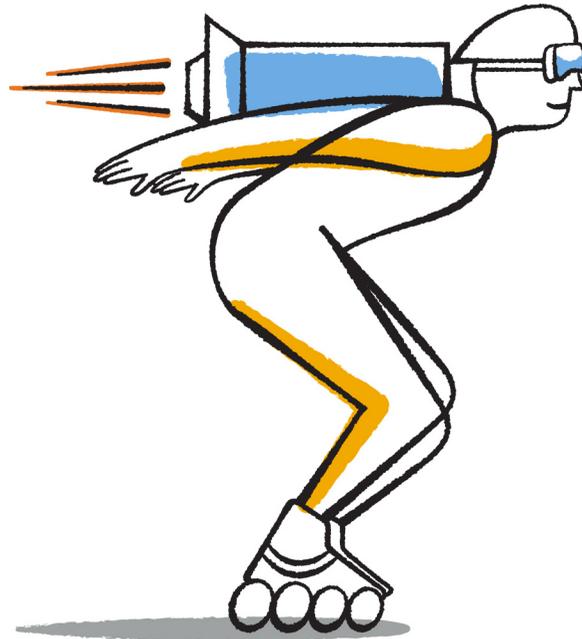




NetApp®

Clustered Data ONTAP® 8.2

Cluster and Vserver 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
Feedback: doccomments@netapp.com

Part number: 215-07997_B0
July 2014

Contents

Deciding whether to use this guide	4
Cluster peering workflow	5
Intercluster Vserver peering workflow	6
Considerations for configuring intercluster LIFs	7
Determining whether to use shared or dedicated ports for intercluster communication	8
Required network information	9
Connecting clusters in a peer relationship	14
Configuring intercluster LIFs to use dedicated intercluster ports	14
Configuring intercluster LIFs to share data ports	19
Creating the cluster peer relationship	23
Connecting Vservers in a peer relationship	25
Creating a Vserver peer relationship	25
Accepting a Vserver peer relationship	26
Where to find additional information	28
Copyright information	29
Trademark information	30
How to send your comments	31
Index	32

Deciding whether to use this guide

This guide describes how to create peer relationships between two clusters and between two Vservers. Peer relationships enable clusters to communicate with other clusters and Vservers to communicate with other Vservers. You must create intercluster peer relationships to replicate data between different geographical locations or between Vservers on different clusters.

You should use this guide if you want to create peer relationships and do not want a lot of conceptual background for the tasks. This guide does not cover every option and provides minimal background information.

This guide is written with the following assumptions:

- You are a cluster administrator.
- The two clusters to be peered must have been set up and must be running any version of clustered Data ONTAP in the same release family.
For example, if one of the clusters is running clustered Data ONTAP 8.2.0, the other cluster can run Data ONTAP 8.2.1, 8.2.2, or later versions in the 8.2 release family.
- The time on the clusters must be synchronized within 300 seconds (five minutes).
Cluster peers can be in different time zones.

If these assumptions are not correct for your situation, or if you want more conceptual background information, you should see the following documentation:

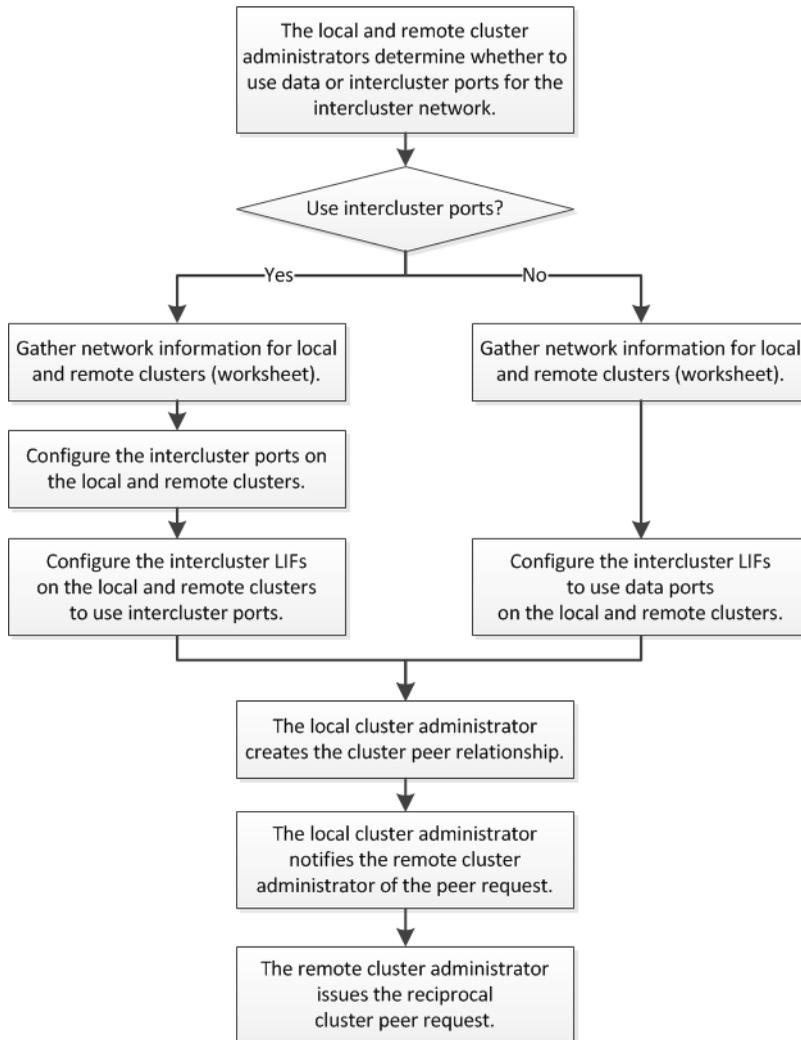
- *Clustered Data ONTAP System Administration Guide for Cluster Administrators*
- *Clustered Data ONTAP Network Management Guide*
- *SnapMirror Configuration and Best Practices Guide for Clustered Data ONTAP (TR-4015)*

Related information

[Documentation on the NetApp Support Site: mysupport.netapp.com](https://mysupport.netapp.com)

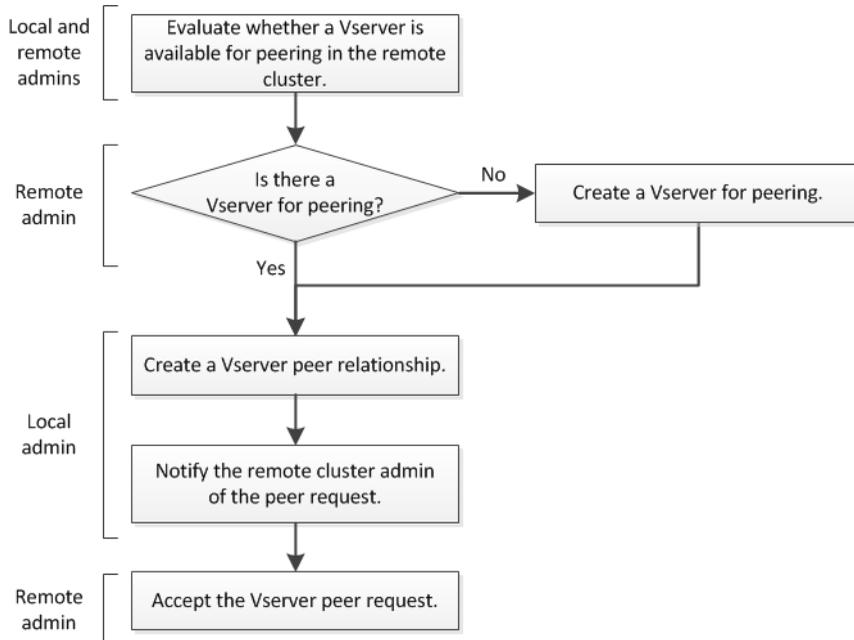
Cluster peering workflow

When configuring a cluster peer relationship, administrators of both the clusters determine whether to use dedicated ports for intercluster communication or to share ports that are also used by the data network, create intercluster ports (if necessary), assign intercluster LIFs to the selected ports, and then create the peer relationship.



Intercluster Vserver peering workflow

When you configure an intercluster peer relationship for a virtual storage server (Vserver), you ensure there is a suitable Vserver on the destination cluster before creating the peer relationship. When the administrator of the remote cluster accepts the peer request, the Vserver peer relationship is established.



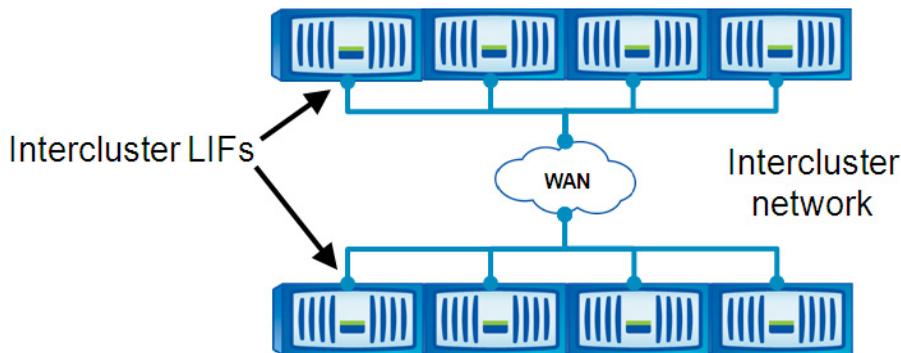
Considerations for configuring intercluster LIFs

The intercluster network uses logical interfaces, or LIFs, that correspond to IP addresses and represent network access points to a node. The intercluster network uses only intercluster LIFs. You assign intercluster LIFs to ports in conformance with requirements to create a supported peer configuration.

- At least one intercluster LIF must be configured on every node in the local cluster and on every node in the remote cluster.
Provisioning intercluster LIFs on only some nodes of the cluster is not a supported configuration.

Tip: To make it easier to identify the node where an intercluster LIF resides, it is a good idea to use a naming convention that indicates the node on which the intercluster LIF was created, followed by `ic#` or `icl#` to indicate the LIF type and number. For example, the second of two intercluster LIFs on the fourth node of cluster01 might be named `cluster01-04_ic02`. The first letter of a LIF name should be either a letter or an underscore.

- Each intercluster LIF requires an IP address dedicated for intercluster communication. The IP addresses assigned to intercluster LIFs can reside in the same subnet as data LIFs or in a different subnet.
- Your cluster peering topology should use *full-mesh connectivity*, meaning that every intercluster LIF on the local cluster should be able to communicate with every intercluster LIF on the remote cluster.



As intercluster LIFs become available or unavailable, the list of active IP addresses used by the cluster can change. The discovery of active IP addresses is automatic in certain events, such as when a node reboots. The creation of a peer relationship requires only one remote cluster address.

Note: If the node hosting the intercluster LIF address goes down and the address becomes unavailable, the cluster peer relationship might not be rediscovered. Therefore, it is a best practice to configure at least one intercluster IP address from each node in both clusters, so that the peer relationship remains stable if a node fails.

Determining whether to use shared or dedicated ports for intercluster communication

When you configure a cluster peer relationship, you assign intercluster LIFs to ports capable of handling intercluster communications. Before selecting the ports, you have to determine whether to use dedicated ports for intercluster communication or to share ports that are also used by the data network.

Network ports in a cluster have roles that define their purpose and their default behavior. Port roles limit the types of LIFs that can be assigned to a port. Intercluster LIFs use only data ports or intercluster ports.

You can assign any Ethernet port type in the system the role of data or intercluster, including physical ports (such as e0e or e0f) and logical ports (such as a VLAN or an interface group).

- A data port is used for data traffic and can be accessed by NFS, CIFS, FC, or iSCSI clients for data requests.
You can create VLANs and interface groups on data ports. VLANs and interface groups have the data port role by default.
- An intercluster port is used exclusively for communication between clusters.
- Each intercluster LIF on a cluster in a peer relationship should be routable to every intercluster LIF of the other cluster.

When you assign the intercluster role to a port, Data ONTAP automatically configures a non-modifiable failover group consisting of intercluster ports for the intercluster LIFs. Data protocols cannot fail over or migrate to an intercluster port.

If you used Data ONTAP operating in 7-Mode and used a dedicated port for replication traffic, it is likely that you will want to use a dedicated network for clustered Data ONTAP. Most of the networking considerations you encountered when using replication in 7-Mode also apply to clustered Data ONTAP.

LAN type and bandwidth

If you have a high-speed network, such as 10 GbE, you might have enough local LAN bandwidth to perform replication using the same 10 GbE ports used for data access.

If you have a 10 GbE network, you should compare your available WAN bandwidth to your LAN bandwidth. If the available WAN bandwidth is significantly less than 10 GbE, you might be limited to the network utilization that the WAN can support.

Data change rate and replication interval

If you plan to use the peer relationship for replication, consider how your available bandwidth will handle the level of client activity during the replication interval. If the WAN bandwidth is similar to

that of the LAN ports and replication will occur during regular client activity, then dedicate Ethernet ports for intercluster communication to avoid contention between replication and the data protocols.

Consider the amount of data that will be replicated in each interval and whether it requires so much bandwidth that it could cause contention with data protocols for shared data ports. If you use the peer relationship for replication and replication is set to occur only when minimal to no client activity occurs, you might be able to use data ports for intercluster replication successfully, even without a 10 GbE LAN connection.

Network utilization by data protocols

If the network utilization generated by the data protocols (CIFS, NFS, iSCSI) is above 50%, then you should dedicate ports for intercluster communication to allow for non-degraded performance if a node failover occurs.

Port availability

Sharing ports for data and intercluster communication eliminates the extra port counts required when using dedicated ports. When data ports for intercluster communication are shared, the intercluster LIFs can be migrated to any other intercluster-capable port on the same node.

Note: If you decide to dedicate ports for intercluster communication, it is a best practice to configure at least two intercluster ports per node. An intercluster LIF cannot fail over to a port on a different node; its failover group contains only intercluster-capable ports on the same node. If you use intercluster ports, Data ONTAP uses only intercluster ports in the failover group for an intercluster LIF. Therefore, if you use intercluster ports, you should configure at least two intercluster ports per node so that there is a port to which the intercluster LIF can fail over.

When physical 10 GbE ports are used for both data and intercluster communication, you can create VLAN ports and use dedicated logical ports for intercluster communication.

Using dedicated ports for intercluster communication requires additional switch ports and cable runs.

Performance characteristics

For best performance, all paths used by intercluster LIFs should have equal performance characteristics. If a node has one intercluster LIF on a slow path and another intercluster LIF on a fast path, performance will be adversely affected because data is multiplexed across the slow and fast paths simultaneously.

Required network information

The administrators of the local and remote clusters both have to gather IP addresses, ports, and other network information for the intercluster LIFs before configuring a cluster peer relationship.

Configure at least one intercluster IP address from each node in both clusters, so that the peer relationship remains stable if a node fails.

If you plan to use intercluster ports, configure at least two intercluster ports per node.

Intercluster LIFs for the local cluster

You must configure at least one intercluster LIF for each node in the cluster. Each intercluster LIF requires an IP address dedicated for intercluster communication.

Type of information		Your values
Node 1	Intercluster LIF name	
	Intercluster LIF port(s)	
	Intercluster LIF IP address	
	Default route gateway	
	Network mask	
Node 2	Intercluster LIF name	
	Intercluster LIF port(s)	
	Intercluster LIF IP address	
	Default route gateway	
	Network mask	
Node 3	Intercluster LIF name	
	Intercluster LIF port(s)	
	Intercluster LIF IP address	
	Default route gateway	
	Network mask	
Node 4	Intercluster LIF name	
	Intercluster LIF port(s)	
	Intercluster LIF IP address	
	Default route gateway	
	Network mask	

Type of information		Your values
Node 5	Intercluster LIF name	
	Intercluster LIF port(s)	
	Intercluster LIF IP address	
	Default route gateway	
	Network mask	
Node 6	Intercluster LIF name	
	Intercluster LIF port(s)	
	Intercluster LIF IP address	
	Default route gateway	
	Network mask	
Node 7	Intercluster LIF name	
	Intercluster LIF port(s)	
	Intercluster LIF IP address	
	Default route gateway	
	Network mask	
Node 8	Intercluster LIF name	
	Intercluster LIF port(s)	
	Intercluster LIF IP address	
	Default route gateway	
	Network mask	

Intercluster LIFs for the remote cluster

You must configure at least one intercluster LIF for each node in the cluster. Each intercluster LIF requires an IP address dedicated for intercluster communication.

12 | Cluster and Vserver Peering Express Guide

Type of information		Your values
Node 1	Intercluster LIF name	
	Intercluster LIF port(s)	
	Intercluster LIF IP address	
	Default route gateway	
	Network mask	
Node 2	Intercluster LIF name	
	Intercluster LIF port(s)	
	Intercluster LIF IP address	
	Default route gateway	
	Network mask	
Node 3	Intercluster LIF name	
	Intercluster LIF port(s)	
	Intercluster LIF IP address	
	Default route gateway	
	Network mask	
Node 4	Intercluster LIF name	
	Intercluster LIF port(s)	
	Intercluster LIF IP address	
	Default route gateway	
	Network mask	
Node 5	Intercluster LIF name	
	Intercluster LIF port(s)	
	Intercluster LIF IP address	
	Default route gateway	
	Network mask	

Type of information		Your values
Node 6	Intercluster LIF name	
	Intercluster LIF port(s)	
	Intercluster LIF IP address	
	Default route gateway	
	Network mask	
Node 7	Intercluster LIF name	
	Intercluster LIF port(s)	
	Intercluster LIF IP address	
	Default route gateway	
	Network mask	
Node 8	Intercluster LIF name	
	Intercluster LIF port(s)	
	Intercluster LIF IP address	
	Default route gateway	
	Network mask	

Connecting clusters in a peer relationship

To connect clusters in a peer relationship, you configure intercluster LIFs by using either shared ports or dedicated ports, and then create the peer relationship.

Before you begin

You must have decided whether to use dedicated ports for intercluster communication or to assign your intercluster LIFs to shared data ports.

About this task

If you are not the administrator of the remote cluster, you must coordinate with the administrator to configure intercluster LIFs on the nodes of the remote cluster and to create the peer relationship. You and the administrator of the remote cluster can configure cluster peering without sharing credential information if you both issue the commands while logged into your respective clusters.

This guide provides two, mutually exclusive procedures for configuring intercluster LIFs: one for using dedicated ports and one for using shared ports. You have to perform one or the other before creating the cluster peer relationship.

Tip: Configuring cluster peering can be very repetitive, because you must configure intercluster LIFs on every node in both clusters to establish full-mesh connectivity. Consider using wildcards in the commands to save time. (For better clarity, the examples do not show the use of wildcards.)

Configuring intercluster LIFs to use dedicated intercluster ports

Configuring intercluster LIFs to use dedicated data ports allows greater bandwidth than using shared data ports on your intercluster networks for cluster peer relationships.

About this task

In this example, a two-node cluster exists in which each node has two data ports, e0e and e0f, which are dedicated for intercluster replication. In your own environment, you would replace the ports, networks, IP addresses, subnet masks, and subnets with those specific to your environment.

To learn more about LIFs and port types, see the *Clustered Data ONTAP Network Management Guide*

Steps

1. Check the role of the ports in the cluster by using the `network port show` command.

Example

```
cluster01::> network port show
```

Node	Port	Role	Link	MTU	Auto-Negot Admin/Oper	Duplex Admin/Oper	Speed(Mbps) Admin/Oper

cluster01-01							
	e0a	cluster	up	1500	true/true	full/full	auto/1000
	e0b	cluster	up	1500	true/true	full/full	auto/1000
	e0c	data	up	1500	true/true	full/full	auto/1000
	e0d	data	up	1500	true/true	full/full	auto/1000
	e0e	data	up	1500	true/true	full/full	auto/1000
	e0f	data	up	1500	true/true	full/full	auto/1000
cluster01-02							
	e0a	cluster	up	1500	true/true	full/full	auto/1000
	e0b	cluster	up	1500	true/true	full/full	auto/1000
	e0c	data	up	1500	true/true	full/full	auto/1000
	e0d	data	up	1500	true/true	full/full	auto/1000
	e0e	data	up	1500	true/true	full/full	auto/1000
	e0f	data	up	1500	true/true	full/full	auto/1000

- Determine whether any of the LIFs are using ports that are dedicated for replication by using the `network interface show` command.

Example

```
cluster01::> network interface show
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home

cluster01						
	cluster_mgmt	up/up	192.168.0.xxx/24	cluster01-01	e0c	true
vs1	vs1_lif1	up/up	192.168.0.151/24	cluster01-01	e0e	true

- If a LIF is using one of the ports dedicated to replication, then assign the LIF to a different home port by using the `network interface modify` command.

The LIF cannot remain on the port you want to dedicate to replication, because intercluster ports cannot host data LIFs.

The `network interface modify` operation is nondisruptive, because the LIF has not yet moved from port e0e. The `network interface modify` command shown below changes the port to which the LIF returns when the `network interface revert` command is issued.

Example

```
cluster01::> network interface modify -vserver vs1 -lif vs_lif1 -home-node cluster01-01 -
home-port e0d
```

```
cluster01::> network interface show
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home

cluster01						
	cluster_mgmt	up/up	192.168.0.xxx/24	cluster01-01	e0c	true
vs1	vs1_lif1	up/up	192.168.0.151/24	cluster01-01	e0e	false

4. Revert the LIF to its new home port by using the `network interface revert` command. Assigning the LIF to a different port by combining the `network interface modify` and `network interface revert` commands avoids the risk that the LIF might fail back to its original port.

Example

```
cluster01::> network interface revert -vserver vs1 -lif vs_lif1

cluster01::> network interface show
Vserver      Logical   Status   Network      Current   Current   Is
Interface    Interface Admin/Oper Address/Mask Node       Port      Home
-----
cluster01
vs1          vs1_lif1 up/up    192.168.0.151/24 cluster01-01 e0d      true
```

5. After all LIFs have been migrated off the ports dedicated for replication, change the role of the port used on each node to `intercluster` by using the `network port modify` command.

Example

```
cluster01::> network port modify -node cluster01-01 -port e0e -role intercluster
cluster01::> network port modify -node cluster01-01 -port e0f -role intercluster
cluster01::> network port modify -node cluster01-02 -port e0e -role intercluster
cluster01::> network port modify -node cluster01-02 -port e0f -role intercluster
```

6. Verify that the roles of the correct ports have been changed to `intercluster` by using the `network port show` command with the `-role intercluster` parameter.

Example

```
cluster01::> network port show -role intercluster
Node  Port  Role           Link MTU   Auto-Negot  Duplex      Speed(Mbps)
      Port  Admin/Oper    Admin/Oper Admin/Oper  Admin/Oper
-----
cluster01-01
e0e   intercluster up    1500 true/true   full/full   auto/1000
e0f   intercluster up    1500 true/true   full/full   auto/1000
cluster01-02
e0e   intercluster up    1500 true/true   full/full   auto/1000
e0f   intercluster up    1500 true/true   full/full   auto/1000
```

7. Create an intercluster LIF on each node in cluster01 by using the `network interface create` command.

Example

This example uses the LIF naming convention `nodename_icl#` for the intercluster LIF.

```
cluster01::> network interface create -vserver cluster01-01 -lif cluster01-01_icl01 -role
intercluster -home-node cluster01-01 -home-port e0e
```

```
-address 192.168.1.201 -netmask 255.255.255.0
Info: Your interface was created successfully; the routing group i192.168.1.0/24 was
created

cluster01::> network interface create -vserver cluster01-02 -lif cluster01-02_icl01 -role
intercluster -home-node cluster01-02 -home-port e0e
-address 192.168.1.202 -netmask 255.255.255.0
Info: Your interface was created successfully; the routing group i192.168.1.0/24 was
created
```

- Verify that the intercluster LIFs are configured for redundancy by using the `network interface show` command with the `-role intercluster` and `-failover` parameters.

Example

The LIFs in this example are assigned the e0e home port on each node. If the e0e port fails, the LIF can fail over to the e0f port because e0f is also assigned the role of intercluster.

The intercluster LIF is assigned to an intercluster port; therefore, a non-modifiable failover group is created automatically, and contains all ports with the intercluster role on that node. Intercluster failover groups are node specific; therefore, if changes are required, they must be managed for each node because different nodes might use different ports for replication.

```
cluster01::> network interface show -role intercluster -failover
-----
Vserver      Logical      Home      Failover      Failover
Interface    Node:Port    Node:Port    Group Usage    Group
-----
cluster01-01
  cluster01-01_icl01  cluster01-01:e0e  system-defined
                                Failover Targets: cluster01-01:e0e,
                                cluster01-01:e0f
cluster01-02
  cluster01-02_icl01  cluster01-02:e0e  system-defined
                                Failover Targets: cluster01-02:e0e,
                                cluster01-02:e0f
```

- Verify that the intercluster LIFs were created properly by using the `network interface show` command.

Example

```
cluster01::> network interface show
-----
Vserver      Logical      Status      Network      Current      Current      Is
Interface    Admin/Oper   Address/Mask Node          Port         Home
-----
cluster01
  cluster_mgmt up/up        192.168.0.xxx/24  cluster01-01  e0c         true
cluster01-01
  cluster01-01_icl01 up/up        192.168.1.201/24  cluster01-01  e0e         true
  clus1        up/up        169.254.xx.xx/24  cluster01-01  e0a         true
  clus2        up/up        169.254.xx.xx/24  cluster01-01  e0b         true
  mgmt1       up/up        192.168.0.xxx/24  cluster01-01  e0c         true
cluster01-02
  cluster01-02_icl01 up/up        192.168.1.202/24  cluster01-02  e0e         true
  clus1        up/up        169.254.xx.xx/24  cluster01-02  e0a         true
  clus2        up/up        169.254.xx.xx/24  cluster01-02  e0b         true
  mgmt1       up/up        192.168.0.xxx/24  cluster01-02  e0c         true
```

- Display routing groups by using the `network routing-group show` command with the `-role intercluster` parameter to determine whether the intercluster network needs intercluster routes.

An intercluster routing group is created automatically for the intercluster LIFs.

Example

```
cluster01::> network routing-group show -role intercluster
```

Vserver	Routing Group	Subnet	Role	Metric
cluster01-01	i192.168.1.0/24	192.168.1.0/24	intercluster	40
cluster01-02	i192.168.1.0/24	192.168.1.0/24	intercluster	40

- Display the routes in the cluster by using the `network routing-group show` command to determine whether intercluster routes are available or you must create them.

Creating a route is required only if the intercluster addresses in both clusters are not on the same subnet and a specific route is needed for communication between the clusters.

Example

In this example, no intercluster routes are available.

```
cluster01::> network routing-group route show
```

Vserver	Routing Group	Destination	Gateway	Metric
cluster01	c192.168.0.0/24	0.0.0.0/0	192.168.0.1	20
cluster01-01	n192.168.0.0/24	0.0.0.0/0	192.168.0.1	10
cluster01-02	n192.168.0.0/24	0.0.0.0/0	192.168.0.1	10

- If communication between intercluster LIFs in different clusters requires routing, create an intercluster route by using the `network routing-groups route create` command.

The intercluster networks apply to each node; therefore, you must create an intercluster route on each node.

Example

In this example, 192.168.1.1 is the gateway address for the 192.168.1.0/24 network.

Note: If the destination is specified as 0.0.0.0/0, then it becomes the default route for the intercluster network.

```
cluster01::> network routing-groups route create -server cluster01-01 -routing-
group i192.168.1.0/24 -destination 0.0.0.0/0 -gateway 192.168.1.1 -metric 40

cluster01::> network routing-groups route create -server cluster01-02 -routing-
group i192.168.1.0/24 -destination 0.0.0.0/0 -gateway 192.168.1.1 -metric 40
```

13. Display the newly created routes by using the `network routing-groups route show` command to confirm that you created the routes correctly.

Although the intercluster routes do not have an assigned role, they are assigned to the routing group `i192.168.1.0/24`, which is assigned the role of `intercluster`. These routes are only used for intercluster communication.

Example

```
cluster01::> network routing-group route show
```

Vserver	Routing Group	Destination	Gateway	Metric
cluster01				
	c192.168.0.0/24	0.0.0.0/0	192.168.0.1	20
cluster01-01	n192.168.0.0/24	0.0.0.0/0	192.168.0.1	10
	i192.168.1.0/24	0.0.0.0/0	192.168.1.1	40
cluster01-02	n192.168.0.0/24	0.0.0.0/0	192.168.0.1	10
	i192.168.1.0/24	0.0.0.0/0	192.168.1.1	40

14. Repeat these steps to configure intercluster networking in the other cluster.
15. Verify that the ports have access to the proper subnets, VLANs, and so on.

Dedicating ports for replication in one cluster does not require dedicating ports in all clusters; one cluster might use dedicated ports, while the other cluster shares data ports for intercluster replication.

Configuring intercluster LIFs to share data ports

Configuring intercluster LIFs to share data ports enables you to use existing data ports to create intercluster networks for cluster peer relationships. Sharing data ports reduces the number of ports you might need for intercluster networking.

Before you begin

You should have reviewed the considerations for sharing data ports and determined that this is an appropriate intercluster networking configuration.

About this task

Creating intercluster LIFs that share data ports involves assigning LIFs to existing data ports and, possibly, creating an intercluster route. In this procedure, a two-node cluster exists in which each node has two data ports, e0c and e0d. These are the two data ports that are shared for intercluster replication. In your own environment, you replace the ports, networks, IP addresses, subnet masks, and subnets with those specific to your environment.

To learn more about LIFs and port types, see the *Clustered Data ONTAP Network Management Guide*

Steps

1. Check the role of the ports in the cluster by using the `network port show` command.

Example

```
cluster01::> network port show
```

Node	Port	Role	Link	MTU	Auto-Negot Admin/Oper	Duplex Admin/Oper	Speed(Mbps) Admin/Oper

cluster01-01	e0a	cluster	up	1500	true/true	full/full	auto/1000
	e0b	cluster	up	1500	true/true	full/full	auto/1000
	e0c	data	up	1500	true/true	full/full	auto/1000
	e0d	data	up	1500	true/true	full/full	auto/1000
cluster01-02	e0a	cluster	up	1500	true/true	full/full	auto/1000
	e0b	cluster	up	1500	true/true	full/full	auto/1000
	e0c	data	up	1500	true/true	full/full	auto/1000
	e0d	data	up	1500	true/true	full/full	auto/1000

2. Create a failover group on each node on both clusters for the data ports you want to share for intercluster communications by using the `network interface failover-groups create` command.

Example

The first of the following commands creates the failover group `failover_cluster01-01_icl01` and assigns port `e0c` to the failover group. The second command adds port `e0d` to the failover group.

```
cluster01::> network interface failover-groups create -failover-group
failover_cluster01-01_icl01 -node cluster01-01 -port e0c

cluster01::> network interface failover-groups create -failover-group
failover_cluster01-01_icl01 -node cluster01-01 -port e0d
```

3. Create an intercluster LIF on each node in cluster01 by using the `network interface create` command.

Example

The following commands create an intercluster LIF on each of the two nodes in cluster01 and assign the intercluster LIF to the failover group you created that contains the data ports.

```

cluster01::> network interface create -vserver cluster01-01 -lif cluster01-01_icl01 -role
intercluster -home-node cluster01-01 -home-port e0c -address 192.168.1.201 -netmask
255.255.255.0 -failover-group failover_cluster01-01_icl01 -failover-policy nextavail
Info: Your interface was created successfully; the routing group i192.168.1.0/24 was
created

cluster01::> network interface create -vserver cluster01-02 -lif cluster01-02_icl01 -role
intercluster -home-node cluster01-02 -home-port e0c -address 192.168.1.202 -netmask
255.255.255.0 -failover-group failover_cluster01-02_icl01 -failover-policy nextavail
Info: Your interface was created successfully; the routing group i192.168.1.0/24 was
created

```

- Verify that the intercluster LIFs were created properly by using the `network interface show` command with the `-role intercluster` parameter.

Example

```

cluster01::> network interface show -role intercluster

```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
cluster01-01	cluster01-01_icl01	up/up	192.168.1.201/24	cluster01-01	e0c	true
cluster01-02	cluster01-02_icl01	up/up	192.168.1.202/24	cluster01-02	e0c	true

- Verify that the intercluster LIFs are configured to be redundant by using the `network interface show` command with the `-role intercluster` and `-failover` parameters.

Example

The LIFs in this example are assigned the e0c port on each node. If the e0c port fails, the LIF can fail over to the e0d port.

```

cluster01::> network interface show -role intercluster -failover

```

Vserver	Logical Interface	Home Node:Port	Failover Group Usage	Failover Group
cluster01-01	cluster01-01_icl01	cluster01-01:e0c	nextavail	failover_cluster01-01_icl01
		Failover Targets:	cluster01-01:e0c, cluster01-01:e0d	
cluster01-02	cluster01-02_icl01	cluster01-02:e0c	nextavail	failover_cluster01-02_icl01
		Failover Targets:	cluster01-02:e0c, cluster01-02:e0d	

- Display routing groups by using the `network routing-group show` command with the `-role intercluster` parameter.

An intercluster routing group is created automatically for the intercluster LIFs.

Example

```

cluster01::> network routing-group show -role intercluster
Routing

```

Vserver	Group	Subnet	Role	Metric
cluster01-01	i192.168.1.0/24	192.168.1.0/24	intercluster	40
cluster01-02	i192.168.1.0/24	192.168.1.0/24	intercluster	40

7. Display the routes in the cluster by using the `network routing-group show` command to determine whether intercluster routes are available or you must create them.

Creating a route is required only if the intercluster addresses in both clusters are not on the same subnet and a specific route is needed for communication between the clusters.

Example

In this example, no intercluster routes are available.

```
cluster01::> network routing-group route show
Routing
Vserver  Group      Destination      Gateway      Metric
-----  -
cluster01
  c192.168.0.0/24
    0.0.0.0/0      192.168.0.1    20
cluster01-01
  n192.168.0.0/24
    0.0.0.0/0      192.168.0.1    10
cluster01-02
  n192.168.0.0/24
    0.0.0.0/0      192.168.0.1    10
```

8. If communication between intercluster LIFs in different clusters requires routing, create an intercluster route by using the `network routing-groups route create` command.

The intercluster networks apply to each node; therefore, you must create an intercluster route on each node.

Example

In this example, 192.168.1.1 is the gateway address for the 192.168.1.0/24 network.

Note: If the destination is specified as 0.0.0.0/0, then it becomes the default route for the intercluster network.

```
cluster01::> network routing-groups route create -server cluster01-01
-routing-group i192.168.1.0/24 -destination 0.0.0.0/0 -gateway 192.168.1.1 -metric 40
cluster01::> network routing-groups route create -server cluster01-02
-routing-group i192.168.1.0/24 -destination 0.0.0.0/0 -gateway 192.168.1.1 -metric 40
```

9. Display the newly created routes by using the `network routing-groups route show` command.

Although the intercluster routes do not have an assigned role, they are assigned to the routing group i192.168.1.0/24, which is assigned the role of intercluster. These routes are only used for intercluster communication.

Example

```
cluster01::> network routing-group route show
Routing
Vserver  Group      Destination      Gateway          Metric
-----  -
cluster01
  c192.168.0.0/24
                0.0.0.0/0       192.168.0.1     20
cluster01-01
  n192.168.0.0/24
                0.0.0.0/0       192.168.0.1     10
  i192.168.1.0/24
                0.0.0.0/0       192.168.1.1     40
cluster01-02
  n192.168.0.0/24
                0.0.0.0/0       192.168.0.1     10
  i192.168.1.0/24
                0.0.0.0/0       192.168.1.1     40
```

10. Repeat these steps on the cluster to which you want to connect.

Creating the cluster peer relationship

You create the cluster peer relationship using a set of intercluster designated logical interfaces to make information about one cluster available to the other cluster for use in cluster peering applications.

Before you begin

You should have the intercluster network configured.

Steps

1. Create the cluster peer relationship from the local cluster by using the `cluster peer create` command, and then notify the administrator of the remote cluster of your peer request.

Unilaterally creating a cluster peer relationship requires the login credentials of the remote cluster administrator. Alternatively, you can create a cluster peer relationship without exchanging credentials by having both cluster administrators issue the `cluster peer create` command from their respective clusters.

When the remote cluster administrator issues the reciprocal cluster peer request, Data ONTAP creates the peer relationship.

Example

In the following example, cluster01 is peered with a remote cluster named cluster02. Cluster02 is a two-node cluster that has one intercluster LIF per node. The IP addresses of the intercluster LIFs created in cluster02 are 192.168.2.203 and 192.168.2.204. These IP addresses are used to create the cluster peer relationship.

```
cluster01::> cluster peer create -peer-addr 192.168.2.203,192.168.2.204
```

When issuing the reciprocal `cluster peer create` command, the administrator of `cluster02` specifies the IP addresses of the intercluster LIFs in `cluster01`.

If DNS is configured to resolve host names for the intercluster IP addresses, you can use host names in the `-peer-addrns` option. It is not likely that intercluster IP addresses frequently change; however, using host names allows intercluster IP addresses to change without having to modify the cluster peer relationship.

2. Display the cluster peer relationship by using the `cluster peer show` command with the `-instance` parameter.

Example

```
cluster01::> cluster peer show -instance
Peer Cluster Name: cluster02
Remote Intercluster Addresses: 192.168.2.203,192.168.2.204
Availability: Available
Remote Cluster Name: cluster02
Active IP Addresses: 192.168.2.203,192.168.2.204
Cluster Serial Number: 1-80-000013
```

3. Preview the health of the cluster peer relationship by using the `cluster peer health show` command.

Example

```
cluster01::> cluster peer health show
Node      cluster-Name      Node-Name
         Ping-Status      RDB-Health Cluster-Health Avail...
-----
cluster01-01
  cluster02
    Data: interface_reachable
    ICMP: interface_reachable true      true      true
  cluster02-02
    Data: interface_reachable
    ICMP: interface_reachable true      true      true
cluster01-02
  cluster02
    Data: interface_reachable
    ICMP: interface_reachable true      true      true
  cluster02-02
    Data: interface_reachable
    ICMP: interface_reachable true      true      true
```

Connecting Vservers in a peer relationship

Cluster administrators can create peer relationships between two Vservers either existing within a cluster or in peered clusters. Vserver peer relationships provide an infrastructure for applications and tasks that require communication between the two peer Vservers.

One Vserver can be peered with multiple Vservers either within a cluster or across clusters.

Note: Vserver names must be unique. When creating a Vserver, use the fully qualified domain name (FQDN) of the Vserver or another convention that ensures unique Vserver names.

For better clarity, the examples in this guide show simple Vserver names, `vserverA` and `vserverB`.

For more information about Vserver peer relationships, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Creating a Vserver peer relationship

Cluster administrators can create an authorization infrastructure for running applications between Vservers by establishing a peer relationship between the Vservers.

Before you begin

- If you want to create an intercluster Vserver peer relationship, both clusters must be peered with each other.
- The admin state of the Vservers to be peered must not be `initializing` or `deleting`.
- The names of Vservers in the peered clusters must be unique across the two clusters.

About this task

When you create a Vserver peer relationship, you can specify the applications that will communicate over the peer relationship. In clustered Data ONTAP 8.2, only SnapMirror is supported as an application over the peer relationship. If you do not specify the application for the peer relationship as `snapmirror`, a Vserver administrator cannot set up SnapMirror relationships between the peered Vservers.

Steps

1. Use the `vserver peer create` command to create a Vserver peer relationship.

Example

```
cluster01::> vserver peer create -vserver vserverA -peer-vserver  
vserverB -applications snapmirror -peer-cluster cluster02
```

```
Info: [Job 43] 'vserver peer create' job queued
```

At this point, the state of the intercluster Vserver peer relationship is *initiated*. A Vserver peer relationship is not established until the cluster administrator of the peered cluster accepts the Vserver peer request.

2. Use the `vserver peer show-all` command to view the status and other details of the Vserver peer relationship.

Example

```
cluster01::> vserver peer show-all
```

Vserver	Peer Vserver	Peer State	Peer Cluster	Peering Applications
vserverA	vserverB	initiated	cluster02	snapmirror

```
cluster02::> vserver peer show-all
```

Vserver	Peer Vserver	Peer State	Peer Cluster	Peering Applications
vserverB	vserverA	pending	cluster01	snapmirror
vserverD	vserverC	peered	cluster02	snapmirror

For more information about these commands, see the man pages.

After you finish

If you have initiated an intercluster Vserver peer relationship, you must inform the cluster administrator of the remote cluster about the Vserver peer request. After the cluster administrator of the remote cluster accepts the Vserver peer request, the Vserver peer relationship is established.

Accepting a Vserver peer relationship

When a cluster administrator creates an intercluster Vserver peer relationship, the cluster administrator of the remote cluster can accept the Vserver peer request to establish the peer relationship.

Steps

1. Use the `vserver peer show` command to view the Vserver peer requests.

Example

```
cluster02::> vserver peer show
```

Vserver	Peer Vserver	Peer State
---------	--------------	------------

-----	-----	-----
vserverB	vserverA	pending

2. Use the `vserver peer accept` command to accept the Vserver peer request and establish the Vserver peer relationship.

Example

```
cluster02::> vserver peer accept -vserver vserverB -peer-vserver
vserverA

Info: [Job 46] 'vserver peer accept' job queued
```

The Vserver peer relationship is established and the state is `peered`.

3. Use the `vserver peer show` command on either of the peered clusters to view the state of the Vserver peer relationship.

Example

```
cluster02::> vserver peer show
      Peer
Vserver  Vserver  Peer
-----  -----  ---
vserverB  vserverA  peered
```

For more information about these commands, see the man pages.

Where to find additional information

There are additional documents to help you learn more about cluster and peering and other related subjects.

All of the following documentation is available from the NetApp Support Site:

<i>Technical Report 4015: SnapMirror Configuration and Best Practices Guide for Clustered Data ONTAP 8.2</i>	Provides information and best practices related to configuring replication in clustered Data ONTAP.
<i>Clustered Data ONTAP Data Protection Guide</i>	Describes how to manage your backup and recover data on clustered systems.
<i>Clustered Data ONTAP Logical Storage Management Guide</i>	Describes how to efficiently manage your logical storage resources on systems running clustered Data ONTAP, using volumes, FlexClone volumes, files and LUNs, FlexCache volumes, deduplication, compression, qtrees, and quotas.
<i>Clustered Data ONTAP Network Management Guide</i>	Describes how to connect your cluster to your Ethernet networks and how to manage logical interfaces (LIFs).
<i>Clustered Data ONTAP System Administration Guide for Cluster Administrators</i>	Describes general system administration for NetApp systems running clustered Data ONTAP.
NetApp Knowledgebase	(A database of articles) Provides troubleshooting information, FAQs, and instructions for a wide range of NetApp products and technologies.

Related information

[Documentation on the NetApp Support Site: mysupport.netapp.com](https://mysupport.netapp.com)

[Technical Report: SnapMirror Configuration and Best Practices Guide for Clustered Data ONTAP: media.netapp.com/documents/tr-4015.pdf](https://media.netapp.com/documents/tr-4015.pdf)

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, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Customer Fitness, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at www.ibm.com/legal/copytrade.shtml.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

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. 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**
 - best practices
 - networking [9](#)
- C**
 - cluster peering
 - where to get additional information about [28](#)
 - cluster peers
 - creating relationships between [23](#)
 - clusters
 - about connecting in peer relationships [14](#)
 - peering workflow flowchart [5](#)
 - requirements for using the Express Guide to configure peering [4](#)
 - configuring intercluster LIFs to use dedicated intercluster [14](#)
- D**
 - data ports
 - configuring intercluster LIFs to share [19](#)
 - dedicated ports
 - determining whether to use for intercluster communication [8](#)
- E**
 - express guides
 - requirements for using the cluster and Vserver peering guide [4](#)
- F**
 - flowcharts
 - cluster peering workflow [5](#)
 - intercluster Vserver peering workflow [6](#)
 - full-mesh connectivity
 - network requirements [7](#)
- I**
 - intercluster communication
 - determining ports for [8](#)
- V**
 - Vserver peering workflow flowchart [6](#)
- intercluster LIFs**
 - configuring to share data ports [19](#)
 - configuring to use dedicated intercluster ports [14](#)
 - defined [7](#)
 - requirements [7](#)
 - worksheets for configuration [9](#)
- intercluster networks**
 - configuring intercluster LIFs for [14](#), [19](#)
- intercluster ports**
 - configuring intercluster LIFs to use dedicated [14](#)
- L**
 - LIFs
 - configuring to share data ports with intercluster [19](#)
 - configuring to use dedicated intercluster ports [14](#)
 - intercluster, defined [7](#)
 - requirements for configuring intercluster [7](#)
- N**
 - network requirements
 - full-mesh connectivity [7](#)
 - worksheets [9](#)
 - networking
 - best practices [9](#)
- P**
 - peer relationships
 - about connecting clusters in [14](#)
 - about connecting Vservers in [25](#)
 - creating cluster [23](#)
 - creating Vserver [25](#)
 - requirements for using the Express Guide to configure for clusters and Vservers [4](#)
 - Vserver naming requirement [25](#)
 - workflow for cluster [5](#)
 - peering
 - cluster, workflow flowchart [5](#)
 - intercluster Vserver, workflow flowchart [6](#)
 - where to get additional information about [28](#)
 - where to get additional information about cluster and [28](#)

ports

- configuring intercluster LIFs to share with data [19](#)
- determining whether to use shared or dedicated, for intercluster communication [8](#)

R**relationships**

- about connecting clusters in peer [14](#)
- about connecting Vservers in peer [25](#)
- creating cluster peer [23](#)
- creating Vserver peer [25](#)
- Vserver naming requirement in peer [25](#)

requirements

- intercluster LIF configuration [7](#)

S**shared ports**

- determining whether to use for intercluster communication [8](#)

V**Vserver peer relationship**

- accepting [26](#)

Vserver peer relationships

- creating [25](#)

Vserver peering

- intercluster workflow flowchart [6](#)

Vservers

- about connecting in a peer relationship [25](#)
- naming requirement for peer relationships [25](#)
- requirements for using the Express Guide to configure peering [4](#)

W**workflows**

- cluster peering, flowchart [5](#)
- intercluster Vserver peering, flowchart [6](#)

worksheets, network information [9](#)