



Cisco® Nexus Switches

Migrating to a Two-Node Switched Cluster

For Cisco Nexus 92300YC switches

July 2019 | 215-13965_BO
doccomments@netapp.com

Contents

Migrating to a two-node switched cluster with Cisco Nexus 92300YC switches	4
Migrating to a switched NetApp cluster environment using Cisco Nexus 92300YC switches	4
Copyright	12
Trademark	13
How to send comments about documentation and receive update notifications	14
Index	15

Migrating to a two-node switched cluster with Cisco Nexus 92300YC switches

You must be aware of certain configuration information, port connections, and cabling requirements when you migrate a two-node switchless cluster, non-disruptively, to a cluster with Cisco Nexus 92300YC cluster switches. The procedure you use depends on whether you have two dedicated cluster-network ports on each controller or a single cluster port on each controller. The process documented works for all nodes using optical or twinax ports but is not supported on this switch if nodes are using onboard 10Gb BASE-T RJ45 ports for the cluster-network ports.

Most systems require two dedicated cluster-network ports on each controller.

Note: After your migration completes, you might need to install the required configuration file to support the Cluster Switch Health Monitor (CSHM) for 92300YC cluster switches. See *Installing the Cluster Switch Health Monitor (CSHM) configuration file for 92300YC switches* in the [Setting up](#) guide.

Migrating to a switched NetApp cluster environment using Cisco Nexus 92300YC switches

If you have an existing two-node switchless cluster environment, you can migrate to a two-node switched cluster environment using Cisco Nexus 92300YC switches to enable you to scale beyond two nodes in the cluster.

Before you begin

Two-node switchless configuration:

- The two-node switchless configuration must be properly set up and functioning.
- The nodes must be running ONTAP 9.6 and later.
- All cluster ports must be in the **up** state.
- All cluster logical interfaces (LIFs) must be in the **up** state and on their home ports.

Cisco Nexus 92300YC switch configuration:

- Both switches must have management network connectivity.
- There must be console access to the cluster switches.
- Nexus 92300YC node-to-node switch and switch-to-switch connections must use twinax or fiber cables.

The *NetApp Hardware Universe* contains more information about cabling.

[Hardware Universe - Switches](#)

- Inter-Switch Link (ISL) cables must be connected to ports 1/65 and 1/66 on both 92300YC switches.
- Initial customization of both the 92300YC switches must be completed. So that the:
 - 92300YC switches are running the latest version of software
 - Reference Configuration Files (RCFs) have been applied to the switches

Any site customization, such as SMTP, SNMP, and SSH must be configured on the new switches.

About this task

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

- The names of the 92300YC switches are *cs1* and *cs2*.
- The names of the cluster SVMs are *node1* and *node2*.
- The names of the LIFs are *node1_clus1* and *node1_clus2* on node 1, and *node2_clus1* and *node2_clus2* on node 2 respectively.
- The *cluster1::*>* prompt indicates the name of the cluster.
- The cluster ports used in this procedure are *e0a* and *e0b*.

The *Hardware Universe* contains the latest information about the actual cluster ports for your platforms.

Hardware Universe

Steps

1. Change the privilege level to advanced, entering **y** when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (**>*) appears.

2. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=xh
```

where *x* is the duration of the maintenance window in hours.

Note: The AutoSupport message notifies technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

Example

The following command suppresses automatic case creation for two hours:

```
cluster1::*> system node autosupport invoke -node * -type all -message MAINT=2h
```

3. Disable all node-facing ports (not ISL ports) on both the new cluster switches *cs1* and *cs2*.

You must not disable the ISL ports.

Example

The following example shows that node-facing ports 1 through 64 are disabled on switch *cs1*:

```
cs1# config
Enter configuration commands, one per line. End with CNTL/Z.
cs1(config)# interface e1/1-64
cs1(config-if-range)# shutdown
```

4. Verify that the ISL and the physical ports on the ISL between the two 92300YC switches *cs1* and *cs2* are **up** on ports 1/65 and 1/66:

```
show port-channel summary
```

Example

The following example shows that the ISL ports are **up** on switch *cs1*:

```
cs1# 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
       b - BFD Session Wait
       S - Switched    R - Routed
       U - Up (port-channel)
       p - Up in delay-lacp mode (member)
       M - Not in use. Min-links not met
-----
Group Port-      Type      Protocol Member Ports
      Channel
-----
1    Po1(SU)     Eth       LACP      Eth1/65(P)  Eth1/66(P)
```

The following example shows that the ISL ports are **up** on switch cs2 :

```
(cs2)# 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
       b - BFD Session Wait
       S - Switched    R - Routed
       U - Up (port-channel)
       p - Up in delay-lacp mode (member)
       M - Not in use. Min-links not met
-----
Group Port-      Type      Protocol Member Ports
      Channel
-----
1    Po1(SU)     Eth       LACP      Eth1/65(P)  Eth1/66(P)
```

5. Display the list of neighboring devices:

show cdp neighbors

This command provides information about the devices that are connected to the system.

Example

The following example lists the neighboring devices on switch cs1:

```
cs1# show cdp neighbors

Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute

Device-ID      Local Intrfce  Hldtme Capability  Platform      Port ID
cs2(FDO220329V5)  Eth1/65      175    R S I s  N9K-C92300YC  Eth1/65
cs2(FDO220329V5)  Eth1/66      175    R S I s  N9K-C92300YC  Eth1/66

Total entries displayed: 2
```

The following example lists the neighboring devices on switch cs2:

```
cs2# show cdp neighbors

Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute

Device-ID      Local Intrfce  Hldtme Capability  Platform      Port ID
cs1(FDO220329KU)  Eth1/65      177    R S I s  N9K-C92300YC  Eth1/65
cs1(FDO220329KU)  Eth1/66      177    R S I s  N9K-C92300YC  Eth1/66

Total entries displayed: 2
```

6. Verify that all cluster ports are **up:**

network port show -ipspace Cluster

Each port should display **up** for Link and **healthy** for Health Status.

Example

```
cluster1::*> network port show -ipspace Cluster

Node: node1

  Port      IPspace      Broadcast Domain Link MTU      Speed(Mbps)  Health
  -----  -----  -----  -----  -----  -----  Admin/Oper  Status
  e0a      Cluster      Cluster      up    9000  auto/10000  healthy
  e0b      Cluster      Cluster      up    9000  auto/10000  healthy

Node: node2

  Port      IPspace      Broadcast Domain Link MTU      Speed(Mbps)  Health
  -----  -----  -----  -----  -----  -----  Admin/Oper  Status
  e0a      Cluster      Cluster      up    9000  auto/10000  healthy
  e0b      Cluster      Cluster      up    9000  auto/10000  healthy

  4 entries were displayed.
```

7. Verify that all cluster LIFs are **up** and operational:

```
network interface show -vserver Cluster
```

Each cluster LIF should display **true** for Is Home and have a Status Admin/Oper of **up/up**

Example

```
cluster1::*> network interface show -vserver Cluster

  Vserver      Logical      Status      Network      Current      Current      Is
  -----  Interface  Admin/Oper  Address/Mask  Node      Port      Home
  -----  -----  -----  -----  -----  -----  -----
  Cluster
    node1_clus1  up/up    169.254.209.69/16  node1      e0a      true
    node1_clus2  up/up    169.254.49.125/16  node1      e0b      true
    node2_clus1  up/up    169.254.47.194/16  node2      e0a      true
    node2_clus2  up/up    169.254.19.183/16  node2      e0b      true

  4 entries were displayed.
```

8. Verify that auto-revert is enabled on all cluster LIFs:

```
network interface show -vserver Cluster -fields auto-revert
```

Example

```
cluster1::*> network interface show -vserver Cluster -fields auto-revert

  Vserver      Logical      Auto-revert
  -----  Interface  -----
  Cluster
    node1_clus1  true
    node1_clus2  true
    node2_clus1  true
    node2_clus2  true

  4 entries were displayed.
```

9. Disconnect the cable from cluster port e0a on node1, and then connect e0a to port 1 on cluster switch cs1, using the appropriate cabling supported by the 92300YC switches.

The *NetApp Hardware Universe* contains more information about cabling.

[Hardware Universe - Switches](#)

10. Disconnect the cable from cluster port e0a on node2, and then connect e0a to port 2 on cluster switch cs1, using the appropriate cabling supported by the 92300YC switches.
11. Enable all node-facing ports on cluster switch cs1.

Example

The following example shows that ports 1/1 through 1/64 are enabled on switch cs1:

```
cs1# config
Enter configuration commands, one per line. End with CNTL/Z.
cs1(config)# interface e1/1-64
cs1(config-if-range)# no shutdown
```

12. Verify that all cluster LIFs are **up**, operational, and display as **true** for **Is Home**:

```
network interface show -vserver Cluster
```

Example

The following example shows that all of the LIFs are **up** on node1 and node2 and that **Is Home** results are **true**:

```
cluster1::*> network interface show -vserver Cluster
      Logical          Status      Network          Current      Current  Is
Vserver  Interface    Admin/Oper  Address/Mask  Node       Port      Home
-----  -----
Cluster
      node1_clus1  up/up      169.254.209.69/16  node1      e0a      true
      node1_clus2  up/up      169.254.49.125/16  node1      e0b      true
      node2_clus1  up/up      169.254.47.194/16  node2      e0a      true
      node2_clus2  up/up      169.254.19.183/16  node2      e0b      true

4 entries were displayed.
```

13. Display information about the status of the nodes in the cluster:

```
cluster show
```

Example

The following example displays information about the health and eligibility of the nodes in the cluster:

```
cluster1::*> cluster show
      Node          Health  Eligibility  Epsilon
-----  -----
node1        true      true        false
node2        true      true        false

2 entries were displayed.
```

14. Disconnect the cable from cluster port e0b on node1, and then connect e0b to port 1 on cluster switch cs2, using the appropriate cabling supported by the 92300YC switches.
15. Disconnect the cable from cluster port e0b on node2, and then connect e0b to port 2 on cluster switch cs2, using the appropriate cabling supported by the 92300YC switches.
16. Enable all node-facing ports on cluster switch cs2.

Example

The following example shows that ports 1/1 through 1/64 are enabled on switch cs2:

```
cs2# config
Enter configuration commands, one per line. End with CNTL/Z.
cs2(config)# interface e1/1-64
cs2(config-if-range)# no shutdown
```

17. Verify that all cluster ports are **up**:

```
network port show -ipspace Cluster
```

Example

The following example shows that all of the cluster ports are **up** on node1 and node2:

```
cluster1::*> network port show -ipspace Cluster

Node: node1
-----+
Port      IPspace      Broadcast Domain Link MTU      Speed(Mbps) Health      Ignore
          Cluster       Cluster          up   9000  auto/10000 healthy    false
e0a      Cluster       Cluster          up   9000  auto/10000 healthy    false
e0b      Cluster       Cluster          up   9000  auto/10000 healthy    false

Node: node2
-----+
Port      IPspace      Broadcast Domain Link MTU      Speed(Mbps) Health      Ignore
          Cluster       Cluster          up   9000  auto/10000 healthy    false
e0a      Cluster       Cluster          up   9000  auto/10000 healthy    false
e0b      Cluster       Cluster          up   9000  auto/10000 healthy    false

4 entries were displayed.
```

18. Verify that all interfaces display **true** for **Is Home**:

```
network interface show -vserver Cluster
```

Note: This might take several minutes to complete.

Example

The following example shows that all LIFs are **up** on node1 and node2 and that **Is Home** results are **true**:

```
cluster1::*> network interface show -vserver Cluster
-----+
Vserver  Logical      Status      Network          Current      Current  Is
          Interface    Admin/Oper  Address/Mask    Node        Port     Home
-----+
Cluster
          node1_clus1  up/up      169.254.209.69/16  node1      e0a      true
          node1_clus2  up/up      169.254.49.125/16  node1      e0b      true
          node2_clus1  up/up      169.254.47.194/16  node2      e0a      true
          node2_clus2  up/up      169.254.19.183/16  node2      e0b      true

4 entries were displayed.
```

19. Verify that both nodes each have one connection to each switch:

```
show cdp neighbors
```

Example

The following example shows the appropriate results for both switches:

```
(csl)# show cdp neighbors
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  S - Supports-STP-Dispute

Device-ID      Local Intrfce  Hldtme Capability  Platform      Port ID
node1          Eth1/1      133      H          FAS2980      e0a
node2          Eth1/2      133      H          FAS2980      e0a
cs2(FDO220329V5)  Eth1/65    175      R S I s  N9K-C92300YC  Eth1/65
cs2(FDO220329V5)  Eth1/66    175      R S I s  N9K-C92300YC  Eth1/66

Total entries displayed: 4
```

```
(cs2)# show cdp neighbors

Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute

Device-ID      Local Intrfce  Hldtme  Capability  Platform    Port ID
node1          Eth1/1      133      H          FAS2980    e0b
node2          Eth1/2      133      H          FAS2980    e0b
cs1(FDO220329KU)  Eth1/65    175      R S I s  N9K-C92300YC  Eth1/65
cs1(FDO220329KU)  Eth1/66    175      R S I s  N9K-C92300YC  Eth1/66

Total entries displayed: 4
```

20. Display information about the discovered network devices in your cluster:

```
network device-discovery show -protocol cdp
```

Example

```
cluster1::*> network device-discovery show -protocol cdp
Node/      Local   Discovered
Protocol   Port    Device (LLDP: ChassisID)  Interface      Platform
-----
node2      /cdp    e0a      cs1                  0/2          N9K-C92300YC
          /cdp    e0b      cs2                  0/2          N9K-C92300YC
node1      /cdp    e0a      cs1                  0/1          N9K-C92300YC
          /cdp    e0b      cs2                  0/1          N9K-C92300YC

4 entries were displayed.
```

21. Verify that the settings are disabled:

```
network options switchless-cluster show
```

Note: It might take several minutes for the command to complete. Wait for the '3 minute lifetime to expire' announcement.

Example

The false output in the following example shows that the configuration settings are disabled:

```
cluster1::*> network options switchless-cluster show
Enable Switchless Cluster: false
```

22. Verify the status of the node members in the cluster:

```
cluster show
```

Example

The following example shows information about the health and eligibility of the nodes in the cluster:

```
cluster1::*> cluster show
Node          Health  Eligibility  Epsilon
-----
node1        true    true        false
node2        true    true        false
```

23. Ensure that the cluster network has full connectivity:

```
cluster ping-cluster -node node-name
```

Example

```
cluster1::*> cluster ping-cluster -node node2
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:
.....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)
```

24. If you suppressed automatic case creation, reenable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Example

```
cluster1::*> system node autosupport invoke -node * -type all -message MAINT=END
```

25. Change the privilege level back to admin:

```
set -privilege admin
```

Copyright

Copyright © 2019 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

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.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account [@NetAppDoc](#).

You can also contact us in the following ways:

- NetApp, Inc., 1395 Crossman Ave., Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

92300YC switches

- migrating from a two-node switchless cluster environment to a two-node switched cluster environment [4](#)

C

cluster environments

- migrating from a switchless two-node cluster to a two-node switched 92300YC cluster environment [4](#)

comments

- how to send feedback about documentation [14](#)

D

documentation

- how to receive automatic notification of changes to [14](#)
- how to send feedback about [14](#)

E

environments

- migrating from a switchless cluster to a two-node switched 92300YC cluster environment [4](#)

F

feedback

- how to send comments about documentation [14](#)

I

information

- how to send feedback about improving documentation [14](#)

S

suggestions

- how to send feedback about documentation [14](#)

switches

- migrating from a two-node switchless cluster environment to a two-node 92300YC switched cluster environment [4](#)

switchless cluster environments

- migrating to a two-node switched 92300YC cluster environment [4](#)

T

Twitter

- how to receive automatic notification of documentation changes [14](#)

two-node switchless cluster environments

- migrating to a two-node switched 92300YC cluster environment [4](#)

two-node switchless clusters

- migrating to a cluster with Cisco Nexus 92300YC switches [4](#)