

Cluster management with the CLI ONTAP 9

NetApp September 19, 2024

This PDF was generated from https://docs.netapp.com/us-en/ontap/system-admin/index.html on September 19, 2024. Always check docs.netapp.com for the latest.

Table of Contents

Cluster management with the CLI	1
Administration overview with the CLI	1
Cluster and SVM administrators	1
Access the cluster by using the CLI (cluster administrators only)	3
Use the ONTAP command-line interface	14
Manage CLI sessions.	28
Cluster management (cluster administrators only).	29
Manage nodes	34
Configure the SP/BMC network	57
Manage nodes remotely using the SP/BMC	63
Manage the cluster time (cluster administrators only)	90
Manage the banner and MOTD	92
Manage jobs and schedule	102
Back up and restore cluster configurations (cluster administrators only).	105
Manage core dumps (cluster administrators only)	114

Cluster management with the CLI

Administration overview with the CLI

You can administer ONTAP systems with the command-line interface (CLI). You can use the ONTAP management interfaces, access the cluster, manage nodes, and much more.

You should use these procedures under the following circumstances:

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

Related information

For details about CLI syntax and usage, see the ONTAP command reference documentation.

Cluster and SVM administrators

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.



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.

Manage access to System Manager

You can enable or disable a web browser's access to 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.

What the cluster management server is

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

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

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

In the CLI, SVMs are displayed as Vservers.

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

Access 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	username

The system responds with the password prompt.

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

Access the cluster using SSH

You can issue SSH requests to an ONTAP 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 VM, 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.

About this task

- You must use an OpenSSH 5.7 or later client.
- 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) public 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 casesensitive. Case mismatch between the user name created in ONTAP and the user name created in AD results in a login failure.

SSH Authentication options

• Beginning 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.

- Beginning with ONTAP 9.4, you can enable SSH multifactor authentication for LDAP and NIS remote users.
- Beginning with ONTAP 9.13.1, you can optionally add certificate validation to the SSH authentication
 process to enhance login security. To do this, associate an X.509 certificate with the public key that an
 account uses. If you log in using SSH with both an SSH public key and an X.509 certificate, ONTAP checks
 the validity of the X.509 certificate before authenticating with the SSH public key. SSH login is refused if
 that certificate is expired or revoked, and the SSH public key is automatically disabled.
- Beginning with ONTAP 9.14.1, ONTAP administrators can add Cisco Duo two-factor authentication to the SSH authentication process to enhance login security. Upon first login after you enable Cisco Duo authentication, users will need to enroll a device to serve as an authenticator for SSH sessions.
- Beginning with ONTAP 9.15.1, administrators can Configure dynamic authorization to provide additional adaptive authentication to SSH users based on the user's trust score.

Steps

- 1. From a host with access to the ONTAP cluster's network, enter the ssh command in one of the following formats:
 - ° ssh username@hostname_or_IP [command]
 - ° ssh -1 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.
```

```
$ ssh -1 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.
```

```
$ 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.
```

Related information

Administrator authentication and RBAC

SSH login security

Beginning 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.



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
```

Enable Telnet or RSH access to the cluster

As a security best practice, Telnet and RSH are disabled by default. To enable the cluster to accept Telnet or RSH requests, you must enable the service in the default management service policy.

Telnet and RSH are not secure protocols; you should consider using SSH to access the cluster. SSH provides a secure remote shell and interactive network session. For more information, refer to Access the cluster using

SSH.

About this task

• ONTAP supports a maximum of 50 concurrent Telnet or 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 incoming connections is higher than 10 per second, the service is temporarily disabled for 60 seconds.

• RSH commands require advanced privileges.

ONTAP 9.6 or later

Steps

1. Confirm that the RSH or Telnet security protocol is enabled:

security protocol show

- a. If the RSH or Telnet security protocol is enabled, continue to the next step.
- b. If the RSH or Telnet security protocol is not enabled, use the following command to enable it:

security protocol modify -application <rsh/telnet> -enabled true

2. Confirm that the management-rsh-server or management-telnet-server service exists on the management LIFs:

network interface show -services management-rsh-server

or

network interface show -services management-telnet-server

- a. If the management-rsh-server or management-telnet-server service exists, continue to the next step.
- b. If the management-rsh-server or management-telnet-server service does not exist, use the following command to add it:

network interface service-policy add-service -vserver cluster1 -policy
default-management -service management-rsh-server

```
network interface service-policy add-service -vserver cluster1 -policy
default-management -service management-telnet-server
```

ONTAP 9.5 or earlier

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.

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* -destination-policy *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/netmask

To allow all IP addresses, you should specify -ip-list 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
```

Access the cluster by using Telnet

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

Telnet and RSH are not secure protocols; you should consider using SSH to access the cluster. SSH provides a secure remote shell and interactive network session. For more information, refer to Access the cluster using SSH.

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.

About this task

• 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.
- RSH commands require advanced privileges.

ONTAP 9.6 or later

Steps

1. Confirm that the Telnet security protocol is enabled:

security protocol show

- a. If the Telnet security protocol is enabled, continue to the next step.
- b. If the Telnet security protocol is not enabled, use the following command to enable it:

security protocol modify -application telnet -enabled true

2. Confirm that the management-telnet-server service exists on the management LIFs:

network interface show -services management-telnet-server

- a. If the management-telnet-server service exists, continue to the next step.
- b. If the management-telnet-server service does not exist, use the following command to add it:

network interface service-policy add-service -vserver cluster1 -policy
default-management -service management-telnet-server

ONTAP 9.5 or earlier

Before you begin

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

• 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.

Steps

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::>
```

Access 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.

Telnet and RSH are not secure protocols; you should consider using SSH to access the cluster. SSH provides a secure remote shell and interactive network session. For more information, refer to Access the cluster using SSH.

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.

About this task

• 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 incoming connections is higher than 10 per second, the service is temporarily disabled for 60 seconds.

• RSH commands require advanced privileges.

ONTAP 9.6 or later

Steps

1. Confirm that the RSH security protocol is enabled:

security protocol show

- a. If the RSH security protocol is enabled, continue to the next step.
- b. If the RSH security protocol is not enabled, use the following command to enable it:

security protocol modify -application rsh -enabled true

2. Confirm that the management-rsh-server service exists on the management LIFs:

network interface show -services management-rsh-server

- a. If the management-rsh-server service exists, continue to the next step.
- b. If the management-rsh-server service does not exist, use the following command to add it:

```
network interface service-policy add-service -vserver cluster1 -policy
default-management -service management-rsh-server
```

ONTAP 9.5 or earlier

Before you begin

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

• 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.

Steps

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

rsh hostname or IP -l username:passwordcommand

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

------

nodel true true

node2 true true

2 entries were displayed.

admin_host$
```

Use the ONTAP command-line interface

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::*>

About the different shells for CLI commands overview (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.

Display 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.



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_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]]
```

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.



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 casesensitive.

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

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

The CLI ignores the text between "#" and the end of the line.

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

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

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

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

Methods of viewing command history and reissuing commands

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

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

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

• A string that matches part of a previous command

For example, if the only volume command you have run is volume show, you can use the redo volume command to reexecute the command.

• The numeric ID of a previous command, as listed by the history command

For example, you can use the redo 4 command to reissue the fourth command in the history list.

• A negative offset from the end of the history list

For example, you can use the redo -2 command to reissue the command that you ran two commands ago.

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

cluster1::> redo -3

Keyboard shortcuts for editing CLI commands

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

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

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

If you want to	Use the following keyboard shortcut
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.	Ctrl-P
With each repetition of the keyboard shortcut, the history cursor moves to the previous entry	Esc-P
	Up arrow

If you want to	Use the following keyboard shortcut
Replace the current content of the command line with the next entry on the history list. With each repetition	Ctrl-N
of the keyboard shortcut, the history cursor moves to the next entry.	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.

Set 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.

Steps

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 NetApp personnel.
Do you wish to continue? (y or n): y
cluster1::*> set -privilege admin
```

Set 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

Steps

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 volume show -volume *tmp* displays a list of all volumes whose names include the string tmp.
!	NOT operator. Indicates a value that is not to be matched; for example, !vs0 indicates not to match the value vs0.
	OR operator. Separates two values that are to be compared; for example, vs0 vs2 matches either vs0 or vs2. You can specify multiple OR statements; for example, $a b^* *c^*$ matches the entry a, any entry that starts with b, and any entry that includes c.
	Range operator. For example, 510 matches any value from 5 to 10, inclusive.
<	Less-than operator. For example, <20 matches any value that is less than 20.
>	Greater-than operator. For example, >5 matches any value that is greater than 5.
<=	Less-than-or-equal-to operator. For example, \leftarrow 5 matches any value that is less than or equal to 5.

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

If you want to parse query characters as literals, you must enclose the characters in double quotes (for example, "<10", "0..100", "*abc*", or "a|b") for the correct results to be returned.

You must enclose raw file names in double quotes to prevent the interpretation of special characters. This also applies to special characters used by the clustershell.

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".

Related information

Keyboard shortcuts for editing CLI commands

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
vsl root vol
               volume true
vs2
      new vol
                volume true
vs2
      root vol
               volume
                            true
. . .
cluster1::>
```

About 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.



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.

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 svm1 -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 svm1 -percent-snapshot-space 0
cluster1::> volume create -vserver svm1 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 CLI commands. These pages are available at the command line and are also published in release-specific *command references*.

At the ONTAP command line, use the man command_name command to display the manual page of the specified command. If you do not specify a command name, the manual 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**.

Refer to the command reference for your version of ONTAP 9 to learn about the admin-level and advancedlevel ONTAP commands available in your release.

Manage 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.

Record a CLI session

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.

Steps

1. Start recording the current CLI session into a file:

system script start

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. When finished, stop recording the session:

system script stop

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	system script start
Stop recording the current CLI session	system script stop
Display information about records of CLI sessions	system script show

If you want to	Use this command
Upload a record of a CLI session to an FTP or HTTP destination	system script upload
Delete a record of a CLI session	system script delete

Related information

ONTAP command reference

Commands for 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 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	system timeout show
Modify the automatic timeout period for CLI sessions	system timeout modify

Related information

ONTAP command reference

Cluster management (cluster administrators only)

Display 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.

Steps

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	
nodel	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
```

Display cluster attributes

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

Steps

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
```

Modify 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.

Steps

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 $\ensuremath{-}\ensuremath{\mathsf{location}}$ 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

Display 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.

Steps

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
```

About 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 Туре 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.

Manage nodes

Add 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.

What you'll need

- If you are adding nodes to a multiple-node cluster, all 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 convert your two-node switchless cluster to a switch-attached cluster using a NetApp supported cluster switch.

The switchless cluster functionality is supported only in a two-node cluster.

- 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 SP automatic configuration enabled, the subnet specified for the SP must have available resources to allow the joining node to use the specified subnet to automatically configure 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 [eOM]:
```

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://<node_mgmt_or_eOM_IP_address>
Otherwise, press Enter to complete cluster setup using the
command line interface:
```



For more information on setting up a cluster using the setup GUI, see the 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.

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

If the ONTAP version running on the new node is different to the version running on the existing cluster, the system reports a System checks Error: Cluster join operation cannot be performed at this time error. This is the expected behavior. To continue, run the add-node -allow-mixed -version-join new_node_name command at the advanced privilege level from an existing node in the cluster.

- 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 information

Mixed version ONTAP clusters

Remove 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

The following conditions must be satisfied before removing nodes from the cluster:

- 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.
 - This might include purging data from an encrypted volume.
- All non-root volumes have been moved from aggregates owned by the node.
- All non-root aggregates have been deleted from the node.
- If the node owns Federal Information Processing Standards (FIPS) disks or self-encrypting disks (SEDs), disk encryption has been removed by returning the disks to unprotected mode.
 - You might also want to sanitize FIPS drives or SEDs.
- Data LIFs have been deleted or relocated from the node.
- Cluster management LIFs have been relocated from the node and the home ports changed.
- All intercluster LIFs have been removed.
 - When you remove intercluster LIFs a warning is displayed that can be ignored.
- Storage failover has been disabled for the node.
- All LIF failover rules have been modified to remove ports on the node.
- All VLANs on the node have been deleted.
- If you have LUNs on the node to be removed, you should modify the Selective LUN Map (SLM) reportingnodes list before you remove the node.

If you do not remove the node and its HA partner from the SLM reporting-nodes list, access to the LUNs previously on the node can be lost even though the volumes containing the LUNs were moved to another node.

It is recommended that you issue an AutoSupport message to notify NetApp technical support that node removal is underway.



You must not perform operations such as cluster remove-node, cluster unjoin, and node rename when an automated ONTAP upgrade is in progress.

About this task

- If you are running a mixed-version cluster, you can remove the last low-version node by using one of the advanced privilege commands beginning with ONTAP 9.3:
 - ONTAP 9.3: cluster unjoin -skip-last-low-version-node-check
 - ONTAP 9.4 and later: cluster remove-node -skip-last-low-version-node-check
- If you unjoin 2 nodes from a 4-node cluster, cluster HA is automatically enabled on the two remaining nodes.



All system and user data, from all disks that are connected to the node, must be made inaccessible to users before removing a node from the cluster. If a node was incorrectly unjoined from a cluster, contact NetApp Support for assistance with options for recovery.

Steps

1. Change the privilege level to advanced:

set -privilege advanced

2. Verify if a node on the cluster holds epsilon:

```
cluster show -epsilon true
```

- 3. If a node on the cluster holds epsilon and that node is going to be unjoined, move epsilon to a node that is not going to be unjoined:
 - a. Move epsilon from the node that is going to be unjoined

```
cluster modify -node <name of node to be unjoined> -epsilon false
```

b. Move epsilon to a node that is not going to be unjoined:

```
cluster modify -node <node name> -epsilon true
```

4. Identify the current master node:

```
cluster ring show
```

The master node is the node that holds processes such as "mgmt", "vldb", "vifmgr", "bcomd", and "crs".

- 5. 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:
 - a. Make the current master node ineligibly to participate in the cluster:

cluster modify -node <node name> -eligibility false

When the master node become ineligible, one of the remaining nodes is elected by the cluster quorum as the new master.

b. Make the previous master node eligible to participate in the cluster again:

```
cluster modify -node <node_name> -eligibility true
```

- 6. Log into the remote node management LIF or the cluster-management LIF on a node other than the one that is being removed.
- 7. Remove the node from the cluster:

For this ONTAP version	Use this command
ONTAP 9.3	cluster unjoin
ONTAP 9.4 and later	cluster remove-node*

If you have a mixed version cluster and you are removing the last lower version node, use the <code>-skip -last-low-version-node-check</code> parameter with these commands.

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 or option (9) Configure Advanced Drive Partitioning 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.

8. If a failure message indicates error conditions, address those conditions and rerun the cluster removenode or cluster unjoin command.

The node is automatically rebooted after it is successfully removed from the cluster.

- 9. 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.
- 10. Return to admin privilege level:

set -privilege admin

11. Repeat the preceding steps to remove the failover partner from the cluster.

Access 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.

What you'll need

• 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.

- You must use a local user account to access the spi web service, domain user accounts are not supported.
- 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:
 - o 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.

Access 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	system console
ONTAP CLI	system node run-console

- 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.

(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.

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

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

Manage node root volumes and root aggregates

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.

Rules governing node root volumes and root aggregates overview

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.

- 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.

- You can move the root volume to another aggregate. See Relocate root volumes to new aggregates.
- The root aggregate is dedicated to the node's root volume only.

ONTAP prevents you from creating other volumes in the root aggregate.

NetApp Hardware Universe

Free 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:

system node coredump show

2. Delete unwanted core dump files from the node:

system node coredump delete

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.

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_namesnapshot_name

7. Exit the nodeshell and return to the clustershell:

exit

Relocate root volumes to new aggregates

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

About this task

Storage failover must be enabled to relocate root volumes. You can use the storage failover modify -node *nodename* -enable true command to enable failover.

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. Set the privilege level to advanced:

set privilege advanced

2. Relocate the root aggregate:

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

 \circ -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.

3. Monitor the progress of the job:

job show -id jobid -instance

Results

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

Start or stop a node overview

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.

Reboot 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

The following example shows that "node1" holds epsilon:

cluster1::*> cluster Node	show 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 <code>-skip-lif-migration</code> 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.

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.

Boot 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	boot_ontap
The ONTAP primary image from the boot device	boot_primary
The ONTAP backup image from the boot device	boot_backup

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

Shut 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

The following example shows that "node1" holds epsilon:

```
cluster1::*> cluster show
                Health Eligibility
                                 Epsilon
Node
----- -----
                                 _____
node1
                true
                     true
                                 true
node2
                true
                     true
                                 false
node3
                true true
                                 false
node4
                true
                                 false
                     true
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 <code>-skip-lif-migration</code> 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.

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

Manage 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.



If an HA pair is using encrypting SAS or NVMe drives (SED, NSE, FIPS), you must follow the instructions in the topic Returning a FIPS drive or SED to unprotected mode for all drives within the HA pair prior to initializing the system (boot options 4 or 9). Failure to do this may result in future data loss if the drives are repurposed.

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:



i

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:

То	Select
Continue to boot the node in normal mode	1) Normal Boot

То	Select	
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 i) This menu option erases all data on the disks of the node and resets your node configuration to the factory default settings. Only select this menu item after the node has been removed from a cluster (unjoined) and is not joined to another cluster. For a node with internal or external disk shelves, the root volume on the internal disks is initialized. If there are no internal disk shelves, then the root volume on the external disks is initialized. For a system running FlexArray Virtualization with internal or external disk shelves, the array LUNs are not initialized. Any native disks on either internal or external shelves are initialized. For a system running FlexArray Virtualization with only array LUNS and no internal or external disk shelves, the root volume on the storage array LUNS are initialized, see Installing FlexArray. 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, see 9) Configure Advanced Drive Partitioning and Disks and aggregates management. 	
Perform aggregate and disk maintenance operations and obtain detailed aggregate and disk information.	5) Maintenance mode boot You exit Maintenance mode by using the halt command.	
Restore the configuration information from the node's root volume to the boot device, such as a PC CompactFlash card	 6) Update flash from backup config 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. 	

То	Select	
Install new software on the node	7) Install new software firstIf 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.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.	
Reboot the node	8) Reboot node	
Unpartition all disks and remove their ownership information or clean the configuration and initialize the system with whole or partitioned disks	 9) Configure Advanced Drive Partitioning 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: 	
	 (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. 	

Display 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.

Steps

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
       Differentiated Services: false
           All-Flash Optimized: true
            Capacity Optimized: false
                 QLC Optimized: false
   All-Flash Select Optimized: false
SAS2/SAS3 Mixed Stack Support: none
```

Modify 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.



You cannot modify node eligibility locally. It must be modified from a different node. Node eligibility also cannot be modified with a cluster HA configuration.



You should avoid setting a node's eligibility to false, except for situations 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.

Steps

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
```

Rename a node

You can change a node's name as required.

Steps

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.



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

Manage single-node clusters

A single-node cluster is a special implementation of a cluster running on a standalone node. Single-node clusters are not recommended because they do not provide redundancy. If the node goes down, data access is lost.



For fault tolerance and nondisruptive operations, it is highly recommended that you configure your cluster with high-availability (HA pairs).

If you choose to configure or upgrade a single-node cluster, you should be aware of the following:

- Root volume encryption is not supported on single-node clusters.
- If you remove 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.
- For single-node clusters, you can specify the configuration backup destination during software setup. After setup, those settings can be modified using ONTAP commands.
- If there are multiple hosts connecting to the node, each host can be configured with a different operating system such as Windows or Linux. If there are multiple paths from the host to the controller, then ALUA must be enabled on the host.

Ways to configure iSCSI SAN hosts with single nodes

You can configure iSCSI SAN hosts to connect directly to a single node or to connect through one or more IP switches. The node can have multiple iSCSI connections to the switch.

Direct-attached single-node configurations

In direct-attached single-node configurations, one or more hosts are directly connected to the node.



Single-network single-node configurations

In single-network single-node configurations, one switch connects a single node to one or more hosts. Because there is a single switch, this configuration is not fully redundant.



Multi-network single-node configurations

In multi-network single-node configurations, two or more switches connect a single node to one or more hosts. Because there are multiple switches, this configuration is fully redundant.



Ways to configure FC and FC-NVMe SAN hosts with single nodes

You can configure FC and FC-NVMe SAN hosts with single nodes through one or more fabrics. N-Port ID Virtualization (NPIV) is required and must be enabled on all FC switches in the fabric. You cannot directly attach FC or FC-NMVE SAN hosts to single nodes without using an FC switch.

Single-fabric single-node configurations

In single-fabric single-node configurations, there is one switch connecting a single node to one or more hosts. Because there is a single switch, this configuration is not fully redundant.

In single-fabric single-node configurations, multipathing software is not required if you only have a single path from the host to the node.

Multifabric single-node configurations

In multifabric single-node configurations, there are two or more switches connecting a single node to one or more hosts. For simplicity, the following figure shows a multifabric single-node configuration with only two fabrics, but you can have two or more fabrics in any multifabric configuration. In this figure, the storage controller is mounted in the top chassis and the bottom chassis can be empty or can have an IOMX module, as it does in this example.

The FC target ports (0a, 0c, 0b, 0d) in the illustrations are examples. The actual port numbers vary depending on the model of your storage node and whether you are using expansion adapters.



Related information

NetApp Technical Report 4684: Implementing and Configuring Modern SANs with NVMe-oF

ONTAP upgrade for single-node cluster

Beginning with ONTAP 9.2, you can use the ONTAP CLI to perform an automated update of a single-node cluster. Because single-node clusters lack redundancy, updates are always disruptive. Disruptive upgrades cannot be performed using System Manager.

Before you begin

You must complete upgrade preparation steps.

Steps

1. Delete the previous ONTAP software package:

cluster image package delete -version <previous package version>

2. Download the target ONTAP software package:

```
cluster image package get -url location
```

```
cluster1::> cluster image package get -url
http://www.example.com/software/9.7/image.tgz
```

Package download completed. Package processing completed.

3. Verify that the software package is available in the cluster package repository:

cluster image package show-repository

cluster1::> cluster image package show-repository
Package Version Package Build Time
-----9.7 M/DD/YYYY 10:32:15

4. Verify that the cluster is ready to be upgraded:

cluster image validate -version <package_version_number>

cluster1::> cluster image validate -version 9.7

WARNING: There are additional manual upgrade validation checks that must be performed after these automated validation checks have completed...

5. Monitor the progress of the validation:

cluster image show-update-progress

- 6. Complete all required actions identified by the validation.
- 7. Optionally, generate a software upgrade estimate:

cluster image update -version <package version number> -estimate-only

The software upgrade estimate displays details about each component to be updated, and the estimated duration of the upgrade.

8. Perform the software upgrade:

cluster image update -version <package_version_number>



If an issue is encountered, the update pauses and prompts you to take corrective action. You can use the cluster image show-update-progress command to view details about any issues and the progress of the update. After correcting the issue, you can resume the update by using the cluster image resume-update command.

9. Display the cluster update progress:

cluster image show-update-progress

The node is rebooted as part of the update and cannot be accessed while rebooting.

10. Trigger a notification:

autosupport invoke -node * -type all -message "Finishing Upgrade"

If your cluster is not configured to send messages, a copy of the notification is saved locally.

Configure the SP/BMC network

Isolate 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.



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.



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 reenable it, the SP might be reconfigured with a different address.

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

• If the SP has never been configured, SP IPv4 network configuration defaults to using IPv4 DHCP, and IPv6 is disabled.

A node that joins the cluster also uses IPv4 DHCP for its SP network configuration by default.

• The system service-processor network modify command enables you to configure a node's SP IP address.

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

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

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

Enable 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.



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 reenable 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.

Configure 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.

What you'll need

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.



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 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.
- 3. Display the SP network configuration and verify the SP setup status by using the system serviceprocessor network show command with the -instance or -field 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::>
```

Modify 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	system service-processor api-service modify with the -port {4915265535} parameter
Renew the SSL and SSH certificates used by the SP API service for internal communication	 For ONTAP 9.5 or later use system service-processor api-service renew-internal-certificate For ONTAP 9.4 and earlier use system service-processor api-service renew-certificates If no parameter is specified, only the host certificates (including the client and server certificates) are renewed. If the -renew-all true parameter is specified, both the host certificates and the root CA certificate are renewed.
comm	

If you want to	Use the following command
Disable or reenable the SP API service	system service-processor api-service modify with the -is-enabled {true false} parameter

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

Manage nodes remotely using the SP/BMC

Manage a node remotely using the SP/BMC overview

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.

The following platforms support BMC instead of SP:

- FAS 8700
- FAS 8300
- FAS27x0
- AFF A800
- AFF A700s
- AFF A400
- AFF A320
- AFF A220
- AFF C190

About 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 from 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.



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):



What the Baseboard Management Controller does

Beginning with ONTAP 9.1, on certain hardware platforms, software is customized to support a new onboard controller in 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.
- 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.



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.



Node operations are not impacted during a BMC reboot.

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.



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:
 - $\,\circ\,$ 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.



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, see the Knowledge Base article xref:./system-admin/ Health Monitor SPAutoUpgradeFailedMajorAlert SP upgrade fails - AutoSupport Message.

- 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
- $\circ\,$ 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 serviceprocessor 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 advance 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.



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

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.



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.

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

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. Beginning with ONTAP 9.9.1, SP user accounts must have the admin role.

User accounts for accessing the SP are managed from ONTAP instead of the SP CLI. A cluster user account 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.

You must specify the -role parameter when creating an SP user account.

- In ONTAP 9.9.1 and later releases, you must specify admin for the -role parameter, and any modifications to an account require the admin role. Other roles are no longer permitted for security reasons.
 - If you are upgrading to ONTAP 9.9.1 or later releases, see Change in user accounts that can access the Service Processor.
 - If you are reverting to ONTAP 9.8 or earlier releases, see Verify user accounts that can access the Service Processor.
- In ONTAP 9.8 and earlier releases, any role can access the SP, but admin is recommended.

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.

Access 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.

What you'll need

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.



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>

Access 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.

Manage 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, fd20:8ble:b255:c09b::/64.

Setting the -allowed-addresses parameter to 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
cluster1::> system service-processor ssh show
  Allowed Addresses: 0.0.0.0/0, ::/0
```

Use 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 help at the SP prompt.	Type system at the BMC prompt.

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

The following example shows the BMC CLI online help.

```
BMC> system

system acp - acp related commands

system battery - battery related commands

system console - connect to the system console

system core - dump the system core and reset

system cpld - cpld commands

system log - print system console logs

system power - commands controlling system power

system reset - reset the system using the selected firmware

system sensors - print environmental sensors status

system service-event - print service-event status

system fru - fru related commands

system watchdog - system watchdog commands

BMC>
```

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

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
```

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

```
BMC> system power help
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
BMC>
```

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	Use this BMC command	Or this ONTAP command
Display available SP commands or subcommands of a specified SP command	help[command]		
Display the current privilege level for the SP CLI	priv show		
Set the privilege level to access the specified mode for the SP CLI	priv set{admin advanced diag}		
Display system date and time	date		date

If you want to	Use this SP command	Use this BMC command	Or this ONTAP command
Display events that are logged by the SP	events{all info newest number oldest number search keyword}		
Display SP status and network configuration information	sp status $[-v -d]$ The $-v$ option displays SP statistics in verbose form. The $-d$ option adds the SP debug log to the display.	bmc status $[-v -d]$ The $-v$ option displays SP statistics in verbose form. The $-d$ option adds the SP debug log to the display.	system service- processor show
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	sp uptime	bmc uptime	
Display system console logs	system log		
Display the SP log archives or the files in an archive	<pre>sp log history show [-archive {latest all archive-name}][-dump {all file- name}]</pre>	<pre>bmc log history show[-archive {latest all archive-name}][-dump {all file-name}]</pre>	
Display the power status for the controller of a node	system power status		system node power show
Display battery information	system battery show		
Display ACP information or the status for expander sensors	system acp[show sensors show]		
List all system FRUs and their IDs	system fru list		
Display product information for the specified FRU	system fru show fru_id		

If you want to	Use this SP command	Use this BMC command	Or this ONTAP command
Display the FRU data history log	system fru log show (advanced privilege level)		
Display the status for the environmental sensors, including their states and current values	system sensors Or system sensors show		system node environment sensors show
Display the status and details for the specified sensor	<pre>system sensors get sensor_name You can obtain sensor_name by using the system sensors or the system sensors show command.</pre>		
Display the SP firmware version information	version		system service- processor image show
Display the SP command history	sp log audit (advanced privilege level)	bmc log audit	
Display the SP debug information	sp log debug (advanced privilege level)	bmc log debug (advanced privilege level)	
Display the SP messages file	sp log messages (advanced privilege level)	bmc log messages (advanced privilege level)	
Display the settings for collecting system forensics on a watchdog reset event, display system forensics information collected during a watchdog reset event, or clear the collected system forensics information	system forensics [show log dump log clear]		
Log in to the system console	system console		system node run- console
	You should press Ctrl-D to	exit the system console ses	sion.

If you want to	Use this SP command	Use this BMC command	Or this ONTAP command
Turn the node on or off, or perform a power-cycle (turning the power off and then back on)	system power on		system node power on (advanced privilege level)
	system power off		
	system power cycle		
	The standby power stays of the power-cycle, a brief partUsing these conclusionUsing these conclusioncause an impro- shutdown) and ONTAP system	on to keep the SP running wi use occurs before power is t ommands to turn off or powe oper shutdown of the node (l is not a substitute for a grad em node halt command.	thout interruption. During turned back on. r-cycle the node might also called a <i>dirty</i> ceful shutdown using the
Create a core dump and reset the node	system core [-f] The -f option forces the creation of a core dump and the reset of the node.		system node coredump trigger (advanced privilege level)
	These commands have the (NMI) button on a node, ca of the core files when haltin ONTAP on the node is hun node shutdown. The ger the system node cored long as the input power to	e same effect as pressing the nusing a dirty shutdown of the ng the node. These comman g or does not respond to con- nerated core dump files are of nump_show command. The st the node is not interrupted.	e Non-maskable Interrupt e node and forcing a dump ds are helpful when mmands such as system displayed in the output of SP stays operational as
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	system reset {primary backup current}		<pre>system node reset with the -firmware {primary backup current} parameter(advanced privilege level) system node reset</pre>
	This operation If no BIOS firmware image The SP stays operational a interrupted.	causes a dirty shutdown of t is specified, the current ima as long as the input power to	the node. ge is used for the reboot. the node is not

If you want to	Use this SP command	Use this BMC command	Or this ONTAP command
Display the status of battery firmware automatic update, or enable or disable battery firmware automatic update upon next SP boot	system battery auto_update[status enable disable] (advanced privilege level)		
Compare the current battery firmware image against a specified firmware image	<pre>system battery verify[image_URL] (advanced privilege level) If image_URL is not specified, the default battery firmware image is used for comparison.</pre>		
Update the battery firmware from the image at the specified location	system battery flash image_URL (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	sp update image_URL image_URL must not exceed 200 characters .	bmc update image_URL image_URL must not exceed 200 characters .	system service- processor image update
Reboot the SP	sp reboot		system service- processor reboot-sp
Erase the NVRAM flash content	system nvram flash clear (advanced privilege level) This command cannot be initiated when the controller power is off (system power off).		
Exit the SP CLI	exit		

About 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 nodel> system sensors			
Sensor Name Current UNC UCR	Unit	Status LCR	LNC
			+
CPU0_Temp_Margin -55.000	degrees C	ok na	na
CPU1_Temp_Margin -56.000	degrees C	ok na	na
In_Flow_Temp 32.000 42.000 52.000	degrees C	ok 0.000	10.000
Out_Flow_Temp 38.000 59.000 68.000	degrees C	ok 0.000	10.000
CPU1_Error 0x0 na na	discrete	0x0180 na	na
CPU1_Therm_Trip 0x0 na	discrete	0x0180 na	na
CPU1_Hot 0x0 na na	discrete	0x0180 na	na
IO_Mid1_Temp 30.000 55.000 64.000	degrees C	ok 0.000	10.000
IO_Mid2_Temp 30.000 55.000 64.000	degrees C	ok 0.000	10.000
CPU_VTT 1.106 1.154 1.174	Volts	ok 1.028	1.048
CPU0_VCC 1.154 1.348 1.368	Volts	ok 0.834	0.844
3.3V 3.323 3.466 3.546	Volts	ok 3.053	3.116
5V 5.002 5.490 5.636	Volts	ok 4.368	4.465
STBY_1.8V 1.794 1.892 1.911	Volts	ok 1.678	1.707

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+
```

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

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

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

Values	Condition of the sensor
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	system service-processor network auto- configuration enable
Disable the SP automatic network configuration for the IPv4 or IPv6 address family of the subnet specified for the SP	system service-processor network auto- configuration disable
Display the SP automatic network configuration	system service-processor network auto- configuration show

If you want to	Run this ONTAP command
Manually configure the SP network for a node, including the following:	system service-processor network modify
 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 	
Display the SP network configuration, including the following:	system service-processor network show
	Displaying complete SP network details requires the
 The configured address family (IPv4 or IPv6) and whether it is enabled 	-instance parameter.
 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	
Modify the SP API service configuration, including the following:	system service-processor api-service modify
 Changing the port used by the SP API service 	(advanced privilege level)
 Enabling or disabling the SP API service 	,

If you want to	Run this ONTAP command
Display the SP API service configuration	system service-processor api-service show (advanced privilege level)
Renew the SSL and SSH certificates used by the SP API service for internal communication	• For ONTAP 9.5 or later: system service- processor api-service renew-internal- certificates
	• For ONTAP 9.4 or earlier: system service- processor api-service renew- certificates
	(advanced privilege level)

Commands for managing the SP firmware image

If you want to	Run this ONTAP command
 Display the details of the currently installed SP firmware image, including the following: 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 	system service-processor image show The -is-current parameter indicates the image (primary or backup) that the SP is currently booted from, not if the installed firmware version is most current.
Enable or disable the SP automatic firmware update	system service-processor image modify 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.

If you want to	Run this	ONTAP command
Manually download an SP firmware image on a node	system node image get	
	i	Before you run the system node image commands, you must set the privilege level to advanced (set -privilege advanced), entering y when prompted to continue.
	The SP fi do not ne unless yc different f	rmware image is packaged with ONTAP. You eed to download the SP firmware manually, ou want to use an SP firmware version that is from the one packaged with ONTAP.
Display the status for the latest SP firmware update triggered from ONTAP, including the following information:	system progres	service-processor image update- s show
 The start and end time for the latest SP firmware update 		
 Whether an update is in progress and the percentage that is complete 		

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	system service-processor ssh add- allowed-addresses
Block the specified IP addresses from accessing the SP	system service-processor ssh remove- allowed-addresses
Display the IP addresses that can access the SP	system service-processor ssh show

Commands for general SP administration

If you want to	Run this ONTAP command
Display general SP information, including the following:	system service-processor show Displaying complete SP information requires the -instance
 The remote management device type 	parameter.
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 	
Reboot the SP on a node	system service-processor reboot-sp
Generate and send an AutoSupport message that includes the SP log files collected from a specified node	system node autosupport invoke-splog
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	system service-processor log show- allocations

Related information

ONTAP command reference

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	system service-processor show
Display/modify the BMC network configuration	system service-processor network show/modify
Reset the BMC	system service-processor reboot-sp

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

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 $\trl+D$ 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.

Command	Function
system power status	Print system power status.
system reset	Reset the system.
system log	Print system console logs
system fru show [id]	Dump all/selected field replaceable unit (FRU) info.

Manage 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 with 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.
 - $\,\circ\,$ 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).



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.

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

If you want to	Use this command
Associate the cluster with an external NTP server without symmetric authentication	cluster time-service ntp server create -server server_name
Associate the cluster with an external NTP server with symmetric authenticationAvailable in ONTAP 9.5 or later	cluster time-service ntp server create -server server_ip_address -key-id key_id
	(i) The key_id must refer to an existing shared key configured with '`cluster time-service ntp key'.
Enable symmetric authentication for an existing NTP serverAn existing NTP server can be modified to enable authentication by adding the required key-id. Available in ONTAP 9.5 or later	cluster time-service ntp server modify -server server_name -key-id key_id
Disable symmetric authentication	cluster time-service ntp server modify -server server_name -is-authentication -enabled false

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

If you want to	Use this command	
Configure a shared NTP key	<pre>cluster time-service ntp key create -ic shared_key_id -type shared_key_type -value shared_key_value</pre>	
	Gi 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	cluster time-service ntp server show	
Modify the configuration of an external NTP server that is associated with the cluster	cluster time-service ntp server modify	
Dissociate an NTP server from the cluster	cluster time-service ntp server delete	
Reset the configuration by clearing all external NTP servers' association with the cluster	cluster time-service ntp server reset	
	This command requires the advanced privilege level.	

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

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

Related information

ONTAP command reference

Manage the banner and MOTD

Manage the banner and MOTD overview

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:
```

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.
```

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.

Create 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 -message "text" parameter to specify the text.
Include newlines (also known as end of lines) in the message	Use the command without the -message or -uri parameter to launch the interactive mode for editing the banner.
Upload content from a location to use for the banner	Use the -uri 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
12345678901234567890123456789012345678901234567890123456789012345678901234
567890
The svm1 SVM is reserved for authorized users only!
cluster1::>
```

The following example displays the banners that have been created:

Related information

Managing the banner

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	<pre>Set a cluster-level banner: 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 (""): security login banner modify -vserver * -message ""

If you want to	Then
Override a banner created by an SVM administrator	<pre>Modify the SVM banner message: 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: security login banner modify -vserver <pre>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: security login banner modify -vserver <pre>svm_name -message ""</pre>
Use the cluster-level banner when the SVM currently uses an SVM-level banner	Set the SVM banner to "-": security login banner modify -vserver <i>svm_name</i> -message "-"

Create 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 -message "text" parameter to specify the text.

If you want to	Then
Include newlines (also known as end of lines)	Use the command without the -message or -uri parameter to launch the interactive mode for editing the MOTD.
Upload content from a location to use for the MOTD	Use the -uri 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 <code>-is-cluster-message-enabled</code> parameter to <code>false</code> 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.

See the security login mote modify command man page for a list of parameters to use to enable the MOTD to display dynamically generated content. Be sure to check the man page specific to your ONTAP version.

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!"
```

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.
                 2 3
                                               5
                                                                   7
0
        1
                                    4
                                                         6
8
123456789012345678901234567890123456789012345678901234567890123456789012345678901234
567890
Welcome to the \n SVM. Your user ID is '\N'. Your last successful login
was \L.
```

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.
```

Manage 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.

For a list of escape sequences that can be used to dynamically generate content for the MOTD, see the command reference.

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	<pre>Set a cluster-level MOTD: security login motd modify -vserver cluster_name { [-message "text"] [- uri ftp_or_http_addr] }</pre>
Change the MOTD for all logins when no SVM-level MOTDs are configured	<pre>Modify the cluster-level MOTD: security login motd modify -vserver cluster_name { [-message "text"] } [-uri ftp_or_http_addr] }</pre>
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 cluster_name -message ""
Have every SVM display the cluster-level MOTD instead of using the SVM-level MOTD	<pre>Set a cluster-level MOTD, then set all SVM-level MOTDs to an empty string with the cluster-level MOTD enabled: 1. security login motd modify -vserver cluster_name { [-message "text"] [-uri ftp_or_http_addr] } 2. security login motd modify { -vserver !"cluster_name" } -message "" -is-cluster-message-enabled true</pre>
Have an MOTD displayed for only selected SVMs, and use no cluster-level MOTD	<pre>Set the cluster-level MOTD to an empty string, then set SVM-level MOTDs for selected SVMs: 1. security login motd modify -vserver cluster_name -message "" 2. security login motd modify -vserver svm_name { [-message "text"] [- uri ftp_or_http_addr] } You can repeat this step for each SVM as needed.</pre>

If you want to	Then	
Use the same SVM-level MOTD for all (data and admin) SVMs	Set the cluster and all SVMs to use the same MOTD:	
	<pre>securit; { [-mes ftp_or_;</pre>	y login motd modify - <i>vserver</i> * sage " <i>text</i> "] [-uri <i>http_addr</i>] }
	i	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	II Set a cluster-level MOTD, but disable its display the cluster:	
	securit cluster uri ftp -messag	y login motd modify -vserver _ <i>name</i> { [-message " <i>text</i> "] [- _ <i>or_http_addr</i>] } -is-cluster e-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 clu for the MC securit; -message	uster and all SVMs to use an empty string DTD: y login motd modify -vserver * e ""
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 exter selectively security !"clust	<pre>nded queries to modify the MOTD y: y login motd modify { -vserver er_name" -message !"" } { [- "toxt"] [-uri</pre>
	ftp_or_	http_addr] }
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 que securit *"Janua	ery to display MOTDs: y login motd show -message ry"*"2015"*
Interactively create an MOTD that includes multiple and consecutive newlines (also known as end of lines, or EOLs)	In the inte followed b terminatin	eractive mode, press the space bar by Enter to create a blank line without ig 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	<pre>Modify the SVM-level MOTD: 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: 1. security login motd modify -vserver svm_name -message "" 2. (For the cluster administrator) security login motd modify -vserver svm_name -is-cluster-message-enabled true 	
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: 1. security login motd modify -vserver svm_name -message "" 2. (For the cluster administrator) security login motd modify -vserver svm_name -is-cluster-message-enabled false 	

Manage jobs and schedule

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.

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

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	job show
Display information about jobs on a per-node basis	job show bynode
Display information about cluster-affiliated jobs	job show-cluster
Display information about completed jobs	job show-completed
Display information about job history	job history show 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.
Display the list of private jobs	job private show (advanced privilege level)
Display information about completed private jobs	job private show-completed (advanced privilege level)
Display information about the initialization state for job managers	job initstate show (advanced privilege level)
Monitor the progress of a job	job watch-progress
Monitor the progress of a private job	job private watch-progress (advanced privilege level)
Pause a job	job pause
Pause a private job	job private pause (advanced privilege level)
Resume a paused job	job resume

If you want to	Use this command		
Resume a paused private job	job private resume (advanced privilege level)		
Stop a job	job stop		
Stop a private job	job private stop (advanced privilege level)		
Delete a job	job delete		
Delete a private job	job private delete (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	job unclaim (advanced privilege level)		

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

Related information

(;

ONTAP command reference

Commands for managing job schedules

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

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

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

If you want to	Use this command		
Display information about all schedules	job schedule show		
Display the list of jobs by schedule	job schedule show-jobs		
Display information about cron schedules	job schedule cron show		
Display information about interval schedules	job schedule interval show		

If you want to	Use this command		
Create a cron schedule	job schedule cron create Beginning with ONTAP 9.10.1, you can include the SVM for your job schedule.		
Create an interval schedule	job schedule interval create You must specify at least one of the following parameters: -days, -hours, -minutes, or -seconds.		
Modify a cron schedule	job schedule cron modify		
Modify an interval schedule	job schedule interval modify		
Delete a schedule	job schedule delete		
Delete a cron schedule	job schedule cron delete		
Delete an interval schedule	job schedule interval delete		

Related information

ONTAP command reference

Back up and restore cluster configurations (cluster administrators only)

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.



Configuration backup files contain configuration information only. They do not include any user data. For information about restoring user data, see Data Protection.

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.

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:	system configuration backup settings modify		
 Specify a remote URL (HTTP, HTTPS, FTP, FTPS, or TFTP) where the configuration backup files will be uploaded in addition to the default locations in the cluster 	When you use HTTPS in the remote URL, use the -validate-certification option to enable or disable digital certificate validation. Certificate validation is disabled by default.		
 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 	i	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.	
Set the password to be used to log in to the remote URL	system configuration backup settings set-password		
If you want to	Use this command		
---	---		
View the settings for the configuration backup schedule	<pre>system configuration backup settings show You set the -instance parameter to view the user name and the number of backups to keep for each schedule.</pre>		

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	system configuration backup create
Copy a configuration backup file from a node to another node in the cluster	system configuration backup copy
Upload a configuration backup file from a node in the cluster to a remote URL (FTP, HTTP, HTTPS, TFTP, or FTPS)	 system configuration backup upload When you use HTTPS in the remote URL, use the -validate-certification option to enable or disable digital certificate validation. Certificate validation is disabled by default. 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. Some web servers might require the installation of an additional module. For more information, see your web server's documentation. Supported URL formats vary by ONTAP release. See the command line help for your ONTAP version.
Download a configuration backup file from a remote URL to a node in the cluster, and, if specified, validate the digital certificate	system configuration backup download When you use HTTPS in the remote URL, use the -validate-certification option to enable or disable digital certificate validation. Certificate validation is disabled by default.

If you want to	Use this command
Rename a configuration backup file on a node in the cluster	system configuration backup rename
View the node and cluster configuration backup files for one or more nodes in the cluster	system configuration backup show
Delete a configuration backup file on a node	system configuration backup deleteThis 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.

Find 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 system configuration backup download command at the advanced privilege level to download it to the recovering node.
On a node in the cluster	a. Use the system configuration backup show 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 system configuration backup copy 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.

Restore 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. Change to the advanced privilege level:

set -privilege advanced

2. 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.

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

3. 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.

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.

4. 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.

5. 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.

Find 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.

If you are recovering the cluster from a configuration backup file, any configuration changes made since the backup was taken will be lost. You must resolve any discrepancies between the configuration backup file and the present configuration after recovery. See Knowledge Base article ONTAP Configuration Backup Resolution Guide for troubleshooting guidance.

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 system configuration backup download command at the advanced privilege level to download it to the recovering node.

If the configuration backup file is located	Then
On a node in the cluster	a. Use the system configuration backup show 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 system configuration backup copy command to copy it to the recovering node.

Restore 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.

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.

If you are recovering the cluster from a configuration backup file, any configuration changes made since the backup was taken will be lost. You must resolve any discrepancies between the configuration backup file and the present configuration after recovery. See the Knowledge Base article ONTAP Configuration Backup Resolution Guide for troubleshooting guidance.

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"
```

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.

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.

```
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

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.

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 reenable cluster HA.
- 11. Use the **storage failover modify** command to reenable 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 Data Protection.

Synchronize 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.

This example synchronizes a node (*node2*) with the rest of the cluster:

Result

The RDB is replicated to the node, and the node becomes eligible to participate in the cluster.

Manage 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

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	system node coredump config modify
Display the configuration settings for core dumps	system node coredump config show
Display basic information about core dumps	system node coredump show
Manually trigger a core dump when you reboot a node	system node reboot with both the -dump and -skip-lif-migration-before-reboot parametersThe skip-lif-migration- before-reboot parameter specifies that LIF migration prior to a reboot will be skipped.
Manually trigger a core dump when you shut down a node	system node halt with both the -dump and-skip-lif-migration-before-shutdownparametersImage: Colspan="2">The skip-lif-migration-before-shutdown parameterspecifies that LIF migration prior to a shutdown will be skipped.
Save a specified core dump	system node coredump save
Save all unsaved core dumps that are on a specified node	system node coredump save-all
Generate and send an AutoSupport message with a core dump file you specify	system node autosupport invoke-core- upload The -uri optional parameter specifies an alternate destination for the AutoSupport message.
Display status information about core dumps	system node coredump status
Delete a specified core dump	system node coredump delete
Delete all unsaved core dumps or all saved core files on a node	system node coredump delete-all

If you want to	Use this command
Display application core dump reports	system node coredump reports show
Delete an application core dump report	system node coredump reports delete

Related information

ONTAP command reference

Copyright information

Copyright © 2024 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.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, 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) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at http://www.netapp.com/TM are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.