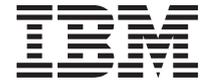


IBM System Storage N series



# Clustered Data ONTAP 8.2 System Administration Guide for Cluster Administrators

SC27-6285-01



# Contents

<b>Preface .....</b>	<b>10</b>
About this guide .....	10
Supported features .....	10
Websites .....	10
Getting information, help, and service .....	11
Before you call .....	11
Using the documentation .....	11
Hardware service and support .....	12
Firmware updates .....	12
How to send your comments .....	12
<b>Differences between cluster and SVM administrators .....</b>	<b>13</b>
<b>Data ONTAP management interface basics .....</b>	<b>14</b>
Accessing the cluster by using the CLI (cluster administrators only) .....	14
Accessing the cluster by using the serial port .....	14
Accessing the cluster by using SSH .....	15
Enabling Telnet or RSH access to the cluster .....	17
Accessing the cluster by using Telnet .....	18
Accessing the cluster by using RSH .....	19
Using the Data ONTAP command-line interface .....	21
Understanding the different shells for CLI commands (cluster administrators only) .....	21
Methods of navigating CLI command directories .....	22
Rules for specifying values in the CLI .....	23
Methods of viewing command history and reissuing commands .....	24
Keyboard shortcuts for editing CLI commands .....	24
Use of administrative privilege levels .....	26
Setting the privilege level in the CLI .....	26
Setting display preferences in the CLI .....	27
Methods of using query operators .....	27
Methods of using extended queries .....	29
Methods of customizing show command output by using fields .....	29
Methods of accessing Data ONTAP man pages .....	30

Managing CLI sessions (cluster administrators only) .....	31
Managing records of CLI sessions .....	31
Managing the automatic timeout period of CLI sessions .....	32
Understanding OnCommand System Manager .....	32
<b>Cluster management basics (cluster administrators only) .....</b>	<b>34</b>
What a cluster is .....	34
Considerations for single node clusters .....	35
What the cluster management server is .....	36
Understanding quorum and epsilon .....	36
What a cluster replication ring is .....	37
Displaying the nodes in a cluster .....	38
Displaying cluster attributes .....	39
Modifying cluster attributes .....	39
Reassigning epsilon to another node in the cluster .....	40
Displaying the status of cluster replication rings .....	41
<b>Managing nodes (cluster administrators only) .....</b>	<b>42</b>
Displaying node attributes .....	42
Modifying node attributes .....	43
Renaming a node .....	43
Adding nodes to the cluster .....	44
Removing nodes from the cluster .....	45
Accessing a node's log, core dump, and MIB files by using a web browser .....	46
Rules governing node root volumes and root aggregates .....	47
Freeing up space on a node's root volume .....	48
Starting and stopping a node .....	49
Rebooting a node at the system prompt .....	49
Booting Data ONTAP at the boot environment prompt .....	50
Rebooting a node remotely .....	50
Shutting down a node .....	51
Managing a node by using the boot menu .....	51
Recovering from a corrupted image of a node's boot device .....	53
Managing a node remotely .....	54
Managing a node remotely by using the Service Processor .....	54
Managing a node remotely by using the Remote LAN Module .....	77
<b>Managing SVMs (cluster administrators only) .....</b>	<b>94</b>
What SVMs are .....	94

About SVMs root volumes .....	96
Types of SVMs .....	96
Why you use SVMs .....	97
Number of SVMs in a cluster .....	98
Creating SVMs .....	98
List of language options .....	99
Language configurations .....	102
Completing the SVM setup worksheet .....	103
Creating SVMs by using the CLI wizard .....	108
Creating SVMs with FlexVol volumes .....	116
Creating SVMs with Infinite Volume .....	118
Considerations for modifying SVMs .....	120
Modifying SVMs .....	122
Delegating administration to SVM administrators .....	123
Displaying information about SVMs .....	126
Renaming SVMs .....	127
Deleting SVMs .....	128
Administering SVMs from the SVM context .....	129
Starting SVMs .....	130
Stopping SVMs .....	131
Restoring the root volume of an SVM .....	132
Promoting a load-sharing mirror copy .....	133
Promoting a data-protection mirror copy .....	134
Promoting a new FlexVol volume .....	135
Controlling and monitoring I/O performance to SVMs by using Storage QoS .....	136
<b>Managing access to the cluster (cluster administrators only) .....</b>	<b>138</b>
Managing user accounts .....	138
Access methods for user accounts .....	139
Authentication methods for user accounts .....	140
Enabling Active Directory domain users to access the cluster .....	141
Commands for managing user accounts .....	142
Managing access-control roles .....	143
Predefined roles for cluster administrators .....	143
Predefined roles for SVM administrators .....	144
Considerations for customizing an access-control role .....	146

Customizing an access-control role to restrict user access to specific commands .....	149
Managing rule settings for user names and passwords in an access-control role .....	151
Considerations for password rule settings .....	152
Commands for managing access-control roles .....	153
Managing firewall service and policies .....	155
Creating a firewall policy and assigning it to a LIF .....	156
Commands for managing firewall service and policies .....	158
Managing SSH security configurations .....	159
Commands for managing SSH security configurations .....	161
Managing public keys .....	161
Commands for managing public keys .....	162
Managing digital certificates for server or client authentication .....	162
Generating and installing a CA-signed digital certificate for server authentication .....	164
Installing a server intermediate certificate .....	166
Providing mutual authentication .....	166
Installing the root certificate of an SSL server .....	169
Commands for managing digital certificates .....	170
Managing access to web services .....	171
Managing the web protocol engine .....	172
Managing web services .....	174
Managing SSL .....	176
Configuring access to web services .....	177
Managing audit settings .....	181
Commands for managing audit settings .....	182
<b>Managing the cluster time (cluster administrators only) .....</b>	<b>183</b>
Commands for managing the cluster time .....	184
<b>Managing licenses (cluster administrators only) .....</b>	<b>186</b>
License types and licensed method .....	186
Commands for managing licenses .....	188
<b>Managing jobs and schedules .....</b>	<b>189</b>
Job categories .....	189
Commands for managing jobs .....	189
Commands for managing job schedules .....	191

<b>Backing up and restoring cluster configurations (cluster administrators only) .....</b>	<b>193</b>
What configuration backup files are .....	193
Managing configuration backups .....	193
How the node and cluster configurations are backed up automatically .....	193
Commands for managing configuration backup schedules .....	194
Commands for managing configuration backup files .....	195
Recovering a node configuration .....	196
Finding a configuration backup file to use for recovering a node .....	196
Restoring the node configuration using a configuration backup file .....	197
Recovering a cluster configuration .....	198
Finding a configuration to use for recovering a cluster .....	198
Restoring a cluster configuration from an existing configuration .....	199
Synchronizing a node with the cluster .....	201
<b>Managing core dumps (cluster administrators only) .....</b>	<b>202</b>
Methods of segmenting core dump files .....	202
Commands for managing core dumps .....	203
Commands for managing core segmenting .....	204
<b>Monitoring the storage system .....</b>	<b>206</b>
Managing event messages .....	206
Setting up the Event Management System .....	207
Commands for managing events .....	209
Managing AutoSupport .....	210
When and where AutoSupport messages are sent .....	211
How AutoSupport creates and sends event-triggered messages .....	212
How AutoSupport On Demand obtains delivery instructions from technical support .....	213
Types of AutoSupport messages and their content .....	214
Structure of AutoSupport messages sent by email .....	215
AutoSupport severity types .....	216
Requirements for using AutoSupport .....	216
Setting up AutoSupport .....	218
Commands for managing AutoSupport .....	220
Information included in the AutoSupport manifest .....	222
What My AutoSupport is .....	222
Troubleshooting AutoSupport .....	222

Monitoring the health of your system .....	227
How health monitoring works .....	227
What health monitors are available .....	229
Receiving system health alerts automatically .....	230
Responding to degraded system health .....	231
Configuring discovery of cluster and management network switches .....	233
Verifying the monitoring of cluster and management network switches ....	234
Commands for monitoring the health of your system .....	235
Using dashboards to display critical system information .....	238
Getting notified of dashboard alarms .....	239
Commands for managing dashboards .....	240
Monitoring cluster performance .....	241
What objects, instances, and counters are .....	241
Decisions to make before you view performance data .....	242
Viewing performance data for a time period .....	243
Viewing continuously updated performance data .....	245
Commands for monitoring clusterSVM performance .....	246
Displaying environmental information .....	248
<b>Managing system performance (cluster administrators only) .....</b>	<b>249</b>
Managing workload performance by using Storage QoS .....	249
How Storage QoS works .....	251
Controlling and monitoring workload performance .....	257
Example: Isolating a workload .....	259
Example: Proactively setting a limit on non-critical workloads .....	261
Example: Proactively setting a limit on workloads in a shared storage infrastructure .....	262
Commands for controlling and monitoring workloads .....	263
Increasing WAFL cache memory .....	265
How Flash Pool aggregates and Flash Cache compare .....	266
Enabling and disabling WAFL external cache .....	267
Caching normal user data blocks .....	267
Caching low-priority user data blocks .....	268
Caching only system metadata .....	268
Displaying the WAFL external cache configuration .....	269
Displaying usage and access information for WAFL external cache .....	269
Preserving the cache in the Flash Cache family of modules .....	270

Improving read performance .....	272
What read reallocation is .....	272
Commands for managing read reallocation .....	273
Improving write performance .....	273
How free space reallocation optimizes free space .....	273
When to enable free space reallocation .....	274
When to use free space reallocation with other reallocation features .....	275
Types of aggregates that free space reallocation can and cannot optimize ..	275
Commands for managing free space reallocation .....	275
<b>Managing peer relationships for data backup and recovery (cluster administrators only) .....</b>	<b>276</b>
Managing cluster peer relationships .....	276
What a cluster peer is .....	276
Connecting one cluster to another cluster in a peer relationship .....	276
Displaying a cluster peer relationship .....	291
Modifying a cluster peer relationship .....	292
Deleting a cluster peering relationship .....	292
Managing jobs on another cluster .....	293
Managing SVM peer relationships .....	295
What an SVM peer relationship is .....	296
States of SVM peer relationships .....	297
Creating an SVM peer relationship .....	298
Accepting an SVM peer relationship .....	300
Rejecting an SVM peer relationship .....	301
Modifying an SVM peer relationship .....	303
Deleting an SVM peer relationship .....	304
Suspending an SVM peer relationship .....	305
Resuming an SVM peer relationship .....	306
Displaying information about SVM peer relationships .....	306
<b>Glossary .....</b>	<b>308</b>
<b>Copyright information .....</b>	<b>315</b>
<b>Trademark information .....</b>	<b>316</b>
<b>Index .....</b>	<b>319</b>

# Preface

---

## About this guide

This document applies to IBM N series systems running Data ONTAP, including systems with gateway functionality. If the terms *Cluster-Mode* or *clustered Data ONTAP* are used in this document, they refer to the Data ONTAP features and functionality designed for clusters, which are different from 7-Mode and prior Data ONTAP 7.1, 7.2, and 7.3 release families.

In this document, the term *gateway* describes IBM N series storage systems that have been ordered with gateway functionality. Gateways support various types of storage, and they are used with third-party disk storage systems—for example, disk storage systems from IBM, HP®, Hitachi Data Systems®, and EMC®. In this case, disk storage for customer data and the RAID controller functionality is provided by the back-end disk storage system. A gateway might also be used with disk storage expansion units specifically designed for the IBM N series models.

The term *filer* describes IBM N series storage systems that either contain internal disk storage or attach to disk storage expansion units specifically designed for the IBM N series storage systems. Filer storage systems do not support using third-party disk storage systems.

## Supported features

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

Information about supported features can also be found on the N series support website (accessed and navigated as described in [Websites](#) on page 10).

## Websites

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

- A listing of currently available N series products and features can be found at the following web page:  
[www.ibm.com/storage/nas/](http://www.ibm.com/storage/nas/)
- The IBM System Storage N series support website requires users to register in order to obtain access to N series support content on the web. To understand how the N series support web

content is organized and navigated, and to access the N series support website, refer to the following publicly accessible web page:

[www.ibm.com/storage/support/nseries/](http://www.ibm.com/storage/support/nseries/)

This web page also provides links to AutoSupport information as well as other important N series product resources.

- IBM System Storage N series products attach to a variety of servers and operating systems. To determine the latest supported attachments, go to the IBM N series interoperability matrix at the following web page:

[www.ibm.com/systems/storage/network/interophome.html](http://www.ibm.com/systems/storage/network/interophome.html)

- For the latest N series hardware product documentation, including planning, installation and setup, and hardware monitoring, service and diagnostics, see the IBM N series Information Center at the following web page:

[publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp](http://publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp)

## Getting information, help, and service

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your IBM N series product, and whom to call for service, if it is necessary.

## Before you call

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

- Check all cables to make sure they are connected.
- Check the power switches to make sure the system is turned on.
- Use the troubleshooting information in your system documentation and use the diagnostic tools that come with your system.
- Refer to the N series support website (accessed and navigated as described in [Websites](#) on page 10) for information on known problems and limitations.

## Using the documentation

The latest versions of N series software documentation, including Data ONTAP and other software products, are available on the N series support website (accessed and navigated as described in [Websites](#) on page 10).

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

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

[publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp](http://publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp)

## Hardware service and support

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

[www.ibm.com/planetwide/](http://www.ibm.com/planetwide/)

## Firmware updates

IBM N series product firmware is embedded in Data ONTAP. As with all devices, ensure that you run the latest level of firmware. Any firmware updates are posted to the N series support website (accessed and navigated as described in [Websites](#) on page 10).

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

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

## How to send your comments

Your feedback helps us to provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, please send them by email to [starpubs@us.ibm.com](mailto:starpubs@us.ibm.com).

Be sure to include the following:

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

# 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. For more information about SVM administrator capabilities, see the *Clustered Data ONTAP System Administration Guide for SVM 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 94

[Predefined roles for cluster administrators](#) on page 143

[Predefined roles for SVM administrators](#) on page 144

## 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.

**Note:** Starting with Data ONTAP 8.2, Element Manager (including ClusterView) is no longer available. You can use the Data ONTAP command-line interface or OnCommand System Manager, a web-based graphical management interface, to manage the cluster. For information about OnCommand System Manager, see the N series support website (accessed and navigated as described in [Websites](#) on page 10).

### Related concepts

[What a cluster is](#) on page 34

[Understanding the different shells for CLI commands \(cluster administrators only\)](#) on page 21

### Related information

[IBM N series support website: www.ibm.com/storage/support/nseries](http://www.ibm.com/storage/support/nseries)

## 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	<b>admin</b>
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

- The Data ONTAP 8.2 release family 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 in-coming 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.
- 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 will result 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::> system services web show
External Web Services: true
                    Status: online
  HTTP Protocol Port: 80
  HTTPS Protocol Port: 443
                    TLSv1 Enabled: true
                    SSLv3 Enabled: true
                    SSLv2 Enabled: false
cluster1::>
```

```
$ ssh -l joe 10.72.137.28 cluster show
Password:
cluster1::> system services web show
External Web Services: true
                    Status: online
  HTTP Protocol Port: 80
  HTTPS Protocol Port: 443
                    TLSv1 Enabled: true
                    SSLv3 Enabled: true
                    SSLv2 Enabled: false
$
```

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::> system services web show
```

```

External Web Services: true
                        Status: online
  HTTP Protocol Port: 80
  HTTPS Protocol Port: 443
                        TLSv1 Enabled: true
                        SSLv3 Enabled: true
                        SSLv2 Enabled: false
cluster1::>

$ ssh -l "DOMAIN1\john" 10.72.137.28 system services web show
Password:
cluster1::> system services web show
External Web Services: true
                        Status: online
  HTTP Protocol Port: 80
  HTTPS Protocol Port: 443
                        TLSv1 Enabled: true
                        SSLv3 Enabled: true
                        SSLv2 Enabled: false
$

```

## 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. However, you can create a new policy by cloning the predefined `mgmt` management firewall policy and then enabling Telnet or RSH under the new policy.

### Steps

1. Use the `system services firewall policy clone` command to create a new management firewall policy based on the `mgmt` management firewall policy.

### Example

```

cluster1::> system services firewall policy clone -policy mgmt
-new-policy-name mgmt1

```

2. Use the `system services firewall policy create` command to enable Telnet or RSH in the new management firewall policy.

**Example**

```
cluster1::> system services firewall policy create -policy mgmt1
-service telnet -action allow -ip-list 0.0.0.0/0
```

```
cluster1::> system services firewall policy create -policy mgmt1
-service rsh -action allow -ip-list 0.0.0.0/0
```

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

**Example**

```
cluster1::> network interface modify -vserver cluster1
-lif cluster_mgmt -firewall-policy mgmt1
```

**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/amd64
login: joe
Password:
cluster1::>
```

## Related concepts

[Managing firewall service and policies](#) on page 155

[Access methods for user accounts](#) on page 139

## 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 system services web  
show  
External Web Services: true  
                          Status: online  
      HTTP Protocol Port: 80  
      HTTPs Protocol Port: 443  
          TLSv1 Enabled: true  
          SSLv3 Enabled: true  
          SSLv2 Enabled: false
```

```
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*. Depending on the task you perform, you might need to use different shells to execute different commands.

- 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.
- The *systemshell* is a low-level shell that is used only for diagnostic and troubleshooting purposes. The *systemshell* is not intended for general administrative purposes. You access the *systemshell* only with guidance from technical support.

## 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]]
```

## Uses of the systemshell and the diagnostic account

A diagnostic account, named “diag”, is provided with your storage system. You can use the `diag` account to perform troubleshooting tasks in the systemshell. The `diag` account and the systemshell are intended only for low-level diagnostic purposes and should be used only with guidance from technical support.

The `diag` account is the only account that can be used to access the systemshell, through the advanced command `system node systemshell`. The `diag` account is locked by default. Before accessing the systemshell for the first time, you must first unlock the `diag` account (`security login unlock`) and then set the `diag` account password (`security login password`). Neither the `diag` account nor the systemshell is intended for general administrative purposes.

## 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 text string, 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::> vsserver create -vsserver vs0 -rootvolume root_vs0 -aggregate myaggr
-ns-switch nis -nm-switch file -language en_US -rootvolume-security-style unix
-comment "My SVM"
cluster1::> vsserver modify -vsserver vs0 -comment ""
```

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

```
cluster1::> security login create -vsserver vs0 -username 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 current active command. You can edit the command by using key combinations. These key combinations are similar to those of the UNIX `tcsh` shell and the Emacs editor.

The following table lists the keyboard shortcuts for editing CLI commands. A caret (^) indicates that you must press the Ctrl key with the specified key.

<b>Edit Command</b>	<b>Action</b>
^b	Move the cursor back one character.
^f	Move the cursor forward one character.
^a	Move the cursor to the beginning of the line.
^e	Move the cursor to the end of the line.
^k	Remove the contents of the edit buffer, from the cursor to the end of the line, and save it in the cut buffer.
^y	Yank the contents of the cut buffer, pushing it into the edit buffer at the cursor.
ESC b	Move the cursor back one word.
ESC f	Move the cursor forward one word.
ESC d	Cut the contents of the edit buffer, beginning at the cursor and continuing to the end of the following word.
^w	Delete the word before the cursor.
^h	Delete the character before the cursor.
Backspace	Delete the character before the cursor.
^d	Delete the character after the cursor.
^p	Replace the current contents of the edit buffer with the previous entry on the history list. For each successive ^p action, the history cursor moves to the previous entry.
^n	Replace the current contents of the edit buffer with the next entry on the history buffer.
Down arrow	Down history.
Up arrow	Up history.
Back arrow	Go backward one character.
Forward arrow	Go forward one character.
^q	TTY start output.
^s	TTY stop output.
^u	Clear the current edit buffer.
^v	Escapes a special mapping for the following character. For instance, to enter a question mark into a command's arguments, press ^v, then press ?.

Edit Command	Action
?	Display context-sensitive help.

## 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

- 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 26

## 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 you are connected to the system through a console connection, the default number of rows is 24. If you are connected to the system through an SSH connection, the number of default rows is determined by the terminal configuration.

- 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.
<	Less-than operator. For example, <code>&lt;20</code> matches any value that is less than 20.
>	Greater-than operator. For example, <code>&gt;5</code> matches any value that is greater than 5.
<=	Less-than-or-equal-to operator. For example, <code>&lt;=5</code> matches any value that is less than or equal to 5.
>=	Greater-than-or-equal-to operator. For example, <code>&gt;=5</code> 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, if you are using a query to identify strings that contain the characters `^.*$`, you must enclose these characters in double quotes ("`^.*$`") 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::>

```

## 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 N series support website (accessed and navigated as described in [Websites](#) on page 10).

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 N series support website (accessed and navigated as described in [Websites](#) on page 10).

### Related information

*IBM N series support website: [www.ibm.com/storage/support/nseries](http://www.ibm.com/storage/support/nseries)*

## 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>

If you want to...	Use this command...
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>

For more information, see the man pages.

## 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>

For more information, see the man pages.

## 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.

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.
- Verify and configure network configuration settings in the storage systems.
- 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.

For more information about System Manager, see the N series support website (accessed and navigated as described in [Websites](#) on page 10).

### **Related information**

*IBM N series support website: [www.ibm.com/storage/support/nseries](http://www.ibm.com/storage/support/nseries)*

## Cluster management basics (cluster administrators only)

---

After a 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.

For information about setting up a cluster and joining nodes to a cluster, see the *Clustered Data ONTAP Software Setup Guide*.

### What a cluster is

You can group pairs of nodes together 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. For details about cluster size limits, see the *IBM System Storage N series Introduction and Planning Guide* at the N series support website (accessed and navigated as described in [Websites](#) on page 10).

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.

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, you must configure cluster high availability (HA). For more information, see the *Clustered Data ONTAP High-Availability Configuration Guide*.

You can create a cluster on a standalone 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.

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. For information about network management for cluster and nodes, see the *Clustered Data ONTAP Network Management Guide*.

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

#### Related concepts

[Understanding quorum and epsilon](#) on page 36

[What the cluster management server is](#) on page 36

**Related tasks**

[Reassigning epsilon to another node in the cluster](#) on page 40

## Considerations for single node clusters

A single node cluster is a special implementation of a cluster running on a standalone 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, by using the Cluster Setup wizard. 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 standalone HA mode. If the node goes offline, clients will not be able to 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 will not be able to recover the cluster unless the cluster configuration backup file is stored at a remote URL.

**Related tasks**

[Adding nodes to the cluster](#) on page 44

[Removing nodes from the cluster](#) on page 45

**Related references**

[Commands for managing configuration backup schedules](#) on page 194

## What the cluster management server is

A cluster management server, also called an *admin SVM*, 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, remote LAN manager, or the 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, you must specify all of the characteristics of the cluster management LIF, including its IP address, netmask, gateway, and port. For more information about creating a cluster, see the *Clustered Data ONTAP Software Setup Guide*.

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 96

## 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, 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*. When 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, if a single link is established between 12 nodes in one room and 12 nodes in another room to compose a 24-node cluster and the link fails, then the group of nodes that holds epsilon would maintain quorum and continue to serve data while the other 12 nodes would stop serving data. However, if the node holding epsilon was unhealthy or offline, then quorum would not be formed, and all of the nodes would stop serving data.

Epsilon is automatically assigned to the first node when the cluster is created. If the node that holds epsilon becomes unhealthy or is taken over by its high availability partner, epsilon does not move to another node but is rather no longer a factor in determining quorum.

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. In a two-node cluster, 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.

## 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 a cluster. There are a number of processes that use the RDB to ensure consistent data across the cluster. These processes include the management application (`mgmt`), volume location database (`vldb`), virtual-interface manager (`vimgr`), and SAN management daemon (`bcomd`).

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 `cluster ring show` displays the status of replication rings and can assist with troubleshooting efforts.

## Displaying the nodes in a cluster

You can display information about the nodes in a cluster and their state.

### Step

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

The command displays the following information:

- Node name
- Whether the node is healthy
- Whether the node is eligible to participate in the cluster
- Whether the node holds epsilon (advanced privilege level or higher only)

### 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
-----
node0          true   true
node1          true   true
node2          true   true
node3          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 wish to continue? (y or 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

After a cluster has been created, you can modify its attributes such as the cluster name, location, and contact information.

### 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 of a cluster must begin with a letter and can include the following special characters: ".", "-", "\_". Any name more than 44 characters in length is truncated.

#### Example of renaming a cluster

The following example renames the current cluster to cluster2:

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

## Reassigning epsilon to another node in the cluster

Only one node in the cluster can hold epsilon. Epsilon gives the holding node an extra fractional voting weight in the quorum.

### About this task

You must follow the steps specifically; otherwise, you can leave the cluster vulnerable to failure or cause data outages.

### Steps

1. If you are currently at the admin privilege level, set the privilege level to advanced by using the `set` command with the `-privilege` parameter.
2. Remove epsilon from the node that holds it currently by using the `cluster modify` command with the `-epsilon` parameter set to `false` for the node.

You can use the `cluster show` command with the `-epsilon` parameter to identify the node that holds epsilon currently.

3. Assign epsilon to another node by using the `cluster modify` command with the `-epsilon` parameter set to `true` for the node.

### Example of reassigning epsilon to another node

The following example removes epsilon from node1 and assigns it to node4:

```
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 show -epsilon *
Node      Health  Eligibility  Epsilon
-----
node1     true    true         true
node2     true    true         false
node3     true    true         false
node4     true    true         false
4 entries were displayed.

cluster1::> cluster modify -node node1 -epsilon false

cluster1::> cluster modify -node node4 -epsilon true

cluster1::> cluster show -epsilon *
Node      Health  Eligibility  Epsilon
-----
node1     true    true         false
node2     true    true         false
node3     true    true         false
```

```
node4      true      true      true
4 entries were displayed.
```

## 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 a remote management device.

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 third-party storage arrays that provide array LUNs for Data ONTAP use. Services and components that are controlled by the node, not by the cluster, can be managed by using the `system node` commands.

A node 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: IBM
      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.

### Step

1. To modify a node's attributes, use the `system node modify` command.

A node's eligibility to participate in the cluster (the `-eligibility` parameter) can be modified by using 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.

### Example of modifying node attributes

The following command modifies the attributes of a node named “node1”. 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 maximum length of a node's name is 47 characters.

### Example

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

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

## Adding nodes to the cluster

After a cluster is created, you can add nodes to it by using the Cluster Setup wizard. You add only one node at a time.

### Before you begin

The following conditions must be met before you add nodes to the cluster:

- 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 cluster, cluster HA must be disabled.  
The *Clustered Data ONTAP High-Availability Configuration Guide* contains information about cluster HA.
- If you are adding a second node to a single-node cluster, the second node must be installed, and the cluster network must be configured.  
The *Hardware and Service Guide* contains details.

### About this task

Nodes must be in even numbers so that they can form HA pairs.

### Steps

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

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

2. Use the Cluster Setup wizard to add the node to the cluster.

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.

For detailed information about the Cluster Setup wizard, see the *Clustered Data ONTAP Software Setup Guide*.

3. If IPv6 is enabled in the cluster, use the `network interface create` command to create the node's node management LIF with an IPv6 address.

### Example

The following command creates node management LIF “mgmt1” with an IPv6 address on node2.

```
cluster1::> network interface create -vserver node2 -lif mgmt1 -role
node-mgmt -home-node node2 -home-port e1a -address 3d20:16fe::101 -
netmask-length 64
```

4. 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. You must also synchronize the time. For detailed information, see the *Clustered Data ONTAP Software Setup Guide*.

## 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.

**Before you begin**

Removing a node from a cluster requires that more than half of the nodes in the cluster be healthy (indicated by `cluster show`).

**Steps**

1. To remove a node from the cluster, use the advanced command `cluster unjoin` from another node in 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.

For information about how to use the `cluster unjoin` command, see the man page.

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.

2. 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.

3. If the node will rejoin the same cluster or join a new cluster, 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.
4. Repeat the preceding steps to remove the failover partner from the cluster.

**After you finish**

If you removed nodes in order to have a single node cluster, you can 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. For more information, see the *Clustered Data ONTAP Network Management Guide*.

**Related tasks**

[Rebooting a node at the system prompt](#) on page 49

## 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 172

[Managing firewall service and policies](#) on page 155

[Managing web services](#) on page 174

[Managing access to web services](#) on page 171

[Access methods for user accounts](#) on page 139

[Managing SSL](#) on page 176

[Managing audit settings](#) on page 181

[Managing core dumps \(cluster administrators only\)](#) on page 202

## Rules governing node root volumes and root aggregates

A node's root volume contains special directories and configuration 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 and reserved for system files, log files, and core files. The directory name is `/mroot`, which is accessible only through the systemshell and with guidance from technical support.

The following rules govern the node's root volume:

- Do not change the preconfigured size for the root volume or modify the content of the root directory, unless technical support instructs you to do so.

The minimum size for a node's root volume depends on the platform model. For information about the minimum size for the root FlexVol volume, see the *IBM System Storage N series Introduction and Planning Guide* at the N series support website (accessed and navigated as described in [Websites](#) on page 10).

Editing configuration files directly in the root directory might result in an adverse impact on the health of the node and possibly the cluster. If you need to modify system configurations, you use Data ONTAP commands to do so.

- 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.

- Do not set the root volume's fractional reserve to any value other than 100%.
- 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 node's root aggregate contains the node's root volume. Starting with Data ONTAP 8.1, new systems are shipped with the root volume in a dedicated, 64-bit root aggregate that contains three disks. By default, a node is set up to use a hard disk drive (HDD) aggregate for the root aggregate. When no HDDs are available, the node is set up to use a solid-state drive (SSD) aggregate for the root aggregate.

The root aggregate must be dedicated to the root volume only. You must not include or create data volumes in the root aggregate.

## 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 `vol0`.

```
node1*> vol status

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

- 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 and stopping a node

You can start a node from the system prompt or boot environment prompt. You can also start a node by using the remote management device (which can be the SP or the RLM depending on the platform model). You can stop a node by halting it at the system prompt.

## 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.

**Step**

1. Reboot a node by using the `system node reboot` command.

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.

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. To access the boot environment prompt from the storage system prompt, use 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>

## Rebooting a node remotely

You can reboot a node remotely by using the remote management device.

### Steps

1. From the administration host, log in to the remote management device of the node you want to reboot by entering the following command:

```
ssh username@IP_for_remote_management_device
```

One of the following remote management device CLI prompts appears, depending on the platform model:

```
SP>
```

```
RLM>
```

2. To power on the node, enter the following command at the CLI prompt for the remote management device:

```
system power on
```

3. To access the system console, enter the following command at the CLI prompt for the remote management device:

```
system console
```

4. If the node does not reboot automatically, enter one of the following commands at the boot environment prompt:

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

## 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.

### Step

1. To shut down a node, use the `system node halt` command.

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` and `-skip-lif-migration` parameters.

### Example of shutting down a node

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** has no effect on systems operating in clustered Data ONTAP.

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

To...	Select...
Continue to boot the node in normal mode	<b>1) Normal Boot</b>
Change the password of the node, which is also the “admin” account password	<b>3) Change Password</b>
Initialize the node's disks and create a root volume for the node	<p><b>4) Clean configuration and initialize all disks</b></p> <p><b>Attention:</b> This menu option erases all data on the disks of the node and resets your node configuration to the factory default settings.</p> <p>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.</p> <p>For a gateway that has a disk shelf, this menu option initializes only the disks on the disk shelf, not the array LUNs. For a gateway that does not have a disk shelf, this menu option initializes the root volume on the storage array.</p>
Perform aggregate and disk maintenance operations and obtain detailed aggregate and disk information.	<p><b>5) Maintenance mode boot</b></p> <p>You exit Maintenance mode by using the <code>halt</code> command.</p>
Restore the configuration information from the node's root volume to the boot device, such as a PC CompactFlash card	<p><b>6) Update flash from backup config</b></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>

To...	Select...
Install new software on a gateway	<p><b>7) Install new software first</b></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 gateway that has no root volume installed. Do <i>not</i> use this menu option to upgrade the Data ONTAP software on either a filer or a gateway.</p>
Reboot the node	<b>8) Reboot node</b>

## Recovering from a corrupted image of a node's boot device

You can recover from a corrupted image of the boot device (such as the CompactFlash card) for a node by using the remote management device.

### Steps

1. Log in to the remote management device by entering the following command at the administration host:

```
ssh username@IP_for_remote_management_device
```

The CLI prompt for the remote management device appears. It can be one of the following, depending on the platform model:

```
SP>
```

```
RLM>
```

2. At the CLI prompt for the remote management device, perform one of the following steps:
  - To reboot the node by using the primary BIOS firmware image, enter the following command:

```
system reset primary
```
  - To reboot the node by using the backup BIOS firmware image, enter the following command:

```
system reset backup
```

The console informs you that the command will cause a “dirty system shutdown” and asks you whether to continue.

3. Enter `y` to continue.

The node shuts down immediately.

## Managing a node remotely

You can manage a node remotely by using a remote management device, which can be the SP or the RLM, depending on the platform model. The device stays operational regardless of the operating state of the node.

The RLM is included in the N6000 series, N7700, and N7900 platforms.

The SP is included in all other supported platform models.

## Managing a node remotely by using the Service Processor

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

The SP provides the following capabilities:

- 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 allows 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 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.

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 `system node autosupport modify` command's `-transport` parameter setting 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.



## Configuring the SP network

Before you can access the SP of a node, the SP network must be configured and enabled. 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.

### Before you begin

To configure 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. For more information about IPv6 configuration, see the *Clustered Data ONTAP Network Management Guide*.

### Steps

1. Configure and enable the SP by using the `system node service-processor network modify` command.
  - The `-address-type` 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 type.
  - 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.
  - 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` specifies the gateway IP address for the SP.

For more information about the `system node service-processor network modify` command, see the man page.
2. Display the SP network configuration to verify the settings by using the `system node service-processor network show` command.

### 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 node service-processor network modify -node local
-address-type IPv4 -enable true -ip-address 192.168.123.98
-netmask 255.255.255.0 -gateway 192.168.123.1

cluster1::> system node 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

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: -
2 entries were displayed.

cluster1::>

```

## Accounts that can access the SP

Cluster user accounts that are created with the `service-processor` application type have access to the SP CLI on any node of the cluster that supports the SP. 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.

**Note:** 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 you configured the SP 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. Enter the following command from the administration host to 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>
```

## Accessing the SP from the serial console

You can access the SP from the serial console to perform monitoring or troubleshooting tasks.

### Steps

1. To access the SP CLI from the serial console, press 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. To exit the SP CLI and return to the serial console, press Ctrl-D and then press Enter.

### Example of accessing the SP CLI from the serial console

The following example shows the result of pressing Ctrl-G from the serial 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 serial 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 serial console.)

```
cluster1::>
```

## Accessing the serial console from the SP

The SP's `system console` command enables you to log in to the serial console from the SP.

### Steps

1. Enter the following command at the SP prompt:  

```
system console
```

The message `Type Ctrl-D to exit` appears.
2. Log in to the console when you are prompted.  

The storage system prompt appears.
3. To exit from the serial console and return to the SP CLI, press `Ctrl-D`.

### Example of accessing the serial console from the SP

The following example shows the result of entering the `system console` command at the SP prompt. The `system node image show` command is entered at the console, followed by pressing `Ctrl-D`, which returns you to the SP prompt.

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

(Log in to the console when you are prompted.)

```
login:
Password:
*****
* This is a SP/RLM console session. Output from the *
* serial console is also mirrored on this session. *
*****
cluster1::>
cluster1::> system node image show
```

(Command output is displayed.)

(Press `Ctrl-D` to exit the serial console and return to the SP CLI.)

```
SP>
```

## Relations among the SP CLI, SP console, and serial console sessions

You can open an SP CLI session to manage a node remotely and a separate SP console session to run Data ONTAP commands remotely. The SP console session mirrors output displayed in a concurrent

serial console session. The SP and the serial console have independent shell environments with independent login authentication.

Understanding how the SP CLI, SP console, and serial console sessions are related helps you manage a node remotely. The following describes the relations 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.
- For security reasons, the SP CLI session and the serial 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 serial console credential. When you access the SP CLI from a serial console session (by pressing Ctrl-G), you are prompted for the SP CLI credential.
- The SP console session and the serial console session have independent shell environments. The SP console session mirrors output that is displayed in a concurrent serial console session. However, the concurrent serial console session does not mirror the SP console session. The SP console session does not mirror output of concurrent SSH sessions.

## 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 node1> 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
rsa - commands for Remote Support Agent
```

```
system - commands to control the system
version - print SP version
```

- 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 node1> 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 at the SP admin privilege level

The SP commands at the admin privilege level enable you to display events, logs, and status information for node power, batteries, sensors, field-replaceable units (FRUs), or the SP itself. The commands also enable you to reboot the node or the SP and create a core dump.

The following SP commands are available at the admin privilege level:

**Note:** Some commands are platform-specific and might not be available on your platform.

If you want to...	Use this command...
Display system date and time	<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>
Exit the SP CLI	<code>exit</code>
Display a list of available commands or subcommands of a specified command	<code>help [<i>command</i>]</code>
Set the privilege level to access the specified mode for the SP CLI	<code>priv set {admin advanced diag}</code>  <b>Attention:</b> You should use advanced or diag commands only under the guidance of technical support.
Display the current privilege level for the SP CLI	<code>priv show</code>

If you want to...	Use this command...
Manage the Remote Support Agent (RSA) if it is installed on the node	<code>rsa</code>
Display the SP log archives or the files in an archive	<code>sp log history show [-archive {latest   all   archive-name}] [-dump {all   file-name}]</code>
Reboot the SP	<code>sp reboot</code>
Display SP status and network configuration information	<code>sp status [-v   -d]</code> <p><b>Note:</b> The <code>-v</code> option displays SP statistics in verbose form. The <code>-d</code> option adds the SP debug log to the display.</p>
Update the SP firmware by using the image at the specified location	<code>sp update image_URL</code> <p><b>Note:</b> <code>image_URL</code> must not exceed 200 characters.</p>
Display the current time, the length of time the system 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 ACP information or the status for expander sensors	<code>system acp [show   sensors show]</code>
Display battery information	<code>system battery show</code>
Log in to the system console	<code>system console</code> <p><b>Note:</b> You use Ctrl-D to exit from the system console and return to the SP CLI.</p>
Create a core dump and reset the node	<code>system core</code> <p><b>Note:</b> This command has the same effect as pressing the Non-maskable Interrupt (NMI) button on a node. The SP stays operational as long as the input power to the node is not interrupted.</p>
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.	<code>system forensics [show   log dump   log clear]</code>

If you want to...	Use this command...
List all system FRUs and their IDs	<code>system fru list</code>
Display product information for the specified FRU	<code>system fru show fru_id</code> <p><b>Note:</b> You can display FRU IDs by using the <code>system fru list</code> command.</p>
Display console logs	<code>system log</code>
Turn the node on or off, or perform a power-cycle (turning the power off and then back on)	<code>system power {on off cycle}</code> <p><b>Note:</b> 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.</p> <p><b>Attention:</b> Using the <code>system power</code> command 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.</p>
Display the status for the power supply	<code>system power status</code>
Reset the node by using the specified BIOS firmware image	<code>system reset {primary backup current}</code> <p><b>Note:</b> The SP stays operational as long as the input power to the node is not interrupted.</p>
Display the status for the environmental sensors, including their states and current values	<code>system sensors</code> <p><b>Note:</b> This command has an equivalent command, <code>system sensors show</code>.</p>
Display the status and details for the specified sensor	<code>system sensors get sensor_name</code> <p><b>Note:</b> You can obtain <code>sensor_name</code> by using the <code>system sensors</code> or the <code>system sensors show</code> command.</p>
Display the SP hardware and firmware version information	<code>version</code>

## Commands for managing a node at the SP advanced privilege level

You can use the SP advanced privilege level to display the SP command history, SP debug file, SP messages file, and data history for field-replaceable units (FRUs). You can also manage the battery firmware and automatic update.

The following SP commands are available only at the advanced privilege level:

If you want to...	Use this command...
Display the SP command history	<code>sp log audit</code>
Display the SP debug information	<code>sp log debug</code>
Display the SP messages file	<code>sp log messages</code>
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>
Update the battery firmware from the image at the specified location	<code>system battery flash <i>image_URL</i></code>  <b>Note:</b> You use this command if the automatic battery firmware upgrade process has failed for some reason.
Compare the current battery firmware image against a specified firmware image	<code>system battery verify [<i>image_URL</i>]</code>  <b>Note:</b> If <i>image_URL</i> is not specified, the default battery firmware image is used for comparison.
Display the FRU data history log	<code>system fru log show</code>

## How to determine the status of a threshold-based SP sensor

Threshold-based sensors take periodic readings of a verity 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 command `system sensors`:


- 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 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. 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 either `nc` (noncritical) or `cr` (critical), 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. `na` means 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 status of a discrete SP sensor

The Status column of the `system sensors` command output in the SL CLI shows the discrete sensors' conditions in hexadecimal values. To interpret the status values of most discrete sensors, you can use the `system sensors get sensor_name` command in the SL CLI.

Discrete sensors do not have thresholds. Their readings (displayed under the Current column in the `system sensors` command output) do not carry actual meanings and thus are ignored by the SP.

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.

While 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. However, 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

<b>0x0180</b>	Timer expired
<b>0x0280</b>	Hard reset
<b>0x0480</b>	Power down
<b>0x0880</b>	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:

<b>0x01xx</b>	220V PSU type
<b>0x02xx</b>	110V PSU type

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

#### Examples of the `system sensors get sensor_name` command output for 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 nodel> 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 nodel> 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]
```

## Troubleshooting a node by using the SP

When you encounter a problem with a node, you can use the SP to display information about the problem, create a core dump, and reboot the node, even if the node's firmware is corrupted.

The following table describes the common SP commands that you can use at the SP prompt to troubleshoot a node:

If this condition occurs...	And you want to...	Enter this command at the SP CLI prompt...
An environmental sensor has reached an abnormal condition.	Display the status for all environmental sensors, their states, and the current values.	<code>system sensors show</code>
	Display the status and details for a specific sensor.	<code>system sensors get <i>sensor_name</i></code>
The node is not responding properly.	Access the system console from the SP.	<code>system console</code>
	Create a core dump and reboot the node.	<code>system core</code>
	Power-cycle the node.	<code>system power cycle</code>
You receive an AutoSupport message indicating an event such as a panic or hardware component failure.	Display what has occurred at the system console.	<code>system log</code>
	Display all events.	<code>events all</code>
	Display a specific number of recent events.	<code>events newest <i>number</i></code>
	Search for specific events regarding <i>keyword</i> .	<code>events search <i>keyword</i></code>
The node firmware is corrupted.	Boot the node by using the backup image of the firmware.	<code>system reset backup</code>
A FRU is malfunctioning.	Display the FRU's product information.	<code>system fru list</code> to list all FRU IDs
		<code>system fru show <i>fru_id</i></code> to display product information for a specific FRU

## Managing the SP with Data ONTAP

You can use Data ONTAP to set up and display the SP configuration, display the SP status, reboot the SP, manage the SP firmware image, and manage access to the SP.

### Methods of managing SP firmware updates

Starting with Data ONTAP 8.2, a baseline SP firmware image is packaged with the Data ONTAP image. By default, the SP automatic update functionality is enabled. You have the option to manually trigger an SP update.

Data ONTAP 8.2 and later releases include an SP firmware image that is called the *baseline image*. You do not need to download the baseline SP firmware image separately. If a new version of the SP firmware becomes subsequently available, you have the option to download it from the N series support website (accessed and navigated as described in [Websites](#) on page 10) and update the SP firmware to the downloaded version without upgrading the Data ONTAP version. For information about manually downloading and updating the SP firmware, see the SP Firmware Download and Installation Instructions on the N series support website (accessed and navigated as described in [Websites](#) on page 10).

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 have failed, you should contact technical support.
  - When you download a version of the SP firmware from the N series support website (accessed and navigated as described in [Websites](#) on page 10) and the downloaded version is newer than the one that the SP is currently running

You have the option to disable the SP automatic update functionality by using the `system node service-processor image modify` command. However, it is best to 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 node 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 node 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 (`-update-type`)
- 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 (`-clear-logs`)

For information about the `system node service-processor image update` command, see the man page.

- Data ONTAP enables you to display the status for the latest SP firmware update by using the `system node 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 automatic or manually triggered.

### Related information

*IBM N series support website: [www.ibm.com/storage/support/nseries](http://www.ibm.com/storage/support/nseries)*

## Restricting SP access to only the specified administration hosts

You can configure the SP to accept SSH requests from only the administration hosts that you specify.

### Step

1. Enter the following command in the nodeshell to specify the administration host or hosts that you want to grant SP access:

```
options sp.ssh.access host_spec
```

You can specify *host\_spec* in the following forms:

- `host[=|!=]host_list`  
*host\_list* is a comma-separated list that includes host names, IP addresses, or IP addresses with a netmask.
- `all` or `*`  
Allows all hosts to access the SP.
- `none` or `-`  
Allows no hosts to access the SP.

The default for *host\_spec* is `*`.

For more information and examples about using this option, see the `na_spaccess(8)` man page in the nodeshell.

**Examples of restricting SP access to only the specified hosts**

The following example grants SP SSH access to the administration host with the specified IP address:

```
node1> options sp.ssh.access host=192.168.123.98
```

The following example grants SP SSH access to two administration hosts, identified by their host names:

```
node1> options sp.ssh.access host=myhost1,myhost2
```

The following example grants SP SSH access to all hosts with their IP address prefix matching 3FFE:81D0:107:2082:

```
node1> options sp.ssh.access host=3FFE:81D0:107:2082::1/64
```

**Configuring automatic logout of idle SSH connections to the SP**

You can configure the automatic logout settings so that an SSH connection to the SP is automatically terminated after the connection has been idle for the number of minutes you specify.

**About this task**

Setting changes for automatic logout of idle SP SSH connections take effect only on SSH sessions that start after the changes.

Automatic logout does not take effect if you access the SP through the serial console.

**Steps**

1. Enter the following command in the nodeshell to enable SSH automatic logout for the SP:

```
options sp.autologout.enable on
```

**Note:** The default is `on`. Setting the option to `off` disables SSH automatic logout for the SP, causing the `sp.autologout.timeout` option to have no effect.

2. Enter the following command in the nodeshell to specify the number of minutes after which an idle SSH connection to the SP is automatically disconnected:

```
options sp.autologout.timeout minutes
```

The default is 60 minutes.

**Example of configuring automatic logout of idle SSH connections to the SP**

The following example configures the SP to automatically disconnect SSH sessions that are idle for 30 minutes or more:

```
node1> options sp.autologout.enable on
node1> options sp.autologout.timeout 30
```

**Data ONTAP commands for managing the SP**

Data ONTAP provides commands for managing the SP, including setting up and displaying the SP network configuration, displaying the current SP status, rebooting the SP, managing the SP firmware image, and managing SSH access to the SP.

You can use the following Data ONTAP commands and nodeshell options to manage the SP:

If you want to...	Use this Data ONTAP command...
<p>Set up or modify the SP network configuration of a node, including the following:</p> <ul style="list-style-type: none"> <li>• The IP address type (IPv4 or IPv6)</li> <li>• Whether the network interface of the specified IP address type should be enabled</li> <li>• If you are using IPv4, whether to use the network configuration from the DHCP server or the network address that you specify</li> <li>• The public IP address for the SP</li> <li>• The netmask for the SP (if using IPv4)</li> <li>• The network prefix-length of the subnet mask for the SP (if using IPv6)</li> <li>• The gateway IP address for the SP</li> </ul>	<pre>system node service-processor network modify</pre>
<p>Display the SP network configuration, including the following:</p> <ul style="list-style-type: none"> <li>• The configured address type (IPv4 or IPv6) and whether it is enabled</li> <li>• The remote management device type</li> <li>• The current SP status and link status</li> <li>• 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</li> </ul>	<pre>system node service-processor network show</pre> <p><b>Note:</b> Displaying complete SP network details requires the <code>-instance</code> parameter.</p>

If you want to...	Use this Data ONTAP command...
<p>Display general SP information, including the following:</p> <ul style="list-style-type: none"> <li>• The remote management device type</li> <li>• The current SP status</li> <li>• Whether the SP network is configured</li> <li>• Network information, such as the public IP address and the MAC address</li> <li>• The SP firmware version and Intelligent Platform Management Interface (IPMI) version</li> <li>• Whether the SP firmware automatic update is enabled</li> </ul>	<p><code>system node service-processor show</code></p> <p><b>Note:</b> Displaying complete SP information requires the <code>-instance</code> parameter.</p>
<p>Reboot the SP on a node and optionally specify the SP firmware image (primary or backup) to use</p>	<p><code>system node service-processor reboot-sp</code></p> <p><b>Attention:</b> 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>
<p>Display the details of the currently installed SP firmware image, including the following:</p> <ul style="list-style-type: none"> <li>• The remote management device type</li> <li>• The partition (primary or backup) that the SP is booted from, its status, and firmware version</li> <li>• Whether the firmware automatic update is enabled and the last update status</li> </ul>	<p><code>system node service-processor image show</code></p> <p><b>Note:</b> The <code>-is-current</code> parameter indicates the partition (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...
Enable or disable the SP automatic firmware update	<pre>system node service-processor image modify</pre> <p><b>Note:</b> 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>
Manually download an SP firmware image on a node	<pre>system node image get</pre> <p><b>Note:</b> 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>
Manually update the SP firmware, by specifying the following: <ul style="list-style-type: none"> <li>• 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.</li> <li>• The installation baseline You can update the SP firmware to the baseline version that is bundled with the currently running version of Data ONTAP.</li> <li>• Whether to update the entire firmware image or only the changed portions</li> <li>• 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</li> </ul>	<pre>system node service-processor image update</pre>

If you want to...	Use this Data ONTAP command...
Display the status for the latest SP firmware update, including the following information: <ul style="list-style-type: none"> <li>• The start and end time for the latest SP firmware update</li> <li>• Whether an update is in progress and the percentage that is complete</li> </ul>	<pre>system node service-processor image update-progress show</pre>
Enable or disable automatic logout of idle SSH connections to the SP	<pre>options sp.autologout.enable</pre> <p><b>Note:</b> This command is available through the nodeshell.</p>
Specify the number of minutes after which an idle SSH connection to the SP is automatically disconnected	<pre>options sp.autologout.timeout</pre> <p><b>Note:</b> This command is available through the nodeshell. For this option to take effect, the <code>sp.autologout.enable</code> option must be set to on.</p>
Restrict SP access to only the specified administration hosts	<pre>options sp.ssh.access</pre> <p><b>Note:</b> This command is available through the nodeshell.</p>

## Disabling SNMP traps for only the SP

You can disable SNMP traps for only the SP and leave SNMP traps for Data ONTAP enabled.

### Step

1. To disable SNMP traps for only the SP, enter the following command in the nodeshell:

```
options sp.snmp.traps off
```

The default is on.

You cannot enable SNMP traps for only the SP when SNMP traps for Data ONTAP is disabled. If you disable SNMP traps for Data ONTAP, SNMP traps for the SP are also disabled.

## Managing a node remotely by using the Remote LAN Module

The Remote LAN Module (RLM) is a remote management device that is provided on the N6000 series, N7700, and N7900 platforms. The RLM provides remote node management capabilities, including remote access, monitoring, troubleshooting, logging, and alerting features.

The RLM stays operational regardless of the operating state of the node. It 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.

The RLM has a single temperature sensor to detect ambient temperature around the RLM board. Data generated by this sensor is not used for any node or RLM environmental policies. It is only used as a reference point that might help you troubleshoot node issues. For example, it might help a remote administrator determine if the node was shut down due to an extreme temperature change.

For instructions about how to cable a node to the RLM, see the procedure about installing or replacing a Remote LAN Module in the *Hardware and Service Guide* for your platform.

- Without the RLM, you can access the node through the serial console or from an Ethernet connection using any supported network interface.  
You use the Data ONTAP CLI to administer the node.
- With the RLM, you can *remotely* access the node through the serial console.  
You use the Data ONTAP CLI to administer the node and the RLM.
- With the RLM, you can also access the node through an Ethernet connection using a secure shell client application.  
You use the RLM CLI to monitor and troubleshoot the node.

If you have a data center configuration where management traffic and data traffic are on separate networks, you can configure the RLM on the management network.

## What the RLM does

The commands in the RLM CLI enable you to remotely access and administer the storage system and diagnose error conditions. Also, the RLM extends AutoSupport capabilities by sending alerts and notifications through an AutoSupport message.

Using the RLM CLI commands, you can perform the following tasks:

- Remotely administer the storage system by using the Data ONTAP CLI through the RLM's system console redirection feature
- Remotely access the storage system and diagnose error conditions, even if the storage system has failed, by performing the following tasks:
  - View the storage system console messages, captured in the RLM's console log
  - View storage system events, captured in the RLM's system event log
  - Initiate a storage system core dump
  - Power-cycle the storage system (or turn it on or off)
  - Reset the storage system
  - Reboot the storage system

The RLM extends AutoSupport capabilities by sending alerts and “down system” or “down filer” notifications through an AutoSupport message when the storage system goes down, regardless of whether the storage system can send AutoSupport messages. Other than generating these messages on behalf of a system that is down, and attaching additional diagnostic information to AutoSupport messages, the RLM has no effect on the storage system's AutoSupport functionality. The AutoSupport configuration settings and message content behavior of the RLM are inherited from Data ONTAP.

**Note:** The RLM does not rely on the `system node autosupport modify` command's `-transport` parameter setting to send notifications. The RLM only uses the Simple Mail Transport Protocol (SMTP) and requires its host's AutoSupport configuration to include mail host information.

In addition to AutoSupport messages, the RLM generates SNMP traps to configured trap hosts for all “down system” or “down filer” events, if SNMP is enabled for the RLM.

The RLM has a nonvolatile memory buffer that stores up to 4,000 system events in a system event log (SEL) to help you diagnose system issues. The event list from the SEL is automatically sent by the RLM to specified recipients in an AutoSupport message. The records contain the following data:

- Hardware events detected by the RLM—for example, system sensor status about power supplies, voltage, or other components
- Errors (generated by the storage system or the RLM) detected by the RLM—for example, a communication error, a fan failure, a memory or CPU error, or a `boot image not found` message
- Critical software events sent to the RLM by the storage system—for example, a system panic, a communication failure, an unexpected boot environment prompt, a boot failure, or a user-triggered “down system” as a result of issuing the `system reset` or `system power cycle` command.

The RLM monitors the storage system console regardless of whether administrators are logged in or connected to the console. When storage system messages are sent to the console, the RLM stores them in the console log. The console log persists as long as the RLM has power from either of the storage system's power supplies. Because the RLM operates with standby power, it remains available even when the storage system is power-cycled or turned off.

Hardware-assisted takeover is available on systems that support the RLM and have the RLM modules set up. For more information about hardware-assisted takeover, see the *Clustered Data ONTAP High-Availability Configuration Guide*.

The RLM supports the SSH protocol for CLI access from UNIX clients and PuTTY for CLI access from PC clients.

## Configuring the RLM network

Before you can access the RLM of a node, the RLM network must be configured and enabled. You can configure the RLM to use IPv4, IPv6, or both. The RLM IPv4 configuration supports static and DHCP addressing, and the RLM IPv6 configuration supports static addressing only.

### Before you begin

To configure IPv6 connections for the RLM, IPv6 must already be configured and enabled for Data ONTAP. The `network options ipv6` commands manage IPv6 settings for Data ONTAP. For more information about IPv6 configuration, see the *Clustered Data ONTAP Network Management Guide*.

## Steps

1. Configure and enable the RLM by using the `system node service-processor network modify` command.
  - The `-address-type` parameter specifies whether the IPv4 or IPv6 configuration of the RLM is to be modified.
  - The `-enable` parameter enables the network interface of the specified IP address type.
  - 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.  
If `-dhcp` is set to `v4`, you do not need to specify the `-ip-address`, `-netmask`, `-prefix-length`, or `-gateway` parameter.
  - The `-ip-address` parameter specifies the public IP address for the RLM.
  - The `-netmask` parameter specifies the netmask for the RLM (if using IPv4.)
  - The `-prefix-length` parameter specifies the network prefix-length of the subnet mask for the RLM (if using IPv6.)
  - The `-gateway` specifies the gateway IP address for the RLM.

For more information about the `system node service-processor network modify` command, see the man page.
2. Display the RLM network configuration to verify the settings by using the `system node service-processor network show` command.

### Example of configuring the RLM network

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

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

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

                Node: node1
            Address Type: IPv4
    Interface Enabled: true
            Type of Device: RLM
                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
```

```

Node: node1
Address Type: IPv6
Interface Enabled: false
Type of Device: RLM
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: -
2 entries were displayed.

cluster1::>

```

## Accounts that can access the RLM

Cluster user accounts that are created with the `service-processor` application type have access to the RLM CLI on any node of the cluster that supports the RLM. RLM user accounts are managed from Data ONTAP and authenticated by password.

User accounts for accessing the RLM are managed from Data ONTAP instead of the RLM CLI. A cluster user account of any role can access the RLM 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 RLM supports only password authentication.

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

**Note:** 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 RLM.

### Related concepts

[Managing user accounts](#) on page 138

[Access methods for user accounts](#) on page 139

## Logging in to the RLM from an administration host

You can log in to the RLM from an administration host to perform administrative tasks remotely, if the host has a Secure Shell client application that supports SSHv2 and your account name is configured with the `service-processor` application type.

### About this task

If the RLM is running firmware version 4.0 or later and is configured to use an IPv4 address, the RLM rejects SSH login requests and suspends all communication with the IP address for 15 minutes

if five SSH login attempts fail repeatedly within 10 minutes. The communication resumes after 15 minutes, and you can try to log in to the RLM again.

### Steps

1. Enter the following command from the UNIX host:

```
ssh username@RLM_IP_address
```

2. When you are prompted, enter the password for *username*.

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

#### Examples of RLM access from an administration host

The following example shows how to log in to the RLM with a user account, “joe”, which has been set up on the storage system to access the RLM.

```
ssh joe@192.168.123.98
```

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

```
ssh joe@fd22:8b1e:b255:202::1234
```

```
ssh joe@fd22:8b1e:b255:202:2a0:98ff:fe01:7d5b
```

### Accessing the serial console from the RLM

The RLM's `system console` command enables you to log in to the serial console from the RLM.

#### Steps

1. Enter the following command at the RLM prompt:

```
system console
```

The message `Type Ctrl-D to exit` appears.

2. Log in to the console when you are prompted.
3. To exit the serial console and return to the RLM CLI, press Ctrl-D.

#### Example of accessing the serial console from the RLM

The following example shows the result of entering the `system console` command at the RLM prompt. The `system node image show` command is entered at the console, followed by Ctrl-D, which returns you to the RLM prompt.

```
RLM> system console
Type Ctrl-D to exit.
```

(Log in to the console when you are prompted.)

```
login:
Password:
*****
* This is a SP/RLM console session. Output from the *
* serial console is also mirrored on this session. *
*****
cluster1::>
cluster1::> system node image show
```

(Command output is displayed.)

(Press Ctrl-D to exit the storage serial console and return to the RLM CLI.)

```
RLM>
```

## Relations among the RLM CLI, RLM console, and serial console sessions

You can open an RLM CLI session to manage a node remotely and a separate RLM console session to run Data ONTAP commands remotely. The RLM console session mirrors output displayed in a concurrent serial console session. The RLM and the serial console have independent shell environments with independent login authentication.

Understanding how the RLM CLI, RLM console, and serial console sessions are related helps you manage a node remotely. The following describes the relations among the sessions:

- Only one administrator can log in to the RLM CLI session at a time; however, the RLM enables you to open both an RLM CLI session and a separate RLM console session simultaneously. The RLM CLI is indicated with the RLM prompt (`RLM>`). From an RLM CLI session, you can use the RLM `system console` command to initiate an RLM console session. At the same time, you can start a separate RLM CLI session through SSH. If you press Ctrl-D to exit from the RLM console session, you automatically return to the RLM CLI session. If an RLM CLI session already exists, a message asks you whether to terminate the existing RLM CLI session. If you enter “y”, the existing RLM CLI session is terminated, enabling you to return from the RLM console to the RLM CLI. This action is recorded in the RLM event log.
- For security reasons, the RLM CLI session and the serial console session have independent login authentication.
 

When you initiate an RLM console session from the RLM CLI (by using the RLM `system console` command), you are prompted for the serial console credential.
- The RLM console session and the serial console session have independent shell environments.

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

## Using online help at the RLM CLI

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

### Steps

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

- `help`
- `?`

### Example

The following example shows the RLM CLI online help:

```
RLM node1> help
date - print date and time
exit - exit from the RLM command line interface
events - print system events and event information
help - print command help
priv - show and set user mode
rlm - commands to control the RLM
rsa - commands for Remote Support Agent
system - commands to control the system
version - print RLM version
```

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

`help RLM_command`

### Example

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

```
RLM node1> 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 at the RLM admin privilege level

You can perform most RLM tasks at the admin privilege level. For example, you can display system events and status information for environmental sensors, reboot the storage system or the RLM, and create a system core dump.

The following RLM commands are available at the admin privilege level:

If you want to...	Use this command at the RLM CLI...
Display system date and time	<code>date</code>
Display storage system events logged by the RLM	<code>events {all   info   newest   oldest   search string }</code>
Exit the RLM CLI	<code>exit</code>
Display a list of available commands or subcommands of a specified command	<code>help [command]</code>
Set the privilege level to access the specified mode	<code>priv set {admin   advanced   diag}</code>
Display the current privilege level	<code>priv show</code>
Reboot the RLM	<code>rlm reboot</code>
Display the RLM environmental sensor status	<code>rlm sensors [-c]</code>  <b>Note:</b> The <code>-c</code> option, which takes a few seconds to display, shows current values rather than cached values.
Display RLM status	<code>rlm status [-v -d]</code>  <b>Note:</b> The <code>-v</code> option displays verbose statistics. The <code>-d</code> option displays RLM debug information.
Update the RLM firmware	<code>rlm update http://path [-f]</code>  <b>Note:</b> The <code>-f</code> option issues a full image update.
Manage the RSA if it is configured for the node	<code>rsa</code>
Log in to the Data ONTAP CLI	<code>system console</code>  <b>Note:</b> Pressing Ctrl-D returns you to the RLM CLI.

If you want to...	Use this command at the RLM CLI...
Dump the system core and reset the system	<pre>system core</pre> <p><b>Note:</b> This command has the same effect as pressing the Non-maskable Interrupt (NMI) button on a storage system. The RLM stays operational as long as input power to the storage system is not interrupted.</p>
Turn on or turn off the storage system, or perform a power-cycle (which turns off system power and then turns it back on)	<pre>system power {on off cycle}</pre> <p><b>Note:</b> Standby power stays on, even when the storage system is off. During power-cycling, a brief pause occurs before power is turned back on.</p> <p><b>Attention:</b> Using the <code>system power</code> command to turn off or power-cycle the storage system might cause an improper shutdown of the system (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.</p>
Display status for each power supply, such as presence, input power, and output power	<pre>system power status</pre>
Reset the storage system using the specified BIOS firmware image	<pre>system reset {primary backup current}</pre> <p><b>Note:</b> The RLM stays operational as long as input power to the storage system is not interrupted.</p>
Display the RLM version information, including hardware and firmware information	<pre>version</pre>

### Commands for managing a node at the RLM advanced privilege level

In addition to using the RLM admin commands, you can use the RLM advanced privilege level to display RLM command history, RLM debug and message files, status of environmental sensors, and RLM statistics.

The following RLM commands are available only at the advanced privilege level:

If you want to display...	Use this command at the RLM CLI...
RLM command history or search for audit logs from the system event log (SEL)	<code>rlm log audit</code>
RLM debug file	<code>rlm log debug</code>
RLM message file	<code>rlm log messages</code>
List of environmental sensors, their states, and their current values	<code>system sensors</code>
RLM statistics	<code>rlm status -v</code>

## Managing the RLM with Data ONTAP

You can use Data ONTAP to set up and display the RLM configuration, display the RLM status, reboot the RLM, manage the RLM firmware image, and manage access to the RLM.

### Data ONTAP commands for managing the RLM

Data ONTAP provides commands for managing the RLM, including setting up and displaying the RLM network configuration, displaying the current RLM status, rebooting the RLM, managing the RLM firmware image, and managing SSH access to the RLM.

You can use the following Data ONTAP commands and nodeshell options to manage the RLM:

If you want to...	Use this Data ONTAP command...
Set up or modify the RLM network configuration of a node, including the following: <ul style="list-style-type: none"> <li>• The IP address type (IPv4 or IPv6)</li> <li>• Whether the network interface of the specified IP address type should be enabled</li> <li>• If you are using IPv4, whether to use the network configuration from the DHCP server or the network address that you specify</li> <li>• The public IP address for the RLM</li> <li>• The netmask for the RLM (if using IPv4)</li> <li>• The network prefix-length of the subnet mask for the RLM (if using IPv6)</li> <li>• The gateway IP address for the RLM</li> </ul>	<pre>system node service-processor network modify</pre>

If you want to...	Use this Data ONTAP command...
<p>Display the RLM network configuration, including the following:</p> <ul style="list-style-type: none"> <li>• The configured address type (IPv4 or IPv6) and whether it is enabled</li> <li>• The remote management device type</li> <li>• The current RLM status and link status</li> <li>• 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</li> </ul>	<pre>system node service-processor network show</pre> <p><b>Note:</b> Displaying complete RLM network details requires the <code>-instance</code> parameter.</p>
<p>Display general RLM information, including the following:</p> <ul style="list-style-type: none"> <li>• The remote management device type</li> <li>• The current RLM status</li> <li>• Whether the RLM network is configured</li> <li>• Network information, such as the public IP address and the MAC address</li> <li>• The RLM firmware version</li> </ul>	<pre>system node service-processor show</pre> <p><b>Note:</b> Displaying complete RLM information requires the <code>-instance</code> parameter.</p>
<p>Reboot the RLM on a node and optionally specify the RLM firmware image (primary or backup) to use</p>	<pre>system node service-processor reboot-sp</pre>
<p>Display the details of the currently installed RLM firmware image, including the following:</p> <ul style="list-style-type: none"> <li>• The remote management device type</li> <li>• The partition (primary or backup) that the RLM is booted from, its status, and firmware version</li> </ul>	<pre>system node service-processor image show</pre> <p><b>Note:</b> The <code>-is-current</code> parameter indicates the partition (primary or backup) that the RLM is currently booted from, not whether the installed firmware version is most current.</p>
<p>Download an RLM firmware image on a node</p>	<pre>system node image get</pre>

If you want to...	Use this Data ONTAP command...
<p>Update the RLM firmware, by specifying the following:</p> <ul style="list-style-type: none"> <li>The RLM firmware package to use You can have the RLM use a specific RLM 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 RLM firmware package) that are available on a node.</li> <li>Whether to update the entire firmware image or only the changed portions</li> <li>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 RLM</li> </ul>	<p><code>system node service-processor image update</code></p> <p><b>Note:</b> The RLM firmware does not support automatic firmware update. Setting the <code>-baseline</code> parameter to <code>true</code> is not supported on a node that has the RLM.</p>
<p>Display the status for the latest RLM firmware update, including the following information:</p> <ul style="list-style-type: none"> <li>The start and end time for the latest RLM firmware update</li> <li>Whether an update is in progress and the percentage that is complete</li> </ul>	<p><code>system node service-processor image update-progress show</code></p>
<p>Enable or disable automatic logout of idle SSH connections to the RLM</p>	<p><code>options rlm.autologout.enable</code></p> <p><b>Note:</b> This command is available through the nodeshell.</p>
<p>Specify the number of minutes after which an idle SSH connection to the RLM is automatically disconnected</p>	<p><code>options rlm.autologout.timeout</code></p> <p><b>Note:</b> This command is available through the nodeshell. For this option to take effect, the <code>rlm.autologout.enable</code> option must be set to <code>on</code>.</p>
<p>Restrict RLM access to only the specified administration hosts</p>	<p><code>options rlm.ssh.access</code></p> <p><b>Note:</b> This command is available through the nodeshell.</p>

## Restricting RLM access to only the specified administration hosts

You can configure the RLM to accept SSH requests from only the administration hosts that you specify.

### Before you begin

Your system must be running RLM firmware 4.1 or later for the RLM access control to be supported. For information about downloading and updating the RLM firmware, see the *Clustered Data ONTAP Upgrade and Revert/Downgrade Guide*.

### Step

1. Enter the following command in the nodeshell to specify the administration host or hosts that you want to grant RLM access:

```
options rlm.ssh.access host_spec
```

You can specify *host\_spec* in the following forms:

- `host[=|!=]host_list`  
*host\_list* is a comma-separated list that includes host names, IP addresses, or IP addresses with a netmask.
- `all` or `*`  
Allows all hosts to access the RLM.
- `none` or `-`  
Allows no hosts to access the RLM.

The default for *host\_spec* is `*`.

For more information and examples about using this option, see the `na_rlmaccess(8)` man page in the nodeshell.

### Examples of restricting RLM access to only the specified hosts

The following example grants RLM SSH access to the administration host with the specified IP address:

```
node1> options rlm.ssh.access host=192.168.123.98
```

The following example grants RLM SSH access to two administration hosts, identified by their host names:

```
node1> options rlm.ssh.access host=myhost1,myhost2
```

The following example grants RLM SSH access to all hosts with their IP address prefix matching `3FFE:81D0:107:2082:`

```
node1> options rlm.ssh.access host=3FFE:81D0:107:2082::1/64
```

## Configuring automatic logout of idle SSH connections to the RLM

You can configure the automatic logout settings so that an SSH connection to the RLM is automatically terminated after the connection has been idle for the number of minutes you specify.

### Before you begin

Your system must be running RLM firmware version 4.1 or later for the automatic logout configuration to be supported. For information about downloading and updating the RLM firmware, see the *Clustered Data ONTAP Upgrade and Revert/Downgrade Guide*.

### About this task

Setting changes for automatic logout of idle RLM SSH connections take effect only on SSH sessions that start after the changes.

Automatic logout does not take effect if you access the RLM through the serial console.

### Steps

1. Enter the following command in the nodeshell to enable SSH automatic logout for the RLM:

```
options rlm.autologout.enable on
```

**Note:** The default is `on`. Setting the option to `off` disables SSH automatic logout for the RLM, causing the `rlm.autologout.timeout` option to have no effect.

2. Enter the following command in the nodeshell to specify the number of minutes after which an idle SSH connection to the RLM is automatically disconnected:

```
options rlm.autologout.timeout minutes
```

The default is 60 minutes.

### Example of configuring automatic logout of idle SSH connections to the RLM

The following example configures the RLM to automatically disconnect SSH sessions that are idle for 30 minutes or more:

```
node1> options rlm.autologout.enable on
node1> options rlm.autologout.timeout 30
```

## RLM and SNMP traps

If SNMP is enabled for the RLM, the RLM generates SNMP traps to configured trap hosts for all "down system" events.

You can enable SNMP traps for both Data ONTAP and the RLM. You can also disable the SNMP traps for only the RLM and leave the SNMP traps for Data ONTAP enabled.

For information about SNMP traps, see the *Clustered Data ONTAP Network Management Guide*.

## Disabling SNMP traps for only the RLM

You can disable SNMP traps for only the RLM and leave SNMP traps for Data ONTAP enabled.

### Step

1. To disable SNMP traps for only the RLM, enter the following command in the nodeshell:

```
options rlm.snmp.traps off
```

The default is on.

You cannot enable SNMP traps for only the RLM when SNMP traps for Data ONTAP is disabled. If you disable SNMP traps for Data ONTAP, SNMP traps for the RLM are also disabled.

## Troubleshooting a node by using the RLM

When you encounter a problem with a node, you can use the RLM to display information about the problem, create a core dump, and reboot the node, even if the node's firmware is corrupted.

The following table describes the RLM commands that you can use to troubleshoot a node:

If this condition occurs...	And you want to...	Enter this command at the RLM CLI prompt...
You receive an AutoSupport message indicating an event such as a panic or hardware component failure.	Display what has occurred at the storage system console.	<code>system log</code>
	Display all events.	<code>events all</code>
	Display a specific number of recent events.	<code>events newest <i>number</i></code>
	Search for specific events in the SEL.	<code>events search <i>string</i></code>

If this condition occurs...	And you want to...	Enter this command at the RLM CLI prompt...
The node is not responding properly.	Access the system console from the RLM.	<code>system console</code>
	Create a core dump and reboot the node.	<code>system core</code>
	Power-cycle the node.	<code>system power cycle</code>
The node firmware is corrupted.	Boot the node by using a backup copy of the node firmware.	<code>system reset backup</code>

## Troubleshooting RLM connection problems

If you are having difficulty connecting to the RLM, you should verify that you are using a secure shell client and that the IP configuration is correct.

### Steps

1. Verify that you are using a secure shell client to connect to the RLM.
2. From the Data ONTAP CLI, verify that the RLM is online and the IP configuration is correct by using the `system node service-processor network show` command.
3. From the administration host, test the network connection for the RLM by using the `ping rlm_IP_address` command.
4. If the ping fails, do one of the following:
  - Verify that the RLM network port on the back of the node is cabled and active. For more information, see the Installation and Setup Instructions for your storage system.
  - Modify the RLM configuration to use a valid IP address by using the `system node service-processor network modify` command.
  - Verify that the administration host has a route to the RLM.
5. Reboot the RLM by using the `system node service-processor reboot-sp` command.

**Note:** It takes approximately one minute for the RLM to reboot.
6. If the RLM does not reboot, repeat Steps 2 through 5. If the RLM still does not reboot, contact technical support for assistance.

## 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.

For more information about SVM administrator capabilities, see the *Clustered Data ONTAP System Administration Guide for SVM 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.

## 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 (SMB 1.0) 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.

## About SVMs root volumes

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 root volume namespace and SAN clients' data access is not dependent on the root volume namespace.

The root volume serves as the entry point to the namespace provided by that SVM. The root volume of an 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.

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

## Types of SVMs

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

A cluster contains the following types of SVMs:

- Admin SVM
- Node SVM
- Data SVM

The cluster setup process automatically creates the admin SVM for the cluster. A node SVM is created when the node joins the cluster. The admin SVM represents the cluster, and node SVM represents the individual nodes of the cluster.

The 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 without regard to physical storage or controller, similar to any storage system. When you use SVMs, they provide benefits such as nondisruptive operation, scalability, security and support unified storage.

SVMs have the following benefits:

- **Nondisruptive operation**  
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 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 by using NAS protocols, such as CIFS and NFS, and block-level data access by using 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 (SMB 1.0) 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 as the SVM administrator has to 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 *IBM System Storage N series Introduction and Planning Guide* at the N series support website (accessed and navigated as described in [Websites](#) on page 10).

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 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.

SVMs with Infinite Volume do not exist in a SAN cluster and they 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.

## Creating SVMs

Cluster administrators can create SVMs with FlexVol volumes and SVMs with Infinite Volume in a cluster to serve data to the clients.

You can use one of the following methods to create SVMs:

- The `vserver setup` command enables you to create fully configured SVMs with FlexVol volumes that can serve data immediately.

**Note:** You cannot create fully configured SVMs with Infinite Volumes by using the `vserver setup` command.

With the `vserver setup` command, you can quickly set up SVMs by following the prompts of the wizard.

- The `vserver create` command enables you to create SVMs with FlexVol Volumes or SVMs with Infinite Volume with the root volume and basic configuration, such as name service switch, name mapping switch, and root volume security style.

You must run the various commands to fully configure the SVMs to serve data after creating SVMs by using the `vserver create` command.

Before you create SVMs, you must understand the various requirements and gather the required information such as language setting option for an SVM.

### Choices

- [List of language options](#) on page 99
- [Language configurations](#) on page 102
- [Completing the SVM setup worksheet](#) on page 103
- [Creating SVMs by using the CLI wizard](#) on page 108
- [Creating SVMs with FlexVol volumes](#) on page 116
- [Creating SVMs with Infinite Volume](#) on page 118

## 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

<b>Language values</b>	<b>Languages</b>
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
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)

<b>Language values</b>	<b>Languages</b>
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)
zh.GBK.UTF-8	Simplified GBK Chinese with UTF-8

Language values	Languages
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.

## Completing the SVM setup worksheet

Before you start the Vserver Setup wizard to create and configure a Storage Virtual Machine (SVM), you must gather the required information to complete the wizard successfully.

**Note:** You can create and configure only SVMs with FlexVol volumes by using the Vserver Setup wizard.

The Vserver Setup wizard has the following subwizards, which you can run after you create an SVM:

- Network setup
- Storage setup
- Services setup
- Data access protocol setup

Each subwizard has its specific requirements, depending on the types of services, protocols, and the protocol traffic.

You can use the following worksheet to record values for the setup process:

### SVM information

Types of information	Your values
<p><i>SVM name</i></p> <p>The name of an SVM can contain alphanumeric characters and the following special characters: ".", "-", and "_". However, the name of an SVM should not start with a number or the following special characters: "." and "-".</p> <p>The maximum number of characters allowed in an SVM name is 47.</p> <p><b>Note:</b> SVM names must be unique. You must use the fully qualified domain name (FQDN) of the SVM or another convention that ensures unique SVM names.</p>	
<p><i>Data protocols</i></p> <p>Protocols that you want to configure or allow on that SVM</p>	
<p><i>Client services</i></p> <p>Services that you want to configure on the SVM</p>	

Types of information	Your values
<p><i>Aggregate name</i></p> <p>Aggregate on which you want to create the root volume for an SVM. The default aggregate name is used if you do not specify one.</p>	
<p><i>Language setting</i></p> <p>The default language 'C.UTF-8 ' is used if you do not specify one.</p> <p>The language is set for an SVM. The language of the SVM determines default language setting for volumes in that SVM.</p> <p><b>Note:</b> The language of an SVM is inherited by its volumes if the language is not specified when creating the volumes.</p> <p>For all the available language options and language configurations, see <a href="#">list of language options</a> on page 99 and <a href="#">language configurations</a> on page 102.</p>	
<p><i>SVM root volume's security style</i></p> <p>Determines the type of permissions that can be used to control data access to a volume</p> <p>For more information about the security styles, see the <i>Clustered Data ONTAP File Access Management Guide for CIFS</i> or <i>Clustered Data ONTAP File Access Management Guide for NFS</i>.</p>	

### Information for creating volumes on the SVM

Types of information	Values
<p><i>Volume name</i></p> <p>The default volume name is used if you do not specify one.</p>	
<p><i>Aggregate name</i></p> <p>Aggregate on which you want to create the volume. The default aggregate name is used if you do not specify one.</p>	
<p><i>Volume size</i></p>	

Types of information	Values
<p><i>Volume junction path</i></p> <p>The default junction path is used if you do not specify one.</p>	

### Information for creating an IP network interface on the SVM

Types of information	Values
<p><i>LIF name</i></p> <p>The default LIF name is used if you do not specify one.</p>	
<p><i>Protocols</i></p> <p>Protocols that can use the LIF</p> <p><b>Note:</b> Protocols that can use the LIF cannot be modified after the LIF is created.</p>	
<p><i>Home node</i></p> <p>Home node is the node on which you want to create a LIF. The default home node is used if you do not specify one.</p>	
<p><i>Home port</i></p> <p>Home port is the port on which you want to create a LIF. The default home port is used if you do not specify one.</p>	
<p><i>IP address</i></p>	
<p><i>Network mask</i></p>	
<p><i>Default gateway IP address</i></p>	

### Information for creating an FC network interface on the SVM

Types of information	Values
<p><i>LIF name</i></p> <p>The default LIF name is used if you do not specify one.</p>	

Types of information	Values
<p><i>Protocols</i></p> <p>Protocols that can use the LIF</p> <p><b>Note:</b> Protocols that can use the LIF cannot be modified after the LIF is created.</p>	
<p><i>Home node</i></p> <p>Home node is the node on which you want to create a LIF. The default home node is used if you do not specify one.</p>	
<p><i>Home port</i></p> <p>Home port is the port on which you want to create a LIF. The default home port is used if you do not specify one.</p>	

### Information for configuring LDAP

Types of information	Values
<i>LDAP server IP address</i>	
<p><i>LDAP server port number</i></p> <p>The default LDAP server port number is used if you do not specify one.</p>	
<i>LDAP server minimum bind authentication level</i>	
<i>Bind domain name and password</i>	
<i>Base domain name</i>	

### Information for configuring NIS

Types of information	Values
<i>NIS domain name</i>	
<i>IP addresses of the NIS servers</i>	

### Information for configuring DNS

Types of information	Values
<i>DNS domain name</i>	

Types of information	Values
<i>IP addresses of the DNS servers</i>	

**Note:** You do not need to enter any information to configure NFS on an SVM. The NFS configuration is created when you specify the protocol value as `nfs`.

### Information for configuring CIFS protocol

Types of information	Values
<i>Domain name</i>	
<p><i>CIFS share name</i></p> <p>The default CIFS share name is used if you do not specify one.</p> <p><b>Note:</b> You must not use space characters or Unicode characters in CIFS share names. You can use alphanumeric characters and any of the following special characters: ! @ # \$ % &amp; ( ) _ ' { } . ~ -.</p>	
<p><i>CIFS share path</i></p> <p>The default CIFS share path is used if you do not specify one.</p>	
<p><i>CIFS access control list</i></p> <p>The default CIFS access control list is used if you do not specify one.</p>	

### Information for configuring iSCSI protocol

Types of information	Values
<p><i>igroup name</i></p> <p>The default igroup name is used if you do not specify one.</p>	
<i>Names of the initiators</i>	
<i>Operating system type of the initiator</i>	
<p><i>LUN name</i></p> <p>The default LUN name is used if you do not specify one.</p>	

Types of information	Values
<i>Volume for LUN</i> Volume that is to be used for the LUN	
<i>LUN size</i>	

### Information for configuring Fibre Channel (FC) protocol (FCoE included)

Types of information	Values
<i>igroup name</i> The default igroup name is used if you do not specify one.	
<i>Worldwide port number (WWPN) of the initiators</i>	
<i>Operating system type of the initiator</i>	
<i>LUN name</i> The default LUN name is used if you do not specify one.	
<i>Volume for LUN</i> Volume that is to be used for the LUN	
<i>LUN size</i>	

## Creating SVMs by using the CLI wizard

You can create and configure Storage Virtual Machines (SVMs) with FlexVol volumes fully to start serving data immediately or with minimal configuration to delegate administration to the SVM administrator by using the `vserver setup` command.

### Before you begin

You must have understood the [requirements and gathered the required information](#) on page 103 before you start the Vserver Setup wizard or any of the subwizards.

### About this task

By using the `vserver setup` command, which launches a CLI wizard, you can perform the following tasks:

- Creating and configuring SVMs fully

- Creating and configuring SVMs with minimal network configuration
- Configuring existing SVMs
  - Setting up a network interface
  - Provisioning storage by creating volumes
  - Configuring services
  - Configuring protocols

**Note:** When you select NDMP as one of the protocols for protocol configuration, NDMP is added to the allowed list of protocols of the SVM. The Vserver Setup wizard does not configure the NDMP protocol.

When you start the Vserver Setup wizard, the following is displayed:

```
cluster1::>vserver setup
Welcome to the Vserver Setup Wizard, which will lead you through
the steps to create a virtual storage server that serves data to clients.

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 your answers to previous questions, and
"exit" if you want to quit the Vserver Setup Wizard. Any changes
you made before typing "exit" will be applied.

You can restart the Vserver Setup Wizard by typing "vserver setup". To accept a
default or omit a question, do not enter a value.

Vserver Setup wizard creates and configures only data Vservers.
If you want to create a Vserver with Infinite Volume use the vserver create
command.
```

## Steps

1. Depending on your requirements, perform the following tasks:

If you want to...	Perform the following task...
Set up an SVM by using the Vserver Setup wizard	<a href="#">Creating a fully configured Vserver</a> on page 110
Set up a network interface for an existing SVM	<a href="#">Setting up a network interface</a> on page 112
Provision storage by creating volumes on an existing SVM	<a href="#">Provisioning storage by creating volumes</a> on page 113
Configure services for an existing SVM	<a href="#">Configuring services</a> on page 114
Configure protocols for an existing SVM	<a href="#">Configuring protocols</a> on page 115

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

You can view the attributes of the SVM in detail by using the `vserver show -instance` command.

**Example**

The following example shows how to display information about all existing SVMs:

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

Vserver	Type	Admin State	Root Volume	Aggregate	Name Service	Name Mapping
vs1.example.com	data	running	root_vol1	aggr1	file	file
cluster1	admin	-	-	-	-	-
cluster1-01	node	-	-	-	-	-
cluster1-02	node	-	-	-	-	-
vs2.example.com	data	running	root_vol2	aggr2	file	file

5 entries were displayed.

**Result**

When an SVM is created, its root volume of 1 GB size is created. When you set up an SVM, 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.

**After you finish**

To delegate the administration to an SVM administrator, you must set up a password, unlock the vsadmin user account, create a LIF for accessing the SVM, and enable the firewall policy for managing the SVM.

For SVMs with FlexVol volumes, 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. For more information about creating load-sharing mirror copy, see the *Clustered Data ONTAP Logical Storage Management Guide*.

**Related tasks**

[Delegating administration to SVM administrators](#) on page 123

[Displaying information about SVMs](#) on page 126

**Related references**

[Commands for managing user accounts](#) on page 142

**Creating fully configured SVMs**

You can create a fully configured Storage Virtual Machine (SVM) with FlexVol volumes or an SVM with minimal configuration by using the `vserver setup` command.

**About this task**

The `vserver setup` command prompts you to create and configure an SVM in the following sequence:

1. Create an SVM
2. Create data volumes
3. Create logical interfaces
4. Configure services
5. Configure protocols

To create and configure the SVMs, you must follow all the prompts of the wizard. To create an SVM with minimal configuration, you can follow the prompts of the required sections and then exit the wizard.

### Steps

1. Use the `vserver setup` command to launch the Vserver Setup wizard.

### Example

The following example shows how to create a fully configured SVM:

```
cluster1::>vserver setup
Welcome ...

Step 1. Create a Vserver.
You can type "back", "exit", or "help" at any question.
.....
Enter the Vserver name: vs2.example.com
.....
```

2. Follow the prompts to complete the Vserver Setup wizard:
  - To accept the default value for a prompt, press Enter.
  - To enter your own value for the prompt, enter the value and then press Enter.
3. Use the `vserver show` command to verify the newly created SVM.

You can view the attributes of the SVM in detail by using the `vserver show -instance` command.

### Example

The following example shows how to display information about all existing SVMs:

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

Vserver Mapping	Type	Admin State	Root Volume	Aggregate	Name Service	Name
vs1.example.com	data	running	root_voll	aggr1	file	file

```

cluster1      admin - - - -
cluster1-01   node - - - -
cluster1-02   node - - - -
vs2.example.com data running root_vol2 aggr2 file file
5 entries were displayed.

```

**Result**

When an SVM is created, its root volume of 1 GB size is created. When you set up an SVM, 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.

**After you finish**

- To delegate the administration to an SVM administrator, you must set up a password, unlock the vsadmin user account, create a LIF for accessing the SVM, and enable the firewall policy for managing the SVM.
- If you want to change the role of the default user vsadmin, you must set the password for vsadmin user before changing the role.
- For SVMs with FlexVol volumes, 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. For more information about creating load-sharing mirror copy, see the *Clustered Data ONTAP Logical Storage Management Guide*.

**Setting up a network interface**

You can set up a network interface on an existing Storage Virtual Machine (SVM) by using the `vserver setup -network` command.

**About this task**

You cannot set up IPv6 network interface by using the Vserver Setup wizard.

**Steps**

1. Use the `vserver setup -network` command to set up a network interface for an SVM.

**Example**

The following example shows how to set up a network interface for an SVM:

```

cluster1::> vserver setup -vserver vs2.example.com -network true

Welcome...

Step 1. Create a Vserver.
You can type "back", "exit", or "help" at any question.

```

```

Choose the Vserver data protocols to be configured {nfs, cifs, fcp,
iscsi, ndmp}
[nfs,cifs,fcps,iscsi,ndmp]:
Choose the Vserver client services to be configured {ldap, nis, dns}:

Vserver vs2.example.com's allowed protocol list has been modified to
nfs,cifs,fcps,iscsi,ndmp

Step 2: Create a logical interface.
You can type "back", "exit", or "help" at any question.

Do you want to create a logical interface? {yes, no} [yes]:

```

2. Follow the prompts to complete the Vserver Setup wizard:

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

### Provisioning storage by creating volumes

You can provision storage for an existing Storage Virtual Machine (SVM) by using the `vserver setup -storage` command.

#### Steps

1. Use the `vserver setup -storage` to provision storage for an SVM:

#### Example

The following example shows how to create volumes for an existing SVM:

```

cluster1::> vserver setup -vserver vs2.example.com -storage true

Welcome....

Step 1. Create a Vserver.
You can type "back", "exit", or "help" at any question.

Choose the Vserver data protocols to be configured {nfs, cifs, fcp,
iscsi, ndmp}
[nfs,cifs,fcps,iscsi,ndmp]:
Choose the Vserver client services to be configured {ldap, nis, dns}:

Vserver vs2.example.com's allowed protocol list has been modified to
nfs,cifs,fcps,iscsi,ndmp

Step 2: Create a data volume
You can type "back", "exit", or "help" at any question.

Do you want to create a data volume? {yes, no} [yes]:
.....

```

2. Follow the prompts to complete the Vserver Setup wizard:
  - To accept the default value for a prompt, press Enter.
  - To enter your own value for the prompt, enter the value and then press Enter.

## Configuring services

You can configure services for an existing SVM by using the `vserver setup -services` command.

### About this task

If you want to configure your SVM to use **ldap** or **nis** as the name service (-ns-switch), you should also include **file** as a name service. Including **file** as a name service enables the SVM user account authentication through the SVM's local administrative repository if the ldap or nis authentication fails.

### Steps

1. Use the `vserver setup -services` command to configure services:

#### Example

The following example shows how to configure services for an SVM:

```
cluster1::> vserver setup -vserver vs2.example.com -services ldap

Welcome ...

Step 1. Create a Vserver.
You can type "back", "exit", or "help" at any question.

Choose the Vserver data protocols to be configured {nfs, cifs, fcp,
iscsi, ndmp}
[nfs,cifs,fcf,iscsi,ndmp]:
Choose the Vserver client services to be configured {ldap, nis, dns}:

Vserver vs2.example.com's allowed protocol list has been modified to
nfs,cifs,fcf,iscsi,ndmp

Step 2: Configure LDAP (Lightweight Directory Access Protocol).
You can type "back", "exit", or "help" at any question.

Do you want to configure LDAP? {yes, no} [yes]:
.....
```

2. Follow the prompts to complete the Vserver Setup wizard:
  - To accept the default value for a prompt, press Enter.
  - To enter your own value for the prompt, enter the value and then press Enter.

## Configuring protocols

You can configure protocols for an existing Storage Virtual Machine (SVM) by using the `vserver setup -protocols` command.

### About this task

You can configure only the protocols that are allowed on an SVM.

When you specify the protocols value as `fc`, you can configure both FC and FCoE for an SVM. When you select NDMP as one of the protocols for protocol configuration, NDMP is added to the allowed list of protocols of the SVM. The Vserver Setup wizard does not configure the NDMP protocol.

### Steps

1. Use the `vserver setup -protocols` to configure the allowed protocols for an SVM.

### Example

The following example shows how to configure protocols:

```
cluster1::> vserver setup -vserver vs2.example.com -protocols iscsi
Welcome....

Step 1. Create a Vserver.
You can type "back", "exit", or "help" at any question.

Choose the Vserver client services to be configured {ldap, nis, dns}:
Vserver vs2.example.com's allowed protocol list has been modified to
nfs,cifs,fc,fcpl,iscsi,ndmp

Step 2: Configure iSCSI.
You can type "back", "exit", or "help" at any question.

Do you want to configure iSCSI? {yes, no} [yes]:
.....
```

2. Follow the prompts to complete the Vserver Setup wizard:
  - To accept the default value for a prompt, press Enter.
  - To enter your own value for the prompt, enter the value and then press Enter.

## Creating SVMs with FlexVol volumes

You can create Storage Virtual Machines (SVMs) with FlexVol volumes with basic configuration by using the `vserver create` command. SVMs with FlexVol volumes can serve data to the NAS clients and SAN hosts.

### Before you begin

- The cluster must have at least one non-root aggregate created by using the `aggr create` command.
- You must have at least 1 GB of space on the aggregate for the SVM root volume.
- You must have synchronized the time across the cluster by configuring and enabling NTP to prevent CIFS creation and authentication failures.

### About this task

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

### Step

1. Use the `vserver create` command with the `is-repository` parameter set to `false`. If you do not use the `is-repository` parameter, then the default value `false` is used.

When you create an SVM, the language is set for the SVM and is inherited by all its volumes if the volumes are created without any language setting. The language of the SVM determines the default language setting for volumes in that SVM.

Language is an optional parameter of the `vserver create` command. If you do not use this parameter, the default value `C.UTF-8` (or `POSIX.UTF-8`) is used. For more information about the available language options and SVM language configurations, see [list of language options](#) on page 99 and [language configurations](#) on page 102.

SVM names must be unique. When creating an SVM, you must use the fully qualified domain name (FQDN) of the SVM or another convention that ensures unique SVM names.

The maximum number of characters allowed in an SVM name is 47. To name an SVM, you can use alphanumeric characters and the following special characters: ".", "-", and "\_". However, the name of an SVM should not start with a number or the following special characters: "." and "-".

If you want to configure your SVM to use `ldap` or `nis` as the name service (`-ns-switch`), you should also include `file` as a name service. Including `file` as a name service enables the SVM user account authentication through the SVM's local administrative repository if the `ldap` or `nis` authentication fails.

For more information about this command, see the man pages.

**Note:** SVM create operation might fail due to any intermediate operation failures such as volume creation failure. As a result, the SVM will be in initializing state. It is best to delete

such SVMs because you cannot perform other SVMs operations on that SVM. For example, you cannot create an SVM peering relationship with SVMs in initializing state.

The following example shows how to create a Storage Virtual Machine (SVM, formerly known as Vserver) vs0.example.com with FlexVol volume:

```
cluster1::>vserver create -vserver vs0.example.com -rootvolume
root_vs0 -aggregate
aggr1 -ns-switch nis -rootvolume-security-style unix -language
C.UTF-8 -is-repository false

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

## Result

When an SVM is created, its root volume of size 1 GB is created.

The SVM root volume does not contain any of the configuration information of the SVM. Optionally, the SVM root volume can also contain user data. However, it is best not to store user data in the root volume.

When you create an SVM, 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.

## After you finish

- You must specify the aggregates for an SVM for all the volume related operations that require aggregate name.
- To delegate the SVM administration to an SVM administrator, you must set up a password and unlock the vsadmin user account.
- To provide data access from the SVM, you must create volumes and LIFs and then configure LIFs, protocols, and services.

You can either use the `vserver setup` command or the relevant protocols and services commands to configure the SVM.

For more information about creating and configuring LIFs, see the *Clustered Data ONTAP Network Management Guide*.

For more information about configuring NAS protocols and services, see the *Clustered Data ONTAP File Access Management Guide for CIFS* and *Clustered Data ONTAP File Access Management Guide for NFS*.

For more information about configuring SAN protocols, see the *Clustered Data ONTAP SAN Administration Guide*.

- 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. For more information about creating load-sharing mirror copy, see the *Clustered Data ONTAP Logical Storage Management Guide*.

### Related concepts

[Managing the cluster time \(cluster administrators only\)](#) on page 183

### Related tasks

[Delegating administration to SVM administrators](#) on page 123

[Creating SVMs by using the CLI wizard](#) on page 108

[Displaying information about SVMs](#) on page 126

[Modifying SVMs](#) on page 122

### Related references

[Commands for managing user accounts](#) on page 142

## Creating SVMs with Infinite Volume

You can create Storage Virtual Machines (SVMs) with Infinite Volume to serve data to the NAS clients by using the `vserver create` command.

### Before you begin

- You must have at least 1 GB of space on the aggregate for the SVM root volume.
- You must have synchronized the time across the cluster by configuring and enabling NTP to prevent CIFS creation and authentication failures.
- You must know what language the SVM will use.  
The language of an SVM with Infinite Volume cannot be modified.

### About this task

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

Although you must specify `mixed`, `unix`, or `ntfs` as the root volume security style, the value is not significant for an SVM with Infinite Volume because the security style is not inherited by the Infinite Volume. The security style applies only to the SVM root volume.

### Step

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

When you create an SVM, the language is set for the SVM and is inherited by its Infinite Volume. Language is an optional parameter. If you do not use this parameter, the default value `C.UTF-8` (or `POSIX.UTF-8`) is used.

The root volume security style is not significant for an SVM with Infinite Volume because the security style is not inherited by the Infinite Volume. The security style applies only to the SVM root volume.

SVM names must be unique. When creating an SVM, you must use the fully qualified domain name (FQDN) of the SVM or another convention that ensures unique SVM names.

The maximum number of characters allowed in an SVM name is 47. To name an SVM, you can use alphanumeric characters and the following special characters: ".", "-", and "\_". However, the name of an SVM should not start with a number or the following special characters: "." and "-".

If you want to configure your SVM to use `ldap` or `nis` as the name service (`-ns-switch`), you should also include `file` as a name service. Including `file` as a name service enables the SVM user account authentication through the SVM's local administrative repository if the `ldap` or `nis` authentication fails.

### Example

The following example illustrates how to create the SVM `vs1` with Infinite Volume:

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

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

### Result

When an SVM is created, its root volume of size 1 GB is created.

The SVM root volume does not contain any of the configuration information of the SVM. Optionally, the SVM root volume can also contain user data. However, it is best not to store user data in the root volume.

When you create an SVM, 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.

### After you finish

- You must specify the aggregates for an SVM for all the volume related operations that require aggregate name.
  - To delegate the SVM administration to an SVM administrator, you must set up a password and unlock the `vsadmin` user account.
  - After you create an SVM, you must create an Infinite Volume for the SVM.
- For more information about creating an Infinite Volume, see the *Clustered Data ONTAP Infinite Volumes Management Guide*.

- You can use the relevant protocols and services commands to configure the SVM.
- To provide data access from the SVM, you must create volumes and LIFs and then configure LIFs, protocols, and services.

You can use the relevant protocols and services commands to configure the SVM.

For more information about creating and configuring LIFs, see the *Clustered Data ONTAP Network Management Guide*.

For more information about configuring NAS protocols and services, see the *Clustered Data ONTAP File Access Management Guide for CIFS* and *Clustered Data ONTAP File Access Management Guide for NFS*.

## Considerations for modifying SVMs

When modifying a Storage Virtual Machine (SVM), a cluster administrator must understand the significance of the SVM attributes such as aggregate list and maximum number of volumes on the SVM. If the attributes such as aggregate list is not set for an SVM, the SVM administrator cannot perform volume operations that require aggregate name on that SVM.

### SVMs with FlexVol volumes

You must be aware of the following attributes and their effects when modifying an SVM with FlexVol volumes:

- Name Service Switch
- Name Mapping Switch
- Snapshot policy
- Antivirus policy
- Quota policy
- Admin state

You can set the admin state of the SVM at the advanced privilege level if the operations such as starting or stopping an SVM fails.

- QoS policy group
- Maximum number of volumes that can be created on the SVM

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.

**Note:** This parameter is effective only when you specify the list of aggregates for an SVM.

- Language  
The language of an SVM is set when the SVM is created. `C.UTF-8` is the default language option. When you modify the SVM's language, the language setting of the existing volumes does not change. When a new volume is created without specifying the language, it inherits the SVM's language.
- List of the aggregates available to create volumes

You must specify the aggregate names for an SVM which allows an SVM administrator to view the list of available aggregates to perform any provisioning operations that require an aggregate name, for example, creating a volume or a FlexClone. When you specify the aggregate names for an SVM, you can perform the limited provisioning operations same as the SVM administrator. However, you can move the volumes and copy the volumes across aggregates.

If you do not specify the aggregate names for an SVM, the SVM administrator cannot perform any provisioning operations that require an aggregate name. As a cluster administrator, you can perform all the operations that require an aggregate name.

- Allowed protocols list

When you specify the list of allowed protocols, the remaining protocols are added to the disallowed protocols list automatically. Only the allowed protocols can be configured to serve data from an SVM.

**Note:** Only the protocols that have been licensed on the cluster can be part of the allowed protocols.

- Disallowed protocols list

The disallowed protocols are not available for configuration and cannot serve data. When you disallow the protocol, you cannot modify the state of the protocol.

If you add NDMP to the disallowed protocols list, you cannot establish NDMP sessions.

For the detailed description of all the parameters, see the man pages.

## SVMs with Infinite Volume

You must be aware of the following attributes and their effects when modifying an SVM with Infinite Volume:

- Name Service Switch
- Name Mapping Switch
- Snapshot policy
- List of the aggregates available to create an Infinite Volume

You must specify the aggregate names for the SVM with Infinite Volume because the aggregates list determines the aggregates that are used by the Infinite Volume when created.

If you do not specify any aggregates for an SVM with Infinite Volume, then the Infinite Volume spans across all the aggregates in the cluster if created by the cluster administrator. However, an SVM administrator sees the empty aggregate list and will not have enough aggregates to create the Infinite Volume.

- Admin state

You can set the admin state of the SVM at the advanced privilege level if the operations such as starting or stopping an SVM fails.

- Allowed protocols list

Only NFS and CIFS are allowed in the protocol list.

- Disallowed protocols list

- If NFS and CIFS are a part of the disallowed protocols list, then the SVM with Infinite Volume cannot serve data.

**Note:** You cannot modify the language, quota policy, and antivirus policy of an SVM with Infinite Volume.

**Note:** You cannot modify the language, quota policy, and maximum number of volumes of an SVM with Infinite Volume.

## Modifying SVMs

You can modify a Storage Virtual Machine (SVM) and its attributes such as maximum number of volumes, aggregate list, and allowed protocols by using the `vserver modify` command.

### Before you begin

You must have understood the [various attributes](#) on page 120 that can be modified for the SVM with FlexVol volumes and SVM with Infinite Volume and the significance of these attributes.

### Steps

1. Use the `vserver modify` command to modify the attributes of an SVM.

#### Example

The following example shows how to modify a Snapshot policy named `daily`, add the comment "Sales team access," modify the quota policy to `poll`, and modify allowed protocols to `nfs`, `cifs`, and `ndmp` for a Storage Virtual Machine (SVM, formerly known as Vserver) named `vs8.example.com`:

```
cluster1::>vserver modify -vserver vs8.example.com -allowed-protocols
nfs,cifs,ndmp -snapshot-policy daily
-comment "Sales team access" -quota-policy poll
```

For more information about this command, see the man pages.

2. Use the `vserver show` command to verify the modified attributes of the SVM.

#### Example

The following example shows how to display the detailed information of the SVM `vs8.example.com`:

```
cluster1::> vserver show -instance -vserver vs8.example.com

                               Vserver: vs8.example.com
                               Vserver Type: data
                               Vserver UUID: 6f181736-33a5-11e2-
bbb6-123478563412
                               Root Volume: root_vs0
```

```

Snapshot Policy: daily
                Comment: Sales team access
                Quota Policy: poll
                .
                .
                Allowed Protocols: nfs, cifs, ndmp
                Disallowed Protocols: fcp, iscsi
Is Vserver with Infinite Volume: false
QoS Policy Group: -

```

### Related tasks

[Displaying information about SVMs](#) on page 126

## Delegating administration to SVM administrators

After setting up a functional Storage Virtual Machine (SVM) or an SVM with basic network configuration, you can optionally delegate the administration of the SVM to an SVM administrator. You can delegate 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. Optional: 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. Optional: If the SVM does not have any assigned aggregates, use the `vserver modify` command to specify aggregates in the aggregates list of an SVM.

### Example

The following example shows how to specify the aggregates `aggr1` and `aggr2` for the SVM `vs1.example.com`:

```
vserver modify -vserver vs1.example.com -aggr-list aggr1,aggr2
```

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

### Example

The following example shows how to specify the maximum number of volumes for the SVM `vs1.example.com`:

```
vserver modify -vserver vs1.example.com -max-volumes 10
```

- Use the `vserver modify` command to allow or disallow protocols for an SVM.

### Example

The following example shows how to disallow protocols for an SVM `vs1.example.com`:

```
vserver modify -vserver vs1.example.com -disallowed-protocols ndmp
```

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

- Depending on the type of protocols, enter the appropriate command to create a management LIF for an SVM:

If you want to...	Then...
Create a new LIF for SVM management	<p>Use the <code>network interface create</code> command.</p> <p><b>Note:</b> A dedicated SVM management LIF is required for SAN protocols, where data and management protocols cannot share the same LIF. SVM management LIF can be created only on data ports. You can use the <code>network port show</code> command to determine the data ports.</p> <p>The following example shows how to create a data LIF <code>lif3</code> for the SVM <code>vs1.example.com</code> to support SVM management:</p> <pre>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>
Use a LIF for NFS, CIFS, and SVM management	<p>Change the firewall policy to <code>mgmt</code> by using the <code>network interface modify</code> command.</p> <p>The following example shows how to modify a data LIF <code>lif1</code> for the SVM <code>vs1.example.com</code> to support SVM management:</p> <pre>network interface modify -vserver vs1.example.com -lif lif1 -firewall-policy mgmt</pre>

- Depending on the type of SVM administrator roles, perform the appropriate action:

---

**If you want to use... Then...**


---

vsadmin, a predefined role that is created and is in the locked state when an SVM is created.

You must set up a password and unlock the user account to delegate the SVM administration.

- a. Use the `security login password` command to set up a password
  - a. Enter a password for the user account.
  - b. Reenter the password to confirm.

The following example shows how to set up a password for the user account vsadmin on the SVM vs1.example.com:

```
cluster1::>security login password -username
vsadmin -vserver vs1.example.com
Please enter a password for user 'vsadmin':
Please enter it again:
```

```
cluster1::>
```

- b. Use the `security login unlock` command to unlock the user account.

The following example shows how to unlock the user account vsadmin for the SVM vs1.example.com:

```
security login unlock -username vsadmin -vserver
vs1.example.com
```

---

Any customized role or other predefined roles, such as vsadmin-volume, vsadmin-protocol, or vsadmin-readonly

- a. Use the `security login create` command to create a user account with a role.
  - a. Enter a password for the user account.
  - b. Reenter the password to confirm.

The following example shows how to create user account vsadmin-monitor with vsadmin-readonly role for the SVM vs1.example.com:

```
cluster1::> security login create -username user1
-application ssh -authmethod password -vserver
vs1.example.com -role vsadmin-readonly
Please enter a password for user 'vsadmin-monitor':
Please enter it again:
```

```
cluster1::>
```

---

For more information about these commands, see the man pages.

## Result

After you assign an 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.

## 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.

For more information about this command, see the man pages.

The following example displays detailed information about all Storage Virtual Machines (SVMs, formerly known as Vservers):

```
cluster1::>vserver show

Vserver          Type      Admin   Root   Aggregate  Name   Name
-----          -
vs1.example.com  data     running root_vol1 aggr1     file   file
cluster1         admin    -       -      -         -     -
cluster1-01      node     -       -      -         -     -
cluster1-02      node     -       -      -         -     -
vs2.example.com  data     running root_vol2 aggr2     file   file
5 entries were displayed.

cluster1::> vserver show -instance

Vserver: vs1.example.com
Vserver Type: data
Vserver UUID: 49294a39-e762-11df-8768-123478563412
Root Volume: root_vol1
Aggregate: aggr1
.
.
Allowed Protocols: nfs
Disallowed Protocols: cifs, fcp, iscsi, ndmp

Vserver: cluster1
Vserver Type: admin
Vserver UUID: 00000000-0000-0000-0000-000000000000
Root Volume: -
Aggregate: -
.
.
```

```

Allowed Protocols: -
Disallowed Protocols: -

cluster1::> vserver show -vserver vs2.example.com

Vserver: vs1
Vserver Type: data
Vserver UUID: ca34e6b2-ddec-11df-b066-123478563412
Root Volume: root_vol2
.
.
Allowed Protocols: iscsi
Disallowed Protocols: nfs, cifs, fcp,
ndmp

```

## 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.

## 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.

## Starting SVMs

You can provide data access from a Storage Virtual Machine (SVM) by starting the SVM. You can start an SVM by using the `vserver start` command.

### About this task

When you start an SVM, the protocols that were stopped either when the SVM was stopped or stopped independently by running commands such as `vserver fcp stop`, will start serving data.

### Step

1. Use the `vserver start` command to start an SVM.

### Example

The following example shows how to start the SVM `vs1.example.com`:

```
cluster1::> vserver start -vserver vs1.example.com
[Job 71] Job succeeded: DONE

cluster1::> vserver show
Vserver          Type      Admin  Root  Aggregate  Name  Name
-----
vs1.example.com  data     running  root_voll  aggr1     file  file
cluster1         admin    -        -        -         -    -
cluster1-01      node    -        -        -         -    -
cluster1-02      node    -        -        -         -    -
```

For more information about `vserver start` command, see the man pages.

## Result

The SVM is in `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. For more information about data policies and JSON requirements for data policies, see the *Clustered Data ONTAP Infinite Volumes Management Guide*.

## Related tasks

[Displaying information about SVMs](#) on page 126

# Stopping SVMs

You can stop a Storage Virtual Machine (SVM) if you need to troubleshoot or delete the SVM, or stop the data access from the SVM by using the `vserver stop` command.

## Before you begin

All clients connected to the SVM must be disconnected.

**Attention:** If any clients are connected to an 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 the HA pair.

When you stop an SVM, other operations such as SnapMirror data transfers continue to run as per the schedule.

## Step

1. Use the `vserver stop` command to stop an SVM.

## Example

The following example shows how to stop the SVM `vs1.example.com`:

```
cluster1::> vserver stop -vserver vs1.example.com

[Job 72] Job succeeded: DONE

cluster1::> vserver show
```

Vserver	Type	Admin State	Root Volume	Aggregate	Name Service	Name Mapping
vs1.example.com	data	stopped	root_voll	aggr1	file	file
cluster1	admin	-	-	-	-	-

```
cluster1-01      node      -      -      -      -      -
cluster1-02      node      -      -      -      -      -
```

For more information about `vserver stop` command, see the man pages.

### Result

SVM is in `stopped` state and stops serving data to clients. The SVM administrator cannot log in to the SVM.

### Related tasks

[Displaying information about SVMs](#) on page 126

## 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 to facilitate data access to the clients.

### About this task

When the SVM root volume becomes unavailable, you can restore the root volume by promoting another volume, which does not have other volumes junctioned to it.

For SVMs with FlexVol volumes, you can promote one of the following volumes as the root volume:

- Load-sharing mirror copy
- Data-protection mirror copy
- A new FlexVol volume

**Note:** If you want to restore the root volume of an SVM with Infinite Volume, you must contact technical support.

Starting from clustered Data ONTAP 8.2, 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 at least 1 GB in size.

### Steps

1. Depending on the type of volume you select for promoting a root volume, perform the appropriate action:

If you want to promote...	Perform the following tasks...
A load-sharing mirror as the root volume of an SVM	<a href="#">Promoting load-sharing mirror copy</a> on page 133

If you want to promote...	Perform the following tasks...
A data-protection mirror as the root volume of an SVM	<a href="#">Promoting data-protection mirror copy</a> on page 134
A new FlexVol volume	<a href="#">Promoting new FlexVol volume</a> on page 135

- Use the `volume mount` command to remount the new root volume.

For more information about these commands, see the man pages.

### Result

When the new volume is promoted as the SVM root volume, the other data volumes get associated with the new SVM root volume.

## 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

- Use the `set -privilege advanced` command to set the privilege level to advanced.
- Use the `snapmirror promote` command to promote the load-sharing mirror copy as the root volume.
- Use the `vol show` command to verify the new root volume of the SVM.
- 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/vol1 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_dst1s
.
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).

### 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 example shows how to promote a 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
      .

```

## 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. Storage QoS is supported on clusters that have up to eight nodes.

### 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 262  
*Managing workload performance by using Storage QoS* on page 249

## 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, firewall service and policies, 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 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, the SP, or the RLM.

**Related concepts**

[Managing rule settings for user names and passwords in an access-control role](#) on page 151

**Related tasks**

[Customizing an access-control role to restrict user access to specific commands](#) on page 149

**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.

You use the `-application` parameter of the `security login` commands to specify 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.
- The SP or RLM (`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. For more information, see the `system services firewall policy man` pages.

**Related concepts**

[Accessing the cluster by using the CLI \(cluster administrators only\)](#) on page 14

**Related references**

[Commands for managing user accounts](#) on page 142

## 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 must be mapped to access-control roles by using the `security login create` command with the `-authmethod` parameter set to `domain`.

In addition, to authenticate Windows Active Directory domain users for cluster access, a tunnel must be set up through a CIFS-enabled SVM.

- LDAP or NIS authentication (`nsswitch`)

To use LDAP or NIS authentication, SVM users must be mapped to SVM 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).

Cluster user accounts cannot use `nsswitch` as an authentication method.

- User password (`password`)
- SSH public key authentication (`publickey`)
- SNMP user-based security model (`usm`)

For more information about the `security login` commands, see the appropriate man pages.

### Related references

[Commands for managing user accounts](#) on page 142

## 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 Active Directory domain users to access the cluster

To enable Active Directory domain users to access the cluster, you must set up an authentication tunnel through a CIFS-enabled Storage Virtual Machine (SVM). You must also create cluster user accounts for the domain users. This functionality requires that CIFS is licensed on the cluster.

### Steps

1. If a tunnel for authenticating Active Directory domain users' cluster access does not exist, complete the following steps:

**Note:** The `security login domain-tunnel show` command displays the authentication tunnel if it has been set up.

- a) Create a CIFS server for an SVM that you will use as an authentication tunnel by using the `vserver cifs create` command.

You can use any data SVM that has a CIFS server created as an authentication tunnel.

For information about CIFS servers, see the *Clustered Data ONTAP File Access Management Guide for CIFS*.

- b) Specify the authentication tunnel by using the `security login domain-tunnel create` command.

You can specify only one authentication tunnel.

2. Create a cluster user account to enable an Active Directory domain user to access the cluster by using the `security login create` command with the `-authmethod` parameter set to `domain`.

Domain authentication supports only `ssh`, `ontapi`, and `http` for the `-application` parameter.

The value of `-username` must be specified in the format of `domainname\username`, where `domainname` is the name of the CIFS domain server.

If you delete the authentication tunnel, subsequent login sessions cannot be authenticated, and Active Directory domain users cannot access the cluster. Open sessions that were authenticated prior to the deletion of the authentication tunnel remain unaffected.

### Example of enabling an Active Directory domain user to access the cluster

The following commands create a CIFS server for the vs0 SVM, specify vs0 as the tunnel for Active Directory domain authentication, and create a cluster user account to enable the Administrator user of the DOMAIN1 domain to access the cluster through SSH:

```
cluster1::> vserver cifs create -vserver vs0 -cifs-server vs0cifs
-domain companyname.example.com
```

```
cluster1::> security login domain-tunnel create -vserver vs0
cluster1::> security login create -vserver cluster1 -username
DOMAIN1\Administrator -application ssh -authmethod domain
```

### Related tasks

[Installing the root certificate of an SSL server](#) on page 169

## 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 a user	<code>security login create</code>
Display information about user accounts	<code>security login show</code>
Display information about SNMP users	<code>security snmpusers</code>
Modify the access-control role of a user's login method	<code>security login modify</code>  <b>Note:</b> It is best to use a single role for all access and authentication methods of a user account.
Delete a user's login method	<code>security login delete</code>
Change a user password	<code>security login password</code>
Lock a user account	<code>security login lock</code>  <b>Note:</b> 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>
Display the tunnel that is used for Active Directory domain user authentication	<code>security login domain-tunnel show</code>

If you want to...	Use this command...
Delete the tunnel that is used for Active Directory domain user authentication	<code>security login domain-tunnel delete</code>

For more information, see the man pages.

## Managing access-control roles

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

You use the `security login role` and `security login role config` commands to manage access-control roles. For information about these commands, see the appropriate man pages.

## 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...
admin	all	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"> <li>set</li> <li>system node autosupport</li> </ul>
	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"> <li>security login password</li> <li>set</li> </ul>
	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.

## 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 an 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 along with the respective capabilities:

<b>SVM Administrator Role Name</b>	<b>Description</b>
vsadmin	<p>This role is the super user role for an SVM and is assigned by default. The SVM administrator with this role has the following capabilities:</p> <ul style="list-style-type: none"> <li>• Managing own user account local password and key information</li> <li>• Managing volumes, quotas, qtrees, Snapshot copies, FlexCache volumes, and files</li> <li>• Managing LUNs</li> <li>• Configuring protocols: NFS, CIFS, iSCSI, and FC (FCoE included)</li> <li>• Configuring services: DNS, LDAP, and NIS</li> <li>• Monitoring jobs</li> <li>• Monitoring network connections and network interface</li> <li>• Monitoring the health of an SVM</li> </ul> <p>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"> <li>• Managing own user account local password and key information</li> <li>• Managing volumes, quotas, qtrees, Snapshot copies, FlexCache volumes, and files</li> <li>• Managing LUNs</li> <li>• Configuring protocols: NFS, CIFS, iSCSI, and FC (FCoE included)</li> <li>• Configuring services: DNS, LDAP, and NIS</li> <li>• Monitoring network interface</li> <li>• Monitoring the health of SVM</li> </ul>

SVM Administrator Role Name	Description
vsadmin-protocol	<p>The SVM administrator with this role has the following capabilities:</p> <ul style="list-style-type: none"> <li>• Managing own user account local password and key information</li> <li>• Configuring protocols: NFS, CIFS, iSCSI, and FC (FCoE included)</li> <li>• Configuring services: DNS, LDAP, and NIS</li> <li>• Managing LUNs</li> <li>• Monitoring network interface</li> <li>• Monitoring the health of an SVM</li> </ul>
vsadmin-backup	<p>The SVM administrator with this role has the following capabilities:</p> <ul style="list-style-type: none"> <li>• Managing NDMP operations</li> <li>• Making a restored volume as read-write</li> <li>• Viewing volumes and LUNs</li> </ul> <p><b>Note:</b> SVM administrator with vsadmin-backup role cannot manage own user account local password and key information.</p>
vsadmin-readonly	<p>The SVM administrator with this role has the following capabilities:</p> <ul style="list-style-type: none"> <li>• Managing own user account local password and key information</li> <li>• Monitoring the health of an SVM</li> <li>• Monitoring network interface</li> <li>• Viewing volumes and LUNs</li> <li>• Viewing services and protocols</li> </ul>

## 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	<code>all</code>	The role can access the specified directory and its subdirectories (if any), and the role can execute all commands in the directory or subdirectories.
	<code>readonly</code>	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.
	<code>none</code>	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 create, modify, delete, or show)	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	
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

- 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 149

## 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 -username 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
-cmddirname "security login role" -access none

cluster1::> security login role show -vserver cluster1 -role sec_login_readonly
(security login role show)
Vserver   Role           Command/      Access
Name      Name           Directory     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 -username 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 138

[Considerations for customizing an access-control role](#) on page 146

### Related references

[Commands for managing user accounts](#) on page 142

[Commands for managing access-control roles](#) on page 153

## 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, underscores (`_`), periods (`.`), or a combination of them.

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-alphanumeric`)
- 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-alphanumeric`)

- 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 153

## 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 into 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 password (the `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.

Regardless of the settings of the `-require-initial-passwd-update` and `-passwd-expiry-time` parameters of the `security login role config modify` command, when the “diag” user enters the systemshell from the clustershell, the systemshell does not require or prompt the “diag” user to change the password.

### Related concepts

[Access methods for user accounts](#) on page 139

### Related tasks

[Customizing an access-control role to restrict user access to specific commands](#) on page 149

## 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"> <li>• The required minimum length of a user name</li> <li>• Whether a mix of alphabetic and numeric characters is required in a user name</li> <li>• The required minimum length of a password</li> <li>• Whether a mix of alphabetic and numeric characters is required in a password</li> <li>• The required number of special characters in a password</li> <li>• Whether users must change their passwords when logging in to their accounts for the first time</li> <li>• The number of previous passwords that cannot be reused</li> <li>• The minimum number of days that must pass between password changes</li> <li>• The number of days after which a password expires</li> <li>• The number of invalid login attempts that triggers the account to be locked automatically</li> <li>• The number of days for which an account is locked if invalid login attempts reach the allowed maximum</li> </ul>	<pre>security login role config modify</pre>
<p>Display user account restrictions and rule settings</p>	<pre>security login role config show</pre>

If you want to...	Use this command...
<p>Reset the following settings to the default value:</p> <ul style="list-style-type: none"> <li>• The required number of special characters in a password (<code>-passwd-min-special-chars 0</code>)</li> <li>• Whether users must change their passwords when logging in to their accounts for the first time (<code>-require-initial-passwd-update disabled</code>)</li> <li>• The number of days after which a password expires (<code>-passwd-expiry-time unlimited</code>)</li> <li>• The number of invalid login attempts that triggers the account to be locked automatically (<code>-max-failed-login-attempts 0</code>)</li> <li>• The number of days for which an account is locked if invalid login attempts reach the allowed maximum (<code>-lockout-duration 0</code>)</li> </ul> <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></p> <p><b>Note:</b> This command requires the advanced privilege level.</p>

For more information, see the man pages for the `security login role` and `security login role config` commands.

## Managing firewall service and policies

Setting up a firewall enhances the security of the storage system and helps you prevent unauthorized access to the system. You can enable, configure, and display information about firewall service and policies.

Firewall policies can be used to control access to only management service protocols such as SSH, HTTP, HTTPS, Telnet, NTP, NDMP, or SNMP, and not data protocols such as NFS or CIFS.

You can manage firewall service and policies in the following ways:

- Enabling or disabling firewall service  
By default, firewall service is enabled.
- Displaying the current configuration about firewall service
- Creating a firewall policy with the specified policy name and network service and putting it into effect for a logical interface

- Creating a new firewall policy that is an exact copy of an existing policy but with a new policy name
- Displaying information about firewall policies
- Modifying the IP addresses and netmasks that are used by a firewall policy
- Changing a LIF's firewall policy
- Deleting a firewall policy that is not being used by a LIF

Starting from Data ONTAP 8.2, you can create firewall policies with IPv6 addresses. For more information about IPv6 addresses, see the *Clustered Data ONTAP Network Management Guide*.

You can use the `system services firewall`, `system services firewall policy`, and `network interface modify` commands to manage firewall. For information about these commands, see the appropriate man pages.

## Creating a firewall policy and assigning it to a LIF

You can create a firewall policy by specifying a policy name, a network service, and one or more IP addresses with their corresponding netmasks. After the policy is created, you can assign the firewall policy to a LIF.

### About this task

- You cannot create a firewall policy with a policy value that is either `cluster`, `data`, `intercluster`, or `mgmt`. These values are defined for the system-defined firewall policies.
- If you want to change the service associated with a firewall policy, you must delete the existing firewall policy. After deleting the firewall policy, create a new firewall policy.
- If IPv6 is enabled on the cluster, you can create firewall policies with IPv6 addresses. Once IPv6 is enabled, `data` and `mgmt` firewall policies show `::/0` address (by default).

### Steps

1. Use the `system-defined firewall policy show` command to view the information about firewall policies.

### Example

The following example shows the system-defined firewall policies with both IPv4 and IPv6 addresses.

```
cluster1::> system services firewall policy show
Policy          Service      Action IP-List
-----
cluster
                dns         allow  0.0.0.0/0
                http        allow  0.0.0.0/0
                https       allow  0.0.0.0/0
```

	ndmp	allow	0.0.0.0/0
	ntp	allow	0.0.0.0/0
	rsh	allow	0.0.0.0/0
	snmp	allow	0.0.0.0/0
	ssh	allow	0.0.0.0/0
	telnet	allow	0.0.0.0/0
data			
	dns	allow	0.0.0.0/0, ::/0
	http	deny	0.0.0.0/0, ::/0
	https	deny	0.0.0.0/0, ::/0
	ndmp	allow	0.0.0.0/0, ::/0
	ntp	deny	0.0.0.0/0, ::/0
	rsh	deny	0.0.0.0/0, ::/0
	snmp	deny	0.0.0.0/0, ::/0
	ssh	deny	0.0.0.0/0, ::/0
	telnet	deny	0.0.0.0/0, ::/0
intercluster			
	dns	deny	0.0.0.0/0
	http	deny	0.0.0.0/0
	https	deny	0.0.0.0/0
	ndmp	allow	0.0.0.0/0
	ntp	deny	0.0.0.0/0
	rsh	deny	0.0.0.0/0
	snmp	deny	0.0.0.0/0
	ssh	deny	0.0.0.0/0
	telnet	deny	0.0.0.0/0
mgmt			
	dns	allow	0.0.0.0/0, ::/0
	http	allow	0.0.0.0/0, ::/0
	https	allow	0.0.0.0/0, ::/0
	ndmp	allow	0.0.0.0/0, ::/0
	ntp	allow	0.0.0.0/0, ::/0
	rsh	deny	0.0.0.0/0, ::/0
	snmp	allow	0.0.0.0/0, ::/0
	ssh	allow	0.0.0.0/0, ::/0
	telnet	deny	0.0.0.0/0, ::/0

4 entries were displayed.

- Use the `system services firewall policy create` command to create a firewall policy.

### Example

The following example creates a policy named `data_https` that uses the HTTPS protocol and enables access from IP addresses on the 10.10 subnet:

```
cluster1::> system services firewall policy create -policy data_https
-service https -action allow -ip-list 10.10.0.0/16
```

- Optional: Use the `system services firewall policy clone` command to create a firewall policy that is an exact copy of an existing policy, but has a new name.

**Example**

The following example demonstrates how you can create a new firewall policy named `mgmt1` from an existing firewall policy named `mgmt`.

```
cluster1::> firewall policy clone -policy mgmt -new-policy-name mgmt1
(system services firewall policy clone)
```

4. Use the `network interface modify` command with the `-firewall-policy` parameter to assign the policy to a LIF.

**Example**

```
cluster1::> network interface modify -vserver vs1 -lif data1 -
firewall-policy data_https
```

For more information, see the `network interface modify` man page.

## Commands for managing firewall service and policies

You can use the `system services firewall` commands to manage firewall service, the `system services firewall policy` commands to manage firewall policies, and the `network interface modify` command to manage firewall for a LIF.

If you want to...	Use this command...
Enable and configure firewall service	<code>system services firewall modify</code>
Display the current configuration for firewall service	<code>system services firewall show</code>
Create a firewall policy or add a service to an existing firewall policy	<code>system services firewall policy create</code>
Put a firewall policy into effect for a LIF	<code>network interface modify</code>  <b>Note:</b> You use the <code>-firewall-policy</code> parameter to modify the firewall policy of a LIF.
Modify the IP addresses and netmasks used by a firewall policy	<code>system services firewall policy modify</code>  <b>Note:</b> You cannot modify the default system-defined firewall policies
Display information about firewall policies	<code>system services firewall policy show</code>

If you want to...	Use this command...
Create a new firewall policy that is an exact copy of an existing policy	<code>system services firewall policy clone</code>
Delete a firewall policy that is not used by a logical interface	<code>system services firewall policy delete</code>

For more information, see the man pages for the `system services firewall`, `system services firewall policy`, and `network interface modify` commands.

## 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>  <b>Note:</b> This command requires the 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.

## 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>

For more information, see the man pages for the `security login publickey` commands.

## Managing digital certificates for server or client authentication

A digital certificate ensures that web communications are transmitted in encrypted form. It also ensures that information is sent privately and unaltered to only the specified server or from the authenticated client. Data ONTAP enables you to generate, install, and manage a self-signed or Certificate Authority (CA) signed digital certificate for server or client authentication.

The following facts apply to digital certificates (sometimes called public key certificates):

- A digital certificate is an electronic document that verifies the owner of a public key.
- A digital certificate can be either self signed (by owner) or CA signed.

Which way to have a digital certificate signed depends on your security requirements and budget. You can obtain a self-signed digital certificate for free, but a digital certificate signed by a trusted CA can incur a considerable expense. A self-signed digital certificate is not as secure as a digital certificate signed by a CA. Therefore, it is not recommended in a production environment. A CA-signed digital certificate helps prevent man-in-the-middle attacks and provides better security protection than a self-signed digital certificate.

- By default, Data ONTAP uses the SHA256 cryptographic hashing function for signing a CSR or digital certificate, and the SHA1 and MD5 cryptographic hashing functions are also supported. Private keys generated by Data ONTAP are 2048-bit by default. Data ONTAP also enables you to generate a 512-bit, 1024-bit, or 1536-bit private key. However, the higher the value, the more secure the key is.

When the cluster or Storage Virtual Machine (SVM) functions as an SSL server, you can manage digital certificates in the following ways:

- Create a self-signed or CA-signed digital certificate (the `server` certificate type)  
To obtain a self-signed digital certificate, you simply create one on the cluster or SVM. Data ONTAP automatically creates a self-signed digital certificate for server authentication of an SVM when you create that SVM.  
To obtain a CA-signed digital certificate, you generate a digital certificate signing request (CSR), which contains a private key and information that identifies you as the applicant. You then send the CSR to a CA electronically to apply for a digital certificate. After the CA sends you the signed digital certificate, you install it with the associated private key on the cluster or SVM.
- Mutually authenticate the server and the client  
For mutual authentication, you create a self-signed root CA certificate for the server (the `root-ca` certificate type), generate a CSR for the client that can be self signed by the server using its root CA or signed by a third-party CA, and install the client certificate on the client.  
If the CSR for the client is signed by a third-party CA, you also need to install the root certificate and each intermediate certificate of the CA that signed the certificate (the `client-ca` certificate type).
- Display information about the installed digital certificates
- Revoke a compromised CA-issued digital certificate
- Delete self-signed or CA-signed digital certificates

When the cluster or SVM functions as a client to an SSL server (which, for example, can be an Active Directory domain controller that supports LDAP over SSL), you can manage digital certificates in the following ways:

- Install the root certificate of the server (the `server-ca` certificate type) on the cluster or SVM  
The root certificate is provided by the server and used for server authentication.
- Display information about the installed root certificate of the server
- Delete the root certificate of the server

Before reverting to a release earlier than Data ONTAP 8.2, all digital certificates except for the `server` type (`security certificate show -type server`) must be deleted. Otherwise, the revert procedure fails.

You use the `security certificate` commands to manage digital certificates. For information about these commands, see the man pages.

## Related tasks

[Configuring access to web services](#) on page 177

## Generating and installing a CA-signed digital certificate for server authentication

When the cluster or Storage Virtual Machine (SVM) functions as an SSL server, you can generate and install a CA-signed digital certificate for server authentication. A CA-signed digital certificate helps prevent man-in-the-middle attacks and provides better security protection than a self-signed digital certificate.

### Steps

1. If you do not already have a certificate signed by a CA, complete the following steps to obtain a CA-signed digital certificate:

- a) Generate a digital certificate signing request (CSR) by using the `security certificate generate-csr` command.

The system displays the CSR output on the console. The output includes a certificate request and a private key.

- b) Copy the certificate request from the CSR output and send it 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 digital certificate for future reference.

For more information, see the `security certificate generate-csr` man page.

2. Install the CA-signed digital certificate by using the `security certificate install` command with the `-type server` parameter.

For more information, see the `security certificate install` man page.

3. Enter the private key when the system prompts you to.

### Examples of generating and installing a CA-signed digital certificate

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-----
MIICrjCCAZYCAQMwaTEQMA4GA1UEAxMHcnRwLmNvbTELMakGAlUEBhMVCVVMxCzAJ
```

```
BgNVBAgTAK5DMQwwCgYDVQQHEwNSVFAXDTALBgNVBAoTBGNVcmUxDALBgNVBASt
BGNVcmUxDzANBgkqhkiG9w0BCQEWADCCASlWdQYJKoZIhvcNAQEBBQADggEPADCC
AQoCggEBAK3azmz6UniYDKVjA4iD3ImclAJ0sst3jPH2VqFwKbR9+srrC7l7yt8
1s3JMDFZVXxv+GmBYWfOuzvMzajR2G7fg6/U2Z9XviXQo0m+FsqYt5H3ZEzhhK6
G8rIEqKPL9yY3RFxfVcwoRn7k/Q9IvKwjlvxywjkVYi jN9o7l9G159jBvmAkKyH0
SXz6iIwGzc8so8jiUm6QdU5viDNBxeo+tkHy12gKDEjy5TGnu0cVVQ56Cx0zYwG
cgg32elgMg3MFUFV+TtAVoPkBibC9AuZfrXfMBJW/IR4mDs+fQL0Q5becWzETCwu
9mY4kPt0YvyJiPXuJmWGl44giQMi6cUCAwEAAaAAMA0GCSqGSIb3DQEBCwUAA4IB
kYz7hzkFpuMibAaCkp54Qrho
-----END CERTIFICATE REQUEST-----
```

Private Key:

```
-----BEGIN RSA PRIVATE KEY-----
MIIBPAIBAAJBAM16ytrK8nQj82UsWeHOeT8gk0BPX+Y5MLycsUdXA7hXhumHNpvF
C6lX2G32Sx8VEa1th94tx+vOEq+UaqHlt0CAwEAAQJBAMZjDw1gmlm3qIr/n8VT
PFnnZnbVcXVM7OtUsgPKw+QCCh9dF1jmuQKeDr+wUMWknlDeGrfhILpzfJGhrLJ
z7UCIQDr8d3gOG7lUyX+BbFmo/N0uAKjS2cvUU+Y8a8pDxGLLwIhANqa99SuS18U
DiPvdaKTj6+EcGuXfCXz+G0rfgTZK8uzAiEARlmmrfYC8KwE9k7A0y1RzBLdUwK9
AvuJdn+/z+H1Bd0CIQDD93P/xpaJETNz53Au49VE5Jba/Jugckrbosd/lSd7nQIG
aEMAzt6qHHT4mndi8Bo8sDGedG2SKx6Qbn2IpuNZ7rc
-----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 digital certificate for the SVM named vs1:

```
cluster1:> security certificate install -vserver vs1 -type server
Please enter Certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIIB8TCCAzugAwIBAwIBADANBgkqhkiG9w0BAQQFADBfMRMwEQYDVQQDEwpuZXRh
cHAuY29tMQswCQYDVQQGEwJVUzEUMAcGALUECBMAMQkwBwYDVQQHEwAxCCTAHHBgNV
BAOTADEJMAcGALUECXMAMQ8wDQYJKoZIhvcNAQkBFgAwHhcNMTAwdi2MTk00TI4
...
-----END CERTIFICATE-----

Please enter Private Key: Press <Enter> when done
-----BEGIN RSA PRIVATE KEY-----
MIIBPAIBAAJBAM16ytrK8nQj82UsWeHOeT8gk0BPX+Y5MLycsUdXA7hXhumHNpvF
C6lX2G32Sx8VEa1th94tx+vOEq+UaqHlt0CAwEAAQJBAMZjDw1gmlm3qIr/n8VT
PFnnZnbVcXVM7OtUsgPKw+QCCh9dF1jmuQKeDr+wUMWknlDeGrfhILpzfJGhrLJ
...
-----END RSA PRIVATE KEY-----

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+zCCBGsGAWIBAgICAQ0wDQYJKoZIhvcNAQEFBQAwwgbsxJDAiBgNVBACGTG1Zh
bGlDZXJ0IFZhbG1kYXRpb24gTmV0d29yazEXMBUGALUEChMOVmFsaUNlcnQsIElu
Yy4xNTAzBgNVBASlTLFZhbG1DZXJ0IENsYXNzIDIGUG9saWN5IFZhbG1kYXRpb24g
...
-----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.

## Installing a server intermediate certificate

You must install the intermediate certificate on the server (the cluster or Storage Virtual Machine (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.

### About this task

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.

### Step

1. Install the intermediate certificate by using the `security certificate install` command.

## Providing mutual authentication

You can configure the server (the cluster or a Storage Virtual Machine (SVM)) to provide mutual authentication for greater security between the server and a group of clients.

### About this task

When using mutual authentication, also called *two-way authentication*, both the server and the client present their certificates to each other and validate their respective identities to each other. To configure mutual authentication using a self-signed root CA certificate, you must create a self-signed root CA certificate, enable client authentication, generate and sign a certificate signing request (CSR) for each user, and install the client certificate on the client side. You must also set up user accounts for them to be authenticated by digital certificates.

You can also provide client authentication using a CSR signed by a third-party CA that is installed on the client and installing intermediate certificates of the CA that signed the certificate.

### Steps

1. Create a self-signed root CA certificate for the server by using the `security certificate create` command.

### Example

The following command creates a root CA certificate for the “vs1” SVM for a software group in the IT department of a company whose custom common name is `lab.companyname.com`:



- Certificate authority (CA)
- Serial number of the certificate
- Server name

You do this for every client that you need to authenticate.

### Example

```
cluster1::> security certificate show -instance -vserver vs1
                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: Thu Jan 17 18:31:47 2013
    Certificate Expiration Date: Fri Jan 17 18:31:47 2014
    Public Key Certificate: -----BEGIN CERTIFICATE-----
                                MIID
+zCCAuOgAwIBAgIEUPhDkjANBgkqhkiG9w0BAQsFADBbMQ8wDQYDVQQDEwZt
                                .
                                .
                                .
```

- c) Sign the CSR with the root CA generated previously by using the `security certificate sign` command.

You do this for every user client that you need to authenticate.

### Example

```
cluster1::> security certificate sign -vserver vs1 -ca
lab.companyname.com -ca-serial 50F84392
```

5. If you have a third-party CA sign the CSR, complete the following steps:
- a) Copy the certificate request from the CSR output and send it 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) 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.

You do this for each certificate.

6. 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`.



## 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 by specifying one of the following types: <ul style="list-style-type: none"> <li>• <code>server</code>, used for server authentication</li> <li>• <code>root-ca</code>, used for signing other certificates by acting as a certificate authority (CA)</li> </ul>	<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>
Install a CA-signed digital certificate and the public key certificate of the root CA by specifying one of the following types: <ul style="list-style-type: none"> <li>• <code>server</code>, the server certificate and its intermediate certificates</li> <li>• <code>client-ca</code>, the public key certificate for the root CA of an SSL client</li> <li>• <code>server-ca</code>, the public key certificate for the root CA of an SSL server to which the cluster or SVM is a client</li> </ul>	<code>security certificate install</code>
Display information about installed digital certificates	<code>security certificate show</code>
Display CA-issued digital certificates	<code>security certificate ca-issued show</code>
Revoke a compromised CA-issued digital certificate	<code>security certificate ca-issued revoke</code>
Delete an installed digital certificate	<code>security certificate delete</code>

For more information, see the man pages for the `security certificate` commands.

## 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*)  
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.

**Note:** Starting with Data ONTAP 8.2, Element Manager (including ClusterView) is no longer available.

### Related concepts

[Understanding OnCommand System Manager](#) on page 32

### Related tasks

[Accessing a node's log, core dump, and MIB files by using a web browser](#) on page 46

## 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:

- Configuring the web protocol engine to control whether remote clients can use HTTP or HTTPS to access web service content
- Specifying whether SSLv3 or SSLv2 should be used for secure web access  
Data ONTAP supports SSLv3 and SSLv2. By default, SSLv3 is enabled and SSLv2 is disabled. If SSL is configured, Transport Layer Security 1.0 (TLSv1.0) is also enabled and cannot be disabled.
- Displaying the configuration and status of web services

You use the `system services web` commands to manage the web protocol engine at the cluster level.

If a 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.

### Related concepts

[Managing SSL](#) on page 176

[Managing web services](#) on page 174

### Related tasks

[Creating a firewall policy and assigning it to a LIF](#) on page 156

[Configuring access to web services](#) on page 177

[Creating a firewall policy and assigning it to a LIF](#) on page 156

## 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"> <li>• Enable or disable the web protocol engine for the cluster</li> <li>• Enable or disable SSLv2 or SSLv3 for the cluster</li> </ul>	<code>system services web modify</code>
Display the configuration of the web protocol engine at the cluster level and determine whether the web protocols are functional throughout the cluster	<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 an LIF	<code>network interface modify</code> You can use the <code>-firewall-policy</code> parameter to modify the firewall policy of an LIF.

For more information, see the man pages.

### Related references

[Commands for managing SSL](#) on page 176

[Commands for managing firewall service and policies](#) on page 158

[Commands for managing digital certificates](#) on page 170

## 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 172

[Managing user accounts](#) on page 138

[Access methods for user accounts](#) on page 139

[Managing SSL](#) on page 176

### Related tasks

[Configuring access to web services](#) on page 177

## 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"> <li>• Enable or disable a web service</li> <li>• Specify whether only HTTPS can be used for accessing a web service</li> </ul>	<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>

For more information, see the man pages.

## 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>  <b>Note:</b> Only a single mount point can exist from a 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>

If you want to...	Use this command...
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>

For more information, see the man pages.

## 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 (SSLv2 or SSLv3) can be used
- Restricting access to only HTTPS requests for a web service

### Related concepts

[Managing the web protocol engine](#) on page 172

[Managing web services](#) on page 174

[Managing digital certificates for server or client authentication](#) on page 162

### Related tasks

[Configuring access to web services](#) on page 177

[Creating a firewall policy and assigning it to a LIF](#) on page 156

## 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>

For more information, see the man pages.

**Related references**

*Commands for managing web services* on page 175

*Commands for managing the web protocol engine* on page 173

*Commands for managing firewall service and policies* on page 158

*Commands for managing digital certificates* on page 170

**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 176

*Managing digital certificates for server or client authentication* on page 162

*Managing the web protocol engine* on page 172

*Managing web services* on page 174

*Managing access-control roles* on page 143

*Access methods for user accounts* on page 139

### Related tasks

*Creating a firewall policy and assigning it to a LIF* on page 156

## 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 unable to connect or failure to establish a connection 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.  <b>Note:</b> 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.  <b>Note:</b> 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.  <b>Note:</b> You use the <code>system services web</code> commands to manage the web protocol engine for the cluster.
Your web browser returns a not found 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.  <b>Note:</b> 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.	<p>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.</p> <p><b>Note:</b> 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.</p>
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.	<p>Ensure that the cluster or SVM has SSL enabled and that the digital certificate is valid.</p> <p><b>Note:</b> 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.</p>
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.	<p>Ensure that the digital certificate associated with the cluster or SVM is signed by a trusted CA.</p> <p><b>Note:</b> 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.</p>

**Related concepts**

*Managing firewall service and policies* on page 155

[Managing the web protocol engine](#) on page 172

[Managing digital certificates for server or client authentication](#) on page 162

[Managing web services](#) on page 174

[Managing user accounts](#) on page 138

[Managing access-control roles](#) on page 143

[Managing SSL](#) on page 176

### Related tasks

[Creating a firewall policy and assigning it to a LIF](#) on page 156

[Generating and installing a CA-signed digital certificate for server authentication](#) on page 164

## Managing audit settings

Audit logging creates a chronological record of management activities. You can specify what types of activities in the management interface are audited.

Data ONTAP enables you to audit two types of requests—set requests and get requests. A set request typically applies to non-display commands, such as creating, modifying, or deleting an object. A get request occurs when information is retrieved and displayed to a management interface. This is the type of request that is issued when you run a `show` command, for instance.

You use the `security audit` commands to manage audit settings. Regardless of the settings for the `security audit` commands, set requests are *always* recorded in the `command-history.log` file of the `/mroot/etc/log/mlog/` directory, and the file is sent by AutoSupport to the specified recipients.

You can also use the `security audit modify` command to specify whether the following requests are also recorded in the `mgwd.log` file of the `/mroot/etc/log/mlog/` directory for technical support and diagnostic purposes:

- Set requests for the CLI
- Set requests for the ONTAP API
- Get requests for the CLI
- Get requests for the ONTAP API

By default, auditing of set requests is enabled (that is, recorded in the `mgwd.log` file), and auditing of get requests is disabled.

The `command-history.log` and `mgwd.log` files are rotated when they reach 100 MB in size, and their previous 34 copies are preserved (with a maximum total of 35 files, respectively).

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 46

## Commands for managing audit settings

You use the `security audit` commands to manage audit settings.

<b>If you want to...</b>	<b>Use this command...</b>
Set preferences for audit logging	<code>security audit modify</code>
Display the current audit settings	<code>security audit show</code>

For more information, see the man pages for the `security audit` commands.

## Managing the cluster time (cluster administrators only)

---

Problems can occur when the cluster time is inaccurate. You can manually set the time zone, date, and time on the cluster. However, it is best to keep your cluster time synchronized automatically by using the Network Time Protocol (NTP) servers.

Data ONTAP enables you to manage the cluster time in the following ways:

- Configuring the NTP servers

On a cluster running Data ONTAP 8.2 and later releases, NTP is always enabled on the cluster. To disable NTP (not recommended), you must contact technical support.

Although NTP is always enabled, for the cluster to synchronize with an external time source, you must configure the NTP servers. The `system services ntp server` commands enable you to manage the NTP servers in the following ways:

- Associating a node with an NTP server  
You can get a list of public NTP time servers from the NTP Public Services page.
- Specifying the preferred NTP server and version for a node  
Data ONTAP 8.0 and 8.0.1 use NTP v4 by default. To address situations where certain time servers support only NTP v3, starting with Data ONTAP 8.0.2, the NTP version to be used for communicating with a newly configured NTP server defaults to v3 instead of v4. The NTP Daemon continues to use the highest supported version (v4 in this case) to communicate with the time servers that were configured prior to Data ONTAP 8.0.2.
- Displaying information about NTP servers that are associated with a node or the cluster
- Dissociating a node from an NTP server

For more information about the `system services ntp server` commands, see the man pages.

- Manually setting the cluster time

Data ONTAP enables you to manually manage the cluster time. The time you set takes effect across all nodes in the cluster. This capability is helpful for the following purposes:

- Ensuring the intra-cluster time consistency  
If no external time server is used, setting the cluster time manually ensures a time setting that is consistent across all nodes in the cluster.
- Manually correcting erroneous cluster time  
Even if an external time server is used, it is possible for the times on the nodes to become significantly incorrect (for example, a node's time has become incorrect after a reboot). In that case, you can manually specify an approximate time for the cluster until NTP can synchronize with an external time server.

The `cluster date` commands enable you to manually manage the cluster time in the following ways:

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

For more information about the `cluster date` commands, see the man pages.

**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

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

## Commands for managing the cluster time

You use the `system services 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 on the cluster:

If you want to...	Use this command...
Associate a node with an NTP server and optionally specify the following options: <ul style="list-style-type: none"> <li>• The preferred NTP server (advanced privilege level)</li> <li>• The version of NTP that is running on the specified NTP server</li> </ul>	<code>system services ntp server create</code>
Modify NTP server options: <ul style="list-style-type: none"> <li>• The preferred NTP server (advanced privilege level)</li> <li>• The NTP version to be used for communicating with a specific time server</li> </ul>	<code>system services ntp server modify</code>
Display information about NTP servers associated with a node or the cluster	<code>system services ntp server show</code>
Dissociate a node from an NTP server	<code>system services ntp server delete</code>

The following commands enable you to manage the cluster time manually:

<b>If you want to...</b>	<b>Use this command...</b>
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>

For more information, see the man pages.

## 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.

When you set up a cluster, the setup wizard prompts you to enter the cluster base license key. 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.

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.

For information about finding license keys for your initial or add-on software orders, contact technical support.

Data ONTAP enables you to manage 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`)

### Related information

*IBM N series support website: [www.ibm.com/storage/support/nseries](http://www.ibm.com/storage/support/nseries)*

## 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.

**Note:** The `sysconfig` command in the nodeshell displays the system serial number of a node.

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 licenses for the cluster.

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"> <li>• License package name and description</li> <li>• License type (<code>site</code>, <code>license</code>, or <code>demo</code>)</li> <li>• Expiration date, if applicable</li> <li>• The cluster or nodes that a package is licensed for</li> <li>• Whether the license was installed prior to Data ONTAP 8.2 (<code>legacy</code>)</li> <li>• Customer ID</li> </ul>	<code>system license show</code>  <b>Note:</b> Some information is displayed only when you use the <code>-instance</code> parameter.
Display all packages that require licenses and their current license status, including the following: <ul style="list-style-type: none"> <li>• The package name</li> <li>• The licensed method</li> <li>• The expiration date, if applicable</li> </ul>	<code>system license status show</code>
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>

For more information, see the man pages for the `system license` commands.

## Managing jobs and schedules

---

A *job* is any asynchronous task. Jobs are typically long-running volume operations such as copy, move, and mirror. You can monitor, pause, stop, and restart jobs, and 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 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.

If you want to...	Use this command...
Display information about all jobs	<code>job show</code>
Display information about jobs on a per-node basis	<code>job show-bynode</code>
Display information about cluster-affiliated jobs	<code>job show-cluster</code>
Display information about completed jobs	<code>job show-completed</code>

If you want to...	Use this command...
Display information about job history	<pre>job history show</pre> <p>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.</p>
Display the list of private jobs	<pre>job private show</pre> <p><b>Note:</b> This command is only available at the advanced privilege level.</p>
Display information about completed private jobs	<pre>job private show-completed</pre> <p><b>Note:</b> This command is only available at the advanced privilege level.</p>
Display information about the initialization state for job managers	<pre>job initstate show</pre> <p><b>Note:</b> This command is only available at the advanced privilege level.</p>
Monitor the progress of a job	<pre>job watch-progress</pre>
Monitor the progress of a private job	<pre>job private watch-progress</pre> <p><b>Note:</b> This command is only available at the advanced privilege level.</p>
Pause a job	<pre>job pause</pre>
Pause a private job	<pre>job private pause</pre> <p><b>Note:</b> This command is only available at the advanced privilege level.</p>
Resume a paused job	<pre>job resume</pre>
Resume a paused private job	<pre>job private resume</pre> <p><b>Note:</b> This command is only available at the advanced privilege level.</p>
Stop a job	<pre>job stop</pre>

If you want to...	Use this command...
Stop a private job	<code>job private stop</code> <b>Note:</b> This command is only available at the advanced privilege level.
Delete a job	<code>job delete</code>
Delete a private job	<code>job private delete</code> <b>Note:</b> This command is only available at the 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> <b>Note:</b> This command is only available at the advanced privilege level.

**Note:** You can use the `event log show` command to determine the outcome of a completed job.

For more information, see the man pages.

## Commands for managing job schedules

Many tasks—for instance, volume snapshots—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 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>

If you want to...	Use this command...
Create an interval schedule	<pre>job schedule interval create</pre> <p>You must specify at least one of the following parameters: -days, -hours, -minutes, or -seconds.</p>
Modify a cron schedule	<pre>job schedule cron modify</pre>
Modify an interval schedule	<pre>job schedule interval modify</pre>
Delete a schedule	<pre>job schedule delete</pre>
Delete a cron schedule	<pre>job schedule cron delete</pre>
Delete an interval schedule	<pre>job schedule interval delete</pre>

For more information, see the man pages.

# 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, 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...
<p>Change the settings for a configuration backup schedule, including:</p> <ul style="list-style-type: none"> <li>• Specifying a remote URL (either HTTP or FTP) where the configuration backup files will be uploaded in addition to the default locations in the cluster</li> <li>• Specifying a user name to be used to log in to the remote URL</li> <li>• Setting the number of backups to keep for each configuration backup schedule</li> </ul>	<pre>system configuration backup settings modify</pre>
<p>Set the password to be used to log in to the remote URL</p>	<pre>system configuration backup settings set-password</pre>
<p>View the settings for the configuration backup schedule</p>	<pre>system configuration backup settings show</pre> <p><b>Note:</b> You set the <code>-instance</code> parameter to view the user name and the number of backups to keep for each schedule.</p>

For more information, see the man pages.

## 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>
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>  <b>Note:</b> 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>  <b>Note:</b> 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.

For more information, see the man pages.

## 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 196
2. [Restoring the node configuration using a configuration backup file](#) on page 197

## 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.
On a node in the cluster	<ol style="list-style-type: none"> <li>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> <li>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.</li> </ol>

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 from 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.

**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 recreation, then you must re-create the cluster again.

**Related tasks**

[Synchronizing a node with the cluster](#) on page 201

## 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 198
2. [Restoring a cluster configuration from an existing configuration](#) on page 199

## 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. Use the `cluster ring show` command at the advanced privilege level 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.
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="475 343 1239 430">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> <li data-bbox="475 447 1239 534">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.</li> </ol>

## 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.

### Steps

1. Halt each node except for the recovering node by using the `system node halt` command.

### Example

```
cluster1::*> system node halt -node node0 -reason "recovering cluster"
Warning: Are you sure you want to halt the node? {y|n}: y
```

2. On the recovering node, use the `system configuration recovery cluster recreate` command at the advanced privilege level 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, with a new UUID, on the recovering node.

3. Use the `cluster identity show` command to verify that the recovering node has a different UUID than the other nodes.
4. Boot each node that needs to be rejoined to the re-created cluster.

You must reboot the nodes one at a time.

5. 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, use the `system configuration recovery cluster rejoin` command at the advanced privilege level to rejoin the target node to the cluster.

### Example

This example rejoins the target node (*node2*) 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) Use the `cluster show` command with the `-eligibility` parameter to verify that the target node is healthy and has formed quorum with the rest of the nodes in the cluster.

The target node must rejoin the re-created cluster before you can rejoin another node.

### 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 synch with the cluster, then you synchronize the node to restore the replicated database (RDB) on the node and bring it into quorum.

### Step

1. Use the `system configuration recovery cluster sync` command at the advanced privilege level to synchronize the node 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
- 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 46

## 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>

If you want to...	Use this command...
Save all unsaved core dumps that are on a specified node	<code>system node coredump save-all</code>
Upload a saved core dump file to a specified location	<code>system node coredump upload</code>
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>
Display application core dump reports	<code>system node coredump reports show</code>
Upload an application core dump report to a specified location	<code>system node coredump reports upload</code>
Delete an application core dump report	<code>system node coredump reports delete</code>

For more information, see the man pages.

## 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, including: <ul style="list-style-type: none"> <li>Whether to automatically segment a core dump file after it is saved</li> <li>Whether to automatically delete the original core dump file after it is segmented</li> </ul>	<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>

If you want to...	Use this command...
Display information about the core segments on a node, for example: <ul style="list-style-type: none"> <li>• The core segment name</li> <li>• Total number of core segments for the full core</li> <li>• The time when the panic occurred that generated the core dump file</li> </ul>	<pre>system node coredump segment show</pre>
Delete a specified core segment from a node	<pre>system node coredump segment delete</pre>
Delete all core segments from a node	<pre>system node coredump segment delete-all</pre>
Displays the status of a core segmenting job, including the following: <ul style="list-style-type: none"> <li>• Job ID</li> <li>• Name of the core dump file that is being segmented</li> <li>• Job status</li> <li>• Percent completed</li> </ul>	<pre>system node coredump segment status</pre>
Cancel a core segmenting job as specified by its job ID	<pre>system node coredump segment stop</pre>

For more information, see the man pages.

## 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)
- 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 or SNMP trap hosts.
- 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 and SNMP trap hosts.

**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 {-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 Threshd	Severity	Destinations	Freq Threshd	Time
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

- 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 critical and above events 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>
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

If you want to...	Use this command...
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>

For more information, see the man pages.

## Managing AutoSupport

AutoSupport is a mechanism that proactively monitors the health of your system and automatically sends messages to IBM 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 N series support website (accessed and navigated as described in [Websites](#) on page 10).

## 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.

**Note:** 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 storage 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 <code>enable</code>

### 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 <code>enable</code>
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 <code>true</code>	Addresses specified in <code>-partner-address</code> Technical support, if <code>-support</code> is set to <code>enable</code>
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 <code>enable</code>

### 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 <code>enable</code> .
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 On Demand 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 <code>enable</code> and the transport protocol is HTTPS
When AutoSupport obtains delivery instructions to resend past AutoSupport messages	Technical support, if <code>-support</code> is set to <code>enable</code> and the transport protocol is HTTPS

### Related concepts

[How AutoSupport On Demand obtains delivery instructions from technical support](#) on page 213

## 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.

## How AutoSupport On Demand obtains delivery instructions from technical support


AutoSupport On Demand periodically communicates with technical support to obtain delivery instructions for sending, resending, and declining AutoSupport messages.

The AutoSupport On Demand client, which runs on each node, periodically polls the AutoSupport On Demand service to obtain delivery instructions. The AutoSupport On Demand service resides in technical support. AutoSupport On Demand is enabled by default, and it automatically communicates with technical support if AutoSupport is configured to send messages to technical support and the transport protocol is HTTPS. The client sends HTTPS requests to the same technical support location to which AutoSupport messages are sent.

The AutoSupport On Demand client does not accept incoming connections.

**Note:** AutoSupport On Demand uses the “autosupport” user account to communicate with technical support. You should not delete this account.

The following illustration shows how AutoSupport On Demand 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.
- 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.

## 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

Type of message	Type of data the message contains
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"> <li><code>test</code> 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>autosupport.to</code> option, so that you can confirm that AutoSupport messages are being received.</li> <li><code>performance</code> sends performance data.</li> <li><code>all</code> 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.</li> </ul>
Triggered by AutoSupport On Demand	<p>AutoSupport On Demand can request new messages or past messages:</p> <ul style="list-style-type: none"> <li>New messages, depending on the type of AutoSupport collection, can be <code>test</code>, <code>all</code>, or <code>performance</code>.</li> <li>Past messages depend on the type of message that is resent.</li> </ul>

## 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

AutoSupport provides flexible delivery options. You can choose the format in which AutoSupport provides information to technical support: HTTPS, HTTP, and SMTP, as well as other proxy configurations that meet your networking and security requirements. All of these protocols run on IPv4 or IPv6, based on the address family to which the name resolves.

### Supported protocols

**Note:** IBM N series systems support HTTPS and SMTP. HTTP is not supported.

If you enable AutoSupport to send messages to technical support, you can use any of the following transport protocols:

Protocol and port	Description
HTTPS on port 443	<p>This is the default protocol. You should use this whenever possible.</p> <p>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>The delivery uses an HTTP PUT request. With PUT, if the request fails during transmission, the request restarts where it left off. If the server receiving the request does not support PUT, the delivery uses an HTTP POST request.</p>
SMTP on port 25	<p>You should use this protocol only if the network connection does not allow HTTPS or HTTP, because SMTP can introduce limitations on message length and line length.</p>

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 would 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.

## Configuration requirements

Use of HTTP or HTTPS protocols requires additional configuration of the 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 username 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), 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

Perform this procedure on each node in your system where you want to configure AutoSupport.

For more information about the following commands, see the man pages.

### 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 following parameters of the `system node autosupport modify` command:
  - a) Set `-support` to `enable`.
  - b) Select a transport protocol for messages to technical support by setting `-transport` to `smtp`, `http`, or `https`.
  - c) If you chose HTTP or HTTPS as the transport protocol and you use a proxy, set `-proxy-url` to the URL of your proxy.
3. If you want your internal support organization or a support partner to receive AutoSupport messages, perform the following actions:
  - a) 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

Set this parameter ...	To this...
<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

- b) Check that addresses are correctly configured by listing the destinations using the `system node autosupport destinations show` command.
4. 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.
  - Set `-from` to the email address that sends the AutoSupport message.
  - Set `-max-smtp-size` to the email size limit of your SMTP server.
5. If you want AutoSupport to specify a fully qualified domain name when it sends connection requests to your SMTP mail server, configure DNS.

For information about configuring DNS, see the *Clustered Data ONTAP Network Management Guide*.

6. 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 <code>true</code> . If you change from <code>false</code> to <code>true</code> , all AutoSupport history and all associated files are deleted.
Stop sending performance data in periodic AutoSupport messages	Set <code>-perf</code> to <code>false</code> .

7. Check the overall configuration using the `system node autosupport show` command with the `-node` parameter.
8. 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 IBM is receiving your AutoSupport messages by checking the email addresses that were specified in the `autosupport.to` option.

An automated response from the IBM mail handler should have been sent to these email addresses.

- 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 223

## Commands for managing AutoSupport

You use the `system node autosupport` commands to change or view AutoSupport configuration, display information about past AutoSupport messages, and send or resend 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 configuration of AutoSupport

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>

If you want to...	Use this command...
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>
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 or resend AutoSupport messages

If you want to...	Use this command...
Retransmit a locally stored AutoSupport message, identified by its AutoSupport sequence number  <b>Note:</b> 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.	<code>system node autosupport history retransmit</code>
Generate and send an AutoSupport message—for example, for testing purposes	<code>system node autosupport invoke</code>  <b>Note:</b> 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.

For more information, see the man pages.

## 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.

## 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 N series support website (accessed and navigated as described in [Websites](#) on page 10) that you can access using a browser. Your system must have AutoSupport enabled and configured so that it sends data back to IBM. You can access My AutoSupport by going to [www.ibm.com/storage/support/nseries/](http://www.ibm.com/storage/support/nseries/).

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.
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 On Demand 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 delivery over SMTP](#) on page 225

[Troubleshooting AutoSupport over HTTP or HTTPS](#) on page 224

## Troubleshooting AutoSupport 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 determined that AutoSupport can generate the message, but not deliver the message over HTTP or HTTPS.

### About this task

IBM N series systems support HTTPS and SMTP. HTTP is not supported.

### Steps

1. At the storage system's CLI, ensure that DNS is enabled and configured correctly by entering the following command:  
`vserver services dns`
2. Read the error for the AutoSupport message by using the `system node autosupport history show` command with the `-seq-num` and `-destination` parameters.
3. At the storage system's CLI, ensure that the system is routing out to the Internet successfully by entering the following command:  
`network traceroute`
4. Use the `system node run` command to run the nodeshell CLI on a specific node, and use the `rdfile` command to read the `/etc/log/mlog/notifyd.log` file.

### Related tasks

[Troubleshooting AutoSupport when messages are not received](#) on page 223

## Troubleshooting AutoSupport 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 determined that AutoSupport can generate the message, but not deliver the message over SMTP.

### About this task

IBM N series systems support HTTPS and SMTP. HTTP is not supported.

### Steps

1. At the storage system CLI, ensure that DNS for the system is enabled and configured correctly by entering the following command:

```
vserver services dns
```

2. At the storage system CLI, ensure that the mail host specified in the configuration is a host that the storage system can communicate with by entering the following command:

```
network ping -node node_name -destination mailhost
```

*mailhost* is the name or IP address of your mail host.

3. Log in to the host designated as the mail host, and ensure that it can serve SMTP requests by entering the following command:

```
netstat -aAn|grep 25
```

25 is the listener SMTP port number.

A message appears, similar to the following text:

```
ff64878c tcp          0          0 *.25    *.*     LISTEN.
```

4. At the storage system CLI, ensure that the system is reaching the mail host successfully by entering the following command:

```
network traceroute
```

5. From some other host, open a Telnet session with the SMTP port by entering the following command:

```
telnet mailhost 25
```

A message similar to the following text is displayed:

```
Trying 192.9.200.16 ...
Connected to filer.
Escape character is '^]'.

```

```
220 filer.yourco.com Sendmail 4.1/SMI-4.1 ready at Thu, 30 Nov 95
10:49:04 PST
```

6. Use the `system node run` command to run the nodeshell CLI on a specific node, and use the `rdfile` command to read the `/etc/log/mlog/notifyd.log` file.

### Related tasks

[Troubleshooting AutoSupport when messages are not received](#) on page 223

## Troubleshooting EMS events about rejected or failed SMTP attempts

If the system attempted to send an AutoSupport email, but the attempt resulted in an EMS event about a rejected or failed SMTP or an unknown user, you can check the relaying configuration for the mail host to determine whether relaying is denied or incorrectly configured.

### About this task

The EMS identifiers for this event are `asup.smtp.fail` and `asup.smtp.reject`.

### Steps

1. From a Windows, UNIX, or Linux host, telnet to port 25 of the mail host by entering the following command:  

```
telnet mailhost 25
```
2. Test whether relaying is denied on the mail host.
  - a) Enter the following commands:  

```
HELO DOMAIN NAME
MAIL FROM: your_email_address
RCPT TO: autosupport@ibm.com
```
  - b) If you receive a message similar to `relaying denied`, contact the mail host vendor because relaying is denied. Otherwise, continue to the next step.
3. Test whether relaying is incorrectly configured on the mail host.
  - a) Enter the following commands:  

```
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.
  - b) If you receive a message similar to `unknown user` or `unknown mailbox`, contact the mail host vendor because relaying is incorrectly configured.

## 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.

### How health monitoring works

Individual health monitors have a set of policies that trigger alerts when certain conditions and 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.

## What health monitors are available

There are several health monitors that monitor different parts of a cluster.

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.
Node connectivity (node-connect)	CIFS nondisruptive operations (CIFS-NDO)	Monitors SMB connections to ensure nondisruptive operations to Hyper-V applications.
	Storage (SAS-connect)	Monitors shelves, disks, and adapters at the node level to ensure that they have 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 to ensure that all shelves always have 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.

### 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 206

## 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.

### 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 if you can resolve the problem or if you 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. Optional: 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
      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 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 health 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 health 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 health cluster-switch create` command to configure discovery of the switch:

**Example**

```
cluster1::> system health 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 health cluster-switch show` command to verify whether 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 N series support website (accessed and navigated as described in [Websites](#) on page 10).

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 233

## 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 drill-down for 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>  <b>Note:</b> 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>  <b>Note:</b> 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	<pre>system health node-connectivity adapter show</pre> <p><b>Note:</b> Use the <code>-instance</code> parameter to display detailed information about each adapter.</p>

### Managing the discovery of cluster and management network switches

If you want to...	Use this command...
Display the switches that the cluster monitors	<pre>system health cluster-switch show</pre>
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.	<pre>system health cluster-switch show- all</pre>
Configure discovery of an undiscovered switch	<pre>system health cluster-switch create</pre>
Modify information about a switch that the cluster monitors (for example, device name, IP address, SNMP version, and community string)	<pre>system health cluster-switch modify</pre>
Disable monitoring of a switch	<pre>system health cluster-switch modify -disable-monitoring</pre>
Display the interval in which the health monitor polls switches to gather information	<pre>system health cluster-switch polling-interval show</pre>
Modify the interval in which the health monitor polls switches to gather information  This command is available at the advanced privilege level.	<pre>system health cluster-switch polling-interval modify</pre>
Disable discovery and monitoring of a switch and delete switch configuration information	<pre>system health cluster-switch delete</pre>
Permanently remove the switch configuration information which is stored in the database (doing so re-enables automatic discovery of the switch)	<pre>system health cluster-switch delete -force</pre>

**Responding to generated alerts**

<b>If you want to...</b>	<b>Use this command...</b>
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>
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**

<b>If you want to...</b>	<b>Use this command...</b>
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**

<b>If you want to...</b>	<b>Use this command...</b>
Display information about health monitors, such as their nodes, names, subsystems, and status	<code>system health config show</code>  <b>Note:</b> Use the <code>-instance</code> parameter to display detailed information about each health monitor.

If you want to...	Use this command...
Display information about the alerts that a health monitor can potentially generate	<pre>system health alert definition show</pre> <p><b>Note:</b> Use the <code>-instance</code> parameter to display detailed information about each alert definition.</p>
Display information about health monitor policies, which determine when alerts are raised	<pre>system health policy definition show</pre> <p><b>Note:</b> 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.</p>

### 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	<pre>system health cluster-switch prepare-to-downgrade</pre>

## 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:

<b>aggregate-used</b>	The following EMS messages are related to this object type: <ul style="list-style-type: none"> <li>• <code>mgmtgwd.aggregate.used.rising</code></li> <li>• <code>mgmtgwd.aggregate.used.falling</code></li> </ul>
<b>cpu-busy</b>	The following EMS messages are related to this object type: <ul style="list-style-type: none"> <li>• <code>mgmtgwd.cpu.busy.rising</code></li> <li>• <code>mgmtgwd.cpu.busy.falling</code></li> </ul>
<b>op-latency</b>	The following EMS messages are related to this object type: <ul style="list-style-type: none"> <li>• <code>mgmtgwd.op.latency.rising</code></li> <li>• <code>mgmtgwd.op.latency.falling</code></li> </ul>
<b>port-problems</b>	The following EMS messages are related to this object type: <ul style="list-style-type: none"> <li>• <code>mgmtgwd.port.problems.rising</code></li> <li>• <code>mgmtgwd.port.problems.falling</code></li> </ul>
<b>port-util</b>	The following EMS messages are related to this object type: <ul style="list-style-type: none"> <li>• <code>mgmtgwd.port.util.rising</code></li> <li>• <code>mgmtgwd.port.util.falling</code></li> </ul>

### 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

You use the `dashboard` commands to configure dashboards, display dashboard information, and display health status for SVMs.

**Note:** The `dashboard health vservers` 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"> <li>• The threshold value that generates a warning or critical alarm for an event</li> <li>• Whether an EMS message is sent when an alarm is generated</li> <li>• The interval at which objects are monitored by the alarm dashboard</li> </ul>	<code>dashboard alarm thresholds modify</code>
Display settings about alarm thresholds	<code>dashboard alarm thresholds show</code>

If you want to...	Use this command...
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.

### 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.



## 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?	You have two choices: <ul style="list-style-type: none"> <li>• 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.</li> <li>• 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.</li> </ul>
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"> <li>• If you collect data for a time period, the default setting shows data for all instances; however, you can specify one or more instances.</li> <li>• If you view continuously updated data and specify any object other than <code>cluster</code>, you must specify an instance.</li> </ul>
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.

For more information about the `statistics` commands, see the man pages.

**Steps**

1. 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 during the same CLI session and do not specify the `-sample-id` parameter, the command overwrites the previous default sample.

2. 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.

3. 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

## 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.

### Step

1. 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:

- What is the average response time for a volume?
- How many operations are completing 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
cluster1: 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
------------	-------	-------	----	-----	-------	-----

## Commands for monitoring clusterSVM performance

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 a time period

Use the following commands to collect data samples and to manage the samples that you collect. You need to collect a data sample before you can use the `statistics show` command.

If you want to...	Use this command...
Start data collection for a sample	<code>statistics start</code>
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

Use the following commands to view performance data. You need to collect a data sample before you can use the `statistics show` command.

If you want to...	Use this command...
View performance data for a time period (a sample)	<code>statistics show</code>  <b>Note:</b> 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>

### Viewing all objects, instances, and counters

Use the `statistics catalog` commands to view information about objects, instances, and counters.

If you want to...	Use this command...
View descriptions of objects	<code>statistics catalog object show</code>

If you want to...	Use this command...
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

Use the `statistics settings` commands to modify settings for the `statistics` commands.

If you want to...	Use this command...
View the settings for the statistics 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

Use the following commands to view advanced performance data about your cluster.

**Note:** The following commands are deprecated and will be removed in a future major release.

If you want to...	Use this command...
View information about SecD RPC usage statistics for the nodes in the cluster	<code>statistics secd show</code>  <b>Note:</b> This command is available at the advanced privilege level. Use this command only as directed by support personnel to help analyze performance and diagnose problems.
View information about the contents of the Open Network Computing Remote Procedure Call (ONC RPC) replay caches for the nodes in the cluster	<code>statistics oncrpc show-replay-cache</code>
View information about the ONC RPC calls performed by the nodes in the cluster	<code>statistics oncrpc show-rpc-calls</code>

## 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.


Storage QoS is supported on clusters that have up to eight nodes.

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 257

**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:

<b>User-defined policy group</b>	Enforces a maximum throughput limit on the storage objects that belong to the policy group by throttling input/output (I/O) requests.
<b>System-defined policy group</b>	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:

<b>User-defined workload</b>	<p>Represents the input/output (I/O) operations from clients to a storage object that belongs to a policy group. Storage objects can be:</p> <ul style="list-style-type: none"> <li>• SVMs with FlexVol volumes</li> <li>• FlexVol volumes</li> <li>• LUNs</li> <li>• Files (typically represent virtual machines)</li> </ul> <p>I/O to storage objects that are not assigned to policy groups belongs to the "User-Default" workload.</p>
<b>System-defined workload</b>	<p>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.</p>

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 a...	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
- FlexCache volumes

**Note:** You can assign the origin volume to a policy group, which controls the origin volume and its FlexCache 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 253.
- You must be familiar with [Rules for assigning storage objects to QoS policy groups](#) on page 254.

### About this task

Storage QoS is supported on clusters that have up to eight nodes.

### 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 263

## 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 pgl
```

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-widl2459    12459    10063      39.31MB/s      410.00us
vol2-widl445     1445     3505       13.69MB/s      437.00us
vol1-widl1344    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-widl1344    11344    6016       23.50MB/s      355.00us
vol3-widl2459    12459    5133       20.05MB/s      2.42ms
vol2-widl445     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 263
- [Commands for assigning storage objects to policy groups](#) on page 263
- [Commands for identifying the storage objects that belong to policy groups](#) on page 264
- [Commands for monitoring policy group performance](#) on page 264
- [Commands for monitoring workload performance](#) on page 264

For more information about these commands, see the man pages.

## Commands for managing policy groups

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

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>

<b>If you want to assign the..</b>	<b>Use this command with the <code>-qos-policy-group</code> parameter...</b>
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

Use a storage object's show command to identify the storage objects that belong to policy groups.

<b>If you want to identify the...</b>	<b>Use this command with the <code>-qos-policy-group</code> parameter...</b>
SVMs with FlexVol volumes that belong to a policy group	<code>vserver show</code>
FlexVol volumes that belong to a policy group	<code>volume show</code>
LUNs that belong to a policy group	<code>lun show</code>
Files that belong to a policy group	<code>volume file show</code>

### Commands for monitoring policy group and workload performance

Use the following commands to monitor policy group and workload performance in terms of IOPS, throughput, and latency.

<b>If you want to view the...</b>	<b>Use this command...</b>
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

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

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>

## Increasing WAFL cache memory

You can increase Write Anywhere File Layout (WAFL) cache memory in a system that has a caching module installed (Performance Acceleration Module (PAM), Flash Cache module, or Flash Cache 2 module). To increase the WAFL cache memory, you use the WAFL external cache, a software component of Data ONTAP.

WAFL external cache provides extra WAFL cache memory to improve the performance of the storage system by reducing the number of disk reads. You can control how user data blocks are cached by changing the mode of operation for a caching module. You can keep the default mode (normal user data blocks) or you can choose metadata mode or low-priority blocks mode.

You should verify that the WAFL external cache functionality is enabled after you install a caching module.

**Note:** WAFL external cache does not require a separate license if your system is running Data ONTAP 8.1 or later.

**Note:** Not all systems have a caching module installed. Therefore, not all systems can utilize the WAFL external cache functionality.

WAFL external cache does not cache data that is stored in a RAID group composed of SSDs.

If you use WAFL external cache on storage systems with a high-availability configuration, you must ensure that the WAFL external cache options are the same on both nodes. Otherwise, a takeover can result in lower performance due to the lack of WAFL external cache on the remaining node.

Besides the Data ONTAP options that you can use to manage WAFL external cache, a diagnostic command is available for sanitizing a caching module. For more information, see the *Diagnostics Guide*.

## How Flash Pool aggregates and Flash Cache compare

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 (or an SSD 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
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
Supported with array LUNs?	No	Yes

For more information about Flash Pool aggregates, see the *Clustered Data ONTAP Physical Storage Management Guide*.

## 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}
```

## 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`

## 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									Reads		Writes		Disk Read	
Usage	Hit	Meta	Miss	Hit	Evict	Inval	Insrt	Chain	Blcks	Chain	Blcks	Replcd		
%	/s	/s	/s	%	/s	/s	/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 265

## 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]`

## Events that abort cache rewarming

After the cache rewarming process starts, some events can abort the entire process and some events can abort the process on specific aggregates.

The following events abort the entire cache rewarming process:

- You add, remove, or move a Flash Cache or Flash Cache 2 module after the WAFL external cache takes the snapshot, 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 `wafiron` process mounts the root aggregate.

The following events 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 `wafiron` process mounts an aggregate.

## 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 267

## 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>  <b>Note:</b> <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 272

## 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 <code>on</code>
Disable free space reallocation on an aggregate	<code>storage aggregate modify</code> with the <code>-free-space-realloc</code> parameter set to <code>off</code>
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 295

## What a cluster peer is

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

## Connecting one cluster to another cluster in a peer relationship

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

### About this task

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

### Steps

1. [Cluster peer network topologies](#) on page 277
2. [What cluster peer intercluster networking is](#) on page 279
3. [Cluster peer intercluster networking requirements](#) on page 279
4. [Considerations when sharing data ports](#) on page 280
5. [Considerations when using dedicated ports](#) on page 281
6. [Configuring intercluster LIFs to share data ports](#) on page 281
7. [Configuring intercluster LIFs to use dedicated intercluster ports](#) on page 285
8. [Creating the cluster peer relationship](#) on page 290


### Cluster peer network topologies

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

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

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

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




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

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

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

- D1a is connected to S1a and S1b
- D1b is connected to S1a and S1b
- D2a is connected to S2a and S2b

- D2b is connected to S2a and S2b
- D3a is connected to S3a and S3b
- D3b is connected to S3a and S3b




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

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

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

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



### What cluster peer intercluster networking is

A cluster peer relationship, that is, two different clusters communicating with each other, requires an intercluster network on which the communication occurs. An intercluster network consists of intercluster logical interfaces (LIFs) that are assigned to network ports.

The intercluster network 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 network only; this is true regardless of whether the intercluster network 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, an intercluster routing group is automatically created on that node too. A gateway address for the intercluster routing group must be defined and the intercluster routing group must be routed to the defined gateway address.

Intercluster LIFs can be assigned to ports that have the role of data, which are the same ports used for CIFS or NFS access, or intercluster LIFs can be assigned to dedicated ports that have the role of intercluster. Each method has its advantages and disadvantages.

### Cluster peer intercluster networking requirements

Your cluster peer intercluster network must fulfill requirements that include synchronized cluster time, number of intercluster LIFs, IP addresses for intercluster LIFs, maximum transmission units, and more.

The following are requirements of cluster peer intercluster networking:

- The time on the clusters that you want to connect using an intercluster network must be synchronized within 300 seconds (5 minutes).

Cluster peers can be in different time zones.

- At least one intercluster LIF must be created on every node in the cluster.
- Every intercluster LIF requires an IP address dedicated for intercluster replication.

**Note:** The IPv6 communication protocol is not supported.

- The correct maximum transmission unit (MTU) value must be used on the network ports that are used for replication.

The network administrator can identify which MTU value to use in the environment. The MTU value should be set to a value that is supported by the network end point to which it is connected. The default value of 1,500 is correct for most environments.

- All paths on a node used for intercluster networking should have equal performance characteristics.
- The intercluster network must provide connectivity among all intercluster LIFs on all nodes in the cluster peers.

Every intercluster LIF on every node in a cluster must be able to connect to every intercluster LIF on every node in the peer cluster.

## Considerations when sharing data ports

When determining whether sharing a data port for intercluster replication is the correct interconnect 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 interconnect network 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, which reduces the LAN network utilization to only that which the WAN is capable of supporting.
- 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.
- If the replication interval is set to perform only after hours, when little or no client activity exists, then using data ports for replication during this time is acceptable, even without a 10-GbE LAN connection.
- 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 interconnect 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 interconnect 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.
- 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.
- If the replication network requires configuration of a maximum transmission unit (MTU) size that differs from the MTU size used on the data network, then you must use physical ports for replication because the MTU size can only be configured on physical 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.

### Before you begin

You should have reviewed the considerations for sharing data ports and determined that this is an appropriate intercluster networking configuration.

### About this task

Creating intercluster LIFs that share data ports involves assigning LIFs to existing data ports and, possibly, creating an intercluster route. In this procedure, a two-node cluster exists in which each node has two data ports, e0c and e0d. These are the two data ports that are shared for intercluster replication. In your own environment, you replace the ports, networks, IP addresses, subnet masks, and subnets with those specific to your environment.

### Steps

1. Check the role of the ports in the cluster by using the `network port show` command.

**Example**

```
cluster01::> network port show
```

Node	Port	Role	Link	MTU	Auto-Negot Admin/Oper	Duplex Admin/Oper	Speed(Mbps) Admin/Oper
-----							
cluster01-01							
	e0a	cluster	up	1500	true/true	full/full	auto/1000
	e0b	cluster	up	1500	true/true	full/full	auto/1000
	e0c	data	up	1500	true/true	full/full	auto/1000
	e0d	data	up	1500	true/true	full/full	auto/1000
cluster01-02							
	e0a	cluster	up	1500	true/true	full/full	auto/1000
	e0b	cluster	up	1500	true/true	full/full	auto/1000
	e0c	data	up	1500	true/true	full/full	auto/1000
	e0d	data	up	1500	true/true	full/full	auto/1000

2. Create a failover group on each node on both clusters for the data ports you want to share for intercluster communications by using the `network interface failover-groups create` command.

**Example**

The first of the following commands creates the failover group `failover_cluster01-01_icl01` and assigns port `e0c` to the failover group. The second command adds port `e0d` to the failover group.

```
cluster01::> network interface failover-groups create -failover-group
failover_cluster01-01_icl01 -node cluster01-01 -port e0c

cluster01::> network interface failover-groups create -failover-group
failover_cluster01-01_icl01 -node cluster01-01 -port e0d
```

3. Create an intercluster LIF on each node in cluster01 by using the `network interface create` command.

**Example**

The following commands create an intercluster LIF on each of the two nodes in cluster01 and assign the intercluster LIF to the failover group you created that contains the data ports.

```
cluster01::> network interface create -vserver cluster01-01 -lif cluster01-01_icl01 -role
intercluster -home-node cluster01-01 -home-port e0c -address 192.168.1.201 -netmask
255.255.255.0 -failover-group failover_cluster01-01_icl01 -failover-policy nextavail
Info: Your interface was created successfully; the routing group i192.168.1.0/24 was
created

cluster01::> network interface create -vserver cluster01-02 -lif cluster01-02_icl01 -role
intercluster -home-node cluster01-02 -home-port e0c -address 192.168.1.202 -netmask
255.255.255.0 -failover-group failover_cluster01-02_icl01 -failover-policy nextavail
Info: Your interface was created successfully; the routing group i192.168.1.0/24 was
created
```

4. 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
-----
cluster01-01
  cluster01-01_ic101
    up/up      192.168.1.201/24  cluster01-01  e0c         true
cluster01-02
  cluster01-02_ic101
    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    Group Usage   Group
-----
cluster01-01
  cluster01-01_ic101  cluster01-01:e0c  nextavail    failover_cluster01-01_ic101
    Failover Targets: cluster01-01:e0c,
    cluster01-01:e0d
cluster01-02
  cluster01-02_ic101  cluster01-02:e0c  nextavail    failover_cluster01-02_ic101
    Failover Targets: cluster01-02:e0c,
    cluster01-02:e0d
```

- Display routing groups by using the `network routing-group show` command with the `-role intercluster` parameter.

An intercluster routing group is created automatically for the intercluster LIFs.

### Example

```
cluster01::> network routing-group show -role intercluster
-----
Routing
Vserver      Group      Subnet      Role      Metric
-----
cluster01-01
  i192.168.1.0/24
    192.168.1.0/24  intercluster 40
cluster01-02
  i192.168.1.0/24
    192.168.1.0/24  intercluster 40
```

- Display the routes in the cluster by using the `network routing-group 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 routing-group route show
Routing
Vserver  Group      Destination      Gateway          Metric
-----  -
cluster01
  c192.168.0.0/24
    0.0.0.0/0      192.168.0.1     20
cluster01-01
  n192.168.0.0/24
    0.0.0.0/0      192.168.0.1     10
cluster01-02
  n192.168.0.0/24
    0.0.0.0/0      192.168.0.1     10
```

8. If communication between intercluster LIFs in different clusters requires routing, create an intercluster route by using the `network routing-groups route create` command.

The intercluster networks apply to each node; therefore, you must create an intercluster route on each node.

**Example**

In this example, 192.168.1.1 is the gateway address for the 192.168.1.0/24 network.

**Note:** If the destination is specified as 0.0.0.0/0, then it becomes the default route for the intercluster network.

```
cluster01::> network routing-groups route create -server cluster01-01
-routing-group i192.168.1.0/24 -destination 0.0.0.0/0 -gateway 192.168.1.1 -metric 40
cluster01::> network routing-groups route create -server cluster01-02
-routing-group i192.168.1.0/24 -destination 0.0.0.0/0 -gateway 192.168.1.1 -metric 40
```

9. Display the newly created routes by using the `network routing-groups route show` command.

Although the intercluster routes do not have an assigned role, they are assigned to the routing group `i192.168.1.0/24`, which is assigned the role of intercluster. These routes are only used for intercluster communication.

**Example**

```
cluster01::> network routing-group route show
Routing
Vserver  Group      Destination      Gateway          Metric
-----  -
cluster01
  c192.168.0.0/24
    0.0.0.0/0      192.168.0.1     20
cluster01-01
  n192.168.0.0/24
    0.0.0.0/0      192.168.0.1     10
  i192.168.1.0/24
    0.0.0.0/0      192.168.1.1     40
cluster01-02
```

```

n192.168.0.0/24
    0.0.0.0/0      192.168.0.1    10
i192.168.1.0/24
    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

In this example, a two-node cluster exists in which each node has two data ports, e0e and e0f, which are dedicated for intercluster replication. In your own environment, you would replace the ports, networks, IP addresses, subnet masks, and subnets with those specific to your environment.

### Steps

- Check the role of the ports in the cluster by using the `network port show` command.

#### Example

```

cluster01::> network port show
Node  Port  Role      Link  MTU  Auto-Negot  Duplex  Speed(Mbps)
-----  ---  ---      ---  ---  ---  ---  ---
cluster01-01
  e0a  cluster  up      1500  true/true  full/full  auto/1000
  e0b  cluster  up      1500  true/true  full/full  auto/1000
  e0c  data    up      1500  true/true  full/full  auto/1000
  e0d  data    up      1500  true/true  full/full  auto/1000
  e0e  data    up      1500  true/true  full/full  auto/1000
  e0f  data    up      1500  true/true  full/full  auto/1000
cluster01-02
  e0a  cluster  up      1500  true/true  full/full  auto/1000
  e0b  cluster  up      1500  true/true  full/full  auto/1000
  e0c  data    up      1500  true/true  full/full  auto/1000
  e0d  data    up      1500  true/true  full/full  auto/1000
  e0e  data    up      1500  true/true  full/full  auto/1000
  e0f  data    up      1500  true/true  full/full  auto/1000
    
```

- Determine whether any of the LIFs are using ports that are dedicated for replication by using the `network interface show` command.

#### Example

```

cluster01::> network interface show
Vserver  Logical  Status  Network  Current  Current  Is
-----  ---  ---  ---  ---  ---  ---
cluster01
vs1      cluster_mgmt  up/up    192.168.0.xxx/24  cluster01-01  e0c    true
vs1      vs1_lif1      up/up    192.168.0.151/24  cluster01-01  e0e    true
    
```

3. If a LIF is using one of the ports dedicated to replication, then migrate the LIF to another port by using the `network interface migrate` command.

This is required because intercluster ports cannot host data LIFs. This migration is nondisruptive, assuming that the other data ports have been configured properly so that clients can access the LIF after migration.

### Example

```
cluster01::> network interface migrate -vserver vs1 -lif vs1_lif1 -dest-node
cluster01-01 -dest-port e0d

cluster01::> network interface show
Vserver      Logical   Status   Network      Current   Current   Is
Interface   Admin/Oper Address/Mask Node       Port      Home
-----
cluster01
vs1          cluster_mgmt up/up    192.168.0.xxx/24 cluster01-01 e0c      true
vs1          vs1_lif1   up/up    192.168.0.151/24 cluster01-01 e0d      false
```

4. If necessary, modify the newly migrated LIFs to the LIF home port by using the `network interface modify` command, so that the new port is the LIF home port.

### Example

```
cluster01::> network interface modify -vserver vserver1 -lif vs1_lif1 -home-node
dpg1-01 -home-port e0d

cluster01::> network interface show
Vserver      Logical   Status   Network      Current   Current   Is
Interface   Admin/Oper Address/Mask Node       Port      Home
-----
vserver1
vs1_lif1     up/up    192.168.0.151/24 cluster01-01 e0d      true
```

5. After all LIFs have been migrated off the ports dedicated for replication, change the role of the port used on each node to `intercluster` by using the `network port modify` command.

### Example

```
cluster01::> network port modify -node cluster01-01 -port e0e -role intercluster
cluster01::> network port modify -node cluster01-01 -port e0f -role intercluster
cluster01::> network port modify -node cluster01-02 -port e0e -role intercluster
cluster01::> network port modify -node cluster01-02 -port e0f -role intercluster
```

6. Verify that the roles of the correct ports have been changed to `intercluster` by using the `network port show` command with the `-role intercluster` parameter.

### Example

```
cluster01::> network port show -role intercluster
Auto-Negot Duplex      Speed(Mbps)
```

Node	Port	Role	Link	MTU	Admin/Oper	Admin/Oper	Admin/Oper
cluster01-01	e0e	intercluster	up	1500	true/true	full/full	auto/1000
	e0f	intercluster	up	1500	true/true	full/full	auto/1000
cluster01-02	e0e	intercluster	up	1500	true/true	full/full	auto/1000
	e0f	intercluster	up	1500	true/true	full/full	auto/1000

7. Create an intercluster LIF on each node in cluster01 by using the `network interface create` command.

### Example

This example uses the LIF naming convention `nodename_icl#` for the intercluster LIF.

```
cluster01::> network interface create -vserver cluster01-01 -lif cluster01-01_icl01 -role
intercluster -home-node cluster01-01 -home-port e0e
-address 192.168.1.201 -netmask 255.255.255.0
Info: Your interface was created successfully; the routing group i192.168.1.0/24 was
created

cluster01::> network interface create -vserver cluster01-02 -lif cluster01-02_icl01 -role
intercluster -home-node cluster01-02 -home-port e0e
-address 192.168.1.202 -netmask 255.255.255.0
Info: Your interface was created successfully; the routing group i192.168.1.0/24 was
created
```

8. 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 because e0f is also assigned the role of intercluster.

The intercluster LIF is assigned to an intercluster port; therefore, a non-modifiable failover group is created automatically, and contains all ports with the intercluster role on that node. Intercluster failover groups are node specific; therefore, if changes are required, they must be managed for each node because different nodes might use different ports for replication.

```
cluster01::> network interface show -role intercluster -failover
Logical          Home          Failover      Failover
Vserver  Interface      Node:Port      Group  Usage      Group
-----
cluster01-01
  cluster01-01_icl01  cluster01-01:e0e  system-defined
                                     Failover Targets: cluster01-01:e0e,
                                     cluster01-01:e0f
cluster01-02
  cluster01-02_icl01  cluster01-02:e0e  system-defined
                                     Failover Targets: cluster01-02:e0e,
                                     cluster01-02:e0f
```

9. 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
cluster01	cluster_mgmt	up/up	192.168.0.xxx/24	cluster01-01	e0c	true
cluster01-01	cluster01-01_icl01	up/up	192.168.1.201/24	cluster01-01	e0e	true
	clus1	up/up	169.254.xx.xx/24	cluster01-01	e0a	true
	clus2	up/up	169.254.xx.xx/24	cluster01-01	e0b	true
	mgmt1	up/up	192.168.0.xxx/24	cluster01-01	e0c	true
cluster01-02	cluster01-02_icl01	up/up	192.168.1.202/24	cluster01-02	e0e	true
	clus1	up/up	169.254.xx.xx/24	cluster01-02	e0a	true
	clus2	up/up	169.254.xx.xx/24	cluster01-02	e0b	true
	mgmt1	up/up	192.168.0.xxx/24	cluster01-02	e0c	true

10. Display routing groups by using the `network routing-group show` command with the `-role intercluster` parameter to determine whether the intercluster network needs intercluster routes.

An intercluster routing group is created automatically for the intercluster LIFs.

**Example**

```
cluster01::> network routing-group show -role intercluster
```

Vserver	Routing Group	Subnet	Role	Metric
cluster01-01	i192.168.1.0/24	192.168.1.0/24	intercluster	40
cluster01-02	i192.168.1.0/24	192.168.1.0/24	intercluster	40

11. Display the routes in the cluster by using the `network routing-group 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 routing-group route show
```

Vserver	Routing Group	Destination	Gateway	Metric
cluster01	c192.168.0.0/24	0.0.0.0/0	192.168.0.1	20
cluster01-01	n192.168.0.0/24	0.0.0.0/0	192.168.0.1	10
cluster01-02				

```
n192.168.0.0/24
0.0.0.0/0      192.168.0.1    10
```

12. If communication between intercluster LIFs in different clusters requires routing, create an intercluster route by using the `network routing-groups route create` command.

The intercluster networks apply to each node; therefore, you must create an intercluster route on each node.

### Example

In this example, 192.168.1.1 is the gateway address for the 192.168.1.0/24 network.

**Note:** If the destination is specified as 0.0.0.0/0, then it becomes the default route for the intercluster network.

```
cluster01::> network routing-groups route create -server cluster01-01 -routing-
group i192.168.1.0/24 -destination 0.0.0.0/0 -gateway 192.168.1.1 -metric 40

cluster01::> network routing-groups route create -server cluster01-02 -routing-
group i192.168.1.0/24 -destination 0.0.0.0/0 -gateway 192.168.1.1 -metric 40
```

13. Display the newly created routes by using the `network routing-groups route show` command to confirm that you created the routes correctly.

Although the intercluster routes do not have an assigned role, they are assigned to the routing group `i192.168.1.0/24`, which is assigned the role of `intercluster`. These routes are only used for intercluster communication.

### Example

```
cluster01::> network routing-group route show
```

Vserver	Routing Group	Destination	Gateway	Metric
cluster01	c192.168.0.0/24	0.0.0.0/0	192.168.0.1	20
cluster01-01	n192.168.0.0/24	0.0.0.0/0	192.168.0.1	10
	i192.168.1.0/24	0.0.0.0/0	192.168.1.1	40
cluster01-02	n192.168.0.0/24	0.0.0.0/0	192.168.0.1	10
	i192.168.1.0/24	0.0.0.0/0	192.168.1.1	40

14. Repeat these steps to configure intercluster networking in the other cluster.
15. 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 designated logical interfaces to make information about one cluster available to the other cluster for use in cluster peering applications.

### Before you begin

You should have the intercluster network configured.

### Steps

1. Create the cluster peer relationship by using the `cluster peer create` command.

#### Example

In the following example, cluster01 is peered with a remote cluster named cluster02. 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 -username admin
Password: *****
```

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.

2. Display the cluster peer relationship by using the `cluster peer show` command with the `-instance` parameter.

#### 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 cluster peer relationship by using the `cluster peer health show` command.

## Example

```

cluster01::> cluster peer health show
Node      cluster-Name      Node-Name
         Ping-Status      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
    
```

## Displaying a cluster peer relationship

You can see if a cluster is connected to another cluster if you want to make use of cluster peer features such as mirroring a volume from one cluster to another.

### Step

1. To display information about a cluster to which you previously connected, use the `cluster peer show` command.

This command displays only basic information about the other cluster. If you want to see more information about the other cluster, use the `cluster peer show -instance` command.

The following example displays basic information about a cluster connected to a cluster named `cluster_a`:

```

cluster_a::>cluster peer show
Peer Cluster Name      Cluster Serial Number LIF Role      Availability
-----
cluster_b              1-80-123456          intercluster Available
    
```

## 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.

### Steps

1. To delete the cluster peering relationship from the cluster of which you are the administrator, use the `cluster peer delete` command.

**Note:** This procedure assumes that you are the administrator of only one of the clusters in the cluster peering relationship.

### Example

The following example deletes the cluster peering relationship with the `cluster 2` cluster from the `cluster1` cluster:

```
cluster1::> cluster peer delete -cluster cluster2
```

2. To delete the cluster peering relationship from the other cluster, an administrator of the other cluster uses the `cluster peer delete` command.

### Example

The following example deletes the cluster peering relationship with the `cluster 1` cluster from the `cluster2` cluster:

```
cluster2::> cluster peer delete -cluster cluster1
```

## Managing jobs on another cluster

From the local cluster, you can manage jobs that are running on the cluster to which the local cluster is connected. This is useful for monitoring and controlling cross cluster applications like a data protection mirror.

### Viewing jobs on another cluster

From the local cluster, you can see jobs that are running on the cluster to which the local cluster is connected. This is useful for monitoring cross cluster applications like a data protection mirror.

#### Step

1. To view the jobs running on a cluster which is connected to the local cluster in a cross cluster relationship, complete the following step.

```
cluster peer job show -cluster cluster_name
```

The following example shows information about jobs running on a cluster named `cluster_a` connected to the local cluster in a cross cluster relationship:

```
cluster_b::cluster peer job show -cluster cluster_a
Job ID      Name                State      Description
-----
Cluster: cluster_a
5           Vol Create         Running   create striped-volume vol0
```

You can show detailed information using the `-instance` parameter of the `cluster peer job show` command. See the `cluster peer job show` command for details.

## Monitoring progress of a job on another cluster

From the local cluster, you can monitor the progress of a job on the cluster to which the local cluster is connected.

#### About this task

The progress of the job is monitored until the job ends, terminates, or you interrupt its progress.

#### Step

1. To monitor a job on a cluster which is connected to the local cluster in a cross cluster relationship, complete the following step.

```
cluster peer job watch-progress -cluster cluster_a -id ID_number -interval integer
```

The following example monitors the progress of a job whose job ID is 15. The progress is updated every 3 seconds.

```
cluster_b::cluster peer job watch-progress -cluster cluster_a -id 15
-interval 3
```

See the `cluster peer job watch-progress` command for details.

## Pausing jobs on another cluster

From the local cluster, you can pause jobs that are running on the cluster to which the local cluster is connected. If a job is consuming too many system resources, you can pause it until there is less demand on the system.

### Step

1. To pause a job running on a cluster which is connected to the local cluster in a cross cluster relationship, complete the following step.

```
cluster peer job pause -cluster cluster_name -id ID_number
```

The following example pauses a job whose job ID is 15:

```
cluster_b::cluster peer job pause -cluster cluster_a -id 15
```

See the `cluster peer job pause` command for details.

## Resuming jobs on another cluster

From the local cluster, you can resume paused jobs on the cluster to which the local cluster is connected.

### Step

1. To resume a paused job on a cluster which is connected to the local cluster in a cross cluster relationship, complete the following step.

```
cluster peer job resume -cluster cluster_name -id ID_number
```

The following example resumes a job whose job ID is 15:

```
cluster_b::cluster peer job resume -cluster cluster_a -id 15
```

See the `cluster peer job resume` command for details.

## Stopping a job on another cluster

From the local cluster, you can stop a job on the cluster to which the local cluster is connected.

### About this task

If you stop a job, you cannot resume it using the `cluster peer job resume` command.

### Step

1. To stop a job on a cluster which is connected to the local cluster in a cluster peer relationship, use the `cluster peer job stop` command.

The following example stops a job whose job ID is 15:

```
cluster_b::cluster peer job stop -cluster cluster_a -id 15
```

See the `cluster peer job stop` command for details.

## Deleting a job on another cluster

From the local cluster, you can delete a job on the cluster to which the local cluster is connected.

### Step

1. To delete a job on a cluster which is connected to the local cluster in a cluster peer relationship, use the `cluster peer job delete` command.

The following example deletes a job whose job ID is 15:

```
cluster_b::cluster peer job delete -cluster cluster_a -id 15
```

See the `cluster peer job delete` command for details.

## Managing SVM peer relationships

A cluster administrator can create and manage data protection mirroring relationships between two Storage Virtual Machines (SVMs, formerly known as Vservers) either existing within a cluster (intracluster) or in the 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 peer relationships
- Deleting SVM peer relationships
- Viewing SVM peer relationships
- Setting up SnapMirror relationship between the volumes of the peered SVMs

**Note:** You cannot set up load-sharing SnapMirror relationship between the volumes of intercluster SVM peers.

An SVM administrator can perform only the following SVM peer management tasks:

- Viewing the SVM peer relationships to identify the peered SVMs
- Setting up SnapMirror relationship such as data protection relationship (DP), SnapVault relationship (XDP), and transition relationship (TDP) between the volumes of the peered SVMs

For more information about setting up peering applications, see the *Clustered Data ONTAP Data Protection Guide*.

**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 276

## What an SVM peer relationship is

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:



SVM peer infrastructure enables you to set up a backup and recovery mechanism between SVMs. You can set up mirroring relationship at volume level between peered SVMs. If a volume of 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.

In clustered Data ONTAP 8.2, only SnapMirror data protection (DP), SnapVault (XDP) and load-sharing relationship (LS) relationships can be set up 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	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	An intercluster SVM peer relationship is requested from the local cluster
peered state on the local and peered clusters	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

SVM peer relationship is in...	When...
rejected state on the local cluster	An intercluster SVM peer relationship is rejected from the peered cluster
suspended state on the local and peered clusters	An intercluster or intracluster SVM peer relationship is suspended from the local or peered cluster
deleted state	An intercluster SVM peer relationship is deleted from any of the peered clusters

## Creating an SVM peer relationship

A cluster administrator can create an 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 existing either in a single cluster (intracluster) or existing 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.
- SVM peer relationship that is in rejected or deleted state must be deleted if you want to re-create the SVM peer relationship between the same SVMs.
- The admin state of the SVMs to be peered must not be `initializing` or `deleting`.
- The names of SVMs in the peered clusters must be unique across the two clusters. If they do not have unique names, you must rename one of the SVMs.

### About this task

Peered clusters can be managed by a single cluster administrator or different cluster administrators. In clustered Data ONTAP 8.2, you can set up only SnapMirror relationships between the peered SVMs. If you do not specify the peering application as `SnapMirror`, an SVM administrator cannot set up SnapMirror relationship between the peered SVMs.

You can create an SVM peer relationship either between SVM with FlexVol volumes or between SVM with Infinite Volumes. You cannot create an SVM peer relationship between SVM with FlexVol volume and SVM with Infinite Volume.

You can create only intercluster SVM peer relationship for SVMs with Infinite Volumes.

### Steps

1. Use the `vserver peer create` command to create an SVM peer relationship.

### Example

The following example illustrates how to create an intercluster SVM peer relationship between vs1.example0.com and vs3.example0.com residing on cluster1 and cluster2 respectively:

```
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
```

At this point, the state of the intercluster SVM peer relationship is initiated. The SVM peer relationship is not established until the cluster administrator of the peered cluster accepts the SVM peer request.

### Example

The following example illustrates how to create an intracluster SVM peer relationship between SVMs vs4.example1.com and vs0.example1.com 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 when the command is executed. Authentication is not needed as the cluster is managed by a single cluster administrator. The state of the SVM peer relationship is peered.

2. Use the `vserver peer show-all` command to view the status and other details of the SVM peer relationship.

### Example

The following example illustrates how to view the status and other details of the SVM peer relationship:

```
cluster1::> vserver peer show-all
```

Vserver	Peer Vserver	Peer State	Peer Cluster	Peering Applications
vs1.example0.com	vs3.example0.com	initiated	Cluster2	snapmirror

```
Cluster2::> vserver peer show-all
```

Vserver	Peer Vserver	Peer State	Peer Cluster	Peering Applications
vs3.example0.com	vs1.example0.com	pending	cluster1	snapmirror
vs4.example1.com	vs0.example1.com	peered	cluster2	snapmirror

For more information about these commands, see the man pages.

**Note:** You can use the `job show -fields state, completion -id` to view the status of the intercluster operations.

### After you finish

If you have initiated an intercluster SVM peer relationship, you must inform the cluster administrator of the remote cluster about the SVM peer request. After the cluster administrator of the remote cluster accepts the SVM peer request, the SVM peer relationship is established.

## 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
```

Vserver	Peer 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
-----          -
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
      Peer
Vserver  Vserver  Peer
-----  -
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
      Peer
Vserver  Vserver  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 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. In clustered Data ONTAP 8.2, the SVM peering relationship can either have SnapMirror or no application. The default value is `snapmirror`.

### About this task

If the value of the application parameter is " ", then a cluster or SVM administrator cannot set up SnapMirror relationship between the peered SVMs.

### Steps

1. Use the `vserver peer modify` command to modify the application on the SVM peer relationship.

### Example

The following example shows how to modify the application on the SVM peer relationship:

```
cluster2::> vserver peer modify -vserver vs4.example1.com -peer-
vserver vs0.example1.com -applications " "
Info: [Job 78] 'vserver peer modify' job queued
```

2. Use the `vserver peer show-all` to view the applications running on the SVM peer relationship.

### Example

The following example shows how to view 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	-

For more information about these commands, see the man pages.

## Deleting an SVM peer relationship

A cluster administrator can delete the SVM peer relationship by using the `vserver peer delete` command when the relationship between two SVMs is no longer needed.

### About this task

When you are deleting an SVM peer relationship, you must delete the SVM peer relationship from both the peered clusters.

### Steps

1. Use the `vserver peer delete` command on both the clusters to delete an SVM peer relationship.

When the SVM peer relationship is deleted from one cluster, the relationship is in deleted state on the other peered cluster.

2. Use the `vserver peer show` command on both the clusters to view if the relationship is deleted.

### Example

The following example shows how to delete an SVM 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
```

```
cluster1::> vserver peer show
There are no Vserver peer relationships.
```

```
cluster2::> vserver peer show
```

Vserver	Peer Vserver	Peer State
vs3.example0.com	vs1.example0.com	deleted
vs4.example1.com	vs0.example1.com	peered

2 entries were displayed.

```
cluster2::> vserver peer delete -vserver vs3.example0.com -peer-
vserver vs1.example0.com
```

```
Info: 'vserver peer delete' command is successful.
```

```
cluster2::> vserver peer show
```

Vserver	Peer Vserver	Peer State
---------	-----------------	---------------

```
-----
vs4.example1.com    vs0.example1.com    peered
```

For more information about these commands, see the man pages.

## 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
      Peer
Vserver  Vserver      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 Cluster    Peering
-----                -
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 gateways (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, 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 computer system.

**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 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.

<b>CIFS share</b>	<ul style="list-style-type: none"> <li>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 <i>share</i>.</li> </ul>
<b>CLI</b>	command-line interface. The storage system prompt is an example of a command-line interface.
<b>client</b>	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.
<b>cluster</b>	<ul style="list-style-type: none"> <li>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.</li> <li>In the Data ONTAP 7.1 release family, a pair of storage systems (sometimes called <i>nodes</i>) configured to serve data for each other if one of the two systems stops functioning.</li> <li>In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an <i>active/active configuration</i>.</li> <li>For some storage array vendors, <i>cluster</i> refers to the hardware component on which host adapters and ports are located. Some storage array vendors refer to this component as a <i>controller</i>.</li> </ul>
<b>cluster Vserver</b>	Former name for a data SVM; see data SVM.
<b>Common Internet File System (CIFS)</b>	Microsoft's file-sharing networking protocol that evolved from SMB.
<b>community</b>	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.
<b>console</b>	The physical or virtual terminal that is used to monitor and control a storage system.
<b>Copy-On-Write (COW)</b>	The technique for creating Snapshot copies without consuming excess disk space.
<b>D</b>	
<b>data SVM</b>	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.

<b>degraded mode</b>	The operating mode of a storage system when a disk in the RAID group fails or the batteries on the NVRAM card are low.
<b>disk ID number</b>	The number assigned by the storage system to each disk when it probes the disks at startup.
<b>disk sanitization</b>	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.
<b>disk shelf</b>	A shelf that contains disk drives and is attached to a storage system.

## E

<b>emulated storage system</b>	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.
<b>Ethernet adapter</b>	An Ethernet interface card.

<b>expansion card</b>	A SCSI card, NVRAM card, network card, hot-swap card, or console card that plugs into a storage system expansion slot. Sometimes called an <i>adapter</i> .
-----------------------	---

<b>expansion slot</b>	The slots on the storage system board into which you insert expansion cards.
-----------------------	--

## F

<b>failed storage system</b>	A physical storage system that has ceased operating. In a high-availability configuration, it remains the failed storage system until a giveback succeeds.
------------------------------	--

<b>Flash Cache module</b>	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

<b>giveback</b>	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.
-----------------	---

<b>global namespace</b>	See <i>namespace</i> .
-------------------------	------------------------

<b>group</b>	In Data ONTAP operating in 7-Mode, a group of users defined in the storage system's <code>/etc/group</code> file.
--------------	---

<b>Group ID (GID)</b>	The number used by UNIX systems to identify groups.
-----------------------	---

## H

<b>HA (high availability)</b>	<ul style="list-style-type: none"> <li>• In Data ONTAP 8.x, the recovery capability provided by a pair of nodes (storage systems), called an <i>HA pair</i>, that are configured to serve data for each other if one of the two nodes stops functioning.</li> <li>• In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an <i>active/active configuration</i>.</li> </ul>
<b>HA pair</b>	<ul style="list-style-type: none"> <li>• 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.</li> <li>• In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an <i>active/active configuration</i>.</li> </ul>
<b>heartbeat</b>	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.
<b>hot swap</b>	The process of adding, removing, or replacing a disk while the storage system is running.
<b>hot swap adapter</b>	An expansion card that makes it possible to add or remove a hard disk with minimal interruption to file system activity.
<b>I</b>	
<b>inode</b>	A data structure containing information about files on a storage system and in a UNIX file system.
<b>interrupt switch</b>	A switch on some storage system front panels used for debugging purposes.
<b>M</b>	
<b>Maintenance mode</b>	An option when booting a storage system from a system boot disk. Maintenance mode provides special commands for troubleshooting hardware and configuration.
<b>N</b>	
<b>namespace</b>	In network-attached storage (NAS) environments, a collection of files and path names to the files.
<b>NDMP</b>	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.
<b>network adapter</b>	An Ethernet, FDDI, or ATM card.
<b>node SVM</b>	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 high-availability 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 storage system or gateway 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 a high-availability 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 an 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

<b>takeover</b>	The emulation of the failed node identity by the takeover node in a high-availability configuration; the opposite of <i>giveback</i> .
<b>takeover mode</b>	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.
<b>takeover node</b>	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.
<b>trap</b>	An asynchronous, unsolicited message sent by an SNMP agent to an SNMP manager indicating that an event has occurred on the storage system.
<b>U</b>	
<b>UID</b>	user identification number.
<b>Unicode</b>	A 16-bit character set standard. It was designed and is maintained by the nonprofit consortium Unicode Inc.
<b>V</b>	
<b>volume</b>	A file system.
<b>Vserver</b>	(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— <i>admin</i> , <i>node</i> , and <i>cluster</i> (“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.
<b>W</b>	
<b>WAFL</b>	Write Anywhere File Layout. A file system designed for the storage system to optimize write performance.
<b>WAFL External Cache</b>	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 <i>WAFL extended cache</i> .
<b>WINS</b>	Windows Internet Name Service.
<b>workgroup</b>	A collection of computers running Windows NT or Windows for Workgroups that is grouped for browsing and sharing.

---

## Copyright and trademark information

Copyright ©1994 - 2014 NetApp, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2014 IBM Corporation. All rights reserved.

US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

No part of this document covered by copyright may be reproduced in any form or by any means— graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

References in this documentation to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's or NetApp's intellectual property rights may be used instead of the IBM or NetApp product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM and NetApp, are the user's responsibility.

No part of this document covered by copyright may be reproduced in any form or by any means— graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT

(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S.A. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

---

## Trademark information

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Mars, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP,

ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

---

## Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe on any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, N.Y. 10504-1785  
U.S.A.

For additional information, visit the web at:  
<http://www.ibm.com/ibm/licensing/contact/>

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

**INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.** Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM web sites are provided for convenience only and do not in any manner serve as an endorsement of those web sites. The materials at those web sites are not part of the materials for this IBM product and use of those web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

# Index

- A**
- access
    - configuring for web services [177](#)
    - customizing an access-control role to restrict user to specified commands [149](#)
    - enabling cluster, for Active Directory domain users [141](#)
    - managing for web services [171](#)
    - restricting RLM to only the specified administration hosts [90](#)
    - restricting SP to only the specified administration hosts [72](#)
    - to the cluster, Introduction to managing (cluster administrators only) [138](#)
  - access methods
    - user account [139](#)
  - access problems
    - troubleshooting web service [178](#)
  - access-control roles
    - commands for managing [153](#)
    - considerations for customizing [146](#)
    - customizing to restrict user access to specified commands [149](#)
    - introduction to managing [143](#)
    - managing rule settings for user names and passwords in [151](#)
    - predefined roles for cluster administrators [143](#)
  - accessing
    - cluster with RSH [19](#)
    - cluster with serial port [14](#)
    - cluster with SSH [15](#)
    - cluster with Telnet [18](#)
    - Data ONTAP man pages [30](#)
    - node log, core dump, and MIB files with a web browser [46](#)
    - the cluster, introduction to using the CLI to (cluster administrators only) [14](#)
    - the SP from an administration host [58](#)
  - accounts
    - commands for managing user [142](#)
    - considerations for password rule settings [152](#)
    - for accessing the RLM [81](#)
    - for accessing the SP [57](#)
    - ways to manage user [138](#)
  - Active Directory domains
    - enabling users to access the cluster [141](#)
  - activities
    - introduction to managing audit settings for management [181](#)
  - admin
    - use of administrative privilege levels [26](#)
  - admin SVM
    - definition of [36](#)
  - administration hosts
    - accessing the SP from [58](#)
    - logging in to the RLM from [81](#)
    - restricting RLM access to only the specified [90](#)
    - restricting SP access to only the specified [72](#)
  - administrative privileges
    - use of levels [26](#)
  - administrators
    - differences between cluster and SVM [13](#)
    - predefined roles for cluster [143](#)
  - advanced
    - use of administrative privilege levels [26](#)
  - aggregates
    - improving write performance [273](#)
    - rules governing node root [47](#)
  - alerts
    - customization of system health [228](#)
    - how they trigger AutoSupport messages and events [229](#)
    - responding to [227](#)
    - triggers used in health monitoring [227](#)
    - See also* health monitoring
  - algorithms
    - key exchange and data encryption, introduction to managing SSH security configuration [159](#)
  - asset tags
    - modifying node attributes [43](#)
  - attributes
    - displaying node [42](#)
    - modifying node [43](#)
  - audit settings
    - commands for managing [182](#)
    - introduction to managing [181](#)
  - authentication
    - behavior when methods include both public key and password [140](#)
    - providing mutual, for the cluster or SVMs [166](#)

- ways to manage digital certificates for server or client [162](#)
- authentication methods
  - for user accounts [140](#)
- authentication tunnels
  - using to enable cluster access for Active Directory domain users [141](#)
- automatic timeout
  - commands for managing period of CLI sessions [32](#)
- AutoSupport
  - about [210](#)
  - commands [220](#)
  - configuring [220](#)
  - content of messages sent by email [215](#)
  - daily message content [214](#)
  - daily messages [211](#)
  - defined [210](#)
  - displaying information [220](#)
  - enabled by default
    - having messages sent to your organization [210](#)
  - enabling and disabling [220](#)
  - event-triggered message content [214](#)
  - event-triggered messages [211](#)
  - history [220](#)
  - mail host support for [216](#)
  - manifest [220](#)
  - manifest, content of [222](#)
  - messages, files collected for [222](#)
  - modifying triggers [220](#)
  - performance message content [214](#)
  - performance messages [211](#)
  - requirements for using [216](#)
  - resending messages [220](#)
  - sending messages [220](#)
  - setup [218](#)
  - severity types [216](#)
  - structure of messages sent by email [215](#)
  - supported transport protocols [216](#)
  - troubleshooting
    - HTTP [224](#)
    - HTTPS [224](#)
  - troubleshooting failure to deliver messages over SMTP [225](#)
  - troubleshooting mail host relaying [226](#)
  - troubleshooting messages [223](#)
  - weekly message content [214](#)
  - weekly messages [211](#)
  - when messages are generated [210](#)
  - when messages are sent [211](#)
  - where messages are sent [211](#)

- Autosupport messages and events
  - how alerts trigger them [229](#)
- AutoSupport messages and events
  - triggered by system health alerts [229](#)
- AutoSupport On Demand
  - communication with technical support [213](#)
- available health monitors [229](#)

## B

- boot devices
  - recovering from a corrupted image of a node's [53](#)
- boot environment prompt
  - booting Data ONTAP from [50](#)
- boot menus
  - managing a node with the [51](#)
- booting
  - Data ONTAP at the boot environment prompt [50](#)

## C

- CA-signed digital certificates
  - generating and installing for server authentication [164](#)
- cache rewarming
  - abort events [271](#)
  - about [270](#)
  - disabling [272](#)
  - enabling [272](#)
  - how it works [270](#)
  - trigger events [271](#)
- caches
  - comparison of Flash Pool and Flash Cache [266](#)
- certificates
  - commands for managing digital [170](#)
  - generating and installing CA-signed digital for server authentication [164](#)
  - installing intermediate [166](#)
  - installing SSL server root [169](#)
  - ways to manage digital, for server or client authentication [162](#)
- ciphers
  - key exchange and data encryption, introduction to managing SSH security configurations [159](#)
- CLI
  - keyboard shortcuts [24](#)
  - methods of navigating command directories [22](#)
  - overview of using Data ONTAP [21](#)
  - rules for specifying values [23](#)
  - sessions, automatic timeout [32](#)

- sessions, records of [31](#)
- setting display preferences in [27](#)
- setting privilege levels [26](#)
- CLI commands
  - for Storage QoS [263](#)
  - introduction to shells for executing [21](#)
- CLI sessions
  - commands for managing automatic timeout period of [32](#)
  - commands for managing records of [31](#)
  - introduction to managing (cluster administrators only) [31](#)
  - recording [31](#)
- CLIs
  - introduction to accessing the cluster by using, cluster administrators only [14](#)
- cluster
  - attributes, displaying [39](#)
  - attributes, modifying [39](#)
  - description of [34](#)
  - management server [36](#)
  - nodes of [38](#)
  - quorum [36](#)
  - replication ring, description of [37](#)
  - replication ring, displaying status [41](#)
- cluster access
  - enabling Telnet or RSH for [17](#)
  - introduction to using the CLI for (cluster administrators only) [14](#)
- cluster administrators
  - introduction to managing access to the cluster [138](#)
  - predefined roles for [143](#)
- cluster configurations
  - automatic backups for [193](#)
  - backing up and restoring [193](#)
  - backup files for [193](#)
  - choosing a configuration for recovering [198](#)
  - commands for managing backup files for [195](#)
  - commands for managing backup schedules for [194](#)
  - managing backups for [193](#)
  - recovering [198](#)
  - restoring using a configuration backup file [199](#)
- cluster peer
  - deleting the relationship [292](#)
- cluster peers
  - creating relationships between [290](#)
  - definition of [279](#)
  - intercluster network requirements for [279](#)
- cluster switch health monitor
  - commands for [235](#)
  - troubleshooting [233](#)
  - verifying switch monitoring [234](#)
  - what it is [229](#)
- cluster time
  - commands for managing [184](#)
  - managing [183](#)
- clusters
  - accessing with RSH [19](#)
  - accessing with SSH [15](#)
  - accessing with Telnet [18](#)
  - adding nodes to [44](#)
  - administrators, definition [13](#)
  - automatic configuration backups for [193](#)
  - backing up and restoring configurations for [193](#)
  - commands for managing configuration backup files for [195](#)
  - commands for managing configuration backup schedules for [194](#)
  - configuration backup files for [193](#)
  - enabling Active Directory domain users to access [141](#)
  - introduction to accessing by using the CLI (cluster administrators only) [14](#)
  - introduction to management basics [34](#)
  - introduction to managing access, cluster administrators only [138](#)
  - introduction to managing CLI sessions (cluster administrators only) [31](#)
  - managing configuration backups for [193](#)
  - providing mutual authentication for [166](#)
  - reassigning epsilon to another node [40](#)
  - recovering configurations for [198](#)
  - removing nodes from [45](#)
  - single node, considerations for [35](#)
  - synchronizing nodes with [201](#)
  - using serial port to access [14](#)
- clustershell
  - introduction to [21](#)
- command directories
  - methods of navigating CLI [22](#)
- command-line interface
  - See CLI
- commands
  - customizing an access-control role to restrict user access to specified [149](#)
  - displaying available for nodeshell [21](#)
  - for displaying environmental sensor information [248](#)
  - for managing a node at the SP admin privilege level [62](#)

- for managing a node at the SP advanced privilege level [65](#)
- for managing access-control roles [153](#)
- for managing audit settings [182](#)
- for managing digital certificates [170](#)
- for managing events [209](#)
- for managing job schedules [191](#)
- for managing jobs [189](#)
- for managing licenses [188](#)
- for managing mount points on the nodes [175](#)
- for managing public keys [162](#)
- for managing records of CLI sessions [31](#)
- for managing SSH security configurations [161](#)
- for managing SSL [176](#)
- for managing the automatic timeout period of CLI sessions [32](#)
- for managing the cluster time [184](#)
- for managing the RLM with Data ONTAP [87](#)
- for managing the SP with Data ONTAP [74](#)
- for managing the web protocol engine [173](#)
- for managing user accounts [142](#)
- introduction to shells for executing CLI [21](#)
- methods of customizing show output by using fields [29](#)
- methods of viewing history and reissuing [24](#)
- system services firewall [158](#)
- configuration backup file
  - finding for recovering node configurations [196](#)
- configuration backup files
  - commands for managing [195](#)
  - definition of [193](#)
  - finding for recovering cluster configurations [198](#)
  - using to restore a cluster configuration [199](#)
  - using to restore node configurations [197](#)
- configuration backup schedules
  - about [193](#)
  - commands for managing [194](#)
- configuration files
  - rules governing node root volumes and root aggregates [47](#)
- configuring intercluster LIFs to use dedicated intercluster [285](#)
- connection problems
  - troubleshooting RLM [93](#)
- console sessions
  - relations among RLM CLI, RLM console, and serial [83](#)
  - relations among SP CLI, SP console, and serial [60](#)
- consoles
  - accessing the serial console from the RLM [82](#)

- accessing the serial console from the SP [60](#)
- accessing the SP from serial [59](#)
- core dump files
  - managing [202](#)
  - methods of segmenting [202](#)
  - node, accessing with a web browser [46](#)
- core dumps
  - commands for managing [203](#)
- core segments
  - commands for managing [204](#)
- counters
  - what they are [241](#)
- cross cluster relationship
  - deleting [292](#)

## D

- daily AutoSupport messages
  - type of data contained within [214](#)
- dashboards
  - about [238](#)
  - commands for managing [240](#)
  - getting notified of alarms [239](#)
  - performing My AutoSupport tasks [222](#)
- data
  - commands for viewing [246](#)
- data encryption algorithms
  - introduction to managing SSH security configurations [159](#)
- Data ONTAP
  - accessing man pages [30](#)
  - booting at the boot environment prompt [50](#)
  - introduction to management interface [14](#)
  - introduction to managing the RLM with [87](#)
  - introduction to managing the SP with [71](#)
  - overview of using the CLI [21](#)
- Data ONTAP commands
  - for managing the RLM [87](#)
  - for managing the SP [74](#)
- data ports
  - configuring intercluster LIFs to share [281](#)
  - considerations when sharing intercluster and [280](#)
- diagnostic
  - use of administrative privilege levels [26](#)
- diagnostic accounts
  - uses of [22](#)
- digital certificates
  - commands for managing [170](#)
  - generating and installing CA-signed for server authentication [164](#)

- installing intermediate [166](#)
- ways to manage for server or client authentication [162](#)
- directories
  - methods of navigating CLI command [22](#)
- discrete SP sensors
  - understanding the status of [67](#)
- display preferences
  - setting in CLI [27](#)
- DNS [224](#)
- domains
  - enabling Active Directory users to access the cluster [141](#)

## E

- eligibility
  - cluster participation, modifying node attributes [43](#)
- email messages
  - structure of for AutoSupport [215](#)
- EMS
  - callhome events
    - EMS [212](#)
    - getting notified of dashboard alarms [239](#)
    - getting notified of system health alerts [230](#)
    - managing event messages [206](#)
    - relationship with event-triggered AutoSupport messages [212](#)
    - unknown user event [226](#)
- encryption algorithms
  - data, introduction to managing SSH security configurations [159](#)
- engines
  - ways to manage the web protocol [172](#)
- environmental sensors
  - displaying information about [248](#)
- epsilon
  - reassigning to another node [40](#)
  - understanding [36](#)
- Event Management System
  - commands for managing events [209](#)
  - setting up [207](#)
  - See also* EMS
- Event Management Systems
  - See* EMS
- event messages
  - managing [206](#)
  - reducing number of [207](#)
- event-triggered AutoSupport messages
  - how AutoSupport creates and sends them [212](#)

- events
  - commands for managing [209](#)
- extended queries
  - methods of using [29](#)

## F

- fields
  - methods of customizing show command output by using [29](#)
- files
  - controlling I/O performance [257](#)
  - methods of segmenting core dump [202](#)
  - rules for assigning to Storage QoS policy groups [254](#)
- firewall policy
  - creating [156](#)
  - putting into effect [156](#)
- firewall service and policies
  - commands for managing [158](#)
- firmware updates
  - methods of managing SP [71](#)
- Flash Cache
  - compared with Flash Pool aggregates [266](#)
- Flash Cache family of modules [265](#)
- Flash Pool aggregates
  - compared with Flash Cache [266](#)
- flexscale.rewarm option [272](#)
- FlexVol volumes
  - controlling I/O performance [257](#)
  - promoting to root for SVMs [132](#)
  - rules for assigning to Storage QoS policy groups [254](#)
  - with SVMs, explained [94](#)
- free space reallocation
  - disabling [275](#)
  - enabling [275](#)
  - how it works [273](#)
  - overview [273](#)
  - supported aggregates [275](#)
  - using with other reallocation features [275](#)
  - viewing status [275](#)
  - when to enable [274](#)

## G

- get requests
  - introduction to managing audit settings for [181](#)

## H

- health monitoring

- commands [235](#)
- commands for managing dashboards [240](#)
- example of responding to degraded health [232](#)
- getting notified of alerts [230](#)
- how alerts trigger AutoSupport messages and events [229](#)
- responding to degraded health [231](#)
- ways to control when alerts occur [228](#)
- ways to respond to alerts [227](#)
- what it is [227](#)
- health monitors
  - about alerts [227](#)
  - availability of [229](#)
  - how they work and about health policies [227](#)
- health policies
  - about alerts and matching conditions [227](#)
- history of commands
  - methods of viewing [24](#)
- hosts
  - accessing the SP from administration [58](#)
  - logging in to the RLM from administration [81](#)
  - restricting RLM access to only the specified administration [90](#)
  - restricting SP access to only the specified administration [72](#)
- how alerts trigger AutoSupport messages and events
  - health monitoring [229](#)

## I

- images
  - recovering from the corruption of a node's boot device [53](#)
- increasing cache memory [265](#)
- Infinite Volumes
  - with SVMs, explained [94](#)
- instances
  - what they are [241](#)
- intercluster LIFs
  - configuring to share data ports [281](#)
  - configuring to use dedicated intercluster ports [285](#)
  - considerations when sharing with data ports [280](#)
- intercluster networking
  - definition of cluster peer [279](#)
- intercluster networks
  - configuring intercluster LIFs for [281](#), [285](#)
  - considerations when sharing data and intercluster ports [280](#)
  - requirements for cluster peer [279](#)
- intercluster ports

- configuring intercluster LIFs to use dedicated [285](#)
- considerations when using dedicated [281](#)
- interfaces
  - introduction to Data ONTAP management [14](#)
  - introduction to managing activity audit settings for management [181](#)
  - overview of using Data ONTAP command line [21](#)

## J

- job categories
  - about [189](#)
- job schedules
  - commands for managing [191](#)
- jobs
  - categories of [189](#)
  - commands for managing [189](#)
  - managing schedules for [189](#)
  - viewing information about [189](#)
- joining nodes
  - to the cluster [44](#)

## K

- key exchange algorithms
  - introduction to managing SSH security configurations [159](#)
- keys
  - ways to manage public [161](#)

## L

- levels
  - use of administrative privilege [26](#)
- license
  - types and licensed method [186](#)
- licenses
  - commands for managing [188](#)
  - managing [186](#)
- LIFs
  - configuring to share data ports with intercluster [281](#)
  - configuring to use dedicated intercluster ports [285](#)
- location
  - modifying node attributes [43](#)
- log files
  - node, accessing with a web browser [46](#)
- LUNs
  - controlling I/O performance [257](#)
  - rules for assigning to Storage QoS policy groups [254](#)

- ## M
- mail hosts
    - supported by AutoSupport [216](#)
  - man pages
    - accessing Data ONTAP [30](#)
  - management basics
    - cluster, introduction to [34](#)
  - management interfaces
    - introduction to Data ONTAP [14](#)
    - introduction to managing activity audit settings [181](#)
  - managing
    - events on cluster, commands for [209](#)
    - licenses [186](#)
  - manifest
    - AutoSupport messages within [222](#)
  - messages
    - configuring EMS [207](#)
    - files collected for AutoSupport [222](#)
    - managing event [206](#)
  - MIB files
    - node, accessing with a web browser [46](#)
  - monitoring
    - commands for managing dashboards [240](#)
    - dashboard [238](#)
    - node connectivity [227](#)
    - switches [227](#)
    - system connectivity [227](#)
    - using system health alerts [228](#)
  - mount points
    - commands for managing the node [175](#)
  - mutual authentication
    - providing for the cluster or SVMs [166](#)
  - My AutoSupport
    - dashboard tasks [222](#)
    - described [222](#)
- ## N
- network
    - configuring the SP [56](#)
  - Network Time Protocol
    - See* NTP
  - networks
    - configuring the RLM [79](#)
  - node attributes
    - modifying [43](#)
  - node configurations
    - automatic backups for [193](#)
    - backup files for [193](#)
    - commands for managing backup files for [195](#)
    - commands for managing backup schedules for [194](#)
    - finding configuration backup files for recovering [196](#)
    - managing backups for [193](#)
    - recovering [196](#)
    - restoring using a configuration backup file [197](#)
  - node connectivity health monitor
    - about [229](#)
    - commands for [235](#)
  - node root aggregates
    - rules governing [47](#)
  - node root volumes
    - rules governing [47](#)
  - nodes
    - accessing log, core dump, and MIB files with a web browser [46](#)
    - adding to the cluster [44](#)
    - automatic configuration backups for [193](#)
    - clusters for single [35](#)
    - commands for managing configuration backup files for [195](#)
    - commands for managing configuration backup schedules for [194](#)
    - commands for managing mount points on [175](#)
    - configuration backup files for [193](#)
    - displaying attributes [42](#)
    - freeing up space on the root volume [48](#)
    - introduction to starting and stopping [49](#)
    - management basics [42](#)
    - managing configuration backups for [193](#)
    - managing core dump files after panic [202](#)
    - managing with the boot menu [51](#)
    - reassigning epsilon [40](#)
    - rebooting at the system prompt [49](#)
    - rebooting remotely [50](#)
    - recovering from a corrupted image of the boot device [53](#)
    - recovering the configuration for [196](#)
    - rejoining to a re-created cluster [199](#)
    - remotely managing [54](#)
    - removing from the cluster [45](#)
    - renaming [43](#)
    - shutting down [51](#)
    - single node cluster
      - See* single-node cluster
    - SP admin privilege level commands for managing the [62](#)
    - SP advanced privilege level commands for managing the [65](#)
    - synchronizing with the cluster [201](#)

- using the RLM to remotely manage [77](#)
- using the SP to remotely manage [54](#)
- nodeshell
  - displaying available commands for [21](#)
  - introduction to [21](#)
- NTP
  - commands for managing the cluster time [184](#)
  - managing the cluster time with [183](#)

**O**

- objects
  - what they are [241](#)
- online help
  - for using RLM CLI [84](#)
  - for using SP CLI [61](#)
- operators
  - methods of using query [27](#)
- output
  - methods of customizing show command by using fields [29](#)
- overview [272](#)
- owners
  - modifying node attributes [43](#)

**P**

- PAM (Performance Acceleration Module) [265](#)
- panics
  - managing core dump files after [202](#)
- password rules
  - considerations for settings [152](#)
- passwords
  - authentication behavior when methods include both public key and [140](#)
  - managing rule settings in access-control role [151](#)
- peer relationships
  - creating cluster [290](#)
- performance
  - controlling workload performance [257](#)
  - data
    - decisions before you view [242](#)
    - viewing continuously [245](#)
    - viewing for a time period [243](#)
    - what objects, instances, and counters are [241](#)
  - improving write performance [273](#)
  - monitoring [241](#)
  - read [272](#)
  - read reallocation [272](#)
- Performance Acceleration Module [265](#)

- performance AutoSupport messages
  - type of data contained within [214](#)
- performance improvements, in storage systems
  - WAFL external cache [265](#)
- policy groups
  - creating [257](#)
  - how maximum throughput works [253](#)
  - maximum number of [257](#)
  - monitoring [257](#)
  - types of [251](#)
  - what they are [251](#)
- ports
  - configuring intercluster LIFs to share with data [281](#)
  - considerations when sharing data and intercluster roles on [280](#)
  - considerations when using dedicated intercluster [281](#)
- predefined roles
  - for cluster administrators [143](#)
- preferences
  - setting display in CLI [27](#)
- privilege levels
  - setting in CLI [26](#)
  - use of administrative [26](#)
- prompts
  - booting Data ONTAP at the boot environment [50](#)
  - overview of Data ONTAP command [21](#)
  - rebooting a node at the system [49](#)
- protocol engines
  - commands for managing the web [173](#)
  - ways to manage the web [172](#)
- public keys
  - authentication behavior when methods include both password and [140](#)
  - commands for managing [162](#)
  - ways to manage [161](#)

## Q

- Quality of Service
  - See* Storage QoS
- queries
  - methods of using extended [29](#)
- query operators
  - methods of using [27](#)
- quorum
  - understanding [36](#)

## R

- read reallocation

- disabling [273](#)
    - enabling [273](#)
    - viewing status [273](#)
  - reallocation
    - free space [273](#)
    - read [272](#)
    - when to use with free space reallocation [275](#)
  - rebooting
    - a node at the system prompt [49](#)
  - records
    - commands for managing CLI session [31](#)
  - recovering
    - cluster configurations [198](#)
    - from a corrupted image of a node's boot device [53](#)
    - node configurations [196](#)
  - reissuing commands
    - methods of [24](#)
  - rejoining
    - nodes to a cluster [199](#)
  - relationships
    - creating cluster peer [290](#)
  - remote
    - node management by using the RLM [77](#)
    - node management by using the SP [54](#)
  - Remote LAN Modules
    - See* RLM
  - remote management
    - of a node [54](#)
  - requirements
    - for using AutoSupport [216](#)
  - requirements for cluster peer intercluster [279](#)
  - respond to
    - alerts [227](#)
    - system health alerts [227](#)
  - restoring
    - cluster configurations [199](#)
    - node configurations [197](#)
  - RLM
    - down filer events [92](#)
    - down system events [92](#)
    - SNMP traps [92](#)
  - RLMs
    - accessing the serial console from [82](#)
    - accounts that can access [81](#)
    - commands for managing a node at the admin privilege level [85](#)
    - commands for managing a node at the advanced privilege level [86](#)
    - commands for troubleshooting a node [92](#)
    - configuring automatic logout of idle SSH connections to [91](#)
    - configuring the network [79](#)
    - Data ONTAP commands for managing [87](#)
    - introduction to [78](#)
    - introduction to managing, with Data ONTAP [87](#)
    - logging in to, from an administration host [81](#)
    - managing a node remotely by using [77](#)
    - relations among RLM CLI, RLM console, and serial console sessions [83](#)
    - remote management
      - RLMs [78](#)
    - restricting access to only the specified administration hosts [90](#)
    - troubleshooting connection problems [93](#)
    - using online help at CLI [84](#)
  - roles
    - commands for managing access control [153](#)
    - considerations for customizing access-control [146](#)
    - considerations for password rule settings [152](#)
    - customizing to restrict user access to specified commands for access control [149](#)
    - introduction to managing access-control [143](#)
    - managing rule settings for user names and passwords in access control [151](#)
    - predefined for cluster administrators [143](#)
  - root aggregates
    - rules governing node [47](#)
  - root certificates
    - installing SSL server [169](#)
  - root volumes
    - freeing up space on node [48](#)
    - promoting a FlexVol volume for SVMs [132](#)
    - rules governing node [47](#)
  - RSH
    - accessing cluster with [19](#)
    - enabling access to the cluster [17](#)
  - rule settings
    - managing for user names and passwords [151](#)
  - rules
    - considerations for password settings [152](#)
  - rules for assigning storage objects to [254](#)
- ## S
- schedules
    - commands for managing job [191](#)
    - managing jobs and [189](#)
  - Secure Sockets Layer
    - See* SSL

- security
  - managing user names and passwords in access-control role [151](#)
- security configurations
  - commands for managing SSH [161](#)
  - introduction to managing SSH [159](#)
- sensors
  - displaying information about environmental [248](#)
  - understanding the status of discrete SP [67](#)
- serial console sessions
  - relations among RLM CLI, RLM console sessions, and [83](#)
  - relations among SP CLI, SP console sessions, and [60](#)
- serial consoles
  - accessing from the RLM [82](#)
  - accessing from the SP [60](#)
  - accessing the SP from [59](#)
- serial ports
  - using to access cluster [14](#)
- server authentication
  - generating and installing a CA-signed digital certificate for [164](#)
- servers
  - installing root certificate of SSL [169](#)
- Service Processors
  - See* SP
- services
  - commands for managing web [175](#)
  - configuring access to web [177](#)
  - requirements for user access to web [174](#)
  - ways to manage web [174](#)
- sessions
  - introduction to managing CLI, (cluster administrators only) [31](#)
  - recording CLI [31](#)
- set requests
  - introduction to managing audit settings for [181](#)
- settings
  - introduction to managing audit [181](#)
  - managing rule, for user names and passwords [151](#)
- setup
  - AutoSupport [218](#)
- severity
  - AutoSupport [216](#)
- shells
  - introduction to CLI command [21](#)
- show command output
  - methods of customizing by using fields [29](#)
- shutting down
  - a node [51](#)
- single node clusters
  - considerations for [35](#)
- single-node cluster
  - description of [34](#)
- SMTP
  - troubleshooting AutoSupport failure to deliver messages over [225](#)
- SP
  - sensors, threshold-based [65](#)
  - SNMP traps [77](#)
- SP sensors
  - understanding the status of discrete [67](#)
- space
  - freeing up on a node's root volume [48](#)
- SPs
  - accessing from an administration host [58](#)
  - accessing from the serial console [59](#)
  - accessing the serial console from [60](#)
  - accounts that can access [57](#)
  - commands for managing a node at the admin privilege level [62](#)
  - commands for managing a node at the advanced privilege level [65](#)
  - commands for troubleshooting a node [70](#)
  - configuring automatic logout of idle SSH connections to [73](#)
  - configuring the network [56](#)
  - Data ONTAP commands for managing [74](#)
  - introduction to managing with Data ONTAP [71](#)
  - managing a node remotely by using [54](#)
  - methods of managing firmware updates [71](#)
  - relations among SP CLI, SP console, and serial console sessions [60](#)
  - restricting access to only the specified administration hosts [72](#)
  - using online help at CLI [61](#)
- SSH
  - accessing cluster with [15](#)
  - commands for managing security configurations [161](#)
  - configuring automatic logout of idle connections to the RLM [91](#)
  - configuring automatic logout of idle connections to the SP [73](#)
  - security configurations, introduction to managing [159](#)
- SSL
  - commands for managing [176](#)
  - managing [176](#)
- SSL servers
  - installing root certificate of [169](#)

- standalone nodes
  - clusters for [35](#)
- statistics
  - See performance
- Storage QoS
  - assigning storage objects to policy groups [257](#)
  - commands [263](#)
  - creating policy groups [257](#)
  - effect on non-throttled workloads [254](#)
  - examples
    - isolating a workload [259](#)
    - setting a limit on all workloads [262](#)
    - setting a proactive limit on non-critical workloads [261](#)
  - how it helps [249](#)
  - how it works [251](#)
  - how maximum throughput works [253](#)
  - how to monitor workload performance [256](#)
  - maximum number of policy groups [257](#)
  - maximum number of workloads [257](#)
  - monitoring policy group performance [257](#)
  - monitoring workload performance [257](#)
  - rules for assigning storage objects to policy groups [254](#)
  - types of policy groups [251](#)
  - types of workloads [251](#)
  - what it is [249](#)
  - workflow [249](#)
- storage systems
  - monitoring the [206](#)
- SVM administrator roles [144](#)
- SVM with FlexVol volumes
  - controlling I/O performance [136](#), [257](#)
  - rules for assigning to Storage QoS policy groups [254](#)
- SVMs
  - about peer relationship [296](#)
  - accepting peer relationship [300](#)
  - admin SVMs [96](#)
  - administering from SVM context [129](#)
  - administrators, definition [13](#)
  - benefits [97](#)
  - complete configuration by using the Vserver Setup wizard [110](#)
  - configuring protocols [115](#)
  - configuring services [114](#)
  - consideration for modifying [120](#)
  - creating peer relationship [298](#)
  - creating, by using vserver create command [116](#), [118](#)
  - data SVMs [96](#)
  - delegating administration [123](#)
  - deleting [128](#)
  - deleting peer relationship [301](#), [304](#)
  - displaying information about [126](#)
  - displaying peer relationship [306](#)
  - information to gather for [103](#)
  - language configurations [102](#)
  - language options [99](#)
  - managing [94](#)
  - managing peer relationships [295](#)
  - modifying [122](#)
  - modifying peer relationship [303](#)
  - node SVMs [96](#)
  - performance [241](#)
  - providing mutual authentication for [166](#)
  - rejecting peer relationship [301](#)
  - renaming [127](#)
  - requirements [103](#)
  - resuming peer relationships [306](#)
  - root volume [96](#)
  - root volume, promoting data-protecting mirrors [134](#)
  - root volume, promoting LS mirrors [133](#)
  - root volume, promoting new FlexVol volume [135](#)
  - starting [130](#)
  - states of peer relationships [297](#)
  - stopping [131](#)
  - suspending peer relationship [305](#)
  - SVMs
    - setup wizard [98](#)
    - vserver create [98](#)
  - types of [96](#)
  - Vserver setup worksheet [103](#)
  - with FlexVol volumes, explained [94](#)
  - with Infinite Volume, explained [94](#)
- SVMs provisioning storage [113](#)
- SVMs setting up network interface [112](#)
- SVMs setup
  - using the Vserver Setup wizard [108](#)
- SVMs with FlexVol volumes
  - creating [116](#)
  - explained [94](#)
  - maximum number of [98](#)
- SVMs with Infinite Volume
  - creating [118](#)
  - explained [94](#)
  - maximum number of [98](#)
- switches
  - monitoring [227](#)
  - troubleshooting discovery of [233](#)
  - verifying the monitoring of using the cluster switch health monitor [234](#)

- synchronizing
    - nodes with the cluster [201](#)
  - system configurations
    - backing up and restoring [193](#)
  - system connectivity health monitor
    - commands for [235](#)
    - what it is [229](#)
  - system health
    - alert customization [228](#)
    - responding to degradation [231](#)
    - See also* health monitoring
  - system health alerts
    - how they trigger AutoSupport messages and events [229](#)
    - responding to [227](#)
  - System Manager
    - about [32](#)
    - supported Data ONTAP versions [32](#)
    - tasks you can perform from [32](#)
  - system panics
    - managing core dump files after [202](#)
  - system prompts
    - rebooting a node at the [49](#)
  - systems
    - monitoring the storage [206](#)
  - systemshell
    - introduction to [21](#)
    - uses of [22](#)
- T**
- technical support
    - how AutoSupport On Demand communicates with [213](#)
  - Telnet
    - accessing cluster with [18](#)
    - enabling access to the cluster [17](#)
  - thresholds
    - displaying information about environmental sensor [248](#)
  - time
    - commands for managing the cluster [184](#)
    - managing the cluster [183](#)
  - transport protocols
    - supported by AutoSupport [216](#)
  - trigger events
    - how AutoSupport responds to [212](#)
  - troubleshooting
    - delivery status of AutoSupport messages [223](#)
    - mail host [226](#)
    - managing core dump files for [202](#)
    - nodes with RLM commands [92](#)
    - nodes with SP commands [70](#)
    - RLM connection problems [93](#)
    - switch discovery for health monitoring [233](#)
    - using systemshell and diagnostic account for [22](#)
    - web service access problems [178](#)
  - two-way authentication
    - See* mutual authentication
- U**
- unjoining nodes
    - from the cluster [45](#)
  - updates
    - methods of managing SP firmware [71](#)
  - user accounts
    - access methods for [139](#)
    - authentication methods [140](#)
    - commands for managing [142](#)
    - considerations for password rule settings [152](#)
    - ways to manage [138](#)
  - user names
    - managing rule settings in access-control role [151](#)
  - users
    - enabling cluster access for Active Directory domain [141](#)
- V**
- values
    - rules for specifying in CLI [23](#)
  - volumes
    - freeing up space on a node's root [48](#)
    - rules governing node root [47](#)
    - SVMs with FlexVol, explained [94](#)
    - SVMs with Infinite, explained [94](#)
  - Vservers
    - See* SVMs
- W**
- WAFL (Write Anywhere File Layout) [265](#)
  - WAFL external cache
    - about [265](#)
    - compared with Flash Pool aggregates [266](#)
    - disabling [267](#)
    - displaying configuration [269](#)
    - displaying usage and access information [269](#)
    - enabling [267](#)

## 331 | System Administration Guide for Cluster Administrators

- low-priority user data blocks [268](#)
- normal user data blocks [267](#)
- rewarming [270](#)
- system metadata cache [268](#)
- ways to respond to alerts
  - health monitoring [227](#)
- web browsers
  - accessing node log, core dump, and MIB files using [46](#)
- web protocol engines
  - commands for managing [173](#)
  - ways to manage [172](#)
- web services
  - commands for managing [175](#)
  - configuring access to [177](#)
  - managing access to [171](#)
  - requirements for user access [174](#)
  - troubleshooting access problems [178](#)
  - ways to manage [174](#)
- weekly AutoSupport messages
  - type of data contained within [214](#)
- workloads
  - controlling performance of [257](#)
  - effect of throttling on non-throttled workloads [254](#)
  - how to monitor performance [256](#)
  - maximum number of [257](#)
  - types of [251](#)
  - what they are [251](#)
- Write Anywhere File Layout (WAFL) [265](#)





NA 210-06389\_A0, Printed in USA

SC27-6285-01

