



Replace switches

Cluster and storage switches

NetApp
April 05, 2024

Table of Contents

- Replace switches 1
 - Replace a NetApp CN1610 cluster switch 1
 - Replace NetApp CN1610 cluster switches with switchless connections 11

Replace switches

Replace a NetApp CN1610 cluster switch

Follow these steps to replace a defective NetApp CN1610 switch in a cluster network. This is a non-disruptive procedure (NDU).

What you'll need

Before you perform the switch replacement, the following conditions must exist before you perform the switch replacement in the current environment and on the replacement switch for existing cluster and network infrastructure:

- The existing cluster must be verified as completely functional, with at least one fully connected cluster switch.
- All of the cluster ports must be **up**.
- All of the cluster logical interfaces (LIFs) must be up and must not have been migrated.
- The ONTAP cluster `ping-cluster -node node1` command must indicate that basic connectivity and larger than PMTU communication are successful on all of the paths.

About this task

You must execute the command for migrating a cluster LIF from the node where the cluster LIF is hosted.

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

- The names of the two CN1610 cluster switches are `cs1` and `cs2`.
- The name of the CN1610 switch that is to be replaced (the defective switch) is `old_cs1`.
- The name of the new CN1610 switch (the replacement switch) is `new_cs1`.
- The name of the partner switch that is not being replaced is `cs2`.

Steps

1. Confirm that the startup configuration file matches the running configuration file. You must save these files locally for use during the replacement.

The configuration commands in the following example are for FASTPATH 1.2.0.7:

Show example

```
(old_cs1) >enable
(old_cs1) #show running-config
(old_cs1) #show startup-config
```

2. Create a copy of the running configuration file.

The command in the following example is for FASTPATH 1.2.0.7:

Show example

```
(old_cs1) #show running-config filename.scr  
Config script created successfully.
```



You can use any file name except CN1610_CS_RCF_v1.2.scr. The file name must have the .scr extension.

3. Save the running configuration file of the switch to an external host in preparation for the replacement.

Show example

```
(old_cs1) #copy nvram:script filename.scr  
scp://<Username>@<remote_IP_address>/path_to_file/filename.scr
```

4. Verify that the switch and ONTAP versions match in the compatibility matrix. See the [NetApp CN1601 and CN1610 Switches](#) page for details.
5. From the [Software Downloads page](#) on the NetApp Support Site, select NetApp Cluster Switches to download the appropriate RCF and FASTPATH versions.
6. Set up a Trivial File Transfer Protocol (TFTP) server with the FASTPATH, RCF, and saved configuration .scr file for use with the new switch.
7. Connect the serial port (the RJ-45 connector labeled “IOIOI” on the right side of the switch) to an available host with terminal emulation.
8. On the host, set the serial terminal connection settings:
 - a. 9600 baud
 - b. 8 data bits
 - c. 1 stop bit
 - d. parity: none
 - e. flow control: none
9. Connect the management port (the RJ-45 wrench port on the left side of the switch) to the same network where your TFTP server is located.
10. Prepare to connect to the network with the TFTP server.

If you are using Dynamic Host Configuration Protocol (DHCP), you do not have to configure an IP address for the switch at this time. The service port is set to use DHCP by default. The network management port is set to none for the IPv4 and IPv6 protocol settings. If your wrench port is connected to a network that has a DHCP server, then the server settings are configured automatically.

To set a static IP address, you should use the serviceport protocol, network protocol, and serviceport ip commands.

Show example

```
(new_cs1) #serviceport ip <ipaddr> <netmask> <gateway>
```

11. Optionally, if the TFTP server is on a laptop, then connect the CN1610 switch to the laptop by using a standard Ethernet cable, and then configure its network port in the same network with an alternate IP address.

You can use the `ping` command to verify the address. If you are unable to establish the connectivity, you should use a nonrouted network, and configure the service port using IP 192.168.x or 172.16.x. You can reconfigure the service port to the production management IP address at a later date.

12. Optionally, verify and install the appropriate versions of the RCF and FASTPATH software for the new switch. If you have verified that the new switch is correctly set up and does not require updates to the RCF and FASTPATH software, you should go to step 13.
 - a. Verify the new switch settings.

Show example

```
(new_cs1) >*enable*  
(new_cs1) #show version
```

- b. Download the RCF to the new switch.

Show example

```
(new_cs1) #copy tftp://<server_ip_address>/CN1610_CS_RCF_v1.2.txt
nvram:script CN1610_CS_RCF_v1.2.scr
Mode.      TFTP
Set Server IP. 172.22.201.50
Path.      /
Filename.....
CN1610_CS_RCF_v1.2.txt
Data Type..... Config Script
Destination Filename.....
CN1610_CS_RCF_v1.2.scr
File with same name already exists.
WARNING:Continuing with this command will overwrite the existing
file.

Management access will be blocked for the duration of the
transfer Are you sure you want to start? (y/n) y

File transfer in progress. Management access will be blocked for
the duration of the transfer. please wait...
Validating configuration script...
(the entire script is displayed line by line)
...
description "NetApp CN1610 Cluster Switch RCF v1.2 - 2015-01-13"
...
Configuration script validated.
File transfer operation completed successfully.
```

c. Verify that the RCF is downloaded to the switch.

Show example

```
(new_cs1) #script list
Configuration Script Nam    Size(Bytes)
-----
CN1610_CS_RCF_v1.1.scr      2191
CN1610_CS_RCF_v1.2.scr      2240
latest_config.scr           2356

4 configuration script(s) found.
2039 Kbytes free.
```

13. Apply the RCF to the CN1610 switch.

Show example

```
(new_cs1) #script apply CN1610_CS_RCF_v1.2.scr
Are you sure you want to apply the configuration script? (y/n) y
...
(the entire script is displayed line by line)
...
description "NetApp CN1610 Cluster Switch RCF v1.2 - 2015-01-13"
...
Configuration script 'CN1610_CS_RCF_v1.2.scr' applied. Note that the
script output will go to the console.
After the script is applied, those settings will be active in the
running-config file. To save them to the startup-config file, you
must use the write memory command, or if you used the reload answer
yes when asked if you want to save the changes.
```

- a. Save the running configuration file so that it becomes the startup configuration file when you reboot the switch.

Show example

```
(new_cs1) #write memory
This operation may take a few minutes.
Management interfaces will not be available during this time.

Are you sure you want to save? (y/n) y

Config file 'startup-config' created successfully.

Configuration Saved!
```

- b. Download the image to the CN1610 switch.

Show example

```
(new_cs1) #copy
tftp://<server_ip_address>/NetApp_CN1610_1.2.0.7.stk active
Mode.      TFTP
Set Server IP.  tftp_server_ip_address
Path.        /
Filename.....
NetApp_CN1610_1.2.0.7.stk
Data Type.   Code
Destination Filename.  active

Management access will be blocked for the duration of the
transfer

Are you sure you want to start? (y/n) y

TFTP Code transfer starting...

File transfer operation completed successfully.
```

- c. Run the new active boot image by rebooting the switch.

The switch must be rebooted for the command in step 6 to reflect the new image. There are two possible views for a response that you might see after you enter the reload command.

Show example

```
(new_cs1) #reload
The system has unsaved changes.
Would you like to save them now? (y/n) y

Config file 'startup-config' created successfully.

Configuration Saved! System will now restart!
.
.
.
Cluster Interconnect Infrastructure

User:admin Password: (new_cs1) >*enable*
```

- d. Copy the saved configuration file from the old switch to the new switch.

Show example

```
(new_cs1) #copy tftp://<server_ip_address>/<filename>.scr  
nvram:script <filename>.scr
```

- e. Apply the previously saved configuration to the new switch.

Show example

```
(new_cs1) #script apply <filename>.scr  
Are you sure you want to apply the configuration script? (y/n) y  
  
The system has unsaved changes.  
Would you like to save them now? (y/n) y  
  
Config file 'startup-config' created successfully.  
  
Configuration Saved!
```

- f. Save the running configuration file to the startup configuration file.

Show example

```
(new_cs1) #write memory
```

14. 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`

x is the duration of the maintenance window in hours.



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

15. On the new switch new_cs1, log in as the admin user, and shut down all of the ports that are connected to the node cluster interfaces (ports 1 through 12).

Show example

```
User:*admin*
Password:
(new_cs1) >*enable*
(new_cs1) #

(new_cs1) config
(new_cs1) (config) interface 0/1-0/12
(new_cs1) (interface 0/1-0/12) shutdown
(new_cs1) (interface 0/1-0/12) exit
(new_cs1) #write memory
```

16. Migrate the cluster LIFs from the ports that are connected to the old_cs1 switch.

You must migrate each cluster LIF from its current node's management interface.

Show example

```
cluster::> set -privilege advanced
cluster::> network interface migrate -vserver <vserver_name> -lif
<Cluster_LIF_to_be_moved> - sourcenode <current_node> -dest-node
<current_node> -dest-port <cluster_port_that_is_UP>
```

17. Verify that all of the cluster LIFs have been moved to the appropriate cluster port on each node.

Show example

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

18. Shut down the cluster ports that are attached to the switch that you replaced.

Show example

```
cluster::*> network port modify -node <node_name> -port
<port_to_admin_down> -up-admin false
```

19. Verify the health of the cluster.

Show example

```
cluster::*> cluster show
```

20. Verify that the ports are down.

Show example

```
cluster::*> cluster ping-cluster -node <node_name>
```

21. On the switch cs2, shut down the ISL ports 13 through 16.

Show example

```
(cs2) config  
(cs2)(config) interface 0/13-0/16  
(cs2)(interface 0/13-0/16) #shutdown  
(cs2) #show port-channel 3/1
```

22. Verify whether the storage administrator is ready for the replacement of the switch.
23. Remove all of the cables from the old_cs1 switch, and then connect the cables to the same ports on the new_cs1 switch.
24. On the cs2 switch, bring up the ISL ports 13 through 16.

Show example

```
(cs2) config  
(cs2)(config) interface 0/13-0/16  
(cs2)(interface 0/13-0/16) #no shutdown
```

25. Bring up the ports on the new switch that are associated with the cluster nodes.

Show example

```
(cs2) config  
(cs2)(config) interface 0/1-0/12  
(cs2)(interface 0/13-0/16) #no shutdown
```

26. On a single node, bring up the cluster node port that is connected to the replaced switch, and then confirm that the link is up.

Show example

```
cluster::*> network port modify -node node1 -port  
<port_to_be_online> -up-admin true  
cluster::*> network port show -role cluster
```

27. Revert the cluster LIFs that are associated with the port in step 25 on the same node.

In this example, the LIFs on node1 are successfully reverted if the “Is Home” column is true.

Show example

```
cluster::*> network interface revert -vserver node1 -lif  
<cluster_lif_to_be_reverted>  
cluster::*> network interface show -role cluster
```

28. If the first node’s cluster LIF is up and is reverted to its home port, repeat steps 25 and 26 to bring up the cluster ports and to revert the cluster LIFs on the other nodes in the cluster.
29. Display information about the nodes in the cluster.

Show example

```
cluster::*> cluster show
```

30. Confirm that the startup configuration file and running configuration file are correct on the replaced switch. This configuration file should match the output in step 1.

Show example

```
(new_cs1) >*enable*  
(new_cs1) #show running-config  
(new_cs1) #show startup-config
```

31. If you suppressed automatic case creation, re-enable it by invoking an AutoSupport message:

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

Replace NetApp CN1610 cluster switches with switchless connections

You can migrate from a cluster with a switched cluster network to one where two nodes are directly connected for ONTAP 9.3 and later.

Review requirements

Guidelines

Review the following guidelines:

- Migrating to a two-node switchless cluster configuration is a nondisruptive operation. Most systems have two dedicated cluster interconnect ports on each node, but you can also use this procedure for systems with a larger number of dedicated cluster interconnect ports on each node, such as four, six or eight.
- You cannot use the switchless cluster interconnect feature with more than two nodes.
- If you have an existing two-node cluster that uses cluster interconnect switches and is running ONTAP 9.3 or later, you can replace the switches with direct, back-to-back connections between the nodes.


What you'll need

- A healthy cluster that consists of two nodes connected by cluster switches. The nodes must be running the same ONTAP release.
- Each node with the required number of dedicated cluster ports, which provide redundant cluster interconnect connections to support your system configuration. For example, there are two redundant ports for a system with two dedicated cluster interconnect ports on each node.

Migrate the switches

About this task

The following procedure removes the cluster switches in a two-node cluster and replaces each connection to the switch with a direct connection to the partner node.



About the examples

The examples in the following procedure show nodes that are using "e0a" and "e0b" as cluster ports. Your

nodes might be using different cluster ports as they vary by system.

Step 1: Prepare for migration

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

```
set -privilege advanced
```

The advanced prompt `*>` appears.

2. ONTAP 9.3 and later supports automatic detection of switchless clusters, which is enabled by default.

You can verify that detection of switchless clusters is enabled by running the advanced privilege command:

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

Show example

The following example output shows if the option is enabled.

```
cluster::*> network options detect-switchless-cluster show
(network options detect-switchless-cluster show)
Enable Switchless Cluster Detection: true
```

If "Enable Switchless Cluster Detection" is `false`, contact NetApp support.

3. 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=<number_of_hours>h
```

where `h` is the duration of the maintenance window in hours. The message notifies technical support of this maintenance task so that they can suppress automatic case creation during the maintenance window.

In the following example, the command suppresses automatic case creation for two hours:

Show example

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


Step 2: Configure ports and cabling

1. Organize the cluster ports on each switch into groups so that the cluster ports in group1 go to cluster switch1 and the cluster ports in group2 go to cluster switch2. These groups are required later in the procedure.

2. Identify the cluster ports and verify link status and health:

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

In the following example for nodes with cluster ports "e0a" and "e0b", one group is identified as "node1:e0a" and "node2:e0a" and the other group as "node1:e0b" and "node2:e0b". Your nodes might be using different cluster ports because they vary by system.



Verify that the ports have a value of `up` for the "Link" column and a value of `healthy` for the "Health Status" column.

Show example

```
cluster::> network port show -ipspace Cluster
Node: node1

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

Node: node2

Ignore
Speed (Mbps) Health
Health
Port IPspace Broadcast Domain Link MTU Admin/Oper Status
Status
-----
-----
e0a Cluster Cluster up 9000 auto/10000 healthy
false
e0b Cluster Cluster up 9000 auto/10000 healthy
false
4 entries were displayed.
```

3. Confirm that all the cluster LIFs are on their home ports.

Verify that the “is-home” column is `true` for each of the cluster LIFs:

```
network interface show -vserver Cluster -fields is-home
```


Show example

```
cluster::*> net int show -vserver Cluster -fields is-home
(network interface show)
vserver  lif           is-home
-----  -
Cluster  node1_clus1   true
Cluster  node1_clus2   true
Cluster  node2_clus1   true
Cluster  node2_clus2   true
4 entries were displayed.
```

If there are cluster LIFs that are not on their home ports, revert those LIFs to their home ports:

```
network interface revert -vserver Cluster -lif *
```

4. Disable auto-revert for the cluster LIFs:

```
network interface modify -vserver Cluster -lif * -auto-revert false
```

5. Verify that all ports listed in the previous step are connected to a network switch:

```
network device-discovery show -port cluster_port
```

The “Discovered Device” column should be the name of the cluster switch that the port is connected to.

Show example

The following example shows that cluster ports "e0a" and "e0b" are correctly connected to cluster switches "cs1" and "cs2".

```
cluster::> network device-discovery show -port e0a|e0b
(network device-discovery show)
Node/      Local  Discovered
Protocol  Port   Device (LLDP: ChassisID)  Interface  Platform
-----  -
node1/cdp
          e0a    cs1                      0/11       BES-53248
          e0b    cs2                      0/12       BES-53248
node2/cdp
          e0a    cs1                      0/9        BES-53248
          e0b    cs2                      0/9        BES-53248
4 entries were displayed.
```

6. Verify the cluster connectivity:

```
cluster ping-cluster -node local
```

7. Verify that the cluster is healthy:

```
cluster ring show
```

All units must be either master or secondary.


8. Set up the switchless configuration for the ports in group 1.



To avoid potential networking issues, you must disconnect the ports from group1 and reconnect them back-to-back as quickly as possible, for example, **in less than 20 seconds**.


a. Disconnect all the cables from the ports in group1 at the same time.

In the following example, the cables are disconnected from port "e0a" on each node, and cluster traffic continues through the switch and port "e0b" on each node:



b. Cable the ports in group1 back-to-back.

In the following example, "e0a" on node1 is connected to "e0a" on node2:



9. The switchless cluster network option transitions from `false` to `true`. This might take up to 45 seconds. Confirm that the switchless option is set to `true`:

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

The following example shows that the switchless cluster is enabled:

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

10. Verify that the cluster network is not disrupted:

```
cluster ping-cluster -node local
```



Before proceeding to the next step, you must wait at least two minutes to confirm a working back-to-back connection on group 1.


11. Set up the switchless configuration for the ports in group 2.



To avoid potential networking issues, you must disconnect the ports from group2 and reconnect them back-to-back as quickly as possible, for example, **in less than 20 seconds**.

- a. Disconnect all the cables from the ports in group2 at the same time.

In the following example, the cables are disconnected from port "e0b" on each node, and cluster traffic continues through the direct connection between the "e0a" ports:



b. Cable the ports in group2 back-to-back.

In the following example, "e0a" on node1 is connected to "e0a" on node2 and "e0b" on node1 is connected to "e0b" on node2:



Step 3: Verify the configuration

1. Verify that the ports on both nodes are correctly connected:

```
network device-discovery show -port cluster_port
```

Show example

The following example shows that cluster ports "e0a" and "e0b" are correctly connected to the corresponding port on the cluster partner:

```
cluster::> net device-discovery show -port e0a|e0b
(network device-discovery show)
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface  Platform
-----
node1/cdp
          e0a    node2                      e0a        AFF-A300
          e0b    node2                      e0b        AFF-A300
node1/lldp
          e0a    node2 (00:a0:98:da:16:44) e0a        -
          e0b    node2 (00:a0:98:da:16:44) e0b        -
node2/cdp
          e0a    node1                      e0a        AFF-A300
          e0b    node1                      e0b        AFF-A300
node2/lldp
          e0a    node1 (00:a0:98:da:87:49) e0a        -
          e0b    node1 (00:a0:98:da:87:49) e0b        -
8 entries were displayed.
```

2. Re-enable auto-revert for the cluster LIFs:

```
network interface modify -vserver Cluster -lif * -auto-revert true
```

3. Verify that all LIFs are home. This might take a few seconds.

```
network interface show -vserver Cluster -lif lif_name
```

Show example

The LIFs have been reverted if the “Is Home” column is `true`, as shown for `node1_clus2` and `node2_clus2` in the following example:

```
cluster::> network interface show -vserver Cluster -fields curr-  
port,is-home  
vserver  lif                curr-port is-home  
-----  
Cluster  node1_clus1         e0a      true  
Cluster  node1_clus2         e0b      true  
Cluster  node2_clus1         e0a      true  
Cluster  node2_clus2         e0b      true  
4 entries were displayed.
```

If any cluster LIFS have not returned to their home ports, revert them manually from the local node:

```
network interface revert -vserver Cluster -lif lif_name
```

4. Check the cluster status of the nodes from the system console of either node:

```
cluster show
```

Show example

The following example shows `epsilon` on both nodes to be `false`:

```
Node  Health  Eligibility Epsilon  
-----  
node1 true    true       false  
node2 true    true       false  
2 entries were displayed.
```

5. Confirm connectivity between the cluster ports:

```
cluster ping-cluster local
```

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

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

For more information, see [NetApp KB Article 1010449: How to suppress automatic case creation during scheduled maintenance windows](#).

7. Change the privilege level back to admin:

```
set -privilege admin
```

Copyright information

Copyright © 2024 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.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, 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) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.