



Technical Report

# SnapMirror Startup Guide for Data ONTAP 8.1 Operating in Cluster-Mode

Larry Touchette, Scott Owens, NetApp  
February 2012 | TR-3999

## **INTERCLUSTER SNAPMIRROR STARTUP GUIDE**

This document is intended as a quick start guide to configure intercluster replication between different clusters of NetApp® Data ONTAP® 8.1 operating in Cluster-Mode. Intracluster replication, such as intracluster data protection relationships and load-sharing mirrors, are not discussed in this document.

## TABLE OF CONTENTS

<b>1</b>	<b>REPLICATION BETWEEN DIFFERENT CLUSTERS .....</b>	<b>3</b>
1.1	<b>CLUSTER PEERING AND INTERCLUSTER NETWORKING .....</b>	<b>3</b>
<b>2</b>	<b>INTERCLUSTER NETWORKING FOR REPLICATION.....</b>	<b>4</b>
2.1	<b>CLUSTER PEER REQUIREMENTS .....</b>	<b>4</b>
2.2	<b>SHARING OR DEDICATING PORTS FOR REPLICATION .....</b>	<b>4</b>
2.3	<b>SHARING DATA PORTS FOR REPLICATION.....</b>	<b>5</b>
2.4	<b>DEDICATING INTERCLUSTER PORTS FOR REPLICATION.....</b>	<b>8</b>
2.5	<b>CONFIGURING CLUSTER PEERS .....</b>	<b>12</b>
<b>3</b>	<b>INTERCLUSTER SNAPMIRROR RELATIONSHIPS.....</b>	<b>13</b>
3.1	<b>CREATING INTERCLUSTER SNAPMIRROR RELATIONSHIPS.....</b>	<b>13</b>
<b>4</b>	<b>REFERENCES .....</b>	<b>15</b>
<b>5</b>	<b>DOCUMENT INFORMATION .....</b>	<b>15</b>

## LIST OF TABLES

Table 1)	Document version control.....	15
----------	-------------------------------	----

# 1 REPLICATION BETWEEN DIFFERENT CLUSTERS

NetApp Data ONTAP 8.1 operating in Cluster-Mode is the first release of Data ONTAP that allows replication between different clusters while operating in Cluster-Mode, providing cross-site disaster recovery capabilities. New Cluster-Mode capabilities have been introduced in Data ONTAP 8.1 for the following purposes:

- **Cluster peering.** The act of connecting two clusters to allow replication to occur between them.
- **Intercluster logical interfaces.** Logical network interfaces used for intercluster communication.
- **Intercluster ports.** Ports dedicated to intercluster replication.

Clusters must be joined in a peer relationship before replication between different clusters is possible. Cluster peering is a one-time operation that must be performed by the cluster administrators. The following steps are required to configure replication between different clusters:

1. Have a clear understanding of cluster peering.
2. Determine whether or not to share ports for data access and intercluster replication.
3. Designate ports for intercluster replication, if dedicating ports.
4. Create intercluster logical interfaces on each node in the clusters.
5. Peer the clusters together.
6. Create NetApp SnapMirror<sup>®</sup> relationships between different clusters.

For additional information about SnapMirror in Cluster-Mode, see NetApp Technical Report 4015 “SnapMirror Configuration and Best Practices Guide for Data ONTAP 8.1 Operating in Cluster-Mode.”

## 1.1 CLUSTER PEERING AND INTERCLUSTER NETWORKING

Cluster peering requires intercluster logical interfaces (LIFs) for communication and replication between two clusters. An intercluster LIF must be created on an intercluster-capable port, which is a port assigned the role of intercluster or a port assigned the role of data. Ports that are used for the intracluster cluster interconnect may not be used for intercluster replication.

It is possible to assign intercluster LIFs to ports that have the role of data, which are the same ports used for CIFS or NFS access. However, to dedicate certain ports for intercluster replication, simply assign intercluster LIFs to dedicated ports that have the role of intercluster, because intercluster ports can only be used for replication. Each method and how to determine which method to use are discussed later in this document.

Any Ethernet port type in the system can have the role of data or intercluster and be used for intercluster replication, which includes physical ports, such as `e0e` or `e0f`, and logical ports, such as a virtual LAN (VLAN) or an interface group (ifgrp).

At least one intercluster LIF must be created on every node in the cluster. Every intercluster LIF requires an IP address, meaning at least one IP address per node must be dedicated to performing intercluster replication. Provisioning intercluster LIFs only on some nodes of the cluster is not a supported configuration. Configuring an intercluster LIF on every node is required to allow intercluster SnapMirror transfers to continue when volume move operations move volumes between nodes. SnapMirror relationships do not have to be modified because volumes are moved to other nodes in the cluster.

Intercluster LIFs are node scoped; therefore, when the port hosting an intercluster LIF fails, the LIF can fail over only to another intercluster-capable port on that node, as defined by the LIF's failover policy. If a node fails while an intercluster SnapMirror transfer is in progress, the transfer automatically continues using an intercluster LIF on the surviving node of the high-availability (HA) pair.

Creating intercluster LIFs defines the intercluster network on which replication occurs between two different clusters. Replication between two clusters can only occur on the intercluster network, even if the

intercluster network is on the same subnet as a data network in the same cluster. It is not possible for other protocols such as Common Internet File System (CIFS), Network File System (NFS), iSCSI, Fibre Channel (FC), or Fibre Channel over Ethernet (FCoE) to use intercluster LIFs or ports. The IP addresses assigned to intercluster LIFs can reside in the same subnet as data LIFs or in a different subnet. When an intercluster LIF is created, an intercluster routing group is automatically created on that node. If the source and destination clusters must use different subnets for intercluster replication, then it is necessary to define a gateway address for the intercluster routing group.

## 2 INTERCLUSTER NETWORKING FOR REPLICATION

### 2.1 CLUSTER PEER REQUIREMENTS

Cluster peer requirements include the following:

- The time on the clusters must be in sync within 300 seconds (five minutes) for peering to be successful. Cluster peers can be in different time zones.
- At least one intercluster LIF must be created on every node in the cluster.
- Every intercluster LIF requires an IP address dedicated for intercluster replication.
- The correct maximum transmission unit (MTU) value must be used on the network ports that are used for replication. The network administrator can identify which MTU value to use in the environment. The default value of 1,500 is correct for most environments.
- All paths on a node used for intercluster replication should have equal performance characteristics.
- The intercluster network must provide connectivity among all intercluster LIFs on all nodes in the cluster peers. Every intercluster LIF on every node in a cluster must be able to connect to every intercluster LIF on every node in the peer cluster.

### 2.2 SHARING OR DEDICATING PORTS FOR REPLICATION

There are a number of configurations and requirements to consider when determining whether to share or dedicate ports for replication; they include:

- LAN type
- Available WAN bandwidth (compared to LAN bandwidth)
- Replication interval
- Change rate
- Number of ports used by the solution

Intercluster LIF failover policies can be manually configured so that intercluster LIFs only use a subset of data ports. However, when dedicating specific ports for intercluster by assigning ports the role of intercluster, Data ONTAP automatically verifies that only intercluster ports are used for replication and does not allow any data protocols to fail over or migrate to the intercluster ports. This is a more rigid and automated method for dedicating ports to intercluster replication.

Consider the following when sharing ports for intercluster replication:

- For a high-speed network such as 10 Gigabit Ethernet (10GbE), a sufficient amount of local LAN bandwidth might be available to perform replication on the same 10GbE ports that are used for data access. In many cases, the available WAN bandwidth is far less than 10GbE, which reduces the LAN network utilization to only that which the WAN is capable of supporting.
- All nodes in the cluster might have to replicate data and share the available WAN bandwidth, making data port sharing more acceptable.

- Sharing ports for data and replication eliminates the extra port counts required to dedicate ports for replication.
- If the replication interval is set to perform only after hours, while minimal to no client activity exists, then simply using data ports for replication during this time is acceptable, even without a 10GbE LAN connection.
- Consider the data change rate and replication interval and whether or not the amount of data that must be replicated on each interval requires enough bandwidth that it might cause contention with data protocols if sharing data ports.
- When data ports for intercluster replication are shared, the intercluster LIFs can be migrated to any other intercluster-capable port on the same node to control the specific data port that is used for replication.

Consider the following when dedicating ports for intercluster replication:

- If the amount of available WAN bandwidth is similar to that of the LAN ports, and the replication interval is such that replication occurs while regular client activity exists, then dedicate Ethernet ports for intercluster replication to avoid contention between replication and the data protocols.
- If the network utilization generated by the data protocols (CIFS, NFS, iSCSI) is such that the network utilization is above 50%, then dedicate ports for replication to allow for nondegraded performance in the event of a node failover.
- When physical 10GbE ports are used for data and replication, VLAN ports can be created for replication and the logical ports dedicated for intercluster replication.
- Consider the data change rate and replication interval and whether or not the amount of data that must be replicated on each interval requires enough bandwidth that it might cause contention with data protocols if sharing data ports.
- If the replication network requires configuration of an MTU size that differs from the MTU size used on the data network, then physical ports must be dedicated for replication because the MTU size can only be configured on physical ports.

## 2.3 SHARING DATA PORTS FOR REPLICATION

### CONFIGURING INTERCLUSTER LIFS TO SHARE DATA PORTS

This section demonstrates how to create intercluster LIFs that share data ports. In this example, a two-node cluster exists in which each node has two data ports, e0c and e0d, which are shared for intercluster replication. Before completing the following steps, replace the ports, networks, IP addresses, subnet masks, and subnets with those specific to the environment.

1. Check the role of the ports in the cluster.

```
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. This example uses the LIF naming convention of <nodename>\_icl# for the intercluster LIF.

```

cluster01::> network int 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 int 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.

```

cluster01::> network int 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

### 4. Verify that the intercluster LIFs are configured to be redundant.

```

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

```

Vserver	Logical Interface	Home Node:Port	Failover Group Usage	Failover Group
cluster01-01	cluster01-01_icl01	cluster01-01:e0c	system-defined	cluster01-01:e0c, cluster01-01:e0d
cluster01-02	cluster01-02_icl01	cluster01-02:e0c	system-defined	cluster01-02:e0c, cluster01-02:e0d

**Note:** 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 role of data. The intercluster LIF is assigned to a data port; therefore, a failover group for the intercluster LIF is automatically created containing all ports with the role of data 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.

#### Best Practice

Verify that all necessary ports have access to the necessary networks or VLANs to allow communication after port failover. Intercluster LIFs can only fail over to other intercluster-capable ports on the same node, as defined in the failover group (in this example, other data ports); they cannot fail over to ports on other nodes. It is a best practice to configure two intercluster ports per node when dedicating ports for intercluster replication.

### 5. The intercluster network might require intercluster routes. An intercluster routing group is automatically created for the intercluster LIFs. Run the `net routing-group show` command to display the routing groups.

```

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.

**Note:** No intercluster routes are available.

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

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

7. If communication between intercluster LIFs in different clusters requires routing, then create an intercluster route. The intercluster networks are scoped to each node; therefore, create an intercluster route on each node. In this example, 192.168.1.1 is the gateway address for the 192.168.1.0/24 network.

**Note:** 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.

```
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
```

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

8. Display the newly created routes.

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

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

**Note:** 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, as shown in the output. These routes are only used for intercluster communication.

- Repeat the steps in this section to configure intercluster networking in the other cluster.

## 2.4 DEDICATING INTERCLUSTER PORTS FOR REPLICATION

### CONFIGURING INTERCLUSTER LIFS TO USE DEDICATED INTERCLUSTER PORTS

This section demonstrates how to designate specific ports as intercluster ports and assign intercluster LIFs to those ports. 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. Before completing the following steps, replace the ports, networks, IP addresses, subnet masks, and subnets with those specific to the environment.

- Check the role of the ports in the cluster.

```
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.

```
cluster01::> network int 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 migrate the LIF to another port because intercluster ports cannot host data LIFs. This migration is nondisruptive, assuming the other data ports have been configured properly so that clients are able to access the LIF after migration.

**Note:** If no LIFs are assigned to the ports, skip this step.

```
cluster01::> network int migrate -vserver vs1 -lif vs1_lif1 -dest-node cluster01-01 -dest-port e0d
```

```
cluster01::> network int 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	e0d	true



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	e0d	false

4. The newly migrated LIFs might need modification to the LIF home port to reflect the new port where the LIF should reside.

```
cluster01::> network int modify -vserver vserver1 -lif vs1_lif1 -home-node dpgl-01 -home-port e0d

cluster01::> network int show
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
vserver1	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.

```
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 role of the correct ports has been changed to intercluster.

```
cluster01::> network port show -role intercluster
```

Node	Port	Role	Link	MTU	Auto-Negot Admin/Oper	Duplex Admin/Oper	Speed (Mbps) 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. This example uses the LIF naming convention <nodename>\_icl# for intercluster LIF.

```
cluster01::> network int 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 int 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
```

8. Verify that the intercluster LIFs are configured for redundancy.

```
cluster01::> network int show -role intercluster -failover
```

Vserver	Logical Interface	Home Node:Port	Failover Group Usage	Failover 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

```

**Note:** 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 failover group is automatically created containing all ports with the intercluster role on that node. In this example, the failover group does not include any data ports. 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.

### Best Practice

Verify that all the necessary ports have access to the necessary networks or VLANs to allow communication after port failover. Intercluster LIFs can only fail over to other intercluster-capable ports on the same node, as defined in the failover group (in this example, other intercluster ports); they cannot fail over to ports on other nodes. It is a best practice to configure two intercluster ports per node when dedicating ports for intercluster replication.

#### 9. Verify that the intercluster LIFs were created properly.

```

cluster01::> network int 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
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

#### 10. The intercluster network might require intercluster routes. An intercluster routing group is automatically created for the intercluster LIFs. Run the network routing-groups show command to display the routing groups.

```

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

11. Display the routes in the cluster.

**Note:** 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, then create an intercluster route. The intercluster networks are scoped to each node; therefore, create an intercluster route on each node. In this example, 192.168.1.1 is the gateway address for the 192.168.1.0/24 network.

**Note:** 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.

```
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
```

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

13. Display the newly created routes.

```
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
```

**Note:** 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, as shown in the output. These routes are only used for intercluster communication.

14. Repeat the steps in this section to configure intercluster networking in the other cluster. 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. In both configurations, verify that the ports have access to the proper subnets, VLANs, and so on.

## 2.5 CONFIGURING CLUSTER PEERS

After the nodes in both clusters are configured with intercluster LIFs, the clusters can be peered together to allow the creation of replication relationships between the clusters.

A single cluster can be peered with up to eight remote clusters. Each cluster peer relationship must be created separately. Once a peer relationship is created, mirroring can be performed in either direction between the two clusters in that peer relationship.

In this example, `cluster01` is peered with a remote cluster named `cluster02`. `cluster02` is also 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` are used in the following `cluster peer create` command example. However, if DNS is configured to resolve hostnames for the intercluster IP addresses, hostnames in the `-peer-addr`s option can be used. It is not likely that intercluster IP addresses frequently change; however, using hostnames allows intercluster IP addresses to be changed without having to modify the cluster peer relationship.

Before completing the following steps, replace the cluster name and IP addresses with those specific to your environment.

1. Run the `cluster peer create` command to peer the clusters. Specify the IP address of at least one intercluster LIF per node from the remote cluster in the `-peer-addr`s option, separated by commas. Provide the remote cluster administrator user name and password when prompted.

```
cluster01::> cluster peer create -peer-addr 192.168.2.203,192.168.2.204 -username
admin
Password: *****
```

2. Display the newly created cluster peer relationship.

```
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
```

### Best Practice

The IP addresses used in the `-peer-addr`s option are listed as remote intercluster addresses. These addresses are used to discover all other intercluster IP addresses capable of communication between the clusters, listed as active IP addresses. The list of active IP addresses might be longer, depending on how many intercluster LIFs exist per node, because they are automatically discovered. As intercluster LIFs become available or unavailable, the list of active IP addresses could change. Active IP addresses discovery might be automatic in certain events, such as when a node reboots. The `-peer-addr`s option requires only one remote cluster address to be provided; however, if that address becomes unavailable in the event that the node hosting that address is down, then the cluster peer relationship could not be rediscovered. Therefore, it is a best practice to use at least one intercluster IP address from each node in the remote cluster, so in the event of a node failure, the peer relationship remains stable.

3. Preview the health of the cluster peer relationship.

```
cluster01::> cluster peer health show

Node          cluster-Name          Node-Name
```

	Ping-Status	RDB-Health	Cluster-Health	Avail...
cluster01-01				
cluster02		cluster02-01		
Data: interface_reachable				
ICMP: interface_reachable	true	true		true
		cluster02-02		
Data: interface_reachable				
ICMP: interface_reachable	true	true		true
cluster01-02				
cluster02		cluster02-01		
Data: interface_reachable				
ICMP: interface_reachable	true	true		true
		cluster02-02		
Data: interface_reachable				
ICMP: interface_reachable	true	true		true

### 3 INTERCLUSTER SNAPMIRROR RELATIONSHIPS

After the cluster peer relationship has been successfully created between the two clusters, create the intercluster SnapMirror relationships.

Always use the hostname of the source and destination clusters to create SnapMirror relationships in Cluster-Mode. The hostnames of the source and destination clusters are exchanged and stored while creating the cluster peer relationship. If the hostname resolution is not configured on the clusters, still create SnapMirror relationships using only the names of the clusters; never use the IP addresses.

Complete the following requirements before creating an intercluster SnapMirror relationship:

- Configure the source and destination nodes for intercluster networking.
- Configure the source and destination clusters in a peer relationship.
- Create a destination NetApp Vserver; volumes cannot exist in Cluster-Mode without a Vserver.
- Verify that the source and destination Vservers have the same language type.
- Create a destination volume with a type of data protection (DP), with a size equal to or greater than that of the source volume.
- Assign a schedule to the SnapMirror relationship to perform periodic updates. If any of the existing schedules are not adequate, a new schedule entry must be created.

#### 3.1 CREATING INTERCLUSTER SNAPMIRROR RELATIONSHIPS

In this startup guide, the SnapMirror relationship is created to replicate FlexVol volume `vol1` from Vserver `vs1` in `cluster01` to FlexVol® volume `vol1` in Vserver `vs5` in `cluster02`. Before completing the following steps, replace the clusters, Vservers, and volumes with those specific to the environment.

**Note:** Vserver creation is beyond the scope of this document. The steps in this section assume that Vserver `vs5` has already been created and configured.

1. Create a destination volume with `-type DP` in the destination Vserver.

**Note:** The size of the volume must be equal to or greater than that of the source volume.

```
cluster02::> vol create -vserver vs5 -volume vol1 -aggregate aggr1 -size 1.24GB -type DP
```

2. Create a schedule to assign to the relationship. In this example, an hourly schedule is created to run at the top of every hour.

```
cluster02::> job schedule cron create -name Hourly_SnapMirror -minute 0
```

### 3. Display the newly created schedule.

```
cluster02::> job schedule cron show
Name          Description
-----
5min          @:00,:05,:10,:15,:20,:25,:30,:35,:40,:45,:50,:55
8hour        @2:15,10:15,18:15
Hourly_SnapMirror
              @:00
avUpdateSchedule @2:00
daily        @0:10
hourly       @:05
weekly       Sun@0:15
```

### 4. Create a SnapMirror relationship with `-type DP` and assign the schedule created in the previous step.

```
cluster02::> snapmirror create -source-path cluster01://vs1/vol1 -destination-path
cluster02://vs5/vol1 -type DP -schedule Hourly_SnapMirror
```

### 5. Review the SnapMirror relationship.

**Note:** The `create` command only creates the SnapMirror relationship; it does not initialize it.

```
cluster02::> snapmirror show

Source          Destination  Mirror          Relationship  Total
Path           Type  Path           State          Status        Progress  Healthy
-----
cluster01://vs1/vol1
                DP    cluster02://vs5/vol1
                Uninitialized Idle          -           -
```

### 6. Initialize the SnapMirror relationship in the destination cluster.

```
cluster02::> snapmirror initialize -destination-path cluster02://vs5/vol1

[Job 105] Job succeeded: SnapMirror: done
```

### 7. Verify the progress of the replication.

```
cluster02::> snapmirror show

Source          Destination  Mirror          Relationship  Total
Path           Type  Path           State          Status        Progress  Healthy
-----
cluster01://vs1/vol1
                DP    cluster02://vs5/vol1
                Uninitialized Transferring  31.8MB     -
```

### 8. Verify the status using the job ID provided for the initialize command.

```
cluster02::> job show -id 105

Job ID Name          Owning
              Vserver      Node          State
-----
105    SnapMirror initialize
              cluster02  cluster02-02  Running
              Description: snapmirror initialize of destination cluster02://vs5/vol1
```

### 9. Periodically check the progress of the replication until `Status = Idle` and `Healthy = true`.

**Note:** Press the up arrow and then Enter to run the same command again.

```
cluster02::> snapmirror show
```

Source Path	Type	Destination Path	Mirror State	Relationship Status	Total Progress	Healthy
cluster01://vs1/vol1	DP	cluster02://vs5/vol1	Snapmirrored	Idle	-	true

The process to perform failover is beyond the scope of this guide. SnapMirror relationships can be failed over using the `snapmirror break` command and resynchronized in either direction using the `snapmirror resync` command. In order for NAS clients to access data in the destination volumes, CIFS shares and NFS export policies must be created in the destination Vserver and assigned to the volumes. For SAN clients, igroups must be created, and LUN mapping must be performed in the destination Vserver. More information about SnapMirror failover will be provided in a later version of this document.

## 4 REFERENCES

For more information, refer to the following documents:

- [Data ONTAP 8.1 Cluster-Mode Data Protection Guide](#)
- NetApp Technical Report 4015 “SnapMirror Configuration and Best Practices Guide for Data ONTAP 8.1 Operating in Cluster-Mode”

## 5 DOCUMENT INFORMATION

This section contains the version control details for this document.

Table 1) Document version control.

Date	Version	Description	Author
02/1/2012	1.1	Minor updates based on approved/edited information from the NetApp Technical Report 4015 “SnapMirror Configuration and Best Practices Guide for Data ONTAP 8.1 Operating in Cluster-Mode”.	Larry Touchette
12/8/2011	1.0	Initial release	Larry Touchette

NetApp provides no representations or warranties regarding the accuracy, reliability or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®



[www.netapp.com](http://www.netapp.com)

© 2012 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, Data ONTAP, FlexVol, and SnapMirror are trademarks or registered trademarks of NetApp, 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. TR-3999-0212