



ONTAP® 9

Cluster and SVM Peering Express Guide

July 2019 | 215-11182_JO
doccomments@netapp.com

Updated for ONTAP 9.6

 **NetApp®**

Contents

Deciding whether to use this guide	4
Prerequisites for cluster peering	5
Cluster and SVM peering workflow	7
Preparing for cluster peering	7
Configuring peer relationships (starting with ONTAP 9.3)	8
Creating intercluster LIFs (starting with ONTAP 9.3)	8
Creating a cluster peer relationship (starting with ONTAP 9.3)	9
Creating SVM peer relationship	10
Configuring peer relationships (ONTAP 9.2 and earlier)	11
Creating intercluster interfaces on all nodes (ONTAP 9.2 or earlier)	11
Creating a cluster peer relationship (ONTAP 9.2 or earlier)	13
Where to find additional information	15
Copyright	16
Trademark	17
How to send comments about documentation and receive update notifications	18

Deciding whether to use the Cluster and SVM Peering Express Guide

This guide describes how cluster administrators create authenticated peer relationships between clusters and SVMs to enable the clusters to communicate with each other so that you can replicate data between volumes in different clusters.

ONTAP System Manager in ONTAP 9.3 simplifies the way that you configure peer relationships between clusters and between SVMs. This guide describes the cluster peering procedure and SVM peering procedure for all ONTAP 9 versions. You should use the appropriate procedure for your version of ONTAP.

You should use this guide if you want to create cluster peer relationships and SVM peer relationships in the following way:

- You are working with clusters running ONTAP 9.
- You want cluster peering relationships that are authenticated.
- 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 ONTAP System Manager, not the ONTAP command-line interface or an automated scripting tool.

If these assumptions are not correct for your situation, you should see the following resources:

- [Cluster and SVM peering](#)
Describes how to use the command-line interface to set up cluster peering relationships and SVM peering relationships.
- [Network and LIF management](#)
Describes how to use the command-line interface to configure subnets, intercluster LIFs, routes, firewall policies, and other networking components.
- [NetApp Documentation: OnCommand Workflow Automation \(current releases\)](#)
OnCommand Workflow Automation enables you to run prepackaged workflows that automate management tasks such as the workflows described in Express Guides.

Related information

[Cluster management using System Manager](#)

Prerequisites for cluster peering

Before you set up cluster peering, you should confirm that the connectivity, port, IP address, subnet, firewall, and cluster-naming requirements are met.

Connectivity requirements

Every intercluster LIF on the local cluster must be able to communicate with every intercluster LIF on the remote cluster.

Although it is not required, it is typically simpler to configure the IP addresses used for intercluster LIFs in the same subnet. The IP addresses can reside in the same subnet as data LIFs, or in a different subnet. The subnet used in each cluster must meet the following requirements:

- The subnet must have enough IP addresses available to allocate to one intercluster LIF per node. For example, in a six-node cluster, the subnet used for intercluster communication must have six available IP addresses.

Each node must have an intercluster LIF with an IP address on the intercluster network.

Intercluster LIFs can have an IPv4 address or an IPv6 address.

Note: ONTAP 9 enables you to migrate your peering networks from IPv4 to IPv6 by optionally allowing both protocols to be present simultaneously on the intercluster LIFs. In earlier releases, all intercluster relationships for an entire cluster were either IPv4 or IPv6. This meant that changing protocols was a potentially disruptive event.

Port requirements

You can use dedicated ports for intercluster communication, or share ports used by the data network. Ports must meet the following requirements:

- All ports that are used to communicate with a given remote cluster must be in the same IPspace. You can use multiple IPspaces to peer with multiple clusters. Pair-wise full-mesh connectivity is required only within an IPspace.
- The broadcast domain that is used for intercluster communication must include at least two ports per node so that intercluster communication can fail over from one port to another port. Ports added to a broadcast domain can be physical network ports, VLANs, or interface groups (ifgrps).
- All ports must be cabled.
- All ports must be in a healthy state.
- The MTU settings of the ports must be consistent.

Firewall requirements

Firewalls and the intercluster firewall policy must allow the following protocols:

- ICMP service
 - TCP to the IP addresses of all the intercluster LIFs over the ports 10000, 11104, and 11105
 - Bidirectional HTTPS between the intercluster LIFs
- Although HTTPS is not required when you set up cluster peering using the CLI, HTTPS is required later if you use ONTAP System Manager to configure data protection.

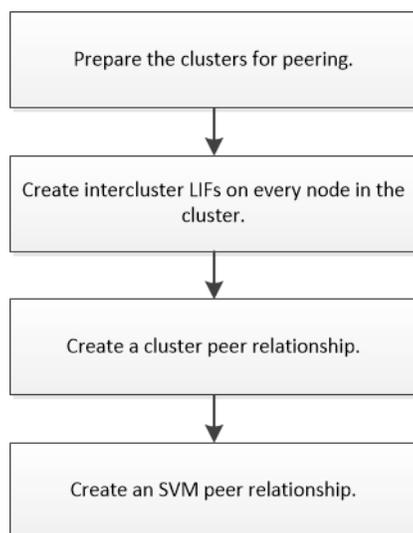
The default **intercluster** firewall policy allows access through the HTTPS protocol and from all IP addresses (0.0.0.0/0). You can modify or replace the policy if necessary.

Related information

[Data protection](#)

Cluster and SVM peering workflow

Setting up a peering relationship involves preparing each cluster for peering, creating intercluster logical interfaces (LIFs) on each node of each cluster, setting up a cluster peer relationship, and then setting up an SVM peering relationship.



If you are running ONTAP 9.2 or earlier, you create an SVM peering relationship while creating a data protection relationship between the source volume and the destination volume.

Preparing for cluster peering

Before creating a cluster peering relationship, you must verify that the time on each cluster is synchronized with an external Network Time Protocol (NTP) server, and determine the subnets, ports, and passphrases that you want to use.

Steps

1. If you are running ONTAP 9.2 or earlier, determine the passphrase that you want to use for each cluster peer relationship.

The passphrase must include at least eight characters.

For the relationship between...	The passphrase is...
Cluster A and Cluster B	

Starting with ONTAP 9.3, you can generate the passphrase from the remote cluster while creating the cluster peer relationship.

[Creating a cluster peer relationship \(starting with ONTAP 9.3\)](#)

2. Identify the subnets, IP addresses, and ports that you will use for intercluster LIFs.

By default, the IP address is automatically selected from the subnet. If you want to specify the IP address manually, you must ensure that the IP address either is already available in the subnet or can be added to the subnet later. Information about subnets is available in the Network tab.

The following table assumes that each cluster has four nodes. If a cluster has more than four nodes, you can record the ports on another piece of paper.

	Cluster A	Cluster B
Subnet (ONTAP 9.2 or earlier)		
IP address (starting with ONTAP 9.3, optional for ONTAP 9.2 or earlier)		
Node 1 port		
Node 2 port		
Node 3 port		
Node 4 port		

Configuring peer relationships (starting with ONTAP 9.3)

A peer relationship defines the network connections that enable clusters and SVMs to exchange data securely. ONTAP 9.3 simplifies the way that you configure peer relationships between clusters and between SVMs.

Creating intercluster LIFs (starting with ONTAP 9.3)

Creating intercluster logical interfaces (LIFs) enables the cluster network to communicate with a node. You must create an intercluster LIF within each IPspace that will be used for peering, on each node in each cluster for which you want to create a peer relationship.

About this task

For example, if you have a four-node cluster that you want to peer with cluster X over IPspace A, and peer with cluster Y over IPspace Y, then you need a total of eight intercluster LIFs; Four on IPspace A (one per node), and four on IPspace Y (one per node).

You must perform this procedure on both clusters for which you want to create a peer relationship.

Steps

1. Click **Configuration > Advanced Cluster Setup**.
2. In the **Setup Advanced Cluster Features** window, click **Proceed** next to the **Cluster Peering** option.
3. Select an IPspace from the **IPspace** list.
4. Enter the IP address, port, network mask, and gateway details of each node.

Intercluster LIF Details per Node

IPspace:

	IP Address	Port	Netmask	Gateway (Optional)	
st150-vs1m-ucs103a	<input type="text" value="10.53.32.1"/>	<input type="text" value="e0d"/>	<input type="text" value="255.255.240.0"/>	<input type="text"/>	<input checked="" type="checkbox"/> Use same net...and gateway
st150-vs1m-ucs103b	<input type="text" value="10.53.32.2"/>	<input type="text" value="e0d"/>			

5. Click **Submit and Continue**.

After you finish

You should enter the cluster details in the Cluster Peering window to continue with cluster peering.

Creating a cluster peer relationship (starting with ONTAP 9.3)

You can create a cluster peer relationship between two clusters by providing a system-generated passphrase and the IP addresses of the intercluster LIFs of the remote cluster.

About this task

Beginning in ONTAP 9.6, cluster peering encryption is enabled by default on all newly created cluster peering relationships. Cluster peering encryption must be enabled manually for peering relationship created prior to upgrading to ONTAP 9.6. Cluster peering encryption is not available for clusters running ONTAP 9.5 or earlier. Therefore, both clusters in the peering relationship must be running ONTAP 9.6 in order to enable cluster peering encryption.

Cluster peering encryption uses the Transport Security Layer (TLS) to secure cross-cluster peering communications for ONTAP features such as SnapMirror and FlexCache.

Steps

1. In the **Target Cluster Intercluster LIF IP addresses** field, enter the IP addresses of the intercluster LIFs of the remote cluster.
2. Generate a passphrase from the remote cluster.
 - a. Specify the management address of the remote cluster.
 - b. Click **Management URL** to launch ONTAP System Manager on the remote cluster.
 - c. Log in to the remote cluster.
 - d. In the **Cluster Peers** window, click **Generate Peering Passphrase**.
 - e. Select the IPspace, validity of the passphrase, and SVM permissions.

You can allow all of the SVMs or selected SVMs for peering. When a SVM peer request is generated, the permitted SVMs are automatically peered with the source SVMs without requiring you to accept the peer relationship from the remote SVMs.

- f. Click **Generate**.

The passphrase information is displayed.

Generate Peering Passphrase



Passphrase generated successfully

Use the following information for peering based on the IPspace "Default":

Intercluster LIF IP Address 172.21.91.12

Passphrase QS7k+laFYJzclV9UMPXvHgWd

Passphrase Validity Valid Until Mon Nov ... America/New_Y

SVM Permissions All

Email passphrase details

Copy passphrase details

Done

- g. Click **Copy passphrase details** or **Email passphrase details**.
 - h. Click **Done**.
3. In the source cluster, enter the generated passphrase that you obtained in Step 2.
 4. Click **Initiate Cluster Peering**.
The cluster peer relationship is successfully created.
 5. Click **Continue**.

After you finish

You should specify the SVM details in the SVM Peering window to continue with the peering process.

Creating SVM peer relationship

The storage virtual machine (SVM) peering enables you to establish a peer relationship between two SVMs for data protection.

Steps

1. Select the initiator SVM.

2. Select the target SVM from the list of permitted SVMs.
3. Click **Initiate SVM Peering**.
4. Click **Continue**.

After you finish

You can view the intercluster LIFs, cluster peer relationship, and SVM peer relationship in the Summary window.

Configuring peer relationships (ONTAP 9.2 and earlier)

A peer relationship defines network connections that enable clusters and SVMs to exchange data securely. You must create a cluster peer relationship before you can create an SVM peer relationship.

Creating intercluster interfaces on all nodes (ONTAP 9.2 or earlier)

Clusters communicate with each other through logical interfaces (LIFs) that are dedicated to intercluster communication. You must create an intercluster LIF within each IPspace that will be used for peering, on each node in each cluster for which you want to create a peer relationship.

Before you begin

You must have identified the subnet and ports, and optionally the IP addresses, that you plan to use for the intercluster LIFs.

About this task

You must perform this procedure on both clusters for which you want to create a peer relationship. For example, if you have a four-node cluster that you want to peer with cluster X over IPspace A, and peer with cluster Y over IPspace Y, then you need a total of eight intercluster LIFs; Four on IPspace A (one per node), and four on IPspace Y (one per node).

Steps

1. Create an intercluster LIF on one node of the source cluster:
 - a. Navigate to the **Network Interfaces** window.
 - b. Click **Create**.
The Create Network Interface dialog box is displayed.
 - c. Enter a name for the intercluster LIF.

Example

You can use “icl01” for the intercluster LIF on the first node, and “icl02” for the intercluster LIF on the second node.

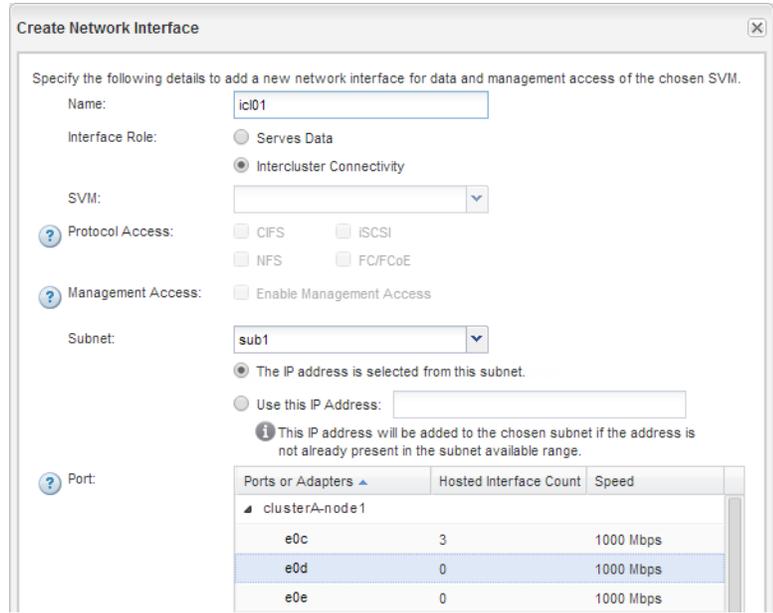
- d. Select **Intercluster Connectivity** as the interface role.
- e. Select the IPspace.
- f. In the **Add Details** dialog box, select **Using a subnet** from the **Assign IP Address** drop-down list, and then select the subnet that you want to use for intercluster communication.

By default, the IP address is automatically selected from the subnet after you click **Create**. If you do not want to use the IP address that is automatically selected, you must manually specify the IP address that the node uses for intercluster communication.

- g. Optional: If you want to manually specify the IP address that the node uses for intercluster communication, select **Use this IP Address**, and type the IP address.

You must ensure that the IP address that you want to use either is already available in the subnet or can be added to the subnet later.

- h. In the **Ports** area, click the node that you are configuring, and select the port that you want to use for this node.
- i. If you decided not to share ports for intercluster communication with data communication, confirm that the selected port displays “0” in the **Hosted Interface Count** column.



- j. Click **Create**.
2. Repeat Step 1 for each node in the cluster.
Each node in the cluster has an intercluster LIF.
 3. Make a note of the IP addresses of the intercluster LIFs so that you can use them later when you create peer relationships with other clusters:

- a. In the **Network Interfaces** window, in the **Role** column, click , clear the **All** check box, and then select **Intercluster**.

The Network Interfaces window displays only intercluster LIFs.

- b. Note down the IP addresses that are listed in the **IP Addresses/WWPN** column, or leave the **Network Interfaces** window open so that you can retrieve the IP addresses later.

You can click the column display icon () to hide the columns that you do not want to view.

Result

All of the nodes in each cluster have intercluster LIFs that can all communicate with each other.

Creating a cluster peer relationship (ONTAP 9.2 or earlier)

You can create a cluster peer relationship between two clusters by entering a predetermined passphrase and the IP addresses of the intercluster LIFs of the remote cluster, and then verifying that the relationship was created successfully.

Before you begin

- You must know the IP addresses of all of the intercluster LIFs of the clusters that you want to peer.
- You must know the passphrase that you will use for each peer relationship.

About this task


You must perform this procedure on each cluster.

Steps

1. From the source cluster, create a cluster peer relationship with the destination cluster.
 - a. Click the **Configurations** tab.
 - b. In the **Cluster Settings** pane, click **Cluster Peers**.
 - c. Click **Create**.
The Create Cluster Peer dialog box is displayed.
 - d. In the **Details of the remote cluster to be peered** area, specify the passphrase that both peers will use to ensure an authenticated cluster peer relationship.
 - e. Enter the IP addresses of all of the intercluster LIFs of the destination cluster (one per node) separated by commas.

- f. Click **Create**.
The authentication status is **pending** because only one cluster has been configured.
2. Switch to the destination cluster, and then create a cluster peer relationship with the source cluster:
 - a. Click the **Configurations** tab.
 - b. In the **Cluster Settings** pane, click **Cluster Peers**.
 - c. Click **Create**.
The Create Cluster Peer dialog box is displayed.

- d. In the **Details of the remote cluster to be peered** area, specify the same passphrase that you specified in step 1 and the IP addresses of the intercluster LIFs of the source cluster, and then click **Create**.



3. From the **Cluster Peers** window of the destination cluster, confirm that the source cluster is **available** and that the authentication status is **ok**.



You might have to click **Refresh** to view the updated information.

The two clusters are in a peer relationship.

4. Switch to the source cluster, and confirm that the destination cluster is **available** and that the authentication status is **ok**.

You might have to click **Refresh** to view the updated information.

After you finish

Create an SVM peer relationship between the source and destination SVMs while creating a data protection relationship between the source volume and the destination volume.

[Volume express backup using SnapVault](#)

[Volume disaster recovery express preparation](#)

Where to find additional information

After you successfully create a cluster peering relationship, you are ready to protect the availability of your data. There are express guides to help you configure data protection, as well as additional guides to do advanced configuration of cluster peering.

Express guides

You can protect your data by using the following express guides:

- [Volume disaster recovery express preparation](#)
Describes how to quickly configure and monitor the SnapMirror relationships between volumes in different clusters.
- [Volume express backup using SnapVault](#)
Describes how to quickly configure an intercluster SnapVault relationship.

Other guides

More information is available in the following guides:

- [Data protection](#)
Describes how to prevent data loss using Snapshot copies and SnapMirror replication to a remote system
- [NetApp Technical Report 4015: SnapMirror Configuration and Best Practices Guide for ONTAP 9.1, 9.2](#)
Describes information and best practices about configuring replication, including cluster peering.
- [ONTAP concepts](#)
Provides conceptual information about cluster peering.

Copyright

Copyright © 2019 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, 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).

Trademark

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., 1395 Crossman Ave., Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277