



Setting up ONTAP in a MetroCluster configuration with array LUNs

ONTAP MetroCluster

NetApp
June 14, 2021

Table of Contents

- Setting up ONTAP in a MetroCluster configuration with array LUNs 1
 - Verifying and configuring the HA state of components in Maintenance mode 1
 - Configuring ONTAP on a system that uses only array LUNs 2
 - Setting up the cluster 8
 - Installing the license for using array LUNs in a MetroCluster configuration 8
 - Configuring FC-VI ports on a X1132A-R6 quad-port card on FAS8020 systems 9
 - Assigning ownership of array LUNs 12
 - Peering the clusters 12
 - Mirroring the root aggregates 13
 - Creating data aggregates on, implementing, and verifying the MetroCluster configuration 13

Setting up ONTAP in a MetroCluster configuration with array LUNs

After connecting the devices in the MetroCluster configuration, you must set up the ONTAP systems to use the storage on the storage array. You must also configure any required ONTAP feature.

Verifying and configuring the HA state of components in Maintenance mode

When configuring a storage system in a MetroCluster configuration, you must make sure that the high-availability (HA) state of the controller module and chassis components is `mcc` or `mcc-2n` so that these components boot properly.

The system must be in Maintenance mode.

This task is not required on systems that are received from the factory.

1. In Maintenance mode, display the HA state of the controller module and chassis:

```
ha-config show
```

The correct HA state depends on your MetroCluster configuration.

Number of controllers in the MetroCluster configuration	HA state for all components should be...
Eight- or four-node MetroCluster FC configuration	<code>mcc</code>
Two-node MetroCluster FC configuration	<code>mcc-2n</code>
MetroCluster IP configuration	<code>mccip</code>

2. If the displayed system state of the controller is not correct, set the HA state for the controller module:

Number of controllers in the MetroCluster configuration	Command
Eight- or four-node MetroCluster FC configuration	<code>ha-config modify controller mcc</code>
Two-node MetroCluster FC configuration	<code>ha-config modify controller mcc-2n</code>
MetroCluster IP configuration	<code>ha-config modify controller mccip</code>

3. If the displayed system state of the chassis is not correct, set the HA state for the chassis:

Number of controllers in the MetroCluster configuration	Command
Eight- or four-node MetroCluster FC configuration	<code>ha-config modify chassis mcc</code>
Two-node MetroCluster FC configuration	<code>ha-config modify chassis mcc-2n</code>
MetroCluster IP configuration	<code>ha-config modify chassis mccip</code>

4. Boot the node to ONTAP:

```
boot_ontap
```

5. Repeat these steps on each node in the MetroCluster configuration.

Configuring ONTAP on a system that uses only array LUNs

If you want to configure ONTAP for use with array LUNs, you must configure the root aggregate and root volume, reserve space for diagnostics and recovery operations, and set up the cluster.

- The ONTAP system must be connected to the storage array.
- The storage array administrator must have created LUNs and presented them to ONTAP.
- The storage array administrator must have configured the LUN security.

You must configure each node that you want to use with array LUNs. If the node is in an HA pair, then you must complete the configuration process on one node before proceeding with the configuration on the partner node.

Steps

1. Power on the primary node and interrupt the boot process by pressing Ctrl-C when you see the following message on the console: Press CTRL-C for special boot menu.
2. Select option 4 (Clean configuration and initialize all disks) on the boot menu.

The list of array LUNs made available to ONTAP is displayed. In addition, the array LUN size required for root volume creation is also specified. The size required for root volume creation differs from one ONTAP system to another.

- If no array LUNs were previously assigned, ONTAP detects and displays the available array LUNs, as shown in the following example:

```

mcc8040-ams1::> disk show NET-1.6 -instance
          Disk: NET-1.6
    Container Type: aggregate
      Owner/Home: mcc8040-ams1-01 / mcc8040-ams1-01
        DR Home: -
Stack ID/Shelf/Bay: - / - / -
          LUN: 0
        Array: NETAPP_INF_1
      Vendor: NETAPP
        Model: INF-01-00
    Serial Number: 60080E50004317B4000003B158E35974
          UID:
60080E50:004317B4:000003B1:58E35974:00000000:00000000:00000000:000000
00:00000000:00000000
          BPS: 512
    Physical Size: 87.50GB
      Position: data
Checksum Compatibility: block
      Aggregate: eseries
        Plex: plex0

Paths:

          LUN  Initiator Side      Target
Side                               Link
Controller      Initiator  ID  Switch Port      Switch
Port            Acc Use  Target Port      TPGN      Speed
I/O KB/s            IOPS
-----
-----
-----
mcc8040-ams1-01    2c                0  mccb6505-ams1:16  mccb6505-
ams1:18          AO  INU  20330080e54317b4  1  4 Gb/S
0                0
mcc8040-ams1-01    2a                0  mccb6505-ams1:17  mccb6505-
ams1:19          ANO RDY  20320080e54317b4  0  4 Gb/S
0                0

Errors:
-
```

- If array LUNs were previously assigned, for example, through the maintenance mode, they are either marked local or partner in the list of the available array LUNs, depending on whether the array LUNs were selected from the node on which you are installing ONTAP or its HA partner:

In this example, array LUNs with index numbers 3 and 6 are marked local because they had been previously assigned from this particular node:

*
* No disks are owned by this node, but array LUNs are assigned.
*
* You can use the following information to verify connectivity from
*
* HBAs to switch ports. If the connectivity of HBAs to switch ports
*
* does not match your expectations, configure your SAN and rescan.
*
* You can rescan by entering 'r' at the prompt for selecting
*
* array LUNs below.

WPN	HBA	HBA WWPN	Switch port	Switch port
-----	---	-----	-----	-----
20190027f88948dd	0e	500a098001baf8e0	vgbr6510s203:25	
2011547feead680	0f	500a098101baf8e0	vgci9710s202:1-17	
201b0027f88948dd	0g	500a098201baf8e0	vgbr6510s203:27	
2012547feead680	0h	500a098301baf8e0	vgci9710s202:1-18	

No native disks were detected, but array LUNs were detected.
You will need to select an array LUN to be used to create the root aggregate and root volume.

The array LUNs visible to the system are listed below. Select one array LUN to be used to create the root aggregate and root volume. **The root volume requires 350.0 GB of space.**

Warning: The contents of the array LUN you select will be erased by ONTAP prior to their use.

Index	Array LUN Name	Model	Vendor	Size	Owner
Checksum	Serial Number				
-----	-----	-----	-----	-----	-----
-----	-----	-----	-----	-----	-----

0	vgci9710s202:2-24.0L19	RAID5	DGC	217.3	GB	
Block	6006016083402B0048E576D7					
1	vgbr6510s203:30.126L20	RAID5	DGC	217.3	GB	
Block	6006016083402B0049E576D7					
2	vgci9710s202:2-24.0L21	RAID5	DGC	217.3	GB	
Block	6006016083402B004AE576D7					
3	vgbr6510s203:30.126L22	RAID5	DGC	405.4	GB	local
Block	6006016083402B004BE576D7					
4	vgci9710s202:2-24.0L23	RAID5	DGC	217.3	GB	
Block	6006016083402B004CE576D7					
5	vgbr6510s203:30.126L24	RAID5	DGC	217.3	GB	
Block	6006016083402B004DE576D7					
6	vgbr6510s203:30.126L25	RAID5	DGC	423.5	GB	local
Block	6006016083402B003CF93694					
7	vgci9710s202:2-24.0L26	RAID5	DGC	423.5	GB	
Block	6006016083402B003DF93694					

3. Select the index number corresponding to the array LUN you want to assign as the root volume.

The array LUN must be of sufficient size to create the root volume.

The array LUN selected for root volume creation is marked local (root).

In the following example, the array LUN with index number 3 is marked for root volume creation:

The root volume will be created on switch 0:5.183L33.

ONTAP requires that 11.0 GB of space be reserved for use in diagnostic and recovery operations. Select one array LUN to be used as spare for diagnostic and recovery operations.

Index	Array LUN Name	Model	Vendor	Size	Owner
Checksum	Serial Number				
0	switch0:5.183L1	SYMMETRIX	EMC	266.1 GB	
Block	600604803436313734316631				
1	switch0:5.183L3	SYMMETRIX	EMC	266.1 GB	
Block	600604803436316333353837				
2	switch0:5.183L31	SYMMETRIX	EMC	266.1 GB	
Block	600604803436313237643666				
3	switch0:5.183L33	SYMMETRIX	EMC	658.3 GB	local (root)
Block	600604803436316263613066				
4	switch0:7.183L0	SYMMETRIX	EMC	173.6 GB	
Block	600604803436313261356235				
5	switch0:7.183L2	SYMMETRIX	EMC	173.6 GB	
Block	600604803436313438396431				
6	switch0:7.183L4	SYMMETRIX	EMC	658.3 GB	
Block	600604803436313161663031				
7	switch0:7.183L30	SYMMETRIX	EMC	173.6 GB	
Block	600604803436316538353834				
8	switch0:7.183L32	SYMMETRIX	EMC	266.1 GB	
Block	600604803436313237353738				
9	switch0:7.183L34	SYMMETRIX	EMC	658.3 GB	
Block	600604803436313737333662				

4. Select the index number corresponding to the array LUN you want to assign for use in diagnostic and recovery options.

The array LUN must be of sufficient size for use in diagnostic and recovery options. If required, you can also select multiple array LUNs with a combined size greater than or equal to the specified size. To select multiple entries, you must enter the comma-separated values of all of the index numbers corresponding to the array LUNs you want to select for diagnostic and recovery options.

The following example shows a list of array LUNs selected for root volume creation and for diagnostic and recovery options:


```

Here is a list of the selected array LUNs
Index Array LUN Name      Model      Vendor      Size      Owner
Checksum Serial Number
-----
2  switch0:5.183L31      SYMMETRIX  EMC        266.1 GB  local
Block      600604803436313237643666
3  switch0:5.183L33      SYMMETRIX  EMC        658.3 GB  local    (root)
Block      600604803436316263613066
4  switch0:7.183L0       SYMMETRIX  EMC        173.6 GB  local
Block      600604803436313261356235
5  switch0:7.183L2       SYMMETRIX  EMC        173.6 GB  local
Block      600604803436313438396431
Do you want to continue (yes|no)?

```



Selecting “no” clears the LUN selection.

- Enter **y** when prompted by the system to continue with the installation process.

The root aggregate and the root volume are created and the rest of the installation process continues.

- Enter the required details to create the node management interface.

The following example shows the node management interface screen with a message confirming the creation of the node management interface:

```
Welcome to node setup.
```

```
You can enter the following commands at any time:
```

```
"help" or "?" - if you want to have a question clarified,  
"back" - if you want to change previously answered questions, and  
"exit" or "quit" - if you want to quit the setup wizard.  
Any changes you made before quitting will be saved.
```

```
To accept a default or omit a question, do not enter a value.
```

```
Enter the node management interface port [e0M]:
```

```
Enter the node management interface IP address: 192.0.2.66
```

```
Enter the node management interface netmask: 255.255.255.192
```

```
Enter the node management interface default gateway: 192.0.2.7
```

```
A node management interface on port e0M with IP address 192.0.2.66 has  
been created.
```

```
This node has its management address assigned and is ready for cluster  
setup.
```

After configuring ONTAP on all of the nodes that you want to use with array LUNs, you should complete the cluster setup process.

[Software setup](#)

Related information

[FlexArray virtualization installation requirements and reference](#)

Setting up the cluster

Setting up the cluster involves setting up each node, creating the cluster on the first node, and joining any remaining nodes to the cluster.

Related information

[Software setup](#)

Installing the license for using array LUNs in a MetroCluster configuration

You must install the V_StorageAttach license on each MetroCluster node that you want to use with array LUNs. You cannot use array LUNs in an aggregate until the license is installed.

- The cluster must be installed.
- You must have the license key for the V_StorageAttach license.

You must use a separate license key for each node on which you want to install the V_StorageAttach license.

Steps

1. Use the `system license add` command to install the V_StorageAttach license.

Repeat this step for each cluster node on which you want to install the license.

2. Use the `system license show` command to verify that the V_StorageAttach license is installed on all required nodes in a cluster.

The following sample output shows that the V_StorageAttach license is installed on the nodes of cluster_A:

```
cluster_A::> system license show
Serial Number: nnnnnnnn
Owner: controller_A_1
Package          Type      Description          Expiration
-----
V_StorageAttach  license  Virtual Attached Storage

Serial Number: 11111111
Owner: controller_A_2
Package          Type      Description          Expiration
-----
V_StorageAttach  license  Virtual Attached Storage
```

Configuring FC-VI ports on a X1132A-R6 quad-port card on FAS8020 systems

If you are using the X1132A-R6 quad-port card on a FAS8020 system, you can enter Maintenance mode to configure the 1a and 1b ports for FC-VI and initiator usage. This is not required on MetroCluster systems received from the factory, in which the ports are set appropriately for your configuration.

This task must be performed in Maintenance mode.



Converting an FC port to an FC-VI port with the `ucadmin` command is only supported on the FAS8020 and AFF 8020 systems. Converting FC ports to FCVI ports is not supported on any other platform.

Steps

1. Disable the ports:

storage disable adapter 1a

storage disable adapter 1b

```
*> storage disable adapter 1a
Jun 03 02:17:57 [controller_B_1:fc.adapter.offlining:info]: Offlining
Fibre Channel adapter 1a.
Host adapter 1a disable succeeded
Jun 03 02:17:57 [controller_B_1:fc.adapter.offline:info]: Fibre Channel
adapter 1a is now offline.
*> storage disable adapter 1b
Jun 03 02:18:43 [controller_B_1:fc.adapter.offlining:info]: Offlining
Fibre Channel adapter 1b.
Host adapter 1b disable succeeded
Jun 03 02:18:43 [controller_B_1:fc.adapter.offline:info]: Fibre Channel
adapter 1b is now offline.
*>
```

2. Verify that the ports are disabled:

ucadmin show

```
*> ucadmin show
      Current  Current  Pending  Pending  Admin
Adapter Mode    Type     Mode     Type     Status
-----
...
1a    fc      initiator -        -        offline
1b    fc      initiator -        -        offline
1c    fc      initiator -        -        online
1d    fc      initiator -        -        online
```

3. Set the a and b ports to FC-VI mode:

ucadmin modify -adapter 1a -type fcvi

The command sets the mode on both ports in the port pair, 1a and 1b (even though only 1a is specified in the command).

```
*> ucaadmin modify -t fcvi 1a
Jun 03 02:19:13 [controller_B_1:ucm.type.changed:info]: FC-4 type has
changed to fcvi on adapter 1a. Reboot the controller for the changes to
take effect.
Jun 03 02:19:13 [controller_B_1:ucm.type.changed:info]: FC-4 type has
changed to fcvi on adapter 1b. Reboot the controller for the changes to
take effect.
```

4. Confirm that the change is pending:

ucaadmin show

```
*> ucaadmin show
      Current  Current  Pending  Pending  Admin
Adapter Mode    Type     Mode     Type     Status
-----
...
1a    fc      initiator -        fcvi     offline
1b    fc      initiator -        fcvi     offline
1c    fc      initiator -        -        online
1d    fc      initiator -        -        online
```

5. Shut down the controller, and then reboot into Maintenance mode.

6. Confirm the configuration change:

ucaadmin show local

```
Node          Adapter  Mode    Type     Mode     Type     Status
-----
...
controller_B_1
      1a      fc      fcvi     -        -        online
controller_B_1
      1b      fc      fcvi     -        -        online
controller_B_1
      1c      fc      initiator -        -        online
controller_B_1
      1d      fc      initiator -        -        online
6 entries were displayed.
```

Assigning ownership of array LUNs

Array LUNs must be owned by a node before they can be added to an aggregate to be used as storage.

- Back-end configuration testing (testing of the connectivity and configuration of devices behind the ONTAP systems) must be completed.
- Array LUNs that you want to assign must be presented to the ONTAP systems.

You can assign ownership of array LUNs that have the following characteristics:

- They are unowned.
- They have no storage array configuration errors, such as the following:
 - The array LUN is smaller than or larger than the size that ONTAP supports.
 - The LDEV is mapped on only one port.
 - The LDEV has inconsistent LUN IDs assigned to it.
 - The LUN is available on only one path.

ONTAP issues an error message if you try to assign ownership of an array LUN with back-end configuration errors that would interfere with the ONTAP system and the storage array operating together. You must fix such errors before you can proceed with array LUN assignment.

ONTAP alerts you if you try to assign an array LUN with a redundancy error: for example, all paths to this array LUN are connected to the same controller or only one path to the array LUN. You can fix a redundancy error before or after assigning ownership of the LUN.

Steps

1. Enter the following command to see the array LUNs that have not yet been assigned to a node:

```
storage disk show -container-type unassigned
```

2. Enter the following command to assign an array LUN to this node:

```
storage disk assign -disk array LUNname -owner nodename
```

If you want to fix a redundancy error after disk assignment instead of before, you must use the `-force` parameter with the `storage disk assign` command.

Related information

[FlexArray virtualization installation requirements and reference](#)

Peering the clusters

The clusters in the MetroCluster configuration must be in a peer relationship so that they can communicate with each other and perform the data mirroring essential to MetroCluster disaster recovery.

Steps

1. Configure intercluster LIFs using the procedure in [Configuring intercluster LIFs](#).
2. Create a cluster peer relationship using the procedure in [Creating a cluster peer relationship](#).

Mirroring the root aggregates

You must mirror the root aggregates in your MetroCluster configuration to ensure data protection.

You must have ensured that the SyncMirror requirements for the MetroCluster configuration with array LUNs are satisfied.

Requirements for a MetroCluster configuration with array LUNs

You must repeat this task for each controller in the MetroCluster configuration.

Steps

1. Use the `storage aggregate mirror` command to mirror the unmirrored root aggregate.

The following command mirrors the root aggregate for controller_A_1:

```
controller_A_1::> storage aggregate mirror aggr0_controller_A_1
```

The root aggregate is mirrored with array LUNs from pool1.

Creating data aggregates on, implementing, and verifying the MetroCluster configuration

You must create data aggregates on each node, implement, and verify the MetroCluster configuration.

1. Create data aggregates on each node:
 - a. Create a mirrored data aggregate on each node using the procedure in [Mirroring the root aggregates](#).
 - b. If desired, create unmirrored data aggregates using the procedure in [Creating a mirrored data aggregate on each node](#).
2. Implement the MetroCluster configuration using the procedure in [Implementing the MetroCluster configuration](#).
3. Configure the MetroCluster FC switches for health monitoring using the procedure in [Configuring the MetroCluster FC switches for health monitoring](#).
4. Check and verify the configuration:
 - a. Check the MetroCluster configuration using the procedure in [Checking the MetroCluster configuration](#).
 - b. Check for MetroCluster configuration errors with Config Advisor using the procedure in [Checking for MetroCluster configuration errors with Config Advisor](#).
 - c. Verify switchover, healing, and switchback using the procedure in [Verifying switchover, healing, and switchback](#).

5. Install the MetroCluster Tiebreaker software using the procedure in [Configuring the MetroCluster Tiebreaker software](#).
6. Set the destination for configuration backup files using the procedure in [Protecting configuration backup files](#).

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.