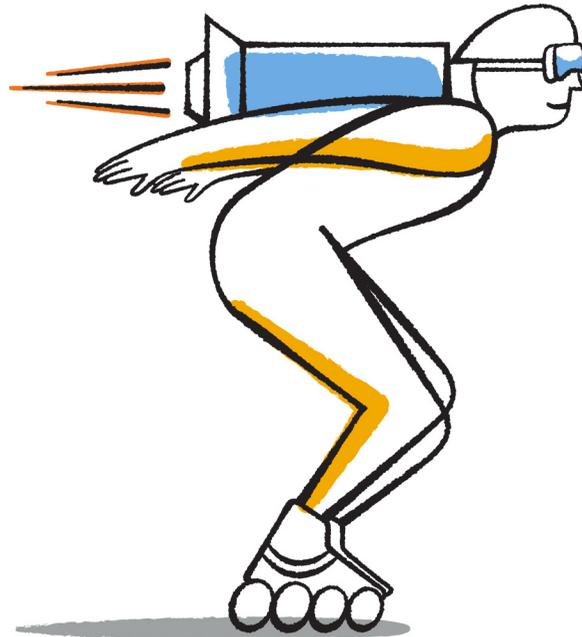




**NetApp®**

## Clustered Data ONTAP® 8.3

### Cluster Peering Express Guide



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  
Web: [www.netapp.com](http://www.netapp.com)  
Feedback: [doccomments@netapp.com](mailto:doccomments@netapp.com)

Part number: 215-09062\_A0  
November 2014



# Contents

|  |           |
|--|-----------|
| <b>Deciding whether to use this guide .....</b>        | <b>4</b>  |
| <b>Prerequisites for cluster peering .....</b>         | <b>5</b>  |
| <b>Cluster peering workflow .....</b>                  | <b>7</b>  |
| Preparing for cluster peering .....                    | 7         |
| Creating intercluster interfaces on all nodes .....    | 9         |
| Creating peer relationships among three clusters ..... | 11        |
| <b>Where to find additional information .....</b>      | <b>14</b> |
| <b>Copyright information .....</b>                     | <b>15</b> |
| <b>Trademark information .....</b>                     | <b>16</b> |
| <b>How to send your comments .....</b>                 | <b>17</b> |
| <b>Index .....</b>                                     | <b>18</b> |

## Deciding whether to use this guide

---

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

By following this guide, you will create three peer relationships among three clusters named Cluster A, Cluster B, and Cluster C. If you want only a single peer relationship between two clusters, you can use this guide and stop after creating the first cluster peer relationship. If you want a different intercluster configuration, you can adjust the workflow as required.

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

- You are working with clusters running Data ONTAP 8.3 or later.
- 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 OnCommand System Manager, not the command-line interface or an automated scripting tool.

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

- [\*Clustered Data ONTAP 8.3 System Administration Guide for Cluster Administrators\*](#)  
Describes how to use the command-line interface to set up various types of cluster peering relationships, such as unauthenticated relationships and relationships with clusters running earlier versions of Data ONTAP.
- [\*Clustered Data ONTAP 8.3 Network Management Guide\*](#)  
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.

# Prerequisites for cluster peering

---

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

## Connectivity requirements

The subnet used in each cluster for intercluster communication must meet the following requirements:

- The subnet must belong to the **Default** IPspace.
- The subnet must belong to the broadcast domain that contains the ports used for intercluster communication.
- The subnet must contain all of the IP addresses used for intercluster LIFs.
- You must have considered whether the subnet will be dedicated to intercluster communication or shared with data communication.
- 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.

The intercluster network must be configured so that cluster peers have *pair-wise full-mesh connectivity*, which means that each pair of clusters in a cluster peer relationship have connectivity among all of their intercluster LIFs.

A cluster's intercluster LIFs must use the same IP address version: all IPv4 addresses or all IPv6 addresses. Similarly, all the intercluster LIFs of the peered clusters must use the same IP addressing version.

## Port requirements

The ports that will be used for intercluster communication must meet the following requirements:

- All the ports must be in the **Default** 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. The ports added to a broadcast domain can be physical network ports, VLANs, or interface groups (ifgrps).
- All the ports must be cabled.
- All the ports must be in a healthy state.

## 6 | Cluster Peering Express Guide

- You must have considered whether the ports used for intercluster communication will be shared with data communication.  
If you want to dedicate ports to intercluster communication, you can create a broadcast domain specifically for intercluster communication.

### Firewall requirements

Firewalls and the intercluster firewall policy must allow the following:

- ICMP service
- TCP to the IP addresses of all of the intercluster LIFs over all of the following ports: 10000, 11104, and 11105
- HTTPS  
Although HTTPS is not required when you set up cluster peering, HTTPS is required later if you use OnCommand System Manager to configure data protection. However, if you use the command-line interface to configure data protection, HTTPS is not required to configure cluster peering or data protection.

The default **intercluster** firewall policy allows access through the HTTPS protocol and from all IP addresses (0.0.0.0/0), but the policy can be altered or replaced.

### Cluster requirements

Clusters must meet the following requirements:

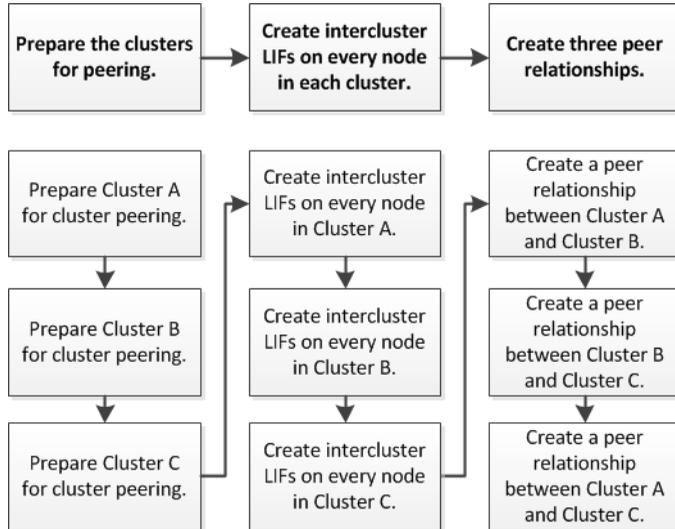
- Each cluster must have a unique name.  
You cannot create a cluster peering relationship with any cluster that has the same name or is in a peer relationship with a cluster of the same name.
- The time on the clusters in a cluster peering relationship must be synchronized within 300 seconds (5 minutes).  
Cluster peers can be in different time zones.
- Each cluster cannot be in a peer relationship with more than 255 clusters.

### Related information

[\*Clustered Data ONTAP 8.3 Data Protection Guide\*](#)

# Cluster peering workflow

Setting up cluster peering typically involves at least three clusters and therefore three peer relationships. After you prepare each cluster and create intercluster logical interfaces (LIFs) on each node, you can set up each peer relationship.



## Preparing for cluster peering

Before setting up cluster peering, you must verify that each cluster's time is synchronized with an external Network Time Protocol (NTP) server and you must plan subnets, ports, and passphrases.

### Before you begin

The clusters must be running Data ONTAP 8.3 or later.

### Steps

1. Decide on the passphrase you will use for each cluster peer relationship.

The passphrase must be at least eight characters and cannot contain spaces. You should use a unique passphrase for each relationship.

| For the relationship between... | The passphrase will be... |
|---------------------------------|---------------------------|
| Cluster A and Cluster B         |                           |
| Cluster B and Cluster C         |                           |
| Cluster A and Cluster C         |                           |

2. On each cluster, verify that the time is synchronized to an external NTP server to ensure that the time on all of the cluster peers is within five minutes of each other:
  - a. 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.
  - b. Expand the **Cluster** hierarchy in the left navigation pane.
  - c. In the navigation pane, click **Configuration > System Tools > DateTime**.
  - d. Confirm that the **Time servers** area contains at least one IP address of an external NTP server.



- e. If the **Time servers** area is blank, click **Edit** and add an NTP server.
3. Gather information about which subnet and ports you will use for intercluster LIFs, and optionally which IP address you will use.

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

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

|                       | Cluster A | Cluster B | Cluster C |
|-----------------------|-----------|-----------|-----------|
| Subnet                |           |           |           |
| IP address (optional) |           |           |           |
| Node 1 port           |           |           |           |
| Node 2 port           |           |           |           |
| Node 3 port           |           |           |           |
| Node 4 port           |           |           |           |

## Creating intercluster interfaces on all nodes

Clusters communicate with each other through logical interfaces (LIFs) that are dedicated to intercluster communication. You must create an intercluster LIF on each node in each cluster that will be in a cluster peer relationship.

### Before you begin

You should know which subnet and ports, and optionally which IP addresses, that you plan to use for the intercluster LIFs.

### About this task

This procedure assumes you are creating peer relationships among three clusters named A, B, and C.

### Steps

1. On Cluster A, navigate to the **Network Interfaces** window:
  - a. Expand the **Cluster** hierarchy in the left navigation pane.
  - b. Click **Configuration > Network**.
  - c. In the **Network** window, select the **Network Interfaces** tab.
2. Create an intercluster LIF on one node in the cluster:
  - a. Click **Create**.
  - b. In the **Create Network Interface** dialog box, specify 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.

- c. Select **Intercluster Connectivity** as the interface role.
- d. Select the subnet that will be used for intercluster communication on this cluster.

By default, an available IP address in that subnet is automatically selected after you click **Create**.
- e. If you want to control which IP address the node uses for intercluster communication, select **Use this IP Address** and type the address.
- f. In the **Ports** area, expand the specific node you are configuring, and select the port that you decided earlier to use for this node.

## 10 | Cluster Peering Express Guide

- g. If you decided earlier not to share ports with data communication, confirm that the selected port displays “0” in the **Hosted Interface Count** column.

Specify the following details to add a new network interface for data and management access of the chosen SVM.

Name:

Interface Role:  Serves Data  Intercluster Connectivity

SVM:

Protocol Access:  CIFS  iSCSI  NFS  FC/FCoE

Management Access:  Enable Management Access

Subnet:

The IP address is selected from this subnet.  
 Use this IP Address:

**i** This IP address will be added to the chosen subnet if the address is not already present in the subnet available range.

Port:

| Ports or Adapters | Hosted Interface Count | Speed     |
|-------------------|------------------------|-----------|
| clusterA-node 1   |                        |           |
| e0c               | 3                      | 1000 Mbps |
| e0d               | 0                      | 1000 Mbps |
| e0e               | 0                      | 1000 Mbps |

h. Click **Create**.

3. Repeat the previous step for each node in the cluster.

Each node in Cluster A has an intercluster LIF.

4. Record the IP addresses of the intercluster LIFs so that you can use them later when you create peer relationships with other clusters:

- a. On the **Network Interfaces** tab, in the Role column, click , clear **All**, and select **Intercluster**.

The window displays only intercluster LIFs.

- b. Record the IP addresses listed in the **IP Addresses** column, or leave the display open so that you can locate the IP addresses later.

| Network   |                 |              |  |
|---|-----------------|--------------|--|
| <span>Subnets</span>   <b>Network Interfaces</b>   Ethernet Ports   Broadcast Domains   FC/   |                 |              |  |
| <span>Create</span>   <span>Edit</span>   <span>Delete</span>   <span>Status</span>   <span>Migrate</span>   <span>Send to Home</span>   <span>Refresh</span> |                 |              |  |
| Interface Name  | IP Address/WWPN | Role         |  |
| icl01   | 10.53.32.1      | Intercluster |  |
| icl02   | 10.53.32.2      | Intercluster |  |

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

5. Repeat the entire procedure for Cluster B.  
Each node in Cluster B has an intercluster LIF.
6. Repeat the entire procedure for Cluster C.  
Each node in Cluster C has an intercluster LIF.

### Result

All nodes in Cluster A, Cluster B, and Cluster C have intercluster LIFs that can all communicate with each other.

## Creating peer relationships among three clusters

You can create a cluster peer relationship between two clusters by entering an agreed-upon passphrase and the IP addresses of the other cluster's intercluster LIFs, and then verifying that the relationship was created successfully. You must perform this procedure on each of the three clusters.

### Before you begin

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

### About this task

This procedure assumes you are creating peer relationships among three clusters named A, B, and C.

### Steps

1. On Cluster A, create a cluster peer relationship with Cluster B.
  - a. Expand the **Cluster** hierarchy in the left navigation pane.
  - b. Click **Configuration > Peers**.

## 12 | Cluster Peering Express Guide

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 Cluster B's intercluster LIFs—one per node—separated by commas.

Create Cluster Peer

For a cluster to communicate with another cluster in a peer relationship, enter a passphrase and the intercluster IP addresses of the peer cluster.  
[Tell me more about cluster peering](#)

**Details of the local cluster**

Cluster Name: clusterA

Intercluster IP Addresses:

|                |              |
|----------------|--------------|
| clusterA-node1 | 10.53.52.120 |
| clusterA-node2 | 10.53.52.121 |

**Details of the remote cluster to be peered**

Passphrase: [redacted]

Intercluster IP Addresses: 10.238.14.33,10.238.14.36

f. Click **Create**.

The authentication status is *pending* because only one cluster has been configured.

2. Switch to Cluster B and create a cluster peer relationship with Cluster A:

- a. Expand the **Cluster** hierarchy in the left navigation pane.
- b. Click **Configuration > Peers**.
- c. Click **Create**.
- d. In the **Create Cluster Peer** dialog box, specify the same passphrase and the IP addresses of Cluster A's intercluster LIFs, and click **Create**.

Create Cluster Peer

For a cluster to communicate with another cluster in a peer relationship, enter a passphrase and the intercluster IP addresses of the peer cluster.  
[Tell me more about cluster peering](#)

**Details of the local cluster**

Cluster Name: clusterB

Intercluster IP Addresses:

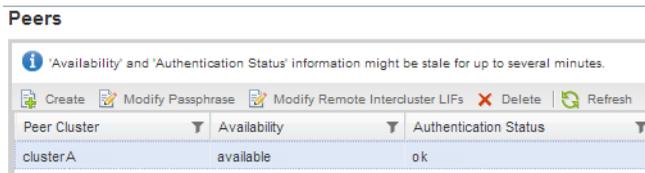
|                |              |
|----------------|--------------|
| clusterB-node1 | 10.238.14.33 |
| clusterB-node2 | 10.238.14.36 |

**Details of the remote cluster to be peered**

Passphrase: [redacted]

Intercluster IP Addresses: 10.53.52.120,10.53.52.121

3. In the **Peers** window of Cluster B, confirm that Cluster A is `available` and its **Authentication Status** is `ok`.



You might need to click **Refresh** to see updated information.

4. Switch to Cluster A, and confirm that Cluster B is `available` and its **Authentication Status** is `ok`.

You might need to click **Refresh** to see updated information.

Cluster A and Cluster B are in a peer relationship.

5. Create a peer relationship between Cluster B and Cluster C:
  - a. On Cluster B, create a cluster peer relationship with Cluster C by entering a passphrase and the IP addresses of Cluster C's intercluster LIFs in the **Create Cluster Peer** dialog box.
  - b. On Cluster C, create a cluster peer relationship with Cluster B by entering the same passphrase and the IP addresses of Cluster B's intercluster LIFs in the **Create Cluster Peer** dialog box.
  - c. On Cluster C, confirm that Cluster B is `available` and its **Authentication Status** is `ok`.
  - d. On Cluster B, confirm that Cluster C is `available` and its **Authentication Status** is `ok`.

Cluster B and Cluster C are in a peer relationship.

6. Create a peer relationship between Cluster A and Cluster C:
  - a. On Cluster A, create a cluster peer relationship with Cluster C by entering a passphrase and the IP addresses of Cluster C's intercluster LIFs in the **Create Cluster Peer** dialog box.
  - b. On Cluster C, create a cluster peer relationship with Cluster A by entering the same passphrase and the IP addresses of Cluster A's intercluster LIFs in the **Create Cluster Peer** dialog box.
  - c. On Cluster C, confirm that Cluster A is `available` and its **Authentication Status** is `ok`.
  - d. On Cluster A, confirm that Cluster C is `available` and its **Authentication Status** is `ok`.

Cluster A and Cluster C are in a peer relationship.

## 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:

- [\*Clustered Data ONTAP 8.3 Volume Disaster Recovery Preparation Express Guide\*](#)  
Describes how to quickly configure and monitor the SnapMirror relationships between volumes in different clusters. Includes instructions for creating Storage Virtual Machine (SVM) peer relationships.
- [\*Clustered Data ONTAP 8.3 Volume Backup Using SnapVault Express Guide\*](#)  
Describes how to quickly configure an intercluster SnapVault relationship. Includes instructions for creating SVM peer relationships.

### Other guides

More information is available in the following guides:

- [\*Clustered Data ONTAP 8.3 System Administration Guide for Cluster Administrators\*](#)  
Provides detailed conceptual and task information about modifying or removing the cluster peering configuration.
- [\*NetApp Technical Report 4015: SnapMirror Configuration and Best Practices Guide for Clustered Data ONTAP\*](#)  
Describes information and best practices about configuring replication, including cluster peering.

## Copyright information

---

Copyright © 1994–2014 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, Go Further, Faster, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

## How to send your comments

---

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to [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.

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

## B

- backup vaults
  - where to get additional information about [14](#)

## C

- cascade configurations
  - networking requirements for cluster peering [5](#)
- cluster peer relationships
  - creating [11](#)
  - preparing for [7](#)
  - requirements for [5](#)
  - requirements for using the Express Guide to configure [4](#)
  - where to get additional information about [14](#)
  - workflow for configuring [7](#)

## clusters

- creating peer relationships [11](#)
- naming requirements for cluster peering [5](#)
- peering workflow flowchart [7](#)
- requirements for using the Express Guide to configure peering [4](#)

## comments

- how to send feedback about documentation [17](#)

## communication

- intercluster, creating interfaces for [9](#)

## configurations

- cascade, for cluster peering [5](#)
- fan-out, for cluster peering [5](#)

## creating

- cluster peer relationships [11](#)
- LIFs [9](#)

## D

### data protection

- where to get additional information about [14](#)

### documentation

- how to send feedback about [17](#)

## E

### express guides

- requirements for using the cluster peering guide [4](#)

## F

### fan-out configurations

- networking requirements for cluster peering [5](#)

### feedback

- how to send comments about documentation [17](#)

### firewalls

- requirements for cluster peering [5](#)

### flowcharts

- cluster peering workflow [7](#)

### full-mesh connectivity

- description [5](#)

## G

### gathering

- network information [7](#)
- passphrases [7](#)

## I

### information

- how to send feedback about improving documentation [17](#)

### intercluster

- communication, creating interfaces [9](#)
- LIFs, creating [9](#)

### interfaces

- See* LIFs

### IP addresses

- requirements for cluster peering [5](#)
- specifying for intercluster LIFs [9](#)

### IPspaces

- requirements for cluster peering [5](#)

## L

### LIFs

- creating [9](#)

## N

### network

- full-mesh connectivity described [5](#)
- requirements for cluster peering [5](#)

network interfaces

*See* LIFs

## P

passphrases

planning [7](#)

peer relationships

cluster, workflow flowchart [7](#)

creating cluster [11](#)

preparing for cluster [7](#)

requirements for clusters [5](#)

requirements for using the Express Guide to

configure for clusters [4](#)

where to get additional information about cluster and SVM [14](#)

ports

intercluster communication, specifying [9](#)

requirements for cluster peering [5](#)

preparing

cluster peer relationships [7](#)

protection

where to get additional information about data [14](#)

## R

requirements

cluster naming when peering [5](#)

firewall for cluster peering [5](#)

IP addresses for cluster peering [5](#)

IPspaces for cluster peering [5](#)

network for cluster peering [5](#)

ports for cluster peering [5](#)

subnets for cluster peering [5](#)

## S

subnets

planning [7](#)

requirements for cluster peering [5](#)

specifying for intercluster communication [9](#)

suggestions

how to send feedback about documentation [17](#)

SVM peer relationships

where to get additional information about [14](#)

## T

time synchronization

synchronizing across clusters [7](#)

## V

vaults

where to get additional information about backup [14](#)

## W

workflows

cluster peering, flowchart [7](#)