status of the policy is changed. You can use the CIFS window to reload the group policy.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols > CIFS**.
4. Click **Domain**.
5. In the Group Policy area, select the group policy configuration that you want to reload and click **Reload**.

Monitoring CIFS

Viewing CIFS domain information

You can view information about the domain controllers and servers that the storage system is connected to.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols > CIFS**.
4. Click **Domain**.
5. 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 184
- [Command buttons](#) on page 184
- [Details area](#) on page 185

Tabs

Configuration tab	Displays the Configuration tab, which enables you to manage your CIFS, sessions.
Domain tab	Displays the Domain tab, which enables you to view, test, and reset your CIFS domain controllers. You can use this tab to manage CIFS group policies configurations.
Symlinks	Displays the Symlinks tab, which enables you to manage symbolic links.

Command buttons

Set up	Starts the CIFS Set Up wizard.
Options	Displays the CIFS Options dialog box, which enables you to modify the CIFS properties.
Refresh	Updates the information in the window.

Details area

- CIFS** Specifies the CIFS session details such as the authentication type.
- Home directories** Specifies home directory paths and the style that determines how you want PC user names mapped to home directory entries.

Related tasks

- [Setting up CIFS](#) on page 180
- [Editing CIFS general properties](#) on page 181
- [Adding home directory paths](#) on page 181
- [Deleting home directory paths](#) on page 182
- [Resetting CIFS domain controllers](#) on page 182

Configuration > 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 modify the NFS settings, such as enabling or disabling NFSv2, 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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols > NFS**.
4. Click **Edit**.
5. Make the necessary changes.
6. Click **Save and Close** to save your changes and close the dialog box.

Related references

[NFS window](#) on page 186

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.
Refresh	Updates the information in the window.

Related tasks

[Editing NFS settings](#) on page 186

Configuration > 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 gigabit 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 3720 - www.ietf.org/rfc/rfc3270.txt](http://www.ietf.org/rfc/rfc3270.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.

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.

For iSCSI sessions that use multiple connections, all of the connections must use interfaces in the same target portal group.

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 an interface is added to the storage system, each network interface is automatically assigned to its own target portal group.

In addition, some storage systems support the use of an iSCSI Target expansion adapter, which contains special network interfaces that offload part of the iSCSI protocol processing. You cannot combine these iSCSI hardware-accelerated interfaces with standard iSCSI storage system interfaces in the same 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 Vserver appears as a single iSCSI target node with one iSCSI node name.

On the storage system, the interface can be an Ethernet port, interface group (ifgrp), 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 maximum number of sessions for a storage system depends on the number of nodes in a cluster and the memory capacity of that storage system. 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.

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 hexadecimal secret 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.
 - 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.
- 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.

Related information

[Data ONTAP documentation on NetApp support site -- support.netapp.com](http://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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols > iSCSI**.
4. In the **Service** tab, click **Edit**.
5. Type an iSCSI alias in the **Target Alias** field and click **OK**.

Related references

[iSCSI window](#) on page 195

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols > iSCSI**.
4. In the **Target Portal Group** tab, click **Create**.
5. 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.
6. Select the interfaces to include in the target portal group and click **Create**.

Related references

[iSCSI window](#) on page 195

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols > iSCSI**.
4. In the **Target Portal Group** tab, select the target portal group that you want to delete and click **Delete**.
5. Select the confirmation check box and click **Delete**.

Related references

[iSCSI window](#) on page 195

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.

About this task

By default, the iSCSI service is enabled on all Ethernet interfaces after you enable the iSCSI license. Disabling an interface requires termination of any outstanding iSCSI connections and sessions currently using that interface.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Network > NIS**.
4. In the **iSCSI Interfaces** area, select the interface on which you want to enable or disable the iSCSI service.
5. Click **Enable** or **Disable**, as appropriate.

Related references

[iSCSI window](#) on page 195

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols > iSCSI**.
4. In the **Initiator Security** tab, click **Add** in the **Initiator Security** area.
5. Specify the initiator name and 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 is optional.

6. Click **OK**.

Related references

[iSCSI window](#) on page 195

Managing iSCSI

Editing default security settings

You can use the **Edit Default Security** dialog box to edit the default security settings for iSCSI initiators connected to the storage system.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols > iSCSI**.
4. In the **Default Security** box on the **Initiator Security** tab, click **Edit**.
5. 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.

6. Click **OK**.

Related references

[iSCSI window](#) on page 195

Editing initiator security

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols > iSCSI**.
4. In the **Initiator Security** tab, select one or more initiators from the initiator list and click **Edit** in the **Initiator Security** area.
5. 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.

6. Click **OK**.
7. Verify the changes you made in the **Initiator Security** tab.

Related references

[iSCSI window](#) on page 195

Changing the default initiator authentication method

You can change the default 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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols > iSCSI**.
4. In the **Initiator Security** tab, click **Edit** in the **Default Security** area.
5. 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.

6. Click **OK**.

Related references

[iSCSI window](#) on page 195

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols > iSCSI**.
4. In the **Initiator Security** tab, select the initiator whose security setting you want change.
5. Click **Set Default** in the **Initiator Security** area, and then click **Set Default** in the confirmation box.

Related references

[iSCSI window](#) on page 195

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.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols > iSCSI**.
4. In the **Target Portal Groups** tab, select the portal group that you want to edit and click **Edit**.

The **Edit** button is disabled for system-defined portal groups.

5. Select the interfaces that you want to add to the portal group and click **Save**.

Related references

[iSCSI window](#) on page 195

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols > iSCSI**.
4. Click either **Start** or **Stop**, as appropriate.
5. If you are disabling the iSCSI service, click **Stop**.

Related references

[iSCSI window](#) on page 195

Monitoring iSCSI

Viewing initiator security information

You can use the **Initiator Security** tab to view the default authentication and all the initiator specific authentication information.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols > iSCSI**.
4. In the **Initiator Security** tab, review the details listed.

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.
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 190

[Creating target portal groups](#) on page 190

[Deleting target portal groups](#) on page 191

[Enabling or disabling the iSCSI service on storage system interfaces](#) on page 191

[Adding the security method for iSCSI initiators](#) on page 192

[Editing default security settings](#) on page 192

[Editing initiator security](#) on page 193

[Changing the default initiator authentication method](#) on page 193

[Setting the default security for iSCSI initiators](#) on page 194

[Editing a target portal group](#) on page 194

[Starting or stopping the iSCSI service](#) on page 195

Configuration > 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. Each Vserver that has a FCP service is a different FC target node.

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 cables. For switch or storage system management, they might be connected to each other or to TCP/IP switches with Ethernet cable.

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.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols**, and then click the Fibre Channel protocol.
4. Click either **Start** or **Stop**, as appropriate.
5. If you are stopping the FC or FCoE service, click **Stop**.

Related references

[FC/FCoE window](#) on page 199

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.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Protocols**, and then click the Fibre Channel protocol.
4. Click **Edit**.
5. Type the new name and click **OK**.

Related references

[FC/FCoE window](#) on page 199

Window descriptions

FC/FCoE window

You can use the FC/FCoE window to start or stop the FC service.

Command buttons

Edit	Opens the Edit Node Name dialog box.
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 197

[Changing an FC or FCoE node name](#) on page 198

Configuration > Security > Users

Managing users

Adding Vserver user accounts

You can add a Vserver user account and specify a login user method to access the storage system.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Security > Users**.
4. Click **Add**.
5. Type the user name for the new user.
6. Type the password that the user uses to connect to the storage system, then confirm the password.
7. Add one or more user login methods, and then click **Add**.

Changing passwords for Vserver user accounts

You can use System Manager to reset the password for a Vserver user account.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Security > Users**.
4. Select the user account for which you want to modify the password and click **Change Password**.
5. In the **Reset Password** dialog box, type the new password, confirm the new password, and then click **Change**.

Editing Vserver user accounts

You can edit a Vserver user account by modifying the user login methods to access the storage system.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Security > Users**.
4. Select the user account that you want to modify, and then click **Edit**.
5. Modify one or more user login methods, and then click **Modify**.

Locking or unlocking Vserver user accounts

You can either lock or unlock Vserver user accounts, as required.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Security > Users**.
4. Select the user account whose account status you want to modify and click either **Lock** or **Unlock**, as appropriate.

Window descriptions

Users window

You can use the Users window to manage user accounts, reset a user's password, or display information about all user accounts.

Command buttons

Add	Opens the Add User dialog box, which enables you to add user accounts.
Edit	Opens the Modify User dialog box, which enables you to modify user login methods.
	Note: It is best to use a single role for all access and authentication methods of a user account.
Change Password	Opens the Reset Password dialog box, which enables you to reset the user password.
Lock	Locks the user account.
Refresh	Updates the information in the window.

Users list

User	Specifies the name of the user account.
Account Locked	Specifies whether the user account is locked.

Details area

The details area specifies the user login methods used by the selected user to access the storage system.

Application	Specifies the access method that a user can use to access the storage system. The supported access methods include the following: • System console (console) • HTTP(S) (http) • Data ONTAP API (ontapi) • Service Processor or RLM (service-processor) • SSH (ssh)
Authentication	Displays the default supported authentication method, which is "password."
Role	Specifies the role of a selected user.

Configuration > Security > Roles

Understanding roles

Predefined roles for Vserver administrators

A Vserver can have its own user and administration authentication domain. You can delegate the management of a Vserver to a Vserver administrator. There are four predefined roles for Vserver administrator.

The following table lists the predefined roles for Vserver administrator along with the description of the roles:

Vserver Administrator Role Name	Description of the role
vsadmin	This role is the super user role for a Vserver. A Vserver administrator with this role has the following capabilities: • Manage own user account local password and public key • Manage volumes, quotas, qtrees, Snapshot copies, FlexCaches, and files. • Manage LUNs • Configure protocols • Configure services • Monitor jobs • Monitor network connections and network interface • Monitor the health of a Vserver vsadmin role is assigned by default.
vsadmin-volume	A Vserver administrator with this role has the following capabilities: • Manage volumes, quotas, qtrees, Snapshot copies, FlexCaches, and files. • Manage LUNs • Configure protocols • Configure services • Monitor network interface • Monitor the health of a Vserver

vsadmin-protocol	A Vserver administrator with this role has the following capabilities: • Configure protocols • Configure services • Manage LUNs • Monitor network interface • Monitor the health of a Vserver
vsadmin-readonly	A Vserver administrator with this role has the following capabilities: • Monitor the health of a Vserver • Monitor network interface • View volumes and LUNs • View services and protocols

All Vserver administrators can manage their own user local password and public key information.

In addition to these predefined roles, you can create customized Vserver administrator roles with combination of available capabilities. You can assign these roles to the Vserver administrators after creating a Vserver and user accounts.

Managing roles

Adding roles

You can add an access-control role and specify the command or command directory that the role's users can access. You can also control the level of access the role has to the command or command directory and specify a query that applies to the command or command directory.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Security > Roles**.
4. Click **Add**.
5. Type the role name and add the role attributes.
6. Click **Add**.

Editing roles

You can modify an access-control role's access to a command or command directory and restrict a user's access to only a specified set of commands.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Security > Roles**.
4. Select the role that you want to modify, and then click **Edit**.
5. Modify role attributes, and then click **Modify**.

Window descriptions

Roles window

You can use the Roles window (**Configuration > Security > Roles**) to manage roles for user accounts.

Command buttons

- Add** Opens the Add Role dialog box, which enables you to create an access-control role and specify the command or command directory that the role's users can access.
- Edit** Opens the Edit Role dialog box, which enables you to add or modify role attributes.
- Refresh** Updates the information in the window.

Roles list

The roles list provides a list of roles that are available to be assigned to users.

Details area

The details area displays the role attributes, such as the command or command directory that the selected role can access, the access level, and the query.

Configuration > Services > NIS

Managing NIS domains

A Network Information Service (NIS) domain provides a directory of hostnames and IP addresses in a network. A Vserver administrator can manage NIS domains by creating, modifying, deleting, or displaying information about them. NIS cannot be configured for the cluster management server.

You can configure multiple NIS domains for a given Vserver, but only one NIS domain can be active on a Vserver at any given time. You can also configure an NIS domain with more than one Vserver.

Adding NIS domains

You can maintain host information centrally using NIS. You can use the Create NIS Domain dialog box to add the NIS domain name of your storage system.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Services > NIS**.
4. Click **Create**.
5. Type the NIS domain name and add one or more NIS servers.
6. Click **Create**.

Editing NIS domains

You can use System Manager to modify NIS domains based on the need for Vserver authentication and authorization.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Services > NIS**.
4. Select the NIS domain and click **Edit**.
5. Make the necessary changes and click **Edit**.

Window descriptions

NIS window

The NIS window enables you to view the current NIS settings for your storage system. You can also change your system's NIS settings.

Command buttons

- Create** Opens the Create NIS Domain dialog box, which enables you to create NIS domains.
- Edit** Opens the Edit NIS Domain dialog box, which enables you to add, delete, or modify NIS servers.
- Delete** Deletes the selected NIS domain.
- Refresh** Updates the information in the window.

Configuration > Services > LDAP

Managing LDAP

Associating LDAP clients with Vservers

You can associate an LDAP client with a Vserver, which enables you to use LDAP as a name service or for name mapping. However, you can modify the LDAP client configuration only at the cluster level.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Services > LDAP**.
4. Click **Edit**.
5. Modify the active LDAP client.
6. Click **OK**.

Deleting active LDAP clients

You can delete an active LDAP client when you do not want a Vserver to be associated with it.

Steps

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

2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Services > LDAP**.
4. Click **Delete**.
5. Select **OK to delete the active LDAP client** and click **Delete**.

Configuration > Services > Kerberos

Managing Kerberos

Editing Kerberos realm configurations

You can use System Manager to modify a Kerberos realm configuration associated with a Vserver.

Steps

1. From the **Home** tab, double-click the appropriate storage system.
2. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Services > Kerberos**.
4. Select the interface and click **Edit**.
5. Make the necessary changes and click **OK**.

Configuration > Services > DNS

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Services > DNS**.
4. Click **Edit**.

5. Either select or clear **Enable DNS**, as appropriate.
6. Click **OK**.

Related references

[DNS window](#) on page 208

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. Ensure that the left navigation pane displays the Vservers hierarchy.
3. In the navigation pane, select the Vserver and click **Configuration > Services > DNS**.
4. Click **Edit**.
5. In the DNS Domains and Name Servers boxes, add or modify the DNS domain names and the IP addresses.
6. Click **Ok**.

Related references

[DNS window](#) on page 208

Window descriptions

DNS window

The DNS window enables you to view the current DNS settings for your system. You can also change your system's DNS settings.

Command buttons

- Edit** Opens the Edit DNS Settings dialog box to enable or disable DNS, or add DNS domain names.
- Refresh** Updates the information in the window.

Related tasks

[Enabling or disabling DNS](#) on page 207

[Adding or editing the DNS domain name](#) on page 208

Configuration > Local Users and Groups > UNIX

Window descriptions

UNIX window

You can use the UNIX window to maintain a list of local UNIX users and groups for each Vserver. You can use local UNIX users and groups for authentication and name mappings.

Tabs

Groups You can use the **Groups** tab to add, edit, and delete UNIX groups that are local to a Vserver.

Users You can use the **Users** tab to add, edit, and delete UNIX users that are local to a Vserver.

Command buttons

Add Group Open the Add Group dialog box, which enables you to create UNIX groups that are local to a Vserver. Local UNIX groups are used with local UNIX users.

Edit Group Opens the Edit Group dialog box, which enables you to edit a group ID.

Delete Deletes the selected group.

Refresh Updates the information in the window.

Groups list

Group Name Specifies the group name.

Group ID Specifies the group ID.

Configuration > Local Users and Groups > Name Mapping

Understanding name mapping

Name mapping concepts

Data ONTAP enables concurrent access to files via the NFS and CIFS protocols and the use of the Kerberos authentication protocol. Although all of these protocols have the concept of a *principal* (a user or group), they use incompatible formats for storing principal information. CIFS stores user and group information on a Windows Domain Controller (primary or backup) by using either NetLanMan (NTLM) or Kerberos authentication. NFS stores user and group information in the

Network Information System (NIS) (or in unix-group/unix-user accounts), with few rules for authentication. NFS supports a single NIS domain per file server, whereas CIFS supports a network of cooperating domains. Kerberos maintains its own database, one per Kerberos realm. Name mappings provide a level of compatibility among these models.

Name mappings are applied in the order in which they appear in the priority list; for example, a name mapping that occurs at position 2 in the priority list is applied before a name mapping that occurs at position 3. Each mapping direction (Kerberos to UNIX, Windows to UNIX, and UNIX to Windows) has its own priority list.

NFS and CIFS use incompatible models to internally represent principal information. Data ONTAP bases its conversions on user and group names. A typical CIFS user might be ENG\jones, where ENG in the domain or workgroup name and jones is the login name of a user. An NFS user name might be jones. Note that CIFS names include a domain, but NFS names do not. Because NFS supports only a single domain per file server, the domain name is implied. A Kerberos user might be jones/myserver@EXAMPLE.COM, where jones is the login name of a user (in Kerberos terms, the *primary*), myserver is a hostname or other information that qualifies the primary (in Kerberos terms, the *instance*), and EXAMPLE.COM is a logical network (in Kerberos terms, the *realm*) served by a single Kerberos database and one or more Key Distribution Centers (KDCs). A KDC is a server that issues Kerberos tickets. Kerberos realms are listed in uppercase.

Data ONTAP can translate CIFS principals and Kerberos principals to NIS principals. It does this by translating a CIFS or Kerberos principal's internal representation into a name, matching the name against a set of rules, and, if a match is found, rewriting the name by using a substitution template and converting the rewritten name into NFS's internal principal representation. A similar sequence is used to convert NIS principals into CIFS principals when sending replies to CIFS clients. Converting NIS principals into Kerberos principals is not required. Only CIFS and Kerberos principals are subject to translation. NIS principals are the default representation.

Note: A Data ONTAP system has a maximum of 256 name mappings.

Name mapping conversion rules

A Data ONTAP system keeps a set of conversion rules for each Vserver. Each rule consists of two pieces: a *pattern* and a *replacement*. Conversions start at the beginning of the appropriate list and perform a substitution based on the first matching rule. The pattern is a UNIX-style regular expression. The replacement is a string containing escape sequences representing subexpressions from the pattern, as in the UNIX sed program.

Regular expressions are case-insensitive when mapping from Windows to UNIX. However, they are case-sensitive for Kerberos-to-UNIX and UNIX-to-Windows mappings.

As an example, the following rule converts the CIFS user named jones in the domain named ENG into the UNIX user named jones.

Pattern	Replacement
ENG\\jones	jones

Note that the backslash is a special character in regular expressions and must be escaped with another backslash.

The caret (^), underscore (_), and ampersand (&) characters can be used as prefixes for digits in replacement patterns. These characters specify uppercase, lowercase, and initial-case transformations, respectively. For instance:

- If the initial pattern is (.+) and the replacement pattern is \1, then the string jOe is mapped to jOe (no change).
- If the initial pattern is (.+) and the replacement pattern is _1, then the string jOe is mapped to joe.
- If the initial pattern is (.+) and the replacement pattern is \^1, then the string jOe is mapped to JOE.
- If the initial pattern is (.+) and the replacement pattern is \&1, then the string jOe is mapped to Joe.

If the character following a backslash-underscore (_), backslash-caret (\^), or backslash-ampersand (\&) sequence is not a digit, then the character following the backslash is used verbatim.

The following example converts any Windows user in the CIFS domain named ENG into a UNIX user with the same name in NIS.

Pattern	Replacement
ENG\\(.+)	\1

The double backslash (\\) matches a single backslash. The parentheses denote a subexpression but do not match any characters themselves. The period matches any single character. The asterisk matches zero or more of the previous expression. In this example, you are matching ENG\ followed by one or more of any character. In the replacement, \1 refers to whatever the first subexpression matched. Assuming the CIFS user ENG\jones, the replacement evaluates to jones; that is, the portion of the name following ENG\.

Note: If you are using the CLI, you must delimit all regular expressions with double quotation marks ("). For instance, to enter the regular expression (.+) in the CLI, type "(.+)\" at the command prompt. Quotation marks are not required in the Web UI.

For further information about regular expressions, see your UNIX system administration documentation, the online UNIX documentation for sed or regex, or *Mastering Regular Expressions*, published by O'Reilly and Associates.

Window descriptions

Name Mapping window

You can use the Name Mapping window to specify the name mapping entries to map users from different platforms.

Command buttons

- | | |
|----------------|--|
| Add | Opens the Add Name Mapping Entry dialog box, which enables you to create a name mapping on a Vserver. |
| Edit | Opens the Edit Name Mapping Entry dialog box, which enables you to edit a name mapping on a Vserver. |
| Delete | Opens the Delete Name Mapping Entries dialog box, which enables you to delete a name mapping entry. |
| Swap | Opens the Swap Name Mapping Entries dialog box, which enables you to swap the selected name mapping entries. |
| Refresh | Updates the information in the window. |

Storage Controllers

Configuration > Ports/Adapters

Window descriptions

Ports/Adapters window

You can use the Ports/Adapters window to manage Ethernet ports and FC/FCoE adapters.

Tabs

- Ethernet Ports** You can use the Ethernet Ports tab to view the ports across a node and create interface groups and VLAN ports.
- FCP Adapters** You can use the FCP Adapters tab to view the ports across a node and edit settings of FC/FCoE adapter.

Command buttons

- Create Ifgroup** Opens the Create Interface Group dialog box, which enables you create interface groups by choosing the ports and determining the use of ports and the network traffic distribution.
- Create VLAN** Opens the Create VLAN dialog box, which enables you to create VLAN by choosing physical interface and adding VLAN tags.
- Edit** Opens the Edit Ethernet Port dialog box, which enables you to modify the network port attributes such as MTU.
- Delete** Opens the Delete VLAN dialog box, which enables you to delete the interface.
- Refresh** Updates the information in the window.

Ports list

- Node** Displays the node on which the port interface group is located.
- Port** Displays the name of port interface group.
- Role** Displays the role of a port.
- Type** Displays the type of the interface.
- Status** Displays the status of a port.

Hosted Interfaces Displays the number of interfaces hosted on a port.

Details area

The area below the ports list displays detailed information about the selected port.

Configuration > NDMP

Window descriptions

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.

Configuration > AutoSupport

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 e-mail address from which e-mail notifications are sent and to add multiple e-mail 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.

HA Configuration

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.

Copyright information

Copyright © 1994–2011 NetApp, Inc. All rights reserved. Printed in the U.S.A.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S.A. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, FAServer, FastStak, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

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 U.S.A. and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the U.S.A. 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 U.S.A. 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 e-mail 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
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support Telephone: +1 (888) 4-NETAPP

Index

A

- aborting SnapMirror transfer 174
- absolute
 - symbolic links 179
- access
 - access-control roles 52
 - user accounts 49
- access-control roles
 - predefined roles for cluster administrators 52
- activating and deactivating
 - quotas 158
- Active Directory
 - LDAP servers, requirements 54
- adding
 - array LUNs to aggregate 47
 - disks to aggregate 42
 - DNS domain name 208
 - initiator 136
 - initiator groups 135
 - licenses 62
 - storage systems 15
- aggregate
 - defined 20
- aggregates
 - 32-bit, 64-bit 25
 - adding array LUNs 47
 - adding disks 28, 42
 - adding smaller disks to 41
 - creating 27
 - creating from spare array LUNs 46
 - creating from spare disks 41
 - deleting 27
 - editing 29
 - mixing array LUNs in 46
 - overview 23
 - RAID type 24
 - upgrading to 64-bit 30
 - viewing information 30
- Aggregates window 31, 32
- alias
 - creating iSCSI 190
- ALUA 132
- array LUNs
 - adding array LUNs to aggregate 47
 - creating aggregate from spare array LUNs 46
- Array LUNs window 47, 48

- authentication
 - iSCSI 189
 - Kerberos 57
- AutoSupport window 214, 215

B

- binding
 - igroup, port set 133
- bringing LUNs online 140

C

- changing default iSCSI initiator security 193
- CHAP
 - defined 189
 - iSCSI authentication 189
- CIFS
 - concepts 178
 - domain controllers
 - resetting 182
 - domain information
 - viewing 184
 - general properties, editing 181
 - group policy
 - enabling 183
 - reloading 183
 - updating 183
 - home directory
 - adding 181
 - deleting 182
 - home directory concepts 124
 - name mapping 209
 - setting up 180
 - supported clients 178
 - supported domain controllers 178
- CIFS share
 - creating 125
- CIFS window 184, 185
- clients
 - CIFS, supported 178
- cluster
 - description of 19
 - epsilon 19
 - quorum 19
- cluster peer relationships

- creating 75
 - deleting 75
 - editing 74
- creating
 - aggregates 27
 - LUNs 134
 - qtrees 147
 - quotas 157
 - volumes 105
- creating aggregate from spare array LUNs 46
- creating aggregate from spare disks 41
- credential caching, about 12
- credentials
 - saving 16
- customization
 - window layout 13

D

- Dashboard window 21
- data compression's interoperability
 - with deduplication 96
- Data ONTAP, with array LUNs 40
- Date/Time window 61
- deduplication
 - changing schedule 112
 - configuration 112
 - FlexVol volume
 - maximum size 104
 - maximum size with deduplication 104
 - maximum volume size 104
 - starting 113
- default initiator security
 - editing 192
- default quota 150
- default quotas 153
- degraded mode 37
- deleting
 - aggregates 27
 - initiator group 136
 - initiators 137
 - licenses 62
 - LUNs 135
 - qtrees 148
 - quotas 157
 - volumes 108
- discovering
 - storage systems 16
- disk
 - speed 35

- types for RAID 32
 - capacity by disk size 34
 - RPM 35
- Disk space hard limit 153
- Disk space soft limit 153
- disks
 - adding disks to aggregate 42
 - adding smaller to aggregate 41
 - creating aggregate from spare disks 41
 - viewing
 - disk information 42
 - viewing information 42
- Disks window 43, 44
- DNS
 - adding domain name 208
 - enabling 207
- DNS window 208
- domain controllers
 - CIFS, supported 178
 - resetting 182
 - viewing information 184
- domain information
 - viewing 184

E

- editing
 - aggregates 29
 - DNS domain name 208
 - domain name 208
 - initiator groups 139
 - initiator security 193
 - initiators 139
 - LUNs 138
 - qtrees 148
 - quotas 158
 - SnapMirror schedule 170
- Editing share
 - general settings 127
 - options 127
 - permissions 127
- efficiency policies window 167
- enabling
 - DNS 207
- Ethernet 187
- example
 - igroup mapping 133
 - port sets 133
- explicit quotas 153
- export policies

- adding rules 163
- concepts 161
- creating 162
- deleting 162
- renaming 163
- export policy rules
 - deleting 162

F

- FC/FCoE window 199
- FCoE
 - converged network adapters 197
 - data center bridging 197
 - Ethernet switch 197
 - traditional FC 197
- FCP
 - changing node name 198
 - defined 197
 - node connection 197
 - nodes defined 197
 - starting and stopping 197
- file
 - FlexClone
 - creating 107
 - hosts 82
 - SnapMirror log 82
 - snapmirror.allow 82
 - snapmirror.conf 82
- Files hard limit 153
- Files soft limit 153
- FlexClone files
 - creating 107
- FlexClone volumes
 - about 92
 - creating 106
 - shared Snapshot copies and 97
 - space guarantees and 99
 - splitting, about 93
- flexible volumes
 - described 20
- FlexVol volumes
 - automatically adding space for 101
 - try_first volume option 101
- free space
 - automatically increasing 101

G

- group policy

- disabling 183
- reloading 183
- updating 183
- group quota 150
- guidelines
 - LUN mapping 132
 - LUN type 130

H

- HA configuration window 215
- HBA 132, 187
- Help, about 9
- home directories
 - CIFS, concepts 124
- home directory
 - about 181
 - deleting 182
- home directory shares
 - creating 127
- hot spares
 - defined 36
 - appropriate 38
 - best practices 38
 - what disks can be used as 38

I

- icons, definitions 13, 14
- igroup 133
- initiator groups
 - adding 135
 - adding initiator 136
 - defined 132
 - deleting 136
 - deleting initiators 137
 - editing 139
 - editing initiators 139
 - name rules 133
 - naming 133
 - ostype of 133
 - requirements for creation 133
 - type 133
 - viewing 142
- iSCSI
 - changing default initiator security 193
 - creating alias 190
 - disabling on interface 191
 - enabling on interface 191
 - explained 187

- how communication sessions work 188

- initiator security

- adding 192

- editing 193

- setting default 194

- viewing 195

- nodes defined 187

- security 189

- target portal groups defined 187

- iSCSI service

- starting 195

- stopping 195

- iSCSI window 195, 196

J

- jobs

- about 76

K

- Kerberos

- authentication 57

- name mapping 209

- Kerberos realm configurations

- creating 57

- deleting 58

- editing 58, 207

L

- LDAP

- Active Directory servers, requirements 54

- using 54

- license

- FC 197

- licenses

- adding 62

- deleting 62

- Licenses window 63

- links

- symbolic 179

- logical transfers 168

- LUN

- host operating system type 130

- limiting access to 133

- mapping guidelines 132

- multiprotocol type 130

- LUN clones

- creating 141

- LUN creation

- host operating system type 130

- LUNs

- bringing online 140

- creating 134

- deleting 135

- editing 138

- guidelines for creating 129

- initiator hosts 132

- resizing 132

- size and type 130

- taking offline 140

- viewing information 142

- LUNs (array)

- Data ONTAP RAID groups with 40

- mixing in an aggregate 46

- RAID groups

- RAID groups

- RAID0

- RAID0 RAID group requirements

- 40

- RAID0 RAID group requirements 40

- relationship to RAID0 aggregates 40

- LUNs window 143, 144

M

- mirrors

- logical transfers 168

- MPIO 132

- Multiprotocol type 130

N

- name mappings

- concepts 209

- conversion rules 210

- name rules

- igroups 133

- namespace window 123, 124

- NDMP window 214

- Network Time Protocol (NTP)

- system time 59

- NFS

- concepts 185

- name mapping 209

- settings

- editing 186

- NIS

- managing domains 205
- NIS window 206
- node
 - description of 19
- node Vserver 19
- nodes
 - FCP 197
 - iSCSI 187

O

- online Help, about 9

P

- peers window 75, 76
- plex
 - defined 20
- port set 133
- port sets
 - deleting 138
 - editing 141
- ports/adapters window 213, 214

Q

- qtree quota 150
- qtrees
 - about 144
 - creating 147
 - defined 20
 - deleting 148
 - deletion, quotas and 155
 - editing 148
 - name restrictions 145
 - options 144
 - renaming, quotas and 155
 - security style 146
 - viewing information 149
 - when to use 145
- Qtrees window 149
- quotas
 - activating and deactivating 158
 - creating 157
 - default 150, 153
 - deleting 157
 - editing 158
 - explicit 153
 - group 150

- hard 150
- managing 153
- qtree 150
- qtree deletion, and 155
- qtree rename and 155
- reinitialization, when required 156
- resizing 159
- security style changes and 155
- soft 150
- threshold 150
- tree 154
- UNIX users and 151
- user 150
- user and group, working with qtrees 154
- viewing information 159
- why you use 150
- Windows users and 151
- Quotas window 160, 161

R

- RAID 0
 - how Data ONTAP uses 37
- RAID disk types 32
- RAID groups
 - definition 23
 - naming convention 25
 - size 25
- RAID protection
 - which device provides 37
- RAID types
 - editing 29
- RAID types Data ONTAP supports
 - for array LUNs 37
 - for disks 37
- RAID-DP 24
- RAID-level mirroring
 - described 20
- RAID4 24
- raw device mapping 132
- RDM 132
- reinitializing quotas 156
- relative
 - symbolic links 179
- removing storage systems 16
- resizing
 - quotas 159
- resizing volumes
 - options 105
- right-sizing 34
- roles 49

- root aggregate 91
- root volume
 - default name 91

S

- schedule deduplication
 - changing 112
- schedules
 - deleting 70
 - editing 69
- schedules window 70, 71
- securing styles
 - changing, quotas and 155
- security styles
 - about 102
- setting time zone 60
- Shares
 - disabling 126
- Shares window 127, 128
- Simple Network Management Protocol (SNMP) 64
- SnapMirror
 - files 82
- SnapMirror Log window 83
- SnapMirror relationship
 - breaking 172
 - creating 168
 - deleting 169
 - initializing 170
 - properties
 - editing 170
 - quiescing 171
 - resuming 172
 - resynchronizing 173
 - reverse resynchronizing 173
 - updating 171
- Snapshot copies
 - automatic, scheduling 116
 - copies list
 - viewing 118
 - creating 110
 - defined 95
 - deleting 109
 - directory, making invisible 117
 - renaming 116
 - reserve
 - setting 108
 - restoring 115
 - scheduling 116
- Snapshot policies

- about 66
- adding 67
- deleting 67
- editing 67
- Snapshot policies window 68
- SNMP
 - about 64
 - agent 64
 - enabling SNMP traps 65
 - MIBs 64
 - options when configuring 64
 - setting information 65
- SNMP (Simple Network Management Protocol) 64
- SNMP window 66
- software efficiency
 - FlexClone 91
 - FlexVol volumes 100
- space guarantees
 - about 97
- space management
 - choosing 98
- spare disks
 - defined 36
 - appropriate 38
 - what disks can be used as 38
- Splitting FlexClone volumes 114
- state, volume
 - changing 111
- storage
 - mixing array LUNs in an aggregate 46
- storage system credentials, saving 16
- storage systems
 - adding 15
 - discovering 16
 - discovery 12
 - monitoring 22
 - removing 16
 - resource management
 - about 11
 - viewing information 22
- storage units
 - types 20
- support bundle
 - creating 15
- symbolic links 179
- SyncMirror
 - plexes 20
- system logging
 - about 12
 - configuring 17

- log levels 12
- System Manager
 - about 11
- system time
 - Network Time Protocol 59
 - NTP 59
- Systems window 17, 18

T

- taking LUNs offline 140
- target portal group
 - deleting 191
- target portal groups
 - about 187
 - creating 190
 - editing 194
- terminology
 - RAID group 37
- thin provisioning
 - using FlexVol volumes 100
- Threshold soft limit 153
- tree quotas 154
- try_first volume option 101

U

- UNIX users, specifying for quotas 151
- user accounts
 - managing 49
- user quota 150
- Users window 51, 201

V

- version information, viewing 17
- viewing
 - aggregate information 30
 - initiator groups 142
 - LUN information 142
 - qtree information 149
 - quota information 159
- Viewing FlexClone hierarchy 117
- viewing initiator security 195
- viewing storage system information 22
- VMware 132
- volume move
 - conflicting operations 26
 - operations supported 26

- operations unsupported 26
- requirements 26
- Volume window 118–120
- volume, Snapshot copies
 - creating 110
- volumes
 - understanding 91
 - about 95
 - automatically adding space for 101
 - changing state 111
 - creating 105
 - deduplication
 - changing schedule 112
 - configuration 112
 - starting 113
 - deleting 108
 - editing properties 110
 - FlexClone
 - creating 106
 - flexible volumes 95
 - resizing
 - options 105
 - Snapshot copies
 - making directory invisible 117
 - renaming 116
 - restoring 115
 - scheduling 116
 - setting reserve 108
 - viewing copies list 118

Vserver

- Admin Vserver 86
- Benefits 86
- Cluster Vserver 86
- Data Vserver 86
- FC traffic 197
- Node Vserver 86
- types of 86
- Vserver administrator roles 202
- Vserver management 85
- Vserver settings
 - editing 87
- Vserver window 90
- Vservers
 - about 85
 - creating 87
 - starting 89
 - stopping 89

W

- widelinks 179

window layout
customization 13

Windows users, specifying for quotas 151