



Clustered Data ONTAP[®] 8.3

System Administration Guide for SVM Administrators



NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.

Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-09168_A0
November 2014

Contents

Understanding SVM administration	6
What SVMs are	6
Why you use SVMs	7
Differences between cluster and SVM administrators	8
Data ONTAP management interface basics	9
Using the Data ONTAP command-line interface	9
Methods of navigating CLI command directories	9
Rules for specifying values in the CLI	10
Methods of viewing command history and reissuing commands	10
Keyboard shortcuts for editing CLI commands	11
Use of administrative privilege levels	12
Setting the privilege level in the CLI	13
Setting display preferences in the CLI	13
Methods of using query operators	14
Methods of using extended queries	15
Methods of customizing show command output by using fields	16
Methods of accessing Data ONTAP man pages	16
Accessing SVMs	18
Access methods for user accounts	18
Authentication methods for user accounts	18
Logging in to an SVM	19
Managing SVM authentication	21
Changing the login password	21
Managing SSH security configurations	22
Commands for managing SSH security configurations	23
Managing public keys	24
Commands for managing public keys	24
Managing digital certificates for server or client authentication	25
Installing a server certificate to authenticate the SVM as an SSL server	26
Installing a client CA or root CA certificate to authenticate an SSL client of the SVM	29

Installing a server CA certificate to authenticate an SSL server to which the SVM is a client	33
Installing a client certificate to authenticate the SVM as an SSL client	34
Replacing an expired digital certificate	36
Commands for managing digital certificates	37
Managing SSL	38
Commands for managing SSL	39
Administering SVMs	40
Identifying the commands that you can execute	41
Displaying ONTAP APIs	42
Managing jobs and schedules	43
Commands for managing jobs	43
Commands for managing job schedules	44
Monitoring SVM performance	45
What objects, instances, and counters are	45
Decisions to make before you view performance data	46
Viewing performance data for a time period	47
Viewing continuously updated performance data	48
Commands for monitoring SVM performance	49
Displaying information about SVMs	50
Displaying information about SVM peer relationships	51
Displaying information about network configuration	52
Monitoring SVMs using dashboard	53
Commands for managing dashboards	53
Data access protocols configuration	54
Commands for configuring data access protocols	55
Data security management	56
Commands for setting up security settings on files and managing tracing ...	56
Services configuration	57
Commands for configuring services	58
Storage management	58
Commands for managing storage	59
LUN management	60
Commands for managing LUNs	60
Backup management	61
Snapshot copy management	61

SnapMirror management	62
NDMP management	63
Commands for managing backup	63
Policy management	64
Commands for managing policies	65
Glossary	66
Copyright information	70
Trademark information	71
How to send your comments	72
Index	73

Understanding SVM administration

SVM administrators can administer Storage Virtual Machines (SVMs) and SVM resources, such as volumes, protocols, and services, depending on the capabilities assigned by the cluster administrator. To administer an SVM efficiently, you must understand what an SVM is, its benefits, and the types of administrators.

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

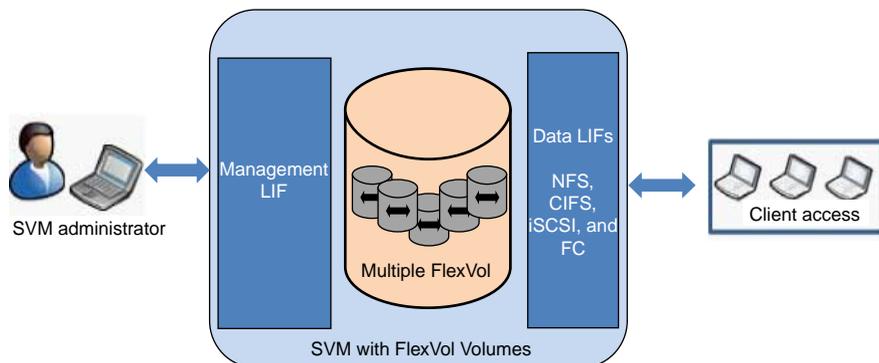
What SVMs are

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

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

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

SVM with FlexVol volumes

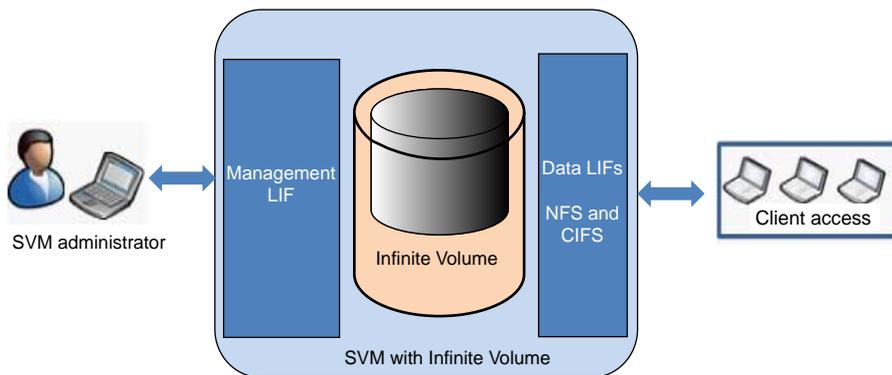


Each SVM with FlexVol volumes in a NAS environment presents a single directory hierarchical view and has a unique namespace. The namespace enables NAS clients to access data without specifying the physical location of the data. The namespace also enables the cluster and SVM administrators to manage distributed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

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

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

SVM with Infinite Volume



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

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

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

Why you use SVMs

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

SVMs provide the following benefits:

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

Note: SVMs with Infinite Volume can serve data only through NFS and CIFS protocols.
- **Easy management of large datasets**
With SVMs with Infinite Volume, management of large and unstructured data is easier because the SVM administrator can manage one data container instead of many.

Differences between cluster and SVM administrators

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

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

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

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

Data ONTAP management interface basics

You administer the Storage Virtual Machine (SVM) by using the Data ONTAP command-line interface (CLI). The CLI provides a command-based mechanism that is similar to the UNIX `tcsh` shell.

Using the Data ONTAP command-line interface

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

The CLI command prompt is represented as `vserver_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, `vserver_name : * >`.

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 the volumes by entering the `volume 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:

```
vsl1::> volume
vsl1::volume> 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 `vol show`. You can also use the Tab key to expand abbreviated commands and to display a command's parameters, including default parameter values.

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

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

Rules for specifying values in the CLI

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

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

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:

```
vs1::> redo -3
```

Keyboard shortcuts for editing CLI commands

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

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

If you want to...	Use the following keyboard shortcut...
Move the cursor back by one character	Ctrl-B
	Back arrow
Move the cursor forward by one character	Ctrl-F
	Forward arrow
Move the cursor back by one word	Esc-B
Move the cursor forward by one word	Esc-F
Move the cursor to the beginning of the line	Ctrl-A
Move the cursor to the end of the line	Ctrl-E
Remove the content of the command line from the beginning of the line to the cursor, and save it in the cut buffer The cut buffer acts like temporary memory, similar to what is called a <i>clipboard</i> in some programs.	Ctrl-U
Remove the content of the command line from the cursor to the end of the line, and save it in the cut buffer	Ctrl-K
Remove the content of the command line from the cursor to the end of the following word, and save it in the cut buffer	Esc-D
Remove the word before the cursor, and save it in the cut buffer	Ctrl-W

If you want to...	Use the following keyboard shortcut...
Yank the content of the cut buffer, and push it into the command line at the cursor	Ctrl-Y
Delete the character before the cursor	Ctrl-H
	Backspace
Delete the character where the cursor is	Ctrl-D
Clear the line	Ctrl-C
Clear the screen	Ctrl-L
Replace the current content of the command line with the previous entry on the history list With each repetition of the keyboard shortcut, the history cursor moves to the previous entry.	Ctrl-P
	Esc-P
	Up arrow
Replace the current content of the command line with the next entry on the history list With each repetition of the keyboard shortcut, the history cursor moves to the next entry.	Ctrl-N
	Esc-N
	Down arrow
Expand a partially entered command or list valid input from the current editing position	Tab
	Ctrl-I
Display context-sensitive help	?
Escape the special mapping for the question mark (“?”) character For instance, to enter a question mark into a command's argument, press Esc and then the “?” character.	Esc-?
Start TTY output	Ctrl-Q
Stop TTY output	Ctrl-S

Use of administrative privilege levels

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

admin

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

advanced

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

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

diagnostic

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

Setting the privilege level in the CLI

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

Step

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

Example of setting the privilege level

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

```
vs1::> set -privilege advanced
Warning: These advanced commands are potentially dangerous; use
them only when directed to do so by technical support.
Do you wish to continue? (y or n): y
vs1::*> set -privilege admin
```

Setting display preferences in the CLI

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

About this task

You can set the following CLI display preferences:

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

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

Step

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

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

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

Example of setting display preferences in the CLI

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

```
vs1::> set -showseparator "," -units GB
vs1::> rows 50
```

Methods of using query operators

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

The following table describes the supported query operators:

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

Operator	Description
..	Range operator. For example, <code>5..10</code> matches any value from 5 to 10, inclusive.
<	Less-than operator. For example, <code><20</code> matches any value that is less than 20.
>	Greater-than operator. For example, <code>>5</code> matches any value that is greater than 5.
<=	Less-than-or-equal-to operator. For example, <code><=5</code> matches any value that is less than or equal to 5.
>=	Greater-than-or-equal-to operator. For example, <code>>=5</code> matches any value that is greater than or equal to 5.
{ <i>query</i> }	Extended query. An extended query must be specified as the first argument after the command name, before any other parameters. For example, the command <code>volume modify {-volume *tmp*} -state offline</code> sets offline all volumes whose names include the string <code>tmp</code> .

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

You can use multiple query operators in one command line. For example, the command `volume show -size >1GB -percent-used <50` displays all volumes that are greater than 1 GB in size and less than 50% utilized.

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:

```
vs1::> 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.

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:

```
vs1::> vservers show -instance

                Vserver Type: data
      Vserver Subtype: default
      Vserver UUID: 4e42c9cf-32f2-11e2-9103-123456789012
      Root Volume: vs1_root
      Aggregate: aggr1
      NIS Domain: -
      Root Volume Security Style: mixed
      LDAP Client: -
      Default Volume Language Code: C.UTF-8
      Snapshot Policy: default
      Comment:
      Quota Policy: default
      ...
      Allowed Protocols: nfs, cifs
      Disallowed Protocols: fcp, iscsi, ndmp
      ...

Press <space> to page down, <return> for next line, or 'q' to quit...
...
vs1::>

vs1::> vservers show -fields allowed-protocols,disallowed-protocols
vservers  allowed-protocols  disallowed-protocols
-----
vs1       nfs,cifs           fcp,iscsi,ndmp

vs1::>
```

Methods of accessing Data ONTAP man pages

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

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

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

Related information

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

Accessing SVMs

As an SVM administrator, you can access SVMs by using different access methods. Your user account can be authenticated by using several authentication methods, as specified by the cluster administrator.

Access methods for user accounts

Depending on how the cluster administrator sets up an SVM user account, an SVM administrator can access the SVM for administration by using certain access methods.

You can access an SVM by using the following access methods:

- SSH
- Data ONTAP APIs

Note: Data ONTAP APIs access method is over HTTPS.

- SNMP

Authentication methods for user accounts

The method used to authenticate an SVM user account depends on the access method used by the cluster administrator to set up the SVM user account.

Your user account can be authenticated by using one of the following authentication methods:

- Network Information Service (NIS) and Lightweight Directory Access Protocol (LDAP)
nsswitch

Note: Clustered Data ONTAP supports only the RFC 2307 schema for LDAP authentication of SVM accounts. It does not support any other schemas, such as Active Directory Identity Management for UNIX (AD-IDMU) and Active Directory Services for UNIX (AD-SFU).

- Windows Active Directory (**domain**)
- User password (**password**)
- SSH public key (**publickey**)
- SNMP user-based security model (**usm**)
- SNMP community strings (**community**)

- SSL certificate authentication(**cert**)

Logging in to an SVM

To manage the SVM resources, an SVM administrator logs in to an SVM by using the user name and password provided by the cluster administrator. The SVM administrator can use an appropriate Secure Shell client application, such as PuTTY for Windows operating system and OpenSSH for UNIX operating system.

Before you begin

You must have the management IP address of the SVM, user name, and password.

About this task

After you log in, you might be able to manage all or some of the following SVM resources depending on the capabilities assigned to your account by the cluster administrator:

- Data access protocols, such as NFS, CIFS, iSCSI, and FC (FCoE included)
- Services, such as NIS, LDAP, and DNS
- Volumes, qtrees, quotas, Snapshot copies, and files
- Data backup with SnapMirror and NDMP
- Data security and policies

You can also monitor the network connection, network interface, LDAP client configuration, and SVM health.

Note: Clustered Data ONTAP supports only the AES and 3DES encryption algorithms (also known as ciphers) for SSH.

Step

1. To log in to an SVM by using SSH application, perform the appropriate action depending on the operating system:

If your host has...	Then...
Windows operating system	<ol style="list-style-type: none"> a. Enter the management IP address of the SVM in the SSH application. b. At the login prompt, enter the user name and password.

If your host has...	Then...
UNIX or Linux operating system	Enter the following command from the client application: <code>ssh vserver_admin_name@vserver_ip_address</code> <i>vserver_admin_name</i> is the user name. <i>vserver_ip_address</i> is the management IP address of the SVM.

Note: If you or the cluster administrator has created a public key for your user account, you do not require a password to log in to the SVM.

Related tasks

Identifying the commands that you can execute on page 41

Managing SVM authentication

As an SVM administrator, you can manage the security aspects of accessing an SVM such as managing your own user accounts and passwords, public keys, digital certificates, and SSL protocol.

You can perform the following tasks to manage the SVM authentication:

- Changing the login password
- Managing public keys
- Managing digital certificates for server or client authentication
- Managing SSL

Changing the login password

After an SVM administrator logs in to the SVM by using the user name and password provided by the cluster administrator, the SVM administrator can change the login password.

About this task

You must remember the following default rules when you change the login password:

- A password cannot contain the user name.
- A password must be at least eight characters long.
- A password must contain at least one letter and one number.
- A password cannot be the same as the last six passwords.

Steps

1. Change the login password by using the `security login password` command.
2. Enter your current password.
3. Enter a new password.
4. Confirm the password by entering the new password again.

Result

Your user account is updated with the new password. You must enter the new password on the subsequent login.

The following example shows how to change a user password:

```
vs1.example.com::~> security login password
Please enter your current password:
Please enter a new password:
Please enter it again:
vs1.example.com::~>
```

Managing SSH security configurations

Managing SSH security configurations involves managing the SSH key exchange algorithms and data encryption algorithms (also known as *ciphers*). Data ONTAP enables you to enable or disable individual SSH key exchange algorithms and ciphers for the Storage Virtual Machine (SVM) according to their SSH security requirements.

Data ONTAP supports the following SSH security configurations for SVMs:

- The following SSH key exchange algorithms are supported and enabled by default:
 - The **diffie-hellman-group-exchange-sha256** SSH key exchange algorithm for SHA-2
 - The **diffie-hellman-group-exchange-sha1**, **diffie-hellman-group14-sha1**, and **diffie-hellman-group1-sha1** SSH key exchange algorithms for SHA-1

SHA-2 algorithms are more secure than SHA-1 algorithms. Data ONTAP, which serves as an SSH server, automatically selects the most secure SSH key exchange algorithm that matches the client. To further enhance SSH security, you can manually disable the SHA-1 algorithms and leave only the SHA-2 algorithm enabled.

- For ciphers, the following counter (CTR) mode and cipher block chaining (CBC) mode of the AES and 3DES symmetric encryptions are supported and enabled by default:
 - **aes256-ctr**
 - **aes192-ctr**
 - **aes128-ctr**
 - **aes256-cbc**
 - **aes192-cbc**
 - **aes128-cbc**
 - **3des-cbc**

The CTR mode ciphers are more secure than the CBC mode ciphers. Among ciphers of the same mode, the higher the key size, the more secure the cipher. Of the ciphers supported by Data ONTAP, **aes256-ctr** is the most secure, and **3des-cbc** is the least secure.

You can manage the SSH key exchange algorithms and ciphers for SVMs in the following ways:

- Display the current configurations of SSH key exchange algorithms and ciphers (`security ssh show`)
The enabled SSH key exchange algorithms are displayed in the order of decreasing security strengths.
The enabled CTR mode ciphers (more secure) are displayed before the CBC mode ciphers (less secure). Within each mode type, the ciphers are displayed in decreasing key size.
- Replace the current configurations of the SSH key exchange algorithms or ciphers with the configuration settings you specify (`security ssh modify`)
- Add SSH key exchange algorithms or ciphers to the current configurations (`security ssh add`)
The added SSH key exchange algorithms or ciphers are enabled.
- Remove the specified SSH key exchange algorithms or ciphers from the current configurations (`security ssh remove`)
The removed SSH key exchange algorithms or ciphers are disabled.
Data ONTAP prevents you from removing all SSH key exchange algorithms or all ciphers from the SVM.

Commands for managing SSH security configurations

You use the `security ssh` commands to manage the SSH security configurations of the Storage Virtual Machine (SVM), including displaying, replacing, adding, and removing the SSH key exchange algorithms and data encryption algorithms (ciphers).

If you want to...	Use this command...
Display the current configurations of the SSH key exchange algorithms and ciphers for the SVM	<code>security ssh show</code>
Replace the current configurations of the SSH key exchange algorithms or ciphers with the configuration settings you specify	<code>security ssh modify</code>
Add SSH key exchange algorithms or ciphers to the current configurations	<code>security ssh add</code>
Remove the specified SSH key exchange algorithms or ciphers from the current configurations of the SVM	<code>security ssh remove</code>

Related information

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

Managing public keys

You can associate, modify, or delete a public key to manage a user's authentication.

You can manage public keys in the following ways:

- Adding a public key by associating an existing public key in a valid OpenSSH format with a user account
Multiple public keys are allowed for a user account.
- Loading a public key from a universal resource identifier (URI), such as FTP or HTTP, and associating it with a user account
You can also overwrite an existing public key with the one you are loading.
- Displaying information about public keys
- Modifying a public key that is associated with a specific user
- Deleting a public key that is associated with a specific user

To create or modify a public key or load a public key from a URI, your user account must be configured with the **publickey** login method.

You use the `security login publickey` commands to manage public keys. For information about these commands, see the appropriate man pages.

Commands for managing public keys

You use the `security login publickey` commands to manage public keys.

If you want to...	Use this command...
Associate an existing public key with a user account	<code>security login publickey create</code>
Load a public key from a URI and associate it with a user	<code>security login publickey load-from-uri</code>
Display information about public keys	<code>security login publickey show</code>
Modify a public key for a specific user	<code>security login publickey modify</code>
Delete a public key for a specific user	<code>security login publickey delete</code>

Related information

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

Managing digital certificates for server or client authentication

A digital certificate ensures that communications are transmitted in encrypted form and that information is sent privately and unaltered to only the specified server or from the authenticated client. You can generate a certificate signing request, create, install, sign, display, revoke, or delete a digital certificate for server or client authentication.

A digital certificate, also called a *public key certificate*, is an electronic document that verifies the owner of a public key. It can be either self signed (by the owner) or Certificate Authority (CA) signed. You can provide server or client authentication by using digital certificates for situations where the Storage Virtual Machine (SVM) is an SSL server or client. When you provide both server and client authentication, you have mutual authentication (also called *two-way authentication*) in which both the server and the client present their certificates to each other for validating their respective identities to each other.

You can manage digital certificates in the following ways (the `security certificate` command family):

- You can create and install a self-signed digital certificate.
- You can generate a digital certificate signing request (CSR) that will be sent to a CA for signing.
- You can sign a digital certificate using a self-signed root CA.
- You can install a CA-signed digital certificate and the public key certificate of the root CA.
- You can display information about created or installed digital certificates.
- You can display digital certificates that are signed by the SVM as the CA.
- You can revoke a digital certificate signed by the SVM as the CA, if the certificate becomes compromised.
- You can delete a digital certificate that is no longer needed.

The following behaviors and default settings apply:

- When the SVM is created, Data ONTAP automatically creates a self-signed digital certificate for authenticating the SVM as a server.
- By default, Data ONTAP uses the SHA256 cryptographic hashing function for signing a CSR or digital certificate.
- By default, private keys generated by Data ONTAP are 2048-bit.

- By default, digital certificates created by Data ONTAP are set to expire in 365 days, but you can specify the expiration setting when you create a digital certificate.
- By default, SSL server authentication is enabled, but SSL client authentication is disabled. The `security ssl modify` command enables or disables SSL authentication of the SVM as an SSL server and that of its client. The `-server-enabled` parameter defaults to `true`, and the `-client-enabled` parameter defaults to `false`. Setting the `-client-enabled` parameter to `true` enables mutual authentication of the server (the SVM) and its client.

When you manage digital certificates, you specify one of the following certificate types (the `-type` parameter of the `security certificate` command family) for server or client authentication:

- **server** is a certificate that authenticates the SVM as an SSL server.
- **client** is a certificate that authenticates the SVM as an SSL client.
- **server-ca** is a root certificate of an SSL server to which the SVM is a client.
- **client-ca** is a root certificate of an SSL client to which the SVM is a server.
- **root-ca** is a self-signed root CA certificate that enables the SVM to act as a CA. When you create a **root-ca** certificate, a **client-ca** certificate and a **server-ca** certificate are also created automatically. When you delete the **root-ca** certificate, the corresponding **client-ca** and **server-ca** certificates are also deleted automatically.

Installing a server certificate to authenticate the SVM as an SSL server

To enable the Storage Virtual Machine (SVM) to be authenticated as an SSL server, you install a digital certificate with the **server** type on the SVM. The certificate you install can be self signed or CA signed.

About this task

When the SVM is created, a self-signed server certificate is created automatically and uses the SVM name as the common name. The corresponding SSL server authentication is enabled and also uses the default common name for the SVM.

If you want the SVM to use a different common name or a CA-signed certificate for server authentication, you can create or install additional server certificates. You can also modify SSL configuration to use a server certificate that you specify.

Steps

1. To create a self-signed digital certificate for server authentication, use the `security certificate create` command with the `-type server` parameter.
2. To use a third-party CA-signed digital certificate for server authentication, complete the following steps:

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

The system displays the CSR output. The output includes a certificate request and a private key. You should keep a copy of the private key.

- b. Copy the certificate request from the CSR output and send it in an electronic form (such as email) to a trusted third-party CA for signing.

After processing your request, the CA sends you the signed digital certificate. You should keep a copy of the private key and the CA-signed digital certificate.

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

- d. Enter the certificate and the private key when you are prompted, and then press Enter.

- e. When Data ONTAP asks you whether you want to install the CA root and intermediate certificates that form the certificate chain of the server certificate, enter Y.

- f. Enter any additional root or intermediate certificates when you are prompted, and then press Enter

You install the certificates of the CA to form a certificate chain of the server certificate. The chain starts with the certificate of the CA that issued the server certificate, and it can range up to the root certificate of the CA. Any missing intermediate certificates will result in the failure of server certificate installation.

After the CA certificates are entered, the certificates chain is installed as `server-chain` along with the `server` certificate type.

3. To use a self CA-signed digital certificate for server authentication (with the SVM being the signing CA), complete the following steps:

- a. Generate a CSR by using the `security certificate generate-csr` command.

The system displays the CSR output. The output includes a certificate request and a private key. You should keep a copy of the private key.

- b. Create a self-signed root CA certificate for the SVM by using the `security certificate create` command with the `-type root-ca` parameter.

- c. Display the root CA certificate by using the `security certificate show` command with the `-instance` and `-type root-ca` parameters.

You will need the following information from the command output for signing the CSR:

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

- d. Sign the CSR with the root CA by using the `security certificate sign` command.

- e. When you are prompted, enter the CSR and then press ENTER.
 - f. Install the self CA-signed digital certificate by using the `security certificate install` command with the `-type server` parameter.
 - g. Enter the certificate and the private key when you are prompted, and then press Enter.
 - h. When Data ONTAP asks you whether you want to install the CA root and intermediate certificates that form the certificate chain of the server certificate, enter N.
4. If you want to modify the SSL configuration to specify the certificate for server authentication, use the `security ssl modify` command with the `-ca` and the `-serial` parameters.

Examples of installing a server certificate to authenticate the SVM as an SSL server

The following example creates a self-signed server certificate for the “vs1” SVM at a company whose custom common name is `lab.companyname.com`. The certificate is for authenticating the “vs1” SVM as an SSL server:

```
vs1::> security certificate create -common-name lab.companyname.com
-type server
```

The following command creates a CSR with a 2048-bit private key for use by the Software group in the IT department of a company whose custom common name is `server1.companyname.com`, located in Sunnyvale, California, USA. The email address of the contact administrator who manages the SVM is `web@companyname.com`. The system displays the CSR and the private key in the output:

```
vs1::> security certificate generate-csr -common-name
server1.companyname.com
-size 2048 -country US -state CA -locality Sunnyvale
-organization IT -unit Software -email-addr web@companyname.com

Certificate Signing Request:
-----BEGIN CERTIFICATE REQUEST-----
MIICrjCCAZYCAQMwATEQMA4GA1UEAxMHcnRwLmNvbTELMakGA1UEBhMCVVMxMCAJ
BgNVBAgTAk5DMQwwCgYDVQQHEwNSVFaxDTALBgNVBAoTBGNvcmUxDTALBgNVBAsT
BGNvcmUxDTANBkqghkiG9w0BCQEWADCCASIdQYJKoZiIhvcNAQEBBQADggEPADCC
...
-----END CERTIFICATE REQUEST-----

Private Key:
-----BEGIN RSA PRIVATE KEY-----
MIIBPAIBAAJBAMl6ytrK8nQj82UsWeHOeT8gk0BPX+Y5MLyCsUdXA7hXhumHNpvF
C6lX2G32Sx8VEa1th94tx+vOEzq+UaqHlt0CAwEAAQJBAMZjDwlgmlm3qIr/n8VT
PFnnZnbVcXVM70tbUsgPKw+QCCh9dF1jmuQKeDr+wUMWkn1DeGrfhILpzfJGHRlJ
...
-----END RSA PRIVATE KEY-----
```

Note: Please keep a copy of your private key and certificate request for future reference.

The following command installs a CA-signed server certificate for the “vs1” SVM. The certificate is for authenticating the “vs1” SVM as an SSL server:

```
vs1::> security certificate install -type server
Please enter Certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIIB8TCCAZugAwIBAwIBADANBgkqhkiG9w0BAQQFADBfMRMwEQYDVQQDEwpuZXRh
cHhAuY29tMQswCQYDVQQGEwJVUzEJMAcGALUECBMAMQkwBwYDVQQHEwAxCTAHBgNV
BAoTADAEJMAcGALUECXMAMQ8wDQYJKoZIhvcNAQkBFgAwHhcNMTAwNDI2MTk0OTI4
...
-----END CERTIFICATE-----

Please enter Private Key: Press <Enter> when done
-----BEGIN RSA PRIVATE KEY-----
MIIBPAIBAAJBAM16ytrK8nQj82UsWeHOeT8gk0BPX+Y5MLyCsUdXA7hXhumHNpvF
C61X2G32Sx8VEalth94tx+vOEzq+UaqHlt0CAwEAAQJBAMZjDwlgmlm3qIr/n8VT
PFnnZnbVcXVM70tbUsgPKw+QCCh9dF1jmuQKeDr+wUMWknlDeGrfhILpzfJGHRlJ
...
-----END RSA PRIVATE KEY-----

Please enter certificates of Certification Authorities (CA) which form the
certificate chain of the server certificate. This starts with the issuing
CA certificate of the server certificate and can range up to the root CA
certificate.

Do you want to continue entering root and/or intermediate certificates {y|
n}: y

Please enter Intermediate Certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIIE+zCCBGsgAwIBAgICAQ0wDQYJKoZIhvcNAQEFBQAwgbsxJDAiBgNVBACGTG1Zh
bG1DZXJ0IFZhbG1kYXRpb24gTmV0d29yazEXMBUGALUEChMOVmFsaUNlcnQsIElu
Yy4xNTAzBgNVBAsTLFZhbG1DZXJ0IENsYXNzIDFgUG9saWN5IFZhbG1kYXRpb24g
...
-----END CERTIFICATE-----

Do you want to continue entering root and/or intermediate certificates {y|
n}: n

Note: You should keep a copy of your certificate and private key for future
reference.
If you revert to an earlier release, the certificate and private key are
deleted.
```

Installing a client CA or root CA certificate to authenticate an SSL client of the SVM

To enable the Storage Virtual Machine (SVM) to authenticate a client that wants to access it, you can install a digital certificate with the **client-ca** type on the SVM for the root certificate of the CA


```
-----END RSA PRIVATE KEY-----
```

Note: Please keep a copy of your certificate request and private key for future reference.

Data ONTAP displays the certificate request and private key and reminds you to copy them to a file for future reference.

4. If you self-sign the CSR, complete the following steps:
 - a. Display the root CA certificate you created in Step 1 by using the `security certificate show` command with the `-instance` and `-type root-ca` parameters.

You will need the following information from the command output for signing the CSR:

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

Example

```
vs1::> security certificate show -instance -type root-ca
      Vserver: vs1
      FQDN or Custom Common Name: lab.companyname.com
      Serial Number of Certificate: 50F84392
      Certificate Authority: lab.companyname.com
      Type of Certificate: root-ca
      Size of Requested Certificate(bits): 2048
      Certificate Start Date: Wed Jun 25 13:29:16 2014
      Certificate Expiration Date: Thu Jun 25 13:29:16 2015
      Public Key Certificate: -----BEGIN CERTIFICATE-----
                               MIID
+zCCAuOgAwIBAgIEUPhDkjANBgkqhkiG9w0BAQsFADBbMQ8wDQYDVQQDEwZt
      .
      .
      .
```

- b. Sign the CSR with the root CA by using the `security certificate sign` command.

The default format (`-format`) for the signed certificate is PEM. If you specify the format to be PKCS12, you can optionally specify the destination to upload the signed certificate by using the `-destination` parameter.
 - c. When you are prompted, enter the CSR and then press ENTER.

Example

```
vs1::> security certificate sign -ca lab.companyname.com -ca-
serial 50F84392
Please enter Certificate Signing Request (CSR): Press <enter> when
done
```

```

-----BEGIN CERTIFICATE REQUEST-----
MIICrTCCAZUCAQAwaDEcMBoGA1UEAxMTQ1NSLlNpZ25pbmdUZXR0LmNvbTELMakG
A1UEBhMCVVMxCTAHBgNVBAGTADEJMAcGA1UEBxMAMQkwBwYDVQQKEwAxCTAHBgNV
BASATADEPMA0GCSqGSIb3DQEJARYAMIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIB
...
-----END CERTIFICATE REQUEST-----

Signed Certificate: :
-----BEGIN CERTIFICATE-----
MIIDmzCCAoOgAwIBAgIEU9e2rzANBgkqhkiG9w0BAQsFADBoMRwwGgYDVQQDExNO
ZXcuQ29tcGFueU5hbWUuY29tMQswCQYDVQQGEwJVUzEJMAcGA1UECBMAMQkwBwYD
VQQHEwAxCTAHBgNVBAoTADEJMAcGA1UECzMAMQ8wDQYJKoZIhvcNAQkBFgAwHhcN
...
-----END CERTIFICATE-----

```

The signed certificate is displayed. You should keep a copy of the certificate.

5. If you have a third-party CA sign the CSR, complete the following steps:
 - a. Send the certificate request from the CSR output (Step 3) in an electronic form (such as email) to a trusted CA for signing.

After processing your request, the CA sends you the signed digital certificate. You should keep a copy of the private key and the CA-signed certificate for future reference.

- b. On the SVM, install the root certificate and each intermediate certificate of the CA that signed the certificate by using the `security certificate install` command with the `-type client-ca` parameter.

Example

```

vsl:~> security certificate install -type client-ca

Please enter Certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIIDNjCCAp+gAwIBAgIQNhIilsXjOKUgodJfTncJVDANBgkqhkiG9w0BAQUFADCB
zjELMAkGA1UEBhMCWkExFTATBgNVBAGTDFdlc3Rlcm4gQ2FwZTESMBAGA1UEBxMJ
Q2FwZSBub3duMR0wGwYDVQQKEwRlUaGF3dGUgQ29uc3VsdGluZyBjYzEoMCYGA1UE
...
-----END CERTIFICATE-----

You should keep a copy of the CA-signed digital certificate for
future reference.

```

6. Provide the self-signed or CA-signed certificate for the user to install on the client.
7. Repeat Step 3 to Step 6 for each client you want to authenticate.

8. If an SVM user is not set up to be authenticated by digital certificates, contact the cluster administrator to have the user account set up for digital certificate authentication.

For SVM user accounts, digital certificate authentication is supported only with the `ontapi` access method.

Installing a server CA certificate to authenticate an SSL server to which the SVM is a client

Sometimes the Storage Virtual Machine (SVM) is a client to another SSL server (which, for example, can be an Active Directory domain controller that supports LDAP over SSL). In this case, you can enable the SVM to authenticate the SSL server by installing the server's root certificate with the `server-ca` type on the SVM.

Before you begin

You must have the root certificate of the SSL server. The root certificate can be self signed by the server or signed by a third-party CA for the server.

Steps

1. Install the root certificate provided by the SSL server by using the `security certificate install` command with the `-type server-ca` parameter.
2. When you are prompted, enter the certificate, and then press Enter.

Data ONTAP reminds you to keep a copy of the certificate for future reference.

Example of installing a server CA certificate of an SSL server

The following example installs an SSL server's CA certificate with the `server-ca` type. The certificate is used for server authentication and is installed on the “vs1” SVM, which serves as a client to the server:

```
vs1::> security certificate install -type server-ca
Please enter Certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIIDNjCCAp+gAwIBAgIQNhIilsXjOKUgodJfTncJVdANBgkqhkiG9w0BAQUFADCB
zjELMAkGA1UEBhMCWkExFTATBgNVBAGTDfDl3R1cm4gQ2FwZTESMBAGA1UEBxMJ
Q2FwZSBub3duMR0wGwYDVQQKEXRuAGF3dGUgQ29uc3VsdGluZyBjYzEoMCMYGA1UE
...
-----END CERTIFICATE-----
```

You should keep a copy of the CA-signed digital certificate for future reference.

Installing a client certificate to authenticate the SVM as an SSL client

To enable an SSL server to authenticate the Storage Virtual Machine (SVM) as an SSL client, you install a digital certificate with the `client` type on the SVM. Then you provide the `client-ca` certificate to the SSL server administrator for installation on the server.

Before you begin

You must have already installed the root certificate of the SSL server on the SVM with the `server-ca` certificate type.

Steps

1. To use a self-signed digital certificate for client authentication, use the `security certificate create` command with the `-type client` parameter.
2. To use a CA-signed digital certificate for client authentication, complete the following steps:
 - a. Generate a digital certificate signing request (CSR) by using the `security certificate generate-csr` command.

Data ONTAP displays the CSR output, which includes a certificate request and private key, and reminds you to copy the output to a file for future reference.
 - b. Send the certificate request from the CSR output in an electronic form (such as email) to a trusted CA for signing.

After processing your request, the CA sends you the signed digital certificate. You should keep a copy of the private key and the CA-signed certificate for future reference.
 - c. Install the CA-signed certificate by using the `security certificate install` command with the `-type client` parameter.
 - d. Enter the certificate and the private key when you are prompted, and then press Enter.
 - e. Enter any additional root or intermediate certificates when you are prompted, and then press Enter.

You install an intermediate certificate on the SVM if a certificate chain that begins at the trusted root CA, and ends with the SSL certificate issued to you, is missing the intermediate certificates. An intermediate certificate is a subordinate certificate issued by the trusted root specifically to issue end-entity server certificates. The result is a certificate chain that begins at the trusted root CA, goes through the intermediate, and ends with the SSL certificate issued to you.
3. Provide the `client-ca` certificate of the SVM to the administrator of the SSL server for installation on the server.

The security certificate show command with the `-instance` and `-type client-ca` parameters displays the `client-ca` certificate information.

Examples of installing a client certificate to authenticate the SVM as an SSL client

The following example creates a self-signed client certificate for the “vs1” SVM at a company whose custom common name is `lab.companyname.com`. The certificate is for authenticating the “vs1” SVM as an SSL client:

```
vs1::> security certificate create -common-name lab.companyname.com
-type client
```

The following command creates a CSR with a 2048-bit private key for use by the Software group in the IT department of a company whose custom common name is `lab.companyname.com`, located in Sunnyvale, California, USA. The email address of the contact administrator who manages the SVM is `web@companyname.com`. The system displays the CSR and the private key on the console:

```
vs1::> security certificate generate-csr -common-name lab.companyname.com
-size 2048 -country US -state CA -locality Sunnyvale -organization IT
-unit Software -email-addr web@companyname.com
```

Certificate Signing Request:

```
-----BEGIN CERTIFICATE REQUEST-----
MIICrjCCAZYCAQMwaTEQMA4GAlUEAxMHcnRwLmNvbTElMAkGA1UEBhMCVVMxMzA1
BgNVBAGTAK5DMQwwCgYDVQQHEwNSVFAXDTALBgNVBAoTBGNvcmUxDTALBgNVBAsT
BGNvcmUxZDZANBgkqhkiG9w0BCQEWADCCASiWDQYJKoZIhvcNAQEBBQADggEPADCC
...
-----END CERTIFICATE REQUEST-----
```

Private Key:

```
-----BEGIN RSA PRIVATE KEY-----
MIIBPAIBAAJBAMl6ytrK8nQj82UsWeHOeT8gk0BPX+Y5MLyCsUdXA7hXhumHNpvF
C6lX2G32Sx8VEalth94tx+vOEzq+UaqHlt0CAwEAAQJBAMZjDwlglmlm3qIr/n8VT
PFnnZnbVcXVM70tbUsgPKw+QCCh9dF1jmuQKeDr+wUMWknlDeGrfhILpzfJGHRlJ
...
-----END RSA PRIVATE KEY-----
```

Note: Please keep a copy of your private key and certificate request for future reference.

The following command installs a CA-signed client certificate for the “vs1” SVM. The certificate is for authenticating the “vs1” SVM as an SSL client:

```
vs1::> security certificate install -type client
Please enter Certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIIB8TCCAZugAwIBAwIBADANBgkqhkiG9w0BAQQFADBfMRMwEQYDVQQDEwpuZXRh
cHAuY29tMQswCQYDVQQGEwJVUzEUMAcGAlUECBMAMQkwBwYDVQQHEwAxcTAHBGNV
BAoTAEJMAcGAlUECXMAMQ8wDQYJKoZIhvcNAQkBFgAwHhcNMTAwNDI2MTk0OTI4
...
```

```
-----END CERTIFICATE-----
```

```
Please enter Private Key: Press <Enter> when done
```

```
-----BEGIN RSA PRIVATE KEY-----
```

```
MIIBPAIBAAJBAM16ytrK8nQj82UsWeHOeT8gk0BPX+Y5MLyCsUdXA7hXhumHNpvF
C61X2G32Sx8VEalTh94tx+vOEzq+UaqHlt0CAwEAAQJBAMZjDWlgmlm3qIr/n8VT
PFnnZnbVcXVM70tbUsgPKw+QCCh9dFljmuQKeDr+wUMWknlDeGrfhILpzfJGhrLJ
```

```
...
```

```
-----END RSA PRIVATE KEY-----
```

Please enter certificates of Certification Authorities (CA) which form the certificate chain of the client certificate. This starts with the issuing CA certificate of the client certificate and can range up to the root CA certificate.

Do you want to continue entering root and/or intermediate certificates {y|n}: y

```
Please enter Intermediate Certificate: Press <Enter> when done
```

```
-----BEGIN CERTIFICATE-----
```

```
MIIE+zCCBGSgAwIBAgICAQ0wDQYJKoZIhvcNAQEFBQAwgbsxJDAiBgNVBAcTG1Zh
bG1DZXJ0IFZhbGlkYXRpb24gTmV0d29yazEXMBUGAlUEChMOVmFsaUNlcnQsIElu
Yy4xNTAzBqNVBAsTLFZhbG1DZXJ0IENsYXNzIDIGUG9saWN5IFZhbGlkYXRpb24g
```

```
...
```

```
-----END CERTIFICATE-----
```

Do you want to continue entering root and/or intermediate certificates {y|n}: n

Note: You should keep a copy of your certificate and private key for future reference.

If you revert to an earlier release, the certificate and private key are deleted.

Replacing an expired digital certificate

Each certificate that you create or install has an expiration date. When it expires, you must replace it with a new certificate so that the corresponding server or client authentication is not disrupted.

About this task

By default, digital certificates created by Data ONTAP are set to expire in 365 days, but you can specify the expiration setting when you create a digital certificate.

Steps

1. Display certificate expiration information by using the `security certificate show` command with the `-fields expiration, expire-days` parameter.

You need the following information when you delete an expired certificate:

- The common name used for the certificate

- Serial number
 - Certificate authority (CA)
 - Certificate type
2. Delete an expired certificate by using the `security certificate delete` command.
 3. Obtain a new certificate with the same common name to replace the certificate that has expired:

If the certificate is this type...	Then follow the steps in...
<code>server</code>	<i>Installing a server certificate to authenticate the cluster or SVM as an SSL server</i> on page 26
<code>client-ca</code>	<i>Installing a client CA or root CA certificate to authenticate an SSL client of the cluster or SVM</i> on page 29
<code>server-ca</code>	<i>Installing a server CA certificate to authenticate an SSL server to which the cluster or SVM is a client</i> on page 33
<code>client</code>	<i>Installing a client certificate to authenticate the cluster or SVM as an SSL client</i> on page 34

Commands for managing digital certificates

You use the `security certificate` commands to manage digital certificates of the Storage Virtual Machine (SVM).

If you want to...	Use this command...
Create and install a self-signed digital certificate with one of the following types: <ul style="list-style-type: none"> • <code>server</code> • <code>root-ca</code> • <code>client</code> 	<code>security certificate create</code>
Generate a digital certificate signing request that you will send to a CA for signing	<code>security certificate generate-csr</code>
Sign a digital certificate using a self-signed root CA	<code>security certificate sign</code>

If you want to...	Use this command...
Install a CA-signed digital certificate and the public key certificate of the root CA with one of the following types: <ul style="list-style-type: none"> • <code>server</code> • <code>client-ca</code> • <code>server-ca</code> • <code>client</code> 	<code>security certificate install</code>
Display information about installed digital certificates	<code>security certificate show</code>
Display digital certificates that are signed by the SVM as the CA	<code>security certificate ca-issued show</code>
Revoke a compromised digital certificate signed by the SVM as the CA	<code>security certificate ca-issued revoke</code>
Delete an installed digital certificate	<code>security certificate delete</code>

Related information

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

Managing SSL

The SSL protocol improves the security of web access by using a digital certificate to establish an encrypted connection between a web server and a browser.

You can manage SSL for a Storage Virtual Machine (SVM) in the following ways:

- Enabling SSL
- Generating and installing a digital certificate and associating it with the SVM
- Displaying the SSL configuration to see whether SSL has been enabled, and, if available, the SSL certificate name

Commands for managing SSL

You use the `security ssl` commands to manage the SSL protocol for a Storage Virtual Machine (SVM).

If you want to...	Use this command...
Enable SSL for an SVM, and associate a digital certificate with it	<code>security ssl modify</code>
Display the SSL configuration and certificate name for an SVM	<code>security ssl show</code>

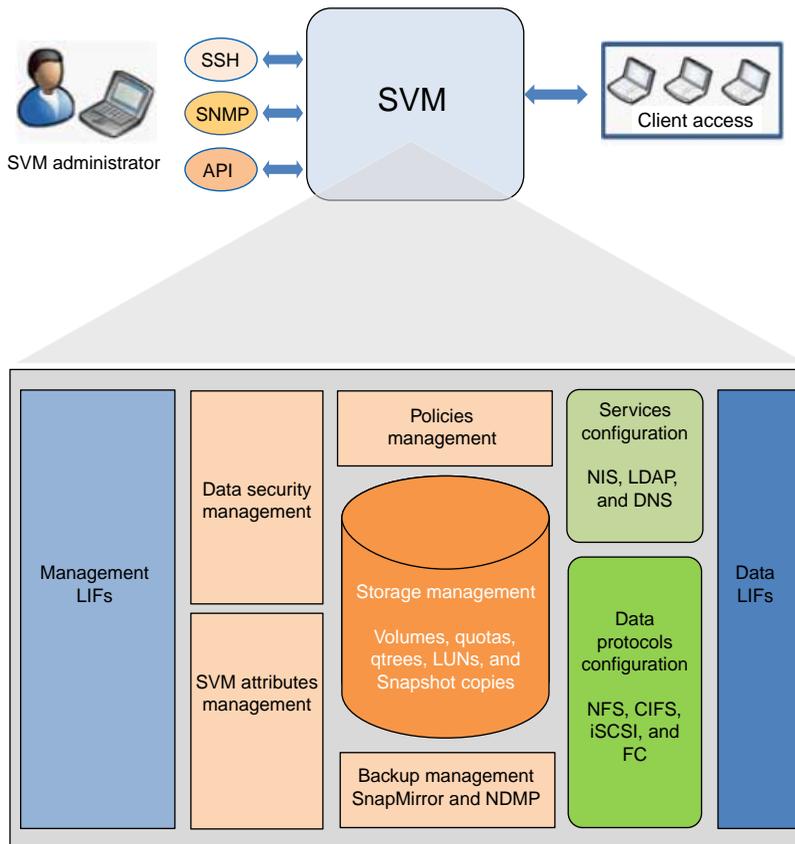
Related information

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

Administering SVMs

Depending on the capabilities assigned by the cluster administrator, an SVM administrator can perform various administration tasks on a Storage Virtual Machine (SVM, formerly known as Vserver). After logging in to the SVM, an SVM administrator can identify the capabilities assigned and the commands that are available for the administration.

The following illustration depicts the SVM administrative components:



You might have all or some of the following administration capabilities:

- Jobs and schedules management
You can manage jobs and schedules related to the SVM.
- Data access protocol configuration
You can configure data access protocols, such as NFS, CIFS, iSCSI, and Fibre Channel (FC) protocol (Fibre Channel over Ethernet included).

- **Policy management**
You can create and manage policies to manage data access from the SVM.
- **Data access security management**
You can set security on the SVM's data without the need of a client.
- **Services configuration**
You can configure services, such as LDAP, NIS, and DNS.
- **Storage management**
You can manage volumes, quotas, qtrees, and files.
- **LUN management**
You can manage LUNs in a SAN environment.
- **Backup management**
You can back up and manage the SVM's data by using SnapMirror technology and NDMP.
- **Monitoring SVM**
You can monitor performance data, network connection, information, and SVM health.

Note: For troubleshooting or modifying SVM configurations, SVM administrators must contact the cluster administrator.

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

Identifying the commands that you can execute

The capabilities to administer an SVM and its resources depend on the capabilities of the user who logs in. After you log in as an SVM administrator, you can identify the commands that you can execute on the SVM.

Steps

1. To identify the available commands, enter the following command:

?

The list of available commands is displayed.

2. To identify the available subcommands within a command, perform the following steps:

- a. Enter the name of the command directory.

- b. At the prompt, enter the following command:

?

The list of available subcommands is displayed.

Example

The following example shows the commands and the volume subcommands that are available for an SVM administrator in the Storage Virtual Machine (SVM, formerly known as Vserver) vs1.example.com:

```
vs1.example.com::> ?
    up                Go up one directory
    dashboard>       Display dashboards
    exit              Quit the CLI session
    .
    .
    .
    volume>          Manage virtual storage, including volumes,
                    snapshots, and mirrors
    vservers>        Manage Vservers
```

```
vs1.example.com::>volume
vs1.example.com::volume> ?
    autosize          Set the autosize settings of the
                    flexible volume.
    clone>            Manage FlexClones
    .
    .
    .
    snapshot>        Manage snapshots
    unmount           Unmount a volume
```

Displaying ONTAP APIs

As an SVM administrator, you can view the Data ONTAP APIs and their corresponding CLI commands by using the `security login role show-ontapi` command to execute administrative functions with a remote program.

Step

1. Use the `security login role show-ontapi` to view the Data ONTAP APIs and their corresponding CLI commands.

Example

The following example shows how to view the Data ONTAP APIs and their corresponding CLI commands for the SVM vs1.example.com:

```

vs1.example.com:~> security login role show-ontapi
ONTAPI
-----
av-get-remedy-info          antivirus remedy show
av-on-access-policy-create  antivirus on-access policy create
av-on-access-policy-delete  antivirus on-access policy delete
av-on-access-policy-get     antivirus on-access policy show
...
...
...
waf1-get-sync-status       volume show
waf1-sync                  volume modify
554 entries were displayed.

```

Managing jobs and schedules

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

Commands for managing jobs

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

If you want to...	Use this command...
Display information about all jobs	<code>job show</code>
Display information about jobs on a per-node basis	<code>job show bynode</code>
Display information about cluster-affiliated jobs	<code>job show-cluster</code>
Display information about completed jobs	<code>job show-completed</code>
Display information about job history	<p><code>job history show</code></p> <p>Up to 25,000 job records are stored for each node in the cluster. Consequently, attempting to display the full job history could take a long time. To avoid potentially long wait times, you should display jobs by node, Storage Virtual Machine (SVM), or record ID.</p>
Display the list of private jobs	<p><code>job private show</code></p> <p>(advanced privilege level)</p>

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

Related information

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

Commands for managing job 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 view job schedules.

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

If you want to...	Use this command...
Display information about all schedules	<code>job schedule show</code>
Display information about cron schedules	<code>job schedule cron show</code>
Display information about interval schedules	<code>job schedule interval show</code>

Related information

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

Monitoring SVM performance

You can view data about your Storage Virtual Machines (SVMs) to monitor SVM performance. For example, you can monitor the performance of volumes by viewing statistics that show throughput and latency.

What objects, instances, and counters are

You can view performance data for specific objects in your cluster. Objects are comprised of instances and counters. Counters provide data about the instances of an object.

An object is any of the following:

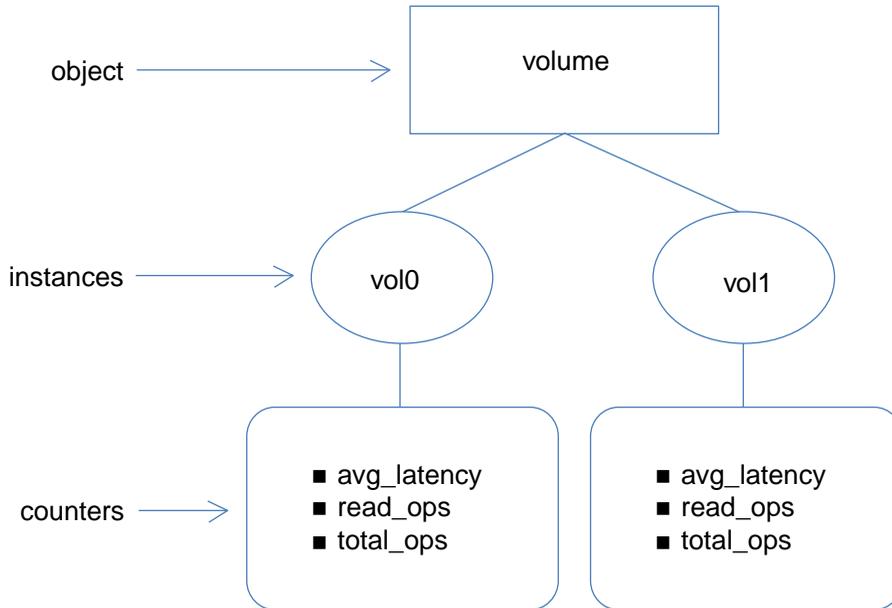
- Logical entities such as LUNs and volumes
- Protocols such as CIFS and NFS

Each object has zero or more instances. For example, the LUN object has an instance for each LUN in your cluster.

A counter is a predefined performance metric that provides data about an object. Examples of data that counters provide include the following:

- The average latency for a volume
- The number of established SMB and SMB2 sessions

The following illustration shows the relationship between an object and its instances and counters. In this illustration, the volume object has two instances: vol0 and vol1. The object's counters provide data about each of these instances. The illustration shows three of the object's counters: avg_latency, read_ops, and total_ops.



Decisions to make before you view performance data

You can view performance data in several ways. You should make a few decisions before you view the data.

You should decide the following before you view performance data:

Decision	Considerations
How do you want to retrieve and display the data?	<p>You have two choices:</p> <ul style="list-style-type: none"> You can collect and view a set of data for a specific time period. If you choose this option, you can view data for several objects and instances at a time. You can view continuously updated data. If you choose this option, you can view data for only one object and one instance at a time.
For which objects do you want to view data?	You need to specify at least one object for which you want to view data.
Do you want data from all counters or from specific counters?	The default setting shows data for all counters in an object; however, you can specify specific counters to get the exact data that you need.

Decision	Considerations
Do you want data for all instances of an object or for specific instances?	<ul style="list-style-type: none"> • If you collect data for a time period, the default setting shows data for all instances; however, you can specify one or more instances. • If you view continuously updated data and specify any object other than cluster, you must specify an instance.

Viewing performance data for a time period

You can monitor SVM performance by collecting and viewing data for a specific time period (a sample). You can view data for several objects and instances at a time.

About this task

You can collect more than one data sample at a time. You can collect more than one sample from the same object at the same time.

Note: You cannot collect and view data for an object that has more than 5,000 instances. If an object has more than 5,000 instances, you need to specify the specific instances for which you want data.

For more information about the `statistics` commands, see the man pages.

Steps

1. Use the `statistics start` command to start collecting data.

If you do not specify the `-sample-id` parameter, the command generates a sample identifier for you and defines this sample as the default sample for the CLI session. If you run this command during the same CLI session and do not specify the `-sample-id` parameter, the command overwrites the previous default sample.

2. Optional: Use the `statistics stop` command to stop collecting data for the sample.

You can view data from the sample if you do not stop data collection. Stopping data collection gives you a fixed sample. Not stopping data collection gives you the ability to get updated data that you can use to compare against previous queries. The comparison can help you identify performance trends.

3. Use the `statistics show` command to view the sample data.

Example: Monitoring NFSv3 performance

The following example shows performance data for the NFSv3 protocol.

The following command starts data collection for a new sample:

```
vs1::> statistics start -object nfsv3 -sample-id nfs_sample
```

The following command shows data from the sample by specifying counters that show the number of successful read and write requests versus the total number of read and write requests:

```
vs1::> statistics show -sample-id nfs_sample -counter read_total|
write_total|read_success|write_success
```

```
Object: nfsv3
Instance: vs1
Start-time: 2/11/2013 15:38:29
End-time: 2/11/2013 15:38:41
Cluster: cluster1
```

Counter	Value
-----	-----
read_success	40042
read_total	40042
write_success	1492052
write_total	1492052

Viewing continuously updated performance data

You can monitor SVM performance by viewing data that continuously updates with the latest status. You can view data for only one object and one instance at a time.

About this task

For more information about the `statistics show-periodic` command, see the man page.

Step

1. Use the `statistics show-periodic` command to view continuously updated performance data.

If you do not specify the `-object` parameter, the command returns summary data for the cluster.

Example: Monitoring volume performance

This example shows how you can monitor volume performance. For example, you might want to monitor volume performance if critical applications run on those volumes. Viewing the performance data can help you answer questions such as:

- What is the average response time for a volume?

- How many operations are completing per second?

The following command shows performance data for a volume by specifying counters that show the number of operations per second and latency:

```
vs1::> statistics show-periodic -object volume -instance vol0 -
counter write_ops|read_ops|total_ops|read_latency|write_latency|
avg_latency
cluster1: volume.vol0: 1/7/2013 20:15:51
  avg      read      total      write      write
  latency  latency read_ops   ops        latency    ops
  -----  -
  202us    218us    0          22         303us     7
  97us     43us     31         71         149us     34
  39us     0us      0          3          0us       0
  152us    0us      0          16         152us     16
  162us    0us      0          342        144us     289
  734us    0us      0          15         0us       0
  49us     0us      0          1          0us       0
cluster: volume.vol0: 1/7/2013 20:16:07
  avg      read      total      write      write
  latency  latency read_ops   ops        latency    ops
  -----  -
Minimums:
  39us     0us      0          1          0us       0
Averages for 7 samples:
  205us    37us     4          67         106us     49
Maximums:
  734us    218us    31         342        303us     289
```

Commands for monitoring SVM performance

You can use the `statistics` commands to display performance data and specify the settings for displaying the data. For more information about these commands, see the man pages.

Collecting data for a sample time period

You can use the following commands to collect data samples and to manage the samples that you collect. You must collect a data sample before you can use the `statistics show` command.

If you want to...	Use this command...
Start data collection for a sample	<code>statistics start</code>
Stop data collection for a sample	<code>statistics stop</code>
View all samples	<code>statistics samples show</code>
Delete a sample	<code>statistics samples delete</code>

Viewing performance data

You can use the following commands to view performance data. You must collect a data sample before you can use the `statistics show` command.

If you want to...	Use this command...
View performance data for a sample time period	<code>statistics show</code> You should limit the scope of this command to only a few objects at a time to avoid a potentially significant impact on system performance.
View continuously updated performance data	<code>statistics show-periodic</code>

Viewing all objects, instances, and counters

Use the `statistics catalog` commands to view information about objects, instances, and counters.

If you want to...	Use this command...
View descriptions of objects	<code>statistics catalog object show</code>
View all instances of an object	<code>statistics catalog instance show</code>
View descriptions of counters in an object	<code>statistics catalog counter show</code>

Displaying information about SVMs

SVM administrators can view the details of Storage Virtual Machine (SVM, formerly known as Vserver) that are assigned by using the `vserver show` command.

Step

1. Enter the appropriate command to view details of the SVM:

If you want to...	Enter the following command...
View basic information about the SVM	<code>vserver show</code>
View detailed information about the SVM	<code>vserver show -instance</code>

Example

The following command displays basic information about the SVM:

```
vs2.example.com:> vserver show
```

Vserver	Type	Subtype	Admin State	Operational State	Root Volume	Aggregate
vs2.example.com		data		default		running
running			root_vol2	aggr1		

The following command displays detailed information about the SVM:

```
vs2.example.com:> vserver show -instance
```

```

Vserver Type: data
Vserver Subtype: default
Vserver UUID: ca34e6b2-ddec-11df-b066-123478563412
Root Volume: root_vol2
.
.
Config Lock: false
IPspace Name: Default
Is Vserver Protected: false

```

Displaying information about SVM peer relationships

Peer Storage Virtual Machines (SVMs) are fully functional SVMs which could be either local or remote. Cluster administrators and SVM administrators can view the peers of the SVM to set up peering applications such as SnapMirror between volumes of the peer SVMs by using the `vserver peer show` command.

About this task

You can also view the status of the SVM peer relationships.

Step

1. Use the `vserver peer show` command to view the peered SVMs and the state of the SVM peer relationship.

Example

The following example shows how to view the information about peered SVMs:

```
vs1.example.com:> vserver peer show
```

Vserver	Peer Vserver	Peer State
vs1.example0.com	vs5.example0.com	peered
vs1.example0.com	vs3.example0.com	peered


```
vs1.example.com lif1 up/up 192.0.2.65/126 node0 e1b false
                lif2 up/up 192.0.2.1/62  node1 e0d false

2 entries were displayed.
```

Monitoring SVMs using dashboard

You can monitor the critical aspects of the SVM, such as the health of the SVM and its volumes, aggregates, network interfaces, ports, and protocols from the dashboard to ensure that the SVM is functional, and data access is nondisruptive.

Related information

[Clustered Data ONTAP 8.3 System Administration Guide for Cluster Administrators](#)

Commands for managing dashboards

The `dashboard` commands are deprecated, but you can still use them to configure dashboards, display dashboard information, and display health status for SVMs.

Note: The `dashboard health vserver` commands support the NFS and CIFS protocols. They do not support the FC and iSCSI protocols.

If you want to...	Use this command...
Display information about general SVM health, including the current operational status, issues, critical alerts, warnings, and informational messages	<code>dashboard health vserver show</code>
Display the health status of aggregates, LIFs, ports, protocols, and volumes in SVMs	<code>dashboard health vserver show-combined</code>
Display the health status of aggregates in SVMs	<code>dashboard health vserver show-aggregate</code>
Display the health status of volumes in SVMs	<code>dashboard health vserver show-volume</code>
Display the health status of LIFs in SVMs	<code>dashboard health vserver show-lif</code>
Display the health status of SVM network ports	<code>dashboard health vserver show-port</code>
Display the health status of protocols in SVMs	<code>dashboard health vserver show-protocol</code>

For more information, see the man pages.

Data access protocols configuration

As an SVM administrator, you can configure an SVM with FlexVol volumes with any combination of supported data access protocols, which are NFS, CIFS, iSCSI, and FC (FCoE included) to serve data. However, you can configure only NFS and CIFS protocols on an SVM with Infinite Volume.

You can configure and manage the following protocols:

- NFS and CIFS protocols for file-level data access.
- iSCSI and FC (FCoE included) protocols for block-level data access.

Note: You can configure and manage only the protocols that are allowed on the SVM by the cluster administrator.

NAS protocols

NFS clients can access data on an SVM by using the NFS protocol. You must configure an NFS server on an SVM to provide data access to its NFS clients. You can set up authentication between the SVM and NFS clients by configuring a network authentication protocol, such as NIS and LDAP.

CIFS clients can access data on an SVM by using the CIFS protocol. You can create multiple CIFS shares for the clients. You can set up authentication between the SVM and CIFS clients by configuring a network authentication protocol, such as Windows Active Directory.

In addition to NFS and CIFS protocols, you can also manage the following:

- Name mappings
You can create and use name mappings to map your UNIX users and groups to Windows users and groups or Windows users and groups to UNIX users and groups.
- Export policies
You can create and use export policies to restrict access to volumes or qtrees for specific clients.
- Locks
You can view and break a lock if it prevents a client's access to the files.

SAN protocols

You must configure the iSCSI protocol on an SVM to export LUNs and transfer block data to the iSCSI initiator hosts.

You must configure the FC (FCoE included) protocol on an SVM to export LUNs and transfer block data to the FC initiator hosts.

Related information

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

Clustered Data ONTAP 8.3 File Access Management Guide for CIFS
Clustered Data ONTAP 8.3 SAN Administration Guide

Commands for configuring data access protocols

SVM administrators can identify the list of commands to configure protocols by navigating to the respective command directories.

To identify the list of commands to configure NAS and SAN protocols, you must navigate to the protocol directory under vserver subdirectory.

Example

The following example shows how to identify the list of NFS protocol commands:

```
vs1.example.com::vserver> ?
audit>          Manage auditing of protocol requests that the
                 Vserver services
cifs>           Manage the CIFS configuration of a Vserver
dashboard>     The dashboard directory
data-policy>   Manage data policy
export-policy> Manage export policies and rules
fcp>          Manage the FCP service on a Vserver
fpolicy>      Manage FPolicy
group-mapping> The group-mapping directory
iscsi>        Manage the iSCSI services on a Vserver
locks>       Manage Client Locks
name-mapping> The name-mapping directory
nfs>         Manage the NFS configuration of a Vserver
peer>        Create and manage Vserver peer relationships
security>    Manage ontap security
services>   The services directory
show        Display Vservers
smtape>     The smtape directory

vs1.example.com::vserver nfs> ?
create      Create an NFS configuration for a Vserver
delete     Delete the NFS configuration of a Vserver
kerberos-config> Manage the Kerberos configuration for an NFS
            server
modify     Modify the NFS configuration of a Vserver
off        Disable the NFS service of a Vserver
on         Enable the NFS service of a Vserver
show      Display the NFS configurations of Vservers
start     Start the NFS service of a Vserver
status    Display the status of the NFS service of a
            Vserver
stop      Stop the NFS service of a Vserver
```

Data security management

As an SVM administrator, you can view and set security on a file or a directory from an SVM without using a client. You can apply security over large directories without significant degradation in performance.

When you set security on a file or a directory from the SVM, you are managing the security settings locally and not from remote clients thus reducing the performance degradation.

A set of security commands acts as a centralized security management tool on the SVM that can handle both CIFS and NFS security information.

You can perform the following tasks to manage security on a file or directory of an SVM:

- Applying files and directory security settings defined in a security policy to an SVM
- Displaying a list of file security jobs
- Managing NTFS file security policies
- Managing file security policies
- Displaying security information of a file or folder

Related information

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

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

Commands for setting up security settings on files and managing tracing

SVM administrators can identify the list of commands to set up security on files and tracing by navigating to the respective command directory.

To identify the list of commands available regarding file security and tracing, you must navigate to the security directory under vservers subdirectory.

Example

The following example shows how to identify the file security and tracing commands:

```
vs1.example.com::vserver security> ?
file-directory>          Manage file security
trace>                   Manage security tracing

vs1.example.com::vserver security file-directory> ?
apply                    Apply files and directory security settings
                        defined in a security policy to a Vserver
```

job>	Manage file security jobs
ntfs>	Manage NTFS file security policies
policy>	Manage file security policies
show	Display file/folder security information

Services configuration

As an SVM administrator, you can configure services such as Network Information Service (NIS), Domain Name Service (DNS), and Lightweight Directory Access Protocol (LDAP) for an SVM. You can configure these services to provide network directory information, authentication, and UNIX compatibility.

Note: The Active Directory service is configured as part of CIFS protocol configuration.

You can configure and manage the following services:

- Network Information Service (NIS)
You can configure NIS domains on an SVM to provide network information and authentication for the data access and management requests.
- Domain Name Service
You can configure DNS servers on an SVM for host-name resolution.
- LDAP services
You can configure LDAP services on an SVM to provide network information and authentication for the data access and management requests.
- Local UNIX users
You can set up UNIX user accounts on an SVM to provide an authentication mechanism for NFS access.
- Local UNIX groups
You can set up local UNIX groups on an SVM along with local UNIX users.
- Local user and groups for Windows
You can enable or disable local Windows users and groups for SMB access on an SVM.
- Netgroups
You can import UNIX netgroups from an FTP or HTTP site that is used by an SVM.

Related information

[NetApp Technical Report 4067: Clustered Data ONTAP NFS Best Practice and Implementation Guide](#)

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

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

Commands for configuring services

SVM administrators can identify the list of commands for configuring the services on an SVM by navigating to the respective command directory.

To identify the list of commands available to configure services, you must navigate to the services directory under vserver subdirectory.

Example

The following example shows how to identify the services commands:

```
vsl.example.com::vserver services> ?
dns>                Manage DNS service
ldap>               Manage LDAP configuration
ndmp>              Manage vserver scoped NDMP
netgroup>          Manage local netgroups
nis-domain>        Manage Network Information Service domains
unix-group>        Manage local UNIX group accounts
unix-user>         Manage local UNIX user accounts
```

Storage management

Storage Virtual Machines (SVMs) represents the logical layer of data storage. SVMs can either contain one or more FlexVol volumes or a single Infinite Volume. The storage space available in an SVM is scalable, thus enabling SVM administrators to provision and manage storage in an SVM.

SVMs with FlexVol volumes can also have quotas and qtrees. SVMs with Infinite Volume cannot have quotas and qtrees. Therefore, you cannot perform the quotas and qtrees related tasks on SVMs with Infinite Volume.

Depending on your capabilities, you can perform the following tasks to manage volumes on an SVM:

- Creating, modifying, renaming, or deleting volumes

You can view the list of aggregates that are available to create volumes by using the `volume create` command with the `aggregate` option. The number of volumes you can create on the SVM is defined by the cluster administrator.

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

- Mounting or unmounting volumes
- Removing junctions from volumes
- Viewing volume status

- Creating quotas, qtrees, and files

Note: You cannot copy or move volumes between aggregates.

Depending on your capabilities, you can manage volume qtrees and volume quotas by performing the following tasks:

- Creating, modifying, renaming, or deleting qtrees
- Viewing qtree status and statistics
- Creating, modifying, renaming, or deleting quota policy and policy rules
- Viewing quota policy and policy rules

Related information

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

Commands for managing storage

SVM administrators can identify the list of commands for managing storage on an SVM by navigating to the respective command directory.

To identify the list of commands available to manage storage, you must navigate to the volume directory.

Example

The following example shows how to identify the storage commands:

```
vs1.example.com::volume> ?
autosize                Set/Display the autosize settings of the
                        flexible volume.
clone>                  Manage FlexClones
create                  Create a new volume
delete                  Delete an existing volume
file>                  File related commands
...
...
...
show-space              Display a list of volumes and their space usage
show-space-old          Display a list of volumes and their space usage
size                    Set/Display the size of flexible volume.
snapshot>              Manage snapshots
unmount                 Unmount a volume
```

LUN management

In a SAN environment, an SVM administrator can provision storage by creating LUNs, igroups, and mapping the LUNs to the igroups. After creating LUNs, SVM administrator can manage their availability, mapping, and accessibility.

Note: SVMs with Infinite Volume cannot have LUNs. Therefore, you cannot perform LUN related tasks on an SVM with Infinite Volume.

Depending on your capabilities, you can perform the following tasks to manage LUNs:

- Creating, modifying, renaming, or deleting LUNs
- Modifying LUN size
- Managing igroups and port sets
- Mapping LUNs to the initiators
- Unmapping LUNs
- Viewing list of LUNs

Related information

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

Commands for managing LUNs

SVM administrators can identify the list of commands for managing LUNs on your SVM by navigating to the respective command directory.

To identify the list of commands available to manage storage, you must navigate to the lun directory.

Example

The following example shows how to identify the lun commands:

```
vs1.example.com::lun> ?
create          Create a new LUN
delete          Delete the LUN
igroup>         Manage initiator groups
map             Map LUN to all the initiators in the group
mapped>         The mapped directory
maxsize         Display the maximum possible size of a LUN on a
                given volume or qtree.
modify          Modify a LUN
move            Move (rename) a LUN
portset>        Manage portsets
```

resize	Changes the size of the LUN to the input value size.
show	Display a list of LUNs
unmap	Remove a previously configured mapping

Backup management

As an SVM administrator, you can back up SVM's data volumes by using Snapshot copy and NDMP technology. You can also set up SnapMirror relationship between volumes of the peered SVMs to protect data volumes of an SVM.

Starting with clustered Data ONTAP 8.2, you can perform tape backup and restore operations for your SVM data by using NDMP and set up SnapMirror relationships between volumes of the peered SVMs. You can create and manage data protection (DP), SnapVault (XDP), and transition (TDP) relationships. You cannot create or manage load-sharing relationship (LS) SnapMirror relationships.

Note: Infinite Volumes do not support NDMP, SnapVault relationships (XDP), transition relationships (TDP), and load-sharing relationships (LS).

Related information

[Clustered Data ONTAP 8.3 Data Protection Guide](#)

[Clustered Data ONTAP 8.3 Data Protection Tape Backup and Recovery Guide](#)

Snapshot copy management

Storage Virtual Machines (SVMs) use Snapshot copy technology to back up the data volumes. The Snapshot copies of the volumes reside within the SVM. As an SVM administrator, you can manage the Snapshot copies and restore files from the Snapshot copies if data is corrupted.

Depending on your capabilities, you can perform the following tasks to manage Snapshot copies of FlexVol volumes of an SVM:

- Creating, modifying, renaming, or deleting Snapshot copies
- Managing Snapshot policies
- Computing reclaimable space for Snapshot copies
- Viewing the list of Snapshot copies
- Restoring files from Snapshot copies

Depending on your capabilities, you can perform the following tasks to manage Snapshot copies of Infinite Volumes of an SVM:

- Creating or deleting Snapshot copies

- Managing Snapshot policies
- Viewing the list of Snapshot copies
- Restoring Snapshot copies

For more information about managing Snapshot copies, see the *Clustered Data ONTAP Data Protection Guide*.

Related information

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

SnapMirror management

As an SVM administrator, you can create and manage SnapMirror relationships with types data protection (DP), SnapVault (XDP), and transition (TDP) between volumes of the peered SVMs to replicate data of the primary SVM. You cannot create or manage load-sharing relationship (LS) SnapMirror relationships.

Depending on your capabilities, you can perform the following tasks to manage SnapMirror relationships of an SVM:

- Creating, modifying, or deleting SnapMirror relationships
- Initializing baseline transfer
- Displaying a list of destinations and SnapMirror relationships
- Managing SnapMirror policies
- Aborting, resuming, and disabling transfer of data
- Starting an incremental transfer of data
- Breaking the SnapMirror relationship to make the destination writable

For more information about SnapMirror operations, see the *Clustered Data ONTAP Data Protection Guide*.

Related information

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

NDMP management

As an SVM administrator, you can perform NDMP operations such as creating and managing NDMP sessions to back up SVM with FlexVol volume's data and restore the data whenever needed. SVMs with Infinite Volume do not support NDMP.

Depending on your capabilities, you can perform the following tasks to manage NDMP sessions of an SVM:

- Enabling and disabling NDMP service
- Terminating the NDMP sessions
- Modifying NDMP properties
- Displaying list of NDMP sessions, properties, and NDMP version

For more information about the NDMP operations, see the *Clustered Data ONTAP Data Protection Tape Backup and Recovery Guide*.

Related information

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

Commands for managing backup

SVM administrators can identify the list of commands for managing backups on an SVM by navigating to the respective command directory.

To identify the list of commands available for:

- Managing Snapshot copies, you must navigate to the snapshot directory under volume directory.
- Managing SnapMirror relationships, you must navigate to the SnapMirror directory.
- Managing NDMP, you must navigate to the ndmp directory under vservers services directory.

Example

The following example shows how to identify the backup commands:

```
vs1.example.com::volume snapshot> ?
autodelete>          Manage snapshot autodelete settings
create               Create a snapshot
delete              Delete a snapshot
modify              Modify snapshot attributes
partial-restore-file Restore part of a file from a snapshot
policy>            Manage snapshot policies
rename              Rename a snapshot
restore-file        Restore a file from a snapshot
show                Display a list of snapshots
```

```

vs1.example.com::snapmirror> ?
  abort                Abort an active transfer
  break                Make SnapMirror destination writable
  create                Create a new SnapMirror relationship
  ...
  ...
  update                Start an incremental transfer

vs1.example.com::vserver services ndmp> ?
  generate-password    Display NDMP password for a user
  kill                 Kill the specified NDMP session
  ...
  ...
  version              Display default NDMP version

```

Policy management

As an SVM administrator, you can create and manage a collection of rules called policies to manage the data access from an SVM. Depending on the capabilities assigned to you, you can create policies such as SnapMirror policy and Snapshot policy.

You can manage the following policies of SVMs:

- Export policies
- File policies
- Quota policies
- SnapMirror policies
- Snapshot copy policies
- Data policies

Each SVM with Infinite Volume has one data policy. When an Infinite Volume contains two or more storage classes, you can use a data policy and its rules to automatically filter incoming data into different storage classes.

Depending on your capabilities, you can perform the following tasks to manage policies of an SVM:

- Creating, renaming, copying, displaying, or deleting export policies
Clustered Data ONTAP 8.3 File Access Management Guide for NFS
- Creating, modifying, displaying, or deleting file policies
Clustered Data ONTAP 8.3 File Access Management Guide for CIFS

Note: SVMs with Infinite Volume do not support file policies.

- Creating, renaming, copying, displaying, or deleting quota policies
Clustered Data ONTAP 8.3 Logical Storage Management Guide

Note: SVMs with Infinite Volume do not support quota policies.

- Creating, renaming, copying, displaying, or deleting SnapMirror policies and rules
Clustered Data ONTAP 8.3 Data Protection Guide
- Creating, renaming, copying, displaying, or deleting Snapshot copy policies and schedules
Clustered Data ONTAP 8.3 Data Protection Guide
- Exporting, importing, and validating data policies in JSON format for SVMs with Infinite Volume.
Clustered Data ONTAP 8.3 Infinite Volumes Management Guide

Commands for managing policies

SVM administrators can identify the list of commands for managing policies on an SVM by navigating to the respective command directory.

To identify the list of commands available to manage policies, you must navigate to the parent directory of the type of policy. For example, if you want to know about SnapMirror policy, you must navigate to the snapmirror policy directory.

Example

The following example shows how to identify the SnapMirror policy commands:

```
vs1.example.com::snapmirror policy> ?
add-rule          Add a new rule to SnapMirror policy
create           Create a new SnapMirror policy
delete           Delete a SnapMirror policy
modify           Modify a SnapMirror policy
modify-rule      Modify an existing rule in SnapMirror
                 policy
remove-rule      Remove a rule from SnapMirror policy
show            Show SnapMirror policies
```

Glossary

administrator

The account that has the required permission to administer a Data ONTAP system.

aggregate

A manageable unit of RAID-protected storage, consisting of one or two plexes, that can contain one traditional volume or multiple FlexVol volumes.

Common Internet File System (CIFS)

Microsoft's file-sharing networking protocol that evolved from SMB.

CIFS share

- In Data ONTAP, a directory or directory structure that has been made available to network users and can be mapped to a drive letter on a CIFS client. Also known simply as a *share*.
- In OnCommand Insight (formerly SANscreen suite), a service exposed from a NAS device to provide file-based storage through the CIFS protocol. CIFS is mostly used for Microsoft Windows clients, but many other operating systems can access CIFS shares as well.

client

A workstation or PC in a client-server architecture; that is, a computer system or process that requests services from and accepts the responses of another computer system or process.

credential

The configuration of a user account name and password that provide administrative privileges on the storage system.

data SVM

Formerly known as data Vserver. In clustered Data ONTAP, a Storage Virtual Machine (SVM) that facilitates data access from the cluster; the hardware and storage resources of the cluster are dynamically shared by data SVMs within a cluster.

domain name server (DNS)

In OnCommand Insight (formerly SANscreen suite), a resource that resolves domain names to their equivalent IP addresses so that IP traffic can be transported to the correct destination. Each domain name is associated with, at a minimum, a primary and a secondary DNS.

FC (Fibre Channel Protocol)

An interface protocol for SCSI transport when mapping block-oriented storage data over Fibre Channel networks.

FlexVol volume

In clustered Data ONTAP, a logical entity contained in a Storage Virtual Machine (SVM, formerly known as Vserver)—referred to as SVM with FlexVol volumes. FlexVol volumes typically hold user data, although they also serve as node or SVM root volumes and metadata containers. A FlexVol volume obtains its storage from a single aggregate.

igroup

initiator group. A collection of unique identifiers, either FC WWPNs (World Wide Port Names) in a SCSI network or iSCSI node names of initiators (hosts) in an IP network, that are given access to LUNs when they are mapped to those LUNs.

initiator

The system component that originates an I/O command over an I/O bus or network. The target is the component that receives this command.

Infinite Volume

In clustered Data ONTAP, a logical entity contained in a Storage Virtual Machine (SVM, formerly known as Vserver)—referred to as SVM with Infinite Volume—that holds user data. An Infinite Volume obtains its storage from multiple aggregates.

iSCSI

Internet Small Computer Systems Interface (iSCSI) protocol. A licensed service on the storage system that enables you to export LUNs to hosts using the SCSI protocol over TCP/IP.

LIF

logical interface. Formerly known as *VIF* (virtual interface) in Data ONTAP GX. A logical network interface, representing a network access point to a node. LIFs currently correspond to IP addresses, but could be implemented by any interconnect. A LIF is generally bound to a physical network port; that is, an Ethernet port. LIFs can fail over to other physical ports (potentially on other nodes) based on policies interpreted by the LIF manager.

Lightweight Directory Access Protocol (LDAP)

A client-server protocol for accessing a directory service.

LUN (Logical Unit Number)

The identifier of an FC or iSCSI logical unit. A logical unit typically corresponds to a storage volume and is represented within a computer operating system as a device.

move (v)

To physically move data and any needed associated configuration of an object from one aggregate to another within a cluster, including within a single node.

namespace

In network-attached storage (NAS) cluster environments, an abstraction layer for data location that provides a single access point for all data in the system. It enables users to access data without specifying the physical location of the data, and enables

administrators to manage distributed data storage as a single file system. Sometimes referred to as *global namespace*.

Network File System (NFS) export

A service exposed from a NAS device to provide file-based storage through the NFS protocol. NFS is mostly used for UNIX-like operating systems, but other operating systems can access NFS exports as well.

policies

The collection of management options, controls, and specifications for directing the automated management of data.

qtree

A special subdirectory of the root of a volume that acts as a virtual subvolume with special attributes.

SAN host

Any storage area network (SAN) device, such as a UNIX or Windows system, that sends requests to other SAN devices in a SAN to perform tasks. To be monitored through Operations Manager console on the OnCommand Unified Manager server, a SAN host must be running the NetApp Host Agent software.

Snapshot copy

An online, read-only copy of an entire file system that protects against accidental deletions or modifications of files without duplicating file contents. Snapshot copies enable users to restore files and to back up the storage system to tape while the storage system is in use.

Storage Virtual Machine (SVM)

(Known as *Vserver* prior to clustered Data ONTAP 8.2.1. The term “Vserver” is still used in CLI displays and `vserver` command syntax.) A virtual machine that provides network access through unique network addresses, that might serve data out of a distinct namespace, and that is separately administrable from the rest of the cluster. There are three types of SVMs—*admin*, *node*, and *data*—but unless there is a specific need to identify the type of SVM, “SVM” usually refers to the data SVM.

throughput

The rate at which data is transferred to or from the storage device, measured in megabytes per second (MBps).

Vserver

(Known as “Storage Virtual Machine (SVM)” in clustered Data ONTAP 8.2.1 and later.) A virtual machine that provides network access through unique network addresses, that might serve data out of a distinct namespace, and that is separately administrable from the rest of the cluster. There are three types of Vservers—*admin*, *node*, and *cluster* (“cluster Vserver” is called “data Vserver” in Data ONTAP 8.2)—but unless there is a specific need to identify the type of Vserver, “Vserver” usually refers to the cluster/data Vserver.

volume

- For Data ONTAP, a logical entity that holds user data that is accessible through one or more of the supported access protocols, including Network File System (NFS), Common Internet File System (CIFS), Fibre Channel (FC), and Internet SCSI (iSCSI). Data ONTAP treats an IBM volume as a disk.
- For IBM, the area on the storage array that is available for a Data ONTAP system or non Data ONTAP host to read data from or write data to. The documentation uses the term *array LUN* to describe this area.

Copyright information

Copyright © 1994–2014 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the NetApp logo, Go Further, Faster, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

- ## A
- accessing
 - Data ONTAP man pages [16](#)
 - admin
 - use of administrative privilege levels [12](#)
 - administrative privileges
 - use of levels [12](#)
 - administrators
 - differences between cluster and SVM [8](#)
 - advanced
 - use of administrative privilege levels [12](#)
 - algorithms
 - key exchange and data encryption, introduction to managing SSH security configuration [22](#)
 - authentication
 - managing digital certificates for server or client [25](#)
- ## B
- benefits
 - of using SVMs [7](#)
- ## C
- CA certificates
 - for server, installing to authenticate SVM as SSL server [33](#)
 - certificates
 - client, installing to authenticate the SVM as an SSL client [34](#)
 - commands for managing digital [37](#)
 - digital, managing for server or client authentication [25](#)
 - replacing an expired digital [36](#)
 - server CA, installing to authenticate SVM as SSL server [33](#)
 - server, installing to authenticate the SVM as an SSL server [26](#)
 - ciphers
 - key exchange and data encryption, introduction to managing SSH security configurations [22](#)
 - CLI
 - methods of navigating command directories [9](#)
 - overview of using Data ONTAP [9](#)
 - rules for specifying values [10](#)
 - setting display preferences in [13](#)
 - setting privilege levels [13](#)
 - CLI commands
 - keyboard shortcuts for editing [11](#)
 - client CA certificates
 - installing to authenticate an SSL client of the SVM [29](#)
 - client certificates
 - installing to authenticate the SVM as an SSL client [34](#)
 - clients
 - managing digital certificates for authentication of servers or [25](#)
 - clusters
 - administrators, definition [8](#)
 - command directories
 - methods of navigating CLI [9](#)
 - command-line interface
 - See* CLI
 - commands
 - CLI, keyboard shortcuts for editing [11](#)
 - for managing dashboards [53](#)
 - for managing digital certificates [37](#)
 - for managing job schedules [44](#)
 - for managing jobs [43](#)
 - for managing public keys [24](#)
 - for managing SSH security configurations [23](#)
 - for managing SSL [39](#)
 - methods of customizing show output by using fields [16](#)
 - methods of viewing history and reissuing [10](#)
 - rules for specifying values in the CLI [10](#)
 - statistics [49](#)
 - using to monitor performance [49](#)
 - comments
 - how to send feedback about documentation [72](#)
 - configuring
 - protocols
 - CIFS [54](#)
 - FC [54](#)
 - iSCSI [54](#)
 - NFS [54](#)
 - counters
 - what they are [45](#)

D

- dashboards
 - commands for managing [53](#)
- data
 - commands for viewing [49](#)
- data encryption algorithms
 - introduction to managing SSH security configurations [22](#)
- Data ONTAP
 - accessing man pages [16](#)
 - introduction to management interface [9](#)
 - overview of using the CLI [9](#)
- diagnostic
 - use of administrative privilege levels [12](#)
- digital certificates
 - commands for managing [37](#)
 - installing a client CA or root CA certificate to authenticate an SSL client of the SVM [29](#)
 - installing a client certificate to authenticate the SVM as an SSL client [34](#)
 - installing a server certificate to authenticate the SVM as an SSL server [26](#)
 - installing server CA certificate to authenticate SVM as SSL server [33](#)
 - managing for server or client authentication [25](#)
 - managing SSL, introduction [38](#)
 - replacing an expired [36](#)
- directories
 - methods of navigating CLI command [9](#)
- display preferences
 - setting in CLI [13](#)
- displaying
 - SVM details [50](#)
- documentation
 - how to send feedback about [72](#)

E

- encryption algorithms
 - data, introduction to managing SSH security configurations [22](#)
- extended queries
 - methods of using [15](#)

F

- feedback
 - how to send comments about documentation [72](#)
- fields

methods of customizing show command output by using [16](#)

- FlexVol volumes
 - with SVMs, explained [6](#)

H

- health monitoring
 - commands for managing dashboards [53](#)
- history of commands
 - methods of viewing [10](#)
- HTTPS
 - managing SSL, introduction [38](#)

I

- Infinite Volumes
 - with SVMs, explained [6](#)
- information
 - how to send feedback about improving documentation [72](#)
- instances
 - what they are [45](#)
- interfaces
 - introduction to Data ONTAP management [9](#)
 - overview of using Data ONTAP command line [9](#)

J

- job schedules
 - commands for managing [44](#)
- jobs
 - commands for managing [43](#)
 - managing schedules for [43](#)
 - viewing information about [43](#)

K

- key exchange algorithms
 - introduction to managing SSH security configurations [22](#)
- keyboard shortcuts
 - for editing CLI commands [11](#)
- keys
 - ways to manage public [24](#)

L

- levels

- use of administrative privilege [12](#)
- LUNs
 - managing [60](#)

M

- man pages
 - accessing Data ONTAP [16](#)
- management interfaces
 - introduction to Data ONTAP [9](#)
- managing
 - qtree [58](#)
 - quotas [58](#)
 - volumes [58](#)
- monitoring
 - commands for managing dashboards [53](#)
- mutual authentication
 - managing digital certificates for server or client authentication [25](#)

O

- objects
 - what they are [45](#)
- operators
 - methods of using query [14](#)
- output
 - methods of customizing show command by using fields [16](#)

P

- parameters
 - rules for specifying values in the CLI [10](#)
- performance
 - data
 - decisions before you view [46](#)
 - viewing continuously [48](#)
 - viewing for a time period [47](#)
 - what objects, instances, and counters are [45](#)
 - monitoring [45](#)
 - monitoring using the statistics commands [49](#)
- preferences
 - setting display in CLI [13](#)
- privilege levels
 - setting in CLI [13](#)
 - use of administrative [12](#)
- prompts
 - overview of Data ONTAP command [9](#)
- public key certificates

- managing digital certificates for server or client authentication [25](#)
- public keys
 - commands for managing [24](#)
 - ways to manage [24](#)

Q

- queries
 - methods of using extended [15](#)
- query operators
 - methods of using [14](#)

R

- reissuing commands
 - methods of [10](#)
- root CA certificates
 - installing to authenticate an SSL client of the SVM [29](#)
- rows command
 - setting display preferences in the CLI [13](#)

S

- schedules
 - commands for managing job [44](#)
 - managing jobs and [43](#)
- Secure Sockets Layer
 - See* SSL
- security certificate commands
 - for managing digital certificates [37](#)
- security configurations
 - commands for managing SSH [23](#)
 - introduction to managing SSH [22](#)
- server CA certificates
 - installing to authenticate SVM as SSL server [33](#)
- server certificates
 - installing to authenticate the SVM as an SSL server [26](#)
- servers
 - managing digital certificates for authentication of clients or [25](#)
- set command
 - setting display preferences in the CLI [13](#)
- shortcuts
 - keyboard, for editing CLI commands [11](#)
- show command output
 - methods of customizing by using fields [16](#)
- SSH

- commands for managing security configurations [23](#)
- security configurations, introduction to managing [22](#)

SSL

- commands for managing [39](#)
- managing digital certificates for server or client authentication [25](#)
- managing, introduction [38](#)

SSL clients

- installing a client certificate to authenticate the SVM as [34](#)
- of the SVM, installing a client CA or root CA certificate to authenticate [29](#)

SSL servers

- installing a server certificate to authenticate the SVM as [26](#)
- installing server CA certificate to authenticate SVM as [33](#)

statistics

- See* performance

suggestions

- how to send feedback about documentation [72](#)

SVM backup

- commands for managing [63](#)

SVM policies

- commands for managing [65](#)

SVMs

- access methods [18](#)
- administration capabilities [40](#)
- administrators, definition [8](#)
- benefits of using [7](#)
- changing password [21](#)
- data security [56](#)
- displaying APIs [42](#)
- displaying information about [50](#)
- displaying peer relationship [51](#)
- DNS
 - configuration [57](#)
- file security commands [56](#)
- identifying the commands [41](#)
- LDAP
 - configuration [57](#)
- local UNIX groups
 - configuration [57](#)
- logging in [19](#)
- LUN commands [60](#)
- managing authentication [21](#)
- managing backups [61](#)

- managing NDMP [63](#)
- managing policies [64](#)
- managing SnapMirror [62](#)
- managing Snapshot copies [61](#)
- monitoring health [53](#)
- NIS

 - configuration [57](#)

- overview of administration [6](#)
- performance [45](#)
- services commands [58](#)
- services configuration [57](#)
- storage commands [59](#)
- user accounts, access methods [18](#)
- user accounts, authentication methods [18](#)
- viewing network configuration [52](#)
- viewing protocol commands [55](#)
- with FlexVol volumes, explained [6](#)
- with Infinite Volume, explained [6](#)

SVMs with FlexVol volumes

- explained [6](#)

SVMs with Infinite Volume

- explained [6](#)

T

two-way authentication

- managing digital certificates for server or client authentication [25](#)

V

values

- rules for specifying in CLI [10](#)

viewing

- SVM details [50](#)

volumes

- SVMs with FlexVol, explained [6](#)
- SVMs with Infinite, explained [6](#)

Vservers

- See* SVMs

W

web access

- managing SSL, introduction [38](#)