



ONTAP® 9

NFS Client Configuration for ESX® Express Guide

February 2018 | 215-11183_F0
doccomments@netapp.com

Updated for ONTAP 9.3

 **NetApp®**

Contents

Deciding whether to use this guide	4
NFS Client Configuration for ESX workflow	5
Verifying that the configuration is supported	5
Completing the NFS client configuration worksheet	6
Installing VSC	8
Adding the storage cluster or SVM to VSC for VMware vSphere	8
Configuring your network for best performance	9
Configuring the ESXi host	9
Configuring host ports and vSwitches	9
Configuring the ESXi host best practice settings	10
Creating an aggregate	10
Deciding where to provision the new volume	11
Creating a new NFS-enabled SVM	11
Creating a new SVM	11
Editing the export policy of the SVM root volume	14
Configuring NFS access to an existing SVM	15
Adding NFS access to an existing SVM	16
Adding and editing the export policy of the SVM root volume	17
Verifying settings on an existing SVM	18
Verifying that NFS is enabled on an existing SVM	18
Adding and editing the export policy of the SVM root volume	18
Provisioning a datastore and creating its containing volume	19
Verifying NFS access from an ESXi host	20
Verifying that the host can write to and read from the NFS share	20
Where to find additional information	21
Copyright information	22
Trademark information	23
How to send comments about documentation and receive update	
notifications	24
Index	25

Deciding whether to use the NFS Client Configuration for ESX Express Guide

This guide describes how to quickly set up NFS access for ESXi hosts to files contained in new volumes.

You should use this guide if you want to configure NFS access for ESXi hosts to a volume in the following way:

- You are working with clusters running ONTAP 9.
- NFS access will be through NFSv3, not NFSv4 or NFSv4.1.
- You want to use best practices, not explore every available option.
- You do not want to read a lot of conceptual background.
- You want to use OnCommand System Manager, not the ONTAP command-line interface or an automated scripting tool.

The UI navigation in OnCommand System Manager 9.3 is different from the UI navigation in previous releases. This guide provides the common steps that you must perform to complete a task in any ONTAP 9 release. If you want the exact steps for navigating to a particular screen or window, you should view the OnCommand System Manager Online Help for your version of ONTAP.

Cluster management using System Manager

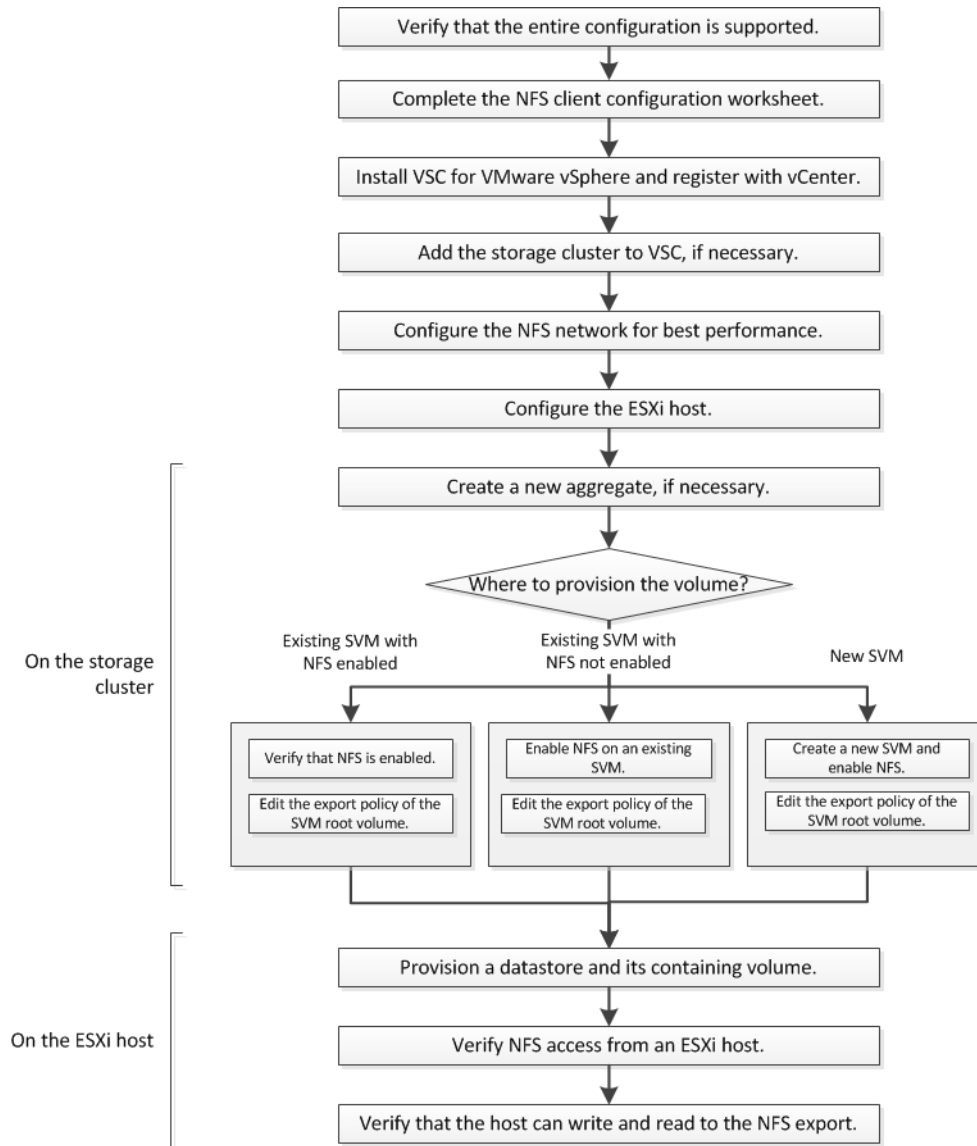
- Your data network uses the default IPspace, the default broadcast domain, and the default failover group.
If your data network is flat, using these default objects ensures that LIFs will fail over correctly in the event of a link failure. If you are not using the default objects, you should refer to the *Network Management Guide* for information on how to configure LIF path failover.
- You want to use Virtual Storage Console to provision a datastore and create a volume.
- You do not want to use the NetApp Plug-In for VMware VAAI.

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

- [*NetApp Technical Report 4068: VMware vSphere 5 on NetApp Clustered Data ONTAP 8.1*](#)
- [*NetApp Technical Report 4333: VMware vSphere 5 on NetApp Clustered Data ONTAP Best Practices Using vSphere Web Client*](#)
- [*NFS management*](#)
- [*NetApp Technical Report 4067: Clustered Data ONTAP Best Practice and NFS Implementation Guide*](#)
- [*NetApp Technical Report 4379: Name Services Best Practice Guide Clustered Data ONTAP*](#)

NFS Client Configuration for ESX workflow

When you make storage available to an ESXi host using NFS, you provision a volume on the storage virtual machine (SVM) using Virtual Storage Console for VMware vSphere and then connect to the NFS export from the ESXi host.



Verifying that the configuration is supported

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

Steps

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

NetApp Interoperability Matrix Tool

- ONTAP software
 - Host computer CPU architecture (for standard rack servers)
 - Specific processor blade model (for blade servers)
 - NFS storage protocol
 - ESXi operating system version
 - Guest operating system type and version
 - Virtual Storage Console for VMware vSphere (VSC) software
 - Windows Server version to run VSC
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.


Completing 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 SVM are already created, record their names here; otherwise, you can create them as required:

Node to own NFS export	
Aggregate name	
SVM name	

NFS export information

Export size	
Export name (optional)	
Export description (optional)	

SVM information

If you are not using an existing SVM, 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:
	Home port:

Installing VSC

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

Before you begin

- You must have administrator credentials on the Windows server.
- You must have administrator credentials on the vCenter Server used to manage the ESXi host.
- The Windows server on which you install VSC must have network connectivity to the ESXi host, to the vCenter Server, and to the storage cluster.

About this task

- You can install VSC on the same Windows server that runs the vCenter Server.
- You cannot install VSC on vCenter Server Appliance (vCSA).
You can install VSC on a separate Windows server or VM and register it with vCSA.
- VSC is not supported on a Windows server with IPv6 enabled.
- You do not need the VSC backup feature license for any tasks in this guide.

Steps

1. Download the version of VSC that is supported for your configuration from the Interoperability Matrix tool.
[NetApp Support](#)
2. Run the installation wizard on the Windows server.
3. After the wizard finishes, complete the web page that is displayed to register VSC with the vCenter Server that manages your ESXi host.

Adding the storage cluster or SVM to VSC for VMware vSphere

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 SVM 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 SVM and then click **OK**.

Configuring 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 ports are best. 1 GbE ports are the minimum.
3. Disable Ethernet flow control for all ports.
4. Enable jumbo frames.
Jumbo frames should have an MTU of 9000 for ESXi hosts and storage systems and 9216 for most switches.
All devices in the data path, including ESXi NICs, storage NICs, and switches, must support jumbo frames. Enabling jumbo frames improves network performance substantially.

Configuring 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.

Configuring 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 network.

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

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).

Configuring 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.

Creating 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 System Manager 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**.

Create Aggregate

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

Name:

Disk Type:

Number of Disks: 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)

Result

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

Deciding where to provision the new volume

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

Choices

- If you want a new SVM, follow the steps that you do for creating an NFS-enabled SVM.
[Creating a new NFS-enabled SVM](#) on page 11
You must choose this option if NFS is not enabled on an existing SVM.
- If you want to provision a volume on an existing SVM 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](#) on page 15
You should choose this option if you created the SVM for SAN access by using the relevant Express Guide.
- If you want to provision a volume on an existing SVM 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](#) on page 18

Creating a new NFS-enabled SVM

Setting up a new SVM involves creating the new SVM, enabling NFS, and editing the default export policy of the SVM root volume. You can then configure NFS access on the ESXi host and verify that NFS is enabled by using Virtual Storage Console.

Creating a new SVM

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

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 SVM 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.

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 SVM 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

1
Enter SVM basic details

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

SVM Name:

? IPspace:

? Volume Type: FlexVol volumes Infinite Volume

An SVM can contain either multiple FlexVol volumes or a single Infinite Volume.
You cannot change the volume type of the SVM after you set it.

? Data Protocols: CIFS NFS iSCSI FC/FCoE

? 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.

? 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.

? Search Domains:

? 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.
 - a. Assign an IP address to the LIF automatically from a subnet you specify or manually enter the address.
 - b. 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: Without a subnet

IP Address: 10.224.107.199 [Change](#)

Port: abccorp_1:e0b [Browse...](#)

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

5. Click **Submit & Continue.**

The following objects are created:

- A data LIF named after the SVM 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.

Result

A new SVM is created with NFS enabled.

Editing the export policy of the SVM root volume

You must add a rule to the default export policy to enable ESXi host access through NFSv3. Without such a rule, all NFS clients are denied access to the storage virtual machine (SVM) and its volumes.

About this task

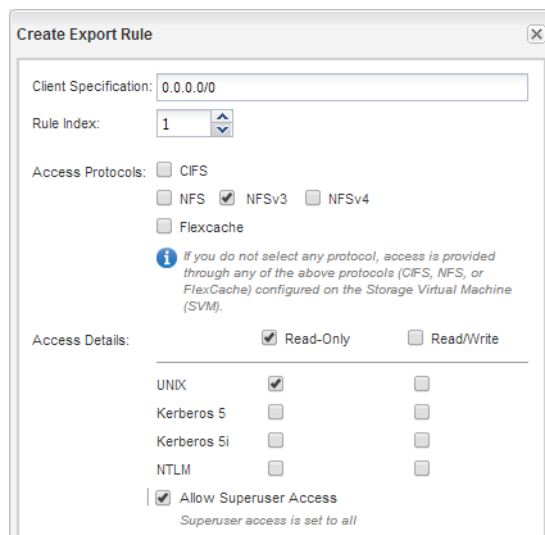
You should open all NFS access in the default export policy, and later restrict the access to individual volumes by creating custom export policies for individual volumes.

You should open all NFS access in the default export policy, and later restrict the access to individual volumes using the export policies managed by Virtual Storage Console.

Steps

- 1.** Navigate to the **SVMs** window.
- 2.** Click the **SVM Settings** tab.
- 3.** In the **Policies** pane, click **Export Policies**.

4. Select the export policy named **default**, which is applied to the SVM root volume.
5. In the lower pane, click **Add**.
6. In the **Create Export Rule** dialog box, create a rule that opens access to all clients for NFS clients.
 - a. In the **Client Specification** field, enter **0.0.0.0/0** so that the rule applies to all clients.
 - b. Retain the default value as **1** for the rule index.
 - c. Select **NFSv3**.
 - d. Clear all the check boxes except the **UNIX** check box under **Read-Only**.
 - e. Click **OK**.



Result

ESXi hosts can now access any volumes created on the SVM.

Configuring NFS access to an existing SVM

Adding NFS access to an existing SVM involves configuring NFS on the SVM and editing the export policy of the SVM root volume. You can then configure NFS access on the ESXi host and verify that NFS is enabled using Virtual Storage Console.

Steps

1. [Adding NFS access to an existing SVM](#) on page 16
2. [Adding and editing the export policy of the SVM root volume](#) on page 17

Adding NFS access to an existing SVM

To add NFS access to an existing SVM, you must create a data LIF and configure the export policy on the SVM root volume.

Before you begin

- You must know which of the following networking components the SVM 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.
- The NFS protocol must be allowed on the SVM.
This is the case if you created the SVM while following another Express Guide to configure a SAN protocol.

Steps

1. Navigate to the **Details** pane where you can configure the protocols of the SVM:
 - a. Select the SVM that you want to configure.
 - b. 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:
 - a. Assign an IP address to the LIF automatically from a subnet you specify or manually enter the address.
 - b. 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: Without a subnet ▼

IP Address: 10.224.107.199 Change

? Port: abccorp_1:e0b Browse...

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

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

Adding and editing the export policy of the SVM root volume

You must add a rule to the default export policy to enable ESXi host access through NFSv3. Without such a rule, all NFS clients are denied access to the storage virtual machine (SVM) and its volumes.

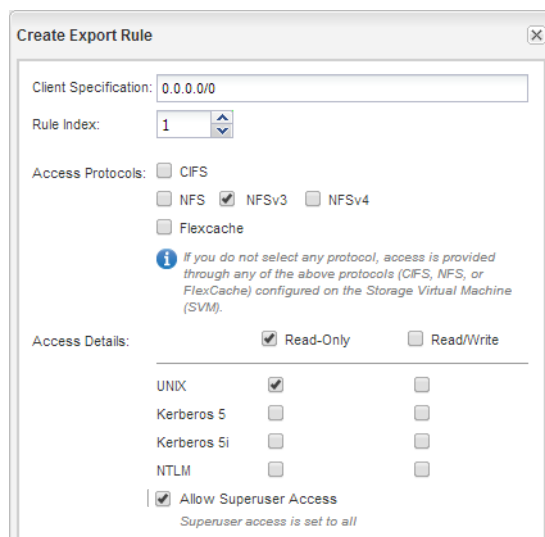
About this task

You should open all NFS access in the default export policy, and later restrict the access to individual volumes by creating custom export policies for individual volumes.

You should open all NFS access in the default export policy, and later restrict the access to individual volumes using the export policies managed by Virtual Storage Console.

Steps

1. Navigate to the **SVMs** window.
2. Click the **SVM Settings** tab.
3. In the **Policies** pane, click **Export Policies**.
4. Select the export policy named **default**, which is applied to the SVM root volume.
5. In the lower pane, click **Add**.
6. In the **Create Export Rule** dialog box, create a rule that opens access to all clients for NFS clients.
 - a. In the **Client Specification** field, enter **0.0.0.0/0** so that the rule applies to all clients.
 - b. Retain the default value as **1** for the rule index.
 - c. Select **NFSv3**.
 - d. Clear all the check boxes except the **UNIX** check box under **Read-Only**.
 - e. Click **OK**.



Result

ESXi hosts can now access any volumes created on the SVM.

Verifying settings on an existing SVM

Verifying the settings on an existing SVM involves verifying that NFS is enabled on the SVM and opening the export policy of the SVM root volume. You can then configure NFS access and verify that NFS is enabled by using Virtual Storage Console.

Verifying that NFS is enabled on an existing SVM

If you choose to use an existing storage virtual machine (SVM), you must verify that NFS is enabled.

Steps

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

After you finish

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

Adding and editing the export policy of the SVM root volume

You must add a rule to the default export policy to enable ESXi host access through NFSv3. Without such a rule, all NFS clients are denied access to the storage virtual machine (SVM) and its volumes.

About this task

You should open all NFS access in the default export policy, and later restrict the access to individual volumes by creating custom export policies for individual volumes.

You should open all NFS access in the default export policy, and later restrict the access to individual volumes using the export policies managed by Virtual Storage Console.

Steps

1. Navigate to the **SVMs** window.
2. Click the **SVM Settings** tab.
3. In the **Policies** pane, click **Export Policies**.
4. Select the export policy named **default**, which is applied to the SVM root volume.
5. In the lower pane, click **Add**.
6. In the **Create Export Rule** dialog box, create a rule that opens access to all clients for NFS clients.
 - a. In the **Client Specification** field, enter **0.0.0.0/0** so that the rule applies to all clients.
 - b. Retain the default value as **1** for the rule index.
 - c. Select **NFSv3**.
 - d. Clear all the check boxes except the **UNIX** check box under **Read-Only**.

e. Click **OK**.

Create Export Rule

Client Specification: 0.0.0.0/0

Rule Index: 1

Access Protocols: CIFS NFS NFSv3 NFSv4 Flexcache

If you do not select any protocol, access is provided through any of the above protocols (CIFS, NFS, or FlexCache) configured on the Storage Virtual Machine (SVM).

Access Details: Read-Only Read/Write

UNIX

Kerberos 5

Kerberos 5i

NTLM

Allow Superuser Access
Superuser access is set to all

Result

ESXi hosts can now access any volumes created on the SVM.

Provisioning a datastore and creating 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 storage virtual machine (SVM) credentials to create the volume.

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**.
4. Provide the required information in the wizard:

NetApp Datastore Provisioning Wizard

1 Name and type

2 Storage system

3 Details

4 Ready to complete

Specify the name and type of datastore you want to provision.

You will be able to select the storage system for your datastore in the next page of this wizard.

Name: *

Type: * NFS VMFS

Back Next Finish Cancel

After you finish

Set storage timeouts for guest operating systems using the Guest OS Tools in VSC.

For an ESXi cluster, make the datastore available to all the hosts in the cluster.

Verifying 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.

Verifying that the host can write to and read from the NFS share

Before using the NFS export, you should verify that the host can write data read from the NFS export.

About this task

If the cluster node on which the NFS export is created can be failed over to its partner node, you should verify reading the data while the node is failed over. This test might not be possible if the cluster is in production use.

Steps

1. On the vSphere Web Client **Home** page, click **Hosts and Clusters**.
2. In the navigation pane, click the **Storage** tab.
3. Expand the datacenter, and then select the new datastore.
4. In the center pane, click **Manage > Files**.
The contents of the datastore are displayed.
5. Create a new folder in the datastore and upload a file to the new folder.
You might need to install the Client Integration Plug-in.
6. Verify that you can access the file you just wrote.
7. Fail over the cluster node containing the NFS export and verify that you can still write and read a file.

If any of the tests fail, verify that the NFS service is running on the storage cluster and check the paths to the export.

8. Optional: If you failed over the cluster node, be sure to give back the node and return all LIFs to their home ports.
9. For an ESXi cluster, view the datastore from each ESXi host in the cluster and verify that the file you uploaded is displayed.

Related information

[High-availability configuration](#)

Where to find additional information

After you have successfully tested NFS client access, you can perform advanced NFS configuration or add SAN access. When protocol access is complete, you should protect the root volume of the SVM. There are express guides, comprehensive guides, and technical reports to help you achieve these goals.

NFS configuration

You can further configure NFS access using the following comprehensive guides and technical reports:

- [NFS management](#)
Describes how to configure and manage file access using the NFS protocol.
- [NetApp Technical Report 4068: VMware vSphere 5 on NetApp Clustered Data ONTAP 8.1](#)
Describes the best practices that should be followed when using ONTAP and VMware vSphere server virtualization environments.
- [NetApp Technical Report 4379: Name Services Best Practice Guide Clustered Data ONTAP](#)
Explains how to configure LDAP, NIS, DNS, and local file configuration for authentication purposes.
- [NetApp Technical Report 4067: Clustered Data ONTAP Best Practice and NFS Implementation Guide](#)
Serves as an NFSv3 and NFSv4 operational guide and provides an overview of ONTAP with a focus on NFSv4.
- [NetApp Technical Report 4333: VMware vSphere 5 on NetApp Clustered Data ONTAP Best Practices Using vSphere Web Client](#)
Describes the best practices that should be followed when using VMware vSphere 5.x with ONTAP and the VMware vSphere Web Client.

Root volume protection

After configuring protocols on the SVM, you should ensure that its root volume is protected:

- [Data protection](#)
Describes how to create load-sharing mirrors on every node of an ONTAP cluster to protect the SVM root volume, which is a NetApp best practice for NAS-enabled SVMs. Also describes how to quickly recover from volume failures or losses by promoting the SVM root volume from a load-sharing mirror.

Copyright information

Copyright © 2018 NetApp, Inc. All rights reserved. Printed in the U.S.

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

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

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

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

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

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

Trademark information

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

[*doccomments@netapp.com*](mailto:doccomments@netapp.com)

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

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

Index

A

- access
 - additional documentation [21](#)
 - verifying NFS [20](#)
- aggregates
 - creating new when provisioning a volume [10](#)
 - selecting for SVM [11](#)

C

- clusters
 - adding to VSC [8](#)
- comments
 - how to send feedback about documentation [24](#)
- configuration
 - NFS client workflow [5](#)
- configuring
 - best practice settings on ESXi host [10](#)
 - NFS client worksheet for [6](#)
 - ports and vSwitches on ESXi host [9](#)

D

- data LIFs
 - using wizard to guide creation [11](#)
- datastores
 - provisioning [19](#)
- discovering
 - storage cluster in VSC [8](#)
- documentation
 - additional information about protocol access [21](#)
 - how to receive automatic notification of changes to [24](#)
 - how to send feedback about [24](#)

E

- ESXi hosts
 - configuring best practice settings [10](#)
 - configuring ports and vSwitches on [9](#)
 - verifying NFS access [20](#)
 - verifying supported configurations [5](#)
- export policies
 - defining for root volumes and enabling NFSv3 [17](#), [18](#)
 - defining for root volumes, enabling NFS [14](#)
- express guides
 - additional documentation [21](#)

F

- feedback
 - how to send comments about documentation [24](#)
- flowcharts
 - NFS client configuration [5](#)

H

- hosts
 - verifying that host can write to and read from [20](#)

I

- information
 - how to send feedback about improving documentation [24](#)
- installing
 - VSC for VMware vSphere [8](#)
- Interoperability Matrix
 - verifying supported configurations using [5](#)

N

- network addresses
 - NFS client worksheet for [6](#)
- network ports
 - configuring for NFS on ESXi host [9](#)
- networks
 - configuring for best NFS performance [9](#)
- new volumes
 - deciding where to provision [11](#)
- NFS
 - additional documentation [21](#)
 - configuring network for best performance [9](#)
 - deciding whether to use this guide to set up for ESXi hosts [4](#)
 - verifying that it is enabled [18](#)
- NFS access
 - adding to existing SVM [16](#)
 - verifying [20](#)
- NFS client
 - configuration workflow [5](#)
 - verifying that configuration is supported [5](#)
- NFS exports
 - verifying that host can write to and read from [20](#)
- NFS shares
 - creating NFS datastore on [19](#)
- NFSv3
 - enabling ESXi host access [17](#), [18](#)

P

- performance
 - configuring NFS networks for best [9](#)
- policies
 - See* export policies
- ports
 - configuring for NFS on ESXi host [9](#)
- provisioning
 - datastores and creating containing NFS exports [19](#)
 - NFS client worksheet for [6](#)

R

- requirements
 - verifying supported configurations [5](#)
- root volumes
 - opening the export policies of [14](#), [17](#), [18](#)

S

- storage clusters
 - adding to VSC [8](#)
- storage configurations
 - NFS client configuration worksheet [6](#)
- suggestions
 - how to send feedback about documentation [24](#)
- supported configurations
 - verifying [5](#)
- SVM
 - verifying that NFS is enabled [18](#)
- SVMs
 - adding to VSC [8](#)
 - creating NFS enabled [11](#)
 - deciding whether to place new volumes in [11](#)
 - enabling NFS access [16](#)
 - using the wizard to guide creation [11](#)
 - using wizard to enable NFS [11](#)

T

- technical reports
 - additional information about file access [21](#)
- Twitter

how to receive automatic notification of documentation changes [24](#)

V

- vCenter Server
 - registering VSC for VMware vSphere with [8](#)
- Virtual Storage Console
 - See* VSC
- VMkernel ports
 - configuring for NFS on ESXi host [9](#)
- VMware vSphere
 - adding storage cluster to VSC for [8](#)
 - verifying supported configurations [5](#)
- volumes
 - creating new aggregates to provide physical storage [10](#)
 - deciding where to provision [11](#)
- VSC
 - adding storage cluster to [8](#)
 - installing for VMware vSphere [8](#)
- vSwitches
 - configuring for NFS on ESXi host [9](#)

W

- workflows
 - NFS client configuration [5](#)
- worksheets
 - NFS client configuration [6](#)