



Configuring hardware for sharing a Brocade 6510 FC fabric during transition

ONTAP MetroCluster

NetApp
June 16, 2021

Table of Contents

- Configuring hardware for sharing a Brocade 6510 FC fabric during transition 1
 - Reviewing Brocade license requirements 1
 - Racking the hardware components 2
 - Cabling the new MetroCluster controllers to the existing FC fabrics 3
 - Configuring switch fabrics sharing between the 7-Mode and clustered MetroCluster configuration 4

Configuring hardware for sharing a Brocade 6510 FC fabric during transition

If your 7-Mode fabric MetroCluster configuration uses Brocade 6510 switches, you can share the existing switch fabrics with the new clustered MetroCluster configuration. Shared switch fabrics means the new MetroCluster configuration does not require a new, separate switch fabric. This temporary configuration is only supported with the Brocade 6510 switch for transition purposes.

- The 7-Mode fabric MetroCluster must be using Brocade 6510 switches.

If the MetroCluster configuration is currently not using Brocade 6510 switches, the switches must be upgraded to Brocade 6510 prior to using this procedure.

- The 7-Mode fabric MetroCluster configuration must be using SAS storage shelves only.

If the existing configuration includes FC storage shelves (such as the DS14mk4 FC), FC switch fabric sharing is not supported.

- The SFPs on the switch ports used by the new, clustered MetroCluster configuration must support 16-Gbps rates.

The existing 7-Mode fabric MetroCluster can remain connected to ports using 8-Gbps or 16-Gbps SFPs.

- On each of the four Brocade 6510 switches, ports 24 through 45 must be available to connect the ports of the new MetroCluster components.
- You should verify that the existing Inter-Switch Links (ISLs) are on ports 46 and 47.
- The Brocade 6510 switches must be running a FOS firmware version that is supported on both the 7-Mode fabric MetroCluster and clustered ONTAP MetroCluster configuration.

After sharing the fabric and completing the MetroCluster configuration, you can transition data from the 7-Mode fabric MetroCluster configuration.

After transitioning the data, you can remove the 7-Mode fabric MetroCluster cabling and, if desired, move the clustered ONTAP MetroCluster cabling to the lower-numbered ports previously used for the 7-Mode MetroCluster cabling. The ports are shown in the section "Reviewing FC switch port assignments for a four node MetroCluster." You must adjust the zoning for the rearranged ports.

[Port assignments for FC switches when using ONTAP 9.1 and later](#)

Related information

[Copy-based transition](#)

Reviewing Brocade license requirements

You need certain licenses for the switches in a MetroCluster configuration. You must install these licenses on all four switches.

The MetroCluster configuration has the following Brocade license requirements:

- Trunking license for systems using more than one ISL, as recommended.
- Extended Fabric license (for ISL distances over 6 km)
- Enterprise license for sites with more than one ISL and an ISL distance greater than 6 km

The Enterprise license includes Brocade Network Advisor and all licenses except for additional port licenses.

You can verify that the licenses are installed by using the `licenseshow` command. If you do not have these licenses, you should contact your sales representative before proceeding.

Racking the hardware components

If you have not received the equipment already installed in cabinets, you must rack the components.

This task must be performed on both MetroCluster sites.

Steps

1. Plan out the positioning of the MetroCluster components.

The rack space depends on the platform model of the controller modules, the switch types, and the number of disk shelf stacks in your configuration.

2. Properly ground yourself.
3. Install the controller modules in the rack or cabinet.

[AFF and FAS Documentation Center](#)

4. Install the FC switches in the rack or cabinet.
5. Install the disk shelves, power them on, and then set the shelf IDs.
 - You must power-cycle each disk shelf.
 - Shelf IDs must be unique for each SAS disk shelf within each MetroCluster DR group (including both sites).
6. Install each FC-to-SAS bridge:

- a. Secure the “L” brackets on the front of the bridge to the front of the rack (flush-mount) with the four screws.

The openings in the bridge “L” brackets are compliant with rack standard ETA-310-X for 19-inch (482.6 mm) racks.

The *ATTO FibreBridge Installation and Operation Manual* for your bridge model contains more information and an illustration of the installation.



For adequate port space access and FRU serviceability, you must leave 1U space below the bridge pair and cover this space with a tool-less blanking panel.

- b. Connect each bridge to a power source that provides a proper ground.
- c. Power on each bridge.



For maximum resiliency, bridges that are attached to the same stack of disk shelves must be connected to different power sources.

The bridge Ready LED might take up to 30 seconds to illuminate, indicating that the bridge has completed its power-on self test sequence.

Cabling the new MetroCluster controllers to the existing FC fabrics

On each controller in the clustered ONTAP MetroCluster configuration, the FC-VI adapter and HBAs must be cabled to specific ports on the existing FC switches.

Steps

1. Cable the FC-VI and HBA ports according to the following table:

Site A		Site B	
Connect this Site A component and port...	FC_switch_A_1 port...	Connect this Site B component and port...	FC_switch_B_1 port...
controller_A_1 FC-VI port 1	32	controller_B_1 FC-VI port 1	32
controller_A_1 HBA port 1	33	controller_B_1 HBA port 1	33
controller_A_1 HBA port 2	34	controller_B_1 HBA port 2	34
controller_A_2 FC-VI port 1	35	controller_B_2 FC-VI port 1	35
controller_A_2 HBA 1	36	controller_B_2 HBA 1	36
controller_A_2 HBA 2	37	controller_B_2 HBA 2	37

2. Cable each FC-SAS bridge in the first switch fabric to the FC switches.

The number of bridges varies depending on the number of SAS storage stacks.

Site A		Site B	
Cable this site A bridge...	FC_switch_A_1 port...	Cable this Site B bridge...	FC_switch_B_1 port...
bridge_A_1_38	38	bridge_B_1_38	38
bridge_A_1_39	39	bridge_B_1_39	39

3. Cable each bridge in the second switch fabric to the FC switches.

The number of bridges varies depending on the number of SAS storage stacks.

Site A		Site B	
Cable this site A bridge...	FC_switch_A_2 port...	Cable this Site B bridge...	FC_switch_B_2 port...
bridge_A_2_38	38	bridge_B_2_38	38
bridge_A_2_39	39	bridge_B_2_39	39

Configuring switch fabrics sharing between the 7-Mode and clustered MetroCluster configuration

To share switch fabrics between the existing 7-Mode fabric MetroCluster and the new MetroCluster configuration, you must set up specific zoning and other settings that are different than an unshared configuration.

This task must be performed on both switch fabrics, one at a time.

Disabling one of the switch fabrics

You must disable one of the switch fabrics so you can modify its configuration. After you complete the configuration and reenble the switch fabric, you will repeat the process on the other fabric.

You must have run the `fmc_dc` utility on the existing 7-Mode fabric MetroCluster configuration and resolved any issues prior to beginning the configuration process.

To ensure continued operation of the MetroCluster configuration, you must not disable the second fabric while the first fabric is disabled.

Steps

1. Disable each of the switches in the fabric:

```
switchCfgPersistentDisable
```

If this command is not available, use the `switchDisable` command.

The following example shows the command issued on `FC_switch_A_1`:

```
FC_switch_A_1:admin> switchCfgPersistentDisable
```

The following example shows the command issued on `FC_switch_B_1`:

```
FC_switch_B_1:admin> switchCfgPersistentDisable
```

2. Ensure that the 7-Mode MetroCluster configuration is functioning correctly using the redundant fabric:

a. Confirm that controller failover is healthy:

cf status

```
node_A> cf status
Controller Failover enabled, node_A is up.
VIA Interconnect is up (link 0 down, link 1 up).
```

b. Confirm that disks are visible:

storage show disk -p

```
node_A> storage show disk -p
```

PRIMARY	PORT	SECONDARY	PORT	SHELF	BAY
-----	----	-----	----	-----	-----
Brocade-6510-2K0GG:5.126L27	B			1	0
Brocade-6510-2K0GG:5.126L28	B			1	1
Brocade-6510-2K0GG:5.126L29	B			1	2
Brocade-6510-2K0GG:5.126L30	B			1	3
Brocade-6510-2K0GG:5.126L31	B			1	4
.					
.					
.					

c. Confirm that the aggregates are healthy:

aggr status

```
node_A> aggr status
```

Aggr State	Status	Options
aggr0 online	raid_dp, aggr mirrored 64-bit	root, nosnap=on

Deleting TI zoning and configuring IOD settings

You must delete the existing TI zoning and reconfigure in-order-delivery (IOD) settings on the switch fabric.

1. Identify the TI zones that are configured on the fabric:

zone --show

The following example shows the zone FCVI_TI_FAB_2.

```
Brocade-6510:admin> zone --show
  Defined TI zone configuration:
  TI Zone Name:    FCVI_TI_FAB_2
  Port List:      1,0; 1,3; 2,0; 2,3
  configured Status: Activated / Failover-Disabled
  Enabled Status: Activated / Failover-Disabled
```

2. Delete the TI zones:

zone --delete zone-name

The following example shows the deletion of zone FCVI_TI_FAB_2.

```
Brocade-6510:admin> zone --delete FCVI_TI_FAB_2
```

3. Confirm that the zones have been deleted:

zone --show

The output should be similar to the following:

```
Brocade-6510:admin> zone --show

  Defined TI zone configuration:
  no TI zone configuration defined
```

4. Save the configuration:

cfgsave

5. Enable in-order-delivery:

iodset

6. Select Advanced Performance Tuning (APT) policy 1, the Port Based Routing Policy:

aptpolicy 1

7. Disable Dynamic Load Sharing (DLS):

dlsreset

8. Verify the IOD settings using the following commands:

iodshow

aptpolicy

dlsshow

The output should be similar to the following:

```
Brocade-6510:admin> iodshow

IOD is set

Brocade-6510:admin> aptpolicy
Current Policy: 1

3 : Default Policy
1: Port Based Routing Policy
2: Device Based Routing Policy (FICON support only)
3: Exchange Based Routing Policy
Brocade-6510:admin> dlsshow

DLS is not set
```

Ensuring ISLs are in the same port group and configuring zoning

You must make sure that the Inter-Switch Links (ISLs) are in the same port group and configure zoning for the MetroCluster configurations to successfully share the switch fabrics.

Steps

1. If the ISLs are not in the same port group, move one of the ISL ports to the same port group as the other one.

You can use any available port except 32 through 45, which are used by the new MetroCluster configuration. The recommended ISL ports are 46 and 47.

2. Follow the steps in the Configuring zoning on a Brocade FC switch section to enable trunking and the QoS zone.

[Configuring zoning on Brocade FC switches](#)

The port numbers when sharing fabrics are different than those shown in the section. When sharing, use ports 46 and 47 for the ISL ports. If you moved your ISL ports, you need to use the procedure in the Configuring the E-ports (ISL ports) on a Brocade FC switch section to configure the ports.

[Configuring the E-ports \(ISL ports\) on a Brocade FC switch](#)

3. Follow the steps in the Configuring the non-E ports on the Brocade switch section to configure the non-E ports.

Configuring the non-E ports on the Brocade switch

4. Do not delete the zones or zone sets that already exist in the backend switches (for the 7-Mode fabric MetroCluster) except the Traffic Isolation (TI) zones in Step 3.
5. Follow the steps in the Configuring the E-ports (ISL ports) on a Brocade FC switch section to add the zones required by the new MetroCluster to the existing zone sets.

Configuring the E-ports (ISL ports) on a Brocade FC switch

The following example shows the commands and system output for creating the zones:

```
Brocade-6510-2K0GG:admin> zonecreate "QOSH2_FCVI_1", "2,32; 2,35; 1,32;
1,35"

Brocade-6510-2K0GG:admin> zonecreate "STOR_A_2_47", "2,33; 2,34; 2,36;
2,37; 1,33; 1,34; 1,36; 1,37; 1,47"

Brocade-6510-2K0GG:admin> zonecreate "STOR_B_2_47", "2,33; 2,34; 2,36;
2,37; 1,33; 1,34; 1,36; 1,37; 2,47"

Brocade-6510-2K0GG:admin> cfgadd config_1_FAB2, "QOSH2_FCVI_1;
STOR_A_2_47; STOR_B_2_47"

Brocade-6510-2K0GG:admin> cfgenable "config_1_FAB2"
You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the
current configuration selected. If the update includes changes
to one or more traffic isolation zones, the update may result in
localized disruption to traffic on ports associated with
the traffic isolation zone changes
Do you want to enable 'config_1_FAB2' configuration (yes, y, no, n):
[no] yes

Brocade-6510-2K0GG:admin> cfsave
You are about to save the Defined zoning configuration. This
action will only save the changes on Defined configuration.
Do you want to save the Defined zoning configuration only? (yes, y, no,
n): [no] yes
Nothing changed: nothing to save, returning ...
Brocade-6510-2K0GG:admin>
```

Reenabling the switch fabric and verify operation

You must enable the FC switch fabric and ensure that the switches and devices are operating correctly.

Steps

1. Enable the switches:

switchCfgPersistentEnable

If this command is not available, the switch should be in the enabled state after the `fastBoot` command is issued.

The following example shows the command issued on `FC_switch_A_1`:

```
FC_switch_A_1:admin> switchCfgPersistentEnable
```

The following example shows the command issued on `FC_switch_B_1`:

```
FC_switch_B_1:admin> switchCfgPersistentEnable
```

2. Verify that the switches are online and all devices are properly logged in:

switchShow

The following example shows the command issued on `FC_switch_A_1`:

```
FC_switch_A_1:admin> switchShow
```

The following example shows the command issued on `FC_switch_B_1`:

```
FC_switch_B_1:admin> switchShow
```

3. Run the `fmc_dc utility` to ensure that the 7-Mode fabric MetroCluster is functioning correctly.

You can ignore errors related to Traffic Isolation (TI) zoning and trunking.

4. Repeat the tasks for the second switch fabric.

Copyright Information

Copyright © 2021 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

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.