



Broadcom®-supported Switches

Migrating to a Two-node Switched Cluster with Broadcom-supported BES-53248 Cluster Switches

For Broadcom-supported BES-53248 switches

July 2019 | 215-13971_B0
doccomments@netapp.com

Contents

Migrating to a two-node switched cluster with Broadcom-supported BES-53248 cluster switches	4
Migrating to a switched NetApp cluster environment using Broadcom-supported BES-53248 cluster 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 Broadcom-supported BES-53248 cluster switches

If you have a two-node switchless cluster, you can migrate, non-disruptively, to a two-node switched cluster that includes Broadcom-supported BES-53248 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.

About this task

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

Ensure that the BES-53248 cluster switch is set up as described in the *Switch Setup and Configuration Guide for Broadcom-supported BES-53248 switches* before starting this migration process.

Setting up and configuring

Note: After your migration completes, you might need to install the required configuration file to support the Cluster Switch Health Monitor (CSHM) for BES-53248 cluster switches.

See *Installing the Cluster Switch Health Monitor (CSHM) configuration file* in the *Setting up and configuring* guide.

Migrating to a switched NetApp cluster environment using Broadcom-supported BES-53248 cluster switches

If you have an existing two-node switchless cluster environment, you can migrate to a two-node switched cluster environment using Broadcom-supported BES-53248 cluster 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.

Broadcom-supported BES-53248 cluster switch configuration:

- The BES-53248 cluster switch must be fully functional on both switches.
- Both switches must have management network connectivity.
- There must be console access to the cluster switches.
- BES-53248 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 0/55 and 0/56 on both BES-53248 switches.
- Initial customization of both the BES-53248 switches must be completed. So that the:
 - BES-53248 switches are running the latest version of software
 - BES-53248 switches have optional licenses installed, if purchased
 - 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 BES-53248 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*.

Note: You must not disable the ISL ports.

Example

The following example shows that node-facing ports 1 through 16 are disabled on switch *cs1*:

```
(cs1)# configure
(cs1)(Config)# interface 0/1-0/16
(cs1)(Interface 0/1-0/16)# shutdown
(cs1)(Interface 0/1-0/16)# exit
(cs1)(Config)# exit
```

- Verify that the ISL and the physical ports on the ISL between the two BES-53248 switches cs1 and cs2 are up:

```
show port-channel
```

Example

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

```
(cs1)# show port-channel 1/1
Local Interface..... 1/1
Channel Name..... Cluster-ISL
Link State..... Up
Admin Mode..... Enabled
Type..... Dynamic
Port channel Min-links..... 1
Load Balance Option..... 7
(Enhanced hashing mode)

Mbr   Device/   Port   Port
Ports Timeout   Speed  Active
-----
0/55  actor/long 100G Full  True
      partner/long
0/56  actor/long 100G Full  True
      partner/long
(cs1) #
```

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

```
(cs2)# show port-channel 1/1
Local Interface..... 1/1
Channel Name..... Cluster-ISL
Link State..... Up
Admin Mode..... Enabled
Type..... Dynamic
Port channel Min-links..... 1
Load Balance Option..... 7
(Enhanced hashing mode)

Mbr   Device/   Port   Port
Ports Timeout   Speed  Active
-----
0/55  actor/long 100G Full  True
      partner/long
0/56  actor/long 100G Full  True
      partner/long
(cs2) #
```

- Display the list of neighboring devices:

```
show isdp 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 isdp neighbors

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge,
                  S - Switch, H - Host, I - IGMP, r - Repeater
Device ID        Intf    Holdtime Capability Platform  Port ID
-----
cs2              0/55   176      R          BES-53248 0/55
cs2              0/56   176      R          BES-53248 0/56
```

The following example lists the neighboring devices on switch cs2:

```
(cs2)# show isdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge,
                  S - Switch, H - Host, I - IGMP, r - Repeater
Device ID      Intf      Holdtime  Capability  Platform  Port ID
-----
cs2            0/55     176      R           BES-53248 0/55
cs2            0/56     176      R           BES-53248 0/56
```

- 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.
```

- 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.
```

- 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
```

8 | Migrating to a two-node switched cluster with Broadcom-supported BES-53248 cluster switches

```
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 BES-53248 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 BES-53248 switches.
11. Enable all node-facing ports on cluster switch cs1.

Example

The following example shows that ports 1 through 16 are enabled on switch cs1:

```
(cs1)# configure
(cs1)(Config)# interface 0/1-0/16
(cs1)(Interface 0/1-0/16)# no shutdown
(cs1)(Interface 0/1-0/16)# exit
(cs1)(Config)# exit
```

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
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is 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 BES-53248 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 BES-53248 switches.
16. Enable all node-facing ports on cluster switch cs2.

Example

The following example shows that ports 1 through 16 are enabled on switch cs2:

```
(cs2)# configure
(cs2)(Config)# interface 0/1-0/16
(cs2)(Interface 0/1-0/16)# no shutdown
(cs2)(Interface 0/1-0/16)# exit
(cs2)(Config)# exit
```

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
-----
Admin/Oper Status Status
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
-----
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.
```

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 isdp neighbors
```

Example

The following example shows the appropriate results for both switches:

```
(cs1)# show isdp neighbors

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge,
                  S - Switch, H - Host, I - IGMP, r - Repeater
Device ID        Intf          Holdtime  Capability  Platform  Port ID
-----
node1            0/1           175       H            FAS2750   e0a
node2            0/2           157       H            FAS2750   e0a
cs2              0/55          178       R            BES-53248 0/55
cs2              0/56          178       R            BES-53248 0/56

(cs2)# show isdp neighbors

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge,
                  S - Switch, H - Host, I - IGMP, r - Repeater
Device ID        Intf          Holdtime  Capability  Platform  Port ID
-----
node1            0/1           137       H            FAS2750   e0b
node2            0/2           179       H            FAS2750   e0b
cs1              0/55          175       R            BES-53248 0/55
cs1              0/56          175       R            BES-53248 0/56
```

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        BES-53248
          e0b   cs2                       0/2        BES-53248
node1     /cdp
          e0a   cs1                       0/1        BES-53248
          e0b   cs2                       0/1        BES-53248

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 using the command:

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

Example

```
cluster1::*> cluster ping-cluster -node local

Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 192.168.168.26 node1 e0a
Cluster node1_clus2 192.168.168.27 node1 e0b
Cluster node2_clus1 192.168.168.28 node2 e0a
Cluster node2_clus2 192.168.168.29 node2 e0b
Local = 192.168.168.28 192.168.168.29
Remote = 192.168.168.26 192.168.168.27
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 1500 byte MTU on 4 path(s):
  Local 192.168.168.28 to Remote 192.168.168.26
  Local 192.168.168.28 to Remote 192.168.168.27
  Local 192.168.168.29 to Remote 192.168.168.26
  Local 192.168.168.29 to Remote 192.168.168.27
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
```

Related information

[Hardware Universe](#)

[Setting up and configuring](#)

[NetApp KB Article 1000275: How to suppress automatic case creation during scheduled maintenance windows](#)

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-transferrable, 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*](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.

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

B

- BES-53248 switches
 - migrating from a two-node switchless cluster environment to a two-node switched cluster environment [4](#)
 - migrating two-node switched clusters [4](#)

C

- cluster environments
 - migrating from a switchless two-node cluster to a two-node switched BES-53248 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 BES-53248 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 BES-53248 switched cluster environment [4](#)
- switchless cluster environments
 - migrating to a two-node switched BES-53248 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 BES-53248 cluster environment [4](#)
- two-node switchless clusters
 - migrating to a two-node switched BES-53248 cluster environment [4](#)