Lenovo

Creating local cloud infrastructure using DM Series hardware

ONTAP® 9

First edition (January 2022)

Contents

Chapter 1. Deciding whether to use this guide

This guide describes how to use Lenovo Storage Manager to configure S3 client access to objects contained in a bucket in an SVM. It includes examples and advanced configuration options.

You should use this guide if you want to configure S3 object storage in the following way:

- You want to provide S3 object storage from an additional DM Series system using ONTAP S3.

  ONTAP deployment is appropriate if you want S3 capabilities on existing clusters without additional hardware and management.
- You want to use ThinkSystem Storage Manager for DM Series, not the ONTAP command line interface or an automated scripting tool.

  Note: If you want the ability to specify which aggregates are used for buckets, you can only do so using the CLI.
- You want to use best practices, not explore every available option.

  Details about command syntax are available from CLI help and ONTAP man pages.
- You do not want to read a lot of conceptual background.

  Additional information about ONTAP technology and interaction with external services is available in Lenovo ThinkSystem Storage Information Center.
- You have cluster administrator privileges, not SVM administrator privileges.

If this guide is not suitable for your situation, you should see the following documentation instead:

- ONTAP command line interfaces

Chapter 2. About S3 support in ONTAP 9

In ONTAP, the underlying architecture for a bucket is a FlexGroup volume—a single namespace that is made up of multiple constituent member volumes but is managed as a single volume.

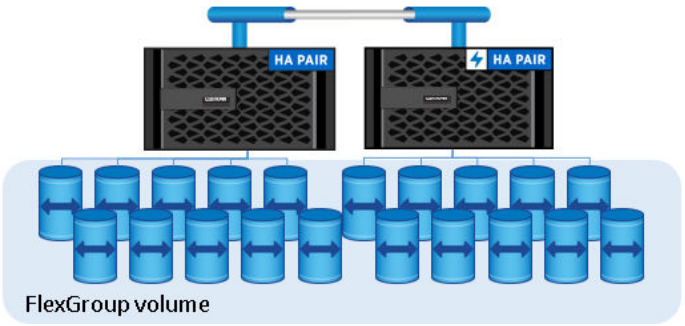

Buckets are only limited by the physical maximums of the underlying hardware, architectural maximums could be higher. Buckets can take advantage of FlexGroup elastic sizing to automatically grow a constituent of a FlexGroup volume if it is running out of space. There is a limit of 1000 buckets per FlexGroup volume, or 1/3 of the FlexGroup volume's capacity (to account for data growth in buckets).

Note: No NAS or SAN protocol access is permitted to the FlexGroup volume that contain S3 buckets, hosted on the secondary DM Series system.

Access to the bucket is provided through authorized users and client applications.

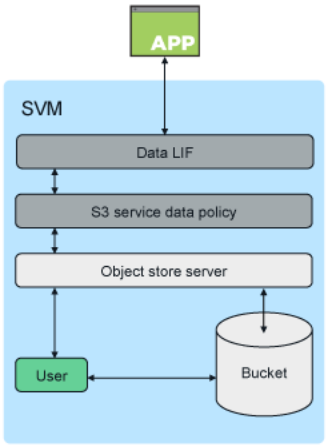

There are three primary use cases for client access to ONTAP S3 services:

- For ONTAP systems using ONTAP S3 as a remote FabricPool capacity (cloud) tier

  The S3 server and bucket containing the capacity tier (for *cold* data) is on a different cluster than the performance tier (for *hot* data).
- For ONTAP systems using ONTAP S3 as a local FabricPool tier

  The S3 server and bucket containing the capacity tier is on the same cluster, but on a different HA pair, as the performance tier.
- For external S3 client apps

  ONTAP S3 serves S3 client apps run on non-Lenovo systems.

We will be focusing on ONTAP S3 being used as a remote FabricPool (cloud) tier for our primary system.

It is a best practice to provide access to ONTAP S3 buckets using HTTPS. When HTTPS is enabled, security certificates are required for proper integration with SSL/TLS. Client users' access and secret keys are then required to authenticate the user with ONTAP S3 as well as authorizing the users' access permissions for operations within ONTAP S3. The client application should also have access to the root CA certificate (the ONTAP S3 server's signed certificate) to be able to authenticate the server and create a secure connection between client and server.

Users are created within the S3-enabled SVM on the secondary system, and their access permissions can be controlled at the bucket or SVM level; that is, they can be granted access to one or more buckets within the SVM.

HTTPS is enabled by default on ONTAP S3 servers. It is possible to disable HTTPS and enable HTTP for client access, in which case authentication using CA certificates is not required. However, when HTTP is enabled and HTTPS is disabled, all communication with the ONTAP S3 server are sent over the network in clear text.

This method will not be covered in this document.

ONTAP version support for S3 object storage

We will be focusing on ONTAP 9.10.1 in this guide since it has extended options for managing the buckets for the secondary system.

ONTAP S3 limitations

The SVM hosting the S3 Object store does not support some standard ONTAP features including.

Unsupported ONTAP functionality:

- Cloud Volumes ONTAP
- FlexCache volumes
- MetroCluster
- NDMP
- SnapMirror Cloud
- SMTape
- Volume clone of the Flexgroup containing ONTAP S3 buckets

  Note: ONTAP 9.10.1 allows SnapMirror to be used to provide disaster recovery of your S3 bucket but we will not be covering that in this guide.

Unsupported object storage functionality:

- Object versioning
- Erasure coding

Chapter 3. 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.

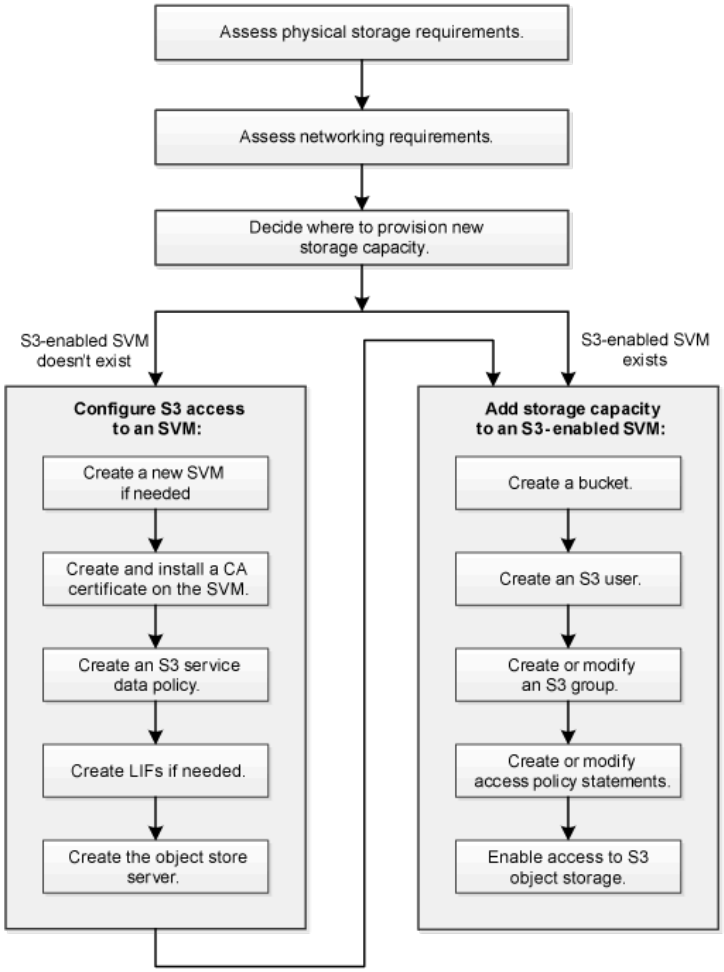

Assessing physical storage requirements

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.

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.

FlexGroup volumes management

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 two-node clusters. We will not be covering this mode for enabling ONTAP S3 in this guide.

Procedure

Step 1. Open DM Series System Manager and go to **Storage → Tiers**.

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

Example

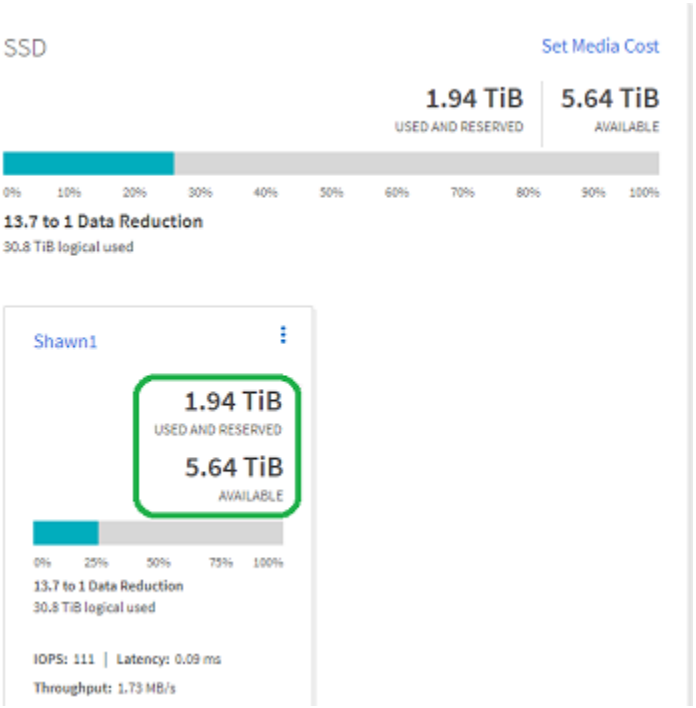

Step 2. If there are no aggregates with sufficient space, create a new aggregate by adding drives to the system and selecting **Storage → Tiers → +Add Local tier**.

Assessing 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

We will be using a local FabricPool (cloud) tier. To do this, you must also configure intercluster LIFs; cluster peering is not required.

For local FabricPool capacity tiers, you must use the system SVM (called “Cluster”), we will be using 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.

Procedure

Step 1. Display the available physical and virtual ports using DM Series System Manager: **Network → Ethernet Ports**

- 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.

Step 2. Determine the available Broadcast Domains for each of the network ports: Go to the Storage Manager for DM Series and select **Network → Overview → Broadcast Domains**.

Each network port will be listed based on the Broadcast Domain in which the port is located.

Step 3. Display available IPspaces: Go to the Storage Manager for DM Series and select **Network → Overview → IPspaces**.

You can use the default IPspace or a custom IPspace.

Deciding 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.

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.

Chapter 4 “Configuring S3 access to an SVM” on page 9

Chapter 5 “Adding storage capacity to an S3-enabled SVM” on page 15

Although S3 can coexist in an SVM with NFS and SMB/CIFS, 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.

Chapter 5 “Adding storage capacity to an S3-enabled SVM” on page 15

Chapter 4. Configuring S3 access to an SVM

Configuring S3 involves creating an SVM or enabling S3 on an existing SVM, then creating an S3 service policy, dedicated data LIFs, and an S3 server.

Creating an SVM and server for S3

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

About this task

Our example will be using a dedicated S3 SVM.

You should first ensure that you have configured DNS for both your local and remote DM Series systems. Go to the Storage Manager for DM Series and select **Cluster → Overview**.

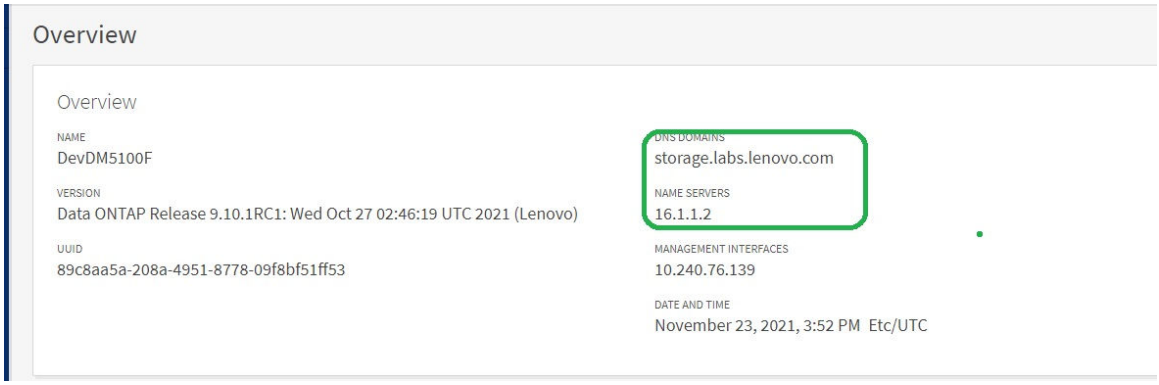

To add a DNS domain, click **More** under **DNS DOMAINS** and select **Edit → +Add**.

To add name server(s), click **More** under **NAME SERVERS** and select **Edit → +Add**.

Procedure

Step 1. Verify that S3 is licensed on your cluster: Go to the Storage Manager for DM Series and select **Cluster → Settings → License**

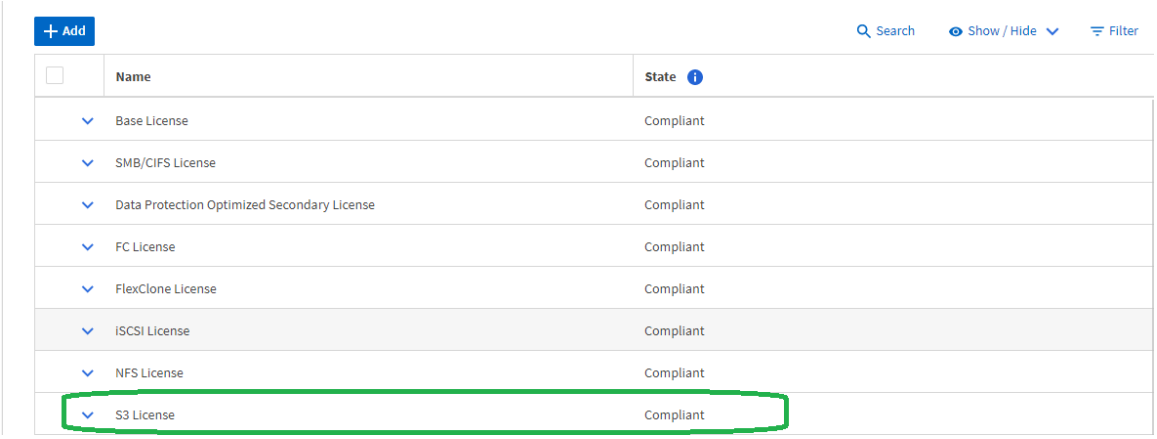

Step 2. Create your S3 enabled SVM.

a. Go to the Storage Manager for DM Series and select **Storage → Storage VMs**.

b. Select **+Add** to create a new SVM.
c. Choose a name for the new SVM.
d. Select **Enable S3** for the data protocol and provide an S3 server name.

Note: This is the name you will use later to connect to the bucket.

e. Verify that **Enable TLS** is selected and the port use is listed.
f. Choose if you want to use an external certificate or a system/ONTAP generated certificate.
g. Provide the IP address to use for the new data LIFs for this S3 SVM.
h. Click **Save**.

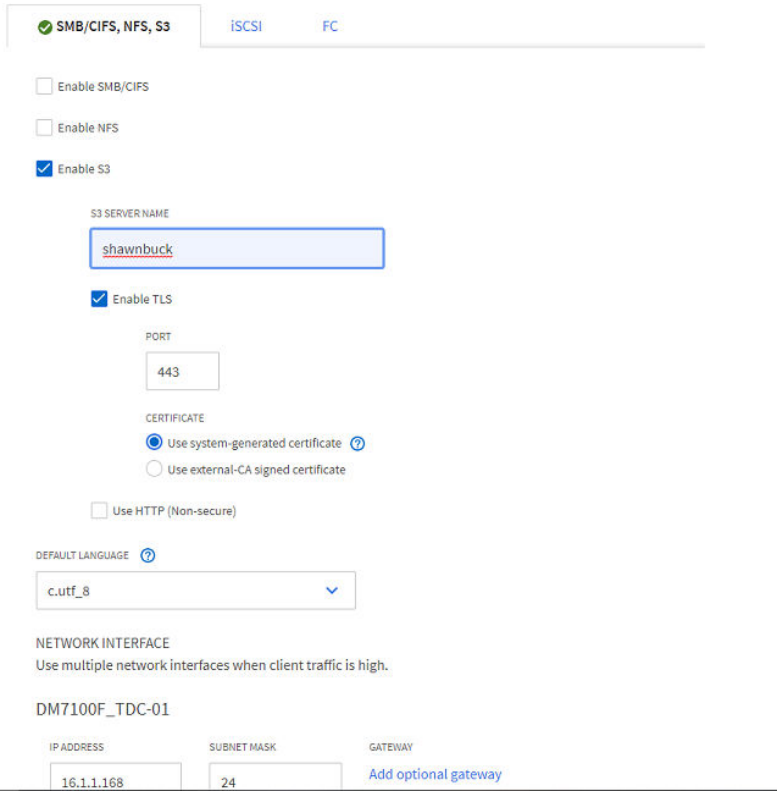

Step 3. Verify the configuration and status of the newly created SVM:

a. Go to the Storage Manager for DM Series and select **Storage → Storage VMs**.
b. Select the SVM that you have created by clicking the name.
c. Verify that the protocol is enabled and the interfaces are created.

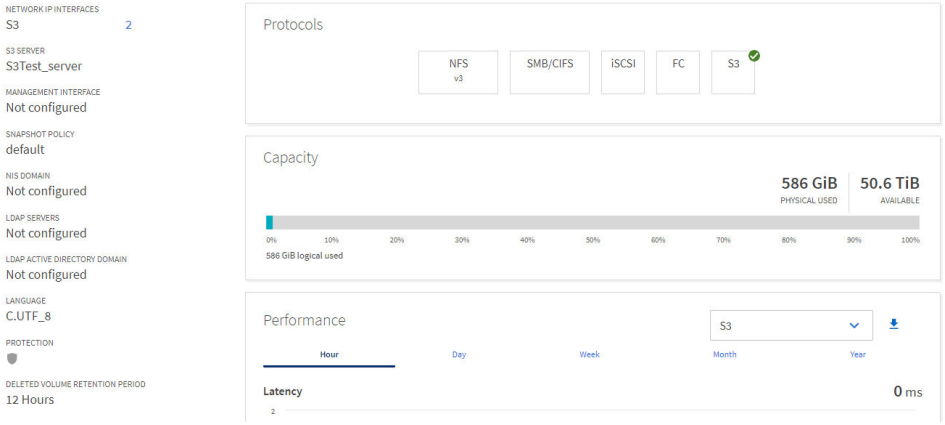

Verifying Certificate Authority

One of the most important features added in ONTAP 9.10.1 is that all features required to create a new S3 SVM hosting a bucket can be performed using System Manager for DM Series. This includes the creation of the Certificate Authority.

Procedure

Step 1. Verify that the certificate that you selected is created:

a. Go to the Storage Manager for DM Series and select **Storage → Storage VMs**.

b. Click the S3 storage VM that was created.

c. Select **Settings**.

d. Select **Security → Certificates**.

e. Select the Client/Server Certificates tab.

The certificates assigned to your S3 SMV will be displayed.

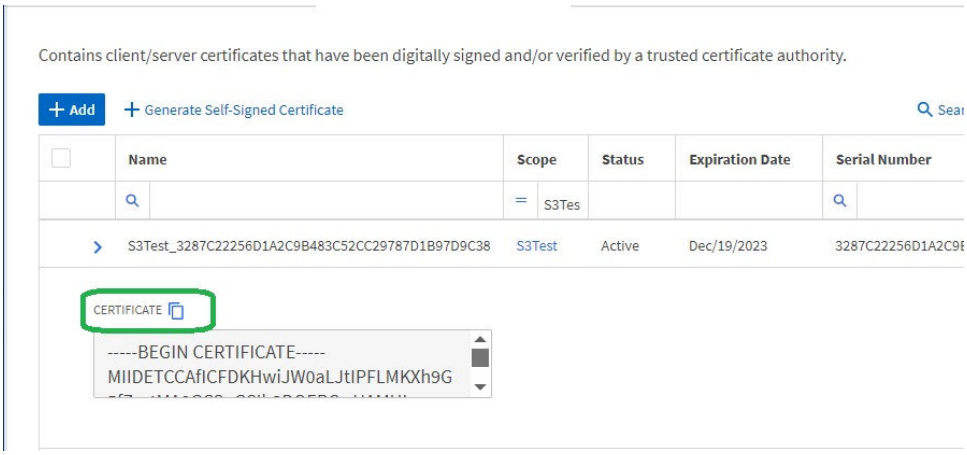

Creating an S3 service data policy

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.

Procedure

Step 1. Change the privilege setting to advanced:
`set -privilege advanced`

Step 2. Create a service data policy:
`network interface service-policy create -vserver` *svm_name* `-policy` *policy_name* `-services` *data-s3-server*

The data-s3-server service policy is the only one required to enable S3 services, although other service policies can be included as needed.

Verifying 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).
- We will create a remote (cloud) tier S3 object store, so we will give data LIFs configured on the secondary system. On the primary system, we will have intercluster LIFs.

Procedure

Step 1. To verify that the data LIFs that were created previously:
   a. Select **Network → Overview**.
   b. Locate the LIFs assigned to that storage VM.
   c. Select **Edit**.
   d. Verfiy that the data service policy is configured correctly.

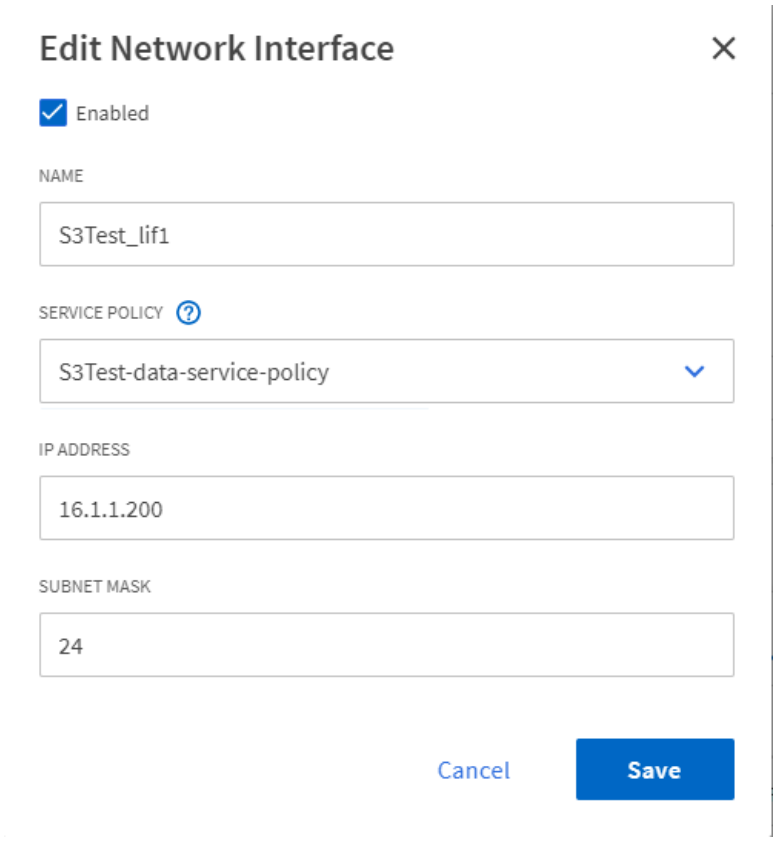

Note: It could also display as S3, NAS if you used the default service policy assigned during the SVM creation and did not create a unique service policy using the instructions above.

Creating intercluster LIFs for remote FabricPool tiering on the primary system

In this section, we will now create the intercluster LIFs on the primary system that will be used to connect to our S3 SVM on the secondary system.

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.

Procedure

Step 1. Go to the Storage Manager for DM Series and do as follows:

a. Select **Network → Overview**.
b. Click **+Add** under the Network Interfaces tab.
c. Select **Intercluster** for the role.
d. Assign a name and a home node for the new LIF.
e. Assign an IP address and a netmas for the new LIF.
f. Click **Save**.

Step 2. Verify that the intercluster LIFs were created:

a. Select **Network → Overview**.
b. Locate the ports that were created and select one..
c. Click **Edit**.

d. Verify that the port is enabled and the role is intercluster.

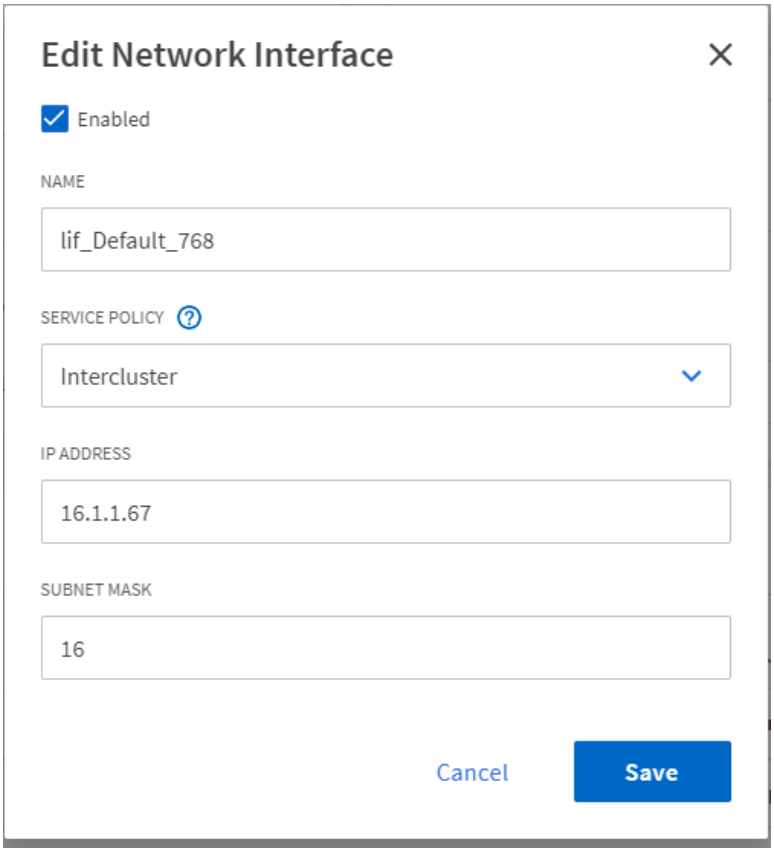

Chapter 5. Adding storage capacity to an S3-enabled SVM

To add storage capacity to an S3-enabled SVM, you must create a bucket to provide a storage container.

Creating a bucket

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

Before you begin

An SVM containing an S3 server must already exist.

About this task

When you create a bucket from the Storage Manager for DM Series, do not select provisioning options:

- ONTAP will select the underlying aggregates and FlexGroup components
  - 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 SVM will have the same underlying FlexGroup volume.
  - During this process, 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.

Procedure

Step 1. Go to the Storage Manager for DM Series.

Step 2. Select **Storage → Buckets**.

Step 3. Click **+Add**.

Step 4. Provide a name for the bucket and select the S3 SVM that you have created previously.

Step 5. Select a capacity and click more options.

Step 6. Click **Use for tiering** to ensure that it uses the low-cost HDD instead of the SSD.

Step 7. Choose if you want to copy permissions from a previous bucket or click on +Add to add new permissions.

Step 8. Click **Save**.

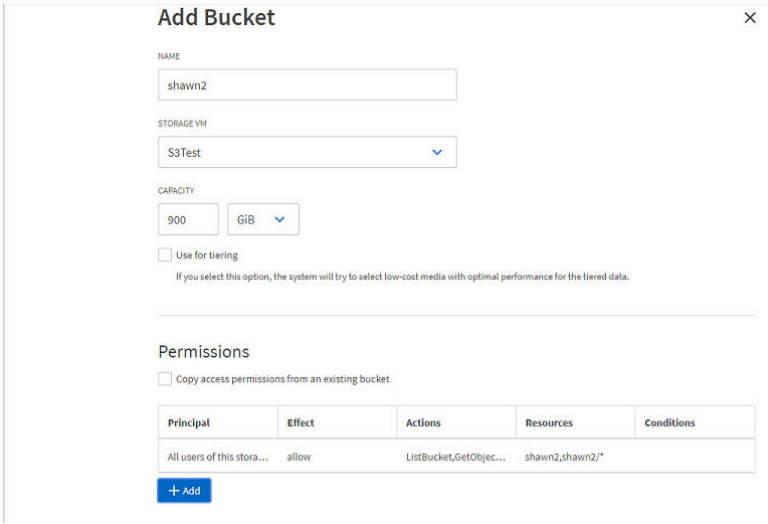

Note: You can also use this option to enable SnapMirror between buckets but we will not cover this topic in this discussion.

Creating an S3 user

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

Before you begin

An S3-enabled SVM must already exist.

About this task

An S3 user can be granted access to any bucket in an SVM but not in multiple SVMs.

When you create an S3 user, an access-key and a secret-key will be generated. They must be shared with the user along with the object store's FQDN and bucket name. S3 users' keys can be displayed 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.

Note: When an object store server is created, a root user (UID 0) is created, a privileged user with access all buckets. Rather than administering ONTAP S3 as root user, it is a best practice to create an admin user role with specific privileges.

Procedure

Step 1. Create an S3 user:
`vserver object-store-server user create -vserver` *svm_name* `-user` *user_name* [`-comment` *text*]

Creating or modifying 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.

Procedure

Step 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.

Creating or modifying access policy statements

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.

Modifying a bucket policy

You can create a bucket policy at the same time that you create the bucket.

About this task

When you click **+Add** in step 3 in "Creating a bucket" on page 15, you will be presented with the following options and you need to configure them.

Procedure

Step 1. Select the type of user. Either all users of this storage VM, all public and anonymous users or a user defined account.

Step 2. Set **EFFECT** to **Allow** or **Deny**.

Step 3. Set the individual permissions.

Step 4. Click **Save**.

Creating or modifying 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 a object storage server group. The policies do not take effect until they are specified in the group definition.

Note: 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` man pages.

Procedure

Step 1. Create an object storage server policy:
`vserver object-store-server policy create -vserver` *svm_name* `-policy` *policy_name* `[-comment` *text*`]`

Step 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
-action	You can specify * to mean all actions, or a list of one or more of the following: `GetObject`, `PutObject`, `DeleteObject`, `ListBucket`,`GetBucketAcl`, `GetObjectAcl`, `ListAllMyBuckets`, `ListBucketMultipartUploads`, and `ListMultipartUploadParts`.
-resource	The bucket and any object it contains. The wildcard characters * and ? 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.

Verifying client access to S3 object storage

After configuring the S3 server and buckets, S3 administrators can make object stores available to client systems, either ONTAP systems with remote FabricPool capacity (cloud) tiers or S3 client apps, or to a local FabricPool capacity tier.

Enabling 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 both local and remote clusters, although cluster peering is not required.

Attaching ONTAP S3 Object Store as Cloud Tier

Once you have all of the information above you are ready to attach the remote DM Series device as the cloud tier for your primary DM Series system. The following steps will walk you throug that process.

Before you begin

You need a local tier defined on your primary system before you start.

Procedure

Step 1. Select to add the cloud tier using the ONTAP S3 bucket from your secondary system.

Step 2. Select **Storage → Tiers**.

Step 3. Select **+Add Cloud Tier → ONTAP S3**.

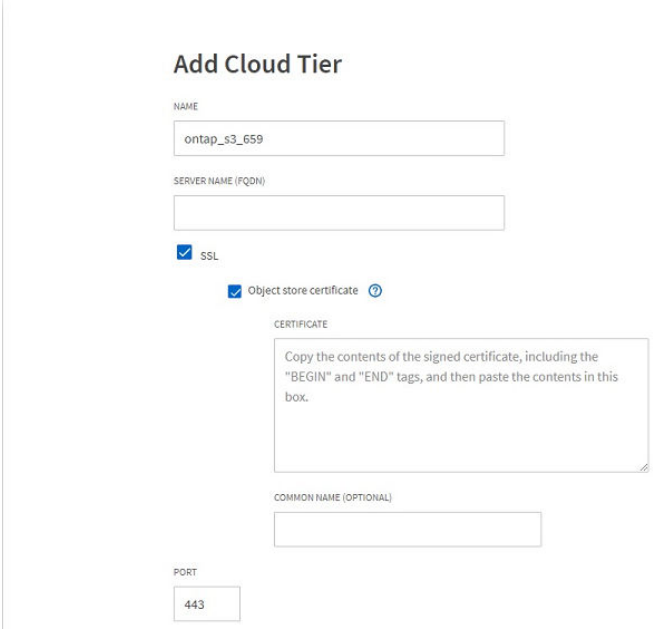

Step 4. Specify the FQDN that you wrote down earlier. Check that the SSL box is enabled and provide the certificate for the S3 SVM. We will go through each one in following steps.

a. To retrieve the FQDN, select **Storage → Storage VMs** on the secondary system and click the SVM name. Then go to **Settings** and click the pencil icon to edit it.

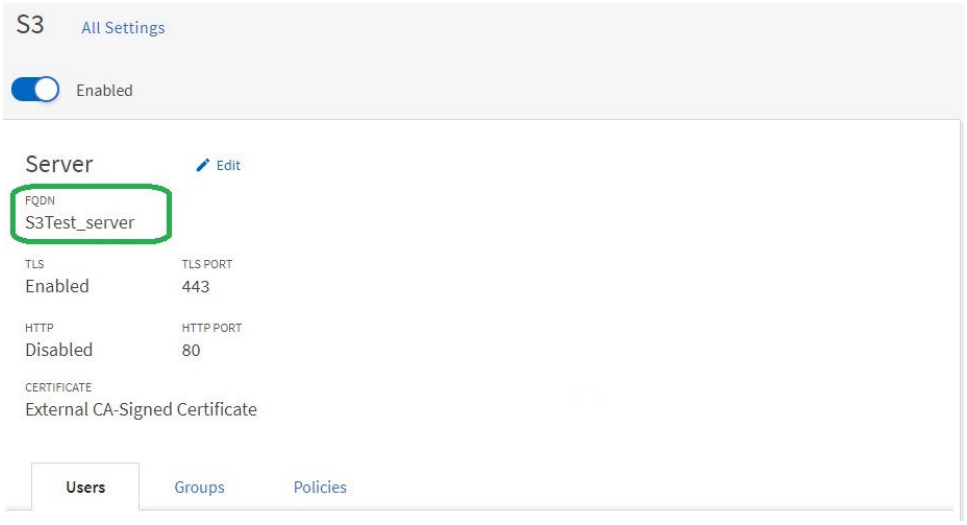

b. Specify the certificate for that SVM. To retrieve the certificate, select **Storage → Storage VMs**. Select the storage VM by clicking it. Select **Settings → Certificates** and click the arrow.

Locate the certificate and click the arrow icon next to the certificate. An option to copy the certificate to the clipboard is displayed.

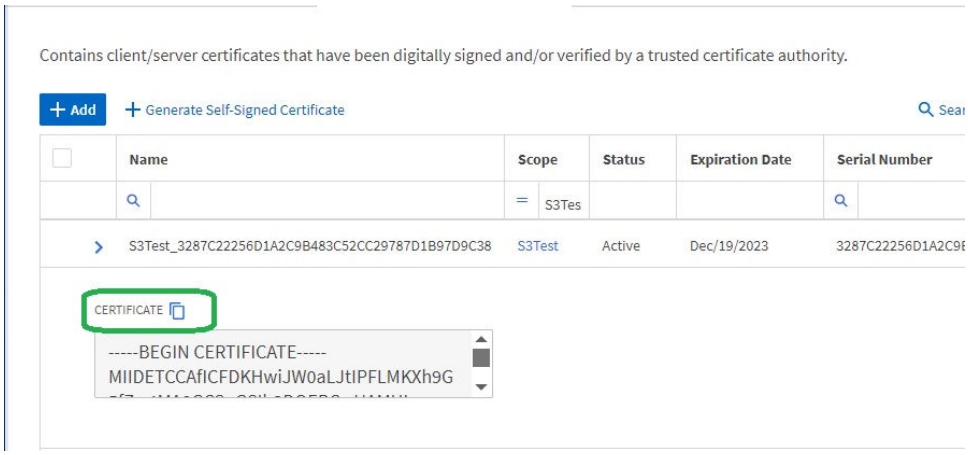

c. Provide the access and secret key associated with the S3 user that you assigned to the bucket. The secret key will not be displayed and if it is lost, you must recreate it, but you can copy the access key off if needed.

d. Click **Save** to complete attaching the cloud tier.

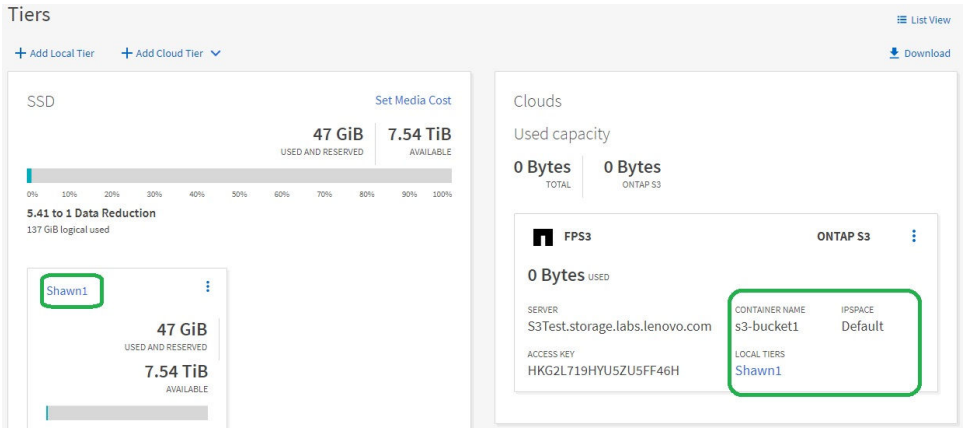

Defining a tiering policy for volumes

Once you have attached the secondary system as your cloud tier, you will need to assign your tiering policy on a per volume basis to determine data locality.

About this task

This tiering policy determines when data will be moved from your local tier to the cloud tier. There are four options for tiering. You should select the policy that matches your data access pattern.

- **Auto**: It is the default setting for the volumes. It is done based on FW default values.
- **Snapshot Only**: The snapshot data is moved, but the root volume data is never sent to the cloud tier.
- **None**: Any volume data is prevented from being moved to the cloud.
- **All**: All data including snapshot data is moved to the cloud

There is one additional consideration that will need to be made and that is the minimum cooling days needed before data will be moved off based on the policy you assign. By default, the minimum cooling days are set to 31. You can change it to be as low as 2 but keep in mind that this means the data will be served from the slower cloud tier until it is moved back to the local tier by being accessed, so you want to a large number of data move operations. To set the minimum cooling days, you will need to use the CLI, for example, `DevDM5100F::*> volume modify -vserver S3TestFP2 -volume Dat3 -tiering-minimum-cooling-days 2 Volume modify successful on volume Dat3 of Vserver S3TestFP2.`

You can change the tiering policy using the Storage Manager for DM Series: Select **Storage → Volumes**. Select the Volume you want to modify and then select **...** and click **Edit Cloud Tier Settings**.

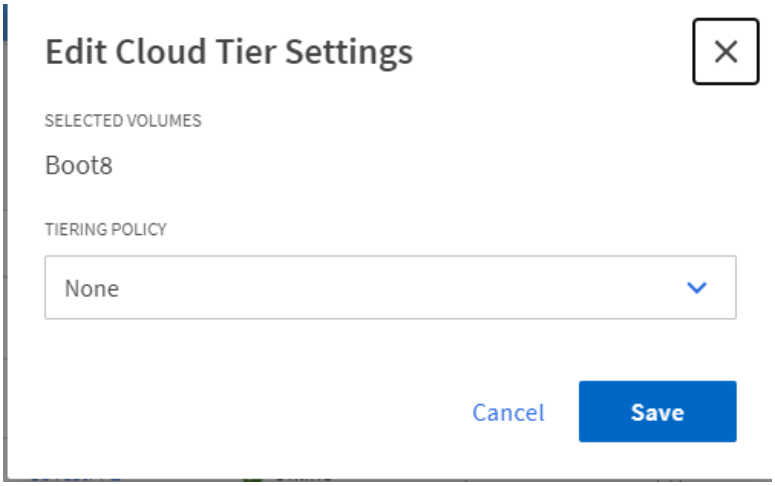

Verifying the cloud tier

Once everything is configured and the cooling period is met, you will see that data is being moved to the cloud tier.

About this task

The cloud tier can be verified by going to **Storage → Tiers**. When you look at the Cloud Tier, it will display the amount of data that has been moved off based on the settings that you assigned.

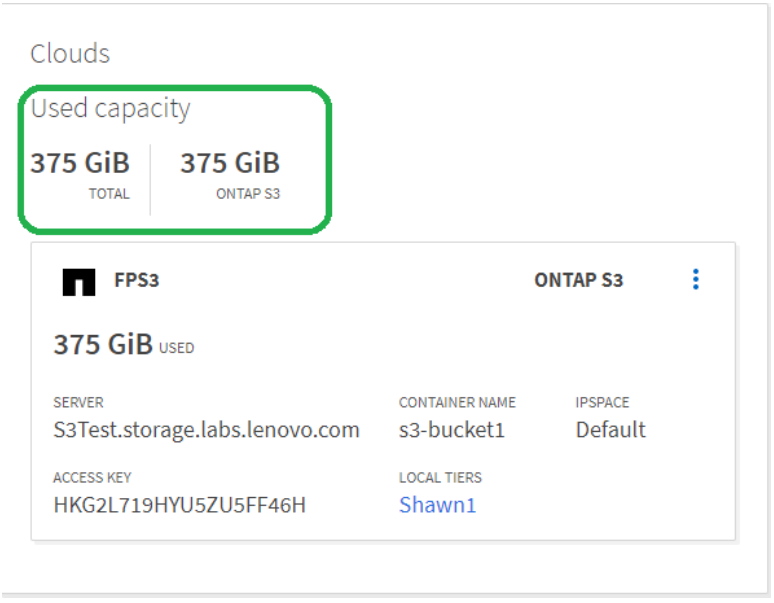

Enabling ONTAP S3 access for local FabricPool tiering

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.

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.

Procedure

Step 1. Create the object store for the local capacity tier:
`storage aggregate object-store config create -object-store-name` *store_name* `-ispace 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.

Example

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

Step 2. Display and verify the object store configuration information:
`storage aggregate object-store config show`

Step 3. Optional: To see how much data in a volume is inactive, follow the steps in Managing Storage Tiers By Using FabricPool.

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

Step 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.

Example

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

Step 5. Display the object store information and verify that the attached object store is available:
`storage aggregate object-store show`

Example

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

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

Enabling 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 AWS Signature Version 4. Earlier 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.

  Note: 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.
- 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. Tags can also be used in condition policy statements for bucket and user policies.

For more information, see the AWS S3 documentation.

Procedure

Step 1. Authenticate the S3 client app with the ONTAP S3 server by entering the S3 server name and the CA certificate.

Step 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

Chapter 6. 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 AFA: 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
Hybrid	value
Capacity-optimized Flash	value
Solid-state drive (SSD) - non-AFA	value
Performance-optimized Flash - SSD (AFA)	extreme, performance, value

Appendix A. Contacting Support

You can contact Support to obtain help for your issue.

You can receive hardware service through a Lenovo Authorized Service Provider. To locate a service provider authorized by Lenovo to provide warranty service, go to https://datacentersupport.lenovo.com/serviceprovider and use filter searching for different countries. For Lenovo support telephone numbers, see https://datacentersupport.lenovo.com/supportphonelist for your region support details.

Appendix B. Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area.

Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document is not an offer and does not provide a license under any patents or patent applications. You can send inquiries in writing to the following:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Trademarks