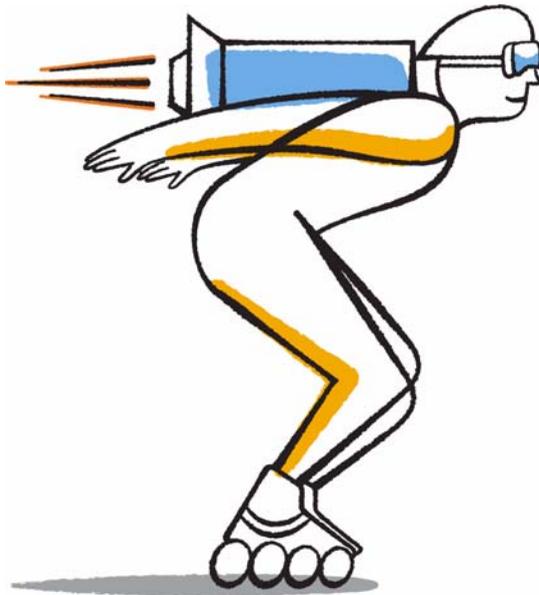




## Data ONTAP<sup>®</sup> 8.1

### Cluster Peering Express Guide For 7-Mode Administrators Learning Cluster-Mode



**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: 215-07236\_A0  
Updated for Data ONTAP 8.1.1  
September 2012



# Contents

<b>Deciding whether to use this guide .....</b>	<b>4</b>
<b>Mode differences: Communication between different geographical locations .....</b>	<b>5</b>
What intercluster LIFs are .....	6
<b>Determining whether to use shared or dedicated ports for intercluster communication .....</b>	<b>8</b>
Required network information .....	9
<b>Connecting clusters in a peer relationship .....</b>	<b>13</b>
Configuring intercluster LIFs to use dedicated intercluster ports .....	13
Configuring intercluster LIFs to share data ports .....	19
Creating the cluster peer relationship .....	23
<b>Where to find additional information .....</b>	<b>25</b>
<b>Copyright information .....</b>	<b>26</b>
<b>Trademark information .....</b>	<b>27</b>
<b>How to send your comments .....</b>	<b>28</b>
<b>Index .....</b>	<b>29</b>

## Deciding whether to use this guide

---

This guide describes how to create a peer relationship between two clusters. A peer relationship enables clusters to communicate. You must create a cluster peer relationship to replicate data between different geographical locations or between two clusters.

You should use this guide if you want to create a peer relationship between two clusters 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 assumes that you have experience administering Data ONTAP operating in 7-Mode but are new to Cluster-Mode.

This guide assumes that the two clusters in the peer relationship that you will create have already been set up and are running the same version of Data ONTAP. The time on the clusters must be synchronized within 300 seconds (five minutes). Cluster peers can be in different time zones.

This guide assumes that you are a cluster administrator. Vserver administrators cannot create cluster peer relationships.

If these assumptions are not correct for your situation, or if you want more conceptual background information, you would be better served by reading the following documentation. This documentation is available from the NetApp Support Site.

- *Data ONTAP System Administration Guide for Cluster-Mode*
- *Data ONTAP Network Management Guide for Cluster-Mode*
- *SnapMirror Configuration and Best Practices Guide for Data ONTAP 8.1 Operating in Cluster-Mode (TR-4015)*

### Related information

*[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)*

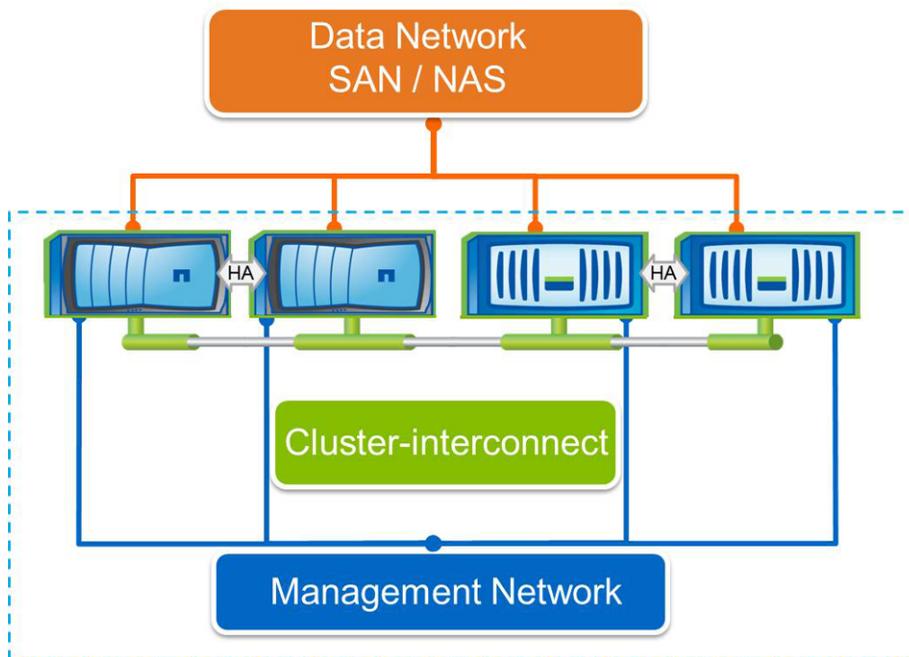
## Mode differences: Communication between different geographical locations

To communicate between different geographical locations in Cluster-Mode, you must establish a peer relationship between a local cluster and a cluster at the other location.

There is no 7-Mode task that corresponds directly to the Cluster-Mode task of setting up cluster peers. The closest comparison would be the task of specifying SnapMirror destination systems by using the `snapmirror.access` option or by editing the `/etc/snapmirror.allow` file.

You set up cluster peering because nodes in Cluster-Mode communicate differently than storage systems do in 7-Mode. In Cluster-Mode, nodes in the same cluster use one method to communicate with each other and other communication methods for other cluster activities:

- The cluster interconnect is used for communication and replication between nodes in the same cluster.
- The data network is used for client access.
- The management network is used for managing nodes, Vservers, and the cluster.
- An intercluster network allows communication between two clusters operating in Cluster-Mode.



When you set up cluster peering, you are configuring the relationship between two clusters over the intercluster network. Cluster peering is necessary for solutions that require data to be sent from one geographical location to another, such as replication for disaster recovery, because all nodes in the

same cluster must be geographically collocated. This means that volumes on nodes in different geographical locations must belong to different clusters. To enable communication between the two clusters, you have to configure the intercluster network.

You can create peer relationships for a single cluster with up to eight remote clusters. Each cluster peer relationship must be created separately. After a peer relationship is created, data can be sent in either direction between the two clusters in that peer relationship. A peer relationship is required for data to be sent in either direction between two clusters.

**Note:** You should only peer to a cluster where administrative control is trusted over a secure network. A peer cluster can back up or copy any data volume from the other cluster.

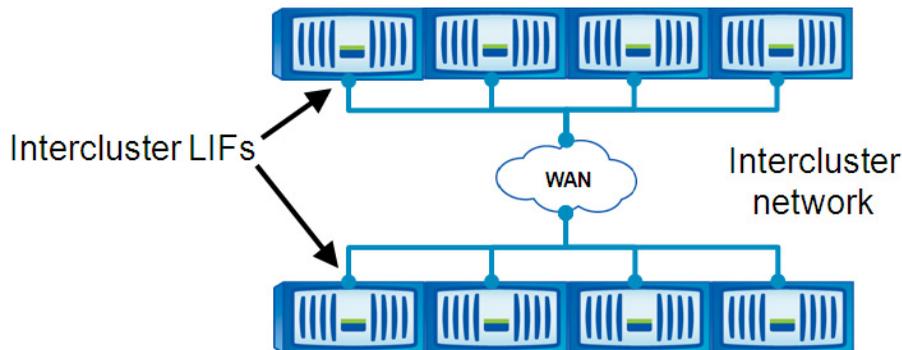
## What intercluster LIFs are

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 as part of the cluster peer configuration process.

The following are requirements for intercluster LIFs:

- 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.
- 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.
- Every intercluster LIF on every node of the local cluster should be able to connect to every intercluster LIF on every node of the remote cluster.

Your cluster peering topology should use full-mesh connectivity. Full-mesh connectivity means that all the intercluster LIFs of one peer cluster can communicate with all of the intercluster LIFs of the other peer 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. However, in the event that the node hosting that address is 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 in the event of a node failure, the peer relationship remains stable.

The port to which you assign an intercluster LIF can be a physical port or a logical port, such as a virtual LAN (VLAN) or an interface group.

**Tip:** Intercluster LIFs fail over only to other ports capable of intercluster communication on the same node. It is a good idea to use a naming convention that indicates the node on which an 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.

## 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, and the port role cannot be modified.
- An intercluster port is used exclusively for communication between clusters.
- Each intercluster port on a cluster in a peer relationship should be routable to every intercluster port of the other cluster.

When you assign the intercluster role to a port, Data ONTAP automatically configures the failover groups for intercluster LIFs to use only intercluster ports. Data protocols cannot fail over or migrate to an intercluster port.

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

### 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 10 GbE, 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 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**

When you configure intercluster communication, 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**

You have to gather IP addresses, ports, and other network information for the intercluster LIFs before configuring a cluster peer relationship.

### **Intercluster LIFs for the local cluster**

You have to 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	
	Intercluster LIF IP address	
	Network mask	
Node 2	Intercluster LIF name	
	Intercluster LIF port	
	Intercluster LIF IP address	
	Network mask	
Node 3	Intercluster LIF name	
	Intercluster LIF port	
	Intercluster LIF IP address	
	Network mask	
Node 4	Intercluster LIF name	
	Intercluster LIF port	
	Intercluster LIF IP address	
	Network mask	
Node 5	Intercluster LIF name	
	Intercluster LIF port	
	Intercluster LIF IP address	
	Network mask	
Node 6	Intercluster LIF name	
	Intercluster LIF port	
	Intercluster LIF IP address	
	Network mask	

Type of information		Your values
Node 7	Intercluster LIF name	
	Intercluster LIF port	
	Intercluster LIF IP address	
	Network mask	
Node 8	Intercluster LIF name	
	Intercluster LIF port	
	Intercluster LIF IP address	
	Network mask	

### Intercluster LIFs for the remote cluster

You have to 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	
	Intercluster LIF IP address	
	Network mask	
Node 2	Intercluster LIF name	
	Intercluster LIF port	
	Intercluster LIF IP address	
	Network mask	
Node 3	Intercluster LIF name	
	Intercluster LIF port	
	Intercluster LIF IP address	
	Network mask	

12 | Data ONTAP 8.1 Cluster Peering Express Guide

Type of information		Your values
Node 4	Intercluster LIF name	
	Intercluster LIF port	
	Intercluster LIF IP address	
	Network mask	
Node 5	Intercluster LIF name	
	Intercluster LIF port	
	Intercluster LIF IP address	
	Network mask	
Node 6	Intercluster LIF name	
	Intercluster LIF port	
	Intercluster LIF IP address	
	Network mask	
Node 7	Intercluster LIF name	
	Intercluster LIF port	
	Intercluster LIF IP address	
	Network mask	
Node 8	Intercluster LIF name	
	Intercluster LIF port	
	Intercluster LIF IP address	
	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. Although the examples in the procedures assume that you have credentials for both clusters, 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 *Data ONTAP Network Management Guide for Cluster-Mode*

### Steps

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

**Example**

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

Speed(Mbps)	Node	Port	Role	Link	MTU	Auto-Negot	Duplex
	Oper					Admin/Oper	Admin/Oper Admin/
-----							
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						
vs1	cluster_mgmt	up/up	192.168.0.xxx/24	cluster01-01	e0c	true
	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
```





```
cluster01-02_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
cluster01-01 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
cluster01-01 clus1      up/up      169.254.xx.xx/24 cluster01-01 e0a      true
cluster01-01 clus2      up/up      169.254.xx.xx/24 cluster01-01 e0b      true
cluster01-01 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
cluster01-02 clus1      up/up      169.254.xx.xx/24 cluster01-02 e0a      true
cluster01-02 clus2      up/up      169.254.xx.xx/24 cluster01-02 e0b      true
cluster01-02 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
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
```

- 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
```

12. 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
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
```

cluster01-02	0.0.0.0/0	192.168.1.1	40
n	192.168.0.0/24		
i	192.168.1.0/24	192.168.0.1	10
	0.0.0.0/0	192.168.1.1	40

14. Repeat Steps 1 through 13 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 *Data ONTAP Network Management Guide for Cluster-Mode*.

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

### Example

This example uses the LIF naming convention of *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 e0c -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 e0c -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
```

3. 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
Logical      Status      Network      Current      Current  Is
Vserver      Interface   Admin/Oper   Address/Mask Node          Port      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
```

4. 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 because e0d is also assigned the data role.

The intercluster LIF is assigned to a data port; therefore, a failover group for the intercluster LIF is created automatically, and contains all ports with the data 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
      Logical          Home          Failover
Failover
Vserver  Interface          Node:Port          Group Usage          Group
-----
cluster01-01
      cluster01-01_icl01 cluster01-01:e0c  system-defined
      Failover Targets: cluster01-01:e0c,
      cluster01-01:e0d
cluster01-02
      cluster01-02_icl01 cluster01-02:e0c  system-defined
      Failover Targets: cluster01-02:e0c,
      cluster01-02:e0d
```

5. 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
```

6. 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
```

cluster01-01	0.0.0.0/0	192.168.0.1	20
n192.168.0.0/24			
cluster01-02	0.0.0.0/0	192.168.0.1	10
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
```

- 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
```

9. Repeat Steps 1 through 8 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 using the `cluster peer create` command.

#### 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 -username admin
Password: *****
```

If DNS is configured to resolve host names for the intercluster IP addresses, you can use host names in the `-peer-addr` 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 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 using the `cluster peer health show` command.

### Example

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

## Where to find additional information

---

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

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

<b><i>SnapMirror Configuration and Best Practices Guide for Data ONTAP 8.1 Operating in Cluster-Mode (TR-4015)</i></b>	Provides information and best practices related to configuring replication in NetApp Data ONTAP software operating in Cluster-Mode.
<b><i>Data ONTAP Data Protection Guide for Cluster-Mode</i></b>	Describes how to manage your backup and recover data on clustered systems with Data ONTAP.
<b><i>Data ONTAP Logical Storage Management Guide for Cluster-Mode</i></b>	Describes how to efficiently manage your logical storage resources on systems running Data ONTAP operating in Cluster-Mode, using volumes, FlexClone volumes, files and LUNs, FlexCache volumes, deduplication, compression, qtrees, and quotas.
<b><i>Data ONTAP Network Management Guide for Cluster-Mode</i></b>	Describes how to connect your cluster to your Ethernet networks and how to manage logical interfaces (LIFs).
<b><i>Data ONTAP System Administration Guide for Cluster-Mode</i></b>	Describes general system administration for NetApp systems running Data ONTAP operating in Cluster-Mode.
<b><i>Data ONTAP 7-Mode to Cluster-Mode Command Map</i></b>	Provides a mapping of 7-Mode commands to Cluster-Mode commands.
<b>NetApp Knowledgebase</b>	(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: support.netapp.com](http://support.netapp.com)*

## Copyright information

---

Copyright © 1994–2012 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, Bypass, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, FAServer, FastStak, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, 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, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, 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](http://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](http://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](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

- A**
  - about this guide [4](#)
  - additional resources [25](#)
- C**
  - cluster peers
    - creating relationships between [23](#)
  - clusters
    - connecting them in peer relationships [13](#)
  - communication
    - between cluster peers in different geographies [5](#)
  - configuring intercluster LIFs to use dedicated intercluster [13](#)
- D**
  - data ports
    - configuring intercluster LIFs to share [19](#)
- I**
  - intercluster communication
    - determining ports for [8](#)
  - intercluster LIFs
    - configuring to share data ports [19](#)
    - configuring to use dedicated intercluster ports [13](#)
    - defined [6](#)
  - intercluster networks
    - configuring intercluster LIFs for [13](#), [19](#)
  - intercluster ports
    - configuring intercluster LIFs to use dedicated [13](#)
- L**
  - LIFs
    - configuring to share data ports with intercluster [19](#)
    - configuring to use dedicated intercluster ports [13](#)
    - intercluster, defined [6](#)
- M**
  - mode differences
    - communication between cluster peers [5](#)
- P**
  - peer relationships
    - connecting clusters [13](#)
    - creating cluster [23](#)
  - ports
    - configuring intercluster LIFs to share with data [19](#)
    - for intercluster communication, shared or dedicated [8](#)
- R**
  - relationships
    - creating cluster peer [23](#)