



Technical FAQ

MetroCluster FC

9.4

Version 1.0

Mike Braden, NetApp
August 2018

Abstract

MetroCluster® Fiber Channel (FC) frequently asked questions. MetroCluster FC is the backend storage fabric and includes MetroCluster stretch configurations.

TABLE OF CONTENTS

1	Overview.....	6
2	General Information	7
	How do I find out what changed between versions in ONTAP?	7
	How can I find out when a specific feature will be supported?	7
	Is there a license or fee for MetroCluster?	7
	What licenses are required for MetroCluster?	7
	Is SyncMirror supported for configurations without MetroCluster?	7
	What is the read and write behavior with plexes in SyncMirror in MetroCluster?	7
	What is the switchover granularity of MetroCluster?.....	8
2.1	Sizing a Solution.....	8
	How can I determine the capacity requirements for MetroCluster FC?	8
	Does SPM support MetroCluster FC?	8
	Does Config Advisor support MetroCluster?.....	8
2.2	Quoting MetroCluster	8
	How do I quote an 8 node MetroCluster configuration?	8
	What specific considerations would be required when quoting the FAS9000 / A700 platforms?	8
2.3	Upgrade of a MetroCluster System	8
	Where do I find information about how to upgrade ONTAP in a MetroCluster configuration?	8
	Where do I find information about how to upgrade a controller?	9
	Where do I find the information for expanding the number of nodes in a MetroCluster configuration?	9
	Where can I find the procedure for aggregate relocation?	9
3	Configuration	9
	Which FAS controller models are supported?	9
	What are the platform mixing rules for the 8-node MetroCluster configuration?.....	9
	Does the 8-node MetroCluster configuration support SAN protocols?	10
	Can I mix NAS and SAN in the 8-node MetroCluster configuration?	10
	What are the front-end network requirements for MetroCluster FC?.....	10
	What shelves are supported with MetroCluster FC?	10
	Where do I find the maximum number of disks?.....	10
	Where do I find the maximum number of disk shelves?	10
	Which storage shelves are supported?.....	10
	How many nodes are supported by MetroCluster FC?	10
	Which back-end storage is supported in a MetroCluster configuration with array LUNs?	10
	For a FlexArray configuration using array LUNs as storage, can the front-end and back-end fabric switches be combined?.....	11
	Can controller types be mixed in a MetroCluster configuration?	11
	Can MetroCluster nodes be mixed with nodes other than MetroCluster in the same cluster?	11
	How can I verify that the correct SFP is used for the 16Gb FC-VI connection?	11
	Are the port efficiency improvements for the FAS8020 controller supported on the AFF8020 controller?.....	11

Can aggregates be configured as RAID DP or RAID 4?	11
Is a RAID-TEC aggregate supported on an MCC?	11
What is the minimum number of disks that a data aggregate can contain?	11
What are the system limits for MetroCluster?	12
Do MetroCluster configurations with array LUNs require FAS shelves for the root aggregate?	12
Does 8-node MetroCluster support mixing FAS and AFF controllers?	12
How does MetroCluster replicate data in an 8-node configuration?	12
Does MetroCluster support clustered Data ONTAP non-disruptive upgrades (NDU)?	12
Can MetroCluster switch over if the clusters are running different minor Data ONTAP versions?	12
3.1 Stretch MetroCluster	13
What are the maximum distances for stretch MetroCluster configurations?	13
3.2 Configuration – Cluster Interconnect Switches	13
Are cluster interconnect switches required in each HA pair?	13
Which cluster interconnect switches are supported for an 8 or 4-node MetroCluster configuration?	13
How many target FC ports are required when deploying AFF-series controllers in a 2-node or 4-node MetroCluster configuration?	13
3.3 Configuration – Switches – General	13
What switches are supported for the storage fabric?	13
Is there any feature or functional differences between a MetroCluster configuration using FC ISL connectivity and FCIP ISL connectivity?	13
Which switches are supported for FCIP ISL connectivity?	14
Are additional licenses required for the Cisco 9250i/Brocade 7840 FCIP switches?	14
Is it required to purchase the switches from NetApp?	14
Do switches ordered through NetApp include SFPs?	14
What is the required latency for the ISLs?	14
Do all the ISLs in both fabrics have to be the same speed and length?	14
What is the maximum distance that is supported between the two clusters?	15
Does MetroCluster FC support shared ISL links between sites?	15
Can the ISL be shared between the two DR groups in an eight-node MetroCluster configuration?	15
Is traffic isolation used for the ISL?	15
Is out-of-order delivery (OOD) supported on the ISL?	15
Is ISL encryption supported?	15
Is ISL compression supported?	15
Can the FC switches be non-disruptively replaced?	15
Can I verify the FCIP ISL latency between the two MetroCluster clusters?	16
Is there an easy way to configure the fabric switches?	16
3.4 Configuration – Switches – Brocade	16
What is included with the Brocade switches that are ordered through NetApp?	16
How is licensing handled for the reuse of existing supported Brocade switches?	16
Is trunking required for Brocade switches?	16
3.5 Configuration – Switches – Cisco	17
Is compression supported on the Cisco MDS 9250i FCIP switch?	17

How many ports are licensed when I purchase the Cisco MDS 9250i FCIP switch through NetApp?	17
What is included with the Cisco switches that are ordered through NetApp?	17
3.6 Configuration – Fibre Channel to SAS Bridges	17
Which configurations require FibreBridges?	17
Which FibreBridges are supported?	17
What is the maximum number of disk shelves when using a FibreBridges?	17
Since only two FC ports on the ATTO 7500N are used, can one ATTO 7500N be used per stack?	17
Can only one FC port be used for ATTO 6500N FibreBridge connectivity?	18
Can two FC ports be used for the ATTO 7500N FibreBridge connectivity?	18
Can I non-disruptively replace an ATTO 6500N with the ATTO 7500N?	18
What firmware version is supported on a FibreBridge?	18
How do I upgrade the firmware on a FibreBridge?	18
Is the MetroCluster stretch configuration supported with the ATTO 7500N?	18
Can the ATTO 6500N and the ATTO 7500N coexist in the same MetroCluster disk stack?	18
Can the ATTO 6500N and the ATTO 7500N coexist in same MetroCluster configuration?	18
Can disk shelves with unmirrored aggregates be connected directly to the FAS or AFF controller?	18
4 Administration and Operation	18
Does switchover require clients to reconnect?	18
Can the cluster contain any data before initializing MetroCluster?	19
Can mirrored aggregates be converted to an unmirrored aggregate?	19
Can unmirrored aggregates be converted to mirrored aggregates?	19
Should unmirrored aggregates be created from pool0?	19
Can data on the remote plex be accessed by the DR partner during normal operation?	19
What happens if all ISLs in a fabric MetroCluster are down?	19
What happens if the links fail in a two-node stretch MetroCluster configuration?	19
How is the DR partner node chosen in an 8-node and 4-node configuration?	19
Is aggregate relocation (ARL) for non-disruptive controller upgrade supported with MetroCluster?	20
Can a 2-node configuration be converted to a 4-node configuration?	20
Can a 4-node configuration be converted to an 8-node configuration?	20
Can a non-MetroCluster node be converted to MetroCluster?	21
Is vol move supported with MetroCluster?	21
Is it possible to transfer data from a volume in a mirrored aggregate to an unmirrored aggregate?	21
Is vol rehost supported with MetroCluster?	21
Where can I find information about other maintenance operations with MetroCluster?	21
How can existing ONTAP customers transition to MetroCluster?	21
Can nodes that were previously running another ONTAP release be reused for MetroCluster?	21
Does System Setup support MetroCluster?	22
How can MetroCluster be monitored?	22
What are the default passwords for hardware shipped from the factory?	22
What configurations are available for two-node stretch MetroCluster?	22
Can a MetroCluster configuration be returned to a configuration without MetroCluster?	22

What happens if the cluster peering network experiences an interruption?	22
4.1 MetroCluster Site Switchover and Switchback	23
How are switchover and switchback initiated?	23
When is automatic switchover initiated in the 2-node, 4-node, and 8-node MetroCluster configurations?	23
Is it possible to perform a switchover test without actually switching over any objects?	23
Does switchover work if the cluster is not in quorum?	23
4.2 MetroCluster Tiebreaker	23
Where can I find more information about Tiebreaker?	23
5 Interoperability	24
Does MetroCluster provide AutoSupport functionality?	24
Does MetroCluster FC support SnapLock Compliance?	24
Is SnapLock Enterprise supported in MetroCluster?	24
Does MetroCluster FC support NetApp Volume Encryption (NVE)?	24
Does MetroCluster FC support unmirrored aggregates?	24
Does MetroCluster FC support FlexGroups?	24
Does MetroCluster FC support FlashPool?	24
Does MetroCluster FC support NSE?	24
Does MetroCluster FC support ADP?	24
Does MetroCluster FC support FlexArray?	24
Does MetroCluster FC support FabricPools?	24
What capabilities does OCUM provide for MetroCluster?	24
Does System Manager support MetroCluster?	25
How do SnapMirror and SnapVault interoperate with MetroCluster?	25

LIST OF TABLES

Table 1) MetroCluster Node Count	10
--	----

LIST OF FIGURES

Figure 1) MetroCluster FC Diagram	6
Figure 2) MetroCluster Stretch Diagram	6

2 General Information

How do I find out what changed between versions in ONTAP?

Answer: The release notes for ONTAP on the documentation site have details about ONTAP changes. The MetroCluster IP SE Technical Presentation contains some of the feature changes in various past releases for MetroCluster.

How can I find out when a specific feature will be supported?

Answer: Refer the NDA Roadmap Presentations for information about what is planned for MetroCluster.

Is there a license or fee for MetroCluster?

Answer: No. MetroCluster is included in Data ONTAP. There is no separate chargeable license or feature. MetroCluster does require additional hardware depending on the platform and storage configuration (FC, IP or stretch).

What licenses are required for MetroCluster?

There are no specific Data ONTAP licenses for MetroCluster or NetApp SyncMirror® software; all the functionality is included in the base Data ONTAP 8.3.x and ONTAP 9.X product licenses. Additional Data ONTAP licenses, such as for data protocols or NetApp SnapMirror® software, should be added as required and must be installed on both clusters. If the equivalent licenses are not available on the destination cluster, switchover does not work. The two clusters must have separate base licenses that generate distinct cluster serial numbers, otherwise, the metrocluster configure command fails.

Brocade or Cisco switch licenses are covered in the Switch section.

Is SyncMirror supported for configurations without MetroCluster?

Answer: Local SyncMirror is supported for aggregate replication within an HA pair in a cluster without MetroCluster. It does not require a PVR or any additional Data ONTAP license. An aggregate mirrored with SyncMirror is written to two plexes that are owned by the same node. FC-VI is supported only with MetroCluster configurations; therefore, local SyncMirror is restricted to HA distances. Aggregates can be mirrored between plexes on the same node (separate stacks) SyncMirror is not supported for Flash Pool aggregates by using advanced disk partitioning (ADP). However, SyncMirror on an All Flash FAS with root disk partitioning is supported for aggregates on the data partitions.

Note: MetroCluster FC does not support advanced disk partitioning (ADP).

For more information about SyncMirror without MetroCluster, see the High Availability Guide, Data Protection Guide, and Commands Reference publications in the ONTAP Product documentation.

What is the read and write behavior with plexes in SyncMirror in MetroCluster?

Answer: By default, when both plexes of an aggregate are online and in sync, data is read from the local plex only, and data is written synchronously to both plexes. If either of the plexes is offline or in the process of being resynchronized, data is read from and written to only the online plex. Either the local or the remote plex could fail (for example, a shelf failure), and data access would continue seamlessly to the surviving plex.

Optionally, a RAID parameter can be set so that reads alternate from the local and the remote plexes. By default, as just described, reads are from the local plex only. The RAID option can be changed non-disruptively, without affecting application I/O.

To set the option to read from alternate plexes, use the following command on each node in the HA pair:

```
storage raid-options modify -node <node-name> -name  
raid.mirror_read_plex_pref-value alternate
```

The above command is only recommended for Stretch configurations. Otherwise, it can lead to ISL degradation.

To set the option back to the default setting, specify `-value local` on the same command.

In the 8-node and 4-node configurations, NetApp strongly recommends that you set the option to the same setting on both nodes of an HA pair. There might be particular workload environments such that different read behavior on each node gives optimal performance. If the setting is different between the nodes, the configured value does not propagate to the other node in the event of an HA failover.

What is the switchover granularity of MetroCluster?

Answer: MetroCluster uses synchronous replication and can achieve a zero RPO and a low RTO (two minutes for planned or automatic switchover and switchback). MetroCluster switchover operates at the site level: All aggregates, volumes, LUNs, and SVMs switch over to the other site. It is not possible to switch over at a smaller granularity such as an individual SVM.

2.1 Sizing a Solution

How can I determine the capacity requirements for MetroCluster FC?

Answer: Currently it is best to use Synergy or Fusion to size for capacity of an HA pair. Then double the storage devices (shelf, disk or SSD) to compensate for the capacity needed to mirror the aggregates. It may also be required to adjust by 5% higher to allow for SyncMirror reserve capacity.

Does SPM support MetroCluster FC?

Answer: Yes, SPM has support for MetroCluster. System Performance Modeler provides the ability to size a system to meet performance requirements.

<https://spm.netapp.com/>

Does Config Advisor support MetroCluster?

Answer: Yes, Config Advisor supports MetroCluster, however, only Network collection is supported. For more information on using Config Advisor refer to the NetApp Field Portal.

2.2 Quoting MetroCluster

How do I quote an 8 node MetroCluster configuration?

Answer: Quote a 4-node configuration, then add the devices to get to the 8-node configuration.

What specific considerations would be required when quoting the FAS9000 / A700 platforms?

The port count of the switches should be considered. There needs to be 4 FCVI from HA pair to switch, 4 HBA from HA pair to switch, 2 for ATTO, 4 for ISL – total of 14 ports. This forces the need to get either additional 12 port licenses (Cisco 9148S or Brocade 6505) or buy the Brocade 6510 24 port FC switch. Given that customers can't intermix their FC switches between Brocade/Cisco, this could be an issue. So, the number of licensed ports on the switches need to be accounted for while quoting for the system.

2.3 Upgrade of a MetroCluster System

Where do I find information about how to upgrade ONTAP in a MetroCluster configuration?

Answer: Refer to the online documentation in the ONTAP 9 Documentation Center.

<http://docs.netapp.com/ontap-9/index.jsp>

See the section titled Installation and Upgrade, then Upgrades, Revert, or Downgrade. Upgrading software in ONTAP clusters contains a section titled [Updating a MetroCluster configuration nondisruptively](#).

Where do I find information about how to upgrade a controller?

The online AFF and FAS System Documentation center contains information about upgrading controllers. It may also be necessary to temporarily increase the node count in the MetroCluster configuration which is described in the MetroCluster Service and Expansion Guide.

<http://docs.netapp.com/platstor/index.jsp>

See Other Platform Procedures, Upgrading controller hardware.

Refer to the NetApp support site product documentation section for platform information.

Refer to the ONTAP Documentation Center.

<http://docs.netapp.com/ontap-9/index.jsp>

See the section titled MetroCluster Configuration, then MetroCluster Service and Expansion.

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.hw-metrocluster-service/home.html>

Where do I find the information for expanding the number of nodes in a MetroCluster configuration?

Answer: Refer to the online documentation in the ONTAP 9 Documentation Center.

<http://docs.netapp.com/ontap-9/index.jsp>

See the section titled MetroCluster Configuration, then MetroCluster Service and Expansion.

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.hw-metrocluster-service/home.html>

Where can I find the procedure for aggregate relocation?

Answer: The NetApp support site product documentation section has the procedure in the FAS or AFF platform sections.

https://library.netapp.com/ecm/ecm_download_file/ECMLP2659356

3 Configuration

Which FAS controller models are supported?

Answer: Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support Site for currently supported controllers.

All controllers across DR groups must be the same model.

FlexArray and FAS controllers that run FlexArray virtualization software are also supported. Mixing FAS and FlexArray controllers in the same MetroCluster configuration is not supported.

What are the platform mixing rules for the 8-node MetroCluster configuration?

Answer: In the 8-node MetroCluster configuration, there are two 2-node HA pairs at each site making up the cluster. The eight controllers do not need to be the same type. However, each DR group must have four matching controllers at each site. All HA pairs must follow the standard platform mixing rules for ONTAP. For more information, see the Hardware Universe.

Does the 8-node MetroCluster configuration support SAN protocols?

Answer: Yes, with ONTAP 9.0 and later. The 2-node and 4-node MetroCluster configurations continue to support SAN configurations.

Can I mix NAS and SAN in the 8-node MetroCluster configuration?

Answer: Yes, with ONTAP 9.0 and later.

What are the front-end network requirements for MetroCluster FC?

Answer: MetroCluster preserves the identity of switched-over objects. LIF IP addresses, FC WWNs and FC WWPNS are identical after switch-over. Therefore, the TCP/IP layer 2 network and SAN Fabrics must span both sites so that clients and hosts can access the same resources by using the same paths at either site.

Note: For customers that are unable to provide a network that spans the site, A FPVR is required.

What shelves are supported with MetroCluster FC?

Answer: Hardware Universe has the correct information for disk shelves supported for MetroCluster.

Where do I find the maximum number of disks?

Answer: Hardware Universe has the correct information for maximum disks supported for MetroCluster.

Where do I find the maximum number of disk shelves?

Answer: Hardware Universe has the correct information for maximum disk shelves supported for MetroCluster.

Which storage shelves are supported?

Answer: Refer to the Interoperability Matrix Tool (IMT) for shelf support for the specific platform and ONTAP version.

Note: ACP (alternate control path) is not supported with MetroCluster.

Note: Mini-SAS cables do provide Glowworm support with some NetApp storage shelves.

How many nodes are supported by MetroCluster FC?

Answer: Currently, 2-node, 4-node and 8-node configurations are available.

Table 1) MetroCluster Node Count

Number of nodes	Added in ONTAP
2-node	8.3
4-node	8.3
8-node	9.1

Note: A 2-node configuration is a single-node cluster at each site.

Which back-end storage is supported in a MetroCluster configuration with array LUNs?

Answer: EMC, HDS, HP 3PAR, and NetApp E-Series storage (including EF-Series) are supported. For the most current list, see the IMT. The configuration must be identical at both sites: the same FAS or FlexArray model, vendor array, disk type, and LUN size. Starting with Data ONTAP 8.3, FAS disk shelves are optional in a ONTAP configuration that includes storage arrays. The root aggregate can be created on array LUNs, and the controllers can boot from those LUNs. Thus, a MetroCluster configuration with

storage arrays can exclusively use array LUNs or a combination of NetApp FAS shelves and array LUNs. The FibreBridge is used only to attach FAS disk shelves.

For a FlexArray configuration using array LUNs as storage, can the front-end and back-end fabric switches be combined?

Answer: No. MetroCluster requires dedicated switches for the back-end fabric switches.

Can controller types be mixed in a MetroCluster configuration?

Answer: In the 2-node and 4-node configurations, all controllers in the MetroCluster DR group must be the identical FAS, AFF, or FlexArray model. In an 8-node configuration each node in a DR group must be identical. The two four node DR groups in an 8-node configuration can vary, e.g. DR group #1 can have one type while DR group #2 can have a different type.

AFF, FAS and FlexArray nodes cannot be mixed in the DR group. DR partner nodes should have the same hardware configuration, including storage and front-end network ports, to allow LIFs to be assigned correctly to ports on the other cluster during switchover.

Can MetroCluster nodes be mixed with nodes other than MetroCluster in the same cluster?

Answer: No, all nodes must be configured for MetroCluster.

How can I verify that the correct SFP is used for the 16Gb FC-VI connection?

Answer: The SFP for the FC-VI 16Gb ports must be the QLogic-branded SFP. For instructions on how to verify the correct SFP, and for information about symptoms if incorrect SFPs are used, see the following KB article: https://kb.netapp.com/app/answers/answer_view/a_id/1015352

Are the port efficiency improvements for the FAS8020 controller supported on the AFF8020 controller?

Answer: Yes. Both the FAS8020 and the AFF8020 support sharing FC-VI and FC connectivity on the same initiator card. These improvements are available only on ONTAP MetroCluster configurations.

Can aggregates be configured as RAID DP or RAID 4?

Answer: Yes, either NetApp RAID DP technology or RAID 4 is supported for aggregates (including the root aggregate) on each node.

Is a RAID-TEC aggregate supported on an MCC?

Answer: Yes, RAID-TEC is supported.

What is the minimum number of disks that a data aggregate can contain?

Answer: For RAID 4 data aggregates, the minimum number of disks is 3 (2 data, 1 parity).

For RAID DP, the minimum number of disks is 5 (3 data, 2 parity). As for any ONTAP configuration, having more drives in an aggregate is likely to give better performance, and aggregates with few drives do not perform optimally.

Does storage need to be added in fixed increments?

Storage added after the initial MetroCluster configuration does not need to be in any set number of shelves. One shelf at a time per site can be added. NetApp recommends that you add storage so that each shelf contains disks in only one pool.

What are the system limits for MetroCluster?

Answer: Detailed limits are documented in a document that is specifically about MetroCluster in the Maximum Configuration Guides on the Field Portal. These limits are the maximum supported for MetroCluster configurations and have been tested to achieve negotiated switchover and switchback completion within two minutes. Note that all SMB sessions experience a disruption on both switchover and switchback in all configurations. These limits are not enforced by ONTAP, but they are the documented supported limits. If the number of volumes per node is exceeded, an EMS message (soft limit) is issued.

Do MetroCluster configurations with array LUNs require FAS shelves for the root aggregate?

Answer: No, MetroCluster can be configured exclusively with array LUNs (including the root aggregates), exclusively with FAS disk shelves, or with a combination of both. If you deploy MetroCluster with both FAS disk shelves and array LUNs, NetApp recommends that you create the root aggregate or volumes on FAS disk shelves and use the array LUNs for data aggregates. With this configuration, the minimum FAS disk shelf requirement is two shelves, or one shelf per site, with 48 drives total. This approach makes six drives available for each node's root aggregate and spares, and the same number of disks for the mirror copy.

Does 8-node MetroCluster support mixing FAS and AFF controllers?

Answer: Yes. Mixing FAS and AFF HA pairs is supported in the 8-node MetroCluster configuration.

How does MetroCluster replicate data in an 8-node configuration?

Answer: Data is replicated between an HA pair on the primary site and an HA pair at the partner site.

Does MetroCluster support clustered Data ONTAP non-disruptive upgrades (NDU)?

Answer: Yes. In the 8-node and 4-node configurations, NDU is used in MetroCluster to upgrade the ONTAP version. This method uses takeover and giveback on the HA pairs at each site to upgrade ONTAP non-disruptively. For a major update, the two sites must be updated in parallel to preserve the DR protection. For minor release updates (for example, from 9.2.0 to 9.2.1), the two clusters can be upgraded independently and still preserve DR protection. However, NetApp recommends that you complete the upgrade of both sites as soon as possible. Automated NDU can be used for minor ONTAP upgrades at each site.

In 2-node configurations, a negotiated switchover is used to update each cluster in turn non-disruptively.

Can MetroCluster switch over if the clusters are running different minor Data ONTAP versions?

Answer: Minor upgrades should be completed across both clusters as soon as possible. If the two clusters are running different minor versions, administrators should not perform unplanned switchover until all nodes have been upgraded to the same release. If an unplanned switchover is required when the clusters have different minor versions, the operation will succeed however there may be some cluster configuration data not replicated. After switchover in this instance, before performing the switchback, it is necessary to heal the root aggregates, boot the recovering nodes, and upgrade them to the same ONTAP release as the surviving nodes. After both clusters are running the same release, the switchback can proceed.

3.1 Stretch MetroCluster

What are the maximum distances for stretch MetroCluster configurations?

Answer: See Hardware Universe or Interoperability Matrix Tools (IMT) for the current distances supported.

3.2 Configuration – Cluster Interconnect Switches

Are cluster interconnect switches required in each HA pair?

Answer: The cluster interconnect switches are optional in the 4-node MetroCluster configurations. In the 8-node configuration cluster interconnect switches are required. NetApp expects that, for cost and simplicity, the majority of four-node MetroCluster configurations consist of two-node switchless clusters. Because the two-node configuration uses single-node clusters, there is no cluster interconnect.

Which cluster interconnect switches are supported for an 8 or 4-node MetroCluster configuration?

Answer: The HA pairs at each site can be configured as switched or switchless in a 4-node configuration. An 8-node configuration requires switches for the cluster interconnect. Switched configurations can use any cluster interconnect switch supported in ONTAP. The cluster interconnect switch must be the same at both sites. For the list of supported cluster interconnect switches, see the Hardware Universe.

How many target FC ports are required when deploying AFF-series controllers in a 2-node or 4-node MetroCluster configuration?

Answer: When deploying MetroCluster on AFF-series controllers, the best practice is to use two target FC ports for optimal performance.

3.3 Configuration – Switches – General

What switches are supported for the storage fabric?

Answer: Four identical switches (same vendor and model) are required for each MetroCluster configuration. Each MetroCluster configuration requires dedicated switches; sharing between multiple MetroCluster configurations or with host or application traffic is not supported. All four switches must be licensed for the same number of ports.

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support Site for a current list of supported switches.

Is there any feature or functional differences between a MetroCluster configuration using FC ISL connectivity and FCIP ISL connectivity?

Answer: Starting in clustered Data ONTAP 8.3.2, NetApp provides FCIP ISL connectivity as an additional option for ISL connectivity, in addition to FC ISL connectivity. Both FC and FCIP MetroCluster configurations provide synchronous replication and continuous availability to protect your data.

There are a few differences between MetroCluster FC and FCIP MetroCluster configurations:

- The Cisco MDS 9250i cannot be connected to more than one MetroCluster cluster.
- Only 10Gbps FCIP ISL connectivity is supported.
- Only one FCIP port is supported.

- Link-write acceleration is not supported.
- Compression is not supported.
- The FCIP ISLs between the MetroCluster clusters cannot be shared.

Which switches are supported for FCIP ISL connectivity?

Answer: The Cisco MDS 9250i is supported for FCIP ISL connectivity. Only 10Gbps FCIP Ethernet ports are supported for ISL connectivity on the Cisco 9250i. FCIP MetroCluster configurations are supported starting in ONTAP 8.3.2. The Brocade 7840 is also supported.

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support Site for a current list of supported switches.

Are additional licenses required for the Cisco 9250i/Brocade 7840 FCIP switches?

Answer: No. When the Cisco 9250i FCIP switches are purchased from NetApp, they come bundled with Enterprise licenses and SFPs. At this time the Brocade 7840 is already on the price list, but components/licenses need to be ordered individually.

Many restrictions apply specifically to Cisco MDS 9250i, please refer to the Interoperability Matrix Tool (IMT) and Config Advisor for feature comparability between 9250i and Brocade 7840.

Is it required to purchase the switches from NetApp?

Answer: Yes, all switches for the back-end storage fabric must be ordered and purchased from NetApp. A customer can reuse existing switches if the switches are on the support matrix and only if they were purchased from NetApp. NetApp provides initial-level support for the switches that are used in MetroCluster.

Do switches ordered through NetApp include SFPs?

Answer: Yes. The appropriate number of shortwave SFPs are included for all Brocade and Cisco switches that are ordered through NetApp. 16Gb SFPs are included for all these switches, except the Cisco 9148, which includes 8Gb SFPs. The Cisco 9148S, however, does include 16Gb SFPs. To connect the switch to ISLs, shortwave or longwave SFPs can be used. Brocade 16Gb ISL SFPs can be ordered separately (X-R000198-R6).

What is the required latency for the ISLs?

Answer: NetApp does not specify the required round-trip latency on the ISLs. Rather, NetApp specifies the maximum supported distance and no active components. For MetroCluster deployments using the Brocade 6505 or 6510 FC switch the maximum distance is 185 miles (300 kilometers). All other configurations using FC or FCIP switches have a maximum distance of 120 miles (200 kilometers). The ISL must conform to the requirements that were stated earlier and must be supported by the FC switch and SFP.

Do all the ISLs in both fabrics have to be the same speed and length?

Answer: All ISLs between the switches should be the same speed, and there should be the same number of ISLs in each fabric (from one to eight). If multiple ISLs are in a fabric, all ISLs in the same fabric must be the same physical length. ISLs in one fabric can be of different physical lengths than the ISLs in the second fabric (for example, if different vendors are used for the two connections). However, NetApp recommends for simplicity and supportability that all ISLs in the same MetroCluster configuration be the same length.

What is the maximum distance that is supported between the two clusters?

Answer: The maximum supported distance is 300 kilometers or 185 miles, depending on the switch that you select. Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support Site for specifics for platforms and switch configurations.

Does MetroCluster FC support shared ISL links between sites?

Answer: No, links shared with non-MetroCluster traffic are not supported.

Can the ISL be shared between the two DR groups in an eight-node MetroCluster configuration?

Answer: Yes. The ISL must be shared with the two DR groups in the eight-node MetroCluster configuration. As of 9.2 ISL sharing is supported on the two-node and four-node MetroCluster configurations allowing multiple MetroCluster systems to share the ISL.

For more information see Considerations for ISLs in the Fabric-attached MetroCluster Installation and Configuration Guide in the ONTAP documentation.

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.dot-mcc-inst-cnfg-fabric/GUID-CEEE2C50-AB13-48BD-86D9-2AF37D69D901.html>

Is traffic isolation used for the ISL?

Answer: No, MetroCluster uses trunking and QoS to balance and consolidate ISL utilization.

Is out-of-order delivery (OOD) supported on the ISL?

Answer: Yes, out-of-order delivery is supported starting ONTAP 9.1.

For more information see Considerations for ISLs in the Fabric-attached MetroCluster Installation and Configuration Guide in the ONTAP documentation.

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.dot-mcc-inst-cnfg-fabric/GUID-CEEE2C50-AB13-48BD-86D9-2AF37D69D901.html>

Is ISL encryption supported?

Answer: Yes. MetroCluster supports ISL encryption with the Brocade 6510 switch. Encryption is also supported for WDM devices.

For more information see the section titled Considerations for using TDM/xWDM equipment with fabric-attached MetroCluster configurations in the Fabric-attached MetroCluster Installation and Configuration Guide in the ONTAP product documentation.

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.dot-mcc-inst-cnfg-fabric/GUID-C54AEB3F-A332-4030-94CD-198284D3762E.html>

Is ISL compression supported?

Answer: No, compression is not supported and should not be enabled on the switches.

Can the FC switches be non-disruptively replaced?

Answer: The FC switch can be non-disruptively replaced if doing an upgrade of the switch from the same vendor. If changing the switch vendor, then it cannot be done non-disruptively.

Can I verify the FCIP ISL latency between the two MetroCluster clusters?

Answer: Yes. MetroCluster provides the ability to verify the FCIP ISL latency between the two MetroCluster clusters using the statistics CLI command. MetroCluster tracks the maximum, minimum, and average latency of the FCIP ISL every 15 minutes for the last four hours.

In addition, latency can be measured on the Cisco 9250i FCIP switch using the ips CLI command.

For more information, see the knowledge base article #1015624.

Note: Average latency between the two MetroCluster clusters must not exceed 2ms.

Is there an easy way to configure the fabric switches?

Answer: Reference configuration files for the Brocade and Cisco switches are available for download from the NetApp Support site. The files contain the required commands to automatically configure the switches. For instructions refer to the ONTAP documentation Fabric-attached MetroCluster® Installation and Configuration Guide.

3.4 Configuration – Switches – Brocade

What is included with the Brocade switches that are ordered through NetApp?

Answer: The Brocade switches that are ordered through the Quote Tool in a new MetroCluster configuration include the following items:

- Four Brocade switches (6505 or 6510)
- 16Gb SFPs
- An Enterprise license bundle:
 - Fabric Vision—includes Fabric Watch and Advanced Performance
 - Monitoring
 - Extended fabric (required if the distance is greater than 3.75 miles [6 kilometers])
 - ISL trunking (used if there are two or more ISLs per fabric)
- Brocade Network Advisor Pro+, which includes the capability for switches to send event messages to AutoSupport and can be used with any of the Brocade switches in your environment

How is licensing handled for the reuse of existing supported Brocade switches?

Answer: If existing switches are reused (for example, a 6510 from an FMC), the extended fabric and trunking licenses can be ordered if the switches do not already include them. NetApp also strongly recommends the Brocade Network Advisor Pro+ license for advanced monitoring and My AutoSupport integration with CallHome.

This license is likely already available with switches that are purchased through NetApp, or it can also be ordered as an add-on. In QuoteEdge, select Brocade Network Advisor (BNA) in the Software section. Be sure to click More to display the entire list. Scroll down the BNA list of licenses to SAN Pro Plus License for 2560 Switch Ports, select this license, and continue the QuoteEdge process. The PN SW-BSW-NTWADV-SAN-PLU is included in the quote.

Is trunking required for Brocade switches?

Answer: If two or more ISL connections are in the fabric, then trunking should be used when connecting the Brocade switches. If trunking is not an option, then OOD must be utilized. OOD is supported starting with ONTAP 9.1. For a single ISL connection, trunking is not required. NetApp recommends a minimum of

two ISLs. You should verify that the xWDM vendor supports trunking on the Brocade FC switch. In cases where the xWDM vendor does not support trunking, you can either use a different xWDM vendor or file a PVR with NetApp to use traffic isolation (maximum of two ISLs per fabric). With 9.1, PVR is no longer needed.

3.5 Configuration – Switches – Cisco

Is compression supported on the Cisco MDS 9250i FCIP switch?

Answer: No. Compression is not supported on the Cisco MDS 9250i FCIP switch.

How many ports are licensed when I purchase the Cisco MDS 9250i FCIP switch through NetApp?

Answer: By default, there are 12 ports licensed on each FCIP switch. This can be increased up to 48 ports in the Quote Tool.

What is included with the Cisco switches that are ordered through NetApp?

Answer: The Cisco switches that are ordered through the Quote Tool include the following items:

- Four Cisco switches:
 - Cisco MDS 9148 in clustered Data ONTAP 8.3.0
 - Cisco MDS 9148S in clustered Data ONTAP 8.3.1
 - Cisco MDS 9250i in clustered Data ONTAP 8.3.2 for FCIP ISL connectivity
- SFP modules:
 - Cisco 9148 is 8Gbps
 - Cisco 9148S and Cisco 9250i are 16Gbps

Note: The Cisco 9148S and 9250i switches bundle the Enterprise package into the same part number.

3.6 Configuration – Fibre Channel to SAS Bridges

Which configurations require FibreBridges?

Answer: All eight-node and four-node configurations with NetApp FAS shelves require FibreBridges. The two-node fabric and two-node stretch configurations with FibreBridge also require them. Only the two-node configuration with SAS optical cables does not use FibreBridges, because the storage connects directly to SAS ports on the two controllers.

Which FibreBridges are supported?

Answer: Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support Site for a current list of supported bridges.

What is the maximum number of disk shelves when using a FibreBridges?

Answer: Refer to the Hardware Universe and the MetroCluster ONTAP Maximum Configuration Guides.

Since only two FC ports on the ATTO 7500N are used, can one ATTO 7500N be used per stack?

Answer: No. Each storage stack requires two ATTO 7500N FibreBridges for redundancy.

Can only one FC port be used for ATTO 6500N FibreBridge connectivity?

Answer: Yes. MetroCluster supports using only one FC port with the ATTO 6500N fabric configuration.

Can two FC ports be used for the ATTO 7500N FibreBridge connectivity?

Answer: Yes. When connecting multiple stacks to the ATTO FibreBridge, the best practice is to use two FC ports for optimal performance.

Can I non-disruptively replace an ATTO 6500N with the ATTO 7500N?

Answer: Yes. the ATTO 6500N can be replaced with the ATTO 7500N without a disruption.

For more information, view the procedure in the MetroCluster Service Guide in the ONTAP documentation.

What firmware version is supported on a FibreBridge?

Answer: Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support Site for a current version of supported firmware for FibreBridges.

How do I upgrade the firmware on a FibreBridge?

Answer: Instructions for upgrading the FibreBridge firmware is available on the NetApp Support site in the MetroCluster Service Guide.

Is the MetroCluster stretch configuration supported with the ATTO 7500N?

Answer: Yes. The ATTO 7500N supports stretch configurations using short-wave (SW) optics.

Can the ATTO 6500N and the ATTO 7500N coexist in the same MetroCluster disk stack?

Answer: Yes. A disk stack can mix the ATTO 6500N and the ATTO 7500N in the same MetroCluster configuration (for example, a different FibreBridge at each end of the disk stack). However, in this configuration, only one FC and one SAS can be used. It is required to use SAS port 1 and FC1 in this configuration. Disable SAS ports 2, 3, 4 and FC2. This behavior is similar to the ATTO 6500N.

Can the ATTO 6500N and the ATTO 7500N coexist in same MetroCluster configuration?

Answer: Yes. One disk stack can connect to a pair of ATTO 6500N FibreBridges and another disk stack can connect to a pair of ATTO 7500N FibreBridges in the same MetroCluster configuration.

Can disk shelves with unmirrored aggregates be connected directly to the FAS or AFF controller?

Answer: MetroCluster direct-attach configurations supports both mirrored and unmirrored aggregates which are directly to the controller. MetroCluster bridge-attach and fabric configurations require all disk shelves to connect to the ATTO bridges for both mirrored and unmirrored shelves. All the disks are required to be visible to all the nodes in the DR group.

4 Administration and Operation

Does switchover require clients to reconnect?

Answer: For NFS and SAN, a reconnection of the client is not required after switchover. For SMB, switchover is disruptive to all SMB applications and clients, including SMB 3.0. This applies to all MetroCluster architectures.

Can the cluster contain any data before initializing MetroCluster?

Answer: No. There cannot be any SVMs existing before running the metrocluster configure command. If data is already on the cluster, the command fails.

Can mirrored aggregates be converted to an unmirrored aggregate?

Answer: Yes. Starting in ONTAP 9.0 unmirrored aggregates are supported. Existing mirrored aggregates can be converted to unmirrored aggregates with the plex delete CLI command. Deleting a mirrored aggregate plex is disallowed when not in switchover when:

- The aggregate is hosting an SVM root volume
- The aggregate is hosting a MetroCluster MDV (metadata volume)
- The aggregate is a MetroCluster root aggregate

Can unmirrored aggregates be converted to mirrored aggregates?

Answer: Yes. Starting in ONTAP 9, MetroCluster supports conversion of unmirrored aggregates to mirrored aggregates. Unmirrored aggregates can be converted to mirrored aggregates with the `mirror <aggregate_name>` CLI command.

Should unmirrored aggregates be created from pool0?

Answer: Unmirrored aggregates should only be built from disks located in pool0 (local to the node). Otherwise all IO would traverse the ISL and may cause congestion and performance problems. Multi-disk panic's can also occur if the ISL connections fail.

Can data on the remote plex be accessed by the DR partner during normal operation?

Answer: No, a volume or LUN is online on only one cluster at a time and cannot be accessed on the other cluster. Client or application access (including NetApp SnapMirror, SnapVault®, and FlexClone® technology) is not possible on the remote cluster unless a switchover is performed. It is not possible to temporarily “break the relationship” so that operations such as SnapMirror, SnapVault, or FlexClone can be performed by the remote cluster. Remember that MetroCluster uses SyncMirror at the aggregate level, not volume-based SnapMirror.

What happens if all ISLs in a fabric MetroCluster are down?

Answer: Each cluster continues to serve its own workloads, writing data to its local plex only. No storage SyncMirror or NVRAM replication occurs as long as there are no available ISLs between the clusters. As soon as at least one link is restored, the NVRAM replication resumes, and each aggregate's remote plex automatically resyncs with the local plex.

What happens if the links fail in a two-node stretch MetroCluster configuration?

Answer: In 2-node fabric MetroCluster, because storage and FC-VI replication occur over the same ISLs, the behavior is the same as in 4-node MetroCluster. The two nodes in the fabric MetroCluster continue to run independently.

In stretch MetroCluster, the storage is connected directly to the nodes, and the FC-VI link for NVRAM replication is a separate connection.

Depending on which link fails, NVRAM replication can be lost, or RAID sync replication can be lost as well. This will lead to loss of DR protection.

How is the DR partner node chosen in an 8-node and 4-node configuration?

Answer: Each node's DR partner is automatically selected during execution of the metrocluster configure command, based on its system ID (NVRAM ID). The lower-ordered node in one HA pair is DR-partnered

with the corresponding lower-ordered node in the other cluster, and so on. The DR partner assignment cannot be changed after the initial configuration. Particularly if the nodes in an HA pair are configured with different storage, check that the two nodes with lower system IDs in each cluster have a matching disk configuration. Perform the same check for the two nodes with higher system IDs.

Adjust the disk ownership before the initial installation (before running the metrocluster configure command) to confirm that the DR partners are assigned correctly. The assignment of DR partners also affects the ports that LIFs use on switchover; in normal circumstances, LIFs switch over to their DR partner node.

For more information about the behavior of LIFs on switchover, see the section Networking and LIF creation guidelines for MetroCluster configurations in the Fabric-attached MetroCluster Installation and Configuration Guide in the ONTAP documentation.

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.dot-mcc-inst-cnfg-fabric/GUID-C23E5DBF-373B-443F-8286-F29C7D62F8D8.html>

Is aggregate relocation (ARL) for non-disruptive controller upgrade supported with MetroCluster?

Answer: Yes. In the 8-node and 4-node configurations, the process is performed one site at a time, not in parallel. Switchover is not used during the process.

All nodes must be upgraded to the same controller model. HA mirroring within the HA pair is disabled as each cluster is being upgraded. NVRAM mirroring between the sites must be disabled for the entire process. With HA and NVRAM mirroring disabled, DR protection is also disabled during the entire upgrade process, although storage replication with SyncMirror continues. After all controllers have been upgraded, NVRAM mirroring is reestablished. Planned switchover is not possible while the upgrade is in process. Unplanned switchover (on disaster) is possible, with the potential for loss of, at most, the last 10 seconds of transactions, due to the unmirrored NVRAM.

The ARL process for non-disruptive controller upgrade is lengthy, and the documented procedure must be followed precisely. See the manual Using Aggregate Relocation to Upgrade Controller Hardware on a Pair of Nodes Running ONTAP, which is available on the NetApp Support site.

In the two-node configuration, controller upgrade is performed by using negotiated switchover.

Can a 2-node configuration be converted to a 4-node configuration?

Answer: Starting in ONTAP 9, 2-node fabric and stretch configurations can be converted to 4-node fabric configuration. The 2-node configuration must be running ONTAP 9.0, as earlier versions are not supported for conversion. Prior to ONTAP 9, conversion from a 2-node configuration to a 4-node configuration is not supported. NetApp strongly recommends that customers who need local HA functionality deploy the 4-node architecture. The procedure is disruptive and requires the clusters at both sites to be powered down during the conversion. The conversion is in-place and does not require a data copy. Conversion from 2-node stretch to 4-node fabric requires additional switches and rewiring.

There are settings (haosc value) which needs to be changed in maintenance mode and that changes how the NVRAM is replicated. If the nodes aren't offline when these changes are made, there will be NVRAM corruption which would result in panic and possibly data corruption.

Can a 4-node configuration be converted to an 8-node configuration?

Answer: Yes. Starting in ONTAP 9, four-node fabric configurations can be converted to an eight-node fabric configuration. The conversion is in-place and does not require a data copy. The four-node configuration must be running ONTAP 9.0, as earlier versions are not supported for conversion. Please refer to MetroCluster installation guide for detailed procedure.

Can a non-MetroCluster node be converted to MetroCluster?

Answer: No, a non-MetroCluster node cannot be converted to MetroCluster. The equipment can be repurposed, but all data needs to be wiped out before using it in a MetroCluster setup.

Is vol move supported with MetroCluster?

Answer: Yes, vol move is supported in the 8-node and 4-node configurations. As in standard ONTAP, volumes move between aggregates in a cluster and stay within the same SVM. MetroCluster does not provide the capability to non-disruptively move volumes from one cluster to another. If data needs to be transferred from one cluster to another, SnapMirror or any other host or application-based copy mechanism can be used.

When a volume is moved, the local and remote plexes in the source and destination aggregates are synchronized with SyncMirror. The VLDB changes are replicated to the secondary cluster over the cluster peering network by using CRS.

Because the controllers belong to different clusters, vol move between MetroCluster controllers across sites is not possible. Vol move is possible only between aggregates that are attached to the same node within the same cluster.

Is it possible to transfer data from a volume in a mirrored aggregate to an unmirrored aggregate?

Answer: Yes. Volumes in a mirrored aggregate can be moved to an unmirrored aggregate using the vol move command. Moving an SVM root volume or MetroCluster MDV to an unmirrored aggregate is not supported.

Is vol rehost supported with MetroCluster?

Answer: Vol rehost is not supported in MetroCluster deployments.

Where can I find information about other maintenance operations with MetroCluster?

Answer: The MetroCluster Service Guide, available on the NetApp Support site under MetroCluster documentation, provides procedures for operations such as updating firmware and replacing FC switches and FibreBridges, and adding and removing storage shelves.

How can existing ONTAP customers transition to MetroCluster?

Answer: Transitioning from ONTAP requires data migration to a separate complete MetroCluster configuration; in-place transition is not available. This requirement means that you cannot convert an existing ONTAP configuration with data into MetroCluster; if any SVMs are present on the cluster, the metrocluster configure command fails. Furthermore, any nodes that are reused must be wiped (zero all disks) and reinitialized to install ONTAP for MetroCluster.

SnapMirror can be used to transition data from a ONTAP cluster to a MetroCluster configuration. However, 32-bit data (volumes or Snapshot copies) cannot be copied with SnapMirror; for more information, see the 7MTT section.

Application and client-level methods, such as rsync or VMware Storage vMotion, can also be used to migrate data from a ONTAP environment to MetroCluster.

Can nodes that were previously running another ONTAP release be reused for MetroCluster?

Answer: Yes, the nodes can be reused, as long as the hardware is supported with a current release of ONTAP (including four identical controllers). In addition, all the required additional hardware for MetroCluster must be in place, including switches, additional storage, FibreBridges if required, FC-VI cards, and FC initiator ports. All nodes must be wiped (zero all disks) and installed freshly with ONTAP.

Does System Setup support MetroCluster?

Answer: Yes, System Setup can be used to create the initial clusters for supported platforms. The CLI is used for the steps that are specific to MetroCluster, such as mirroring the aggregates, creating the MetroCluster configuration, and configuring the health monitors. You can perform other steps, such as cluster peering and creating intercluster LIFs, by using OnCommand System Manager or the CLI. For detailed steps to set up MetroCluster, see the MetroCluster Installation and Configuration Guide in the NetApp product documentation

How can MetroCluster be monitored?

Answer: MetroCluster includes health monitors for the components, including switches and FibreBridges as applicable. OnCommand Unified Manager and OnCommand Performance Manager also support MetroCluster.

What are the default passwords for hardware shipped from the factory?

Answer: The preset passwords for controllers and the preset SP, cluster interconnect switches, back-end fabric switches, FibreBridges, and cluster peering passphrases are listed in the MetroCluster Installation Express Guide.

What configurations are available for two-node stretch MetroCluster?

Answer: There are two 2-node stretch MetroCluster configurations. The direct-attached configuration connects the controllers and disk shelves directly by using optical or copper SAS cables. The bridge-attached configuration connects the controllers and disk shelves by using ATTO FibreBridges. Mixing direct-attached and bridge-attached SAS connections is not supported.

For more information, see the Optical SAS SE Technical Presentation on the Field Portal.

Can a MetroCluster configuration be returned to a configuration without MetroCluster?

Answer: MetroCluster can be returned to a standard clustered Data ONTAP with the assistance of NetApp Support. Returning to clustered Data ONTAP without MetroCluster:

- Stops all SyncMirror replication
- Deletes the remote plexes
- Changes the SVMs back to the nonreplicated type
- Deletes CRS replicated objects in the RDB
- Removes the DR protection and ability to switch over

It does not remove or alter any user data in the SVMs, however; all data is accessible from the original HA pair after the switchover.

Note: Decommissioning MetroCluster is a one-way operation. If you want to subsequently reenable MetroCluster, you must first zero all drives and reinitialize all the nodes.

What happens if the cluster peering network experiences an interruption?

Answer: Cluster configuration changes are logged in the metadata volume (MDV) on the cluster that is performing the update. Changes are not committed to the local RDB until the logging is complete. Updates are then propagated near synchronously to the other cluster's RDB over the cluster peering network. If changes cannot be propagated because of temporary errors in the cluster peering network, after connectivity is restored, the changes are automatically sent to the other cluster by replaying the stored changes in the MDV.

If an unplanned switchover is necessary when the cluster peering network is down and there are unpropagated RDB updates, the updates are processed from the mirrored copy of the MDV at the

surviving site. For resiliency, NetApp recommends redundant networks for the cluster peering network. If the cluster peering network is unavailable, health monitor alerts and EMS messages are sent, and planned switchover is not possible until at least one link is restored.

4.1 MetroCluster Site Switchover and Switchback

How are switchover and switchback initiated?

Answer: Switchover and switchback operations are initiated by the CLI. There is one command to switch over. By default, the switchover command executes a planned switchover; switchover on disaster is initiated by using a flag on the switchover command. Three commands are required to perform a switchback. Switchback is always a negotiated operation. Switchover and switchback operations are described in the MetroCluster Management and Disaster Recovery Guide in the NetApp product documentation.

When is automatic switchover initiated in the 2-node, 4-node, and 8-node MetroCluster configurations?

Answer: In 2-node configurations, hardware failure or panic of one of the nodes triggers an automatic switchover. In a 4-node and 8-node configurations, Automatic Switchover is performed when all the nodes in an HA-group go down. Automatic switchover occurs only if the local storage (in particular, the root aggregate of the failed node) is available. If a disaster affects all the hardware, including the storage, then automatic switchover is not performed.

For 2-node, 4-node, and 8-node configurations, you must run the Tiebreaker software in automatic switchover mode in the event of a site disaster.

Is it possible to perform a switchover test without actually switching over any objects?

Answer: Yes. MetroCluster includes an option (-simulate) on the switchover and switchback commands in advanced mode that performs essentially a dry run for the switchover and switchback operations. It invokes all the processing and health checks to verify that the operation can be performed, including testing for any conditions that can cause the operation to be vetoed. It terminates before any operations are committed and reports success or failure.

Does switchover work if the cluster is not in quorum?

Answer: The switchover completes if the cluster being switched over (the down site) is out of quorum. However, if the cluster doing the switchover (the site being switched over to) is out of quorum, the switchover fails.

Note: In the event of an outage, Tiebreaker initiates a switchover regardless of the quorum status of the clusters. If the cluster that is being switched over is out of quorum, however, the switchover will not complete.

4.2 MetroCluster Tiebreaker

Where can I find more information about Tiebreaker?

Answer: Refer to the ONTAP documentation. The MetroCluster Configuration section contains the Tiebreaker Software Installation and Configuration Guide.

5 Interoperability

Does MetroCluster provide AutoSupport functionality?

Answer: Yes, MetroCluster sends information using AutoSupport and is available in the AutoSupport section of the NetApp support site.

Does MetroCluster FC support SnapLock Compliance?

Answer: No, SnapLock Compliance is not supported.

Is SnapLock Enterprise supported in MetroCluster?

Answer: Yes. Starting in ONTAP 9 SnapLock Enterprise is supported in MetroCluster.

Does MetroCluster FC support NetApp Volume Encryption (NVE)?

Answer: Yes, starting with ONTAP 9.1 NetApp Volume Encryption is supported.

Does MetroCluster FC support unmirrored aggregates?

Answer: Yes. Currently unmirrored aggregates are supported in MetroCluster FC.

Does MetroCluster FC support FlexGroups?

Answer: No, FlexGroups are not supported.

Does MetroCluster FC support FlashPool?

Answer: Yes, FlashPool is supported. For high-performance environments using AFF is a preferred solution

MetroCluster cannot use Flash Pool aggregates that use storage pools for their cache storage. In addition, SyncMirror cannot use Flash Pool aggregates that use storage pools for their cache storage.

Does MetroCluster FC support NSE?

Answer: No, NetApp Storage Encryption (NSE) is not supported. NSE supports "self-encrypting" disks (SEDs).

Does MetroCluster FC support ADP?

Answer: No, Advanced Drive Partitioning (ADP) is not supported for any MetroCluster FC or MetroCluster stretch configurations.

Does MetroCluster FC support FlexArray?

Answer: Yes, FlexArray is supported for MetroCluster FC. FlexArray allows for use of LUNs from E-Series or specific 3rd party vendors for storage instead of using NetApp FAS disk shelves.

Does MetroCluster FC support FabricPools?

Answer: No. FabricPools, using ONTAP with an object store, are not currently supported.

What capabilities does OCUM provide for MetroCluster?

Answer: Starting in version 6.2, OnCommand Unified Manager provides a graphical display of the MetroCluster topology. It also provides monitoring and alerting for all MetroCluster components, including the ISL and IP links and the switches, bridges, and storage. It is intended to supplement and augment

existing enterprise monitoring tools. OCUM receives alerts from the MetroCluster health monitor and displays probable causes, effects, and suggested corrective actions.

Does System Manager support MetroCluster?

Answer: OnCommand System Manager can be used for normal cluster configuration such as creating and displaying cluster peering relationships, SVMs, volumes, LUNs, and data access protocols. System Manager does not support discovery or display of MetroCluster relationships or configuration specific to MetroCluster; however, some events that are related to MetroCluster are displayed. System Manager does not provide the capability to initiate switchover and switchback operations

How do SnapMirror and SnapVault interoperate with MetroCluster?

Answer: Volumes in MetroCluster can be a SnapMirror or a SnapVault source or target. To allow replication to continue after a switchover or a switchback, peering relationships with the clusters without MetroCluster must be created for both MetroCluster clusters. Any replication tasks that are in progress when a switchover or a switchback occurs terminate. You can run SnapMirror or SnapVault operations only from the cluster that owns the volume. Volumes are online only at one cluster, and the other cluster cannot access the volumes unless a switchover is performed.

If MetroCluster volumes are a SnapMirror or SnapVault source, the peering relationships are automatically updated on switchover or switchback, and replication resumes automatically at the next scheduled time. Manually initiated replication operations must be explicitly restarted. For example, consider a volume in an SVM on cluster A that is being replicated with SnapMirror to a cluster outside the MetroCluster configuration. If cluster A switches over to cluster B, the volume is then replicated from cluster B to the destination cluster. On switchback, the replication returns to its normal operation, from cluster A to the destination cluster.

Another case is when the MetroCluster volumes are a SnapMirror or SnapVault destination (for example, volumes from a cluster outside the MetroCluster configuration replicated to a destination volume on one of the MetroCluster clusters). In that case, the peering relationship must be verified and re-created if necessary after every switchover or switchback. Also, all such SnapMirror and SnapVault relationships must be manually re-created after every switchover or switchback. A SnapMirror or SnapVault rebaseline is not required. A WFA workflow, "Re-create SnapMirror and SnapVault protection after MetroCluster switchover and switchback," is available to automate the re-creation of such relationships.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2018 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.