NetApp

Configure ONTAP tools

ONTAP tools for VMware vSphere 9.11

NetApp
December 18, 2025

Table of Contents

Configure ONTAP tools

Workflow for configuring ONTAP tools

Configuring ONTAP tools involves configuring your storage systems, creating a storage capability profile, provisioning the datastore, and optionally configuring SRA for disaster recovery.
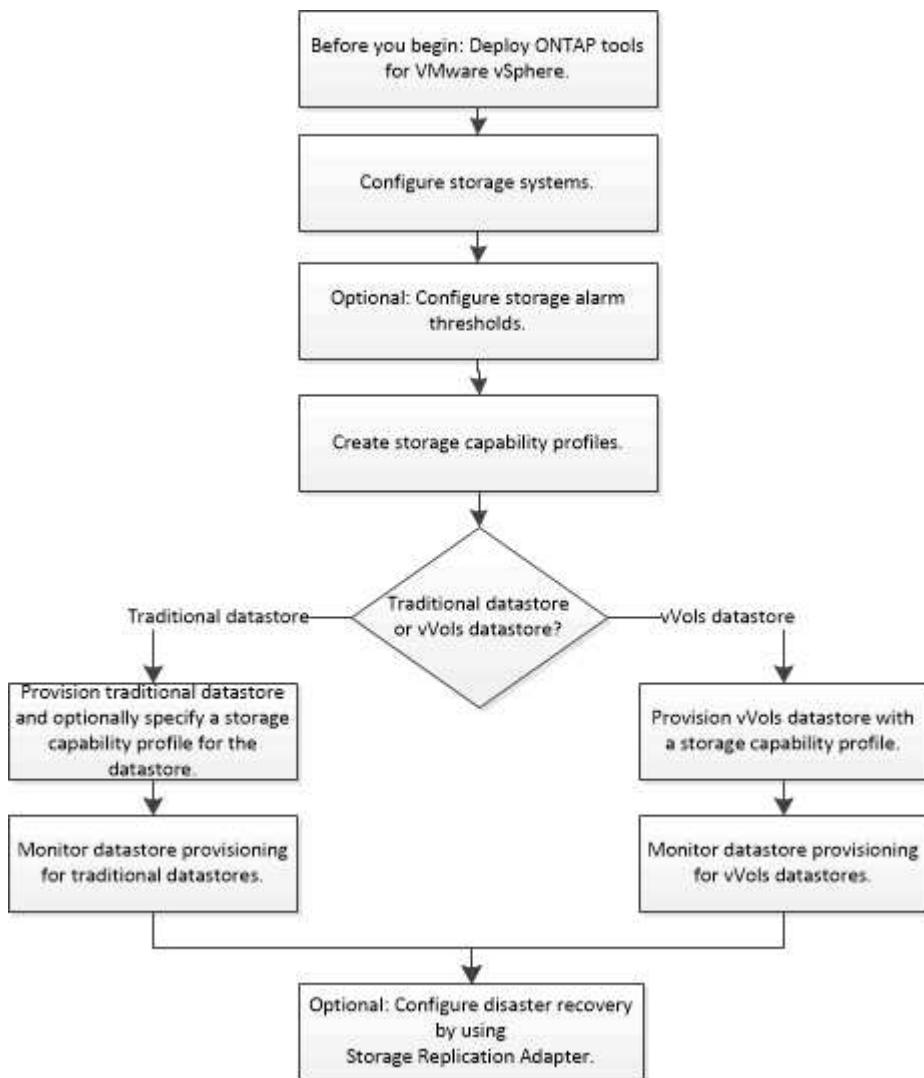
Configure ESXi settings

Configure ESXi server multipathing and timeout settings

ONTAP tools for VMware vSphere checks and sets the ESXi host multipath settings and HBA timeout settings that work best with NetApp storage systems.

About this task

This process might take a long time, depending on your configuration and system load. The task progress is displayed in the Recent Tasks panel. As the tasks are completed, the host status Alert icon is replaced by the

Normal icon or the Pending Reboot icon.

Steps

1. From the VMware vSphere Web Client Home page, click **vCenter** > **Hosts**.
2. Right-click a host, and then select **Actions** > **NetApp VSC** > **Set Recommended Values**.
3. In the NetApp Recommended Settings dialog box, select the values that work best with your system.

   The standard, recommended values are set by default.
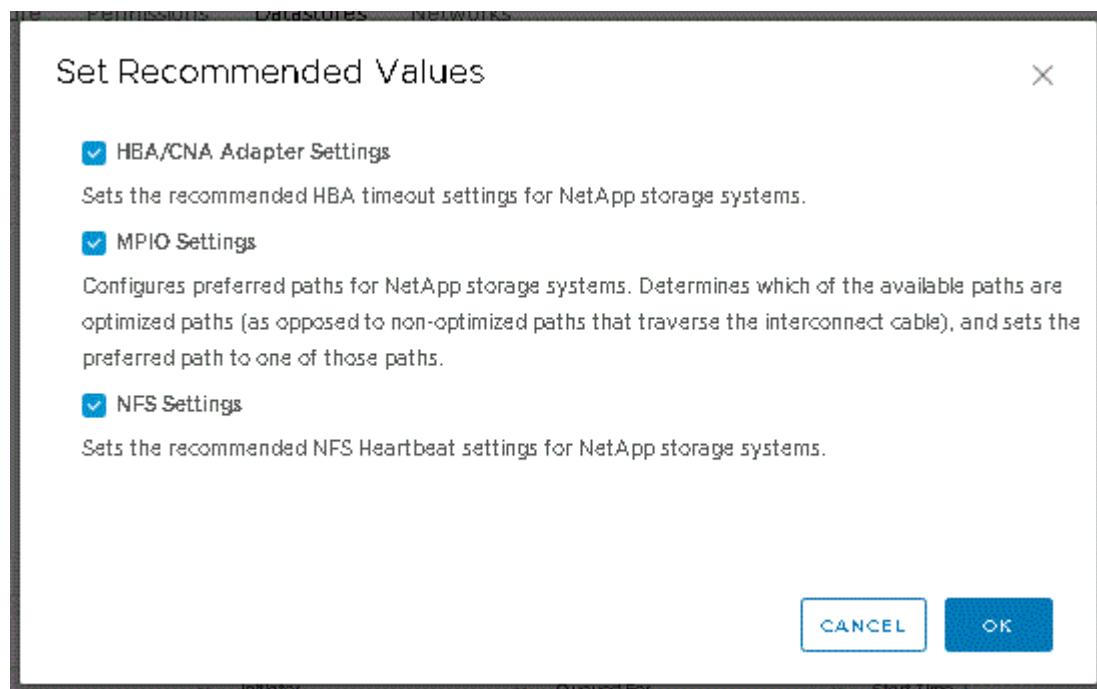
4. Click **OK**.

ESXi host values set using ONTAP tools

You can set timeouts and other values on the ESXi hosts using ONTAP tools for VMware vSphere to ensure best performance and successful failover. The values that ONTAP tools sets are based on internal NetApp testing.

You can set the following values on an ESXi host:

ESXi advanced configuration

- **VMFS3.HardwareAcceleratedLocking**

  You should set this value to 1.

- **VMFS3.EnableBlockDelete**

  You should set this value to 0.

NFS settings

- **Net.TcpipHeapSize**

  If you are using vSphere 6.0 or later, you should set this value to 32.

- **Net.TcpipHeapMax**

  You should set this value to 512MB for vSphere 6.X releases, except for 6.5U3 and 6.7U3 releases. For vSphere 6.5U3 and 6.7U3 releases set this value to 1024MB. If you are using vSphere 7.0 or later versions, you should set this value to 1024MB.

- **NFS.MaxVolumes**

  If you are using vSphere 6.0 or later, you should set this value to 256.

- **NFS41.MaxVolumes**

  If you are using vSphere 6.0 or later, you should set this value to 256.

- **NFS.MaxQueueDepth**

  If you are using the vSphere 6.0 or later version of ESXi host, then you should set this value to 128 or higher to avoid queuing bottlenecks.

  For vSphere versions prior to 6.0, you should set this value to 64.

- **NFS.HeartbeatMaxFailures**

  You should set this value to 10 for all NFS configurations.

- **NFS.HeartbeatFrequency**

  You should set this value to 12 for all NFS configurations.

- **NFS.HeartbeatTimeout**

  You should set this value to 5 for all NFS configurations.

FC/FCoE settings

- **Path selection policy**

  You should set this value to "RR" (round robin) when FC paths with ALUA are used.

  You should set this value to "FIXED" for all other configurations.

  Setting this value to "RR" helps to provide load balancing across all of the active/optimized paths. The value "FIXED" is used for older, non-ALUA configurations and helps to prevent proxy I/O.

- **Disk.QFullSampleSize**

  You should set this value to 32 for all configurations. Setting this value helps to prevent I/O errors.

- **Disk.QFullThreshold**

You should set this value to 8 for all configurations. Setting this value helps prevent I/O errors.

- **Emulex FC HBA timeouts**

  Use the default value.

- **QLogic FC HBA timeouts**

  Use the default value.

iSCSI settings

- **Path selection policy**

  You should set this value to "RR" for all iSCSI paths.

  Setting this value to "RR" helps to provide load balancing across all of the active/optimized paths.

- **Disk.QFullSampleSize**

  You should set this value to 32 for all configurations. Setting this value helps to prevent I/O errors.

- **Disk.QFullThreshold**

  You should set this value to 8 for all configurations. Setting this value helps prevent I/O errors.

Configure guest operating systems

Configure guest operating system scripts

The ISO images of the guest operating system (OS) scripts are mounted on ONTAP® tools for VMware vSphere server. To use the guest OS scripts to set the storage timeouts for virtual machines, you must mount the scripts from the vSphere Client.

Operating System Type	60-second timeout settings	190-second timeout settings
Linux	https://<appliance_ip>:8143/vsc/public/writable/linux_gos_timeout-install.iso	https://<appliance_ip>:8143/vsc/public/writable/linux_gos_timeout_190-install.iso
Windows	https://<appliance_ip>:8143/vsc/public/writable/windows_gos_timeout.iso	https://<appliance_ip>:8143/vsc/public/writable/windows_gos_timeout_190.iso
Solaris	https://<appliance_ip>:8143/vsc/public/writable/solaris_gos_timeout-install.iso	https://<appliance_ip>:8143/vsc/public/writable/solaris_gos_timeout_190-install.iso

You should install the script from the copy of the VSC instance that is registered to the vCenter Server that manages the virtual machine. If your environment includes multiple vCenter Servers, you should select the server that contains the virtual machine for which you want to set the storage timeout values.

You should log in to the virtual machine, and then run the script to set the storage timeout values.

Set timeout values for Windows guest operating systems

The guest operating system (OS) timeout scripts set the SCSI I/O timeout settings for Windows guest operating systems. You can specify either a 60-second timeout or a 190-second timeout. You must reboot the Windows guest OS for the settings to take effect.

What you will need

You must have mounted the ISO image containing the Windows script.

Steps

1. Access the console of the Windows virtual machine, and log in to an account with Administrator privileges.

2. If the script does not automatically start, open the CD drive, and then run the `windows_gos_timeout.reg` script.

   The Registry Editor dialog is displayed.

3. Click **Yes** to continue.

   The following message is displayed:

   ```
   The keys and values contained in 'D:\windows_gos_timeout.reg' have been
   successfully added to the registry.`
   ```

4. Reboot the Windows guest OS.

5. Unmount the ISO image.

Set timeout values for Solaris guest operating systems

The guest operating system (OS) timeout scripts set the SCSI I/O timeout settings for Solaris 10. You can specify either a 60-second timeout or a 190-second timeout.

What you will need

You must have mounted the ISO image containing the Solaris script.

Steps

1. Access the console of the Solaris virtual machine, and log in to an account with root privileges.

2. Run the `solaris_gos_timeout-install.sh` script.

   For Solaris 10, a message similar to the following is displayed:

   ```
   Setting I/O Timeout for /dev/s-a - SUCCESS!
   ```

3. Unmount the ISO image.

Set timeout values for Linux guest operating systems

The guest operating system (OS) timeout scripts set the SCSI I/O timeout settings for versions 4, 5, 6, and 7 of Red Hat Enterprise Linux and versions 9, 10, and 11 of SUSE Linux Enterprise Server. You can specify either a 60-second timeout or a 190-second timeout. You must run the script each time you upgrade to a new version of Linux.

What you will need

You must have mounted the ISO image containing the Linux script.

Steps

1. Access the console of the Linux virtual machine, and log in to an account with root privileges.

2. Run the `linux_gos_timeout-install.sh` script.

   For Red Hat Enterprise Linux 4 or SUSE Linux Enterprise Server 9, a message similar to the following is displayed:

   ```
   Restarting udev... this may take a few seconds.
   ```

   ```
   Setting I/O Timeout (60s) for /dev/sda - SUCCESS!
   ```

   For Red Hat Enterprise Linux 5, Red Hat Enterprise Linux 6, and Red Hat Enterprise Linux 7 a message similar to the following is displayed:

   ```
   patching file /etc/udev/rules.d/50-udev.rules
   ```

   ```
   Hunk #1 succeeded at 333 (offset 13 lines).
   ```

   ```
   Restarting udev... this may take a few seconds.
   ```

   ```
   Starting udev: [ OK ]
   ```

   ```
   Setting I/O Timeout (60s) for /dev/sda - SUCCESS!
   ```

   For SUSE Linux Enterprise Server 10 or SUSE Linux Enterprise Server 11, a message similar to the following is displayed:

```
patching file /etc/udev/rules.d/50-udev-default.rules
```

```
Hunk #1 succeeded at 114 (offset 1 line).
```

```
Restarting udev ...this may take a few seconds.
```

```
Updating all available device nodes in /dev:   done
```

3. Unmount the ISO image.

Requirements for registering VSC in multiple vCenter Servers environment

If you are using ONTAP tools for VMware vSphere in an environment where a single VMware vSphere HTML5 client is managing multiple vCenter Server instances, you must register one instance of VSC with each vCenter Server so that there is a 1:1 pairing between VSC and the vCenter Server. Doing this enables you to manage all of the servers running vCenter 6.0 or later in both linked mode and non-linked mode from a single vSphere HTML5 client.

> ℹ️ If you want to use VSC with a vCenter Server, then you must have set up or registered one VSC instance for every vCenter Server instance that you want to manage. Each registered VSC instance must be of the same version.

Linked mode is installed automatically during the vCenter Server deployment. Linked mode uses Microsoft Active Directory Application Mode (ADAM) to store and synchronize data across multiple vCenter Server systems.

Using the vSphere HTML5 client to perform VSC tasks across multiple vCenter Servers requires the following:

- Each vCenter Server in the VMware inventory that you want to manage must have a single VSC server registered with it in a unique 1:1 pairing.

  For example, you can have VSC server A registered to vCenter Server A, VSC server B registered to vCenter Server B, VSC server C registered to vCenter Server C, and so on.

  You **cannot** have VSC server A registered to both vCenter Server A and vCenter Server B.

  If a VMware inventory includes a vCenter Server that does not have a VSC server registered to it, but there are one or more vCenter Servers that are registered with VSC, then you can view the instances of VSC and perform VSC operations for the vCenter Servers that have VSC registered.

- You must have the VSC-specific View privilege for each vCenter Server that is registered to the single sign-on (SSO).

You must also have the correct RBAC permissions.

When you are performing a task that requires you to specify a vCenter Server, the **vCenter Server** drop-down box displays the available vCenter Servers in alphanumeric order. The default vCenter Server is always the first server in the drop-down list.

If the location of the storage is known (for example, when you use the Provisioning wizard and the datastore is on a host managed by a specific vCenter Server), the vCenter Server list is displayed as a read-only option. This happens only when you use the right-click option to select an item in the vSphere Web Client.

VSC warns you when you attempt to select an object that it does not manage.

You can filter storage systems based on a specific vCenter Server from the VSC summary page. A summary page appears for every VSC instance that is registered with a vCenter Server. You can manage the storage systems that are associated with a specific VSC instance and vCenter Server, but you should keep the registration information for each storage system separate if you are running multiple instances of VSC.

Configure the ONTAP tools preferences file

Set IPv4 or IPv6 using the preferences file

The preferences files contain settings that control ONTAP tools for VMware vSphere operations. Under most circumstances, you do not have to modify the settings in these files. It is helpful to know which preference files Virtual Storage Console (VSC) uses.

VSC has several preference files. These files include entry keys and values that determine how VSC performs various operations. The following are some of the preference files that VSC uses:

`/opt/netapp/vscserver/etc/kamino/kaminoprefs.xml`

`/opt/netapp/vscserver/etc/vsc/vscPreferences.xml`

You might have to modify the preferences files in certain situations. For example, if you use iSCSI or NFS and the subnet is different between your ESXi hosts and your storage system, you have to modify the preferences files. If you do not modify the settings in the preferences file, datastore provisioning fails because VSC cannot mount the datastore.

There is a new option added to the preference file `kaminoprefs.xml` that you can set to enable support for IPv4 or IPv6 for all storage systems added to VSC.

- The `default.override.option.provision.mount.datastore.address.family` parameter has been added to the `kaminoprefs.xml` preference file to set a preferred data LIF protocol for datastore provisioning.

  This preference is applicable for all of the storage systems added to VSC.

- The values for the new option are `IPv4`, `IPv6`, and `NONE`.

- By default the value is set to `NONE`.

Value	Description

NONE	• Provisioning happens using the same IPv6 or IPv4 address type of data LIF as the type of cluster or SVM management LIF used for adding the storage.
	• If the same IPv6 or IPv4 address type of data LIF is not present in the SVM, then the provisioning happens through the other type of data LIF, if available.
IPv4	• Provisioning happens using the IPv4 data LIF in the selected SVM.
	• If the SVM does not have an IPv4 data LIF, then the provisioning happens through the IPv6 data LIF, if it is available in the SVM.
IPv6	• Provisioning happens using the IPv6 data LIF in the selected SVM.
	• If the SVM does not have an IPv6 data LIF, then the provisioning happens through the IPv4 data LIF, if it is available in the SVM.

To configure the IPv4 or IPv6 using the user interface, see the following sections:

- Add different subnets
- Enable datastore mounting across different subnets

Add different subnets

You can use the ONTAP tools interface or REST APIs to add different subnets of ESXi hosts. This enables you to either allow or restrict the subnets for datastore mount operation after provisioning storage systems. If you do not add subnets of ESXi hosts then ONTAP tools blocks datastore mount operation for those subnets.

Steps

1. Log in to your vCenter Server instance and access ONTAP tools.
2. On the homepage, click **Settings** > **Manage Subnet Access**.
3. In the Manage Subnet Access dialog box, click **Selected** option in Allowed subnets for NFS Subnets Access.
4. Enter the values for the required subnets, and then click **ADD**.
5. Select either **None** or **Selected** for Restricted subnets.
6. Repeat the above steps for iSCSI Subnets Access, and click **Apply**.

Enable datastore mounting across different subnets

If you use iSCSI or NFS and the subnet is different between your ESXi hosts and your

storage system, you have to modify the ONTAP tools for VMware vSphere preferences files. If you do not modify the preferences file, then datastore provisioning fails because Virtual Storage Console (VSC) cannot mount the datastore.

About this task

When datastore provisioning fails, VSC logs the following error messages:

`Unable to continue. No ip addresses found when cross-referencing kernel ip addresses and addresses on the controller.

Unable to find a matching network to NFS mount volume to these hosts.`

Steps

1. Log in to your vCenter Server instance.

2. Launch the maintenance console using your unified appliance virtual machine.

   [Maintenance Console of ONTAP tools for VMware vSphere](#)

3. Enter `4` to access the Support and Diagnostics option.

4. Enter `2` to access the Access Diagnostic Shell option.

5. Enter `vi /opt/netapp/vscserver/etc/kamino/kaminoprefs.xml` to update the `kaminoprefs.xml` file.

6. Update the `kaminoprefs.xml` file.

If you use…	Do this…
iSCSI	Change the value of the entry key `default.allow.iscsi.mount.networks` from ALL to the value of your ESXi host networks.
NFS	Change the value of the entry key `default.allow.nfs.mount.networks` from ALL to the value of your ESXi host networks.

The preferences file includes sample values for these entry keys.

> (i) The value "ALL" does not mean all networks. The "ALL" value enables all of the matching networks, between the host and the storage system, to be used for mounting datastores. When you specify host networks, then you can enable mounting only across the specified subnets.

7. Save and close the `kaminoprefs.xml` file.

Regenerate an SSL certificate for Virtual Storage Console

The SSL certificate is generated when you install Virtual Storage Console (VSC). The distinguished name (DN) that is generated for the SSL certificate might not be a common

name (CN) that the client machines recognize. By changing the keystore and private key passwords, you can regenerate the certificate and create a site-specific certificate.

About this task

You can enable remote diagnostic using the maintenance console and generate site-specific certificate.

Virtual Storage Console: Implementing CA signed certificates

Steps

1. Log in to the maintenance console.

2. Enter `1` to access the Application Configuration menu.

3. In the Application Configuration menu, enter `3` to stop the VSC service.

4. Enter `7` to regenerate SSL certificate.

Configure storage systems

Overview of storage systems for ONTAP tools

You should add storage systems to ONTAP tools and set default credentials, if required, by using the VSC interface.

ONTAP tools for VMware vSphere provides a single mechanism to discover storage systems and to set the storage credentials. The credentials provide the ONTAP permissions that are required to enable Virtual Storage Console (VSC) users to perform tasks by using the storage systems.

Before VSC can display and manage the storage resources, VSC must discover the storage systems. As part of the discovery process, you must supply the ONTAP credentials for your storage systems. These are the privileges (or roles) that are associated with the user name and password pair that is assigned to each storage system. These user name and password pairs use ONTAP role-based access control (RBAC) and must be set up from within ONTAP. You cannot change these credentials from within VSC. You can define ONTAP RBAC roles by using ONTAP System Manager.

> ⓘ   If you log in as an administrator, you automatically have all of the privileges for that storage system.

When you add a storage system to VSC, you must supply an IP address for the storage system and the user name and password pair that is associated with that system. You can set up default credentials that VSC will use during the storage system discovery process, or you can manually enter credentials when the storage system is discovered. The details of the storage system that is added to VSC are automatically pushed to the extensions that you enable in your deployment. You do not have to manually add storage to VASA Provider and Storage Replication Adapter (SRA). Both VSC and SRA support the addition of credentials at the cluster level and storage virtual machine (SVM) level. VASA Provider supports only cluster-level credentials for adding storage systems.

If your environment includes multiple vCenter Server instances, when you add a storage system to VSC from the Storage Systems page, the Add Storage System dialog box displays a vCenter Server box where you can specify to which vCenter Server instance the storage system is to be added. If you add a storage system by right-clicking a datacenter name, you do not have the option to specify a vCenter Server instance because the server is already associated with that datacenter.

Discovery happens in one of the following ways. In each case, you must supply credentials for any newly discovered storage system.

- When the VSC service starts, VSC begins its automatic background discovery process.
- You can click the REDISCOVER All button in the **Storage Systems** page, or on a host or datacenter to select it from the **Actions** menu (**Actions** > **Netapp ONTAP tools** > **Update Host and Storage Data**). You can also click **DISCOVER** on the **Getting Started** tab of the 'Overview' section.

All of the VSC features require specific permissions to perform tasks. You can limit what users can do based on the credentials that are associated with the ONTAP role. All of the users that have the same storage system user name and password pair share the same set of storage system credentials and can perform the same operations.

Set default credentials for storage systems

You can use ONTAP tools for VMware vSphere to set default credentials for a storage system in your vCenter Server.

What you will need

You must have selected the vCenter Server that you want to use for creating default credentials.

About this task

If you set up default credentials for storage systems, Virtual Storage Console (VSC) uses these credentials to log in to a storage system that VSC has just discovered. If the default credentials do not work, you must manually log in to the storage system. VSC and SRA support addition of storage system credentials at the cluster level or the SVM level. But VASA Provider will only work with cluster level credentials.

Steps

1. On the ONTAP tools Home page, click **Settings** > **Administrative Settings** > **Configure Default Credentials for Storage System**.

2. In the Storage System Default Credentials dialog box, enter the user name and password for the storage system.

   Storage controller credentials are assigned in ONTAP based on the user name and password pair. The storage controller credentials can either be the administrator account or a custom account that uses role-based access control (RBAC).

   You cannot use VSC to change the roles that are associated with the user name and password pair of the storage controller. To modify or create a new ONTAP user role for use with ONTAP tools, you can use System Manager.

   See the "Configuring user roles and privileges" section.

3. Click **OK** to save the default credentials.

If you updated the storage system credentials because a storage system reported "Authentication Failure" status, you should click the **REDISCOVER ALL** option available on the Storage Systems page. When you do this, VSC tries to connect to the storage system by using the new credentials.

Add storage systems to VSC

You can manually add storage system to Virtual Storage Console (VSC).

About this task

Each time you start Virtual Storage Console (VSC) or select the **REDISCOVER All** option, VSC automatically discovers the available storage systems.

Steps

1. Add a storage system to VSC by using the ONTAP tools home page:
   - Click **Storage Systems** > **Add**.
   - Click **Overview** > **Getting Started**, and then click **ADD** button under Add Storage System.
2. In the **Add Storage System** dialog box, enter the management IP address and credentials for that storage system.

   You can also add storage systems using the IPv6 address of the cluster or SVM. You can also change the defaults for TLS and the port number in this dialog box.

   When you add storage from the VSC Storage System page, you must also specify the vCenter Server instance where the storage will be located. The Add Storage System dialog box provides a drop-down list of the available vCenter Server instances. VSC does not display this option if you are adding storage to a datacenter that is already associated with a vCenter Server instance.

3. Click **OK** after you have added all of the required information.

Discover storage systems and hosts

When you first run Virtual Storage Console (VSC) in a vSphere Client, VSC discovers the ESXi hosts, their LUNs and NFS exports, and the NetApp storage systems that own those LUNs and exports.

What you will need

- All of the ESXi hosts must be powered on and connected.
- All the storage virtual machines (SVMs) to be discovered must be running, and each cluster node must have at least one data LIF configured for the storage protocol in use (NFS, iSCSI or FC).

About this task

You can discover new storage systems or update information about existing storage systems to obtain the latest capacity and configuration information at any time. You can also modify the credentials that VSC uses to log in to the storage systems.

While discovering the storage systems, VSC collects information from the ESXi hosts that are managed by the vCenter Server instance.

Steps

1. From the vSphere Client Home page, select **Hosts and Clusters**.
2. Right-click the required datacenter, and then select **NetApp VSC** > **Update Host and Storage Data** .

VSC displays a Confirm dialog box that informs you that this operation might take a long time.

3. Click **OK**.

4. Select the discovered storage controllers that have the status "Authentication Failure", and then click **ACTIONS** > **Modify**.

5. Fill in the required information in the Modify Storage System dialog box.

6. Repeat steps 4 and 5 for all storage controllers with "Authentication Failure" status.

After the discovery process is complete, perform the following:

- Use VSC to configure ESXi host settings for hosts that display the Alert icon in the Adapter Settings column, the MPIO Settings column, or the NFS Settings column.

- Provide the storage system credentials.

Refresh the storage system display

You can use the update feature that is provided by ONTAP® tools for VMware vSphere to refresh the information about storage systems and to force Virtual Storage Console (VSC) to discover storage systems.

About this task

The "refresh" option is useful if you changed the default credentials for the storage systems after receiving an authentication error. You should always perform an update operation if you changed the storage system credentials after the storage system reported an "Authentication Failure Status". During the update operation, VSC tries to connect to the storage system by using the new credentials.

Depending on your system setup, this task can take a long time to complete.

Steps

1. On the VMware vSphere Client Home page, click **Storage Systems**.

2. Start the update:

If this location is…	Click…
Virtual Storage Console	The **REDISCOVER ALL** icon.
Datacenter	Right-click the datacenter, and then click **NetApp VSC** > **Update Host and Storage Data**.

3. In the Update Host and Storage Data dialog box, click **OK**.

The discovery might take few minutes depending on the number of hosts and storage systems in your datacenter. This discovery operation works in the background.

4. Click **OK** in the Success dialog box.

Configure alarm thresholds

You can use VSC to set alarms to notify you when volume thresholds and aggregate thresholds reach the set limits.

Steps

1. From the ONTAP tools Home page, click **Settings**.
2. Click **Unified Appliance Settings**.
3. Specify the percent values for the **Nearly full threshold (%)** field and the **Full threshold (%)** field for both the volume alarm thresholds and the aggregate alarm thresholds.

   While setting the values, you must keep the following information in mind:

   ◦ Clicking **Reset** resets the thresholds to the previous values.

     Clicking **Reset** does not reset the thresholds to the default values of 80 percent for "Nearly full" and 90 percent for "Full".

   ◦ There are two ways to set the values:

     ▪ You can use the up and down arrows next to the values to adjust the threshold values.
     ▪ You can slide the arrows on the track bar below the values to adjust the threshold values.

   ◦ The lowest value that you can set for the **Full threshold (%)** field for volumes and aggregates is 6 percent.

4. After specifying the required values, click **Apply**.

   You must click **Apply** for both volume alarm and aggregate alarm.

Configure user roles and privileges

You can configure new user roles for managing storage systems using the JSON file provided with ONTAP tools and ONTAP System Manager.

What you'll need

- You should have downloaded the ONTAP Privileges file from ONTAP tools using https://{virtual_appliance_IP}:9083/vsc/config/VSC_ONTAP_User_Privileges.zip.
- You should have configured ONTAP 9.8P1 or later storage.
- You should have logged in with administrator privileges for the storage system.

Steps

1. Unzip the downloaded `https://{virtual_appliance_IP}:9083/vsc/config/VSC_ONTAP_User_Privileges.zip` file.
2. Access ONTAP System Manager.
3. Click **CLUSTER** > **Settings** > **Users and Roles**.
4. Click **Add User**.

5. In the Add User dialog box, select **Virtualization products**.

6. Click **Browse** to select and upload the ONTAP Privileges JSON file.

   The PRODUCT field is auto populated.

7. Select the required capability from the PRODUCT CAPABILITY drop-down menu.

   The **ROLE** field is auto populated based on the product capability selected.

8. Enter the required username and password.

9. Select the privileges (Discovery, Create Storage, Modify Storage, Destroy Storage) required for the user, and then click **Add**.

The new role and user is added and you can see the detailed privileges under the role that you have configured.

> ⓘ The uninstall operation does not remove VSC roles but removes the localized names for the VSC-specific privileges and appends the prefix to "XXX missing privilege" them. This behavior happens because the vCenter Server does not provide an option to remove privileges. When you reinstall VSC or upgrade to a newer version of VSC, all of the standard VSC roles and VSC-specific privileges are restored.

Configure storage capability profiles

Overview of storage capability profiles

VASA Provider for ONTAP allows you to create storage capability profiles and map them to your storage. This helps you maintain consistency across the storage. You can also use VASA Provider to check for compliance between the storage and the storage capability profiles.

A storage capability is a set of storage system attributes that identifies a specific level of storage performance, storage efficiency, and other capabilities such as encryption for the storage object that is associated with the storage capability.

For traditional datastores, you can use a storage capability profile to create datastores consistently with common attributes, and assign QoS policy to them. During provisioning VSC displays clusters, SVMs, and aggregates that match the storage capability profile. You can generate a storage capability profile from existing traditional datastores by using the **GLOBAL AUTO-GENERATE PROFILES** option from the Storage Mapping menu. After the profile is created, you can use VSC to monitor the compliance of datastores with the profile.

When used with vVols datastores, the provisioning wizard can use multiple storage capability profiles to create different FlexVol volumes in the datastore. You can use the VM storage policy to automatically create vVols for a virtual machine in appropriate FlexVol volumes as defined. For example, you can create profiles for common storage classes (such as for performance limits and other capabilities like encryption or FabricPool). You can later create VM storage policies in vCenter Server representing business classes of virtual machines and link these to the appropriate storage capability profile by name (for example Production, Test, HR).

When used with vVols, the storage capability profile is also used to set the storage performance for the individual virtual machine and place it on the FlexVol volume in the vVols datastore that best satisfies the performance requirement. You can specify QoS policy with minimum and/or maximum IOPS for performance. You can use the default policies when you initially provision a virtual machine, or change your VM storage

policy later if your business requirements change. The default storage capability profiles for this release of ONTAP tools:

- AFF_Thick

- FAS_MAX20

- FAS_Default

- AFF_Default

- AFF_Tiering

- AFF_Encrypted

- AFF_Encrypted_Tiering

- AFF_Encrypted_Min50

- Platinum

- Bronze

The vCenter Server then associates the storage capability of a LUN or volume with the datastore that is provisioned on that LUN or volume. This enables you to provision a virtual machine in a datastore that matches the storage profile of the virtual machine and to ensure that all of the datastores in a datastore cluster have the same storage service levels.

With ONTAP tools, you can configure every virtual volume (vVols) datastore with a new storage capability profile that supports the provisioning of virtual machines with varying IOPS requirements on the same vVols datastore. While executing the VM provisioning workflow with IOPS requirement, all of the vVols datastores are listed in the compatible datastore list.

> (i) When you try to provision or modify virtual machines for vCenter Server earlier than 6.5, only the vVols datastores that contain storage capability profiles with performance set to "MAX_IOPS" are listed in the compatible datastore list. The remaining vVols datastores are listed in the incompatible datastore list. You can ignore this classification and select any vVols datastore from the incompatible datastore list to provision or modify the virtual machine.

Considerations for creating and editing storage capability profiles

You should be aware of the considerations for creating and editing storage capability profiles.

- You can configure Min IOPS only on AFF systems.

- You can configure QoS metrics at a virtual volume (vVols) datastore level.

  This capability provides greater flexibility in assigning varied QoS metrics for different VMDKs of the same virtual machine that is provisioned on a virtual datastore.

- You can configure storage capability profiles for both FAS and AFFdatastores.

  For FAS and AFF systems, you can configure space reserve to be either thick or thin.

- You can use storage capability profiles to provide encryption for your datastores.

- You cannot modify existing storage capability profiles (created prior to 7.2 version) after upgrading from an earlier version of the ONTAP tools for VMware vSphere to the latest version of the ONTAP tools.

  The legacy storage capability profiles are retained for backward compatibility. If the default templates are not in use, then during the upgrade to the latest version of the ONTAP tools, the existing templates are

overridden to reflect the new QoS metrics and tiering policies related to the performance of the storage capability profiles.

- You cannot modify or use the legacy storage capability profiles to provision new virtual datastores or VM storage policies.
- You must use new storage capability profiles for all new datastores.

Create storage capability profiles

You can use VSC to manually create storage capability profiles, automatically generate a profile based on the capabilities of a datastore, or modify a profile to meet your requirements.

What you'll need

You must have registered your VASA Provider instance with ONTAP tools for VMware vSphere.

After setting up a profile, you can edit the profile as required.

Steps

1. On the ONTAP tools Home page, click **Storage Capability Profiles**.
2. Create a profile or edit an existing profile, as required:

If you want to…	Do this…
Create a profile	Click CREATE .
Edit an existing profile	Click the profile that you want to modify from the profiles listed on the **Storage Capability Profiles** page.

> ⓘ To view the values that are associated with an existing profile, you can click the profile name in the Storage Capabilities Profile page. VASA Provider then displays the Summary page for that profile.

1. Complete the pages in the Create Storage Capability Profile wizard to set up a profile or to edit values to modify an existing profile.

   Most of the fields in this wizard are self-explanatory. The following table describes some of the fields for which you might require guidance.

Field	Explanation

Identifying multiple profiles	You can use the **Description** field on the Name and Description tab to describe the purpose of the storage capability profile. Providing a good description is useful because it is a good practice to set up different profiles based on the applications that are being used. For example, a business-critical application requires a profile with capabilities that support higher performance, such as an AFF platform. A datastore that is used for testing or training might use a profile with a lower performance FAS platform, and enable all of the storage efficiency capabilities and tiering to control costs. If you have enabled "linked" mode for your vCenter Servers, then you must select the vCenter Server for which you are creating the storage capability profile.
Platform	You can select your storage system to have either the AFF or FAS platform type. The options on the subsequent screens are updated based on your selection of the type of storage system.
Protocol	You can select from the available protocols listed based on the platform selected for the storage system. While configuring virtual machines you can configure VM storage policies with storage capability profile and the protocol field filters datastores based on specific protocol. The field 'Any' allows you to work with all protocols.

Performance	You can set traditional QoS policies for your storage system by using the Performance tab.
	• When you select **None**, a QoS policy with no limit (infinite) is applied to a data VVol.
	• When you select **QoS Policy Group**, then a traditional QoS policy is applied to a VVol.
	You can set the value for **Max IOPS** and **Min IOPS** which enables you to use the QoS functionality. If you select Infinite IOPS, the Max IOPS field is disabled. When applied for a traditional datastore, a QoS policy with "Max IOPS" value is created and assigned to a FlexVol volume. When used with a vVols datastore, a QoS policy with Max IOPS and Min IOPS values is created for each data vVols datastore.
	NOTE:
	◦ Max IOPS and Min IOPS can also be applied to the FlexVol volume for a traditional datastore.
	◦ You must ensure that the performance metrics are not also set separately at an storage virtual machine (SVM) level, an aggregate level, or a FlexVol volume level.

Storage Attributes	The storage attributes that you can enable in this tab depend on the storage type that you select in the Personality tab.
	• If you select FAS storage, you can configure space reserve (thick or thin), enable deduplication, compression, and encryption. The tiering attribute is disabled because this attribute is not applicable to FAS storage. • If you select AFF storage, you can enable encryption and tiering. Deduplication and compression are enabled by default for AFF storage and cannot be disabled. The tiering attribute enables the use of volumes that are part of a FabricPool-enabled aggregate (supported by VASA Provider for AFF systems with ONTAP 9.4 and later). You can configure one of the following policies for the tiering attribute: • None: Prevents volume data from being moved to the capacity tier • Snapshot: Moves user data blocks of volume Snapshot copies that are not associated with the active file system to the capacity tier

2. Review your selections on the Summary page, and then click **OK**.

   After you create a profile, you can return to the Storage Mapping page to view which profiles match which datastores.

Generate storage capability profiles automatically

VASA Provider for ONTAP enables you to automatically generate storage capability profiles for existing traditional datastores. When you select the auto-generate option for a datastore, VASA Provider creates a profile that contains the storage capabilities that are used by that datastore.

What you will need

• You must have registered your VASA Provider instance with Virtual Storage Console (VSC).

• VSC must have discovered your storage.

About this task

After you create a storage capability profile, you can modify the profile to include more capabilities. The Create storage capability profile wizard provides information about the capabilities that you can include in a profile.

Steps

1. From the ONTAP tools home page, click **Storage Mapping**.
2. Select the datastore from the available list.
3. From the Actions menu, select **Auto-generate**.
4. When the auto-generate process finishes, refresh the screen to view information about the new profile.

   The new profile is listed in the Associated profile column. The name of the new profile is based on the resources in the profile. You can rename the profile, if required.

Configure datastores

Provision traditional datastores

Provisioning a datastore creates a logical container for your virtual machines and their virtual machine disks (VMDKs). You can provision a datastore, and then attach the datastore to a single host, to all of the hosts in a cluster, or to all of the hosts in a datacenter.

What you will need

- To provision a datastore on an SVM that is directly connected to Virtual Storage Console (VSC), you must have added the SVM to VSC by using a user account that has the appropriate privileges, not the default vsadmin user account or vsadmin role.

  You can also provision a datastore by adding a cluster.

- You must ensure that the subnet details of all the networks to which the ESXi host is connected is entered in the kaminoprefs.xml.

  See "Enabling datastore mounting across different subnets".

- If you use NFS or iSCSI, and the subnet is different between your ESXi hosts and your storage system, then the NFS or iSCSI settings in the kaminoprefs preferences file must include ESXi host subnet masks.

  This preference file is also applicable to vVols datastore creation. See *Enable datastore mounting across different subnets* and *Configure the VSC preferences files* for more information.

- If you have enabled VASA Provider and you want to specify storage capability profiles for your NFS datastores or VMFS datastores, then you must have created one or more storage capability profiles.

- To create an NFSv4.1 datastore, you must have enabled NFSv4.1 at the SVM level.

The **Provision Datastore** option enables you to specify a storage capability profile for the datastore. Storage capability profiles help in specifying consistent service level objectives (SLOs) and simplify the provisioning process. You can specify a storage capability profile only if you have enabled VASA Provider. The ONTAP tools for VMware vSphere supports the following protocols:

- NFSv3 and NFSv4.1
- VMFS5 and VMFS6

VSC can create a datastore on either an NFS volume or a LUN:

- For an NFS datastore, VSC creates an NFS volume on the storage system, and then updates the export policies.

- For a VMFS datastore, VSC creates a new volume (or uses an existing volume, if you selected that option), and then creates a LUN and an igroup.

> ⓘ
> - ONTAP tools supports provisioning of VMFS5 and VMFS6 datastores up to the maximum VMFS LUN and volume size of 64TB when used with ASA and approved AFF systems running ONTAP 9.8 and later.
>
>   On other platforms the maximum LUN size supported is 16TB.
>
> - VMware does not support NFSv4.1 with datastore clusters.

- For Kerberos authentication, you will need the following:
  ◦ Windows machine with Active Directory (AD)
  ◦ Domain Name Server (DNS)
  ◦ Key Distribution Center (KDC)
  ◦ ONTAP Storage System (Cluster) with Kerberos configured
  ◦ ESXi host with Kerberos configured

If a storage capability profile is not specified during provisioning, you can later use the Storage Mapping page to map a datastore to a storage capability profile. You can apply storage QoS settings, throughput ceiling (Max IOPS) and throughput floor (Min IOPS) on data VMDK files of virtual machines provisioned on FlexGroup backed datastore. QoS settings can be applied either at datastore level or at individual virtual machine level by right clicking the datastore. The right click option is available only on those datastores or virtual machines that are backed by FlexGroup datastore. After the QoS is applied to a datastore, any pre-existing datastore or virtual machine QoS settings are overridden. QoS settings cannot be applied at a datastore level or at a virtual machine level for datastores that are provisioned on direct SVM's, because ONTAP does not support QoS at SVM management level.

Steps

1. You can access the datastore provisioning wizard using one of the following:

If you select from …	Perform the following…
vSphere Client home page	a. Click **Hosts and Clusters**. b. In the navigation pane, select the datacenter on which you want to provision the datastore. c. To specify the hosts to mount the datastore, see the next step.
ONTAP tools home page	a. Click **Overview**. b. Click **Getting Started** tab. c. Click **Provision** button. d. Click **Browse** to select the destination to provision the datastore as per the next step.

2. Specify the hosts on which you want to mount the datastore.

To make the datastore available to…	Do this…
All of the hosts in a datacenter	Right-click a datacenter, and then select **NetApp ONTAP tools** > **Provision Datastore**.
All of the hosts in a cluster	Right-click a host cluster, and then select **NetApp ONTAP tools** > **Provision Datastore**.
A single host	Right-click a host, and select **NetApp ONTAP tools** > **Provision Datastore**.

3. Complete the fields in the New Datastore dialog box to create the datastore.

   Most of the fields in the dialog box are self-explanatory. The following table describes some of the fields for which you might require guidance.

Section	Description
General	The General section of the New Datastore provisioning dialog box provides options to enter the destination, name, size, type, and protocol for the new datastore. You can select either **NFS** or **VMFS** type of protocol to configure a traditional datastore. For NFS, you can either select NFS3 or NFS 4.1. This release allows you to configure a VMFS datastore of maximum size 64TB. You can select the option "Distribute datastore data across the ONTAP cluster" to provision a FlexGroup volume on the storage system. Selecting this option automatically deselects the checkbox "Use Storage Capability Profile for provisioning". For FlexGroup datastore provisioning, ONTAP cluster that are 9.8 and above only are listed for selection. The vVols datastore type is used to configure a vVols datastore. If VASA Provider is enabled, then you can also decide whether to use storage capability profiles. The **Datastore cluster** option is available only for traditional datastores. You should use the **Advanced** option to specify **VMFS5** or **VMFS6** file system.
Kerberos authentication	If you have selected NFS 4.1 in the **General** page, select the security level. Kerberos authentication is supported only for Flexvols.

Storage system	You can select one of the listed storage capability profiles if you have selected the option in the General section. • If you are provisioning a FlexGroup datastore, then the storage capability profile for this datastore is not supported. The system-recommended values for the storage system and storage virtual machine are populated for ease. But you can modify the values if required. • For Kerberos authentication, the storage systems enabled for Kerberos are listed.
Storage attributes	By default, VSC populates the recommended values for **Aggregates** and **Volumes** options. You can customize the values based on your requirements. Aggregate selection is not supported for FlexGroup datastores as ONTAP manages the aggregate selection. The **Space reserve** option available under **Advanced** menu is also populated to give optimum results. (Optional) You can specify the initiator group name in the **Change initiator group name** field. • A new initiator group will be created with this name if one does not already exist. • The protocol name will be appended to the specified initiator group name. • If an existing igroup is found with the selected initiators, the igroup will be renamed with the provided name and will be reused. • If you do not specify an igroup name, igroup will be created with the default name.
Summary	You can review the summary of the parameters you specified for the new datastore. The field "Volume Style" enables you to differentiate the type of datastore created. The "Volume Style" can be either "FlexVol" or "FlexGroup".

ⓘ A FlexGroup that is part of a traditional datastore cannot shrink below the existing size but can grow by 120% maximum. Default snapshots are enabled on these FlexGroup volumes. . In the Summary section, click **Finish**.

Related information

Map datastores to storage capability profiles

You can map the datastores that are associated with VASA Provider for ONTAP to storage capability profiles. You can assign a profile to a datastore that is not associated with a storage capability profile.

What you will need

- You must have registered your VASA Provider instance with ONTAP® tools for VMware vSphere.
- Virtual Storage Console (VSC) must have already discovered your storage.

You can map traditional datastore with a storage capability profile or change the storage capability profile that is associated with a datastore. VASA Provider does *not* display any virtual volume (VVol) datastores on the Storage Mappings page. All the datastores that are referred to in this task are traditional datastores.

Steps

1. From the ONTAP tools Home page, click **Storage Mapping**.

   From the Storage Mapping page, you can determine the following information:

   - The vCenter Server that is associated with the datastore
   - How many profiles match the datastore

     The Storage Mapping page displays only traditional datastores. This page does not display any VVol datastores or qtree datastores.

   - Whether the datastore is currently associated with a profile

     A datastore can match multiple profiles, but a datastore can be associated with only one profile.

   - Whether the datastore is compliant with the profile that is associated with it

2. To map a storage capability profile to a datastore or to change the existing profile of a datastore, select the datastore.

   To locate specific datastores or other information on the Storage Mapping page, you can enter a name or a partial string in the search box. VSC displays the search results in a dialog box. To return to the full display, you should remove the text from the search box, and then click **Enter**.

3. From the Actions menu, select **Assign matching profile**.

4. Select the profile that you want to map to the datastore from the list of matching profiles that is provided in

the **Assign profile to datastore** dialog box, and then click **OK** to map the selected profile to the datastore.

5. Refresh the screen to verify the new assignment.

Assign QoS policies

The provisioning of FlexGroup datastores does not support assigning storage capability profiles to the datastores. But you can assign QoS policies to virtual machines that are created on FlexGroup backed datastores.

About this task

The QoS policies can be applied either at a virtual machine level or a datastore level. The QoS policies are required for a datastore to configure throughput (Max and Min IOPS) thresholds. When you set QoS on a datastore it is applied to the virtual machines residing on the datastore and not on the FlexGroup volume. But if you set QoS on all the virtual machines in a datastore, then any individual QoS settings for the virtual machines are overridden. This is applicable only to the virtual machines available in the datastore and not to any migrated or added virtual machines. If you want to apply QoS to newly added or migrated virtual machines of a particular datastore, then you have to manually set the QoS values.

> ⓘ You cannot apply QoS settings at a datastore or virtual machine level for datastores that are provisioned on direct storage VM's because ONTAP does not support QoS at storage VM management level.

Steps

1. On the ONTAP tools homepage, click **Menu** > **Host and Clusters**.

2. Right-click the required datastore or virtual machine and click **NetApp ONTAP tools** > **Assign QoS**.

3. In the Assign QoS dialog box, enter values the required IOPS values, and click **Apply**.

Verify datastore compliance with the mapped storage capability profile

You can quickly verify whether your datastores are compliant with the storage capability profiles that are mapped to the datastores.

What you will need

- You must have registered your VASA Provider instance with ONTAP® tools for VMware vSphere (VSC).
- VSC must have discovered your storage.

Steps

1. From the ONTAP tools Home page, click **Storage Mapping**.

2. Review the information in the Compliance Status column to identify non-compliant datastores and review the alerts for non-compliance reason.

> ⓘ When you click the **COMPLIANCE CHECK** button, VSC performs a rediscovery operation for all of the storage, which might take few minutes.

If a datastore is no longer compliant with its profile, then the Compliance Status column displays an alert stating the reason for non-compliance. For example, a profile might require compression. If that setting has

been changed on the storage, compression is no longer used, and the datastore is non-compliant.

When you discover a datastore that is not compliant with its profile, you can modify the settings on the volume backing the datastore to make the datastore compliant, or you can assign a new profile to the datastore.

You can modify the settings from the Storage Capability Profile page.

Provision vVols datastores

You can provision a vVols datastore using the Provision Datastore wizard only if VASA Provider is enabled in your ONTAP tools.

What you will need

- You should ensure that the subnet details of all the networks to which the ESXi hosted is connected is entered in the Kaminoprefs.xml.

  See **Enabling datastore mounting across different subnets** section.

- You should configure similar replication policy and schedule on the datastores at both the source and target sites for reverse replication to be successful.

The Provision datastore menu enables you to specify a storage capability profile for the datastore, which helps in specifying consistent service level objectives (SLOs) and simplifies the provisioning process. You can specify a storage capability profile only if you have enabled VASA Provider.

FlexVol volumes that are used as backing storage are displayed on the vVols dashboard only if they are running ONTAP 9.5 or later. You should not use the vCenter Server New Datastore wizard to provision vVols datastores.

- You must use cluster credentials to create vVols datastores.

  You cannot use SVM credentials to create vVols datastores.

- VASA Provider does not support the cloning of a virtual machine that is hosted on the vVols datastore of one protocol to another datastore with a different protocol.

- You should have completed cluster pairing and SVM pairing both on the source and destination sites.

About this task

> ⓘ The 9.10 release of ONTAP tools supports creating vVols datastores with vmdk size greater than 16TB for All SAN Array (ASA) type ONTAP 9.9.1 or later storage platforms.

Steps

1. From the vSphere Client home page, click **Hosts and Clusters**.

2. In the navigation pane, select the datacenter on which you want to provision the datastore.

3. Specify the hosts on which you want to mount the datastore.

To make the datastore available to…	Do this…

All of the hosts in a datacenter	Right-click a datacenter, and then select **NetApp VSC** > **Provision Datastore**.
All of the hosts in a cluster	Right-click a cluster, and then select **NetApp VSC** > **Provision Datastore**.
A single host	Right-click a host, and then select **NetApp VSC** > **Provision Datastore**.

4. Complete the fields in the New Datastore dialog box to create the datastore.

    Most of the fields in the dialog box are self-explanatory. The following table describes some of the fields for which you might require guidance.

Section	Description
General	The General section of the New Datastore dialog box provides options to enter the location, name, description, type, and protocol for the new datastore. The vVols datastore type is used to configure a vVols datastore. ⓘ If you are provisioning iSCSI vVols datastore for vVols replication, then before creating vVols datastore at the target site, you need to perform SnapMirror update and cluster rediscovery.
Storage system	This section enables you to select whether you want the vVols datastore to have either replication enabled or disabled. Only asynchronous type replication profile is allowed for this release. You can then select one or more storage capability profiles listed. The system recommended values of paired **Storage system** and **Storage VM** are populated automatically. The recommended values are populated only if they are paired in ONTAP. You can modify these values if required. **Note:** While creating FlexVol volumes in ONTAP, you should ensure to create them with the attributes you wish to select in the storage capability profile. Both read write and data protection FlexVol volumes should have similar attributes. After FlexVol volumes are created and SnapMirror is initialized in ONTAP, you should run a storage rediscovery in VSC to be able to see the new volumes.

Storage attributes	You should select the schedule for SnapMirror and the required FlexVol volume from the existing list. This schedule should be similar to the one selected in the VM Storage Policies page. The user should have created FlexVol volumes on ONTAP with SnapMirror that are listed. You can select the default storage capability profile to be used for creating vVols using the **Default storage capability profile** option. By default all the volumes are set to maximum Autogrow size to 120 % and default Snapshots are enabled on these volumes. **Note:** • A FlexVol volume that is part of a vVols datastore cannot shrink below the existing size but can grow by 120% maximum. Default snapshots are enabled on this FlexVol volume. • The minimum size of FlexVol volume that you should create is 5GB.

5. In the Summary section, click **Finish**.

Result

A Replication group is created in the backend when a vVols datastore is configured.

Related information

Analyze performance data using the vVols dashboard

Rebalance vVols datastores

ONTAP tools supports a command to rebalance FlexVol volumes in your datacenter. The main goal is to enable even space utilization among FlexVol volumes. ONTAP tools redistributes vVols among existing volumes based on space usage, thin provisioning, LUN count, and storage capability profiles.

The rebalancing of vVols datastore is performed by LUN move or file move. The criteria considered during vVols rebalancing are as follows:

• NFS vVol datastores are not supported

• Existing FlexVol volumes will not be resized and neither will new FlexVol volumes be added

• Only FlexVol volumes that have same storage capability or volume attributes are rebalanced

• FlexVol volumes with highest space utilization are considered for rebalancing

• All vVols associated with a virtual machine are moved to the same FlexVol volumes

• LUN and File count limit is retained

• Rebalance is not performed if the delta between the FlexVol volumes space utilization is 10%

The rebalance command removes empty FlexVol volumes to provide space for other datastores. Thus, the command enables you to remove unwanted FlexVol volumes so that they can be removed from the datastore. The command intends to move all the vVols associated with a virtual machine to same FlexVol volume. There is a precheck performed by the command before rebalance is started to minimize failures. But even with successful precheck, the rebalance operation might fail for one or more vVols. When this happens, then there is no rollback of the rebalance operation. So, vVols associated with a virtual machine might be placed on different FlexVol volumes and will result in warning logs.

> ℹ️
> - Parallel datastore and virtual machine operations are not supported.
> - You must perform cluster rediscovery operation after every vVols rebalance operation completes.
> - During vVols rebalance operation, if large number of vVols datastores are identified, then the transfer operation times out after the set default value.
>   - If this occurs, then you should modify the `vvol.properties` file to set the value to `offtap.operation.timeout.period.seconds=29700` and restart VASA Provider service.
> - If a FlexVol volume has Snapshots, then during the vVols rebalance operation, the vVols are not correctly rebalanced due to insufficient details on the space utilization.
> - You can set the VASA Provider property *enable.update.vvol.through.discovery* to true to get consistent data between ONTAP tools and ONTAP, when timeout occurs during container rebalance operation.