



---

# Replacing a Cisco® Nexus 5010 switch with a Nexus 5020 switch

Use this procedure to replace a Cisco Nexus 5010 switch with a Nexus 5020 switch with or without an expansion module in a cluster.

## Before you begin

The following conditions must exist before performing the switch replacement in the current environment and on the replacement switch:

- The cluster switches supported in this procedure are:
  - Nexus 5010 with or without an expansion module
  - Nexus 5020
- The Data ONTAP and NX-OS versions supported in this procedure are on the [Cisco Ethernet Switch description page](#).
- The current 5010 switch infrastructure requires you to verify that:
  - The existing cluster is fully functional (no defective NICs, switch ports, or connectors).
  - All cluster ports are up.
  - All cluster logical interfaces (LIFs) are up and have not been migrated.
  - Running `cluster ping-cluster -node node1` reports that basic connectivity and larger than PMTU communication is successful on all paths.
- The new 5020 switch infrastructure must have:
  - Management connectivity on both cluster switches
  - Console access to the cluster switches is in place
  - Disabled switch ports for node connectivity (ports 1 to 32)
  - Enabled switch ports for ISL connectivity (ports 33 to 40)
  - Completed initial configuration and customization of both cluster switches
  - Both cluster switches loaded with the latest reference configuration file (RCF)
  - Cluster switches running the required version of NX-OS

## About this task

The examples in this procedure use the following switch and node nomenclature:

- Cluster switch information:
  - The 5010 node ports are e1/1 through e1/12 (and e2/1 through e2/6 if you have an expansion module).
  - The 5010 Inter-Switch Link (ISL) ports are e1/13 through e1/20.
  - The 5020 node ports are e1/1 through e1/32.
  - The 5020 ISL ports are e1/33 through e1/40.
- Node information for the examples used in this procedure:
  - The name of the cluster is cluster.
  - c5010-1 is the name of the first 5010 switch to be replaced; c5010-2 is the name of the second switch to be replaced.
  - c5020-2 is the name of the 5020 switch to replace c5010-2; c5020-1 is the name of the 5020 switch to replace c5010-1.
  - The node names are node1 and node2.
  - clus1 and clus2 are cluster LIFs.

- The first cluster port in is e1a and the second cluster port is e2a. Refer to the *Hardware Universe* for the actual cluster ports supported on your platform.
- The Storage Virtual Machine (SVM, formerly known as Vserver) name is node1.

This procedure has two parts:

1. Steps 1 to 16 migrate ports from the first 5010 to the first 5020.
2. Steps 23 to 39 migrate ports from the second 5010 to the second 5020.

### Steps

1. Before you begin, use `run -node local cdpd show-neighbors` to view data about the configuration of the cluster switch configuration.

### Example

This example shows sample output from the command:

```
cluster::> run -node local cdpd show-neighbors
Local Remote Remote Remote Hold Remote
Port Device Interface Platform Time Capability
-----
e1a icea-4(SSII1329. Ethernet1/21 N5K-C5020P-BFS 170 SI
e1b icea-3(SSII1410. Ethernet1/21 N5K-C5020P-BF 157 SI
```

2. Switch to advanced mode and use the `statistics show-periodic` command to display information about the current cluster traffic and data.

### Example

This example shows sample output from the command (the output has been truncated to fit this page):

```
cluster::> set -privilege advanced
cluster::*> statistics show-periodic -node node1 -summary true -iterations 10
node1: cluster.cluster: 1/7/2014 09:10:49
  cpu  cpu    total          fcache          total    total data ...
  avg busy    ops  nfs-ops cifs-ops    ops spin-ops    recv    sent busy ...
-----
  0%  15%      0          0          0          0          0          5.62MB  24.7KB  0% ...
  0%  15%      0          0          0          0          0          5.26MB  31.4KB  0% ...
  0%  19%      0          0          0          0          0          5.35MB  27.8KB  0% ...
  0%  18%      0          0          0          0          0          5.23MB  15.4KB  0% ...
  0%  19%      0          0          0          0          0          5.29MB  15.8KB  0% ...
  0%  19%      0          0          0          0          0          5.31MB  18.1KB  0% ...
  0%  20%      0          0          0          0          0          5.27MB  21.5KB  0% ...
  0%  17%      0          0          0          0          0          5.27MB  21.1KB  0% ...
  0%  17%      0          0          0          0          0          5.33MB  38.6KB  0% ...
  0%  15%      0          0          0          0          0          5.33MB  32.1KB  0% ...
node1: cluster.cluster: 1/7/2014 09:11:10
  cpu  cpu    total          fcache          total    total data ...
  avg busy    ops  nfs-ops cifs-ops    ops spin-ops    recv    sent busy ...
-----
Minimums:
  0%  15%      0          0          0          0          0          5.23MB  15.4KB  0% ...
Averages for 10 samples:
  0%  17%      0          0          0          0          0          5.33MB  24.7KB  0% ...
Maximums:
  0%  20%      0          0          0          0          0          5.62MB  38.6KB  0% ...
```

3. Disable ports 1 to 32 on both Nexus 5020 cluster switches. Do not disable the ISL ports.

### Example

This example shows ports 1 to 32 being shut down on both 5020 switches, c5020-1 and c5020-2:

```
c5020-1# configure
c5020-1(config)# interface ethernet 1/1-32
c5020-1(config-if-range)# shutdown

c5020-2# configure
c5020-2(config)# interface ethernet 1/1-32
c5020-2(config-if-range)# shutdown
```

4. Verify that all cluster ports are transmitting and receiving traffic through the console or SSH connection.

### Example

This example verifies that both ports e1a and e2a in your cluster are transmitting and receiving traffic.

```
cluster::*> system node run -node * ifstat e1a
4 entries were acted on.
Node: node1
-- interface e1a (13 days, 21 hours, 41 minutes, 53 seconds) --

RECEIVE
Frames/second:    9953 | Bytes/second:    64604k | Errors/minute:    0
Discards/minute: 0 | Total frames:    2565m | Total bytes:      18894g
Total errors:    0 | Total discards:  0 | Multi/broadcast: 108k
No buffers:     0 | Non-primary u/c: 0 | LRO segments:    0
LRO bytes:      0 | Tag drop:        0 | Vlan tag drop:   0
Vlan untag drop: 0 | Vlan forwards:   0 | Vlan broadcasts: 0
.
.
.
cluster::*> system node run -node * ifstat e2a
1 entry was acted on.
Node: node1
-- interface e2a (13 days, 21 hours, 42 minutes, 7 seconds) --

RECEIVE
Frames/second:    9937 | Bytes/second:    64603k | Errors/minute:    0
Discards/minute: 0 | Total frames:    2579m | Total bytes:      18973g
Total errors:    0 | Total discards:  0 | Multi/broadcast: 83599
No buffers:     0 | Non-primary u/c: 0 | LRO segments:    0
LRO bytes:      0 | Tag drop:        0 | Vlan tag drop:   0
Vlan untag drop: 0 | Vlan forwards:   0 | Vlan broadcasts: 0
Vlan unicasts:  0 | Mac octets:     18984g | UCast pkts:      2579m
```

5. Obtain baseline statistics on each node over an SSH session and monitor `sysstat` in the remaining steps of this procedure to ensure that traffic is not affected.

This step is run and monitored for the duration of the switch replacement. At no time should traffic be impacted for more than 60 seconds.

### Example

```
host# ssh admin@node1 "run local sysstat 5"
```

6. Ensure that all cluster LIFs are up and operational on the existing 5010 switches, and that none of the cluster LIFs have migrated to other cluster LIFs, by using `network port show` and `network interface show`.

## Example

This example shows verification that all the cluster LIFs are up and that none of the cluster LIFs have migrated to other cluster LIFs. Make sure that the Is Home column value is true. If the Is Home is false the cluster LIF has migrated to another port on the node.

```
cluster::*> network port show -role cluster
```

Node	Port	Role	Link	MTU	Auto-Negot Admin/Oper	Duplex Admin/Oper	Speed (Mbps) Admin/Oper
-----							
node1	e1a	cluster	up	9000	true/true	full/full	auto/10000
	e2a	cluster	up	9000	true/true	full/full	auto/10000
node2	e1a	cluster	up	9000	true/true	full/full	auto/10000
	e2a	cluster	up	9000	true/true	full/full	auto/10000

4 entries were displayed.

```
cluster::*> network interface show -role cluster
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
-----						
node1	clus1	up/up	10.10.0.1/24	node1	e1a	true
	clus2	up/up	10.10.0.2/24	node1	e2a	true
node2	clus1	up/up	10.10.0.3/24	node2	e1a	true
	clus2	up/up	10.10.0.4/24	node2	e2a	true

4 entries were displayed.

7. Use the system node run ifstat command to verify that both ports e1a and e2a are transmitting and receiving traffic.

## Example

Here is the portion of an example from the output of the system node run ifstat on ports e1a and e2a:

```
cluster::*> system node run -node node1 ifstat e1a
4 entries were acted on.
Node: node1
-- interface e1a (13 days, 21 hours, 41 minutes, 53 seconds) --

RECEIVE
Frames/second:    9953 | Bytes/second:    64604k | Errors/minute:    0
Discards/minute: 0 | Total frames:    2565m | Total bytes:    18894g
Total errors:    0 | Total discards: 0 | Multi/broadcast: 108k
No buffers:      0 | Non-primary u/c: 0 | LRO segments:    0
LRO bytes:       0 | Tag drop:        0 | Vlan tag drop:   0
Vlan untag drop: 0 | Vlan forwards:  0 | Vlan broadcasts: 0
.
.
.

cluster::*> system node run -node node1 ifstat e2a
1 entry was acted on.
Node: node1

-- interface e2a (13 days, 21 hours, 42 minutes, 7 seconds) --

RECEIVE
Frames/second:    9937 | Bytes/second:    64603k | Errors/minute:    0
Discards/minute: 0 | Total frames:    2579m | Total bytes:    18973g
Total errors:    0 | Total discards: 0 | Multi/broadcast: 83599
No buffers:      0 | Non-primary u/c: 0 | LRO segments:    0
```

```
LRO bytes:          0 | Tag drop:          0 | Vlan tag drop:     0
Vlan untag drop:    0 | Vlan forwards:     0 | Vlan broadcasts:  0
Vlan unicasts:      0 | Mac octets:        18984g | UCast pkts:       2579m
.
.
.
```

- On node1's console, migrate the cluster traffic from the cluster port on c5010-1 to the cluster port on c5010-2; confirm that traffic is not flowing through e1a but is flowing through e2a, and that the cluster LIF has migrated.  
Monitor the traffic as shown in Step 3.

**Example**

This example shows traffic being migrated from the cluster port on c5010-1 to the cluster port on c5010-2 and shows verification that the traffic is not flowing through e1a.:

```
cluster::*> network interface migrate -vserver node1 -lif clus1 -
source-node node1 -dest-node node1 -dest-port e2a

cluster::*> network interface show -role cluster
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
node1	clus1	up/up	10.10.0.1/24	node1	e2a	false
	clus2	up/up	10.10.0.2/24	node1	e2a	true
node2	clus1	up/up	10.10.0.3/24	node2	e1a	true
	clus2	up/up	10.10.0.4/24	node2	e2a	true

4 entries were displayed.

- Use the network interface show -role cluster -fields home-port command to display information about the LIFs that have the home port or interface group that you specify.

**Example**

This example shows port e1a on node1 being shut down and shows that action being verified:

```
cluster::*> network interface show -role cluster -fields home-port
```

Vserver	lif	home-port
node1	clus1	e1a
node1	clus2	e2a
node2	clus1	e1a
node2	clus2	e2a

- Shut down the port that was just migrated in Step 6 and confirm the action has been completed.

**Example**

This example shows port e1a on node1 being shut down and shows that action being verified:

```
cluster::*> network port modify -node node1 -port e1a -up-admin false
cluster::*> network port show -role cluster
```

Node	Port	Role	Link	MTU	Auto-Negot Admin/Oper	Duplex Admin/Oper	Speed (Mbps) Admin/Oper
node1	e1a	cluster	down	9000	true/true	full/full	auto/10000

```

node2 e2a cluster up 9000 true/true full/full auto/10000
      e1a cluster up 9000 true/true full/full auto/10000
      e2a cluster up 9000 true/true full/full auto/10000
4 entries were displayed.

```

11. Repeat steps 8 through 10 for the remaining nodes connected to c5010-1, starting with node1's HA partner.
12. At this time, all traffic should be running locally on c5010-2; now disable the ISL ports on c5010-2.  
Monitor traffic on all nodes as described in Step 3.

### Example

This example shows ISL ports 13 through 20 being shut down on the c5010-2 switch:

```

c5010-2# configure
c5010-2(config)#interface ethernet 1/13-20
c5010-2(config-if-range)#shutdown

```

13. Remove all node-facing ISL cables from the c5010-1 switch and connect them to the c5020-1 switch.  
Connect the node-facing ports to the same ports; the new ISL should be on ports 13 through 20 on c5010-2 and ports 33 through 40 on c5020-1.
14. Enable the ISL ports on the c5010-2 switch (ports 13 to 20) and verify that the ISL is working between c5010-2 and c5020-1.

### Example

This example shows ports 13 to 20 on c5010-2 being enabled and the `show port-channel summary` command being used to verify that the ISLs are functioning on both the c5010-2 and c5020-1 switch.

Look for an `S` and `U` in the Port-Channel column to indicate that the port-channel is switched and up. The ports in the Members port column should have a `P` after them to indicate that they are up.

```

c5010-2# configure
c5010-2(config)#interface ethernet 1/13-20
c5010-2(config-if-range)#no shutdown

c5010-2#show port-channel summary

Flags: D - Down          P - Up in port-channel (members)
       I - Individual    H - Hot-standby (LACP only)
       s - Suspended     r - Module-removed
       S - Switched      R - Routed
       U - Up (port-channel)
       M - Not in use. Min-links not met
-----
Group Port-      Type      Protocol  Member Ports
Channel
-----
1      Po1(SU)     Eth       LACP      Eth1/13(P) Eth1/14(P) Eth1/15(P)
                               Eth1/16(P) Eth1/17(P) Eth1/18(P)
                               Eth1/19(P) Eth1/20(P)

c5020-1#show port-channel summary

Flags: D - Down          P - Up in port-channel (members)
       I - Individual    H - Hot-standby (LACP only)
       s - Suspended     r - Module-removed
       S - Switched      R - Routed
       U - Up (port-channel)
       M - Not in use. Min-links not met
-----
Group Port-      Type      Protocol  Member Ports
Channel
-----

```

```

-----
1      Po1(SU)      Eth      LACP      Eth1/33(P)  Eth1/34(P)  Eth1/35(P)
                                Eth1/36(P)  Eth1/37(P)  Eth1/38(P)
                                Eth1/39(P)  Eth1/40(P)

```

15. Enable the node-facing ports on the c5020-1 switch.

### Example

These commands configure and enable ports 1 through 32 on the c5020-1 switch:

```

c5020-1# configure
c5020-1(config)#interface ethernet 1/1-32
c5020-1(config-if-range)#no shutdown

```

16. Use the system node run ifstat command to verify that both ports e1a and e2a are transmitting and receiving traffic.

### Example

Here is the portion of an example from the output of the system node run ifstat on ports e1a and e2a:

```

cluster::*> system node run -node node1 ifstat e1a
4 entries were acted on.
Node: node1
-- interface e1a (13 days, 21 hours, 41 minutes, 53 seconds) --

RECEIVE
Frames/second:    9953 | Bytes/second:    64604k | Errors/minute:    0
Discards/minute: 0 | Total frames:    2565m | Total bytes:      18894g
Total errors:    0 | Total discards: 0 | Multi/broadcast: 108k
No buffers:      0 | Non-primary u/c: 0 | LRO segments:    0
LRO bytes:       0 | Tag drop:        0 | Vlan tag drop:   0
Vlan untag drop: 0 | Vlan forwards:  0 | Vlan broadcasts: 0
.
.
.
cluster::*> system node run -node node1 ifstat e2a
1 entry was acted on.

Node: node1

-- interface e2a (13 days, 21 hours, 42 minutes, 7 seconds) --

RECEIVE
Frames/second:    9937 | Bytes/second:    64603k | Errors/minute:    0
Discards/minute: 0 | Total frames:    2579m | Total bytes:      18973g
Total errors:    0 | Total discards: 0 | Multi/broadcast: 83599
No buffers:      0 | Non-primary u/c: 0 | LRO segments:    0
LRO bytes:       0 | Tag drop:        0 | Vlan tag drop:   0
Vlan untag drop: 0 | Vlan forwards:  0 | Vlan broadcasts: 0
Vlan unicasts:  0 | Mac octets:     18984g | UCast pkts:      2579m

```

17. On node1, enable port e1a and revert traffic back to its home cluster port. Verify that this process is successful and that both cluster ports are transmitting and receiving.

### Example

This example enables port e1a, reverts traffic back to its home cluster port, and verifies the result:

```

cluster::*> network port modify -node node1 -port e1a -up-admin true
cluster::*> network interface revert -vserver node1 -lif clus1

cluster::*> network port show -role cluster
                                Auto-Negot Duplex      Speed (Mbps)

```

```

Node   Port   Role      Link   MTU Admin/Oper Admin/Oper Admin/Oper
-----
node1
      e1a   cluster  up     9000 true/true  full/full  auto/10000
      e2a   cluster  up     9000 true/true  full/full  auto/10000
node2
      e1a   cluster  up     9000 true/true  full/full  auto/10000
      e2a   cluster  up     9000 true/true  full/full  auto/10000
4 entries were displayed.

cluster::*> network interface show -role cluster

Vserver Logical Status Network Current Current Is
         Interface Admin/Oper Address/Mask Node      Port   Home
-----
node1
      clus1 up/up  10.10.0.1/24 node1    e1a   true
      clus2 up/up  10.10.0.2/24 node1    e2a   true
node2
      clus1 up/up  10.10.0.3/24 node2    e1a   true
      clus2 up/up  10.10.0.4/24 node2    e2a   true
4 entries were displayed.

cluster::*> system node run -node node1 ifstat e1a
cluster::*> system node run -node node1 ifstat e2a

```

- Use `cluster show` to display information about the nodes in the cluster.

#### Example

This example shows that the node health for `node1` and `node2` in this cluster is true:

```

cluster::*> cluster show
Node           Health Eligibility  Epsilon
-----
node1          true   true          false
node2          true   true          false

```

- Use the `event log show` command with the `-messagename` parameter to display quorum in the latest event logs:

#### Example

This example displays the event log output for `node1` and `node2` with the quorum status:

```

cluster::*> event log show -messagename *quorum* | *offline*
Time           Node           Severity Event
-----
8/27/2013 16:21:26 node1          INFORMATIONAL scsiblade.in.quorum: The scsi-blade on
this node established quorum with the other nodes in the cluster.
8/27/2013 16:21:17 node2          INFORMATIONAL scsiblade.in.quorum: The scsi-blade on
this node established quorum with the other nodes in the cluster.

```

- Use the `spm show -state !running` command to verify that there are no processes running:

#### Example

```

cluster::*> spm show -state !running
There are no entries matching your query.

```

- Repeat Steps 14 through 20 for the remaining nodes on the cluster, starting with `node1`'s HA partner.
- Use the `system node run ifstat` command to verify that both ports `e1a` and `e2a` are transmitting and receiving traffic.



## Example

Here is the portion of an example from the output of the `system node run ifstat` on ports `e1a` and `e2a`:

```
cluster::*> system node run -node node1 ifstat e1a
4 entries were acted on.
Node: node1
-- interface e1a (13 days, 21 hours, 41 minutes, 53 seconds) --

RECEIVE
Frames/second:    9953 | Bytes/second:    64604k | Errors/minute:    0
Discards/minute: 0 | Total frames:    2565m | Total bytes:      18894g
Total errors:    0 | Total discards: 0 | Multi/broadcast: 108k
No buffers:      0 | Non-primary u/c: 0 | LRO segments:    0
LRO bytes:       0 | Tag drop:        0 | Vlan tag drop:   0
Vlan untag drop: 0 | Vlan forwards:  0 | Vlan broadcasts: 0
.
.
.
cluster::*> system node run -node node1 ifstat e2a
1 entry was acted on.

Node: node1

-- interface e2a (13 days, 21 hours, 42 minutes, 7 seconds) --

RECEIVE
Frames/second:    9937 | Bytes/second:    64603k | Errors/minute:    0
Discards/minute: 0 | Total frames:    2579m | Total bytes:      18973g
Total errors:    0 | Total discards: 0 | Multi/broadcast: 83599
No buffers:      0 | Non-primary u/c: 0 | LRO segments:    0
LRO bytes:       0 | Tag drop:        0 | Vlan tag drop:   0
Vlan untag drop: 0 | Vlan forwards:  0 | Vlan broadcasts: 0
Vlan unicasts:   0 | Mac octets:     18984g | UCast pkts:      2579m
```

23. Before you begin the second part of the procedure, use `run -node local cdpd show-neighbors` to view data about the configuration of the cluster switch configuration.

## Example

This example shows sample output from the command:

```
cluster::*> run -node local cdpd show-neighbors
Local Remote Remote Remote Hold Remote
Port Device Interface Platform Time Capability
-----
e1a icea-4(SSII1329. Ethernet1/21 N5K-C5020P-BFS 170 SI
e1b icea-3(SSII1410. Ethernet1/21 N5K-C5020P-BF 157 SI
```

24. Use the statistics `show-periodic` command to display information about the current cluster traffic and data.

## Example

This example shows sample output from the command (the output has been truncated to fit this page):

```
cluster::*> set -privilege advanced
cluster::*> statistics show-periodic -node node1 -summary true -iterations 10
node1: cluster.cluster: 1/7/2014 09:10:49
cpu cpu total fcache total total data ...
avg busy ops nfs-ops cifs-ops ops spin-ops recv sent busy ...
-----
0% 15% 0 0 0 0 0 5.62MB 24.7KB 0% ...
0% 15% 0 0 0 0 0 5.26MB 31.4KB 0% ...
0% 19% 0 0 0 0 0 5.35MB 27.8KB 0% ...
0% 18% 0 0 0 0 0 5.23MB 15.4KB 0% ...
```

```

0% 19%      0      0      0      0      0      0  5.29MB  15.8KB  0% ...
0% 19%      0      0      0      0      0      0  5.31MB  18.1KB  0% ...
0% 20%      0      0      0      0      0      0  5.27MB  21.5KB  0% ...
0% 17%      0      0      0      0      0      0  5.27MB  21.1KB  0% ...
0% 17%      0      0      0      0      0      0  5.33MB  38.6KB  0% ...
0% 15%      0      0      0      0      0      0  5.33MB  32.1KB  0% ...
node1: cluster.cluster: 1/7/2014 09:11:10
  cpu  cpu      total      fcache      total      total data ...
  avg  busy     ops    nfs-ops  cifs-ops    ops  spin-ops  recv    sent  busy ...
-----
Minimums:
  0% 15%      0      0      0      0      0      0  5.23MB  15.4KB  0% ...
Averages for 10 samples:
  0% 17%      0      0      0      0      0      0  5.33MB  24.7KB  0% ...
Maximums:
  0% 20%      0      0      0      0      0      0  5.62MB  38.6KB  0% ...

```

25. To begin the second part of the switch replacement process, use node1's console to migrate cluster traffic off the cluster port on c5010-2 to the cluster port on c5020-1, and confirm that traffic is not flowing through e2a but through e1a, and that the cluster LIF has migrated.

Monitor the traffic as described in Step 5.

### Example

```

cluster::*> network interface migrate -vserver node1 -lif clus2 -
source-node node1 -dest-node node1 -dest-port e1a

cluster::*> system node run -node node1 ifstat ifstat e1a
4 entries were acted on.
Node: node1
-- interface e1a (13 days, 21 hours, 41 minutes, 53 seconds) --

RECEIVE
Frames/second:    9953 | Bytes/second:    64604k | Errors/minute:    0
Discards/minute: 0 | Total frames:    2565m | Total bytes:      18894g
Total errors:    0 | Total discards: 0 | Multi/broadcast: 108k
No buffers:      0 | Non-primary u/c: 0 | LRO segments:    0
LRO bytes:       0 | Tag drop:        0 | Vlan tag drop:   0
Vlan untag drop: 0 | Vlan forwards:  0 | Vlan broadcasts: 0
.
.
.
cluster::*> system node run -node node1 ifstat ifstat e2a
1 entry was acted on.

Node: node1

-- interface e2a (13 days, 21 hours, 42 minutes, 7 seconds) --

RECEIVE
Frames/second:    9937 | Bytes/second:    64603k | Errors/minute:    0
Discards/minute: 0 | Total frames:    2579m | Total bytes:      18973g
Total errors:    0 | Total discards: 0 | Multi/broadcast: 83599
No buffers:      0 | Non-primary u/c: 0 | LRO segments:    0
LRO bytes:       0 | Tag drop:        0 | Vlan tag drop:   0
Vlan untag drop: 0 | Vlan forwards:  0 | Vlan broadcasts: 0
Vlan unicasts:   0 | Mac octets:     18984g | UCast pkts:      2579m

cluster::*> network interface show -role cluster

Vserver      Logical      Status      Network      Current      Current Is
Interface    Admin/Oper  Address/Mask Node          Port         Home
-----
node1
      clus1      up/up      10.10.0.1/24  node1      e1a         true

```

```

node2      clus2      up/up      10.10.0.2/24      node1      e1a      false
node2      clus1      up/up      10.10.0.3/24      node2      e1a      true
node2      clus2      up/up      10.10.0.4/24      node2      e2a      true
4 entries were displayed.

```

26. Shut down the port that was just migrated in Step 25 and confirm the action has been completed.

### Example

This example shuts down port e2a on node1 and verifies that it is shut down:

```

cluster::*> network port modify -node node1 -port e2a -up-admin false

cluster::*> network port show -role cluster

```

Node	Port	Role	Link	MTU	Auto-Negot Admin/Oper	Duplex Admin/Oper	Speed (Mbps) Admin/Oper
node1	e1a	cluster	up	9000	true/true	full/full	auto/10000
	e2a	cluster	down	9000	true/true	full/full	auto/10000
node2	e1a	cluster	up	9000	true/true	full/full	auto/10000
	e2a	cluster	up	9000	true/true	full/full	auto/10000

4 entries were displayed.

27. Repeat Steps 25 and 26 on the remaining nodes on the cluster, starting with node1's HA partner.
28. At this time, all traffic should be running on locally on the c5020-1 switch; now disable the ISL ports on the switch. Monitor traffic on all nodes as described in Step 5.

### Example

The commands shown in the following example configure and disable ports 33 through 40 on c5020-1:

```

c5020-1# configure
c5020-1(config)#interface ethernet 1/33-40
c5020-1(config-if-range)#shutdown

```

29. Remove all node-facing ISL cables from c5010-2 and connect them to c5020-2. Connect the node-facing ports to the same ports, and the new ISL should be on ports 33 to 40 on both c5020-1 and c5020-2.
30. Enable the ISL ports on the c5020-1 switch (ports 33 to 40) and verify that the ISLs are enabled between the c5020-1 and c5020-2 switches.

### Example

This example shows ports 33 through 40 being enabled on the c5020-2 switch, and verifies that the ISL being verified as functioning on both the c5010-2 and c5020-1 switches.

```

c5020-1# configure
c5020-1(config)#interface ethernet 1/33-40
c5020-1(config-if-range)#no shutdown

c5020-1#show port-channel summary

```

Flags: D - Down            P - Up in port-channel (members)  
I - Individual        H - Hot-standby (LACP only)  
s - Suspended        r - Module-removed  
S - Switched        R - Routed  
U - Up (port-channel)  
M - Not in use. Min-links not met

```

-----
Group Port-      Type      Protocol  Member Ports
Channel
-----
1      Pol(SU)    Eth       LACP      Eth1/33(P)  Eth1/34(P)  Eth1/35(P)
                                Eth1/36(P)  Eth1/37(P)  Eth1/38(P)
                                Eth1/39(P)  Eth1/40(P)

c5020-2#show port-channel summary

Flags:  D - Down          P - Up in port-channel (members)
        I - Individual   H - Hot-standby (LACP only)
        s - Suspended    r - Module-removed
        S - Switched     R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met

-----
Group Port-      Type      Protocol  Member Ports
Channel
-----
1      Pol(SU)    Eth       LACP      Eth1/33(P)  Eth1/34(P)  Eth1/35(P)
                                Eth1/36(P)  Eth1/37(P)  Eth1/38(P)
                                Eth1/39(P)  Eth1/40(P)

```

31. Enable the node-facing ports on the c5020-2 switch.

#### Example

The following commands configure and enable ports 1 through 32 on c5020-2:

```

c5020-2# configure
c5020-2(config)#interface ethernet 1/1-32
c5020-2(config-if-range)#no shutdown

```

32. Use the network interface show -role cluster -fields home-port command to display information about the LIFs that have the home port or interface group that you specify.

#### Example

This example shows port e1a on node1 being shut down and shows that action being verified:

```

cluster::*> network interface show -role cluster -fields home-port
Vserver    lif    home-port
-----
node1      clus1  e1a
node1      clus2  e2a
node2      clus1  e1a
node2      clus2  e2a

```

33. On node1, enable port e2a and revert traffic back to its home cluster port. Verify that this process is successful and that both cluster ports are transmitting and receiving. There is a chance that cluster traffic may take a new route other than the route previously taken.

#### Example

```

cluster::*> network port modify -node node1 -port e2a -up-admin true
cluster::*> network interface revert -vserver node1 -lif clus2

cluster::*> network port show -role cluster

Node  Port  Role      Link  MTU  Auto-Negot  Duplex  Speed (Mbps)
-----
node1  e1a  cluster  up    9000 true/true   full/full  auto/10000

```

```

node2    e2a    cluster    up    9000    true/true    full/full    auto/10000
node2    e1a    cluster    up    9000    true/true    full/full    auto/10000
node2    e2a    cluster    up    9000    true/true    full/full    auto/10000
4 entries were displayed.

cluster::*> network interface show -role cluster

Vserver      Logical      Status      Network      Current      Current      Is
Interface    Admin/Oper   Address/Mask Node          Port         Home
-----
node1
           clus1      up/up      10.10.0.1/24  node1        e1a          true
           clus2      up/up      10.10.0.2/24  node1        e2a          true
node2
           clus1      up/up      10.10.0.3/24  node2        e1a          true
           clus2      up/up      10.10.0.4/24  node2        e2a          true
4 entries were displayed.

cluster::*> system node run -node node1 ifstat e1a
4 entries were acted on.
Node: node1
-- interface e1a (13 days, 21 hours, 41 minutes, 53 seconds) --

RECEIVE
Frames/second:    9953 | Bytes/second:    64604k | Errors/minute:    0
Discards/minute: 0 | Total frames:    2565m | Total bytes:      18894g
Total errors:    0 | Total discards: 0 | Multi/broadcast: 108k
No buffers:      0 | Non-primary u/c: 0 | LRO segments:    0
LRO bytes:       0 | Tag drop:        0 | Vlan tag drop:   0
Vlan untag drop: 0 | Vlan forwards:   0 | Vlan broadcasts: 0
.
.
.
cluster::*> system node run -node node1 ifstat e2a
1 entry was acted on.
Node: node1
-- interface e2a (13 days, 21 hours, 42 minutes, 7 seconds) --

RECEIVE
Frames/second:    9937 | Bytes/second:    64603k | Errors/minute:    0
Discards/minute: 0 | Total frames:    2579m | Total bytes:      18973g
Total errors:    0 | Total discards: 0 | Multi/broadcast: 83599
No buffers:      0 | Non-primary u/c: 0 | LRO segments:    0
LRO bytes:       0 | Tag drop:        0 | Vlan tag drop:   0
Vlan untag drop: 0 | Vlan forwards:   0 | Vlan broadcasts: 0
Vlan unicasts:   0 | Mac octets:      18984g | UCast pkts:      2579m
.
.
.

```

34. Use `cluster show` to display information about the nodes in the cluster.

#### Example

This example shows that the node health for node1 and node2 in this cluster is true:

```

cluster::*> cluster show
Node      Health Eligibility  Epsilon
-----
node1     true   true           false
node2     true   true           false

```

35. Use the event `log show` command with the `-messagename` parameter to display quorum in the latest event logs:

## Example

This example displays the event log output for node1 and node2 with the quorum status:

```
cluster::*> event log show -messagename *quorum* | *offline*
Time          Node          Severity      Event
-----
8/27/2013 16:21:26 node1          INFORMATIONAL scsiblade.in.quorum: The scsi-blade on
this node established quorum with the other nodes in the cluster.
8/27/2013 16:21:17 node2          INFORMATIONAL scsiblade.in.quorum: The scsi-blade on
this node established quorum with the other nodes in the cluster.
```

36. Use the `spm show -state !running` command to verify that there are no processes running:

## Example

```
cluster::*> spm show -state !running
There are no entries matching your query.
```

37. Repeat Steps 33 through 36 for the remaining nodes on the cluster starting with node1's HA partner.
38. To verify the results of the procedure, use `run -node local cdpd show-neighbors` to view data about the configuration of the cluster switch configuration.

## Example

This example shows sample output from the command:

```
cluster::*> run -node local cdpd show-neighbors
Local Remote      Remote      Remote      Hold Remote
Port Device      Interface   Platform    Time Capability
-----
el1a  icea-4(SS11329. Ethernet1/21 N5K-C5020P-BFS 170 SI
el1b  icea-3(SS11410. Ethernet1/21 N5K-C5020P-BF 157 SI
```

39. Switch to advanced mode and use the `statistics show-periodic` command to display information about the current cluster traffic and data.

## Example

This example shows sample output from the command (the output has been truncated to fit on this page):

```
cluster::*> set -privilege advanced
cluster::*> statistics show-periodic -node node1 -summary true -iterations 10
node1: cluster.cluster: 1/7/2014 09:10:49
  cpu  cpu    total          fcache          total    total data ...
  avg busy  ops  nfs-ops cifs-ops  ops spin-ops  recv  sent busy ...
-----
  0% 15%    0      0      0      0      0      5.62MB 24.7KB 0% ...
  0% 15%    0      0      0      0      0      5.26MB 31.4KB 0% ...
  0% 19%    0      0      0      0      0      5.35MB 27.8KB 0% ...
  0% 18%    0      0      0      0      0      5.23MB 15.4KB 0% ...
  0% 19%    0      0      0      0      0      5.29MB 15.8KB 0% ...
  0% 19%    0      0      0      0      0      5.31MB 18.1KB 0% ...
  0% 20%    0      0      0      0      0      5.27MB 21.5KB 0% ...
  0% 17%    0      0      0      0      0      5.27MB 21.1KB 0% ...
  0% 17%    0      0      0      0      0      5.33MB 38.6KB 0% ...
  0% 15%    0      0      0      0      0      5.33MB 32.1KB 0% ...
node1: cluster.cluster: 1/7/2014 09:11:10
  cpu  cpu    total          fcache          total    total data ...
  avg busy  ops  nfs-ops cifs-ops  ops spin-ops  recv  sent busy ...
-----
Minimums:
```

0%	15%	0	0	0	0	0	5.23MB	15.4KB	0%	...
Averages for 10 samples:										
0%	17%	0	0	0	0	0	5.33MB	24.7KB	0%	...
Maximums:										
0%	20%	0	0	0	0	0	5.62MB	38.6KB	0%	...

## Related information

[Cisco Ethernet Switch description page](#)

## 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

## 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](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.