



NFS configuration for ESXi using VSC

System Manager Classic

NetApp
September 05, 2025

Table of Contents

NFS configuration for ESXi using VSC	1
NFS configuration for ESXi using VSC overview	1
NFS Client Configuration for ESXi workflow	1
Verify that the configuration is supported	2
Complete the NFS client configuration worksheet	3
Install	5
Add the storage cluster to VSC	5
Configure your network for best performance	6
Configure the ESXi host	6
Create an aggregate	7
Decide where to provision the new volume	8
Create a new NFS-enabled SVM	8
Add NFS access to an existing SVM	11
Verify that NFS is enabled on an existing SVM	12
Provision a datastore and create its containing volume	12
Verify NFS access from an ESXi host	13
Deploy the NFS Plug-in for VMware VAAI	13
Mount datastores on hosts	14

NFS configuration for ESXi using VSC

NFS configuration for ESXi using VSC overview

Using the ONTAP System Manager classic interface (ONTAP 9.7 and earlier), you can quickly set up NFS access for ESXi hosts to datastores using ONTAP volumes.

Use this procedure if:

- You are using a supported version of Virtual Storage Console for VMware vSphere (VSC) to provision a datastore and create a volume.
 - Beginning with VSC 7.0, VSC is part of the [ONTAP Tools for VMware vSphere](#) virtual appliance, which includes VSC, vStorage APIs for Storage Awareness (VASA) Provider, and Storage Replication Adapter (SRA) for VMware vSphere capabilities.
 - Be sure to check the [NetApp Interoperability Matrix Tool](#) to confirm compatibility between your current ONTAP and VSC releases.
- Your data network uses the default IPspace, the default broadcast domain, and the default failover group.


If your data network is flat, these default objects prescribe that LIFs will fail over correctly in the event of a link failure. If you are not using the default objects, you should refer to [Network Management](#) for information about how to configure LIF path failover.

- You want to use the Plug-In for VMware VAAI.
 - VMware vStorage APIs for Array Integration (VAAI) enable you to perform copy offload and space reservations. The Plug-In for VMware VAAI uses this to improve host performance because operations do not need to go through the ESXi host, thereby taking advantage of space- and time-efficient cloning in ONTAP.
 - Using VMware VAAI for datastore provisioning is a best practice.
 - The NFS Plug-In for VMware VAAI is available from the [NetApp Support](#) site.
- NFS access will be through NFSv3 and NFSv4 for use with VMware VAAI.

For more information, see [TR-4597: VMware vSphere for ONTAP](#) and the documentation for your VSC release.

NFS Client Configuration for ESXi workflow

When you make storage available to an ESXi host using NFS, you provision a volume on the using for and then connect to the NFS export from the ESXi host.



Verify that the configuration is supported

For reliable operation, you must verify that the entire configuration is supported. The lists the supported configurations for NFS and for Virtual Storage Console.

Steps

1. Go to the to verify that you have a supported combination of the following components:

[NetApp Interoperability Matrix Tool](#)

- ONTAP software
- NFS storage protocol
- ESXi operating system version

- Guest operating system type and version
- for (VSC) software
- NFS Plug-In for VAAI

2. Click the configuration name for the selected configuration.

Details for that configuration are displayed in the Configuration Details window.

3. Review the information in the following tabs:

- Notes

Lists important alerts and information that are specific to your configuration.

- Policies and Guidelines

Provides general guidelines for all NAS configurations.


Complete the NFS client configuration worksheet

You require network addresses and storage configuration information to perform NFS client configuration tasks.

Target network addresses

You require a subnet with two IP addresses for NFS data LIFs for each node in the cluster. There should be two separate networks for high availability. The specific IP addresses are assigned by ONTAP when you create the LIFs as part of creating the SVM.

If possible, separate network traffic on separate physical networks or on VLANs.



Subnet for LIFs:

Node or LIF with port to switch	IP address	Network mask	Gateway	VLAN ID	Home port
Node 1 / LIF to switch 1					
Node 2 / LIF to switch 1					
Node 3 / LIF to switch 1					
Node 4 / LIF to switch 1					
Node 1 / LIF to switch 2					
Node 2 / LIF to switch 2					
Node 3 / LIF to switch 2					
Node 4 / LIF to switch 2					

Storage configuration

If the aggregate and are already created, record their names here; otherwise, you can create them as required:

Node to own NFS export

Aggregate name

name

NFS export information

Export size

Export name (optional)

Export description (optional)

SVM information

If you are not using an existing , you require the following information to create a new one:

SVM name	
Aggregate for SVM root volume	
SVM user name (optional)	
SVM password (optional)	
SVM management LIF (optional)	
	Subnet:
	IP address:
	Network mask:
	Gateway:
	Home node:

Install

Virtual Storage Console for automates many of the configuration and provisioning tasks required to use storage with an ESXi host. is a plug-in to vCenter Server.

Before you begin

You must have administrator credentials on the vCenter Server used to manage the ESXi host.

About this task

- Virtual Storage Console is installed as a virtual appliance that includes Virtual Storage Console, vStorage APIs for Storage Awareness (VASA) Provider, and Storage Replication Adapter (SRA) for VMware vSphere capabilities.

Steps

- Download the version of that is supported for your configuration, as shown in the Interoperability Matrix tool.

[NetApp Support](#)

- Deploy the virtual appliance and configure it following the steps in *Deployment and Setup Guide*.

Add the storage cluster to VSC

Before you can provision the first datastore to an ESXi host in your Datacenter, you must add the cluster or a specific storage virtual machine (SVM) to Virtual Storage Console for

VMware vSphere. Adding the cluster enables you to provision storage on any SVM in the cluster.

Before you begin

You must have administrator credentials for the storage cluster or the that is being added.

About this task

Depending on your configuration, the cluster might have been discovered automatically, or might have already been added.

Steps

1. Log in to the vSphere Web Client.
2. Select **Virtual Storage Console**.
3. Select **Storage Systems** and then click the **Add** icon.
4. In the **Add Storage System** dialog box, enter the host name and administrator credentials for the storage cluster or and then click **OK**.

Configure your network for best performance

Ethernet networks vary greatly in performance. You can maximize the performance of the network by selecting specific configuration values.

Steps

1. Connect the host and storage ports to the same network.
It is best to connect to the same switches.
2. Select the highest speed ports available.
10 GbE or faster ports are best. 1 GbE ports are the minimum.
3. Enable jumbo frames if desired and supported by your network.

Jumbo frames should have an MTU of 9000 for ESXi hosts and storage systems, and 9216 for most switches. All network devices in the data path — including ESXi NICs, storage NICs, and switches — must support jumbo frames and should be configured for their maximum MTU values.

For more information, see [Check the network settings on the data switches](#) and the switch vendor documentation.

Configure the ESXi host

Configuring the ESXi host involves configuring ports and vSwitches, and using ESXi host best practice settings. After verifying that these settings are correct, you can then create an aggregate and decide where to provision the new volume.

Configure host ports and vSwitches

The ESXi host requires network ports for the NFS connections to the storage cluster.

About this task

It is recommended that you use IP Hash as the NIC teaming policy, which requires a single VMkernel port on a single vSwitch.

The host ports and storage cluster ports used for NFS must have IP addresses in the same subnet.

This task lists the high-level steps for configuring the ESXi host. If you require more detailed instructions, see the VMware publication *Storage* for your version of ESXi.

VMware

Steps

1. Log in to the vSphere Client, and then select the ESXi host from the inventory pane.
2. On the **Manage** tab, click **Networking**.
3. Click **Add Networking**, and then select **VMkernel** and **Create a vSphere standard switch** to create the VMkernel port and vSwitch.
4. Configure jumbo frames for the vSwitch (MTU size of 9000, if used).

Configure the ESXi host best practice settings

You must ensure that the ESXi host best practice settings are correct so that the ESXi host can correctly manage the loss of an NFS connection or a storage.

Steps

1. From the VMware vSphere Web Client **Home** page, click **vCenter > Hosts**.
2. Right-click the host, and then select **Actions > NetApp VSC > Set Recommended Values**.
3. In the **NetApp Recommended Settings** dialog box, ensure that all of the options are selected, and then click **OK**.

MPIO Settings do not apply to NFS. However, if you use other protocols, you should ensure that all options are selected.

The vCenter Web Client displays the task progress.

Create an aggregate

If you do not want to use an existing aggregate, you can create a new aggregate to provide physical storage to the volume which you are provisioning.

About this task

If you have an existing aggregate that you want to use for the new volume, you can skip this procedure.

Steps

1. Enter the URL `https://IP-address-of-cluster-management-LIF` in a web browser and log in to using your cluster administrator credential.
2. Navigate to the **Aggregates** window.
3. Click **Create**.
4. Follow the instructions on the screen to create the aggregate using the default RAID-DP configuration, and

then click **Create**.

To create an aggregate, select a disk type then specify the number of disks.

Name: aggr2

Disk Type: SAS

Number of Disks: 8 Max: 8 (excluding 1 hot spare), min: 5 for RAID-DP

RAID Configuration: RAID-DP; RAID group size of 16 disks

New Usable Capacity: 4.968 TB (Estimated)

Results

The aggregate is created with the specified configuration and added to the list of aggregates in the Aggregates window.

Decide where to provision the new volume

Before you create an NFS volume, you must decide whether to place it in an existing and, if so, how much configuration the requires. This decision determines your workflow.

Procedure

- If you want a new , follow the steps that you do for creating an NFS-enabled on an existing SVM.

[Creating a new NFS-enabled SVM](#)

You must choose this option if NFS is not enabled on an existing SVM.

- If you want to provision a volume on an existing that has NFS enabled but not configured, follow the steps that you do for configuring NFS access to an existing SVM.

[Configuring NFS access to an existing SVM](#)

This is the case if you followed this procedure to create the SVM.

- If you want to provision a volume on an existing that is fully configured for NFS access, follow the steps that you do for verifying settings on an existing SVM.

[Verifying settings on an existing SVM](#)

Create a new NFS-enabled SVM

Setting up a new SVM involves creating the new and enabling NFS. You can then configure NFS access on the ESXi host and verify that NFS is enabled for ESXi by using Virtual Storage Console.

Before you begin

- Your network must be configured and the relevant physical ports must be connected to the network.
- You must know which of the following networking components the will use:
 - The node and the specific port on that node where the data logical interface (LIF) will be created
 - The subnet from which the data LIF's IP address will be provisioned, or optionally the specific IP address you want to assign to the data LIF
- Any external firewalls must be appropriately configured to allow access to network services.

About this task

You can use a wizard that guides you through the process of creating the SVM, configuring DNS, creating a data LIF, and enabling NFS.

Steps

1. Navigate to the **SVMs** window.
2. Click **Create**.
3. In the **Storage Virtual Machine (SVM) Setup** window, create the SVM:

- a. Specify a unique name for the SVM.

The name must either be a fully qualified domain name (FQDN) or follow another convention that ensures unique names across a cluster.

- b. Select **NFS** for the data protocol.

If you plan to use additional protocols on the same SVM, you should select them even if you do not want to configure them immediately.

- c. Keep the default language setting, C.UTF-8.

This language is inherited by the volume that you create later, and a volume's language cannot be changed.

- d. **Optional:** If you enabled the CIFS protocol, change the security style to **UNIX**.

Selecting the CIFS protocol sets the security style to NTFS by default.

- e. **Optional:** Select the root aggregate to contain the root volume.

The aggregate that you select for the root volume does not determine the location of the data volume.

Storage Virtual Machine (SVM) Setup



Enter SVM basic details

SVM Details

① Specify a unique name and the data protocols for the SVM

SVM Name:

② IPspace:

③ Data Protocols: CIFS NFS iSCSI FC/FCoE NVMe

④ Default Language:

The language of the SVM specifies the default language encoding setting for the SVM and its volumes. Using a setting that incorporates UTF-8 character encoding is recommended.

⑤ Security Style:

Root Aggregate:

f. Optional: In the **DNS Configuration** area, ensure that the default DNS search domain and name servers are the ones that you want to use for this SVM.

DNS Configuration

Specify the DNS domain and name servers. DNS details are required to configure CIFS protocol.

⑥ Search Domains:

⑦ Name Servers:

g. Click **Submit & Continue**.

The SVM is created, but protocols are not yet configured.

4. In the **Data LIF Configuration** section of the **Configure CIFS/NFS protocol** page, specify the details of the first data LIF of the first datastore.

- Assign an IP address to the LIF automatically from a subnet you specify or manually enter the address.
- Click **Browse** and select a node and port that will be associated with the LIF.

 **Data LIF Configuration**

Retain the CIFS data LIF's configuration for NFS clients.

Data Interface details for CIFS

Assign IP Address:

IP Address: 10.224.107.199 [Change](#)

Port:

Do not enter any information to provision a volume. You can provision datastores later using

5. Click **Submit & Continue**.

The following objects are created:

- A data LIF named after the with the suffix “_nfs_lif1”
- An NFS server

6. For all other protocol configuration pages that are displayed, click **Skip**, and then configure the protocol later.
7. When the **SVM Administration** page is displayed, configure or defer configuring a separate administrator for this SVM:
 - Click **Skip**, and then configure an administrator later if required.
 - Enter the requested information, and then click **Submit & Continue**.
8. Review the **Summary** page, record any information that you might require later, and then click **OK**.

NFS clients need to know the IP address of the data LIF.

Results

A new SVM is created with NFS enabled.

Add NFS access to an existing SVM

To add NFS access to an existing SVM, you must first create a data logical interface (LIF). You can then configure NFS access on the ESXi host and verify that NFS is enabled for ESXi using Virtual Storage Console.

Before you begin

- You must know which of the following networking components the will use:
 - The node and the specific port on that node where the data LIF will be created
 - The subnet from which the data LIF's IP address will be provisioned, or optionally the specific IP address you want to assign to the data LIF
- Any external firewalls must be appropriately configured to allow access to network services.
- The NFS protocol must be allowed on the SVM.

This is the case if you did not follow this procedure to create the SVM while configuring a different protocol.

Steps


1. Navigate to the **Details** pane where you can configure the protocols of the SVM:

- Select the SVM that you want to configure.
- In the **Details** pane, next to **Protocols**, click **NFS**.

Protocols: NFS FC/FCoE

2. In the **Configure NFS protocol** dialog box, create a data LIF:

- Assign an IP address to the LIF automatically from a subnet you specify or manually enter the address.
- Click **Browse** and select a node and port that will be associated with the LIF.



Do not enter any information to provision a volume. You can provision datastores later using the Virtual Storage Console.

3. Click **Submit & Close**, and then click **OK**.

Verify that NFS is enabled on an existing SVM

If you choose to use an existing SVM, you must first verify that NFS is enabled on the SVM. You can then configure NFS access and verify that NFS is enabled for ESXi by using ESXi by using Virtual Storage Console.

Steps

1. Navigate to the **SVMs** window.
2. Click the **SVM Settings** tab.
3. In the **Protocols** pane, click **NFS**.
4. Verify that NFS is displayed as enabled.

If NFS is not enabled, you must enable it or create a new SVM.

Provision a datastore and create its containing volume

A datastore contains virtual machines and their VMDKs on the ESXi host. The datastore on the ESXi host is provisioned on a volume on the storage cluster.

Before you begin

Virtual Storage Console for VMware vSphere (VSC) must be installed and registered with the vCenter Server that manages the ESXi host.

VSC must have sufficient cluster or credentials to create the volume on the specified SVM.

About this task

VSC automates the datastore provisioning, including creating a volume on the specified SVM.

Steps

1. From the vSphere Web Client **Home** page, click **Hosts and Clusters**.
2. In the navigation pane, expand the datacenter where you want to provision the datastore.
3. Right-click the ESXi host, and then select **NetApp VSC > Provision Datastore**.

Alternatively, you can right-click the cluster when provisioning to make the datastore available to all hosts in the cluster.

4. Provide the required information in the wizard:



Verify NFS access from an ESXi host

After you have provisioned a datastore, you can verify that the ESXi host has NFS access by creating a virtual machine on the datastore and powering it on.

Steps

1. From the vSphere Web Client **Home** page, click **Hosts and Clusters**.
2. In the navigation pane, expand the datacenter to locate the datastore you previously created.
3. Click **Create a new virtual machine** and provide the required information in the wizard.

To verify NFS access, you should select the datacenter, ESXi host, and datastore that you previously created.

The virtual machine appears in the vSphere Web Client inventory.

4. Power on the virtual machine.

Deploy the NFS Plug-in for VMware VAAI

The plug-in is a software library that integrates the VMware Virtual Disk Libraries that are installed on the ESXi host. Downloading and installing the NFS Plug-In for VMware VAAI enables you to improve the performance of cloning operations by using the copy offload and space reservation options.

About this task

To provide consistent access to the virtual machines residing on the ESXi host on which you are installing the NFS plug-in, you can migrate virtual machines or install the NFS plug-in during planned maintenance.

Steps

1. Download the NFS Plug-In for VMware VAAI.

[NetApp Support](#)

You should download the online bundle (NetAppNasPlugIn.vib) of the most recent plug-in

2. Verify that VAAI is enabled on each ESXi host.

In VMware vSphere 5.0 and later, VAAI is enabled by default.

3. In Virtual Storage Console, go to **Tools > NFS VAAI Tools**.

4. Click **Select File** to upload the NetAppNasPlugIn.vib file.

5. Click **Upload**.

You see an uploaded successfully message.

6. Click **Install on host**.

7. Select the ESXi hosts on which you want to install the plug-in, click **Install**, and then click **OK**.

8. Reboot the ESXi host to enable the plug-in.

After installing the plug-in, you must reboot the ESXi host before installation is complete.

You do not need to reboot the storage system.

Mount datastores on hosts

Mounting a datastore gives a host access to storage. When datastores are provisioned by , they are automatically mounted to the host or cluster. You might need to mount a datastore on a host after you add the host to your VMware environment.

Steps

1. From the vSphere Web Client **Home** page, click **Hosts and Clusters**:
2. In the navigation pane, expand the datacenter that contains the host:
3. Right-click the host, and then select **NetApp VSC > Mount Datastores**.
4. Select the datastores that you want to mount, and then click **OK**.

Related information

[Virtual Storage Console, VASA Provider, and Storage Replication Adapter for VMware vSphere Administration for 9.6 release](#)

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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