



Clustered Data ONTAP® 8.3

System Administration Guide

February 2016 | 215-10556_B0
doccomments@netapp.com

Updated for 8.3.2

 **NetApp®**

Contents

Differences between cluster and SVM administrators	10
Data ONTAP management interface basics	11
Accessing the cluster by using the CLI	11
Accessing the cluster by using the serial port	11
Accessing the cluster by using SSH	11
Enabling Telnet or RSH access to the cluster	13
Accessing the cluster by using Telnet	14
Accessing the cluster by using RSH	15
Using the Data ONTAP command-line interface	16
Understanding the different shells for CLI commands	16
Methods of navigating CLI command directories	18
Rules for specifying values in the CLI	18
Methods of viewing command history and reissuing commands	19
Keyboard shortcuts for editing CLI commands	19
Use of administrative privilege levels	21
Setting the privilege level in the CLI	21
Setting display preferences in the CLI	21
Methods of using query operators	22
Methods of using extended queries	23
Methods of customizing show command output by using fields	24
Understanding positional parameters	24
Methods of accessing Data ONTAP man pages	26
Managing CLI sessions	26
Managing records of CLI sessions	26
Managing the automatic timeout period of CLI sessions	27
Accessing the cluster by using OnCommand System Manager browser-based graphic interface	27
Understanding OnCommand System Manager	28
Ways to manage access to OnCommand System Manager	29
Cluster management basics (cluster administrators only)	30
What a cluster is	30
Considerations for single-node clusters	31
What the cluster management server is	31
Understanding quorum and epsilon	32
What a cluster replication ring is	33
Displaying information about the nodes in a cluster	33
Displaying cluster attributes	34
Modifying cluster attributes	34
Displaying the status of cluster replication rings	35
Managing nodes	36
Displaying node attributes	36

Modifying node attributes	36
Renaming a node	37
Adding nodes to the cluster	37
Removing nodes from the cluster	39
Accessing a node's log, core dump, and MIB files by using a web browser	41
Accessing the system console of a node	42
Rules governing node root volumes and root aggregates	43
Freeing up space on a node's root volume	44
Starting or stopping a node	45
Rebooting a node at the system prompt	45
Booting Data ONTAP at the boot environment prompt	46
Shutting down a node	47
Managing a node by using the boot menu	48
Managing a node remotely by using the Service Processor	49
Understanding the SP	49
Configuring the SP network	51
Methods of managing SP firmware updates	56
Accessing the SP	57
Using online help at the SP CLI	62
Commands for managing a node remotely	62
Commands for managing the SP from Data ONTAP	70
Managing SVMs	74
What SVMs are	74
How SVMs root volumes are used for data access	76
Types of SVMs	76
Why you use SVMs	76
Number of SVMs in a cluster	77
Setting up SVMs	78
Planning to create SVMs	78
Guidelines for creating SVMs	80
Creating SVMs with FlexVol volumes	84
Creating SVMs with Infinite Volume	87
Configuring SVMs	88
Creating user accounts for SVM administrators	89
Creating local user accounts for SVM administrators	89
Providing access to NIS or LDAP user accounts for SVM administrators ...	90
Providing access to Active Directory users or groups for SVM administrators	91
Delegating administration to SVM administrators	91
Modifying SVMs	94
Modifying aggregates for SVMs	94
Modifying protocols for SVMs	95
Commands for modifying SVM attributes	96
Displaying information about SVMs	97
Renaming SVMs	98

Deleting SVMs	99
Starting SVMs	100
Stopping SVMs	101
Administering SVMs from the SVM context	102
Restoring the root volume of an SVM	103
Promoting a load-sharing mirror copy	103
Promoting a data-protection mirror copy	104
Promoting a new FlexVol volume	105
Creating a new root volume on an SVM with Infinite Volume	106
Controlling and monitoring I/O performance to SVMs by using Storage QoS	107
Administering SVMs	109
SVM access and authentication methods	110
Logging in to an SVM	111
Changing the login password	112
Displaying information about SVMs	112
Displaying information about network configuration	113
Where to find additional information for SVM administration	114
Managing access to the cluster	116
Managing user accounts	116
Access methods for user accounts	117
Authentication methods for user accounts	117
Enabling AD users and groups to access the cluster and SVMs	118
Enabling NIS or LDAP users to access the cluster	120
Commands for managing user accounts	122
Managing access-control roles	123
Predefined roles for cluster administrators	123
Predefined roles for SVM administrators	124
Considerations for customizing an access-control role	126
Customizing an access-control role to restrict user access to specific commands	129
Managing rule settings for user names and passwords in an access- control role	130
Considerations for password rule settings	131
Commands for managing access-control roles	132
Managing SSH security configurations	134
Commands for managing SSH security configurations	136
Managing public keys	136
Commands for managing public keys	137
Managing digital certificates for server or client authentication	137
Installing a server certificate to authenticate the cluster or SVM as an SSL server	139
Installing a client CA or root CA certificate to authenticate an SSL client of the cluster or SVM	141
Installing a server CA certificate to authenticate an SSL server to which the cluster or SVM is a client	144

Installing a client certificate to authenticate the cluster or SVM as an SSL client	145
Installing KMIP certificates to mutually authenticate the cluster and an external key manager	147
Replacing an expired digital certificate	149
Commands for managing digital certificates	150
Managing access to web services	151
Managing the web protocol engine	152
Managing web services	154
Managing SSL	156
Configuring access to web services	157
Managing audit settings for management activities	160
Forwarding the command history log to a destination	161
Commands for managing audit settings for management activities	162
Managing the cluster time	163
Commands for managing the cluster time	164
Managing the banner and MOTD	165
Creating a banner	166
Managing the banner	168
Creating an MOTD	169
Managing the MOTD	170
Managing licenses	173
License types and licensed method	174
Commands for managing licenses	175
Managing jobs and schedules	177
Job categories	177
Commands for managing jobs	177
Commands for managing job schedules	178
Backing up and restoring cluster configurations	180
What configuration backup files are	180
Managing configuration backups	180
How the node and cluster configurations are backed up automatically	180
Commands for managing configuration backup schedules	181
Commands for managing configuration backup files	181
Recovering a node configuration	182
Finding a configuration backup file to use for recovering a node	182
Restoring the node configuration using a configuration backup file	183
Recovering a cluster configuration	184
Finding a configuration to use for recovering a cluster	184
Restoring a cluster configuration from an existing configuration	185
Synchronizing a node with the cluster	187
Managing core dumps	188
Methods of segmenting core dump files	188
Commands for managing core dumps	189
Commands for managing core segmenting	190

Monitoring the storage system	192
Managing event messages	192
Setting up the Event Management System	193
Commands for managing events	194
Managing AutoSupport	195
When and where AutoSupport messages are sent	196
How AutoSupport creates and sends event-triggered messages	198
Types of AutoSupport messages and their content	199
What AutoSupport subsystems are	200
AutoSupport size and time budgets	200
Files sent in event-triggered AutoSupport messages	201
Log files sent in AutoSupport messages	202
Files sent in weekly AutoSupport messages	202
How AutoSupport obtains delivery instructions from technical support	203
Structure of AutoSupport messages sent by email	204
AutoSupport severity types	204
Requirements for using AutoSupport	205
Setting up AutoSupport	206
Uploading core dump files	208
Uploading performance archive files	209
Getting AutoSupport message descriptions	210
Commands for managing AutoSupport	211
Information included in the AutoSupport manifest	212
AutoSupport case suppression during scheduled maintenance windows ...	212
What My AutoSupport is	213
Troubleshooting AutoSupport	213
Monitoring the health of your system	218
How health monitoring works	218
Available cluster health monitors	220
Receiving system health alerts automatically	221
Responding to degraded system health	221
Configuring discovery of cluster and management network switches	223
Verifying the monitoring of cluster and management network switches	224
Commands for monitoring the health of your system	225
Using dashboards to display critical system information	228
Getting notified of dashboard alarms	228
Commands for managing dashboards	229
Monitoring cluster performance	230
What objects, instances, and counters are	231
How performance presets define what data is collected	232
Collecting and viewing performance data for objects	232
Filtering the statistics	233
Sorting the statistics	234
Importing a performance preset configuration (cluster administrators only)	235

Monitoring cluster performance using advanced methods	235
Commands for monitoring cluster performance	238
Displaying environmental information	241
Managing system performance	242
Managing workload performance by using Storage QoS	242
How Storage QoS works	243
Controlling and monitoring workload performance	247
Example: Isolating a workload	249
Example: Proactively setting a limit on non-critical workloads	250
Example: Proactively setting a limit on workloads in a shared storage infrastructure	251
Commands for controlling and monitoring workloads	251
Histogram-based predictions in RAVE	254
Increasing WAFL cache memory	254
Comparison of Flash Pool aggregates and Flash Cache	254
Enabling and disabling WAFL external cache	255
Caching normal user data blocks	256
Caching low-priority user data blocks	257
Caching only system metadata	257
Enabling and disabling the caching of blocks from random write workloads	258
Enabling and disabling the caching of readahead data	258
Displaying the WAFL external cache configuration	258
Displaying usage and access information for WAFL external cache	259
Preserving the cache in the Flash Cache family of modules	259
Improving read performance	262
What read reallocation is	262
Commands for managing read reallocation	262
Improving write performance	263
How free space reallocation optimizes free space	263
When to enable free space reallocation	263
When to use free space reallocation with other reallocation features	264
Types of aggregates that free space reallocation can and cannot optimize .	264
Commands for managing free space reallocation	264
Managing peer relationships for data backup and recovery	265
Managing cluster peer relationships	265
What a cluster peer is	265
Commands for managing cluster peer relationships	265
Cluster peer security using authentication passphrases	267
Connecting one cluster to another cluster in a peer relationship	267
Modifying a cluster peer relationship	286
Deleting a cluster peering relationship	287
Managing SVM peer relationships	287
What an SVM peer relationship is	288
States of SVM peer relationships	289

Creating an SVM peer relationship	289
Accepting an SVM peer relationship	292
Rejecting an SVM peer relationship	293
Modifying the peering application on an SVM peer relationship	294
Deleting an SVM peer relationship	295
Suspending an SVM peer relationship	296
Resuming an SVM peer relationship	296
Displaying information about SVM peer relationships	297
Glossary	299
Copyright information	306
Trademark information	307
How to send comments about documentation and receive update notifications	308
Index	309

Differences between cluster and SVM administrators

Cluster administrators administer the entire cluster and the Storage Virtual Machines (SVMs, formerly known as Vservers) it contains. SVM administrators administer only their own data SVMs.

Cluster administrators can administer the entire cluster and its resources. They can also set up data SVMs and delegate SVM administration to SVM administrators. The specific capabilities that cluster administrators have depend on their access-control roles. By default, a cluster administrator with the “admin” account name or role name has all capabilities for managing the cluster and SVMs.

SVM administrators can administer only their own SVM storage and network resources, such as volumes, protocols, LIFs, and services. The specific capabilities that SVM administrators have depend on the access-control roles that are assigned by cluster administrators.

Note: The Data ONTAP command-line interface (CLI) continues to use the term *Vserver* in the output, and `vserver` as a command or parameter name has not changed.

Related concepts

[Managing SVMs \(cluster administrators only\)](#) on page 74

[Predefined roles for cluster administrators](#) on page 123

[Predefined roles for SVM administrators](#) on page 124

Data ONTAP management interface basics

You can administer the cluster by using the Data ONTAP command-line interface (CLI) or the web interface. The CLI provides a command-based mechanism that is similar to the UNIX tcsh shell. The web interface enables you to use a web browser to manage the cluster.

Related concepts

What a cluster is on page 30

Understanding the different shells for CLI commands (cluster administrators only) on page 16

Accessing the cluster by using the CLI (cluster administrators only)

You can access the cluster by using the serial console, SSH, Telnet, or RSH. These protocols enable you to access the cluster to run CLI commands.

Accessing the cluster by using the serial port

You can access the cluster directly from a console that is attached to a node's serial port.

Steps

1. At the console, press Enter.

The system responds with the login prompt.

2. At the login prompt, do one of the following:

To access the cluster with...	Enter the following account name...
The default cluster account	admin
An alternative administrative user account	<i>username</i>

The system responds with the password prompt.

3. Enter the password for the admin or administrative user account, and then press Enter.

Accessing the cluster by using SSH

You can issue SSH requests to the cluster to perform administrative tasks. SSH is enabled by default.

Before you begin

- You must have a user account that is configured to use **ssh** as an access method. The `-application` parameter of the `security login` commands specifies the access method for a user account. For more information, see the `security login` man pages.
- If you use an Active Directory (AD) domain user account to access the cluster, an authentication tunnel for the cluster must have been set up through a CIFS-enabled Storage Virtual Machine (SVM), and your AD domain user account must also have been added to the cluster with **ssh** as an access method and **domain** as the authentication method.
- If you use IPv6 connections, IPv6 must already be configured and enabled on the cluster, and firewall policies must already be configured with IPv6 addresses.

The network options `ipv6 show` command displays whether IPv6 is enabled. The `system services firewall policy show` command displays firewall policies.

About this task

- Data ONTAP supports OpenSSH client version 5.4p1 and OpenSSH server version 5.4p1. Only the SSH v2 protocol is supported; SSH v1 is not supported.
- Data ONTAP supports a maximum of 64 concurrent SSH sessions per node. If the cluster management LIF resides on the node, it shares this limit with the node management LIF. If the rate of incoming connections is higher than 10 per second, the service is temporarily disabled for 60 seconds.
- Data ONTAP supports only the AES and 3DES encryption algorithms (also known as *ciphers*) for SSH. AES is supported with 128, 192, and 256 bits in key length. 3DES is 56 bits in key length as in the original DES, but it is repeated three times.
- If you want to access the Data ONTAP CLI from a Windows host, you can use a third-party utility such as PuTTY.
- If you use a Windows Active Directory (AD) user name to log in to Data ONTAP, you should use the same uppercase or lowercase letters that were used when the AD user name and domain name were created in Data ONTAP. AD user names and domain names are not case sensitive. However, Data ONTAP user names are case sensitive. Case mismatch between the user name created in Data ONTAP and the user name created in AD results in a login failure.

Step

1. From an administration host, enter the `ssh` command in one of the following formats:

- `ssh username@hostname_or_IP [command]`
- `ssh -l username hostname_or_IP [command]`

If you are using an AD domain user account, you must specify *username* in the format of `domainname\AD_accountname` (with double backslashes after the domain name) or `"domainname\AD_accountname"` (enclosed in double quotation marks and with a single backslash after the domain name).

hostname_or_IP is the host name or the IP address of the cluster management LIF or a node management LIF. Using the cluster management LIF is recommended. You can use an IPv4 or IPv6 address.

command is not required for SSH-interactive sessions.

Examples of SSH requests

The following examples show how the user account named “joe” can issue an SSH request to access a cluster whose cluster management LIF is 10.72.137.28:

```
$ ssh joe@10.72.137.28
Password:
cluster1:~> cluster show
Node           Health  Eligibility
-----
node1          true    true
node2          true    true
```

```
2 entries were displayed.
```

```
cluster1::>
```

```
$ ssh -l joe 10.72.137.28 cluster show
Password:
Node                Health  Eligibility
-----
node1               true   true
node2               true   true
2 entries were displayed.

$
```

The following examples show how the user account named “john” from the domain named “DOMAIN1” can issue an SSH request to access a cluster whose cluster management LIF is 10.72.137.28:

```
$ ssh DOMAIN1\john@10.72.137.28
Password:
cluster1::> cluster show
Node                Health  Eligibility
-----
node1               true   true
node2               true   true
2 entries were displayed.

cluster1::>
```

```
$ ssh -l "DOMAIN1\john" 10.72.137.28 cluster show
Password:
Node                Health  Eligibility
-----
node1               true   true
node2               true   true
2 entries were displayed.

$
```

Enabling Telnet or RSH access to the cluster

Telnet and RSH are disabled in the predefined management firewall policy (**mgmt**). To enable the cluster to accept Telnet or RSH requests, you must create a new management firewall policy that has Telnet or RSH enabled and then associate the new policy with the cluster management LIF.

About this task

Data ONTAP prevents you from changing predefined firewall policies, but you can create a new policy by cloning the predefined **mgmt** management firewall policy and then enabling Telnet or RSH under the new policy. However, Telnet and RSH are not secure protocols, so you should consider using SSH to access the cluster. SSH provides a secure remote shell and interactive network session.

Steps

1. Use the `system services firewall policy clone` command to create a new management firewall policy based on the **mgmt** management firewall policy.
2. Use the `system services firewall policy create` command to enable Telnet or RSH in the new management firewall policy.

3. Use the `network interface modify` command to associate the new policy with the cluster management LIF.

Related information

Clustered Data ONTAP 8.3.2 man page: `system services firewall policy clone` - Clone an existing firewall policy

Clustered Data ONTAP 8.3.2 man page: `system services firewall policy create` - Create a firewall policy entry for a network service

Clustered Data ONTAP 8.3.2 man page: `network interface modify` - Modify a logical interface

Accessing the cluster by using Telnet

You can issue Telnet requests to the cluster to perform administrative tasks. Telnet is disabled by default.

Before you begin

The following conditions must be met before you can use Telnet to access the cluster:

- You must have a cluster local user account that is configured to use Telnet as an access method. The `-application` parameter of the `security login` commands specifies the access method for a user account. For more information, see the `security login` man pages.
- Telnet must already be enabled in the management firewall policy that is used by the cluster or node management LIFs so that Telnet requests can go through the firewall. By default, Telnet is disabled. The `system services firewall policy show` command with the `-service telnet` parameter displays whether Telnet has been enabled in a firewall policy. For more information, see the `system services firewall policy` man pages.
- If you use IPv6 connections, IPv6 must already be configured and enabled on the cluster, and firewall policies must already be configured with IPv6 addresses. The `network options ipv6 show` command displays whether IPv6 is enabled. The `system services firewall policy show` command displays firewall policies.

About this task

- Telnet is not a secure protocol. You should consider using SSH to access the cluster. SSH provides a secure remote shell and interactive network session.
- Data ONTAP supports a maximum of 50 concurrent Telnet sessions per node. If the cluster management LIF resides on the node, it shares this limit with the node management LIF. If the rate of in-coming connections is higher than 10 per second, the service is temporarily disabled for 60 seconds.
- If you want to access the Data ONTAP CLI from a Windows host, you can use a third-party utility such as PuTTY.

Step

1. From an administration host, enter the following command:

```
telnet hostname_or_IP
```

hostname_or_IP is the host name or the IP address of the cluster management LIF or a node management LIF. Using the cluster management LIF is recommended. You can use an IPv4 or IPv6 address.

Example of a Telnet request

The following example shows how the user named “joe”, who has been set up with Telnet access, can issue a Telnet request to access a cluster whose cluster management LIF is 10.72.137.28:

```
admin_host$ telnet 10.72.137.28
Data ONTAP
login: joe
Password:
cluster1::>
```

Related concepts

[Access methods for user accounts](#) on page 117

Related information

[Clustered Data ONTAP 8.3 Network Management Guide](#)

Accessing the cluster by using RSH

You can issue RSH requests to the cluster to perform administrative tasks. RSH is not a secure protocol and is disabled by default.

Before you begin

The following conditions must be met before you can use RSH to access the cluster:

- You must have a cluster local user account that is configured to use RSH as an access method. The `-application` parameter of the `security login` commands specifies the access method for a user account. For more information, see the `security login` man pages.
- RSH must already be enabled in the management firewall policy that is used by the cluster or node management LIFs so that RSH requests can go through the firewall. By default, RSH is disabled. The `system services firewall policy show` command with the `-service rsh` parameter displays whether RSH has been enabled in a firewall policy. For more information, see the `system services firewall policy` man pages.
- If you use IPv6 connections, IPv6 must already be configured and enabled on the cluster, and firewall policies must already be configured with IPv6 addresses. The `network options ipv6 show` command displays whether IPv6 is enabled. The `system services firewall policy show` command displays firewall policies.

About this task

- RSH is not a secure protocol. You should consider using SSH to access the cluster. SSH provides a secure remote shell and interactive network session.
- Data ONTAP supports a maximum of 50 concurrent RSH sessions per node. If the cluster management LIF resides on the node, it shares this limit with the node management LIF. If the rate of in-coming connections is higher than 10 per second, the service is temporarily disabled for 60 seconds.

Step

1. From an administration host, enter the following command:

```
rsh hostname_or_IP -l username:password command
```

hostname_or_IP is the host name or the IP address of the cluster management LIF or a node management LIF. Using the cluster management LIF is recommended. You can use an IPv4 or IPv6 address.

command is the command you want to execute over RSH.

Example of an RSH request

The following example shows how the user named “joe”, who has been set up with RSH access, can issue an RSH request to run the `cluster show` command:

```
admin_host$ rsh 10.72.137.28 -l joe:password cluster show

Node                Health  Eligibility
-----
node1                true   true
node2                true   true
2 entries were displayed.

admin_host$
```

Using the Data ONTAP command-line interface

The Data ONTAP command-line interface (CLI) provides a command-based view of the management interface. You enter commands at the storage system prompt, and command results are displayed in text.

The CLI command prompt is represented as `cluster_name::>`.

If you set the privilege level (that is, the `-privilege` parameter of the `set` command) to **advanced**, the prompt includes an asterisk (*), for example, `cluster_name::*>`.

Understanding the different shells for CLI commands (cluster administrators only)

The cluster has three different shells for CLI commands, the *clustershell*, the *nodeshell*, and the *systemshell*. The shells are for different purposes, and they each have a different command set.

- The *clustershell* is the native shell that is started automatically when you log in to the cluster. It provides all the commands you need to configure and manage the cluster. The *clustershell* CLI help (triggered by `?` at the *clustershell* prompt) displays available *clustershell* commands. The `man command_name` command in the *clustershell* displays the man page for the specified *clustershell* command.
- The *nodeshell* is a special shell for commands that take effect only at the node level. The *nodeshell* is accessible through the `system node run` command. The *nodeshell* CLI help (triggered by `?` or `help` at the *nodeshell* prompt) displays available *nodeshell* commands. The `man command_name` command in the *nodeshell* displays the man page for the specified *nodeshell* command. Many commonly used *nodeshell* commands and options are tunneled or aliased into the *clustershell* and can be executed also from the *clustershell*.

- The systemshell is a low-level shell that is used only for diagnostic and troubleshooting purposes. The systemshell and the associated “diag” account are intended for low-level diagnostic purposes. Their access requires the diagnostic privilege level and is reserved only for technical support to perform troubleshooting tasks.

Access of nodeshell commands and options in the clustershell

Nodeshell commands and options are accessible through the nodeshell (`system node run -node nodename`). Many commonly used nodeshell commands and options are tunneled or aliased into the clustershell and can be executed also from the clustershell.

Nodeshell options that are supported in the clustershell can be accessed by using the `vserver options clustershell` command. To see these options, you can do one of the following:

- Query the clustershell CLI with `vserver options -vserver nodename_or_clustername -option-name ?`
- Access the `vserver options` man page in the clustershell CLI with `man vserver options`

If you enter a nodeshell or legacy command or option in the clustershell, and the command or option has an equivalent clustershell command, Data ONTAP informs you of the clustershell command to use.

If you enter a nodeshell or legacy command or option that is not supported in the clustershell, Data ONTAP informs you of the “not supported” status for the command or option.

Related tasks

[Displaying available nodeshell commands](#) on page 17

Displaying available nodeshell commands

You can obtain a list of available nodeshell commands by using the CLI help from the nodeshell.

Steps

1. To access the nodeshell, enter the following command at the clustershell's system prompt:

```
system node run -node {nodename|local}
```

`local` is the node you used to access the cluster.

Note: The `system node run` command has an alias command, `run`.

2. Enter the following command in the nodeshell to see the list of available nodeshell commands:

```
[commandname] help
```

`commandname` is the name of the command whose availability you want to display. If you do not include `commandname`, the CLI displays all available nodeshell commands.

You enter `exit` or type `Ctrl-d` to return to the clustershell CLI.

Example of displaying available nodeshell commands

The following example accesses the nodeshell of a node named `node2` and displays information for the nodeshell command `environment`:

```
cluster1::> system node run -node node2
Type 'exit' or 'Ctrl-D' to return to the CLI

node2> environment help
Usage: environment status |
      [status] [shelf [<adapter>[.<shelf-number>]]] |
      [status] [shelf_log] |
```

```
[status] [shelf_stats] |
[status] [shelf_power_status] |
[status] [chassis [all | list-sensors | Temperature | PSU 1 |
PSU 2 | Voltage | SYS FAN | NVRAM6-temperature-3 | NVRAM6-
battery-3]]
```

Related concepts

[Access of nodeshell commands and options in the clustershell](#) on page 17

Methods of navigating CLI command directories

Commands in the CLI are organized into a hierarchy by command directories. You can run commands in the hierarchy either by entering the full command path or by navigating through the directory structure.

When using the CLI, you can access a command directory by typing the directory's name at the prompt and then pressing Enter. The directory name is then included in the prompt text to indicate that you are interacting with the appropriate command directory. To move deeper into the command hierarchy, you type the name of a command subdirectory followed by pressing Enter. The subdirectory name is then included in the prompt text and the context shifts to that subdirectory.

You can navigate through several command directories by entering the entire command. For example, you can display information about disk drives by entering the `storage disk show` command at the prompt. You can also run the command by navigating through one command directory at a time, as shown in the following example:

```
cluster1::> storage
cluster1::storage> disk
cluster1::storage disk> show
```

You can abbreviate commands by entering only the minimum number of letters in a command that makes the command unique to the current directory. For example, to abbreviate the command in the previous example, you can enter `st d sh`. You can also use the Tab key to expand abbreviated commands and to display a command's parameters, including default parameter values.

You can use the `top` command to go to the top level of the command hierarchy, and the `up` command or `..` command to go up one level in the command hierarchy.

Note: Commands and command options preceded by an asterisk (*) in the CLI can be executed only at the advanced privilege level or higher.

Rules for specifying values in the CLI

Most commands include one or more required or optional parameters. Many parameters require you to specify a value for them. A few rules exist for specifying values in the CLI.

- A value can be a number, a Boolean specifier, a selection from an enumerated list of predefined values, or a text string.
Some parameters can accept a comma-separated list of two or more values. Comma-separated lists of values do not need to be in quotation marks (" "). Whenever you specify text, a space, or a query character (when not meant as a query or text starting with a less-than or greater-than symbol), you must enclose the entity in quotation marks.
- The CLI interprets a question mark ("?") as the command to display help information for a particular command.
- Some text that you enter in the CLI, such as command names, parameters, and certain values, is not case-sensitive.

For example, when you enter parameter values for the `vserver cifs` commands, capitalization is ignored. However, most parameter values, such as the names of nodes, Storage Virtual Machines (SVMs), aggregates, volumes, and logical interfaces, are case-sensitive.

- If you want to clear the value of a parameter that takes a string or a list, you specify an empty set of quotation marks ("") or a dash ("-").
- The hash sign (“#”), also known as the pound sign, indicates a comment for a command-line input; if used, it should appear after the last parameter in a command line.
The CLI ignores the text between “#” and the end of the line.

In the following example, an SVM is created with a text comment. The SVM is then modified to delete the comment:

```
cluster1::> vserver create -vserver vs0 -subtype default -rootvolume root_vs0
-aggregate aggr1 -rootvolume-security-style unix -language C.UTF-8 -is-repository
false -ipSPACE ipSPACEA -comment "My SVM"
cluster1::> vserver modify -vserver vs0 -comment ""
```

In the following example, a command-line comment that uses the “#” sign indicates what the command does.

```
cluster1::> security login create -vserver vs0 -user-or-group-name new-admin
-application ssh -authmethod password #This command creates a new user account
```

Methods of viewing command history and reissuing commands

Each CLI session keeps a history of all commands issued in it. You can view the command history of the session that you are currently in. You can also reissue commands.

To view the command history, you can use the `history` command.

To reissue a command, you can use the `redo` command with one of the following arguments:

- A string that matches part of a previous command
For example, if the only `volume` command you have run is `volume show`, you can use the `redo volume` command to reexecute the command.
- The numeric ID of a previous command, as listed by the `history` command
For example, you can use the `redo 4` command to reissue the fourth command in the history list.
- A negative offset from the end of the history list
For example, you can use the `redo -2` command to reissue the command that you ran two commands ago.

For example, to redo the command that is third from the end of the command history, you would enter the following command:

```
cluster1::> redo -3
```

Keyboard shortcuts for editing CLI commands

The command at the current command prompt is the active command. Using keyboard shortcuts enables you to edit the active command quickly. These keyboard shortcuts are similar to those of the UNIX `tcsh` shell and the Emacs editor.

The following table lists the keyboard shortcuts for editing CLI commands. “Ctrl-” indicates that you press and hold the Ctrl key while typing the character specified after it. “Esc-” indicates that you press and release the Esc key and then type the character specified after it.

If you want to...	Use the following keyboard shortcut...
Move the cursor back by one character	Ctrl-B
	Back arrow
Move the cursor forward by one character	Ctrl-F
	Forward arrow
Move the cursor back by one word	Esc-B
Move the cursor forward by one word	Esc-F
Move the cursor to the beginning of the line	Ctrl-A
Move the cursor to the end of the line	Ctrl-E
Remove the content of the command line from the beginning of the line to the cursor, and save it in the cut buffer The cut buffer acts like temporary memory, similar to what is called a <i>clipboard</i> in some programs.	Ctrl-U
Remove the content of the command line from the cursor to the end of the line, and save it in the cut buffer	Ctrl-K
Remove the content of the command line from the cursor to the end of the following word, and save it in the cut buffer	Esc-D
Remove the word before the cursor, and save it in the cut buffer	Ctrl-W
Yank the content of the cut buffer, and push it into the command line at the cursor	Ctrl-Y
Delete the character before the cursor	Ctrl-H
	Backspace
Delete the character where the cursor is	Ctrl-D
Clear the line	Ctrl-C
Clear the screen	Ctrl-L
Replace the current content of the command line with the previous entry on the history list With each repetition of the keyboard shortcut, the history cursor moves to the previous entry.	Ctrl-P
	Esc-P
	Up arrow
Replace the current content of the command line with the next entry on the history list With each repetition of the keyboard shortcut, the history cursor moves to the next entry.	Ctrl-N
	Esc-N
	Down arrow
Expand a partially entered command or list valid input from the current editing position	Tab
	Ctrl-I
Display context-sensitive help	?
Escape the special mapping for the question mark (“?”) character For instance, to enter a question mark into a command's argument, press Esc and then the “?” character.	Esc-?
Start TTY output	Ctrl-Q

If you want to...	Use the following keyboard shortcut...
Stop TTY output	Ctrl-S

Use of administrative privilege levels

Data ONTAP commands and parameters are defined at three privilege levels: *admin*, *advanced*, and *diagnostic*. The privilege levels reflect the skill levels required in performing the tasks.

admin

Most commands and parameters are available at this level. They are used for common or routine tasks.

advanced

Commands and parameters at this level are used infrequently, require advanced knowledge, and can cause problems if used inappropriately.

You use advanced commands or parameters only with the advice of support personnel.

diagnostic

Diagnostic commands and parameters are potentially disruptive. They are used only by support personnel to diagnose and fix problems.

Setting the privilege level in the CLI

You can set the privilege level in the CLI by using the `set` command. Changes to privilege level settings apply only to the session you are in. They are not persistent across sessions.

Step

1. To set the privilege level in the CLI, use the `set` command with the `-privilege` parameter.

Example of setting the privilege level

The following example sets the privilege level to advanced and then to admin:

```
cluster1::> set -privilege advanced
Warning: These advanced commands are potentially dangerous; use
them only when directed to do so by technical support.
Do you wish to continue? (y or n): y
cluster1::*> set -privilege admin
```

Related references

[Use of administrative privilege levels](#) on page 21

Setting display preferences in the CLI

You can set display preferences for a CLI session by using the `set` command and `rows` command. The preferences you set apply only to the session you are in. They are not persistent across sessions.

About this task

You can set the following CLI display preferences:

- The privilege level of the command session
- Whether confirmations are issued for potentially disruptive commands

- Whether `show` commands display all fields
- The character or characters to use as the field separator
- The default unit when reporting data sizes
- The number of rows the screen displays in the current CLI session before the interface pauses output
If the preferred number of rows is not specified, it is automatically adjusted based on the actual height of the terminal. If the actual height is undefined, the default number of rows is 24.
- The default Storage Virtual Machine (SVM) or node
- Whether a continuing command should stop if it encounters an error

Step

1. To set CLI display preferences, use the `set` command.

To set the number of rows the screen displays in the current CLI session, you can also use the `rows` command.

For more information, see the man pages for the `set` command and `rows` command.

Example of setting display preferences in the CLI

The following example sets a comma to be the field separator, sets **GB** as the default data-size unit, and sets the number of rows to 50:

```
cluster1::> set -showseparator "," -units GB
cluster1::> rows 50
```

Methods of using query operators

The management interface supports queries and UNIX-style patterns and wildcards to enable you to match multiple values in command-parameter arguments.

The following table describes the supported query operators:

Operator	Description
*	Wildcard that matches all entries. For example, the command <code>volume show -volume *tmp*</code> displays a list of all volumes whose names include the string <code>tmp</code> .
!	NOT operator. Indicates a value that is not to be matched; for example, <code>!vs0</code> indicates not to match the value <code>vs0</code> .
	OR operator. Separates two values that are to be compared; for example, <code>vs0 vs2</code> matches either <code>vs0</code> or <code>vs2</code> . You can specify multiple OR statements; for example, <code>a b* *c*</code> matches the entry <code>a</code> , any entry that starts with <code>b</code> , and any entry that includes <code>c</code> .
..	Range operator. For example, <code>5..10</code> matches any value from 5 to 10, inclusive.

Operator	Description
<	Less-than operator. For example, <20 matches any value that is less than 20.
>	Greater-than operator. For example, >5 matches any value that is greater than 5.
<=	Less-than-or-equal-to operator. For example, <=5 matches any value that is less than or equal to 5.
>=	Greater-than-or-equal-to operator. For example, >=5 matches any value that is greater than or equal to 5.
{ <i>query</i> }	Extended query. An extended query must be specified as the first argument after the command name, before any other parameters. For example, the command <code>volume modify {-volume *tmp*} -state offline</code> sets offline all volumes whose names include the string <code>tmp</code> .

If you want to parse query characters as literals, you must enclose the characters in double quotes (for example, “^”, “:”, “*”, or “\$”) for the correct results to be returned.

You can use multiple query operators in one command line. For example, the command `volume show -size >1GB -percent-used <50 -vserver !vs1` displays all volumes that are greater than 1 GB in size, less than 50% utilized, and not in the Storage Virtual Machine (SVM) named “vs1”.

Methods of using extended queries

You can use extended queries to match and perform operations on objects that have specified values.

You specify extended queries by enclosing them within curly brackets ({}). An extended query must be specified as the first argument after the command name, before any other parameters. For example, to set offline all volumes whose names include the string `tmp`, you run the command in the following example:

```
cluster1::> volume modify {-volume *tmp*} -state offline
```

Extended queries are generally useful only with `modify` and `delete` commands. They have no meaning in `create` or `show` commands.

The combination of queries and modify operations is a useful tool. However, it can potentially cause confusion and errors if implemented incorrectly. For example, using the `system node image modify` command to set a node's default software image automatically sets the other software image not to be the default. The command in the following example is effectively a null operation:

```
cluster1::> system node image modify {-isdefault true} -isdefault false
```

This command sets the current default image as the non-default image, then sets the new default image (the previous non-default image) to the non-default image, resulting in the original default settings being retained. To perform the operation correctly, you can use the command in the following example:

```
cluster1::> system node image modify {-iscurrent false} -isdefault true
```

Methods of customizing show command output by using fields

When you use the `-instance` parameter with a `show` command to display details, the output can be lengthy and include more information than you need. The `-fields` parameter of a `show` command enables you to display only the information you specify.

For example, running `volume show -instance` is likely to result in several screens of information. You can use `volume show -fields fieldname[,fieldname...]` to customize the output so that it includes only the specified field or fields (in addition to the default fields that are always displayed.) You can use `-fields ?` to display valid fields for a `show` command.

The following example shows the output difference between the `-instance` parameter and the `-fields` parameter:

```
cluster1::> volume show -instance
                                Vserver Name: cluster1-1
                                Volume Name: vol0
                                Aggregate Name: aggr0
                                Volume Size: 348.3GB
                                Volume Data Set ID: -
Volume Master Data Set ID: -
                                Volume State: online
                                Volume Type: RW
                                Volume Style: flex
                                ...
                                Space Guarantee Style: volume
                                Space Guarantee in Effect: true
                                ...
Press <space> to page down, <return> for next line, or 'q' to quit...
...
cluster1::>

cluster1::> volume show -fields space-guarantee,space-guarantee-enabled

vserver  volume space-guarantee space-guarantee-enabled
-----
cluster1-1 vol0   volume           true
cluster1-2 vol0   volume           true
vs1      root_vol        volume           true
vs2      new_vol         volume           true
vs2      root_vol         volume           true
...
cluster1::>
```

Understanding positional parameters

You can take advantage of the positional parameter functionality of the Data ONTAP CLI to increase efficiency in command input. You can query a command to identify parameters that are positional for the command.

What a positional parameter is

- A positional parameter is a parameter that does not require you to specify the parameter name before specifying the parameter value.
- A positional parameter can be interspersed with nonpositional parameters in the command input, as long as it observes its relative sequence with other positional parameters in the same command, as indicated in the `command_name ?` output.
- A positional parameter can be a required or optional parameter for a command.
- A parameter can be positional for one command but nonpositional for another.

Attention: Using the positional parameter functionality in scripts is not recommended, especially when the positional parameters are optional for the command or have optional parameters listed before them.

How to identify a positional parameter

You can identify a positional parameter in the `command_name ?` command output. A positional parameter has square brackets surrounding its parameter name, in one of the following formats:

- `[-parameter_name] parameter_value` shows a required parameter that is positional.
- `[[[-parameter_name] parameter_value]` shows an optional parameter that is positional.

For example, when displayed as the following in the `command_name ?` output, the parameter is positional for the command it appears in:

- `[-lif] <lif-name>`
- `[[[-lif] <lif-name>]`

However, when displayed as the following, the parameter is nonpositional for the command it appears in:

- `-lif <lif-name>`
- `[-lif <lif-name>]`

Examples of using positional parameters

In the following example, the `volume create ?` output shows that three parameters are positional for the command: `-volume`, `-aggregate`, and `-size`.

```
cluster1::> volume create ?
-vserver <vserver name>           Vserver Name
[-volume] <volume name>           Volume Name
[-aggregate] <aggregate name>     Aggregate Name
[[-size] {<integer>[KB|MB|GB|TB|PB]]} Volume Size
[ -state {online|restricted|offline|force-online|force-offline|mixed} ]
                                   Volume State (default: online)
[ -type {RW|DP|DC} ]               Volume Type (default: RW)
[ -policy <text> ]                 Export Policy
[ -user <user name> ]              User ID
...
[ -space-guarantee|-s {none|volume} ] Space Guarantee Style (default: volume)
[ -percent-snapshot-space <percent> ] Space Reserved for Snapshot Copies
...
```

In the following example, the `volume create` command is specified without taking advantage of the positional parameter functionality:

```
cluster1::> volume create -vserver svml -volume vol1 -aggregate aggr1 -size 1g
-percent-snapshot-space 0
```

The following examples use the positional parameter functionality to increase the efficiency of the command input. The positional parameters are interspersed with nonpositional parameters in the `volume create` command, and the positional parameter values are specified without the parameter names. The positional parameters are specified in the same sequence indicated by the `volume create ?` output. That is, the value for `-volume` is specified before that of `-aggregate`, which is in turn specified before that of `-size`.

```
cluster1::> volume create vol2 aggr1 1g -vserver svml -percent-snapshot-space 0
cluster1::> volume create -vserver svml vol3 -snapshot-policy default aggr1 -nvfail off 1g
-space-guarantee none
```

Methods of accessing Data ONTAP man pages

Data ONTAP manual (`man`) pages explain how to use Data ONTAP commands. They are available at the command line and on the NetApp Support Site.

The `man command_name` command displays the man page of the specified command. If you do not specify a command name, the man page index is displayed. You can use the `man man` command to view information about the `man` command itself. You can exit a man page by entering `q`.

The *Clustered Data ONTAP Commands: Manual Page Reference* is a compilation of man pages for the admin-level and advanced-level Data ONTAP commands. It is available on the NetApp Support Site.

Related information

[NetApp Support Site: mysupport.netapp.com](http://mysupport.netapp.com)

Managing CLI sessions (cluster administrators only)

You can create a log for a CLI session and upload it to a specified destination to keep as a record. In addition, you can specify the automatic timeout period of a CLI session to have the session automatically disconnected after the number of minutes specified by the timeout value has elapsed.

Managing records of CLI sessions

You can record a CLI session into a file with a specified name and size limit, then upload the file to an FTP or HTTP destination. You can also display or delete files in which you previously recorded CLI sessions.

A record of a CLI session ends when you stop the recording or end the CLI session, or when the file reaches the specified size limit. The default file size limit is 1 MB. The maximum file size limit is 2 GB.

Recording a CLI session is useful, for example, if you are troubleshooting an issue and want to save detailed information or if you want to create a permanent record of space usage at a specific point in time.

Recording a CLI session

You can use the `system script start` and `system script stop` commands to record a CLI session.

Steps

1. To start recording the current CLI session into a file, use the `system script start` command.
For more information about using the `system script start` command, see the man page.
Data ONTAP starts recording your CLI session into the specified file.
2. Proceed with your CLI session.
3. To stop recording the session, use the `system script stop` command.
For more information about using the `system script stop` command, see the man page.
Data ONTAP stops recording your CLI session.

Commands for managing records of CLI sessions

You use the `system script` commands to manage records of CLI sessions.

If you want to...	Use this command...
Start recording the current CLI session in to a specified file	<code>system script start</code>
Stop recording the current CLI session	<code>system script stop</code>
Display information about records of CLI sessions	<code>system script show</code>
Upload a record of a CLI session to an FTP or HTTP destination	<code>system script upload</code>
Delete a record of a CLI session	<code>system script delete</code>

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Managing the automatic timeout period of CLI sessions

The timeout value specifies how long a CLI session remains idle before being automatically terminated. The CLI timeout value is cluster-wide. That is, every node in a cluster uses the same CLI timeout value.

By default, the automatic timeout period of CLI sessions is 30 minutes.

You can manage the timeout value for CLI sessions by using the `system timeout` commands.

Commands for managing the automatic timeout period of CLI sessions

You use the `system timeout` commands to manage the automatic timeout period of CLI sessions.

If you want to...	Use this command...
Display the automatic timeout period for CLI sessions	<code>system timeout show</code>
Modify the automatic timeout period for CLI sessions	<code>system timeout modify</code>

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Accessing the cluster by using OnCommand System Manager browser-based graphic interface

If you prefer to use a graphic interface instead of the command-line interface for accessing and managing the cluster, you can do so by using OnCommand System Manager, which is included with Data ONTAP as a web service, enabled by default, and accessible by using a browser.

Before you begin

You must have a cluster user account configured with the `admin` role and the `http`, `ontapi`, and `console` application types.

Steps

1. Point the web browser to the cluster management LIF in one of the following formats:

- `https://cluster-mgmt-LIF` (if using IPv4)
- `https://[cluster-mgmt-LIF]` (if using IPv6)

`cluster-mgmt-LIF` is the IP address of the cluster management LIF.

Only HTTPS is supported for browser access of OnCommand System Manager.

If the cluster uses a self-signed digital certificate, the browser might display a warning indicating that the certificate is not trusted. You can either acknowledge the risk to continue the access or install a Certificate Authority (CA) signed digital certificate on the cluster for server authentication.

2. Log in to OnCommand System Manager by using your cluster administrator credential.

Related concepts

[Managing access to web services](#) on page 151

Related tasks

[Accessing a node's log, core dump, and MIB files by using a web browser](#) on page 41

[Installing a server certificate to authenticate the cluster or SVM as an SSL server](#) on page 139

Understanding OnCommand System Manager

System Manager is a graphical management interface that enables you to manage storage systems and storage objects (such as disks, volumes, and aggregates) and perform common management tasks related to storage systems from a web browser. As a cluster administrator, you can use System Manager to administer the entire cluster and its resources.

Important: System Manager is no longer available as an executable file and is now included with Data ONTAP as a web service, enabled by default, and accessible by using a browser.

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, FCoE, and iSCSI for block access.
- Create and configure network components such as subnets, broadcast domains, data and management interfaces, and interface groups.
- Set up and manage mirroring and vaulting relationships.
- Perform cluster management, storage node management, and Storage Virtual Machine (SVM, formerly known as Vserver) management operations.
- Create and configure SVMs, manage storage objects associated with SVMs, and manage SVM services.
- Monitor and manage HA configurations in a cluster.
- Configure Service Processors to remotely log in, manage, monitor, and administer the node, regardless of the state of the node.

For more information about System Manager, see the NetApp Support Site.

Related information

[NetApp Support Site: mysupport.netapp.com](http://mysupport.netapp.com)

Ways to manage access to OnCommand System Manager

You can enable or disable a web browser's access to OnCommand System Manager. You can also view the System Manager log.

You can control a web browser's access to System Manager by using `vserver services web modify -name sysmgr -vserver cluster_name -enabled [true|false]`.

System Manager logging is recorded in the `/mroot/etc/log/mlog/sysmgr.log` files of the node that hosts the cluster management LIF at the time System Manager is accessed. You can view the log files by using a browser. The System Manager log is also included in AutoSupport messages.

Cluster management basics (cluster administrators only)

The cluster administrator can set up a new cluster by using System Setup. After the cluster is created, the cluster administrator can display the cluster status and attributes, rename the cluster, or assign epsilon to another node in the cluster.

Related information

[NetApp Documentation: System Setup \(current releases\)](#)

What a cluster is

A cluster consists of one or more nodes grouped together as (HA pairs) to form a scalable cluster. Creating a cluster enables the 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.

- The maximum number of nodes within a cluster depends on the platform model and licensed protocols.
- Each node in the cluster can view and manage 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, spans the cluster.
- The nodes in a cluster communicate over a dedicated, physically isolated and secure Ethernet network. The cluster logical interfaces (LIFs) on each node in the cluster must be on the same subnet.
- When new nodes are added to a cluster, there is no need to update clients to point to the new nodes. The existence of the new nodes is transparent to the clients.
- If you have a two-node cluster (a single HA pair), you must configure cluster high availability (HA).
- You can create a cluster on a stand-alone node, called a single-node cluster. This configuration does not require a cluster network, and enables you to use the cluster ports to serve data traffic. However, nondisruptive operations are not supported on single-node clusters.

Related concepts

[Understanding quorum and epsilon](#) on page 32

[What the cluster management server is](#) on page 31

[Considerations for single-node clusters](#) on page 31

Related information

[NetApp Hardware Universe](#)

[Clustered Data ONTAP 8.3 High-Availability Configuration Guide](#)

[Clustered Data ONTAP 8.3 Network Management Guide](#)

Considerations for single-node clusters

A single-node cluster is a special implementation of a cluster running on a stand-alone node. You can deploy a single-node cluster if your workload only requires a single node, but does not need nondisruptive operations.

For example, you could deploy a single-node cluster to provide data protection for a remote office. In this scenario, the single-node cluster would use SnapMirror and SnapVault to replicate the site's data to the primary data center.

In a single-node cluster, the HA mode is set to `non_HA`, which enables the node to use all of the nonvolatile memory (NVRAM) on the NVRAM card. In addition, single-node clusters do not use a cluster network, and you can use the cluster ports as data ports that can host data LIFs.

Single-node clusters are typically configured when the cluster is set up. However, you can also remove nodes from an existing cluster to create a single-node cluster.

The following features and operations are not supported for single-node clusters:

- Storage failover and cluster HA
Single-node clusters operate in a stand-alone HA mode. If the node goes offline, clients cannot access data stored in the cluster.
- Any operation that requires more than one node
For example, you cannot move volumes or perform most copy operations.
- Infinite Volumes
Infinite Volumes must contain aggregates from at least two nodes.
- Storing cluster configuration backups in the cluster
By default, the configuration backup schedule creates backups of the cluster configuration and stores them on different nodes throughout the cluster. However, if the cluster consists of a single node and you experience a disaster in which the node becomes inaccessible, you cannot recover the cluster unless the cluster configuration backup file is stored at a remote URL.

Related tasks

[Adding nodes to the cluster](#) on page 37

[Removing nodes from the cluster](#) on page 39

Related references

[Commands for managing configuration backup schedules](#) on page 181

Related information

[NetApp Documentation: System Setup \(current releases\)](#)

What the cluster management server is

The cluster management server, also called an *adminSVM*, is a specialized Storage Virtual Machine (SVM) implementation that presents the cluster as a single manageable entity. In addition to serving as the highest-level administrative domain, the cluster management server owns resources that do not logically belong with a data SVM.

The cluster management server is always available on the cluster. You can access the cluster management server through the console or cluster management LIF.

Upon failure of its home network port, the cluster management LIF automatically fails over to another node in the cluster. Depending on the connectivity characteristics of the management protocol you are using, you might or might not notice the failover. If you are using a connectionless protocol (for example, SNMP) or have a limited connection (for example, HTTP), you are not likely to notice the failover. However, if you are using a long-term connection (for example, SSH), then you will have to reconnect to the cluster management server after the failover.

When you create a cluster, all of the characteristics of the cluster management LIF are configured, including its IP address, netmask, gateway, and port.

Unlike a data SVM or node SVM, a cluster management server does not have a root volume or host user volumes (though it can host system volumes). Furthermore, a cluster management server can only have LIFs of the cluster management type.

If you run the `vserver show` command, the cluster management server appears in the output listing for that command.

Related concepts

[Types of SVMs](#) on page 76


Understanding quorum and epsilon

Quorum and epsilon are important measures of cluster health and function that together indicate how clusters address potential communications and connectivity challenges.

Quorum is a precondition for a fully functioning cluster. When a cluster is in quorum, a simple majority of nodes are healthy and can communicate with each other. When quorum is lost, the cluster loses the ability to accomplish normal cluster operations. Only one collection of nodes can have quorum at any one time because all of the nodes collectively share a single view of the data. Therefore, if two non-communicating nodes are permitted to modify the data in divergent ways, it is no longer possible to reconcile the data into a single data view.

Each node in the cluster participates in a voting protocol that elects one node *master*; each remaining node is a *secondary*. The master node is responsible for synchronizing information across the cluster. When quorum is formed, it is maintained by continual voting. If the master node goes offline and the cluster is still in quorum, a new master is elected by the nodes that remain online.

Because there is the possibility of a tie in a cluster that has an even number of nodes, one node has an extra fractional voting weight called *epsilon*. If the connectivity between two equal portions of a large cluster fails, the group of nodes containing epsilon maintains quorum, assuming that all of the nodes are healthy. For example, the following illustration shows a four-node cluster in which two of the nodes have failed. However, because one of the surviving nodes holds epsilon, the cluster remains in quorum even though there is not a simple majority of healthy nodes.



Epsilon is automatically assigned to the first node when the cluster is created. If the node that holds epsilon becomes unhealthy, takes over its high-availability partner, or is taken over by its high-availability partner, then epsilon is automatically reassigned to a healthy node in a different HA pair.

Taking a node offline can affect the ability of the cluster to remain in quorum. Therefore, Data ONTAP issues a warning message if you attempt an operation that will either take the cluster out of quorum or else put it one outage away from a loss of quorum. You can disable the quorum warning

messages by using the `cluster quorum-service options modify` command at the advanced privilege level.

In general, assuming reliable connectivity among the nodes of the cluster, a larger cluster is more stable than a smaller cluster. The quorum requirement of a simple majority of half the nodes plus epsilon is easier to maintain in a cluster of 24 nodes than in a cluster of two nodes.

A two-node cluster presents some unique challenges for maintaining quorum. Two-node clusters use *cluster HA*, in which neither node holds epsilon; instead, both nodes are continuously polled to ensure that if one node fails, the other has full read-write access to data, as well as access to logical interfaces and management functions.

Related information

[Clustered Data ONTAP 8.3 High-Availability Configuration Guide](#)

What a cluster replication ring is

A *replication ring* is a set of identical processes running on all nodes in the cluster.

The basis of clustering is the replicated database (RDB). An instance of the RDB is maintained on each node in the cluster. Several processes use the RDB to ensure consistent data across the cluster. The processes include the management application (`mgmt`), volume location database (`vldb`), virtual interface manager (`vi_fmgr`), SAN management daemon (`bcomd`), and configuration replication service (`crs`).

For instance, the `vldb` replication ring for a given cluster consists of all instances of `vldb` running in the cluster.

RDB replication requires healthy cluster links among all nodes in the cluster. If the cluster network fails in whole or in part, file services can become unavailable. The advanced command `cluster ring show` displays the status of replication rings and can assist with troubleshooting efforts.

Displaying information about the nodes in a cluster

You can display node names, whether the nodes are healthy, and whether they are eligible to participate in the cluster. At the advanced privilege level, you can also display whether a node holds epsilon.

Step

1. To display information about the nodes in a cluster, use the `cluster show` command.

If you want the output to show whether a node holds epsilon, run the command at the advanced privilege level.

Examples of displaying the nodes in a cluster

The following example displays information about all nodes in a four-node cluster:

```
cluster1::> cluster show
Node           Health  Eligibility
-----
node1          true    true
node2          true    true
node3          true    true
node4          true    true
```

The following example displays detailed information about the node named “node1” at the advanced privilege level:

```
cluster1::> set -privilege advanced
Warning: These advanced commands are potentially dangerous; use them only
when
directed to do so by support personnel.
Do you want to continue? {y|n}: y

cluster1::*> cluster show -node node1

      Node: node1
Node UUID: a67f9f34-9d8f-11da-b484-000423b6f094
      Epsilon: false
Eligibility: true
      Health: true
```

Displaying cluster attributes

You can display a cluster's unique identifier (UUID), name, serial number, location, and contact information.

Step

1. To display a cluster's attributes, use the `cluster identity show` command.

Example of displaying cluster attributes

The following example displays the name, serial number, location, and contact information of a cluster.

```
cluster1::> cluster identity show

      Cluster UUID: 1cd8a442-86d1-11e0-ae1c-123478563412
      Cluster Name: cluster1
Cluster Serial Number: 1-80-123456
      Cluster Location: Sunnyvale
      Cluster Contact: jsmith@example.com
```

Modifying cluster attributes

You can modify a cluster's attributes, such as the cluster name, location, and contact information as needed.

About this task

You cannot change a cluster's UUID, which is set when the cluster is created.

Step

1. To modify cluster attributes, use the `cluster identity modify` command.

The `-name` parameter specifies the name of the cluster. The `cluster identity modify man` page describes the rules for specifying the cluster's name.

The `-location` parameter specifies the location for the cluster.

The `-contact` parameter specifies the contact information such as a name or e-mail address.

Example of renaming a cluster

The following command renames the current cluster (“cluster1”) to “cluster2”:

```
cluster1::> cluster identity modify -name cluster2
```

Related information

Clustered Data ONTAP 8.3.2 man page: [cluster identity modify - Modify the cluster's attributes](#)

Displaying the status of cluster replication rings

You can display the status of cluster replication rings to help you diagnose cluster-wide problems. If your cluster is experiencing problems, support personnel might ask you to perform this task to assist with troubleshooting efforts.

Step

1. To display the status of cluster replication rings, use the `cluster ring show` command at the advanced privilege level.

Example of displaying cluster ring-replication status

The following example displays the status of the VLDB replication ring on a node named node0:

```
cluster1::> set -privilege advanced
Warning: These advanced commands are potentially dangerous; use
them only when
        directed to do so by support personnel.
Do you wish to continue? (y or n): y

cluster1:*> cluster ring show -node node0 -unitname vldb
      Node: node0
      Unit Name: vldb
      Status: master
      Epoch: 5
      Master Node: node0
      Local Node: node0
      DB Epoch: 5
      DB Transaction: 56
      Number Online: 4
      RDB UUID: e492d2c1-fc50-11e1-bae3-123478563412
```

Managing nodes (cluster administrators only)

A *node* is a controller in a cluster. You can display information about a node, set node attributes, rename a node, add or remove a node, or start or stop a node. You can also manage a node remotely by using the Service Processor (SP).

A node is connected to other nodes in the cluster over a cluster network. It is also connected to the disk shelves that provide physical storage for the Data ONTAP system or to 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 SVM represents a node in the cluster. If you run the `vserver show` command, the output includes node SVMs in the list.

Displaying node attributes

You can display the attributes of one or more nodes in the cluster, for example, the name, owner, location, model number, serial number, how long the node has been running, health state, and eligibility to participate in a cluster.

Step

1. To display the attributes of a specified node or about all nodes in a cluster, use the `system node show` command.

Example of displaying information about a node

The following example displays detailed information about node1:

```
cluster1::> system node show -node node1
      Node: node1
      Owner: Eng IT
      Location: Lab 5
      Model: model_number
      Serial Number: 12345678
      Asset Tag: -
      Uptime: 23 days 04:42
      NVRAM System ID: 118051205
      System ID: 0118051205
      Vendor: NetApp
      Health: true
      Eligibility: true
```

Modifying node attributes

You can modify the attributes of a node as needed. The attributes you can modify include the node's owner information, location information, asset tag, and eligibility to participate in the cluster.

About this task

A node's eligibility to participate in the cluster can be modified at the advanced privilege level by using the `-eligibility` parameter of the `system node modify` or `cluster modify` command. If you set a node's eligibility to **false**, the node becomes inactive in the cluster.

Attention: You should avoid setting a node's eligibility to `false` except for situations such as restoring the node configuration or prolonged node maintenance. If you set a node to be ineligible, it stops serving SAN data until the node is reset to eligible and rebooted. NAS data access to the node might also be impacted when the node is ineligible.

Step

1. Use the `system node modify` command to modify a node's attributes.

Example of modifying node attributes

The following command modifies the attributes of the “node1” node. The node's owner is set to “Joe Smith” and its asset tag is set to “js1234”:

```
cluster1::> system node modify -node node1 -owner "Joe Smith"
-assettag js1234
```

Renaming a node

You can change a node's name as needed.

Step

1. To rename a node, use the `system node rename` command.

The `-newname` parameter specifies the new name for the node. The `system node rename man` page describes the rules for specifying the node name.

If you want to rename multiple nodes in the cluster, you must run the command for each node individually.

Example of renaming a node

The following command renames node “node1” to “node1a”:

```
cluster1::> system node rename -node node1 -newname node1a
```

Related information

[Clustered Data ONTAP 8.3.2 man page: system node rename - Rename a node](#)

Adding nodes to the cluster

After a cluster is created, you can expand it by adding nodes to it. You add only one node at a time.

Before you begin

- If you are adding nodes to a multiple-node cluster, more than half of the existing nodes in the cluster must be healthy (indicated by `cluster show`).
- If you are adding nodes to a two-node switchless cluster, you must have installed and configured the cluster management and interconnect switches before adding additional nodes. The switchless cluster functionality is supported only in a two-node cluster.

When a cluster contains or grows to more than two nodes, cluster HA is not required and is disabled automatically.

- If you are adding a second node to a single-node cluster, the second node must have been installed, and the cluster network must have been configured.
- If the cluster has the SP automatic configuration enabled, the subnet specified for the SP to use must have available resources for the joining node.
A node that joins the cluster uses the specified subnet to perform automatic configuration for the SP.
- You must have gathered the following information for the new node's node management LIF:
 - Port
 - IP address
 - Netmask
 - Default gateway

About this task

Nodes must be in even numbers so that they can form HA pairs. After you start to add a node to the cluster, you must complete the process. The node must be part of the cluster before you can start to add another node.

Steps

1. Power on the node that you want to add to the cluster.

The node boots, and the Node Setup wizard starts on the console.

```
Welcome to node setup.

You can enter the following commands at any time:
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the setup wizard.
    Any changes you made before quitting will be saved.

To accept a default or omit a question, do not enter a value.

Enter the node management interface port [e0c]:
```

2. Exit the Node Setup wizard:

```
exit
```

The Node Setup wizard exits, and a login prompt appears, warning that you have not completed the setup tasks.

3. Log in to the admin account by using the admin user name.

4. Start the Cluster Setup wizard:

```
cluster setup
```

```
::> cluster setup

Welcome to the cluster setup wizard.

You can enter the following commands at any time:
```

```
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the cluster setup wizard.
Any changes you made before quitting will be saved.
```

```
You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.
```

```
Do you want to create a new cluster or join an existing cluster?
{create, join}:
```

5. Join the node to the cluster:

```
join
```

6. Follow the prompts to set up the node and join it to the cluster:

- To accept the default value for a prompt, press Enter.
- To enter your own value for a prompt, enter the value, and then press Enter.

7. Repeat the preceding steps for each additional node that you want to add.

After you finish

After adding nodes to the cluster, you should enable storage failover for each HA pair.

Related concepts

[Considerations for the SP network configuration](#) on page 51

Related information

[Clustered Data ONTAP 8.3 High-Availability Configuration Guide](#)

[Clustered Data ONTAP: Adding a second controller module to create an HA pair](#)

[Clustered Data ONTAP Setup Guide for Cisco Switches](#)

Removing nodes from the cluster

You can remove unwanted nodes from the cluster. You can remove only one node at a time. After you remove a node, you must also remove its failover partner. Removing a node causes its data to become inaccessible or erased.

Before you begin

More than half of the nodes in the cluster must be healthy (indicated by `cluster show`).

All data on the node that you want to remove must have been evacuated.

About this task

Attention: Removing a node from the cluster (`cluster unjoin`) makes all system and user data from the disks that are connected to the node inaccessible to user access. After removing a node from the cluster, if you need to join the node back to the same cluster, you should contact technical support for assistance about restoring data.

Steps

1. If Storage Encryption is enabled on the cluster and the self-encrypting disks (SEDs) are operating in the Federal Information Processing Standard (FIPS) compliance mode, disable the FIPS

compliance mode so that the disks attached to the node can be sanitized and repurposed after the node is removed from the cluster:

- a. Display the SEDs that are currently in the FIPS compliance mode by using the `storage encryption disk show` command with the `-fips` parameter.
 - b. Disable the FIPS compliance mode by using the `storage encryption disk modify` command with the `-disk disk_path_name` and `-fips-key-id 0x0` parameters at the advanced level.
 - c. Display the SEDs again and verify that the modified SEDs are no longer in the FIPS compliance mode by using the `storage encryption disk show` command with the `-fips` parameter.
2. If the node you want to remove is the current master node, reboot the node by using the `system node reboot` command to enable another node in the cluster to be elected as the master node.
The master node is the node that holds processes such as `mgmt`, `vldb`, `vifmgr`, `bcomd`, and `crs`. The `cluster ring show advanced` command shows the current master node.
 3. Use the advanced command `cluster unjoin` from another node in the cluster to remove a node from the cluster.

The system informs you of the following:

- You must also remove the node's failover partner from the cluster.
- After the node is removed and before it can rejoin a cluster, you must use boot menu option **(4) Clean configuration and initialize all disks** to erase the node's configuration and initialize all disks.

A failure message is generated if you have conditions that you must address before removing the node. For example, the message might indicate that the node has shared resources that you must remove or that the node is in a cluster HA configuration or storage failover configuration that you must disable.

If the node is the quorum master, the cluster will briefly lose and then return to quorum. This quorum loss is temporary and does not affect any data operations.

4. If a failure message indicates error conditions, address those conditions and rerun the `cluster unjoin` command.
The node is automatically rebooted after it is successfully removed from the cluster.
5. If the node is being repurposed, do the following after the node is rebooted:
 - a. During the boot process, press Ctrl-C to display the boot menu when prompted to do so.
 - b. Select boot menu option **(4) Clean configuration and initialize all disks** to erase the node's configuration and initialize all disks.
6. Repeat the preceding steps to remove the failover partner from the cluster.

After you finish

If you removed nodes to have a single-node cluster, you should modify the cluster ports to serve data traffic by modifying the cluster ports to be data ports, and creating data LIFs on the data ports.

Related tasks

[Rebooting a node at the system prompt](#) on page 45

Related information*Clustered Data ONTAP 8.3 Network Management Guide*

Accessing a node's log, core dump, and MIB files by using a web browser

The Service Processor Infrastructure (**spi**) web service is enabled by default to enable a web browser to access the log, core dump, and MIB files of a node in the cluster. The files remain accessible even when the node is down, provided that the node is taken over by its partner.

Before you begin

- The cluster management LIF must be up.
You can use the management LIF of the cluster or a node to access the **spi** web service. However, using the cluster management LIF is recommended.
The `network interface show` command displays the status of all LIFs in the cluster.
- If your user account does not have the “admin” role (which has access to the **spi** web service by default), your access-control role must be granted access to the **spi** web service.
The `vserver services web access show` command shows what roles are granted access to which web services.
- If you are not using the “admin” user account (which includes the **http** access method by default), your user account must be set up with the **http** access method.
The `security login show` command shows user accounts' access and login methods and their access-control roles.
- If you want to use HTTPS for secure web access, SSL must be enabled and a digital certificate must be installed.
The `system services web show` command displays the configuration of the web protocol engine at the cluster level.

About this task

The **spi** web service is enabled by default, and the service can be disabled manually (`vserver services web modify -vserver * -name spi -enabled false`).

The “admin” role is granted access to the **spi** web service by default, and the access can be disabled manually (`services web access delete -vserver cluster_name -name spi -role admin`).

Steps

1. Point the web browser to the **spi** web service URL in one of the following formats:
 - `http://cluster-mgmt-LIF/spi/`
 - `https://cluster-mgmt-LIF/spi/`

cluster-mgmt-LIF is the IP address of the cluster management LIF.
2. When prompted by the browser, enter your user account and password.
After your account is authenticated, the browser displays links to the `/mroot/etc/log/`, `/mroot/etc/crash/`, and `/mroot/etc/mib/` directories of each node in the cluster.

Related concepts

[Managing the web protocol engine](#) on page 152

[Managing web services](#) on page 154

[Managing access to web services](#) on page 151

[Access methods for user accounts](#) on page 117

[Managing SSL](#) on page 156

[Managing audit settings for management activities](#) on page 160

[Managing core dumps \(cluster administrators only\)](#) on page 188

Related information

[Clustered Data ONTAP 8.3 Network Management Guide](#)

Accessing the system console of a node

If a node is hanging at the boot menu or the boot environment prompt, you can access it only through the system console (also called the *serial console*). You can access the system console of a node from an SSH connection to the node's SP or to the cluster.

About this task

Both the SP and Data ONTAP offer commands that enable you to access the system console. However, from the SP, you can access only the system console of its own node. From the cluster, you can access the system console of any node in the cluster.

Steps

1. Access the system console of a node:

If you are in the...	Enter this command...
SP CLI of the node	<code>system console</code>
Data ONTAP CLI	<code>system node run-console</code>

2. Log in to the system console when you are prompted to do so.
3. To exit the system console, press Ctrl-D.

Examples of accessing the system console

The following example shows the result of entering the `system console` command at the “SP node2” prompt. The system console indicates that node2 is hanging at the boot environment prompt. The `boot_ontap` command is entered at the console to boot the node to Data ONTAP. Ctrl-D is then pressed to exit the console and return to the SP.

```
SP node2> system console
Type Ctrl-D to exit.

LOADER>
LOADER> boot_ontap
...
*****
*
* Press Ctrl-C for Boot Menu. *
*
*****
...

```

(Ctrl-D is pressed to exit the system console.)

```
Connection to 123.12.123.12 closed.
SP node2>
```

The following example shows the result of entering the `system node run-console` command from Data ONTAP to access the system console of node2, which is hanging at the boot environment prompt. The `boot_ontap` command is entered at the console to boot node2 to Data ONTAP. Ctrl-D is then pressed to exit the console and return to Data ONTAP.

```
cluster1::> system node run-console -node node2
Pressing Ctrl-D will end this session and any further sessions you
might open on top of this session.
Type Ctrl-D to exit.

LOADER>
LOADER> boot_ontap
...
*****
*                                     *
* Press Ctrl-C for Boot Menu.        *
*                                     *
*****
...

```

(Ctrl-D is pressed to exit the system console.)

```
Connection to 123.12.123.12 closed.
cluster1::>
```

Related concepts

[Relationship among the SP CLI, SP console, and system console sessions](#) on page 60

Related tasks

[Accessing the cluster by using SSH](#) on page 11

[Accessing the SP from an administration host](#) on page 58

[Booting Data ONTAP at the boot environment prompt](#) on page 46

Rules governing node root volumes and root aggregates

A node's root volume contains special directories and files for that node. The root aggregate contains the root volume. A few rules govern a node's root volume and root aggregate.

A node's root volume is a FlexVol volume that is installed at the factory or by setup software. It is reserved for system files, log files, and core files. The directory name is `/mroot`, which is accessible only through the systemshell by technical support. The minimum size for a node's root volume depends on the platform model.

- The following rules govern the node's root volume:
 - Unless technical support instructs you to do so, do not modify the configuration or content of the root volume.
 - Do not store user data in the root volume.

Storing user data in the root volume increases the storage giveback time between nodes in an HA pair.

- 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.
- The root aggregate is dedicated to the node's root volume only.
Data ONTAP prevents you from creating other volumes in the root aggregate.

Related information

[NetApp Hardware Universe](#)

Freeing up space on a node's root volume

A warning message appears when a node's root volume has become full or almost full. The node cannot operate properly when its root volume is full. You can free up space on a node's root volume by deleting core dump files, packet trace files, and root volume Snapshot copies.

Steps

1. Display the node's core dump files and their names by using the `system node coredump show` command.
2. Delete unwanted core dump files from the node by using the `system node coredump delete` command.
3. Access the nodeshell by entering the following command:
system node run -node *nodename*
nodename is the name of the node whose root volume space you want to free up.
4. Switch to the nodeshell advanced privilege level by entering the following command in the nodeshell:
priv set advanced
5. Display and delete the node's packet trace files through the nodeshell:
 - a. Display all files in the node's root volume by entering the following command:
ls /etc/
 - b. If any packet trace files (`*.trc`) are in the node's root volume, delete them individually by entering the following command:
rm /etc/*file_name*.trc
6. Identify and delete the node's root volume Snapshot copies through the nodeshell:
 - a. Identify the root volume name by entering the following command:
vol status
The root volume is indicated by the word "root" in the "Options" column of the `vol status` command output.

Example

In the following example, the root volume is `vol10`.

```
node1*> vol status
```

Volume	State	Status	Options
vol0	online	raid_dp, flex 64-bit	root, nvfail=on

- b. Display root volume Snapshot copies by entering the following command:

```
snap list root_vol_name
```

- c. Delete unwanted root volume Snapshot copies by entering the following command:

```
snap delete root_vol_name snapshot_name
```

7. Exit the nodeshell and return to the clustershell by entering the following command:

```
exit
```

Starting or stopping a node

You might need to start or stop a node for maintenance or troubleshooting reasons. You can do so from the Data ONTAP CLI, the boot environment prompt, or the SP CLI.

Using the SP CLI command `system power off` or `system power cycle` to turn off or power-cycle a node might cause an improper shutdown of the node (also called a *dirty shutdown*) and is not a substitute for a graceful shutdown using the Data ONTAP `system node halt` command.

Related concepts

[Managing a node remotely by using the Service Processor](#) on page 49

Related information

[Clustered Data ONTAP 8.3 High-Availability Configuration Guide](#)

Rebooting a node at the system prompt

You can reboot a node in normal mode from the system prompt. A node is configured to boot from the boot device, such as a PC CompactFlash card.

Steps

1. If the cluster contains four or more nodes, verify that the node to be rebooted does not hold epsilon:

- a. Set the privilege level to advanced:

```
set -privilege advanced
```

- b. Determine which node holds epsilon:

```
cluster show
```

Example

The following example shows that “node1” holds epsilon:

```
cluster1::*> cluster show
```

Node	Health	Eligibility	Epsilon
node1	true	true	true

```
node2          true    true    false
node3          true    true    false
node4          true    true    false
4 entries were displayed.
```

- c. If the node to be rebooted holds epsilon, then remove epsilon from the node:

```
cluster modify -node node_name -epsilon false
```

- d. Assign epsilon to a different node that will remain up:

```
cluster modify -node node_name -epsilon true
```

- e. Return to the admin privilege level:

```
set -privilege admin
```

2. Use the `system node reboot` command to reboot the node.

If you do not specify the `-skip-lif-migration` parameter, the command attempts to migrate data and cluster management LIFs synchronously to another node prior to the reboot. If the LIF migration fails or times out, the rebooting process is aborted, and Data ONTAP displays an error to indicate the LIF migration failure.

Example

```
cluster1::> system node reboot -node node1 -reason "software upgrade"
```

The node begins the reboot process. The Data ONTAP login prompt appears, indicating that the reboot process is complete.

Booting Data ONTAP at the boot environment prompt

You can boot the current release or the backup release of Data ONTAP when you are at the boot environment prompt of a node.

Steps

1. Access the boot environment prompt from the storage system prompt by using the `system node halt` command.

The storage system console displays the boot environment prompt.

2. At the boot environment prompt, enter one of the following commands:

To boot...	Enter...
The current release of Data ONTAP	<code>boot_ontap</code>
The Data ONTAP primary image from the boot device	<code>boot_primary</code>
The Data ONTAP backup image from the boot device	<code>boot_backup</code>

If you are unsure about which image to use, you should use `boot_ontap` in the first instance.

Related tasks

[Accessing the system console of a node](#) on page 42

Shutting down a node

You can shut down a node if it becomes unresponsive or if support personnel direct you to do so as part of troubleshooting efforts.

Steps

1. If the cluster contains four or more nodes, verify that the node to be shut down does not hold epsilon:

- a. Set the privilege level to advanced:

```
set -privilege advanced
```

- b. Determine which node holds epsilon:

```
cluster show
```

Example

The following example shows that “node1” holds epsilon:

```
cluster1::*> cluster show
Node           Health  Eligibility  Epsilon
-----
node1          true    true         true
node2          true    true         false
node3          true    true         false
node4          true    true         false
4 entries were displayed.
```

- c. If the node to be shut down holds epsilon, then remove epsilon from the node:

```
cluster modify -node node_name -epsilon false
```

- d. Assign epsilon to a different node that will remain up:

```
cluster modify -node node_name -epsilon true
```

- e. Return to the admin privilege level:

```
set -privilege admin
```

2. Use the `system node halt` command to shut down the node.

If you do not specify the `-skip-lif-migration` parameter, the command attempts to migrate data and cluster management LIFs synchronously to another node prior to the shutdown. If the LIF migration fails or times out, the shutdown process is aborted, and Data ONTAP displays an error to indicate the LIF migration failure.

You can manually trigger a core dump with the shutdown by using both the `-dump` parameter.

Example

The following example shuts down the node named “node1” for hardware maintenance:

```
cluster1::> system node halt -node node1 -reason 'hardware
maintenance'
```

Managing a node by using the boot menu

You can use the boot menu to correct configuration problems of a node, reset the admin password, initialize disks, reset node configuration, and restore node configuration information back to the boot device.

Steps

1. Reboot the node to access the boot menu by using the `system node reboot` command at the system prompt.

The node begins the reboot process.

2. During the reboot process, press Ctrl-C to display the boot menu when prompted to do so.

The node displays the following options for the boot menu:

```
(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
Selection (1-8)?
```

Note: Boot menu option (2) **Boot without /etc/rc** is obsolete and takes no effect on the system.

3. Select one of the following options by entering the corresponding number:

To...	Select...
Continue to boot the node in normal mode	1) Normal Boot
Change the password of the node, which is also the “admin” account password	3) Change Password
Initialize the node's disks and create a root volume for the node	4) Clean configuration and initialize all disks Attention: This menu option erases all data on the disks of the node and resets your node configuration to the factory default settings. You select this menu option after the node has unjoined the cluster and before it rejoins another cluster. This menu option reboots the node before initializing the disks. For a node that uses array, this menu option initializes only the disks on the disk shelf, not the array LUNs. For a node that does not have a disk shelf, this menu option initializes the root volume on the storage array. If the node you want to initialize has disks that are partitioned for root-data partitioning, the disks must be unpartitioned before the node can be initialized.
Perform aggregate and disk maintenance operations and obtain detailed aggregate and disk information.	5) Maintenance mode boot You exit Maintenance mode by using the <code>halt</code> command.

To...	Select...
Restore the configuration information from the node's root volume to the boot device, such as a PC CompactFlash card	<p>6) Update flash from backup config</p> <p>Data ONTAP stores some node configuration information on the boot device. When the node reboots, the information on the boot device is automatically backed up onto the node's root volume. If the boot device becomes corrupted or needs to be replaced, you use this menu option to restore the configuration information from the node's root volume back to the boot device.</p>
Install new software on the node	<p>7) Install new software first</p> <p>If the Data ONTAP software on the boot device does not include support for the storage array that you want to use for the root volume, you can use this menu option to obtain a version of the software that supports your storage array and install it on the node.</p> <p>This menu option is only for installing a newer version of Data ONTAP software on a node that has no root volume installed. Do <i>not</i> use this menu option to upgrade Data ONTAP.</p>
Reboot the node	<p>8) Reboot node</p>

Related tasks

[Accessing the system console of a node](#) on page 42

Related information

[Clustered Data ONTAP 8.3 Physical Storage Management Guide](#)

Managing a node remotely by using the Service Processor

You can manage a node remotely by using the Service Processor (SP), a remote management device that is included in all supported platform models. The SP stays operational regardless of the operating state of the node.

Data ONTAP-v platforms, such as Data ONTAP Edge, do not have the SP. You use the Data ONTAP-v Administration Tool (dvsadmin) for remote management of Data ONTAP-v platforms.

Understanding the SP

The Service Processor (SP) is a remote management device that enables you to access, monitor, and troubleshoot a node remotely.

The key capabilities of the SP include the following:

- The SP enables you to access a node remotely to diagnose, shut down, power-cycle, or reboot the node, regardless of the state of the node controller.

The SP is powered by a standby voltage, which is available as long as the node has input power to at least one of its power supplies.

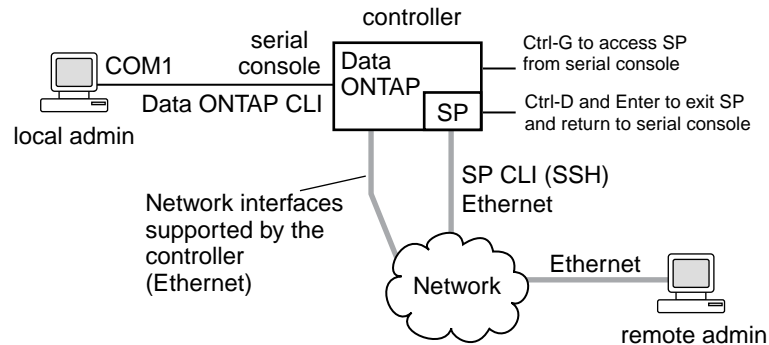
You can log in to the SP by using a Secure Shell client application from an administration host. You can then use the SP CLI to monitor and troubleshoot the node remotely. In addition, you can use the SP to access the serial console and run Data ONTAP commands remotely.

You can access the SP from the serial console or access the serial console from the SP. The SP enables you to open both an SP CLI session and a separate console session simultaneously.

For instance, when a temperature sensor becomes critically high or low, Data ONTAP triggers the SP to shut down the motherboard gracefully. The serial console becomes unresponsive, but you can still press Ctrl-G on the console to access the SP CLI. You can then use the `system power on` or `system power cycle` command from the SP to power on or power-cycle the node.

- The SP monitors environmental sensors and logs events to help you take timely and effective service actions.
The SP monitors environmental sensors such as the node temperatures, voltages, currents, and fan speeds. When an environmental sensor has reached an abnormal condition, the SP logs the abnormal readings, notifies Data ONTAP of the issue, and sends alerts and “down system” notifications as necessary through an AutoSupport message, regardless of whether the node can send AutoSupport messages.
The SP also logs events such as boot progress, Field Replaceable Unit (FRU) changes, events generated by Data ONTAP, and SP command history. You can manually invoke an AutoSupport message to include the SP log files that are collected from a specified node.
Other than generating these messages on behalf of a node that is down and attaching additional diagnostic information to AutoSupport messages, the SP has no effect on the AutoSupport functionality. The AutoSupport configuration settings and message content behavior are inherited from Data ONTAP.
Note: The SP does not rely on the `-transport` parameter setting of the `system node autosupport modify` command to send notifications. The SP only uses the Simple Mail Transport Protocol (SMTP) and requires its host’s AutoSupport configuration to include mail host information.
If SNMP is enabled, the SP generates SNMP traps to configured trap hosts for all “down system” events.
- The SP has a nonvolatile memory buffer that stores up to 4,000 events in a system event log (SEL) to help you diagnose issues.
The SEL stores each audit log entry as an audit event. It is stored in onboard flash memory on the SP. The event list from the SEL is automatically sent by the SP to specified recipients through an AutoSupport message.
The SEL contains the following information:
 - Hardware events detected by the SP—for example, sensor status about power supplies, voltage, or other components
 - Errors detected by the SP—for example, a communication error, a fan failure, or a memory or CPU error
 - Critical software events sent to the SP by the node—for example, a panic, a communication failure, a boot failure, or a user-triggered “down system” as a result of issuing the `SP system reset` or `system power cycle` command
- The SP monitors the serial console regardless of whether administrators are logged in or connected to the console.
When messages are sent to the console, the SP stores them in the console log. The console log persists as long as the SP has power from either of the node power supplies. Because the SP operates with standby power, it remains available even when the node is power-cycled or turned off.
- Hardware-assisted takeover is available if the SP is configured.
- The SP API service enables Data ONTAP to communicate with the SP over the network.
The service enhances Data ONTAP management of the SP by supporting network-based functionality such as using the network interface for the SP firmware update, enabling a node to access another node’s SP functionality or system console, and uploading the SP log from another node.
You can modify the configuration of the SP API service by changing the port the service uses, renewing the SSL and SSH certificates that are used by the service for internal communication, or disabling the service entirely.

The following diagram illustrates access to Data ONTAP and the SP of a node. The SP interface is accessed through the Ethernet port (indicated by a wrench icon on the rear of the chassis):



Related information

[Clustered Data ONTAP 8.3 High-Availability Configuration Guide](#)

Configuring the SP network

Configuring the SP network enables you to access the SP by using its IP address. You can set up the SP automatic network configuration at the cluster level (recommended) or manually configure the SP network for individual nodes. You can also modify the SP API service that Data ONTAP uses to access the SP over the network.

Considerations for the SP network configuration

You can enable cluster-level, automatic network configuration for the SP (recommended). You can also leave the SP automatic network configuration disabled (the default) and manage the SP network configuration manually at the node level. A few considerations exist for each case.

The SP automatic network configuration enables the SP to use address resources (including the IP address, subnet mask, and gateway address) from the specified subnet to set up its network automatically. With the SP automatic network configuration, you do not need to manually assign IP addresses for the SP of each node. By default, the SP automatic network configuration is disabled; this is because enabling the configuration requires that the subnet to be used for the configuration be defined in the cluster first.

If you enable the SP automatic network configuration, the following scenarios and considerations apply:

- If the SP has never been configured, the SP network is configured automatically based on the subnet specified for the SP automatic network configuration.
- If the SP was previously configured manually, or if the existing SP network configuration is based on a different subnet, the SP network of all nodes in the cluster are reconfigured based on the subnet that you specify in the SP automatic network configuration.
The reconfiguration could result in the SP being assigned a different address, which might have an impact on your DNS configuration and its ability to resolve SP host names. As a result, you might need to update your DNS configuration.
- A node that joins the cluster uses the specified subnet to configure its SP network automatically.
- The `system service-processor network modify` command does not enable you to change the SP IP address.

When the SP automatic network configuration is enabled, the command only allows you to enable or disable the SP network interface.

- If the SP automatic network configuration was previously enabled, disabling the SP network interface results in the assigned address resource being released and returned to the subnet.
- If you disable the SP network interface and then reenable it, the SP might be reconfigured with a different address.

If the SP automatic network configuration is disabled (the default), the following scenarios and considerations apply:

- If the SP has never been configured, SP IPv4 network configuration defaults to using IPv4 DHCP, and IPv6 is disabled.
A node that joins the cluster also uses IPv4 DHCP for its SP network configuration by default.
- The `system service-processor network modify` command enables you to configure a node's SP IP address.
A warning message appears when you attempt to manually configure the SP network with addresses that are allocated to a subnet. Ignoring the warning and proceeding with the manual address assignment might result in a scenario with duplicate addresses.

If the SP automatic network configuration is disabled after having been enabled previously, the following scenarios and considerations apply:

- If the SP automatic network configuration has the IPv4 address family disabled, the SP IPv4 network defaults to using DHCP, and the `system service-processor network modify` command enables you to modify the SP IPv4 configuration for individual nodes.
- If the SP automatic network configuration has the IPv6 address family disabled, the SP IPv6 network is also disabled, and the `system service-processor network modify` command enables you to enable and modify the SP IPv6 configuration for individual nodes.

Related tasks

[Enabling the SP automatic network configuration](#) on page 52

[Configuring the SP network manually](#) on page 53

Enabling the SP automatic network configuration

Enabling the SP to use automatic network configuration is preferred over manually configuring the SP network. Because the SP automatic network configuration is cluster wide, you do not need to manually manage the SP network for individual nodes.

Before you begin

- The subnet you want to use for the SP automatic network configuration must already be defined in the cluster and must have no resource conflicts with the SP network interface.
The `network subnet show` command displays subnet information for the cluster.
The parameter that forces subnet association (the `-force-update-lif-associations` parameter of the `network subnet` commands) is supported only on network LIFs and not on the SP network interface.
- If you want to use IPv6 connections for the SP, IPv6 must already be configured and enabled for Data ONTAP.
The `network options ipv6 show` command displays the current state of IPv6 settings for Data ONTAP.

Steps

1. Specify the IPv4 or IPv6 address family and name for the subnet that you want the SP to use by using the `system service-processor network auto-configuration enable` command.

2. Display the SP automatic network configuration by using the `system service-processor network auto-configuration show` command.
3. If you subsequently want to disable or reenble the SP IPv4 or IPv6 network interface for all nodes that are in quorum, use the `system service-processor network modify` command with the `-address-family [IPv4|IPv6]` and `-enable [true|false]` parameters.

When the SP automatic network configuration is enabled, you cannot modify the SP IP address for a node that is in quorum. You can only enable or disable the SP IPv4 or IPv6 network interface.

If a node is out of quorum, you can modify the node's SP network configuration, including the SP IP address, by running `system service-processor network modify` from the node and confirming that you want to override the SP automatic network configuration for the node.

However, when the node joins the quorum, the SP automatic reconfiguration takes place for the node based on the specified subnet.

Related concepts

[Considerations for the SP network configuration](#) on page 51

Related information

[Clustered Data ONTAP 8.3 Network Management Guide](#)

Configuring the SP network manually

If you do not have automatic network configuration set up for the SP, you must manually configure a node's SP network for the SP to be accessible by using an IP address.

Before you begin

If you want to use IPv6 connections for the SP, IPv6 must already be configured and enabled for Data ONTAP. The `network options ipv6` commands manage IPv6 settings for Data ONTAP.

About this task

You can configure the SP to use IPv4, IPv6, or both. The SP IPv4 configuration supports static and DHCP addressing, and the SP IPv6 configuration supports static addressing only.

If the SP automatic network configuration has been set up, you do not need to manually configure the SP network for individual nodes, and the `system service-processor network modify` command allows you to only enable or disable the SP network interface.

Steps

1. Configure the SP network for a node by using by using the `system service-processor network modify` command.
 - The `-address-family` parameter specifies whether the IPv4 or IPv6 configuration of the SP is to be modified.
 - The `-enable` parameter enables the network interface of the specified IP address family.
 - The `-dhcp` parameter specifies whether to use the network configuration from the DHCP server or the network address that you provide.
You can enable DHCP (by setting `-dhcp` to `v4`) only if you are using IPv4. You cannot enable DHCP for IPv6 configurations.
 - The `-ip-address` parameter specifies the public IP address for the SP.

A warning message appears when you attempt to manually configure the SP network with addresses that are allocated to a subnet. Ignoring the warning and proceeding with the manual address assignment might result in a duplicate address assignment.

- The `-netmask` parameter specifies the netmask for the SP (if using IPv4.)
 - The `-prefix-length` parameter specifies the network prefix-length of the subnet mask for the SP (if using IPv6.)
 - The `-gateway` parameter specifies the gateway IP address for the SP.
2. Repeat Step 1 to configure the SP network for each node in the cluster.
 3. Display the SP network configuration and verify the SP setup status by using the `system service-processor network show` command with the `-instance` or `-fields setup-status` parameters.

The SP setup status for a node can be one of the following:

- `not-setup` – Not configured
- `succeeded` -- Configuration succeeded
- `in-progress` -- Configuration in progress
- `failed` -- Configuration failed

Example of configuring the SP network

The following example configures the SP of a node to use IPv4, enables the SP, and displays the SP network configuration to verify the settings:

```
cluster1::> system service-processor network modify -node local
-address-family IPv4 -enable true -ip-address 192.168.123.98
-netmask 255.255.255.0 -gateway 192.168.123.1

cluster1::> system service-processor network show -instance -node local

                Node: node1
                Address Type: IPv4
                Interface Enabled: true
                Type of Device: SP
                  Status: online
                Link Status: up
                DHCP Status: none
                IP Address: 192.168.123.98
                MAC Address: ab:cd:ef:fe:ed:02
                Netmask: 255.255.255.0
                Prefix Length of Subnet Mask: -
                Router Assigned IP Address: -
                Link Local IP Address: -
                Gateway IP Address: 192.168.123.1
                Time Last Updated: Thu Apr 10 17:02:13 UTC 2014
                Subnet Name: -
                Enable IPv6 Router Assigned Address: -
                SP Network Setup Status: succeeded
                SP Network Setup Failure Reason: -

                Node: node1
                Address Type: IPv6
                Interface Enabled: false
                Type of Device: SP
                  Status: online
                Link Status: disabled
                DHCP Status: none
                IP Address: -
                MAC Address: ab:cd:ef:fe:ed:02
                Netmask: -
                Prefix Length of Subnet Mask: -
                Router Assigned IP Address: -
                Link Local IP Address: -
```

```

Gateway IP Address: -
Time Last Updated: Thu Apr 10 17:02:13 UTC 2014
Subnet Name: -
Enable IPv6 Router Assigned Address: -
SP Network Setup Status: not-setup
SP Network Setup Failure Reason: -

2 entries were displayed.

cluster1::>
    
```

Related concepts

[Considerations for the SP network configuration](#) on page 51

Related tasks

[Enabling the SP automatic network configuration](#) on page 52

Related information

[Clustered Data ONTAP 8.3 Network Management Guide](#)

Modifying the SP API service configuration

The SP API is a secure network API that enables Data ONTAP to communicate with the SP over the network. You can change the port used by the SP API service, renew the certificates the service uses for internal communication, or disable the service entirely. You need to modify the configuration only in rare situations.

About this task

- The SP API service uses port **50000** by default.
You can change the port value if, for example, you are in a network setting where port **50000** is used for communication by another networking application, or you want to differentiate between traffic from other applications and traffic generated by the SP API service.
- The SSL and SSH certificates used by the SP API service are internal to the cluster and not distributed externally.
In the unlikely event that the certificates are compromised, you can renew them.
- The SP API service is enabled by default.
You only need to disable the SP API service in rare situations, such as in a private LAN where the SP is not configured or used and you want to disable the service.
If the SP API service is disabled, the API does not accept any incoming connections. In addition, functionality such as network-based SP firmware updates and network-based SP “down system” log collection becomes unavailable. The system switches to using the serial interface.

Steps

1. Switch to the advanced privilege level by using the `set -privilege advanced` command.
2. Modify the SP API service configuration:

If you want to...	Use the following command...
Change the port used by the SP API service	<code>system service-processor api-service modify with the -port {49152..65535} parameter</code>

If you want to...	Use the following command...
Renew the SSL and SSH certificates used by the SP API service for internal communication	<pre>system service-processor api-service renew-certificates</pre> <p>If no parameter is specified, only the host certificates (including the client and server certificates) are renewed.</p> <p>If the <code>-renew-all true</code> parameter is specified, both the host certificates and the root CA certificate are renewed.</p>
Disable or reenable the SP API service	<pre>system service-processor api-service modify with the -is-enabled {true false} parameter</pre>

3. Display the SP API service configuration by using the `system service-processor api-service show` command.

Methods of managing SP firmware updates

Data ONTAP includes an SP firmware image that is called the *baseline image*. If a new version of the SP firmware becomes subsequently available, you have the option to download it and update the SP firmware to the downloaded version without upgrading the Data ONTAP version.

Data ONTAP offers the following methods for managing SP firmware updates:

- The SP automatic update functionality is enabled by default, allowing the SP firmware to be automatically updated in the following scenarios:
 - When you upgrade to a new version of Data ONTAP

The Data ONTAP upgrade process automatically includes the SP firmware update, provided that the SP firmware version bundled with Data ONTAP is newer than the SP version running on the node.

Note: Data ONTAP detects a failed SP automatic update and triggers a corrective action to retry the SP automatic update up to three times. If all three retries fail, you should contact technical support.
 - When you download a version of the SP firmware from the NetApp Support Site and the downloaded version is newer than the one that the SP is currently running
 - When you downgrade or revert to an earlier version of Data ONTAP

The SP firmware is automatically updated to the newest compatible version that is supported by the Data ONTAP version you reverted or downgraded to. A manual SP firmware update is not required.

You have the option to disable the SP automatic update functionality by using the `system service-processor image modify` command. However, it is recommended that you leave the functionality enabled. Disabling the functionality can result in suboptimal or nonqualified combinations between the Data ONTAP image and the SP firmware image.

- Data ONTAP enables you to trigger an SP update manually and specify how the update should take place by using the `system service-processor image update` command.

You can specify the following options:

- The SP firmware package to use (`-package`)

You can update the SP firmware to a downloaded package by specifying the package file name. The `system image package show` command displays all package files (including the files for the SP firmware package) that are available on a node.
- Whether to use the baseline SP firmware package for the SP update (`-baseline`)

You can update the SP firmware to the baseline version that is bundled with the currently running version of Data ONTAP.

- Whether to update the entire firmware image or only the changed portions, using the network interface or the serial interface (**-update-type**)
- If updating the entire firmware image, whether to also reset log settings to the factory defaults and clear contents of all logs maintained by the SP, including the event logs, IPMI logs, and forensics logs (**-clear-logs**)
- Data ONTAP enables you to display the status for the latest SP firmware update triggered from Data ONTAP by using the `system service-processor image update-progress show` command.

Any existing connection to the SP is terminated when the SP firmware is being updated. This is the case whether the SP firmware update is automatically or manually triggered.

Related information

[NetApp Downloads: System Firmware and Diagnostics](#)

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

When the SP uses the network interface for firmware updates

An SP firmware update that is triggered from Data ONTAP with the SP running version 1.5, 2.5, 3.1, or later supports using an IP-based file transfer mechanism over the SP network interface.

An SP firmware update over the network interface is faster than an update over the serial interface. It reduces the maintenance window during which the SP firmware is being updated, and it is also nondisruptive to Data ONTAP operation. The SP versions that support this capability are included with Data ONTAP 8.3 and later. They are also available on the NetApp Support Site and can be installed on controllers that are running a compatible version of Data ONTAP.

When you are running SP version 1.5, 2.5, 3.1, or later, the following firmware upgrade behaviors apply:

- An SP firmware update that is *automatically* triggered by Data ONTAP defaults to using the network interface for the update; however, the SP automatic update switches to using the serial interface for the firmware update if one of the following conditions occurs:
 - The SP network interface is not configured or not available.
 - The IP-based file transfer fails.
 - The SP API service is disabled.
- An SP firmware update that is *manually* triggered from ONTAP (`system service-processor image update`) uses the network interface for the firmware update only when the `-update-type` parameter is set to **network-full**. Otherwise, the serial interface is used for the firmware update.

Regardless of the SP version you are running, an SP firmware update triggered from the SP CLI always uses the SP network interface for the update.

Related information

[NetApp Downloads: System Firmware and Diagnostics](#)

Accessing the SP

A cluster user account that is configured with the **service-processor** application type can access the SP from an administration host (using an SSH connection) or the serial console. Data ONTAP

enables you to restrict SSH access of the SP to only the administration hosts with specific IP addresses.

The `system timeout` commands manage the settings for the automatic timeout period of CLI sessions, including console and SSH connections to the SP.

Related concepts

Managing the automatic timeout period of CLI sessions on page 27

Accounts that can access the SP

When you try to access the SP, you are prompted for credential. Cluster user accounts that are created with the `service-processor` application type have access to the SP CLI on any node of the cluster. SP user accounts are managed from Data ONTAP and authenticated by password.

User accounts for accessing the SP are managed from Data ONTAP instead of the SP CLI. A cluster user account of any role can access the SP if it is created with the `-application` parameter of the `security login create` command set to `service-processor` and the `-authmethod` parameter set to `password`. The SP supports only password authentication.

By default, the cluster user account named “admin” includes the `service-processor` application type and has access to the SP.

Data ONTAP prevents you from creating user accounts with names that are reserved for the system (such as “root” and “naroot”). You cannot use a system-reserved name to access the cluster or the SP.

You can display current SP user accounts by using the `-application service-processor` parameter of the `security login show` command.

Accessing the SP from an administration host

You can log in to the SP of a node from an administration host to perform node management tasks remotely.

Before you begin

The following conditions must be met:

- The administration host you use to access the SP must support SSHv2.
- Your user account must already be set up for accessing the SP.
To access the SP, your user account must have been created with the `-application` parameter of the `security login create` command set to `service-processor` and the `-authmethod` parameter set to `password`.

About this task

If the SP is configured to use an IPv4 or IPv6 address, and if five SSH login attempts from a host fail consecutively within 10 minutes, the SP rejects SSH login requests and suspends the communication with the IP address of the host for 15 minutes. The communication resumes after 15 minutes, and you can try to log in to the SP again.

Data ONTAP prevents you from creating or using system-reserved names (such as “root” and “naroot”) to access the cluster or the SP.

Steps

1. From the administration host, log in to the SP:
`ssh username@SP_IP_address`
2. When you are prompted, enter the password for `username`.

The SP prompt appears, indicating that you have access to the SP CLI.

Examples of SP access from an administration host

The following example shows how to log in to the SP with a user account `joe`, which has been set up to access the SP.

```
[admin_host]$ ssh joe@192.168.123.98
joe@192.168.123.98's password:
SP>
```

The following examples show how to use the IPv6 global address or IPv6 router-advertised address to log in to the SP on a node that has SSH set up for IPv6 and the SP configured for IPv6.

```
[admin_host]$ ssh joe@fd22:8b1e:b255:202::1234
joe@fd22:8b1e:b255:202::1234's password:
SP>
```

```
[admin_host]$ ssh joe@fd22:8b1e:b255:202:2a0:98ff:fe01:7d5b
joe@fd22:8b1e:b255:202:2a0:98ff:fe01:7d5b's password:
SP>
```

Related concepts

[Relationship among the SP CLI, SP console, and system console sessions](#) on page 60

Related tasks

[Accessing the system console of a node](#) on page 42

Accessing the SP from the system console

You can access the SP from the system console (also called *serial console*) to perform monitoring or troubleshooting tasks.

Steps

1. Access the SP CLI from the system console by pressing Ctrl-G at the prompt.
2. Log in to the SP CLI when you are prompted.
The SP prompt appears, indicating that you have access to the SP CLI.
3. Exit the SP CLI and return to the system console by pressing Ctrl-D or entering `exit`, and then press Enter.

Example of accessing the SP CLI from the system console

The following example shows the result of pressing Ctrl-G from the system console to access the SP CLI. The `help system power` command is entered at the SP prompt, followed by pressing Ctrl-D and then Enter to return to the system console.

```
cluster1::>
```

(Press Ctrl-G to access the SP CLI.)

```
Switching console to Service Processor
Service Processor Login:
Password:
SP>
SP> help system power
system power cycle - power the system off, then on
system power off - power the system off
system power on - power the system on
system power status - print system power status
SP>
```

(Press Ctrl-D and then Enter to return to the system console.)

```
cluster1::>
```

Relationship among the SP CLI, SP console, and system console sessions

You can open an SP CLI session to manage a node remotely and open a separate SP console session to access the console of the node. The SP console session mirrors output displayed in a concurrent system console session. The SP and the system console have independent shell environments with independent login authentication.

Understanding how the SP CLI, SP console, and system console sessions are related helps you manage a node remotely. The following describes the relationship among the sessions:

- Only one administrator can log in to the SP CLI session at a time; however, the SP enables you to open both an SP CLI session and a separate SP console session simultaneously.

The SP CLI is indicated with the SP prompt (`SP>`). From an SP CLI session, you can use the `SP system console` command to initiate an SP console session. At the same time, you can start a separate SP CLI session through SSH. If you press Ctrl-D to exit from the SP console session, you automatically return to the SP CLI session. If an SP CLI session already exists, a message asks you whether to terminate the existing SP CLI session. If you enter “y”, the existing SP CLI session is terminated, enabling you to return from the SP console to the SP CLI. This action is recorded in the SP event log.

In a Data ONTAP CLI session that is connected through SSH, you can switch to the system console of a node by running the Data ONTAP `system node run-console` command from another node.

- For security reasons, the SP CLI session and the system console session have independent login authentication.

When you initiate an SP console session from the SP CLI (by using the `SP system console` command), you are prompted for the system console credential. When you access the SP CLI from a system console session (by pressing Ctrl-G), you are prompted for the SP CLI credential.
- The SP console session and the system console session have independent shell environments.

The SP console session mirrors output that is displayed in a concurrent system console session. However, the concurrent system console session does not mirror the SP console session. The SP console session does not mirror output of concurrent SSH sessions.

Managing the IP addresses that can access the SP

By default, the SP accepts SSH connection requests from administration hosts of any IP addresses. You can configure the SP to accept SSH connection requests from only the administration hosts that

have the IP addresses you specify. The changes you make apply to SSH access to the SP of any nodes in the cluster.

Steps

1. Grant SP access to only the IP addresses you specify by using the `system service-processor ssh add-allowed-addresses` command with the `-allowed-addresses` parameter.
 - The value of the `-allowed-addresses` parameter must be specified in the format of `address/netmask`, and multiple `address/netmask` pairs must be separated by commas, for example, `10.98.150.10/24, fd20:8b1e:b255:c09b::/64`.
Setting the `-allowed-addresses` parameter to `0.0.0.0/0, ::/0` enables all IP addresses to access the SP (the default).
 - When you change the default by limiting SP access to only the IP addresses you specify, Data ONTAP prompts you to confirm that you want the specified IP addresses to replace the “allow all” default setting (`0.0.0.0/0, ::/0`).
 - The `system service-processor ssh show` command displays the IP addresses that can access the SP.
2. If you want to block a specified IP address from accessing the SP, use the `system service-processor ssh remove-allowed-addresses` command with the `-allowed-addresses` parameter.

If you block all IP addresses from accessing the SP, the SP becomes inaccessible from any administration hosts.

Examples of managing the IP addresses that can access the SP

The following examples show the default setting for SSH access to the SP, change the default by limiting SP access to only the specified IP addresses, remove the specified IP addresses from the access list, and then restore SP access for all IP addresses:

```
cluster1::> system service-processor ssh show
  Allowed Addresses: 0.0.0.0/0, ::/0

cluster1::> system service-processor ssh add-allowed-addresses -allowed-addresses
192.168.1.202/24, 192.168.10.201/24

Warning: The default "allow all" setting (0.0.0.0/0, ::/0) will be replaced
with your changes. Do you want to continue? {y|n}: y

cluster1::> system service-processor ssh show
  Allowed Addresses: 192.168.1.202/24, 192.168.10.201/24

cluster1::> system service-processor ssh remove-allowed-addresses -allowed-addresses
192.168.1.202/24, 192.168.10.201/24

Warning: If all IP addresses are removed from the allowed address list, all IP
addresses will be denied access. To restore the "allow all" default,
use the "system service-processor ssh add-allowed-addresses
-allowed-addresses 0.0.0.0/0, ::/0" command. Do you want to continue?
{y|n}: y

cluster1::> system service-processor ssh show
  Allowed Addresses: -

cluster1::> system service-processor ssh add-allowed-addresses -allowed-addresses
0.0.0.0/0, ::/0

cluster1::> system service-processor ssh show
  Allowed Addresses: 0.0.0.0/0, ::/0
```

Using online help at the SP CLI

The SP online help displays the SP CLI commands and options when you enter the question mark (?) or `help` at the SP prompt.

Steps

1. To display help information for the SP commands, enter one of the following at the SP prompt:

- `help`
- `?`

Example

The following example shows the SP CLI online help:

```
SP> help
date - print date and time
exit - exit from the SP command line interface
events - print system events and event information
help - print command help
priv - show and set user mode
sp - commands to control the SP
system - commands to control the system
version - print SP version
```

2. To display help information for the option of an SP command, enter the following command at the SP prompt:

```
help SP_command
```

Example

The following example shows the SP CLI online help for the SP `events` command:

```
SP> help events
events all - print all system events
events info - print system event log information
events newest - print newest system events
events oldest - print oldest system events
events search - search for and print system events
```

Commands for managing a node remotely

You can manage a node remotely by accessing its SP and running SP CLI commands to perform node-management tasks. For several commonly performed remote node-management tasks, you can also use Data ONTAP commands from another node in the cluster. Some SP commands are platform-specific and might not be available on your platform.

If you want to...	Use this SP command...	Or this Data ONTAP command ...
Display available SP commands or subcommands of a specified SP command	<code>help [command]</code>	
Display the current privilege level for the SP CLI	<code>priv show</code>	

If you want to...	Use this SP command...	Or this Data ONTAP command ...
Set the privilege level to access the specified mode for the SP CLI	<code>priv set {admin advanced diag}</code>	
Display system date and time	<code>date</code>	<code>date</code>
Display events that are logged by the SP	<code>events {all info newest <i>number</i> oldest <i>number</i> search <i>keyword</i>}</code>	
Display SP status and network configuration information	<code>sp status [-v -d]</code> The <code>-v</code> option displays SP statistics in verbose form. The <code>-d</code> option adds the SP debug log to the display.	<code><i>system service-processor show</i></code>
Display the length of time the SP has been up and the average number of jobs in the run queue over the last 1, 5, and 15 minutes	<code>sp uptime</code>	
Display system console logs	<code>system log</code>	
Display the SP log archives or the files in an archive	<code>sp log history show [-archive {latest all <i>archive-name</i>}] [-dump {all <i>file-name</i>}]</code>	
Display the power status for the controller of a node	<code>system power status</code>	<code><i>system node power show</i></code>
Display battery information	<code>system battery show</code>	
Display ACP information or the status for expander sensors	<code>system acp [show sensors show]</code>	
List all system FRUs and their IDs	<code>system fru list</code>	
Display product information for the specified FRU	<code>system fru show <i>fru_id</i></code>	
Display the FRU data history log	<code>system fru log show</code> (advanced privilege level)	
Display the status for the environmental sensors, including their states and current values	<code>system sensors</code> or <code>system sensors show</code>	<code><i>system node environment sensors show</i></code>
Display the status and details for the specified sensor	<code>system sensors get <i>sensor_name</i></code> You can obtain <code><i>sensor_name</i></code> by using the <code>system sensors</code> or the <code>system sensors show</code> command.	

If you want to...	Use this SP command...	Or this Data ONTAP command ...
Display the SP firmware version information	version	<code>system service-processor image show</code>
Display the SP command history	sp log audit (advanced privilege level)	
Display the SP debug information	sp log debug (advanced privilege level)	
Display the SP messages file	sp log messages (advanced privilege level)	
Display the settings for collecting system forensics on a watchdog reset event, display system forensics information collected during a watchdog reset event, or clear the collected system forensics information	system forensics [show log dump log clear]	
Log in to the system console	system console	system node run-console
	You use Ctrl-D to exit the system console session.	
Turn the node on or off, or perform a power-cycle (turning the power off and then back on)	system power on	<code>system node power on</code> (advanced privilege level)
	system power off	
	system power cycle	
	The standby power stays on to keep the SP running without interruption. During the power-cycle, a brief pause occurs before power is turned back on. Attention: Using these commands to turn off or power-cycle the node might cause an improper shutdown of the node (also called a <i>dirty shutdown</i>) and is not a substitute for a graceful shutdown using the Data ONTAP <code>system node halt</code> command.	
Create a core dump and reset the node	system core [-f] The -f option forces the creation of a core dump and the reset of the node.	<code>system node coredump trigger</code> (advanced privilege level)
	These commands have the same effect as pressing the Non-maskable Interrupt (NMI) button on a node, causing a dirty shutdown of the node and forcing a dump of the core files when halting the node. These commands are helpful when Data ONTAP on the node is hung or does not respond to commands such as <code>system node shutdown</code> . The generated core dump files are displayed in the output of the <code>system node coredump show</code> command. The SP stays operational as long as the input power to the node is not interrupted.	

If you want to...	Use this SP command...	Or this Data ONTAP command ...
Reboot the node with an optionally specified BIOS firmware image (primary, backup, or current) to recover from issues such as a corrupted image of the node's boot device	<code>system reset {primary backup current}</code>	<code>system node reset</code> with the <code>-firmware {primary backup current}</code> parameter (advanced privilege level)
	<p>Attention: This operation causes a dirty shutdown of the node. If no BIOS firmware image is specified, the current image is used for the reboot. The SP stays operational as long as the input power to the node is not interrupted.</p>	
Display the status of battery firmware automatic update, or enable or disable battery firmware automatic update upon next SP boot	<code>system battery auto_update [status enable disable]</code> (advanced privilege level)	
Compare the current battery firmware image against a specified firmware image	<code>system battery verify [image_URL]</code> (advanced privilege level) If <code>image_URL</code> is not specified, the default battery firmware image is used for comparison.	
Update the battery firmware from the image at the specified location	<code>system battery flash image_URL</code> (advanced privilege level) You use this command if the automatic battery firmware upgrade process has failed for some reason.	
Update the SP firmware by using the image at the specified location	<code>sp update image_URL</code> <code>image_URL</code> must not exceed 200 characters.	<code>system service-processor image update</code>
Reboot the SP	<code>sp reboot</code>	<code>system service-processor reboot-sp</code>
	<p>Attention: You should avoid booting the SP from the backup image. Booting from the backup image is reserved for troubleshooting and recovery purposes only. It might require that the SP automatic firmware update be disabled, which is not a recommended setting. You should contact technical support before attempting to boot the SP from the backup image.</p>	
Erase the NVRAM flash content	<code>system nvram flash clear</code> (advanced privilege level) This command cannot be initiated when the controller power is off (<code>system power off</code>).	

If you want to...	Use this SP command...	Or this Data ONTAP command ...
Exit the SP CLI	<code>exit</code>	

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Understanding the threshold-based SP sensor readings and status values of the system sensors command output

Threshold-based sensors take periodic readings of a variety of system components. The SP compares the reading of a threshold-based sensor against its preset threshold limits that define a component's acceptable operating conditions. Based on the sensor reading, the SP displays the sensor state to help you monitor the condition of the component.

Examples of threshold-based sensors include sensors for the system temperatures, voltages, currents, and fan speeds. The specific list of threshold-based sensors depends on the platform.

Threshold-based sensors have the following thresholds, displayed in the output of the `SP system sensors` command:

- Lower critical (LCR)
- Lower noncritical (LNC)
- Upper noncritical (UNC)
- Upper critical (UCR)

A sensor reading between LNC and LCR or between UNC and UCR means that the component is showing signs of a problem and a system failure might occur as a result. Therefore, you should plan for component service soon.

A sensor reading below LCR or above UCR means that the component is malfunctioning and a system failure is about to occur. Therefore, the component requires immediate attention.

The following diagram illustrates the severity ranges that are specified by the thresholds:



You can find the reading of a threshold-based sensor under the `Current` column in the `system sensors` command output. The `system sensors get sensor_name` command displays additional details for the specified sensor. As the reading of a threshold-based sensor crosses the noncritical and critical threshold ranges, the sensor reports a problem of increasing severity. When the reading exceeds a threshold limit, the sensor's status in the `system sensors` command output changes from `ok` to `nc` (noncritical) or `cr` (critical) depending on the exceeded threshold, and an event message is logged in the SEL event log.

Some threshold-based sensors do not have all four threshold levels. For those sensors, the missing thresholds show `na` as their limits in the `system sensors` command output, indicating that the particular sensor has no limit or severity concern for the given threshold and the SP does not monitor the sensor for that threshold.

Example of the `system sensors` command output

The following example shows the information displayed by the `system sensors` command in the SP CLI:

```
SP node1> system sensors
```

Sensor Name	Current	Unit	Status	LCR	LNC	UNC	UCR
CPU0_Temp_Margin	-55.000	degrees C	ok	na	na	-5.000	0.000
CPU1_Temp_Margin	-56.000	degrees C	ok	na	na	-5.000	0.000
In_Flow_Temp	32.000	degrees C	ok	0.000	10.000	42.000	52.000
Out_Flow_Temp	38.000	degrees C	ok	0.000	10.000	59.000	68.000
PCI_Slot_Temp	40.000	degrees C	ok	0.000	10.000	56.000	65.000
NVMEM_Bat_Temp	32.000	degrees C	ok	0.000	10.000	55.000	64.000
LM56_Temp	38.000	degrees C	ok	na	na	49.000	58.000
CPU0_Error	0x0	discrete	0x0180	na	na	na	na
CPU0_Therm_Trip	0x0	discrete	0x0180	na	na	na	na
CPU0_Hot	0x0	discrete	0x0180	na	na	na	na
CPU1_Error	0x0	discrete	0x0180	na	na	na	na
CPU1_Therm_Trip	0x0	discrete	0x0180	na	na	na	na
CPU1_Hot	0x0	discrete	0x0180	na	na	na	na
IO_Mid1_Temp	30.000	degrees C	ok	0.000	10.000	55.000	64.000
IO_Mid2_Temp	30.000	degrees C	ok	0.000	10.000	55.000	64.000
CPU_VTT	1.106	Volts	ok	1.028	1.048	1.154	1.174
CPU0_VCC	1.154	Volts	ok	0.834	0.844	1.348	1.368
CPU1_VCC	1.086	Volts	ok	0.834	0.844	1.348	1.368
1.0V	0.989	Volts	ok	0.941	0.951	1.057	1.067
1.05V	1.048	Volts	ok	0.980	0.999	1.106	1.125
1.1V	1.096	Volts	ok	1.028	1.038	1.154	1.174
1.2V	1.203	Volts	ok	1.125	1.135	1.261	1.280
1.5V	1.513	Volts	ok	1.436	1.455	1.571	1.591
1.8V	1.754	Volts	ok	1.664	1.703	1.896	1.935
2.5V	2.543	Volts	ok	2.309	2.356	2.621	2.699
3.3V	3.323	Volts	ok	3.053	3.116	3.466	3.546
5V	5.002	Volts	ok	4.368	4.465	5.490	5.636
STBY_1.8V	1.794	Volts	ok	1.678	1.707	1.892	1.911
...							

Example of the `system sensors get sensor_name` command output for a threshold-based sensor

The following example shows the result of entering `system sensors get sensor_name` in the SP CLI for the threshold-based sensor 5V:

```
SP node1> system sensors get 5V
```

```
Locating sensor record...
Sensor ID       : 5V (0x13)
Entity ID      : 7.97
Sensor Type (Analog) : Voltage
Sensor Reading  : 5.002 (+/- 0) Volts
Status         : ok
Lower Non-Recoverable : na
Lower Critical   : 4.246
Lower Non-Critical : 4.490
Upper Non-Critical : 5.490
Upper Critical   : 5.758
Upper Non-Recoverable : na
Assertion Events  :
Assertions Enabled : lnc- lcr- ucr+
Deassertions Enabled : lnc- lcr- ucr+
```

Understanding the discrete SP sensor status values of the `system sensors` command output

Discrete sensors do not have thresholds. Their readings, displayed under the `Current` column in the SP CLI `system sensors` command output, do not carry actual meanings and thus are ignored by the SP. The `Status` column in the `system sensors` command output displays the status values of discrete sensors in hexadecimal format.

Examples of discrete sensors include sensors for the fan, power supply unit (PSU) fault, and system fault. The specific list of discrete sensors depends on the platform.

You can use the SP CLI `system sensors get sensor_name` command for help with interpreting the status values for most discrete sensors. The following examples show the results of entering `system sensors get sensor_name` for the discrete sensors `CPU0_Error` and `IO_Slot1_Present`:

```
SP node1> system sensors get CPU0_Error
Locating sensor record...
Sensor ID           : CPU0_Error (0x67)
Entity ID           : 7.97
Sensor Type (Discrete): Temperature
States Asserted     : Digital State
                    [State Deasserted]
```

```
SP node1> system sensors get IO_Slot1_Present
Locating sensor record...
Sensor ID           : IO_Slot1_Present (0x74)
Entity ID           : 11.97
Sensor Type (Discrete): Add-in Card
States Asserted     : Availability State
                    [Device Present]
```

Although the `system sensors get sensor_name` command displays the status information for most discrete sensors, it does not provide status information for the `System_FW_Status`, `System_Watchdog`, `PSU1_Input_Type`, and `PSU2_Input_Type` discrete sensors. You can use the following information to interpret these sensors' status values.

System_FW_Status

The `System_FW_Status` sensor's condition appears in the form of `0xAABB`. You can combine the information of `AA` and `BB` to determine the condition of the sensor.

`AA` can have one of the following values:

- 01** System firmware error
- 02** System firmware hang
- 04** System firmware progress

`BB` can have one of the following values:

- 00** System software has properly shut down
- 01** Memory initialization in progress
- 02** NVMEM initialization in progress (when NVMEM is present)
- 04** Restoring memory controller hub (MCH) values (when NVMEM is present)
- 05** User has entered Setup
- 13** Booting the operating system or LOADER
- 1F**

	BIOS is starting up
20	LOADER is running
21	LOADER is programming the primary BIOS firmware. You must not power down the system.
22	LOADER is programming the alternate BIOS firmware. You must not power down the system.
2F	Data ONTAP is running
60	SP has powered off the system
61	SP has powered on the system
62	SP has reset the system
63	SP watchdog power cycle
64	SP watchdog cold reset

For instance, the System_FW_Status sensor status 0x042F means "system firmware progress (04), Data ONTAP is running (2F)."

System_Watchdog

The System_Watchdog sensor can have one of the following conditions:

0x0080

The state of this sensor has not changed

0x0081

Timer interrupt

0x0180

Timer expired

0x0280

Hard reset

0x0480

Power down

0x0880

Power cycle

For instance, the System_Watchdog sensor status 0x0880 means a watchdog timeout occurs and causes a system power cycle.

PSU1_Input_Type and PSU2_Input_Type

For direct current (DC) power supplies, the PSU1_Input_Type and PSU2_Input_Type sensors do not apply. For alternating current (AC) power supplies, the sensors' status can have one of the following values:

0x01xx

220V PSU type

0x02xx

110V PSU type

For instance, the PSU1_Input_Type sensor status 0x0280 means that the sensor reports that the PSU type is 110V.

Commands for managing the SP from Data ONTAP

Data ONTAP provides commands for managing the SP, including the SP network configuration, SP firmware image, SSH access to the SP, and general SP administration.

Commands for managing the SP network configuration

If you want to...	Use this Data ONTAP command...
Enable the SP automatic network configuration for the SP to use the IPv4 or IPv6 address family of the specified subnet	<code>system service-processor network auto-configuration enable</code>
Disable the SP automatic network configuration for the IPv4 or IPv6 address family of the subnet specified for the SP	<code>system service-processor network auto-configuration disable</code>
Display the SP automatic network configuration	<code>system service-processor network auto-configuration show</code>
Manually configure the SP network for a node, including the following: <ul style="list-style-type: none"> The IP address family (IPv4 or IPv6) Whether the network interface of the specified IP address family should be enabled If you are using IPv4, whether to use the network configuration from the DHCP server or the network address that you specify The public IP address for the SP The netmask for the SP (if using IPv4) The network prefix-length of the subnet mask for the SP (if using IPv6) The gateway IP address for the SP 	<code>system service-processor network modify</code>

If you want to...	Use this Data ONTAP command...
<p>Display the SP network configuration, including the following:</p> <ul style="list-style-type: none"> • The configured address family (IPv4 or IPv6) and whether it is enabled • The remote management device type • The current SP status and link status • Network configuration, such as IP address, MAC address, netmask, prefix-length of subnet mask, router-assigned IP address, link local IP address, and gateway IP address • The time the SP was last updated • The name of the subnet used for SP automatic configuration • Whether the IPv6 router-assigned IP address is enabled • SP network setup status • Reason for the SP network setup failure 	<p><i>system service-processor network show</i></p> <p>Displaying complete SP network details requires the <code>-instance</code> parameter.</p>
<p>Modify the SP API service configuration, including the following:</p> <ul style="list-style-type: none"> • Changing the port used by the SP API service • Enabling or disabling the SP API service 	<p><i>system service-processor api-service modify</i></p> <p>(advanced privilege level)</p>
<p>Display the SP API service configuration</p>	<p><i>system service-processor api-service show</i></p> <p>(advanced privilege level)</p>
<p>Renew the SSL and SSH certificates used by the SP API service for internal communication</p>	<p><i>system service-processor api-service renew-certificates</i></p> <p>(advanced privilege level)</p>

Commands for managing the SP firmware image

If you want to...	Use this Data ONTAP command...
<p>Display the details of the currently installed SP firmware image, including the following:</p> <ul style="list-style-type: none"> • The remote management device type • The image (primary or backup) that the SP is booted from, its status, and firmware version • Whether the firmware automatic update is enabled and the last update status 	<p><i>system service-processor image show</i></p> <p>The <code>-is-current</code> parameter indicates the image (primary or backup) that the SP is currently booted from, not whether the installed firmware version is most current.</p>

If you want to...	Use this Data ONTAP command...
<p>Enable or disable the SP automatic firmware update</p>	<p><i>system service-processor image modify</i></p> <p>By default, the SP firmware is automatically updated with the update of Data ONTAP or when a new version of the SP firmware is manually downloaded. Disabling the automatic update is not recommended because doing so can result in suboptimal or nonqualified combinations between the Data ONTAP image and the SP firmware image.</p>
<p>Manually download an SP firmware image on a node</p>	<p><i>system node image get</i></p> <p>The SP firmware image is packaged with Data ONTAP. You do not need to download the SP firmware manually, unless you want to use an SP firmware version that is different from the one packaged with Data ONTAP.</p>
<p>Manually update the SP firmware, by specifying the following:</p> <ul style="list-style-type: none"> • The SP firmware package to use You can have the SP use a specific SP firmware package by specifying the package file name. The <code>system node image package show</code> command displays all package files (including the files for the SP firmware package) that are available on a node. • The installation baseline You can update the SP firmware to the baseline version that is bundled with the currently running version of Data ONTAP. • Whether to update the entire firmware image or only the changed portions, using the network interface or the serial interface • If updating the entire firmware image, whether to also reset log settings to the factory default and clear contents of all logs maintained by the SP, including the event logs, IPMI logs, and forensics logs 	<p><i>system service-processor image update</i></p>
<p>Display the status for the latest SP firmware update triggered from Data ONTAP, including the following information:</p> <ul style="list-style-type: none"> • The start and end time for the latest SP firmware update • Whether an update is in progress and the percentage that is complete 	<p><i>system service-processor image update-progress show</i></p>

Commands for managing SSH access to the SP

If you want to...	Use this Data ONTAP command...
Grant SP access to only the specified IP addresses	<code>system service-processor ssh add-allowed-addresses</code>
Block the specified IP addresses from accessing the SP	<code>system service-processor ssh remove-allowed-addresses</code>
Display the IP addresses that can access the SP	<code>system service-processor ssh show</code>

Commands for general SP administration

If you want to...	Use this Data ONTAP command...
Display general SP information, including the following: <ul style="list-style-type: none"> • The remote management device type • The current SP status • Whether the SP network is configured • Network information, such as the public IP address and the MAC address • The SP firmware version and Intelligent Platform Management Interface (IPMI) version • Whether the SP firmware automatic update is enabled 	<code>system service-processor show</code> Displaying complete SP information requires the <code>-instance</code> parameter.
Reboot the SP on a node and optionally specify the SP firmware image (primary or backup) to use	<code>system service-processor reboot-sp</code> Attention: You should avoid booting the SP from the backup image. Booting from the backup image is reserved for troubleshooting and recovery purposes only. It might require that the SP automatic firmware update be disabled, which is not a recommended setting. You should contact Technical Support before attempting to boot the SP from the backup image.
Generate and send an AutoSupport message that includes the SP log files collected from a specified node	<code>system node autosupport invoke-splog</code>
Display the allocation map of the collected SP log files in the cluster, including the sequence numbers for the SP log files that reside in each collecting node	<code>system service-processor log show-allocations</code>

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Managing SVMs (cluster administrators only)

Cluster administrators can manage and administer the Storage Virtual Machines (SVMs, formerly known as Vservers) within a cluster. A cluster must have at least one SVM to serve data to the clients. Therefore, a cluster administrator must create and manage SVMs.

Cluster administrators can either choose to perform SVM administration tasks in addition to the SVM management tasks or delegate the administration of the SVMs to SVM administrators.

To manage and administer SVMs, you must understand what an SVM is, its benefits such as nondisruptive operation and scalability, and the associated management tasks.

A cluster administrator can perform the following SVM management tasks:

- Creating SVMs
- Modifying SVMs
- Deleting SVMs
- Renaming SVMs
- Administering SVMs from the SVM context
- Starting and stopping SVMs

Note: Both cluster administrators and SVM administrators can view information about SVMs.

Note: The Data ONTAP command-line interface (CLI) continues to use the term *Vserver* in the output, and `vserver` as a command or parameter name has not changed.

Related concepts

[Administering SVMs](#) on page 109

What SVMs are

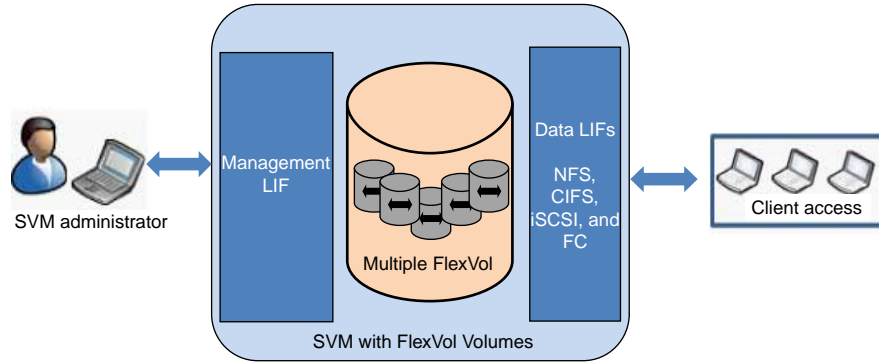
Storage Virtual Machines (SVMs, formerly known as Vservers) contain data volumes and one or more LIFs through which they serve data to the clients. Starting with clustered Data ONTAP 8.1.1, SVMs can either contain one or more FlexVol volumes, or a single Infinite Volume.

SVMs securely isolate the shared virtualized data storage and network, and each SVM appears as a single dedicated server to the clients. Each SVM has a separate administrator authentication domain and can be managed independently by its SVM administrator.

In a cluster, SVMs facilitate data access. A cluster must have at least one SVM to serve data. SVMs use the storage and network resources of the cluster. However, the volumes and LIFs are exclusive to the SVM. Multiple SVMs can coexist in a single cluster without being bound to any node in a cluster. However, they are bound to the physical cluster on which they exist.

A cluster can have one or more SVMs with FlexVol volumes and SVMs with Infinite Volume.

SVM with FlexVol volumes

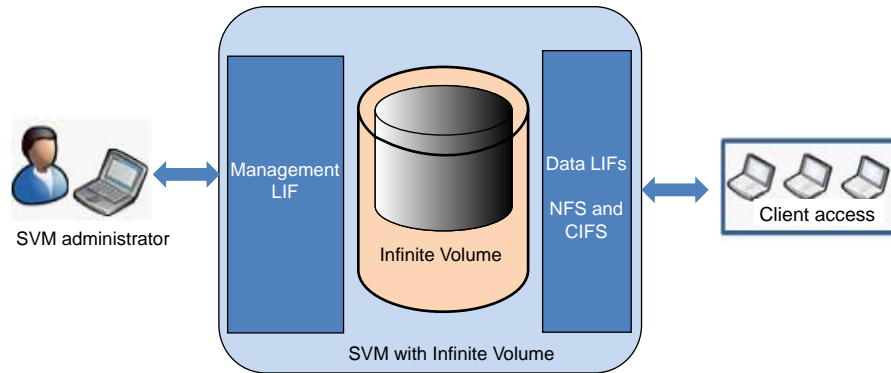


Each SVM with FlexVol volumes in a NAS environment presents a single directory hierarchical view and has a unique namespace. The namespace enables NAS clients to access data without specifying the physical location of the data. The namespace also enables the cluster and SVM administrators to manage distributed data storage as a single directory with multiple levels of hierarchy.

The volumes within each NAS SVM are related to each other through junctions and are mounted on junction paths. These junctions present the file system in each volume. The root volume of the SVM is a FlexVol volume that resides at the top level of the namespace hierarchy; additional volumes are mounted to the SVM root volume to extend the namespace. As volumes are created for the SVM, the root volume of the SVM contains junction paths.

SVMs with FlexVol volumes can contain files and LUNs. They provide file-level data access by using NFS and CIFS protocols for the NAS clients, and block-level data access by using iSCSI and Fibre Channel (FC) (FCoE included) for SAN hosts.

SVM with Infinite Volume



SVMs with Infinite Volume can contain only one Infinite Volume to serve data. Each SVM with Infinite Volume includes only one junction path, which has a default value of `/NS`. The junction provides a single mount point for the large namespace provided by the SVM with Infinite Volume. You cannot add more junctions to an SVM with Infinite Volume. However, you can increase the size of the Infinite Volume.

SVMs with Infinite Volume can contain only files. They provide file-level data access by using NFS and CIFS protocols. SVMs with Infinite Volume cannot contain LUNs and do not provide block-level data access.

Note: The Data ONTAP command-line interface (CLI) continues to use the term `Vserver` in the output, and `vserver` as a command or parameter name has not changed.

How SVMs root volumes are used for data access

Every Storage Virtual Machine (SVM) has a root volume that contains the paths where the data volumes are junctioned into the namespace. NAS clients' data access is dependent on the health of the root volume in the namespace and SAN clients' data access is independent of the root volume's health in the namespace.

The root volume serves as the entry point to the namespace provided by that SVM. The root volume of the SVM is a FlexVol volume that resides at the top level of the namespace hierarchy and contains the directories that are used as mount points, the paths where data volumes are junctioned into the namespace.

In the unlikely event that the root volume of an SVM namespace is unavailable, NAS clients cannot access the namespace hierarchy and therefore cannot access data in the namespace. For this reason, it is best to create a load-sharing mirror copy for the root volume on each node of the cluster so that the namespace directory information remains available in the event of a node outage or failover.

You should not store user data in the root volume of an SVM. The root volume of the SVM should be used for junction paths, and user data should be stored in non-root volumes of the SVM.

Types of SVMs

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

A cluster contains the following types of SVMs:

- **Admin SVM**
The cluster setup process automatically creates the admin SVM for the cluster. The admin SVM represents the cluster.
- **Node SVM**
A node SVM is created when the node joins the cluster, and the node SVM represents the individual nodes of the cluster.
- **System SVM (advanced)**
A system SVM is automatically created for cluster-level communications in an IPspace.
- **Data SVM**
A data SVM represents the data serving SVMs. After the cluster setup, a cluster administrator must create data SVMs and add volumes to these SVMs to facilitate data access from the cluster. A cluster must have at least one data SVM to serve data to its clients.

Note: Unless otherwise specified, the term SVM refers to data (data-serving) SVM, which applies to both SVMs with FlexVol volumes and SVMs with Infinite Volume.

In the CLI, SVMs are displayed as Vservers.

Why you use SVMs

Storage Virtual Machines (SVMs, formerly known as Vservers) provide data access to clients regardless of the physical storage or controller, similar to any storage system. SVMs provide benefits such as nondisruptive operations, scalability, security, and unified storage.

SVMs provide the following benefits:

- **Multi-tenancy**
SVM is the fundamental unit of secure multi-tenancy, which enables partitioning of the storage infrastructure so that it appears as multiple independent storage systems. These partitions isolate the data and management.

- **Nondisruptive operations**
SVMs can operate continuously and nondisruptively for as long as they are needed. SVMs help clusters to operate continuously during software and hardware upgrades, addition and removal of nodes, and all administrative operations.
- **Scalability**
SVMs meet on-demand data throughput and the other storage requirements.
- **Security**
Each SVM appears as a single independent server, which enables multiple SVMs to coexist in a cluster while ensuring no data flows among them.
- **Unified storage**
SVMs can serve data concurrently through multiple data access protocols. SVMs provide file-level data access through NAS protocols, such as CIFS and NFS, and block-level data access through SAN protocols, such as iSCSI and FC (FCoE included). SVMs can serve data to SAN and NAS clients independently at the same time.

Note: SVMs with Infinite Volume can serve data only through NFS and CIFS protocols.
- **Delegation of management**
Each SVM can have its own user and administration authentication. SVM administrators can manage the SVMs that they are authorized to access. However, SVM administrators have privileges assigned by the cluster administrators.
- **Easy management of large datasets**
With SVMs with Infinite Volume, management of large and unstructured data is easier because the SVM administrator can manage one data container instead of many.

Number of SVMs in a cluster

The number of SVMs that you can create in a cluster depends on the number of nodes and how the LIFs are configured and used in your cluster.

The maximum number of nodes within a cluster depends on the platform model and licensed protocols. For details about the number of nodes in a cluster, see the *Hardware Universe* at hww.netapp.com.

The following table lists the recommended number of SVMs in a cluster based on the number of LIFs configured:

SVMs with protocol type Nodes	Nodes per cluster						SVM configuration
	1	2	4	6	8	10-24	
SVMs with NFS/CIFS protocol: one single LIF for data and management	125	250	500	750	1000	1000	Each SVM with one active IP LIF for data and management, and one IP LIF reserved for failover.
SVMs with FC/FCoE protocol: one LIF for data and one LIF for management	125	250	250	250	250	NA	Each SVM with two FC/FCoE LIFs on each node of the cluster and an IP LIF dedicated for management.

SVMs with protocol type Nodes	Nodes per cluster						SVM configuration
	1	2	4	6	8	10-24	
SVMs with iSCSI protocol: one LIF for data and one LIF for management	125	125	165	190	200	NA	Each SVM with one iSCSI LIF on each node of the cluster and an IP LIF dedicated for management.

You should not create SVMs with Infinite Volume if the cluster contains SVMs with SAN configuration.

SVMs with Infinite Volume cannot span more than 10 nodes of a NAS cluster.

Note: The number of SVMs in a cluster might not be same if the cluster has a combination of SVMs with different protocols.

Setting up SVMs

Different subtypes of SVMs provide data access, data protection against disasters, and data high availability. You must set up at least one data access SVM per cluster, which involves planning the setup, understanding requirements, completing the setup worksheet, and creating and configuring the SVM.

Planning to create SVMs

You must understand the various aspects of SVMs that help you plan the SVM setup process. You must also understand the purpose for creating the SVM and understand details such as the type of volume the SVM should contain, the environment for data access, and network segregation.

SVMs for data access

You must create one or more SVMs to serve data from the cluster. SVMs can contain either one or more FlexVol Volumes or a single Infinite Volume to serve data to the clients.

You can also use OnCommand System Manager to create and configure SVMs with FlexVol volumes and SVMs with Infinite Volume.

[Clustered Data ONTAP 8.3 Cluster Management Using OnCommand System Manager](#)

You must understand the following SVM attributes when creating SVMs with FlexVol volumes for data access:

SVM attributes	Possible values	Description
SVM subtype	default	Data SVMs that are created for data access must have the <code>subtype</code> option set to default .
Volumes	FlexVol volumes	The <code>-is-repository</code> option must be set to false , which is the default value.
Protocols	NFS, CIFS, iSCSI, and FC	Data SVMs can serve data to NAS clients and SAN hosts simultaneously.

SVM attributes	Possible values	Description
Root volume security style	<ul style="list-style-type: none"> mixed for NFS and CIFS ntfs for CIFS unix for NFS, iSCSI and FC 	Based on the clients, you can select the appropriate security style for the SVM root volume.
IPspace	Any available IPspace	IPspace defines a secured dedicated network path for each SVM. You can segregate the network for the SVM by assigning the IPspace when creating the SVM.

You must understand the following SVM attributes when creating SVMs with Infinite Volume for data access:

SVM attributes	Possible values	Description
SVM subtype	default	Data SVMs that are created for data access must have the <code>subtype</code> option set to default .
Volumes	Infinite Volume	The <code>-is-repository</code> option must be set to true . The default value is false .
Protocols	<ul style="list-style-type: none"> NFS NFSv3, pNFS, and NFSv4.1 are supported. CIFS SMB 1.0 is supported. 	Data SVMs can serve data to NAS clients only.
Root volume security style	<ul style="list-style-type: none"> mixed for NFS and CIFS ntfs for CIFS unix for NFS 	Based on the clients, you can select the appropriate security style for the SVM root volume.
IPspace	Any available IPspace	IPspace defines a secured dedicated network path for each SVM. You can segregate the network for the SVM by assigning the IPspace when creating the SVM.

SVMs for MetroCluster configuration

You can create source and destination Storage Virtual Machines (SVMs) for a MetroCluster configuration for providing synchronous disaster recovery and high availability of data.

Note: Only SVMs with FlexVol volumes are supported for a MetroCluster configuration.

You must understand the following SVM attributes when creating SVMs for the MetroCluster configuration:

SVM attributes	Possible values	Description
SVM subtype	sync-source and sync-destination	Source and destination SVMs that are created for a MetroCluster configuration have the <code>subtype</code> option set to sync-source and sync-destination , respectively.
Volumes	FlexVol volumes	You do not have to specify the <code>-is-repository</code> option because the default value is false .
Protocols	NFS, CIFS, iSCSI, and FC	Source and destination SVMs support all the protocols.
Root volume security style	<ul style="list-style-type: none"> • mixed for NFS and CIFS • ntfs for CIFS • unix for NFS, iSCSI, and FC 	The root volumes of the source and destination SVMs have an identical security style.
IPspace	Any available IPspace	<p>IPspace defines a secured dedicated network path for each SVM.</p> <p>You can segregate the network for the SVM by assigning the IPspace when creating the SVM.</p> <p>Source and destination SVMs must belong to IPspaces with the same name.</p>

Guidelines for creating SVMs

There are naming guidelines, and language and IPspace considerations that you should understand for creating an SVM successfully.

SVM naming guidelines

- SVM names must be unique across clusters.
You must use the fully qualified domain name (FQDN) of the SVM or another convention that ensures unique SVM names across clusters.
- SVM names can have a maximum of 47 characters.
However, SVMs in a MetroCluster configuration can have a maximum of only 41 characters.
- SVM names are case-sensitive and can contain alphanumeric characters.
SVM names can contain a period (.), a hyphen (-), or an underscore (_), but must not start with a hyphen, period, or number.

Language considerations

The default language setting is **C.UTF-8 (POSIX.UTF-8)**, which is inherited by all of its volumes. You can specify a different language when creating the volume.

When you modify the language of the SVM, the default language setting is modified; the language setting of the existing volumes is not modified.

You should append .UTF-8 for the language encoding values. For example, for the en_US language, the recommended format is en_US.UTF-8.

Note: You cannot modify the language of an SVM with Infinite Volume later.

IPspace considerations

You must assign an IPspace to an SVM when creating the SVM. You cannot modify or remove the IPspace for the SVM later.

List of language options

When you create Storage Virtual Machine (SVM), the language is set for the SVM. The language of the SVM determines the default language setting for volumes in that SVM. You can modify the language of an SVM.

You can specify the language for a volume when creating a volume and it can be different from the language of an SVM. If you do not specify the language for a volume then it inherits the language setting of its SVM. After the volume is created, you cannot modify the language of a volume. Therefore, you must be aware of the available language options.

The following table lists the various available language options that helps you choose and enter the correct value when creating an SVM or volume:

Language values	Languages
c	POSIX
C.UTF-8	POSIX with UTF-8
ar	Arabic
ar.UTF-8	Arabic with UTF-8
cs	Czech
cs.UTF-8	Czech with UTF-8
da	Danish
da.UTF-8	Danish with UTF-8
de	German
de.UTF-8	German with UTF-8
en	English
en.UTF-8	English with UTF-8
en_us	English (US)
en_US.UTF-8	US English with UTF-8
es	Spanish
es.UTF-8	Spanish with UTF-8
fi	Finnish
fi.UTF-8	Finnish with UTF-8
fr	French
fr.UTF-8	French with UTF-8
he	Hebrew
he.UTF-8	Hebrew with UTF-8

Language values	Languages
hr	Croatian
hr.UTF-8	Croatian with UTF-8
hu	Hungarian
hu.UTF-8	Hungarian with UTF-8
it	Italian
it.UTF-8	Italian with UTF-8
ja_v1	Japanese euc-j
ja_v1.UTF-8	Japanese euc-j with UTF-8
ja_jp.pck_v2	Japanese PCK (sjis)
ja_JP.PCK_v2.UTF-8	Japanese PCK sjis with UTF-8
ko	Korean
ko.UTF-8	Korean with UTF-8
no	Norwegian
no.UTF-8	Norwegian with UTF-8
nl	Dutch
nl.UTF-8	Dutch with UTF-8
pl	Polish
pl.UTF-8	Polish with UTF-8
pt	Portuguese
pt.UTF-8	Portuguese with UTF-8
ro	Romanian
ro.UTF-8	Romanian with UTF-8
ru	Russian
ru.UTF-8	Russian with UTF-8
sk	Slovak
sk.UTF-8	Slovak with UTF-8
sl	Slovenian
sl.UTF-8	Slovenian with UTF-8
sv	Swedish
sv.UTF-8	Swedish with UTF-8
tr	Turkish
tr.UTF-8	Turkish with UTF-8
zh	Simplified Chinese
zh.UTF-8	Simplified Chinese with UTF-8
zh.GBK	Simplified Chinese (GBK)

Language values	Languages
zh.GBK.UTF-8	Simplified GBK Chinese with UTF-8
zh_TW	Traditional Chinese euc-tw
zh_TW.UTF-8	Traditional Chinese euc-tw with UTF-8
zh_TW.BIG5	Traditional Chinese Big 5
zh_TW.BIG5.UTF-8	Traditional Chinese Big 5 with UTF-8

Language configurations

The language configuration of a Storage Virtual Machine (SVM) or a volume must match the client's language configuration for the file names to appear correctly. If there is a mismatch in the language configuration, then some file names might contain incorrect characters.

The following table helps you identify the language configuration for various clients depending on the client encoding types:

Clients protocol	Client encoding type	Language configuration
CIFS running on Win95/98/ME	ISO 8859-1	Match non-UTF-8 client locale. Do not append UTF-8 that is 'en_US'.
CIFS running on WinNT 3.1+	UCS-2	Unless other clients use non-UTF-8 locale, match UTF-8 client locale. Append UTF-8 that is 'en_US.UTF-8'. When other clients use non-UTF-8 locales, match non-UTF-8 client locale. Do not append UTF-8 that is 'en_US'.
NFSv2/3	Non-UTF-8 client locale	Match non-UTF-8 client locale. Do not append UTF-8 that is 'en_US'.
NFSv4	UTF-8	Unless other clients use non-UTF-8 locale, match UTF-8 client locale. Append UTF-8 that is 'en_US.UTF-8'. When other clients use non-UTF-8 locales, match non-UTF-8 client locale. Do not append UTF-8 that is 'en_US'.
FC or iSCSI		UTF-8 preferred, C/POSIX is acceptable.

Note: The default language setting for an SVM is C.UTF-8.

Creating SVMs with FlexVol volumes

A cluster must have at least one or more SVMs to provide data access. You can create SVMs with FlexVol volumes to provide data access to NAS clients and SAN hosts. You can also create SVMs with FlexVol volumes for disaster recovery and high availability.

Before you begin

You must have reviewed the planning requirements and understood the guidelines:

- [Planning to create SVMs](#)
- [Guidelines for creating SVMs](#)

You must ensure that the following requirements are met:

- The cluster must have at least one non-root aggregate with sufficient space.
- There must be at least 1 GB of space on the aggregate for the SVM root volume.
- The cluster must be synchronized by configuring and enabling NTP to prevent CIFS creation and authentication failures.
- If you want to assign IPspace, you must have created the IPspace.
- For MetroCluster configuration, the local and remote sites must be configured.

About this task

You can create multiple SVMs simultaneously either by using different SSH sessions or by using a script.

Note: It is best to create not more than five SVMs simultaneously to avoid any performance degradation.

Choices

- [Creating SVMs with FlexVol volumes for data access](#) on page 84
- [Creating SVMs for a MetroCluster configuration](#) on page 86

Related concepts

[Managing the cluster time \(cluster administrators only\)](#) on page 163

Related tasks

[Delegating administration to SVM administrators](#) on page 91

[Displaying information about SVMs](#) on page 97

Creating SVMs with FlexVol volumes for data access

You must create at least one SVM in a cluster to provide data access to the clients.

Steps

1. Create an SVM:

```
vserver create -vserver vserver_name -rootvolume root_volume_name -
aggregate aggregate_name -rootvolume-security-style {unix|ntfs|mixed} -
language language -ipspace ipspace_name
```

2. Verify the configuration and status of the newly created SVM:

```
vserver show -vserver vserver_name
```

The `Vserver Operational State` field must display the `running` state. If it displays the `initializing` state, it means that some intermediate operation such as root volume creation failed, and you must delete the SVM and re-create it.

Examples

The following command creates an SVM for data access in the IPspace `ipspaceA`:

```
cluster1::>vserver create -vserver vs0.example.com -rootvolume
root_vs0 -aggregate aggr1
-rootvolume-security-style unix -language C.UTF-8 -ipspace ipspaceA

[Job 2059] Job succeeded:
Vserver creation completed
```

The following command shows that an SVM was created with a root volume of 1 GB, and it was started automatically and is in `running` state. The root volume has a default export policy that does not include any rules, so the root volume is not exported upon creation. 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.

```
cluster1::> vserver show -vserver vs0.example.com
                                Vserver: vs0.example.com
                                Vserver Type: data
                                Vserver Subtype: default
                                Vserver UUID: b8375669-19b0-11e5-
b9d1-00a0983d9736
                                Root Volume: root_vs0
                                Aggregate: aggr1
                                NIS Domain: -
                                Root Volume Security Style: unix
                                LDAP Client: -
                                Default Volume Language Code: C.UTF-8
                                Snapshot Policy: default
                                Comment:
                                Quota Policy: default
                                List of Aggregates Assigned: -
                                Limit on Maximum Number of Volumes allowed: unlimited
                                Vserver Admin State: running
                                Vserver Operational State: running
                                Vserver Operational State Stopped Reason: -
                                Allowed Protocols: nfs, cifs, fcp, iscsi,
ndmp
                                Disallowed Protocols: -
                                Is Vserver with Infinite Volume: false
                                QoS Policy Group: -
                                Config Lock: false
                                IPspace Name: ipspaceA
```

After you finish

You must configure data access for the data-serving SVM.

Related tasks

[Configuring SVMs](#) on page 88

Creating SVMs for a MetroCluster configuration

You can create SVMs for a MetroCluster configuration to provide synchronous disaster recovery and high availability of data on clusters that are set up for a MetroCluster configuration.

Before you begin

- The two clusters must be in a MetroCluster configuration.
- Aggregates must be available and online in both the clusters.
- If required, IPspaces with the same names must be created on both the clusters.

About this task

When you create an SVM on one of the clusters in a MetroCluster configuration, the SVM is created as the source SVM, and the partner SVM is automatically created with the same name with the “-mc” suffix on the partner cluster.

In a MetroCluster configuration, you can create 64 SVMs on a cluster. A MetroCluster configuration supports 128 SVMs.

Steps

1. Use the `vserver create` command with the `is-repository` parameter set to **false**.

Example

The following command creates the SVM with the subtype **sync-source** on the local site and the SVM with the subtype **sync-destination** subtype on the partner site:

```
cluster_A::>vserver create -vserver vs4 -rootvolume vs4_root -
aggregate aggr1
-rootvolume-security-style mixed -is-repository false
[Job 196] Job succeeded:
Vserver creation completed
```

The SVM **vs4** is created on the local site and **vs4-mc** is created on the partner site.

2. View the newly created SVMs.
 - On the local cluster, use the `metrocluster vserver show` command to verify the SVMs. The following command displays the partner SVMs and their configuration state:

```
cluster_A::> metrocluster vserver show
```

Cluster	Vserver	Partner Vserver	Configuration State
cluster_A	vs4	vs4-mc	healthy
cluster_B	vs1	vs1-mc	healthy

- From the local and partner clusters, use the `vserver show` command to verify the state of the newly configured SVMs.

The following command displays the administrative and operational states of the SVMs:

```
cluster_A::> vserver show
```

Vserver	Type	Subtype	Admin State	Operational State	Root Volume	Aggregate
vs4	data	sync-source	running	running	vs4_root	aggr1

```
cluster_B::> vserver show
```

Vserver	Type	Subtype	Admin State	Operational State	Root Volume	Aggregate
vs4-mc	data	sync-destination	running	stopped	vs4_root	aggr1

SVM creation might fail if any intermediate operations such as root volume creation fail, and the SVM will be in the `initializing` state. You must delete the SVM and re-create it.

Result

The SVMs for the MetroCluster configuration are created with a root volume of size 1 GB. The operational state of the sync-source SVM is in `running` state and sync-destination SVM is in the `stopped` state.

Creating SVMs with Infinite Volume

A cluster must have at least one or more SVMs to provide data access from the cluster. You can create SVMs with Infinite Volume to provide data access to large data in the NAS environment.

Before you begin

You must have reviewed the planning requirements and understood the guidelines:

- [Planning to create SVMs](#)
- [Guidelines for creating SVMs](#)

You must ensure that the following requirements are met:

- The cluster must have at least one non-root aggregate with sufficient space.
- There must be at least 1 GB of space on the aggregate for the SVM root volume.
- The cluster must be synchronized by configuring and enabling NTP to prevent CIFS creation and authentication failures.
- If you want to assign IPspace, you must have created the IPspace.

About this task

You can create a maximum of five SVMs simultaneously either by using different SSH sessions or by using a script.

Note: It is best to create not more than five SVMs simultaneously to avoid any performance degradation.

You cannot modify the language of an SVM with Infinite Volume.

Steps

1. Use the `vserver create` command with the `is-repository` parameter set to `true`.

Example

The following command creates the Storage Virtual Machine (SVM, formerly known as Vserver) with Infinite Volume named `vs0.example.com`:

```
cluster1::>vserver create -vserver vs1.example.com -rootvolume
root_vs0 -aggregate
aggr1 -rootvolume-security-style unix -language C.UTF-8 -snapshot-
policy default -is-repository true

[Job 2061] Job succeeded:
Vserver creation completed
```

2. Use the `vserver show` command to verify the status of the newly created SVM.

Example

```
cluster1::> vserver show
```

Vserver	Type	Subtype	Admin State	Operational State	Root Volume	Aggregate
cluster1	admin	-	-	-	-	-
cluster1-01	node	-	-	-	-	-
vs1.example.com	data	default	running	running	root_vs0	aggr1

SVM creation might fail if intermediate operations such as root volume creation fail, and the SVM will be in the `initializing` state. You must delete the SVM and re-create it.

Result

The SVM is created with a root volume of 1 GB, and it is started automatically and is in `running` state. 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.

Related concepts

[Managing the cluster time \(cluster administrators only\)](#) on page 163

Related tasks

[Delegating administration to SVM administrators](#) on page 91

[Displaying information about SVMs](#) on page 97

Configuring SVMs

After you create SVMs, you must provision storage, configure the network, services, and protocols to facilitate data access to the clients.

About this task

This procedure provides only high-level information about the configuration tasks that you have to perform after creating the SVM. Detailed information about the configuration tasks is available in other clustered Data ONTAP documentation.

[NetApp Documentation: Data ONTAP 8 \(current releases\)](#)

Steps

1. Specify the aggregates for the SVM for all the volume-related operations that require an aggregate name.
2. Set up a password and unlock the `vsadmin` user account for delegating the SVM administration.
3. Provide data access to the SVM by performing the following steps:
 - a. Set up the network interface, such as creating LIFs and routes.
 - b. Provision storage by creating volumes.

- c. Configure the services, such as LDAP, NIS, and DNS.
 - d. Configure the protocols, such as NFS, CIFS, iSCSI, and FC.
4. Create a load-sharing mirror copy for the root volume on each node of the cluster so that the namespace directory information remains available in the event of a node outage or failover.

Related tasks

[Delegating administration to SVM administrators](#) on page 91

[Modifying aggregates for SVMs](#) on page 94

Related information

[Clustered Data ONTAP 8.3 Network Management Guide](#)

[Clustered Data ONTAP 8.3 Logical Storage Management Guide](#)

[Clustered Data ONTAP 8.3 CIFS File Access Reference Guide](#)

[Clustered Data ONTAP 8.3 NFS File Access Reference Guide](#)

[Clustered Data ONTAP 8.3 SAN Administration Guide](#)

[Clustered Data ONTAP 8.3 SVM Root Volume Protection Express Guide](#)

Creating user accounts for SVM administrators

You must create SVM user accounts for local users, or provide access to NIS and LDAP users, and Active Directory (AD) users or groups for any SVM predefined roles or any customized role to delegate the SVM administration to the SVM administrator.

Choices

- [Creating local user accounts](#) on page 89
- [Providing access to NIS or LDAP user accounts](#) on page 90
- [Providing access to Active Directory users or groups](#) on page 91

Creating local user accounts

You can create local user accounts for SVM administrators to access the SVMs with user password, SSH public key, and SSL certificate authentication methods.

Before you begin

- For SSH public key authentication, you must have created the public/private key pair on the client.
You must have configured the public key on the cluster by using the `security login publickey create` command.
- For SSL certificate authentication, you must have created and installed the self-signed digital certificate.

About this task

You can select the application, authentication method, and predefined role or customized role for the SVM local user account.

Step

1. Use the `security login create` command to create the SVM local user account for any SVM role.

Example

The following command creates the SVM role with password as the authentication method:

```
cluster1::>security login create -user-or-group-name user1 -application ssh -authmethod password -role vsadmin-backup -vserver vs1.example.com
```

The following command creates the SVM role with public key as the authentication method:

```
cluster1::>security login create -user-or-group-name user2 -application ssh -authmethod publickey -role vsadmin-backup -vserver vs2.example.com
```

The following command creates the SVM role with certificate as the authentication method:

```
cluster1::>security login create -user-or-group-name user3 -application ontapi -authmethod cert -role vsadmin-backup -vserver vs3.example.com
```

Providing access to NIS or LDAP user accounts

You can provide access to NIS or LDAP user accounts for SVM administrators to access the SVMs with the LDAP or NIS authentication method.

Before you begin

- NIS or LDAP user accounts must be created on the NIS or LDAP server.
- The SVM must be configured to lookup the NIS or LDAP server.

About this task

You can select the application and role for the SVM user account.

You can create the SVM user account in LDAP or NIS directory, or as a local user in the local administrative repository of the SVM. Depending on where you create the SVM user account, you must include the relevant name service source when you configure the name service switch (`ns-switch`).

Steps

1. Use the `security login create` command to create the SVM user account for any SVM role.

Example

The following command creates the SVM role with the `nsswitch` authentication method:

```
cluster1::>security login create -user-or-group-name user1 -application ssh -authmethod nsswitch -role vsadmin-backup -vserver vs1.example.com
```

2. Use the `vserver services name-service` command to add **ldap**, **nis**, **files** as a name service source for the SVM authentication.

Including **files** as a name service enables SVM user account authentication through the SVM local administrative repository if the user is not found in the LDAP or NIS servers.

Example

The following command adds the name service source for the SVM:

```
cluster1::>vserver services name-service ns-switch create -vserver vs1.example.com -sources ldap,nis,files -database password -enabled true
```

Related information

[Clustered Data ONTAP 8.3 NFS File Access Reference Guide](#)

Providing access to Active Directory users or groups

You can provide access to Active Directory (AD) user or group accounts for SVM administrators to access the SVMs with the Windows Active Directory authentication method.

Before you begin

- AD users or groups must be created on the AD server.
- AD computer accounts for SVMs must be created.

About this task

The AD group provides centralized privilege level control for user accounts and supports only SSH and Data ONTAP API access methods. Any user belonging to the AD group can access the SVM with a role assigned to the group.

Step

1. Use the `security login create` command to grant access to AD users or groups:

- AD user accounts

The following command creates an account with the user name `guest` in `DOMAIN1`, the application `ssh`, the authentication method `domain`, and the access-control role `vsadmin` for the SVM `vs0`:

```
cluster1::> security login create -user-or-group-name DOMAIN1\guest -application ssh -authmethod domain -role vsadmin -vserver vs0.example.com
```

- AD group accounts

The following command creates an account with the AD group name `adgroup` in `DOMAIN1`, the application `ssh`, the authentication method `domain`, and the access-control role `vsadmin` for the SVM `vs1`:

```
cluster1::> security login create -user-or-group-name DOMAIN1\adgroup -application ssh -authmethod domain -role vsadmin -vserver vs1.example.com
```

Delegating administration to SVM administrators

After setting up a functional Storage Virtual Machine (SVM) with basic network configuration, you can optionally delegate the administration of the SVM to the SVM administrator. You can delegate the SVM administration by creating and assigning user accounts either with predefined roles or customized roles.

Before you begin

If you want to delegate the SVM administration with any customized roles, you must have created customized roles by using the `security login role create` command.

Steps

1. Use the `vserver show -fields aggr-list` command to verify if the SVM has any aggregates assigned.

Note: If no aggregates are assigned to the SVM, the SVM administrator cannot create volumes.

2. If the SVM does not have any assigned aggregates, use the `vserver add-aggregates` command to specify aggregates in the aggregates list of the SVM.

Note: You cannot assign the root aggregate “aggr0” for any SVM.

Example

The following command specifies the aggregate `aggr1` for the SVM `vs1.example.com`:

```
cluster1::> vserver add-aggregates -vserver vs1.example.com -  
aggregates aggr1
```

3. For the SVM administrator administering an SVM with FlexVol volume, use the `vserver modify` command with the `max-volumes` option to specify the maximum number of volumes that the SVM administrator can create on that SVM.

Example

The following command specifies the maximum number of volumes for the SVM `vs1.example.com`:

```
cluster1::> vserver modify -vserver vs1.example.com -max-volumes 10
```

4. Use the `vserver add-protocols` or `vserver remove-protocols` command to specify the protocols for the SVM.

Example

The following command specifies the CIFS protocol for the SVM `vs1.example.com`:

```
cluster1::> vserver add-protocols -vserver vs1.example.com -protocols  
cifs
```

Only the specified protocols are available for configuration and data access.

5. For SVM management, create a new LIF or use one of the data LIFs.

Important: You cannot use data LIFs configured for SAN protocols for SVM management.

If you want to...	Then...
<p>Create a new LIF for SVM management</p>	<p>a. Identify the IPspace assigned to the SVM by using the <code>vserver show</code> command:</p> <pre data-bbox="711 338 1377 426">cluster1::> vserver show -vserver vs1.example.com -fields ipspace vserver ipspace ----- vs1.example.com ipspace1</pre> <p>b. Select a port from the same IPspace as the SVM by using the <code>network port show</code> command:</p> <pre data-bbox="711 527 1377 674">cluster1::> network port show -ipspace ipspace1 Node Port IPspace Broadcast Domain Link MTU Speed (Mbps) ----- cluster1-01 e0c ipspace1 192.0.2.120/24 up 1500 auto/1000</pre> <p>c. Create a LIF by using the <code>network interface create</code> command.</p> <p>The following command creates the data LIF <code>lif3</code> on the port <code>e0c</code> that belongs to IPspace <code>ipspace1</code> for the SVM <code>vs1.example.com</code> belonging to <code>ipspace1</code>:</p> <pre data-bbox="711 863 1377 951">cluster1::> network interface create -vserver vs1.example.com -lif lif3 - data-protocol none -role data -home-node node1-01 -home-port e0c -address 192.0.2.129 -netmask 255.255.255.128</pre>
<p>Use a LIF for NFS, CIFS, and SVM management</p>	<p>Change the firewall policy to <code>mgmt</code> by using the <code>network interface modify</code> command.</p> <p>The following command modifies the data LIF <code>lif1</code> for the SVM <code>vs1.example.com</code> to support SVM management:</p> <pre data-bbox="672 1129 1377 1188">cluster1::>network interface modify -vserver vs1.example.com -lif lif1 - firewall-policy mgmt</pre>
<p>6. Depending on the type of SVM administrator roles, perform the appropriate action:</p>	
If you want to use...	Then...
<p><code>vsadmin</code>, a predefined role that is created and is in the locked state when the SVM is created</p>	<p>a. Set up a password by using the <code>security login password</code> command:</p> <ol style="list-style-type: none"> i. Enter a password for the user account. ii. Confirm the password by reentering it. <p>The following command sets up a password for the user account <code>vsadmin</code> on the SVM <code>vs1.example.com</code>:</p> <pre data-bbox="711 1598 1377 1717">cluster1::>security login password -username vsadmin -vserver vs1.example.com Please enter a password for user 'vsadmin': Please enter it again: cluster1::></pre> <p>b. Unlock the user account by using the <code>security login unlock</code> command.</p> <p>The following command unlocks the user account <code>vsadmin</code> for the SVM <code>vs1.example.com</code>:</p> <pre data-bbox="711 1871 1377 1934">cluster1::> security login unlock -username vsadmin -vserver vs1.example.com</pre>

If you want to use...	Then...
Any customized role or other predefined roles, such as vsadmin-volume, vsadmin-protocol, or vsadmin-readonly	<p>Create a user account with a role by using the <code>security login create</code> command:</p> <ol style="list-style-type: none"> a. Enter a password for the user account. b. Confirm the password by reentering it. <p>The following command creates the user account <code>user1</code> with <code>vsadmin-readonly</code> role for the SVM <code>vs1.example.com</code>:</p> <pre data-bbox="667 499 1377 636">cluster1::> security login create -user-or-group-name user1 -application ssh -authmethod password -vserver vs1.example.com -role vsadmin-readonly Please enter a password for user 'user1': Please enter it again: cluster1::></pre>

Result

After you assign the SVM to an SVM administrator, the SVM administrator can log in to the SVM by using the user name, password, and the management IP address.

Modifying SVMs

You can specify the aggregates for a Storage Virtual Machine (SVM), allow or disallow protocols on the SVM, and modify other attributes of the SVM, such as the language and policies.

Related tasks

[Displaying information about SVMs](#) on page 97

Modifying aggregates for SVMs

You must specify the aggregates for an SVM so that SVM administrators can view the list of available aggregates before performing provisioning operations that require an aggregate name—for example, creating a volume. You can add or remove aggregates to specify the aggregates list for the SVM.

About this task

As a cluster administrator, you can perform all the operations that require an aggregate name. When you specify the aggregate for the SVM, you can perform the same limited provisioning operations as the SVM administrator. However, you can move the volumes and copy the volumes across aggregates.

Steps

1. Add or remove aggregates for the SVM:
 - Use the `vserver add-aggregates` command to add aggregates to the list of aggregates for the SVM.

The following command adds the aggregates `aggr0` and `aggr1` to the list of aggregates for the SVM `vs1.example.com`:

```
cluster1::> vserver add-aggregates -vserver vs1.example.com -
aggregates aggr0, aggr1
```

- Use the `vserver remove-aggregates` command to remove aggregates from the list of aggregates for the SVM.

The following command removes `aggr0` from the list of aggregates for the SVM `vs1.example.com`:

```
cluster1::> vserver remove-aggregates -vserver vs1.example.com -
aggregates aggr0
```

2. Use the `vserver show -instance` command to view the list of aggregates assigned to the SVM.

Example

The following command shows the list of aggregates assigned to the SVM `vs1.example.com`:

```
cluster1::> vserver show -instance -vserver vs1.example.com

                Vserver: vs1.example.com
                Vserver Type: data
                Vserver Subtype: default
                ...
                ...
                List of Aggregates Assigned: aggr1
                ...
                ...
```

Result

The modified aggregates list is available for the provisioning operations for the SVM.

Modifying protocols for SVMs

When an SVM is created, all protocols are allowed for configuration on that SVM. You can modify the list of protocols allowed for configuration on the SVM by adding or removing the protocols.

About this task

When you add protocols for the SVM, these protocols are available for configuration.

When you remove protocols, these protocols are disallowed for configuration on the SVM, and data access through these protocols is stopped.

Steps

1. Add or remove protocols for the SVM:

- Use the `vserver add-protocols` command to allow protocols for the SVM.

The following command adds the NFS and CIFS protocols to the list of aggregates for the SVM `vs1.example.com`:

```
cluster1::> vserver add-protocols -vserver vs1.example.com -
protocols nfs, cifs
```

- Use the `vserver remove-protocols` command to disallow protocols for the SVM.

The following command removes the NFS protocol from the the list of allowed protocols for the SVM `vs1.example.com`:

```
cluster1::> vserver remove-protocols -vserver vs1.example.com -
protocols nfs
```

2. Use the `vserver show-protocols` command to view the list of allowed protocols for the SVM.

Example

The following command lists the protocols for all the SVMs:

```
cluster1::> vserver show-protocols

Vserver Name          Protocols
-----
...
...
vs2.example1.com      fcp
vs1.example1.com      cifs fcp, iscsi, ndmp
...
...
```

Commands for modifying SVM attributes

You can modify the attributes of an SVM such as the language, maximum number of volumes, Snapshot policy, and quota policy by using the `vserver modify` command.

If you want to modify...	Use the following command...	Applicable to SVMs with FlexVol volumes ?	Applicable to SVMs with Infinite Volume ?	Notes
Language	<code>vserver modify -language</code>	Yes	No	When you modify the SVM language, the language setting of the existing volumes in the SVM does not change.
Snapshot policy	<code>vserver modify -snapshot-policy</code>	Yes	Yes	NA
Quota policy	<code>vserver modify -quota-policy</code>	Yes	No	NA
QoS policy group	<code>vserver modify -qos-policy-group</code>	Yes	No	NA

If you want to modify...	Use the following command...	Applicable to SVMs with FlexVol volumes ?	Applicable to SVMs with Infinite Volume ?	Notes
Maximum number of volumes	<code>vserver modify -max-volumes</code>	Yes	No	When the value is set to unlimited , which is the default value, any number of volumes can be created on that SVM. If you specify the value as 0 , then volumes cannot be created on that SVM. Therefore, you must specify a value so that the SVM administrator can create volumes.

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Displaying information about SVMs

Cluster administrators can view the configuration information about one or more Storage Virtual Machines (SVMs, formerly known as Vservers) by using the `vserver show` command.

Step

1. Enter the appropriate command to view information about the SVMs:

If you want to...	Enter the following command...
View basic information about all the SVMs	<code>vserver show</code>
View detailed information about all the SVMs	<code>vserver show -instance</code>
View information about an SVM	<code>vserver show -vserver <i>Vserver_name</i></code> <i>Vserver_name</i> is the name of the SVM.

Example

The following command displays basic information about all SVMs:

```
cluster1::>vserver show
```

Vserver	Type	Subtype	Admin State	Operational State	Root Volume	Aggregate
cluster1	admin	-	-	-	-	-
cluster1-01	node	-	-	-	-	-

```

cluster1-02      node      -      -      -      -      -
vs0.example.com  data      default  running  running  root_vs0  aggr1
vs2.example.com  data      data      default  default  running
running          root_vol2  aggr1

```

The following command displays detailed information about all the SVMs:

```

cluster1::> vserver show -instance

          Vserver: cluster1
          Vserver Type: admin
          Vserver Subtype:
          Vserver UUID: 00000000-0000-0000-0000-000000000000
          Root Volume: -
          Aggregate: -
          .
          .
          List of Aggregates Assigned: -
          Limit on Maximum Number of Volumes allowed: -
          Vserver Admin State: -
          Vserver Operational State: -
          Vserver Operational State Stopped Reason: -
          .
          .
          Config Lock: false
          IPspace Name: Default
          Is Vserver Protected: -

          Vserver: vs0.example.com
          Vserver Type: data
          Vserver Subtype: default
          Vserver UUID: 49294a39-e762-11df-8768-123478563412
          Root Volume: root_vs0
          Aggregate: aggr1
          .
          .
          List of Aggregates Assigned: aggr0_2, aggr1
          Limit on Maximum Number of Volumes allowed: 10
          Vserver Admin State: running
          Vserver Operational State: running
          Vserver Operational State Stopped Reason: -
          .
          .
          Config Lock: false
          IPspace Name: Default
          Is Vserver Protected: false

```

The following command displays detailed information about the SVM vs2.example.com

```

cluster1::> vserver show -vserver vs2.example.com

          Vserver: vs2.example.com
          Vserver Type: data
          Vserver Subtype: default
          Vserver UUID: ca34e6b2-ddec-11df-b066-123478563412
          Root Volume: root_vol2
          .
          .
          Config Lock: false
          IPspace Name: Default
          Is Vserver Protected: false

```

Renaming SVMs

You can rename a Storage Virtual Machine (SVM) by using the `vserver rename` command. For example, you can rename an SVM when you want the SVM to have a unique name. You cannot rename a node or admin SVM by using the `vserver rename` command.

Before you begin

The SVM being renamed must not be in an SVM peer relationship.

Steps

1. Use the `vserver rename` command to rename an SVM.

Example

The following example shows how to rename the SVM named vs1.example.com as vs2.example.com:

```
Cluster1::> vserver rename -vserver vs1.example.com -newname
vs2.example.com
```

For more information about this command, see the man pages.

2. Use the `vserver show` command to view the changes in the SVM's name.

Deleting SVMs

You can delete Storage Virtual Machines (SVMs) that are no longer needed from the cluster by using the `vserver delete` command.

Before you begin

1. You must have deleted the SVM peer relationship associated with the SVM.
2. You must have disabled Snapshot copies, and DP and LS mirrors for all volumes.
3. If you are using LUNs, you must have unmapped the LUNs, taken them offline, and deleted them.
4. You must have deleted all the igroups that belong to the SVM manually.
5. You must have unmounted all volumes on the SVM, taken them offline, and deleted them including the root volume of the SVM.
6. You must have deleted CIFS server.
7. You must have deleted any customized user accounts and roles associated with the SVM.
8. You must have stopped the SVM.

About this task

When you delete an SVM, the following objects associated with the SVM are also deleted automatically:

- LIFs, LIF failover groups, and LIF routing groups
- Export policies
- Sis policies

You cannot recover any SVM related information after deleting an SVM.

If you delete an SVM that is configured to use Kerberos, or modify an SVM to use a different service principal name (SPN), SVM's original service principal name is not automatically deleted or disabled from 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.

If you need to move data from a first SVM to a second SVM before you delete the first SVM, you can use SnapMirror commands. For more information about SnapMirror, see the *Clustered Data ONTAP Data Protection Guide*.

Step

1. Use the `vserver delete` command to delete an SVM.

Example

The following example shows how to delete the SVM named vs1.example.com:

```
cluster1::> vserver delete -vserver vs1.example.com
```

For more information about this command, see the man pages.

Note: SVM delete operation might fail due to any intermediate operation failures. As a result, the SVM will be in deleting state. It is best to delete such SVMs because you cannot perform other SVM operations on that SVM. For example, you cannot create an SVM peering relationship with SVMs in deleting state.

Starting SVMs

You can provide data access from a Storage Virtual Machine (SVM) by starting the SVM. You can use the `vserver start` command to start an SVM.

About this task

When you start the SVM, the protocols that were stopped will start serving data. The protocol might have been stopped either independently by running commands such as `vserver fcp stop` or when the SVM was stopped.

You cannot start an SVM during a storage failover (SFO) if the resources of that SVM are part of an HA pair.

Steps

1. Use the `vserver start` command to start the SVM.

Example

The following command starts the SVM vs0.example.com:

```
cluster1::> vserver start -vserver vs0.example.com
[Job 21] Job succeeded: DONE
```

2. If the SVM start operation fails with an internal error message, perform the following steps:
 - a. Use the `vserver show` command to verify the admin state of the SVM.
 - b. If the admin state of the SVM is in the **starting** state, run the `vserver start` command again.

Example

The following commands restart the SVM that failed to start:

```
cluster1::> vserver start vs1.example.com
[Job 43] Job is queued: Vserver Start.

Error: command failed: [Job 43] Job failed: Internal Error

cluster1::> vserver show
```

Vserver	Type	Subtype	Admin State	Operational State	Root Volume	Aggregate
cluster1	admin	-	-	-	-	-
cluster1-01	node	-	-	-	-	-
vs0.example.com	data	default	running	running	root_vs0	aggr1

```
vs1.example.com data default starting stopped root_vs1 aggr1
cluster1::> vserver start vs1.example.com
[Job 48] Job succeeded: DONE
```

Result

The SVM is in the **running** state and starts serving data to clients.

When you start an SVM with Infinite Volume, its data policy is automatically re-imported and its JSON format is checked.

Related tasks

[Displaying information about SVMs](#) on page 97

Related information

[Clustered Data ONTAP 8.3 Infinite Volumes Management Guide](#)

Stopping SVMs

You can use the `vserver stop` command to stop a Storage Virtual Machine (SVM) when you want to troubleshoot or delete the SVM, or stop data access from the SVM.

Before you begin

All clients connected to the SVM must be disconnected.

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

About this task

You cannot stop an SVM during a storage failover (SFO) if the resources of that SVM are part of an HA pair.

When you stop the SVM, other operations such as SnapMirror data transfers continue to run as per the schedule.

The SVM administrator cannot log in to the SVM when the SVM is in the **stopped** state.

Steps

1. Use the `vserver stop` command to stop the SVM.

Example

The following command stops the SVM `vs0.example.com`:

```
cluster1::> vserver stop -vserver vs0.example.com
[Job 41] Job succeeded: DONE
```

2. If the SVM stop operation fails with an internal error message, perform the following steps:
 - a. Use the `vserver show` command to verify the admin state of the SVM
 - b. If the admin state of the SVM is in the **stopping** state, run the `vserver stop` command again.

Example

The following commands stop the SVM that failed to stop:

```
cluster1::> vserver stop -vserver vs1.example.com
[Job 45] Job is queued: Vserver Stop.

Error: command failed: [Job 45] Job failed: Internal Error

cluster1::> vserver show
Vserver          Type      Subtype      Admin      Operational      Root
-----          -
cluster1         admin     -            -          -                -
cluster1-01      node     -            -          -                -
vs0.example.com  data     default      stopped   stopped          root_vs0
vs1.example.com  data     default      stopping  running          root_vs1
Aggregate
-----
aggr1
aggr1

cluster1::> vserver stop -vserver vs1.example.com
[Job 49] Job succeeded: DONE
```

Result

The SVM is in **stopped** state and stops serving data to clients.

Related tasks

[Displaying information about SVMs](#) on page 97

Administering SVMs from the SVM context

You can administer a Storage Virtual Machine (SVM) and its resources from the context of an SVM by using the `vserver context` command.

About this task

After you switch to the SVM context, your capabilities will be same as that of the SVM administrator. If you do not specify the user name while executing the `vserver context` command, then you will have capabilities same as that of the default SVM administrator (`vsadmin`). If you specify the user name, then you will have capabilities same as that of the role of the user name.

If you want to switch from one SVM to another, you must exit from the first SVM.

Steps

1. Use the `vserver context` command to enter into the SVM context.

Example

The following example shows how to switch the context from cluster to SVM `vs1.example.com`:

```
cluster1::> vserver context -vserver vs1.example.com -username
vsadmin-volume

Info: Use 'exit' command to return.

vs1.example.com::>
```

For more information about `vserver context` command, see the man pages.

You can use a role of another SVM administrator by specifying the `-username` option.

You are in the context of SVM `vs1.example.com`. Your capabilities will be same as that of the `vsadmin-volume` role.

2. Enter the command you want to run from the SVM context.

Example

The following example shows how to view the volumes that belong to the SVM vs1.example.com from the SVM vs1.example.com context:

```
vs1.example.com::> vol show
(volume show)
Vserver      Volume      Aggregate   State      Type  Size  Available  Used%
-----
vs1.example.com  root_voll  aggr3      online    RW   1GB   972.5MB   5%
vs1.example.com  voll       aggr1      online    RW   20MB  18.88MB  5%
```

3. Type **exit** at the SVM prompt to exit from the SVM context.

Restoring the root volume of an SVM

If the root volume of a Storage Virtual Machine (SVM) becomes unavailable, clients cannot mount the root of the namespace. In such cases, you must restore the root volume by promoting another volume or creating a new root volume to facilitate data access to the clients.

Before you begin

SVM root volume must be protected by using the load-sharing mirror copy or data-protection copy.

About this task

You can promote any volume that does not have other volumes junctioned to it.

When a new volume is promoted as the SVM root volume, the data volumes are associated with the new SVM root volume.

Choices

- For SVMs with FlexVol volumes, promote one of the following volumes to restore the root volume:
 - Load-sharing mirror copy
Promoting a load-sharing mirror copy
 - Data-protection mirror copy
Promoting a data-protection mirror copy
 - New FlexVol volume
Promoting a new FlexVol volume
- For SVMs with Infinite Volume, create a new root volume.
Creating a new root volume on an SVM with Infinite Volume

After you finish

You must mount the new root volume by using the `volume mount` command.

Promoting a load-sharing mirror copy

You can promote a load-sharing mirror copy to restore the root volume of a Storage Virtual Machine (SVM).

Steps

1. Use the `set -privilege advanced` command to set the privilege level to advanced.

2. Use the `snapmirror promote` command to promote the load-sharing mirror copy as the root volume.
 3. Use the `vol show` command to verify the new root volume of the SVM.
 4. Use the `vol rename` command to rename the volume that was promoted as the root volume.
- For more information about these commands, see the man pages.

The following example shows how to promote a load-sharing mirror copy `vol_dstls` as the root volume of the SVM `vs1.example.com`:

```
cluster1::> set -privilege advanced

Warning: These advanced commands are potentially dangerous; use them only
when directed to do so by technical support.
Do you want to continue? {y|n}: y

cluster1::*> snapmirror promote -destination-path vs1.example.com:vol_dstls

Warning: Promote will delete the read-write volume cluster1://
vs1.example.com/voll and replace it with cluster1://vs1.example.com/
vol_dstls.
Do you want to continue? {y|n}: y
[Job 489] Job succeeded: SnapMirror: done

cluster1::*> volume show -volume vol_dstls -instance

      Vserver Name: vs1.example.com
      Volume Name: vol_dstls
      .
      .
      Junction Path: /
      .
      Vserver Root Volume: true
      .
      .
```

Promoting a data-protection mirror copy

You can use a data-protection mirror copy to restore the root volume of a Storage Virtual Machine (SVM).

Steps

1. Use the `snapmirror break` command to break the SnapMirror relationship.
 2. Use the `set -privilege advanced` command to set the privilege level to advanced.
 3. Use the `volume make-vsroot` command to promote the data-protection mirror copy as the root volume.
 4. Use the `volume show` command to verify the new root volume of the SVM.
 5. Use the `volume rename` command to rename the volume that was promoted as the root volume.
- For more information about these commands, see the man pages.

The following example shows how to promote a data-protection mirror copy `vol_dstdp` as the root volume of the SVM `vs1.example.com`:

```
cluster1::
> snapmirror break -destination-path vs1.example.com:vol_dstdp
[Job 521] Job succeeded: SnapMirror Break Succeeded
```



```

cluster1::> set -privilege advanced

Warning: These advanced commands are potentially dangerous; use them only
when directed to do so by technical support.
Do you want to continue? {y|n}: y

cluster1::*> volume make-vsroot -volume vol_dstdp -vserver vs1.example.com
[Job 522] Job succeeded: DONE

cluster1::*> volume show -volume vol_dstdp -instance

      Vserver Name: vs1.example.com
      Volume Name: vol_dstdp
      .
      .
      Junction Path: /
      .
      Vserver Root Volume: true
      .
      .

```

Promoting a new FlexVol volume

You can create and use a new FlexVol volume to restore the root volume of a Storage Virtual Machine (SVM).

About this task

Starting from clustered Data ONTAP 8.2, the SVM root volume is created with 1 GB size to prevent any failures when mounting any volume in the SVM root volume due to lack of space or inodes. Therefore, if you are promoting a new FlexVol volume, it should be minimum 1 GB in size.

Steps

1. Use the `set -privilege advanced` command to set the privilege level to advanced.
2. Use the `volume create` command to create a new FlexVol volume of 1 GB size.
3. Use the `volume make-vsroot` command to promote the FlexVol volume as the root volume.
4. Use the `volume show` command to verify the new root volume of the SVM.
5. Use the `volume rename` command to rename the volume that was promoted as the root volume.

For more information about these commands, see the man pages.

The following command promotes the FlexVol volume `new_rootvol` as the root volume of the SVM `vs1.example.com`:

```

cluster1::> set -privilege advanced

Warning: These advanced commands are potentially dangerous; use them only
when directed to do so by technical support.
Do you want to continue? {y|n}:

cluster1::*> vol create -vserver vs3 -volume new_rootvol -aggregate aggr0 -
size 1GB
(volume create)

cluster1::*> volume make-vsroot -vserver vs1.example.com -volume new_rootvol

cluster1::*> volume show -volume new_rootvol -instance

      Vserver Name: vs1.example.com
      Volume Name: new_rootvol
      .
      .
      Junction Path: /

```

```

      .
      Vserver Root Volume: true
      .
      .

```

Creating a new root volume on an SVM with Infinite Volume

If the root volume for a Storage Virtual Machine (SVM) with Infinite Volume becomes unavailable, you must create a new root volume on a specific aggregate to replace the unavailable root volume. You cannot promote an existing volume as the root volume.

About this task

When you create a new root volume for an SVM with Infinite Volume, the old, unavailable root volume is automatically deleted.

The new root volume is automatically created at a size that is appropriate for an SVM with Infinite Volume. For an SVM with Infinite Volume, the root volume is not required to be 1 GB in size.

Steps

1. Set advanced privilege by using the `set -privilege` command.
2. Create a new root volume for the SVM by using the `volume make-vsroot` command with the `-aggregate` parameter.

A new root volume is created, and the old, unavailable root volume is deleted.

3. Mount the Infinite Volume by using the `volume mount` command.

The following example shows how to create a new root volume for the SVM with Infinite Volume `vs0` and mount Infinite Volume `repo_vol`:

```

cluster1::> set -privilege advanced

Warning: These advanced commands are potentially dangerous; use them
only when directed to do so by technical support.
Do you want to continue? {y|n}: y

cluster1::*> volume make-vsroot -vserver vs0 -volume vs0_root1 -
aggregate aggr1

Warning: Creating a new Vserver root volume will delete the old
Vserver root
volume "vs0_root".
Do you want to continue? {y|n}: y

cluster1::*> volume show -volume vs0_root1 -instance

                Vserver Name: vs0
                Volume Name: vs0_root1
                ...
                Junction Path: /
                ...
                Vserver Root Volume: true
                ...

cluster1::*> volume mount -vserver vs0 -volume repo_vol -junction-
path /NS -active true -policy-override false
[Job 69] Job succeeded: Volume "repo_vol" on Vserver "vs0" is
mounted at "/NS".

cluster1::*> volume show -vserver vs0 -volume repo_vol -fields
state, junction-path

```

```

vserver volume  state  junction-path
-----
vs0      repo_vol  online /NS

```

Controlling and monitoring I/O performance to SVMs by using Storage QoS

You can control the input/output (I/O) performance to Storage Virtual Machines (SVMs) with FlexVol volumes by assigning these SVMs to Storage QoS policy groups. For example, you may want to control I/O performance to ensure that workloads achieve specific performance objectives or you may want to control I/O performance so that you can throttle a workload that negatively impacts other workloads.

About this task

Policy groups enforce maximum throughput limits (for example, 100 MB/s). However, you can create a policy group without specifying a maximum throughput, so that you can monitor performance before you control the workload.

You can also assign FlexVol volumes, LUNs, and files to policy groups.

Following are requirements for assigning an SVM to a policy group:

- The SVM that you assign must be the same SVM to which the policy group belongs. You specify this SVM when you create the policy group.
- If you assign an SVM to a policy group, then you cannot also assign the storage objects contained by that SVM to a policy group.

Steps

1. Use the `qos policy-group create` command to create a policy group.

Example

The following command creates policy group `pg-vs1` with a maximum throughput of 5,000 IOPS.

```

cluster1::> qos policy-group create pg-vs1 -vserver vs1 -max-
throughput 5000iops

```

2. Use the `vserver modify` command with the `-qos-policy-group` parameter to assign an SVM to a policy group.

Example

The following command assigns the SVM `vs1` to policy group `pg-vs1`.

```

cluster1::> vserver modify -vserver vs1 -qos-policy-group pg-vs1

```

3. Use the `qos statistics` commands to view performance data.

Example

The following command shows the performance of policy groups.

```
cluster1::> qos statistics performance show
Policy Group          IOPS          Throughput    Latency
-----
-total-              12316         47.76MB/s    1264.00us
pg_app2              7216          28.19MB/s    420.00us
pg_vs1              5008          19.56MB/s     2.45ms
_System-Best-Effort    62           13.36KB/s     4.13ms
_System-Background    30            0KB/s         0ms
```

4. If necessary, use the `qos policy-group modify` command to adjust the policy group's maximum throughput limit.

Example

The following command modifies the maximum throughput for policy group `pg-vs1` to 4,500 IOPS.

```
cluster1::> qos policy-group modify pg-vs1 -max-throughput 4500iops
```

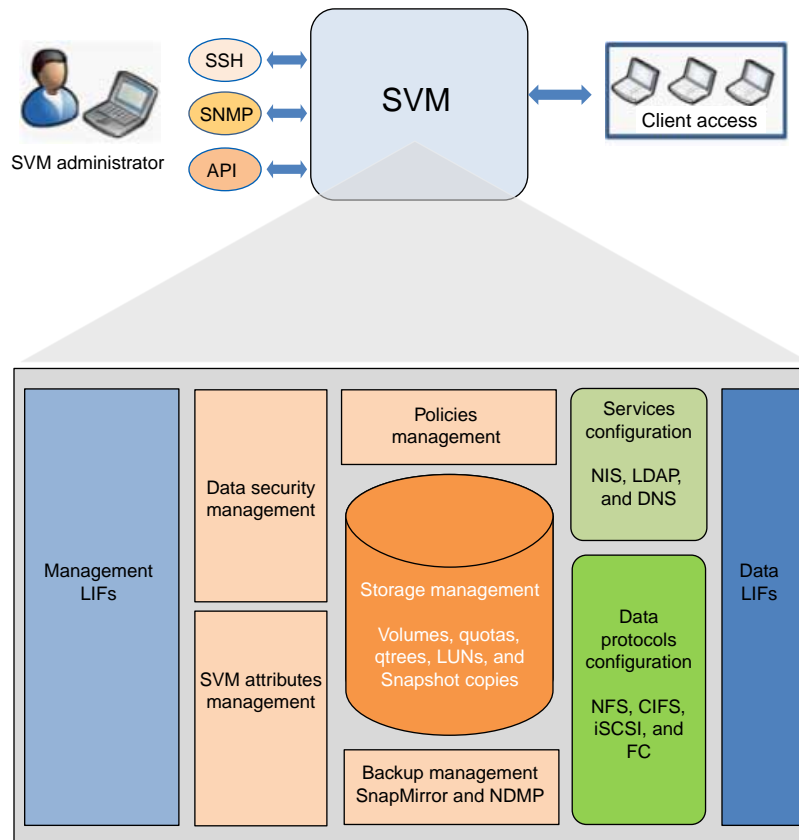
Related concepts

Example: Proactively setting a limit on workloads in a shared storage infrastructure on page 251
Managing workload performance by using Storage QoS on page 242

Administering SVMs

Depending on the capabilities assigned by the cluster administrator, an SVM administrator can perform various administration tasks on a Storage Virtual Machine (SVM, formerly known as Vserver). After logging in to the SVM, an SVM administrator can identify the capabilities assigned and the commands that are available for the administration.

The following illustration depicts the SVM administrative components:



You might have all or some of the following administration capabilities:

- **Jobs and schedules management**
You can manage jobs and schedules related to the SVM.
- **Data access protocol configuration**
You can configure data access protocols, such as NFS, CIFS, iSCSI, and Fibre Channel (FC) protocol (Fibre Channel over Ethernet included).
- **Policy management**
You can create and manage policies to manage data access from the SVM.
- **Data access security management**
You can set security on the SVM's data without the need of a client.
- **Services configuration**
You can configure services, such as LDAP, NIS, and DNS.
- **Storage management**
You can manage volumes, quotas, qtrees, and files.

- LUN management
You can manage LUNs in a SAN environment.
- Backup management
You can back up and manage the SVM's data by using SnapMirror technology and NDMP.
- Monitoring SVM
You can monitor performance data, network connection, information, and SVM health.

Note: For troubleshooting or modifying SVM configurations, SVM administrators must contact the cluster administrator.

Note: The Data ONTAP command-line interface (CLI) continues to use the term *Vserver* in the output, and `vserver` as a command or parameter name has not changed.

SVM access and authentication methods

As an SVM administrator, you must be aware of the access methods and authentication methods that SVM user accounts use to access an SVM.

Access methods for user accounts

Depending on how the cluster administrator sets up an SVM user account, an SVM administrator can access the SVM for administration by using the following access methods:

- SSH
- Data ONTAP APIs

Note: Data ONTAP APIs access method is over HTTPS.

- SNMP

Authentication methods for user accounts

The method used to authenticate an SVM user account depends on the access method used by the cluster administrator to set up the SVM user account.

Your user account can be authenticated by using one of the following authentication methods:

- Network Information Service (NIS) and Lightweight Directory Access Protocol (LDAP) (**nsswitch**)

Note: Clustered Data ONTAP supports only the RFC 2307 schema for LDAP authentication of SVM accounts. It does not support any other schemas, such as Active Directory Identity Management for UNIX (AD-IDMU) and Active Directory Services for UNIX (AD-SFU).

- Windows Active Directory (**domain**)
- User password (**password**)
- SSH public key (**publickey**)
- SNMP user-based security model (**usm**)
- SNMP community strings (**community**)
- SSL certificate authentication (**cert**)

Logging in to an SVM

To manage the SVM resources, an SVM administrator logs in to an SVM by using the user name and password provided by the cluster administrator. The SVM administrator can use an appropriate Secure Shell client application, such as PuTTY for Windows operating system and OpenSSH for UNIX operating system.

Before you begin

You must have the management IP address of the SVM, user name, and password.

About this task

After you log in, you might be able to manage all or some of the following SVM resources depending on the capabilities assigned to your account by the cluster administrator:

- Data access protocols, such as NFS, CIFS, iSCSI, and FC (FCoE included)
- Services, such as NIS, LDAP, and DNS
- Volumes, qtrees, quotas, Snapshot copies, and files
- Data backup with SnapMirror and NDMP
- Data security and policies

You can also monitor the network connection, network interface, LDAP client configuration, and SVM health.

Note: Clustered Data ONTAP supports only the AES and 3DES encryption algorithms (also known as ciphers) for SSH.

Step

1. To log in to an SVM by using SSH application, perform the appropriate action depending on the operating system:

If your host has...	Then...
Windows operating system	<ol style="list-style-type: none"> a. Enter the management IP address of the SVM in the SSH application. b. At the login prompt, enter the user name and password.
UNIX or Linux operating system	Enter the following command from the client application: <code>ssh vserver_admin_name@vserver_ip_address</code> <i>vserver_admin_name</i> is the user name. <i>vserver_ip_address</i> is the management IP address of the SVM.

Note: If you or the cluster administrator has created a public key for your user account, you do not require a password to log in to the SVM.

Changing the login password

After an SVM administrator logs in to the SVM by using the user name and password provided by the cluster administrator, the SVM administrator can change the login password.

About this task

You must remember the following default rules when you change the login password:

- A password cannot contain the user name.
- A password must be at least eight characters long.
- A password must contain at least one letter and one number.
- A password cannot be the same as the last six passwords.

Steps

1. Change the login password by using the `security login password` command.
2. Enter your current password.
3. Enter a new password.
4. Confirm the password by entering the new password again.

Result

Your user account is updated with the new password. You must enter the new password on the subsequent login.

The following example shows how to change a user password:

```
vs1.example.com::> security login password
Please enter your current password:
Please enter a new password:
Please enter it again:
vs1.example.com::>
```

Displaying information about SVMs

SVM administrators can view the details of a Storage Virtual Machine (SVM, formerly known as Vserver) that are assigned to them by using the `vserver show` command.

Step

1. View details about an SVM:
 - For basic information, enter the `vserver show` command.
 - For detailed information, enter the `vserver show -instance` command.

Example

The following command displays basic information about the SVM:


```
vs1.example.com lif1      up/up      192.0.2.65/126  node0      e1b      false
                  lif2      up/up      192.0.2.1/62   node1      e0d      false
2 entries were displayed.
```

Where to find additional information for SVM administration

You can find details of the SVM administration tasks in the other clustered Data ONTAP documentation.

Depending on the administrative task you want to perform on the SVM, you can see the appropriate documentation.

If you want to...		See the following documentation...
Manage jobs and schedules		Managing jobs and schedules on page 177
Configure data access protocols		<ul style="list-style-type: none"> Clustered Data ONTAP 8.3 NFS File Access Reference Guide Clustered Data ONTAP 8.3 CIFS File Access Reference Guide Clustered Data ONTAP 8.3 SAN Administration Guide
Manage policies	Export policies	Clustered Data ONTAP 8.3 NFS File Access Reference Guide
	File policies	Clustered Data ONTAP 8.3 CIFS File Access Reference Guide SVMs with Infinite Volume do not support file policies.
	Quota policies	Clustered Data ONTAP 8.3 Logical Storage Management Guide SVMs with Infinite Volume do not support quota policies.
	SnapMirror policies and rules	Clustered Data ONTAP 8.3 Data Protection Guide
	Snapshot copy policies and schedules	Clustered Data ONTAP 8.3 Data Protection Guide
	Data policies in JSON format for SVMs with Infinite Volume	Clustered Data ONTAP 8.3 Infinite Volumes Management Guide
Manage data access security		<ul style="list-style-type: none"> Clustered Data ONTAP 8.3 NFS File Access Reference Guide Clustered Data ONTAP 8.3 CIFS File Access Reference Guide

If you want to...	See the following documentation...
Configure services	<ul style="list-style-type: none"> • <i>Clustered Data ONTAP 8.3 NFS File Access Reference Guide</i> • <i>Clustered Data ONTAP 8.3 CIFS File Access Reference Guide</i>
Manage storage	<i>Clustered Data ONTAP 8.3 Logical Storage Management Guide</i>
Manage LUNs	<i>Clustered Data ONTAP 8.3 SAN Administration Guide</i>
Manage data protection and backup	<i>Clustered Data ONTAP 8.3 Data Protection Guide</i>
Monitor SVM performance	<i>Monitoring cluster performance</i> on page 230
Manage SVM authentication	<i>Managing access to the cluster (cluster administrators only)</i> on page 116

Managing access to the cluster (cluster administrators only)

You can control access to the cluster and enhance security by managing user accounts, access-control roles and their password rules, public keys, digital certificates, web services, and audit settings.

Managing user accounts

You can create, modify, lock, unlock, or delete a cluster or Storage Virtual Machine (SVM) user account, reset a user's password, or display information for all user accounts.

You can manage cluster or SVM user accounts in the following ways:

- Creating a login method for a user by specifying the user's account name, associated SVM, the access method, and the authentication method
You can optionally specify the access-control role the user is assigned and add a comment about the user account.
The maximum number of cluster user accounts you can create is 100. This limit includes the Active Directory domain user accounts that are added to the cluster. There is no limit to the number of SVM user accounts you can create for an SVM.
- Displaying users' login information, such as the account name, allowed access method, authentication method, access-control role, account comment, and account status
- Displaying the user account that is used for logging in to the cluster in the current console or SSH session
- Displaying information about SNMP users, including the account name, the associated SVM, authentication method, hexadecimal engine ID, authentication protocol, privacy protocol, and security group
- Modifying the access-control role that is associated with a user's login method
It is best to use a single role for all 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

You use the `security login` commands to manage user accounts. You use the `security snmpusers` command to display information about SNMP users. For more information about these commands, see the appropriate man pages.

Note: The system prevents you from creating or using accounts with names that are reserved for the system (such as “root” and “naroot”). You cannot use a system-reserved name to access the cluster, an SVM, or the SP.

Related concepts

[Managing rule settings for user names and passwords in an access-control role](#) on page 130

Related tasks

[Customizing an access-control role to restrict user access to specific commands](#) on page 129

Access methods for user accounts

Data ONTAP provides several methods that you can use to specify how a user account can access the storage system.

The `-application` parameter of the `security login` commands specifies the method that a user can use to access the storage system. The supported access methods include the following:

- System console (**console**)
- HTTP or HTTPS (**http**)
- Data ONTAP API (**ontapi**)
- RSH (**rsh**)
RSH is disabled by default.
- Service Processor (**service-processor**)
- SNMP (**snmp**)
- SSH (**ssh**)
- Telnet (**telnet**)
Telnet is disabled by default.

Storage Virtual Machine (SVM) user accounts cannot use **console**, **rsh**, **service-processor**, or **telnet** as an access method.

If a firewall is enabled, the access method you use must also be added in the firewall policy to allow the access requests to go through the firewall. The `system services firewall policy show` command displays firewall policies.

Related concepts

[Accessing the cluster by using the CLI \(cluster administrators only\)](#) on page 11

Related references

[Commands for managing user accounts](#) on page 122

Related information

[Clustered Data ONTAP 8.3.2 man page: system services firewall policy show - Show firewall policies](#)

Authentication methods for user accounts

Data ONTAP provides several methods that you can use to specify how a user account is authenticated.

The `-authmethod` parameter of the `security login` commands specifies how a user account is authenticated. The following authentication methods are supported:

- SSL certificate authentication (**cert**)
- SNMP community strings (**community**)
- Windows Active Directory authentication (**domain**)
For Windows Active Directory authentication, a CIFS server must be created for the Storage Virtual Machine (SVM), and Windows domain users or groups must be mapped to access-control

roles by using the `security login create` command with the `-authmethod` parameter set to `domain` for the cluster and SVM access.

In addition, to authenticate Windows Active Directory domain users or groups for cluster access, a tunnel must be set up through a CIFS-enabled SVM.

- LDAP or NIS authentication (**nsswitch**)

LDAP and NIS authentication is supported for cluster and SVM user accounts. To use LDAP or NIS for authentication, the cluster or SVM must be configured for LDAP or NIS, and users must be mapped to access-control roles by using the `security login create` command with the `-authmethod` parameter set to `nsswitch`.

Data ONTAP supports only the RFC 2307 schema for LDAP authentication of SVM accounts. It does not support any other schemas, such as Active Directory Identity Management for UNIX (AD-IDMU) and Active Directory Services for UNIX (AD-SFU). Also, Data ONTAP supports only the MD5 and DES password encryption mechanisms for LDAP authentication of user accounts.

- User password (**password**)
- SSH public key authentication (**publickey**)
- SNMP user-based security model (**usm**)

Related tasks

[Enabling AD users and groups to access the cluster and SVMs](#) on page 118

[Enabling NIS or LDAP users to access the cluster](#) on page 120

Related references

[Commands for managing user accounts](#) on page 122

Authentication behavior when methods include both public key and password

When a user uses SSH to access the cluster or a Storage Virtual Machine (SVM) and the user account is configured with both the **publickey** and **password** authentication methods (the `-authmethod` parameter of the `security login` commands), the user is authenticated first with the public key.

If the public key authentication fails, the following occurs:

- Data ONTAP prompts the user to enter a password for authentication.
- If the password expiration functionality (the `-passwd-expiry-time` parameter of the `security login role config modify` command) is enabled and the user password has expired, Data ONTAP prompts the user to change the password before allowing the user to access the account.

Enabling AD users and groups to access the cluster and SVMs

You can enable the Active Directory (AD) domain users and groups to access the cluster and SVMs. Granting an AD group the access enables all AD users in that group to access the cluster or the specified SVM.

Before you begin

- The AD users or groups that you want to grant access must exist on the AD server.
- The cluster time must be kept within five minutes of the time on the AD domain controller (preferably using the same NTP servers) to enable users and groups of that domain to access the cluster or SVM.

Steps

1. If you are setting up AD users or groups for cluster access, complete one of the following steps:
 - If the cluster already has a data SVM with a CIFS server created, you can use that data SVM as an authentication tunnel by using the `security login domain-tunnel create` command with the `-vserver` parameter set to that data SVM.
The `security login domain-tunnel show` command displays the specified authentication tunnel.
 - If the cluster does not have a data SVM with a CIFS server created, you can use any data SVM in the cluster and join it to a domain by using the `vserver active-directory create` command with the `-vserver` parameter set to the data SVM.
Joining a data SVM to a domain does not create a CIFS server or require a CIFS license. However, it enables the authentication of AD users and groups at the SVM or cluster level.
2. Grant an AD user or group access to the cluster or SVM by using the `security login create` command with the `-authmethod` parameter set to `domain`.

The value of the `-user-or-group-name` parameter must be specified in the format of `domainname\username`, where `domainname` is the name of the CIFS domain server and `username` is the AD user or group that you want to grant access.

AD user authentication and AD group authentication support only `ssh` and `ontapi` for the `-application` parameter.

If the authentication tunnel is deleted, AD login sessions cannot be authenticated by the cluster, and AD users and groups cannot access the cluster. Open sessions that were authenticated prior to the deletion of the authentication tunnel remain unaffected.

Examples of enabling an AD user or group to access the cluster or SVM

The following example specifies the “vs1” data SVM as the tunnel that the cluster will use for authenticating an AD user or group, and then displays the authentication tunnel:

```
cluster1::> security login domain-tunnel create -vserver vs1
cluster1::> security login domain-tunnel show
Tunnel Vserver: vs1
```

The following command enables the “Administrator” AD user of the “DOMAIN1” domain to access the cluster through SSH:

```
cluster1::> security login create -vserver cluster1
-user-or-group-name DOMAIN1\Administrator -application ssh
-authmethod domain
```

The following command enables all users of the “group1” AD group in the “DOMAIN1” domain to access the cluster through SSH:

```
cluster1::> security login create -vserver cluster1
-user-or-group-name DOMAIN1\group1 -application ssh
-authmethod domain
```

The following command enables the “Administrator” AD user of the “DOMAIN1” domain to access the “vs1” SVM through SSH:

```
cluster1::> security login create -vserver vs1
-user-or-group-name DOMAIN1\Administrator -application ssh
-authmethod domain
```

The following command enables all users of the “group1” AD group in the “DOMAIN1” domain to access the “vs2” SVM through SSH:

```
cluster1::> security login create -vserver vs2
-user-or-group-name DOMAIN1\group1 -application ssh
-authmethod domain
```

Related concepts

[Managing the cluster time \(cluster administrators only\)](#) on page 163

Related tasks

[Installing a server CA certificate to authenticate an SSL server to which the cluster or SVM is a client](#) on page 144

Related information

[Clustered Data ONTAP 8.3 CIFS File Access Reference Guide](#)

Enabling NIS or LDAP users to access the cluster

If you store your user database on an LDAP or a NIS server, you can enable those users to access the cluster (admin SVM) by configuring LDAP or NIS for the cluster, including LDAP or NIS in the name service sources, and adding the `nsswitch` authentication method to the user accounts.

Steps

1. To configure LDAP for the cluster:
 - a. Create an LDAP client configuration for the cluster by using the `vserver services name-service ldap client create` command and specifying the `-vserver` parameter with the cluster name.

Example

The following example creates an LDAP client configuration named “corp” that makes anonymous binds to the LDAP servers with the IP addresses “172.160.0.100” and “172.16.0.101” for the cluster named “cluster1”:

```
cluster1::> vserver services name-service ldap client create
-vserver cluster1 -client-config corp -servers
172.16.0.100,172.16.0.101
```

The `vserver services name-service ldap client show` command displays the LDAP client configuration.

- b. Enable LDAP on the cluster by using the `vserver services name-service ldap create` command with the `-vserver cluster_name` and `-client-enabled true` parameters.

Example

The following example associates the “corp” LDAP client configuration with the “cluster1” cluster and enables the LDAP client on the cluster:

```
cluster1::> vserver services name-service ldap create
-vserver cluster1 -client-config corp -client-enabled true
```


The `vserver services name-service ldap show` command displays the LDAP configuration.

2. To configure a NIS domain for the cluster, use the `vserver services name-service nis-domain create` command with the `-vserver cluster_name` parameter.

Example

The following example creates a NIS domain configuration for the “cluster1” cluster. The NIS domain is named “nisdomain”, which is active upon creation and uses an NIS server with the IP address “192.0.2.180”:

```
cluster1::> vserver services name-service nis-domain create
-vserver cluster1 -domain nisdomain -active true -servers 192.0.2.180
```

The `vserver services name-service nis-domain show` command displays the NIS domain configuration.

3. Modify the look-up order of the name service sources as necessary for the cluster by using the `vserver services name-service ns-switch modify` command with the `-database passwd`, `-vserver cluster_name`, and `-sources` parameters.

The `passwd` database supports `files`, `ldap`, and `nis` as valid name service sources. The sequence in which the sources are specified in the `-sources` parameter determines their look-up order.

Example

The following example modifies the order of the name service sources for the `passwd` database on the “cluster1” cluster. The order of looking up the sources to use for the cluster user authentication is local files, then LDAP, then NIS:

```
cluster1::> vserver services name-service ns-switch modify
-vserver cluster1 -database passwd -sources files,ldap,nis
```

The `vserver services name-service ns-switch show` command displays the name service switch configuration.

4. Configure `nsswitch` as an authentication method for an LDAP or a NIS user account to access the cluster by using the `security login create` command with the `-authmethod nsswitch` and `-vserver cluster_name` parameters.

Example

The following example enables the user named “john” in the LDAP or NIS server to access the “cluster1” cluster by using SSH:

```
cluster1::> security login create -vserver cluster1
-user-or-group-name john -application ssh -authmethod nsswitch
```

The `security login show` command displays user login methods.

Related information

[Clustered Data ONTAP 8.3 NFS File Access Reference Guide](#)

[Clustered Data ONTAP 8.3.2 man page: vserver services name-service ldap client create - Create an LDAP client configuration](#)

[Clustered Data ONTAP 8.3.2 man page: vserver services name-service ldap create - Create an LDAP configuration](#)

Clustered Data ONTAP 8.3.2 man page: vserver services name-service nis-domain create - Create a NIS domain configuration

Clustered Data ONTAP 8.3.2 man page: vserver services name-service ns-switch modify - Change a Name Service Switch table entry

Clustered Data ONTAP 8.3.2 man page: security login create - Add a login method

Commands for managing user accounts

You use the `security login` and `security snmpusers` commands to manage user accounts.

If you want to...	Use this command...
Create a login method for user access of the cluster or an SVM with a supported application type or authentication method	<code>security login create</code>
Display information about user accounts created for accessing the cluster or an SVM	<code>security login show</code>
Display the user account that is used for logging in to the cluster in the current console or SSH session	<code>security login whoami</code>
Display information about SNMP users	<code>security snmpusers</code>
Modify the access-control role of a user's login method created for accessing the cluster or an SVM	<code>security login modify</code> Note: It is best to use a single role for all access and authentication methods of a user account.
Delete a user login method created for accessing the cluster or an SVM	<code>security login delete</code>
Change a user password	<code>security login password</code> Note: When you change the password of your own user account, Data ONTAP prompts you to enter your old password first. If you do not have the old password, you can ask another cluster administrator of the “admin” role to run this command to reset your password.
Lock a user account	<code>security login lock</code> Note: Data ONTAP requires that at least one cluster user account with the “admin” role capability and the <code>console</code> application type remain unlocked.
Unlock a user account	<code>security login unlock</code>
Specify a CIFS-enabled Storage Virtual Machine (SVM) that you want to use as the tunnel for authenticating Active Directory domain users' cluster access	<code>security login domain-tunnel create</code>
Modify the tunnel that is used for Active Directory domain user authentication	<code>security login domain-tunnel modify</code>

If you want to...	Use this command...
Display the tunnel that is used for Active Directory domain user authentication	<code>security login domain-tunnel show</code>
Delete the tunnel that is used for Active Directory domain user authentication	<code>security login domain-tunnel delete</code>

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Managing access-control roles

Role-based access control (RBAC) limits users' administrative access to the level granted for their role, enabling you to manage users by the role they are assigned to. Data ONTAP provides several predefined roles. You can also create additional access-control roles, modify them, delete them, or specify account restrictions for users of a role.

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
- Modifying an access-control role's account restrictions and settings for user names and passwords
- Displaying the current settings for the restrictions on an access-control role or user account
- Displaying Data ONTAP APIs and their corresponding CLI commands

Data ONTAP prevents you from modifying predefined roles.

You use the `security login role` and `security login role config` commands to manage access-control roles.

Predefined roles for cluster administrators

Data ONTAP provides several predefined roles for cluster user accounts. You can also create additional roles.

The following table describes the Data ONTAP predefined roles and their levels of access to command directories:

This role...	Has this level of access...	To the following command directory or directories...
<code>admin</code>	<code>all</code>	All command directories (DEFAULT)

This role...	Has this level of access...	To the following command directory or directories...
autosupport	all	<ul style="list-style-type: none"> • set • system node autosupport
	none	All other command directories (DEFAULT)
backup	all	vserver services ndmp
	readonly	volume
	none	All other command directories (DEFAULT)
readonly	all	<ul style="list-style-type: none"> • security login password • set
	none	security
	readonly	All other command directories (DEFAULT)
none	none	All command directories (DEFAULT)

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

Note: The predefined `autosupport` role includes a predefined `autosupport` account. The `autosupport` role and `autosupport` account are used by AutoSupport OnDemand. Data ONTAP prevents you from modifying or deleting the `autosupport` account. It also prevents you from adding additional user accounts to the `autosupport` role.

Predefined roles for SVM administrators

The five predefined roles for an SVM administrator are: `vsadmin`, `vsadmin-volume`, `vsadmin-protocol`, `vsadmin-backup`, and `vsadmin-readonly`. In addition to these predefined roles, you can create customized SVM administrator roles by assigning a set of capabilities.

Each SVM can have its own user and administration authentication domain. You can delegate the administration of an SVM to an SVM administrator after creating the SVM and user accounts.

Note: SVMs with Infinite Volume do not support quotas, qtrees, and LUNs. Therefore, an SVM administrator cannot perform the tasks related to quotas, qtrees, and LUNs on an SVM with Infinite Volume.

The following table lists the predefined roles for an SVM administrator and the respective capabilities:

SVM administrator role name	Description
vsadmin	<p>This role is the superuser role for the SVM and is assigned by default. The SVM administrator with this role has the following capabilities:</p> <ul style="list-style-type: none"> • Managing own user account local password and key information • Managing volumes, quotas, qtrees, Snapshot copies, and files • Managing LUNs • Configuring protocols: NFS, CIFS, iSCSI, and FC (FCoE included) • Configuring services: DNS, LDAP, and NIS • Monitoring jobs • Monitoring network connections and network interface • Monitoring the health of the SVM <p>The vsadmin role is assigned by default.</p>
vsadmin-volume	<p>The SVM administrator with this role has the following capabilities:</p> <ul style="list-style-type: none"> • Managing own user account local password and key information • Managing volumes, quotas, qtrees, Snapshot copies, and files • Managing LUNs • Configuring protocols: NFS, CIFS, iSCSI, and FC (FCoE included) • Configuring services: DNS, LDAP, and NIS • Monitoring network interface • Monitoring the health of the SVM
vsadmin-protocol	<p>The SVM administrator with this role has the following capabilities:</p> <ul style="list-style-type: none"> • Managing own user account local password and key information • Configuring protocols: NFS, CIFS, iSCSI, and FC (FCoE included) • Configuring services: DNS, LDAP, and NIS • Managing LUNs • Monitoring network interface • Monitoring the health of the SVM

SVM administrator role name	Description
vsadmin-backup	<p>The SVM administrator with this role has the following capabilities:</p> <ul style="list-style-type: none"> • Managing own user account local password and key information • Managing NDMP operations • Making a restored volume read/write • Managing SnapMirror relationships and Snapshot copies • Viewing volumes and network information
vsadmin-readonly	<p>The SVM administrator with this role has the following capabilities:</p> <ul style="list-style-type: none"> • Managing own user account local password and key information • Monitoring the health of the SVM • Monitoring network interface • Viewing volumes and LUNs • Viewing services and protocols

Considerations for customizing an access-control role

Data ONTAP provides predefined access-control roles for cluster and Storage Virtual Machine (SVM) administrators. You can create additional access-control roles for the cluster or an SVM and customize their access to certain commands or command directories. Several considerations apply when you customize a role for specific access needs.

Syntax considerations

- An access-control role must include one or more rules (specified by the `security login role create` command) that include the following elements:
 - SVM name (`-vserver`)
This is the name of the admin SVM (the cluster) or data SVM that the role belongs to.
 - Role name (`-role`)
 - Capability (`-cmddirname`)
The capability is a command (*intrinsic* or *nonintrinsic*) or command directory for which you want to specify an access level for the role.
In the context of customizing a role, an *intrinsic command* is any command that ends with `create`, `modify`, `delete`, or `show`. All other commands are called *nonintrinsic commands*.
 - Access level (`-access`)
The access level can be **all**, **readonly**, or **none**.
How you specify the access level depends on whether the granted capability is a command or a command directory, and if it is a command, whether the command is intrinsic or nonintrinsic.

- When you specify a role's access for a command directory, the access by default applies to all the subdirectories and all the commands in the directory and subdirectories:

If the capability you grant to a role is...	And the access level you specify is...	Then the effect is...
A command directory	all	The role can access the specified directory and its subdirectories (if any), and the role can execute all commands in the directory or subdirectories.
	readonly	The role has read-only access to the specified directory and its subdirectories (if any). This combination results in the role's access to only the <code>show</code> command in the specified directory and subdirectories. All other commands in the directory are not accessible to the role.
	none	The role has no access to the specified directory, its subdirectories, or commands.

For example, the following command grants the “vol_role” role of the “vs1” SVM **all** access to the `volume` directory, all its subdirectories, and the commands in the directory and subdirectories:

```
security login role create -vserver vs1 -role vol_role -cmddirname "volume" -
access all
```

- Subdirectory access, if specified, overrides parent directory access.
If a parent directory has an access level and its subdirectory is specified with a different access level, the access level specified for the subdirectory overrides that of the parent directory.
For example, the following commands grant the “vol_role” role of the “vs1” SVM **all** access to the commands in the `volume` directory and subdirectories, except for the `volume snapshot` subdirectory, to which the role is restricted to **readonly** access:

```
security login role create -vserver vs1 -role vol_role -cmddirname "volume" -
access all

security login role create -vserver vs1 -role vol_role -cmddirname "volume
snapshot" -access readonly
```

- The access level you can specify for a command depends on whether the command is intrinsic or nonintrinsic:

If the capability you grant to a role is...	And the access level you specify is...	Then the effect is...
An intrinsic command (a command ending with <code>create</code> , <code>modify</code> , <code>delete</code> , or <code>show</code>)	all	An invalid combination. You cannot specify an access level on an intrinsic command; you must specify the access level on the <i>directory</i> of an intrinsic command.
	readonly	
	none	

If the capability you grant to a role is...	And the access level you specify is...	Then the effect is...
A nonintrinsic command	all	The role can execute the specified command.
	readonly	An invalid combination. You cannot grant readonly access at the command level; you must specify it at the <i>directory</i> level.
	none	The role has no access to the specified command.

For example, the following command enables the “ssl_role” role of the “vs1” SVM to access the `security ssl show` command but no other commands in the `security ssl` directory:

```
security login role create -vserver vs1 -role ssl_role -cmddirname "security
ssl" -access readonly
```

In the following example, the first four commands use command directories to restrict the access of the “login_role” role of the “cluster1” cluster to the `security login show intrinsic` command, and the last two commands grant the role additional access to the `security login password` and `security login role show-ontapi nonintrinsic` commands. The role has no access to other commands in the `security login` directory:

```
security login role create -vserver cluster1 -role login_role -cmddirname
"security login" -access readonly

security login role create -vserver cluster1 -role login_role -cmddirname
"security login domain-tunnel" -access none

security login role create -vserver cluster1 -role login_role -cmddirname
"security login publickey" -access none

security login role create -vserver cluster1 -role login_role -cmddirname
"security login role" -access none

security login role create -vserver cluster1 -role login_role -cmddirname
"security login password" -access all

security login role create -vserver cluster1 -role login_role -cmddirname
"security login role show-ontapi" -access all
```

- For a customized role, the commands and command directories for which you do not specify an access level have the default level of **none**, and the role has no access to unspecified commands or command directories.

General considerations

- Data ONTAP prevents you from modifying predefined roles.
- It is recommended that you grant a customized role **all** access to the `security login password` command to enable users of the role to modify their passwords. For example, the following command grants the “guest_role” role of the “vs1” SVM the capability to modify account passwords:

```
security login role create -vserver vs1 -role guest_role -cmddirname "security
login password" -access all
```

- You cannot grant an SVM role any access to a command or command directory that is available to only the cluster administrator. For example, you cannot grant an SVM role the access to the `system license` directory or its commands, because the capability for managing licenses is available to only the cluster

administrator. For information about whether the SVM administrator has access to a specific command, see the man pages.

Related tasks

[Customizing an access-control role to restrict user access to specific commands](#) on page 129

Customizing an access-control role to restrict user access to specific commands

The cluster administrator can restrict a user's access to only specific commands by customizing an access-control role with specified commands and mapping the user account to the role.

Steps

1. Create a customized access-control role that is restricted to only the specified command or commands by using the `security login role create` command with the `-cmddirname` parameter.

The `security login role show` command displays the commands that a role can access.

2. Create a login method for a user account and map it to the customized role by using the `security login create` command with the `-role` parameter.

Examples of customizing an access-control role to restrict user account access

The following example creates an access-control role named “vol_snapshot”, which has access to only the `volume snapshot` commands, and a “vs1” Storage Virtual Machine (SVM, formerly known as Vserver) user account named “snapshot_admin”, which is assigned the “vol_snapshot” role. The user has full access to the `volume snapshot` commands, as defined by the role. The user can use SSH to access the SVM and a password for authentication.

```
cluster1::> security login role create -vserver vs1 -role vol_snapshot
-cmddirname "volume snapshot"

cluster1::> security login role show -vserver vs1 -role vol_snapshot
Vserver      Role      Command/      Access
Name         Name      Directory      Query Level
-----
vs1          vol_snapshot  DEFAULT      none
vs1          vol_snapshot  volume snapshot  all
2 entries were displayed.

cluster1::> security login create -vserver vs1 -user-or-group-name snapshot_admin
-application ssh -authmethod password -role vol_snapshot

Please enter a password for user 'snapshot_admin':
Please enter it again:

cluster1::>
```

The following example creates an access-control role name “sec_login_readonly”. The role is customized to have read-only access to the `security login` directory but no access to the `security login domain-tunnel`, `security login publickey`, or `security login role` subdirectories. As a result, the role can access only the `security login show` command. A cluster user account named “new_admin” is then created and assigned the “sec_login_readonly” role. The user can use the console to access the cluster and a password for authentication.

```
cluster1::> security login role create -vserver cluster1 -role sec_login_readonly
-cmddirname "security login" -access readonly

cluster1::> security login role create -vserver cluster1 -role sec_login_readonly
-cmddirname "security login domain-tunnel" -access none

cluster1::> security login role create -vserver cluster1 -role sec_login_readonly
-cmddirname "security login publickey" -access none
```

```

cluster1::> security login role create -vserver cluster1 -role sec_login_readonly
-cmdirname "security login role" -access none

cluster1::> security login role show -vserver cluster1 -role sec_login_readonly
(security login role show)

```

Vserver	Role Name	Command/Directory	Access Query Level
cluster1	sec_login_readonly	DEFAULT	none
cluster1	sec_login_readonly	security login	readonly
cluster1	sec_login_readonly	security login domain-tunnel	none
cluster1	sec_login_readonly	security login publickey	none
cluster1	sec_login_readonly	security login role	none

```

5 entries were displayed.

cluster1::> security login create -vserver cluster1 -user-or-group-name new_admin
-application console -authmethod password -role sec_login_readonly

Please enter a password for user 'new_admin':
Please enter it again:

cluster1::>

```

Related concepts

[Managing user accounts](#) on page 116

[Considerations for customizing an access-control role](#) on page 126

Related references

[Commands for managing user accounts](#) on page 122

[Commands for managing access-control roles](#) on page 132

Managing rule settings for user names and passwords in an access-control role

The default rules for user names and passwords apply to users of all access-control roles. You can modify the rule settings of user names and passwords for a specific role to enhance user account security.

Following are the default rules for user names:

- A user name must be at least three characters long.
- A user name can contain letters, numbers, special characters, or a combination of them. For a local user name (that is, a user name that is configured with the **password** or **publickey** authentication method), the following additional rules about special characters apply:
 - Only the following characters are supported:


```

_ . -
          
```
 - The user names cannot begin with a hyphen (-).
- A user name that is configured with the **password** authentication method cannot be longer than 16 characters.
- A user name that is configured with the **snmp** application type cannot be longer than 32 characters.

Following are the default rules for passwords:

- A password cannot contain the user name.
- A password must be at least eight characters long.
- A password must contain at least one letter and one number.

- A password cannot be the same as the last six passwords.

To enhance user account security, you can use parameters of the `security login role config modify` command to modify the following settings of an access-control role:

- Rule settings for user names:
 - The required minimum length of a user name (`-username-minlength`)
 - Whether a mix of alphabetic and numeric characters is required in a user name (`-username-alphanum`)
- Rule settings for passwords:
 - The required minimum length of a password (`-passwd-minlength`)
 - Whether a mix of alphabetic and numeric characters is required in a password (`-passwd-alphanum`)
 - The required number of special characters in a password (`-passwd-min-special-chars`)
 - Whether users must change their passwords when logging in to their accounts for the first time (`-require-initial-passwd-update`)
Users can make initial password changes only through SSH or serial-console connections.
 - The number of previous passwords that cannot be reused (`-disallowed-reuse`)
 - The minimum number of days that must pass between password changes (`-change-delay`)
 - The number of days after which a password expires (`-passwd-expiry-time`)
- Rule settings about invalid login attempts:
 - The number of invalid login attempts that triggers the account to be locked automatically (`-max-failed-login-attempts`)
When the number of a user's invalid login attempts reaches the value specified by this parameter, the user's account is locked automatically.
The `security login unlock` command unlocks a user account.
 - The number of days for which an account is locked if invalid login attempts reach the allowed maximum (`-lockout-duration`)

You can display the current settings for the rules by using the `security login role config show` command. For information about the `security login role config` commands and the default settings, see the man pages.

Related references

[Commands for managing access-control roles](#) on page 132

Considerations for password rule settings

Some password rule settings require that users of a role change their passwords. To enable users to change passwords, the user accounts must have a proper access method, and their role must have the privilege to run the password reset command.

Users of a role are required to change their passwords in either of the following situations:

- The role's password settings require that users change their passwords when logging in to their accounts for the first time.
This setting is defined by the `-require-initial-passwd-update` parameter of the `security login role config modify` command.

- The role is set up to have user passwords expire by a certain time. This setting is defined by the `-passwd-expiry-time` parameter of the `security login role config modify` command.

To enable users to change their passwords, the following conditions must be met:

- Users must be granted SSH or console access. Passwords can be changed by their account users only through SSH or console connections. The `-application` parameter of the `security login modify` command grants a user the specified access method.

Note: Console access is not supported for Storage Virtual Machine (SVM) user accounts.
- Users' role must have the privilege to run the command for changing the password (`security login password` command). The `-cmddirname` parameter of the `security login role modify` command grants a role the privilege to run a command or command directory.

Related concepts

[Access methods for user accounts](#) on page 117

Related tasks

[Customizing an access-control role to restrict user access to specific commands](#) on page 129

Commands for managing access-control roles

You use the `security login role` commands to control the level of access users in a role have to the system. You use the `security login role config` commands to manage rule settings of user names and passwords for a role to enhance user account security.

If you want to...	Use this command...
Create an access-control role and specify the command or command directory that the role can access	<code>security login role create</code>
Modify the command or command directory that an access-control role can access	<code>security login role modify</code>
Display information about access-control roles	<code>security login role show</code>
Display Data ONTAP APIs and their corresponding CLI commands	<code>security login role show-ontapi</code>
Delete an access-control role	<code>security login role delete</code>

If you want to...	Use this command...
<p>Modify the following account restrictions and rule settings for an access-control role:</p> <ul style="list-style-type: none"> • The required minimum length of a user name • Whether a mix of alphabetic and numeric characters is required in a user name • The required minimum length of a password • Whether a mix of alphabetic and numeric characters is required in a password • The required number of special characters in a password • Whether users must change their passwords when logging in to their accounts for the first time • The number of previous passwords that cannot be reused • The minimum number of days that must pass between password changes • The number of days after which a password expires • The number of invalid login attempts that triggers the account to be locked automatically • The number of days for which an account is locked if invalid login attempts reach the allowed maximum 	<p><i>security login role config modify</i></p>
<p>Display user account restrictions and rule settings</p>	<p><i>security login role config show</i></p>

If you want to...	Use this command...
<p>Reset the following settings to their default values:</p> <ul style="list-style-type: none"> The required number of special characters in a password (<code>-passwd-min-special-chars 0</code>) Whether users must change their passwords when logging in to their accounts for the first time (<code>-require-initial-passwd-update disabled</code>) The number of days after which a password expires (<code>-passwd-expiry-time unlimited</code>) The number of invalid login attempts that triggers the account to be locked automatically (<code>-max-failed-login-attempts 0</code>) The number of days for which an account is locked if invalid login attempts reach the allowed maximum (<code>-lockout-duration 0</code>) <p>Data ONTAP prompts you to run this command if you revert to Data ONTAP 8.1.2 or earlier.</p>	<p><code>security login role config reset</code> (advanced privilege level)</p>

Related information

Clustered Data ONTAP 8.3 Commands: Manual Page Reference

Managing SSH security configurations

Managing SSH security configurations involves managing the SSH key exchange algorithms and data encryption algorithms (also known as *ciphers*). Data ONTAP enables you to enable or disable individual SSH key exchange algorithms and ciphers for the cluster or Storage Virtual Machines (SVMs) according to their SSH security requirements.

Data ONTAP supports the following SSH security configurations for the cluster and SVMs:

- The following SSH key exchange algorithms are supported and enabled by default:
 - The `diffie-hellman-group-exchange-sha256` SSH key exchange algorithm for SHA-2
 - The `diffie-hellman-group-exchange-sha1`, `diffie-hellman-group14-sha1`, and `diffie-hellman-group1-sha1` SSH key exchange algorithms for SHA-1

SHA-2 algorithms are more secure than SHA-1 algorithms. Data ONTAP, which serves as an SSH server, automatically selects the most secure SSH key exchange algorithm that matches the client. To further enhance SSH security, you can manually disable the SHA-1 algorithms and leave only the SHA-2 algorithm enabled.

- For ciphers, the following counter (CTR) mode and cipher block chaining (CBC) mode of the AES and 3DES symmetric encryptions are supported and enabled by default:
 - `aes256-ctr`

- `aes192-ctr`
- `aes128-ctr`
- `aes256-cbc`
- `aes192-cbc`
- `aes128-cbc`
- `3des-cbc`

The CTR mode ciphers are more secure than the CBC mode ciphers. Among ciphers of the same mode, the higher the key size, the more secure the cipher. Of the ciphers supported by Data ONTAP, `aes256-ctr` is the most secure, and `3des-cbc` is the least secure.

You can manage the SSH key exchange algorithms and ciphers for the cluster and SVMs in the following ways:

- Display the current configurations of SSH key exchange algorithms and ciphers (`security ssh show`)
The enabled SSH key exchange algorithms are displayed in the order of decreasing security strengths.
The enabled CTR mode ciphers (more secure) are displayed before the CBC mode ciphers (less secure). Within each mode type, the ciphers are displayed in decreasing key size.
- Replace the current configurations of the SSH key exchange algorithms or ciphers with the configuration settings you specify (`security ssh modify`)
If you modify the SSH key exchange algorithm or cipher configurations for the cluster, the changes apply also to all subsequently created SVMs.
- Add SSH key exchange algorithms or ciphers to the current configurations (`security ssh add`)
The added SSH key exchange algorithms or ciphers are enabled.
If you add SSH key exchange algorithms or ciphers to the cluster configuration, the changes apply also to all subsequently created SVMs.
- Remove the specified SSH key exchange algorithms or ciphers from the current configurations (`security ssh remove`)
The removed SSH key exchange algorithms or ciphers are disabled.
If you remove SSH key exchange algorithms or ciphers from the cluster configuration, the changes apply also to all subsequently created SVMs.
Data ONTAP prevents you from removing all SSH key exchange algorithms or all ciphers from the cluster or an SVM.
- Restore the configurations of SSH key exchange algorithms and ciphers of the cluster and all SVMs to the settings that were used prior to Data ONTAP 8.2.1 (`security ssh prepare-to-downgrade`, available at the advanced privilege level)
If you downgrade or revert to a release earlier than Data ONTAP 8.2.1, Data ONTAP prompts you to run this command to reset the SSH security configurations of the cluster and all SVMs to the following default settings of the earlier release:
 - All supported SHA-1 and SHA-2 SSH key exchange algorithms are reset to the enabled state.
 - All CTR and CBC modes of data encryption algorithms are reset to the enabled state.

Attention: Do not execute the `security ssh prepare-to-downgrade` command except for the downgrade or revert to a release earlier than Data ONTAP 8.2.1. Otherwise, the SSH configuration functionality is permanently disabled and Data ONTAP does not enable you to manage the SSH configuration settings.

Commands for managing SSH security configurations

You use the `security ssh` commands to manage the SSH security configurations of a cluster or Storage Virtual Machines (SVMs), including displaying, replacing, adding, removing, and restoring the SSH key exchange algorithms and data encryption algorithms (ciphers).

If you want to...	Use this command...
Display the current configurations of the SSH key exchange algorithms and ciphers for the cluster and SVMs	<code>security ssh show</code>
Replace the current configurations of the SSH key exchange algorithms or ciphers for the cluster or SVM with the configuration settings you specify	<code>security ssh modify</code>
Add SSH key exchange algorithms or ciphers to the current configurations for the cluster or SVM	<code>security ssh add</code>
Remove the specified SSH key exchange algorithms or ciphers from the current configurations of the cluster or SVM	<code>security ssh remove</code>
Restore the configurations of SSH key exchange algorithms and ciphers of the cluster and all SVMs to the settings that were used prior to Data ONTAP 8.2.1	<code>security ssh prepare-to-downgrade</code> (advanced privilege level) Data ONTAP prompts you to run this command if you downgrade or revert to a release earlier than Data ONTAP 8.2.1.

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Managing public keys

You can associate, modify, or delete a public key to manage a user's authentication.

You can manage public keys in the following ways:

- Adding a public key by associating an existing public key in a valid OpenSSH format with a user account
Multiple public keys are allowed for a user account.
- Loading a public key from a universal resource identifier (URI), such as FTP or HTTP, and associating it with a user account
You can also overwrite an existing public key with the one you are loading.
- Displaying information about public keys
- Modifying a public key that is associated with a specific user
- Deleting a public key that is associated with a specific user

To create or modify a public key or load a public key from a URI, your user account must be configured with the **publickey** login method (`security login create` command with the `-authmethod` parameter set to **publickey**).

You use the `security login publickey` commands to manage public keys. For information about these commands, see the appropriate man pages.

Commands for managing public keys

You use the `security login publickey` commands to manage public keys.

If you want to...	Use this command...
Associate an existing public key with a user account	<code>security login publickey create</code>
Load a public key from a URI and associate it with a user	<code>security login publickey load-from-uri</code>
Display information about public keys	<code>security login publickey show</code>
Modify a public key for a specific user	<code>security login publickey modify</code>
Delete a public key for a specific user	<code>security login publickey delete</code>

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Managing digital certificates for server or client authentication

A digital certificate ensures that communications are transmitted in encrypted form and that information is sent privately and unaltered to only the specified server or from the authenticated client. You can generate a certificate signing request, create, install, sign, display, revoke, or delete a digital certificate for server or client authentication.

A digital certificate, also called a *public key certificate*, is an electronic document that verifies the owner of a public key. It can be either self signed (by the owner) or Certificate Authority (CA) signed. You can provide server or client authentication by using digital certificates for situations where the cluster or Storage Virtual Machine (SVM) is an SSL server or client. When you provide both server and client authentication, you have mutual authentication (also called *two-way authentication*) in which both the server and the client present their certificates to each other for validating their respective identities to each other.

You can manage digital certificates in the following ways (the `security certificate` command family):

- You can create and install a self-signed digital certificate.
- You can generate a digital certificate signing request (CSR) that will be sent to a CA for signing.
- You can sign a digital certificate using a self-signed root CA.
- You can install a CA-signed digital certificate and the public key certificate of the root CA.
- You can display information about created or installed digital certificates.
- You can display digital certificates that are signed by the cluster or SVM as the CA.
- You can revoke a digital certificate signed by the cluster or SVM as the CA, if the certificate becomes compromised.
- You can delete a digital certificate that is no longer needed.

The following behaviors and default settings apply:

- When the cluster or SVM is created, Data ONTAP automatically creates a self-signed digital certificate for authenticating the cluster or SVM as a server.
- By default, Data ONTAP uses the SHA256 cryptographic hashing function for signing a CSR or digital certificate.
- By default, private keys generated by Data ONTAP are 2048-bit.
- By default, digital certificates created by Data ONTAP are set to expire in 365 days, but you can specify the expiration setting when you create a digital certificate.
- By default, SSL server authentication is enabled, but SSL client authentication is disabled. The `security ssl modify` command enables or disables SSL authentication of the cluster or SVM as an SSL server and that of its client. The `-server-enabled` parameter defaults to **true**, and the `-client-enabled` parameter defaults to **false**. Setting the `-client-enabled` parameter to **true** enables mutual authentication of the server (the cluster or SVM) and its client.

When you manage digital certificates, the following types and subtype apply:

- The certificate type (the `-type` parameter) specifies what a certificate is used for. The certificate type can be one of the following:
 - **server** is a certificate that authenticates the cluster or SVM as an SSL server.
 - **client** is a certificate that authenticates the cluster or SVM as an SSL client.
 - **server-ca** is a root certificate of an SSL server to which the cluster or SVM is a client.
 - **client-ca** is a root certificate of an SSL client to which the cluster or SVM is a server.
 - **root-ca** is a self-signed root CA certificate that enables the cluster or SVM to act as a CA. When you create a **root-ca** certificate, a **client-ca** certificate and a **server-ca** certificate are also created automatically. When you delete the **root-ca** certificate, the corresponding **client-ca** and **server-ca** certificates are also deleted automatically.
- The certificate subtype of Key Management Interoperability Protocol (KMIP) (the `-subtype kmip-cert` parameter), along with the **client** and **server-ca** types, specifies that the certificate is used for mutually authenticating the cluster and an external key manager such as a KMIP server. KMIP certificates are supported at the cluster level but not the data SVM level. At the cluster level, they are supported with the installation, deletion, and display operations (`security certificate install`, `security certificate delete`, and `security certificate show`, respectively) but not the creation operation (`security certificate create`).

Related concepts

[Managing the web protocol engine](#) on page 152

Related tasks

[Configuring access to web services](#) on page 157

Installing a server certificate to authenticate the cluster or SVM as an SSL server

To enable the cluster or Storage Virtual Machine (SVM) to be authenticated as an SSL server, you install a digital certificate with the **server** type on the cluster or SVM. The certificate you install can be self signed or CA signed.

About this task

When the cluster or SVM is created, a self-signed server certificate is created automatically and uses the cluster or SVM name as the common name. The corresponding SSL server authentication is enabled and also uses the default common name for the cluster or SVM.

If you want the cluster or SVM to use a different common name or a CA-signed certificate for server authentication, you can create or install additional server certificates. You can also modify SSL configuration to use a server certificate that you specify.

Steps

1. To create a self-signed digital certificate for server authentication, use the `security certificate create` command with the `-type server` parameter.
2. To use a third-party CA-signed digital certificate for server authentication, complete the following steps:
 - a. Generate a digital certificate signing request (CSR) by using the `security certificate generate-csr` command.

The system displays the CSR output. The output includes a certificate request and a private key. You should keep a copy of the private key.
 - b. Copy the certificate request from the CSR output and send it in an electronic form (such as email) to a trusted third-party CA for signing.

After processing your request, the CA sends you the signed digital certificate. You should keep a copy of the private key and the CA-signed digital certificate.
 - c. Install the third-party CA-signed digital certificate by using the `security certificate install` command with the `-type server` parameter.
 - d. Enter the certificate and the private key when you are prompted, and then press Enter.
 - e. When Data ONTAP asks you whether you want to install the CA root and intermediate certificates that form the certificate chain of the server certificate, enter Y.
 - f. Enter any additional root or intermediate certificates when you are prompted, and then press Enter.

You install the certificates of the CA to form a certificate chain of the server certificate. The chain starts with the certificate of the CA that issued the server certificate, and it can range up to the root certificate of the CA. Any missing intermediate certificates will result in the failure of server certificate installation.

After the CA certificates are entered, the certificates chain is installed as **server-chain** along with the **server** certificate type.
3. To use a self CA-signed digital certificate for server authentication (with the cluster or SVM being the signing CA), complete the following steps:
 - a. Generate a CSR by using the `security certificate generate-csr` command.

The system displays the CSR output. The output includes a certificate request and a private key. You should keep a copy of the private key.

- b. Create a self-signed root CA certificate for the cluster or SVM by using the `security certificate create` command with the `-type root-ca` parameter.
 - c. Display the root CA certificate by using the `security certificate show` command with the `-instance` and `-type root-ca` parameters.
You will need the following information from the command output for signing the CSR:
 - Certificate authority (CA)
 - Serial number of the certificate
 - d. Sign the CSR with the root CA by using the `security certificate sign` command.
 - e. When you are prompted, enter the CSR and then press ENTER.
 - f. Install the self CA-signed digital certificate by using the `security certificate install` command with the `-type server` parameter.
 - g. Enter the certificate and the private key when you are prompted, and then press Enter.
 - h. When Data ONTAP asks you whether you want to install the CA root and intermediate certificates that form the certificate chain of the server certificate, enter N.
4. If you want to modify the SSL configuration to specify the certificate for server authentication, use the `security ssl modify` command with the `-ca` and the `-serial` parameters.

Examples of installing a server certificate to authenticate the cluster or SVM as an SSL server

The following example creates a self-signed server certificate for the “vs1” SVM at a company whose custom common name is `lab.companyname.com`. The certificate is for authenticating the “vs1” SVM as an SSL server:

```
cluster1::> security certificate create -vserver vs1 -common-name
lab.companyname.com -type server
```

The following command creates a CSR with a 2048-bit private key for use by the Software group in the IT department of a company whose custom common name is `server1.companyname.com`, located in Sunnyvale, California, USA. The email address of the contact administrator who manages the SVM is `web@companyname.com`. The system displays the CSR and the private key in the output:

```
cluster1::> security certificate generate-csr -common-name
server1.companyname.com
-size 2048 -country US -state CA -locality Sunnyvale
-organization IT -unit Software -email-addr web@companyname.com
```

```
Certificate Signing Request:
-----BEGIN CERTIFICATE REQUEST-----
MIICrjCCAZYCAQMwateQMA4GA1UEAxMHcnRwLmNvbTELMakGAlUEBhMCVVMxCzAJ
BgNVBAGTAK5DMQwwCgYDVQQHEwNSVFaxDTALBgNVBAoTBGNvcmlUxDTALBgNVBAsT
BGNvcmlUxDTALBgNVBAsTANBgkqhkiG9w0BCQEWADEwDQYJKoZIhvcNAQEBBQADggEPADCC
...
-----END CERTIFICATE REQUEST-----
```

```
Private Key:
-----BEGIN RSA PRIVATE KEY-----
MIIBPAIBAAJBAM16ytrK8nQj82UsWeHOeT8gk0BPX+Y5MLyCsUdXA7hXhumHNpvF
C61X2G32Sx8VEalth94tx+vOEzq+UaqHlt0CAwEAAQJBAMZjDw1gmlm3qIr/n8VT
PFnnZnbVcXVM70tbUsgPKw+QCCh9dF1jmuQKeDr+wUMWkn1DeGrfhILpzfJGHRlLJ
...
-----END RSA PRIVATE KEY-----
```

Note: Please keep a copy of your private key and certificate request for future reference.

The following command installs a CA-signed server certificate for the “vs1” SVM. The certificate is for authenticating the “vs1” SVM as an SSL server:

```
cluster1::> security certificate install -vserver vs1 -type server

Please enter Certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIIB8TCCA ZugAwIBAwIBADANBgkqhkiG9w0BAQQFADBFMRMwEQYDVQQDEwpuZXRh
cHAuY29tMQswCQYDVQQGEwJVUzEJMAcGALUECBMAMQkwBwYDVQQHEwAxCTAHBG9NV
BAoTADAEJMAcGALUECXMAMQ8wDQYJKoZIhvcNAQkBFgAwHhcNMTAwNDI2MTk0OTI4
...
-----END CERTIFICATE-----

Please enter Private Key: Press <Enter> when done
-----BEGIN RSA PRIVATE KEY-----
MIIBPAIBAAJBAM16ytrK8nQj82UsWeHOeT8gk0BPX+Y5MLyCsUdXA7hXhumHNpvF
C6lX2G32Sx8VEal th94tx+vOEzq+UaqHlt0CAwEAAQJBAMzjDwlgmlm3qIr/n8VT
PFnnZnbVcXVM70tbUsgPKw+QCCh9dF1jmuQKeDr+wUMWkn1DeGrfhILpzfJGHRlJ
...
-----END RSA PRIVATE KEY-----

Please enter certificates of Certification Authorities (CA) which form the
certificate chain of the server certificate. This starts with the issuing
CA certificate of the server certificate and can range up to the root CA
certificate.

Do you want to continue entering root and/or intermediate certificates {y|
n}: y

Please enter Intermediate Certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIIE+zCCBGsG AwIBAgICAQ0wDQYJKoZIhvcNAQEFBQAwgbsxJDAiBgNVBACGTG1Z
h bG1DZXJ0IFZhbGlkYXRpb24gTmV0d29yazEXMBUGALUEChMOVmFsaUNlcnQsIElu
Yy4xNNTAzBgNVBAsTTFZhbG1DZXJ0IENsYXNzIDUgUG9saWN5IFZhbGlkYXRpb24g
...
-----END CERTIFICATE-----

Do you want to continue entering root and/or intermediate certificates {y|
n}: n

Note: You should keep a copy of your certificate and private key for future
reference.
If you revert to an earlier release, the certificate and private key are
deleted.
```

Related tasks

[Installing a client CA or root CA certificate to authenticate an SSL client of the cluster or SVM](#) on page 141

Installing a client CA or root CA certificate to authenticate an SSL client of the cluster or SVM

To enable the cluster or Storage Virtual Machine (SVM) to authenticate a client that wants to access it, you can install a digital certificate with the `client-ca` type on the cluster or SVM for the root

certificate of the CA that signed the client's certificate signing request (CSR). You can also create a root CA certificate with the `root-ca` type on the cluster or SVM to self-sign the CSR for the client.

Before you begin

Enabling SSL client authentication requires that SSL server authentication be enabled (the default). The `security ssl show` command displays the configuration setting.

Steps

1. If the cluster or SVM will be the CA that signs the client certificate, and a self-signed root CA certificate for the cluster or SVM does not yet exist, create one by using the `security certificate create` command with the `-type root-ca` parameter.

Example

The following command creates a root CA certificate for the “vs1” SVM whose custom common name is `lab.companyname.com`:

```
cluster1::> security certificate create -vserver vs1 -common-name
lab.companyname.com -type root-ca
```

2. Enable SSL client authentication on the cluster or SVM by using the `security ssl modify` command with the `-client-enabled` parameter set to `true`.
3. Generate a CSR for the client you want to authenticate by using the `security certificate generate-csr` command.

Example

The following command generates a CSR for a client whose custom common name is `vs1admin`:

```
cluster1::> security certificate generate-csr -common-name vs1admin
```

```
Certificate Signing Request :
-----BEGIN CERTIFICATE REQUEST-----
MIICojCCAYoCAQAwXTERMA8GA1UEAxMIdnMxYWRtaW4xCzAJBgNVBAYTA1VTMQkw
BwYDVQQIEwAxCTAHBgNVBACjTADEJMACGALUEChMAMQkwBwYDVQQLEwAxZDZANBgkq
hkiG9w0BCQEWADECCASjLDxLnV3uZu8WNghpbIL98QP4oxwFu7G0HQsOle03HMazOFyvW
...
-----END CERTIFICATE REQUEST-----
```

```
Private Key :
-----BEGIN RSA PRIVATE KEY-----
MIIEowIBAAKCAQEAvgiF1PmYy1Vtmkf6I8+mRXOy/m+3m/O1sEjUILLbopzTlTu92
igqEzDY4W6q7KorKcSa2x/Zn6IRlqxKrQbvUAJvAUDhcV7bn9NAzv9JE1j/6+0RY
IVR6Hr6QnCRSsjLLDxBnV3uZu8WNghpbIL98QP4oxwFu7G0HQsOle03HMazOFyvW
...
-----END RSA PRIVATE KEY-----
```

Note: Please keep a copy of your certificate request and private key for future reference.

Data ONTAP displays the certificate request and private key and reminds you to copy them to a file for future reference.

4. If you self-sign the CSR, complete the following steps:
 - a. Display the root CA certificate you created in Step 1 by using the `security certificate show` command with the `-instance` and `-type root-ca` parameters.

You will need the following information from the command output for signing the CSR:

- Certificate authority (CA)
- Serial number of the certificate

Example

```
cluster1::> security certificate show -instance -vserver vs1 -type
root-ca

                                Vserver: vs1
                                FQDN or Custom Common Name: lab.companyname.com
                                Serial Number of Certificate: 50F84392
                                Certificate Authority: lab.companyname.com
                                Type of Certificate: root-ca
                                Size of Requested Certificate(bits): 2048
                                Certificate Start Date: Wed Jun 25 13:29:16 2014
                                Certificate Expiration Date: Thu Jun 25 13:29:16 2015
                                Public Key Certificate: -----BEGIN CERTIFICATE-----
                                MIID
+zCCAuOgAwIBAgIEUPhDkjANBgkqhkiG9w0BAQsFADBbMQ8wDQYDVQQDEWZt
.
.
.
```

- Sign the CSR with the root CA by using the `security certificate sign` command.

The default format (`-format`) for the signed certificate is PEM. If you specify the format to be PKCS12, you can optionally specify the destination to upload the signed certificate by using the `-destination` parameter.

- When you are prompted, enter the CSR and then press ENTER.

Example

```
cluster1::> security certificate sign -vserver vs1 -ca
lab.companyname.com -ca-serial 50F84392

Please enter Certificate Signing Request (CSR): Press <enter> when
done

-----BEGIN CERTIFICATE REQUEST-----
MIICrTCCAZUCAQAwaDEcMBoGAlUEAxMTQ1NSLlNpZ25pbmdUZXR0LmNvbTELMakG
AlUEBhMCMVVMxCTAHBGnVBAGTADEJMACGAlUEBxMAMQkwBwYDVQQKEwAxCTAHBGnV
BASATADEPMA0GCSqGSIb3DQEJARYAMIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIB
...
-----END CERTIFICATE REQUEST-----

Signed Certificate: :
-----BEGIN CERTIFICATE-----
MIIDmzCCAoOgAwIBAgIEU9e2rzANBgkqhkiG9w0BAQsFADBoMRwwGgYDVQQDExNO
ZXcuQ29tcGFueU5hbWUuY29tMQswCQYDVQQGEwJVUzEJMACGAlUECBMAMQkwBwYD
VQQHEwAxCTAHBGnVBAGTADEJMACGAlUECXMAMQ8wDQYJKoZIhvcNAQkBFgAwHhcN
...
-----END CERTIFICATE-----
```

The signed certificate is displayed. You should keep a copy of the certificate.

- If you have a third-party CA sign the CSR, complete the following steps:
 - Send the certificate request from the CSR output (Step 3) in an electronic form (such as email) to a trusted CA for signing.

After processing your request, the CA sends you the signed digital certificate. You should keep a copy of the private key and the CA-signed certificate for future reference.

- b. On the cluster or SVM, install the root certificate and each intermediate certificate of the CA that signed the certificate by using the `security certificate install` command with the `-type client-ca` parameter.

Example

```
cluster1::> security certificate install -vserver vs1 -type client-
ca

Please enter Certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIIDNjCCAp+gAwIBAgIQNhIilsXjOKUgodJfTNcJVDANBgkqhkiG9w0BAQUFADCB
zjELMAkGALUEBhMCWkExFTATBgNVBAGTDFdlc3Rlcm4gQ2FwZTESMBAGALUEBxMJ
Q2FwZSBub3duMR0wGwYDVQQKEXRUAaGF3dGUgQ29uc3VsdGluZyBjYzEoMCYGA1UE
...
-----END CERTIFICATE-----

You should keep a copy of the CA-signed digital certificate for
future reference.
```

6. Provide the self-signed or CA-signed certificate for the user to install on the client.
7. Repeat Step 3 to Step 6 for each client you want to authenticate.
8. If users are not set up to be authenticated by digital certificates, add users individually by using the `security login create` command with the `-authmethod` parameter set to `cert`.

For cluster user accounts, digital certificate authentication is supported only with the `http` and `ontapi` access methods (`-application`). For SVM user accounts, digital certificate authentication is supported only with the `ontapi` access method.

The `security login show` command displays user login methods.

Related tasks

[Installing a server certificate to authenticate the cluster or SVM as an SSL server](#) on page 139

Installing a server CA certificate to authenticate an SSL server to which the cluster or SVM is a client

Sometimes the cluster or Storage Virtual Machine (SVM) is a client to another SSL server (which, for example, can be an Active Directory domain controller that supports LDAP over SSL). In this case, you can enable the cluster or SVM to authenticate the SSL server by installing the server's root certificate with the `server-ca` type on the cluster or SVM.

Before you begin

You must have the root certificate of the SSL server. The root certificate can be self signed by the server or signed by a third-party CA for the server.

Steps

1. Install the root certificate provided by the SSL server by using the `security certificate install` command with the `-type server-ca` parameter.
2. When you are prompted, enter the certificate, and then press Enter.

Data ONTAP reminds you to keep a copy of the certificate for future reference.

Example of installing a server CA certificate of an SSL server

The following example installs an SSL server's CA certificate with the `server-ca` type. The certificate is used for server authentication and is installed on the “vs1” SVM, which serves as a client to the server:

```
cluster1::> security certificate install -vserver vs1 -type server-
ca

Please enter Certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIIDNjCCAp+gAwIBAgIQNhIilsXjOKUgodJfTncJVdANBgkqhkiG9w0BAQUFADCB
zjELMAkGALUEBhMCWkExFTATBgNVBAGTDfDlc3Rlcm4gQ2FwZTESMBAGALUEBxMJ
Q2FwZSBub3duMR0wGwYDVQQKEXRUaGF3dGUgQ29uc3VsdGluZyBjYzEoMCMYGA1UE
...
-----END CERTIFICATE-----

You should keep a copy of the CA-signed digital certificate for
future reference.
```

Related tasks

[Installing a client certificate to authenticate the cluster or SVM as an SSL client](#) on page 145

Installing a client certificate to authenticate the cluster or SVM as an SSL client

To enable an SSL server to authenticate the cluster or Storage Virtual Machine (SVM) as an SSL client, you install a digital certificate with the `client` type on the cluster or SVM. Then you provide the `client-ca` certificate to the SSL server administrator for installation on the server.

Before you begin

You must have already installed the root certificate of the SSL server on the cluster or SVM with the `server-ca` certificate type.

Steps

1. To use a self-signed digital certificate for client authentication, use the `security certificate create` command with the `-type client` parameter.
2. To use a CA-signed digital certificate for client authentication, complete the following steps:
 - a. Generate a digital certificate signing request (CSR) by using the `security certificate generate-csr` command.

Data ONTAP displays the CSR output, which includes a certificate request and private key, and reminds you to copy the output to a file for future reference.
 - b. Send the certificate request from the CSR output in an electronic form (such as email) to a trusted CA for signing.

After processing your request, the CA sends you the signed digital certificate. You should keep a copy of the private key and the CA-signed certificate for future reference.
 - c. Install the CA-signed certificate by using the `security certificate install` command with the `-type client` parameter.
 - d. Enter the certificate and the private key when you are prompted, and then press Enter.

- e. Enter any additional root or intermediate certificates when you are prompted, and then press Enter.

You install an intermediate certificate on the cluster or SVM if a certificate chain that begins at the trusted root CA, and ends with the SSL certificate issued to you, is missing the intermediate certificates. An intermediate certificate is a subordinate certificate issued by the trusted root specifically to issue end-entity server certificates. The result is a certificate chain that begins at the trusted root CA, goes through the intermediate, and ends with the SSL certificate issued to you.

3. Provide the `client-ca` certificate of the cluster or SVM to the administrator of the SSL server for installation on the server.

The `security certificate show` command with the `-instance` and `-type client-ca` parameters displays the `client-ca` certificate information.

Examples of installing a client certificate to authenticate the cluster or SVM as an SSL client

The following example creates a self-signed client certificate for the “vs1” SVM at a company whose custom common name is `lab.companyname.com`. The certificate is for authenticating the “vs1” SVM as an SSL client:

```
cluster1::> security certificate create -vserver vs1 -common-name
lab.companyname.com -type client
```

The following command creates a CSR with a 2048-bit private key for use by the Software group in the IT department of a company whose custom common name is `lab.companyname.com`, located in Sunnyvale, California, USA. The email address of the contact administrator who manages the SVM is `web@companyname.com`. The system displays the CSR and the private key on the console:

```
cluster1::> security certificate generate-csr -common-name
lab.companyname.com
-size 2048 -country US -state CA -locality Sunnyvale -organization IT
-unit Software -email-addr web@companyname.com
```

```
Certificate Signing Request:
-----BEGIN CERTIFICATE REQUEST-----
MIICrjCCAZYCAQMwaTEQMA4GA1UEAxMHcnRwLmNvbTELMakGAlUEBhMCVVMxCzAJ
BgNVBAGTAK5DMQwwCgYDVQQHEwNSVFAXDTALBgNVBAoTBGNvcmUxDTALBgNVBAsT
BGNvcmUxDzANBgkqhkiG9w0BCQEWADCCASlWdQYJKoZIhvcNAQEBBQADggEPADCC
...
-----END CERTIFICATE REQUEST-----
```

```
Private Key:
-----BEGIN RSA PRIVATE KEY-----
MIIBPAIBAAJBAM16ytrK8nQj82UsWeHOeT8gk0BPX+Y5MLyCsUdXA7hXhumHNpvF
C61X2G32Sx8VEalth94tx+vOEzq+UaqHlt0CAwEAAQJBAMZjDwlgmlm3qIr/n8VT
PFnnZnbVcXVM7OtbUsgPKw+QCCh9dF1jmuQKeDr+wUMWknldEGrfhILpzfJGHRlJ
...
-----END RSA PRIVATE KEY-----
```

Note: Please keep a copy of your private key and certificate request for future reference.

The following command installs a CA-signed client certificate for the “vs1” SVM. The certificate is for authenticating the “vs1” SVM as an SSL client:

```
cluster1::> security certificate install -vserver vs1 -type client

Please enter Certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIIB8TCCAZugAwIBAwIBADANBgkqhkiG9w0BAQQFADBfMRMwEQYDVQQDEwpuZXRh
cHAuY29tMQswCQYDVQQGEwJVUzEJMAcGAlUECBMAMQkwBwYDVQQHEwAxCTAHBgNV
```

```

BAoTAEJMAcGAlUECXMAMQ8wDQYJKoZIhvcNAQkBFgAwHhcNMTAwNDI2MTk0OTI4
...
-----END CERTIFICATE-----

Please enter Private Key: Press <Enter> when done
-----BEGIN RSA PRIVATE KEY-----
MIIBPAIBAAJBAMl6ytrK8nQj82UsWeHOeT8gk0BPX+Y5MLyCsUdXA7hXhumHNpvF
C61X2G32Sx8VEalth94tx+vOEzq+UaqHlt0CAwEAAQJBAMZjDwlglmlm3qIr/n8VT
PFnnZnbVcXVM7OtbUsgPKw+QCCh9dF1jmuQKeDr+wUMWknlDeGrfhlLpzfJGHRlJ
...
-----END RSA PRIVATE KEY-----

Please enter certificates of Certification Authorities (CA) which form the
certificate chain of the client certificate. This starts with the issuing
CA certificate of the client certificate and can range up to the root CA
certificate.

Do you want to continue entering root and/or intermediate certificates {y|
n}: y

Please enter Intermediate Certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIIE+zCCBGSgAwIBAgICAQ0wDQYJKoZIhvcNAQEFBQAwbG1DZXJ0IFZhbG1kYXRpb24gTmV0d29yazEXMBUGAlUEChMOVmFsaUNlcnQsIElu
Yy4xNTAzBzBzNVBAsTLFZhbG1DZXJ0IENsYXNzIDFgUG9saWN5IFZhbG1kYXRpb24g
...
-----END CERTIFICATE-----

Do you want to continue entering root and/or intermediate certificates {y|
n}: n

Note: You should keep a copy of your certificate and private key for future
reference.
If you revert to an earlier release, the certificate and private key are
deleted.
```

Related tasks

[Installing a server CA certificate to authenticate an SSL server to which the cluster or SVM is a client](#) on page 144

Installing KMIP certificates to mutually authenticate the cluster and an external key manager

Mutually authenticating the cluster and an external key manager such as a Key Management Interoperability Protocol (KMIP) server enables the key manager to communicate with the cluster by using KMIP over SSL. You do so when an application or certain functionality (for example, the Data ONTAP-v platform or the Storage Encryption functionality) requires secure keys to ensure secure data access.

Before you begin

You must have obtained the root certificate from the external key manager. This root certificate is for authenticating the key manager as an SSL server, and it can be self signed by the server or signed by a third-party CA for the server.

About this task

You can configure for the cluster to mutually authenticate with up to four KMIP servers.

Steps

1. Install a KMIP certificate to authenticate the cluster as an SSL client to an external KMIP server:

- a. If you do not have a signed certificate to use for KMIP authentication for the cluster, generate a digital certificate signing request (CSR) by using the `security certificate generate-csr` command.

Data ONTAP displays the CSR output, which includes a certificate request and private key, and reminds you to copy the output to a file for future reference.

- b. Send the CSR to a third-party CA for signing, or sign the CSR with the cluster's root CA certificate by using the `security certificate sign` command.

The cluster's root CA certificate can be displayed by using the `security certificate show` command with the `-instance` and `-type root-ca` parameters.

You should keep a copy of the private key and the signed certificate for future reference.

- c. Use the `security certificate install` command with the `-type client` and `-subtype kmip-cert` parameters to install the signed KMIP certificate to authenticate the cluster.
- d. Enter the certificate and the private key when you are prompted, and then press Enter.
- e. Enter any additional root or intermediate certificates when you are prompted, and then press Enter.

Example

```
cluster1::> security certificate install -type client -subtype kmip-cert
-vserver cluster1

Please enter Certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIIFYjCCBLKgAwIBAgIQelt1/RNlMt1xn+/KLgtvOzANBgkqhkiG9w0BAQUFADCB
tTELMaKGA1UEBhMCVVMxZzAVBgNVBAoTDlZlcm1TaWduLkCBJmMuMR8wHQYDVQQL
ExxZWZxJpU2lnbiBUcncVzdCBOZXR3b3JrMTswOQYDVQQLEzJUZXJtcyBvZi1c2Ug
...
-----END CERTIFICATE-----

Please enter Private Key: Press <Enter> when done
-----BEGIN RSA PRIVATE KEY-----
MIIEogIBAAKCAQEAlX+NLxxxtOJUM3dhpHOn3Z+76wGGax0jYOGzC5AQhigB8buL
Y2bfrs4Jttv6FeFVzfnpoWV9bZzgHiqzyRJkxv960YEVHPI3i5x+HDPDg3qRhWRS
Zqfia35wJ8rxMydzydO7tvMtZyzVS8EpCl3rEMC+gMG+uVr0e4KJvamOoyAQ7c6z
...
-----END RSA PRIVATE KEY-----

Please enter certificates of Certification Authorities (CA) which form the
certificate chain of the client certificate. This starts with the issuing
CA certificate of the client certificate and can range up to the root CA
certificate.

Do you want to continue entering root and/or intermediate certificates {y|n}: n

You should keep a copy of the private key and the CA-signed digital
certificate for future reference.

cluster1::>
```

2. Install a KMIP certificate to authenticate a KMIP server as an SSL server to the cluster:

- a. Use the `security certificate install` command with the `-type server-ca`, `-subtype kmip-cert`, and `-kmip-server-ip ip_address` parameters to install a KMIP certificate for the KMIP server.

If you want to authenticate multiple KMIP servers that are in the same subnet, you specify the `-kmip-server-ip` parameter with the address of the subnet.

For example, `-kmip-server-ip 10.53.27.55` specifies the KMIP server at the **10.53.27.55** IP address, and `-kmip-server-ip 10.53.0.0` specifies all KMIP servers that are in the **10.53.0.0** subnet.

- b. When you are prompted, enter the certificate, and then press Enter.

Data ONTAP reminds you to keep a copy of the certificate for future reference.

Example

```
cluster1::> security certificate install -type server-ca -subtype kmip-cert
-kmip-server-ip 10.53.0.0 -vserver cluster1

Please enter Certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIICPDCCAaUCEDyRMcsf9tAbDpq40ES/Er4wDQYJKoZIhvcNAQEFBQAwXzELMAkG
2JhucwNhkcV8sEVAbkSdjbCxlRhLQ2pRdKkkirWmnWXbj9T/UWZYB2oK0z5XqcJ
2HUw19JlYDln1khVdWk/kfVIC0dpImmClr7JyDiGSnosxliAU5rfGW/D/xwzoiQ
...
-----END CERTIFICATE-----

You should keep a copy of the CA-signed digital certificate for future
reference.

cluster1::>
```

3. Display the KMIP certificates that are installed on the cluster by using the `security certificate show` command with the `-subtype kmip-cert` parameter.

Example

```
cluster1::> security certificate show -subtype kmip-cert

Vserver      Serial Number      Common Name      Type
-----
cluster1
      7B5B75FD136532    cert.servername.com      server-ca
      Certificate Authority: Secure Server CA
      Expiration Date: Mon Oct 23 16:59:59 2017

cluster1
      050E37C2DB56E5    www.example.com      client
      Certificate Authority: www.example.com
      Expiration Date: Mon Apr 30 14:14:46 2018

2 entries were displayed.

cluster1::>
```

Related tasks

[Installing a client CA or root CA certificate to authenticate an SSL client of the cluster or SVM](#) on page 141

Replacing an expired digital certificate

Each certificate that you create or install has an expiration date. When it expires, you must replace it with a new certificate so that the corresponding server or client authentication is not disrupted.

About this task

By default, digital certificates created by Data ONTAP are set to expire in 365 days, but you can specify the expiration setting (up to 10 years) when you create a digital certificate.

Steps

1. Display certificate expiration information by using the `security certificate show` command with the `-fields expiration, expire-days` parameter.

You need the following information when you delete an expired certificate:

- The SVM name
 - The common name used for the certificate
 - Serial number
 - Certificate authority (CA)
 - Certificate type or subtype
2. Delete an expired certificate by using the `security certificate delete` command.
 3. Obtain a new certificate with the same common name to replace the certificate that has expired:

If the certificate is...	Then follow the steps in...
The <code>server</code> type	<i>Installing a server certificate to authenticate the cluster or SVM as an SSL server</i> on page 139
The <code>client-ca</code> type	<i>Installing a client CA or root CA certificate to authenticate an SSL client of the cluster or SVM</i> on page 141
The <code>server-ca</code> type	<i>Installing a server CA certificate to authenticate an SSL server to which the cluster or SVM is a client</i> on page 144
The <code>client</code> type	<i>Installing a client certificate to authenticate the cluster or SVM as an SSL client</i> on page 145
The <code>kmip-cert</code> subtype	<i>Installing KMIP certificates to mutually authenticate the cluster and an external key manager</i> on page 147

Commands for managing digital certificates

You use the `security certificate` commands to manage digital certificates of the cluster or Storage Virtual Machine (SVM).

If you want to...	Use this command...
Create and install a self-signed digital certificate with one of the following types: <ul style="list-style-type: none"> • <code>server</code> • <code>root-ca</code> • <code>client</code> 	<code>security certificate create</code>
Generate a digital certificate signing request that you will send to a CA for signing	<code>security certificate generate-csr</code>
Sign a digital certificate using a self-signed root CA	<code>security certificate sign</code>

If you want to...	Use this command...
Install a CA-signed digital certificate and the public key certificate of the root CA with one of the following types, or install the kmip-cert subtype of certificate for a client or server-ca certificate on the cluster: <ul style="list-style-type: none"> • server • client-ca • server-ca • client 	<code>security certificate install</code>
Display information about installed digital certificates	<code>security certificate show</code>
Display digital certificates that are signed by the cluster or SVM as the CA	<code>security certificate ca-issued show</code>
Revoke a compromised digital certificate signed by the cluster or SVM as the CA	<code>security certificate ca-issued revoke</code>
Delete an installed digital certificate	<code>security certificate delete</code>

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Managing access to web services

A web service is an application that users can access by using HTTP or HTTPS. The cluster administrator can set up the web protocol engine, configure SSL, enable a web service, and enable users of a role to access a web service.

Data ONTAP supports the following web services:

- **Service Processor Infrastructure (`spi`)**
 This service makes a node's log, core dump, and MIB files available for HTTP or HTTPS access through the cluster management LIF or a node management LIF. The default setting is **enabled**.
 Upon a request to access a node's log files or core dump files, the **spi** web service automatically creates a mount point from a node to another node's root volume where the files reside. You do not need to manually create the mount point.
- **Data ONTAP Classic (`compat`)**
 This service provides an alternative interface to the **spi** web service. The default setting is **enabled**.
 When both the **spi** and **compat** web services are enabled, a node's log files and core dump files are available for HTTP or HTTPS access through a node management LIF.
- **Data ONTAP APIs (`ontapi`)**
 This service enables you to run Data ONTAP APIs to execute administrative functions with a remote program. The default setting is **enabled**.
 This service might be required for some external management tools. For example, if you use OnCommand System Manager, you should leave this service enabled.
- **Data ONTAP Discovery (`disco`)**

This service enables off-box management applications to discover the cluster in the network. The default setting is **enabled**.

- **Support Diagnostics (`supdiag`)**
This service controls access to a privileged environment on the system to assist problem analysis and resolution. The default setting is **disabled**. You should enable this service only when directed by technical support.
- **System Manager (`sysmgr`)**
This service controls the availability of OnCommand System Manager, which is included with Data ONTAP. The default setting is **enabled**. This service is supported only on the cluster.

Related concepts

[Understanding OnCommand System Manager](#) on page 28

Related tasks

[Accessing the cluster by using OnCommand System Manager browser-based graphic interface](#) on page 27

[Accessing a node's log, core dump, and MIB files by using a web browser](#) on page 41

Managing the web protocol engine

You can configure the web protocol engine on the cluster to control whether web access is allowed and what SSL versions can be used. You can also display the configuration settings for the web protocol engine.

You can manage the web protocol engine at the cluster level in the following ways:

- You can specify whether remote clients can use HTTP or HTTPS to access web service content (the `system services web modify` command with the `-external` parameter).
- You can specify whether SSLv3 should be used for secure web access (the `system services web modify` command with the `-sslv3-enabled` parameter).
By default, SSLv3 is enabled. Transport Layer Security 1.0 (TLSv1.0) is enabled and cannot be disabled.
- You can enable the compliance with the Federal Information Processing Standard (FIPS) 140-2 for secure web services (HTTPS) (the `system services web modify` command with the `-ssl-fips-enabled` parameter) and display whether FIPS 140-2 compliance is enabled and online (`system services web show`).
By default, FIPS 140-2 compliance is disabled. When FIPS 140-2 compliance is enabled, SSLv3 is disabled, and only TLSv1 remains enabled. Data ONTAP prevents you from enabling SSLv3 when FIPS 140-2 compliance is enabled.
If you enable FIPS 140-2 and then subsequently disable it, SSLv3 remains disabled, but TLSv1 is always enabled.
- You can display the configuration of web protocols (`system services web show`).

If the firewall is enabled, the firewall policy for the logical interface (LIF) to be used for web services must be set up to allow HTTP or HTTPS access.

If you use HTTPS for web service access, SSL for the cluster or Storage Virtual Machine (SVM) that offers the web service must also be enabled, and you must provide a digital certificate for the cluster or SVM.

In MetroCluster configurations, the setting changes you make for the web protocol engine on a cluster are not replicated on the partner cluster.

Related concepts[Managing SSL](#) on page 156[Managing web services](#) on page 154**Related tasks**[Configuring access to web services](#) on page 157**Related information**[Clustered Data ONTAP 8.3 Network Management Guide](#)**Commands for managing the web protocol engine**

You use the `system services web` commands to manage the web protocol engine. You use the `system services firewall policy create` and `network interface modify` commands to allow web access requests to go through the firewall.

If you want to...	Use this command...
Configure the web protocol engine at the cluster level: <ul style="list-style-type: none"> • Enable or disable the web protocol engine for the cluster • Enable or disable SSLv3 for the cluster • Enable or disable FIPS 140-2 compliance for secure web services (HTTPS) 	<code>system services web modify</code>
Display the configuration of the web protocol engine at the cluster level, determine whether the web protocols are functional throughout the cluster, and display whether FIPS 140-2 compliance is enabled and online	<code>system services web show</code>
Display the configuration of the web protocol engine at the node level and the activity of web service handling for the nodes in the cluster	<code>system services web node show</code>
Create a firewall policy or add HTTP or HTTPS protocol service to an existing firewall policy to allow web access requests to go through firewall	<code>system services firewall policy create</code> Setting the <code>-service</code> parameter to <code>http</code> or <code>https</code> enables web access requests to go through firewall.
Associate a firewall policy with a LIF	<code>network interface modify</code> You can use the <code>-firewall-policy</code> parameter to modify the firewall policy of a LIF.

Related references[Commands for managing SSL](#) on page 156[Commands for managing digital certificates](#) on page 150**Related information**[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Managing web services

You can enable or disable a web service for the cluster or a Storage Virtual Machine (SVM), display the settings for web services, and control whether users of a role can access a web service.

You can manage web services for the cluster or an SVM in the following ways:

- Enabling or disabling a specific web service
- Specifying whether access to a web service is restricted to only encrypted HTTP (SSL)
- Displaying the availability of web services
- Allowing or disallowing users of a role to access a web service
- Displaying the roles that are permitted to access a web service

For a user to access a web service, all of the following conditions must be met:

- The user must be authenticated.
For instance, a web service might prompt for a user name and password. The user's response must match a valid account.
- The user must be set up with the correct access method.
Authentication only succeeds for users with the correct access method for the given web service. For the Data ONTAP API web service (`ontapi`), users must have the `ontapi` access method. For all other web services, users must have the `http` access method.
Note: You use the `security login` commands to manage users' access methods and authentication methods.
- The web service must be configured to allow the user's access-control role.
Note: You use the `vserver services web access` commands to control a role's access to a web service.

If a firewall is enabled, the firewall policy for the LIF to be used for web services must be set up to allow HTTP or HTTPS.

If you use HTTPS for web service access, SSL for the cluster or SVM that offers the web service must also be enabled, and you must provide a digital certificate for the cluster or SVM.

Related concepts

[Managing the web protocol engine](#) on page 152

[Managing user accounts](#) on page 116

[Access methods for user accounts](#) on page 117

[Managing SSL](#) on page 156

Related tasks

[Configuring access to web services](#) on page 157

Commands for managing web services

You use the `vserver services web` commands to manage the availability of web services for the cluster or a Storage Virtual Machine (SVM). You use the `vserver services web access` commands to control a role's access to a web service.

If you want to...	Use this command...
Configure a web service for the cluster or an SVM: <ul style="list-style-type: none"> • Enable or disable a web service • Specify whether only HTTPS can be used for accessing a web service 	<code>vserver services web modify</code>
Display the configuration and availability of web services for the cluster or an SVM	<code>vserver services web show</code>
Authorize a role to access a web service on the cluster or an SVM	<code>vserver services web access create</code>
Display the roles that are authorized to access web services on the cluster or an SVM	<code>vserver services web access show</code>
Prevent a role from accessing a web service on the cluster or an SVM	<code>vserver services web access delete</code>

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Commands for managing mount points on the nodes

The `spi` web service automatically creates a mount point from one node to another node's root volume upon a request to access the node's log files or core files. Although you do not need to manually manage mount points, you can do so by using the `system node root-mount` commands.

If you want to...	Use this command...
Manually create a mount point from one node to another node's root volume	<code>system node root-mount create</code> Only a single mount point can exist from one node to another.
Display existing mount points on the nodes in the cluster, including the time a mount point was created and its current state	<code>system node root-mount show</code>
Delete a mount point from one node to another node's root volume and force connections to the mount point to close	<code>system node root-mount delete</code>

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Managing SSL

The SSL protocol improves the security of web access by using a digital certificate to establish an encrypted connection between a web server and a browser.

You can manage SSL for the cluster or a Storage Virtual Machine (SVM) in the following ways:

- Enabling SSL
- Generating and installing a digital certificate and associating it with the cluster or SVM
- Displaying the SSL configuration to see whether SSL has been enabled, and, if available, the SSL certificate name
- Setting up firewall policies for the cluster or SVM, so that web access requests can go through
- Defining which SSL versions can be used
- Restricting access to only HTTPS requests for a web service

Related concepts

[Managing the web protocol engine](#) on page 152

[Managing web services](#) on page 154

[Managing digital certificates for server or client authentication](#) on page 137

Related tasks

[Configuring access to web services](#) on page 157

Related information

[Clustered Data ONTAP 8.3 Network Management Guide](#)

Commands for managing SSL

You use the `security ssl` commands to manage the SSL protocol for the cluster or a Storage Virtual Machine (SVM).

If you want to...	Use this command...
Enable SSL for the cluster or an SVM, and associate a digital certificate with it	<code>security ssl modify</code>
Display the SSL configuration and certificate name for the cluster or an SVM	<code>security ssl show</code>

Related references

[Commands for managing web services](#) on page 155

[Commands for managing the web protocol engine](#) on page 153

[Commands for managing digital certificates](#) on page 150

Related information

[Clustered Data ONTAP 8.3 Network Management Guide](#)

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Configuring access to web services

Configuring access to web services allows authorized users to use HTTP or HTTPS to access the service content on the cluster or a Storage Virtual Machine (SVM).

Steps

1. If a firewall is enabled, ensure that HTTP or HTTPS access is set up in the firewall policy for the LIF that will be used for web services:

Note: You can check whether a firewall is enabled by using the `system services firewall show` command.

- a. To verify that HTTP or HTTPS is set up in the firewall policy, use the `system services firewall policy show` command.

You set the `-service` parameter of the `system services firewall policy create` command to `http` or `https` to enable the policy to support web access.

- b. To verify that the firewall policy supporting HTTP or HTTPS is associated with the LIF that provides web services, use the `network interface show` command with the `-firewall-policy` parameter.

You use the `network interface modify` command with the `-firewall-policy` parameter to put the firewall policy into effect for a LIF.

2. To configure the cluster-level web protocol engine and make web service content accessible, use the `system services web modify` command.
3. If you plan to use secure web services (HTTPS), enable SSL and provide digital certificate information for the cluster or SVM by using the `security ssl modify` command.
4. To enable a web service for the cluster or SVM, use the `vserver services web modify` command.

You must repeat this step for each service that you want to enable for the cluster or SVM.

5. To authorize a role to access web services on the cluster or SVM, use the `vserver services web access create` command.

The role that you grant access must already exist. You can display existing roles by using the `security login role show` command or create new roles by using the `security login role create` command.

6. For a role that has been authorized to access a web service, ensure that its users are also configured with the correct access method by checking the output of the `security login show` command.

To access the Data ONTAP API web service (`ontapi`), a user must be configured with the `ontapi` access method. To access all other web services, a user must be configured with the `http` access method.

Note: You use the `security login create` command to add an access method for a user.

Related concepts

[Managing SSL](#) on page 156

[Managing digital certificates for server or client authentication](#) on page 137

[Managing the web protocol engine](#) on page 152

[Managing web services](#) on page 154

[Managing access-control roles](#) on page 123

[Access methods for user accounts](#) on page 117

Related information

[Clustered Data ONTAP 8.3 Network Management Guide](#)

Troubleshooting web service access problems

Configuration errors cause web service access problems to occur. You can address the errors by ensuring that the LIF, firewall policy, web protocol engine, web services, digital certificates, and user access authorization are all configured correctly.

The following table helps you identify and address web service configuration errors:

This access problem...	Occurs because of this configuration error...	To address the error...
Your web browser returns an <code>unable to connect</code> or <code>failure to establish a connection</code> error when you try to access a web service.	Your LIF might be configured incorrectly.	Ensure that you can ping the LIF that provides the web service. Note: You use the <code>network ping</code> command to ping a LIF. For information about network configuration, see the <i>Clustered Data ONTAP Network Management Guide</i> .
	Your firewall might be configured incorrectly.	Ensure that a firewall policy is set up to support HTTP or HTTPS and that the policy is assigned to the LIF that provides the web service. Note: You use the <code>system services firewall policy</code> commands to manage firewall policies. You use the <code>network interface modify</code> command with the <code>-firewall-policy</code> parameter to associate a policy with a LIF.
	Your web protocol engine might be disabled.	Ensure that the web protocol engine is enabled so that web services are accessible. Note: You use the <code>system services web</code> commands to manage the web protocol engine for the cluster.
Your web browser returns a <code>not found</code> error when you try to access a web service.	The web service might be disabled.	Ensure that each web service that you want to allow access to is enabled individually. Note: You use the <code>vserver services web modify</code> command to enable a web service for access.

This access problem...	Occurs because of this configuration error...	To address the error...
The web browser fails to log in to a web service with a user's account name and password.	The user cannot be authenticated, the access method is not correct, or the user is not authorized to access the web service.	Ensure that the user account exists and is configured with the correct access method and authentication method. Also, ensure that the user's role is authorized to access the web service. Note: You use the <code>security login</code> commands to manage user accounts and their access methods and authentication methods. Accessing the Data ONTAP API web service requires the <code>ontapi</code> access method. Accessing all other web services requires the <code>http</code> access method. You use the <code>vserver services web access</code> commands to manage a role's access to a web service.
You connect to your web service with HTTPS, and your web browser indicates that your connection is interrupted.	You might not have SSL enabled on the cluster or Storage Virtual Machine (SVM) that provides the web service.	Ensure that the cluster or SVM has SSL enabled and that the digital certificate is valid. Note: You use the <code>security ssl</code> commands to manage SSL configuration for HTTP servers and the <code>security certificate show</code> command to display digital certificate information.
You connect to your web service with HTTPS, and your web browser indicates that the connection is untrusted.	You might be using a self-signed digital certificate.	Ensure that the digital certificate associated with the cluster or SVM is signed by a trusted CA. Note: You use the <code>security certificate generate-csr</code> command to generate a digital certificate signing request and the <code>security certificate install</code> command to install a CA-signed digital certificate. You use the <code>security ssl</code> commands to manage the SSL configuration for the cluster or SVM that provides the web service.

Related concepts

[Managing the web protocol engine](#) on page 152

[Managing digital certificates for server or client authentication](#) on page 137

[Managing web services](#) on page 154

[Managing user accounts](#) on page 116

[Managing access-control roles](#) on page 123

[Managing SSL](#) on page 156

Related tasks

[Installing a server certificate to authenticate the cluster or SVM as an SSL server](#) on page 139

Related information

[Clustered Data ONTAP 8.3 Network Management Guide](#)

Managing audit settings for management activities

Audit logging creates a chronological record of management activities that take place on the management interface. You can forward the command history log to destinations that you specify. You can also control what requests are recorded in the `mgwd.log` file for technical support. You can display audit log files by using a web browser.

Data ONTAP logs management activities that are performed on the cluster, for example, what request was issued, the user who triggered the request, the user's access method, and the time of the request. The management activities can be one of the following types:

- Set requests, which typically apply to nondisplay commands or operations
This is the type of requests that are issued when you run a `create`, `modify`, or `delete` command, for instance.
- Get requests, which retrieve information and display it in the management interface
This is the type of requests that are issued when you run a `show` command, for instance.

Data ONTAP logs management activities in the following files of a node:

- The `/mroot/etc/log/mlog/command-history.log` file logs only set requests.
This file logs all set requests but not get requests. This behavior is not affected by the `security audit modify` command settings.
This file is sent by AutoSupport to the specified recipients. The content can also be forwarded to up to 10 destinations that you specify (`cluster log-forwarding create`), for example, a Splunk or syslog server for monitoring, analysis, or backup purposes.
- The `/mroot/etc/log/mlog/mgwd.log` file logs only set requests by default.
This file is used for technical support and diagnostic purposes. By default, it logs all set requests but not get requests. However, the `security audit modify` command enables you to control whether the following requests are logged in the file:
 - Set requests sent from the Data ONTAP CLI (`-cliset`)
 - Set requests sent from the Data ONTAP APIs (`-ontapiset`)
 - Get requests sent from the Data ONTAP CLI (`-cliget`)
 - Get requests sent from the Data ONTAP APIs (`-ontapiget`)
- The `/mroot/etc/log/auditlog` file logs nodeshell CLI commands.
This file logs the commands that are executed in the nodeshell. AutoSupport messages include the nodeshell audit log files of a node and those of the partner node. The audit log also includes the remote client address from failed login attempts through services such as RSH, SSH, console, telnet, and ontapi.

The `command-history.log` and `mgwd.log` files are rotated daily. The rotation also occurs when they reach 100 MB in size, and their previous 34 copies are preserved (with a maximum total of 35 files, respectively).

The `auditlog` file is rotated weekly. The rotation also occurs if the file reaches 10 MB in size. Up to six `auditlog` files are preserved.

You can display the content of the `/mroot/etc/log/` directory by using a web browser.

Related tasks

[Accessing a node's log, core dump, and MIB files by using a web browser](#) on page 41

Forwarding the command history log to a destination

You can forward the command history log (content of `/mroot/etc/log/mlog/command-history.log`) to a maximum of 10 destinations that you specify. For example, you can forward the log to a Splunk or syslog server for monitoring, analysis, or backup purposes.

Steps

1. For each destination that you want to forward the command history log to, specify the destination IP address or host name by using the `cluster log-forwarding create` command.

Before the command creates a record for the specified destination, it pings the destination host to verify connectivity. If the ping is unsuccessful, the command fails with an error. Although not recommended, using the `-force` parameter with the command bypasses the connectivity verification.

The `cluster log-forwarding modify` command enables you to modify a previously specified destination.

The `cluster log-forwarding delete` command enables you to delete a previously specified destination.

Example

```
cluster1::> cluster log-forwarding create -destination 192.168.123.96
-port 514 -facility user

cluster1::> cluster log-forwarding create -destination 192.168.123.98
-port 514 -facility user
```

2. Display the destination records by using the `cluster log-forwarding show` command.

Example

```
cluster1::> cluster log-forwarding show

Destination Host          Port    Syslog
-----
192.168.123.96           514     user
192.168.123.98           514     user
2 entries were displayed.

cluster1::>
```

Related tasks

[Accessing a node's log, core dump, and MIB files by using a web browser](#) on page 41

Related information

[Clustered Data ONTAP 8.3.2 man page: cluster log-forwarding create - Create a log forwarding destination](#)

[Clustered Data ONTAP 8.3.2 man page: cluster log-forwarding modify - Modify log forwarding destination settings](#)

Clustered Data ONTAP 8.3.2 man page: cluster log-forwarding delete - Delete a log forwarding destination

Clustered Data ONTAP 8.3.2 man page: cluster log-forwarding show - Display log forwarding destinations

Commands for managing audit settings for management activities

You use the `security audit` commands to manage what management activities are logged in the `mgwd.log` file. You use the `cluster log-forwarding` commands to manage destinations for forwarding the command history log.

If you want to...	Use this command...
Specify whether set or get requests from the Data ONTAP CLI or APIs should be logged in the <code>/mroot/etc/log/mlog/mgwd.log</code> file for technical support	<code>security audit modify</code>
Display the settings of the <code>/mroot/etc/log/mlog/mgwd.log</code> file	<code>security audit show</code>
Specify a destination for the command history log (<code>/mroot/etc/log/mlog/command-history.log</code>) to be forwarded to	<code>cluster log-forwarding create</code>
Modify a destination for the command history log	<code>cluster log-forwarding modify</code>
Delete a destination for the command history log	<code>cluster log-forwarding delete</code>
Show the configured destinations for the command history log	<code>cluster log-forwarding show</code>

Related information

Clustered Data ONTAP 8.3 Commands: Manual Page Reference

Managing the cluster time (cluster administrators only)

Problems can occur when the cluster time is inaccurate. Although Data ONTAP enables you to manually set the time zone, date, and time on the cluster, you should configure the Network Time Protocol (NTP) servers to synchronize the cluster time.

NTP is always enabled. However, configuration is still required for the cluster to synchronize with an external time source. Data ONTAP enables you to manage the cluster's NTP configuration in the following ways:

- You can associate a maximum of 10 external NTP servers with the cluster (`cluster time-service ntp server create`).
 - For redundancy and quality of time service, you should associate at least three external NTP servers with the cluster.
 - You can specify an NTP server by using its IPv4 or IPv6 address or fully qualified host name.
 - You can manually specify the NTP version (v3 or v4) to use. By default, Data ONTAP automatically selects the NTP version that is supported for a given external NTP server. If the NTP version you specify is not supported for the NTP server, time exchange cannot take place.
 - At the advanced privilege level, you can specify an external NTP server that is associated with the cluster to be the primary time source for correcting and adjusting the cluster time.
- You can display the NTP servers that are associated with the cluster (`cluster time-service ntp server show`).
- You can modify the cluster's NTP configuration (`cluster time-service ntp server modify`).
- You can disassociate the cluster from an external NTP server (`cluster time-service ntp server delete`).
- At the advanced privilege level, you can reset the configuration by clearing all external NTP servers' association with the cluster (`cluster time-service ntp server reset`).

A node that joins a cluster automatically adopts the NTP configuration of the cluster.

In addition to using NTP, Data ONTAP also enables you to manually manage the cluster time. This capability is helpful when you need to correct erroneous time (for example, a node's time has become significantly incorrect after a reboot). In that case, you can specify an approximate time for the cluster until NTP can synchronize with an external time server. The time you manually set takes effect across all nodes in the cluster.

You can manually manage the cluster time in the following ways:

- You can set or modify the time zone, date, and time on the cluster (`cluster date modify`).
- You can display the current time zone, date, and time settings of the cluster (`cluster date show`).

Note: Job schedules do not adjust to manual cluster date and time changes. These jobs are scheduled to run based on the current cluster time when the job was created or when the job most recently ran. Therefore, if you manually change the cluster date or time, you must use the job

show and job history show commands to verify that all scheduled jobs are queued and completed according to your requirements.

Related information

[Network Time Protocol \(NTP\) Support](#)

Commands for managing the cluster time

You use the `cluster time-service ntp server` commands to manage the NTP servers for the cluster. You use the `cluster date` commands to manage the cluster time manually.

The following commands enable you to manage the NTP servers for the cluster:

If you want to...	Use this command...
Associate the cluster with an external NTP server	<code>cluster time-service ntp server create</code>
Display information about the NTP servers that are associated with the cluster	<code>cluster time-service ntp server show</code>
Modify the configuration of an external NTP server that is associated with the cluster	<code>cluster time-service ntp server modify</code>
Dissociate an NTP server from the cluster	<code>cluster time-service ntp server delete</code>
Reset the configuration by clearing all external NTP servers' association with the cluster	<code>cluster time-service ntp server reset</code> Note: This command requires the advanced privilege level.

The following commands enable you to manage the cluster time manually:

If you want to...	Use this command...
Set or modify the time zone, date, and time	<code>cluster date modify</code>
Display the time zone, date, and time settings for the cluster	<code>cluster date show</code>

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Managing the banner and MOTD

Data ONTAP enables you to configure a login banner or a message of the day (MOTD) to communicate administrative information to CLI users of the cluster or Storage Virtual Machine (SVM).

A banner is displayed in a console session (for cluster access only) or an SSH session (for cluster or SVM access) before a user is prompted for authentication such as a password. For example, you can use the banner to display a warning message such as the following to someone who attempts to log in to the system:

```
$ ssh admin@cluster1-01

This system is for authorized users only. Your IP Address has been
logged.

Password:

cluster1::> _
```

An MOTD is displayed in a console session (for cluster access only) or an SSH session (for cluster or SVM access) after a user is authenticated but before the clustershell prompt appears. For example, you can use the MOTD to display a welcome or informational message such as the following that only authenticated users will see:

```
$ ssh admin@cluster1-01

Password:

Greetings. This system is running Data ONTAP Release 8.3.1.
Your user name is 'admin'. Your last login was Wed Apr 08 16:46:53 2015
from 10.72.137.28.

cluster1::> _
```

You can create or modify the content of the banner or MOTD by using the `security login banner modify` or `security login motd modify` command, respectively, in the following ways:

- You can use the CLI interactively or noninteractively to specify the text to use for the banner or MOTD.
 - The interactive mode, launched when the command is used without the `-message` or `-uri` parameter, enables you to use newlines (also known as end of lines) in the message.
 - The noninteractive mode, which uses the `-message` parameter to specify the message string, does not support newlines.
- You can upload content from an FTP or HTTP location to use for the banner or MOTD.
- You can configure the MOTD to display dynamic content.
 - Examples of what you can configure the MOTD to display dynamically include the following:
 - Cluster name, node name, or SVM name
 - Cluster date and time
 - Name of the user logging in
 - Last login for the user on any node in the cluster

- Login device name or IP address
- Operating system name
- Software release version
- Effective cluster version string

The `security login motd modify` man page describes the escape sequences that you can use to enable the MOTD to display dynamically generated content.

The banner does not support dynamic content.

You can manage the banner and MOTD at the cluster or SVM level:

- The following facts apply to the banner:
 - The banner configured for the cluster is also used for all SVMs that do not have a banner message defined.
 - An SVM-level banner can be configured for each SVM.
If a cluster-level banner has been configured, it is overridden by the SVM-level banner for the given SVM.
- The following facts apply to the MOTD:
 - By default, the MOTD configured for the cluster is also enabled for all SVMs.
 - Additionally, an SVM-level MOTD can be configured for each SVM.
In this case, users logging in to the SVM will see two MOTDs, one defined at the cluster level and the other at the SVM level.
 - The cluster-level MOTD can be enabled or disabled on a per-SVM basis by the cluster administrator.
If the cluster administrator disables the cluster-level MOTD for an SVM, a user logging in to the SVM does not see the cluster-level MOTD.

Creating a banner

You can create a banner to display a message to someone who attempts to access the cluster or SVM. The banner is displayed in a console session (for cluster access only) or an SSH session (for cluster or SVM access) before a user is prompted for authentication.

Steps

1. Use the `security login banner modify` command to create a banner for the cluster or SVM:

If you want to...	Then...
Specify a message that is a single line	Use the <code>-message "text"</code> parameter to specify the text.
Include newlines (also known as end of lines) in the message	Use the command without the <code>-message</code> or <code>-uri</code> parameter to launch the interactive mode for editing the banner.
Upload content from a location to use for the banner	Use the <code>-uri</code> parameter to specify the content's FTP or HTTP location.

The maximum size for a banner is 2,048 bytes, including newlines.

A banner created by using the `-uri` parameter is static. It is not automatically refreshed to reflect subsequent changes of the source content.

The banner created for the cluster is displayed also for all SVMs that do not have an existing banner. Any subsequently created banner for an SVM overrides the cluster-level banner for that SVM. Specifying the `-message` parameter with a hyphen within double quotes ("`-`") for the SVM resets the SVM to use the cluster-level banner.

2. Verify that the banner has been created by displaying it with the `security login banner show` command.

Specifying the `-message` parameter with an empty string ("`''`") displays banners that have no content.

Specifying the `-message` parameter with "`-`" displays all (admin or data) SVMs that do not have a banner configured.

Examples of creating banners

The following example uses the noninteractive mode to create a banner for the "cluster1" cluster:

```
cluster1::> security login banner modify -message "Authorized users only!"
cluster1::>
```

The following example uses the interactive mode to create a banner for the "svm1" SVM:

```
cluster1::> security login banner modify -vserver svm1
Enter the message of the day for Vserver "svm1".
Max size: 2048. Enter a blank line to terminate input. Press Ctrl-C to abort.
0          1          2          3          4          5          6          7          8
1234567890123456789012345678901234567890123456789012345678901234567890
The svm1 SVM is reserved for authorized users only!

cluster1::>
```

The following example displays the banners that have been created:

```
cluster1::> security login banner show
Vserver: cluster1
Message
-----
Authorized users only!

Vserver: svm1
Message
-----
The svm1 SVM is reserved for authorized users only!

2 entries were displayed.

cluster1::>
```

Related tasks

[Managing the banner](#) on page 168

Related information

[Clustered Data ONTAP 8.3.2 man page: security login banner modify - Modify the login banner message](#)

[Clustered Data ONTAP 8.3.2 man page: security login banner show - Display the login banner message](#)

Managing the banner

You can manage the banner at the cluster or SVM level. The banner configured for the cluster is also used for all SVMs that do not have a banner message defined. A subsequently created banner for an SVM overrides the cluster banner for that SVM.

Choices

- Manage the banner at the cluster level:

If you want to...	Then...
Create a banner to display for all CLI login sessions	Set a cluster-level banner: <code>security login banner modify -vserver <i>cluster_name</i> { [-message "text"] [-uri <i>ftp_or_http_addr</i>] }</code>
Remove the banner for all (cluster and SVM) logins	Set the banner to an empty string (""): <code>security login banner modify -vserver * -message ""</code>
Override a banner created by an SVM administrator	Modify the SVM banner message: <code>security login banner modify -vserver <i>svm_name</i> { [-message "text"] [-uri <i>ftp_or_http_addr</i>] }</code>

- Manage the banner at the SVM level:

Specifying `-vserver svm_name` is not required in the SVM context.

If you want to...	Then...
Override the banner supplied by the cluster administrator with a different banner for the SVM	Create a banner for the SVM: <code>security login banner modify -vserver <i>svm_name</i> { [-message "text"] [-uri <i>ftp_or_http_addr</i>] }</code>
Suppress the banner supplied by the cluster administrator so that no banner is displayed for the SVM	Set the SVM banner to an empty string for the SVM: <code>security login banner modify -vserver <i>svm_name</i> -message ""</code>
Use the cluster-level banner when the SVM currently uses an SVM-level banner	Set the SVM banner to "-": <code>security login banner modify -vserver <i>svm_name</i> -message "-"</code>

Related tasks

[Creating a banner](#) on page 166

Related information

[Clustered Data ONTAP 8.3.2 man page: security login banner modify - Modify the login banner message](#)

[Clustered Data ONTAP 8.3.2 man page: security login banner show - Display the login banner message](#)

Creating an MOTD

You can create a message of the day (MOTD) to communicate information to authenticated CLI users. The MOTD is displayed in a console session (for cluster access only) or an SSH session (for cluster or SVM access) after a user is authenticated but before the `clustershell` prompt appears.

Steps

1. Use the `security login motd modify` command to create an MOTD for the cluster or SVM:

If you want to...	Then...
Specify a message that is a single line	Use the <code>-message "text"</code> parameter to specify the text.
Include newlines (also known as end of lines)	Use the command without the <code>-message</code> or <code>-uri</code> parameter to launch the interactive mode for editing the MOTD.
Upload content from a location to use for the MOTD	Use the <code>-uri</code> parameter to specify the content's FTP or HTTP location.

The maximum size for an MOTD is 2,048 bytes, including newlines.

The `security login motd modify man` page describes the escape sequences that you can use to enable the MOTD to display dynamically generated content.

An MOTD created by using the `-uri` parameter is static. It is not automatically refreshed to reflect subsequent changes of the source content.

An MOTD created for the cluster is displayed also for all SVM logins by default, along with an SVM-level MOTD that you can create separately for a given SVM. Setting the `-is-cluster-message-enabled` parameter to **false** for an SVM prevents the cluster-level MOTD from being displayed for that SVM.

2. Verify that the MOTD has been created by displaying it with the `security login motd show` command.

Specifying the `-message` parameter with an empty string (`""`) displays MOTDs that are unconfigured or have no content.

Examples of creating MOTDs

The following example uses the noninteractive mode to create an MOTD for the “cluster1” cluster:

```
cluster1::> security login motd modify -message "Greetings!"
cluster1::>
```

The following example uses the interactive mode to create an MOTD for the “svm1” SVM that uses escape sequences to display dynamically generated content:

```
cluster1::> security login motd modify -vserver svm1

Enter the message of the day for Vserver "svm1".
Max size: 2048. Enter a blank line to terminate input. Press Ctrl-C to abort.
0          1          2          3          4          5          6          7          8
1234567890123456789012345678901234567890123456789012345678901234567890
Welcome to the \n SVM. Your user ID is '\N'. Your last successful login was \L.

cluster1::>
```

The following example displays the MOTDs that have been created:

```

cluster1::> security login motd show
Vserver: cluster1
Is the Cluster MOTD Displayed?: true
Message
-----
Greetings!

Vserver: svm1
Is the Cluster MOTD Displayed?: true
Message
-----
Welcome to the \n SVM. Your user ID is '\N'. Your last successful login was \L.

2 entries were displayed.

cluster1::>

```

Related tasks

[Managing the MOTD](#) on page 170

Related information

[Clustered Data ONTAP 8.3.2 man page: security login motd modify - Modify the message of the day](#)

[Clustered Data ONTAP 8.3.2 man page: security login motd show - Display the message of the day](#)

Managing the MOTD

You can manage the message of the day (MOTD) at the cluster or SVM level. By default, the MOTD configured for the cluster is also enabled for all SVMs. Additionally, an SVM-level MOTD can be configured for each SVM. The cluster-level MOTD can be enabled or disabled for each SVM by the cluster administrator.

Choices

- Manage the MOTD at the cluster level:

If you want to...	Then...
Create an MOTD for all logins when there is no existing MOTD	Set a cluster-level MOTD: <code>security login motd modify -vserver <i>cluster_name</i> { [-message "<i>text</i>"] [-uri <i>ftp_or_http_addr</i>] }</code>
Change the MOTD for all logins when no SVM-level MOTDs are configured	Modify the cluster-level MOTD: <code>security login motd modify -vserver <i>cluster_name</i> { [-message "<i>text</i>"] [-uri <i>ftp_or_http_addr</i>] }</code>
Remove the MOTD for all logins when no SVM-level MOTDs are configured	Set the cluster-level MOTD to an empty string (""): <code>security login motd modify -vserver <i>cluster_name</i> -message ""</code>

If you want to...	Then...
Have every SVM display the cluster-level MOTD instead of using the SVM-level MOTD	<p>Set a cluster-level MOTD, then set all SVM-level MOTDs to an empty string with the cluster-level MOTD enabled:</p> <ol style="list-style-type: none"> <code>security login motd modify -vserver <i>cluster_name</i> { [-message "text"] [-uri <i>ftp_or_http_addr</i>] }</code> <code>security login motd modify { -vserver !"cluster_name" } -message "" -is-cluster-message-enabled true</code>
Have an MOTD displayed for only selected SVMs, and use no cluster-level MOTD	<p>Set the cluster-level MOTD to an empty string, then set SVM-level MOTDs for selected SVMs:</p> <ol style="list-style-type: none"> <code>security login motd modify -vserver <i>cluster_name</i> -message ""</code> <code>security login motd modify -vserver <i>svm_name</i> { [-message "text"] [-uri <i>ftp_or_http_addr</i>] }</code> <p>You can repeat this step for each SVM as needed.</p>
Use the same SVM-level MOTD for all (data and admin) SVMs	<p>Set the cluster and all SVMs to use the same MOTD:</p> <pre>security login motd modify -vserver * { [-message "text"] [-uri <i>ftp_or_http_addr</i>] }</pre> <p>Note: If you use the interactive mode, the CLI prompts you to enter the MOTD individually for the cluster and each SVM. You can paste the same MOTD into each instance when you are prompted to.</p>
Have a cluster-level MOTD optionally available to all SVMs, but do not want the MOTD displayed for cluster logins	<p>Set a cluster-level MOTD, but disable its display for the cluster:</p> <pre>security login motd modify -vserver <i>cluster_name</i> { [-message "text"] [-uri <i>ftp_or_http_addr</i>] } -is-cluster-message-enabled false</pre>
Remove all MOTDs at the cluster and SVM levels when only some SVMs have both cluster-level and SVM-level MOTDs	<p>Set the cluster and all SVMs to use an empty string for the MOTD:</p> <pre>security login motd modify -vserver * -message ""</pre>
Modify the MOTD only for the SVMs that have a non-empty string, when other SVMs use an empty string, and when a different MOTD is used at the cluster level	<p>Use extended queries to modify the MOTD selectively:</p> <pre>security login motd modify { -vserver !"cluster_name" -message !"" } { [-message "text"] [-uri <i>ftp_or_http_addr</i>] }</pre>
Display all MOTDs that contain specific text (for example, “January” followed by “2015”) anywhere in a single or multiline message, even if the text is split across different lines	<p>Use a query to display MOTDs:</p> <pre>security login motd show -message *"January"*"2015"*</pre>
Interactively create an MOTD that includes multiple and consecutive newlines (also known as end of lines, or EOLs)	<p>In the interactive mode, press the space bar followed by Enter to create a blank line without terminating the input for the MOTD.</p>

- Manage the MOTD at the SVM level:

Specifying `-vserver svm_name` is not required in the SVM context.

If you want to...	Then...
Use a different SVM-level MOTD, when the SVM already has an existing SVM-level MOTD	Modify the SVM-level MOTD: <pre>security login motd modify -vserver svm_name { [-message "text"] [-uri ftp_or_http_addr] }</pre>
Use only the cluster-level MOTD for the SVM, when the SVM already has an SVM-level MOTD	Set the SVM-level MOTD to an empty string, then have the cluster administrator enable the cluster-level MOTD for the SVM: <ol style="list-style-type: none"> 1. <code>security login motd modify -vserver svm_name -message ""</code> 2. (For the cluster administrator) <code>security login motd modify -vserver svm_name -is-cluster-message-enabled true</code>
Not have the SVM display any MOTD, when both the cluster-level and SVM-level MOTDs are currently displayed for the SVM	Set the SVM-level MOTD to an empty string, then have the cluster administrator disable the cluster-level MOTD for the SVM: <ol style="list-style-type: none"> 1. <code>security login motd modify -vserver svm_name -message ""</code> 2. (For the cluster administrator) <code>security login motd modify -vserver svm_name -is-cluster-message-enabled false</code>

Related tasks

[Creating an MOTD](#) on page 169

Related information

[Clustered Data ONTAP 8.3.2 man page: security login motd modify - Modify the message of the day](#)

[Clustered Data ONTAP 8.3.2 man page: security login motd show - Display the message of the day](#)

Managing licenses (cluster administrators only)

A license is a record of one or more software entitlements. Installing license keys, also known as *license codes*, enables you to use certain features or services on your cluster. Data ONTAP enables you to manage feature licenses and monitor feature usage and license entitlement risk.

Each cluster requires a cluster base license key, which you can install either during or after the cluster setup. Some features require additional licenses. Data ONTAP feature licenses are issued as *packages*, each of which contains multiple features or a single feature. A package requires a license key, and installing the key enables you to access all features in the package. Data ONTAP prevents you from installing a feature license before a cluster base license key is installed.

Starting with Data ONTAP 8.2, all license keys are 28 characters in length. Licenses installed prior to Data ONTAP 8.2 continue to work in Data ONTAP 8.2 and later releases. However, if you need to reinstall a license (for example, you deleted a previously installed license and want to reinstall it in Data ONTAP 8.2 or later, or you perform a controller replacement procedure for a node in a cluster running Data ONTAP 8.2 or later), Data ONTAP requires that you enter the license key in the 28-character format.

You can find license keys for your initial or add-on software orders at the NetApp Support Site under **My Support > Software Licenses** (login required). If you cannot locate your license keys from the Software Licenses page, contact your sales or support representative.

Data ONTAP enables you to manage feature licenses in the following ways:

- Add one or more license keys (`system license add`)
- Display information about installed licenses (`system license show`)
- Display the packages that require licenses and their current license status on the cluster (`system license status show`)
- Delete a license from the cluster or a node whose serial number you specify (`system license delete`)
The cluster base license is required for the cluster to operate. Data ONTAP does not enable you to delete it.
- Display or remove expired or unused licenses (`system license clean-up`)

Data ONTAP enables you to monitor feature usage and license entitlement risk in the following ways:

- Display a summary of feature usage in the cluster on a per-node basis (`system feature-usage show-summary`)
The summary includes counter information such as the number of weeks a feature was in use and the last date and time the feature was used.
- Display feature usage status in the cluster on a per-node and per-week basis (`system feature-usage show-history`)
The feature usage status can be **not-used**, **configured**, or **in-use**. If the usage information is not available, the status shows **not-available**.
- Display the status of license entitlement risk for each license package (`system license entitlement-risk show`)
The risk status can be **low**, **medium**, **high**, **unlicensed**, or **unknown**. The risk status is also included in the AutoSupport message. License entitlement risk does not apply to the base license package.
The license entitlement risk is evaluated by using a number of factors, which might include but are not limited to the following:

- Each package's licensing state
- The type of each license, its expiry status, and the uniformity of the licenses across the cluster
- Usage for the features associated with the license package

If the evaluation process determines that the cluster has a license entitlement risk, the command output also suggests a corrective action.

Related information

[NetApp KB Article 3013749: Data ONTAP 8.2 and 8.3 Licensing Overview and References](#)

[NetApp KB Article 1014509: How to verify Data ONTAP Software Entitlements and related License Keys using the Support Site](#)

[NetApp: Data ONTAP Entitlement Risk Status](#)

License types and licensed method

Understanding license types and the licensed method helps you manage the licenses in a cluster.

License types

A package can have one or more of the following types of license installed in the cluster. The `system license show` command displays the installed license type or types for a package.

- **Standard license (`license`)**

A standard license is a node-locked license. It is issued for a node with a specific system serial number (also known as a *controller serial number*). A standard license is valid only for the node that has the matching serial number.

Installing a standard, node-locked license entitles a node to the licensed functionality. For the cluster to use licensed functionality, at least one node must be licensed for the functionality. It might be out of compliance to use licensed functionality on a node that does not have an entitlement for the functionality.

Data ONTAP 8.2 and later releases treat a license installed prior to Data ONTAP 8.2 as a standard license. Therefore, in Data ONTAP 8.2 and later releases, all nodes in the cluster automatically have the standard license for the package that the previously licensed functionality is part of. The `system license show` command with the `-legacy yes` parameter indicates such licenses.
- **Site license (`site`)**

A site license is not tied to a specific system serial number. When you install a site license, all nodes in the cluster are entitled to the licensed functionality. The `system license show` command displays site licenses under the cluster serial number.

If your cluster has a site license and you remove a node from the cluster, the node does not carry the site license with it, and it is no longer entitled to the licensed functionality. If you add a node to a cluster that has a site license, the node is automatically entitled to the functionality granted by the site license.
- **Evaluation license (`demo`)**

An evaluation license is a temporary license that expires after a certain period of time (indicated by the `system license show` command). It enables you to try certain software functionality without purchasing an entitlement. It is a cluster-wide license, and it is not tied to a specific serial number of a node.

If your cluster has an evaluation license for a package and you remove a node from the cluster, the node does not carry the evaluation license with it.

Licensed method

It is possible to install both a cluster-wide license (the **site** or **demo** type) and a node-locked license (the **license** type) for a package. Therefore, an installed package can have multiple license types in the cluster. However, to the cluster, there is only one *licensed method* for a package. The **licensed method** field of the `system license status show` command displays the entitlement that is being used for a package. The command determines the licensed method as follows:

- If a package has only one license type installed in the cluster, the installed license type is the licensed method.
- If a package does not have any licenses installed in the cluster, the licensed method is **none**.
- If a package has multiple license types installed in the cluster, the licensed method is determined in the following priority order of the license type—**site**, **license**, and **demo**.

For example:

- If you have a site license, a standard license, and an evaluation license for a package, the licensed method for the package in the cluster is **site**.
- If you have a standard license and an evaluation license for a package, the licensed method for the package in the cluster is **license**.
- If you have only an evaluation license for a package, the licensed method for the package in the cluster is **demo**.

Commands for managing licenses

You use the `system license` commands to manage feature licenses for the cluster. You use the `system feature-usage` commands to monitor feature usage.

If you want to...	Use this command...
Add one or more licenses	<code>system license add</code>
Display information about installed licenses, for example: <ul style="list-style-type: none"> • License package name and description • License type (site, license, or demo) • Expiration date, if applicable • The cluster or nodes that a package is licensed for • Whether the license was installed prior to Data ONTAP 8.2 (legacy) • Customer ID 	<code>system license show</code> <p>Note: Some information is displayed only when you use the <code>-instance</code> parameter.</p>
Display all packages that require licenses and their current license status, including the following: <ul style="list-style-type: none"> • The package name • The licensed method • The expiration date, if applicable 	<code>system license status show</code>

If you want to...	Use this command...
Delete the license of a package from the cluster or a node whose serial number you specify	<code>system license delete</code>
Display or remove expired or unused licenses	<code>system license clean-up</code>
Display summary of feature usage in the cluster on a per-node basis	<code>system feature-usage show-summary</code>
Display feature usage status in the cluster on a per-node and per-week basis	<code>system feature-usage show-history</code>
Display the status of license entitlement risk for each license package	<code>system license entitlement-risk show</code> Note: Some information is displayed only when you use the <code>-detail</code> and <code>-instance</code> parameters.

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Managing jobs and schedules

A *job* is any asynchronous task that is managed by the Job Manager. Jobs are typically long-running volume operations such as copy, move, and mirror. You can monitor, pause, stop, and restart jobs, and you can configure them to run on specified schedules.

Job categories

There are three categories of jobs that you can manage: server-affiliated, cluster-affiliated, and private.

A job can be in any of the following categories:

Server-Affiliated jobs

These jobs are queued by the management framework to a specific node to be run.

Cluster-Affiliated jobs

These jobs are queued by the management framework to any node in the cluster to be run.

Private jobs

These jobs are specific to a node and do not use the replicated database (RDB) or any other cluster mechanism. The commands that manage private jobs require the advanced privilege level or higher.

Commands for managing jobs

Jobs are placed into a job queue and run in the background when resources are available. If a job is consuming too many cluster resources, you can stop it or pause it until there is less demand on the cluster. You can also monitor and restart jobs.

When you enter a command that invokes a job, typically, the command informs you that the job has been queued and then returns to the CLI command prompt. However, some commands instead report job progress and do not return to the CLI command prompt until the job has been completed. In these cases, you can press Ctrl-C to move the job to the background.

If you want to...	Use this command...
Display information about all jobs	<i>job show</i>
Display information about jobs on a per-node basis	<i>job show bynode</i>
Display information about cluster-affiliated jobs	<i>job show-cluster</i>
Display information about completed jobs	<i>job show-completed</i>
Display information about job history	<i>job history show</i> Up to 25,000 job records are stored for each node in the cluster. Consequently, attempting to display the full job history could take a long time. To avoid potentially long wait times, you should display jobs by node, Storage Virtual Machine (SVM), or record ID.

If you want to...	Use this command...
Display the list of private jobs	<code>job private show</code> (advanced privilege level)
Display information about completed private jobs	<code>job private show-completed</code> (advanced privilege level)
Display information about the initialization state for job managers	<code>job initstate show</code> (advanced privilege level)
Monitor the progress of a job	<code>job watch-progress</code>
Monitor the progress of a private job	<code>job private watch-progress</code> (advanced privilege level)
Pause a job	<code>job pause</code>
Pause a private job	<code>job private pause</code> (advanced privilege level)
Resume a paused job	<code>job resume</code>
Resume a paused private job	<code>job private resume</code> (advanced privilege level)
Stop a job	<code>job stop</code>
Stop a private job	<code>job private stop</code> (advanced privilege level)
Delete a job	<code>job delete</code>
Delete a private job	<code>job private delete</code> (advanced privilege level)
Disassociate a cluster-affiliated job with an unavailable node that owns it, so that another node can take ownership of that job	<code>job unclaim</code> (advanced privilege level)

Note: You can use the `event log show` command to determine the outcome of a completed job.

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Commands for managing job schedules

Many tasks—for instance, volume Snapshot copies—can be configured to run on specified schedules. Schedules that run at specific times are called *cron* schedules (similar to UNIX `cron` schedules). Schedules that run at intervals are called *interval* schedules. You use the `job schedule` commands to manage job schedules.

Job schedules do not adjust to manual changes to the cluster date and time. These jobs are scheduled to run based on the current cluster time when the job was created or when the job most recently ran. Therefore, if you manually change the cluster date or time, you should use the `job show` and `job history show` commands to verify that all scheduled jobs are queued and completed according to your requirements.

If the cluster is part of a MetroCluster configuration, then the job schedules on both clusters must be identical. Therefore, if you create, modify, or delete a job schedule, you must perform the same operation on the remote cluster.

If you want to...	Use this command...
Display information about all schedules	<code>job schedule show</code>
Display the list of jobs by schedule	<code>job schedule show-jobs</code>
Display information about cron schedules	<code>job schedule cron show</code>
Display information about interval schedules	<code>job schedule interval show</code>
Create a cron schedule	<code>job schedule cron create</code>
Create an interval schedule	<code>job schedule interval create</code> You must specify at least one of the following parameters: -days, -hours, -minutes, or -seconds.
Modify a cron schedule	<code>job schedule cron modify</code>
Modify an interval schedule	<code>job schedule interval modify</code>
Delete a schedule	<code>job schedule delete</code>
Delete a cron schedule	<code>job schedule cron delete</code>
Delete an interval schedule	<code>job schedule interval delete</code>

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Backing up and restoring cluster configurations (cluster administrators only)

Backing up the cluster configuration enables you to restore the configuration of any node or the cluster in the event of a disaster or emergency.

What configuration backup files are

Configuration backup files are archive files (.7z) that contain information for all configurable options that are necessary for the cluster, and the nodes within it, to operate properly.

These files store the local configuration of each node, plus the cluster-wide replicated configuration. You use configuration backup files to back up and restore the configuration of your cluster.

There are two types of configuration backup files:

Node configuration backup file

Each healthy node in the cluster includes a node configuration backup file, which contains all of the configuration information and metadata necessary for the node to operate healthy in the cluster.

Cluster configuration backup file

These files include an archive of all of the node configuration backup files in the cluster, plus the replicated cluster configuration information (the replicated database, or RDB file). Cluster configuration backup files enable you to restore the configuration of the entire cluster, or of any node in the cluster. The cluster configuration backup schedules create these files automatically and store them on several nodes in the cluster.

Note: Configuration backup files contain configuration information only. They do not include any user data. For information about restoring user data, see the *Clustered Data ONTAP Data Protection Guide*.

Managing configuration backups

The configuration backup schedules automatically create configuration backup files for each node in the cluster, and for the cluster itself. You can change some of the settings for these schedules, and you can create configuration backup files manually.

How the node and cluster configurations are backed up automatically

Three separate schedules automatically create cluster and node configuration backup files and replicate them among the nodes in the cluster.

The configuration backup files are automatically created according to the following schedules:

- Every 8 hours
- Daily
- Weekly

At each of these times, a node configuration backup file is created on each healthy node in the cluster. All of these node configuration backup files are then collected in a single cluster configuration backup file along with the replicated cluster configuration and saved on one or more nodes in the cluster.

For single-node clusters (including Data ONTAP Edge systems), you can specify the configuration backup destination during software setup. After setup, those settings can be modified using Data ONTAP commands.

Commands for managing configuration backup schedules

You use the `system configuration backup settings` commands to manage configuration backup schedules.

These commands are available at the advanced privilege level.

If you want to...	Use this command...
Change the settings for a configuration backup schedule: <ul style="list-style-type: none"> Specify a remote URL (either HTTP or FTP) where the configuration backup files will be uploaded in addition to the default locations in the cluster Specify a user name to be used to log in to the remote URL Set the number of backups to keep for each configuration backup schedule 	<code>system configuration backup settings modify</code>
Set the password to be used to log in to the remote URL	<code>system configuration backup settings set-password</code>
View the settings for the configuration backup schedule	<code>system configuration backup settings show</code> <p>Note: You set the <code>-instance</code> parameter to view the user name and the number of backups to keep for each schedule.</p>

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Commands for managing configuration backup files

You use the `system configuration backup` commands to manage cluster and node configuration backup files.

These commands are available at the advanced privilege level.

If you want to...	Use this command...
Create a new node or cluster configuration backup file	<code>system configuration backup create</code>
Copy a configuration backup file from a node to another node in the cluster	<code>system configuration backup copy</code>

If you want to...	Use this command...
Upload a configuration backup file from a node in the cluster to a remote URL (either HTTP or FTP)	<code>system configuration backup upload</code> Note: The web server to which you are uploading the configuration backup file must have PUT operations enabled. For more information, see your web server's documentation.
Download a configuration backup file from a remote URL to a node in the cluster	<code>system configuration backup download</code>
Rename a configuration backup file on a node in the cluster	<code>system configuration backup rename</code>
View the node and cluster configuration backup files for one or more nodes in the cluster	<code>system configuration backup show</code>
Delete a configuration backup file on a node	<code>system configuration backup delete</code> Note: This command deletes the configuration backup file on the specified node only. If the configuration backup file also exists on other nodes in the cluster, it remains on those nodes.

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Recovering a node configuration

You recover a node's configuration using a configuration backup file if the node, its root volume, or any of its configuration information is lost or corrupted.

Steps

1. [Finding a configuration backup file to use for recovering a node](#) on page 182
2. [Restoring the node configuration using a configuration backup file](#) on page 183

Finding a configuration backup file to use for recovering a node

You use a configuration backup file located at a remote URL or on a node in the cluster to recover a node configuration.

About this task

You can use either a cluster or node configuration backup file to restore a node configuration.

Step

1. Make the configuration backup file available to the node for which you need to restore the configuration.

If the configuration backup file is located...	Then...
--	---------

At a remote URL	Use the <code>system configuration backup download</code> command at the advanced privilege level to download it to the recovering node.
-----------------	--

If the configuration backup file is located...	Then...
On a node in the cluster	<ol style="list-style-type: none"> <li data-bbox="665 283 1377 409">a. Use the <code>system configuration backup show</code> command at the advanced privilege level to view the list of configuration backup files available in the cluster that contains the recovering node's configuration. <li data-bbox="665 420 1377 514">b. If the configuration backup file you identify does not exist on the recovering node, then use the <code>system configuration backup copy</code> command to copy it to the recovering node.

If you previously re-created the cluster, you should choose a configuration backup file that was created after the cluster recreation. If you must use a configuration backup file that was created prior to the cluster recreation, then after recovering the node, you must re-create the cluster again.

Restoring the node configuration using a configuration backup file

You restore the node configuration using the configuration backup file that you identified and made available to the recovering node.

About this task

You should only perform this task to recover from a disaster that resulted in the loss of the node's local configuration files.

Steps

1. If the node is healthy, then at the advanced privilege level of a different node, use the `cluster modify` command with the `-node` and `-eligibility` parameters to mark it ineligible and isolate it from the cluster.

If the node is not healthy, then you should skip this step.

Example

This example modifies node2 to be ineligible to participate in the cluster so that its configuration can be restored:

```
cluster1::*> cluster modify -node node2 -eligibility false
```

2. Use the `system configuration recovery node restore` command at the advanced privilege level to restore the node's configuration from a configuration backup file. If the node lost its identity, including its name, then you should use the `-nodename-in-backup` parameter to specify the node name in the configuration backup file.

Example

This example restores the node's configuration using one of the configuration backup files stored on the node:

```
cluster1::*> system configuration recovery node restore -backup
cluster1.8hour.2011-02-22.18_15_00.7z
```

```
Warning: This command overwrites local configuration files with
files contained in the specified backup file. Use this
command only to recover from a disaster that resulted
in the loss of the local configuration files.
The node will reboot after restoring the local configuration.
Do you want to continue? {y|n}: y
```

The configuration is restored, and the node reboots.

3. If you marked the node ineligible, then use the `system configuration recovery cluster sync` command to mark the node as eligible and synchronize it with the cluster.
4. If you are operating in a SAN environment, use the `system node reboot` command to reboot the node and reestablish SAN quorum.

After you finish

If you previously re-created the cluster, and if you are restoring the node configuration by using a configuration backup file that was created prior to that cluster re-creation, then you must re-create the cluster again.

Related tasks

[Synchronizing a node with the cluster](#) on page 187

Recovering a cluster configuration

If cluster-wide quorum does not exist, then you recover the cluster configuration by finding a configuration to use for recovery, re-creating the cluster, and then rejoining each node to it.

Steps

1. [Finding a configuration to use for recovering a cluster](#) on page 184
2. [Restoring a cluster configuration from an existing configuration](#) on page 185

Finding a configuration to use for recovering a cluster

You use the configuration from either a node in the cluster or a cluster configuration backup file to recover a cluster.

Steps

1. Choose a type of configuration to recover the cluster.
 - A node in the cluster

If the cluster consists of more than one node, and one of the nodes has a cluster configuration from when the cluster was in the desired configuration, then you can recover the cluster using the configuration stored on that node.

In most cases, the node containing the replication ring with the most recent transaction ID is the best node to use for restoring the cluster configuration. The `cluster ring show` command at the advanced privilege level enables you to view a list of the replicated rings available on each node in the cluster.
 - A cluster configuration backup file

If you cannot identify a node with the correct cluster configuration, or if the cluster consists of a single node, then you can use a cluster configuration backup file to recover the cluster.

Recovering a cluster from a configuration backup file might require the use of diagnostic tools to resolve discrepancies between the backup and the cluster configuration present on the recovering node. Therefore, if you use a configuration backup file, you should plan to contact technical support.
2. If you chose to use a cluster configuration backup file, then make the file available to the node you plan to use to recover the cluster.

If the configuration backup file is located...	Then...
At a remote URL	Use the <code>system configuration backup download</code> command at the advanced privilege level to download it to the recovering node.
On a node in the cluster	<ol style="list-style-type: none"> <li data-bbox="667 369 1377 453">a. Use the <code>system configuration backup show</code> command at the advanced privilege level to find a cluster configuration backup file that was created when the cluster was in the desired configuration. <li data-bbox="667 474 1377 590">b. If the cluster configuration backup file is not located on the node you plan to use to recover the cluster, then use the <code>system configuration backup copy</code> command to copy it to the recovering node.

Restoring a cluster configuration from an existing configuration

To restore a cluster configuration from an existing configuration after a cluster failure, you re-create the cluster using the cluster configuration that you chose and made available to the recovering node, and then rejoin each additional node to the new cluster.

About this task

You should only perform this task to recover from a disaster that resulted in the loss of the cluster's configuration.

Attention: If you are re-creating the cluster from a configuration backup file, you must contact technical support to resolve any discrepancies between the configuration backup file and the configuration present in the cluster.

Steps

1. Disable storage failover for each HA pair:

```
storage failover modify -node node_name -enabled false
```

You only need to disable storage failover once for each HA pair. When you disable storage failover for a node, storage failover is also disabled on the node's partner.

2. Halt each node except for the recovering node:

```
system node halt -node node_name -reason "text"
```

Example

```
cluster1::*> system node halt -node node0 -reason "recovering cluster"
Warning: Are you sure you want to halt the node? {y|n}: y
```

3. Set the privilege level to advanced:

```
set -privilege advanced
```

4. On the recovering node, use the `system configuration recovery cluster recreate` command to re-create the cluster.

Example

This example re-creates the cluster using the configuration information stored on the recovering node:

```
cluster1::*> configuration recovery cluster recreate -from node

Warning: This command will destroy your existing cluster. It will
rebuild a new single-node cluster consisting of this node
and its current configuration. This feature should only be
used to recover from a disaster. Do not perform any other
recovery operations while this operation is in progress.
Do you want to continue? {y|n}: y
```

A new cluster is created on the recovering node.

5. If you are re-creating the cluster from a configuration backup file, verify that the cluster recovery is still in progress:

```
system configuration recovery cluster show
```

You do not need to verify the cluster recovery state if you are re-creating the cluster from a healthy node.

Example

```
cluster1::*> system configuration recovery cluster show
Recovery Status: in-progress
Is Recovery Status Persisted: false
```

6. Boot each node that needs to be rejoined to the re-created cluster.

You must reboot the nodes one at a time.

7. For each node that needs to be joined to the re-created cluster, do the following:
 - a. From a healthy node on the re-created cluster, rejoin the target node:

```
system configuration recovery cluster rejoin -node node_name
```

Example

This example rejoins the “node2” target node to the re-created cluster:

```
cluster1::*> system configuration recovery cluster rejoin -node
node2

Warning: This command will rejoin node "node2" into the local
cluster, potentially overwriting critical cluster
configuration files. This command should only be used
to recover from a disaster. Do not perform any other
recovery operations while this operation is in progress.
This command will cause node "node2" to reboot.
Do you want to continue? {y|n}: y
```

The target node reboots and then joins the cluster.

- b. Verify that the target node is healthy and has formed quorum with the rest of the nodes in the cluster:

```
cluster show -eligibility true
```

The target node must rejoin the re-created cluster before you can rejoin another node.

Example

```
cluster1::*> cluster show -eligibility true
Node           Health Eligibility Epsilon
-----
node0          true   true        false
node1          true   true        false
2 entries were displayed.
```

8. If you re-created the cluster from a configuration backup file, set the recovery status to be complete:


```
system configuration recovery cluster modify -recovery-status complete
```
9. Return to the admin privilege level:


```
set -privilege admin
```
10. If the cluster consists of only two nodes, use the `cluster ha modify` command to reenables cluster HA.
11. Use the `storage failover modify` command to reenables storage failover for each HA pair.

After you finish

If the cluster has SnapMirror peer relationships, then you also need to re-create those relationships. For more information, see the *Clustered Data ONTAP Data Protection Guide*.

Synchronizing a node with the cluster

If cluster-wide quorum exists, but one or more nodes are out of sync with the cluster, then you must synchronize the node to restore the replicated database (RDB) on the node and bring it into quorum.

Step

1. From a healthy node, use the `system configuration recovery cluster sync` command at the advanced privilege level to synchronize the node that is out of sync with the cluster configuration.

Example

This example synchronizes a node (*node2*) with the rest of the cluster:

```
cluster1::*> system configuration recovery cluster sync -node node2

Warning: This command will synchronize node "node2" with the cluster
configuration, potentially overwriting critical cluster
configuration files on the node. This feature should only be
used to recover from a disaster. Do not perform any other
recovery operations while this operation is in progress. This
command will cause all the cluster applications on node
"node2" to restart, interrupting administrative CLI and Web
interface on that node.
Do you want to continue? {y|n}: y
All cluster applications on node "node2" will be restarted. Verify
that the cluster applications go online.
```

Result

The RDB is replicated to the node, and the node becomes eligible to participate in the cluster.

Managing core dumps (cluster administrators only)

When a node panics, a core dump occurs and the system creates a core dump file that technical support can use to troubleshoot the problem. You can configure or display core dump attributes. You can also save, display, segment, upload, or delete a core dump file.

You can manage core dumps in the following ways:

- Configuring core dumps and displaying the configuration settings
- Displaying basic information, the status, and attributes of core dumps
Core dump files and reports are stored in the `/mroot/etc/crash/` directory of a node. You can display the directory content by using the `system node coredump` commands or a web browser.
- Saving the core dump content and uploading the saved file to a specified location or to technical support
Data ONTAP prevents you from initiating the saving of a core dump file during a takeover, an aggregate relocation, or a giveback.
- Deleting core dump files that are no longer needed

A core dump file can be very large and time-consuming to upload. You must not further compress a core dump file. However, you can segment the file in the following ways:

- Configure the automatic segmenting of core dump files
- Manually segment a core dump file and manage the core segments

Related tasks

[Accessing a node's log, core dump, and MIB files by using a web browser](#) on page 41

Methods of segmenting core dump files

A core dump file can be very large, making it time consuming to upload to technical support when you need to. Segmenting the core dump file enables you to upload only the needed portion instead of the entire file.

You can segment a saved core dump file into a maximum of three core segments:

This core segment...	Contains system information from the memory of...
Primary core segment	Data ONTAP and the systemshell
Caching module core segment	Flash Cache family of modules
NVRAM core segment	NVRAM

Segmenting the core dump file enables you to upload a portion of the file as you need to. For instance, instead of uploading the entire core dump file to technical support for a core dump analysis, you can upload only the primary core segment of the file, and if necessary, upload the caching module core segment or NVRAM core segment later.

By using the `system node coredump segment config` commands, you can configure the automatic segmenting of the core dump file in the following ways:

- Specify whether to automatically segment a core dump file after it is saved
The default setting for automatic segmenting is system dependent.
- Specify whether to automatically delete the original core dump file after it is segmented
By default, automatic deletion of the original core dump file is disabled.
- Display the current configuration of the automatic segmenting of core dump files

By using the `system node coredump segment` commands, you can manually manage the segmenting of a core dump file in the following ways:

- Manually schedule a core segmenting job to segment a specified core dump file on a node into core segments and specify whether the original core dump file is to be deleted after the core segmenting is complete
- Display information about core segments
- Delete a specified core segment or all segments from a node
- Display the status of a core segmenting job
- Cancel a core segmenting job as specified by its job ID

Commands for managing core dumps

You use the `system node coredump config` commands to manage the configuration of core dumps, the `system node coredump` commands to manage the core dump files, and the `system node coredump reports` commands to manage application core reports.

If you want to...	Use this command...
Configure core dumps	<code>system node coredump config modify</code>
Display the configuration settings for core dumps	<code>system node coredump config show</code>
Display basic information about core dumps	<code>system node coredump show</code>
Manually trigger a core dump when you reboot a node	<code>system node reboot</code> with both the <code>-dump</code> and <code>-skip-lif-migration</code> parameters
Manually trigger a core dump when you shut down a node	<code>system node halt</code> with both the <code>-dump</code> and <code>-skip-lif-migration</code> parameters
Save a specified core dump	<code>system node coredump save</code>
Save all unsaved core dumps that are on a specified node	<code>system node coredump save-all</code>
Generate and send an AutoSupport message with a core dump file you specify	<code>system node autosupport invoke-core-upload</code> Note: The <code>-uri</code> optional parameter specifies an alternate destination for the AutoSupport message.
Display status information about core dumps	<code>system node coredump status</code>
Delete a specified core dump	<code>system node coredump delete</code>
Delete all unsaved core dumps or all saved core files on a node	<code>system node coredump delete-all</code>

If you want to...	Use this command...
Display application core dump reports	<code>system node coredump reports show</code>
Delete an application core dump report	<code>system node coredump reports delete</code>

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Commands for managing core segmenting

You use the `system node coredump segment config` commands to manage the automatic segmenting of core dump files. You use the `system node coredump segment` commands to manage core segments.

If you want to...	Use this command...
Configure the automatic segmenting of core dump files for a node: <ul style="list-style-type: none"> Whether to automatically segment a core dump file after it is saved Whether to automatically delete the original core dump file after it is segmented 	<code>system node coredump segment config modify</code>
Show the current configuration of automatic core segmenting	<code>system node coredump segment config show</code>
Manually start segmenting a specified core dump file on a node into core segments and specify whether the original core dump file is to be deleted after the core segmenting is complete	<code>system node coredump segment start</code>
Display information about the core segments on a node: <ul style="list-style-type: none"> The core segment name Total number of core segments for the full core The time when the panic occurred that generated the core dump file 	<code>system node coredump segment show</code>
Delete a specified core segment from a node	<code>system node coredump segment delete</code>
Delete all core segments from a node	<code>system node coredump segment delete-all</code>
Display the status of a core segmenting job: <ul style="list-style-type: none"> Job ID Name of the core dump file that is being segmented Job status Percent completed 	<code>system node coredump segment status</code>

If you want to...	Use this command...
Cancel a core segmenting job as specified by its job ID	<code>system node coredump segment stop</code>

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Monitoring the storage system

You can use event messages, the AutoSupport feature, dashboards, statistics, and environmental component sensors to monitor the storage system.

The cluster administrator can perform all system monitoring tasks. The Storage Virtual Machine (SVM) administrator can perform only the following monitoring tasks:

- Display the SVM health dashboard (by using the `dashboard health vserver show` commands)
The `dashboard health vserver show` commands are deprecated, but you can still use them to monitor the system.
- Manage and obtain performance data (by using the `statistics` commands)

Managing event messages

The Event Management System (EMS) collects and displays information about events that occur on your cluster. You can manage the event destination, event route, mail history records, and SNMP trap history records. You can also configure event notification and logging.

Event messages for high-severity events appear on your system console, and are written to the cluster's event log. An event message consists of the following elements:

- Message name
- Severity level
Possible values include the following, listed in decreasing order of urgency:
 - EMERGENCY (the cluster is unusable)
 - ALERT (action must be taken immediately to prevent system failure)
 - CRITICAL
 - ERROR
 - WARNING
 - NOTICE (a normal but significant condition has occurred)
 - INFORMATIONAL
 - DEBUG
- Description
- Corrective action, if applicable

You can manage the following event capabilities:

- Event destination
Specifies the destination to which events are sent. Destinations can be email addresses, SNMP trap hosts, or syslog servers.
- Event route
Specifies which events generate notifications. An event route is a mapping between events and their destinations. An event route includes information about severity, destinations, and notification thresholds.

- **Event notification and logging**
Specifies the email “from” address, the email “to” address, and whether events are displayed on the console.
- **Mail history records**
A list of emailed event notifications.
- **SNMP trap history records**
A list of event notifications that have been sent to SNMP trap hosts. For information about SNMP traps, see the *Clustered Data ONTAP Network Management Guide*.

Setting up the Event Management System

You can configure EMS to reduce the number of event messages that you receive, and to set up the event destinations and the event routes for a particular event severity.

Steps

1. Display the mail server settings:

```
event config show
```

Example

```
cluster1::> event config show
Mail From: admin@localhost
Mail Server: localhost
```

2. Optional: If necessary, change the mail server settings to meet your requirements:

```
event config modify -mailserver name -mailfrom email address
```

Example

The following example shows how to change the mail server and display the results:

```
cluster1::> event config modify -mailserver mailhost.example.com
-mailfrom admin@node1-example.com
cluster1::> event config show
Mail From: admin@node1-example.com
Mail Server: mailhost.example.com
```

3. Create the destination for events by using the `event destination create` command.

You can send events to email addresses, SNMP trap hosts, and syslog servers.

Example

The following command creates an email destination, sends all important events to the specified email address, and displays the results:

```
cluster1::> event destination create -name test_dest -mail me@example.com
cluster1::> event destination show
```

Name	Mail Dest.	SNMP Dest.	Syslog Dest.	Hide Params
allevents	-	-	-	false
asup	-	-	-	false

```

criticals - - - false
pager - - - false
test_dest me@example.com - - - false
traphost - - - false

```

4. Use the event route `add-destinations` command to define the severity level of messages to receive.

You should set up event routes for critical and above events.

Example

The following example sends all critical, alert, and emergency events to the `test_dest` event destination:

```

cluster1::> event route add-destinations -messagename * -severity <=CRITICAL -
destinations test_dest

```

5. To display all critical and above events, enter the following command:

```
event route show -severity <=CRITICAL
```

Example

The following example shows the events with critical and above severity levels:

```

cluster1::> event route show -severity <=CRITICAL

```

Message	Severity	Destinations	Freq Threshd	Time Threshd
adminapi.time.zoneDiff	ALERT	test_dest	0	3600
api.engine.killed	CRITICAL	test_dest	0	0
app.log.alert	ALERT	test_dest	0	0
app.log.crit	CRITICAL	test_dest	0	0
app.log.emerg	EMERGENCY	test_dest	0	0

6. If you are still getting too many event messages, use the `-timethreshold` parameter to specify how often events are sent to the destination.

Example

For example, the following event is sent to the destinations no more than once per hour:

```

cluster1::> event route modify -messagename adminapi.time.zoneDiff
-timethreshold 3600

```

Result

When you have completed these steps, all events with a severity of critical or above are automatically sent to the destination specified in the event route.

Commands for managing events

You can use specific Data ONTAP commands in the event family for managing events on your cluster.

The following table lists commands for managing events:

If you want to...	Use this command...
Create an event destination	<code>event destination create</code>
Display information about event destinations	<code>event destination show</code>

If you want to...	Use this command...
Modify an event destination	<code>event destination modify</code>
Delete an event destination	<code>event destination delete</code>
Modify an event route or the frequency of event notifications	<code>event route modify</code>
Add an existing destination or destinations to an event route	<code>event route add-destinations</code>
Specify the severity level for an event route	<code>event route add-destinations</code> with the <code>-messagename</code> parameter
Remove a destination or destinations from an event route	<code>event route remove-destinations</code>
Display information about event routes	<code>event route show</code>
Display the event log	<code>event log show</code>
Display the corrective action for an event	<code>event log show</code> with the <code>-instance</code> parameter
Display the configuration for event notification and logging	<code>event config show</code>
Modify the configuration for event notification and logging	<code>event config modify</code>
Display information about event occurrences	<code>event status show</code>
Display mail-history records	<code>event mailhistory show</code>
Delete mail-history records	<code>event mailhistory delete</code>
Display a list of event notifications that have been sent to SNMP traps	<code>event snmphistory show</code>
Delete an SNMP trap-history record	<code>event snmphistory delete</code>

Related information

Clustered Data ONTAP 8.3 Commands: Manual Page Reference

Managing AutoSupport

AutoSupport is a mechanism that proactively monitors the health of your system and automatically sends messages to NetApp technical support, your internal support organization, and a support partner. Although AutoSupport messages to technical support are enabled by default, you must set the correct options and have a valid mail host to have messages sent to your internal support organization.

Only the cluster administrator can perform AutoSupport management. The Storage Virtual Machine (SVM) administrator has no access to AutoSupport.

AutoSupport is enabled by default when you configure your storage system for the first time. AutoSupport begins sending messages to technical support 24 hours after AutoSupport is enabled. You can shorten the 24-hour period by upgrading or reverting the system, modifying the AutoSupport configuration, or changing the system time to be something other than a 24-hour period.

Note: You can disable AutoSupport at any time, but you should leave it enabled. Enabling AutoSupport can significantly help speed problem determination and resolution should a problem

occur on your storage system. By default, the system collects AutoSupport information and stores it locally, even if you disable AutoSupport.

For more information about AutoSupport, see the NetApp Support Site.

Related information

The NetApp Support Site: mysupport.netapp.com

When and where AutoSupport messages are sent

AutoSupport sends messages to different recipients, depending on the type of message. Learning when and where AutoSupport sends messages can help you understand messages that you receive through email or view on the My AutoSupport web site.

Unless specified otherwise, settings in the following tables are parameters of the `system node autosupport modify` command.

Event-triggered messages

When events occur on the system that require corrective action, AutoSupport automatically sends an event-triggered message.

When the message is sent	Where the message is sent
AutoSupport responds to a trigger event in the EMS	Addresses specified in <code>-to</code> and <code>-noteto</code> . (Only critical, service-affecting events are sent.) Addresses specified in <code>-partner-address</code> Technical support, if <code>-support</code> is set to enable

Scheduled messages

AutoSupport automatically sends several messages on a regular schedule.

When the message is sent	Where the message is sent
Daily (by default, sent between 12:00 a.m. and 1:00 a.m. as a log message)	Addresses specified in <code>-partner-address</code> Technical support, if <code>-support</code> is set to enable
Daily (by default, sent between 12:00 a.m. and 1:00 a.m. as a performance message), if the <code>-perf</code> parameter is set to true	Addresses specified in <code>-partner-address</code> Technical support, if <code>-support</code> is set to enable
Weekly (by default, sent Sunday between 12:00 a.m. and 1:00 a.m.)	Addresses specified in <code>-partner-address</code> Technical support, if <code>-support</code> is set to enable

Manually triggered messages

You can manually initiate or resend an AutoSupport message.

When the message is sent	Where the message is sent
You manually initiate a message using the <code>system node autosupport invoke</code> command	If a URI is specified using the <code>-uri</code> parameter in the <code>system node autosupport invoke</code> command, the message is sent to that URI. If <code>-uri</code> is omitted, the message is sent to the addresses specified in <code>-to</code> and <code>-partner-address</code> . The message is also sent to technical support if <code>-support</code> is set to enable .
You manually initiate a message using the <code>system node autosupport invoke-core-upload</code> command	<p>If a URI is specified using the <code>-uri</code> parameter in the <code>system node autosupport invoke-core-upload</code> command, the message is sent to that URI, and the core dump file is uploaded to the URI.</p> <p>If <code>-uri</code> is omitted in the <code>system node autosupport invoke-core-upload</code> command, the message is sent to technical support, and the core dump file is uploaded to the technical support site.</p> <p>Both scenarios require that <code>-support</code> is set to enable and <code>-transport</code> is set to https or http.</p> <p>Due to the large size of core dump files, the message is not sent to the addresses specified in the <code>-to</code> and <code>-partner-addresses</code> parameters.</p>
You manually initiate a message using the <code>system node autosupport invoke-performance-archive</code> command	<p>If a URI is specified using the <code>-uri</code> parameter in the <code>system node autosupport invoke-performance-archive</code> command, the message is sent to that URI, and the performance archive file is uploaded to the URI.</p> <p>If <code>-uri</code> is omitted in the <code>system node autosupport invoke-performance-archive</code>, the message is sent to technical support, and the performance archive file is uploaded to the technical support site.</p> <p>Both scenarios require that <code>-support</code> is set to enable and <code>-transport</code> is set to https or http.</p> <p>Due to the large size of performance archive files, the message is not sent to the addresses specified in the <code>-to</code> and <code>-partner-addresses</code> parameters.</p>
You manually resend a past message using the <code>system node autosupport history retransmit</code> command	Only to the URI that you specify in the <code>-uri</code> parameter of the <code>system node autosupport history retransmit</code> command

Messages triggered by technical support

Technical support can request messages from AutoSupport using the AutoSupport OnDemand feature.

When the message is sent	Where the message is sent
When AutoSupport obtains delivery instructions to generate new AutoSupport messages	Addresses specified in <code>-partner-address</code> Technical support, if <code>-support</code> is set to enable and <code>-transport</code> is set to https
When AutoSupport obtains delivery instructions to resend past AutoSupport messages	Technical support, if <code>-support</code> is set to enable and <code>-transport</code> is set to https
When AutoSupport obtains delivery instructions to generate new AutoSupport messages that upload core dump or performance archive files	Technical support, if <code>-support</code> is set to enable and <code>-transport</code> is set to https . The core dump or performance archive file is uploaded to the technical support site.

Related concepts

How AutoSupport OnDemand obtains delivery instructions from technical support on page 203

How AutoSupport creates and sends event-triggered messages

AutoSupport creates event-triggered AutoSupport messages when the EMS processes a trigger event. An event-triggered AutoSupport message alerts recipients to problems that require corrective action and contains only information that is relevant to the problem. You can customize what content to include and who receives the messages.

AutoSupport uses the following process to create and send event-triggered AutoSupport messages:

1. When the EMS processes a trigger event, EMS sends AutoSupport a request.
A trigger event is an EMS event with an AutoSupport destination and a name that begins with a `callhome.` prefix.
2. AutoSupport creates an event-triggered AutoSupport message.
AutoSupport collects basic and troubleshooting information from subsystems that are associated with the trigger to create a message that includes only information that is relevant to the trigger event.
A default set of subsystems is associated with each trigger. However, you can choose to associate additional subsystems with a trigger by using the `system node autosupport trigger modify` command.
3. AutoSupport sends the event-triggered AutoSupport message to the recipients defined by the `system node autosupport modify` command with the `-to`, `-noteto`, `-partner-address`, and `-support` parameters.
You can enable and disable delivery of AutoSupport messages for specific triggers by using the `system node autosupport trigger modify` command with the `-to` and `-noteto` parameters.

Example of data sent for a specific event

The `storage shelf PSU failed` EMS event triggers a message that contains basic data from the Mandatory, Log Files, Storage, RAID, HA, Platform, and Networking subsystems and troubleshooting data from the Mandatory, Log Files, and Storage subsystems.

You decide that you want to include data about NFS in any AutoSupport messages sent in response to a future `storage shelf PSU failed` event. You enter the following command to enable troubleshooting-level data for NFS for the `callhome.shlf.ps.fault` event:

```
cluster1:~> system node autosupport trigger modify -node node1 -
autosupport-message shlf.ps.fault -troubleshooting-additional nfs
```

Note that the `callhome.` prefix is dropped from the `callhome.shlf.ps.fault` event when you use the `system node autosupport trigger` commands, or when referenced by AutoSupport and EMS events in the CLI.

Types of AutoSupport messages and their content

AutoSupport messages contain status information about supported subsystems. Learning what AutoSupport messages contain can help you interpret or respond to messages that you receive in email or view on the My AutoSupport web site.

Type of message	Type of data the message contains
Event-triggered	Files containing context-sensitive data about the specific subsystem where the event occurred
Daily	Log files
Performance	Performance data sampled during the previous 24 hours
Weekly	Configuration and status data
Triggered by the <code>system node autosupport invoke</code> command	<p>Depends on the value specified in the <code>-type</code> parameter :</p> <ul style="list-style-type: none"> • test sends a user-triggered message with some basic data. This message also triggers an automated email response from technical support to any specified email addresses, using the <code>-to</code> option, so that you can confirm that AutoSupport messages are being received. • performance sends performance data. • all sends a user-triggered message with a complete set of data similar to the weekly message, including troubleshooting data from each subsystem. Technical support typically requests this message.
Triggered by the <code>system node autosupport invoke-core-upload</code> command	Core dump files for a node
Triggered by the <code>system node autosupport invoke-performance-archive</code> command	Performance archive files for a specified period of time

Type of message	Type of data the message contains
Triggered by AutoSupport OnDemand	<p>AutoSupport OnDemand can request new messages or past messages:</p> <ul style="list-style-type: none"> New messages, depending on the type of AutoSupport collection, can be test, all, or performance. Past messages depend on the type of message that is resent. <p>Autosupport OnDemand can request new messages that upload the following files to the NetApp support site:</p> <ul style="list-style-type: none"> Core dump Performance archive

What AutoSupport subsystems are

Each subsystem provides basic and troubleshooting information that AutoSupport uses for its messages. Each subsystem is also associated with trigger events that allow AutoSupport to collect from subsystems only information that is relevant to the trigger event.

AutoSupport collects context-sensitive content. You can view information about subsystems and trigger events by using the `system node autosupport trigger show` command.

AutoSupport size and time budgets

AutoSupport collects information, organized by subsystem, and enforces a size and time budget on content for each subsystem. As storage systems grow, AutoSupport budgets provide control over the AutoSupport payload, which in turn provides scalable delivery of AutoSupport data.

AutoSupport stops collecting information and truncates the AutoSupport content if the subsystem content exceeds its size or time budget. If the content cannot be truncated easily (for example, binary files), AutoSupport omits the content.

You should modify the default size and time budgets only with guidance from technical support. The CLI for AutoSupport size and time budgets is a diagnostic privilege command set.

Subsystem	Size budget (bytes)	Time budget (seconds)
asup_ems	2097152	60
av	2097152	60
cifs	2097152	60
cluster_peer	2097152	60
cps	2097152	60
crs	2097152	60
dedupe	10485760	120
fpolicy	2097152	60
ha	2097152	60
hm	2097152	60
kcs	2097152	60
kernel	2097152	60
log_files	5242880	120

Subsystem	Size budget (bytes)	Time budget (seconds)
mandatory	unlimited	unlimited
memevt	2097152	60
metrocluster	8388608	300
mhost	3670016	120
mot	2097152	60
ndmp	2097152	60
ndu	2097152	60
networking	4194304	60
nfs	2097152	60
nht	2097152	60
nwd	2097152	60
performance	20971520	300
performance_asup	52428800	300
platform	2097152	60
raid	2097152	60
repository	2097152	60
san	2097152	120
san_fcp	2097152	60
san_ffdc	367001600	120
secd	2097152	60
snapmirror	2097152	60
splog_before_spreboot	512000	30
splog_downcontroller	512000	30
splog_latest	512000	30
splog_rnode	2097152	60
splog_rnode_down	2097152	60
splog_rnode_latest	2097152	60
storage	10485760	180
vscan	2097152	60
vserver_dr	2097152	60
waf1	23068672	300

Files sent in event-triggered AutoSupport messages

Event-triggered AutoSupport messages only contain basic and troubleshooting information from subsystems that are associated with the event that caused AutoSupport to generate the message. The specific data helps NetApp support and support partners troubleshoot the problem.

AutoSupport uses the following criteria to control content in event-triggered AutoSupport messages:

- Which subsystems are included
Data is grouped into subsystems, including common subsystems, such as Log Files, and specific subsystems, such as RAID. Each event triggers a message that contains only the data from specific subsystems.
- The detail level of each included subsystem
Data for each included subsystem is provided at a basic or troubleshooting level.

You can view all possible events and determine which subsystems are included in messages about each event using the `system node autosupport trigger show` command with the `-instance` parameter.

In addition to the subsystems that are included by default for each event, you can add additional subsystems at either a basic or a troubleshooting level using the `system node autosupport trigger modify` command.

Log files sent in AutoSupport messages

AutoSupport messages can contain several key log files that enable technical support staff to review recent system activity.

All types of AutoSupport messages might include the following log files when the Log Files subsystem is enabled:

Log file	Amount of data included from the file
<ul style="list-style-type: none"> • Log files from the <code>/mroot/etc/log/mlog/</code> directory • The MESSAGES log file 	<p>Only new lines added to the logs since the last AutoSupport message up to a specified maximum. This ensures that AutoSupport messages have unique, relevant—not overlapping—data. (Log files from partners are the exception; for partners, the maximum allowed data is included.)</p>
<ul style="list-style-type: none"> • Log files from the <code>/mroot/etc/log/shelflog/</code> directory • Log files from the <code>/mroot/etc/log/acp/</code> directory • Event Management System (EMS) log data 	<p>The most recent lines of data up to a specified maximum.</p>

The content of AutoSupport messages can change between releases of Data ONTAP.

Files sent in weekly AutoSupport messages

Weekly AutoSupport messages contain additional configuration and status data that is useful to track changes in your system over time.

The following information is sent in weekly AutoSupport messages:

- Basic information about every subsystem
- Contents of selected `/mroot/etc` directory files
- Log files
- Output of commands that provide system information
- Additional information, including replicated database (RDB) information, service statistics, and more

How AutoSupport OnDemand obtains delivery instructions from technical support

AutoSupport OnDemand periodically communicates with technical support to obtain delivery instructions for sending, resending, and declining AutoSupport messages as well as uploading large files to the NetApp support site. AutoSupport OnDemand enables AutoSupport messages to be sent on-demand instead of waiting for the weekly AutoSupport job to run.

AutoSupport OnDemand consists of the following components:

- AutoSupport OnDemand client that runs on each node
- AutoSupport OnDemand service that resides in technical support

The AutoSupport OnDemand client periodically polls the AutoSupport OnDemand service to obtain delivery instructions from technical support. For example, technical support can use the AutoSupport OnDemand service to request that a new AutoSupport message be generated. When the AutoSupport OnDemand client polls the AutoSupport OnDemand service, the client obtains the delivery instructions and sends the new AutoSupport message on-demand as requested.

AutoSupport OnDemand is enabled by default. However AutoSupport OnDemand relies on some AutoSupport settings to continue communicating with technical support. AutoSupport OnDemand automatically communicates with technical support when the following requirements are met:


- AutoSupport is enabled.
- AutoSupport is configured to send messages to technical support.
- AutoSupport is configured to use the HTTPS transport protocol.

The AutoSupport OnDemand client sends HTTPS requests to the same technical support location to which AutoSupport messages are sent. The AutoSupport OnDemand client does not accept incoming connections.

Note: AutoSupport OnDemand uses the “autosupport” user account to communicate with technical support. Data ONTAP prevents you from deleting this account.

If you want to have AutoSupport OnDemand disabled, but keep AutoSupport enabled and configured to send messages to technical support by using the HTTPS transport protocol, contact technical support.

The following illustration shows how AutoSupport OnDemand sends HTTPS requests to technical support to obtain delivery instructions.



The delivery instructions can include requests for AutoSupport to do the following:

- Generate new AutoSupport messages.
Technical support might request new AutoSupport messages to help triage issues.

- Generate new AutoSupport messages that upload core dump files or performance archive files to the NetApp support site.
Technical support might request core dump or performance archive files to help triage issues.
- Retransmit previously generated AutoSupport messages.
This request automatically happens if a message was not received due to a delivery failure.
- Disable delivery of AutoSupport messages for specific trigger events.
Technical support might disable delivery of data that is not used.

Structure of AutoSupport messages sent by email

When an AutoSupport message is sent by email, the message has a standard subject, a brief body, and a large attachment in 7z file format that contains the data.

Note: If AutoSupport is configured to hide private data, certain information, such as the hostname, is omitted or masked in the header, subject, body, and attachments.

Subject

The subject line of messages sent by the AutoSupport mechanism contains a text string that identifies the reason for the notification. The format of the subject line is as follows:

HA Group Notification from *System_Name (Message) Severity*

- *System_Name* is either the hostname or the system ID, depending on the AutoSupport configuration

Body

The body of the AutoSupport message contains the following information:

- Date and timestamp of the message
- Version of Data ONTAP on the node that generated the message
- System ID, serial number, and hostname of the node that generated the message
- AutoSupport sequence number
- SNMP contact name and location, if specified
- System ID and hostname of the HA partner node

Attached files

The key information in an AutoSupport message is contained in files that are compressed into a 7z file called `body.7z` and attached to the message.

The files contained in the attachment are specific to the type of AutoSupport message.

AutoSupport severity types

AutoSupport messages have severity types that help you understand the purpose of each message—for example, to draw immediate attention to a critical problem, or only to provide information.

Messages have one of the following severities:

- Critical: critical conditions
- Error: error conditions
- Warning: warning conditions

- Notice: normal but significant condition
- Info: informational message
- Debug: debug-level messages

If your internal support organization receives AutoSupport messages via email, the severity appears in the subject line of the email message.

Requirements for using AutoSupport

You should use HTTPS for delivery of AutoSupport messages to provide the best security and to support all of the latest AutoSupport features. Although AutoSupport supports HTTP and SMTP for delivery of AutoSupport messages, HTTPS is recommended.

Supported protocols

All of these protocols run on IPv4 or IPv6, based on the address family to which the name resolves.

Protocol and port	Description
HTTPS on port 443	<p>This is the default protocol. You should use this whenever possible.</p> <p>This protocol supports AutoSupport OnDemand and uploads of large files. The certificate from the remote server is validated against the root certificate, unless you disable validation.</p> <p>The delivery uses an HTTP PUT request. With PUT, if the request fails during transmission, the request restarts where it stopped. If the server receiving the request does not support PUT, the delivery uses an HTTP POST request.</p>
HTTP on port 80	<p>This protocol is preferred over SMTP.</p> <p>This protocol supports uploads of large files, but not AutoSupport OnDemand.</p> <p>The delivery uses an HTTP PUT request. With PUT, if the request fails during transmission, the request restarts where it stopped. If the server receiving the request does not support PUT, the delivery uses an HTTP POST request.</p>
SMTP on port 25 or another port	<p>You should use this protocol only if the network connection does not allow HTTPS or HTTP.</p> <p>The default port value is 25, but you can configure AutoSupport to use a different port.</p> <p>Keep the following limitations in mind when using SMTP:</p> <ul style="list-style-type: none"> • AutoSupport OnDemand and uploads of large files are not supported. • Data is not encrypted. SMTP sends data in clear text, making text in the AutoSupport message easy to intercept and read. • Limitations on message length and line length can be introduced.

If you configure AutoSupport with specific email addresses for your internal support organization, or a support partner organization, those messages are always sent by SMTP.

For example, if you use the recommended protocol to send messages to technical support and you also want to send messages to your internal support organization, your messages will be transported using both HTTPS and SMTP, respectively.

AutoSupport limits the maximum file size for each protocol. The default setting for HTTP and HTTPS transfers is 10 MB. The default setting for SMTP transfers is 5 MB. If the size of the AutoSupport message exceeds the configured limit, AutoSupport delivers as much of the message as possible. You can edit the maximum size by modifying AutoSupport configuration. See the `system node autosupport modify` man page for more information.

Note: AutoSupport automatically overrides the maximum file size limit for the HTTPS and HTTP protocols when you generate and send AutoSupport messages that upload core dump or performance archive files to the NetApp support site or a specified URI. The automatic override applies only when you upload files by using the `system node autosupport invoke-core-upload` or the `system node autosupport invoke-performance-archive` commands.

Configuration requirements

Depending on your network configuration, use of HTTP or HTTPS protocols may require additional configuration of a proxy URL. If you use HTTP or HTTPS to send AutoSupport messages to technical support and you have a proxy, you must identify the URL for that proxy. If the proxy uses a port other than the default port, which is 3128, you can specify the port for that proxy. You can also specify a user name and password for proxy authentication.

If you use SMTP to send AutoSupport messages either to your internal support organization or to technical support, you must configure an external mail server. The storage system does not function as a mail server; it requires an external mail server at your site to send mail. The mail server must be a host that listens on the SMTP port (25) or another port, and it must be configured to send and receive 8-bit Multipurpose Internet Mail Extensions (MIME) encoding. Example mail hosts include a UNIX host running an SMTP server such as the sendmail program and a Windows server running the Microsoft Exchange server. You can have one or more mail hosts.

Setting up AutoSupport

You can control whether and how AutoSupport information is sent to technical support and your internal support organization, and then test that the configuration is correct.

About this task

You must perform this procedure on each node in your cluster where you want to configure AutoSupport.

By default, AutoSupport is enabled on each node to send messages to technical support by using the HTTPS transport protocol.

Steps

1. Ensure that AutoSupport is enabled by setting the `-state` parameter of the `system node autosupport modify` command to **enable**.
2. If you want technical support to receive AutoSupport messages, set the `-support` parameter of the `system node autosupport modify` command to **enable**.

You must enable this option if you want to enable AutoSupport to work with AutoSupport OnDemand or if you want to upload large files, such as core dump and performance archive files, to technical support or a specified URI.

3. If technical support is enabled to receive AutoSupport messages, specify what transport protocol to use for the messages by choosing one of the following options:

If you want to...	Then set the following parameters of the <code>system node autosupport modify</code> command...
Use the default HTTPS protocol	<p>a. Set <code>-transport</code> to https.</p> <p>b. If you use a proxy, set <code>-proxy-url</code> to the URL of your proxy.</p> <p>This configuration supports communication with AutoSupport OnDemand and uploads of large files.</p>
Use HTTP that is preferred over SMTP	<p>a. Set <code>-transport</code> to http.</p> <p>b. If you use a proxy, set <code>-proxy-url</code> to the URL of your proxy.</p> <p>This configuration supports uploads of large files, but not AutoSupport OnDemand.</p>
Use SMTP	<p>Set <code>-transport</code> to smtp.</p> <p>This configuration does not support AutoSupport OnDemand or uploads of large files.</p>

4. If you want your internal support organization or a support partner to receive AutoSupport messages, perform the following actions:
- Identify the recipients in your organization by setting the following parameters of the `system node autosupport modify` command:

Set this parameter..	To this...
<code>-to</code>	Up to five comma-separated individual email addresses or distribution lists in your internal support organization that will receive key AutoSupport messages
<code>-noteto</code>	Up to five comma-separated individual email addresses or distribution lists in your internal support organization that will receive a shortened version of key AutoSupport messages designed for cell phones and other mobile devices
<code>-partner-address</code>	Up to five comma-separated individual email addresses or distribution lists in your support partner organization that will receive all AutoSupport messages

- Check that addresses are correctly configured by listing the destinations using the `system node autosupport destinations show` command.
5. If you are sending messages to your internal support organization or you chose SMTP transport for messages to technical support, configure SMTP by setting the following parameters of the `system node autosupport modify` command:
- Set `-mail-hosts` to one or more mail hosts, separated by commas. You can set a maximum of five. You can configure a port value for each mail host by specifying a colon and port number after the mail host name: for example, `mymailhost.example.com:5678`, where 5678 is the port for the mail host.
 - Set `-from` to the email address that sends the AutoSupport message.

6. Configure DNS.
7. Optional: Change the following settings:

If you want to do this...	Then set the following parameters of the <code>system node autosupport modify</code> command...
Hide private data by removing, masking, or encoding sensitive data in the messages	Set <code>-remove-private-data</code> to true . If you change from false to true , all AutoSupport history and all associated files are deleted.
Stop sending performance data in periodic AutoSupport messages	Set <code>-perf</code> to false .

8. Check the overall configuration using the `system node autosupport show` command with the `-node` parameter.
9. Test that AutoSupport messages are being sent and received:
 - a. Use the `system node autosupport invoke` command with the `-type` parameter set to **test**.

Example

```
cluster1::> system node autosupport invoke -type test -node node1
```

- b. Confirm that NetApp is receiving your AutoSupport messages:


```
system node autosupport history -node local
```

The status of the latest outgoing AutoSupport message should eventually change to `sent-successful` for all appropriate protocol destinations.
- c. Optional: Confirm that the AutoSupport message is being sent to your internal support organization or to your support partner by checking the email of any address that you configured for the `-to`, `-noteto`, or `-partner-address` parameters of the `system node autosupport modify` command.

Related tasks

[Troubleshooting AutoSupport when messages are not received](#) on page 213

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

[Clustered Data ONTAP 8.3 Network Management Guide](#)

Uploading core dump files

You can generate and send an AutoSupport message that contains a core dump file for a node. By default, NetApp technical support receives the message, and the core dump file is uploaded to the NetApp support site. You can specify an alternate destination for the message and upload.

Before you begin

- You must have set up AutoSupport with the following settings:
 - AutoSupport is enabled on the node.
 - AutoSupport is configured to send messages to technical support.

- AutoSupport is configured to use the HTTP or HTTPS transport protocol. The SMTP transport protocol is not supported when sending messages that include large files, such as core dump files.

Steps

1. View the core dump files for a node by using the `system node coredump show` command.

Example

In the following example, core dump files are displayed for the local node:

```
cluster1::> system node coredump show -node local
Node:Type Core Name Saved Panic Time
-----
node:kernel
core.4073000068.2013-09-11.15_05_01.nz true 9/11/2013 15:05:01
```

2. Generate an AutoSupport message and upload a core dump file by using the `system node autosupport invoke-core-upload` command.

Example

In the following example, an AutoSupport message is generated and sent to the default location, which is technical support, and the core dump file is uploaded to the default location, which is the NetApp support site:

```
cluster1::> system node autosupport invoke-core-upload -core-filename
core.4073000068.2013-09-11.15_05_01.nz -node local
```

In the following example, an AutoSupport message is generated and sent to the location specified in the URI, and the core dump file is uploaded to the URI:

```
cluster1::> system node autosupport invoke-core-upload -uri https//
files.company.com -core-filename
core.4073000068.2013-09-11.15_05_01.nz -node local
```

Related tasks

[Setting up AutoSupport](#) on page 206

Uploading performance archive files

You can generate and send an AutoSupport message that contains a performance archive. By default, NetApp technical support receives the AutoSupport message, and the performance archive is uploaded to the NetApp support site. You can specify an alternate destination for the message and upload.

Before you begin

- You must have set up AutoSupport with the following settings:
 - AutoSupport is enabled on the node.
 - AutoSupport is configured to send messages to technical support.
 - AutoSupport is configured to use the HTTP or HTTPS transport protocol. The SMTP transport protocol is not supported when sending messages that include large files, such as performance archive files.

About this task

You must specify a start date for the performance archive data that you want to upload. Most storage systems retain performance archives for two weeks, enabling you to specify a start date up to two weeks ago. For example, if today is January 15, you can specify a start date of January 2. If you need help identifying the start date for the performance archive that you want to upload, contact technical support.

Step

1. Generate an AutoSupport message and upload the performance archive file by using the `system node autosupport invoke-performance-archive` command.

Example

In the following example, 4 hours of performance archive files from January 12, 2015 are added to an AutoSupport message and uploaded to the default location, which is the NetApp support site:

```
cluster1::> system node autosupport invoke-performance-archive -node
local -start-date 1/12/2015 13:42:09 -duration 4h
```

In the following example, 4 hours of performance archive files from January 12, 2015 are added to an AutoSupport message and uploaded to the location specified by the URI:

```
cluster1::> system node autosupport invoke-core-upload -uri https://
files.company.com -core-filename
core.4073000068.2013-09-11.15_05_01.nz -node local
```

Related tasks

[Setting up AutoSupport](#) on page 206

Getting AutoSupport message descriptions

The descriptions of the AutoSupport messages that you receive are available through the online AutoSupport Message Matrices page.

Steps

1. Go to the AutoSupport Message Matrices page: support.netapp.com/NOW/knowledge/docs/olio/autosupport/matrices/
2. On the AutoSupport Message Matrices page under Select a Release, select your version of Data ONTAP and click **View Matrix**.

The Syslog Translator page appears with all AutoSupport message descriptions listed alphabetically by subject line.

Commands for managing AutoSupport

You use the `system node autosupport` commands to change or view AutoSupport configuration, display information about previous AutoSupport messages, and send, resend or cancel an AutoSupport message.

Configure AutoSupport

If you want to...	Use this command...
Control whether any AutoSupport messages are sent	<code>system node autosupport modify</code> with the <code>-state</code> parameter
Control whether AutoSupport messages are sent to technical support	<code>system node autosupport modify</code> with the <code>-support</code> parameter
Set up AutoSupport or modify the configuration of AutoSupport	<code>system node autosupport modify</code>
Enable and disable AutoSupport messages to your internal support organization for individual trigger events, and specify additional subsystem reports to include in messages sent in response to individual trigger events	<code>system node autosupport trigger modify</code>

Display information about the AutoSupport configuration

If you want to...	Use this command...
Display the AutoSupport configuration	<code>system node autosupport show</code> with the <code>-node</code> parameter
View a summary of all addresses and URLs that receive AutoSupport messages	<code>system node autosupport destinations show</code>
Display which AutoSupport messages are sent to your internal support organization for individual trigger events	<code>system node autosupport trigger show</code>

Display information about past AutoSupport messages

If you want to...	Use this command...
Display information about one or more of the 50 most recent AutoSupport messages	<code>system node autosupport history show</code>
Display information about recent AutoSupport messages generated to upload core dump or performance archive files to the technical support site or a specified URI	<code>system node autosupport history show-upload-details</code>
View the information in the AutoSupport messages including the name and size of each file collected for the message along with any errors	<code>system node autosupport manifest show</code>

Send, resend, or cancel AutoSupport messages

If you want to...	Use this command...
Retransmit a locally stored AutoSupport message, identified by its AutoSupport sequence number Note: If you retransmit an AutoSupport message, and if support already received that message, the support system will not create a duplicate case. If, on the other hand, support did not receive that message, then the AutoSupport system will analyze the message and create a case, if necessary.	<pre>system node autosupport history retransmit</pre>
Generate and send an AutoSupport message—for example, for testing purposes	<pre>system node autosupport invoke</pre> <p>Note: Use the <code>-force</code> parameter to send a message even if AutoSupport is disabled. Use the <code>-uri</code> parameter to send the message to the destination you specify instead of the configured destination.</p>
Cancel an AutoSupport message	<pre>system node autosupport history cancel</pre>

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Information included in the AutoSupport manifest

The AutoSupport manifest provides you with a detailed view of the files collected for each AutoSupport message. The AutoSupport manifest also includes information about collection errors when AutoSupport cannot collect the files it needs.

The AutoSupport manifest includes the following information:

- Sequence number of the AutoSupport message
- Which files AutoSupport included in the AutoSupport message
- Size of each file, in bytes
- Status of the AutoSupport manifest collection
- Error description, if AutoSupport failed to collect one or more files

You can view the AutoSupport manifest by using the `system node autosupport manifest show` command.

The AutoSupport manifest is included with every AutoSupport message and presented in XML format, which means that you can either use a generic XML viewer to read it or view it using the My AutoSupport portal.

AutoSupport case suppression during scheduled maintenance windows

AutoSupport case suppression enables you to stop unnecessary cases from being created by AutoSupport messages that are triggered during scheduled maintenance windows.

To suppress AutoSupport cases, you must manually invoke an AutoSupport message with a specially formatted text string: `MAINT=xh`. `x` is the duration of the maintenance window in units of hours.

Related information

[NetApp KB Article: 1010449: How to suppress automatic case creation during scheduled maintenance windows](#)

What My AutoSupport is

My AutoSupport is a web-based application, working in conjunction with AutoSupport, that presents information enabling you to easily analyze data to model and optimize your storage infrastructure.

My AutoSupport is a web-based application hosted on the NetApp Support Site at mysupport.netapp.com that you can access using a browser. Your system must have AutoSupport enabled and configured so that it sends data back to . You can access My AutoSupport by going to <http://support.netapp.com/NOW/asuphome/>.

My AutoSupport provides a dashboard from which you can perform the following actions:

- Generate reports and export them to PDF or CSV files
- View information about configurations, performance, system health, installed software, and storage efficiency
- Access system and AutoSupport tools

Troubleshooting AutoSupport

If you do not receive AutoSupport messages, you can check a number of settings to resolve the problem.

Troubleshooting AutoSupport when messages are not received

If the system does not send the AutoSupport message, you can determine whether that is because AutoSupport cannot generate the message or cannot deliver the message.

Steps

1. Check delivery status of the messages by using the `system node autosupport history show` command.
2. Read the status.

This status	Means
initializing	The collection process is starting. If this state is temporary, all is well. However, if this state persists, there is an issue.
collection-failed	AutoSupport cannot create the AutoSupport content in the spool directory. You can view what AutoSupport is trying to collect by entering the <code>system node autosupport history show -detail</code> command.
collection-in-progress	AutoSupport is collecting AutoSupport content. You can view what AutoSupport is collecting by entering the <code>system node autosupport manifest show</code> command.
queued	AutoSupport messages are queued for delivery, but not yet delivered.
transmitting	AutoSupport is currently delivering messages.
sent-successful	AutoSupport successfully delivered the message. You can find out where AutoSupport delivered the message by entering the <code>system node autosupport history show -delivery</code> command.

This status	Means
ignore	AutoSupport has no destinations for the message. You can view the delivery details by entering the <code>system node autosupport history show -delivery</code> command.
re-queued	AutoSupport tried to deliver messages, but the attempt failed. As a result, AutoSupport placed the messages back in the delivery queue for another attempt. You can view the error by entering the <code>system node autosupport history show</code> command.
transmission-failed	AutoSupport failed to deliver the message the specified number of times and stopped trying to deliver the message. You can view the error by entering the <code>system node autosupport history show</code> command.
ondemand-ignore	The AutoSupport message was processed successfully, but the AutoSupport OnDemand service chose to ignore it.

3. Perform one of the following actions:

For this status	Do this
initializing or collection-failed	Contact technical support because AutoSupport cannot generate the message.
ignore, re-queued, or transmission failed	Check that destinations are correctly configured for SMTP, HTTP, or HTTPS because AutoSupport cannot deliver the message.

Related tasks

[Troubleshooting AutoSupport message delivery over SMTP](#) on page 215

[Troubleshooting AutoSupport message delivery over HTTP or HTTPS](#) on page 214

Troubleshooting AutoSupport message delivery over HTTP or HTTPS

If the system does not send the expected AutoSupport message and you are using HTTP or HTTPS, you can check a number of settings to resolve the problem.

Before you begin

You should have confirmed basic network connectivity and DNS lookup:

- Your network management LIF must be up for operational and administrative status.
- You must be able to ping a functioning host on the same subnet from the cluster management LIF (not a LIF on any of the nodes).
- You must be able to ping a functioning host outside the subnet from the cluster management LIF.
- You must be able to ping a functioning host outside the subnet from the cluster management LIF using the name of the host (not the IP address).

About this task

These steps are for cases when you have determined that AutoSupport can generate the message, but cannot deliver the message over HTTP or HTTPS.

If you encounter errors or cannot complete a step in this procedure, determine and address the problem before proceeding to the next step.

Steps

1. Verify the status of the node management LIF:

```
network interface show -home-node local -role node-mgmt -fields
vserver,lif,status-oper,status-admin,address,role
```

The `status-oper` and `status-admin` fields should return `up`.

2. Record the SVM name, the LIF name, and the LIF IP address for later use.
3. Ensure that DNS is enabled and configured correctly:

```
vserver services name-service dns show
```

4. Address any errors returned by the AutoSupport message:

```
system node autosupport history show -node * -fields node,seq-
num,destination,last-update,status,error
```

If the error refers to a problem with the digital certificate, contact technical support.

5. Confirm that the cluster can access both the servers it needs and the Internet successfully:

- a. `network traceroute -node local -destination network_management_LIF`

- b. `network traceroute -node local -destination support.netapp.com`

- c. `system node autosupport show -fields proxy-url`

- d. `network traceroute -node local -destination proxy_url`

If any of these routes is not functioning, try the same route from a functioning host on the same subnet as the cluster, using the “traceroute” or “tracert” utility found on most third-party network clients. This assists you in determining whether the issue is in your network configuration or your cluster configuration.

6. If you are using HTTPS for your AutoSupport transport protocol, ensure that HTTPS traffic can exit your network:

- a. Configure a web client on the same subnet as the cluster management LIF.

Ensure that all configuration parameters are the same values as for the AutoSupport configuration, including using the same proxy server, user name, password, and port.

- b. Access `https://support.netapp.com` with the web client.

The access should be successful. If not, ensure that all firewalls are configured correctly to allow HTTPS and DNS traffic, and that the proxy server is configured correctly.

Related tasks

[Troubleshooting AutoSupport when messages are not received](#) on page 213

Related information

[Clustered Data ONTAP 8.3 Network Management Guide](#)

Troubleshooting AutoSupport message delivery over SMTP

If the system cannot deliver AutoSupport messages over SMTP, you can check a number of settings to resolve the problem.

Before you begin

You should have confirmed basic network connectivity and DNS lookup:

- Your network management LIF must be up for operational and administrative status.

- You must be able to ping a functioning host on the same subnet from the cluster management LIF (not a LIF on any of the nodes).
- You must be able to ping a functioning host outside the subnet from the cluster management LIF.
- You must be able to ping a functioning host outside the subnet from the cluster management LIF using the name of the host (not the IP address).

About this task

These steps are for cases when you have determined that AutoSupport can generate the message, but cannot deliver the message over SMTP.

If you encounter errors or cannot complete a step in this procedure, determine and address the problem before proceeding to the next step.

All commands are entered at the Data ONTAP command-line interface, unless otherwise specified.

Steps

1. Verify the status of the node management LIF:

```
network interface show -home-node local -role node-mgmt -fields
vserver,lif,status-oper,status-admin,address,role
```

The `status-oper` and `status-admin` fields should return `up`.

2. Record the SVM name, the LIF name, and the LIF IP address for later use.
3. Ensure that DNS is enabled and configured correctly:

```
vserver services name-service dns show
```

4. Display all of the servers configured to be used by AutoSupport:

```
system node autosupport show -fields mail-hosts
```

Record all server names displayed.

5. For each server displayed by the previous step, and `support.netapp.com`, ensure that the server or URL can be reached by the node:

```
network traceroute -node local -destination server_name
```

If any of these routes is not functioning, try the same route from a functioning host on the same subnet as the cluster, using the “traceroute” or “tracert” utility found on most third-party network clients. This assists you in determining whether the issue is in your network configuration or your cluster configuration.

6. Log in to the host designated as the mail host, and ensure that it can serve SMTP requests:

```
netstat -aAn|grep 25
```

25 is the listener SMTP port number.

A message similar to the following text is displayed:

```
ff64878c tcp          0          0  *.25      *.*      LISTEN.
```

7. From some other host, open a Telnet session with the SMTP port of the mail host:

```
telnet mailhost 25
```

A message similar to the following text is displayed:


```
220 filer.yourco.com Sendmail 4.1/SMI-4.1 ready at Thu, 30 Nov 2014
10:49:04 PST
```

8. At the telnet prompt, ensure that a message can be relayed from your mail host:

```
HELO domain_name
MAIL FROM: your_email_address
RCPT TO: autosupport@netapp.com
```

domain_name is the domain name of your network.

If an error is returned saying that relaying is denied, relaying is not enabled on the mail host. Contact your system administrator.

9. At the telnet prompt, send a test message:

```
DATA
SUBJECT: TESTING
THIS IS A TEST
```

.

Note: Ensure that you enter the last period (.) on a line by itself. The period indicates to the mail host that the message is complete.

If an error is returned, your mail host is not configured correctly. Contact your system administrator.

10. From the Data ONTAP command-line interface, send an AutoSupport test message to a trusted email address that you have access to:

```
system node autosupport invoke -node local -type test
```

11. Find the sequence number of the attempt:

```
system node autosupport history show -node local -destination smtp
```

Find the sequence number for your attempt based on the timestamp. It is probably the most recent attempt.

12. Display the error for your test message attempt:

```
system node autosupport history show -node local -seq-num seq_num -
fields error
```

If the error displayed is `Login denied`, your SMTP server is not accepting send requests from the cluster management LIF. If you do not want to change to using HTTPS as your transport protocol, contact your site network administrator to configure the SMTP gateways to address this issue.

If this test succeeds but the same message sent to `mailto:autosupport@netapp.com` does not, ensure that SMTP relay is enabled on all of your SMTP mail hosts, or use HTTPS as a transport protocol.

If even the message to the locally administered email account does not succeed, confirm that your SMTP servers are configured to forward attachments with both of these characteristics:

- The “7z” suffix
- The “application/x-7x-compressed” MIME type.

Related tasks

[Troubleshooting AutoSupport when messages are not received](#) on page 213

Related information

[Clustered Data ONTAP 8.3 Network Management Guide](#)

Monitoring the health of your system

Health monitors proactively monitor certain critical conditions in your cluster and raise alerts if they detect a fault or risk. If there are active alerts, the system health status reports a degraded status for the cluster. The alerts include the information that you need to respond to degraded system health.

If the status is degraded, you can view details about the problem, including the probable cause and recommended recovery actions. After you resolve the problem, the system health status automatically returns to OK.

The system health status reflects multiple separate health monitors. A degraded status in an individual health monitor causes a degraded status for the overall system health.

Data ONTAP supports the following cluster switches for system health monitoring in your cluster:

- NetApp CN1601
- NetApp CN1610
- Cisco Nexus 5010
- Cisco Nexus 5020
- Cisco Nexus 5596
- Cisco Catalyst 2960-24TT-L

How health monitoring works

Individual health monitors have a set of policies that trigger alerts when certain conditions occur. Understanding how health monitoring works can help you respond to problems and control future alerts.

Health monitoring consists of the following components:

- Individual health monitors for specific subsystems, each of which has its own health status
For example, the Storage subsystem has a node connectivity health monitor.
- An overall system health monitor that consolidates the health status of the individual health monitors
A degraded status in any single subsystem results in a degraded status for the entire system. If no subsystems have alerts, the overall system status is OK.

Each health monitor is made up of the following key elements:

- Alerts that the health monitor can potentially raise
Each alert has a definition, which includes details such as the severity of the alert and its probable cause.
- Health policies that identify when each alert is triggered
Each health policy has a rule expression, which is the exact condition or change that triggers the alert.

A health monitor continuously monitors and validates the resources in its subsystem for condition or state changes. When a condition or state change matches a rule expression in a health policy, the health monitor raises an alert. An alert causes the subsystem's health status and the overall system health status to become degraded.

Ways to respond to system health alerts

When a system health alert occurs, you can acknowledge it, learn more about 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 did not automatically clear it.
- Suppress an alert to prevent it from affecting the health status of a subsystem.
Suppressing is useful when you understand a problem. After you suppress an alert, it can still occur, but the subsystem health displays as “ok-with-suppressed.” when the suppressed alert occurs.

System health alert customization

You can control which alerts a health monitor generates by enabling and disabling the system health policies that define when alerts are triggered. This enables you to customize the health monitoring system for your particular environment.

You can learn the name of a policy either by displaying detailed information about a generated alert or by displaying policy definitions for a specific health monitor, node, or alert ID.

Disabling health policies is different from suppressing alerts. When you suppress an alert, it does not affect the subsystem's health status, but the alert can still occur.

If you disable a policy, the condition or state that is defined in its policy rule expression no longer triggers an alert.

Example of an alert that you want to disable

For example, suppose an alert occurs that is not useful to you. You use the `system health alert show -instance` command to obtain the Policy ID for the alert. You use the policy ID in the `system health policy definition show` command to view information about the policy. After reviewing the rule expression and other information about the policy, you decide to disable the policy. You use the `system health policy definition modify` command to disable the policy.

How health alerts trigger AutoSupport messages and events

System health alerts trigger AutoSupport messages and events in the Event Management System (EMS), enabling you to monitor the health of the system using AutoSupport messages and the EMS in addition to using the health monitoring system directly.

Your system sends an AutoSupport message within five minutes of an alert. The AutoSupport message includes all alerts generated since the previous AutoSupport message, except for alerts that duplicate an alert for the same resource and probable cause within the previous week.

Some alerts do not trigger AutoSupport messages. An alert does not trigger an AutoSupport message if its health policy disables the sending of AutoSupport messages. For example, a health policy might disable AutoSupport messages by default because AutoSupport already generates a message when the problem occurs. You can configure policies to not trigger AutoSupport messages by using the `system health policy definition modify` command.

You can view a list of all of the alert-triggered AutoSupport messages sent in the previous week using the `system health autosupport trigger history show` command.

Alerts also trigger the generation of events to the EMS. An event is generated each time an alert is created and each time an alert is cleared.

Available cluster health monitors

There are several health monitors that monitor different parts of a cluster. Health monitors help you to recover from errors within Data ONTAP subsystems by detecting events, sending alerts to you, and deleting events as they clear.

Health monitor name (identifier)	Subsystem name (identifier)	Purpose
Cluster switch (cluster-switch)	Switch (Switch-Health)	Monitors cluster network switches and management network switches for temperature, utilization, interface configuration, redundancy (cluster network switches only), and fan and power supply operation. The cluster switch health monitor communicates with switches through SNMP. SNMPv2c is the default setting.
MetroCluster Fabric	Switch	Monitors the MetroCluster configuration back-end fabric topology and detects misconfigurations such as incorrect cabling and zoning, and ISL failures.
MetroCluster Health	Interconnect, RAID, and storage	Monitors FC-VI adapters, FC initiator adapters, left-behind aggregates and disks, and inter-cluster ports
Node connectivity (node-connect)	CIFS nondisruptive operations (CIFS-NDO)	Monitors SMB connections for nondisruptive operations to Hyper-V applications.
	Storage (SAS-connect)	Monitors shelves, disks, and adapters at the node level for appropriate paths and connections.
System	not applicable	Aggregates information from other health monitors.

Health monitor name (identifier)	Subsystem name (identifier)	Purpose
System connectivity (system-connect)	Storage (SAS-connect)	Monitors shelves at the cluster level for appropriate paths to two HA clustered nodes.

Receiving system health alerts automatically

You can manually view system health alerts by using the `system health alert show` command. However, you should subscribe to specific Event Management System (EMS) messages to automatically receive notifications when a health monitor generates an alert.

About this task

The following procedure shows you how to set up notifications for all `hm.alert.raised` messages and all `hm.alert.cleared` messages.

All `hm.alert.raised` messages and all `hm.alert.cleared` messages include an SNMP trap. The names of the SNMP traps are `HealthMonitorAlertRaised` and `HealthMonitorAlertCleared`. For information about SNMP traps, see the *Clustered Data ONTAP Network Management Guide*.

Steps

1. Use the `event destination create` command to define the destination to which you want to send the EMS messages.

Example

```
cluster1::> event destination create -name health_alerts -mail
admin@example.com
```

2. Use the `event route add-destinations` command to route the `hm.alert.raised` message and the `hm.alert.cleared` message to a destination.

Example

```
cluster1::> event route add-destinations -messagename hm.alert* -
destinations health_alerts
```

Related concepts

[Managing event messages](#) on page 192

Related information

[Clustered Data ONTAP 8.3 Network Management Guide](#)

Responding to degraded system health

When your system's health status is degraded, you can show alerts, read about the probable cause and corrective actions, show information about the degraded subsystem, and resolve the problem.

Suppressed alerts are also shown so that you can modify them and see whether they have been acknowledged.

About this task

You can discover that an alert was generated by viewing an AutoSupport message or an EMS event, or by using the `system health` commands.

Steps

1. Use the `system health alert show` command to view the alerts that are compromising the system's health.
2. Read the alert's probable cause, possible effect, and corrective actions to determine whether you can resolve the problem or need more information.
3. If you need more information, take any of the following actions:
 - Use the `system health alert show -instance` command to view additional information available for the alert.
 - Use the specific commands in the `system health` command directory for the affected subsystem to investigate the problem.

Example

For example, if a disk has a problem, use the `system health node-connectivity disk` command to get more information about the disk.

4. Use the `system health alert modify` command with the `-acknowledge` parameter to indicate that you are working on a specific alert.
5. Take corrective action to resolve the problem as described by the `Corrective Actions` field in the alert.

The corrective actions might include rebooting the system.

When the problem is resolved, the alert is automatically cleared. If the subsystem has no other alerts, the health of the subsystem changes to `OK`. If the health of all subsystems is `OK`, the overall system health status changes to `OK`.

6. Use the `system health status show` command to confirm that the system health status is `OK`.

If the system health status is not `OK`, repeat this procedure.

Example of responding to degraded system health

By reviewing a specific example of degraded system health caused by a shelf that lacks two paths to a node, you can see what the CLI displays when you respond to an alert.

After starting Data ONTAP, you check the system health and you discover that the status is degraded:

```
cluster1::>system health status show
      Status
      -----
      degraded
```

You show alerts to find out where the problem is, and see that shelf 2 does not have two paths to node1:

```
cluster1::>system health alert show
      Node: node1
      Resource: Shelf ID 2
      Severity: Major
```

```

Indication Time: Mon Nov 10 16:48:12 2013
Probable Cause: Disk shelf 2 does not have two paths to controller
node1.
Possible Effect: Access to disk shelf 2 via controller node1 will be
lost with a single hardware component failure (e.g.
cable, HBA, or IOM failure).
Corrective Actions: 1. Halt controller node1 and all controllers attached to disk shelf 2.
2. Connect disk shelf 2 to controller node1 via two paths following the
rules in the Universal SAS and ACP Cabling Guide.
3. Reboot the halted controllers.
4. Contact support personnel if the alert persists.

```

You display details about the alert to get more information, including the alert ID:

```

cluster1::>system health alert show -monitor node-connect -alert-id DualPathToDiskShelf_Alert
instance
      Node: node1
      Monitor: node-connect
      Alert ID: DualPathToDiskShelf_Alert
      Alerting Resource: 50:05:0c:c1:02:00:0f:02
      Subsystem: SAS-connect
      Indication Time: Mon Mar 21 10:26:38 2011
      Perceived Severity: Major
      Probable Cause: Connection_establishment_error
      Description: Disk shelf 2 does not have two paths to controller node1.
      Corrective Actions: 1. Halt controller node1 and all controllers attached to disk shelf 2.
2. Connect disk shelf 2 to controller node1 via two paths following
the rules in the Universal SAS and ACP Cabling Guide.
3. Reboot the halted controllers.
4. Contact support personnel if the alert persists.
      Possible Effect: Access to disk shelf 2 via controller node1 will be lost with a single
hardware component failure (e.g. cable, HBA, or IOM failure).
      Acknowledge: false
      Suppress: false
      Policy: DualPathToDiskShelf_Policy
      Acknowledger: -
      Suppressor: -
      Additional Information: Shelf uuid: 50:05:0c:c1:02:00:0f:02
      Shelf id: 2
      Shelf Name: 4d.shelf2
      Number of Paths: 1
      Number of Disks: 6
      Adapter connected to IOMA:
      Adapter connected to IOMB: 4d
      Alerting Resource Name: Shelf ID 2

```

You acknowledge the alert to indicate that you are working on it.

```

cluster1::>system health alert modify -node node1 -alert-id DualPathToDiskShelf_Alert -
acknowledge true

```

You fix the cabling between shelf 2 and node1, and then reboot the system. Then you check system health again, and see that the status is OK:

```

cluster1::>system health status show
Status
-----
OK

```

Configuring discovery of cluster and management network switches

The cluster switch health monitor automatically attempts to discover your cluster and management network switches using the Cisco Discovery Protocol (CDP). You must configure the health monitor if it cannot automatically discover a switch or if you do not want to use CDP for automatic discovery.

About this task

The system `cluster-switch show` command lists the switches that the health monitor discovered. If you do not see a switch that you expected to see in that list, then the health monitor cannot automatically discover it.

Steps

1. If you want to use CDP for automatic discovery, do the following; otherwise, go to Step 2:

- a. Ensure that the Cisco Discovery Protocol (CDP) is enabled on your switches.
Refer to your switch documentation for instructions.
 - b. Run the following command on each node in the cluster to verify whether CDP is enabled or disabled:

```
run -node node_name -command options cdpd.enable
```


If CDP is enabled, go to step d. If CDP is disabled, go to step c.
 - c. Run the following command to enable CDP:

```
run -node node_name -command options cdpd.enable on
```


Wait five minutes before you go to the next step.
 - d. Use the `system cluster-switch show` command to verify whether Data ONTAP can now automatically discover the switches.
2. If the health monitor cannot automatically discover a switch, use the `system cluster-switch create` command to configure discovery of the switch:

Example

```
cluster1::> system cluster-switch create -device switch1 -address
192.0.2.250 -snmp-version SNMPv2c -community cshml! -model NX5020 -
type cluster-network
```

Wait five minutes before you go to the next step.

3. Use the `system cluster-switch show` command to verify that Data ONTAP can discover the switch for which you added information.

After you finish

Verify that the health monitor can monitor your switches.

Verifying the monitoring of cluster and management network switches

The cluster switch health monitor automatically attempts to monitor the switches that it discovers; however, monitoring might not happen automatically if the switches are not configured correctly. You should verify that the health monitor is properly configured to monitor your switches.

Steps

1. Use the `system health cluster-switch show` command to identify the switches that the cluster switch health monitor discovered.

If the `Model` column displays the value `OTHER`, then Data ONTAP cannot monitor the switch. Data ONTAP sets the value to `OTHER` if a switch that it automatically discovers is not supported for health monitoring.

Note: If a switch does not display in the command output, then you must configure discovery of the switch.

2. Upgrade to the latest supported switch software and reference the configuration file (RCF) from the [Cisco Ethernet Switch page](#).

The community string in the switch's RCF must match the community string that the health monitor is configured to use. By default, the health monitor uses the community string `cshml!`.

If you need to change information about a switch that the cluster monitors, you can modify the community string that the health monitor uses by using the `system health cluster-switch modify` command.

3. Verify that the switch's management port is connected to the management network.

This connection is required to perform SNMP queries.

Related tasks

[Configuring discovery of cluster and management network switches](#) on page 223

Commands for monitoring the health of your system

You can use the `system health` commands to display information about the health of system resources, to respond to alerts, and to configure future alerts. Using the CLI commands enables you to view in-depth information about how health monitoring is configured. For more information, see the man pages for the commands.

Displaying the status of system health

If you want to...	Use this command...
Display the health status of the system, which reflects the overall status of individual health monitors	<code>system health status show</code>
Display the health status of subsystems for which health monitoring is available	<code>system health subsystem show</code>

Displaying the status of cluster connectivity

If you want to...	Use this command...
Display the status of shelves from the cluster-level view	<code>system health system-connectivity shelf show</code>
Display detailed information about each shelf, including the shelf's UUID and ID, its connected nodes, and the number of paths to the shelf	<code>system health system-connectivity shelf show-instance</code>

Displaying the status of node connectivity

If you want to...	Use this command...
Display the status of shelves from the node-level view, along with other information, such as the owner node, shelf name, and how many disks and paths the shelf has	<code>system health node-connectivity shelf show</code> Use the <code>-instance</code> parameter to display detailed information about each shelf.
Display the status of disks, along with other information, such as the owner node, disk name and bay number, and the number of paths to the disk	<code>system health node-connectivity disk show</code> Use the <code>-instance</code> parameter to display detailed information about each disk.

If you want to...	Use this command...
Display the status of adapters, along with other information, such as the owner node, whether they are used and enabled, and the number of shelves attached	<code>system health node-connectivity adapter show</code> Use the <code>-instance</code> parameter to display detailed information about each adapter.

Managing the discovery of cluster and management network switches

If you want to...	Use this command...
Display the switches that the cluster monitors	<code>system health cluster-switch show</code>
Display the switches that the cluster currently monitors, including switches that you deleted (shown in the Reason column in the command output), and configuration information that you need for network access to the cluster and management network switches This command is available at the advanced privilege level.	<code>system health cluster-switch show-all</code>
Configure discovery of an undiscovered switch	<code>system health cluster-switch create</code>
Modify information about a switch that the cluster monitors (for example, device name, IP address, SNMP version, and community string)	<code>system health cluster-switch modify</code>
Disable monitoring of a switch	<code>system health cluster-switch modify -disable-monitoring</code>
Display the interval in which the health monitor polls switches to gather information	<code>system health cluster-switch polling-interval show</code>
Modify the interval in which the health monitor polls switches to gather information This command is available at the advanced privilege level.	<code>system health cluster-switch polling-interval modify</code>
Disable discovery and monitoring of a switch and delete switch configuration information	<code>system health cluster-switch delete</code>
Permanently remove the switch configuration information which is stored in the database (doing so reenables automatic discovery of the switch)	<code>system health cluster-switch delete -force</code>

Responding to generated alerts

If you want to...	Use this command...
Display information about generated alerts, such as the resource and node where the alert was triggered, and the alert's severity and probable cause	<code>system health alert show</code>
Display information about each generated alert	<code>system health alert show -instance</code>
Indicate that someone is working on an alert	<code>system health alert modify</code>

If you want to...	Use this command...
Acknowledge an alert	<code>system health alert modify - acknowledge</code>
Suppress a subsequent alert so that it does not affect the health status of a subsystem	<code>system health alert modify - suppress</code>
Delete an alert that was not automatically cleared	<code>system health alert delete</code>
Display information about the AutoSupport messages that alerts triggered within the last week—for example, to determine if an alert triggered an AutoSupport message	<code>system health autosupport trigger history show</code>

Configuring future alerts

If you want to...	Use this command...
Enable or disable the policy that controls whether a specific resource state raises a specific alert	<code>system health policy definition modify</code>

Displaying information about how health monitoring is configured

If you want to...	Use this command...
Display information about health monitors, such as their nodes, names, subsystems, and status	<code>system health config show</code> Note: Use the <code>-instance</code> parameter to display detailed information about each health monitor.
Display information about the alerts that a health monitor can potentially generate	<code>system health alert definition show</code> Note: Use the <code>-instance</code> parameter to display detailed information about each alert definition.
Display information about health monitor policies, which determine when alerts are raised	<code>system health policy definition show</code> Note: Use the <code>-instance</code> parameter to display detailed information about each policy. Use other parameters to filter the list of alerts—for example, by policy status (enabled or not), health monitor, alert, and so on.

Downgrading software versions between minor releases

When downgrading a cluster from Data ONTAP 8.2.1 to Data ONTAP 8.2, Data ONTAP clears the history of deleted switches. It also sets the `Model` column to `OTHER` if a switch that it automatically discovers is not supported for health monitoring.

If you want to...	Use this command...
Downgrade the software version command	<code>system health cluster-switch prepare-to-downgrade</code>

Using dashboards to display critical system information

Dashboards provide visibility into critical aspects of your cluster, including Storage Virtual Machine (SVM, formerly known as Vserver) health, system and cluster performance, and storage space utilization. You can also configure alarm thresholds and view information about alarms.

You can configure alarm thresholds for the following:

- Aggregate utilization (aggregate-used)
- Average client latency of NFS and CIFS operations (op-latency)
- CPU utilization (cpu-busy)
- Packet error ratio (port-problems)
- Port utilization (port-util)

For example, you can modify the warning and critical alarm thresholds for space used on aggregates. You might set the warning threshold to 50% and the critical threshold to 60%. The cluster generates an "over threshold" alarm when the value exceeds the configured threshold. In addition, the Event Management System (EMS) generates a message when an alarm is generated or cleared, if you configured it to do so.

Getting notified of dashboard alarms

You can view dashboard alarms by using the `dashboard alarm show` command. You can also subscribe to specific Event Management System (EMS) messages to receive notifications of dashboard alarms.

Before you begin

You must have used the `dashboard alarm thresholds modify` command to specify that the EMS sends a message when an alarm is generated.

About this task

The EMS generates messages for dashboard alarms when the threshold value is equal or greater than the critical threshold (rising) and when the threshold value is less than the warning value (falling). You need to route EMS messages for the object type for which you want alarm notifications:

aggregate-used

The following EMS messages are related to this object type:

- `mgmtgwd.aggregate.used.rising`
- `mgmtgwd.aggregate.used.falling`

cpu-busy

The following EMS messages are related to this object type:

- `mgmtgwd.cpu.busy.rising`
- `mgmtgwd.cpu.busy.falling`

op-latency

The following EMS messages are related to this object type:

- `mgmtgwd.op.latency.rising`

- `mgmtgwd.op.latency.falling`

port-problems

The following EMS messages are related to this object type:

- `mgmtgwd.port.problems.rising`
- `mgmtgwd.port.problems.falling`

port-util

The following EMS messages are related to this object type:

- `mgmtgwd.port.util.rising`
- `mgmtgwd.port.util.falling`

Steps

1. Use the `event destination create` command to define the destination to which you want to send the EMS messages.

Example

```
cluster1::> event destination create -name dashboard_alarms -mail
admin@example.com
```

2. Use the `event route add-destinations` command to route EMS messages to a destination.

Example

The following example specifies that aggregate utilization messages go to the destination named `dashboard_alarms`.

```
cluster1::> event route add-destinations -messagename
mgmtgwd.aggregate.used* -destinations dashboard_alarms
```

Example

The following example specifies that all dashboard alarm messages go to the destination named `dashboard_alarms`.

```
cluster1::> event route add-destinations -messagename
mgmtgwd.aggregate.used*,mgmtgwd.port.problems*,mgmtgwd.op.latency*,
mgmtgwd.port.util*,mgmtgwd.cpu.busy* -destinations dashboard_alarms
```

Commands for managing dashboards

The `dashboard` commands are deprecated, but you can still use them to configure dashboards, display dashboard information, and display health status for SVMs.

Note: The `dashboard health vserver` commands support the NFS and CIFS protocols. They do not support the FC and iSCSI protocols.

If you want to...	Use this command...
Configure the following cluster-wide alarm settings: <ul style="list-style-type: none"> • The threshold value that generates a warning or critical alarm for an event • Whether an EMS message is sent when an alarm is generated • The interval at which objects are monitored by the alarm dashboard 	<code>dashboard alarm thresholds modify</code>
Display settings about alarm thresholds	<code>dashboard alarm thresholds show</code>
Display information about alarms whose values exceed the configured threshold value	<code>dashboard alarm show</code>
Display information about system and cluster performance	<code>dashboard performance show</code>
Display information about storage space utilization and trend	<code>dashboard storage show</code>
Display information about general SVM health, including the current operational status, issues, critical alerts, warnings, and informational messages	<code>dashboard health vserver show</code>
Display the health status of aggregates, LIFs, ports, protocols, and volumes in SVMs	<code>dashboard health vserver show-combined</code>
Display the health status of aggregates in SVMs	<code>dashboard health vserver show-aggregate</code>
Display the health status of volumes in SVMs	<code>dashboard health vserver show-volume</code>
Display the health status of LIFs in SVMs	<code>dashboard health vserver show-lif</code>
Display the health status of SVM network ports	<code>dashboard health vserver show-port</code>
Display the health status of protocols in SVMs	<code>dashboard health vserver show-protocol</code>

For more information, see the man pages.

Monitoring cluster performance

You can view data about your cluster to monitor cluster performance. For example, you can monitor the performance of volumes by viewing statistics that show throughput and latency. You can view statistics by using preconfigured commands associated with performance presets or by using advanced commands that require more knowledge.

You can upload a performance archive to the NetApp support site or a specified URI by using the `system node autosupport invoke-performance-archive` command. After you upload a performance archive to the NetApp support site, you can view the data by using the SmartSolve tool that is available on the NetApp support site.

Related tasks

[Uploading performance archive files](#) on page 209

What objects, instances, and counters are

You can view performance data for specific objects in your cluster. Objects are comprised of instances and counters. Counters provide data about the instances of an object.

An object is any of the following:

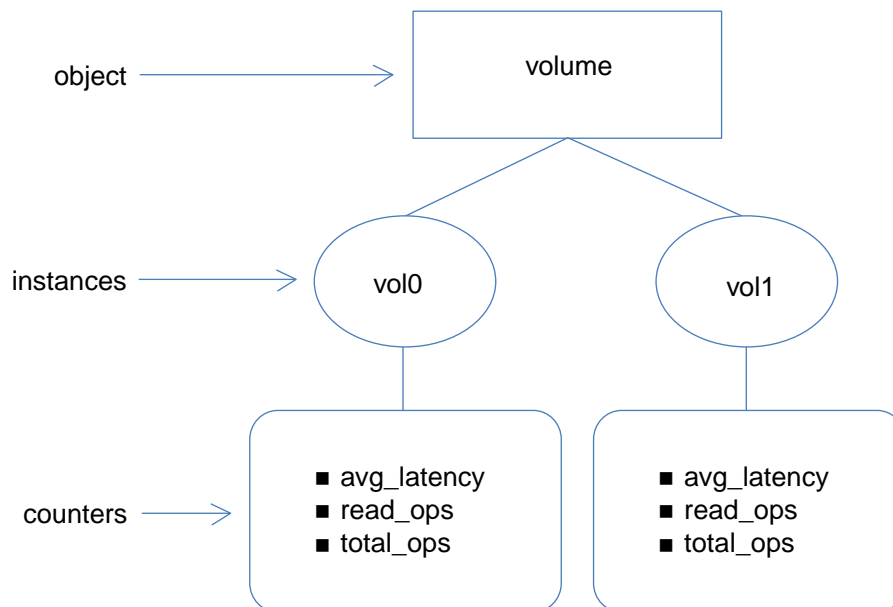
- Physical entities such as disks, processors, and ports
- Logical entities such as LUNs, volumes, and workloads
- Protocols such as CIFS, NFS, iSCSI, and FC

Each object has zero or more instances. For example, the LUN object has an instance for each LUN in your cluster.

A counter is a predefined performance metric that provides data about an object. Examples of data that counters provide include the following:

- Disk capacity
- The average latency for a volume
- The number of established SMB and SMB2 sessions

The following illustration shows the relationship between an object and its instances and counters. In this illustration, the volume object has two instances: vol0 and vol1. The object's counters provide data about each of these instances. The illustration shows three of the object's counters: `avg_latency`, `read_ops`, and `total_ops`.



How performance presets define what data is collected

Performance presets define what counters collect data for objects and whether any of the data is added to performance archives. Performance presets that are associated with `statistics` commands simplify how you collect and view performance data for objects.

The system uses many performance presets. Some performance presets are associated with `statistics` commands, and other performance presets are used for internal processes, such as collecting data for performance archives.

Performance presets include the following information:

- Name of the performance preset
- Name of the objects for which the performance preset collects data
- Whether a sample period of collected data is added to performance archives
- Name of the counters to collect data for each object

Performance presets that are associated with `statistics` commands enable you to collect and view performance data for objects by using a simplified command. For example, you can use performance presets with the `statistics start` and `statistics show-periodic` commands by using the `-preset` parameter.

Related concepts

[Monitoring cluster performance using advanced methods](#) on page 235

Related tasks

[Collecting and viewing performance data for objects](#) on page 232

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Collecting and viewing performance data for objects

You can collect and view performance data for an object by using simplified `statistics` commands. The performance preset associated with the command defines what counters collect data for the object, and you use the `statistics` command to define how many iterations of data to collect.

About this task

This procedure shows an example of collecting and displaying performance data for instances of a volume in a cluster by using the `statistics volume show` command. A performance preset is associated with the `statistics volume show` command and defines what counters collect data for instances of a volume in a cluster.

The `statistics` command family includes many commands that are associated with performance presets to simplify how you collect and view performance data for different objects.

Step

1. Collect and view performance data for instances of a volume by using the `statistics volume show` command.

By default, data is returned for the top 25 instances of a volume to a maximum of 100 instances of a volume.

Example

In the following example, performance data for instances of a volume is collected for one five-second iteration and displayed. By default, data is collected for 5 seconds for each iteration:

```
cluster1::> statistics volume show -iterations 1
```

In the following example, performance data for instances of a volume is continuously collected and displayed, until you manually stop the data collection by pressing Ctrl-C:

```
cluster1::> statistics volume show -iterations 0
```

Related concepts

[How performance presets define what data is collected](#) on page 232

Related references

[Commands for monitoring cluster performance](#) on page 238

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Filtering the statistics

You can use filters to help you track resource utilization for a specific object, or narrow down the amount of statistics that are collected. For example, collecting statistics from 30,000 LUNs in a system could take a long time to complete, and you could narrow down the amount of statistics that are collected and the time it takes to complete by filtering statistics by volume name.

About this task

The examples in this procedure show how you can filter performance data by volume.

Steps

1. Use the `statistics volume show` command to view performance data for a volume.

The output includes the most important counters and sorts the results by the most active instances of the object you specify.

Example

In this example, the output shows the volumes reporting the highest IOP values:

```
cluster1::> statistics volume show
cluster1 : 12/31/2013
16:00:04

          *Total Read Write Other  Read Write
Latency  Volume Vserver Aggregate  Ops  Ops  Ops  Ops  (Bps) (Bps)
(us)
-----
vol0     -   aggr0           58  13   15   29 9585  3014
39
vol1     -   aggr0_n0        56   0   11   45 8192 28826  47
```

2. Filter the output by volume name.

Example

In this example, the output shows statistics for vol1:

```
cluster1::> statistics volume show -volume
vol1
cluster1 : 12/31/2013
16:00:04

*Total Read Write Other   Read Write
Latency
Volume Vserver Aggregate   Ops  Ops  Ops  Ops  (Bps) (Bps)
(us)
-----
vol1      -   aggr0_n0      56   0   11   45   8192 28826   47
```

Sorting the statistics

You can sort the statistics by any counter to diagnose a performance issue or identify a hot spot. For example, you might want to collect volume statistics and sort by total operations to get a list of most active volumes.

Steps

1. Switch to the advanced privilege level by using the `set -privilege advanced` command.
2. View the sample data by using the `statistics show` command.

Example

In this example, the output shows the volumes by volume name, read operations, and write operations, and sorts the output by `read_ops` counter:

```
::>statistics show -object volume -counter read_ops|write_ops -
filter "vserver_name=vs1,read_ops>15" -sort-key read_ops -sort-order
ascending

Object: volume
Instance: vol1
Start-time: 05/23/2014 4:00 PM
End-time: 05/23/2014 4:10 PM
Cluster: cluster1
Number of Constituents: 1 (complete_aggregation)
Counter                               Value
-----
read_ops                               20
write_ops                               90

Object: volume
Instance: vol2
Start-time: 05/23/2014 4:00 PM
End-time: 05/23/2014 4:10 PM
Cluster: cluster1
Number of Constituents: 1 (complete_aggregation)
Counter                               Value
-----
read_ops                               40
write_ops                               30
```

Importing a performance preset configuration (cluster administrators only)

You can create a custom performance preset or modify a writable performance preset by importing a performance preset configuration in XML file format. You can also use this method to modify what data is collected and stored in performance archives.

Before you begin

- You must have cluster administrator privileges to perform this task.
- You must have a performance preset configuration in XML file format.
Technical support can help you create the XML file of performance preset definitions.

About this task

Data ONTAP includes a performance archive that automatically collects and stores performance statistics at predefined times. With the help of technical support, you can modify what data is collected for the performance archive by importing a performance preset.

You cannot modify read-only performance presets. You can only modify performance presets that have the `read-only` parameter set to `false`.

Steps

1. Switch to the advanced privilege level by using the `set -privilege advanced` command.
2. Import the performance preset configuration by using the `statistics preset import` command.

Example

In the following example, a performance preset configuration is imported from the NetApp support site:

```
cluster1::*> statistics preset import -source-uri http://
www.netapp.com/support/
nfs_monitor.xml -comment "New NFS Monitor preset."
```

Monitoring cluster performance using advanced methods

You can collect and view performance data using the `statistics start`, `statistics stop`, and `statistics show` commands. The commands let you specify how to collect and display performance data and require advanced privilege and knowledge of objects and counters.

If you want to use commands that are preconfigured to collect and display data about specific objects, use the simplified `statistics` commands, such as the `statistics volume show` command.

Related concepts

[How performance presets define what data is collected](#) on page 232

Decisions to make before you view performance data

You can view performance data in several ways. You should make a few decisions before you view the data.

You should decide the following before you view performance data:

Decision	Considerations
How do you want to retrieve and display the data?	<p>You have two choices:</p> <ul style="list-style-type: none"> You can collect and view a set of data for a specific time period. If you choose this option, you can view data for several objects and instances at a time. You can view continuously updated data. If you choose this option, you can view data for only one object and one instance at a time when you use the <code>statistics show-periodic</code> command.
For which objects do you want to view data?	You need to specify at least one object for which you want to view data.
Do you want data from all counters or from specific counters?	The default setting shows data for all counters in an object; however, you can specify specific counters to get the exact data that you need.
Do you want data for all instances of an object or for specific instances?	<ul style="list-style-type: none"> If you collect data for a time period, the default setting shows data for all instances; however, you can specify one or more instances. If you view continuously updated data and specify any object other than <code>cluster</code>, you must specify an instance when you use the <code>statistics show-periodic</code> command.
Do you want data for the entire cluster or do you want to scope the data?	The default setting shows data for the entire cluster; however, you can scope the data to a specific SVM or a specific node.

Viewing performance data for a time period

You can monitor cluster performance by collecting and viewing data for a specific time period (a sample). You can view data for several objects and instances at a time.

About this task

You can collect more than one data sample at a time. You can collect more than one sample from the same object at the same time.

Note: You cannot collect and view data for an object that has more than 5,000 instances. If an object has more than 5,000 instances, you need to specify the specific instances for which you want data. This applies to all `statistics` commands, including `statistics` views.

For more information about the `statistics` commands, see the man pages.

Steps

1. Switch to the advanced privilege level by using the `set -privilege advanced` command.
2. Use the `statistics start` command to start collecting data.

If you do not specify the `-sample-id` parameter, the command generates a sample identifier for you and defines this sample as the default sample for the CLI session. If you run this command

again during the same CLI session and do not specify the `-sample-id` parameter, the command can overwrite the previous default sample. You are prompted to confirm whether to overwrite the previous default sample.

- Optional: Use the `statistics stop` command to stop collecting data for the sample.

You can view data from the sample if you do not stop data collection. Stopping data collection gives you a fixed sample. Not stopping data collection gives you the ability to get updated data that you can use to compare against previous queries. The comparison can help you identify performance trends.

- Use the `statistics show` command to view the sample data.

Example: Monitoring NFSv3 performance

The following example shows performance data for the NFSv3 protocol.

The following command starts data collection for a new sample:

```
cluster1::*> statistics start -object nfsv3 -sample-id nfs_sample
```

The following command shows data from the sample by specifying counters that show the number of successful read and write requests versus the total number of read and write requests:

```
cluster1::*> statistics show -sample-id nfs_sample -counter
read_total|write_total|read_success|write_success
```

```
Object: nfsv3
Instance: vs1
Start-time: 2/11/2013 15:38:29
End-time: 2/11/2013 15:38:41
Cluster: cluster1
```

Counter	Value
read_success	40042
read_total	40042
write_success	1492052
write_total	1492052

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Viewing continuously updated performance data

You can monitor cluster performance by viewing data that continuously updates with the latest status. You can view data for only one object and one instance at a time.

About this task

For more information about the `statistics show-periodic` command, see the man page.

Note: The `statistics show-periodic` command is deprecated, but you can still use it to view performance data.

You can also use the `statistics show` command with the `-tab` parameter to display continuously updated data.

Steps

1. Switch to the advanced privilege level by using the `set -privilege advanced` command.
2. Use the `statistics show-periodic` command to view continuously updated performance data.

If you do not specify the `-object` parameter, the command returns summary data for the cluster.

Example: Monitoring volume performance

This example shows how you can monitor volume performance. For example, you might want to monitor volume performance if critical applications run on those volumes. Viewing the performance data can help you answer questions such as the following:

- What is the average response time for a volume?
- How many operations are being completed per second?

The following command shows performance data for a volume by specifying counters that show the number of operations per second and latency:

```
cluster1::*> statistics show-periodic -object volume -instance
vol0 -counter write_ops|read_ops|total_ops|read_latency|
write_latency|avg_latency
cluster1: volume.vol0: 1/7/2013 20:15:51
  avg      read      total      write      write
  latency  latency read_ops   ops        latency    ops
-----
  202us    218us     0          22         303us     7
   97us    43us     31         71         149us    34
   39us     0us     0           3           0us     0
  152us     0us     0           16         152us    16
  162us     0us     0          342         144us   289
  734us     0us     0           15           0us     0
   49us     0us     0           1           0us     0
cluster: volume.vol0: 1/7/2013 20:16:07
  avg      read      total      write      write
  latency  latency read_ops   ops        latency    ops
-----
Minimums:
  39us     0us     0           1           0us     0
Averages for 7 samples:
  205us    37us     4           67         106us    49
Maximums:
  734us    218us    31          342         303us   289
```

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Commands for monitoring cluster performance

You can use the `statistics` commands to display performance data and specify the settings for displaying the data. For more information about these commands, see the man pages.

Collecting data for objects by using performance presets

You can use the following commands to collect and view data samples for objects. Each of the following commands is associated with a performance preset that defines what counters collect data for the object. It is noted when one of the following commands requires advanced privilege.

If you want to...	Use this command...
Collect and view data samples for aggregates	<code>statistics aggregate show</code>
Collect and view data samples for caches	<code>statistics cache flash-pool show</code>
Collect and view data samples for disks	<code>statistics disk show</code> This command is available at the advanced privilege level.
Collect and view data samples for logical interfaces (LIFs)	<code>statistics lif show</code>
Collect and view data samples for LUNs	<code>statistics lun show</code>
Collect and view data samples for each node in the cluster	<code>statistics node show</code>
Collect and view data samples for FC ports	<code>statistics port fcp show</code>
Collect and view data samples for a cluster	<code>statistics system show</code>
Collect and view data samples for volumes	<code>statistics volume show</code>
Collect and view data samples for Storage Virtual Machines (SVMs)	<code>statistics vserver show</code>
Collect and view data samples for QoS workloads	<code>statistics workload show</code> This command is available at the advanced privilege level.

Viewing performance presets

You can use the following commands to view performance presets. The following commands require advanced privilege.

If you want to...	Use this command...
View all performance presets	<code>statistics preset show</code>
View details of each performance preset	<code>statistics preset detail show</code>

Creating, modifying, and deleting performance presets

You can use the following commands to create, modify, and delete performance presets. The following commands require advanced privilege.

Note: You cannot modify or delete read-only performance presets. You can only modify and delete performance presets that have the `read-only` parameter set to `false`.

If you want to...	Use this command...
Modify properties of a performance preset, such as the name or comment	<code>statistics preset modify</code>
Create or modify a performance preset by importing an XML file of definitions	<code>statistics preset import</code> Contact technical support for help with creating the XML file of definitions.
Delete a performance preset	<code>statistics preset delete</code>

Collecting data for a sample time period by using advanced methods

You can use the following commands to collect data samples and to manage the samples that you collect. You must collect a data sample before you can use the `statistics show` command. The following commands require advanced privilege.

If you want to...	Use this command...
Start data collection for a sample	<code>statistics start</code> You can use a performance preset to collect data by using the <code>-preset</code> parameter.
Stop data collection for a sample	<code>statistics stop</code>
View all samples	<code>statistics samples show</code>
Delete a sample	<code>statistics samples delete</code>

Viewing performance data by using advanced methods

You can use the following commands to view performance data. You must collect a data sample before you can use the `statistics show` command. The following commands require advanced privilege.

If you want to...	Use this command...
View performance data for a sample time period	<code>statistics show</code> You should limit the scope of this command to only a few objects at a time to avoid a potentially significant impact on system performance.
View continuously updated performance data	<code>statistics show-periodic</code> This command is deprecated, but you can still use it to view performance data.

Viewing all objects, instances, and counters

You can use the `statistics catalog` commands to view information about objects, instances, and counters. The following commands require advanced privilege.

If you want to...	Use this command...
View descriptions of objects	<code>statistics catalog object show</code>
View all instances of an object	<code>statistics catalog instance show</code>
View descriptions of counters in an object	<code>statistics catalog counter show</code>

Managing settings for the statistics commands

You can use the `statistics settings` commands to modify settings for the `statistics` commands. The following commands require advanced privilege.

If you want to...	Use this command...
View the settings for the <code>statistics</code> commands	<code>statistics settings show</code>
Modify whether the commands display rate statistics in rates per second.	<code>statistics settings modify</code>

Viewing advanced performance data

Although the following commands are deprecated, you can currently still use them to view advanced performance data about your cluster. The following commands require advanced privilege.

If you want to...	Use this command...
View information about SecD RPC usage statistics for the nodes in the cluster	<pre>statistics secd show</pre> <p>Use this command only as directed by support personnel to help analyze performance and diagnose problems.</p>
View information about the ONC RPC calls performed by the nodes in the cluster	<pre>statistics oncrpc show-rpc-calls</pre>

Related tasks

[Uploading performance archive files](#) on page 209

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Displaying environmental information

Sensors help you monitor the environmental components of your system. The information you can display about environmental sensors include their type, name, state, value, and threshold warnings.

Step

1. To display information about environmental sensors, use the `system node environment sensors show` command.

Managing system performance (cluster administrators only)

You can use several features to improve system performance. Only the cluster administrator can manage system performance. The SVM administrator cannot perform these tasks.

Managing workload performance by using Storage QoS


Storage QoS (Quality of Service) can help you manage risks around meeting your performance objectives. You use Storage QoS to limit the throughput to workloads and to monitor workload performance. You can reactively limit workloads to address performance problems and you can proactively limit workloads to prevent performance problems.

A workload represents the input/output (I/O) operations to one of the following kinds of storage objects:

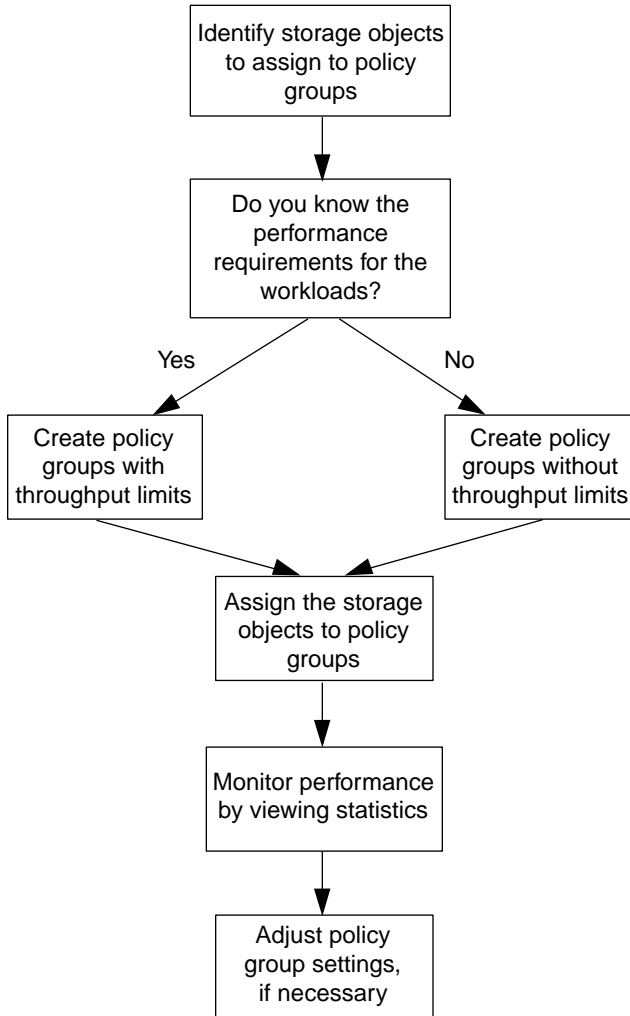
- SVMs with FlexVol volumes
- FlexVol volumes
- LUNs
- Files (typically represent virtual machines)

You assign a storage object to a policy group to control and monitor a workload. You can monitor workloads without controlling them.

The following illustration shows an example environment before and after using Storage QoS. On the left, workloads compete for cluster resources to transmit I/O. These workloads get "best effort" performance, which means you have less performance predictability (for example, a workload might get such good performance that it negatively impacts other workloads). On the right are the same workloads assigned to policy groups. The policy groups enforce a maximum throughput limit.



The following workflow shows how you use Storage QoS to control and monitor workloads:

**Related tasks**

[Controlling and monitoring workload performance](#) on page 247

How Storage QoS works

Storage QoS controls workloads that are assigned to policy groups by throttling and prioritizing client operations (SAN and NAS data requests) and system operations.

What policy groups are

A policy group is comprised of one or more workloads and a performance limit that applies collectively to all workloads in the policy group. There are two types of policy groups:

User-defined policy group

Enforces a maximum throughput limit on the storage objects that belong to the policy group by throttling input/output (I/O) requests.

System-defined policy group

Manages internal work that the cluster performs.

You can view performance data for both types of policy groups. The names of system-defined policy groups start with an underscore.

What workloads are

A workload represents work that the cluster performs. There are two types of workloads:

User-defined workload

Represents the input/output (I/O) operations from clients to a storage object that belongs to a policy group. Storage objects can be:

- SVMs with FlexVol volumes
- FlexVol volumes
- LUNs
- Files (typically represent virtual machines)

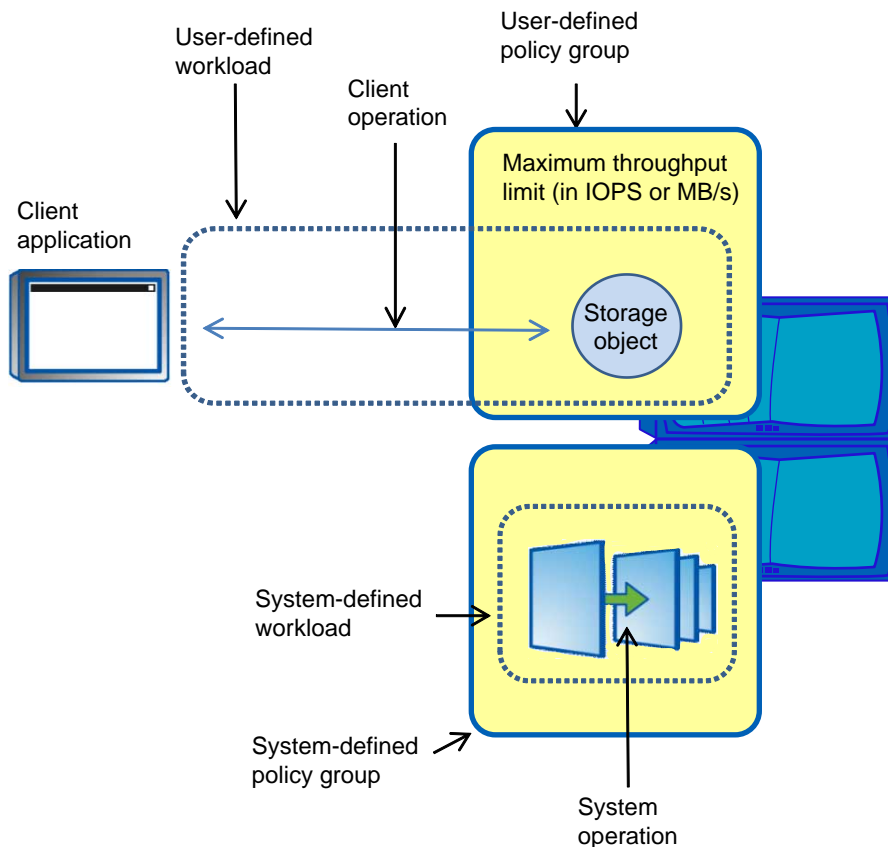
I/O to storage objects that are not assigned to policy groups belongs to the "User-Default" workload.

System-defined workload

Represents internal work that the cluster performs. Storage QoS controls specific system operations to prevent them from interfering with client operations. Examples include storage efficiency operations and data replication operations.

You can view performance data for both types of workloads. The names of system-defined workloads start with an underscore.

The following illustration shows a user-defined policy group and a system-defined policy group. The user-defined policy group controls the user-defined workload, which represents the client operations from the application to the storage object. The system-defined policy group controls the system-defined workload, which represents the internal system operations that the cluster performs.



How the maximum throughput limit works

You can specify one service-level objective for a Storage QoS policy group: a maximum throughput limit. A maximum throughput limit, which you define in terms of IOPS or MBps, specifies the throughput that the workloads in the policy group cannot collectively exceed.

When you specify a maximum throughput for a policy group, Storage QoS controls client operations to ensure that the combined throughput for all workloads in the policy group does not exceed the specified maximum throughput.

For example, assume that you create the policy group “`untested_apps`” and specify a maximum throughput of 300 MBps. You assign three volumes to the policy group. The combined throughput to those three volumes cannot exceed 300 MBps.

Note: The combined throughput to the workloads in a policy group might exceed the specified limit by up to 10%. A deviation might occur if you have a workload that experiences rapid changes in throughput (sometimes called a *bursty workload*).

Note the following about specifying a maximum throughput:

- A throughput limit applies to all clients that access a storage object.
- Do not set the limit too low, because you might underutilize the cluster.
- Consider the minimum amount of throughput that you want to reserve for workloads that do not have limits.
For example, you can ensure that your critical workloads get the throughput that they need by limiting noncritical workloads.
- You might want to provide room for growth.
For example, if you see an average utilization of 500 IOPS, you might specify a limit of 1,000 IOPS.

How throttling a workload can affect non-throttled workload requests from the same client

In some situations, throttling a workload (I/O to a storage object) can affect the performance of non-throttled workloads if the I/O requests are sent from the same client.

If a client sends I/O requests to multiple storage objects and some of those storage objects belong to Storage QoS policy groups, performance to the storage objects that do not belong to policy groups might be degraded. Performance is affected because resources on the client, such as buffers and outstanding requests, are shared.

For example, this might affect a configuration that has multiple applications or virtual machines running on the same host.

This behavior is likely to occur if you set a low maximum throughput limit and there are a high number of I/O requests from the client.

If this occurs, you can increase the maximum throughput limit or separate the applications so they do not contend for client resources.

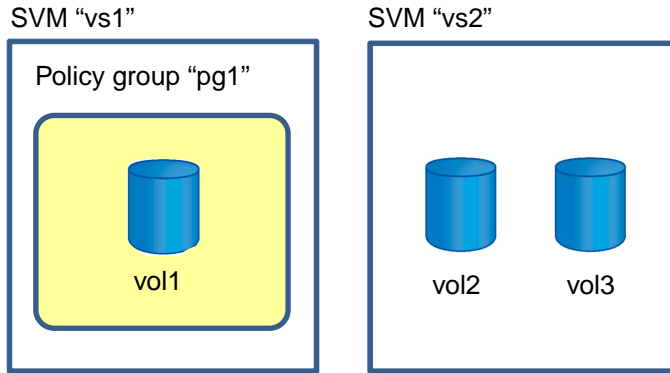
Rules for assigning storage objects to policy groups

You should be aware of rules that dictate how you can assign storage objects to Storage QoS policy groups.

Storage objects and policy groups must belong to the same SVM

A storage object must be contained by the Storage Virtual Machine (SVM) to which the policy group belongs. You specify the SVM to which the policy group belongs when you create the policy group. Multiple policy groups can belong to the same SVM.

In the following illustration, the policy group pg1 belongs to SVM vs1. You cannot assign volumes vol2 or vol3 to policy group pg1 because those volumes are contained by a different SVM.




Nested storage objects cannot belong to policy groups

You cannot assign a storage object to a policy group if its containing object or its child objects belong to a policy group. The following table lists the restrictions.

If you assign the...	Then you cannot assign...
SVM to a policy group	Any storage objects contained by the SVM to a policy group
Volume to a policy group	The volume's containing SVM or any child LUNs or files to a policy group
LUN to a policy group	The LUN's containing volume or SVM to a policy group
File to a policy group	The file's containing volume or SVM to a policy group

In the following illustration, the SVM vs3 is assigned to policy group pg2. You cannot assign volumes vol4 or vol5 to a policy group because an object in the storage hierarchy (SVM vs3) is assigned to a policy group.



Some types of volumes not supported with Storage QoS

You can assign FlexVol volumes to policy groups. Infinite Volumes are not supported with Storage QoS.

The following FlexVol volume variations are not supported with Storage QoS:

- Data protection mirrors
- Load-sharing mirrors
- Node root volumes

How to monitor workload performance when using Storage QoS

To determine an appropriate throughput limit, you should monitor performance from the cluster. You should not use a tool on the host to monitor performance. A host can report different results than the cluster.

Storage QoS limits I/O to and from the cluster. The rate of I/O that the cluster experiences can be different from what an application experiences. For example, reads from the application can go to the file system buffer cache and not to the cluster.

Due to this behavior, you should monitor performance from the cluster and not from a host-side tool.

Supported number of Storage QoS policy groups and workloads

You can create up to 3,500 policy groups per cluster. You can assign up to 10,000 storage objects to those policy groups. Assigning a storage object to a policy group creates a workload. There are no other limits.

Controlling and monitoring workload performance

You control and monitor workload performance to address performance problems and to proactively limit workloads that have defined performance targets.

Before you begin

- You must be familiar with [How the maximum throughput limit works](#) on page 245.
- You must be familiar with [Rules for assigning storage objects to QoS policy groups](#) on page 245.

Steps

1. Identify the storage objects that you want to assign to Storage QoS policy groups.

A best practice is to assign the same type of storage object to all policy groups.

2. Use the `qos policy-group create` command to create a new policy group or use the `qos policy-group modify` command to modify an existing policy group.

You can specify a maximum throughput limit when you create the policy group or you can wait until after you monitor the workload. Monitoring the workload first can help you identify the limit that you need to set. If you do not specify a maximum throughput, the workloads get best-effort performance.

Example

The following command creates policy group `pg-vs1` with a maximum throughput of 5,000 IOPS.

```
cluster1::> qos policy-group create pg-vs1 -vserver vs1 -max-throughput 5000iops
```

Example

The following command creates policy group `pg-app2` without a maximum throughput.

```
cluster1::> qos policy-group create pg-app2 -vserver vs2
```

- To assign a storage object to a policy group, use the create or modify command for the SVM, volume, LUN, or file.

Example

The following command assigns the SVM vs1 to policy group pg-vs1.

```
cluster1::> vserver modify -vserver vs1 -qos-policy-group pg-vs1
```

Example

The following command creates the volume app2 and assigns it to policy group pg-app2.

```
cluster1::> volume create -vserver vs2 -volume app2 -aggregate aggr2 -qos-policy-group pg-app2
```

- To identify whether you are meeting your performance objectives, use the `qos statistics` commands to monitor policy group and workload performance.

You should monitor performance from the cluster. You should not use a tool on the host to monitor performance.

Example

The following command shows the performance of policy groups.

```
cluster1::> qos statistics performance show
Policy Group          IOPS          Throughput    Latency
-----
-total-              12316         47.76MB/s    1264.00us
pg_app2              7216          28.19MB/s    420.00us
pg_vs1               5008          19.56MB/s    2.45ms
_System-Best-Effort    62           13.36KB/s    4.13ms
_System-Background    30            0KB/s        0ms
```

Example

The following command shows the performance of workloads.

```
cluster1::> qos statistics workload performance show
Workload             ID           IOPS          Throughput    Latency
-----
-total-              -            12320         47.84MB/s    1215.00us
app2-wid7967         7967         7219          28.20MB/s    319.00us
vs1-wid12279         12279        5026          19.63MB/s    2.52ms
_USERSPACE_APPS      14           55            10.92KB/s    236.00us
_Scan_Backgro...     5688         20            0KB/s        0ms
```

- If necessary, use the `qos policy-group modify` command to adjust the policy group's maximum throughput limit.

Example

The following command modifies the maximum throughput for policy group pg-app2 to 20 MB/s.

```
cluster1::> qos policy-group modify pg-app2 -max-throughput 20mb/s
```

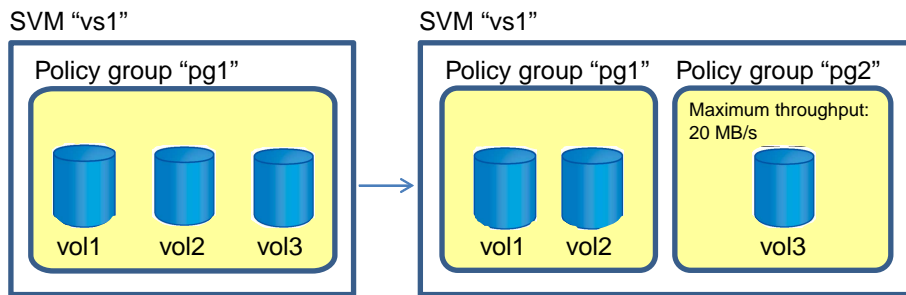
Related references

[Commands for controlling and monitoring workloads](#) on page 251

Example: Isolating a workload

You might have a workload that gets better performance than necessary, which affects the performance of other workloads. To address this problem, you use Storage QoS to throttle the workload, which frees cluster resources for other workloads. In this example, the workloads are at the volume level.

The following illustration shows three volumes. You place each volume in policy group pg1, but you do not set a maximum throughput because you want to monitor the workloads first. When you monitor the workloads, you find that vol3 is getting better performance than other workloads. To limit the workload's resource consumption, you move vol3 to policy group pg2. This should allow the other workloads to speed up.



Using the CLI to isolate a workload

The following command creates a policy group without a maximum throughput.

```
cluster1::> qos policy-group create pg1 -vserver vs1
```

The following command assigns three existing volumes to the policy group.

```
cluster1::> volume modify vol1,vol2,vol3 -vserver vs1 -qos-policy-group pg1
```

The following command displays performance data for the workloads.

```
cluster1::> qos statistics workload performance show
Workload          ID      IOPS      Throughput      Latency
-----
-total-          -      16645      64.77MB/s      411.00us
vol3-wid12459    12459   10063      39.31MB/s      410.00us
vol2-wid1445     1445    3505       13.69MB/s      437.00us
vol1-wid11344    11344    3007       11.75MB/s      277.00us
_USERSPACE_APPS  14       40         26.40KB/s       8.68ms
_Scan_Backgro..  5688    30         0KB/s           0ms
```

The vol3 workload is getting such good performance that other workloads cannot meet your performance objectives. You decide to move that workload to a new policy group that has a maximum throughput.

The following command creates a policy group with a maximum throughput.

```
cluster1::> qos policy-group create pg2 -vserver vs1 -max-throughput 20mb/s
```

The following command assigns vol3 to the new policy group.

```
cluster1::> volume modify vol3 -vserver vs1 -qos-policy-group pg2
```

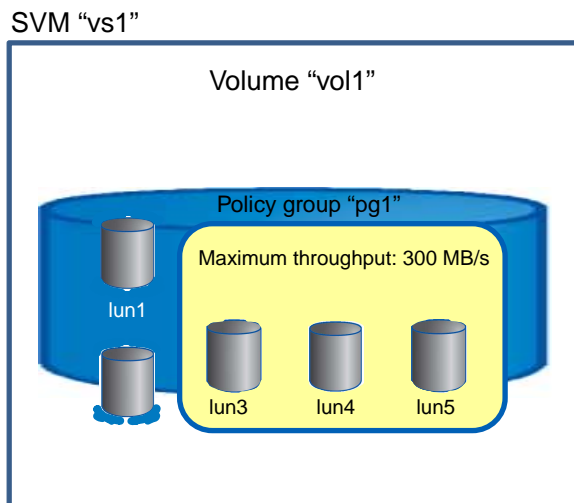
Displaying performance data for the workloads shows that limiting vol3 has allowed the other workloads to get better performance.

```
cluster1::> qos statistics workload performance show
Workload          ID      IOPS      Throughput      Latency
-----
-total-          -      15691      61.17MB/s      1001.00us
vol1-wid11344    11344    6016      23.50MB/s      355.00us
vol3-wid12459    12459    5133      20.05MB/s      2.42ms
vol2-wid1445     1445     4462      17.43MB/s      253.00us
_USERSPACE_APPS  14       50        204.20KB/s     355.00us
_Scan_Backgro..  5688     30        0KB/s          0ms
```

Example: Proactively setting a limit on non-critical workloads

You might want to ensure that your critical workloads get the best performance possible, so you use Storage QoS to limit the throughput to non-critical workloads. In this example, the workloads are at the LUN level.

The following illustration shows five LUNs in volume vol1. lun1 and lun2 are used for critical applications. lun3, lun4, and lun5 are used for non-critical applications. You want lun1 and lun2 to get best effort performance, so you limit lun3, lun4, and lun5 by assigning them to a policy group with a maximum throughput limit.



Using the CLI to set a limit on non-critical workloads

The following command creates a policy group with a maximum throughput of 300 MB/s.

```
cluster1::> qos policy-group create pg1 -vserver vs1 -max-
throughput 300mb/s
```

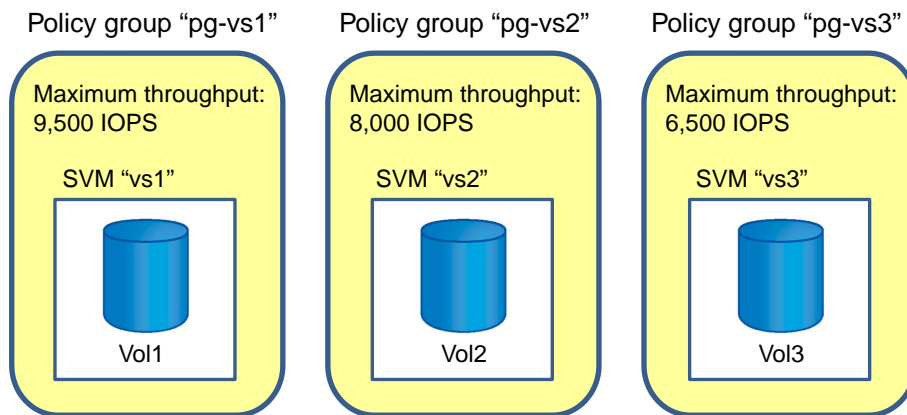
The following commands assign three new LUNs to the policy group.

```
cluster1::> lun create -vserver vs1 -volume vol1 -lun lun3 -size
50GB -ostype windows_2008 -qos-policy-group pg1
cluster1::> lun create -vserver vs1 -volume vol1 -lun lun4 -size
50GB -ostype windows_2008 -qos-policy-group pg1
cluster1::> lun create -vserver vs1 -volume vol1 -lun lun5 -size
50GB -ostype windows_2008 -qos-policy-group pg1
```

Example: Proactively setting a limit on workloads in a shared storage infrastructure

If you have a shared storage infrastructure, you might need to ensure that each workload does not get better performance than necessary. In this example, you use Storage QoS policy groups to set a limit on each workload, all of which are at the Storage Virtual Machine (SVM) level.

The following illustration shows three SVMs assigned to three separate policy groups. You assign each SVM to a policy group because you know the performance objectives for each workload and you do not want one tenant taking system resources from other tenants.



Using the CLI to set a limit on workloads in a shared storage infrastructure

The following commands create three policy groups with maximum throughput limits.

```
cluster1::> qos policy-group create pg-vs1 -vserver vs1 -max-
throughput 9500iops
cluster1::> qos policy-group create pg-vs2 -vserver vs2 -max-
throughput 8000iops
cluster1::> qos policy-group create pg-vs3 -vserver vs3 -max-
throughput 6500iops
```

The following commands assign three existing SVMs to the policy groups.

```
cluster1::> vserver modify -vserver vs1 -qos-policy-group pg-vs1
cluster1::> vserver modify -vserver vs2 -qos-policy-group pg-vs2
cluster1::> vserver modify -vserver vs3 -qos-policy-group pg-vs3
```

Commands for controlling and monitoring workloads

You can use commands to manage Storage QoS policy groups, assign storage objects to policy groups, identify the storage objects that belong to policy groups, and monitor workload and policy group performance.

- [Commands for managing policy groups](#) on page 252

- [Commands for assigning storage objects to policy groups](#) on page 252
- [Commands for identifying the storage objects that belong to policy groups](#) on page 252
- [Commands for monitoring policy group performance](#) on page 253
- [Commands for monitoring workload performance](#) on page 253

For more information about these commands, see the man pages.

Commands for managing policy groups

You use the `qos policy-group` commands to manage policy groups. You use policy groups to control and monitor workload performance.

If you want to...	Use this command...
Create a policy group	<code>qos policy-group create</code>
Modify a policy group	<code>qos policy-group modify</code>
Rename a policy group	<code>qos policy-group rename</code>
View all user-defined policy groups	<code>qos policy-group show</code>
Delete a policy group	<code>qos policy-group delete</code>

Commands for assigning storage objects to policy groups

You use a storage object's `create` command or `modify` command to assign a storage object to a policy group. You assign a storage object to a policy group to control and monitor workload performance.

Note: To remove a storage object from a policy group, set the `-qos-policy-group` parameter to `none`.

If you want to assign the..	Use this command with the <code>-qos-policy-group</code> parameter..
SVM with FlexVol volumes to a policy group	<code>vserver modify</code>
New FlexVol volume to a policy group	<code>volume create</code>
Existing FlexVol volume to a policy group	<code>volume modify</code>
New FlexClone volume to a policy group	<code>volume clone create</code>
New LUN to a policy group	<code>lun create</code>
Existing LUN to a policy group	<code>lun modify</code>
File to a policy group	<code>volume file modify</code>
New clone of a file or LUN to a policy group	<code>volume file clone create</code>

Commands for identifying the storage objects that belong to policy groups

You use a storage object's `show` command to identify the storage objects that belong to policy groups.

If you want to identify the...	Use this command with the <code>-qos-policy-group</code> parameter..
SVMs with FlexVol volumes that belong to a policy group	<code>vserver show</code>

If you want to identify the...	Use this command with the <code>-qos-policy-group</code> parameter...
FlexVol volumes that belong to a policy group	<code>volume show</code>
LUNs that belong to a policy group	<code>lun show</code>

Commands for monitoring policy group and workload performance

You use the following commands to monitor policy group and workload performance in terms of IOPS, throughput, and latency.

If you want to view the...	Use this command...
Collective performance of all workloads in a policy group	<code>qos statistics performance show</code>
Performance of individual workloads	<code>qos statistics workload performance show</code>

Commands for advanced monitoring of policy group performance

You use the following commands to view advanced performance data for policy groups. These commands show the collective performance of all workloads in a policy group.

If you want to view data about...	Use this command...
The client load as it enters the cluster, in terms of request size, read percentage, and concurrency	<code>qos statistics characteristics show</code>
Latency across Data ONTAP subsystems, which helps to determine why response time is slow	<code>qos statistics latency show</code>
CPU utilization	<code>qos statistics resource cpu show</code>
Disk utilization, in terms of the percentage of time spent on the disk during read and write operations	<code>qos statistics resource disk show</code>

Commands for advanced monitoring of workload performance

You use the following commands to view advanced performance data for individual workloads.

If you want to view data about...	Use this command...
The client load as it enters the cluster, in terms of request size, read percentage, and concurrency	<code>qos statistics workload characteristics show</code>
Latency across Data ONTAP subsystems, which helps to determine why response time is slow	<code>qos statistics workload latency show</code>
CPU utilization	<code>qos statistics workload resource cpu show</code>
Disk utilization, in terms of the percentage of time spent on the disk during read and write operations	<code>qos statistics workload resource disk show</code>

Histogram-based predictions in RAVE

Starting in Data ONTAP 8.2.1, the speculative read-ahead engine (RAVE) in WAFL can capture additional temporal data about previous user read requests to a file and use this information to intelligently speculate on future read requests. In prior releases, speculation was based only on information from the current user I/O.

For clustered systems, you can enable histogram-based predictions as part of the QoS read-ahead settings, and then attach them to a QoS workload. The `qos settings read-ahead create | modify read_ahead_setting_name -use-histogram true | false` command enables or disables the functionality. The `qos workload modify -read-ahead read_ahead_setting_name -workload workload_name` command attaches the read-ahead setting to any workload.

Increasing WAFL cache memory

You can increase Write Anywhere File Layout (WAFL) cache memory in a system that has a caching module installed (Performance Acceleration Module (PAM), Flash Cache module, or Flash Cache 2 module). To increase the WAFL cache memory, you use the WAFL external cache, a software component of Data ONTAP.

The WAFL external cache provides extra WAFL cache memory to improve the performance of the storage system by reducing the number of disk reads. You can control how user data blocks are cached by changing the mode of operation for a caching module. You can keep the default mode (normal user data blocks) or you can choose metadata mode or low-priority blocks mode.

You should verify that the WAFL external cache functionality is enabled after you install a caching module.

Note: WAFL external cache does not require a separate license if your system is running Data ONTAP 8.1 or later.

Note: Not all systems have a caching module installed. Therefore, not all systems can use the WAFL external cache functionality.

The WAFL external cache does not cache data that is stored in a RAID group composed of SSDs.

If you use WAFL external cache on storage systems on high-availability (HA) configured storage systems, you must ensure that the WAFL external cache options are the same on both nodes. Otherwise, a takeover can result in lower performance due to the lack of WAFL external cache on the remaining node.

In addition to the Data ONTAP options that you can use to manage WAFL external cache, a diagnostic command is available for sanitizing a caching module.

Comparison of Flash Pool aggregates and Flash Cache

Both the Flash Pool technology and the family of Flash Cache modules (Flash Cache and Flash Cache 2) provide a high-performance cache to increase storage performance. However, there are differences between the two technologies that you should understand before choosing between them.

You can employ both technologies on the same system. However, data stored in volumes associated with a Flash Pool aggregate is not cached by Flash Cache.

Criteria	Flash Pool aggregate	Flash Cache
Scope	A specific aggregate	All aggregates assigned to a node
Caching types supported	Read and write	Read

Criteria	Flash Pool aggregate	Flash Cache
Cached data availability during and after takeover events	Cached data is available and unaffected by either planned or unplanned takeover events.	Cached data is not available during takeover events. After giveback for a planned takeover, previously cached data that is still valid is re-cached automatically.
PCIe slot on storage controller required?	No	Yes
Compressed blocks cached?	Yes	No
Supported with array LUNs?	No	Yes

Enabling and disabling WAFL external cache

You can enable or disable the WAFL external cache functionality for a storage system that has a caching module installed (Performance Acceleration Module, Flash Cache module, or Flash Cache 2 module). You should verify that the WAFL external cache functionality is enabled after you install a caching module.

About this task

The `flexscale.enable` option enables or disables the WAFL external cache functionality. If your storage system does not have a caching module installed, the `flexscale.enable` option enables or disables the Predictive Cache Statistics (PCS). PCS is supported on platforms that support caching modules.

WAFL external cache does not require a separate license if your system is running Data ONTAP 8.1 or later. PCS does not require a license.

This command is available through the nodeshell. You access the nodeshell by using the `system node run` command. For more information, see the man page.

Steps

1. To verify whether the WAFL external cache is enabled or disabled, enter the following command:

```
options flexscale.enable
```

2. To enable or disable the WAFL external cache, enter the following command:

```
options flexscale.enable {on|off}
```

Changing the cache size emulated by PCS

If Predictive Cache Statistics (PCS) is enabled on your storage system, you can change the size of cache emulated by PCS to better predict how a caching module would handle your workloads. The size of cache allowed depends on the amount of memory in the controller.

Before you begin

PCS must be enabled.

About this task

The `flexscale.pcs_size` option changes the cache size emulated by PCS. This option is used only when PCS is enabled (setting `flexscale.enable` to `on`), and your storage system does not have a caching module installed.

This command is available through the nodeshell. You access the nodeshell by using the `system node run` command. For more information, see the man page.

Steps

1. Enter the following command to view the current cache size emulated by PCS:

```
options flexscale.pcs_size
```

The default value is chosen automatically based on the amount of memory in the controller.

2. Enter the following command to change the size of cache emulated by PCS:

```
options flexscale.pcs_size integer
```

You can specify integers between 16 and 16383.

Related tasks

[Enabling and disabling WAFL external cache](#) on page 255

Enabling and disabling high resolution sampling for the cache emulated by PCS

If Predictive Cache Statistics (PCS) is enabled on your storage system, you can enable or disable high resolution sampling for the cache emulated by PCS. High resolution sampling might improve measurements of workloads with small hot spots.

Before you begin

PCS must be enabled.

About this task

The `flexscale.pcs_high_res` option enables or disables high resolution sampling for the cache emulated by PCS. This option is used only when PCS is enabled (setting `flexscale.enable` to `on`), and your storage system does not have a caching module installed.

This command is available through the nodeshell. You access the nodeshell by using the `system node run` command. For more information, see the man page.

Step

1. Enter the following command to enable or disable high resolution sampling for the PCS cache:

```
options flexscale.pcs_high_res {on|off}
```

The default value is `off` and should suffice for most workloads. Measurements of workloads with small hot spots might improve when you set the value to `on`.

Related tasks

[Enabling and disabling WAFL external cache](#) on page 255

Caching normal user data blocks

If you cache normal user data blocks, the WAFL external cache interprets this setting as the buffer cache policy of `keep` and saves normal user data blocks in the external cache.

About this task

This command is available through the nodeshell. You access the nodeshell by using the `system node run` command. For more information, see the man page.

Step

1. To enable or disable caching for normal user data blocks, enter the following command:

```
options flexscale.normal_data_blocks {on|off}
```

The default value is **on**.

When the `flexscale.normal_data_blocks` option is set to **on**, the WAFL external cache interprets this setting as the buffer cache policy of **keep** and saves normal user data blocks in the external cache.

If this option is set to **off**, only metadata blocks are cached.

Caching low-priority user data blocks

You can cache low-priority user data blocks that are not normally stored by WAFL external cache. Low-priority blocks include blocks read in large sequential scans that are not normally reused, and blocks that have been written to the storage system through the iSCSI, NFS, or CIFS protocols.

About this task

Caching low-priority user data blocks is useful if you have workloads that fit within WAFL external cache memory and if the workloads consist of either write followed by read or large sequential reads.

You can cache low-priority user data blocks (setting `flexscale.lopri_blocks` to **on**) only if you also cache normal user data blocks (by setting `flexscale.normal_data_blocks` to **on**).

This command is available through the nodeshell. You access the nodeshell by using the `system node run` command. For more information, see the man page.

Step

1. To control whether low-priority user data blocks are cached, enter the following command:

```
options flexscale.lopri_blocks {on|off}
```

The default value is **off**.

Setting the option to **on** caches low-priority user data blocks.

Caching only system metadata

If the working set of the storage system is very large, such as a large e-mail server, you can cache only system metadata in WAFL external cache memory by turning off both normal user data block caching and low-priority user data block caching.

About this task

When you cache only system metadata, with both `flexscale.normal_data_blocks` and `flexscale.lopri_blocks` set to **off**, WAFL external cache interprets this setting as the buffer cache policy of **reuse** and does not save normal data blocks or low-priority blocks in the external cache.

These commands are available through the nodeshell. You access the nodeshell by using the `system node run` command. For more information, see the man page.

Steps

1. Enter the following command to turn off normal user data block caching:

```
options flexscale.normal_data_blocks off
```

2. Enter the following command to turn off low-priority user data block caching:

```
options flexscale.lopri_blocks off
```

Enabling and disabling the caching of blocks from random write workloads

Write blocks are cached in memory until Data ONTAP writes the blocks to disk. After write blocks are written to disk, you can have write blocks that Data ONTAP considers to be random moved from memory to the WAFL extended cache to improve read cache rates for certain workloads.

About this task

Workloads that read data blocks soon after Data ONTAP writes random blocks to disk might benefit from enabling the caching of blocks from random write workloads.

This command is available through the nodeshell. You access the nodeshell by using the `system node run` command. For more information, see the man page.

Step

1. Enter the following command to enable or disable the caching of blocks from random write workloads:

```
options flexscale.random_write_through_blocks {on|off}
```

Enabling and disabling the caching of readahead data

Sequentially read blocks (readahead data) are typically evicted from the buffer cache quickly to make space for randomly read blocks. You can add sequentially read data to the WAFL external cache to improve read response times for sequential data that clients read repetitively.

About this task

This command is available through the nodeshell. You access the nodeshell by using the `system node run` command. For more information, see the man page.

Step

1. Enter the following command to enable or disable the caching of readahead data:

```
options flexscale.readahead_blocks {on|off}
```

Displaying the WAFL external cache configuration

Data ONTAP enables you to display configuration information for WAFL external cache.

About this task

This command is available through the nodeshell. You access the nodeshell by using the `system node run` command. For more information, see the man page.

Step

1. Enter the following command:

```
stats show -p flexscale
```

Displaying usage and access information for WAFL external cache

You can display usage and access information for WAFL external cache, have output produced periodically, and terminate the output after a specified number of iterations.

About this task

This command is available through the nodeshell. You access the nodeshell by using the `system node run` command. For more information, see the man page.

Step

1. Enter the following command:

```
stats show -p flexscale-access [-i interval] [-n num]
```

- If no options are used, a single one-second snapshot of statistics is used.
- `-i interval` specifies that output is to be produced periodically, with an interval of `interval` seconds between each set of output.
- `-n num` terminates the output after `num` number of iterations, when the `-i` option is also used.
If no `num` value is specified, the output runs forever until a user issues a break.
- Press Ctrl-c to interrupt output.

Example

The following example shows sample output from the `stats show -p flexscale-access` command:

Cache Usage	Hit %	Meta /s	Miss /s	Hit %	Evict /s	Inval /s	Insrt /s	Reads Chain /s	Blcks /s	Writes Chain /s	Blcks /s	Disk Replcd /s	Read /s
0	581	0	83	87	0	604	13961	579	581	218	13960	552	
0	777	0	133	85	0	121	21500	773	777	335	21494	744	
0	842	0	81	91	0	1105	23844	837	842	372	23845	812	
0	989	0	122	89	0	0	23175	981	989	362	23175	960	

Example

The following command displays access and usage information for WAFL external cache once every 10 seconds for 5 times:

```
stats show -p flexscale-access -i 10 -n 5
```

Preserving the cache in the Flash Cache family of modules

The system does not serve data from a Flash Cache or Flash Cache 2 module when a node is shutdown. However, the WAFL external cache preserves the cache during a graceful shutdown and can serve "warm" data after giveback.

The WAFL external cache can preserve the cache in Flash Cache modules during a graceful shutdown. It preserves the cache through a process called "cache rewarming," which helps to maintain system performance after a graceful shutdown. For example, you might shut down a system to add hardware or upgrade software.

Cache rewarming is enabled by default if you have a Flash Cache or Flash Cache 2 module installed. Cache rewarming is available when both nodes in an HA pair are running Data ONTAP 8.1 or later.

Related concepts

[Increasing WAFL cache memory](#) on page 254

How cache rewarming works

WAFL external cache initiates the cache rewarming process during a reboot or a takeover and giveback. The process keeps the cache in Flash Cache and Flash Cache 2 modules "warm."

When a storage system powers down, the WAFL external cache takes a snapshot of the data in Flash Cache and Flash Cache 2 modules. When the system powers up, it uses the snapshot to rebuild the cache. After the process completes, the system can read data from the cache.

In an HA configuration, cache rewarming is more successful when minimal changes are made to data during takeover and giveback. When you initiate takeover and giveback, the takeover partner maintains a log of data written to the down partner's storage. If there are changes to a large amount of the data that is stored in the cache, then the cache rewarming process has more data to rewarm when the node comes back online. As a result, the cache may require additional warming time.

Note: Cache rewarming does not work if the WAFL external cache functionality is disabled.

Events that initiate cache rewarming

You can initiate cache rewarming when you shut down a node or when you initiate takeover and giveback.

The following commands initiate cache rewarming:

- `system node halt`
- `storage failover takeover ([-ofnode] | [-bynode]) node -option takeover_option`
- `cf takeover [-node]`
- `cf takeover [-f]`

Events that do not initiate cache rewarming

WAFL external cache does not initiate cache rewarming if the storage system crashes, if there is a sudden loss of power, or if you run certain commands.

The following commands do not initiate cache rewarming:

- `system node halt -dump`
- `system node reboot -dump`
- `cf forcetakeover [-f]`

Cache rewarming abort events and error messages

After the cache rewarming process starts, some events can abort the entire process and some events can abort the process on specific aggregates. Cache rewarming errors and important events are reported entirely through EMS messages.

Events that abort the cache rewarming process

- You add, remove, or move a Flash Cache or Flash Cache 2 module after the WAFL external cache creates the Snapshot copy, but before it rebuilds the cache.
- The takeover node crashes.
- The local node crashes as the WAFL external cache rebuilds the cache.

- After a node reboots, it shuts down before the WAFL external cache can rebuild the cache.
- You initiate a SnapRestore operation on the node's root aggregate before the WAFL external cache rebuilds the cache.
- The `wafliiron` process mounts the root aggregate.

Events that abort cache rewarming on the affected aggregate

- You initiate a SnapRestore operation on an aggregate before the WAFL external cache rebuilds the cache.
- An aggregate does not come online within 20 minutes after the WAFL external cache starts to rebuild the cache.
- The `wafliiron` process mounts the aggregate.

EMS messages that indicate cache rewarming failures

Specific EMS messages indicate exactly what occurred during the cache rewarming process, but the following are some examples of EMS messages that you might see:

- `extCache.rw.snap.canceled`
- `extCache.rw.snap.canceled`
- `extCache.rw.replay.canceled`
- `extCache.rw.replay.timeout`
- `extCache.rw.terminated`
- `extCache.rw.configChange`
- `extCache.rw.timeout`

These EMS messages are informational to explain why the rewarm process did not complete, or to help technical support interpret and explain the events. Usually you do not have to do anything about them, but, in the event there is a corrective action, you can find it in the EMS log.

Enabling and disabling cache rewarming

Cache "rewarming" is enabled by default if a Flash Cache or Flash Cache 2 module is installed. You can disable and then re-enable cache rewarming, if necessary. You should do this only under the guidance of technical support.

Before you begin

You can enable cache rewarming if the following is true:

- A Flash Cache or Flash Cache 2 module is installed.
- The WAFL external cache functionality is enabled.

About this task

Cache rewarming works at the node level. To ensure that cache rewarming works during a takeover and giveback, enable it on all nodes.

These commands are available through the nodeshell. You access the nodeshell by using the `system node run` command. For more information, see the man page.

Step

1. Enter one of the following commands:

If you want to...	Use this command:
Disable cache rewarming	<code>options flexscale.rewarm off</code>
Enable cache rewarming	<code>options flexscale.rewarm on</code>

Related tasks

[Enabling and disabling WAFL external cache](#) on page 255

Improving read performance

You can improve the read performance of your storage system by enabling read reallocation on volumes. Read reallocation is disabled by default.

What read reallocation is

For workloads that perform a mixture of random writes and large and multiple sequential reads, read reallocation improves file layout and sequential read performance. You can enable read reallocation on FlexVol volumes and Infinite Volumes.

Read reallocation analyzes the parts of the file that are read sequentially. If the associated blocks are not already largely contiguous, Data ONTAP updates the layout by rewriting those blocks to another location on disk. The rewrite improves the layout, thus improving the sequential read performance the next time that section of the file is read. However, read reallocation might result in a higher load on the storage system.

Read reallocation is not supported on compressed volumes and FlexCache volumes.

Commands for managing read reallocation

Use the `volume modify` and `volume show` commands to manage read reallocation.

If you want to...	Use this command...
Enable read reallocation on a volume	<code>volume modify</code> with the <code>-read-realloc</code> parameter set to <code>on</code> or <code>space-optimized</code> Note: <code>space-optimized</code> conserves space if you have Snapshot copies, but it can result in degraded read performance of Snapshot copies. <code>space-optimized</code> also rearranges the shared blocks in a deduplicated volume, where as <code>on</code> does not.
Disable read reallocation on a volume	<code>volume modify</code> with the <code>-read-realloc</code> parameter set to <code>off</code>
Identify whether read reallocation is enabled or disabled on volumes	<code>volume show -fields read-realloc</code>

For more information, see the man pages.

Improving write performance

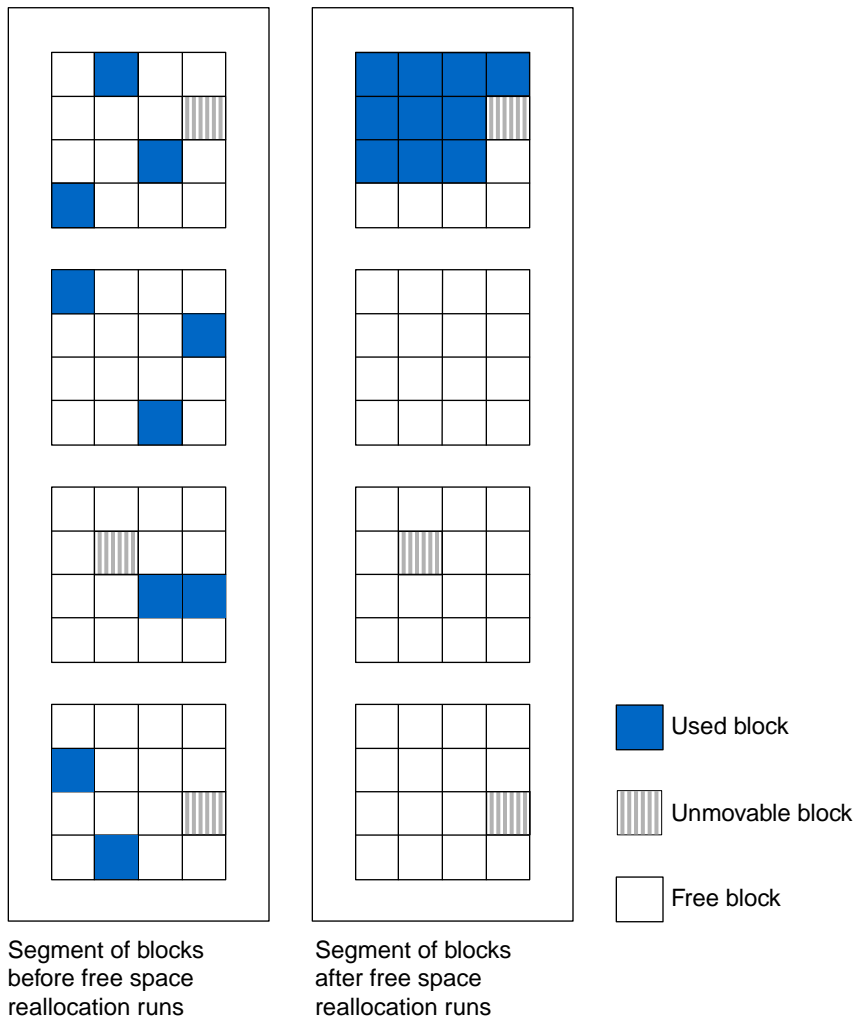
You can enable free space reallocation on aggregates to improve write performance. Free space reallocation improves write performance by optimizing the free space within an aggregate. Free space reallocation is disabled by default.

How free space reallocation optimizes free space

Free space reallocation optimizes the free space in an aggregate immediately before Data ONTAP writes data to the blocks in that aggregate.

Before Data ONTAP writes data to a segment of blocks in an aggregate, free space reallocation evaluates the layout of those blocks. If the layout is not optimal, the free space reallocation function rearranges the blocks. Rearranging the blocks increases the amount of contiguous free space available in the aggregate, which improves the performance of Data ONTAP writes to those blocks.

The following graphic illustrates how free space reallocation optimizes the free space in a segment of blocks:



When to enable free space reallocation

Free space reallocation works best on workloads that perform a mixture of small random overwrites and sequential or random reads. You can expect additional CPU utilization when you enable free

space reallocation. You should not enable free space reallocation if your storage system has sustained, high CPU utilization.

Note: You can use the `statistics show-periodic` command to monitor CPU utilization.

For best results, you should enable free space reallocation when you create a new aggregate. If you enable free space reallocation on an existing aggregate, there might be a period where Data ONTAP performs additional work to optimize free space. This additional work can temporarily impact system performance.

When to use free space reallocation with other reallocation features

When you enable free space reallocation, you should also consider enabling read reallocation. Free space reallocation and read reallocation are complementary technologies that optimize data layout. Read reallocation optimizes the system for sequential reads, while free space reallocation optimizes for writes.

Related concepts

What read reallocation is on page 262

Types of aggregates that free space reallocation can and cannot optimize

Free space reallocation optimizes the free space in specific types of aggregates.

Free space reallocation optimizes free space in the following:

- Aggregates that provide storage to FlexVol volumes or Infinite Volumes
- The HDD RAID groups in an aggregate

Free space reallocation does not optimize free space in the following:

- The SSD RAID groups in an aggregate
- Read-only volumes such as load-sharing or data protection mirrors

Commands for managing free space reallocation

Use the `storage aggregate modify` and `storage aggregate show` commands to manage free space reallocation.

If you want to...	Use this command...
Enable free space reallocation on an aggregate	<code>storage aggregate modify</code> with the <code>-free-space-realloc</code> parameter set to on
Disable free space reallocation on an aggregate	<code>storage aggregate modify</code> with the <code>-free-space-realloc</code> parameter set to off
Identify whether free space reallocation is enabled or disabled on aggregates	<code>storage aggregate show -fields free-space-realloc</code>

For more information, see the man pages.

Managing peer relationships for data backup and recovery (cluster administrators only)

Establishing peer relationships between two clusters or two Storage Virtual Machines (SVMs) enables you to back up and recover the data on the clusters or SVMs.

Managing cluster peer relationships

You can create data protection mirroring relationships from one cluster to another and you can manage the jobs on a remote cluster from another cluster if you have cluster peer relationships.

Related concepts

[Managing SVM peer relationships](#) on page 287

What a cluster peer is

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

Commands for managing cluster peer relationships

There are specific Data ONTAP commands for managing cluster peer relationships.

If you want to...	Use this command...
Create an authenticated cluster peer relationship	<code>cluster peer create</code> This creates an authenticated cluster peer relationship by default. If you want to create an unauthenticated cluster peer relationship, you use this command with the <code>-no-authentication</code> parameter and unauthenticated cluster peer relationships must also be allowed by the cluster peer policy.
Create an authenticated cluster peer relationship with an extended authentication offer	<code>cluster peer create</code> with the <code>-offer-expiration</code> parameter. This is useful if the second cluster in the relationship cannot be authenticated in the default time of one hour.
Create an unauthenticated cluster peer relationship	<code>cluster peer create</code> with the <code>-no-authentication</code> parameter. Unauthenticated cluster peer relationships must also be allowed by the cluster peer policy.
Create an authenticated cluster peer relationship using a specific IPspace	<code>cluster peer create</code> with the <code>ipspace</code> parameter.
Delete a cluster peer relationship	<code>cluster peer delete</code>
Reverting a cluster peer relationship's IPspace designation back to the Default IPspace	<code>cluster peer modify</code> with the <code>ipspace</code> parameter set to Default .
Modify a cluster peer relationship	<code>cluster peer modify</code>

If you want to...	Use this command...
Change an IPspace for a cluster in a cluster peer relationship	<p><i>cluster peer modify</i> with the <code>ipspace</code> parameter.</p> <p>This is useful if you have an existing cluster peer relationship and you want to use intercluster LIFs in a designated IPspace.</p>
Initiate an intercluster connectivity test	<i>cluster peer ping</i>
Display information about the cluster peer relationship	<i>cluster peer show</i>
Display TCP connection information for a cluster peer	<i>cluster peer connection show</i>
Display health information of the nodes in a cluster peer relationship from the local cluster perspective	<i>cluster peer health show</i>
Display information about outstanding authentication offers to peer clusters	<i>cluster peer offer show</i>
Disable an existing cluster peer relationship	<i>cluster peer modify</i> with the <code>-auth-status-admin</code> parameter set to revoked .
Update a cluster peer relationship to use authentication with a different passphrase	<p><i>cluster peer modify</i> with the <code>-auth-status-admin</code> parameter set to use-authentication.</p> <p>You can use this command to perform one of the following tasks:</p> <ul style="list-style-type: none"> • Add authentication to an existing unauthenticated cluster peer relationship • Add authentication to a revoked cluster peer relationship • Change a passphrase for an authenticated cluster peer relationship
Reestablish a disabled cluster peer relationship with no-authentication	<i>cluster peer modify</i> with the <code>-auth-status-admin</code> parameter set to no-authentication .
Modify an outstanding authentication offer to a peer cluster	<p><i>cluster peer offer modify</i> with the <code>-offer-expiration</code> parameter.</p> <p>You can change when the authentication offer expires if you determine that there is not enough time to authenticate the cluster peer relationship before the offer expires.</p>
Cancel an outstanding authentication offer to a peer cluster	<i>cluster peer offer cancel</i>
Display whether unauthenticated cluster peer relationships can exist and what the minimum passphrase character length is	<i>cluster peer policy show</i>

If you want to...	Use this command...
Modify whether unauthenticated cluster peer relationships can exist and change the minimum passphrase character length	<code>cluster peer policy modify</code>

Related information

[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Cluster peer security using authentication passphrases

When creating a cluster peer relationship, a passphrase is used by the administrators of the two clusters to authenticate the relationship. This ensures that the cluster to which you send data is the cluster to which you intend to send data.

A part of the cluster peer creation process is to use a passphrase to authenticate the cluster peers to each other. The passphrase is used when creating the relationship from the first cluster to the second and, again, when creating the relationship from the second cluster to the first. The passphrase is not exchanged on the network by Data ONTAP, but each cluster in the cluster peer relationship recognizes the passphrase when Data ONTAP creates the cluster peer relationship.

When you create the cluster peer relationship from the first cluster to the second, the first cluster waits for the administrator of the second cluster to create the cluster peer relationship. The administrator of the second cluster must create the cluster peer relationship before the waiting period expires, one hour by default, but can be shortened. If the cluster peer relationship is not created from the second cluster to the first before the waiting period expires, the cluster peer relationship is not created and the administrators must start again.

Related tasks

[Creating the cluster peer relationship](#) on page 284

Related references

[Commands for managing cluster peer relationships](#) on page 265

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.

Steps

1. [What cluster peer intercluster connectivity is](#) on page 268
2. [Supported cluster peer network topologies](#) on page 268
3. [Comparison of designated and undesignated intercluster connectivity](#) on page 271
4. [Prerequisites for cluster peering](#) on page 272
5. [Considerations when sharing data ports](#) on page 273
6. [Considerations when using dedicated ports](#) on page 273
7. [Configuring intercluster LIFs to share data ports](#) on page 274
8. [Configuring intercluster LIFs to use dedicated intercluster ports](#) on page 276
9. [Configuring intercluster LIFs to use intercluster ports in their own networks](#) on page 279

10. [Creating the cluster peer relationship](#) on page 284

What cluster peer intercluster connectivity is

You should know about the interfaces and ports that you put together to create a cluster peer intercluster connection and how they are used. Knowing this information might reduce the amount of time you use to create the cluster peer intercluster connectivity.

Cluster peer intercluster connectivity consists of intercluster logical interfaces (LIFs) that are assigned to network ports. The intercluster connection on which replication occurs between two different clusters is defined when the intercluster LIFs are created. Replication between two clusters can occur on the intercluster connection only; this is true regardless of whether the intercluster connectivity is on the same subnet as a data network in the same cluster.

The IP addresses assigned to intercluster LIFs can reside in the same subnet as data LIFs or in a different subnet. When an intercluster LIF is created, it uses routes that belong to the System SVM that the intercluster LIF is in.

Related concepts

[Considerations when sharing data ports](#) on page 273

[Considerations when using dedicated ports](#) on page 273

Supported cluster peer network topologies

To provide data protection, all of the intercluster LIFs of one cluster must be able to communicate with all of the intercluster LIFs of the cluster peer using *pair-wise full-mesh* connectivity. You need to understand how this connectivity works for different cluster topologies.


Pair-wise full-mesh connectivity applies only to the two clusters in the peer relationship. All of the intercluster LIFs of an IPspace in one cluster must be able to communicate with all of the intercluster LIFs of an IPspace in the other cluster.

Using the concept of *pair-wise full-mesh* connectivity helps you to build more complex cluster peer topologies. Understanding how this connectivity works for two cluster, cluster cascade, and cluster fan-out or fan-in topologies will help you to create viable intercluster networks without adding intercluster networks that are unnecessary.

Intercluster networking between two clusters

Creating an intercluster network between two clusters is the basic cluster peer configuration. For example, you want to create an intercluster network between two clusters, Cluster A and Cluster B. Cluster A has two intercluster LIFs, A1 and A2, in its Default IPspace and Cluster B has two intercluster LIFs, B1 and B2, in its Default IPspace. The LIFs are connected as follows:

- A1 communicates with B1 and
- A1 communicates with B2 and
- A2 communicates with B1 and
- A2 communicates with B2



Intercluster networking in a cluster cascade


When you connect three clusters in a cascade, all of the intercluster LIFs of the primary cluster must be able to communicate with all of the intercluster LIFs of the secondary cluster. Likewise, all of the intercluster LIFs of the secondary cluster must be able to communicate with all of the intercluster LIFs of the tertiary cluster. You do not need to create an intercluster network between the primary cluster and the tertiary cluster if you do not want to connect the two clusters in a cluster peer relationship.

For example, you want to create an intercluster network between Cluster A and Cluster B and an intercluster network between Cluster B and Cluster C. Cluster A has two intercluster LIFs, A1 and A2, in its Default IPspace, Cluster B has two intercluster LIFs, B1 and B2, in its Default IPspace, and Cluster C has two intercluster LIFs, C1 and C2, in its Default IPspace. The intercluster LIFs between Cluster A and Cluster B are connected as follows:

- A1 communicates with B1 and
- A1 communicates with B2 and
- A2 communicates with B1 and
- A2 communicates with B2 and


The intercluster LIFs between Cluster B and Cluster C are connected as follows:

- B1 communicates with C1 and
- B1 communicates with C2 and
- B2 communicates with C1 and
- B2 communicates with C2



You might have a cluster cascade configured in which you want the tertiary cluster to connect to the primary cluster if something happens to the secondary cluster. An example is if you have a disaster recovery relationship between the primary cluster and the secondary cluster, and a backup relationship between the secondary cluster and the tertiary cluster, and you want the tertiary cluster to communicate with the primary cluster if something happens to the secondary cluster. If this configuration is what you want, then the intercluster LIFs of the tertiary cluster must be able to communicate with all of the intercluster LIFs of the primary cluster. Therefore, in addition to the connections previously mentioned, you would also have the following intercluster LIF connections between Cluster C and Cluster A:

- A1 communicates with C1 and
- A1 communicates with C2 and
- A2 communicates with C1 and
- A2 communicates with C2



Intercluster networking in a cluster fan-out or fan-in

When you connect clusters in a fan-out or fan-in configuration, the intercluster LIFs of each cluster that connects to the primary cluster must be able to communicate with all of the intercluster LIFs of the primary cluster. You do not need to connect intercluster LIFs between the remote clusters if the remote clusters do not need to communicate with each other.


For example, you want to create an intercluster network between Cluster A and Cluster B and an intercluster network between Cluster A and Cluster C. Cluster A has two intercluster LIFs, A1 and A2, in its Default IPspace, Cluster B has two intercluster LIFs, B1 and B2, in its Default IPspace, and Cluster C has two intercluster LIFs, C1 and C2, in its Default IPspace. The intercluster LIFs between Cluster A and Cluster B are connected as follows:

- A1 communicates with B1 and
- A1 communicates with B2 and
- A2 communicates with B1 and
- A2 communicates with B2

The intercluster LIFs between Cluster A and Cluster C are connected as follows:

- A1 communicates with C1 and
- A1 communicates with C2 and
- A2 communicates with C1 and
- A2 communicates with C2

Cluster B is not connected to Cluster C.



If you do want a cluster peer relationship between two remote clusters in addition to the fan-in or fan-out configuration, then use the concept of *pair-wise full-mesh* connectivity to create an intercluster network between them.

Comparison of designated and undesigned intercluster connectivity

Unlike undesigned intercluster connectivity between clusters in their respective default IPspaces, designated intercluster connectivity uses a non-default IPspace to contain the interactions that a cluster has with its peer. Understanding how connectivity differs helps you decide if you want to use designated intercluster connectivity.


Undesignated intercluster connectivity operates within the default IPspaces of the clusters. Because of the pair-wise full-mesh connectivity requirement within an IPspace, all of the intercluster LIFs of the two clusters in the cluster peer relationship must be able to connect to each other. This connectivity requirement does not allow separation of connectivity that storage service providers might need. You would have to introduce hardware, such as a router, into the network to separate intercluster connectivity between clusters.

Designated intercluster connectivity operates within a specified non-default IPspace on at least one of the clusters in the peer relationship. The requirement of pair-wise full-mesh connectivity still exists, but the connectivity is within the IPspace that each cluster defines for the peering relationship and not across IPspaces. This keeps the intercluster connectivity isolated from different IPspaces that a cluster might have with other peers and also reduces the scope of the full-mesh connectivity requirement. In this way, a storage service provider can control the intercluster connectivity separation that might be needed.

As an example, you want Cluster A to have a peering relationship with two other clusters, Cluster B and Cluster C. Additionally, you want the connectivity between Cluster A and Cluster B to be separate from the connectivity between Cluster A and Cluster C. To do this, you can create two IPspaces on Cluster A. The first IPspace, called `ipspaceAB`, will contain the intercluster LIFs of Cluster A that you want to use to communicate with Cluster B. Likewise, the second IPspace, called `ipspaceAC`, will contain the intercluster LIFs of Cluster A that you want to use to communicate with Cluster C.

The intercluster LIFs on Cluster B and Cluster C are contained in the default IPspaces of their respective clusters. It is not a requirement that these intercluster LIFs be in the default IPspace. These intercluster LIFs can be in the default IPspace or in their own designated IPspaces.

For the peering relationship between Cluster A, Cluster B, and Cluster C, Cluster A would use only those LIFs in IPspace `ipspaceAB` to communicate with Cluster B, and only those LIFs in IPspace `ipspaceAC` to communicate with Cluster C. There is no requirement that intercluster LIFs in Cluster B communicate with any other intercluster LIFs on Cluster A other than those in IPspace `ipspaceAB`, and no requirement that intercluster LIFs in Cluster C communicate with any other intercluster LIFs on Cluster A other than those in IPspace `ipspaceAC`.



Prerequisites for cluster peering

Before you set up cluster peering, you should confirm that the IPspace, connectivity, port, IP address, subnet, firewall, and cluster-naming requirements are met.

Connectivity requirements

The subnet used in each cluster for intercluster communication must meet the following requirements:

- The subnet must belong to the broadcast domain that contains the ports used for intercluster communication.
- IP addresses used for intercluster LIFs do not need to be in the same subnet, but having them in the same subnet is a simpler configuration.
- You must have considered whether the subnet will be dedicated to intercluster communication or shared with data communication.
- The subnet must have enough IP addresses available to allocate to one intercluster LIF per node. For example, in a six-node cluster, the subnet used for intercluster communication must have six available IP addresses.

The intercluster network must be configured so that cluster peers have *pair-wise full-mesh connectivity* within the applicable IPspace, which means that each pair of clusters in a cluster peer relationship has connectivity among all of their intercluster LIFs.

A cluster's intercluster LIFs must use the same IP address version: all IPv4 addresses or all IPv6 addresses. Similarly, all of the intercluster LIFs of the peered clusters must use the same IP addressing version.

Port requirements

The ports that will be used for intercluster communication must meet the following requirements:

- The broadcast domain that is used for intercluster communication must include at least two ports per node so that intercluster communication can fail over from one port to another. The ports added to a broadcast domain can be physical network ports, VLANs, or interface groups (ifgrps).
- All of the ports must be cabled.
- All of the ports must be in a healthy state.
- The MTU settings of the ports must be consistent.
- You must have considered whether the ports used for intercluster communication will be shared with data communication.

Firewall requirements

Firewalls and the intercluster firewall policy must allow the following:

- ICMP service
- TCP to the IP addresses of all of the intercluster LIFs over all of the following ports: 10000, 11104, and 11105
- HTTPS

Although HTTPS is not required when you set up cluster peering, HTTPS is required later if you use OnCommand System Manager to configure data protection. However, if you use the

command-line interface to configure data protection, HTTPS is not required to configure cluster peering or data protection.

The default **intercluster** firewall policy allows access through the HTTPS protocol and from all IP addresses (0.0.0.0/0), but the policy can be altered or replaced.

Cluster requirements

Clusters must meet the following requirements:

- Each cluster must have a unique name.
You cannot create a cluster peering relationship with any cluster that has the same name or is in a peer relationship with a cluster of the same name.
- The time on the clusters in a cluster peering relationship must be synchronized within 300 seconds (5 minutes).
Cluster peers can be in different time zones.

Considerations when sharing data ports

When determining whether sharing a data port for intercluster replication is the correct intercluster network solution, you should consider configurations and requirements such as LAN type, available WAN bandwidth, replication interval, change rate, and number of ports.

Consider the following aspects of your network to determine whether sharing data ports is the best intercluster connectivity solution:

- For a high-speed network, such as a 10-Gigabit Ethernet (10-GbE) network, a sufficient amount of local LAN bandwidth might be available to perform replication on the same 10-GbE ports that are used for data access.
In many cases, the available WAN bandwidth is far less than 10 GbE LAN bandwidth .
- All nodes in the cluster might have to replicate data and share the available WAN bandwidth, making data port sharing more acceptable.
- Sharing ports for data and replication eliminates the extra port counts required to dedicate ports for replication.
- The maximum transmission unit (MTU) size of the replication network will be the same size as that used on the data network.
- Consider the data change rate and replication interval and whether the amount of data that must be replicated on each interval requires enough bandwidth that it might cause contention with data protocols if sharing data ports.
- When data ports for intercluster replication are shared, the intercluster LIFs can be migrated to any other intercluster-capable port on the same node to control the specific data port that is used for replication.

Considerations when using dedicated ports

When determining whether using a dedicated port for intercluster replication is the correct intercluster network solution, you should consider configurations and requirements such as LAN type, available WAN bandwidth, replication interval, change rate, and number of ports.

Consider the following aspects of your network to determine whether using a dedicated port is the best intercluster network solution:

- If the amount of available WAN bandwidth is similar to that of the LAN ports and the replication interval is such that replication occurs while regular client activity exists, then you should dedicate Ethernet ports for intercluster replication to avoid contention between replication and the data protocols.

- If the network utilization generated by the data protocols (CIFS, NFS, and iSCSI) is such that the network utilization is above 50 percent, then you should dedicate ports for replication to allow for nondegraded performance if a node failover occurs.
- When physical 10 GbE ports are used for data and replication, you can create VLAN ports for replication and dedicate the logical ports for intercluster replication.
The bandwidth of the port is shared between all VLANs and the base port.
- Consider the data change rate and replication interval and whether the amount of data that must be replicated on each interval requires enough bandwidth that it might cause contention with data protocols if sharing data ports.

Configuring intercluster LIFs to share data ports

Configuring intercluster LIFs to share data ports enables you to use existing data ports to create intercluster networks for cluster peer relationships. Sharing data ports reduces the number of ports you might need for intercluster networking.

About this task

Creating intercluster LIFs that share data ports involves assigning LIFs to existing data ports. In this procedure, a two-node cluster exists in which each node has two data ports, e0c and e0d, and these data ports are in the default IPspace. These are the two data ports that are shared for intercluster replication. You must configure intercluster LIFs on the peer cluster before you can create cluster peer relationships. In your own environment, you replace the ports, networks, IP addresses, subnet masks, and subnets with those specific to your environment.

Steps

1. List the ports in the cluster by using the `network port show` command:

Example

```
cluster01::> network port show
```

Node	Port	IPspace	Broadcast Domain	Link	MTU	Speed (Mbps) Admin/Oper

cluster01-01						
	e0a	Cluster	Cluster	up	1500	auto/1000
	e0b	Cluster	Cluster	up	1500	auto/1000
	e0c	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
cluster01-02						
	e0a	Cluster	Cluster	up	1500	auto/1000
	e0b	Cluster	Cluster	up	1500	auto/1000
	e0c	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000

2. Create an intercluster LIF on the admin SVM cluster01 by using the `network interface create` command.

Example

This example uses the LIF naming convention `adminSVMname_icl#` for the intercluster LIF:

```
cluster01::> network interface create -vserver cluster01 -lif cluster01_icl01 -role
intercluster
-home-node cluster01-01 -home-port e0c -address 192.168.1.201 -netmask 255.255.255.0

cluster01::> network interface create -vserver cluster01 -lif cluster01_icl02 -role
intercluster
-home-node cluster01-02 -home-port e0c -address 192.168.1.202 -netmask 255.255.255.0
```

- Verify that the intercluster LIFs were created properly by using the `network interface show` command with the `-role intercluster` parameter:

Example

```
cluster01::> network interface show -role intercluster
Vserver Logical Status Network Current Current Is
Interface Admin/Oper Address/Mask Node Port Home
-----
cluster01
cluster01_icl01
up/up 192.168.1.201/24 cluster01-01 e0c true
cluster01_icl02
up/up 192.168.1.202/24 cluster01-02 e0c true
```

- Verify that the intercluster LIFs are configured to be redundant by using the `network interface show` command with the `-role intercluster` and `-failover` parameters.

Example

The LIFs in this example are assigned the e0c port on each node. If the e0c port fails, the LIF can fail over to the e0d port.

```
cluster01::> network interface show -role intercluster -failover
Vserver Logical Home Failover Failover
Interface Node:Port Policy Group
-----
cluster01
cluster01_icl01 cluster01-01:e0c local-only 192.168.1.201/24
Failover Targets: cluster01-01:e0c,
cluster01-01:e0d
cluster01_icl02 cluster01-02:e0c local-only 192.168.1.201/24
Failover Targets: cluster01-02:e0c,
cluster01-02:e0d
```

- Display the routes in the cluster by using the `network route show` command to determine whether intercluster routes are available or you must create them.

Creating a route is required only if the intercluster addresses in both clusters are not on the same subnet and a specific route is needed for communication between the clusters.

Example

In this example, no intercluster routes are available:

```
cluster01::> network route show
Vserver Destination Gateway Metric
-----
Cluster
0.0.0.0/0 192.168.0.1 20
cluster01
0.0.0.0/0 192.168.0.1 10
```

- If communication between intercluster LIFs in different clusters requires routing, create an intercluster route by using the `network route create` command.

The gateway of the new route should be on the same subnet as the intercluster LIF.

Example

In this example, 192.168.1.1 is the gateway address for the 192.168.1.0/24 network. If the destination is specified as 0.0.0.0/0, then it becomes a default route for the intercluster network.

```
cluster01::> network route create -vserver cluster01
-destination 0.0.0.0/0 -gateway 192.168.1.1 -metric 40
```

- Verify that you created the routes correctly by using the `network route show` command.

Example

```
cluster01::> network route show
Vserver  Destination      Gateway          Metric
-----  -
Cluster
cluster01  0.0.0.0/0        192.168.0.1     20
           0.0.0.0/0        192.168.0.1     10
           0.0.0.0/0        192.168.1.1     40
```

- Repeat these steps on the cluster to which you want to connect.

Configuring intercluster LIFs to use dedicated intercluster ports

Configuring intercluster LIFs to use dedicated data ports allows greater bandwidth than using shared data ports on your intercluster networks for cluster peer relationships.

About this task

Creating intercluster LIFs that use dedicated ports involves creating a failover group for the dedicated ports and assigning LIFs to those ports. In this procedure, a two-node cluster exists in which each node has two data ports that you have added, e0e and e0f. These ports are ones you will dedicate for intercluster replication and currently are in the default IPspace. These ports will be grouped together as targets for the intercluster LIFs you are configuring. You must configure intercluster LIFs on the peer cluster before you can create cluster peer relationships. In your own environment, you would replace the ports, networks, IP addresses, subnet masks, and subnets with those specific to your environment.

Steps

- List the ports in the cluster by using `network port show` command.

Example

```
cluster01::> network port show
Node  Port      IPspace      Broadcast Domain Link  MTU      Speed (Mbps)
-----  -
cluster01-01
e0a    Cluster   Cluster      Cluster      up    1500    auto/1000
e0b    Cluster   Cluster      Cluster      up    1500    auto/1000
e0c    Default   Default      Default      up    1500    auto/1000
e0d    Default   Default      Default      up    1500    auto/1000
e0e    Default   Default      Default      up    1500    auto/1000
e0f    Default   Default      Default      up    1500    auto/1000
cluster01-02
e0a    Cluster   Cluster      Cluster      up    1500    auto/1000
e0b    Cluster   Cluster      Cluster      up    1500    auto/1000
e0c    Default   Default      Default      up    1500    auto/1000
e0d    Default   Default      Default      up    1500    auto/1000
e0e    Default   Default      Default      up    1500    auto/1000
e0f    Default   Default      Default      up    1500    auto/1000
```

- Determine whether any of the LIFs are using ports that are dedicated for replication by using the `network interface show` command.

Example

Ports e0e and e0f do not appear in the following output; therefore, they do not have any LIFs located on them:

```
cluster01::> network interface show -fields home-port,curr-port
vserver lif      home-port  curr-port
-----  -
Cluster cluster01-01_clus1  e0a       e0a
Cluster cluster01-01_clus2  e0b       e0b
Cluster cluster01-02_clus1  e0a       e0a
Cluster cluster01-02_clus2  e0b       e0b
```

```

cluster01
  cluster_mgmt          e0c      e0c
cluster01
  cluster01-01_mgmt1   e0c      e0c
cluster01
  cluster01-02_mgmt1   e0c      e0c
    
```

3. If a LIF is using a port that you want dedicated to intercluster connectivity, migrate the LIF to a different port.
 - a. Migrate the LIF to another port by using the `network interface migrate` command.

Example

The following example assumes that the data LIF named `cluster01_data01` uses port `e0e` and you want only an intercluster LIF to use that port:

```

cluster01::> network interface migrate -vserver cluster01
-lif cluster01_data01 -dest-node cluster01-01 -dest-port e0d
    
```

- b. You might need to modify the migrated LIF home port to reflect the new port where the LIF should reside by using the `network interface modify` command:

Example

```

cluster01::> network interface modify -vserver cluster01
-lif cluster01_data01 -home-node cluster01-01 -home-port e0d
    
```

4. Group the ports that you will use for the intercluster LIFs by using the `network interface failover-groups create` command.

Example

```

cluster01::> network interface failover-groups create -vserver cluster01
-failover-group intercluster01 -targets cluster01-01:e0e,cluster01-01:e0f,
cluster01-02:e0e,cluster01-02:e0f
    
```

5. Display the failover-group that you created by using the `network interface failover-groups show` command.

Example

```

cluster01::> network interface failover-groups show
Vserver      Group      Failover
-----
Cluster
cluster01    Cluster    cluster01-01:e0a, cluster01-01:e0b,
cluster01-02:e0a, cluster01-02:e0b
cluster01    Default    cluster01-01:e0c, cluster01-01:e0d,
cluster01-02:e0c, cluster01-02:e0d,
cluster01-01:e0e, cluster01-01:e0f
cluster01-02:e0e, cluster01-02:e0f
cluster01    intercluster01    cluster01-01:e0e, cluster01-01:e0f
cluster01-02:e0e, cluster01-02:e0f
    
```

6. Create an intercluster LIF on the admin SVM `cluster01` by using the `network interface create` command.

Example

This example uses the LIF naming convention `adminSVMname_ic1#` for the intercluster LIF:

```
cluster01::> network interface create -vserver cluster01 -lif cluster01_icl01 -role
intercluster -home-node cluster01-01 -home-port e0e
-address 192.168.1.201 -netmask 255.255.255.0 -failover-group intercluster01

cluster01::> network interface create -vserver cluster01 -lif cluster01_icl02 -role
intercluster -home-node cluster01-02 -home-port e0e
-address 192.168.1.202 -netmask 255.255.255.0 -failover-group intercluster01
```

- Verify that the intercluster LIFs were created properly by using the `network interface show` command.

Example

```
cluster01::> network interface show
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
Cluster	cluster01-01_clus_1	up/up	192.168.0.xxx/24	cluster01-01	e0a	true
	cluster01-01_clus_2	up/up	192.168.0.xxx/24	cluster01-01	e0b	true
	cluster01-02_clus_1	up/up	192.168.0.xxx/24	cluster01-01	e0a	true
	cluster01-02_clus_2	up/up	192.168.0.xxx/24	cluster01-01	e0b	true
cluster01	cluster_mgmt	up/up	192.168.0.xxx/24	cluster01-01	e0c	true
	cluster01_icl01	up/up	192.168.1.201/24	cluster01-01	e0e	true
	cluster01_icl02	up/up	192.168.1.202/24	cluster01-02	e0e	true
	cluster01-01_mgmt1	up/up	192.168.0.xxx/24	cluster01-01	e0c	true
	cluster01-02_mgmt1	up/up	192.168.0.xxx/24	cluster01-02	e0c	true

- Verify that the intercluster LIFs are configured for redundancy by using the `network interface show` command with the `-role intercluster` and `-failover` parameters.

Example

The LIFs in this example are assigned the e0e home port on each node. If the e0e port fails, the LIF can fail over to the e0f port.

```
cluster01::> network interface show -role intercluster -failover
```

Vserver	Logical Interface	Home Node:Port	Failover Policy	Failover Group
cluster01-01	cluster01-01_icl01	cluster01-01:e0e	local-only	intercluster01
		Failover Targets:	cluster01-01:e0e,	cluster01-01:e0f
	cluster01-01_icl02	cluster01-02:e0e	local-only	intercluster01
		Failover Targets:	cluster01-02:e0e,	cluster01-02:e0f

- Display the routes in the cluster by using the `network route show` command to determine whether intercluster routes are available or you must create them.

Creating a route is required only if the intercluster addresses in both clusters are not on the same subnet and a specific route is needed for communication between the clusters.

Example

In this example, no intercluster routes are available:

```
cluster01::> network route show
```

Vserver	Destination	Gateway	Metric
Cluster	0.0.0.0/0	192.168.0.1	20
cluster01	0.0.0.0/0	192.168.0.1	10

10. If communication between intercluster LIFs in different clusters requires routing, create an intercluster route by using the `network route create` command.

The gateway of the new route should be on the same subnet as the intercluster LIF.

Example

In this example, 192.168.1.1 is the gateway address for the 192.168.1.0/24 network. If the destination is specified as 0.0.0.0/0, then it becomes a default route for the intercluster network.

```
cluster01::> network route create -vserver cluster01
-destination 0.0.0.0/0 -gateway 192.168.1.1 -metric 40
```

11. Verify that you created the routes correctly by using the `network route show` command.

Example

```
cluster01::> network route show
Vserver Destination Gateway Metric
-----
Cluster
cluster01 0.0.0.0/0 192.168.0.1 20
           0.0.0.0/0 192.168.0.1 10
           0.0.0.0/0 192.168.1.1 40
```

12. Repeat these steps to configure intercluster networking in the peer cluster.
13. Verify that the ports have access to the proper subnets, VLANs, and so on.
Dedicating ports for replication in one cluster does not require dedicating ports in all clusters; one cluster might use dedicated ports, while the other cluster shares data ports for intercluster replication.

Configuring intercluster LIFs to use intercluster ports in their own networks

You might need to direct intercluster traffic over a designated network. For example, you might want to connect to different clusters that are not reachable in the default IPspace . You can do this by moving ports to their own IPspaces and configuring intercluster LIFs.

About this task

In this procedure, a two-node cluster exists in which each node has two ports that you want to use for cluster peer relationships: e0e and e0f. These ports are ones you move from the default IPspace to their own IPspace. In the examples, these ports are configured with intercluster LIFs only, but you could configure ports to share data LIFs as well. In your own environment, you would replace the ports, networks, IP addresses, subnet masks, and subnets with those specific to your environment.

Steps

1. List the ports in the cluster by using `network port show` command:

Example

```
cluster01::> network port show
Node Port IPspace Broadcast Domain Link MTU Speed (Mbps)
Admin/Oper
-----
cluster01-01
e0a Cluster Cluster up 1500 auto/1000
e0b Cluster Cluster up 1500 auto/1000
e0c Default Default up 1500 auto/1000
e0d Default Default up 1500 auto/1000
e0e Default Default up 1500 auto/1000
e0f Default Default up 1500 auto/1000
cluster01-02
e0a Cluster Cluster up 1500 auto/1000
```

e0b	Cluster	Cluster	up	1500	auto/1000
e0c	Default	Default	up	1500	auto/1000
e0d	Default	Default	up	1500	auto/1000
e0e	Default	Default	up	1500	auto/1000
e0f	Default	Default	up	1500	auto/1000

2. Create a nondefault IPspace on the cluster on which you want to segregate the intercluster network by using the `network ipspace create` command:

Example

```
cluster01::> network ipspace create -ipspace ipspace-IC1
```

3. Determine whether any of the LIFs are using ports that are dedicated for replication by using the `network interface show` command.

Example

Ports e0e and e0f do not appear in the following output; therefore, they do not have any LIFs located on them:

```
cluster01::> network interface show -fields home-port,curr-port
vserver lif                home-port  curr-port
-----
Cluster cluster01-01_clus1 e0a       e0a
Cluster cluster01-01_clus2 e0b       e0b
Cluster cluster01-02_clus1 e0a       e0a
Cluster cluster01-02_clus2 e0b       e0b
cluster01
  cluster_mgmt             e0c       e0c
cluster01
  cluster01-01_mgmt1       e0c       e0c
cluster01
  cluster01-02_mgmt1       e0c       e0c
```

4. If a LIF is using a port that you want dedicated to intercluster connectivity, migrate the LIF to a different port.
 - a. Migrate the LIF to another port by using the `network interface migrate` command.

Example

The following example assumes that the data LIF named `cluster01_data01` uses port e0e and you want only an intercluster LIF to use that port.

```
cluster01::> network interface migrate -vserver cluster01
-lif cluster01_data01 -dest-node cluster01-01 -dest-port e0d
```

- b. You might need to modify the migrated LIF home port to reflect the new port where the LIF should reside by using the `network interface modify` command.

Example

```
cluster01::> network interface modify -vserver cluster01
-lif cluster01_data01 -home-node cluster01-01 -home-port e0d
```

5. Remove ports e0e and e0f from the default broadcast domain by using the `network port broadcast-domain remove-ports` command.

Ports must be removed from the broadcast domain before being added to another broadcast domain because a port cannot be in more than one broadcast domain at one time.

Example

```
cluster01::> network port broadcast-domain remove-ports -broadcast-domain Default -ports cluster01-01:e0e,cluster01-01:e0f,cluster01-02:e0e,cluster01-02:e0f
```

6. Verify that the ports are unassigned by using the network port show command.

Ports that are not assigned to a broadcast domain display - in the Broadcast Domain column.

Example

```
cluster01::> network port show
```

(Mbps)	Node	Port	IPspace	Broadcast Domain	Link	MTU	Speed Admin/Oper

cluster01-01							
1000		e0a	Cluster	Cluster	up	9000	auto/
1000		e0b	Cluster	Cluster	up	9000	auto/
1000		e0c	Default	Default	up	1500	auto/
1000		e0d	Default	Default	up	1500	auto/
1000		e0e	Default	-	up	1500	auto/
1000		e0f	Default	-	up	1500	auto/
1000		e0g	Default	Default	up	1500	auto/
cluster01-02							
1000		e0a	Cluster	Cluster	up	9000	auto/
1000		e0b	Cluster	Cluster	up	9000	auto/
1000		e0c	Default	Default	up	1500	auto/
1000		e0d	Default	Default	up	1500	auto/
1000		e0e	Default	-	up	1500	auto/
1000		e0f	Default	-	up	1500	auto/
1000		e0g	Default	Default	up	1500	auto/

7. Create the broadcast domain in the “ip-space-IC1” IPspace for the ports you want to dedicate to intercluster operations by using the network port broadcast-domain create command.

Part of the process of creating the broadcast domain for the ports is assigning the unassigned ports to the broadcast domain.

Example

This example creates the “ip-space-IC1-bd” broadcast domain in the “ip-space-IC1” IPspace:

```
cluster01::> network port broadcast-domain create -ip-space ip-space-IC1 -broadcast-domain ip-space-IC1-bd -mtu 1500 -ports cluster01-01:e0e,cluster01-01:e0f,cluster01-02:e0e,cluster01-02:e0f
```

- Optional: Group the ports that you want to use for the intercluster LIFs by using the `network interface failover-groups create` command.

If the intercluster connectivity requirements allow for the intercluster LIFs to use any port in the broadcast domain, then you do not need to create a separate failover group.

Example

```
cluster01::> network interface failover-groups create -vserver cluster01
-failover-group intercluster01 -targets cluster01-01:e0e,cluster01-01:e0f,
cluster01-02:e0e,cluster01-02:e0f
```

- Verify that the broadcast domain was created and the ports were assigned by using the `network port broadcast-domain show` command:

Example

```
cluster01::> network port broadcast-domain show
IPspace Broadcast
Name Domain Name MTU Port List Update
Details Status
-----
Cluster Cluster 9000
cluster01-01:e0a complete
cluster01-01:e0b complete
cluster01-02:e0a complete
cluster01-02:e0b complete
Default Default 1500
cluster01-01:e0c complete
cluster01-01:e0d complete
cluster01-01:e0f complete
cluster01-01:e0g complete
cluster01-02:e0c complete
cluster01-02:e0d complete
cluster01-02:e0f complete
cluster01-02:e0g complete
ipspace-IC1
  ipspace-IC1-bd 1500
cluster01-01:e0e complete
cluster01-01:e0f complete
cluster01-02:e0e complete
cluster01-02:e0f complete
```

- Display the failover-group that you created by using the `network interface failover-groups show` command:

Example

```
cluster01::> network interface failover-groups show
Vserver Group Failover
Targets
-----
Cluster Cluster
cluster01-01:e0a, cluster01-01:e0b,
cluster01-02:e0a, cluster01-02:e0b
cluster01 Default
cluster01-01:e0c, cluster01-01:e0d,
cluster01-02:e0c, cluster01-02:e0d,
cluster01-01:e0e, cluster01-01:e0f
cluster01-02:e0e, cluster01-02:e0f
ipspace-IC1 ipspace-IC-bd
cluster01-01:e0e, cluster01-01:e0f
cluster01-02:e0e, cluster01-02:e0f
```

11. Create an intercluster LIF on the system SVM ipspace-IC1 by using the `network interface create` command.

Example

This example uses the LIF naming convention `adminSVMname_icl#` for the intercluster LIF:

```
cluster01::> network interface create -vserver ipspace-IC1 -lif cluster01_icl01 -role
intercluster -home-node cluster01-01 -home-port e0e
-address 192.168.1.201 -netmask 255.255.255.0 -failover-group intercluster01

cluster01::> network interface create -vserver ipspace-IC1 -lif cluster01_icl02 -role
intercluster -home-node cluster01-02 -home-port e0e
-address 192.168.1.202 -netmask 255.255.255.0 -failover-group intercluster01
```

12. Verify that the intercluster LIFs were created properly by using the `network interface show` command:

Example

```
cluster01::> network interface show
Vserver Logical Status Network Current Current Is
----- Interface Admin/Oper Address/Mask Node Port Home
-----
Cluster
cluster01-01 cluster01-01_clus_1 up/up 192.168.0.xxx/24 cluster01-01 e0a true
cluster01-01 cluster01-01_clus_2 up/up 192.168.0.xxx/24 cluster01-01 e0b true
cluster01-02 cluster01-02_clus_1 up/up 192.168.0.xxx/24 cluster01-01 e0a true
cluster01-02 cluster01-02_clus_2 up/up 192.168.0.xxx/24 cluster01-01 e0b true
cluster01 cluster_mgmt up/up 192.168.0.xxx/24 cluster01-01 e0c true
cluster01 cluster01_icl01 up/up 192.168.1.201/24 cluster01-01 e0e true
cluster01 cluster01_icl02 up/up 192.168.1.202/24 cluster01-02 e0e true
cluster01 cluster01-01_mgmt1 up/up 192.168.0.xxx/24 cluster01-01 e0c true
cluster01 cluster01-02_mgmt1 up/up 192.168.0.xxx/24 cluster01-02 e0c true
```

13. Verify that the intercluster LIFs are configured for redundancy by using the `network interface show` command with the `-role intercluster` and `-failover` parameters.

Example

The LIFs in this example are assigned the e0e home port on each node. If the e0e port fails, the LIF can fail over to the e0f port.

```
cluster01::> network interface show -role intercluster -failover
Vserver Logical Home Failover Failover
----- Interface Node:Port Policy Group
-----
cluster01-01
cluster01-01 cluster01-01_icl01 cluster01-01:e0e local-only intercluster01
Failover Targets: cluster01-01:e0e,
cluster01-01:e0f
cluster01-01 cluster01-01_icl02 cluster01-02:e0e local-only intercluster01
Failover Targets: cluster01-02:e0e,
cluster01-02:e0f
```

14. Display the routes in the cluster by using the `network route show` command to determine whether intercluster routes are available or you must create them.

Creating a route is required only if the intercluster addresses in both clusters are not on the same subnet and a specific route is needed for communication between the clusters.

Example

In this example, no intercluster routes are available:

```
cluster01::> network route show
Vserver  Destination      Gateway      Metric
-----  -
Cluster
cluster01 0.0.0.0/0        192.168.0.1  20
cluster01 0.0.0.0/0        192.168.0.1  10
```

15. If communication between intercluster LIFs in different clusters requires routing, create an intercluster route by using the `network route create` command.

The gateway of the new route should be on the same subnet as the intercluster LIF.

Example

In this example, 192.168.1.1 is the gateway address for the 192.168.1.0/24 network. If the destination is specified as 0.0.0.0/0, then it becomes a default route for the intercluster network.

```
cluster01::> network route create -vserver cluster01
-destination 0.0.0.0/0 -gateway 192.168.1.1 -metric 40
```

16. Verify that you created the routes correctly by using the `network route show` command:

Example

```
cluster01::> network route show
Vserver  Destination      Gateway      Metric
-----  -
Cluster
cluster01 0.0.0.0/0        192.168.0.1  20
cluster01 0.0.0.0/0        192.168.0.1  10
cluster01 0.0.0.0/0        192.168.1.1  40
```

17. Repeat these steps to configure intercluster networking in the peer cluster.

The peer cluster can have its intercluster LIFs in its Default IPspace, or any other IPspace, as long as there is connectivity between the intercluster LIFs of the two clusters.

18. Verify that the ports have access to the proper subnets, VLANs, and so on.

Dedicating ports for replication in one cluster does not require dedicating ports in all clusters; one cluster might use dedicated ports, while the other cluster shares data ports for intercluster replication.

Creating the cluster peer relationship

You create the cluster peer relationship using a set of intercluster logical interfaces to make information about one cluster available to the other cluster for use in cluster peering applications.

Before you begin

- Intercluster LIFs should be created in the IPspaces of both clusters you want to peer.
- You should ensure that the intercluster LIFs of the clusters can route to each other.
- If there are different administrators for each cluster, the passphrase used to authenticate the cluster peer relationship should be agreed upon.

About this task

If you created intercluster LIFs in a nondefault IPspace, you need to designate the IPspace when you create the cluster peer.

Steps

1. Create the cluster peer relationship on each cluster by using the `cluster peer create` command.

The passphrase that you use is not displayed as you type it.

If you created a nondefault IPspace to designate intercluster connectivity, you use the `ipspace` parameter to select that IPspace.

Example

In the following example, cluster01 is peered with a remote cluster named cluster02. Cluster01 is a two-node cluster that has one intercluster LIF per node. The IP addresses of the intercluster LIFs created in cluster01 are 192.168.2.201 and 192.168.2.202. Similarly, cluster02 is a two-node cluster that has one intercluster LIF per node. The IP addresses of the intercluster LIFs created in cluster02 are 192.168.2.203 and 192.168.2.204. These IP addresses are used to create the cluster peer relationship.

```
cluster01::> cluster peer create -peer-addr
192.168.2.203,192.168.2.204
Please type the passphrase:
Please type the passphrase again:
```

```
cluster02::> cluster peer create -peer-addr
192.168.2.201,192.168.2.202
Please type the passphrase:
Please type the passphrase again:
```

If DNS is configured to resolve host names for the intercluster IP addresses, you can use host names in the `-peer-addr` option. It is not likely that intercluster IP addresses frequently change; however, using host names allows intercluster IP addresses to change without having to modify the cluster peer relationship.

Example

In the following example, an IPspace called IP01A was created on cluster01 for intercluster connectivity. The IP addresses used in the previous example are used in this example to create the cluster peer relationship.

```
cluster01::> cluster peer create -peer-addr
192.168.2.203,192.168.2.204
-ipspace IP01A
Please type the passphrase:
Please type the passphrase again:
```

```
cluster02::> cluster peer create -peer-addr
192.168.2.201,192.168.2.202
Please type the passphrase:
Please type the passphrase again:
```

2. Display the cluster peer relationship by using the `cluster peer show` command with the `-instance` parameter.

Displaying the cluster peer relationship verifies that the relationship was established successfully.

Example

```
cluster01::> cluster peer show -instance
Peer Cluster Name: cluster02
Remote Intercluster Addresses: 192.168.2.203,192.168.2.204
Availability: Available
Remote Cluster Name: cluster02
Active IP Addresses: 192.168.2.203,192.168.2.204
Cluster Serial Number: 1-80-000013
```

3. Preview the health of the nodes in the peer cluster by using the `cluster peer health show` command.

Previewing the health checks the connectivity and status of the nodes on the peer cluster.

Example

```
cluster01::> cluster peer health show
Node      cluster-Name      Node-Name      RDB-Health      Cluster-Health      Avail...
-----
cluster01-01
  cluster02
    Data: interface_reachable
    ICMP: interface_reachable      true      true      true
  cluster02-02
    Data: interface_reachable
    ICMP: interface_reachable      true      true      true
cluster01-02
  cluster02
    Data: interface_reachable
    ICMP: interface_reachable      true      true      true
  cluster02-02
    Data: interface_reachable
    ICMP: interface_reachable      true      true      true
```

Related tasks

[Configuring intercluster LIFs to share data ports](#) on page 274

[Configuring intercluster LIFs to use dedicated intercluster ports](#) on page 276

Modifying a cluster peer relationship

You can modify a cluster peer relationship if the name of the cluster you connected to, the logical interface you used, or the IP address you used when creating the cluster peer relationship changes. For example, the IP address of the cluster you used when creating the relationship changed.

Step

1. To change the configuration of a cluster peer relationship, use the `cluster peer modify` command.

The following example changes the IP address of the cluster peer configuration of a cluster named `cluster_b` to `172.19.7.3`:

```
node::> cluster peer modify -cluster cluster_b -stable-addr
172.19.7.3
```

Deleting a cluster peering relationship

You can delete a cluster peering relationship if the relationship is no longer needed. You must delete the cluster peering relationship from each of the clusters in the relationship.

Before you begin

All Storage Virtual Machine (SVM) peer relationships between the two cluster peers must have been deleted.

About this task

This procedure assumes that you are the administrator of only one of the clusters in the cluster peering relationship.

Steps

1. Delete the cluster peering relationship from the cluster of which you are the administrator by using the `cluster peer delete` command.

Example

The following example deletes the cluster peering relationship with the `cluster2` cluster from the `cluster1` cluster:

```
cluster1::> cluster peer delete -cluster cluster2
```

2. Ask the administrator of the other cluster to delete the cluster peering relationship from the other cluster by using the `cluster peer delete` command.

Example

The following example deletes the cluster peering relationship with the `cluster1` cluster from the `cluster2` cluster:

```
cluster2::> cluster peer delete -cluster cluster1
```

Related tasks

[Deleting an SVM peer relationship](#) on page 295

Managing SVM peer relationships

A cluster administrator can create and manage SVM peer relationships between two Storage Virtual Machines (SVMs, formerly known as Vservers) either existing within a cluster (intracluster) or in peered clusters (intercluster) to provide an infrastructure for peering applications, such as SnapMirror.

Peered clusters and peered SVMs can be managed either by the same cluster administrator or different cluster administrators.

The cluster administrator can perform the following SVM peer management tasks:

- Creating SVM peer relationships
- Accepting SVM peer relationships
- Rejecting SVM peer relationships

- Suspending SVM peer relationships
- Resuming SVM peer relationships
- Modifying SVM peering applications on the SVM peer relationships
- Deleting SVM peer relationships
- Viewing SVM peer relationships
- Setting up SnapMirror relationships between volumes of the peered SVMs
 - Note:** You cannot set up a load-sharing SnapMirror relationship between volumes of intercluster SVM peers.

An SVM administrator can perform only the following SVM peer management tasks:

- Viewing SVM peer relationships to identify the peered SVMs
- Setting up SnapMirror relationships, such as a data protection relationship (DP), SnapVault relationship (XDP), and transition relationship (TDP), between volumes of the peered SVMs

Note: The Data ONTAP command-line interface (CLI) continues to use the term *Vserver* in the output, and `vserver` as a command or parameter name has not changed.

Related concepts

[Managing cluster peer relationships](#) on page 265

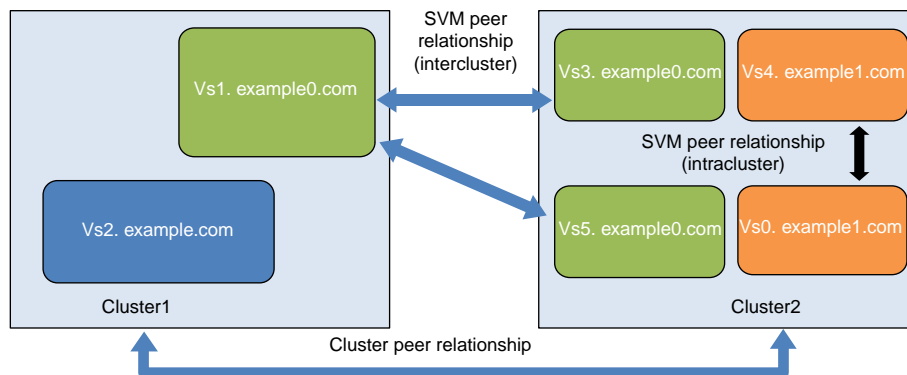
Related information

[Clustered Data ONTAP 8.3 Data Protection Guide](#)

What an SVM peer relationship is

An SVM peer relationship is an authorization infrastructure that enables a cluster administrator to set up peering applications such as SnapMirror relationships between SVMs either existing within a cluster (intracluster) or in the peered clusters (intercluster). Only a cluster administrator can set up SVM peer relationships.

The following illustration shows the intercluster and intracluster SVM peer relationships:



The SVM peer infrastructure enables you to set up a backup and recovery mechanism between SVMs. You can set up a mirroring relationship at the volume level between peered SVMs. If a volume in the SVM becomes unavailable, the cluster administrator or SVM administrator can configure the respective mirrored volume of the peered SVM to serve data.

One SVM can be peered with multiple SVMs within a cluster or across clusters.

You can set up only SnapMirror data protection (DP) and SnapVault (XDP) relationships by using the SVM peer infrastructure.

States of SVM peer relationships

SVM peer relationships can be in different states depending on the operation performed on the SVM peer relationship. You must be aware of the states of the SVM peer relationship to perform other operations such as SnapMirror data transfer between peered SVMs.

The following table lists the different states of an SVM peer relationship and helps you understand when an SVM peer relationship is in a particular state:

SVM peer relationship is in...	When...
initializing state on the local cluster	<ul style="list-style-type: none"> The local cluster is communicating with the peer cluster for initializing the SVM peer relationship
initiated state on the local cluster pending state on the peered cluster	<ul style="list-style-type: none"> An intercluster SVM peer relationship is requested from the local cluster
peered state on the local and peered clusters	<ul style="list-style-type: none"> An intercluster SVM peer relationship is accepted from the peered cluster An intracluster SVM peer relationship is established An intercluster or intracluster SVM peer relationship is resumed
rejected state on the local cluster	<ul style="list-style-type: none"> An intercluster SVM peer relationship is rejected from the peered cluster
suspended state on the local and peered clusters	<ul style="list-style-type: none"> An intercluster or intracluster SVM peer relationship is suspended from the local or peered cluster

Creating an SVM peer relationship

A cluster administrator can create a Storage Virtual Machine (SVM) peer relationship to provide an authorization infrastructure for running SVM peering applications between two SVMs by using the `vserver peer create` command. You can create an SVM peer relationship between two SVMs either in a single cluster (intracluster) or in peered clusters (intercluster).

Before you begin

- If you want to create an intercluster SVM peer relationship, you must have ensured that both the clusters are peered with each other.
- The names of the SVMs in the peered clusters must be unique across the clusters to be peered and any other clusters with which either of the clusters are individually peered.
If the SVMs do not have unique names, you must rename one of the SVMs by using the `vserver rename` command.

For example, consider two clusters, cluster A and cluster B, that are peered with cluster C. Clusters A and cluster B must not have SVMs with identical names even though cluster A and

cluster B are not peered. You must rename one of the SVMs, if there are SVMs with identical names.

- The admin state of the SVMs to be peered must not be in `initializing` or `deleting` state.
- If any previously attempted SVM peer relationship between the same SVMs is in the `rejected` state, you must have deleted these SVM peer relationships.

About this task

- Peered clusters can be managed by a single cluster administrator or different cluster administrators.
- You can specify the applications that will communicate over the peer relationship when you create an SVM peer relationship.
If you do not specify the application for the peer relationship, such as `snapmirror`, an SVM administrator cannot perform operations related to the applications between the peered SVMs.
- For SVMs with FlexVol volumes, you can create intercluster and intracluster SVM peer relationships.
- For SVMs with Infinite Volume, you can create only intercluster SVM peer relationships.
- You cannot create an SVM peer relationship between SVMs with FlexVol volumes and SVMs with Infinite Volume.
- You can create multiple SVM peer relationships simultaneously either by using different SSH sessions or by using a script.

Note: It is best to create not more than five SVM peer relationships simultaneously to avoid any performance degradation.

Choices

- [Creating an intercluster SVM peer relationship](#) on page 290
- [Creating an intracluster SVM peer relationship](#) on page 291

Creating an intercluster SVM peer relationship

You can create intercluster SVM peer relationships between two clusters to provide the infrastructure for use cases such as intercluster volume SnapMirror configurations. You can create intercluster SVM peer relationships between two clusters to provide the infrastructure for use cases such as intercluster volume SnapMirror configurations and SVM disaster recovery.

Before you begin

The two clusters must already be peered.

Steps

1. Use the `vserver peer create` command to create an SVM peer relationship.

Example

The following command creates an intercluster SVM peer relationship between `vs1.example0.com` (on cluster1) and `vs3.example0.com` (on cluster2):

```
cluster1::> vserver peer create -vserver vs1.example0.com -peer-vserver vs3.example0.com -
applications snapmirror -peer-cluster cluster2
Info: [Job 43] 'vserver peer create' job queued
```

The intercluster SVM peer relationship is in `initiated` state.

2. Use the `vserver peer show-all` command to view the status and other details of the SVM peer relationship.

Example

```
cluster1::> vserver peer show-all
Vserver      Peer      Peer      Peer      Peering
-----      -
vs1.example0.com  vs3.example0.com  initiated  Cluster2  snapmirror
```

For more information about this command, see the man pages.

After you finish

You must inform the cluster administrator of the peered cluster about the SVM peer request for the authentication to be completed.

The SVM peer relationship is not established until the cluster administrator of the peered cluster accepts the SVM peer request.

Related tasks

[Accepting an SVM peer relationship](#) on page 292

Creating an intracluster SVM peer relationship

You can create SVM peer relationships between SVMs within a cluster for operations such as backup of SVM data within a cluster.

About this task

You cannot create intracluster SVM peer relationships for SVMs with Infinite Volumes.

Steps

1. Use the `vserver peer create` command to create an SVM peer relationship.

Example

The following command creates an intracluster SVM peer relationship between the SVMs `vs4.example1.com` and `vs0.example1.com`, both residing on `cluster2`:

```
cluster2::> vserver peer create -vserver vs4.example1.com -peer-vserver vs0.example1.com -
applications snapmirror
Info: 'vserver peer create' command is successful.
```

An intracluster SVM peer relationship is created and is in `peered` state. Authentication is not required because the cluster is managed by a single cluster administrator.

2. Use the `vserver peer show-all` command to view the status and other details of the SVM peer relationship.

Example

```
cluster2::> vserver peer show-all
Peer
Vserver      Vserver      Peer State  Peer Cluster  Peering
-----
vs4.example1.com vs0.example1.com peered   cluster2     snapmirror
vs0.example1.com vs4.example1.com peered   cluster2     snapmirror
```

Accepting an SVM peer relationship

When a cluster administrator creates an intercluster SVM peer relationship, the cluster administrator of the remote cluster can accept the SVM peer request to establish the peer relationship between the SVMs by using the `vserver peer accept` command.

About this task

Peered clusters can be managed by a single administrator or different cluster administrators. If a single cluster administrator is managing the peered clusters, the cluster administrator has to accept the SVM peer request on the peered cluster. If different administrators are managing the peered clusters, the cluster administrator who initiates the SVM peer request has to notify the cluster administrator of the peered cluster about the incoming SVM peer request through any channel such as email.

Steps

1. Use the `vserver peer show` command to view the SVM peer requests.

Example

The following example shows how to view the SVM peer requests on cluster2:

```
cluster2::> vserver peer show
Peer
Vserver      Vserver      Peer State
-----
vs3.example0.com vs1.example0.com pending
```

2. Use the `vserver peer accept` command to accept the SVM peer request and establish the SVM peer relationship.

Example

The following example shows how to accept an incoming SVM peer request to establish an SVM peer relationship between `vs1.example0.com` and `vs3.example0.com` on cluster1 and cluster2 respectively:

```
cluster2::> vserver peer accept -vserver vs3.example0.com -peer-
vserver vs1.example0.com

Info: [Job 46] 'vserver peer accept' job queued
```

The SVM peer relationship is established and state is peered.

3. Use the `vserver peer show` command on either of the peered clusters to view the state of the SVM peer relationship.

Example

The following example shows how to view to state of the SVM peer relationships:

```
cluster2::> vserver peer show
Vserver          Peer          Peer
                Vserver      State
-----
vs3.example0.com vs1.example0.com  peered
```

For more information about these commands, see the man pages.

Result

A cluster or SVM administrator can establish peering applications such as SnapMirror between the peered SVMs.

Rejecting an SVM peer relationship

When a cluster administrator creates an intercluster SVM peer relationship, the cluster administrator of the peered cluster can reject the SVM peer request to prevent peer relationship between the SVM by using the `vserver peer reject` command.

About this task

If the SVM peer request is initiated with an unauthorized SVM, then the cluster administrator of the peered cluster can reject the relationship. Other peering operations cannot be performed on the rejected peering relationship.

Steps

1. Use the `vserver peer show` command to view the SVM peer requests on the peered cluster.

Example

The following example shows how to view the SVM peer requests on cluster2:

```
cluster2::> vserver peer show
Vserver          Peer          Peer
                Vserver      State
-----
vs5.example0.com vs1.example0.com  pending
```

2. Use the `vserver peer reject` command to reject the SVM peer request.

Example

The following example illustrates how to reject an incoming SVM peer request between `vs1.example0.com` and `vs5.example0.com` on cluster1 and cluster2 respectively:

```
cluster2::> vserver peer reject -vserver vs5.example0.com -peer-
vserver vs1.example0.com

Info: [Job 48] 'vserver peer reject' job queued
```

The SVM peer relationship is in rejected state.

3. Use the `vserver peer show` command on the cluster from which the SVM peer request was created to view the state of the SVM peer relationship.

Example

The following example shows how to view to state of the SVM peer relationships:

```
cluster1::> vserver peer show
Vserver          Peer          Peer
-----          -
vs1.example0.com vs5.example0.com rejected
```

4. Use the `vserver peer delete` command to delete the rejected SVM peer requests because when you create the SVM relationship between the same SVM again, it fails.

Example

The following example shows how to delete the rejected SVM peer requests:

```
cluster1::> vserver peer delete -vserver vs1.example0.com -peer-
vserver vs5.example0.com

Info: 'vserver peer delete' command is successful.
```

For more information about these commands, see the man pages.

Modifying the peering application on an SVM peer relationship

A cluster administrator can modify an SVM peering application running on the SVM peer relationship by using the `vserver peer modify` command. The SVM peering relationship can have SnapMirror, FileCopy, or no application.

About this task

The SVM peer relationship must have peering application as **snapmirror** for all SnapMirror operations between the peered SVMs or **file-copy** for all the FileCopy related operations between the peered SVMs.

Steps

1. Use the `vserver peer modify` command to modify the application on the SVM peer relationship.

Example

The following command modifies the application on the SVM peer relationship:

```
cluster2::>vserver peer modify -vserver vs4.example.com -peer-vserver vs0.example.com -
applications snapmirror

Warning: The following applications were enabled between Vserver "vs4.example.com" and
peer Vserver "vs0.example.com": file-copy, snapmirror. The following applications will be
removed: file-copy. Any operations related to the removed application in the context of
this Vserver peer relationship will be disabled.
Do you want to continue? {y|n}: y

Info: 'vserver peer modify' command is successful.
```

2. Use the `vserver peer show-all` to view the applications running on the SVM peer relationship.

Example

The following command displays the applications running on the SVM peer relationship:

```
cluster2::> vserver peer show-all
```

Vserver	Vserver	State	Peer Cluster	Applications
-----	-----	-----	-----	-----
vs4.example1.com	vs0.example1.com	peered	cluster2	snapmirror

Deleting an SVM peer relationship

A cluster administrator can delete the Storage Virtual Machine (SVM) peer relationship by using the `vserver peer delete` command when the relationship between two SVMs is no longer required.

Before you begin

The SnapMirror relationship defined on the SVM peer relationship must be deleted.

About this task

If one of the peered clusters is running clustered Data ONTAP 8.2 or 8.2.1, then you must delete the SVM peer relationship from both the peered clusters.

You can delete multiple SVM peer relationships simultaneously either by using different SSH sessions or by using a script.

Note: It is best to delete not more than five SVM peer relationships simultaneously to avoid any performance degradation.

Steps

1. Use the `vserver peer delete` command on one of the peered clusters to delete the SVM peer relationship.

Example

The following command deletes the SVM peer relationship from both the clusters:

```
cluster1::> vserver peer delete -vserver vs1.example0.com -peer-
vserver vs3.example0.com

Info: [Job 47] 'vserver peer delete' job queued
```

2. If the `vserver peer delete` command fails due to unavailability of one of the peered clusters, choose one of the following actions:
 - Establish the network connectivity between the two clusters and use the `vserver peer delete` command to delete the SVM peer relationship (recommended).
 - Use the `vserver peer delete` command with the `-force` option on both the local and peered clusters to delete the SVM peer relationship if the cluster peer relationship is not reestablished.
3. Use the `vserver peer show` command on both the clusters to verify that the deleted SVM peer relationship is not displayed.

Example

```
cluster1::> vserver peer show
```

Vserver	Peer Vserver	Peer State
-----	-----	-----
vs1.example0.com	vs3.example0.com	peered

4. If any SVM peer relationship is in the `deleted` state, delete that SVM peer relationship again by using the `vserver peer delete` command.

Related tasks

[Connecting one cluster to another cluster in a peer relationship](#) on page 267

Suspending an SVM peer relationship

A cluster administrator can suspend an established SVM peer relationship whenever needed by using the `vserver peer suspend` command. For example, during the maintenance period, you might want to suspend the SVM peer relationship.

About this task

When you suspend the SVM peer relationship, any SnapMirror data transfer that was initiated before suspending an SVM peer relationship is not affected and the operation is completed. Any data transfer that was scheduled to run during suspension period will not get initiated.

Steps

1. Use the `vserver peer suspend` command on either of the peered cluster to suspend an active SVM peer relationship.

Example

The following example shows how to suspend an SVM peer relationship:

```
cluster2::> vserver peer suspend -vserver vs4.example1.com -peer-
vserver vs0.example1.com

Info: [Job 50] 'vserver peer suspend' job queued
```

The SVM peer relationship is in suspended state.

2. Use the `vserver peer show` command to verify the status of the SVM peer relationship.

Example

The following example shows how to verify the status of the SVM peer relationship:

```
cluster2::> vserver peer show
Vserver          Peer          Peer
-----          -
vs4.example1.com vs0.example1.com suspended
```

For more information about these commands, see the man pages.

Resuming an SVM peer relationship

A cluster administrator can resume a suspended SVM peer relationship by using the `vserver peer resume` command. For example, after the maintenance is complete, you can resume the suspended SVM peering relationship.

About this task

Any SnapMirror data transfer that was scheduled to run during the suspension period will not get initiated when you resume the SVM peer relationship. You must manually initiate the data transfer.

Steps

1. Use the `vserver peer resume` command to resume a suspended SVM peer relationship from either of the peered clusters.

Example

The following example shows how to resume a suspended SVM peer relationship:

```
cluster1::> vserver peer resume -vserver vs4.example1.com -peer-
vserver vs0.example1.com

Info: [Job 76] 'vserver peer resume' job queued
```

The SVM peer relationship is in peered state.

2. Use the `vserver peer show` command to verify the status of the SVM peer relationship.

Example

The following example shows how to verify the status of the SVM peer relationship:

```
cluster1::> vserver peer show

Vserver                Peer                Peer
-----                -Vserver            -State
vs4.example1.com      vs0.example1.com    peered
```

For more information about these commands, see the man pages.

Displaying information about SVM peer relationships

Peer Storage Virtual Machines (SVMs) are fully functional SVMs which could be either local or remote. Cluster administrators and SVM administrators can view the peers of the SVM to set up peering applications such as SnapMirror between volumes of the peer SVMs by using the `vserver peer show` command.

About this task

You can also view the status of the SVM peer relationships and the applications running on the peer relationship.

Step

1. Use the appropriate command to view the details of SVM peer relationships:

If you want to view information about...**Enter the following command...**

Peered SVMs and the peer state

```
vserver peer show
```

The following example shows how to view the information about the peered Storage Virtual Machines (SVMs, formerly known as Vservers):

```
cluster1::> vserver peer show

Vserver          Peer          Peer
-----          -
vs1.example0.com vs3.example0.com peered
vs1.example0.com vs5.example0.com rejected
2 entries were displayed.
```

The applications running on the SVM peer relationship

```
vserver peer show-all
```

The following example shows how to view the information about the peered SVMs:

```
cluster1::> vserver peer show-all

Vserver          Peer          Peer          Peer          Peering
Applications     Vserver       State         Cluster
-----          -
vs1.example0.com vs5.example0.com peered        cluster2
snapmirror
```

For more information about this command, see the man pages.

Glossary

A

ACL

Access control list.

active/active configuration

- In the Data ONTAP 7.2 and 7.3 release families, a pair of storage systems or V-Series systems (sometimes called *nodes*) configured to serve data for each other if one of the two systems stops functioning. Also sometimes referred to as *active/active pairs*.
- In the Data ONTAP 8.x release family, this functionality is referred to as a *high-availability (HA) configuration* or an *HA pair*.
- In the Data ONTAP 7.1 release family and earlier releases, this functionality is referred to as a *cluster*.

address resolution

The procedure for determining an address corresponding to the address of a LAN or WAN destination.

admin SVM

Formerly known as admin Vserver. In clustered Data ONTAP, a Storage Virtual Machine (SVM) that has overall administrative access to all objects in the cluster, including all objects owned by other SVMs, but does not provide data access to clients or hosts.

administration host

A client computer that is used to manage a storage system through a Telnet or Remote Shell connection.

Application Program Interface (API)

A language and message format used by an application program to communicate with the operating system or some other system, control program, or communications protocol.

authentication

The process of verifying the identity of a user who is logging in to a secured system or network.

AutoSupport

An integrated technology that triggers email messages from the customer site to technical support or another specified email recipient when there are any failures in Unified Manager services. These messages contain information such as feature usage metrics, configuration and user settings, system health, and so on.

B

big-endian

A binary data format for storage and transmission in which the most significant byte comes first.

C

caching module

A Flash Cache 2, Flash Cache, or Performance Acceleration Module (PAM) PCIe-based, memory module that optimizes the performance of random read-intensive workloads by functioning as an intelligent external read cache. This hardware works in tandem with the WAFL External Cache software component of Data ONTAP.

CIFS share

- In Data ONTAP, a directory or directory structure that has been made available to network users and can be mapped to a drive letter on a CIFS client. Also known simply as a *share*.
- In OnCommand Insight (formerly SANscreen suite), a service exposed from a NAS device to provide file-based storage through the CIFS protocol. CIFS is mostly used for Microsoft Windows clients, but many other operating systems can access CIFS shares as well.

CLI

command-line interface. The storage system prompt is an example of a command-line interface.

client

A workstation or PC in a client-server architecture; that is, a computer system or process that requests services from and accepts the responses of another computer system or process.

cluster

- In clustered Data ONTAP 8.x, a group of connected nodes (storage systems) that share a namespace and that you can manage as a single virtual server or multiple virtual servers, providing performance, reliability, and scalability benefits.
- In the Data ONTAP 7.1 release family and earlier releases, a pair of storage systems (sometimes called *nodes*) configured to serve data for each other if one of the two systems stops functioning.
- In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an *active/active configuration*.
- For some storage array vendors, *cluster* refers to the hardware component on which host adapters and ports are located. Some storage array vendors refer to this component as a *controller*.

cluster Vserver

Former name for a data SVM; see data SVM.

Common Internet File System (CIFS)

Microsoft's file-sharing networking protocol that evolved from SMB.

community

A logical relationship between an SNMP agent and one or more SNMP managers. A community is identified by name, and all members of the community have the same access privileges.

console

The physical or virtual terminal that is used to monitor and control a storage system.

Copy-On-Write (COW)

The technique for creating Snapshot copies without consuming excess disk space.

D**data SVM**

Formerly known as data Vserver. In clustered Data ONTAP, a Storage Virtual Machine (SVM) that facilitates data access from the cluster; the hardware and storage resources of the cluster are dynamically shared by data SVMs within a cluster.

degraded mode

The operating mode of a storage system when a disk in the RAID group fails or the batteries on the NVRAM card are low.

disk ID number

The number assigned by the storage system to each disk when it probes the disks at startup.

disk sanitization

A multiple write process for physically obliterating existing data on specified disks in such a manner that the obliterated data is no longer recoverable by known means of data recovery.

disk shelf

A shelf that contains disk drives and is attached to a storage system.

E

emulated storage system

A software copy of a failed storage system that is hosted by its takeover storage system. The emulated storage system appears to users and administrators to be a functional version of the failed storage system. For example, it has the same name as the failed storage system.

Ethernet adapter

An Ethernet interface card.

expansion card

A SCSI card, NVRAM card, network card, hot-swap card, or console card that plugs into a storage system expansion slot. Sometimes called an *adapter*.

expansion slot

The slots on the storage system board into which you insert expansion cards.

F

failed storage system

A physical storage system that has ceased operating. In a high-availability configuration, it remains the failed storage system until a giveback succeeds.

Flash Cache module

A PCIe-based, solid state memory module that optimizes the performance of random read-intensive workloads by functioning as an intelligent external read cache. The Flash Cache 2 module is the successor of the Flash Cache module, which is the successor of the Performance Acceleration Module (PAM). This hardware works in tandem with the WAFL External Cache software component of Data ONTAP.

G

giveback

The technology that enables two storage systems to return control of each other's data after the issues that caused a controller failover are resolved.

global namespace

See *namespace*.

group

In Data ONTAP operating in 7-Mode, a group of users defined in the storage system's `/etc/group` file.

Group ID (GID)

The number used by UNIX systems to identify groups.

H**HA (high availability)**

- In Data ONTAP 8.x, the recovery capability provided by a pair of nodes (storage systems), called an *HA pair*, that are configured to serve data for each other if one of the two nodes stops functioning.
- In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an *active/active configuration*.

HA pair

- In Data ONTAP 8.x, a pair of nodes whose controllers are configured to serve data for each other if one of the two nodes stops functioning.
Depending on the system model, both controllers can be in a single chassis, or one controller can be in one chassis and the other controller can be in a separate chassis.
- In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an *active/active configuration*.

heartbeat

A repeating signal transmitted from one storage system to the other that indicates that the storage system is in operation. Heartbeat information is also stored on disk.

hot swap

The process of adding, removing, or replacing a disk while the storage system is running.

hot swap adapter

An expansion card that makes it possible to add or remove a hard disk with minimal interruption to file system activity.

I**inode**

A data structure containing information about files on a storage system and in a UNIX file system.

interrupt switch

A switch on some storage system front panels used for debugging purposes.

M**Maintenance mode**

An option when booting a storage system from a system boot disk. Maintenance mode provides special commands for troubleshooting hardware and configuration.

N**namespace**

In network-attached storage (NAS) environments, a collection of files and path names to the files.

NDMP

Network Data Management Protocol. A protocol that allows storage systems to communicate with backup applications and provides capabilities for controlling the robotics of multiple tape backup devices.

network adapter

An Ethernet, FDDI, or ATM card.

node SVM

Formerly known as node Vserver. In clustered Data ONTAP, a Storage Virtual Machine (SVM) that is restricted to operation in a single node of the cluster at any one time, and provides administrative access to some objects owned by that node. A node SVM does not provide data access to clients or hosts.

normal mode

The state of a storage system when there is no takeover in the HA configuration.

NVMEM

nonvolatile memory.

NVRAM cache

Nonvolatile RAM in a storage system, used for logging incoming write data and NFS requests. Improves system performance and prevents loss of data in case of a storage system or power failure.

NVRAM card

An adapter that contains the storage system's NVRAM cache.

NVRAM mirror

A synchronously updated copy of the contents of the storage system NVRAM (nonvolatile random access memory) contents kept on the partner storage system.

P

PAM (Performance Acceleration Module)

A PCIe-based, DRAM memory module that optimizes the performance of random read-intensive workloads by functioning as an intelligent external read cache. This hardware is the predecessor of the Flash Cache module and works in tandem with the WAFL External Cache software component of Data ONTAP.

panic

A serious error condition causing the system that is running Data ONTAP to halt. Similar to a software crash in the Windows system environment.

parity disk

The disk on which parity information is stored for a RAID4 disk drive array. In RAID groups using RAID-DP protection, two parity disks store the parity and double-parity information. Used to reconstruct data in failed disk blocks or on a failed disk.

partner mode

The method you use to communicate through the command-line interface with a virtual storage system during a takeover.

partner node

From the point of view of the local node (storage system), the other node in an HA configuration.

Performance Acceleration Module (PAM)

See *PAM (Performance Acceleration Module)*.

POST

Power-on self-tests. The tests run by a storage system after the power is turned on.

Q

qtree

A special subdirectory of the root of a volume that acts as a virtual subvolume with special attributes.

R

RAID

Redundant Array of Independent Disks. A technique that protects against disk failure by computing parity information based on the contents of all the disks in a storage array. Storage systems use either RAID4, which stores all parity information on a single disk, or RAID-DP, which stores all parity information on two disks.

RAID disk scrubbing

The process in which a system reads each disk in the RAID group and tries to fix media errors by rewriting the data to another disk area.

S**serial adapter**

An expansion card for attaching a terminal as the console on some storage system models.

serial console

An ASCII or ANSI terminal attached to a storage system's serial port. Used to monitor and manage storage system operations.

SFO

See *storage failover (SFO)*.

SID

Security identifier used by the Windows operating system.

Snapshot copy

An online, read-only copy of an entire file system that protects against accidental deletions or modifications of files without duplicating file contents. Snapshot copies enable users to restore files and to back up the storage system to tape while the storage system is in use.

storage failover (SFO)

In clustered Data ONTAP, the method of ensuring data availability by transferring the data service of a failed node to another node in an HA pair. Transfer of data service is often transparent to users and applications. In Data ONTAP 7.2 and later, and in Data ONTAP operating in 7-Mode, the failover method is called *controller failover*.

Storage Virtual Machine (SVM)

(Known as *Vserver* prior to clustered Data ONTAP 8.2.1. The term “Vserver” is still used in CLI displays and `vserver` command syntax.) A virtual machine that provides network access through unique network addresses, that might serve data out of a distinct namespace, and that is separately administrable from the rest of the cluster. There are three types of SVMs—*admin*, *node*, and *data*—but unless there is a specific need to identify the type of SVM, “SVM” usually refers to the data SVM.

SVM

(Storage Virtual Machine; known as *Vserver* prior to clustered Data ONTAP 8.2.1. The term “Vserver” is still used in CLI displays and `vserver` command syntax.) A virtual machine that provides network access through unique network addresses, that might serve data out of a distinct namespace, and that is separately administrable from the rest of the cluster. There are three types of SVMs—*admin*, *node*, and *data*—but unless there is a specific need to identify the type of SVM, “SVM” usually refers to the data SVM.

T**takeover**

The emulation of the failed node identity by the takeover node in an HA pair; the opposite of *giveback*.

takeover mode

The method you use to interact with a node (storage system) when it has taken over its partner. The console prompt indicates when the node is in takeover mode.

takeover node

A node (storage system) that remains in operation after the other node stops working and that hosts a virtual node that manages access to the failed node disk shelves and network connections. The takeover node maintains its own identity and the virtual node maintains the failed node identity.

trap

An asynchronous, unsolicited message sent by an SNMP agent to an SNMP manager indicating that an event has occurred on the storage system.

U**UID**

user identification number.

Unicode

A 16-bit character set standard. It was designed and is maintained by the nonprofit consortium Unicode Inc.

V**volume**

A file system.

Vserver

(Known as “Storage Virtual Machine (SVM)” in clustered Data ONTAP 8.2.1 and later.)
A virtual machine that provides network access through unique network addresses, that might serve data out of a distinct namespace, and that is separately administrable from the rest of the cluster. There are three types of Vservers—*admin*, *node*, and *cluster* (“cluster Vserver” is called “data Vserver” in Data ONTAP 8.2)—but unless there is a specific need to identify the type of Vserver, “Vserver” usually refers to the cluster/data Vserver.

W**WAFL**

Write Anywhere File Layout. A file system designed for the storage system to optimize write performance.

WAFL External Cache

On a storage system that has a Performance Acceleration Module (PAM), Flash Cache, or Flash Cache 2 module installed, this cache improves storage system performance by reducing the number of disk reads. Sometimes referred to as *WAFL extended cache*.

WINS

Windows Internet Name Service.

workgroup

A collection of computers running Windows NT or Windows for Workgroups that is grouped for browsing and sharing.

Copyright information

Copyright © 1994–2016 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email to docomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

A

access

- configuring for web services [157](#)
- customizing an access-control role to restrict user to specified commands [129](#)
- introduction to managing cluster (cluster administrators only) [116](#)
- managing for web services [151](#)
- nodeshell commands and options in the clustershell [17](#)
- providing access to AD users or groups for SVM administrators [91](#)
- providing access to NIS or LDAP user accounts for SVM administrators [90](#)
- to OnCommand System Manager, ways to manage [29](#)

access methods

- for SVMs [110](#)
- user account [117](#)

access problems

- troubleshooting web service [158](#)

access-control roles

- commands for managing [132](#)
- considerations for customizing [126](#)
- customizing to restrict user access to specified commands [129](#)
- introduction to managing [123](#)
- managing rule settings for user names and passwords in [130](#)
- predefined roles for cluster administrators [123](#)

accessing

- accounts that can access the SP [58](#)
- cluster by using Telnet [14](#)
- cluster with RSH [15](#)
- cluster with serial port [11](#)
- cluster with SSH [11](#)
- Data ONTAP man pages [26](#)
- enabling AD users and groups to access the cluster and SVMs [118](#)
- enabling NIS or LDAP users to access the cluster [120](#)
- managing the IP addresses that can access the SP [60](#)
- node log, core dump, and MIB files with a web browser [41](#)
- remote, understanding the SP [49](#)
- SP from an administration host [58](#)
- SPs, introduction to [57](#)
- system console of a node [42](#)
- the cluster by using OnCommand System Manager browser-based graphic interface [27](#)
- the cluster, introduction to using the CLI to (cluster administrators only) [11](#)

accounts

- access methods for user [117](#)
- commands for managing user [122](#)
- considerations for password rule settings [131](#)
- for accessing the SP [58](#)
- ways to manage user [116](#)

Active Directory

See AD

AD

- enabling users and groups to access the cluster and SVMs [118](#)

AD users or groups

- providing access for SVM administrators [91](#)

adding

- aggregates for SVMs [94](#)
- protocols for SVMs [95](#)

admin

- use of administrative privilege levels [21](#)

admin SVM

- definition of [31](#)

admin SVMs

- described [76](#)

administration

- delegating SVM [91](#)

administration hosts

- accessing the SP from [58](#)
- accessing the SP, introduction to [57](#)
- managing the IP addresses that can access the SP [60](#)

administrative privileges

- use of levels [21](#)

administrators

- differences between cluster and SVM [10](#)
- predefined roles for cluster [123](#)

advanced

- use of administrative privilege levels [21](#)

aggregates

- improving write performance [263](#)
- modifying the aggregates list for SVMs [94](#)
- rules governing node root [43](#)

alerts

- customization of system health [219](#)
- how they trigger AutoSupport messages and events [220](#)
- introduction to monitoring system health [218](#)
- responding to [219](#)
- triggers used in health monitoring [218](#)

algorithms

- key exchange and data encryption, introduction to managing SSH security configuration [134](#)

allowing

- protocols on the SVMs [95](#)

application types

- accounts that can access the SP [58](#)

archive files

- performance, uploading to technical support [209](#)

asset tags

- modifying node attributes [36](#)

attachment

- AutoSupport messages [201](#)

attributes

- displaying node [36](#)
- modifying cluster [34](#)
- modifying node [36](#)

audit logs

- forwarding the command history to a destination [161](#)

- audit settings
 - commands for managing, for management activities [162](#)
 - managing for management activities [160](#)
- auditlog files
 - managing audit settings for management activities [160](#)
- authentication
 - behavior when methods include both public key and password [118](#)
 - cluster peer relationships [267](#)
 - managing digital certificates for server or client [137](#)
- authentication methods
 - enabling NIS or LDAP users to access the cluster [120](#)
 - for accessing SVMs [110](#)
 - for user accounts [117](#)
- automatic configurations
 - enabling the SP network [52](#)
- automatic timeout
 - commands for managing period of CLI sessions [27](#)
- AutoSupport
 - collection size and time budgets [200](#)
 - commands for managing [211](#)
 - configuring [211](#)
 - content [201](#), [202](#)
 - content of messages sent by email [204](#)
 - daily message content [199](#)
 - defined [195](#)
 - displaying information [211](#)
 - enabling and disabling [211](#)
 - event-triggered [201](#)
 - event-triggered message content [199](#)
 - events [201](#)
 - files [202](#)
 - getting message descriptions [210](#)
 - history [211](#)
 - introduction to subsystems [200](#)
 - log files included in messages [202](#)
 - mail host support for
 - transport protocols [205](#)
 - manifest [211](#)
 - manifest, content of [212](#)
 - message content [202](#)
 - Message Matrices [210](#)
 - messages, files collected for [212](#)
 - modifying triggers [211](#)
 - performance message content [199](#)
 - requirements for
 - support for AutoSupport, mail host [205](#)
 - requirements for using [195](#), [205](#)
 - resending messages [211](#)
 - sending messages [211](#)
 - sending messages with core dump files [208](#)
 - sending messages with performance archive files [209](#)
 - setting up [206](#)
 - severity types [204](#)
 - structure of messages sent by email [204](#)
 - subsystems [201](#)
 - supported transport protocols [205](#)
 - troubleshooting message delivery over HTTP or HTTPS [214](#)
 - troubleshooting message delivery over SMTP [215](#)

- troubleshooting when messages are not received [213](#)
- uploading core dump files [208](#)
- uploading performance archive files [209](#)
- weekly message content [199](#)
- weekly messages [202](#)
- when and where messages are sent [196](#)
- AutoSupport cases
 - suppression of [212](#)
- Autosupport messages and events
 - how alerts trigger them [220](#)
- AutoSupport messages and events
 - triggered by system health alerts [220](#)
- AutoSupport OnDemand
 - how it communicates with technical support [203](#)
- available health monitors
 - for clusters [220](#)

B

- backup files
 - commands for managing cluster and node configuration [181](#)
- banners
 - creating [166](#)
 - managing at the cluster and SVM level [168](#)
 - managing the MOTD and, introduction [165](#)
- benefits
 - of using SVMs [76](#)
- blocks
 - caching from random write workloads [258](#)
- boot environment
 - accessing the system console of a node [42](#)
- boot environment prompt
 - booting Data ONTAP from [46](#)
- boot menu
 - accessing the system console of a node [42](#)
- boot menus
 - managing a node with the [48](#)
- booting
 - Data ONTAP at the boot environment prompt [46](#)
- browser-based
 - graphic interface, accessing the cluster by using OnCommand System Manager [27](#)
- budgets
 - AutoSupport size and time [200](#)

C

- CA certificates
 - for server, installing to authenticate cluster or SVM as SSL server [144](#)
- cache rewarming
 - abort events and error messages [260](#)
 - about [259](#)
 - disabling [261](#)
 - enabling [261](#)
 - how it works [260](#)
 - trigger events [260](#)
- caches
 - caching blocks from random write workloads [258](#)
 - caching sequentially read data [258](#)
 - changing the cache size emulated by PCS [255](#)

- changing the sampling resolution of PCS [256](#)
 - comparison of Flash Pool and Flash Cache [254](#)
 - introduction to increasing WAFL external memory [254](#)
- callhome events [201](#)
- cascade configurations
 - networking requirements for cluster peering [272](#)
- cases, AutoSupport
 - suppression during maintenance windows [212](#)
- certificates
 - client, installing to authenticate the cluster or SVM as an SSL client [145](#)
 - commands for managing digital [150](#)
 - digital, managing for server or client authentication [137](#)
 - KMIP, installing to mutually authenticate the cluster and an external key manager [147](#)
 - modifying the SP API configuration [55](#)
 - replacing an expired digital [149](#)
 - server CA, installing to authenticate cluster or SVM as SSL server [144](#)
 - server, installing to authenticate the cluster or SVM as an SSL server [139](#)
- ciphers
 - key exchange and data encryption, introduction to managing SSH security configurations [134](#)
- CLI
 - methods of navigating command directories [18](#)
 - overview of using Data ONTAP [16](#)
 - rules for specifying values [18](#)
 - sessions, automatic timeout [27](#)
 - sessions, records of [26](#)
 - setting display preferences in [21](#)
 - setting privilege levels [21](#)
 - understanding positional parameters [24](#)
- CLI commands
 - for Storage QoS [251](#)
 - introduction to shells for executing [16](#)
 - keyboard shortcuts for editing [19](#)
- CLI sessions
 - commands for managing automatic timeout period of [27](#)
 - commands for managing records of [27](#)
 - introduction to managing (cluster administrators only) [26](#)
 - recording [26](#)
- client CA certificates
 - installing to authenticate an SSL client of the cluster or SVM [141](#)
- client certificates
 - installing KMIP certificates to mutually authenticate the cluster and an external key manager [147](#)
 - installing to authenticate the cluster or SVM as an SSL client [145](#)
- clients
 - managing digital certificates for authentication of servers or [137](#)
- CLIs
 - introduction to accessing the cluster by using, cluster administrators only [11](#)
- cluster
 - attributes, displaying [34](#)
 - replication ring, displaying status [35](#)
- cluster access
 - enabling Telnet or RSH for [13](#)
 - introduction to using the CLI for (cluster administrators only) [11](#)
- cluster administrators
 - differences between SVM administrators and [10](#)
 - introduction to managing cluster access [116](#)
 - predefined roles for [123](#)
- cluster banners
 - creating [166](#)
- cluster configurations
 - automatic backups for [180](#)
 - backing up and restoring [180](#)
 - backup files for [180](#)
 - choosing a configuration for recovering [184](#)
 - commands for managing backup files for [181](#)
 - commands for managing backup schedules for [181](#)
 - managing backups for [180](#)
 - recovering [184](#)
 - restoring using a configuration backup file [185](#)
- cluster management servers
 - definition of [31](#)
- cluster MOTDs
 - creating [169](#)
- cluster peer relationships
 - authentication of [267](#)
 - commands for managing [265](#)
 - requirements for [272](#)
- cluster peers
 - creating relationships between [284](#)
 - definition of intercluster networking [268](#)
 - deleting the relationship [287](#)
 - supported network topologies [268](#)
- cluster switch health monitor
 - commands for [225](#)
 - troubleshooting [223](#)
 - verifying switch monitoring [224](#)
- cluster switch health monitors
 - about [220](#)
- cluster time
 - commands for managing [164](#)
 - managing [163](#)
- cluster user accounts
 - for accessing the SP [58](#)
- cluster-level banners
 - managing [168](#)
- cluster-level MOTDs
 - managing [170](#)
- clusters
 - accessing by using Telnet [14](#)
 - accessing with RSH [15](#)
 - accessing with SSH [11](#)
 - adding nodes to [37](#)
 - automatic configuration backups for [180](#)
 - backing up and restoring configurations for [180](#)
 - commands for managing configuration backup files for [181](#)
 - commands for managing configuration backup schedules for [181](#)
 - configuration backup files for [180](#)
 - description of [30](#)
 - displaying information about the nodes in [33](#)
 - enabling NIS or LDAP users to access the [120](#)

- installing KMIP certificates to mutually authenticate an external key manager and [147](#)
- introduction to accessing by using the CLI (cluster administrators only) [11](#)
- introduction to management basics [30](#)
- introduction to managing access to (cluster administrators only) [116](#)
- introduction to managing CLI sessions (cluster administrators only) [26](#)
- introduction to managing nodes in [36](#)
- managing configuration backups for [180](#)
- modifying attributes of [34](#)
- modifying contact information of [34](#)
- modifying location of [34](#)
- naming requirements for cluster peering [272](#)
- recovering configurations for [184](#)
- removing nodes from [39](#)
- renaming [34](#)
- replication ring, description of [33](#)
- single-node, considerations for [31](#)
- synchronizing nodes with [187](#)
- understanding quorum and epsilon [32](#)
- using serial port to access [11](#)
- clustershell
 - access of nodeshell commands and options [17](#)
 - introduction to [16](#)
- command directories
 - methods of navigating CLI [18](#)
- command history logs
 - forwarding to a destination [161](#)
- command input
 - understanding positional parameters [24](#)
- command-history.log files
 - commands for managing audit settings for management activities [162](#)
 - forwarding the command history log to a destination [161](#)
 - managing audit settings for management activities [160](#)
- command-line interface
 - See* CLI
- commands
 - CLI, keyboard shortcuts for editing [19](#)
 - customizing an access-control role to restrict user access to specified [129](#)
 - displaying available for nodeshell [17](#)
 - for displaying environmental sensor information [241](#)
 - for managing a node remotely [62](#)
 - for managing access-control roles [132](#)
 - for managing audit settings for management activities [162](#)
 - for managing cluster peer relationships [265](#)
 - for managing core dumps [189](#)
 - for managing dashboards [229](#)
 - for managing digital certificates [150](#)
 - for managing events [194](#)
 - for managing job schedules [178](#)
 - for managing jobs [177](#)
 - for managing licenses [175](#)
 - for managing mount points on the nodes [155](#)
 - for managing public keys [137](#)
 - for managing records of CLI sessions [27](#)
 - for managing SSH security configurations [136](#)
 - for managing SSL [156](#)
 - for managing the automatic timeout period of CLI sessions [27](#)
 - for managing the cluster time [164](#)
 - for managing the SP from Data ONTAP [70](#)
 - for managing the web protocol engine [153](#)
 - for managing user accounts [122](#)
 - for monitoring cluster and SVM performance data [238](#)
 - introduction to shells for executing CLI [16](#)
 - methods of customizing show output by using fields [24](#)
 - methods of viewing history and reissuing [19](#)
 - modifying the SVM attributes [96](#)
 - rules for specifying values in the CLI [18](#)
 - system health [225](#)
- comments
 - how to send feedback about documentation [308](#)
- configuration backup file
 - finding for recovering node configurations [182](#)
- configuration backup files
 - commands for managing [181](#)
 - definition of [180](#)
 - finding for recovering cluster configurations [184](#)
 - using to restore a cluster configuration [185](#)
 - using to restore node configurations [183](#)
- configuration backup schedules
 - about [180](#)
 - commands for managing [181](#)
- configuration files
 - rules governing node root volumes and root aggregates [43](#)
- configurations
 - cascade, for cluster peering [272](#)
 - considerations for SP network [51](#)
 - enabling the SP automatic network [52](#)
 - fan-out, for cluster peering [272](#)
 - modifying the SP API service [55](#)
- configurations, network
 - viewing for an SVM [113](#)
- configuring
 - AutoSupport, commands for [211](#)
 - SVMs [88](#)
- connectivity
 - comparison of designated and undesignated intercluster [271](#)
 - supported cluster peer network topologies [268](#)
- console sessions
 - relationship among SP CLI, SP console, and system [60](#)
- consoles
 - accessing the SP from system [59](#)
 - accessing the SP, introduction to [57](#)
 - system, accessing node [42](#)
- contact information
 - modifying cluster [34](#)
- core dump files
 - managing [188](#)
 - methods of segmenting [188](#)
 - node, accessing with a web browser [41](#)
 - sending AutoSupport messages that contain [208](#)
 - uploading to technical support [208](#)
- core dumps

- commands for managing [189](#)
 - core segments
 - commands for managing [190](#)
 - counters
 - importing performance presets for [235](#)
 - what they are [231](#)
 - creating
 - intercluster SVM peer relationships [290](#)
 - intracluster SVM peer relationships [291](#)
 - local user accounts for SVM administrators [89](#)
 - SVM peer relationships [289](#)
 - SVMs for MetroCluster configuration [86](#)
 - SVMs with FlexVol volumes [84](#)
 - SVMs with Infinite Volume [87](#)
 - creating SVMs
 - for data access [84](#)
 - cross cluster relationships
 - deleting [287](#)
- ## D
- daily AutoSupport messages
 - type of data contained within [199](#)
 - when and where they are sent [196](#)
 - dashboards
 - about [228](#)
 - commands for managing [229](#)
 - getting notified of alarms [228](#)
 - performing My AutoSupport tasks [213](#)
 - data
 - caching sequentially read blocks [258](#)
 - performance, collecting and viewing for objects [232](#)
 - performance, commands for monitoring cluster and SVM [238](#)
 - performance, viewing cluster and SVM continuously updated [237](#)
 - SVM performance, introduction to monitoring using advanced methods [235](#)
 - data access
 - creating SVMs with FlexVol volumes for [84](#)
 - data encryption algorithms
 - introduction to managing SSH security configurations [134](#)
 - Data ONTAL CLI
 - understanding positional parameters [24](#)
 - Data ONTAP
 - accessing man pages [26](#)
 - booting at the boot environment prompt [46](#)
 - introduction to management interface [11](#)
 - overview of using the CLI [16](#)
 - Data ONTAP commands
 - for managing the SP [70](#)
 - Data ONTAP-v platforms
 - installing KMIP certificates to mutually authenticate the cluster and an external key manager [147](#)
 - data ports
 - configuring intercluster LIFs to share [274](#)
 - considerations when sharing intercluster and [273](#)
 - data SVMs
 - described [76](#)
 - data, performance
 - viewing cluster and SVM for a time period [236](#)
 - dedicated ports
 - considerations when using for intercluster replication [273](#)
 - degraded system health
 - example of responding to [222](#)
 - delegating
 - SVM administration [91](#)
 - deleting
 - SVM peer relationships [295](#)
 - diagnostic
 - use of administrative privilege levels [21](#)
 - digital certificates
 - commands for managing [150](#)
 - installing a client CA or root CA certificate to authenticate an SSL client of the cluster or SVM [141](#)
 - installing a client certificate to authenticate the cluster or SVM as an SSL client [145](#)
 - installing a server certificate to authenticate the cluster or SVM as an SSL server [139](#)
 - installing server CA certificate to authenticate cluster or SVM as SSL server [144](#)
 - managing for server or client authentication [137](#)
 - managing SSL, introduction [156](#)
 - replacing an expired [149](#)
 - directories
 - methods of navigating CLI command [18](#)
 - disallowing
 - protocols on the SVMs [95](#)
 - discrete SP sensors
 - understanding status values of, in system sensors command output [67](#)
 - display preferences
 - setting in CLI [21](#)
 - displaying
 - SVM details [97, 112](#)
 - documentation
 - additional information about SVM administration [114](#)
 - how to receive automatic notification of changes to [308](#)
 - how to send feedback about [308](#)
 - dump files
 - uploading to technical support [208](#)
- ## E
- eligibility
 - cluster node, displaying information about [33](#)
 - cluster participation, modifying node attributes [36](#)
 - email messages
 - structure of for AutoSupport [204](#)
 - EMS
 - callhome events
 - EMS [198](#)
 - commands for managing events [194](#)
 - data in AutoSupport messages [202](#)
 - getting notified of dashboard alarms [228](#)
 - getting notified of system health alerts [221](#)
 - managing event messages [192](#)
 - relationship with event-triggered AutoSupport messages [198](#)
 - EMS messages
 - cache rewarming [260](#)
 - encryption algorithms

- data, introduction to managing SSH security configurations [134](#)
- entitlement risks
 - managing licenses [173](#)
- environmental sensors
 - displaying information about [241](#)
- epsilon
 - displaying information about whether cluster nodes hold [33](#)
 - understanding cluster [32](#)
- error messages
 - cache rewarming [260](#)
- evaluation license
 - license types and licensed method [174](#)
- Event Management System
 - commands for managing events [194](#)
 - setting up [193](#)
 - See also* EMS
- Event Management Systems
 - See* EMS
- event messages
 - managing [192](#)
 - reducing number of [193](#)
- event-triggered AutoSupport messages
 - how AutoSupport creates and sends them [198](#)
 - when and where they are sent [196](#)
- events
 - AutoSupport messages [201](#)
 - commands for managing [194](#)
- extended queries
 - methods of using [23](#)
- external caches
 - introduction to increasing WAFL memory [254](#)
- external key managers
 - installing KMIP certificates to mutually authenticate the cluster and [147](#)

F

- fabric health monitors
 - about [220](#)
- fan-out configurations
 - networking requirements for cluster peering [272](#)
- feedback
 - how to send comments about documentation [308](#)
- fields
 - methods of customizing show command output by using [24](#)
- files
 - commands for managing cluster and node configuration backup [181](#)
 - controlling I/O performance [247](#)
 - methods of segmenting core dump [188](#)
 - rules for assigning to Storage QoS policy groups [245](#)
- filters
 - viewing statistics by using [233](#)
- FIPS 140-2
 - introduction to managing the web protocol engine [152](#)
- firewall policies
 - enabling Telnet or RSH access to the cluster [13](#)
- firewalls
 - requirements for cluster peering [272](#)

- firmware
 - system, understanding status values of discrete SP sensors in system sensors command output [67](#)
- firmware updates
 - methods of managing SP [56](#)
 - when the SP uses the network interface for [57](#)
- Flash Cache
 - compared with Flash Pool aggregates [254](#)
- Flash Cache modules
 - introduction to increasing WAFL cache memory [254](#)
- Flash Pool aggregates
 - compared with Flash Cache [254](#)
- flexscale.rewarm option [261](#)
- FlexVol volumes
 - controlling I/O performance [247](#)
 - promoting to root for SVMs [103](#)
 - rules for assigning to Storage QoS policy groups [245](#)
 - with SVMs, explained [74](#)
- free space reallocation
 - disabling [264](#)
 - enabling [264](#)
 - how it works [263](#)
 - overview [263](#)
 - supported aggregates [264](#)
 - using with other reallocation features [264](#)
 - viewing status [264](#)
 - when to enable [263](#)
- full-mesh connectivity
 - description [272](#)

G

- get requests
 - managing audit settings for management activities [160](#)
- groups
 - Active Directory, providing access for SVM administrators [91](#)
 - AD access of the cluster and SVMs, enabling users and [118](#)
- guidelines
 - for creating SVMs [80](#)

H

- health
 - cluster node, displaying information about [33](#)
- health monitoring
 - commands [225](#)
 - commands for managing dashboards [229](#)
 - example of responding to degraded health [222](#)
 - getting notified of alerts [221](#)
 - how alerts trigger AutoSupport messages and events [220](#)
 - introduction to system [218](#)
 - responding to degraded health [221](#)
 - ways to control when alerts occur [219](#)
 - ways to respond to alerts [219](#)
- health monitors
 - about alerts [218](#)
 - available cluster [220](#)
 - how they work and about health policies [218](#)

health policies
 about alerts and matching conditions [218](#)

histogram-based predictions in RAVE
 support for [254](#)

history logs
 forwarding the command, to a destination [161](#)

history of commands
 methods of viewing [19](#)

hosts
 accessing the SP from administration [58](#)

how alerts trigger AutoSupport messages and events
 health monitoring [220](#)

HTTP
 troubleshooting AutoSupport message delivery over [214](#)

HTTPS
 managing SSL, introduction [156](#)
 troubleshooting AutoSupport message delivery over [214](#)

I

Infinite Volumes
 creating a new volume as root volume [103](#)
 with SVMs, explained [74](#)

information
 how to send feedback about improving documentation [308](#)

instances
 what they are [231](#)

intercluster connectivity
 comparison of designated and undesignated [271](#)

intercluster LIFs
 configuring to share data ports [274](#)
 configuring to use dedicated intercluster ports [276](#)
 configuring to use dedicated ports in their own networks [279](#)
 considerations when sharing with data ports [273](#)

intercluster networking
 definition of cluster peer [268](#)

intercluster networks
 configuring intercluster LIFs for [274](#), [276](#)
 configuring LIFs in [279](#)
 considerations when sharing data and intercluster ports [273](#)
 supported topologies [268](#)

intercluster ports
 configuring intercluster LIFs to use dedicated [276](#)
 considerations when using dedicated [273](#)

interfaces
 graphic, accessing the cluster by using OnCommand System Manager browser-based [27](#)
 introduction to Data ONTAP management [11](#)
 overview of using Data ONTAP command line [16](#)

IP addresses
 managing access to the SP [60](#)
 requirements for cluster peering [272](#)

IPspace
 considerations when creating SVMs [80](#)

IPspaces
 requirements for cluster peering [272](#)
 using for cluster peer connectivity [279](#)

J

job categories
 about [177](#)

job schedules
 commands for managing [178](#)

jobs
 categories of [177](#)
 commands for managing [177](#)
 managing schedules for [177](#)
 viewing information about [177](#)

joining nodes
 to the cluster [37](#)

K

key exchange algorithms
 introduction to managing SSH security configurations [134](#)

Key Management Interoperability Protocol
See KMIP

key managers
 installing KMIP certificates to mutually authenticate the cluster and external [147](#)

keyboard shortcuts
 for editing CLI commands [19](#)

keys
 managing licenses [173](#)
 ways to manage public [136](#)

KMIP certificates
 installing to mutually authenticate the cluster and an external key manager [147](#)

L

language considerations
 for creating SVMs [80](#)

language settings
 modifying for SVMs [96](#)

LDAP
 enabling users to access the cluster [120](#)

LDAP user accounts
 providing access for SVM administrators [90](#)

levels
 use of administrative privilege [21](#)

licenses
 commands for managing [175](#)
 managing [173](#)
 types and licensed method [174](#)

LIFs
 configuring to use dedicated intercluster ports [276](#)
 configuring to use dedicated ports in their own intercluster networks [279](#)

LIFs, intercluster
 configuring to share data ports [274](#)

local users
 creating accounts for SVM administrators [89](#)

location
 modifying cluster [34](#)
 modifying node attributes [36](#)

log files
 AutoSupport messages, included in [202](#)

- node, accessing with a web browser [41](#)
- logins
 - enabling NIS or LDAP users to access the cluster [120](#)
- logs
 - commands for managing audit settings for management activities [162](#)
 - forwarding the command history to a destination [161](#)
 - ways to manage access to OnCommand System Manager [29](#)
- LUNs
 - controlling I/O performance [247](#)
 - rules for assigning to Storage QoS policy groups [245](#)

M

- mail host support for AutoSupport [205](#)
- mail hosts
 - supported by AutoSupport [205](#)
- maintenance windows
 - suppression of AutoSupport cases [212](#)
- man pages
 - accessing Data ONTAP [26](#)
- management activities
 - commands for managing audit settings for [162](#)
 - managing audit settings for [160](#)
- management basics
 - cluster, introduction to [30](#)
- management devices
 - remote, understanding the SP [49](#)
- management interfaces
 - introduction to Data ONTAP [11](#)
- managing
 - events on cluster, commands for [194](#)
- manifest
 - AutoSupport messages within [212](#)
- manually triggered AutoSupport messages
 - type of data contained within [199](#)
 - when and where they are sent [196](#)
- messages
 - configuring EMS [193](#)
 - files collected for AutoSupport [212](#)
 - managing event [192](#)
- messages of the day
 - See* MOTDs
- messages, AutoSupport
 - suppression of cases during maintenance windows [212](#)
- methods
 - user account access [117](#)
- MetroCluster configuration
 - naming guidelines for SVMs
 - guidelines for creating [80](#)
 - starting SVMs in [100](#)
- MetroCluster configuration for creating SVMs [86](#)
- mgwd.log files
 - commands for managing audit settings for management activities [162](#)
 - managing audit settings for management activities [160](#)
- MIB files

- node, accessing with a web browser [41](#)
- modifying
 - aggregates list for SVMs [94](#)
 - protocols for SVMs [95](#)
 - SVM attributes [94](#)
 - SVM peer relationships [294](#)
- modifying SVM attributes [96](#)
- monitoring
 - commands for managing dashboards [229](#)
 - dashboard [228](#)
 - node connectivity, introduction to [218](#)
 - remote, understanding the SP [49](#)
 - switches, introduction to [218](#)
 - system connectivity, introduction to [218](#)
 - using system health alerts [219](#)
- MOTDs
 - creating [169](#)
 - managing at the cluster and SVM level [170](#)
 - managing the banner and, introduction [165](#)
- mount points
 - commands for managing the node [155](#)
- mutual authentication
 - installing KMIP certificates for the cluster and an external key manager [147](#)
 - managing digital certificates for server or client authentication [137](#)
- My AutoSupport
 - dashboard tasks [213](#)
 - described [213](#)

N

- name services
 - enabling NIS or LDAP users to access the cluster [120](#)
- names
 - modifying cluster [34](#)
- names of nodes
 - changing [37](#)
- network
 - configuring the SP manually [53](#)
 - configuring the SP, introduction to [51](#)
 - full-mesh connectivity described [272](#)
 - requirements for cluster peering [272](#)
- network configurations
 - considerations for SP [51](#)
 - viewing for an SVM [113](#)
- network interfaces
 - when the SP uses for firmware updates [57](#)
- Network Time Protocol
 - See* NTP
- networks
 - enabling the SP automatic configuration [52](#)
- new features
 - support for histogram-based predictions in RAVE [254](#)
- NIS
 - enabling users to access the cluster [120](#)
- NIS user accounts
 - providing access for SVM administrators [90](#)
- node attributes
 - modifying [36](#)
- node configurations

- automatic backups for [180](#)
- backup files for [180](#)
- commands for managing backup files for [181](#)
- commands for managing backup schedules for [181](#)
- finding configuration backup files for recovering [182](#)
- managing backups for [180](#)
- recovering [182](#)
- restoring using a configuration backup file [183](#)
- node connectivity health monitor
 - commands for [225](#)
- node connectivity health monitors
 - about [220](#)
- node names
 - changing [37](#)
- node root aggregates
 - rules governing [43](#)
- node root volumes
 - rules governing [43](#)
- node SVMs
 - described [76](#)
- nodes
 - accessing log, core dump, and MIB files with a web browser [41](#)
 - adding to the cluster [37](#)
 - automatic configuration backups for [180](#)
 - commands for managing configuration backup files for [181](#)
 - commands for managing configuration backup schedules for [181](#)
 - commands for managing mount points on [155](#)
 - commands for managing remotely [62](#)
 - configuration backup files for [180](#)
 - considerations for clusters with single [31](#)
 - displaying attributes [36](#)
 - displaying information about cluster [33](#)
 - freeing up space on the root volume [44](#)
 - introduction to managing [36](#)
 - introduction to starting or stopping [45](#)
 - managing configuration backups for [180](#)
 - managing core dump files after panic [188](#)
 - managing remotely by using the SP, introduction to [49](#)
 - managing with the boot menu [48](#)
 - rebooting at the system prompt [45](#)
 - recovering the configuration for [182](#)
 - rejoining to a re-created cluster [185](#)
 - removing from the cluster [39](#)
 - renaming [37](#)
 - shutting down [47](#)
 - single-node cluster
 - See* single-node clusters
 - synchronizing with the cluster [187](#)
- nodeshell
 - access of commands and options in the clustershell [17](#)
 - displaying available commands for [17](#)
 - introduction to [16](#)
- nsswitch
 - enabling NIS or LDAP users to access the cluster [120](#)
- NTP
 - commands for managing the cluster time [164](#)
 - managing the cluster time with [163](#)

O

- objects
 - collecting and viewing performance data for [232](#)
 - how performance presets define data collection for [232](#)
 - importing performance presets for [235](#)
 - what they are [231](#)
- OnCommand System Manager
 - accessing the cluster by using browser-based graphic interface [27](#)
 - ways to manage access to [29](#)
- online help
 - for using SP CLI [62](#)
- operators
 - methods of using query [22](#)
- output
 - methods of customizing show command by using fields [24](#)
- overview [262](#)
- owners
 - modifying node attributes [36](#)

P

- pair-wise full-mesh connectivity
 - supported cluster peer network topologies [268](#)
- PAMs
 - introduction to increasing WAFL cache memory [254](#)
- panics
 - managing core dump files after [188](#)
- parameters
 - positional, understanding [24](#)
 - rules for specifying values in the CLI [18](#)
- passphrases
 - cluster peer authentication using [267](#)
- password rules
 - considerations for settings [131](#)
- passwords
 - authentication behavior when methods include both public key and [118](#)
 - managing rule settings in access-control role [130](#)
- PCS
 - changing the cache size [255](#)
 - changing the sampling resolution [256](#)
 - improving workload measurements [256](#)
- peer relationships
 - authentication of cluster [267](#)
 - commands for managing cluster [265](#)
 - creating between SVMs [289](#)
 - creating cluster [284](#)
 - creating intracluster SVM [291](#)
 - deleting cluster [287](#)
 - deleting SVM [295](#)
 - for SVMs, described [288](#)
 - intercluster, creating on SVMs [290](#)
 - managing SVM [287](#)
 - modifying SVM [294](#)
 - requirements for clusters [272](#)
 - states of SVMs in [289](#)
- performance
 - controlling workload performance [247](#)
 - data

- decisions before you view [235](#)
 - definition of performance presets [232](#)
 - performance preset definition [232](#)
 - what objects, instances, and counters are [231](#)
 - improving write performance [263](#)
 - introduction to monitoring cluster [230](#)
 - read [262](#)
 - read reallocation [262](#)
 - Performance Acceleration Modules
 - introduction to increasing WAFL cache memory [254](#)
 - performance archive files
 - how performance presets define what data is collected [232](#)
 - modifying what data is collected [235](#)
 - sending AutoSupport messages that contain [209](#)
 - uploading to technical support [209](#)
 - performance AutoSupport messages
 - type of data contained within [199](#)
 - uploading performance archives to technical support [209](#)
 - when and where they are sent [196](#)
 - performance data
 - collecting and viewing for objects [232](#)
 - commands for monitoring cluster and SVM [238](#)
 - SVM, introduction to monitoring using advanced methods [235](#)
 - uploading to technical support [209](#)
 - viewing cluster and SVM continuously updated [237](#)
 - viewing cluster and SVM for a time period [236](#)
 - performance presets
 - collecting performance data for objects [232](#)
 - definition of [232](#)
 - importing collection definitions [235](#)
 - planning
 - to create SVMs [78](#)
 - policy groups
 - creating [247](#)
 - how maximum throughput works [245](#)
 - maximum number of [247](#)
 - monitoring [247](#)
 - rules for assigning storage objects to [245](#)
 - types of [243](#)
 - what they are [243](#)
 - ports
 - configuring intercluster LIFs to use dedicated intercluster [276](#)
 - considerations when sharing data and intercluster roles on [273](#)
 - considerations when using dedicated intercluster [273](#)
 - modifying the SP API service configuration [55](#)
 - requirements for cluster peering [272](#)
 - ports, data
 - configuring intercluster LIFs to share [274](#)
 - ports, dedicated
 - configuring intercluster LIFs to use in their own networks [279](#)
 - positional parameters
 - understanding [24](#)
 - predefined roles
 - for cluster administrators [123](#)
 - predictions
 - support for histogram-based, in RAVE [254](#)
 - Predictive Cache Statistics
 - See* PCS
 - preferences
 - setting display in CLI [21](#)
 - presets
 - performance, importing data collection definitions [235](#)
 - privilege levels
 - setting in CLI [21](#)
 - use of administrative [21](#)
 - prompts
 - booting Data ONTAP at the boot environment [46](#)
 - overview of Data ONTAP command [16](#)
 - rebooting a node at the system [45](#)
 - protocol engines
 - commands for managing the web [153](#)
 - introduction to managing web [152](#)
 - protocols
 - adding or removing for the SVMs [95](#)
 - public key certificates
 - managing digital certificates for server or client authentication [137](#)
 - public keys
 - authentication behavior when methods include both password and [118](#)
 - commands for managing [137](#)
 - ways to manage [136](#)
- ## Q
- QoS policy groups
 - modifying for SVMs [96](#)
 - Quality of Service
 - See* Storage QoS
 - queries
 - methods of using extended [23](#)
 - query operators
 - methods of using [22](#)
 - quorum
 - understanding cluster [32](#)
 - quota policies
 - modifying for SVMs [96](#)
- ## R
- random write workloads
 - caching blocks from [258](#)
 - RAVE
 - support for histogram-based predictions in [254](#)
 - RBAC
 - considerations for customizing an access-control role [126](#)
 - RDBs
 - cluster replication ring, description of [33](#)
 - read reallocation
 - disabling [262](#)
 - enabling [262](#)
 - viewing status [262](#)
 - readahead data
 - enabling caching of [258](#)
 - reallocation
 - free space [263](#)
 - read [262](#)

- when to use with free space reallocation [264](#)
 - rebooting
 - a node at the system prompt [45](#)
 - records
 - commands for managing CLI sessions [27](#)
 - recovering
 - cluster configurations [184](#)
 - node configurations [182](#)
 - reissuing commands
 - methods of [19](#)
 - rejoining
 - nodes to a cluster [185](#)
 - relationships
 - authentication of cluster peer [267](#)
 - commands for managing cluster peer [265](#)
 - creating cluster peer [284](#)
 - creating intercluster SVM peer [290](#)
 - creating intracluster SVM peer [291](#)
 - creating SVM peer [289](#)
 - deleting cluster peer [287](#)
 - deleting SVM peer [295](#)
 - peering SVMs, described [288](#)
 - release notes
 - support for histogram-based predictions in RAVE [254](#)
 - remote access
 - enabling NIS or LDAP users to access the cluster [120](#)
 - system console of a node [42](#)
 - remote management
 - introduction to using the SP for node [49](#)
 - remote management devices
 - understanding the SP [49](#)
 - removing
 - aggregates for SVMs [94](#)
 - protocols for SVMs [95](#)
 - renaming
 - clusters [34](#)
 - replicated database
 - See* RDB
 - replication rings
 - cluster, description of [33](#)
 - requirements
 - cluster naming when peering [272](#)
 - firewall for cluster peering [272](#)
 - for using AutoSupport
 - to upload large files [205](#)
 - IP addresses for cluster peering [272](#)
 - IPspaces for cluster peering [272](#)
 - network for cluster peering [272](#)
 - ports for cluster peering [272](#)
 - subnets for cluster peering [272](#)
 - respond to
 - alerts [219](#)
 - system health alerts [219](#)
 - restoring
 - cluster configurations [185](#)
 - node configurations [183](#)
 - the root volume of SVMs [103](#)
 - roles
 - commands for managing access control [132](#)
 - considerations for customizing access-control [126](#)
 - considerations for password rule settings [131](#)
 - customizing to restrict user access to specified commands for access control [129](#)
 - introduction to managing access-control [123](#)
 - managing rule settings for user names and passwords in access control [130](#)
 - predefined for cluster administrators [123](#)
 - predefined for SVM administrators [124](#)
 - RBAC
 - managing access-control roles, introduction to [123](#)
 - role-based access control
 - managing access-control roles, introduction to [123](#)
 - root aggregates
 - rules governing node [43](#)
 - root CA certificates
 - installing to authenticate an SSL client of the cluster or SVM [141](#)
 - root volumes
 - creating a new root volume for Infinite Volumes [103](#)
 - freeing up space on node [44](#)
 - promoting a FlexVol volume for SVMs [103](#)
 - promoting a new FlexVol volume for SVMs [105](#)
 - restoring for SVMs with Infinite Volume [106](#)
 - rules governing node [43](#)
 - SVMs, understanding [76](#)
 - rows command
 - setting display preferences in the CLI [21](#)
 - RSH
 - accessing cluster with [15](#)
 - enabling access to the cluster [13](#)
 - rule settings
 - managing for user names and passwords [130](#)
 - rules
 - considerations for password settings [131](#)
- ## S
- sampling resolution
 - changing for the cache emulated by PCS [256](#)
 - schedules
 - commands for managing job [178](#)
 - managing jobs and [177](#)
 - secure data access
 - installing KMIP certificates to mutually authenticate the cluster and an external key manager [147](#)
 - secure keys
 - installing KMIP certificates to mutually authenticate the cluster and an external key manager [147](#)
 - Secure Sockets Layer
 - See* SSL
 - security
 - managing user names and passwords in access-control role [130](#)
 - security certificate commands
 - for managing digital certificates [150](#)
 - security configurations
 - commands for managing SSH [136](#)
 - introduction to managing SSH [134](#)
 - sensors
 - displaying information about environmental [241](#)

- understanding readings and status values of threshold-based SP, in system sensors command output [66](#)
 - understanding status values of discrete SP, in system sensors command output [67](#)
- serial consoles
 - accessing node [42](#)
 - accessing the SP from [59](#)
 - accessing the SP, introduction to [57](#)
- serial ports
 - using to access cluster [11](#)
- server CA certificates
 - installing KMIP certificates to mutually authenticate the cluster and an external key manager [147](#)
 - installing to authenticate cluster or SVM as SSL server [144](#)
- server certificates
 - installing to authenticate the cluster or SVM as an SSL server [139](#)
- servers
 - cluster management, definition of [31](#)
 - managing digital certificates for authentication of clients or [137](#)
- Service Processors
 - See* SPs
- services
 - commands for managing web [155](#)
 - configuring access to web [157](#)
 - requirements for user access to web [154](#)
 - ways to manage web [154](#)
- sessions
 - introduction to managing CLI, (cluster administrators only) [26](#)
 - recording CLI [26](#)
- set command
 - setting display preferences in the CLI [21](#)
- set requests
 - managing audit settings for management activities [160](#)
- setting up
 - SVMs [78](#)
- settings
 - managing rule, for user names and passwords [130](#)
- settings, audit
 - managing for management activities [160](#)
- setup
 - AutoSupport [206](#)
- severity
 - AutoSupport [204](#)
- shells
 - introduction to CLI command [16](#)
- shortcuts
 - keyboard, for editing CLI commands [19](#)
- show command output
 - methods of customizing by using fields [24](#)
- shutting down
 - a node [47](#)
- single-node clusters
 - considerations for [31](#)
 - description of [30](#)
- site license
 - license types and licensed method [174](#)
- size budgets
 - AutoSupport [200](#)
- SMTP
 - troubleshooting AutoSupport message delivery over [215](#)
- Snapshot policies
 - modifying for SVMs [96](#)
- software entitlements
 - managing licenses [173](#)
- sorting data
 - to diagnose or identify issues [234](#)
- space
 - freeing up on a node's root volume [44](#)
- SPs
 - accessing from an administration host [58](#)
 - accessing from the system console [59](#)
 - accessing the system console of a node [42](#)
 - accessing, introduction to [57](#)
 - accounts that can access [58](#)
 - commands for managing a node remotely [62](#)
 - commands for managing from Data ONTAP [70](#)
 - configuring the network manually [53](#)
 - configuring the network, introduction to [51](#)
 - considerations for the SP network configuration [51](#)
 - enabling the automatic network configuration [52](#)
 - managing a node remotely using, introduction to [49](#)
 - managing the IP addresses that can access [60](#)
 - methods of managing firmware updates [56](#)
 - modifying the API service configuration [55](#)
 - relationship among SP CLI, SP console, and system console sessions [60](#)
 - understanding [49](#)
 - understanding readings and status values of threshold-based sensors in system sensors command output [66](#)
 - understanding status values of discrete sensors in system sensors command output [67](#)
 - using online help at CLI [62](#)
 - when the network interface is used for firmware updates [57](#)
- SSH
 - accessing cluster with [11](#)
 - accessing the SP using, introduction to [57](#)
 - commands for managing security configurations [136](#)
 - security configurations, introduction to managing [134](#)
- ssh requests
 - managing the IP addresses that can access the SP [60](#)
- SSL
 - commands for managing [156](#)
 - installing KMIP certificates to mutually authenticate the cluster and an external key manager [147](#)
 - introduction to managing the web protocol engine [152](#)
 - managing digital certificates for server or client authentication [137](#)
 - managing, introduction [156](#)
- SSL clients
 - installing a client certificate to authenticate the cluster or SVM as [145](#)
 - of the cluster or SVM, installing a client CA or root CA certificate to authenticate [141](#)
- SSL servers

- installing a server certificate to authenticate the cluster or SVM as [139](#)
- installing server CA certificate to authenticate cluster or SVM as [144](#)
- SSLv3
 - introduction to managing the web protocol engine [152](#)
- stand-alone nodes
 - considerations for clusters with [31](#)
- standard licenses
 - license types and licensed method [174](#)
- starting
 - SVMs
 - starting [100](#)
- statistics
 - filtering [233](#)
 - sorting [234](#)
 - See also* performance
- statistics commands
 - for monitoring cluster and SVM performance [238](#)
- stopping
 - SVMs
 - stopping [101](#)
- Storage Encryption functionality
 - installing KMIP certificates to mutually authenticate the cluster and an external key manager [147](#)
- Storage QoS
 - assigning storage objects to policy groups [247](#)
 - commands [251](#)
 - creating policy groups [247](#)
 - effect on non-throttled workloads [245](#)
 - examples
 - isolating a workload [249](#)
 - setting a limit on all workloads [251](#)
 - setting a proactive limit on non-critical workloads [250](#)
 - how it helps [242](#)
 - how it works [243](#)
 - how maximum throughput works [245](#)
 - how to monitor workload performance [247](#)
 - maximum number of policy groups [247](#)
 - maximum number of workloads [247](#)
 - monitoring policy group performance [247](#)
 - monitoring workload performance [247](#)
 - rules for assigning storage objects to policy groups [245](#)
 - types of policy groups [243](#)
 - types of workloads [243](#)
 - what it is [242](#)
 - workflow [242](#)
- storage systems
 - introduction to monitoring [192](#)
- subnets
 - enabling the SP automatic network configuration [52](#)
 - requirements for cluster peering [272](#)
- subsystems
 - AutoSupport [201](#)
 - AutoSupport, collection size and time budgets [200](#)
 - introduction to AutoSupport [200](#)
- suggestions
 - how to send feedback about documentation [308](#)
- SVM administrators
 - differences between cluster administrators and [10](#)
 - predefined roles for [124](#)
- SVM banners [166](#)
- SVM MOTDs
 - creating [169](#)
- SVM start fails
 - troubleshooting [100](#)
- SVM stop fails
 - troubleshooting [101](#)
- SVM with FlexVol volumes
 - controlling I/O performance [107](#), [247](#)
- SVM-level banners
 - managing [168](#)
- SVM-level MOTDs
 - managing [170](#)
- SVMs
 - accepting peer relationship [292](#)
 - access and authentication methods [110](#)
 - additional information about administering [114](#)
 - administering from SVM context [102](#)
 - administration capabilities [109](#)
 - benefits of using [76](#)
 - changing password [112](#)
 - configuring [88](#)
 - creating AD user or groups accounts [91](#)
 - creating for a MetroCluster configuration [86](#)
 - creating intercluster peer relationships [290](#)
 - creating intracluster peer relationships [291](#)
 - creating local user accounts [89](#)
 - creating NIS or LDAP user accounts [90](#)
 - creating peer relationships [289](#)
 - creating users accounts for providing SVM administrators access to [89](#)
 - definition of admin [31](#)
 - delegating administration [91](#)
 - deleting [99](#)
 - deleting peer relationship [293](#)
 - deleting peer relationships [295](#)
 - displaying information about [97](#), [112](#)
 - displaying peer relationship [297](#)
 - guidelines for creating [80](#)
 - introduction to monitoring performance [230](#)
 - introduction to monitoring performance using advanced methods [235](#)
 - language configurations [83](#)
 - language options [81](#)
 - logging in [111](#)
 - managing [74](#)
 - managing peer relationships [287](#)
 - modifying [94](#)
 - modifying the attributes of [96](#)
 - modifying the peer relationship [294](#)
 - peer relationship described [288](#)
 - planning to create [78](#)
 - promoting a new FlexVol volume as root volume [105](#)
 - rejecting peer relationship [293](#)
 - renaming [98](#)
 - resuming peer relationships [296](#)
 - root volume, promoting data-protecting mirrors [104](#)
 - root volume, promoting LS mirrors [103](#)
 - setting up [78](#)
 - starting [100](#)
 - states of peer relationships [289](#)
 - stopping [101](#)

- suspending peer relationship [296](#)
 - types of [76](#)
 - understanding root volumes [76](#)
 - viewing network configuration [113](#)
 - with FlexVol volumes, explained [74](#)
 - with Infinite Volume, explained [74](#)
 - SVMs with FlexVol volumes
 - creating [84](#)
 - creating for data access [84](#)
 - explained [74](#)
 - maximum number of [77](#)
 - rules for assigning to Storage QoS policy groups [245](#)
 - SVMs with Infinite Volume
 - creating [87](#)
 - creating a new root volume for [106](#)
 - explained [74](#)
 - maximum number of [77](#)
 - switches
 - introduction to monitoring health of [218](#)
 - troubleshooting discovery of [223](#)
 - verifying the monitoring of using the cluster switch health monitor [224](#)
 - synchronizing
 - nodes with the cluster [187](#)
 - syslog servers
 - forwarding the command history log to a destination [161](#)
 - sysmgr web service
 - ways to manage access to OnCommand System Manager [29](#)
 - system configurations
 - backing up and restoring [180](#)
 - system connectivity health monitor
 - commands for [225](#)
 - system connectivity health monitors
 - about [220](#)
 - system console sessions
 - relationship among SP CLI, SP console sessions, and [60](#)
 - system consoles
 - accessing node [42](#)
 - accessing the SP from [59](#)
 - system health
 - alert customization [219](#)
 - example of responding to degraded [222](#)
 - introduction to monitoring [218](#)
 - monitoring using AutoSupport [195](#)
 - responding to degraded [221](#)
 - system health alerts
 - how they trigger AutoSupport messages and events [220](#)
 - responding to [219](#)
 - system health monitoring
 - commands [225](#)
 - System Manager
 - about [28](#)
 - OnCommand, accessing the cluster by using browser-based graphic interface [27](#)
 - OnCommand, ways to manage access to [29](#)
 - tasks you can perform from [28](#)
 - system panics
 - managing core dump files after [188](#)
 - system performance
 - introduction to managing [242](#)
 - system prompts
 - rebooting a node at the [45](#)
 - system sensors command output
 - understanding readings and status values of threshold-based SP sensors in [66](#)
 - understanding status values of discrete SP sensors in [67](#)
 - system SVMs
 - described [76](#)
 - systems
 - introduction to monitoring [192](#)
 - systemshell
 - introduction to [16](#)
- ## T
- technical support
 - how AutoSupport OnDemand communicates with [203](#)
 - Telnet
 - accessing the cluster by using [14](#)
 - enabling access to the cluster [13](#)
 - threshold-based SP sensors
 - understanding readings and status values of, in system sensors command output [66](#)
 - thresholds
 - displaying information about environmental sensor [241](#)
 - time
 - commands for managing the cluster [164](#)
 - managing the cluster [163](#)
 - time budgets
 - AutoSupport [200](#)
 - TLSv1
 - introduction to managing the web protocol engine [152](#)
 - topologies
 - supported cluster peer network [268](#)
 - transport protocols
 - supported by AutoSupport [205](#)
 - trigger events
 - how AutoSupport responds to [198](#)
 - triggers
 - commands for modifying AutoSupport [211](#)
 - troubleshooting
 - AutoSupport message delivery over HTTP or HTTPS [214](#)
 - AutoSupport message delivery over SMTP [215](#)
 - creating SVMs for a MetroCluster configuration fails [86](#)
 - creating SVMs with FlexVol volumes fails [84](#)
 - creating SVMs with Infinite Volume fails [87](#)
 - delivery status of AutoSupport messages [213](#)
 - managing core dump files for [188](#)
 - remote, understanding the SP [49](#)
 - switch discovery for health monitoring [223](#)
 - web service access problems [158](#)
 - when starting the SVM fails [100](#)
 - when stopping the SVM fails [101](#)
 - tunnels
 - enabling AD users and groups to access the cluster and SVMs [118](#)

- Twitter
 - how to receive automatic notification of documentation changes [308](#)
- two-node clusters
 - description of [30](#)
- two-way authentication
 - installing KMIP certificates to mutually authenticate the cluster and an external key manager [147](#)
 - managing digital certificates for server or client authentication [137](#)

U

- unjoining nodes
 - from the cluster [39](#)
- updates
 - methods of managing SP firmware [56](#)
- uploading
 - core dump files [208](#)
 - performance archive files [209](#)
- user accounts
 - access methods for [117](#)
 - authentication and access methods for SVMs [110](#)
 - authentication methods for [117](#)
 - commands for managing [122](#)
 - considerations for password rule settings [131](#)
 - creating for SVM administrators [89](#)
 - for accessing the SP [58](#)
 - local, creating for SVM administrators [89](#)
 - NIS or LDAP, providing access for SVM administrators [90](#)
 - ways to manage [116](#)
- user names
 - managing rule settings in access-control role [130](#)
- users
 - Active Directory, providing access for SVM administrators [91](#)
 - AD access of the cluster and SVMs, enabling groups and [118](#)
 - enabling NIS or LDAP to access the cluster [120](#)
- UUIDs
 - displaying information about cluster node [33](#)

V

- values
 - rules for specifying in CLI [18](#)
- viewing
 - SVM details [97](#), [112](#)
- volumes
 - creating a new root volume for SVMs with Infinite Volume [106](#)
 - freeing up space on a node's root [44](#)
 - modifying number of volumes in SVMs [96](#)
 - promoting a new FlexVol volume as root volume [105](#)
 - rules governing node root [43](#)
 - SVMs with FlexVol, explained [74](#)

- SVMs with Infinite, explained [74](#)
- Vservers
 - See* SVMs

W

- WAFL external cache
 - caching blocks from random write workloads [258](#)
 - caching sequentially read data [258](#)
 - compared with Flash Pool aggregates [254](#)
 - disabling [255](#)
 - displaying configuration [258](#)
 - displaying usage and access information [259](#)
 - enabling [255](#)
 - low-priority user data blocks [257](#)
 - normal user data blocks [256](#)
 - rearming [259](#)
 - system metadata cache [257](#)
- WAFLcache memory
 - introduction to increasing [254](#)
- watchdogs
 - system, understanding status values of discrete SP sensors in system sensors command output [67](#)
- ways to respond to alerts
 - health monitoring [219](#)
- web access
 - managing SSL, introduction [156](#)
- web browsers
 - accessing node log, core dump, and MIB files using [41](#)
- web protocol engines
 - commands for managing [153](#)
 - introduction to managing [152](#)
- web services
 - accessing the cluster by using OnCommand System Manager browser-based graphic interface [27](#)
 - commands for managing [155](#)
 - configuring access to [157](#)
 - managing access to [151](#)
 - requirements for user access [154](#)
 - troubleshooting access problems [158](#)
 - ways to manage [154](#)
- weekly AutoSupport messages
 - type of data contained within [199](#)
 - when and where they are sent [196](#)
- workloads
 - caching blocks from random write [258](#)
 - controlling performance of [247](#)
 - effect of throttling on non-throttled workloads [245](#)
 - how to monitor performance [247](#)
 - maximum number of [247](#)
 - types of [243](#)
 - what they are [243](#)
- Write Anywhere File Layout
 - See* WAFL
- write workloads
 - caching blocks from random [258](#)