



Manage LIFs for all SAN protocols

ONTAP 9

NetApp
February 20, 2026

Table of Contents

- Manage LIFs for all SAN protocols 1
 - Manage LIFs for all SAN protocols 1
 - Configure an NVMe LIF in ONTAP 1
 - What to know before moving a SAN LIF 2
 - Remove a SAN LIF from a port set 2
 - Move a SAN LIF 3
 - Delete a LIF in a SAN environment 4
 - SAN LIF requirements for adding nodes to a cluster 6
 - Configure iSCSI LIFs to return FQDN to host iSCSI SendTargets Discovery Operation 6

Manage LIFs for all SAN protocols

Manage LIFs for all SAN protocols

Initiators must use Multipath I/O (MPIO) and asymmetric logical unit access (ALUA) for failover capability for clusters in a SAN environment. If a node fails, LIFs do not migrate or assume the IP addresses of the failed partner node. Instead, the MPIO software, using ALUA on the host, is responsible for selecting the appropriate paths for LUN access through LIFs.

You need to create one or more iSCSI paths from each node in an HA pair, using logical interfaces (LIFs) to allow access to LUNs that are serviced by the HA pair. You should configure one management LIF for every storage virtual machine (SVM) supporting SAN.

Direct connect or the use of Ethernet switches is supported for connectivity. You must create LIFs for both types of connectivity.

- You should configure one management LIF for every storage virtual machine (SVM) supporting SAN. You can configure two LIFs per node, one for each fabric being used with FC and to separate Ethernet networks for iSCSI.

After LIFs are created, they can be removed from port sets, moved to different nodes within a storage virtual machine (SVM), and deleted.

Related information

- [Configure LIFs overview](#)
- [Create a LIF](#)

Configure an NVMe LIF in ONTAP

Certain requirements must be met when configuring NVMe LIFs.

Before you begin

NVMe must be supported by the FC adapter on which you create the LIF. Supported adapters are listed in [Hardware Universe](#).

About this task

Beginning with ONTAP 9.12.1 and later, you can configure two NVMe LIFs per node on a maximum of 12 nodes. In ONTAP 9.11.1 and earlier, you can configure two NVMe LIFs per node on a maximum of two nodes.

The following rules apply when creating an NVMe LIF:

- NVMe can be the only data protocol on data LIFs.
- You should configure one management LIF for every SVM that supports SAN.
- For ONTAP 9.5 and later, you must configure an NVMe LIF on the node containing the namespace and on node's HA partner.
- For ONTAP 9.4 only:
 - NVMe LIFs and namespaces must be hosted on the same node.

- Only one NVMe data LIF can be configured per SVM.

Steps

1. Create the LIF:

```
network interface create -vserver <SVM_name> -lif <LIF_name> -role  
<LIF_role> -data-protocol {fc-nvme|nvme-tcp} -home-node <home_node>  
-home-port <home_port>
```



NVME/TCP is available beginning with ONTAP 9.10.1 and later.

2. Verify that the LIF was created:

```
network interface show -vserver <SVM_name>
```

After creation, NVMe/TCP LIFs listen for discovery on port 8009.

Related information

- [network interface](#)

What to know before moving a SAN LIF

You only need to perform a LIF movement if you are changing the contents of your cluster, for example, adding nodes to the cluster or deleting nodes from the cluster. If you perform a LIF movement, you do not have to re-zone your FC fabric or create new iSCSI sessions between the attached hosts of your cluster and the new target interface.

You cannot move a SAN LIF using the `network interface move` command. SAN LIF movement must be performed by taking the LIF offline, moving the LIF to a different home node or port, and then bringing it back online in its new location. Asymmetric Logical Unit Access (ALUA) provides redundant paths and automatic path selection as part of any ONTAP SAN solution. Therefore, there is no I/O interruption when the LIF is taken offline for the movement. The host simply retries and then moves I/O to another LIF.

Using LIF movement, you can nondisruptively do the following:

- Replace one HA pair of a cluster with an upgraded HA pair in a way that is transparent to hosts accessing LUN data
- Upgrade a target interface card
- Shift the resources of a storage virtual machine (SVM) from one set of nodes in a cluster to another set of nodes in the cluster

Remove a SAN LIF from a port set

If the LIF you want to delete or move is in a port set, you must remove the LIF from the port set before you can delete or move the LIF.

About this task

You need to do Step 1 in the following procedure only if one LIF is in the port set. You cannot remove the last LIF in a port set if the port set is bound to an initiator group. Otherwise, you can start with Step 2 if multiple LIFs are in the port set.

Steps

1. If only one LIF is in the port set, use the `lun igroup unbind` command to unbind the port set from the initiator group.



When you unbind an initiator group from a port set, all of the initiators in the initiator group have access to all target LUNs mapped to the initiator group on all network interfaces.

```
cluster1::>lun igroup unbind -vserver vs1 -igroup ig1
```

Learn more about `lun igroup unbind` in the [ONTAP command reference](#).

2. Use the `lun portset remove` command to remove the LIF from the port set.

```
cluster1::> port set remove -vserver vs1 -portset ps1 -port-name lif1
```

Learn more about `lun portset remove` in the [ONTAP command reference](#).

Move a SAN LIF

If a node needs to be taken offline, you can move a SAN LIF to preserve its configuration information, such as its WWPN, and avoid rezoning the switch fabric. Because a SAN LIF must be taken offline before it is moved, host traffic must rely on host multipathing software to provide nondisruptive access to the LUN. You can move SAN LIFs to any node in a cluster, but you cannot move the SAN LIFs between storage virtual machines (SVMs).

Before you begin

If the LIF is a member of a port set, the LIF must have been removed from the port set before the LIF can be moved to a different node.

About this task

The destination node and physical port for a LIF that you want to move must be on the same FC fabric or Ethernet network. If you move a LIF to a different fabric that has not been properly zoned, or if you move a LIF to an Ethernet network that does not have connectivity between iSCSI initiator and target, the LUN will be inaccessible when you bring it back online.

Steps

1. View the administrative and operational status of the LIF:

```
network interface show -vserver vserver_name
```

Learn more about `network interface show` in the [ONTAP command reference](#).

2. Change the status of the LIF to `down` (offline):

```
network interface modify -vserver vserver_name -lif LIF_name -status-admin down
```

Learn more about `network interface modify` in the [ONTAP command reference](#).

3. Assign the LIF a new node and port:

```
network interface modify -vserver vserver_name -lif LIF_name -home-node node_name -home-port port_name
```

4. Change the status of the LIF to `up` (online):

```
network interface modify -vserver vserver_name -lif LIF_name -status-admin up
```

Learn more about `up` in the [ONTAP command reference](#).

5. Verify your changes:

```
network interface show -vserver vserver_name
```

Delete a LIF in a SAN environment

Before you delete a LIF, you should ensure that the host connected to the LIF can access the LUNs through another path.

Before you begin

If the LIF you want to delete is a member of a port set, you must first remove the LIF from the port set before you can delete the LIF.

System Manager

Delete a LIF with ONTAP System Manager (9.7 and later).

Steps

1. In System Manager, click **Network > Overview**, and then select **Network Interfaces**.
2. Select the storage VM from which you want to delete the LIF.
3. Click  and select **Delete**.

CLI

Delete a LIF with the ONTAP CLI.

Steps

1. Verify the name of the LIF and current port to be deleted:

```
network interface show -vserver vs1
```

2. Delete the LIF:

```
network interface delete
```

```
network interface delete -vserver vs1 -lif lif1
```

Learn more about `network interface delete` in the [ONTAP command reference](#).

3. Verify that you deleted the LIF:

```
network interface show
```

```
network interface show -vserver vs1
```

```
Logical Status      Network              Current   Current Is
Vserver Interface  Admin/Oper Address/Mask  Node      Port
Home
-----
vs1
      lif2          up/up         192.168.2.72/24  node-01   e0b
true
      lif3          up/up         192.168.2.73/24  node-01   e0b
true
```

Learn more about `network interface show` in the [ONTAP command reference](#).

SAN LIF requirements for adding nodes to a cluster

You need to be aware of certain considerations when adding nodes to a cluster.

- You must create LIFs on the new nodes as appropriate before you create LUNs on those new nodes.
- You must discover those LIFs from the hosts as dictated by the host stack and protocol.
- You must create LIFs on the new nodes so that the LUN and volume movements are possible without using the cluster interconnect network.

Configure iSCSI LIFs to return FQDN to host iSCSI SendTargets Discovery Operation

Beginning with ONTAP 9, iSCSI LIFs can be configured to return a Fully Qualified Domain Name (FQDN) when a host OS sends an iSCSI SendTargets Discovery Operation. Returning a FQDN is useful when there is a Network Address Translation (NAT) device between the host OS and the storage service.

About this task

IP addresses on one side of the NAT device are meaningless on the other side, but FQDNs can have meaning on both sides.



The FQDN value interoperability limit is 128 characters on all host OS.

Steps

1. Change the privilege setting to advanced:

```
set -privilege advanced
```

2. Configure iSCSI LIFs to return FQDN:

```
vserver iscsi interface modify -vserver SVM_name -lif iscsi_LIF_name  
-sendtargets_fqdn FQDN
```

In the following example, the iSCSI LIFs are configured to return storagehost-005.example.com as the FQDN.

```
vserver iscsi interface modify -vserver vs1 -lif vs1_iscsi1 -sendtargets-fqdn  
storagehost-005.example.com
```

3. Verify that sendtargets is the FQDN:

```
vserver iscsi interface show -vserver SVM_name -fields sendtargets-fqdn
```

In this example, storagehost-005.example.com is displayed in the sendtargets-fqdn output field.

```
cluster::vserver*> vserver iscsi interface show -vserver vs1 -fields
sendtargets-fqdn
vserver lif          sendtargets-fqdn
-----
vs1      vs1_iscsi1  storagehost-005.example.com
vs1      vs1_iscsi2  storagehost-006.example.com
```

Related information

[ONTAP command reference](#)

Copyright information

Copyright © 2026 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.