



Configure

ONTAP 9

NetApp
April 29, 2024

Table of Contents

- Configure 1
 - About the S3 configuration process 1
 - Configure S3 access to an SVM 5
 - Add storage capacity to an S3-enabled SVM 19
 - Create or modify access policy statements 34
 - Enable client access to S3 object storage 44
 - Storage service definitions 47


Configure

About the S3 configuration process

S3 configuration workflow

Configuring S3 involves assessing physical storage and networking requirements, and then choosing a workflow that is specific to your goal—configuring S3 access to a new or existing SVM, or adding a bucket and users to an existing SVM that is already fully configured for S3 access.

When you configure S3 access to a new storage VM using System Manager, you are prompted to enter certificate and networking information, and the storage VM and S3 object storage server are created in a single operation.



Assess physical storage requirements

Before provisioning S3 storage for clients, you must ensure that there is sufficient space in existing aggregates for the new object store. If there is not, you can add disks to existing aggregates or create new aggregates of the desired type and location.

About this task

When you create an S3 bucket in an S3-enabled SVM, a FlexGroup volume is automatically created to support the bucket. You can let ONTAP select the underlying aggregates and FlexGroup components automatically (the default) or you can select the underlying aggregates and FlexGroup components yourself.

If you decide to specify the aggregates and FlexGroup components — for example, if you have specific

performance requirements for the underlying disks — you should make sure that your aggregate configuration conforms to best practice guidelines for provisioning a FlexGroup volume. Learn more:

- [FlexGroup volumes management](#)
- [NetApp Technical Report 4571-a: NetApp ONTAP FlexGroup Volume Top Best Practices](#)

If you are serving buckets from Cloud Volumes ONTAP, it is strongly recommended that you manually select the underlying aggregates to ensure that they are using one node only. Using aggregates from both nodes can impact performance, because the nodes will be in geographically separated availability zones and hence susceptible to latency issues. Learn about [creating buckets for Cloud Volumes ONTAP](#).

You can use the ONTAP S3 server to create a local FabricPool capacity tier; that is, in the same cluster as the performance tier. This can be useful, for example, if you have SSD disks attached to one HA pair and you want to tier *cold* data to HDD disks in another HA pair. In this use case, the S3 server and the bucket containing the local capacity tier should therefore be in a different HA pair than the performance tier. Local tiering is not supported on one-node and two-node clusters.

Steps

1. Display available space in existing aggregates:

```
storage aggregate show
```

If there is an aggregate with sufficient space or requisite node location, record its name for your S3 configuration.

```
cluster-1::> storage aggregate show
Aggregate      Size Available Used% State  #Vols  Nodes  RAID Status
-----
aggr_0         239.0GB   11.13GB   95% online    1 node1  raid_dp, normal
aggr_1         239.0GB   11.13GB   95% online    1 node1  raid_dp, normal
aggr_2         239.0GB   11.13GB   95% online    1 node2  raid_dp, normal
aggr_3         239.0GB   11.13GB   95% online    1 node2  raid_dp, normal
aggr_4         239.0GB   238.9GB   95% online    5 node3  raid_dp, normal
aggr_5         239.0GB   239.0GB   95% online    4 node4  raid_dp, normal
6 entries were displayed.
```

2. If there are no aggregates with sufficient space or requisite node location, add disks to an existing aggregate by using the `storage aggregate add-disks` command, or create a new aggregate by using the `storage aggregate create` command.

Assess networking requirements

Before providing S3 storage to clients, you must verify that networking is correctly

configured to meet the S3 provisioning requirements.

Before you begin

The following cluster networking objects must be configured:

- Physical and logical ports
- Broadcast domains
- Subnets (if required)
- IPspaces (as required, in addition to the default IPspace)
- Failover groups (as required, in addition to the default failover group for each broadcast domain)
- External firewalls

About this task

For remote FabricPool capacity (cloud) tiers and remote S3 clients, you must use a data SVM and configure data LIFs. For FabricPool cloud tiers, you must also configure intercluster LIFs; cluster peering is not required.

For local FabricPool capacity tiers, you must use the system SVM (called “Cluster”), but you have two options for LIF configuration:

- You can use the cluster LIFs.

In this option, no further LIF configuration is required, but there will be an increase in traffic on the cluster LIFs. Also, the local tier will not be accessible to other clusters.

- You can use data and intercluster LIFs.

This option requires additional configuration, including enabling the LIFs for the S3 protocol, but the local tier will also be accessible as a remote FabricPool cloud tier to other clusters.

Steps

1. Display the available physical and virtual ports:

```
network port show
```

- When possible, you should use the port with the highest speed for the data network.
- All components in the data network must have the same MTU setting for best performance.

2. If you are planning to use a subnet name to allocate the IP address and network mask value for a LIF, verify that the subnet exists and has sufficient addresses available:

```
network subnet show
```

Subnets contain a pool of IP addresses that belong to the same layer 3 subnet. Subnets are created by using the `network subnet create` command.

3. Display available IPspaces:

```
network ipspace show
```

You can use the default IPspace or a custom IPspace.

4. If you want to use IPv6 addresses, verify that IPv6 is enabled on the cluster:

```
network options ipv6 show
```

If required, you can enable IPv6 by using the `network options ipv6 modify` command.

Decide where to provision new S3 storage capacity

Before you create a new S3 bucket, you must decide whether to place it in a new or existing SVM. This decision determines your workflow.

Choices

- If you want to provision a bucket in a new SVM or an SVM that is not enabled for S3, complete the steps in the following topics.

[Create an SVM for S3](#)

[Create a bucket for S3](#)

Although S3 can coexist in an SVM with NFS and SMB, you might choose to create a new SVM if one of the following is true:

- You are enabling S3 on a cluster for the first time.
- You have existing SVMs in a cluster in which you do not want to enable S3 support.
- You have one or more S3-enabled-SVMs in a cluster, and you want another S3 server with different performance characteristics. After enabling S3 on the SVM, proceed to provision a bucket.
- If you want to provision the initial bucket or an additional bucket on an existing S3-enabled SVM, complete the steps in the following topic.

[Create a bucket for S3](#)

Configure S3 access to an SVM

Create an SVM for S3

Although S3 can coexist with other protocols in an SVM, you might want to create a new SVM to isolate the namespace and workload.

About this task

If you are only providing S3 object storage from an SVM, the S3 server does not require any DNS configuration. However, you might want to configure DNS on the SVM if other protocols are used.

When you configure S3 access to a new storage VM using System Manager, you are prompted to enter certificate and networking information, and the storage VM and S3 object storage server are created in a single operation.

Example 1. Steps

System Manager

You should be prepared to enter the S3 server name as a Fully Qualified Domain Name (FQDN), which clients will use for S3 access. The S3 server FQDN must not begin with a bucket name.


You should be prepared to enter IP addresses for interface role Data.

If you are using an external-CA signed certificate, you will be prompted to enter it during this procedure; you also have the option to use a system-generated certificate.

1. Enable S3 on a storage VM.

- a. Add a new storage VM: Click **Storage > Storage VMs**, then click **Add**.

If this is a new system with no existing storage VMs: Click **Dashboard > Configure Protocols**.

If you are adding an S3 server to an existing storage VM: Click **Storage > Storage VMs**, select a storage VM, click **Settings**, and then click  under **S3**.

- b. Click **Enable S3**, then enter the S3 Server Name.

- c. Select the certificate type.

Whether you select system-generated certificate or one of your own, it will be required for client access.

- d. Enter the network interfaces.

2. If you selected the system-generated certificate, you see the certificate information when the new storage VM creation is confirmed. Click **Download** and save it for client access.

- The secret key will not be displayed again.
- If you need the certificate information again: Click **Storage > Storage VMs**, select the storage VM, and click **Settings**.

CLI

1. Verify that S3 is licensed on your cluster:

```
system license show -package s3
```

If it is not, contact your sales representative.

2. Create an SVM:

```
vserver create -vserver <svm_name> -subtype default -rootvolume  
<root_volume_name> -aggregate <aggregate_name> -rootvolume-security  
-style unix -language C.UTF-8 -data-services <data-s3-server>  
-ipSPACE <ipSPACE_name>
```

- Use the UNIX setting for the `-rootvolume-security-style` option.

- Use the default C.UTF-8 -language option.
- The ipspace setting is optional.

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

```
vserver show -vserver <svm_name>
```

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

Examples

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

```
cluster-1::> vserver create -vserver svm1.example.com -rootvolume  
root_svm1 -aggregate aggr1 -rootvolume-security-style unix -language  
C.UTF-8 -data-services _data-s3-server_ -ipspace ipspaceA
```

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

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

```

cluster-1::> vserver show -vserver svm1.example.com
                                Vserver: svm1.example.com
                                Vserver Type: data
                                Vserver Subtype: default
                                Vserver UUID: b8375669-19b0-11e5-b9d1-
00a0983d9736

                                Root Volume: root_svm1
                                Aggregate: aggr1
                                NIS Domain: -
                                Root Volume Security Style: unix
                                LDAP Client: -
                                Default Volume Language Code: C.UTF-8
                                Snapshot Policy: default
                                Comment:
                                Quota Policy: default
                                List of Aggregates Assigned: -
                                Limit on Maximum Number of Volumes allowed: unlimited
                                Vserver Admin State: running
                                Vserver Operational State: running
                                Vserver Operational State Stopped Reason: -
                                Allowed Protocols: nfs, cifs
                                Disallowed Protocols: -
                                QoS Policy Group: -
                                Config Lock: false
                                IPspace Name: ipspaceA

```

Create and install a CA certificate on the SVM

A Certificate Authority (CA) certificate is required to enable HTTPS traffic from S3 clients to the S3-enabled SVM.

About this task

Although it is possible to configure an S3 server to use HTTP only, and although it is possible to configure clients without a CA certificate requirement, it is a best practice to secure HTTPS traffic to ONTAP S3 servers with a CA certificate.

A CA certificate is not necessary for a local tiering use case, where IP traffic is going over cluster LIFs only.

The instructions in this procedure will create and install an ONTAP self-signed certificate. CA certificates from third-party vendors are also supported; see the administrator authentication documentation for more information.

Administrator authentication and RBAC

See the `security certificate` man pages for additional configuration options.

Steps

1. Create a self-signed digital certificate:

```
security certificate create -vserver svm_name -type root-ca -common-name  
ca_cert_name
```

The `-type root-ca` option creates and installs a self-signed digital certificate to sign other certificates by acting as a certificate authority (CA).

The `-common-name` option creates the SVM's Certificate Authority (CA) name and will be used when generating the certificate's complete name.

The default certificate size is 2048 bits.

Example

```
cluster-1::> security certificate create -vserver svm1.example.com -type  
root-ca -common-name svm1_ca
```

The certificate's generated name for reference:
svm1_ca_159D1587CE21E9D4_svm1_ca

When the certificate's generated name is displayed; be sure to save it for later steps in this procedure.

2. Generate a certificate signing request:

```
security certificate generate-csr -common-name s3_server_name  
[additional_options]
```

The `-common-name` parameter for the signing request must be the S3 server name (FQDN).

You can provide the location and other detailed information about the SVM if desired.

You are prompted to keep a copy of your certificate request and private key for future reference.

3. Sign the CSR using SVM_CA to generate S3 Server's certificate:

```
security certificate sign -vserver svm_name -ca ca_cert_name -ca-serial  
ca_cert_serial_number [additional_options]
```

Enter the command options that you used in previous steps:

- `-ca` — the common name of the CA that you entered in Step 1.
- `-ca-serial` — the CA serial number from Step 1. For example, if the CA certificate name is svm1_ca_159D1587CE21E9D4_svm1_ca, the serial number is 159D1587CE21E9D4.

By default, the signed certificate will expire in 365 days. You can select another value, and specify other signing details.

When prompted, copy and enter the certificate request string you saved in Step 2.

A signed certificate is displayed; save it for later use.

4. Install the signed certificate on the S3-enabled SVM:

```
security certificate install -type server -vserver svm_name
```

When prompted, enter the certificate and private key.

You have the option to enter intermediate certificates if a certificate chain is desired.

When the private key and the CA-signed digital certificate are displayed; save them for future reference.

5. Get the public key certificate:

```
security certificate show -vserver svm_name -common-name ca_cert_name -type  
root-ca -instance
```

Save the public key certificate for later client-side configuration.

Example

```
cluster-1::> security certificate show -vserver svm1.example.com -common  
-name svm1_ca -type root-ca -instance  
  
Name of Vserver: svm1.example.com  
FQDN or Custom Common Name: svm1_ca  
Serial Number of Certificate: 159D1587CE21E9D4  
Certificate Authority: svm1_ca  
Type of Certificate: root-ca  
(DEPRECATED)-Certificate Subtype: -  
Unique Certificate Name: svm1_ca_159D1587CE21E9D4_svm1_ca  
Size of Requested Certificate in Bits: 2048  
Certificate Start Date: Thu May 09 10:58:39 2020  
Certificate Expiration Date: Fri May 08 10:58:39 2021  
Public Key Certificate: -----BEGIN CERTIFICATE-----  
MIIDZ ...==  
-----END CERTIFICATE-----  
  
Country Name: US  
State or Province Name:  
Locality Name:  
Organization Name:  
Organization Unit:  
Contact Administrator's Email Address:  
Protocol: SSL  
Hashing Function: SHA256  
Self-Signed Certificate: true  
Is System Internal Certificate: false
```

Create an S3 service data policy

You can create service policies for S3 data and management services. An S3 service data policy is required to enable S3 data traffic on LIFs.

About this task

An S3 service data policy is required if you are using data LIFs and intercluster LIFs. It is not required if you are using cluster LIFs for the local tiering use case.

When a service policy is specified for a LIF, the policy is used to construct a default role, failover policy, and data protocol list for the LIF.

Although multiple protocols can be configured for SVMs and LIFs, it is a best practice for S3 to be the only protocol when serving object data.

Steps

1. Change the privilege setting to advanced:

```
set -privilege advanced
```

2. Create a service data policy:

```
network interface service-policy create -vserver svm_name -policy policy_name  
-services data-core,data-s3-server
```

The `data-core` and `data-s3-server` services are the only ones required to enable ONTAP S3, although other services can be included as needed.

Create data LIFs

If you created a new SVM, the dedicated LIFs you create for S3 access should be data LIFs.

Before you begin

- The underlying physical or logical network port must have been configured to the administrative `up` status.
- If you are planning to use a subnet name to allocate the IP address and network mask value for a LIF, the subnet must already exist.

Subnets contain a pool of IP addresses that belong to the same layer 3 subnet. They are created using the `network subnet create` command.

- The LIF service policy must already exist.

About this task

- You can create both IPv4 and IPv6 LIFs on the same network port.
- If you have a large number of LIFs in your cluster, you can verify the LIF capacity supported on the cluster by using the `network interface capacity show` command and the LIF capacity supported on each node by using the `network interface capacity details show` command (at the advanced privilege level).
- If you are enabling remote FabricPool capacity (cloud) tiering, you must also configure intercluster LIFs.

Steps

1. Create a LIF:

```
network interface create -vserver svm_name -lif lif_name -service-policy
service_policy_names -home-node node_name -home-port port_name {-address
IP_address -netmask IP_address | -subnet-name subnet_name} -firewall-policy
data -auto-revert {true|false}
```

- `-home-node` is the node to which the LIF returns when the `network interface revert` command is run on the LIF.

You can also specify whether the LIF should automatically revert to the home-node and home-port with the `-auto-revert` option.

- `-home-port` is the physical or logical port to which the LIF returns when the `network interface revert` command is run on the LIF.
- You can specify an IP address with the `-address` and `-netmask` options, or you enable allocation from a subnet with the `-subnet_name` option.
- When using a subnet to supply the IP address and network mask, if the subnet was defined with a gateway, a default route to that gateway is added automatically to the SVM when a LIF is created using that subnet.
- If you assign IP addresses manually (without using a subnet), you might need to configure a default route to a gateway if there are clients or domain controllers on a different IP subnet. The `network route create` man page contains information about creating a static route within an SVM.
- For the `-firewall-policy` option, use the same default data as the LIF role.

You can create and add a custom firewall policy later if desired.



Beginning with ONTAP 9.10.1, firewall policies are deprecated and wholly replaced with LIF service policies. For more information, see [Configure firewall policies for LIFs](#).

- `-auto-revert` allows you to specify whether a data LIF is automatically reverted to its home node under circumstances such as startup, changes to the status of the management database, or when the network connection is made. The default setting is `false`, but you can set it to `true` depending on network management policies in your environment.
- The `-service-policy` option specifies the data and management services policy you created and any other policies you need.

2. If you want to assign an IPv6 address in the `-address` option:

- a. Use the `network ndp prefix show` command to view the list of RA prefixes learned on various interfaces.

The `network ndp prefix show` command is available at the advanced privilege level.

- b. Use the format `prefix:id` to construct the IPv6 address manually.

`prefix` is the prefix learned on various interfaces.

For deriving the `id`, choose a random 64-bit hexadecimal number.

3. Verify that the LIF was created successfully by using the `network interface show` command.
4. Verify that the configured IP address is reachable:

To verify an...	Use...
IPv4 address	<code>network ping</code>
IPv6 address	<code>network ping6</code>

Examples

The following command shows how to create an S3 data LIF that is assigned with the `my-S3-policy` service policy:

```
network interface create -vserver svm1.example.com -lif lif2 -home-node  
node2 -homeport e0d -service-policy my-S3-policy -subnet-name ipspace1
```

The following command shows all the LIFs in cluster-1. Data LIFs `datalif1` and `datalif3` are configured with IPv4 addresses, and `datalif4` is configured with an IPv6 address:

```
cluster-1::> network interface show
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Is Port
Home					
-----	-----	-----	-----	-----	-----
cluster-1					
cluster_mgmt	up/up	192.0.2.3/24	node-1	e1a	
true					
node-1					
clus1	up/up	192.0.2.12/24	node-1	e0a	
true					
clus2	up/up	192.0.2.13/24	node-1	e0b	
true					
mgmt1	up/up	192.0.2.68/24	node-1	e1a	
true					
node-2					
clus1	up/up	192.0.2.14/24	node-2	e0a	
true					
clus2	up/up	192.0.2.15/24	node-2	e0b	
true					
mgmt1	up/up	192.0.2.69/24	node-2	e1a	
true					
vs1.example.com					
datalif1	up/down	192.0.2.145/30	node-1	e1c	
true					
vs3.example.com					
datalif3	up/up	192.0.2.146/30	node-2	e0c	
true					
datalif4	up/up	2001::2/64	node-2	e0c	
true					

5 entries were displayed.

Create intercluster LIFs for remote FabricPool tiering

If you are enabling remote FabricPool capacity (cloud) tiering using ONTAP S3, you must configure intercluster LIFs. You can configure intercluster LIFs on ports shared with the data network. Doing so reduces the number of ports you need for intercluster networking.

Before you begin

- The underlying physical or logical network port must have been configured to the administrative `up` status.
- The LIF service policy must already exist.

About this task

Intercluster LIFs are not required for local Fabric pool tiering or for serving external S3 apps.

Steps

1. List the ports in the cluster:

```
network port show
```

The following example shows the network ports in `cluster01`:

```
cluster01::> network port show
```

(Mbps)						Speed
Node	Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper
-----	-----	-----	-----	-----	-----	
cluster01-01						
	e0a	Cluster	Cluster	up	1500	auto/1000
	e0b	Cluster	Cluster	up	1500	auto/1000
	e0c	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
cluster01-02						
	e0a	Cluster	Cluster	up	1500	auto/1000
	e0b	Cluster	Cluster	up	1500	auto/1000
	e0c	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000

2. Create intercluster LIFs on the system SVM:

```
network interface create -vserver Cluster -lif LIF_name -service-policy  
default-intercluster -home-node node -home-port port -address port_IP -netmask  
netmask
```

The following example creates intercluster LIFs `cluster01_icl01` and `cluster01_icl02`:

```
cluster01::> network interface create -vserver Cluster -lif  
cluster01_icl01 -service-  
policy default-intercluster -home-node cluster01-01 -home-port e0c  
-address 192.168.1.201  
-netmask 255.255.255.0  
  
cluster01::> network interface create -vserver Cluster -lif  
cluster01_icl02 -service-  
policy default-intercluster -home-node cluster01-02 -home-port e0c  
-address 192.168.1.202  
-netmask 255.255.255.0
```

3. Verify that the intercluster LIFs were created:

```
network interface show -service-policy default-intercluster
```

```
cluster01::> network interface show -service-policy default-intercluster
```

	Logical	Status	Network	Current	
Current Is					
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
-----	-----	-----	-----	-----	-----
cluster01	cluster01_icl01	up/up	192.168.1.201/24	cluster01-01	e0c
true	cluster01_icl02	up/up	192.168.1.202/24	cluster01-02	e0c
true					

4. Verify that the intercluster LIFs are redundant:

```
network interface show -service-policy default-intercluster -failover
```

The following example shows that the intercluster LIFs cluster01_icl01 and cluster01_icl02 on the e0c port will fail over to the e0d port.

```
cluster01::> network interface show -service-policy default-intercluster -failover
```

	Logical	Home	Failover	Failover
Vserver	Interface	Node:Port	Policy	Group
-----	-----	-----	-----	-----
cluster01	cluster01_icl01	cluster01-01:e0c	local-only	
192.168.1.201/24				
		Failover Targets: cluster01-01:e0c,		
		cluster01-01:e0d		
	cluster01_icl02	cluster01-02:e0c	local-only	
192.168.1.201/24				
		Failover Targets: cluster01-02:e0c,		
		cluster01-02:e0d		

Create the S3 object store server

The ONTAP object store server manages data as S3 objects, as opposed to file or block storage provided by ONTAP NAS and SAN servers.

Before you begin

You should be prepared to enter the S3 server name as a Fully Qualified Domain Name (FQDN), which clients will use for S3 access. The FQDN must not begin with a bucket name.

You should have a self-signed CA certificate (created in previous steps) or a certificate signed by an external CA vendor. A CA certificate is not necessary for a local tiering use case, where IP traffic is going over cluster LIFs only.

About this task

When an object store server is created, a root user with UID 0 is created. No access key or secret key is generated for this root user. The ONTAP administrator must run the `object-store-server users regenerate-keys` command to set the access key and secret key for this user.



As a NetApp best practice, do not use this root user. Any client application that uses the access key or secret key of the root user has full access to all buckets and objects in the object store.

See the `vserver object-store-server` man pages for additional configuration and display options.


Example 2. Steps

System Manager

Use this procedure if you are adding an S3 server to an existing storage VM. To add an S3 server to a new storage VM, see [Create a storage SVM for S3](#).

You should be prepared to enter IP addresses for interface role Data.

1. Enable S3 on an existing storage VM.

- a. Select the storage VM: click **Storage > Storage VMs**, select a storage VM, click **Settings**, and then click  under **S3**.
- b. Click **Enable S3**, then enter the S3 Server Name.
- c. Select the certificate type.

Whether you select system-generated certificate or one of your own, it will be required for client access.

- d. Enter the network interfaces.

2. If you selected the system-generated certificate, you see the certificate information when the new storage VM creation is confirmed. Click **Download** and save it for client access.

- The secret key will not be displayed again.
- If you need the certificate information again: click **Storage > Storage VMs**, select the storage VM, and click **Settings**.

CLI

1. Create the S3 server:

```
vserver object-store-server create -vserver svm_name -object-store-server  
s3_server_fqdn -certificate-name server_certificate_name -comment text  
[additional_options]
```

You can specify additional options when creating the S3 server or at any time later.

- If you are configuring local tiering, the SVM name can either be a data SVM or system SVM (cluster) name.
- The certificate name should be the name of the server certificate (end user or leaf certificate), and not server CA certificate (intermediate or root CA certificate).
- HTTPS is enabled by default on port 443. You can change the port number with the `-secure -listener-port` option.

When HTTPS is enabled, CA certificates are required for correct integration with SSL/TLS.

- HTTP is disabled by default. When enabled, the server listens on port 80. You can enable it with the `-is-http-enabled` option, or change the port number with the `-listener-port` option.

When HTTP is enabled, the request and responses are sent over the network in clear text.

2. Verify that S3 is configured:

```
vserver object-store-server show
```

Example

This command verifies the configuration values of all object storage servers:

```
cluster1::> vservers object-store-server show

Vserver: vs1

Object Store Server Name: s3.example.com
Administrative State: up
Listener Port For HTTP: 80
Secure Listener Port For HTTPS: 443
HTTP Enabled: false
HTTPS Enabled: true
Certificate for HTTPS Connections: svm1_ca
Comment: Server comment
```

Add storage capacity to an S3-enabled SVM

Create a bucket

S3 objects are kept in *buckets*. They are not nested as files inside a directory inside other directories.

Before you begin

A storage VM containing an S3 server must already exist.

About this task

- Beginning with ONTAP 9.14.1, automatic resizing has been enabled on S3 FlexGroup volumes when buckets are created on them. This eliminates excessive capacity allocation during bucket creation on existing and new FlexGroup volumes. FlexGroup volumes are resized to a minimum required size based on the following guidelines. The minimum required size is the total size of all the S3 buckets in a FlexGroup volume.
 - Beginning with ONTAP 9.14.1, if an S3 FlexGroup volume is created as part of a new bucket creation, the FlexGroup volume is created with the minimum required size.
 - If an S3 FlexGroup volume was created prior to ONTAP 9.14.1, the first bucket created or deleted subsequent to ONTAP 9.14.1 resizes the FlexGroup volume to the minimum required size.
 - If an S3 FlexGroup volume was created prior to ONTAP 9.14.1, and already had the minimum required size, the creation or deletion of a bucket subsequent to ONTAP 9.14.1 maintains the size of the S3 FlexGroup volume.
- Storage service levels are predefined adaptive Quality of Service (QoS) policy groups, with *value*, *performance*, and *extreme* default levels. Instead of one of the default storage service levels, you can also define a custom QoS policy group and apply it to a bucket. For more information about storage service definitions, see [Storage service definitions](#). For more information about performance management, see [Performance management](#). Beginning with ONTAP 9.8, when you provision storage, QoS is enabled by default. You can disable QoS or choose a custom QoS policy during the provisioning process or at a later time.

- If you are configuring local capacity tiering, you create buckets and users in a data storage VM, not in the system storage VM where the S3 server is located.
- For remote client access, you must configure buckets in an S3-enabled storage VM. If you create a bucket in a storage VM that is not S3-enabled, it will only be available for local tiering.
- Beginning with ONTAP 9.14.1, you can [create a bucket on a mirrored or unmirrored aggregate in a MetroCluster configuration](#).
- For the CLI, when you create a bucket, you have two provisioning options:
 - Let ONTAP select the underlying aggregates and FlexGroup components (default)
 - ONTAP creates and configures a FlexGroup volume for the first bucket by automatically selecting the aggregates. It will automatically select the highest service level available for your platform, or you can specify the storage service level. Any additional buckets you add later in the storage VM will have the same underlying FlexGroup volume.
 - Alternatively, you can specify whether the bucket will be used for tiering, in which case ONTAP tries to select low-cost media with optimal performance for the tiered data.
 - You select the underlying aggregates and FlexGroup components (requires advanced privilege command options): You have the option to manually select the aggregates on which the bucket and containing FlexGroup volume must be created, and then specifying the number of constituents on each aggregate. When adding additional buckets:
 - If you specify aggregates and constituents for a new bucket, a new FlexGroup will be created for the new bucket.
 - If you do not specify aggregates and constituents for a new bucket, the new bucket will be added to an existing FlexGroup. See [FlexGroup volumes management](#) for more information.

When you specify aggregates and constituents when creating a bucket, no QoS policy groups, default or custom, are applied. You can do so later with the `vserver object-store-server bucket modify` command.

See [vserver object-store-server bucket modify](#) for more information.

Note: If you are serving buckets from Cloud Volumes ONTAP, you should use the CLI procedure. It is strongly recommended that you manually select the underlying aggregates to ensure that they are using one node only. Using aggregates from both nodes can impact performance, because the nodes will be in geographically separated availability zones and hence susceptible to latency issues.

Create S3 buckets with the ONTAP CLI

1. If you plan to select aggregates and FlexGroup components yourself, set the privilege level to advanced (otherwise, admin privilege level is sufficient): `set -privilege advanced`
2. Create a bucket:

```
vserver object-store-server bucket create -vserver svm_name -bucket
bucket_name [-size integer[KB|MB|GB|TB|PB]] [-comment text]
[additional_options]
```

The storage VM name can be either a data storage VM or `Cluster` (the system storage VM name) if you are configuring local tiering.

If you specify no options, ONTAP creates an 800GB bucket with the service level set to the highest level

available for your system.

If you want ONTAP to create a bucket based on performance or usage, use one of the following options:

- service level

Include the `-storage-service-level` option with one of the following values: `value`, `performance`, or `extreme`.

- tiering

Include the `-used-as-capacity-tier true` option.

If you want to specify the aggregates on which to create the underlying FlexGroup volume, use the following options:

- The `-aggr-list` parameter specifies the list of aggregates to be used for FlexGroup volume constituents.

Each entry in the list creates a constituent on the specified aggregate. You can specify an aggregate multiple times to have multiple constituents created on the aggregate.

For consistent performance across the FlexGroup volume, all of the aggregates must use the same disk type and RAID group configurations.

- The `-aggr-list-multiplier` parameter specifies the number of times to iterate over the aggregates that are listed with the `-aggr-list` parameter when creating a FlexGroup volume.

The default value of the `-aggr-list-multiplier` parameter is 4.

3. Add a QoS policy group if needed:

```
vserver object-store-server bucket modify -bucket bucket_name -qos-policy  
-group qos_policy_group
```

4. Verify bucket creation:

```
vserver object-store-server bucket show [-instance]
```

Example

The following example creates a bucket for storage VM `vs1` of size 1TB and specifying the aggregate:

```
cluster-1::*> vserver object-store-server bucket create -vserver  
svml.example.com -bucket testbucket -aggr-list aggr1 -size 1TB
```

Create S3 buckets with System Manager

1. Add a new bucket on an S3-enabled storage VM.
 - a. Click **Storage > Buckets**, then click **Add**.
 - b. Enter a name, select the storage VM, and enter a size.

- If you click **Save** at this point, a bucket is created with these default settings:
 - No users are granted access to the bucket unless any group policies are already in effect.



You should not use the S3 root user to manage ONTAP object storage and share its permissions, because it has unlimited access to the object store. Instead, create a user or group with administrative privileges that you assign.

- A Quality of Service (performance) level that is the highest available for your system.
- Click **Save** to create a bucket with these default values.

Configure additional permissions and restrictions

You can click **More Options** to configure settings for object locking, user permissions, and performance level when you configure the bucket, or you can modify these settings later.

If you intend to use the S3 object store for FabricPool tiering, consider selecting **Use for tiering** (use low-cost media with optimal performance for the tiered data) rather than a performance service level.

If you want to enable versioning for your objects for later recovery, select **Enable Versioning**. Versioning is enabled by default if you are enabling object locking on the bucket. For information about object versioning, see the [Using versioning in S3 buckets for Amazon](#).

Beginning with 9.14.1, object locking is supported on S3 buckets. S3 object locking requires a standard SnapLock license. This license is included with [ONTAP One](#). Prior to ONTAP One, the SnapLock license was included in the Security and Compliance bundle. The Security and Compliance bundle is no longer offered but is still valid. Although not currently required, existing customers can choose to [upgrade to ONTAP One](#). If you are enabling object locking on a bucket, you should [verify that a SnapLock license is installed](#). If a SnapLock license is not installed, you must [install](#) it before you can enable object locking. When you have verified that the SnapLock license is installed, to protect objects in your bucket from getting deleted or overwritten, select **Enable object locking**. Locking can be enabled on either all or specific versions of objects, and only when the SnapLock compliance clock is initialized for the cluster nodes. Follow these steps:

1. If the SnapLock compliance clock is not initialized on any node of the cluster, the **Initialize SnapLock Compliance Clock** button appears. Click **Initialize SnapLock Compliance Clock** to initialize the SnapLock compliance clock on the cluster nodes.
2. Select **Governance** mode to activate a time-based lock that allows *Write once, read many (WORM)* permissions on the objects. Even in *Governance* mode, the objects can be deleted by administrator users with specific permissions.
3. Select **Compliance** mode if you want to assign stricter rules of deletion and update on the objects. In this mode of object locking, the objects can be expired only on the completion of the specified retention period. Unless a retention period is specified, the objects remain locked indefinitely.
4. Specify the retention tenure for the lock in days or years if you want the locking to be effective for a certain period.



Locking is applicable to versioned and non-versioned S3 buckets. Object locking is not applicable to NAS objects.

You can configure protection and permission settings, and performance service level for the bucket.



You must have already created user and groups before configuring the permissions.

For information, see [Create mirror for new bucket](#).

Verify access to the bucket

On S3 client applications (whether ONTAP S3 or an external third-party application), you can verify your access to the newly created bucket by entering the following:

- The S3 server CA certificate.
- The user's access key and secret key.
- The S3 server FQDN name and bucket name.

Create a bucket on a mirrored or unmirrored aggregate in a MetroCluster configuration

Beginning with ONTAP 9.14.1, you can provision a bucket on a mirrored or unmirrored aggregate in MetroCluster FC and IP configurations.

About this task

- By default, buckets are provisioned on mirrored aggregates.
- The same provisioning guidelines outlined in [Create a bucket](#) apply to creating a bucket in a MetroCluster environment.
- The following S3 object storage features are **not** supported in MetroCluster environments:
 - S3 SnapMirror
 - S3 bucket lifecycle management
 - S3 object lock in **Compliance** mode



S3 object lock in **Governance** mode is supported.

- Local FabricPool tiering

Before you begin

An SVM containing an S3 server must already exist.

Process to create buckets

CLI

1. If you plan to select aggregates and FlexGroup components yourself, set the privilege level to advanced (otherwise, admin privilege level is sufficient): `set -privilege advanced`
2. Create a bucket:

```
vserver object-store-server bucket create -vserver <svm_name> -bucket  
<bucket_name> [-size integer[KB|MB|GB|TB|PB]] [-use-mirrored-aggregates  
true/false]
```

Set the `-use-mirrored-aggregates` option to `true` or `false` depending on whether you want to use a mirrored or unmirrored aggregate.



By default, the `-use-mirrored-aggregates` option is set to `true`.

- The SVM name must be a data SVM.
- If you specify no options, ONTAP creates an 800GB bucket with the service level set to the highest level available for your system.
- If you want ONTAP to create a bucket based on performance or usage, use one of the following options:
 - **service level**

Include the `-storage-service-level` option with one of the following values: `value`, `performance`, or `extreme`.
 - **tiering**

Include the `-used-as-capacity-tier true` option.
- If you want to specify the aggregates on which to create the underlying FlexGroup volume, use the following options:
 - The `-aggr-list` parameter specifies the list of aggregates to be used for FlexGroup volume constituents.

Each entry in the list creates a constituent on the specified aggregate. You can specify an aggregate multiple times to have multiple constituents created on the aggregate.

For consistent performance across the FlexGroup volume, all of the aggregates must use the same disk type and RAID group configurations.
 - The `-aggr-list-multiplier` parameter specifies the number of times to iterate over the aggregates that are listed with the `-aggr-list` parameter when creating a FlexGroup volume.

The default value of the `-aggr-list-multiplier` parameter is 4.

3. Add a QoS policy group if needed:

```
vserver object-store-server bucket modify -bucket bucket_name -qos-policy  
-group qos_policy_group
```

4. Verify bucket creation:

```
vserver object-store-server bucket show [-instance]
```

Example

The following example creates a bucket for SVM vs1 of size 1TB on a mirrored aggregate:

```
cluster-1::*> vserver object-store-server bucket create -vserver  
svml.example.com -bucket testbucket -size 1TB -use-mirrored-aggregates  
true
```


System Manager

1. Add a new bucket on an S3-enabled storage VM.
 - a. Click **Storage > Buckets**, then click **Add**.
 - b. Enter a name, select the storage VM, and enter a size.

By default, the bucket is provisioned on a mirrored aggregate. If you want to create a bucket on an unmirrored aggregate, select **More Options** and uncheck the **Use the SyncMirror tier** box under **Protection** as shown in the following image:

Add bucket ✕

NAME

 To use this bucket from a remote cluster, configure S3 service on storage VM "vs1".

FOLDER (OPTIONAL)

Browse

Specify the folder to map to this bucket. [Know more](#)

CAPACITY

Size

GB

▼

☐ Use tiering

If you select this option, the system will try to select low-cost media with optimal performance for the tiered data.

☐ Enable versioning

Versioning-enabled buckets allow you to recover objects that were accidentally deleted or overwritten. After versioning is enabled, it can't be disabled. However, you can suspend versioning.

PERFORMANCE SERVICE LEVEL

Value

▼

Not sure? [Get help selecting type](#)

Permissions

☐ Copy access permissions from an existing bucket

Principal	Effect	Actions	Resources	Conditions
All users of this stor...	allow	ListBucket	*	

+ Add

Object locking

☐ Enable object locking

Object locking utilizes the "Write Once, Read Many" (WORM) model in which objects or their versions are protected from being deleted or overwritten during the specified retention period.

Protection

☒ Use the S3x33n3r3n3t3n

Save

Cancel

- If you click **Save** at this point, a bucket is created with these default settings:
 - No users are granted access to the bucket unless any group policies are already in effect.



You should not use the S3 root user to manage ONTAP object storage and share its permissions, because it has unlimited access to the object store. Instead, create a user or group with administrative privileges that you assign.

- A Quality of Service (performance) level that is the highest available for your system.
- You can click **More Options** to configure user permissions and performance level when you configure the bucket, or you can modify these settings later.
 - You must have already created user and groups before using **More Options** to configure their permissions.

- If you intend to use the S3 object store for FabricPool tiering, consider selecting **Use for tiering** (use low-cost media with optimal performance for the tiered data) rather than a performance service level.
2. On S3 client apps – another ONTAP system or an external 3rd-party app – verify access to the new bucket by entering the following:
 - The S3 server CA certificate.
 - The user's access key and secret key.
 - The S3 server FQDN name and bucket name.

Create a bucket lifecycle management rule

Beginning with ONTAP 9.13.1, you can create lifecycle management rules to manage object lifecycles in your S3 buckets. You can define deletion rules for specific objects in a bucket, and through these rules, expire those bucket objects. This enables you to meet retention requirements and manage overall S3 object storage efficiently.



If object locking is enabled for your bucket objects, the lifecycle management rules for object expiration will not be applied on locked objects. For information about object locking, see [Create a bucket](#).

Before you begin

An S3-enabled SVM containing an S3 server and a bucket must already exist. See [Create an SVM for S3](#) for more information.

About this task

When creating your lifecycle management rules, you can apply the following deletion actions to your bucket objects:

- Deletion of current versions - This action expires objects identified by the rule. If versioning is enabled on the bucket, S3 makes all expired objects unavailable. If versioning is not enabled, this rule deletes the objects permanently. The CLI action is `Expiration`.
- Deletion of non-current versions - This action specifies when S3 can permanently remove non-current objects. The CLI action is `NoncurrentVersionExpiration`.
- Deletion of expired delete markers - This action deletes expired object delete markers. In versioning-enabled buckets, objects with a delete markers become the current versions of the objects. The objects are not deleted, and no action can be performed on them. These objects become expired when there are no current versions associated with them. The CLI action is `Expiration`.
- Deletion of incomplete multipart uploads - This action sets a maximum time (in days) that you want to allow multipart uploads to remain in progress. Following which, they are deleted. The CLI action is `AbortIncompleteMultipartUpload`.

The procedure you follow depends on the interface that you use. With ONTAP 9.13.1, you need to use the CLI. Beginning with ONTAP 9.14.1, you can also use System Manager.

Manage lifecycle management rules with the CLI

Beginning with ONTAP 9.13.1, you can use the ONTAP CLI to create lifecycle management rules to expire objects in your S3 buckets.

Before you begin

For the CLI, you need to define the required fields for each expiration action type when creating a bucket lifecycle management rule. These fields can be modified after initial creation. The following table displays the unique fields for each action type.

Action type	Unique fields
NonCurrentVersionExpiration	<ul style="list-style-type: none">• <code>-non-curr-days</code> - Number of days after which non-current versions will be deleted• <code>-new-non-curr-versions</code> - Number of latest non-current versions to be retained
Expiration	<ul style="list-style-type: none">• <code>-obj-age-days</code> - Number of days since creation, after which current version of objects can be deleted• <code>-obj-exp-date</code> - Specific date when the objects should expire• <code>-expired-obj-del-markers</code> - Cleanup object delete markers
AbortIncompleteMultipartUpload	<ul style="list-style-type: none">• <code>-after-initiation-days</code> - Number of days of initiation, after which upload can be aborted

In order for the bucket lifecycle management rule to only be applied to a specific subset of objects, admins must set each filter when creating the rule. If these filters are not set when creating the rule, the rule will be applied to all objects within the bucket.

All filters can be modified after initial creation *except* for the following:

- `-prefix`
- `-tags`
- `-obj-size-greater-than`
- `-obj-size-less-than`

Steps

1. Use the `vserver object-store-server bucket lifecycle-management-rule create` command with required fields for your expiration action type to create your bucket lifecycle management rule.

Example

The following command creates a `NonCurrentVersionExpiration` bucket lifecycle management rule:

```
vserver object-store-server bucket lifecycle-management-rule create
-vserver <svm_name> -bucket <bucket_name> -rule-id <rule_name> -action
NonCurrentVersionExpiration -index <lifecycle_rule_index_integer> -is
-enabled {true|false} -prefix <object_name> -tags <text> -obj-size-greater
-than {<integer>[KB|MB|GB|TB|PB]} -obj-size-less-than
{<integer>[KB|MB|GB|TB|PB]} -new-non-curr-versions <integer> -non-curr
-days <integer>
```

Example

The following command creates an Expiration bucket lifecycle management rule:

```
vserver object-store-server bucket lifecycle-management-rule create
-vserver <svm_name> -bucket <bucket_name> -rule-id <rule_name> -action
Expiration -index <lifecycle_rule_index_integer> -is-enabled {true|false}
-prefix <object_name> -tags <text> -obj-size-greater-than
{<integer>[KB|MB|GB|TB|PB]} -obj-size-less-than
{<integer>[KB|MB|GB|TB|PB]} -obj-age-days <integer> -obj-exp-date
<"MM/DD/YYYY HH:MM:SS"> -expired-obj-del-marker {true|false}
```

Example


The following command creates an AbortIncompleteMultipartUpload bucket lifecycle management rule:

```
vserver object-store-server bucket lifecycle-management-rule create
-vserver <svm_name> -bucket <bucket_name> -rule-id <rule_name> -action
AbortIncompleteMultipartUpload -index <lifecycle_rule_index_integer> -is
-enabled {true|false} -prefix <object_name> -tags <text> -obj-size-greater
-than {<integer>[KB|MB|GB|TB|PB]} -obj-size-less-than
{<integer>[KB|MB|GB|TB|PB]} -after-initiation-days <integer>
```

Manage lifecycle management rules with System Manager

Beginning with ONTAP 9.14.1, you can expire S3 objects by using System Manager. You can add, edit, and delete lifecycle management rules for your S3 objects. Additionally, you can import a lifecycle rule created for one bucket and utilize it for the objects in another bucket. You can disable an active rule and enable it later.

Add a lifecycle management rule

1. Click **Storage > Buckets**.
2. Select the bucket for which you want to specify the expiration rule.
3. Click the  icon and select **Manage lifecycle rules**.
4. Click **Add > Lifecycle rule**.
5. On the Add a lifecycle rule page, add the name of the rule.
6. Define the scope of the rule, whether you want it to apply to all the objects in the bucket or on specific

objects. If you want to specify objects, add at least one of the following filter criteria:

- a. **Prefix:** Specify a prefix of the object key names to which the rule should apply. Typically, it is the path or folder of the object. You can enter one prefix per rule. Unless a valid prefix is provided, the rule applies to all the objects in a bucket.
- b. **Tags:** Specify up to three key and value pairs (tags) for the objects to which the rule should apply. Only valid keys are used for filtering. The value is optional. However, if you add values, ensure that you add only valid values for the corresponding keys.
- c. **Size:** You can limit the scope between the minimum and maximum sizes of the objects. You can enter either or both the values. The default unit is MiB.

7. Specify the action:


- a. **Expire the current version of objects:** Set a rule to make all current objects permanently unavailable after a specific number of days since their creation, or on a specific date. This option is unavailable if the **Delete expired object delete markers** option is selected.
- b. **Permanently delete noncurrent versions:** Specify the number of days after which the version becomes non-current, and thereafter can be deleted, and the number of versions to retain.
- c. **Delete expired object delete markers:** Select this action to delete objects with expired delete markers, that is delete markers without an associated current object.



This option becomes unavailable when you select the **Expire the current version of objects** option that automatically deletes all objects after the retention period. This option also becomes unavailable when object tags are used for filtering.

- d. **Delete incomplete multipart uploads:** Set the number of days after which incomplete multipart uploads are to be deleted. If the multipart uploads that are in progress fail within the specified retention period, you can delete the incomplete multipart uploads. This option becomes unavailable when object tags are used for filtering.
- e. Click **Save**.


Import a lifecycle rule

1. Click **Storage > Buckets**.
2. Select the bucket for which you want to import the expiration rule.
3. Click the  icon and select **Manage lifecycle rules**.
4. Click **Add > Import a rule**.
5. Select the bucket from which you want to import the rule. The lifecycle management rules defined for the selected bucket appear.
6. Select the rule that you want to import. You have the option to select one rule at a time, with the default selection being the first rule.
7. Click **Import**.

Edit, delete, or disable a rule

You can only edit the lifecycle management actions associated with the rule. If the rule was filtered with object tags, then the **Delete expired object delete markers** and **Delete incomplete multipart uploads** options are unavailable.

When you delete a rule, that rule will no longer apply to previously associated objects.

1. Click **Storage > Buckets**.
2. Select the bucket for which you want to edit, delete, or disable the lifecycle management rule.
3. Click the  icon and select **Manage lifecycle rules**.
4. Select the required rule. You can edit and disable one rule at a time. You can delete multiple rules at once.
5. Select **Edit**, **Delete**, or **Disable**, and complete the procedure.

Create an S3 user

User authorization is required on all ONTAP object stores to restrict connectivity to authorized clients.

Before you begin.

An S3-enabled storage VM must already exist.

About this task

An S3 user can be granted access to any bucket in a storage VM. When you create an S3 user, an access key and a secret key are also generated for the user. They should be shared with the user along with the FQDN of the object store and bucket name. An S3 users' keys can be viewed with the `vserver object-store-server user show` command.

You can grant specific access permissions to S3 users in a bucket policy or an object server policy.



When you create a new object store server, ONTAP creates a root user (UID 0), which is a privileged user with access to all buckets. Rather than administering ONTAP S3 as the root user, NetApp recommends that an admin user role be created with specific privileges.

CLI

1. Create an S3 user:

```
vserver object-store-server user create -vserver svm_name -user user_name  
-comment [-comment text] -key-time-to-live time
```


- Adding a comment is optional.
- Beginning with ONTAP 9.14.1, you can define the period of time for which the key will be valid in the `-key-time-to-live` parameter. You can add the retention period in this format, to indicate the period after which the access key expires:
`P[<integer>D]T[<integer>H][<integer>M][<integer>S] | P<integer>W`
For example, if you want to enter a retention period of one day, two hours, three minutes, and four seconds, enter the value as `P1DT2H3M4S`. Unless specified, the key is valid for an indefinite period of time.

The below example creates a user with name `sm_user1` on storage VM `vs0`, with a key retention period of one week.

```
vserver object-store-server user create -vserver vs0 -user  
sm_user1 -key-time-to-live P1W
```

- ### 2. Be sure to save the access key and secret key. They will be required for access from S3 clients.

System Manager

1. Click **Storage > Storage VMs**. Select the storage VM to which you need to add a user, select **Settings** and then click  under S3.
2. To add a user, click **Users > Add**.
3. Enter a name for the user.
4. Beginning with ONTAP 9.14.1, you can specify the retention period of the access keys that get created for the user. You can specify the retention period in days, hours, minutes, or seconds, after which the keys automatically expire. By default, the value is set to 0 that indicates that the key is indefinitely valid.
5. Click **Save**. The user is created, and an access key and a secret key are generated for the user.
6. Download or save the access key and secret key. They will be required for access from S3 clients.

Next steps

- [Create or modify S3 groups](#)

Create or modify S3 groups

You can simplify bucket access by creating groups of users with appropriate access authorizations.

Before you begin

S3 users in an S3-enabled SVM must already exist.

About this task

Users in an S3 group can be granted access to any bucket in an SVM but not in multiple SVMs. Group access permissions can be configured in two ways:


- At the bucket level

After creating a group of S3 users, you specify group permissions in bucket policy statements and they apply only to that bucket.

- At the SVM level

After creating a group of S3 users, you specify object server policy names in the group definition. Those policies determine the buckets and access for the group members.

System Manager

1. Edit the storage VM: click **Storage > storage VMs**, click the storage VM, click **Settings** and then click  under S3.
2. Add a group: select **Groups**, then select **Add**.
3. Enter a group name and select from a list of users.
4. You can select an existing group policy or add one now, or you can add a policy later.

CLI

1. Create an S3 group:

```
vserver object-store-server group create -vserver svm_name -name group_name  
-users user_name\(s\) [-policies policy_names] [-comment text\]
```

The `-policies` option can be omitted in configurations with only one bucket in an object store; the group name can be added to the bucket policy.

The `-policies` option can be added later with the `vserver object-store-server group modify` command after object storage server policies are created.

Regenerate keys and modify their retention period

Access keys and secret keys are automatically generated during user creation for enabling S3 client access. You can regenerate keys for a user if a key is expired or compromised.

For information about generation of access keys, see [Create an S3 user](#).



CLI

1. Regenerate access and secret keys for a user by running the `vserver object-store-server user regenerate-keys` command.
2. By default, generated keys are valid indefinitely. Beginning with 9.14.1, you can modify their retention period, after which the keys automatically expire. You can add the retention period in this format:
`P[<integer>D]T[<integer>H][<integer>M][<integer>S] | P<integer>W`
For example, if you want to enter a retention period of one day, two hours, three minutes, and four seconds, enter the value as `P1DT2H3M4S`.

```
vserver object-store-server user regenerate-keys -vserver svm_name  
-user user -key-time-to-live 0
```

3. Save the access and secret keys. They will be required for access from S3 clients.

System Manager

1. Click **Storage > Storage VMs** and then select the storage VM.
2. In the **Settings** tab, click  in the **S3** tile.
3. In the **Users** tab, verify that there is no access key, or the key has expired for the user.
4. If you need to regenerate the key, click  next to the user, then click **Regenerate Key**.
5. By default, generated keys are valid for an indefinite amount of time. Beginning with 9.14.1, you can modify their retention period, after which the keys automatically expire. Enter the retention period in days, hours, minutes, or seconds.
6. Click **Save**. The key is regenerated. Any change in the key retention period takes effect immediately.
7. Download or save the access key and secret key. They will be required for access from S3 clients.

Create or modify access policy statements

About bucket and object store server policies

User and group access to S3 resources is controlled by bucket and object store server policies. If you have a small number of users or groups, controlling access at the bucket level is probably sufficient, but if you have many users and groups, it is easier to control access at the object store server level.

Modify a bucket policy

You can add access rules to the default bucket policy. The scope of its access control is the containing bucket, so it is most appropriate when there is a single bucket.

Before you begin

An S3-enabled storage VM containing an S3 server and a bucket must already exist.

You must have already created users or groups before granting permissions.

About this task

You can add new statements for new users and groups, or you can modify the attributes of existing statements. For more options, see the `vserver object-store-server bucket policy man` pages.

User and group permissions can be granted when the bucket is created or as needed later. You can also modify the bucket capacity and QoS policy group assignment.

Beginning with ONTAP 9.9.1, if you plan to support AWS client object tagging functionality with the ONTAP S3 server, the actions `GetObjectTagging`, `PutObjectTagging`, and `DeleteObjectTagging` need to be allowed using the bucket or group policies.

The procedure you follow depends on the interface that you use—System Manager or the CLI:

System Manager

Steps

1. Edit the bucket: click **Storage > Buckets**, click the desired bucket, and then click **Edit**.

When adding or modifying permissions, you can specify the following parameters:

- **Principal**: the user or group to whom access is granted.
- **Effect**: allows or denies access to a user or group.
- **Actions**: permissible actions in the bucket for a given user or group.
- **Resources**: paths and names of objects within the bucket for which access is granted or denied.

The defaults **bucketname** and **bucketname/*** grant access to all objects in the bucket. You can also grant access to single objects; for example, **bucketname/*_readme.txt**.

- **Conditions** (optional): expressions that are evaluated when access is attempted. For example, you can specify a list of IP addresses for which access will be allowed or denied.



Beginning with ONTAP 9.14.1, you can specify variables for the bucket policy in the **Resources** field. These variables are placeholders that are replaced with contextual values when the policy is evaluated. For example, If `${aws:username}` is specified as a variable for a policy, then this variable is replaced with the request context username, and the policy action can be performed as configured for that user.

CLI

Steps

1. Add a statement to a bucket policy:

```
vserver object-store-server bucket policy add-statement -vserver svm_name
-bucket bucket_name -effect {allow|deny} -action object_store_actions
-principal user_and_group_names -resource object_store_resources [-sid
text] [-index integer]
```

The following parameters define access permissions:

-effect	The statement may allow or deny access
-action	You can specify * to mean all actions, or a list of one or more of the following: GetObject, PutObject, DeleteObject, ListBucket, GetBucketAcl, GetObjectAcl, ListBucketMultipartUploads, and ListMultipartUploadParts.

-principal	<p>A list of one or more S3 users or groups.</p> <ul style="list-style-type: none"> • A maximum of 10 users or groups can be specified. • If an S3 group is specified, it must be in the form <code>group/group_name</code>. • * can be specified to mean public access; that is, access without an access-key and secret-key. • If no principal is specified, all S3 users in the storage VM are granted access.
-resource	<p>The bucket and any object it contains. The wildcard characters * and ? can be used to form a regular expression for specifying a resource. For a resource, you can specify variables in a policy. These are policy variables are placeholders that are replaced with the contextual values when the policy is evaluated.</p>

You can optionally specify a text string as comment with the `-sid` option.

Examples

The following example creates an object store server bucket policy statement for the storage VM `svm1.example.com` and `bucket1` which specifies allowed access to a `readme` folder for object store server user `user1`.

```
cluster1::> vsserver object-store-server bucket policy statement create
-vserver svm1.example.com -bucket bucket1 -effect allow -action
GetObject,PutObject,DeleteObject,ListBucket -principal user1 -resource
bucket1/readme/* -sid "fullAccessToReadmeForUser1"
```

The following example creates an object store server bucket policy statement for the storage VM `svm1.example.com` and `bucket1` which specifies allowed access to all objects for object store server group `group1`.

```
cluster1::> vsserver object-store-server bucket policy statement create
-vserver svm1.example.com -bucket bucket1 -effect allow -action
GetObject,PutObject,DeleteObject,ListBucket -principal group/group1
-resource bucket1/* -sid "fullAccessForGroup1"
```

Beginning with ONTAP 9.14.1, you can specify variables for a bucket policy. The following example creates a server bucket policy statement for the storage VM `svm1` and `bucket1`, and specifies `${aws:username}` as a variable for a policy resource. When the policy is evaluated, the policy variable is replaced with the request context username, and the policy action can be performed as configured for that user. For example, when the following policy statement is evaluated, `${aws:username}` is replaced with the user performing the S3 operation. If a user `user1` performs the operation, that user is granted access to `bucket1` as `bucket1/user1/*`.

```
cluster1::> object-store-server bucket policy statement create -vserver  
svm1 -bucket bucket1 -effect allow -action * -principal - -resource  
bucket1,bucket1/${aws:username}/*##
```

Create or modify an object store server policy

You can create policies that can apply to one or more buckets in an object store. Object store server policies can be attached to groups of users, thereby simplifying the management of resource access across multiple buckets.

Before you begin

An S3-enabled SVM containing an S3 server and a bucket must already exist.

About this task

You can enable access policies at the SVM level by specifying a default or custom policy in an object storage server group. The policies do not take effect until they are specified in the group definition.



When you use object storage server policies, you specify principals (that is, users and groups) in the group definition, not in the policy itself.

There are three read-only default policies for access to ONTAP S3 resources:

- FullAccess
- NoS3Access
- ReadOnlyAccess

You can also create new custom policies, then add new statements for new users and groups, or you can modify the attributes of existing statements. For more options, see the `vserver object-store-server policy` [command reference](#).


Beginning with ONTAP 9.9.1, if you plan to support AWS client object tagging functionality with the ONTAP S3 server, the actions `GetObjectTagging`, `PutObjectTagging`, and `DeleteObjectTagging` need to be allowed using the bucket or group policies.

The procedure you follow depends on the interface that you use—System Manager or the CLI:

System Manager

Use System Manager to create or modify an object store server policy

Steps

1. Edit the storage VM: click **Storage > storage VMs**, click the storage VM, click **Settings** and then click  under S3.
2. Add a user: click **Policies**, then click **Add**.
 - a. Enter a policy name and select from a list of groups.
 - b. Select an existing default policy or add a new one.

When adding or modifying a group policy, you can specify the following parameters:

- Group: the groups to whom access is granted.
- Effect: allows or denies access to one or more groups.
- Actions: permissible actions in one or more buckets for a given group.
- Resources: paths and names of objects within one or more buckets for which access is granted or denied.

For example:

- * grants access to all buckets in the storage VM.
 - **bucketname** and **bucketname/*** grant access to all objects in a specific bucket.
 - **bucketname/readme.txt** grants access to an object in a specific bucket.
- c. If desired, add statements to existing policies.

CLI

Use the CLI to create or modify an object store server policy

Steps

1. Create an object storage server policy:

```
vserver object-store-server policy create -vserver svm_name -policy policy_name [-comment text]
```

2. Create a statement for the policy:

```
vserver object-store-server policy statement create -vserver svm_name -policy policy_name -effect {allow|deny} -action object_store_actions -resource object_store_resources [-sid text]
```

The following parameters define access permissions:

-effect	The statement may allow or deny access
---------	--

<code>-action</code>	You can specify <code>*</code> to mean all actions, or a list of one or more of the following: <code>GetObject</code> , <code>PutObject</code> , <code>DeleteObject</code> , <code>ListBucket</code> , <code>GetBucketAcl</code> , <code>GetObjectAcl</code> , <code>ListAllMyBuckets</code> , <code>ListBucketMultipartUploads</code> , and <code>ListMultipartUploadParts</code> .
<code>-resource</code>	The bucket and any object it contains. The wildcard characters <code>*</code> and <code>?</code> can be used to form a regular expression for specifying a resource.

You can optionally specify a text string as comment with the `-sid` option.

By default, new statements are added to the end of the list of statements, which are processed in order. When you add or modify statements later, you have the option to modify the statement's `-index` setting to change the processing order.

Configure S3 access for external directory services

Beginning with ONTAP 9.14.1, services for external directories have been integrated with ONTAP S3 object storage. This integration simplifies user and access management through external directory services.

You can provide user groups belonging to an external directory service with access to your ONTAP object storage environment. Lightweight Directory Access Protocol (LDAP) is an interface for communicating with directory services, such as Active Directory, that provide a database and services for identity and access management (IAM). To provide access, you need to configure LDAP groups in your ONTAP S3 environment. After you have configured access, the group members have permissions to ONTAP S3 buckets. For information about LDAP, see [Overview of using LDAP](#).

You can also configure Active Directory user groups for fast bind mode, so that user credentials can be validated and third-party and open-source S3 applications can be authenticated over LDAP connections.

Before you begin

Ensure the following before configuring LDAP groups and enabling the fast bind mode for group access:

1. An S3-enabled storage VM containing an S3 server has been created. See [Create an SVM for S3](#).
2. A bucket has been created in that storage VM. See [Create a bucket](#).
3. DNS is configured on the storage VM. See [Configure DNS services](#).
4. A self-signed root certification authority (CA) certificate of the LDAP server is installed on the storage VM. See [Install the self-signed root CA certificate on the SVM](#).
5. An LDAP client is configured with TLS enabled on the SVM. See [Create an LDAP client configuration](#) and [Associate the LDAP client configuration with SVMs for information](#).

Configure S3 access for external directory services

1. Specify LDAP as the *name service database* of the SVM for the group and password to LDAP:

```
ns-switch modify -vserver <vserver-name> -database group -sources
files,ldap
ns-switch modify -vserver <vserver-name> -database passwd -sources
files,ldap
```

For more information about this command, see the [vserver services name-service ns-switch modify](#) command.

2. Create an object store bucket policy statement with the `principal` set to the LDAP group to which you want to grant access:

```
object-store-server bucket policy statement create -bucket <bucket-name>
-effect allow -principal nasgroup/<ldap-group-name> -resource <bucket-
name>, <bucket-name>/*
```

Example: The following example creates a bucket policy statement for `buck1`. The policy allows access for the LDAP group `group1` to the resource (bucket and its objects) `buck1`.

```
vserver object-store-server bucket policy add-statement -bucket buck1
-effect allow -action
GetObject,PutObject,DeleteObject,ListBucket,GetBucketAcl,GetObjectAcl,Li
stBucketMultipartUploads,ListMultipartUploadParts,
ListBucketVersions,GetObjectTagging,PutObjectTagging,DeleteObjectTagging
,GetBucketVersioning,PutBucketVersioning -principal nasgroup/group1
-resource buck1, buck1/*
```

3. Verify that a user from the LDAP group `group1` is able to perform S3 operations from the S3 client.

Use LDAP fast bind mode for authentication

1. Specify LDAP as the *name service database* of the SVM for the group and password to LDAP:

```
ns-switch modify -vserver <vserver-name> -database group -sources
files,ldap
ns-switch modify -vserver <vserver-name> -database passwd -sources
files,ldap
```

For more information about this command, see the [vserver services name-service ns-switch modify](#) command.

2. Ensure that an LDAP user accessing the S3 bucket has permissions defined in the bucket-policies. For

more information, see [Modify a bucket policy](#).

3. Verify that a user from the LDAP group can perform the following operations:

a. Configure the access key on the S3 client in this format:

"NTAPFASTBIND" + base64-encode(user-name:password)

Example: "NTAPFASTBIND" + base64-encode(ldapuser:password), which results in

NTAPFASTBINDbGRhcHVzZXI6cGFzc3dvcmQ=



The S3 client might prompt for a secret key. In the absence of a secret key, any password of at least 16 characters can be entered.

b. Perform basic S3 operations from the S3 client for which the user has permissions.

Enable LDAP or domain users to generate their own S3 access keys

Beginning with ONTAP 9.14.1, as an ONTAP administrator, you can create custom roles and grant them to local or domain groups or Lightweight Directory Access Protocol (LDAP) groups, so that the users belonging to those groups can generate their own access and secret keys for S3 client access.

You have to perform a few configuration steps on your storage VM, so that the custom role can be created and assigned to the user that invokes the API for access key generation.

Before you begin

Ensure the following:

1. An S3-enabled storage VM containing an S3 server has been created. See [Create an SVM for S3](#).
2. A bucket has been created in that storage VM. See [Create a bucket](#).
3. DNS is configured on the storage VM. See [Configure DNS services](#).
4. A self-signed root certification authority (CA) certificate of the LDAP server is installed on the storage VM. See [Install the self-signed root CA certificate on the SVM](#).
5. An LDAP client is configured with TLS enabled on the storage VM. See [Create an LDAP client configuration](#) and .
6. Associate the client configuration with the Vserver. See [Associate the LDAP client configuration with SVMs and vservice services name-service ldap create](#).
7. If you are using a data storage VM, create a management network interface (LIF) and on the VM, and also a service policy for the LIF. See the [network interface create](#) and [network interface service-policy create](#) commands.

Configure users for access key generation

1. Specify LDAP as the *name service database* of the storage VM for the group and password to LDAP:

```
ns-switch modify -vserver <vserver-name> -database group -sources
files,ldap
ns-switch modify -vserver <vserver-name> -database passwd -sources
files,ldap
```

For more information about this command, see the [vserver services name-service ns-switch modify](#) command.

2. Create a custom role with access to S3 user REST API endpoint:

```
security login rest-role create -vserver <vserver-name> -role <custom-role-name> -api "/api/protocols/s3/services/*/users" -access <access-type>
```

In this example, the `s3-role` role is generated for users on the storage VM `svm-1`, to which all access rights, read, create, and update are granted.

```
security login rest-role create -vserver svm-1 -role s3role -api  
"/api/protocols/s3/services/*/users" -access all
```

For more information about this command, see the [security login rest-role create](#) command.

3. Create an LDAP user group with the security login command and add the new custom role for accessing the S3 user REST API endpoint. For more information about this command, see the [security login create](#) command.

```
security login create -user-or-group-name <ldap-group-name> -application  
http -authentication-method nsswitch -role <custom-role-name> -is-ns  
-switch-group yes
```

In this example, the LDAP group `ldap-group-1` is created in `svm-1`, and the custom role `s3role` is added to it for accessing the API endpoint, along with enabling LDAP access in the fast bind mode.

```
security login create -user-or-group-name ldap-group-1 -application http  
-authentication-method nsswitch -role s3role -is-ns-switch-group yes  
-second-authentication-method none -vserver svm-1 -is-ldap-fastbind yes
```

For more information, see [Use LDAP fast bind for nsswitch authentication](#).

Adding the custom role to the domain or LDAP group allows users in that group a limited access to the ONTAP `/api/protocols/s3/services/{svm.uid}/users` endpoint. By invoking the API, the domain or LDAP group users can generate their own access and secret keys to access the S3 client. They can generate the keys for only themselves and not for other users.

As an S3 or LDAP user, generate your own access keys

Beginning with ONTAP 9.14.1, you can generate your own access and secret keys for accessing S3 clients, if your administrator has granted you the role to generate your own keys. You can generate keys for only yourself by using the following ONTAP REST API endpoint.

HTTP method and endpoint

This REST API call uses the following method and endpoint. For information about the other methods of this endpoint, see the reference [API documentation](#).

HTTP method	Path
POST	/api/protocols/s3/services/{svm.uuid}/users

Curl example

```
curl
--request POST \
--location "https://$FQDN_IP /api/protocols/s3/services/{svm.uuid}/users "
\
--include \
--header "Accept: */*" \
--header "Authorization: Basic $BASIC_AUTH"
--data '{"name": "_name_"}'
```

JSON output example

```
{
  "records": [
    {
      "access_key":
"Pz3SB54G2B_6dsXQPrA5HrTPcf478qoAW6_Xx6qyqZ948AgZ_7YfCf_9nO87YoZmskxx3cq41
U2JAH2M3_fs321B4rkzS3a_oC5_8u7D8j_45N8OsBCBPWGD_1d_ccfq",
      "_links": {
        "next": {
          "href": "/api/resourcelink"
        },
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "user-1",
      "secret_key":
"A20_tDhC_cux2C2BmtL45bXB_a_Q65c_96FsAcOdo14Az8V31jBKDTc0uCL62Bh559gPB8s9r
rn0868QrF38_1dsV2u1_9H2tSf3qQ5xp9NT259C6z_GiZQ883Qn63X1"
    }
  ],
  "num_records": "1"
}
```

Enable client access to S3 object storage

Enable ONTAP S3 access for remote FabricPool tiering

For ONTAP S3 to be used as a remote FabricPool capacity (cloud) tier, the ONTAP S3 administrator must provide information about the S3 server configuration to the remote

ONTAP cluster administrator.

About this task

The following S3 server information is required to configure FabricPool cloud tiers:

- server name (FQDN)
- bucket name
- CA certificate
- access key
- password (secret access key)

In addition, the following networking configuration is required:

- There must be an entry for the remote ONTAP S3 server's hostname in the DNS server configured for the admin SVM, including the S3 server's FQDN name and the IP addresses on its LIFs.
- Intercluster LIFs must be configured on the local cluster, although cluster peering is not required.

See the FabricPool documentation about configuring ONTAP S3 as a cloud tier.

[Managing Storage Tiers By Using FabricPool](#)

Enable ONTAP S3 access for local FabricPool tiering

For ONTAP S3 to be used as a local FabricPool capacity tier, you must define an object store based on the bucket you created, and then attach the object store to a performance tier aggregate to create a FabricPool.

Before you begin

You must have the ONTAP S3 server name and a bucket name, and the S3 server must have been created using cluster LIFs (with the `-vserver Cluster` parameter).

About this task

The object-store configuration contains information about the local capacity tier, including the S3 server and bucket names and authentication requirements.

An object-store configuration once created must not be reassociated with a different object-store or bucket. You can create multiple buckets for local tiers, but you cannot create multiple object stores in a single bucket.

A FabricPool license is not required for a local capacity tier.

Steps

1. Create the object store for the local capacity tier:

```
storage aggregate object-store config create -object-store-name store_name
-ipspace Cluster -provider-type ONTAP_S3 -server S3_server_name -container
-name bucket_name -access-key access_key -secret-password password
```

- The `-container-name` is the S3 bucket you created.
- The `-access-key` parameter authorizes requests to the ONTAP S3 server.

- The `-secret-password` parameter (secret access key) authenticates requests to the ONTAP S3 server.
- You can set the `-is-certificate-validation-enabled` parameter to `false` to disable certificate checking for ONTAP S3.

```
cluster1::> storage aggregate object-store config create
-object-store-name MyLocalObjStore -ipSpace Cluster -provider-type
ONTAP_S3 -server s3.example.com
-container-name bucket1 -access-key myS3key -secret-password myS3pass
```

2. Display and verify the object store configuration information:

```
storage aggregate object-store config show
```

3. Optional: To see how much data in a volume is inactive, follow the steps in [Determining how much data in a volume is inactive by using inactive data reporting](#).

Seeing how much data in a volume is inactive can help you decide which aggregate to use for FabricPool local tiering.

4. Attach the object store to an aggregate:

```
storage aggregate object-store attach -aggregate aggr_name -object-store-name
store_name
```

You can use the `allow-flexgroup` **true** option to attach aggregates that contain FlexGroup volume constituents.

```
cluster1::> storage aggregate object-store attach
-aggregate aggr1 -object-store-name MyLocalObjStore
```

5. Display the object store information and verify that the attached object store is available:

```
storage aggregate object-store show
```

```
cluster1::> storage aggregate object-store show
```

Aggregate	Object Store Name	Availability State
-----	-----	-----
aggr1	MyLocalObjStore	available

Enable client access from an S3 app

For S3 client apps to access the ONTAP S3 server, the ONTAP S3 administrator must provide configuration information to the S3 user.

Before you begin

The S3 client app must be capable of authenticating with the ONTAP S3 server using the following AWS signature versions:

- Signature Version 4, ONTAP 9.8 and later
- Signature Version 2, ONTAP 9.11.1 and later

Other signature versions are not supported by ONTAP S3.

The ONTAP S3 administrator must have created S3 users and granted them access permissions, as an individual users or as a group member, in the bucket policy or the object storage server policy.

The S3 client app must be capable of resolving the ONTAP S3 server name, which requires that ONTAP S3 administrator provide the S3 server name (FQDN) and IP addresses for the S3 server's LIFs.

About this task

To access an ONTAP S3 bucket, a user on the S3 client app enters information provided by the ONTAP S3 administrator.

Beginning with ONTAP 9.9.1, the ONTAP S3 server supports the following AWS client functionality:

- user-defined object metadata

A set of key-value pairs can be assigned to objects as metadata when they are created using PUT (or POST). When a GET/HEAD operation is performed on the object, the user-defined metadata is returned along with the system metadata.

- object tagging

A separate set of key-value pairs can be assigned as tags for categorizing objects. Unlike metadata, tags are created and read with REST APIs independently of the object, and they implemented when objects are created or any time after.



To enable clients to get and put tagging information, the actions `GetObjectTagging`, `PutObjectTagging`, and `DeleteObjectTagging` need to be allowed using the bucket or group policies.

For more information, see the AWS S3 documentation.

Steps

1. Authenticate the S3 client app with the ONTAP S3 server by entering the S3 server name and the CA certificate.
2. Authenticate a user on the S3 client app by entering the following information:
 - S3 server name (FQDN) and bucket name
 - the user's access key and secret key

Storage service definitions

ONTAP includes predefined storage services that are mapped to corresponding minimum performance factors.

The actual set of storage services available in a cluster or SVM is determined by the type of storage that makes up an aggregate in the SVM.

The following table shows how the minimum performance factors are mapped to the predefined storage services:

Storage service	Expected IOPS (SLA)	Peak IOPS (SLO)	Minimum volume IOPS	Estimated latency	Are expected IOPS enforced?
value	128 per TB	512 per TB	75	17 ms	On AFF: Yes Otherwise: No
performance	2048 per TB	4096 per TB	500	2 ms	Yes
extreme	6144 per TB	12288 per TB	1000	1 ms	Yes

The following table defines the available storage service level for each type of media or node:

Media or node	Available storage service level
Disk	value
Virtual machine disk	value
FlexArray LUN	value
Hybrid	value
Capacity-optimized Flash	value
Solid-state drive (SSD) - non-AFF	value
Performance-optimized Flash - SSD (AFF)	extreme, performance, value

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

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