



FC configuration for ESXi using VSC

System Manager Classic

NetApp
September 05, 2025

Table of Contents

FC configuration for ESXi using VSC	1
FC configuration for ESXi using VSC overview	1
FC configuration workflow	1
Verify that the FC configuration is supported	2
Complete the FC configuration worksheet	3
Install Virtual Storage Console	5
Add the storage cluster or SVM to VSC for VMware vSphere	5
Update the HBA driver, firmware, and BIOS	6
Configure the ESXi host best practice settings	6
Create an aggregate	6
Decide where to provision the volume	7
Zone the FC switches by the host and LIF WWPNs	11
Provision a datastore and create its containing LUN and volume	11
Verify that the host can write to and read from the LUN	12

FC configuration for ESXi using VSC

FC configuration for ESXi using VSC overview

Using the ONTAP System Manager classic interface (ONTAP 9.7 and earlier), you can quickly set up the FC service on a storage virtual machine (SVM), provision a LUN, and make the LUN available as a datastore using an FC HBA on an ESX host computer.

Use this procedure if:

- You are using a supported version of Virtual Storage Console for VMware vSphere (VSC) to configure storage settings for your ESX host and to provision the datastores.
 - 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 network uses IPv4 addressing.
- You are using traditional FC HBAs on ESXi 5.x and traditional FC switches.

This procedure does not cover FCoE.

- You have at least two FC target ports available on each node in the cluster.

Onboard FC and UTA2 (also called “CNA”) ports, as well as some adapters are configurable. Configuring those ports is done in the ONTAP CLI and is not covered in this procedure.


- You are not configuring FC SAN boot.
- You are creating datastores on the host.

This procedure does not cover raw device mapping (RDM) disks or using N-port ID virtualization (NPIV) to provide FC directly to VMs.

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

FC configuration workflow

When you make storage available to a host using FC, you provision a volume and LUN on the storage virtual machine (SVM), and then connect to the LUN from the host.



Verify that the FC configuration is supported

For reliable operation, you must verify that the entire FC configuration is supported.

Steps

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

components:

- ONTAP software
- Host computer CPU architecture (for standard rack servers)
- Specific processor blade model (for blade servers)
- FC host bus adapter (HBA) model and driver, firmware, and BIOS versions
- Storage protocol (FC)
- ESXi operating system version
- Guest operating system type and version
- Virtual Storage Console (VSC) for VMware vSphere 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 SAN configurations.

Complete the FC configuration worksheet

You require FC initiator and target WWPNs and storage configuration information to perform FC configuration tasks.

FC host WWPNs

Port	WWPN
Initiator (host) port connected to FC switch 1	
Initiator (host) port connected to FC switch 2	

FC target WWPNs

You require two FC data LIFs for each node in the cluster. The WWPNs are assigned by ONTAP when you create the LIFs as part of creating the storage virtual machine (SVM).

LIF	WWPN
Node 1 LIF with port connected to FC switch 1	

LIF	WWPN
Node 2 LIF with port connected to FC switch 1	
Node 3 LIF with port connected to FC switch 1	
Node 4 LIF with port connected to FC switch 1	
Node 1 LIF with port connected to FC switch 2	
Node 2 LIF with port connected to FC switch 2	
Node 3 LIF with port connected to FC switch 2	
Node 4 LIF with port connected to FC 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 LUN
Aggregate name
SVM name

LUN information

LUN size
LUN name (optional)
LUN description (optional)

SVM information

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

SVM name
SVM IPspace
Aggregate for SVM root volume
SVM user name (optional)

SVM name	
SVM password (optional)	
SVM management LIF (optional)	
	Subnet:
	IP address:
	Network mask:
	Gateway:
	Home node:
	Home port:

Install Virtual Storage Console

Virtual Storage Console for VMware vSphere automates many of the configuration and provisioning tasks required to use NetApp FC storage with an ESXi host. Virtual Storage Console 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 Virtual Storage Console 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 the *Deployment and Setup Guide*.

Add 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**.

Update the HBA driver, firmware, and BIOS

If the FC host bus adapters (HBAs) in the ESX host are not running supported driver, firmware, and BIOS versions, you must update them.

Before you begin

You must have identified the supported driver, firmware, and BIOS versions for your configuration from the [NetApp Interoperability Matrix Tool](#).

About this task

Drivers, firmware, BIOS, and HBA utilities are provided by the HBA vendors.

Steps

1. List the installed HBA driver, firmware, and BIOS versions using the ESXi host console commands for your version of ESXi.
2. Download and install the new driver, firmware, and BIOS as needed from the HBA vendor's support site.

Installation instructions and any required installation utilities are available with the download.

Configure the ESXi host best practice settings

You must ensure that the host multipathing and best practice settings are correct so that the ESXi host can correctly manage the loss of an FC path or a storage failover event.

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**.

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.

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**.

The screenshot shows the 'Create Aggregate' dialog box. The 'Name' field is set to 'aggr2'. The 'Disk Type' field is set to 'SAS' with a 'Browse' button. The 'Number of Disks' field is set to '8' with a dropdown menu showing 'Max: 8 (excluding 1 hot spare), min: 5 for RAID-DP'. The 'RAID Configuration' field shows 'RAID-DP; RAID group size of 16 disks' with a 'Change' link. The 'New Usable Capacity' field shows '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 volume

Before you provision a volume to contain your LUNs, you need to decide whether to add the volume to an existing storage virtual machine (SVM) or to create a new SVM for the volume. You might also need to configure FC on an existing SVM.

About this task

If an existing SVM is already configured with the needed protocols and has LIFs that can be accessed from the host, it is easier to use the existing SVM.

You can create a new SVM to separate data or administration from other users of the storage cluster. There is no advantage to using separate SVMs just to separate different protocols.

Procedure

- If you want to provision volumes on an SVM that is already configured for FC, you must verify that the FC service is running.

[Verifying that the FC service is running on an existing SVM](#)

- If you want to provision volumes on an existing SVM that has FC enabled but not configured, configure iSCSI on the existing SVM.

Configuring FC on an existing SVM

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

- If you want to provision volumes on a new SVM, create the SVM.

Creating a new SVM

Verify that the FC service is running on an existing SVM


If you choose to use an existing storage virtual machine (SVM), you must verify that the FC service is running on the SVM by using ONTAP System Manager. You must also verify that FC logical interfaces (LIFs) are already created.

Before you begin

You must have selected an existing SVM on which you plan to create a new LUN.

Steps

1. Navigate to the **SVMs** window.
2. Select the required SVM.
3. Click the **SVM Settings** tab.
4. In the **Protocols** pane, click **FC/FCoE**.
5. Verify that the FC service is running.



The screenshot shows the SVM Settings window for an SVM named 'SVM1'. The top bar includes buttons for Edit, Start, Stop, and Refresh. The main area displays the following information:

FC/FCoE Interfaces:			
Network Interface	WWPN	Current Port	Status
FC_1_1	20:10:00:a0:98:29:18:76	innovate-01:0b	Enabled
FC_2_1	20:11:00:a0:98:29:18:76	innovate-02:0b	Enabled
FC_1_2	20:04:00:a0:98:29:18:76	innovate-02:0a	Enabled
FC_2_2	20:03:00:a0:98:29:18:76	innovate-01:0a	Enabled

The FC/FCoE service status is shown as "FC/FCoE service is running".

If the FC service is not running, start the FC service or create a new SVM.

6. Verify that there are at least two FC LIFs listed for each node.

If there are fewer than two FC LIFs per node, update the FC configuration on the SVM or create a new SVM for FC.

Configure FC on an existing SVM

You can configure FC on an existing storage virtual machine (SVM). The FC protocol

must already be enabled but not configured on the SVM. This information is intended for SVMs for which you are configuring multiple protocols, but have not yet configured FC.

Before you begin

Your FC fabric must be configured and the desired physical ports must be connected to the fabric.

Steps

1. Navigate to the **SVMs** window.
2. Select the SVM that you want to configure.
3. In the **SVMDetails** pane, verify that **FC/FCoE** is displayed with a gray background, which indicates that the protocol is enabled but not fully configured.

If **FC/FCoE** is displayed with a green background, the SVM is already configured.



4. Click the **FC/FCoE** protocol link with the gray background.

The Configure FC/FCoE Protocol window is displayed.

5. Configure the FC service and LIFs from the **Configure FC/FCoE protocol** page:
 - a. Select the **Configure Data LIFs for FC** check box.
 - b. Enter 2 in the **LIFs per node** field.

Two LIFs are required for each node, to ensure availability and data mobility.

- c. Ignore the optional **Provision a LUN for FCP storage** area, because the LUN is provisioned by Virtual Storage Console for VMware vSphere in a later step.
- d. Click **Submit & Close**.

6. Review the **Summary** page, record the LIF information, and then click **OK**.

Create a new SVM

The storage virtual machine (SVM) provides the FC target through which a host accesses LUNs. When you create the SVM, you also create logical interfaces (LIFs) that provide paths to the LUN. You can create an SVM to separate the data and administration functions of a user from those of the other users in a cluster.

Before you begin

- Your FC fabric must be configured and the desired physical ports must be connected to the fabric.

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 the IPspace that the SVM will belong to.

If the cluster does not use multiple IPspaces, the “Default” IPspace is used.

- c. Keep the default volume type selection.

Only FlexVol volumes are supported with SAN protocols.

- d. Select all of the protocols that you have licenses for and that you might use on the SVM, even if you do not want to configure all of the protocols immediately.

Selecting both NFS and CIFS when you create the SVM enables these two protocols to share the same LIFs. Adding these protocols later does not allow them to share LIFs.

If CIFS is one of the protocols you selected, then the security style is set to NTFS. Otherwise, the security style is set to UNIX.

- e. Keep the default language setting C.UTF-8.

- f. Select the desired root aggregate to contain the SVM root volume.

The aggregate for the data volume is selected separately in a later step.

- g. Click **Submit & Continue**.

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

4. If the **Configure CIFS/NFS protocol** page appears because you enabled CIFS or NFS, click **Skip** and then configure CIFS or NFS later.

5. If the **Configure iSCSI protocol** page appears because you enabled iSCSI, click **Skip** and then configure iSCSI later.

6. Configure the FC service and LIFs from the **Configure FC/FCoE protocol** page:

- a. Select the **Configure Data LIFs for FC** check box.

- b. Enter 2 in the **LIFs per node** field.

Two LIFs are required for each node to ensure availability and data mobility.

- c. Skip the optional **Provision a LUN for FCP storage** area because the LUN is provisioned by Virtual Storage Console for VMware vSphere in a later step.

- d. Click **Submit & Continue**.

7. When the **SVM Administration** appears, configure or defer configuring a separate administrator for this SVM:

- Click **Skip** and configure an administrator later if desired.

- Enter the requested information, and then click **Submit & Continue**.

8. Review the **Summary** page, record the LIF information, and then click **OK**.

Zone the FC switches by the host and LIF WWPNs

Zoning the FC switches enables the hosts to connect to the storage and limits the number of paths. You zone the switches using the management interface of the switches.

Before you begin

- You must have administrator credentials for the switches.
- You must know the WWPN of each host initiator port and of each FC LIF for the storage virtual machine (SVM) in which you created the LUN.

About this task

For details about zoning your switches, see the switch vendor's documentation.

You must zone by WWPN, not by physical port. Each initiator port must be in a separate zone with all of its corresponding target ports.

LUNs are mapped to a subset of the initiators in the igroup to limit the number of paths from the host to the LUN.

- By default, ONTAP uses Selective LUN Map to make the LUN accessible only through paths on the node owning the LUN and its HA partner.
- You still must zone all of the FC LIFs on every node for LUN mobility in case the LUN is moved to another node in the cluster.
- When moving a volume or a LUN, you must modify the Selective LUN Map reporting-nodes list before moving.

The following illustration shows a host connected to a four-node cluster. There are two zones, one zone indicated by the solid lines and one zone indicated by the dashed lines. Each zone contains one initiator from the host and a LIF from each storage node.

You must use the WWPNs of the target LIFs, not the WWPNs of the physical FC ports on the storage nodes. The LIF WWPNs are all in the range `2x:xx:00:a0:98:xx:xx:xx`, where x is any hexadecimal digit. The physical port WWPNs are all in the range `50:0a:09:8x:xx:xx:xx:xx`.

Steps

1. Log in to the FC switch administration program, and then select the zoning configuration option.
2. Create a new zone that includes the first initiator and all of the FC LIFs that connect to the same FC switch as the initiator.
3. Create additional zones for each FC initiator in the host.
4. Save the zones, and then activate the new zoning configuration.

Provision a datastore and create its containing LUN and volume

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

About this task

VSC automates the datastore provisioning, including creating a LUN and 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:



- Select **VMFS** as the datastore type.
- Select **FC/FCoE** as the VMFS protocol.
- Select **None** as the Storage Capability Profile.
- Select the box for **Thin provision**.
- Select the **Create new volume** check box.

Verify that the host can write to and read from the LUN

Before using the LUN, you should verify that the host can write data to the LUN and read it back.

About this task

If the cluster node on which the LUN 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. **Optional:** Fail over the cluster node containing the LUN and verify that you can still write and read a file.

If any of the tests fail, verify that the FC service is running on the storage cluster and check the FC paths to the LUN and FC switch zoning.

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

[ONTAP 9 High-Availability Configuration](#)

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.