



ONTAP® 9

System Administration Reference

April 2020 | 215-11148_2020-04_en-us
doccomments@netapp.com

Updated for ONTAP 9.7

 **NetApp**®

Contents

Deciding whether to use this guide	7
Cluster and SVM administrators	8
What the cluster management server is	8
Types of SVMs	8
ONTAP management interface basics	10
Accessing the cluster by using the CLI	10
Accessing the cluster by using the serial port	10
Accessing the cluster by using SSH	10
SSH login security	13
Enabling Telnet or RSH access to the cluster	14
Accessing the cluster by using Telnet	15
Accessing the cluster by using RSH	16
Using the ONTAP command-line interface	17
Understanding the different shells for CLI commands	17
Methods of navigating CLI command directories	19
Rules for specifying values in the CLI	19
Methods of viewing command history and reissuing commands	20
Keyboard shortcuts for editing CLI commands	20
Use of administrative privilege levels	22
Setting the privilege level in the CLI	22
Setting display preferences in the CLI	22
Methods of using query operators	23
Methods of using extended queries	24
Methods of customizing show command output by using fields	25
Understanding positional parameters	25
Methods of accessing ONTAP man pages	27
Managing CLI sessions	27
Managing records of CLI sessions	27
Managing the automatic timeout period of CLI sessions	28
Accessing a cluster by using the ONTAP System Manager browser-based graphic interface	28
Understanding ONTAP System Manager	29
Ways to manage access to ONTAP System Manager	30
Cluster management basics (cluster administrators only)	31
Displaying information about the nodes in a cluster	31
Displaying cluster attributes	32
Modifying cluster attributes	32
Displaying the status of cluster replication rings	33
Understanding quorum and epsilon	33
What system volumes are	34
Managing nodes	35

Displaying node attributes	35
Modifying node attributes	35
Renaming a node	36
Adding nodes to the cluster	36
Removing nodes from the cluster	38
Accessing a node's log, core dump, and MIB files by using a web browser	40
Accessing the system console of a node	41
Rules governing node root volumes and root aggregates	43
Freeing up space on a node's root volume	43
Relocating root volumes to new aggregates	44
Starting or stopping a node	45
Rebooting a node at the system prompt	45
Booting ONTAP at the boot environment prompt	46
Shutting down a node	47
Managing a node by using the boot menu	48
Managing a node remotely using the SP/BMC	49
Understanding the SP	49
What the Baseboard Management Controller does	51
Configuring the SP/BMC network	52
Methods of managing SP/BMC firmware updates	57
Accessing the SP/BMC	59
Using online help at the SP/BMC CLI	63
Commands for managing a node remotely	65
Commands for managing the SP from ONTAP	72
ONTAP commands for BMC management	76
BMC CLI commands	77
Managing audit logging for management activities	78
Changes to audit logging in ONTAP 9	78
How ONTAP implements audit logging	78
Forwarding the audit log to a destination	79
Commands for managing audit settings for management activities	80
Managing the cluster time	81
Commands for managing the cluster time	82
Managing the banner and MOTD	84
Creating a banner	85
Managing the banner	86
Creating an MOTD	87
Managing the MOTD	88
Managing licenses	91
License types and licensed method	92
Commands for managing licenses	93
Managing jobs and schedules	95
Job categories	95
Commands for managing jobs	95
Commands for managing job schedules	96

Backing up and restoring cluster configurations	98
What configuration backup files are	98
Managing configuration backups	98
How the node and cluster configurations are backed up automatically	98
Commands for managing configuration backup schedules	99
Commands for managing configuration backup files	99
Recovering a node configuration	100
Finding a configuration backup file to use for recovering a node	100
Restoring the node configuration using a configuration backup file	101
Recovering a cluster configuration	102
Finding a configuration to use for recovering a cluster	102
Restoring a cluster configuration from an existing configuration	103
Synchronizing a node with the cluster	105
Managing core dumps	107
Commands for managing core dumps	107
Monitoring a storage system	109
Managing AutoSupport	109
When and where AutoSupport messages are sent	109
How AutoSupport creates and sends event-triggered messages	111
Types of AutoSupport messages and their content	112
What AutoSupport subsystems are	113
AutoSupport size and time budgets	114
Files sent in event-triggered AutoSupport messages	114
Log files sent in AutoSupport messages	114
Files sent in weekly AutoSupport messages	115
How AutoSupport obtains delivery instructions from technical support	115
Structure of AutoSupport messages sent by email	116
AutoSupport severity types	117
Requirements for using AutoSupport	117
Setting up AutoSupport	119
Uploading core dump files	122
Uploading performance archive files	123
Getting AutoSupport message descriptions	124
Commands for managing AutoSupport	124
Information included in the AutoSupport manifest	126
AutoSupport case suppression during scheduled maintenance windows ...	126
What Active IQ is	126
Troubleshooting AutoSupport	127
Monitoring the health of your system	132
How health monitoring works	132
Available cluster health monitors	134
Receiving system health alerts automatically	134
Responding to degraded system health	135
Configuring discovery of cluster and management network switches	137
Verifying the monitoring of cluster and management network switches	138

Commands for monitoring the health of your system	138
Displaying environmental information	141
Managing access to web services	142
Managing the web protocol engine	142
Commands for managing the web protocol engine	144
Configuring SAML authentication for web services	144
Configuring SAML authentication	145
Disabling SAML authentication	146
Troubleshooting issues with SAML configuration	147
Managing web services	148
Commands for managing web services	149
Commands for managing mount points on the nodes	149
Managing SSL	150
Commands for managing SSL	151
Configuring access to web services	151
Troubleshooting web service access problems	152
Verifying the identity of remote servers using certificates	155
Verifying digital certificates are valid using OCSP	155
Viewing default certificates for TLS-based applications	157
Mutually authenticating the cluster and a KMIP server	157
Generating a certificate signing request for the cluster	158
Installing a CA-signed server certificate for the cluster	159
Installing a CA-signed client certificate for the KMIP server	159
Logical space reporting and enforcement for volumes	161
What logical space reporting shows	161
Logical space reporting limitations	161
What logical space enforcement does	162
Controlling how space available and space consumed is displayed to end users ...	162
Copyright	163
Trademark	164
How to send comments about documentation and receive update notifications	165

Deciding whether to use the System Administration Reference

This guide describes how to administer ONTAP systems. It shows you how to use the CLI interface, how to access the cluster, how to manage nodes, and much more.

You should use this guide under the following circumstances:

- You want to understand the range of ONTAP administrator capabilities.
- You want to use the command-line interface (CLI), not ONTAP System Manager or an automated scripting tool.

If you want to use ONTAP System Manager to administer your system, you should choose the following documentation:

- [*Cluster management using System Manager*](#)

If you require additional configuration or conceptual information, you should choose among the following documentation:

- Conceptual background for clusters and SVMs
[*ONTAP concepts*](#)
- Administrator authentication and authorization
[*Administrator authentication and RBAC*](#)
- Network and LIF management
[*Network and LIF management*](#)
- Command reference
[*ONTAP 9 commands*](#)
- Automation of management tasks
[*NetApp Documentation: OnCommand Workflow Automation \(current releases\)*](#)

Cluster and SVM administrators

Cluster administrators administer the entire cluster and the storage virtual machines (SVMs, formerly known as Vservers) it contains. SVM administrators administer only their own data SVMs.

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

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

Note: The 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 the cluster management server is

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

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

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

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

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

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

Types of SVMs

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

A cluster contains the following types of SVMs:

- Admin SVM
 - The cluster setup process automatically creates the admin SVM for the cluster. The admin SVM represents the cluster.
- Node SVM

A node SVM is created when the node joins the cluster, and the node SVM represents the individual nodes of the cluster.

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

Note: Unless otherwise specified, the term SVM refers to a data (data-serving) SVM.

In the CLI, SVMs are displayed as Vservers.

ONTAP management interface basics

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

Related concepts

[Understanding the different shells for CLI commands \(cluster administrators only\)](#) on page 17

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

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

Accessing the cluster by using the serial port

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

Steps

1. At the console, press Enter.

The system responds with the login prompt.

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

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

The system responds with the password prompt.

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

Accessing the cluster by using SSH

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

Before you begin

- You must have a user account that is configured to use **ssh** as an access method. The `-application` parameter of the `security login` commands specifies the access method for a user account. The `security login` man pages contain additional information.
- 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

- Only the SSH v2 protocol is supported; SSH v1 is not supported.
- ONTAP supports a maximum of 64 concurrent SSH sessions per node.
If the cluster management LIF resides on the node, it shares this limit with the node management LIF.
If the rate of incoming connections is higher than 10 per second, the service is temporarily disabled for 60 seconds.
- ONTAP supports only the AES and 3DES encryption algorithms (also known as *ciphers*) for SSH.
AES is supported with 128, 192, and 256 bits in key length. 3DES is 56 bits in key length as in the original DES, but it is repeated three times.
- When FIPS mode is on, SSH clients should negotiate with Elliptic Curve Digital Signature Algorithm (ECDSA) host key algorithms for the connection to be successful.
- If you want to access the ONTAP CLI from a Windows host, you can use a third-party utility such as PuTTY.
- If you use a Windows AD user name to log in to ONTAP, you should use the same uppercase or lowercase letters that were used when the AD user name and domain name were created in ONTAP.
AD user names and domain names are not case-sensitive. However, ONTAP user names are case-sensitive. Case mismatch between the user name created in ONTAP and the user name created in AD results in a login failure.
- Starting with ONTAP 9.3, you can enable SSH multifactor authentication for local administrator accounts.
When SSH multifactor authentication is enabled, users are authenticated by using a public key and a password.
- Starting with ONTAP 9.4, you can enable SSH multifactor authentication for LDAP and NIS remote users.

Step

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

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

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

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

command is not required for SSH-interactive sessions.

Examples of SSH requests

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

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

cluster1::>
```

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

$
```

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

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

cluster1::>
```

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

$
```

The following example shows how the user account named “joe” can issue an SSH MFA request to access a cluster whose cluster management LIF is 10.72.137.32:

```
$ ssh joe@10.72.137.32
Authenticated with partial success.
Password:
cluster1::> cluster show
Node                Health  Eligibility
-----
node1                true   true
```

```
node2                true    true
2 entries were displayed.

cluster1::>
```

Related information

[Administrator authentication and RBAC](#)

SSH login security

Starting with ONTAP 9.5, you can view information about previous logins, unsuccessful attempts to log in, and changes to your privileges since your last successful login.

Security-related information is displayed when you successfully log in as an SSH admin user. You are alerted about the following conditions:

- The last time your account name was logged in.
- The number of unsuccessful login attempts since the last successful login.
- Whether the role has changed since the last login (for example, if the admin account's role changed from "admin" to "backup.")
- Whether the add, modify, or delete capabilities of the role were modified since the last login.

Note: If any of the information displayed is suspicious, you should immediately contact your security department.

To obtain this information when you login, the following prerequisites must be met:

- Your SSH user account must be provisioned in ONTAP.
- Your SSH security login must be created.
- Your login attempt must be successful.

Restrictions and other considerations for SSH login security

The following restrictions and considerations apply to SSH login security information:

- The information is available only for SSH-based logins.
- For group-based admin accounts, such as LDAP/NIS and AD accounts, users can view the SSH login information if the group of which they are a member is provisioned as an admin account in ONTAP.
However, alerts about changes to the role of the user account cannot be displayed for these users. Also, users belonging to an AD group that has been provisioned as an admin account in ONTAP cannot view the count of unsuccessful login attempts that occurred since the last time they logged in.
- The information maintained for a user is deleted when the user account is deleted from ONTAP.
- The information is not displayed for connections to applications other than SSH.

Examples of SSH login security information

The following examples demonstrate the type of information displayed after you login.

- This message is displayed after each successful login:

```
Last Login : 7/19/2018 06:11:32
```

- These messages are displayed if there have been unsuccessful attempts to login since the last successful login:

```
Last Login : 4/12/2018 08:21:26
Unsuccessful login attempts since last login - 5
```

- These messages are displayed if there have been unsuccessful attempts to login and your privileges were modified since the last successful login:

```
Last Login : 8/22/2018 20:08:21
Unsuccessful login attempts since last login - 3
Your privileges have changed since last login
```

Enabling Telnet or RSH access to the cluster

As a security best practice, 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

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

Perform the following steps to enable Telnet or RSH access to the clusters:

Steps

1. Enter the advanced privilege mode:


```
set advanced
```
2. Enable a security protocol (RSH or Telnet):


```
security protocol modify -application security_protocol -enabled true
```
3. Create a new management firewall policy based on the `mgmt` management firewall policy:


```
system services firewall policy clone -policy mgmt -new-policy-name policy-name
```
4. Enable Telnet or RSH in the new management firewall policy:


```
system services firewall policy create -policy policy-name -service security_protocol -action allow -ip-list ip_address/port
```

To allow all IP addresses, you should specify

```
-ip-list 0.0.0.0/0
.
```
5. Associate the new policy with the cluster management LIF:

```
network interface modify -vserver cluster_management_LIF -lif
cluster_mgmt -firewall-policy policy-name
```

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.
- 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 ONTAP CLI from a Windows host, you can use a third-party utility such as PuTTY.

Step

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

```
telnet hostname_or_IP
```

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

Example of a Telnet request

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

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

Related information

[Network and LIF management](#)

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.
- ONTAP supports a maximum of 50 concurrent RSH sessions per node. If the cluster management LIF resides on the node, it shares this limit with the node management LIF. If the rate of in-coming connections is higher than 10 per second, the service is temporarily disabled for 60 seconds.

Step

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

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

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

command is the command you want to execute over RSH.

Example of an RSH request

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

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

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

admin_host$
```

Using the ONTAP command-line interface

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

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

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

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

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

- The clustershell is the native shell that is started automatically when you log in to the cluster. It provides all the commands you need to configure and manage the cluster. The clustershell CLI help (triggered by `?` at the clustershell prompt) displays available clustershell commands. The `man command_name` command in the clustershell displays the man page for the specified clustershell command.
- The nodeshell is a special shell for commands that take effect only at the node level. The nodeshell is accessible through the `system node run` command. The nodeshell CLI help (triggered by `?` or `help` at the nodeshell prompt) displays available nodeshell commands. The `man command_name` command in the nodeshell displays the man page for the specified nodeshell command. Many commonly used nodeshell commands and options are tunneled or aliased into the clustershell and can be executed also from the clustershell.
- The systemshell is a low-level shell that is used only for diagnostic and troubleshooting purposes. The systemshell and the associated “diag” account are intended for low-level diagnostic purposes. Their access requires the diagnostic privilege level and is reserved only for technical support to perform troubleshooting tasks.

Access of nodeshell commands and options in the clustershell

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

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

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

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

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

Related tasks

[Displaying available nodeshell commands](#) on page 18

Displaying available nodeshell commands

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

Steps

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

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

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

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

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

```
[commandname] help
```

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

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

Example of displaying available nodeshell commands

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

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

node2> environment help
Usage: environment status |
      [status] [shelf [<adapter>[.<shelf-number>]]] |
      [status] [shelf_log] |
      [status] [shelf_stats] |
      [status] [shelf_power_status] |
      [status] [chassis [all | list-sensors | Temperature | PSU 1 |
      PSU 2 | Voltage | SYS FAN | NVRAM6-temperature-3 | NVRAM6-
      battery-3]]
```

Related concepts

[Access of nodeshell commands and options in the clustershell](#) on page 18

Methods of navigating CLI command directories

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

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

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

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

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

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

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

Rules for specifying values in the CLI

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

- A value can be a number, a Boolean specifier, a selection from an enumerated list of predefined values, or a text string.
Some parameters can accept a comma-separated list of two or more values. Comma-separated lists of values do not need to be in quotation marks (" "). Whenever you specify text, a space, or a query character (when not meant as a query or text starting with a less-than or greater-than symbol), you must enclose the entity in quotation marks.
- The CLI interprets a question mark ("?") as the command to display help information for a particular command.
- Some text that you enter in the CLI, such as command names, parameters, and certain values, is not case-sensitive.
For example, when you enter parameter values for the `vserver cifs` commands, capitalization is ignored. However, most parameter values, such as the names of nodes, storage virtual machines (SVMs), aggregates, volumes, and logical interfaces, are case-sensitive.
- If you want to clear the value of a parameter that takes a string or a list, you specify an empty set of quotation marks ("") or a dash ("-").

- The hash sign (“#”), also known as the pound sign, indicates a comment for a command-line input; if used, it should appear after the last parameter in a command line.
The CLI ignores the text between “#” and the end of the line.

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

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

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

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

Methods of viewing command history and reissuing commands

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

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

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

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

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

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

Keyboard shortcuts for editing CLI commands

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

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

If you want to...	Use the following keyboard shortcut...
Move the cursor back by one character	Ctrl-B
	Back arrow

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

Use of administrative privilege levels

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

admin

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

advanced

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

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

diagnostic

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

Setting the privilege level in the CLI

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

Step

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

Example of setting the privilege level

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

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

Related references

[Use of administrative privilege levels](#) on page 22

Setting display preferences in the CLI

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

About this task

You can set the following CLI display preferences:

- The privilege level of the command session
- Whether confirmations are issued for potentially disruptive commands
- Whether `show` commands display all fields
- The character or characters to use as the field separator
- The default unit when reporting data sizes

- The number of rows the screen displays in the current CLI session before the interface pauses output
If the preferred number of rows is not specified, it is automatically adjusted based on the actual height of the terminal. If the actual height is undefined, the default number of rows is 24.
- The default storage virtual machine (SVM) or node
- Whether a continuing command should stop if it encounters an error

Step

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

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

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

Example of setting display preferences in the CLI

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

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

Methods of using query operators

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

The following table describes the supported query operators:

Operator	Description
*	Wildcard that matches all entries. For example, the command <code>volume show -volume *tmp*</code> displays a list of all volumes whose names include the string <code>tmp</code> .
!	NOT operator. Indicates a value that is not to be matched; for example, <code>!vs0</code> indicates not to match the value <code>vs0</code> .
	OR operator. Separates two values that are to be compared; for example, <code>vs0 vs2</code> matches either <code>vs0</code> or <code>vs2</code> . You can specify multiple OR statements; for example, <code>a b* *c*</code> matches the entry <code>a</code> , any entry that starts with <code>b</code> , and any entry that includes <code>c</code> .
..	Range operator. For example, <code>5..10</code> matches any value from 5 to 10, inclusive.
<	Less-than operator. For example, <code><20</code> matches any value that is less than 20.
>	Greater-than operator. For example, <code>>5</code> matches any value that is greater than 5.

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

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

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

Methods of using extended queries

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

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

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

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

The combination of queries and modify operations is a useful tool. However, it can potentially cause confusion and errors if implemented incorrectly. For example, using the (advanced privilege) `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 as given in the following example::

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


Methods of customizing show command output by using fields

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

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

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

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

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

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

Understanding positional parameters

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

What a positional parameter is

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

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

How to identify a positional parameter

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

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

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

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

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

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

Examples of using positional parameters

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

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

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

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

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

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

Methods of accessing ONTAP man pages

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

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

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

Related information

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

Managing CLI sessions (cluster administrators only)

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

Managing records of CLI sessions

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

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

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

Recording a CLI session

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

Steps

1. To start recording the current CLI session into a file, use the `system script start` command.
For more information about using the `system script start` command, see the man page.
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.
ONTAP stops recording your CLI session.

Commands for managing records of CLI sessions

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

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

Related information

[ONTAP 9 commands](#)

Managing the automatic timeout period of CLI sessions

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

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

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

Commands for managing the automatic timeout period of CLI sessions

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

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

Related information

[ONTAP 9 commands](#)

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

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

Before you begin

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

- You must have enabled cookies and site data in the browser.

About this task

You can use a cluster management LIF or node management LIF to access ONTAP System Manager. For uninterrupted access to ONTAP System Manager, you should use a cluster management LIF.

Steps

1. Point the web browser to the IP address of the cluster management LIF:

- If you are using IPv4: `https://cluster-mgmt-LIF`
- If you are using IPv6: `https://[cluster-mgmt-LIF]`

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

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

2. Optional: If you have configured an access banner by using the CLI, then read the message that is displayed in the **Warning** dialog box, and choose the required option to proceed.

This option is not supported on systems on which Security Assertion Markup Language (SAML) authentication is enabled.

- If you do not want to continue, click **Cancel**, and close the browser.
- If you want to continue, click **OK** to navigate to the ONTAP System Manager login page.

3. Log in to ONTAP System Manager by using your cluster administrator credentials.

Related concepts

[Managing access to web services](#) on page 142

Related tasks

[Accessing a node's log, core dump, and MIB files by using a web browser](#) on page 40

Understanding ONTAP System Manager

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

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

Note: The name of System Manager has changed from previous versions. Versions 9.5 and earlier were named OnCommand System Manager. Versions 9.6 and later are now called ONTAP System Manager.

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

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

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

Related information

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

Ways to manage access to ONTAP System Manager

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

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

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

Cluster management basics (cluster administrators only)

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

Displaying information about the nodes in a cluster

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

Step

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

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

Examples of displaying the nodes in a cluster

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

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

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

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

cluster1::*> cluster show -node node1

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

Displaying cluster attributes

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

Step

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

Example of displaying cluster attributes

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

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

Modifying cluster attributes

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

About this task

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

Step

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

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

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

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

Example of renaming a cluster

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

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


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
```


Understanding quorum and epsilon

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

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

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

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



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

Taking a node offline can affect the ability of the cluster to remain in quorum. Therefore, ONTAP issues a warning message if you attempt an operation that will either take the cluster out of quorum or else put it one outage away from a loss of quorum. You can disable the quorum warning messages by using the `cluster quorum-service options modify` command at the advanced privilege level.

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

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

What system volumes are

System volumes are FlexVol volumes that contain special metadata, such as metadata for file services audit logs. These volumes are visible in the cluster so that you can fully account for storage use in your cluster.

System volumes are owned by the cluster management server (also called the admin SVM), and they are created automatically when file services auditing is enabled.

You can view system volumes by using the `volume show` command, but most other volume operations are not permitted. For example, you cannot modify a system volume by using the `volume modify` command.

This example shows four system volumes on the admin SVM, which were automatically created when file services auditing was enabled for a data SVM in the cluster:

```
cluster1::> volume show -vserver cluster1
Vserver   Volume                               Aggregate      State      Type      Size  Available  Used%
-----
cluster1  MDV_aud_1d0131843d4811e296fc123478563412
          aggr0                                online        RW         2GB      1.90GB    5%
cluster1  MDV_aud_8be27f813d7311e296fc123478563412
          root_vs0                              online        RW         2GB      1.90GB    5%
cluster1  MDV_aud_9dc4ad503d7311e296fc123478563412
          aggr1                                 online        RW         2GB      1.90GB    5%
cluster1  MDV_aud_a4b887ac3d7311e296fc123478563412
          aggr2                                 online        RW         2GB      1.90GB    5%
4 entries were displayed.
```

Managing nodes

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

A node is connected to other nodes in the cluster over a cluster network. It is also connected to the disk shelves that provide physical storage for the ONTAP system or to storage arrays that provide array LUNs for ONTAP use. Services and components that are controlled by the node, not by the cluster, can be managed by using the `system node` commands.

A node SVM represents a node in the cluster. If you run the `vserver show` command, the output includes node SVMs in the list.

Displaying node attributes

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

Step

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

Example of displaying information about a node

The following example displays detailed information about node1:

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

Modifying node attributes

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

About this task

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

Attention: You should avoid setting a node's eligibility to `false`, except for situation such as restoring the node configuration or prolonged node maintenance. SAN and NAS data access to the node might be impacted when the node is ineligible.

Step

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

Example of modifying node attributes

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

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

Renaming a node

You can change a node's name as required.

Step

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

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

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

Note: Node name cannot be “all” because “all” is a system reserved name.

Example of renaming a node

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 expand it by adding nodes to it. You add only one node at a time.

Before you begin

- If you are adding nodes to a multiple-node cluster, more than half of the existing nodes in the cluster must be healthy (indicated by `cluster show`).
- If you are adding nodes to a two-node switchless cluster, you must have installed and configured the cluster management and interconnect switches before adding additional nodes. The switchless cluster functionality is supported only in a two-node cluster. When a cluster contains or grows to more than two nodes, cluster HA is not required and is disabled automatically.

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

About this task

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

Steps

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

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

```
Welcome to node setup.

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

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

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

2. Exit the Node Setup wizard:

exit

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

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

4. Start the Cluster Setup wizard:

cluster setup

```
::> cluster setup

Welcome to the cluster setup wizard.

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

You can return to cluster setup at any time by typing "cluster setup".
```

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

```
Use your web browser to complete cluster setup by accessing
https://10.63.11.29
```

```
Otherwise, press Enter to complete cluster setup using the
command line interface:
```

Note: For more information on setting up a cluster using the setup GUI, see the ONTAP System Manager online help.

5. Press Enter to use the CLI to complete this task. When prompted to create a new cluster or join an existing one, enter `join`.

Example

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

6. Follow the prompts to set up the node and join it to the cluster:
 - To accept the default value for a prompt, press Enter.
 - To enter your own value for a prompt, enter the value, and then press Enter.
7. Repeat the preceding steps for each additional node that you want to add.

After you finish

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

Related concepts

[Considerations for the SP/BMC network configuration](#) on page 53

Related information

[High-availability configuration](#)

Removing nodes from the cluster

You can remove unwanted nodes from a cluster, one node at a time. After you remove a node, you must also remove its failover partner. If you are removing a node, then its data becomes inaccessible or erased.

Before you begin

- More than half of the nodes in the cluster must be healthy.
- All of the data on the node that you want to remove must have been evacuated.
- All intercluster LIFs must have been removed.

Note: When you remove intercluster LIFs a warning is displayed that can be ignored.

Note: You must not perform operations such as `cluster unjoin` and `node rename` when the automated upgrade is in progress. If you are planning to perform `cluster unjoin` and `node rename` operations, then before initiating an upgrade, you must remove all of the ONTAP images from the package repository by using the `cluster image package delete` command.

About this task

Beginning with ONTAP 9.3, in the case of a mixed-version cluster, you can unjoin the last low-version node from the cluster using the advance privilege `cluster unjoin -skip-last-low-version-node-check` command.

Attention: You must remove a node from the cluster using the `cluster unjoin` command to make all of the system and user data from the disks that are connected to the node inaccessible to users. After removing a node from the cluster, if you need to join the node back to the same cluster, you should contact technical support for assistance about restoring data.

Steps

1. Change the privilege level to advanced:

```
set -privilege advanced
```

2. If Storage Encryption is enabled on the cluster and the self-encrypting disks (SEDs) are operating you must set the data-id-key to 0x0.

- a. Set the data-id-key to 0x0:

```
storage encryption disk modify -disk [disk_path_name] -data-key-id 0x0
```

- b. Optional: If Storage Encryption is enabled and the SEDs are operating in the Federal Information Processing Standard (FIPS) compliance mode, in addition to setting the data-id-key you must also disable the FIPS compliance mode.

This enables the disks that are attached to the node to be sanitized and repurposed after the node is removed from the cluster.

- Display the SEDs that are currently in the FIPS compliance mode:

```
storage encryption disk show -fips
```

- Disable the FIPS compliance mode:

```
storage encryption disk modify disk [disk_path_name] -fips-key-id 0x0
```

- Verify that the modified SEDs are no longer in the FIPS compliance mode:

```
storage encryption disk show -fips
```

3. If the node you want to remove is the current master node, then enable another node in the cluster to be elected as the master node by changing the master node's cluster eligibility to **false**:

```
cluster modify -eligibility false
```

The master node is the node that holds processes such as “mgmt”, “vldb”, “vifmgr”, “bcomd”, and “crs”. The `cluster ring show advanced` command shows the current master node.

Example

```
cluster::*> cluster modify -node node1 -eligibility false
```

4. From another node in the cluster, unjoin the node from the node you want to remove:

```
cluster unjoin -node nodename
```

If you have a mixed version cluster and you are unjoining the last low version node, use the `cluster unjoin -skip-last-low-version-node-check` command.

The system informs you of the following:

- You must also remove the node’s failover partner from the cluster.

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

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

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

5. 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.

6. If you are repurposing the node, erase the node configuration and initialize all disks:
 - a. During the boot process, press Ctrl-C to display the boot menu when prompted to do so.
 - b. Select the boot menu option **(4) Clean configuration and initialize all disks**.
7. Return to admin privilege level:


```
set -privilege admin
```
8. Repeat the preceding steps to remove the failover partner from the cluster.

After you finish

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

Related tasks

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

Related information

[Network and LIF management](#)

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

[How ONTAP implements audit logging](#) on page 78

[Managing the web protocol engine](#) on page 142

[Managing web services](#) on page 148

[Managing access to web services](#) on page 142

[Managing SSL](#) on page 150

[Managing core dumps \(cluster administrators only\)](#) on page 107

Related information

[Network and LIF management](#)

Accessing the system console of a node

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

About this task

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

Steps

1. Access the system console of a node:

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

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

Examples of accessing the system console

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

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

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

```

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

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

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

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

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

```

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

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

Related concepts

[Relationship among the SP CLI, SP console, and system console sessions](#) on page 62

Related tasks

[Accessing the cluster by using SSH](#) on page 10

[Accessing the SP/BMC from an administration host](#) on page 60

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

Rules governing node root volumes and root aggregates

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

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

- The following rules govern the node's root volume:
 - Unless technical support instructs you to do so, do not modify the configuration or content of the root volume.
 - Do not store user data in the root volume.
Storing user data in the root volume increases the storage giveback time between nodes in an HA pair.
 - Starting with ONTAP 9.0, you can move the root volume to another aggregate.
[Relocating root volumes to new aggregates](#) on page 44
- The root aggregate is dedicated to the node's root volume only.
ONTAP prevents you from creating other volumes in the root aggregate.

Related information

[NetApp Hardware Universe](#)

Freeing up space on a node's root volume

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

Steps

1. Display the node's core dump files and their names by using the `system node coredump show` command.
2. Delete unwanted core dump files from the node by using the `system node coredump delete` command.

3. Access the nodeshell:

```
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 from 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:

```
ls /etc/
```

- b. If any packet trace files (*.trc) are in the node's root volume, delete them individually:

```
rm /etc/log/packet_traces/file_name.trc
```

6. Identify and delete the node's root volume Snapshot copies through the nodeshell:

- a. Identify the root volume name:

```
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:

```
snap list root_vol_name
```

- c. Delete unwanted root volume Snapshot copies:

```
snap delete root_vol_name snapshot_name
```

7. Exit the nodeshell and return to the clustershell:

```
exit
```

Relocating root volumes to new aggregates

The root replacement procedure migrates the current root aggregate to another set of disks without disruption.

About this task

You can change the location of the root volume to a new aggregate in the following scenarios:

- When the root aggregates are not on the disk you prefer
- When you want to rearrange the disks connected to the node
- When you are performing a shelf replacement of the EOS disk shelves

Steps

1. Relocate the root aggregate:

```
system node migrate-root -node nodename -disklist disklist -raid-type raid-type
```

-node

Specifies the node that owns the root aggregate that you want to migrate.

-disklist

Specifies the list of disks on which the new root aggregate will be created. All disks must be spares and owned by the same node. The minimum number of disks required is dependent on the RAID type.

-raid-type

Specifies the RAID type of the root aggregate. The default value is **raid-dp**. This is the only type supported in advanced mode.

2. Monitor the progress of the job:

```
job show -id jobid -instance
```

Result

If all of the pre-checks are successful, the command starts a root volume replacement job and exits.

Starting or stopping a node

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

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

Related concepts

[Managing a node remotely using the SP/BMC](#) on page 49

Related information

[High-availability configuration](#)

Rebooting a node at the system prompt

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

Steps

1. If the cluster contains four or more nodes, verify that the node to be rebooted does not hold epsilon:
 - a. Set the privilege level to advanced:


```
set -privilege advanced
```
 - b. Determine which node holds epsilon:


```
cluster show
```

Example

The following example shows that “node1” holds epsilon:

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

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

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

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

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

- e. Return to the admin privilege level:

```
set -privilege admin
```

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

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

Example

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

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

Booting ONTAP at the boot environment prompt

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

Steps

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

The storage system console displays the boot environment prompt.

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

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

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

Related tasks

[Accessing the system console of a node](#) on page 41

Shutting down a node

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

Steps

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

- a. Set the privilege level to advanced:

```
set -privilege advanced
```

- b. Determine which node holds epsilon:

```
cluster show
```

Example

The following example shows that “node1” holds epsilon:

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

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

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

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

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

- e. Return to the admin privilege level:

```
set -privilege admin
```

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

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

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

Example

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

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

Managing a node by using the boot menu

You can use the boot menu to correct configuration problems on a node, reset the admin password, initialize disks, reset the node configuration, and restore the 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.
(9) Configure Advanced Drive Partitioning
Selection (1-9)?
```

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

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

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

To...	Select...
Restore the configuration information from the node's root volume to the boot device, such as a PC CompactFlash card	<p>6) Update flash from backup config</p> <p>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 must use this menu option to restore the configuration information from the node's root volume back to the boot device.</p>
Install new software on the node	<p>7) Install new software first</p> <p>If the 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 ONTAP software on a node that has no root volume installed. Do <i>not</i> use this menu option to upgrade ONTAP.</p>
Reboot the node	<p>8) Reboot node</p>
Unpartition all disks and remove their ownership information or clean the configuration and initialize the system with whole or partitioned disks	<p>9) Configure Advanced Drive Partitioning</p> <p>Beginning with ONTAP 9.2, the Advanced Drive Partitioning option provides additional management features for disks that are configured for root-data or root-data-data partitioning. The following options are available from Boot Option 9:</p> <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <pre>(9a) Unpartition all disks and remove their ownership information. (9b) Clean configuration and initialize system with partitioned disks. (9c) Clean configuration and initialize system with whole disks. (9d) Reboot the node. (9e) Return to main boot menu.</pre> </div>

Related tasks

[Accessing the system console of a node](#) on page 41

Related information

[Disk and aggregate management](#)

Managing a node remotely using the SP/BMC

You can manage a node remotely using an onboard controller, called a Service Processor (SP) or Baseboard Management Controller (BMC). This remote management controller is included in all current platform models. The controller stays operational regardless of the operating state of the node.

Most AFF, FAS, and V-Series platforms have an SP, except for the AFF A700s, which has a BMC.

Understanding the SP

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

The key capabilities of the SP include the following:

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

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

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

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

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

- The SP monitors environmental sensors and logs events to help you take timely and effective service actions.

The SP monitors environmental sensors such as the node temperatures, voltages, currents, and fan speeds. When an environmental sensor has reached an abnormal condition, the SP logs the abnormal readings, notifies ONTAP of the issue, and sends alerts and “down system” notifications as necessary through an AutoSupport message, regardless of whether the node can send AutoSupport messages.

The SP also logs events such as boot progress, Field Replaceable Unit (FRU) changes, events generated by ONTAP, and SP command history. You can manually invoke an AutoSupport message to include the SP log files that are collected from a specified node.

Other than generating these messages on behalf of a node that is down and attaching additional diagnostic information to AutoSupport messages, the SP has no effect on the AutoSupport functionality. The AutoSupport configuration settings and message content behavior are inherited from ONTAP.

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

If SNMP is enabled, the SP generates SNMP traps to configured trap hosts for all “down system” events.

- The SP has a nonvolatile memory buffer that stores up to 4,000 events in a system event log (SEL) to help you diagnose issues.

The SEL stores each audit log entry as an audit event. It is stored in onboard flash memory on the SP. The event list from the SEL is automatically sent by the SP to specified recipients through an AutoSupport message.

The SEL contains the following information:

- Hardware events detected by the SP—for example, sensor status about power supplies, voltage, or other components
- Errors detected by the SP—for example, a communication error, a fan failure, or a memory or CPU error
- Critical software events sent to the SP by the node—for example, a panic, a communication failure, a boot failure, or a user-triggered “down system” as a result of issuing the SP `system reset` or `system power cycle` command

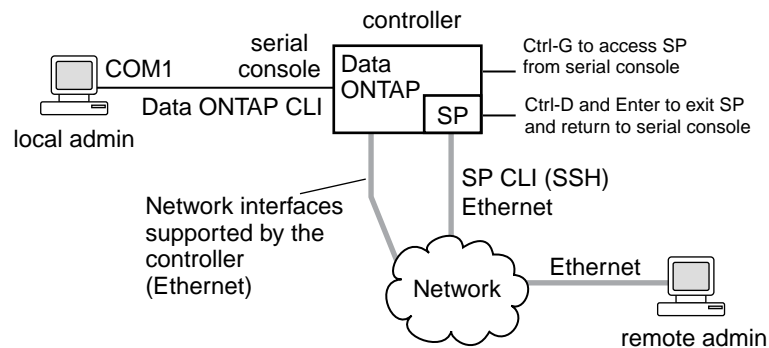
- The SP monitors the serial console regardless of whether administrators are logged in or connected to the console.

When messages are sent to the console, the SP stores them in the console log. The console log persists as long as the SP has power from either of the node power supplies. Because the SP

operates with standby power, it remains available even when the node is power-cycled or turned off.

- Hardware-assisted takeover is available if the SP is configured.
- The SP API service enables ONTAP to communicate with the SP over the network. The service enhances ONTAP management of the SP by supporting network-based functionality such as using the network interface for the SP firmware update, enabling a node to access another node's SP functionality or system console, and uploading the SP log from another node. You can modify the configuration of the SP API service by changing the port the service uses, renewing the SSL and SSH certificates that are used by the service for internal communication, or disabling the service entirely.

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



Related information

[High-availability configuration](#)

What the Baseboard Management Controller does

ONTAP 9.1 supports a new compact AFF storage controller, the AFF A700s. The ONTAP software is customized to support a new onboard controller in the AFF A700s called the Baseboard Management Controller (BMC). The BMC has command-line interface (CLI) commands you can use to manage the device remotely.

The BMC works similarly to the Service Processor (SP) and uses many of the same commands. The BMC allows you to do the following:

- Configure the BMC network settings.
- Access a node remotely and perform node management tasks such as diagnose, shut down, power-cycle, or reboot the node.

There are some differences between the SP and BMC:

- The BMC completely controls the environmental monitoring of power supply elements, cooling elements, temperature sensors, voltage sensors, and current sensors. The BMC reports sensor information to ONTAP through IPMI.
- Some of the high-availability (HA) and storage commands are different as the AFF A700s platform is only supported in an HA configuration.
- The BMC does not send AutoSupport messages.

Automatic firmware updates are also available when running ONTAP 9.2 GA or later with the following requirements:

- BMC firmware revision 1.15 or later must be installed.

Note: A manual update is required to upgrade BMC firmware from 1.12 to 1.15 or later.

- BMC automatically reboots after a firmware update is completed.

Note: Node operations are not impacted during a BMC reboot.

Related references

[ONTAP commands for BMC management](#) on page 76

[BMC CLI commands](#) on page 77

Configuring the SP/BMC network


Configuring the SP network enables you to access the SP by using its IP address. You can set up the SP automatic network configuration at the cluster level (recommended) or manually configure the SP network for individual nodes.

You can modify the SP API service that ONTAP uses to access the SP over the network, however the BMC does not have this capability.

Isolating management network traffic

It is a best practice to configure SP/BMC and the e0M management interface on a subnet dedicated to management traffic. Running data traffic over the management network can cause performance degradation and routing problems.

The management Ethernet port on most storage controllers (indicated by a wrench icon on the rear of the chassis) is connected to an internal Ethernet switch. The internal switch provides connectivity to SP/BMC and to the e0M management interface, which you can use to access the storage system via TCP/IP protocols like Telnet, SSH, and SNMP.



If you plan to use both the remote management device and e0M, you must configure them on the same IP subnet. Since these are low-bandwidth interfaces, the best practice is to configure SP/BMC and e0M on a subnet dedicated to management traffic.

If you cannot isolate management traffic, or if your dedicated management network is unusually large, you should try to keep the volume of network traffic as low as possible. Excessive ingress broadcast or multicast traffic may degrade SP/BMC performance.

Note: Some storage controllers, such as the AFF A800, have two external ports, one for BMC and the other for e0M. For these controllers, there is no requirement to configure BMC and e0M on the same IP subnet.

Considerations for the SP/BMC network configuration

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

Note: This topic applies to both the SP and the BMC.

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

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

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

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

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

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

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

Related tasks

[Enabling the SP/BMC automatic network configuration](#) on page 54

[Configuring the SP/BMC network manually](#) on page 55

Enabling the SP/BMC automatic network configuration

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

Before you begin

Note: This task applies to both the SP and the BMC.

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

Steps

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

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

If a node is out of quorum, you can modify the node's SP network configuration, including the SP IP address, by running `system service-processor network modify` from the node and confirming that you want to override the SP automatic network configuration for the node. However, when the node joins the quorum, the SP automatic reconfiguration takes place for the node based on the specified subnet.

Related concepts

Considerations for the SP/BMC network configuration on page 53

Related information

Network and LIF management

Configuring the SP/BMC network manually

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

Before you begin

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

About this task

Note: This task applies to both the SP and the BMC.

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

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

Steps

1. Configure the SP network for a node by using by using the `system service-processor network modify` command.
 - The `-address-family` parameter specifies whether the IPv4 or IPv6 configuration of the SP is to be modified.
 - The `-enable` parameter enables the network interface of the specified IP address family.
 - The `-dhcp` parameter specifies whether to use the network configuration from the DHCP server or the network address that you provide.
You can enable DHCP (by setting `-dhcp` to `v4`) only if you are using IPv4. You cannot enable DHCP for IPv6 configurations.
 - The `-ip-address` parameter specifies the public IP address for the SP.
A warning message appears when you attempt to manually configure the SP network with addresses that are allocated to a subnet. Ignoring the warning and proceeding with the manual address assignment might result in a duplicate address assignment.
 - The `-netmask` parameter specifies the netmask for the SP (if using IPv4.)
 - The `-prefix-length` parameter specifies the network prefix-length of the subnet mask for the SP (if using IPv6.)
 - The `-gateway` parameter specifies the gateway IP address for the SP.
2. Configure the SP network for the remaining nodes in the cluster by repeating the step 1 on page 55.
3. Display the SP network configuration and verify the SP setup status by using the `system service-processor network show` command with the `-instance` or `-fields setup-status` parameters.

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

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

Example of configuring the SP network

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

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

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

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

1 entries were displayed.

cluster1::>
```

Related concepts

[Considerations for the SP/BMC network configuration](#) on page 53

Related tasks

[Enabling the SP/BMC automatic network configuration](#) on page 54

Related information

[Network and LIF management](#)

Modifying the SP API service configuration

The SP API is a secure network API that enables ONTAP to communicate with the SP over the network. You can change the port used by the SP API service, renew the certificates the service uses

for internal communication, or disable the service entirely. You need to modify the configuration only in rare situations.

About this task

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

Steps

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

If you want to...	Use the following command...
Change the port used by the SP API service	<code>system service-processor api-service modify with the -port {49152..65535} parameter</code>
Renew the SSL and SSH certificates used by the SP API service for internal communication	<code>system service-processor api-service renew-certificates</code> If no parameter is specified, only the host certificates (including the client and server certificates) are renewed. If the <code>-renew-all true</code> parameter is specified, both the host certificates and the root CA certificate are renewed.
Disable or reenable the SP API service	<code>system service-processor api-service modify with the -is-enabled {true false} parameter</code>

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

Methods of managing SP/BMC firmware updates

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

Note: This topic applies to both the SP and the BMC.

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 ONTAP

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

Note: ONTAP detects a failed SP automatic update and triggers a corrective action to retry the SP automatic update up to three times. If all three retries fail, you should contact technical support.

- When you download a version of the SP firmware from the NetApp Support Site and the downloaded version is newer than the one that the SP is currently running
- When you downgrade or revert to an earlier version of ONTAP
The SP firmware is automatically updated to the newest compatible version that is supported by the ONTAP version you reverted or downgraded to. A manual SP firmware update is not required.

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

- ONTAP enables you to trigger an SP update manually and specify how the update should take place by using the `system service-processor image update` command.
You can specify the following options:
 - The SP firmware package to use (**-package**)
You can update the SP firmware to a downloaded package by specifying the package file name. The `system image package show` command displays all package files (including the files for the SP firmware package) that are available on a node.
 - Whether to use the baseline SP firmware package for the SP update (**-baseline**)
You can update the SP firmware to the baseline version that is bundled with the currently running version of ONTAP.
 - Whether to update the entire firmware image or only the changed portions, using the network interface or the serial interface (**-update-type**)
 - If updating the entire firmware image, whether to also reset log settings to the factory defaults and clear contents of all logs maintained by the SP, including the event logs, IPMI logs, and forensics logs (**-clear-logs**)

Note: If you use some of the more advanced update options or parameters, the BMC's configuration settings may be temporarily cleared. After reboot, it can take up to 10 minutes for ONTAP to restore the BMC configuration.
- ONTAP enables you to display the status for the latest SP firmware update triggered from ONTAP by using the `system service-processor image update-progress show` command.

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

Related information

[NetApp Downloads: System Firmware and Diagnostics](#)
[ONTAP 9 commands](#)

When the SP/BMC uses the network interface for firmware updates

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

Note: This topic applies to both the SP and the BMC.

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

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

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

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

Related information

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

Accessing the SP/BMC

A cluster user account that is configured with the **service-processor** application type can access the SP from an administration host (using an SSH connection) or the serial console. ONTAP enables you to restrict SSH access of the SP to only the administration hosts with specific IP addresses.

Note: This topic applies to both the SP and the BMC.

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

Related concepts

[Managing the automatic timeout period of CLI sessions](#) on page 28

Accounts that can access the SP

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

User accounts for accessing the SP are managed from 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.

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/BMC 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

Note: This task applies to both the SP and the BMC.

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

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

Steps

1. From the administration host, log in to the SP:

```
ssh username@SP_IP_address
```

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

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

Examples of SP access from an administration host

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

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

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

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

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

Related concepts

[Relationship among the SP CLI, SP console, and system console sessions](#) on page 62

Related tasks

[Accessing the system console of a node](#) on page 41

Accessing the SP/BMC from the system console

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

About this task

This task applies to both the SP and the BMC.

Steps

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

Example of accessing the SP CLI from the system console

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

```
cluster1::>
```

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

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

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

```
cluster1::>
```

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

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

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

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

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

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

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

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

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

Managing the IP addresses that can access the SP

By default, the SP accepts SSH connection requests from administration hosts of any IP addresses. You can configure the SP to accept SSH connection requests from only the administration hosts that have the IP addresses you specify. The changes you make apply to SSH access to the SP of any nodes in the cluster.

Steps

1. Grant SP access to only the IP addresses you specify by using the `system service-processor ssh add-allowed-addresses` command with the `-allowed-addresses` parameter.
 - The value of the `-allowed-addresses` parameter must be specified in the format of `address/netmask`, and multiple `address/netmask` pairs must be separated by commas, for example, `10.98.150.10/24, fd20:8b1e:b255:c09b::/64`.

Setting the `-allowed-addresses` parameter to `0.0.0.0/0, ::/0` enables all IP addresses to access the SP (the default).
 - When you change the default by limiting SP access to only the IP addresses you specify, ONTAP prompts you to confirm that you want the specified IP addresses to replace the “allow all” default setting (`0.0.0.0/0, ::/0`).

- The `system service-processor ssh show` command displays the IP addresses that can access the SP.
2. If you want to block a specified IP address from accessing the SP, use the `system service-processor ssh remove-allowed-addresses` command with the `-allowed-addresses` parameter.

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

Examples of managing the IP addresses that can access the SP

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

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

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

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

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

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

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

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

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

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

Using online help at the SP/BMC CLI

The online help displays the SP/BMC CLI commands and options.

About this task

This task applies to both the SP and the BMC.

Steps

1. To display help information for the SP/BMC commands, enter the following:

To access SP help...	To access BMC help...
Type <code>help</code> at the SP prompt.	Type <code>system</code> at the BMC prompt.

Example

The following example shows the SP CLI online help.

```

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

```

Example

The following example shows the BMC CLI online help.

```

BMC> system
Usage: system cmd [option]

Support cmd list
  help - display this help message
  fw - platform firmware related feature
  log - log related feature
  reset - reset the system
  console - connect to the system console
  core - dump the system core and reset
  power - system power related feature
  fru - FRU related feature

Support option list
  help: display the help message of the cmd

```

2. To display help information for the option of an SP/BMC command, enter

help

before or after the SP/BMC command.

Example

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

```

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

```

Example

The following example shows the BMC CLI online help for the BMC `system power` command.

```

BMC> system power help
Usage: system power [option]

Support option list
  help: print this help message
  cycle: power cycle the system, then on
  off: power the system off
  on: power the system on
  status: system power status

```


Commands for managing a node remotely

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

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

If you want to...	Use this SP command...	Or this ONTAP command ...
Display the status for the environmental sensors, including their states and current values	<code>system sensors</code> or <code>system sensors show</code>	<code>system node environment sensors show</code>
Display the status and details for the specified sensor	<code>system sensors get <i>sensor_name</i></code> You can obtain <i>sensor_name</i> by using the <code>system sensors</code> or the <code>system sensors show</code> command.	
Display the SP firmware version information	<code>version</code>	<code>system service-processor image show</code>
Display the SP command history	<code>sp log audit</code> (advanced privilege level)	
Display the SP debug information	<code>sp log debug</code> (advanced privilege level)	
Display the SP messages file	<code>sp log messages</code> (advanced privilege level)	
Display the settings for collecting system forensics on a watchdog reset event, display system forensics information collected during a watchdog reset event, or clear the collected system forensics information	<code>system forensics [show log dump log clear]</code>	
Log in to the system console	<code>system console</code>	<code>system node run-console</code>
	You should press Ctrl-D to exit the system console session.	
Turn the node on or off, or perform a power-cycle (turning the power off and then back on)	<code>system power on</code>	<code>system node power on</code> (advanced privilege level)
	<code>system power off</code>	
	<code>system power cycle</code>	
	The standby power stays on to keep the SP running without interruption. During the power-cycle, a brief pause occurs before power is turned back on. Attention: Using these commands to turn off or power-cycle the node might cause an improper shutdown of the node (also called a <i>dirty shutdown</i>) and is not a substitute for a graceful shutdown using the ONTAP <code>system node halt</code> command.	

If you want to...	Use this SP command...	Or this ONTAP command ...
Create a core dump and reset the node	<code>system core [-f]</code> The <code>-f</code> option forces the creation of a core dump and the reset of the node.	<code>system node coredump trigger</code> (advanced privilege level)
	These commands have the same effect as pressing the Non-maskable Interrupt (NMI) button on a node, causing a dirty shutdown of the node and forcing a dump of the core files when halting the node. These commands are helpful when ONTAP on the node is hung or does not respond to commands such as <code>system node shutdown</code> . The generated core dump files are displayed in the output of the <code>system node coredump show</code> command. The SP stays operational as long as the input power to the node is not interrupted.	
Reboot the node with an optionally specified BIOS firmware image (primary, backup, or current) to recover from issues such as a corrupted image of the node's boot device	<code>system reset {primary backup current}</code>	<code>system node reset</code> with the <code>-firmware {primary backup current}</code> parameter (advanced privilege level) <code>system node reset</code>
	<p>Attention: This operation causes a dirty shutdown of the node. If no BIOS firmware image is specified, the current image is used for the reboot. The SP stays operational as long as the input power to the node is not interrupted.</p>	
Display the status of battery firmware automatic update, or enable or disable battery firmware automatic update upon next SP boot	<code>system battery auto_update [status enable disable]</code> (advanced privilege level)	
Compare the current battery firmware image against a specified firmware image	<code>system battery verify [image_URL]</code> (advanced privilege level) If <code>image_URL</code> is not specified, the default battery firmware image is used for comparison.	
Update the battery firmware from the image at the specified location	<code>system battery flash image_URL</code> (advanced privilege level) You use this command if the automatic battery firmware upgrade process has failed for some reason.	
Update the SP firmware by using the image at the specified location	<code>sp update image_URL</code> <code>image_URL</code> must not exceed 200 characters.	<code>system service-processor image update</code>

If you want to...	Use this SP command...	Or this ONTAP command ...
Reboot the SP	<code>sp reboot</code>	<code>system service-processor reboot-sp</code>
	<p>Attention: You should avoid booting the SP from the backup image. Booting from the backup image is reserved for troubleshooting and recovery purposes only. It might require that the SP automatic firmware update be disabled, which is not a recommended setting. You should contact technical support before attempting to boot the SP from the backup image.</p>	
Erase the NVRAM flash content	<code>system nvram flash clear</code> (advanced privilege level) This command cannot be initiated when the controller power is off (<code>system power off</code>).	
Exit the SP CLI	<code>exit</code>	

Related information

[ONTAP 9 commands](#)

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

Threshold-based sensors take periodic readings of a variety of system components. The SP compares the reading of a threshold-based sensor against its preset threshold limits that define a component's acceptable operating conditions.

Based on the sensor reading, the SP displays the sensor state to help you monitor the condition of the component.

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


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

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

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

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

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



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

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

Example of the `system sensors` command output

The following example shows some of 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
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
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 sensor_name` command output for a threshold-based sensor

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

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

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

```
Upper Non-Recoverable : na
Assertion Events      :
Assertions Enabled    : lnc- lcr- ucr+
Deassertions Enabled : lnc- lcr- ucr+
```

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

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

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

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

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

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

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

System_FW_Status

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

`AA` can have one of the following values:

Values	Condition of the sensor
01	System firmware error
02	System firmware hang
04	System firmware progress

`BB` can have one of the following values:

Values	Condition of the sensor
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	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), 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

Values	Condition of the sensor
0x0081	Timer interrupt
0x0180	Timer expired
0x0280	Hard reset
0x0480	Power down
0x0880	Power cycle

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

PSU1_Input_Type and PSU2_Input_Type

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

Values	Condition of the sensor
--------	-------------------------

0x01 xx	220V PSU type
0x02 xx	110V PSU type

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

Commands for managing the SP from ONTAP

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

Commands for managing the SP network configuration

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

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

Commands for managing the SP firmware image

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

If you want to...	Run this ONTAP command...
<p>Enable or disable the SP automatic firmware update</p>	<pre>system service-processor image modify</pre> <p>By default, the SP firmware is automatically updated with the update of 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 ONTAP image and the SP firmware image.</p>
<p>Manually download an SP firmware image on a node</p>	<pre>system node image get</pre> <p>Note: Before you run the <code>system node image</code> commands, you must set the privilege level to advanced (<code>set -privilege advanced</code>), entering y when prompted to continue.</p> <p>The SP firmware image is packaged with 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 ONTAP.</p>
<p>Manually update the SP firmware, by specifying the following:</p> <ul style="list-style-type: none"> • The SP firmware package to use You can have the SP use a specific SP firmware package by specifying the package file name. The <code>system node image package show</code> command displays all package files (including the files for the SP firmware package) that are available on a node. • The installation baseline You can update the SP firmware to the baseline version that is bundled with the currently running version of ONTAP. • Whether to update the entire firmware image or only the changed portions, using the network interface or the serial interface • If updating the entire firmware image, whether to also reset log settings to the factory default and clear contents of all logs maintained by the SP, including the event logs, IPMI logs, and forensics logs 	<pre>system service-processor image update</pre>

If you want to...	Run this ONTAP command...
Display the status for the latest SP firmware update triggered from ONTAP, including the following information: <ul style="list-style-type: none"> The start and end time for the latest SP firmware update Whether an update is in progress and the percentage that is complete 	<pre>system service-processor image update-progress show</pre>

Commands for managing SSH access to the SP

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

Commands for general SP administration

If you want to...	Run this ONTAP command...
Display general SP information, including the following: <ul style="list-style-type: none"> The remote management device type The current SP status Whether the SP network is configured Network information, such as the public IP address and the MAC address The SP firmware version and Intelligent Platform Management Interface (IPMI) version Whether the SP firmware automatic update is enabled 	<pre>system service-processor show</pre> <p>Displaying complete SP information requires the <code>-instance</code> parameter.</p>
Reboot the SP on a node and optionally specify the SP firmware image (primary or backup) to use	<pre>system service-processor reboot-sp</pre> <p>Attention: You should avoid booting the SP from the backup image. Booting from the backup image is reserved for troubleshooting and recovery purposes only. It might require that the SP automatic firmware update be disabled, which is not a recommended setting. You should contact Technical Support before attempting to boot the SP from the backup image.</p>

If you want to...	Run this ONTAP command...
Generate and send an AutoSupport message that includes the SP log files collected from a specified node	<code>system node autosupport invoke-splog</code>
Display the allocation map of the collected SP log files in the cluster, including the sequence numbers for the SP log files that reside in each collecting node	<code>system service-processor log show-allocations</code>

Related information

[ONTAP 9 commands](#)

ONTAP commands for BMC management

These ONTAP commands are supported on the Baseboard Management Controller (BMC).

The BMC uses some of the same commands as the Service Processor (SP). The following SP commands are supported on the BMC.

If you want to...	Use this command
Display the BMC information	<code>system service-processor show</code>
Display/modify the BMC network configuration	<code>system service-processor network show/modify</code>
Reset the BMC	<code>system service-processor reboot-sp</code>
Display/modify the details of the currently installed BMC firmware image	<code>system service-processor image show/modify</code>
Update BMC firmware	<code>system service-processor image update</code>
Display the status for the latest BMC firmware update	<code>system service-processor image update-progress show</code>
Enable the automatic network configuration for the BMC to use an IPv4 or IPv6 address on the specified subnet	<code>system service-processor network auto-configuration enable</code>
Disable the automatic network configuration for an IPv4 or IPv6 address on the subnet specified for the BMC	<code>system service-processor network auto-configuration disable</code>
Display the BMC automatic network configuration	<code>system service-processor network auto-configuration show</code>

For commands that are not supported by the BMC firmware, the following error message is returned.

```
::> Error: Command not supported on this platform.
```

BMC CLI commands

You can log into the BMC using SSH. The following commands are supported from the BMC command line.

Command	Function
system	Display a list of all commands.
system console	Connect to the system's console. Use ESC + t to exit the session.
system core	Dump the system core and reset.
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.
system reset	Reset the system.
system fw ver	View all firmware versions that the BMC can read.
system fw upgrade	Update the BMC firmware.
system log console	Dump host console logs.
system log sel	Dump system event (SEL) logs.
system fru show [id]	Dump all/selected field replaceable unit (FRU) info.

Managing audit logging for management activities

ONTAP audit logging creates a record of activities that take place on the management interface. The audit logs can be used for monitoring, analysis, and backup purposes.

Changes to audit logging in ONTAP 9

Beginning in ONTAP 9, the `command-history.log` file is replaced by `audit.log`, and the `mgwd.log` file no longer contains audit information. If you are upgrading to ONTAP 9, you should review any scripts or tools that refer to the legacy files and their contents.

After upgrade to ONTAP 9, existing `command-history.log` files are preserved. They are rotated out (deleted) as new `audit.log` files are rotated in (created).

Tools and scripts that check the `command-history.log` file might continue to work, because a soft link from `command-history.log` to `audit.log` is created at upgrade. However, tools and scripts that check the `mgwd.log` file will fail, because that file no longer contains audit information.

In addition, audit logs in ONTAP 9 and later no longer include the following entries because they are not considered useful and cause unnecessary logging activity:

- Internal commands run by ONTAP (that is, where `username=root`)
- Command aliases (separately from the command they point to)

Beginning in ONTAP 9, you can transmit the audit logs securely to external destinations using the TCP and TLS protocols.

How ONTAP implements audit logging

Management activities recorded in the audit log are included in standard AutoSupport reports, and certain logging activities are included in EMS messages. You can also forward the audit log to destinations that you specify, and you can display audit log files by using the CLI or a web browser.

ONTAP logs management activities that are performed on the cluster, for example, what request was issued, the user who triggered the request, the user's access method, and the time of the request.

The management activities can be one of the following types:

- Set requests, which typically apply to non-display commands or operations
 - These requests are issued when you run a `create`, `modify`, or `delete` command, for instance.
 - Set requests are logged by default.
- Get requests, which retrieve information and display it in the management interface
 - These requests are issued when you run a `show` command, for instance.
 - Get requests are not logged by default, but you can use the `security audit modify` command to control whether get requests sent from the ONTAP CLI (`-cliget`) or from the ONTAP APIs (`-ontapiget`) are logged in the file.

ONTAP records management activities in the `/mroot/etc/log/mlog/audit.log` file of a node. Commands from the three shells for CLI commands—the clustershell, the nodeshell, and the non-interactive systemshell (interactive systemshell commands are not logged)—as well as API

commands are logged here. Audit logs include timestamps to show whether all nodes in a cluster are time synchronized.

The `audit.log` file is sent by the AutoSupport tool to the specified recipients. You can also forward the content securely to external destinations that you specify; for example, a Splunk or a syslog server.

The `audit.log` file is rotated daily. The rotation also occurs when it reaches 100 MB in size, and the previous 48 copies are preserved (with a maximum total of 49 files). When the audit file performs its daily rotation, no EMS message is generated. If the audit file rotates because its file size limit is exceeded, an EMS message is generated.

You can use the `security audit log show` command to display audit entries for individual nodes or merged from multiple nodes in the cluster. You can also display the content of the `/mroot/etc/log/mlog` directory on a single node by using a web browser.

Related tasks

[Accessing a node's log, core dump, and MIB files by using a web browser](#) on page 40

Forwarding the audit log to a destination

You can forward the audit log to a maximum of 10 destinations that you specify by using the `cluster log-forwarding create` command. For example, you can forward the log to a Splunk or syslog server for monitoring, analysis, or backup purposes.

About this task

If the `cluster log-forwarding create` command cannot ping the destination host to verify connectivity, the command fails with an error. Although not recommended, using the `-force` parameter with the command bypasses the connectivity verification.

You can configure transmission security options when forwarding log files:

- Protocols for sending messages to the destination
You can select one of the following `-protocol` values:
 - **udp-unencrypted**: User Datagram Protocol with no security (default)
 - **tcp-unencrypted**: Transmission Control Protocol with no security
 - **tcp-encrypted**: Transmission Control Protocol with Transport Layer Security (TLS)
- Verification of destination server identity
When you set the `-verify-server` parameter to **true**, the identity of the log forwarding destination is verified by validating its certificate. You can set the value to **true** only when you select the **tcp-encrypted** value in the `-protocol` field.

Steps

1. For each destination that you want to forward the audit log to, specify the destination IP address or host name and any security options.

Example

```
cluster1::> cluster log-forwarding create -destination 192.168.123.96
-port 514 -facility user

cluster1::> cluster log-forwarding create -destination 192.168.123.98
-port 514 -protocol tcp-encrypted -facility user
```

2. Verify that the destination records are correct by using the `cluster log-forwarding show` command.

Example

```
cluster1::> cluster log-forwarding show
```

Destination Host	Port	Protocol	Verify Server	Syslog Facility
192.168.123.96	514	udp-unencrypted	false	user
192.168.123.98	514	tcp-encrypted	true	user

2 entries were displayed.

Related information

[ONTAP 9 commands](#)

Commands for managing audit settings for management activities

You use the `security audit` commands to manage which management activities are logged in the `audit.log` file. You use the `cluster log-forwarding` commands to manage destinations for forwarding the audit log to.

If you want to...	Use this command...
Specify that get requests from the ONTAP CLI or APIs should be recorded in the audit log (the <code>audit.log</code> file), in addition to default set requests	<code>security audit modify</code>
Display the settings of the audit log	<code>security audit show</code>
Display audit entries merged from multiple nodes in the cluster	<code>security audit log show</code>
Specify a forwarding destination for the audit log and security measures for its transmission	<code>cluster log-forwarding create</code>
Modify a destination for the audit log	<code>cluster log-forwarding modify</code>
Delete a destination for the audit log	<code>cluster log-forwarding delete</code>
Show the configured destinations for the audit log	<code>cluster log-forwarding show</code>

Related information

[ONTAP 9 commands](#)

Managing the cluster time (cluster administrators only)

Problems can occur when the cluster time is inaccurate. Although ONTAP enables you to manually set the time zone, date, and time on the cluster, you should configure the Network Time Protocol (NTP) servers to synchronize the cluster time.

Beginning in ONTAP 9.5, you can configure your NTP server with symmetric authentication.

NTP is always enabled. However, configuration is still required for the cluster to synchronize with an external time source. ONTAP enables you to manage the cluster's NTP configuration in the following ways:

- You can associate a maximum of 10 external NTP servers with the cluster (`cluster time-service ntp server create`).
 - For redundancy and quality of time service, you should associate at least three external NTP servers with the cluster.
 - You can specify an NTP server by using its IPv4 or IPv6 address or fully qualified host name.
 - You can manually specify the NTP version (v3 or v4) to use. By default, ONTAP automatically selects the NTP version that is supported for a given external NTP server. If the NTP version you specify is not supported for the NTP server, time exchange cannot take place.
 - At the advanced privilege level, you can specify an external NTP server that is associated with the cluster to be the primary time source for correcting and adjusting the cluster time.
- You can display the NTP servers that are associated with the cluster (`cluster time-service ntp server show`).
- You can modify the cluster's NTP configuration (`cluster time-service ntp server modify`).
- You can disassociate the cluster from an external NTP server (`cluster time-service ntp server delete`).
- At the advanced privilege level, you can reset the configuration by clearing all external NTP servers' association with the cluster (`cluster time-service ntp server reset`).

A node that joins a cluster automatically adopts the NTP configuration of the cluster.

In addition to using NTP, ONTAP also enables you to manually manage the cluster time. This capability is helpful when you need to correct erroneous time (for example, a node's time has become significantly incorrect after a reboot). In that case, you can specify an approximate time for the cluster until NTP can synchronize with an external time server. The time you manually set takes effect across all nodes in the cluster.

You can manually manage the cluster time in the following ways:

- You can set or modify the time zone, date, and time on the cluster (`cluster date modify`).
- You can display the current time zone, date, and time settings of the cluster (`cluster date show`).

Note: Job schedules do not adjust to manual cluster date and time changes. These jobs are scheduled to run based on the current cluster time when the job was created or when the job most recently ran. Therefore, if you manually change the cluster date or time, you must use the job

`show` and `job history show` commands to verify that all scheduled jobs are queued and completed according to your requirements.

Related information

[Network Time Protocol \(NTP\) Support](#)

Commands for managing the cluster time

You use the `cluster time-service ntp server` commands to manage the NTP servers for the cluster. You use the `cluster date` commands to manage the cluster time manually.

Beginning in ONTAP 9.5, you can configure your NTP server with symmetric authentication.

The following commands enable you to manage the NTP servers for the cluster:

If you want to...	Use this command...
Associate the cluster with an external NTP server without symmetric authentication	<code>cluster time-service ntp server create -server <i>server_name</i></code>
Associate the cluster with an external NTP server with symmetric authentication Available in ONTAP 9.5 or later	<code>cluster time-service ntp server create -server <i>server_ip_address</i> -key-id <i>key_id</i></code> Note: The <i>key_id</i> must refer to an existing shared key configured with 'cluster time-service ntp key'.
Enable symmetric authentication for an existing NTP server An existing NTP server can be modified to enable authentication by adding the required key-id. Available in ONTAP 9.5 or later	<code>cluster time-service ntp server modify -server <i>server_name</i> -key-id <i>key_id</i></code>
Disable symmetric authentication	<code>cluster time-service ntp server modify -server <i>server_name</i> -authentication disabled</code>
Configure a shared NTP key	<code>cluster time-service ntp key create -id <i>shared_key_id</i> -type <i>shared_key_type</i> -value <i>shared_key_value</i></code> Note: Shared keys are referred to by an ID. The ID, its type, and value must be identical on both the node and the NTP server
Display information about the NTP servers that are associated with the cluster	<code>cluster time-service ntp server show</code>
Modify the configuration of an external NTP server that is associated with the cluster	<code>cluster time-service ntp server modify</code>
Dissociate an NTP server from the cluster	<code>cluster time-service ntp server delete</code>

If you want to...	Use this command...
Reset the configuration by clearing all external NTP servers' association with the cluster	<pre>cluster time-service ntp server reset</pre> <p>Note: This command requires the advanced privilege level.</p>

The following commands enable you to manage the cluster time manually:

If you want to...	Use this command...
Set or modify the time zone, date, and time	<pre>cluster date modify</pre>
Display the time zone, date, and time settings for the cluster	<pre>cluster date show</pre>

Related information

[ONTAP 9 commands](#)

Managing the banner and MOTD

ONTAP enables you to configure a login banner or a message of the day (MOTD) to communicate administrative information to CLI users of the cluster or storage virtual machine (SVM).

A banner is displayed in a console session (for cluster access only) or an SSH session (for cluster or SVM access) before a user is prompted for authentication such as a password. For example, you can use the banner to display a warning message such as the following to someone who attempts to log in to the system:

```
$ ssh admin@cluster1-01
This system is for authorized users only. Your IP Address has been logged.
Password:
cluster1::> _
```

An MOTD is displayed in a console session (for cluster access only) or an SSH session (for cluster or SVM access) after a user is authenticated but before the clustershell prompt appears. For example, you can use the MOTD to display a welcome or informational message such as the following that only authenticated users will see:

```
$ ssh admin@cluster1-01
Password:
Greetings. This system is running ONTAP 9.0.
Your user name is 'admin'. Your last login was Wed Apr 08 16:46:53 2015 from
10.72.137.28.
cluster1::> _
```

You can create or modify the content of the banner or MOTD by using the `security login banner modify` or `security login motd modify` command, respectively, in the following ways:

- You can use the CLI interactively or noninteractively to specify the text to use for the banner or MOTD.
 - The interactive mode, launched when the command is used without the `-message` or `-uri` parameter, enables you to use newlines (also known as end of lines) in the message.
 - The noninteractive mode, which uses the `-message` parameter to specify the message string, does not support newlines.
- You can upload content from an FTP or HTTP location to use for the banner or MOTD.
- You can configure the MOTD to display dynamic content.
 - Examples of what you can configure the MOTD to display dynamically include the following:
 - Cluster name, node name, or SVM name
 - Cluster date and time
 - Name of the user logging in
 - Last login for the user on any node in the cluster
 - Login device name or IP address
 - Operating system name

- Software release version
- Effective cluster version string

The `security login motd modify man` page describes the escape sequences that you can use to enable the MOTD to display dynamically generated content.

The banner does not support dynamic content.

You can manage the banner and MOTD at the cluster or SVM level:

- The following facts apply to the banner:
 - The banner configured for the cluster is also used for all SVMs that do not have a banner message defined.
 - An SVM-level banner can be configured for each SVM.
If a cluster-level banner has been configured, it is overridden by the SVM-level banner for the given SVM.
- The following facts apply to the MOTD:
 - By default, the MOTD configured for the cluster is also enabled for all SVMs.
 - Additionally, an SVM-level MOTD can be configured for each SVM.
In this case, users logging in to the SVM will see two MOTDs, one defined at the cluster level and the other at the SVM level.
 - The cluster-level MOTD can be enabled or disabled on a per-SVM basis by the cluster administrator.
If the cluster administrator disables the cluster-level MOTD for an SVM, a user logging in to the SVM does not see the cluster-level MOTD.

Creating a banner

You can create a banner to display a message to someone who attempts to access the cluster or SVM. The banner is displayed in a console session (for cluster access only) or an SSH session (for cluster or SVM access) before a user is prompted for authentication.

Steps

1. Use the `security login banner modify` command to create a banner for the cluster or SVM:

If you want to...	Then...
Specify a message that is a single line	Use the <code>-message "text"</code> parameter to specify the text.
Include newlines (also known as end of lines) in the message	Use the command without the <code>-message</code> or <code>-uri</code> parameter to launch the interactive mode for editing the banner.
Upload content from a location to use for the banner	Use the <code>-uri</code> parameter to specify the content's FTP or HTTP location.

The maximum size for a banner is 2,048 bytes, including newlines.

A banner created by using the `-uri` parameter is static. It is not automatically refreshed to reflect subsequent changes of the source content.

The banner created for the cluster is displayed also for all SVMs that do not have an existing banner. Any subsequently created banner for an SVM overrides the cluster-level banner for that

SVM. Specifying the `-message` parameter with a hyphen within double quotes ("`-`") for the SVM resets the SVM to use the cluster-level banner.

2. Verify that the banner has been created by displaying it with the `security login banner show` command.

Specifying the `-message` parameter with an empty string ("`''`") displays banners that have no content.

Specifying the `-message` parameter with "`-`" displays all (admin or data) SVMs that do not have a banner configured.

Examples of creating banners

The following example uses the noninteractive mode to create a banner for the "cluster1" cluster:

```
cluster1::> security login banner modify -message "Authorized users only!"
cluster1::>
```

The following example uses the interactive mode to create a banner for the "svm1" SVM:

```
cluster1::> security login banner modify -vserver svm1

Enter the message of the day for Vserver "svm1".
Max size: 2048. Enter a blank line to terminate input. Press Ctrl-C to abort.
0          1          2          3          4          5          6          7          8
1234567890123456789012345678901234567890123456789012345678901234567890
The svm1 SVM is reserved for authorized users only!

cluster1::>
```

The following example displays the banners that have been created:

```
cluster1::> security login banner show
Vserver: cluster1
Message
-----
Authorized users only!

Vserver: svm1
Message
-----
The svm1 SVM is reserved for authorized users only!

2 entries were displayed.

cluster1::>
```

Related tasks

[Managing the banner](#) on page 86

Managing the banner

You can manage the banner at the cluster or SVM level. The banner configured for the cluster is also used for all SVMs that do not have a banner message defined. A subsequently created banner for an SVM overrides the cluster banner for that SVM.

Choices

- Manage the banner at the cluster level:

If you want to...	Then...
Create a banner to display for all CLI login sessions	Set a cluster-level banner: <pre>security login banner modify -vserver cluster_name { [-message "text"] [-uri ftp_or_http_addr] }</pre>
Remove the banner for all (cluster and SVM) logins	Set the banner to an empty string (""): <pre>security login banner modify -vserver * -message ""</pre>
Override a banner created by an SVM administrator	Modify the SVM banner message: <pre>security login banner modify -vserver svm_name { [-message "text"] [-uri ftp_or_http_addr] }</pre>

- Manage the banner at the SVM level:

Specifying `-vserver svm_name` is not required in the SVM context.

If you want to...	Then...
Override the banner supplied by the cluster administrator with a different banner for the SVM	Create a banner for the SVM: <pre>security login banner modify -vserver svm_name { [-message "text"] [-uri ftp_or_http_addr] }</pre>
Suppress the banner supplied by the cluster administrator so that no banner is displayed for the SVM	Set the SVM banner to an empty string for the SVM: <pre>security login banner modify -vserver svm_name - message ""</pre>
Use the cluster-level banner when the SVM currently uses an SVM-level banner	Set the SVM banner to "-": <pre>security login banner modify -vserver svm_name - message "-"</pre>

Related tasks

[Creating a banner](#) on page 85

Creating an MOTD

You can create a message of the day (MOTD) to communicate information to authenticated CLI users. The MOTD is displayed in a console session (for cluster access only) or an SSH session (for cluster or SVM access) after a user is authenticated but before the clustershell prompt appears.

Steps

1. Use the `security login motd modify` command to create an MOTD for the cluster or SVM:

If you want to...	Then...
Specify a message that is a single line	Use the <code>-message "text"</code> parameter to specify the text.
Include newlines (also known as end of lines)	Use the command without the <code>-message</code> or <code>-uri</code> parameter to launch the interactive mode for editing the MOTD.
Upload content from a location to use for the MOTD	Use the <code>-uri</code> parameter to specify the content's FTP or HTTP location.

The maximum size for an MOTD is 2,048 bytes, including newlines.

The `security login motd modify` man page describes the escape sequences that you can use to enable the MOTD to display dynamically generated content.

An MOTD created by using the `-uri` parameter is static. It is not automatically refreshed to reflect subsequent changes of the source content.

An MOTD created for the cluster is displayed also for all SVM logins by default, along with an SVM-level MOTD that you can create separately for a given SVM. Setting the `-is-cluster-message-enabled` parameter to `false` for an SVM prevents the cluster-level MOTD from being displayed for that SVM.

2. Verify that the MOTD has been created by displaying it with the `security login motd show` command.

Specifying the `-message` parameter with an empty string (`""`) displays MOTDs that are not configured or have no content.

Examples of creating MOTDs

The following example uses the noninteractive mode to create an MOTD for the “cluster1” cluster:

```
cluster1::> security login motd modify -message "Greetings!"
cluster1::>
```

The following example uses the interactive mode to create an MOTD for the “svm1” SVM that uses escape sequences to display dynamically generated content:

```
cluster1::> security login motd modify -vserver svm1

Enter the message of the day for Vserver "svm1".
Max size: 2048. Enter a blank line to terminate input. Press Ctrl-C to abort.
0          1          2          3          4          5          6          7          8
12345678901234567890123456789012345678901234567890123456789012345678901234567890
Welcome to the \n SVM. Your user ID is '\N'. Your last successful login was \L.

cluster1::>
```

The following example displays the MOTDs that have been created:

```
cluster1::> security login motd show
Vserver: cluster1
Is the Cluster MOTD Displayed?: true
Message
-----
Greetings!

Vserver: svm1
Is the Cluster MOTD Displayed?: true
Message
-----
Welcome to the \n SVM. Your user ID is '\N'. Your last successful login was \L.

2 entries were displayed.

cluster1::>
```

Related tasks

[Managing the MOTD](#) on page 88

Managing the MOTD

You can manage the message of the day (MOTD) at the cluster or SVM level. By default, the MOTD configured for the cluster is also enabled for all SVMs. Additionally, an SVM-level MOTD can be

configured for each SVM. The cluster-level MOTD can be enabled or disabled for each SVM by the cluster administrator.

Choices

- Manage the MOTD at the cluster level:

If you want to...	Then...
Create an MOTD for all logins when there is no existing MOTD	Set a cluster-level MOTD: security login motd modify -vserver <i>cluster_name</i> { [-message "text"] [-uri <i>ftp_or_http_addr</i>] }
Change the MOTD for all logins when no SVM-level MOTDs are configured	Modify the cluster-level MOTD: security login motd modify -vserver <i>cluster_name</i> { [-message "text"] [-uri <i>ftp_or_http_addr</i>] }
Remove the MOTD for all logins when no SVM-level MOTDs are configured	Set the cluster-level MOTD to an empty string (""): security login motd modify -vserver <i>cluster_name</i> -message ""
Have every SVM display the cluster-level MOTD instead of using the SVM-level MOTD	Set a cluster-level MOTD, then set all SVM-level MOTDs to an empty string with the cluster-level MOTD enabled: <ol style="list-style-type: none"> security login motd modify -vserver <i>cluster_name</i> { [-message "text"] [-uri <i>ftp_or_http_addr</i>] } security login motd modify { -vserver !"cluster_name" } -message "" -is-cluster-message-enabled true
Have an MOTD displayed for only selected SVMs, and use no cluster-level MOTD	Set the cluster-level MOTD to an empty string, then set SVM-level MOTDs for selected SVMs: <ol style="list-style-type: none"> security login motd modify -vserver <i>cluster_name</i> -message "" security login motd modify -vserver <i>svm_name</i> { [-message "text"] [-uri <i>ftp_or_http_addr</i>] } <p>You can repeat this step for each SVM as needed.</p>
Use the same SVM-level MOTD for all (data and admin) SVMs	Set the cluster and all SVMs to use the same MOTD: security login motd modify -vserver * { [-message "text"] [-uri <i>ftp_or_http_addr</i>] } Note: If you use the interactive mode, the CLI prompts you to enter the MOTD individually for the cluster and each SVM. You can paste the same MOTD into each instance when you are prompted to.
Have a cluster-level MOTD optionally available to all SVMs, but do not want the MOTD displayed for cluster logins	Set a cluster-level MOTD, but disable its display for the cluster: security login motd modify -vserver <i>cluster_name</i> { [-message "text"] [-uri <i>ftp_or_http_addr</i>] } -is-cluster-message-enabled false
Remove all MOTDs at the cluster and SVM levels when only some SVMs have both cluster-level and SVM-level MOTDs	Set the cluster and all SVMs to use an empty string for the MOTD: security login motd modify -vserver * -message ""

If you want to...	Then...
Modify the MOTD only for the SVMs that have a non-empty string, when other SVMs use an empty string, and when a different MOTD is used at the cluster level	Use extended queries to modify the MOTD selectively: <pre>security login motd modify { - vserver !"cluster_name" -message !"" } { [- message "text"] [-uri ftp_or_http_addr] }</pre>
Display all MOTDs that contain specific text (for example, “January” followed by “2015”) anywhere in a single or multiline message, even if the text is split across different lines	Use a query to display MOTDs: <pre>security login motd show -message *"January"*"2015"*</pre>
Interactively create an MOTD that includes multiple and consecutive newlines (also known as end of lines, or EOLs)	In the interactive mode, press the space bar followed by Enter to create a blank line without terminating the input for the MOTD.

- Manage the MOTD at the SVM level:

Specifying `-vserver svm_name` is not required in the SVM context.

If you want to...	Then...
Use a different SVM-level MOTD, when the SVM already has an existing SVM-level MOTD	Modify the SVM-level MOTD: <pre>security login motd modify -vserver svm_name { [- message "text"] [-uri ftp_or_http_addr] }</pre>
Use only the cluster-level MOTD for the SVM, when the SVM already has an SVM-level MOTD	Set the SVM-level MOTD to an empty string, then have the cluster administrator enable the cluster-level MOTD for the SVM: <ol style="list-style-type: none"> 1. <code>security login motd modify -vserver svm_name -message ""</code> 2. (For the cluster administrator) <code>security login motd modify -vserver svm_name -is-cluster-message-enabled true</code>
Not have the SVM display any MOTD, when both the cluster-level and SVM-level MOTDs are currently displayed for the SVM	Set the SVM-level MOTD to an empty string, then have the cluster administrator disable the cluster-level MOTD for the SVM: <ol style="list-style-type: none"> 1. <code>security login motd modify -vserver svm_name -message ""</code> 2. (For the cluster administrator) <code>security login motd modify -vserver svm_name -is-cluster-message-enabled false</code>

Related tasks

[Creating an MOTD](#) on page 87

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. ONTAP enables you to manage feature licenses and monitor feature usage and license entitlement risk.

All license keys are 28 characters in length. Licenses installed prior to ONTAP 8.2 continue to work in 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 ONTAP 8.2 or later, or you perform a controller replacement procedure for a node in a cluster running ONTAP 8.2 or later), ONTAP requires that you enter the license key in the 28-character format.

You can find license keys for your initial or add-on software orders at the NetApp Support Site under **My Support > Software Licenses** (login required). If you cannot locate your license keys from the Software Licenses page, contact your sales or support representative.

In ONTAP 9.0 and 9.1, you must install a cluster base license key for each cluster, which you can install either during or after the cluster setup. Starting with ONTAP 9.2, installing the cluster base license key is optional. If this license key is not provided, a default cluster base license key is installed during the cluster create process with a Cluster Serial Number of 180000011. If you require a different or unique Cluster Serial Number for your operational requirements, you should install a cluster base license key.

ONTAP enables you to manage feature licenses in the following ways:

- Add one or more license keys (`system license add`)
- Display information about installed licenses (`system license show`)
- Display the packages that require licenses and their current license status on the cluster (`system license status show`)
- Delete a license from the cluster or a node whose serial number you specify (`system license delete`)
- Display or remove expired or unused licenses (`system license clean-up`)

ONTAP enables you to monitor feature usage and license entitlement risk in the following ways:

- Display a summary of feature usage in the cluster on a per-node basis (`system feature-usage show-summary`)
The summary includes counter information such as the number of weeks a feature was in use and the last date and time the feature was used.
- Display feature usage status in the cluster on a per-node and per-week basis (`system feature-usage show-history`)
The feature usage status can be **not-used**, **configured**, or **in-use**. If the usage information is not available, the status shows **not-available**.
- Display the status of license entitlement risk for each license package (`system license entitlement-risk show`)
The risk status can be **low**, **medium**, **high**, **unlicensed**, or **unknown**. The risk status is also included in the AutoSupport message. License entitlement risk does not apply to the base license package.
The license entitlement risk is evaluated by using a number of factors, which might include but are not limited to the following:
 - Each package's licensing state

- The type of each license, its expiry status, and the uniformity of the licenses across the cluster
- Usage for the features associated with the license package

If the evaluation process determines that the cluster has a license entitlement risk, the command output also suggests a corrective action.

Related information

[NetApp Knowledgebase Answer 1002749: Data ONTAP 8.2 and 8.3 Licensing Overview and References](#)

[NetApp Knowledgebase Answer 1032121: How to verify Data ONTAP Software Entitlements and related License Keys using the Support Site](#)

[NetApp Knowledgebase Answer 1000928: FAQ: Licensing updates in Data ONTAP 9.2](#)

[NetApp: Data ONTAP Entitlement Risk Status](#)

License types and licensed method

Understanding license types and the licensed method helps you manage the licenses in a cluster.

License types

A package can have one or more of the following types of license installed in the cluster. The `system license show` command displays the installed license type or types for a package.

- **Standard license (`license`)**

A standard license is a node-locked license. It is issued for a node with a specific system serial number (also known as a *controller serial number*). A standard license is valid only for the node that has the matching serial number.

Installing a standard, node-locked license entitles a node to the licensed functionality. For the cluster to use licensed functionality, at least one node must be licensed for the functionality. It might be out of compliance to use licensed functionality on a node that does not have an entitlement for the functionality.

ONTAP releases treat a license installed prior to Data ONTAP 8.2 as a standard license. Therefore, in ONTAP releases, all nodes in the cluster automatically have the standard license for the package that the previously licensed functionality is part of. The `system license show` command with the `-legacy yes` parameter indicates such licenses.
- **Site license (`site`)**

A site license is not tied to a specific system serial number. When you install a site license, all nodes in the cluster are entitled to the licensed functionality. The `system license show` command displays site licenses under the cluster serial number.

If your cluster has a site license and you remove a node from the cluster, the node does not carry the site license with it, and it is no longer entitled to the licensed functionality. If you add a node to a cluster that has a site license, the node is automatically entitled to the functionality granted by the site license.
- **Evaluation license (`demo`)**

An evaluation license is a temporary license that expires after a certain period of time (indicated by the `system license show` command). It enables you to try certain software functionality without purchasing an entitlement. It is a cluster-wide license, and it is not tied to a specific serial number of a node.

If your cluster has an evaluation license for a package and you remove a node from the cluster, the node does not carry the evaluation license with it.

Licensed method

It is possible to install both a cluster-wide license (the **site** or **demo** type) and a node-locked license (the **license** type) for a package. Therefore, an installed package can have multiple license types in the cluster. However, to the cluster, there is only one *licensed method* for a package. The **licensed method** field of the `system license status show` command displays the entitlement that is being used for a package. The command determines the licensed method as follows:

- If a package has only one license type installed in the cluster, the installed license type is the licensed method.
- If a package does not have any licenses installed in the cluster, the licensed method is **none**.
- If a package has multiple license types installed in the cluster, the licensed method is determined in the following priority order of the license type—**site**, **license**, and **demo**.

For example:

- If you have a site license, a standard license, and an evaluation license for a package, the licensed method for the package in the cluster is **site**.
- If you have a standard license and an evaluation license for a package, the licensed method for the package in the cluster is **license**.
- If you have only an evaluation license for a package, the licensed method for the package in the cluster is **demo**.

Commands for managing licenses

You use the `system license` commands to manage feature licenses for the cluster. You use the `system feature-usage` commands to monitor feature usage.

If you want to...	Use this command...
Add one or more licenses	<code>system license add</code>
Display information about installed licenses, for example: <ul style="list-style-type: none"> • License package name and description • License type (site, license, or demo) • Expiration date, if applicable • The cluster or nodes that a package is licensed for • Whether the license was installed prior to Data ONTAP 8.2 (legacy) • Customer ID 	<code>system license show</code> Note: 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"> • The package name • The licensed method • The expiration date, if applicable 	<code>system license status show</code>

If you want to...	Use this command...
Delete the license of a package from the cluster or a node whose serial number you specify	<code>system license delete</code>
Display or remove expired or unused licenses	<code>system license clean-up</code>
Display summary of feature usage in the cluster on a per-node basis	<code>system feature-usage show-summary</code>
Display feature usage status in the cluster on a per-node and per-week basis	<code>system feature-usage show-history</code>
Display the status of license entitlement risk for each license package	<code>system license entitlement-risk show</code> Note: Some information is displayed only when you use the <code>-detail</code> and <code>-instance</code> parameters.

Related information

[ONTAP 9 commands](#)

Managing jobs and schedules

A *job* is any asynchronous task that is managed by the Job Manager. Jobs are typically long-running volume operations such as copy, move, and mirror. You can monitor, pause, stop, and restart jobs, and you can configure them to run on specified schedules.

Job categories

There are three categories of jobs that you can manage: server-affiliated, cluster-affiliated, and private.

A job can be in any of the following categories:

Server-Affiliated jobs

These jobs are queued by the management framework to a specific node to be run.

Cluster-Affiliated jobs

These jobs are queued by the management framework to any node in the cluster to be run.

Private jobs

These jobs are specific to a node and do not use the replicated database (RDB) or any other cluster mechanism. The commands that manage private jobs require the advanced privilege level or higher.

Commands for managing jobs

Jobs are placed into a job queue and run in the background when resources are available. If a job is consuming too many cluster resources, you can stop it or pause it until there is less demand on the cluster. You can also monitor and restart jobs.

When you enter a command that invokes a job, typically, the command informs you that the job has been queued and then returns to the CLI command prompt. However, some commands instead report job progress and do not return to the CLI command prompt until the job has been completed. In these cases, you can press Ctrl-C to move the job to the background.

If you want to...	Use this command...
Display information about all jobs	<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>
Display information about job history	<code>job history show</code> Up to 25,000 job records are stored for each node in the cluster. Consequently, attempting to display the full job history could take a long time. To avoid potentially long wait times, you should display jobs by node, storage virtual machine (SVM), or record ID.

If you want to...	Use this command...
Display the list of private jobs	<code>job private show</code> (advanced privilege level)
Display information about completed private jobs	<code>job private show-completed</code> (advanced privilege level)
Display information about the initialization state for job managers	<code>job initstate show</code> (advanced privilege level)
Monitor the progress of a job	<code>job watch-progress</code>
Monitor the progress of a private job	<code>job private watch-progress</code> (advanced privilege level)
Pause a job	<code>job pause</code>
Pause a private job	<code>job private pause</code> (advanced privilege level)
Resume a paused job	<code>job resume</code>
Resume a paused private job	<code>job private resume</code> (advanced privilege level)
Stop a job	<code>job stop</code>
Stop a private job	<code>job private stop</code> (advanced privilege level)
Delete a job	<code>job delete</code>
Delete a private job	<code>job private delete</code> (advanced privilege level)
Disassociate a cluster-affiliated job with an unavailable node that owns it, so that another node can take ownership of that job	<code>job unclaim</code> (advanced privilege level)

Note: You can use the `event log show` command to determine the outcome of a completed job.

Related information

[ONTAP 9 commands](#)

Commands for managing job schedules

Many tasks—for instance, volume Snapshot copies—can be configured to run on specified schedules. Schedules that run at specific times are called *cron* schedules (similar to UNIX `cron` schedules). Schedules that run at intervals are called *interval* schedules. You use the `job schedule` commands to manage job schedules.

Job schedules do not adjust to manual changes to the cluster date and time. These jobs are scheduled to run based on the current cluster time when the job was created or when the job most recently ran. Therefore, if you manually change the cluster date or time, you should use the `job show` and `job history show` commands to verify that all scheduled jobs are queued and completed according to your requirements.

If the cluster is part of a MetroCluster configuration, then the job schedules on both clusters must be identical. Therefore, if you create, modify, or delete a job schedule, you must perform the same operation on the remote cluster.

If you want to...	Use this command...
Display information about all schedules	<code>job schedule show</code>
Display the list of jobs by schedule	<code>job schedule show-jobs</code>
Display information about cron schedules	<code>job schedule cron show</code>
Display information about interval schedules	<code>job schedule interval show</code>
Create a cron schedule	<code>job schedule cron create</code>
Create an interval schedule	<code>job schedule interval create</code> You must specify at least one of the following parameters: <code>-days</code> , <code>-hours</code> , <code>-minutes</code> , or <code>-seconds</code> .
Modify a cron schedule	<code>job schedule cron modify</code>
Modify an interval schedule	<code>job schedule interval modify</code>
Delete a schedule	<code>job schedule delete</code>
Delete a cron schedule	<code>job schedule cron delete</code>
Delete an interval schedule	<code>job schedule interval delete</code>

Related information

[ONTAP 9 commands](#)

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 *Data Protection Power Guide*.

Related information

[Data protection](#)

Managing configuration backups

The configuration backup schedules automatically create configuration backup files for each node in the cluster, and for the cluster itself. You can change some of the settings for these schedules, and you can create configuration backup files manually.

How the node and cluster configurations are backed up automatically

Three separate schedules automatically create cluster and node configuration backup files and replicate them among the nodes in the cluster.

The configuration backup files are automatically created according to the following schedules:

- Every 8 hours
- Daily
- Weekly

At each of these times, a node configuration backup file is created on each healthy node in the cluster. All of these node configuration backup files are then collected in a single cluster

configuration backup file along with the replicated cluster configuration and saved on one or more nodes in the cluster.

For single-node clusters (including Data ONTAP Edge systems), you can specify the configuration backup destination during software setup. After setup, those settings can be modified using ONTAP commands.

Commands for managing configuration backup schedules

You can use the `system configuration backup settings` commands to manage configuration backup schedules.

These commands are available at the advanced privilege level.

If you want to...	Use this command...
Change the settings for a configuration backup schedule: <ul style="list-style-type: none"> Specify a remote URL (either HTTP, HTTPS, or FTP) where the configuration backup files will be uploaded in addition to the default locations in the cluster Specify a user name to be used to log in to the remote URL Set the number of backups to keep for each configuration backup schedule 	<pre>system configuration backup settings modify</pre> <p>When you use HTTPS in the remote URL, use the <code>-validate-certification</code> option to enable or disable digital certificate validation. Certificate validation is disabled by default.</p> <p>Note: The web server to which you are uploading the configuration backup file must have PUT operations enabled for HTTP and POST operations enabled for HTTPS. For more information, see your web server's documentation.</p>
Set the password to be used to log in to the remote URL	<pre>system configuration backup settings set-password</pre>
View the settings for the configuration backup schedule	<pre>system configuration backup settings show</pre> <p>Note: You set the <code>-instance</code> parameter to view the user name and the number of backups to keep for each schedule.</p>

Related information

[ONTAP 9 commands](#)

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	<pre>system configuration backup create</pre>
Copy a configuration backup file from a node to another node in the cluster	<pre>system configuration backup copy</pre>

If you want to...	Use this command...
Upload a configuration backup file from a node in the cluster to a remote URL (either HTTP, HTTPS, or FTP)	<pre>system configuration backup upload</pre> <p>When you use HTTPS in the remote URL, use the <code>-validate-certification</code> option to enable or disable digital certificate validation. Certificate validation is disabled by default.</p> <p>Note: The web server to which you are uploading the configuration backup file must have PUT operations enabled for HTTP and POST operations enabled for HTTPS. For more information, see your web server's documentation.</p>
Download a configuration backup file from a remote URL to a node in the cluster, and, if specified, validate the digital certificate	<pre>system configuration backup download</pre> <p>When you use HTTPS in the remote URL, use the <code>-validate-certification</code> option to enable or disable digital certificate validation. Certificate validation is disabled by default.</p>
Rename a configuration backup file on a node in the cluster	<pre>system configuration backup rename</pre>
View the node and cluster configuration backup files for one or more nodes in the cluster	<pre>system configuration backup show</pre>
Delete a configuration backup file on a node	<pre>system configuration backup delete</pre> <p>Note: This command deletes the configuration backup file on the specified node only. If the configuration backup file also exists on other nodes in the cluster, it remains on those nodes.</p>

Related information

[ONTAP 9 commands](#)

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 100
2. [Restoring the node configuration using a configuration backup file](#) on page 101

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"> 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. b. If the configuration backup file you identify does not exist on the recovering node, then use the <code>system configuration backup copy</code> command to copy it to the recovering node.

If you previously re-created the cluster, you should choose a configuration backup file that was created after the cluster recreation. If you must use a configuration backup file that was created prior to the cluster recreation, then after recovering the node, you must re-create the cluster again.

Restoring the node configuration using a configuration backup file

You restore the node configuration using the configuration backup file that you identified and made available to the recovering node.

About this task

You should only perform this task to recover from a disaster that resulted in the loss of the node's local configuration files.

Steps

1. If the node is healthy, then at the advanced privilege level of a different node, use the `cluster modify` command with the `-node` and `-eligibility` parameters to mark it ineligible and isolate it from the cluster.

If the node is not healthy, then you should skip this step.

Example

This example modifies node2 to be ineligible to participate in the cluster so that its configuration can be restored:

```
cluster1::*> cluster modify -node node2 -eligibility false
```

2. Use the `system configuration recovery node restore` command at the advanced privilege level to restore the node's configuration from a configuration backup file.

If the node lost its identity, including its name, then you should use the `-nodename-in-backup` parameter to specify the node name in the configuration backup file.

Example

This example restores the node's configuration using one of the configuration backup files stored on the node:

```
cluster1::*> system configuration recovery node restore -backup
cluster1.8hour.2011-02-22.18_15_00.7z

Warning: This command overwrites local configuration files with
files contained in the specified backup file. Use this
command only to recover from a disaster that resulted
in the loss of the local configuration files.
The node will reboot after restoring the local configuration.
Do you want to continue? {y|n}: y
```

The configuration is restored, and the node reboots.

3. If you marked the node ineligible, then use the `system configuration recovery cluster sync` command to mark the node as eligible and synchronize it with the cluster.
4. If you are operating in a SAN environment, use the `system node reboot` command to reboot the node and reestablish SAN quorum.

After you finish

If you previously re-created the cluster, and if you are restoring the node configuration by using a configuration backup file that was created prior to that cluster re-creation, then you must re-create the cluster again.

Related tasks

[Synchronizing a node with the cluster](#) on page 105

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 102
2. [Restoring a cluster configuration from an existing configuration](#) on page 103

Finding a configuration to use for recovering a cluster

You use the configuration from either a node in the cluster or a cluster configuration backup file to recover a cluster.

Steps

1. Choose a type of configuration to recover the cluster.
 - A node in the cluster

If the cluster consists of more than one node, and one of the nodes has a cluster configuration from when the cluster was in the desired configuration, then you can recover the cluster using the configuration stored on that node.

In most cases, the node containing the replication ring with the most recent transaction ID is the best node to use for restoring the cluster configuration. The `cluster ring show` command at the advanced privilege level enables you to view a list of the replicated rings available on each node in the cluster.
 - A cluster configuration backup file

If you cannot identify a node with the correct cluster configuration, or if the cluster consists of a single node, then you can use a cluster configuration backup file to recover the cluster.

Recovering a cluster from a configuration backup file might require the use of diagnostic tools to resolve discrepancies between the backup and the cluster configuration present on the recovering node. Therefore, if you use a configuration backup file, you should plan to contact technical support.

2. If you chose to use a cluster configuration backup file, then make the file available to the node you plan to use to recover the cluster.

If the configuration backup file is located...	Then...
At a remote URL	Use the <code>system configuration backup download</code> command at the advanced privilege level to download it to the recovering node.
On a node in the cluster	<ol style="list-style-type: none"> 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. b. If the cluster configuration backup file is not located on the node you plan to use to recover the cluster, then use the <code>system configuration backup copy</code> command to copy it to the recovering node.

Restoring a cluster configuration from an existing configuration

To restore a cluster configuration from an existing configuration after a cluster failure, you re-create the cluster using the cluster configuration that you chose and made available to the recovering node, and then rejoin each additional node to the new cluster.

About this task

You should only perform this task to recover from a disaster that resulted in the loss of the cluster's configuration.

Attention: If you are re-creating the cluster from a configuration backup file, you must contact technical support to resolve any discrepancies between the configuration backup file and the configuration present in the cluster.

Steps

1. Disable storage failover for each HA pair:

```
storage failover modify -node node_name -enabled false
```

You only need to disable storage failover once for each HA pair. When you disable storage failover for a node, storage failover is also disabled on the node's partner.

2. Halt each node except for the recovering node:

```
system node halt -node node_name -reason "text"
```

Example

```
cluster1::*> system node halt -node node0 -reason "recovering cluster"
Warning: Are you sure you want to halt the node? {y|n}: y
```

3. Set the privilege level to advanced:

```
set -privilege advanced
```

4. On the recovering node, use the `system configuration recovery cluster recreate` command to re-create the cluster.

Example

This example re-creates the cluster using the configuration information stored on the recovering node:

```
cluster1::*> configuration recovery cluster recreate -from node

Warning: This command will destroy your existing cluster. It will
        rebuild a new single-node cluster consisting of this node
        and its current configuration. This feature should only be
        used to recover from a disaster. Do not perform any other
        recovery operations while this operation is in progress.
Do you want to continue? {y|n}: y
```

A new cluster is created on the recovering node.

5. If you are re-creating the cluster from a configuration backup file, verify that the cluster recovery is still in progress:

```
system configuration recovery cluster show
```

You do not need to verify the cluster recovery state if you are re-creating the cluster from a healthy node.

Example

```
cluster1::*> system configuration recovery cluster show
Recovery Status: in-progress
Is Recovery Status Persisted: false
```

6. Boot each node that needs to be rejoined to the re-created cluster.

You must reboot the nodes one at a time.

7. For each node that needs to be joined to the re-created cluster, do the following:
 - a. From a healthy node on the re-created cluster, rejoin the target node:

```
system configuration recovery cluster rejoin -node node_name
```

Example

This example rejoins the “node2” target node to the re-created cluster:

```
cluster1::*> system configuration recovery cluster rejoin -node
node2

Warning: This command will rejoin node "node2" into the local
        cluster, potentially overwriting critical cluster
        configuration files. This command should only be used
        to recover from a disaster. Do not perform any other
        recovery operations while this operation is in progress.
        This command will cause node "node2" to reboot.
Do you want to continue? {y|n}: y
```

The target node reboots and then joins the cluster.

- b. Verify that the target node is healthy and has formed quorum with the rest of the nodes in the cluster:

```
cluster show -eligibility true
```

The target node must rejoin the re-created cluster before you can rejoin another node.

Example

```
cluster1::*> cluster show -eligibility true
Node           Health Eligibility Epsilon
-----
node0          true   true        false
node1          true   true        false
2 entries were displayed.
```

8. If you re-created the cluster from a configuration backup file, set the recovery status to be complete:
`system configuration recovery cluster modify -recovery-status complete`
9. Return to the admin privilege level:
`set -privilege admin`
10. If the cluster consists of only two nodes, use the `cluster ha modify` command to reenables cluster HA.
11. Use the `storage failover modify` command to reenables storage failover for each HA pair.

After you finish

If the cluster has SnapMirror peer relationships, then you also need to re-create those relationships. For more information, see the *Data Protection Power Guide*.

Related information

[Data protection](#)

Synchronizing a node with the cluster

If cluster-wide quorum exists, but one or more nodes are out of sync with the cluster, then you must synchronize the node to restore the replicated database (RDB) on the node and bring it into quorum.

Step

1. From a healthy node, use the `system configuration recovery cluster sync` command at the advanced privilege level to synchronize the node that is out of sync with the cluster configuration.

Example

This example synchronizes a node (*node2*) with the rest of the cluster:

```
cluster1::*> system configuration recovery cluster sync -node node2

Warning: This command will synchronize node "node2" with the cluster
configuration, potentially overwriting critical cluster
configuration files on the node. This feature should only be
used to recover from a disaster. Do not perform any other
recovery operations while this operation is in progress. This
command will cause all the cluster applications on node
"node2" to restart, interrupting administrative CLI and Web
interface on that node.
Do you want to continue? {y|n}: y
All cluster applications on node "node2" will be restarted. Verify
that the cluster applications go online.
```

Result

The RDB is replicated to the node, and the node becomes eligible to participate in the cluster.

Managing core dumps (cluster administrators only)

When a node panics, a core dump occurs and the system creates a core dump file that technical support can use to troubleshoot the problem. You can configure or display core dump attributes. You can also save, display, segment, upload, or delete a core dump file.

You can manage core dumps in the following ways:

- Configuring core dumps and displaying the configuration settings
- Displaying basic information, the status, and attributes of core dumps
Core dump files and reports are stored in the `/mroot/etc/crash/` directory of a node. You can display the directory content by using the `system node coredump` commands or a web browser.
- Saving the core dump content and uploading the saved file to a specified location or to technical support
ONTAP prevents you from initiating the saving of a core dump file during a takeover, an aggregate relocation, or a giveback.
- Deleting core dump files that are no longer needed

Note: AFF A220, AFF A800, FAS2720, FAS2750, and later systems store core dumps on their boot device. When NetApp Volume Encryption (NVE) or NetApp Storage Encryption (NSE) is enabled on these systems, the core dump is also encrypted.

Related tasks

[Accessing a node's log, core dump, and MIB files by using a web browser](#) on page 40

Related information

[NetApp Knowledgebase Answer 1030587: How to upload a clustered Data ONTAP, ONTAP 9 core file for analysis](#)

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

If you want to...	Use this command...
Save a specified core dump	<code>system node coredump save</code>
Save all unsaved core dumps that are on a specified node	<code>system node coredump save-all</code>
Generate and send an AutoSupport message with a core dump file you specify	<code>system node autosupport invoke-core-upload</code> Note: The <code>-uri</code> optional parameter specifies an alternate destination for the AutoSupport message.
Display status information about core dumps	<code>system node coredump status</code>
Delete a specified core dump	<code>system node coredump delete</code>
Delete all unsaved core dumps or all saved core files on a node	<code>system node coredump delete-all</code>
Display application core dump reports	<code>system node coredump reports show</code>
Delete an application core dump report	<code>system node coredump reports delete</code>

Related information

[ONTAP 9 commands](#)

Monitoring a storage system

You can use event messages, the AutoSupport feature, statistics, and environmental component sensors to monitor a storage system.

The cluster administrator can perform all system monitoring tasks.

Managing AutoSupport

AutoSupport is a mechanism that proactively monitors the health of your system and automatically sends messages to NetApp technical support, your internal support organization, and a support partner. Although AutoSupport messages to technical support are enabled by default, you must set the correct options and have a valid mail host to have messages sent to your internal support organization.

Only the cluster administrator can perform AutoSupport management. The storage virtual machine (SVM) administrator has no access to AutoSupport.

AutoSupport is enabled by default when you configure your storage system for the first time. AutoSupport begins sending messages to technical support 24 hours after AutoSupport is enabled. You can shorten the 24-hour period by upgrading or reverting the system, modifying the AutoSupport configuration, or changing the system time to be something other than a 24-hour period.

Note: You can disable AutoSupport at any time, but you should leave it enabled. Enabling AutoSupport can significantly help speed problem determination and resolution should a problem occur on your storage system. By default, the system collects AutoSupport information and stores it locally, even if you disable AutoSupport.

For more information about AutoSupport, see the NetApp Support Site.

Related information

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

When and where AutoSupport messages are sent

AutoSupport sends messages to different recipients, depending on the type of message. Learning when and where AutoSupport sends messages can help you understand messages that you receive through email or view on the Active IQ (formerly known as My AutoSupport) web site.

Unless specified otherwise, settings in the following tables are parameters of the `system node autosupport modify` command.

Event-triggered messages

When events occur on the system that require corrective action, AutoSupport automatically sends an event-triggered message.

When the message is sent	Where the message is sent
AutoSupport responds to a trigger event in the EMS	Addresses specified in <code>-to</code> and <code>-noteto</code> . (Only critical, service-affecting events are sent.) Addresses specified in <code>-partner-address</code> Technical support, if <code>-support</code> is set to enable

Scheduled messages

AutoSupport automatically sends several messages on a regular schedule.

When the message is sent	Where the message is sent
Daily (by default, sent between 12:00 a.m. and 1:00 a.m. as a log message)	Addresses specified in <code>-partner-address</code> Technical support, if <code>-support</code> is set to enable
Daily (by default, sent between 12:00 a.m. and 1:00 a.m. as a performance message), if the <code>-perf</code> parameter is set to true	Addresses specified in <code>-partner-address</code> Technical support, if <code>-support</code> is set to enable
Weekly (by default, sent Sunday between 12:00 a.m. and 1:00 a.m.)	Addresses specified in <code>-partner-address</code> Technical support, if <code>-support</code> is set to enable

Manually triggered messages

You can manually initiate or resend an AutoSupport message.

When the message is sent	Where the message is sent
You manually initiate a message using the <code>system node autosupport invoke</code> command	If a URI is specified using the <code>-uri</code> parameter in the <code>system node autosupport invoke</code> command, the message is sent to that URI. If <code>-uri</code> is omitted, the message is sent to the addresses specified in <code>-to</code> and <code>-partner-address</code> . The message is also sent to technical support if <code>-support</code> is set to enable .
You manually initiate a message using the <code>system node autosupport invoke-core-upload</code> command	If a URI is specified using the <code>-uri</code> parameter in the <code>system node autosupport invoke-core-upload</code> command, the message is sent to that URI, and the core dump file is uploaded to the URI. If <code>-uri</code> is omitted in the <code>system node autosupport invoke-core-upload</code> command, the message is sent to technical support, and the core dump file is uploaded to the technical support site. Both scenarios require that <code>-support</code> is set to enable and <code>-transport</code> is set to https or http . Due to the large size of core dump files, the message is not sent to the addresses specified in the <code>-to</code> and <code>-partner-addresses</code> parameters.

When the message is sent	Where the message is sent
You manually initiate a message using the <code>system node autosupport invoke-performance-archive</code> command	<p>If a URI is specified using the <code>-uri</code> parameter in the <code>system node autosupport invoke-performance-archive</code> command, the message is sent to that URI, and the performance archive file is uploaded to the URI.</p> <p>If <code>-uri</code> is omitted in the <code>system node autosupport invoke-performance-archive</code>, the message is sent to technical support, and the performance archive file is uploaded to the technical support site.</p> <p>Both scenarios require that <code>-support</code> is set to enable and <code>-transport</code> is set to https or http.</p> <p>Due to the large size of performance archive files, the message is not sent to the addresses specified in the <code>-to</code> and <code>-partner-addresses</code> parameters.</p>
You manually resend a past message using the <code>system node autosupport history retransmit</code> command	Only to the URI that you specify in the <code>-uri</code> parameter of the <code>system node autosupport history retransmit</code> command

Messages triggered by technical support

Technical support can request messages from AutoSupport using the AutoSupport OnDemand feature.

When the message is sent	Where the message is sent
When AutoSupport obtains delivery instructions to generate new AutoSupport messages	Addresses specified in <code>-partner-address</code> Technical support, if <code>-support</code> is set to enable and <code>-transport</code> is set to https
When AutoSupport obtains delivery instructions to resend past AutoSupport messages	Technical support, if <code>-support</code> is set to enable and <code>-transport</code> is set to https
When AutoSupport obtains delivery instructions to generate new AutoSupport messages that upload core dump or performance archive files	Technical support, if <code>-support</code> is set to enable and <code>-transport</code> is set to https . The core dump or performance archive file is uploaded to the technical support site.

Related concepts

[How AutoSupport OnDemand obtains delivery instructions from technical support](#) on page 115

How AutoSupport creates and sends event-triggered messages

AutoSupport creates event-triggered AutoSupport messages when the EMS processes a trigger event. An event-triggered AutoSupport message alerts recipients to problems that require corrective action and contains only information that is relevant to the problem. You can customize what content to include and who receives the messages.

AutoSupport uses the following process to create and send event-triggered AutoSupport messages:

1. When the EMS processes a trigger event, EMS sends AutoSupport a request.

A trigger event is an EMS event with an AutoSupport destination and a name that begins with a `callhome.` prefix.

2. AutoSupport creates an event-triggered AutoSupport message.
AutoSupport collects basic and troubleshooting information from subsystems that are associated with the trigger to create a message that includes only information that is relevant to the trigger event.
A default set of subsystems is associated with each trigger. However, you can choose to associate additional subsystems with a trigger by using the `system node autosupport trigger modify` command.
3. AutoSupport sends the event-triggered AutoSupport message to the recipients defined by the `system node autosupport modify` command with the `-to`, `-noteto`, `-partner-address`, and `-support` parameters.
You can enable and disable delivery of AutoSupport messages for specific triggers by using the `system node autosupport trigger modify` command with the `-to` and `-noteto` parameters.

Example of data sent for a specific event

The `storage shelf PSU failed` EMS event triggers a message that contains basic data from the Mandatory, Log Files, Storage, RAID, HA, Platform, and Networking subsystems and troubleshooting data from the Mandatory, Log Files, and Storage subsystems.

You decide that you want to include data about NFS in any AutoSupport messages sent in response to a future `storage shelf PSU failed` event. You enter the following command to enable troubleshooting-level data for NFS for the `callhome.shlf.ps.fault` event:

```
cluster1::> system node autosupport trigger modify -node node1 -
autosupport-message shlf.ps.fault -troubleshooting-additional nfs
```

Note that the `callhome.` prefix is dropped from the `callhome.shlf.ps.fault` event when you use the `system node autosupport trigger` commands, or when referenced by AutoSupport and EMS events in the CLI.

Types of AutoSupport messages and their content

AutoSupport messages contain status information about supported subsystems. Learning what AutoSupport messages contain can help you interpret or respond to messages that you receive in email or view on the Active IQ (formerly known as 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"> • test sends a user-triggered message with some basic data. This message also triggers an automated email response from technical support to any specified email addresses, using the <code>-to</code> option, so that you can confirm that AutoSupport messages are being received. • performance sends performance data. • all sends a user-triggered message with a complete set of data similar to the weekly message, including troubleshooting data from each subsystem. Technical support typically requests this message.
Triggered by the <code>system node autosupport invoke-core-upload</code> command	Core dump files for a node
Triggered by the <code>system node autosupport invoke-performance-archive</code> command	Performance archive files for a specified period of time
Triggered by AutoSupport OnDemand	<p>AutoSupport OnDemand can request new messages or past messages:</p> <ul style="list-style-type: none"> • New messages, depending on the type of AutoSupport collection, can be test, all, or performance. • Past messages depend on the type of message that is resent. <p>AutoSupport OnDemand can request new messages that upload the following files to the NetApp Support Site at mysupport.netapp.com:</p> <ul style="list-style-type: none"> • Core dump • Performance archive

What AutoSupport subsystems are

Each subsystem provides basic and troubleshooting information that AutoSupport uses for its messages. Each subsystem is also associated with trigger events that allow AutoSupport to collect from subsystems only information that is relevant to the trigger event.

AutoSupport collects context-sensitive content. You can view information about subsystems and trigger events by using the `system node autosupport trigger show` command.

AutoSupport size and time budgets

AutoSupport collects information, organized by subsystem, and enforces a size and time budget on content for each subsystem. As storage systems grow, AutoSupport budgets provide control over the AutoSupport payload, which in turn provides scalable delivery of AutoSupport data.

AutoSupport stops collecting information and truncates the AutoSupport content if the subsystem content exceeds its size or time budget. If the content cannot be truncated easily (for example, binary files), AutoSupport omits the content.

You should modify the default size and time budgets only with guidance from technical support. You can also review the default size and time budgets of the subsystems by using the `autosupport manifest show` command.

Files sent in event-triggered AutoSupport messages

Event-triggered AutoSupport messages only contain basic and troubleshooting information from subsystems that are associated with the event that caused AutoSupport to generate the message. The specific data helps NetApp support and support partners troubleshoot the problem.

AutoSupport uses the following criteria to control content in event-triggered AutoSupport messages:

- Which subsystems are included
Data is grouped into subsystems, including common subsystems, such as Log Files, and specific subsystems, such as RAID. Each event triggers a message that contains only the data from specific subsystems.
- The detail level of each included subsystem
Data for each included subsystem is provided at a basic or troubleshooting level.

You can view all possible events and determine which subsystems are included in messages about each event using the `system node autosupport trigger show` command with the `-instance` parameter.

In addition to the subsystems that are included by default for each event, you can add additional subsystems at either a basic or a troubleshooting level using the `system node autosupport trigger modify` command.

Log files sent in AutoSupport messages

AutoSupport messages can contain several key log files that enable technical support staff to review recent system activity.

All types of AutoSupport messages might include the following log files when the Log Files subsystem is enabled:

Log file	Amount of data included from the file
<ul style="list-style-type: none"> • Log files from the <code>/mroot/etc/log/mlog/</code> directory • The MESSAGES log file 	<p>Only new lines added to the logs since the last AutoSupport message up to a specified maximum. This ensures that AutoSupport messages have unique, relevant—not overlapping—data. (Log files from partners are the exception; for partners, the maximum allowed data is included.)</p>

Log file	Amount of data included from the file
<ul style="list-style-type: none"> • Log files from the <code>/mroot/etc/log/shelflog/</code> directory • Log files from the <code>/mroot/etc/log/acp/</code> directory • Event Management System (EMS) log data 	The most recent lines of data up to a specified maximum.

The content of AutoSupport messages can change between releases of ONTAP.

Files sent in weekly AutoSupport messages

Weekly AutoSupport messages contain additional configuration and status data that is useful to track changes in your system over time.

The following information is sent in weekly AutoSupport messages:

- Basic information about every subsystem
- Contents of selected `/mroot/etc` directory files
- Log files
- Output of commands that provide system information
- Additional information, including replicated database (RDB) information, service statistics, and more

How AutoSupport OnDemand obtains delivery instructions from technical support

AutoSupport OnDemand periodically communicates with technical support to obtain delivery instructions for sending, resending, and declining AutoSupport messages as well as uploading large files to the NetApp support site. AutoSupport OnDemand enables AutoSupport messages to be sent on-demand instead of waiting for the weekly AutoSupport job to run.

AutoSupport OnDemand consists of the following components:

- AutoSupport OnDemand client that runs on each node
- AutoSupport OnDemand service that resides in technical support

The AutoSupport OnDemand client periodically polls the AutoSupport OnDemand service to obtain delivery instructions from technical support. For example, technical support can use the AutoSupport OnDemand service to request that a new AutoSupport message be generated. When the AutoSupport OnDemand client polls the AutoSupport OnDemand service, the client obtains the delivery instructions and sends the new AutoSupport message on-demand as requested.

AutoSupport OnDemand is enabled by default. However, AutoSupport OnDemand relies on some AutoSupport settings to continue communicating with technical support. AutoSupport OnDemand automatically communicates with technical support when the following requirements are met:


- AutoSupport is enabled.
- AutoSupport is configured to send messages to technical support.
- AutoSupport is configured to use the HTTPS transport protocol.

The AutoSupport OnDemand client sends HTTPS requests to the same technical support location to which AutoSupport messages are sent. The AutoSupport OnDemand client does not accept incoming connections.

Note: AutoSupport OnDemand uses the “autosupport” user account to communicate with technical support. ONTAP prevents you from deleting this account.

If you want to have AutoSupport OnDemand disabled, but keep AutoSupport enabled and configured to send messages to technical support by using the HTTPS transport protocol, contact technical support.

The following illustration shows how AutoSupport OnDemand sends HTTPS requests to technical support to obtain delivery instructions.



The delivery instructions can include requests for AutoSupport to do the following:

- Generate new AutoSupport messages.
Technical support might request new AutoSupport messages to help triage issues.
- Generate new AutoSupport messages that upload core dump files or performance archive files to the NetApp support site.
Technical support might request core dump or performance archive files to help triage issues.
- Retransmit previously generated AutoSupport messages.
This request automatically happens if a message was not received due to a delivery failure.
- Disable delivery of AutoSupport messages for specific trigger events.
Technical support might disable delivery of data that is not used.

Structure of AutoSupport messages sent by email

When an AutoSupport message is sent by email, the message has a standard subject, a brief body, and a large attachment in 7z file format that contains the data.

Note: If AutoSupport is configured to hide private data, certain information, such as the hostname, is omitted or masked in the header, subject, body, and attachments.

Subject

The subject line of messages sent by the AutoSupport mechanism contains a text string that identifies the reason for the notification. The format of the subject line is as follows:

HA Group Notification from *System_Name* (*Message*) *Severity*

- *System_Name* is either the hostname or the system ID, depending on the AutoSupport configuration

Body

The body of the AutoSupport message contains the following information:

- Date and timestamp of the message
- Version of 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 an emergency problem, or only to provide information.

Messages have one of the following severities:

- **Alert:** Alert messages indicate that a next-higher level event might occur if you do not take some action.
You must take an action against alert messages within 24 hours.
- **Emergency:** Emergency messages are displayed when a disruption has occurred.
You must take an action against emergency messages immediately.
- **Error:** Error conditions indicate what might happen if you ignore.
- **Notice:** Normal but significant condition.
- **Info:** Informational message provides details about the issue, which you can ignore.
- **Debug:** Debug-level messages provide instructions you should perform.

If your internal support organization receives AutoSupport messages through email, the severity appears in the subject line of the email message.

Requirements for using AutoSupport

You should use HTTPS for delivery of AutoSupport messages to provide the best security and to support all of the latest AutoSupport features. Although AutoSupport supports HTTP and SMTP for delivery of AutoSupport messages, HTTPS is recommended.

Supported protocols

All of these protocols run on IPv4 or IPv6, based on the address family to which the name resolves.

Protocol and port	Description
HTTPS on port 443	<p>This is the default protocol. You should use this whenever possible.</p> <p>This protocol supports AutoSupport OnDemand and uploads of large files. The certificate from the remote server is validated against the root certificate, unless you disable validation.</p> <p>The delivery uses an HTTP PUT request. With PUT, if the request fails during transmission, the request restarts where it stopped. If the server receiving the request does not support PUT, the delivery uses an HTTP POST request.</p>
HTTP on port 80	<p>This protocol is preferred over SMTP.</p> <p>This protocol supports uploads of large files, but not AutoSupport OnDemand.</p> <p>The delivery uses an HTTP PUT request. With PUT, if the request fails during transmission, the request restarts where it stopped. If the server receiving the request does not support PUT, the delivery uses an HTTP POST request.</p>
SMTP on port 25 or another port	<p>You should use this protocol only if the network connection does not allow HTTPS or HTTP.</p> <p>The default port value is 25, but you can configure AutoSupport to use a different port.</p> <p>Keep the following limitations in mind when using SMTP:</p> <ul style="list-style-type: none"> • AutoSupport OnDemand and uploads of large files are not supported. • Data is not encrypted. SMTP sends data in clear text, making text in the AutoSupport message easy to intercept and read. • Limitations on message length and line length can be introduced.

If you configure AutoSupport with specific email addresses for your internal support organization, or a support partner organization, those messages are always sent by SMTP.

For example, if you use the recommended protocol to send messages to technical support and you also want to send messages to your internal support organization, your messages will be transported using both HTTPS and SMTP, respectively.

AutoSupport limits the maximum file size for each protocol. The default setting for HTTP and HTTPS transfers is 25 MB. The default setting for SMTP transfers is 5 MB. If the size of the AutoSupport message exceeds the configured limit, AutoSupport delivers as much of the message as possible. You can edit the maximum size by modifying AutoSupport configuration. See the `system node autosupport modify man` page for more information.

Note: AutoSupport automatically overrides the maximum file size limit for the HTTPS and HTTP protocols when you generate and send AutoSupport messages that upload core dump or performance archive files to the NetApp support site or a specified URI. The automatic override applies only when you upload files by using the `system node autosupport invoke-core-upload` or the `system node autosupport invoke-performance-archive` commands.

Configuration requirements

Depending on your network configuration, use of HTTP or HTTPS protocols may require additional configuration of a proxy URL. If you use HTTP or HTTPS to send AutoSupport messages to technical support and you have a proxy, you must identify the URL for that proxy. If the proxy uses a

port other than the default port, which is 3128, you can specify the port for that proxy. You can also specify a user name and password for proxy authentication.

If you use SMTP to send AutoSupport messages either to your internal support organization or to technical support, you must configure an external mail server. The storage system does not function as a mail server; it requires an external mail server at your site to send mail. The mail server must be a host that listens on the SMTP port (25) or another port, and it must be configured to send and receive 8-bit Multipurpose Internet Mail Extensions (MIME) encoding. Example mail hosts include a UNIX host running an SMTP server such as the sendmail program and a Windows server running the Microsoft Exchange server. You can have one or more mail hosts.

Setting up AutoSupport

You can control whether and how AutoSupport information is sent to technical support and your internal support organization, and then test that the configuration is correct.

About this task

In ONTAP 9.5 and later releases, you can enable AutoSupport and modify its configuration on all nodes of the cluster simultaneously. When a new node joins the cluster, the node inherits the AutoSupport cluster configuration automatically. You do not have to update the configuration on each node separately.

Note: Starting with ONTAP 9.5, the scope of the `system node autosupport modify` command is cluster-wide. The AutoSupport configuration is modified on all nodes in the cluster, even when the `-node` option is specified. The option is ignored, but it has been retained for CLI backward compatibility.

In ONTAP 9.4 and earlier releases, the scope of the "system node autosupport modify" command is specific to the node. The AutoSupport configuration should be modified on each node in your cluster.

By default, AutoSupport is enabled on each node to send messages to technical support by using the HTTPS transport protocol.

Steps

1. Ensure that AutoSupport is enabled by using the following command:

Example

```
system node autosupport modify -state enable
```

2. If you want technical support to receive AutoSupport messages, use the following command:

Example

```
system node autosupport modify -support enable
```

You must enable this option if you want to enable AutoSupport to work with AutoSupport OnDemand or if you want to upload large files, such as core dump and performance archive files, to technical support or a specified URL.

3. If technical support is enabled to receive AutoSupport messages, specify which transport protocol to use for the messages.

You can choose from the following options:

If you want to...	Then set the following parameters of the <code>system node autosupport modify</code> command...
Use the default HTTPS protocol	<p>a. Set <code>-transport</code> to https.</p> <p>b. If you use a proxy, set <code>-proxy-url</code> to the URL of your proxy.</p> <p>This configuration supports communication with AutoSupport OnDemand and uploads of large files.</p>
Use HTTP that is preferred over SMTP	<p>a. Set <code>-transport</code> to http.</p> <p>b. If you use a proxy, set <code>-proxy-url</code> to the URL of your proxy.</p> <p>This configuration supports uploads of large files, but not AutoSupport OnDemand.</p>
Use SMTP	<p>Set <code>-transport</code> to smtp.</p> <p>This configuration does not support AutoSupport OnDemand or uploads of large files.</p>

4. If you want your internal support organization or a support partner to receive AutoSupport messages, perform the following actions:
 - 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
<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.
5. If you are sending messages to your internal support organization or you chose SMTP transport for messages to technical support, configure SMTP by setting the following parameters of the `system node autosupport modify` command:
 - Set `-mail-hosts` to one or more mail hosts, separated by commas. You can set a maximum of five. You can configure a port value for each mail host by specifying a colon and port number after the mail host name: for example, `mymailhost.example.com:5678`, where 5678 is the port for the mail host.
 - Set `-from` to the email address that sends the AutoSupport message.

6. Configure DNS.
7. Optional: Add command options if you want to change specific settings:

If you want to do this...	Then set the following parameters of the <code>system node autosupport modify</code> command...
Hide private data by removing, masking, or encoding sensitive data in the messages	Set <code>-remove-private-data</code> to true . If you change from false to true , all AutoSupport history and all associated files are deleted.
Stop sending performance data in periodic AutoSupport messages	Set <code>-perf</code> to false .

8. Check the overall configuration by using the `system node autosupport show` command with the `-node` parameter.
9. Verify the AutoSupport operation by using the `system node autosupport check show` command.
If any problems are reported, use the `system node autosupport check show-details` command to view more information.
10. Test that AutoSupport messages are being sent and received:
 - a. Use the `system node autosupport invoke` command with the `-type` parameter set to **test**.

Example

```
cluster1::> system node autosupport invoke -type test -node node1
```

- b. Confirm that NetApp is receiving your AutoSupport messages:
`system node autosupport history -node local`
The status of the latest outgoing AutoSupport message should eventually change to `sent-successful` for all appropriate protocol destinations.
- c. Optional: Confirm that the AutoSupport message is being sent to your internal support organization or to your support partner by checking the email of any address that you configured for the `-to`, `-noteto`, or `-partner-address` parameters of the `system node autosupport modify` command.

Related tasks

[Troubleshooting AutoSupport when messages are not received](#) on page 127

Related information

[ONTAP 9 commands](#)

[Network and LIF management](#)

Uploading core dump files

When a core dump file is saved, an event message is generated. If the AutoSupport service is enabled and configured to send messages to NetApp support, an AutoSupport message is transmitted, and an automated email acknowledgement is sent to you.

Before you begin

- You must have set up AutoSupport with the following settings:
 - AutoSupport is enabled on the node.
 - AutoSupport is configured to send messages to technical support.
 - AutoSupport is configured to use the HTTP or HTTPS transport protocol.
The SMTP transport protocol is not supported when sending messages that include large files, such as core dump files.

About this task

You can also upload the core dump file through the AutoSupport service over HTTPS by using the `system node autosupport invoke-core-upload` command, if requested by NetApp support.

[NetApp Knowledgebase Answer 1029870: How to upload a file to NetApp](#)

Steps

1. View the core dump files for a node by using the `system node coredump show` command.

Example

In the following example, core dump files are displayed for the local node:

```
cluster1::> system node coredump show -node local
Node:Type Core Name Saved Panic Time
-----
node:kernel
core.4073000068.2013-09-11.15_05_01.nz true 9/11/2013 15:05:01
```

2. Generate an AutoSupport message and upload a core dump file by using the `system node autosupport invoke-core-upload` command.

Example

In the following example, an AutoSupport message is generated and sent to the default location, which is technical support, and the core dump file is uploaded to the default location, which is the NetApp support site:

```
cluster1::> system node autosupport invoke-core-upload -core-filename
core.4073000068.2013-09-11.15_05_01.nz -node local
```

In the following example, an AutoSupport message is generated and sent to the location specified in the URI, and the core dump file is uploaded to the URI:

```
cluster1::> system node autosupport invoke-core-upload -uri https://
files.company.com -core-filename
core.4073000068.2013-09-11.15_05_01.nz -node local
```

Related tasks

[Setting up AutoSupport](#) on page 119

Uploading performance archive files

You can generate and send an AutoSupport message that contains a performance archive. By default, NetApp technical support receives the AutoSupport message, and the performance archive is uploaded to the NetApp support site. You can specify an alternate destination for the message and upload.

Before you begin

- You must have set up AutoSupport with the following settings:
 - AutoSupport is enabled on the node.
 - AutoSupport is configured to send messages to technical support.
 - AutoSupport is configured to use the HTTP or HTTPS transport protocol.
The SMTP transport protocol is not supported when sending messages that include large files, such as performance archive files.

About this task

You must specify a start date for the performance archive data that you want to upload. Most storage systems retain performance archives for two weeks, enabling you to specify a start date up to two weeks ago. For example, if today is January 15, you can specify a start date of January 2. If you need help identifying the start date for the performance archive that you want to upload, contact technical support.

Step

1. Generate an AutoSupport message and upload the performance archive file by using the `system node autosupport invoke-performance-archive` command.

Example

In the following example, 4 hours of performance archive files from January 12, 2015 are added to an AutoSupport message and uploaded to the default location, which is the NetApp support site:

```
cluster1::> system node autosupport invoke-performance-archive -node
local -start-date 1/12/2015 13:42:09 -duration 4h
```

In the following example, 4 hours of performance archive files from January 12, 2015 are added to an AutoSupport message and uploaded to the location specified by the URI:

```
cluster1::> system node autosupport invoke-core-upload -uri https://
files.company.com -core-filename
core.4073000068.2013-09-11.15_05_01.nz -node local
```

Related tasks

[Setting up AutoSupport](#) on page 119

Getting AutoSupport message descriptions

The descriptions of the AutoSupport messages that you receive are available through the online AutoSupport Message Matrices page.

Steps

1. Go to the AutoSupport Message Matrices page.
support.netapp.com/NOW/knowledge/docs/olio/autosupport/matrices/
2. On the AutoSupport Message Matrices page under Select a Release, select your version of ONTAP, and then click **View Matrix**.

The Syslog Translator page appears with all AutoSupport message descriptions listed alphabetically by subject line.

Commands for managing AutoSupport

You use the `system node autosupport` commands to change or view AutoSupport configuration, display information about previous AutoSupport messages, and send, resend or cancel an AutoSupport message.

Configure AutoSupport

If you want to...	Use this command...
Control whether any AutoSupport messages are sent	<code>system node autosupport modify with the -state parameter</code>
Control whether AutoSupport messages are sent to technical support	<code>system node autosupport modify with the -support parameter</code>
Set up AutoSupport or modify the configuration of AutoSupport	<code>system node autosupport modify</code>
Enable and disable AutoSupport messages to your internal support organization for individual trigger events, and specify additional subsystem reports to include in messages sent in response to individual trigger events	<code>system node autosupport trigger modify</code>

Display information about the AutoSupport configuration

If you want to...	Use this command...
Display the AutoSupport configuration	<code>system node autosupport show with the -node parameter</code>
View a summary of all addresses and URLs that receive AutoSupport messages	<code>system node autosupport destinations show</code>
Display which AutoSupport messages are sent to your internal support organization for individual trigger events	<code>system node autosupport trigger show</code>
Display status of AutoSupport configuration as well as delivery to various destinations	<code>system node autosupport check show</code>

If you want to...	Use this command...
Display detailed status of AutoSupport configuration as well as delivery to various destinations	<code>system node autosupport check show-details</code>

Display information about past AutoSupport messages

If you want to...	Use this command...
Display information about one or more of the 50 most recent AutoSupport messages	<code>system node autosupport history show</code>
Display information about recent AutoSupport messages generated to upload core dump or performance archive files to the technical support site or a specified URI	<code>system node autosupport history show-upload-details</code>
View the information in the AutoSupport messages including the name and size of each file collected for the message along with any errors	<code>system node autosupport manifest show</code>

Send, resend, or cancel AutoSupport messages

If you want to...	Use this command...
Retransmit a locally stored AutoSupport message, identified by its AutoSupport sequence number Note: If you retransmit an AutoSupport message, and if support already received that message, the support system will not create a duplicate case. If, on the other hand, support did not receive that message, then the AutoSupport system will analyze the message and create a case, if necessary.	<code>system node autosupport history retransmit</code>
Generate and send an AutoSupport message—for example, for testing purposes	<code>system node autosupport invoke</code> Note: Use the <code>-force</code> parameter to send a message even if AutoSupport is disabled. Use the <code>-uri</code> parameter to send the message to the destination you specify instead of the configured destination.
Cancel an AutoSupport message	<code>system node autosupport history cancel</code>

Related information

[ONTAP 9 commands](#)

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 Active IQ (formerly known as My AutoSupport) portal.

AutoSupport case suppression during scheduled maintenance windows

AutoSupport case suppression enables you to stop unnecessary cases from being created by AutoSupport messages that are triggered during scheduled maintenance windows.

To suppress AutoSupport cases, you must manually invoke an AutoSupport message with a specially formatted text string: `MAINT=xh`. `x` is the duration of the maintenance window in units of hours.

Related information

[NetApp KB Article: 1010449: How to suppress automatic case creation during scheduled maintenance windows](#)

What Active IQ is

Active IQ (formerly known as My AutoSupport) is a web-based application that works in conjunction with AutoSupport. It presents information that enables you to easily analyze data to model and optimize your storage infrastructure.

Active IQ is a web-based application hosted on the NetApp Support Site at mysupport.netapp.com that you can access using a browser. Your system must have AutoSupport enabled and configured so that it sends data back to NetApp.

[NetApp Active IQ](#)

Active IQ 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 OnDemand service chose to ignore it.

3. Perform one of the following actions:

For this status	Do this
initializing or collection-failed	Contact technical support because AutoSupport cannot generate the message.
ignore, re-queued, or transmission failed	Check that destinations are correctly configured for SMTP, HTTP, or HTTPS because AutoSupport cannot deliver the message.

Related tasks

Troubleshooting AutoSupport message delivery over SMTP on page 129

Troubleshooting AutoSupport message delivery over HTTP or HTTPS on page 128

Troubleshooting AutoSupport message delivery over HTTP or HTTPS

If the system does not send the expected AutoSupport message and you are using HTTP or HTTPS, you can check a number of settings to resolve the problem.

Before you begin

You should have confirmed basic network connectivity and DNS lookup:

- Your node management LIF must be up for operational and administrative status.
- You must be able to ping a functioning host on the same subnet from the cluster management LIF (not a LIF on any of the nodes).
- You must be able to ping a functioning host outside the subnet from the cluster management LIF.
- You must be able to ping a functioning host outside the subnet from the cluster management LIF using the name of the host (not the IP address).

About this task

These steps are for cases when you have determined that AutoSupport can generate the message, but cannot deliver the message over HTTP or HTTPS.

If you encounter errors or cannot complete a step in this procedure, determine and address the problem before proceeding to the next step.

Steps

1. Verify the status of the node management LIF:

```
network interface show -home-node local -role node-mgmt -fields
vserver,lif,status-oper,status-admin,address,role
```

The `status-oper` and `status-admin` fields should return `up`.

2. Record the SVM name, the LIF name, and the LIF IP address for later use.
3. Ensure that DNS is enabled and configured correctly:

```
vserver services name-service dns show
```

4. Address any errors returned by the AutoSupport message:

```
system node autosupport history show -node * -fields node,seq-
num,destination,last-update,status,error
```

If the error refers to a problem with the digital certificate, contact technical support.

5. Confirm that the cluster can access both the servers it needs and the Internet successfully:

- a. `network traceroute -node local -destination network_management_LIF`

- b. `network traceroute -node local -destination support.netapp.com`

- c. `system node autosupport show -fields proxy-url`

- d. `network traceroute -node local -destination proxy_url`

If any of these routes is not functioning, try the same route from a functioning host on the same subnet as the cluster, using the “traceroute” or “tracert” utility found on most third-party network

clients. This assists you in determining whether the issue is in your network configuration or your cluster configuration.

6. If you are using HTTPS for your AutoSupport transport protocol, ensure that HTTPS traffic can exit your network:
 - a. Configure a web client on the same subnet as the cluster management LIF.

Ensure that all configuration parameters are the same values as for the AutoSupport configuration, including using the same proxy server, user name, password, and port.
 - b. Access `https://support.netapp.com` with the web client.

The access should be successful. If not, ensure that all firewalls are configured correctly to allow HTTPS and DNS traffic, and that the proxy server is configured correctly.

Related tasks

[Troubleshooting AutoSupport when messages are not received](#) on page 127

Related information

[Network and LIF management](#)

Troubleshooting AutoSupport message delivery over SMTP

If the system cannot deliver AutoSupport messages over SMTP, you can check a number of settings to resolve the problem.

Before you begin

You should have confirmed basic network connectivity and DNS lookup:

- Your node management LIF must be up for operational and administrative status.
- You must be able to ping a functioning host on the same subnet from the cluster management LIF (not a LIF on any of the nodes).
- You must be able to ping a functioning host outside the subnet from the cluster management LIF.
- You must be able to ping a functioning host outside the subnet from the cluster management LIF using the name of the host (not the IP address).

About this task

These steps are for cases when you have determined that AutoSupport can generate the message, but cannot deliver the message over SMTP.

If you encounter errors or cannot complete a step in this procedure, determine and address the problem before proceeding to the next step.

All commands are entered at the ONTAP command-line interface, unless otherwise specified.

Steps

1. Verify the status of the node management LIF:


```
network interface show -home-node local -role node-mgmt -fields
vserver,lif,status-oper,status-admin,address,role
```

The `status-oper` and `status-admin` fields should return `up`.
2. Record the SVM name, the LIF name, and the LIF IP address for later use.
3. Ensure that DNS is enabled and configured correctly:

```
vserver services name-service dns show
```

4. Display all of the servers configured to be used by AutoSupport:

```
system node autosupport show -fields mail-hosts
```

Record all server names displayed.

5. For each server displayed by the previous step, and `support.netapp.com`, ensure that the server or URL can be reached by the node:

```
network traceroute -node local -destination server_name
```

If any of these routes is not functioning, try the same route from a functioning host on the same subnet as the cluster, using the “traceroute” or “tracert” utility found on most third-party network clients. This assists you in determining whether the issue is in your network configuration or your cluster configuration.

6. Log in to the host designated as the mail host, and ensure that it can serve SMTP requests:

```
netstat -aAn|grep 25
```

25 is the listener SMTP port number.

A message similar to the following text is displayed:

```
ff64878c tcp          0          0 *.25    *.*     LISTEN.
```

7. From some other host, open a Telnet session with the SMTP port of the mail host:

```
telnet mailhost 25
```

A message similar to the following text is displayed:

```
220 filer.yourco.com Sendmail 4.1/SMI-4.1 ready at Thu, 30 Nov 2014
10:49:04 PST
```

8. At the telnet prompt, ensure that a message can be relayed from your mail host:

```
HELO domain_name
```

```
MAIL FROM: your_email_address
```

```
RCPT TO: autosupport@netapp.com
```

domain_name is the domain name of your network.

If an error is returned saying that relaying is denied, relaying is not enabled on the mail host. Contact your system administrator.

9. At the telnet prompt, send a test message:

```
DATA
```

```
SUBJECT: TESTING
```

```
THIS IS A TEST
```

```
.
```

Note: Ensure that you enter the last period (.) on a line by itself. The period indicates to the mail host that the message is complete.

If an error is returned, your mail host is not configured correctly. Contact your system administrator.

10. From the ONTAP command-line interface, send an AutoSupport test message to a trusted email address that you have access to:

```
system node autosupport invoke -node local -type test
```

- Find the sequence number of the attempt:

```
system node autosupport history show -node local -destination smtp
```

Find the sequence number for your attempt based on the timestamp. It is probably the most recent attempt.

- Display the error for your test message attempt:

```
system node autosupport history show -node local -seq-num seq_num -
fields error
```

If the error displayed is `Login denied`, your SMTP server is not accepting send requests from the cluster management LIF. If you do not want to change to using HTTPS as your transport protocol, contact your site network administrator to configure the SMTP gateways to address this issue.

If this test succeeds but the same message sent to `mailto:autosupport@netapp.com` does not, ensure that SMTP relay is enabled on all of your SMTP mail hosts, or use HTTPS as a transport protocol.

If even the message to the locally administered email account does not succeed, confirm that your SMTP servers are configured to forward attachments with both of these characteristics:

- The “7z” suffix
- The “application/x-7x-compressed” MIME type.

Related tasks

[Troubleshooting AutoSupport when messages are not received](#) on page 127

Related information

[Network and LIF management](#)

Troubleshooting the AutoSupport subsystem

The `system node check show` commands can be used to verify and troubleshoot any issues related to the AutoSupport configuration and delivery.

Step

- Use the following commands to display the status of the AutoSupport subsystem.

Use this command...	To do this...
<code>system node autosupport check show</code>	Display overall status of the AutoSupport subsystem, such as the status of AutoSupport HTTP or HTTPS destination, AutoSupport SMTP destinations, AutoSupport OnDemand Server, and AutoSupport configuration
<code>system node autosupport check show-details</code>	Display detailed status of the AutoSupport subsystem, such as detailed descriptions of errors and the corrective actions

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.

For details on how ONTAP supports cluster switches for system health monitoring in your cluster, you can refer to the *Hardware Universe*.

[Supported switches in the Hardware Universe](#)

For details on the causes of Cluster Switch Health Monitor (CSHM) AutoSupport messages, and the necessary actions required to resolve these alerts, you can refer to the *Knowledge Base*.

[Cluster Switch Health Monitor \(CSHM\) AutoSupport messages](#)

How health monitoring works

Individual health monitors have a set of policies that trigger alerts when certain conditions occur. Understanding how health monitoring works can help you respond to problems and control future alerts.

Health monitoring consists of the following components:

- Individual health monitors for specific subsystems, each of which has its own health status
For example, the Storage subsystem has a node connectivity health monitor.
- An overall system health monitor that consolidates the health status of the individual health monitors
A degraded status in any single subsystem results in a degraded status for the entire system. If no subsystems have alerts, the overall system status is OK.

Each health monitor is made up of the following key elements:

- Alerts that the health monitor can potentially raise
Each alert has a definition, which includes details such as the severity of the alert and its probable cause.
- Health policies that identify when each alert is triggered
Each health policy has a rule expression, which is the exact condition or change that triggers the alert.

A health monitor continuously monitors and validates the resources in its subsystem for condition or state changes. When a condition or state change matches a rule expression in a health policy, the health monitor raises an alert. An alert causes the subsystem's health status and the overall system health status to become degraded.

Ways to respond to system health alerts

When a system health alert occurs, you can acknowledge it, learn more about it, repair the underlying condition, and prevent it from occurring again.

When a health monitor raises an alert, you can respond in any of the following ways:

- Get information about the alert, which includes the affected resource, alert severity, probable cause, possible effect, and corrective actions.
- Get detailed information about the alert, such as the time when the alert was raised and whether anyone else has acknowledged the alert already.
- Get health-related information about the state of the affected resource or subsystem, such as a specific shelf or disk.
- Acknowledge the alert to indicate that someone is working on the problem, and identify yourself as the “Acknowledger.”
- Resolve the problem by taking the corrective actions provided in the alert, such as fixing cabling to resolve a connectivity problem.
- Delete the alert, if the system did not automatically clear it.
- Suppress an alert to prevent it from affecting the health status of a subsystem. Suppressing is useful when you understand a problem. After you suppress an alert, it can still occur, but the subsystem health displays as “ok-with-suppressed.” when the suppressed alert occurs.

System health alert customization

You can control which alerts a health monitor generates by enabling and disabling the system health policies that define when alerts are triggered. This enables you to customize the health monitoring system for your particular environment.

You can learn the name of a policy either by displaying detailed information about a generated alert or by displaying policy definitions for a specific health monitor, node, or alert ID.

Disabling health policies is different from suppressing alerts. When you suppress an alert, it does not affect the subsystem's health status, but the alert can still occur.

If you disable a policy, the condition or state that is defined in its policy rule expression no longer triggers an alert.

Example of an alert that you want to disable

For example, suppose an alert occurs that is not useful to you. You use the `system health alert show -instance` command to obtain the Policy ID for the alert. You use the policy ID in the `system health policy definition show` command to view information about the policy. After reviewing the rule expression and other information about the policy, you decide to disable the policy. You use the `system health policy definition modify` command to disable the policy.

How health alerts trigger AutoSupport messages and events

System health alerts trigger AutoSupport messages and events in the Event Management System (EMS), enabling you to monitor the health of the system using AutoSupport messages and the EMS in addition to using the health monitoring system directly.

Your system sends an AutoSupport message within five minutes of an alert. The AutoSupport message includes all alerts generated since the previous AutoSupport message, except for alerts that duplicate an alert for the same resource and probable cause within the previous week.

Some alerts do not trigger AutoSupport messages. An alert does not trigger an AutoSupport message if its health policy disables the sending of AutoSupport messages. For example, a health policy might disable AutoSupport messages by default because AutoSupport already generates a message when the problem occurs. You can configure policies to not trigger AutoSupport messages by using the `system health policy definition modify` command.

You can view a list of all of the alert-triggered AutoSupport messages sent in the previous week using the `system health autosupport trigger history show` command.

Alerts also trigger the generation of events to the EMS. An event is generated each time an alert is created and each time an alert is cleared.

Available cluster health monitors

There are several health monitors that monitor different parts of a cluster. Health monitors help you to recover from errors within ONTAP systems by detecting events, sending alerts to you, and deleting events as they clear.

Health monitor name (identifier)	Subsystem name (identifier)	Purpose
Cluster switch (cluster-switch)	Switch (Switch-Health)	Monitors cluster network switches and management network switches for temperature, utilization, interface configuration, redundancy (cluster network switches only), and fan and power supply operation. The cluster switch health monitor communicates with switches through SNMP. SNMPv2c is the default setting. Note: Beginning with ONTAP 9.2, this monitor can detect and report when a cluster switch has rebooted since the last polling period.
MetroCluster Fabric	Switch	Monitors the MetroCluster configuration back-end fabric topology and detects misconfigurations such as incorrect cabling and zoning, and ISL failures.
MetroCluster Health	Interconnect, RAID, and storage	Monitors FC-VI adapters, FC initiator adapters, left-behind aggregates and disks, and inter-cluster ports
Node connectivity (node-connect)	CIFS nondisruptive operations (CIFS-NDO)	Monitors SMB connections for nondisruptive operations to Hyper-V applications.
	Storage (SAS-connect)	Monitors shelves, disks, and adapters at the node level for appropriate paths and connections.
System	not applicable	Aggregates information from other health monitors.
System connectivity (system-connect)	Storage (SAS-connect)	Monitors shelves at the cluster level for appropriate paths to two HA clustered nodes.

Receiving system health alerts automatically

You can manually view system health alerts by using the `system health alert show` command. However, you should subscribe to specific Event Management System (EMS) messages to automatically receive notifications when a health monitor generates an alert.

About this task

The following procedure shows you how to set up notifications for all `hm.alert.raised` messages and all `hm.alert.cleared` messages.

All `hm.alert.raised` messages and all `hm.alert.cleared` messages include an SNMP trap. The names of the SNMP traps are `HealthMonitorAlertRaised` and `HealthMonitorAlertCleared`. For information about SNMP traps, see the *Network Management Guide*.

Steps

1. Use the `event destination create` command to define the destination to which you want to send the EMS messages.

Example

```
cluster1::> event destination create -name health_alerts -mail
admin@example.com
```

2. Use the `event route add-destinations` command to route the `hm.alert.raised` message and the `hm.alert.cleared` message to a destination.

Example

```
cluster1::> event route add-destinations -messagename hm.alert* -
destinations health_alerts
```

Related information

[Network and LIF management](#)

Responding to degraded system health

When your system's health status is degraded, you can show alerts, read about the probable cause and corrective actions, show information about the degraded subsystem, and resolve the problem. Suppressed alerts are also shown so that you can modify them and see whether they have been acknowledged.

About this task

You can discover that an alert was generated by viewing an AutoSupport message or an EMS event, or by using the `system health` commands.

Steps

1. Use the `system health alert show` command to view the alerts that are compromising the system's health.
2. Read the alert's probable cause, possible effect, and corrective actions to determine whether you can resolve the problem or need more information.
3. If you need more information, use the `system health alert show -instance` command to view additional information available for the alert.
4. Use the `system health alert modify` command with the `-acknowledge` parameter to indicate that you are working on a specific alert.
5. Take corrective action to resolve the problem as described by the `Corrective Actions` field in the alert.

The corrective actions might include rebooting the system.

When the problem is resolved, the alert is automatically cleared. If the subsystem has no other alerts, the health of the subsystem changes to OK. If the health of all subsystems is OK, the overall system health status changes to OK.

6. Use the `system health status show` command to confirm that the system health status is OK.

If the system health status is not OK, repeat this procedure.

Example of responding to degraded system health

By reviewing a specific example of degraded system health caused by a shelf that lacks two paths to a node, you can see what the CLI displays when you respond to an alert.

After starting ONTAP, you check the system health and you discover that the status is degraded:

```
cluster1::>system health status show
      Status
      -----
      degraded
```

You show alerts to find out where the problem is, and see that shelf 2 does not have two paths to node1:

```
cluster1::>system health alert show
      Node: node1
      Resource: Shelf ID 2
      Severity: Major
      Indication Time: Mon Nov 10 16:48:12 2013
      Probable Cause: Disk shelf 2 does not have two paths to controller
                     node1.
      Possible Effect: Access to disk shelf 2 via controller node1 will be
                     lost with a single hardware component failure (e.g.
                     cable, HBA, or IOM failure).
      Corrective Actions: 1. Halt controller node1 and all controllers attached to disk shelf 2.
                        2. Connect disk shelf 2 to controller node1 via two paths following the
                           rules in the Universal SAS and ACP Cabling Guide.
                        3. Reboot the halted controllers.
                        4. Contact support personnel if the alert persists.
```

You display details about the alert to get more information, including the alert ID:

```
cluster1::>system health alert show -monitor node-connect -alert-id DualPathToDiskShelf_Alert
-instance
      Node: node1
      Monitor: node-connect
      Alert ID: DualPathToDiskShelf_Alert
      Alerting Resource: 50:05:0c:c1:02:00:0f:02
      Subsystem: SAS-connect
      Indication Time: Mon Mar 21 10:26:38 2011
      Perceived Severity: Major
      Probable Cause: Connection_establishment_error
      Description: Disk shelf 2 does not have two paths to controller node1.
      Corrective Actions: 1. Halt controller node1 and all controllers attached to disk shelf 2.
                        2. Connect disk shelf 2 to controller node1 via two paths following
                           the rules in the Universal SAS and ACP Cabling Guide.
                        3. Reboot the halted controllers.
                        4. Contact support personnel if the alert persists.
      Possible Effect: Access to disk shelf 2 via controller node1 will be lost with a single
                     hardware component failure (e.g. cable, HBA, or IOM failure).
      Acknowledge: false
      Suppress: false
      Policy: DualPathToDiskShelf_Policy
      Acknowledger: -
      Suppressor: -
      Additional Information: Shelf uuid: 50:05:0c:c1:02:00:0f:02
                           Shelf id: 2
                           Shelf Name: 4d.shelf2
                           Number of Paths: 1
                           Number of Disks: 6
                           Adapter connected to IOMA:
                           Adapter connected to IOMB: 4d
      Alerting Resource Name: Shelf ID 2
```

You acknowledge the alert to indicate that you are working on it.


```
cluster1::>system health alert modify -node node1 -alert-id DualPathToDiskShelf_Alert -
acknowledge true
```

You fix the cabling between shelf 2 and node1, and then reboot the system. Then you check system health again, and see that the status is OK:

```
cluster1::>system health status show
Status
-----
OK
```

Configuring discovery of cluster and management network switches

The cluster switch health monitor automatically attempts to discover your cluster and management network switches using the Cisco Discovery Protocol (CDP). You must configure the health monitor if it cannot automatically discover a switch or if you do not want to use CDP for automatic discovery.

About this task

The `system cluster-switch show` command lists the switches that the health monitor discovered. If you do not see a switch that you expected to see in that list, then the health monitor cannot automatically discover it.

Steps

1. If you want to use CDP for automatic discovery, do the following; otherwise, go to step 2 on page 137:
 - a. Ensure that the Cisco Discovery Protocol (CDP) is enabled on your switches.
Refer to your switch documentation for instructions.
 - b. Run the following command on each node in the cluster to verify whether CDP is enabled or disabled:


```
run -node node_name -command options cdpd.enable
```

 If CDP is enabled, go to step d. If CDP is disabled, go to step c.
 - c. Run the following command to enable CDP:


```
run -node node_name -command options cdpd.enable on
```

 Wait five minutes before you go to the next step.
 - d. Use the `system cluster-switch show` command to verify whether ONTAP can now automatically discover the switches.
2. If the health monitor cannot automatically discover a switch, use the `system cluster-switch create` command to configure discovery of the switch:

Example

```
cluster1::> system cluster-switch create -device switch1 -address
192.0.2.250 -snmp-version SNMPv2c -community cshml! -model NX5020 -
type cluster-network
```

Wait five minutes before you go to the next step.

3. Use the `system cluster-switch show` command to verify that 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 ONTAP cannot monitor the switch. 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 NetApp Support Site.

[Cisco Ethernet Switch page](#)

The community string in the switch's RCF must match the community string that the health monitor is configured to use. By default, the health monitor uses the community string `cshml!`.

If you need to change information about a switch that the cluster monitors, you can modify the community string that the health monitor uses by using the `system health cluster-switch modify` command.

3. Verify that the switch's management port is connected to the management network.

This connection is required to perform SNMP queries.

Related tasks

[Configuring discovery of cluster and management network switches](#) on page 137

Commands for monitoring the health of your system

You can use the `system health` commands to display information about the health of system resources, to respond to alerts, and to configure future alerts. Using the CLI commands enables you to view in-depth information about how health monitoring is configured. The man pages for the commands contain more information.

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 node connectivity

If you want to...	Use this command...
Display details about connectivity from the node to the storage shelf, including port information, HBA port speed, I/O throughput, and the rate of I/O operations per second	<code>storage shelf show -connectivity</code> Use the <code>-instance</code> parameter to display detailed information about each shelf.
Display information about drives and array LUNs, including the usable space, shelf and bay numbers, and owning node name	<code>storage disk show</code> Use the <code>-instance</code> parameter to display detailed information about each drive.
Display detailed information about storage shelf ports, including port type, speed, and status	<code>storage port show</code> Use the <code>-instance</code> parameter to display detailed information about each adapter.

Managing the discovery of cluster and management network switches

If you want to...	Use this command...
Display the switches that the cluster monitors	<code>system cluster-switch show</code>
Display the switches that the cluster currently monitors, including switches that you deleted (shown in the Reason column in the command output), and configuration information that you need for network access to the cluster and management network switches This command is available at the advanced privilege level.	<code>system cluster-switch show-all</code>
Configure discovery of an undiscovered switch	<code>system cluster-switch create</code>
Modify information about a switch that the cluster monitors (for example, device name, IP address, SNMP version, and community string)	<code>system cluster-switch modify</code>
Disable monitoring of a switch	<code>system cluster-switch modify -disable-monitoring</code>
Display the interval in which the health monitor polls switches to gather information	<code>system cluster-switch polling-interval show</code>
Modify the interval in which the health monitor polls switches to gather information This command is available at the advanced privilege level.	<code>system cluster-switch polling-interval modify</code>
Disable discovery and monitoring of a switch and delete switch configuration information	<code>system cluster-switch delete</code>
Permanently remove the switch configuration information which is stored in the database (doing so reenables automatic discovery of the switch)	<code>system cluster-switch delete -force</code>
Enable automatic logging to send with AutoSupport messages This command is available at the advanced privilege level.	<code>system cluster-switch log</code>

Responding to generated alerts

If you want to...	Use this command...
Display information about generated alerts, such as the resource and node where the alert was triggered, and the alert's severity and probable cause	<code>system health alert show</code>
Display information about each generated alert	<code>system health alert show -instance</code>
Indicate that someone is working on an alert	<code>system health alert modify</code>
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 whether an alert triggered an AutoSupport message	<code>system health autosupport trigger history show</code>

Configuring future alerts

If you want to...	Use this command...
Enable or disable the policy that controls whether a specific resource state raises a specific alert	<code>system health policy definition modify</code>

Displaying information about how health monitoring is configured

If you want to...	Use this command...
Display information about health monitors, such as their nodes, names, subsystems, and status	<code>system health config show</code> Note: Use the <code>-instance</code> parameter to display detailed information about each health monitor.
Display information about the alerts that a health monitor can potentially generate	<code>system health alert definition show</code> Note: Use the <code>-instance</code> parameter to display detailed information about each alert definition.
Display information about health monitor policies, which determine when alerts are raised	<code>system health policy definition show</code> Note: Use the <code>-instance</code> parameter to display detailed information about each policy. Use other parameters to filter the list of alerts—for example, by policy status (enabled or not), health monitor, alert, and so on.

Downgrading software versions between minor releases

When downgrading a cluster from Data ONTAP 8.2.1 to Data ONTAP 8.2, ONTAP clears the history of deleted switches. It also sets the `Model` column to `OTHER` if a switch that it automatically discovers is not supported for health monitoring.

If you want to...	Use this command...
Downgrade the software version command	<code>system cluster-switch prepare-to-downgrade</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 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.

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. `
- **ONTAP APIs (`ontapi`)**
This service enables you to run 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 ONTAP System Manager, you should leave this service enabled.
- **Data ONTAP Discovery (`disco`)**
This service enables off-box management applications to discover the cluster in the network. The default setting is **enabled**.
- **Support Diagnostics (`supdiag`)**
This service controls access to a privileged environment on the system to assist problem analysis and resolution. The default setting is **disabled**. You should enable this service only when directed by technical support.
- **System Manager (`sysmgr`)**
This service controls the availability of ONTAP System Manager, which is included with ONTAP. The default setting is **enabled**. This service is supported only on the cluster.

Related concepts

[Understanding ONTAP System Manager](#) on page 29

Related tasks

[Accessing a cluster by using the ONTAP System Manager browser-based graphic interface](#) on page 28

[Accessing a node's log, core dump, and MIB files by using a web browser](#) on page 40

Managing the web protocol engine

You can configure the web protocol engine on the cluster to control whether web access is allowed and what SSL versions can be used. You can also display the configuration settings for the web protocol engine.

You can manage the web protocol engine at the cluster level in the following ways:

- You can specify whether remote clients can use HTTP or HTTPS to access web service content by using the `system services web modify` command with the `-external` parameter.

- You can specify whether SSLv3 should be used for secure web access by using the `security config modify` command with the `-supported-protocol` parameter. By default, SSLv3 is disabled. Transport Layer Security 1.0 (TLSv1.0) is enabled and it can be disabled if needed.
- You can enable Federal Information Processing Standard (FIPS) 140-2 compliance mode for cluster-wide control plane web service interfaces.

Note: By default, FIPS 140-2 compliance mode is disabled.

When FIPS 140-2 compliance mode is disabled

You can enable FIPS 140-2 compliance mode by setting the `is-fips-enabled` parameter to `true` for the `security config modify` command, and then using the `security config show` command to confirm the online status.

When FIPS 140-2 compliance mode is enabled

Both TLSv1 and SSLv3 are disabled and only TLSv1.1 and TLSv1.2 remain enabled. ONTAP prevents you from enabling both TLSv1 and SSLv3 when FIPS 140-2 compliance mode is enabled. If you enable FIPS 140-2 compliance mode and then subsequently disable it, TLSv1 and SSLv3 remain disabled, but either TLSv1.2 or both TLSv1.1 and TLSv1.2 are enabled depending on the previous configuration.

- You can display the configuration of cluster-wide security by using the `system security config show` command.

If the firewall is enabled, the firewall policy for the logical interface (LIF) to be used for web services must be set up to allow HTTP or HTTPS access.

If you use HTTPS for web service access, SSL for the cluster or storage virtual machine (SVM) that offers the web service must also be enabled, and you must provide a digital certificate for the cluster or SVM.

In MetroCluster configurations, the setting changes you make for the web protocol engine on a cluster are not replicated on the partner cluster.

Related concepts

[Managing SSL](#) on page 150

[Managing web services](#) on page 148

Related tasks

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

Related information

[Network and LIF management](#)

Commands for managing the web protocol engine

You use the `system services web` commands to manage the web protocol engine. You use the `system services firewall policy create` and `network interface modify` commands to allow web access requests to go through the firewall.

If you want to...	Use this command...
Configure the web protocol engine at the cluster level: <ul style="list-style-type: none"> • Enable or disable the web protocol engine for the cluster • Enable or disable SSLv3 for the cluster • Enable or disable FIPS 140-2 compliance for secure web services (HTTPS) 	<code>system services web modify</code>
Display the configuration of the web protocol engine at the cluster level, determine whether the web protocols are functional throughout the cluster, and display whether FIPS 140-2 compliance is enabled and online	<code>system services web show</code>
Display the configuration of the web protocol engine at the node level and the activity of web service handling for the nodes in the cluster	<code>system services web node show</code>
Create a firewall policy or add HTTP or HTTPS protocol service to an existing firewall policy to allow web access requests to go through firewall	<code>system services firewall policy create</code> Setting the <code>-service</code> parameter to <code>http</code> or <code>https</code> enables web access requests to go through firewall.
Associate a firewall policy with a LIF	<code>network interface modify</code> You can use the <code>-firewall-policy</code> parameter to modify the firewall policy of a LIF.

Related references

[Commands for managing SSL](#) on page 151

Related information

[ONTAP 9 commands](#)

Configuring SAML authentication for web services

Starting with ONTAP 9.3, you can configure multifactor authentication (MFA) for web services by using Security Assertion Markup Language (SAML) authentication. You can use SAML authentication for Service Processor Infrastructure (spi), ONTAP APIs, and ONTAP System Manager.

When you configure SAML authentication, users are authenticated by an external Identity Provider (IdP). The IdP is a third-party software such as Microsoft Active Directory Federated Services (ADFS) IdP or open source Shibboleth IdP. The ONTAP cluster acts as the SAML Service Provider (SP) host. Authentication is performed by exchanging metadata between the IdP and SP.

Related information

[Administrator authentication and RBAC](#)

Configuring SAML authentication

Starting with ONTAP 9.3, you can configure Security Assertion Markup Language (SAML) authentication for web services. When SAML authentication is configured and enabled, users are authenticated by an external Identity Provider (IdP) instead of the directory service providers such as Active Directory and LDAP.

Before you begin

- You must have configured the IdP for SAML authentication.
- You must have the IdP URI.

About this task

- SAML authentication applies only to the `http` and `ontapi` applications. The `http` and `ontapi` applications are used by the following web services: Service Processor Infrastructure, ONTAP APIs, or ONTAP System Manager.
- SAML authentication is applicable only for accessing the admin SVM.

Steps

1. Create a SAML configuration so that ONTAP can access the IdP metadata:

```
security saml-sp create -idp-uri idp_uri -sp-host ontap_host_name
```

idp_uri is the FTP or HTTP address of the IdP host from where the IdP metadata can be downloaded.

ontap_host_name is the host name or IP address of the SAML service provider host, which in this case is the ONTAP system. By default, the IP address of the cluster-management LIF is used.

You can optionally provide the ONTAP server certificate information. By default, the ONTAP web server certificate information is used.

Example

```
cluster_12::> security saml-sp create -idp-uri https://
scspr0235321001.gdl.englab.netapp.com/idp/shibboleth -verify-metadata-
server false
```

Warning: This restarts the web server. Any HTTP/S connections that are active

will be disrupted.

Do you want to continue? {y|n}: y

[Job 179] Job succeeded: Access the SAML SP metadata using the

URL:

https://10.63.56.150/saml-sp/Metadata

Configure the IdP and Data ONTAP users for the same directory server domain to ensure that users are the same for different authentication methods. See the "security login show" command for the Data ONTAP user configuration.

The URL to access the ONTAP host metadata is displayed.

2. From the IdP host, configure the IdP with the ONTAP host metadata.

For more information about configuring the IdP, see the IdP documentation.

3. Enable SAML configuration:

```
security saml-sp modify -is-enabled true
```

Any existing user that accesses the `http` or `ontapi` application is automatically configured for SAML authentication.

4. If you want to create users for the `http` or `ontapi` application after SAML is configured, specify SAML as the authentication method for the new users.

- a. Create a login method for new users with SAML authentication:

```
security login create -user-or-group-name user_name -application [http
| ontapi] -authentication-method saml -vserver svm_name
```

Example

```
cluster_12::> security login create -user-or-group-name admin1 -
application http -authentication-method saml -vserver cluster_12
```

- b. Verify that the user entry is created:

```
security login show
```

Example

```
cluster_12::> security login show

Vserver: cluster_12

User/Group Authentication Acct Second
Authentication Name Application Method Role Name Locked Method
-----
admin console password admin no none
admin http password admin no none
admin http saml admin - none
admin ontapi password admin no none
admin ontapi saml admin - none
admin service-processor password admin no none
admin ssh password admin no none
admin1 http password backup no none
admin1 http saml backup - none
```

Related information

[ONTAP 9 commands](#)

Disabling SAML authentication

You can disable SAML authentication when you want to stop authenticating web users by using an external Identity Provider (IdP). When SAML authentication is disabled, the configured directory service providers such as Active Directory and LDAP are used for authentication.

Before you begin

You must be logged in from the console.

Steps

1. Disable SAML authentication:

```
security saml-sp modify -is-enabled false
```

- Optional: If you no longer want to use SAML authentication or if you want to modify the IdP, delete the SAML configuration:

```
security saml-sp delete
```

Troubleshooting issues with SAML configuration

If configuring Security Assertion Markup Language (SAML) authentication fails, you can manually repair each node on which the SAML configuration failed and recover from the failure. During the repair process, the web server is restarted and any active HTTP connections or HTTPS connections are disrupted.

About this task

When you configure SAML authentication, ONTAP applies SAML configuration on a per-node basis. When you enable SAML authentication, ONTAP automatically tries to repair each node if there are configuration issues. If there are issues with SAML configuration on any node, you can disable SAML authentication and then reenable SAML authentication. There can be situations when SAML configuration fails to apply on one or more nodes even after you reenable SAML authentication. You can identify the node on which SAML configuration has failed and then manually repair that node.

Steps

- Log in to the advanced privilege level:
- Identify the node on which SAML configuration failed:

```
set -privilege advanced
```

```
security saml-sp status show -instance
```

Example

```
cluster_12::*> security saml-sp status show -instance

                Node: node1
                Update Status: config-success
                Database Epoch: 9
                Database Transaction Count: 997
                Error Text:
SAML Service Provider Enabled: false
                ID of SAML Config Job: 179

                Node: node2
                Update Status: config-failed
                Database Epoch: 9
                Database Transaction Count: 997
                Error Text: SAML job failed, Reason: Internal
error. Failed to receive the SAML IDP Metadata file.
SAML Service Provider Enabled: false
                ID of SAML Config Job: 180
2 entries were displayed.
```

- Repair the SAML configuration on the failed node:

```
security saml-sp repair -node node_name
```

Example

```
cluster_12::*> security saml-sp repair -node node2

Warning: This restarts the web server. Any HTTP/S connections that
are active
```

```

    will be disrupted.
Do you want to continue? {y|n}: y
[Job 181] Job is running.
[Job 181] Job success.

```

The web server is restarted and any active HTTP connections or HTTPS connections are disrupted.

4. Verify that SAML is successfully configured on all of the nodes:

```
security saml-sp status show -instance
```

Example

```

cluster_12::*> security saml-sp status show -instance

                Node: node1
        Update Status: config-success
        Database Epoch: 9
Database Transaction Count: 997
                Error Text:
SAML Service Provider Enabled: false
        ID of SAML Config Job: 179

                Node: node2
        Update Status: config-success
        Database Epoch: 9
Database Transaction Count: 997
                Error Text:
SAML Service Provider Enabled: false
        ID of SAML Config Job: 180
2 entries were displayed.

```

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 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 142

[Managing SSL](#) on page 150

Related tasks

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

Commands for managing web services

You use the `vserver services web` commands to manage the availability of web services for the cluster or a storage virtual machine (SVM). You use the `vserver services web access` commands to control a role's access to a web service.

If you want to...	Use this command...
Configure a web service for the cluster or an SVM: <ul style="list-style-type: none"> • Enable or disable a web service • Specify whether only HTTPS can be used for accessing a web service 	<code>vserver services web modify</code>
Display the configuration and availability of web services for the cluster or an SVM	<code>vserver services web show</code>
Authorize a role to access a web service on the cluster or an SVM	<code>vserver services web access create</code>
Display the roles that are authorized to access web services on the cluster or an SVM	<code>vserver services web access show</code>
Prevent a role from accessing a web service on the cluster or an SVM	<code>vserver services web access delete</code>

Related information

[ONTAP 9 commands](#)

Commands for managing mount points on the nodes

The `spi` web service automatically creates a mount point from one node to another node's root volume upon a request to access the node's log files or core files. Although you do not need to

manually manage mount points, you can do so by using the `system node root-mount` commands.

If you want to...	Use this command...
Manually create a mount point from one node to another node's root volume	<code>system node root-mount create</code> Only a single mount point can exist from one node to another.
Display existing mount points on the nodes in the cluster, including the time a mount point was created and its current state	<code>system node root-mount show</code>
Delete a mount point from one node to another node's root volume and force connections to the mount point to close	<code>system node root-mount delete</code>

Related information

[ONTAP 9 commands](#)

Managing SSL

The SSL protocol improves the security of web access by using a digital certificate to establish an encrypted connection between a web server and a browser.

You can manage SSL for the cluster or a storage virtual machine (SVM) in the following ways:

- Enabling SSL
- Generating and installing a digital certificate and associating it with the cluster or SVM
- Displaying the SSL configuration to see whether SSL has been enabled, and, if available, the SSL certificate name
- Setting up firewall policies for the cluster or SVM, so that web access requests can go through
- Defining which SSL versions can be used
- Restricting access to only HTTPS requests for a web service

Related concepts

[Managing the web protocol engine](#) on page 142

[Managing web services](#) on page 148

Related tasks

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

Related information

[Network and LIF management](#)

Commands for managing SSL

You use the `security ssl` commands to manage the SSL protocol for the cluster or a storage virtual machine (SVM).

If you want to...	Use this command...
Enable SSL for the cluster or an SVM, and associate a digital certificate with it	<code>security ssl modify</code>
Display the SSL configuration and certificate name for the cluster or an SVM	<code>security ssl show</code>

Related references

[Commands for managing web services](#) on page 149

[Commands for managing the web protocol engine](#) on page 144

Related information

[Network and LIF management](#)

[ONTAP 9 commands](#)

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 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 150

[Managing the web protocol engine](#) on page 142

[Managing web services](#) on page 148

Related information

[Network and LIF management](#)

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...
<p>Your web browser returns an unable to connect or failure to establish a connection error when you try to access a web service.</p>	<p>Your LIF might be configured incorrectly.</p>	<p>Ensure that you can ping the LIF that provides the web service.</p> <p>Note: You use the <code>network ping</code> command to ping a LIF. For information about network configuration, see the <i>Network Management Guide</i>.</p>
	<p>Your firewall might be configured incorrectly.</p>	<p>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.</p> <p>Note: You use the <code>system services firewall policy</code> commands to manage firewall policies. You use the <code>network interface modify</code> command with the <code>- firewall-policy</code> parameter to associate a policy with a LIF.</p>
	<p>Your web protocol engine might be disabled.</p>	<p>Ensure that the web protocol engine is enabled so that web services are accessible.</p> <p>Note: You use the <code>system services web</code> commands to manage the web protocol engine for the cluster.</p>
<p>Your web browser returns a not found error when you try to access a web service.</p>	<p>The web service might be disabled.</p>	<p>Ensure that each web service that you want to allow access to is enabled individually.</p> <p>Note: You use the <code>vserver services web modify</code> command to enable a web service for access.</p>
<p>The web browser fails to log in to a web service with a user's account name and password.</p>	<p>The user cannot be authenticated, the access method is not correct, or the user is not authorized to access the web service.</p>	<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>Note: You use the <code>security login</code> commands to manage user accounts and their access methods and authentication methods. Accessing the 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>

This access problem...	Occurs because of this configuration error...	To address the error...
<p>You connect to your web service with HTTPS, and your web browser indicates that your connection is interrupted.</p>	<p>You might not have SSL enabled on the cluster or storage virtual machine (SVM) that provides the web service.</p>	<p>Ensure that the cluster or SVM has SSL enabled and that the digital certificate is valid.</p> <p>Note: You use the <code>security ssl</code> commands to manage SSL configuration for HTTP servers and the <code>security certificate show</code> command to display digital certificate information.</p>
<p>You connect to your web service with HTTPS, and your web browser indicates that the connection is untrusted.</p>	<p>You might be using a self-signed digital certificate.</p>	<p>Ensure that the digital certificate associated with the cluster or SVM is signed by a trusted CA.</p> <p>Note: You use the <code>security certificate generate-csr</code> command to generate a digital certificate signing request and the <code>security certificate install</code> command to install a CA-signed digital certificate. You use the <code>security ssl</code> commands to manage the SSL configuration for the cluster or SVM that provides the web service.</p>

Related concepts

[Managing the web protocol engine](#) on page 142

[Managing web services](#) on page 148

[Managing SSL](#) on page 150

Related information

[Network and LIF management](#)

Verifying the identity of remote servers using certificates

ONTAP supports security certificate features to verify the identity of remote servers.

ONTAP software enables secure connections using these digital certificate features and protocols:

- Online Certificate Status Protocol (OCSP) validates the status of digital certificate requests from ONTAP services using SSL and Transport Layer Security (TLS) connections. This feature is disabled by default.
- A default set of trusted root certificates is included with ONTAP software.
- Key Management Interoperability Protocol (KMIP) certificates enable mutual authentication of a cluster and a KMIP server.

Verifying digital certificates are valid using OCSP

Beginning with ONTAP 9.2, Online Certificate Status Protocol (OCSP) enables ONTAP applications that use Transport Layer Security (TLS) communications to receive digital certificate status when OCSP is enabled. You can enable or disable OCSP certificate status checks for specific applications at any time. By default, OCSP certificate status checking is disabled.

Before you begin

These commands must be performed at the advanced privilege level.

About this task

OCSP supports the following applications:

- AutoSupport
- Event Management System (EMS)
- LDAP over TLS
- Key Management Interoperability Protocol (KMIP)
- Audit Logging
- FabricPool

Steps

1. Set the privilege level to advanced: `set -privilege advanced`.
2. To enable or disable OCSP certificate status checks for specific ONTAP applications, use the appropriate command.

If you want OCSP certificate status checks for some applications to be...	Use the command...
Enabled	<code>security config ocspl enable -app <i>app name</i></code>
Disabled	<code>security config ocspl disable -app <i>app name</i></code>

Viewing default certificates for TLS-based applications

Beginning with ONTAP 9.2, ONTAP provides a default set of trusted root certificates for ONTAP applications using Transport Layer Security (TLS).

Before you begin

The default certificates are installed only on the admin SVM during its creation, or during an upgrade to ONTAP 9.2.

About this task

The current applications that act as a client and require certificate validation are AutoSupport, EMS, LDAP, Audit Logging, FabricPool, and KMIP.

When certificates expire, an EMS message is invoked that requests the user to delete the certificates. The default certificates can only be deleted at the advanced privilege level.

Note: Deleting the default certificates may result in some ONTAP applications not functioning as expected (for example, AutoSupport and Audit Logging).

Step

1. You can view the default certificates that are installed on the admin SVM by using the security certificate show command: `security certificate show -type server-ca`.

Example

```
Cluster1::> security certificate show -type server-ca
Vserver   Serial Number   Common Name   Type
-----
Cluster1  01              AAACertificateServices  server-ca
Certificate Authority: AAA Certificate Services
Expiration Date: Sun Dec 31 18:59:59 2028

Cluster1  570A119742C4E3  ActalisAuthenticationRootCA  server-ca
Certificate Authority: Actalis Authentication Root CA
Expiration Date: Sun Sep 22 07:22:02 2030

Cluster1  01              AddTrustExternalCARoot  server-ca
Certificate Authority: AddTrust External CA Root
Expiration Date: Sat May 30 06:48:38 2020
```

Mutually authenticating the cluster and a KMIP server

Mutually authenticating the cluster and an external key manager such as a Key Management Interoperability Protocol (KMIP) server enables the key manager to communicate with the cluster by using KMIP over SSL. You do so when an application or certain functionality (for example, the Storage Encryption functionality) requires secure keys to provide secure data access.

Generating a certificate signing request for the cluster

You can use the security certificate `generate-csr` command to generate a certificate signing request (CSR). After processing your request, the certificate authority (CA) sends you the signed digital certificate.

Before you begin

You must be a cluster administrator or SVM administrator to perform this task.

Steps

1. Generate a CSR:

```
security certificate generate-csr -common-name FQDN_or_common_name -size
512|1024|1536|2048 -country country -state state -locality locality -
organization organization -unit unit -email-addr email_of_contact -hash-
function SHA1|SHA256|MD5
```

For complete command syntax, see the man pages.

Example

The following command creates a CSR with a 2,048-bit private key generated by the SHA256 hashing function for use by the Software group in the IT department of a company whose custom common name is `server1.companyname.com`, located in Sunnyvale, California, USA. The email address of the SVM contact administrator is `web@example.com`. The system displays the CSR and the private key in the output.

```
cluster1::>security certificate generate-csr -common-name
server1.companyname.com -size 2048 -country US -state California -
locality Sunnyvale -organization IT -unit Software -email-addr
web@example.com -hash-function SHA256
Certificate Signing Request :
-----BEGIN CERTIFICATE REQUEST-----
MIIBGjCBxQIBADBGMRQwEgYDVQQDEwtleGFtcGx1LmNvbTELMakGAlUEBhMCVVMx
CTAHBgNVBAgTADEJMAcGAlUEBxMAMQkwBwYDVQQKEwAxCTAHBgNVBAStADEPMA0G
CSqGSIb3DQEJARYAMFwwDQYJKoZIhvcNAQEBBQADSwAwSAJBAPXFanNoJApTlnzS
xOcxixqImRRGZCR7tVmTYyqPSuTvfhVtwDJbmXuj6U3alwoUsb13wfEvQnHVFNCi
2ninsJ8CAwEAAaAAMA0GCSqGSIb3DQEBwUAA0EA6EagLfs05+4g+ejiRKKTUPQO
UqOUEoKuvxhOvPC2w7b//fNSFsFHvXloqEOhYECn/NX9h8mbphCom5YZ4OfnKw==
-----END CERTIFICATE REQUEST-----
Private Key :
24 | Administrator Authentication and RBAC Power Guide
-----BEGIN RSA PRIVATE KEY-----
MIIBOwIBAAJBAPXFanNoJApTlnzSxOcxixqImRRGZCR7tVmTYyqPSuTvfhVtwDJb
mXuj6U3alwoUsb13wfEvQnHVFNCi2ninsJ8CAwEAAQJAWt2A0+bW3FKezEuIrQlu
KoMyRYK455wtMk8BrOyJfhYsB20B28eifjJvRWdTOBEav99M7cEzgPv+p5kaZTm
gQIhAPsp+j1hrUXSRj979LIJY0sNez397i7ViFXWQScx/ehAiEA+oDbOooWlVvu
xj4aitxVBu6ByVckyU8LbsferNsZwD8CIQCbz1/ENvmlJ/P7N9Exj2NctEYxd0Q5
cwBZ5NfZeMBpwQIhAPk0KWQSLadGfsK0077itF+h9FGFNHbtuNTrVq4vPW3nAiAA
peMBQgEv28y2r8D4dkYzxcXmjzJluUSZS9c/wS6fA==
-----END RSA PRIVATE KEY-----
Note: Please keep a copy of your certificate request and private key
for future reference.
```

2. Copy the certificate request from the CSR output, and then send it in electronic form (such as email) to a trusted third-party CA for signing.

After processing your request, the CA sends you the signed digital certificate. You should keep a copy of the private key and the CA-signed digital certificate.

Installing a CA-signed server certificate for the cluster

To enable an SSL server to authenticate the cluster or storage virtual machine (SVM) as an SSL client, you install a digital certificate with the client type on the cluster or SVM. Then you provide the client-ca certificate to the SSL server administrator for installation on the server.

Before you begin

You must have already installed the root certificate of the SSL server on the cluster or SVM with the `server-ca` certificate type.

Steps

1. To use a self-signed digital certificate for client authentication, use the `security certificate create` command with the `type client` parameter.

2. To use a CA-signed digital certificate for client authentication, complete the following steps:

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

ONTAP displays the CSR output, which includes a certificate request and private key, and reminds you to copy the output to a file for future reference.

- b. Send the certificate request from the CSR output in an electronic form (such as email) to a trusted CA for signing.

You should keep a copy of the private key and the CA-signed certificate for future reference.

After processing your request, the CA sends you the signed digital certificate.

- c. Install the CA-signed certificate by using the `security certificate install` command with the `-type client` parameter.

- d. Enter the certificate and the private key when you are prompted, and then press **Enter**.

- e. Enter any additional root or intermediate certificates when you are prompted, and then press **Enter**.

You install an intermediate certificate on the cluster or SVM if a certificate chain that begins at the trusted root CA, and ends with the SSL certificate issued to you, is missing the intermediate certificates. An intermediate certificate is a subordinate certificate issued by the trusted root specifically to issue end-entity server certificates. The result is a certificate chain that begins at the trusted root CA, goes through the intermediate certificate, and ends with the SSL certificate issued to you.

3. Provide the `client-ca` certificate of the cluster or SVM to the administrator of the SSL server for installation on the server.

The `security certificate show` command with the `-instance` and `-type client-ca` parameters displays the `client-ca` certificate information.

Installing a CA-signed client certificate for the KMIP server

The certificate subtype of Key Management Interoperability Protocol (KMIP) (the `-subtype kmip-cert` parameter), along with the `client` and `server-ca` types, specifies that the certificate is used for mutually authenticating the cluster and an external key manager, such as a KMIP server.

About this task

Install a KMIP certificate to authenticate a KMIP server as an SSL server to the cluster.

Steps

1. Use the `security certificate install` command with the `-type server-ca` and `-subtype kmip-cert` parameters to install a KMIP certificate for the KMIP server.
2. When you are prompted, enter the certificate, and then press Enter.
ONTAP reminds you to keep a copy of the certificate for future reference.

Example

```
cluster1::> security certificate install -type server-ca -subtype kmip-cert
-vserver cluster1

Please enter Certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIICPDCCAaUCEDyRMcsf9tAbDpq40ES/Er4wDQYJKoZIhvcNAQEFBQAwXzELMAkG
2JhucwNhkcV8sEVAbkSdjbCxlRhLQ2pRdKkkirWmnWxbj9T/UWZYB2oK0z5XqcJ
2HUw19JlYDln1khVdWk/kfVIC0dpImmClr7JyDiGSnoscx1IaU5rfGW/D/xwzoiQ
...
-----END CERTIFICATE-----

You should keep a copy of the CA-signed digital certificate for future
reference.

cluster1::>
```


Logical space reporting and enforcement for volumes

Beginning with ONTAP 9.4, you can allow the logical space used in a volume and the amount of remaining storage space to be displayed to users in System Manager. Beginning with 9.5, you can limit the amount of logical space consumed by users.

The percentage of logical space used is shown in the Logical Space Used (%) column of the Volumes on SVM page in System Manager.

The Logical Space Reporting Column shows you if the logical space reporting feature is enabled.

What logical space reporting shows

When you enable the logical space reporting feature in ONTAP, System Manager displays the amount of logical used and available space in addition to the total space on a volume.

When used, logical space reporting shows the following columns to users in System Manager:

- **Used:** shows the total afs-physical-used space and savings in the volume
- **Available:** shows the amount of physical space currently available on the volume
- **Logical Space Used (%):** shows the percentage of current afs-logical-used with the provisioned size (excluding snap reserve of the volume)
- **Logical Space:** shows whether the logical space reporting feature is enabled
- **Total:** shows the amount of used and available space

The Total column can be greater than the provisioned space on the volume; this discrepancy occurs because the Total column includes any block savings achieved through duplication, compression, and other space-saving capabilities.

The Logical Space Used (%) column can be greater than 100%, as afs-logical-used includes savings of the volume. The afs-logical-used space of a volume does not include snap spill as used space. The afs-physical-used space of a volume includes snap spill as used space.

Note:

Logical Space Reporting is disabled by default.

Logical space reporting limitations

You should be aware that the logical space reporting feature has certain limitations.

The following features do not support logical space reporting:

- FlexGroups
- Logical Replication engine with Storage Efficiency (LRSE)

What logical space enforcement does

When the logical space enforcement feature is enabled, ONTAP counts the logical-used blocks in a volume to determine the amount of available space that is still available in that volume. If there is no space available in a volume, the system returns an ENOSPC (out-of-space) error message.

Logical space enforcement ensures that users are notified when a volume is full or nearly full. Beginning with ONTAP 9.5, logical space enforcement returns three types of alerts to inform you about the available space on a volume:

- `Monitor.vol.full.inc.sav`: This alert is triggered when 98% of the logical space in the volume has been used.
- `Monitor.vol.nearFull.inc.sav`: This alert is triggered when 95% of the logical space in the volume has been used.
- `Vol.log.overalloc.inc.sav`: This alert is triggered when the logical space used in the volume is greater than the total size of the volume.
This alert tells you that adding to the size of the volume might not create available space since that space will already be consumed by overallocated logical blocks.

Note: Total (logical space) should be equal to provisioned space excluding snap-reserve of the volume with logical space enforcement.

Controlling how space available and space consumed is displayed to end users

You can use the command-line interface in ONTAP 9.5 to enable System Manager to display information about used space and available space on volumes to users.

Step

1. Allow the logical space columns to be visible by setting the

```
-is-space-reporting-logical  
value to
```

```
true  
:
```

```
volume modify -vserver SVM-name -volume volume-name -size volume-size -  
is-space-reporting-logical true
```

The Logical Space Used (%) and the Logical Space columns are visible in System Manager.

Copyright

Copyright © 2020 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

[*doccomments@netapp.com*](mailto:doccomments@netapp.com)

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 1395 Crossman Ave., Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277