



Technical FAQ

NetApp MetroCluster ONTAP 8.3.x and ONTAP 9

Eddie Berin, NetApp
Version 3.4
October 2016

Abstract

This document provides answers to frequently asked questions about NetApp® MetroCluster™ software in the NetApp ONTAP® 8.3.x and ONTAP 9 operating system.

Data Classification

NetApp Confidential - Limited Use

Version History

Version	Date	Document Version History
Version 1.1	October 2014	Initial release
Version 1.1.1	November 2014	Minor additions and typos corrected
Version 1.1.2	December 2014	Additions for FlexArray, SSD, disk assignment
Version 1.1.3	December 2014	SnapMirror and SnapVault clarification, added more info on CRS
Versions 1.1.4, 1.1.4a	January 2015	Clarification on 16Gb SFP support for Brocade switches Additional transition guidelines FC initiator speed same speed requirement
Version 1.1.5 Version 1.1.5a	February 2015	Tiebreaker availability Minor clarifications around SFPs and HBA selection Clarification on LUN requirements for FlexArray
Version 1.1.6 Versions 1.1.6a, b	February 2015 March 2015	More SFP and SyncMirror clarifications, minor corrections Minor corrections, clarification of controller upgrade (ARL) process
Versions 1.1.7, 1.1.7a, b, c	March 2015	FlexArray and V-Series terminology clarification, shelf ID guidelines Flash Pool correction, ISL compression not supported NVRAM replication in ARL process, OCUM polling Ordering less than 4 shelves per site Brocade switch licensing
Versions 1.1.8, 1.1.8a	April 2015	Adding storage, small tweaks to graphics Information about DR partner assignment, removed Quote Tool enforcement of 4 shelf per site minimum FlexArray LUN requirement, updated components list diagram
Version 1.9	May 2015	FlexArray LUN requirement, relationship of serial ID to selection of DR partner
Version 2.0	June 2015	Updated for Data ONTAP 8.3.1
Version 2.1	September 2015	Minor additions and corrections
Version 2.2	November 2015	Minor corrections
Version 2.3	November 2015	Updated for Data ONTAP 8.3.2 RC
Version 2.4	December 2015	Minor additions and clarifications
Version 2.5	February 2016	Updated for Data ONTAP 8.3.2 GA
Version 2.5.1	March 2016	Revised distance for Stretch MetroCluster for ATTO 7500N between clusters
Version 2.5.2	April 2016	Added FC best practice for AFF MetroCluster deployments. Fixed image for slot assignment when sharing FC-VI and FC ports on the FAS8020.
Version 3.0	August 2016	Updated for ONTAP 9

Version	Date	Document Version History
Version 3.4	10/11/2016	Various corrections

TABLE OF CONTENTS

Version History	2
1 Overview.....	10
1.1 MetroCluster Release Summary	10
What was released for MetroCluster in Data ONTAP 8.3.0?	10
What was released for MetroCluster in Data ONTAP 8.3.1?	10
What was released for MetroCluster in Data ONTAP 8.3.2?	10
What was released for MetroCluster in ONTAP 9?	10
Does the Data ONTAP 8.3.0 or ONTAP 9.0 release affect 7-Mode stretch MetroCluster or fabric MetroCluster?	11
Is MetroCluster a chargeable feature or a separate product?	11
Which licenses are required for MetroCluster?	11
What are the key benefits of MetroCluster over SMC and FMC?	11
How does MetroCluster 8.3.x compare with SMC and FMC?	11
Are all clustered Data ONTAP features supported in MetroCluster?	13
What are the maximum distances for stretch MetroCluster configurations?	13
2 Requirements.....	14
What are the basic MetroCluster configurations?	14
What is the maximum supported distance between sites in MetroCluster?	14
What is the minimum supported eight-node MetroCluster configuration?	15
What is the minimum supported four-node MetroCluster configuration?	17
What is the minimum supported two-node stretch MetroCluster configuration in Data ONTAP?	19
What is required for the two-node fabric MetroCluster configuration in clustered Data ONTAP?	20
What is required for the two-node configuration with FibreBridges?	21
What is required for the two-node configuration with optical SAS cables?	21
How many target FC ports are required when deploying AFF-series controllers in a two-node or four-node MetroCluster configuration?	21
What does the cabling and architecture look like in the four-node configuration?	22
What does the cabling and architecture look like in the eight-node configuration?	22
What two-node configurations are available in Data ONTAP 8.3.1 or later?	23
How does the MetroCluster FCIP configuration differ from the FC configuration?	25
What are the front-end network requirements for MetroCluster?	26
Do the IMT and Config Advisor support MetroCluster on clustered Data ONTAP 8.3.x/ ONTAP 9.0?	26
Does the NetApp AutoSupport® tool provide support for MetroCluster 8.3.x?	26
How can I size a MetroCluster 8.3.x configuration?	26
Can the ISL be shared between the two DR groups in an eight-node MetroCluster configuration?	26
2.1 Controllers.....	26

Which FAS controller models are supported?	26
What are the platform mixing rules for the 8-node MetroCluster configuration?	27
Does the 8-node MetroCluster configuration support SAN protocols?	27
Can I mix NAS and SAN in the 8-node MetroCluster configuration?	27
Are cluster interconnect switches required in each HA pair?	27
Which are the supported configurations of the FAS controllers?	27
Which cluster interconnect switches are supported for an 8 or 4-node MetroCluster configuration?	28
Which back-end storage is supported in a MetroCluster configuration with array LUNs?	28
If the storage is purely array LUNs, does MetroCluster support combining the front-end and back-end fabric switches?	28
How many target and initiator ports are required in a MetroCluster configuration with array LUNs?	28
Can controller types be mixed in a MetroCluster configuration?	29
Can MetroCluster nodes be mixed with nodes other than MetroCluster in the same cluster?	29
How can I verify that the correct SFP is used for the 16Gb FC-VI connection?	30
Are the port efficiency improvements for the FAS8020 controller supported on the AFF8020 controller?	30
2.2 Storage	30
Which storage shelves are supported?	30
What are the disk requirements for NetApp SyncMirror with MetroCluster?	30
What is the recommended way to assign the IDs to the shelves?	30
How many disks and shelves are required?	30
Can aggregates be configured as RAID DP or RAID 4?	32
What is the minimum number of disks that a data aggregate can contain?	32
Can the cluster contain any data before initializing MetroCluster?	32
Does MetroCluster support unmirrored aggregates?	32
Can mirrored aggregates be converted to an unmirrored aggregate?	32
Can unmirrored aggregates be converted to mirrored aggregates?	33
What is the minimum number of aggregates required in MetroCluster?	33
Why are two data aggregates on each cluster recommended to configure MetroCluster?	33
Does storage need to be added in fixed increments?	34
What is the maximum spindle count for MetroCluster?	34
What are the system limits for MetroCluster?	34
Do MetroCluster configurations with array LUNs require FAS shelves for the root aggregate?	34
How are LUNs assigned in a MetroCluster configuration with storage arrays?	35
Does MetroCluster support advanced disk partitioning for the root aggregate?	36
Does MetroCluster support SSD advanced disk partitioning for Flash Pool?	36
Does SMC support SSD partitioning for Flash Pool?	36
Does a MetroCluster All Flash FAS configuration require a PVR?	36
What is the maximum number of SSDs/HDDs in a MetroCluster configuration?	36
Are SSD-only aggregates supported in a MetroCluster configuration without requiring a PVR?	37
How much Flash Pool cache can be configured in a MetroCluster configuration?	37
2.3 FC Switches and ISLs (Fabric Configurations Only)	37
Which switches are supported for the storage fabric?	37

Are there any feature or functional differences between a MetroCluster configuration using FC ISL connectivity and FCIP ISL connectivity?	38
Which switches are supported for FCIP ISL connectivity?	38
Are any additional licenses required for the Cisco 9250i/Brocade 7840 FCIP switches?	38
Can I verify the FCIP ISL latency between the two MetroCluster clusters?	38
How do I create a quote for an FCIP MetroCluster deployment?	39
Can I get more information about the Cisco MDS 9250i FCIP switch?	41
Is compression supported on the Cisco MDS 9250i FCIP switch?	41
What are the steps to configure the Cisco MDS 9250i FCIP switch for MetroCluster?	41
How many ports are licensed when I purchase the Cisco MDS 9250i FCIP switch through NetApp?	41
Is it required to purchase the switches from NetApp?	41
Do switches ordered through NetApp include SFPs?	41
What is included with the Brocade switches that are ordered through NetApp?	41
How is licensing handled for the reuse of existing supported Brocade switches?	42
What is included with the Cisco switches that are ordered through NetApp?	42
Can MetroCluster 8.3.x share the same switches and ISLs with an FMC during transition?	42
Is the 8Gb Brocade 6510 supported?	42
What are the requirements for the ISLs between switches?	42
Is trunking required for Brocade switches?	43
How do I check the required Cisco and Brocade LW SFP and xWDM support?	43
What is the required latency for the ISLs?	44
Do all the ISLs in both fabrics have to be the same speed and length?	44
Can separate MetroCluster configurations share the same switches and ISLs?	44
What is the maximum distance that is supported between the two clusters?	44
What distances are supported using SW optics between the switch and the xWDM?	44
What distances are supported using LW optics between the switch and the xWDM?	45
Is traffic isolation used?	45
Is out-of-order delivery (OOD) supported on the ISL?	45
Is ISL encryption supported?	45
Is ISL compression supported?	45
Can the FC switch be shared with all four controllers in an eight-node MetroCluster configuration?	45
2.4 FibreBridges	45
Which configurations require FibreBridges?	45
Which FibreBridges are supported?	45
How many FibreBridges are required in a MetroCluster configuration?	46
How many shelves are supported in a FibreBridge-attached stack with MetroCluster?	46
Is SNMP supported for the FibreBridge?	47
What are the disk shelf limits in an All Flash FAS (AFF) configuration when using the ATTO 7500N? ...	47
What are some of the key advantages of the ATTO 7500N compared to the ATTO 6500N?	47
Since only two FC ports on the ATTO 7500N are used, can one ATTO 7500N be used per stack?	47
What is the factory default configuration of the disk shelves between the four SAS ports?	48
What is the factory default configuration of the ports on the ATTO 6500N and 7500N?	48

Can only one FC port be used for ATTO 6500N FibreBridge connectivity?	48
Can two FC ports be used for the ATTO 7500N FibreBridge connectivity?.....	48
Can I nondisruptively replace an ATTO 6500N with the ATTO 7500N?	48
What firmware version is supported on ATTO 7500N?	48
What should I be aware of when deploying the ATTO 7500N?	48
How do I configure the management port on the ATTO 7500N?.....	49
How do I upgrade the firmware on ATTO 7500N?.....	49
Is the MetroCluster stretch configuration supported with the ATTO 7500N?.....	49
What distances are supported using SW optics with the ATTO 7500N?	49
Can the ATTO 6500N and the ATTO 7500N coexist in the same MetroCluster disk stack?	49
Can the ATTO 6500N and the ATTO 7500N coexist in same MetroCluster configuration?	49
Is the four-port 8Gbps FC card (X1132A-EN-R6-C) supported with the ATTO 7500N?	49
What disk shelf models are supported with the ATTO 6500N and the ATTO 7500N?	49
Can disk shelves with unmirrored aggregates be connected directly to the FAS or AFF controller?	49
2.5 SAN Host Support.....	50
Which SAN OSs are supported?	50
What SAN host connectivity (front end) is required?	50
2.6 Configuring MetroCluster	50
Which components are required in a MetroCluster configuration?	50
3 Architecture.....	51
3.1 Replication	51
Which types of replication are used in MetroCluster?.....	51
Must all aggregates be mirrored in a MetroCluster configuration?	52
Is SyncMirror supported for configurations without MetroCluster?	52
Is synchronous or semi-synchronous SnapMirror available in Data ONTAP 8.3.x?	52
What is the read and write behavior with plexes in SyncMirror (including MetroCluster)?.....	52
What is the switchover granularity of MetroCluster?.....	52
Can data on the remote plex be accessed by the DR partner during normal operation?	53
What happens if all ISLs in a fabric MetroCluster are down?	53
What happens if the links fail in a two-node stretch MetroCluster configuration?.....	53
What happens if the cluster peering network experiences an interruption?.....	53
Does epsilon play a role in MetroCluster availability?.....	53
How is the MetroCluster DR partner assigned in the 8-node MetroCluster configuration?.....	53
Does eight-node MetroCluster support mixing FAS and AFF controllers?	54
How does MetroCluster replicate data in an eight-node configuration?.....	54
3.2 Nondisruptive Operations.....	54
Does MetroCluster support clustered Data ONTAP nondisruptive upgrades (NDU)?	54
Can MetroCluster switch over if the clusters are running different minor Data ONTAP versions?.....	54
Is aggregate relocation (ARL) for nondisruptive controller upgrade supported with MetroCluster?	54
Can a two-node configuration be converted to a four-node configuration?	54
Can a four-node configuration be converted to an eight-node configuration?	55

Is vol move supported with MetroCluster?	55
Is it possible to transfer data from a volume in a mirrored aggregate to an unmirrored aggregate?	55
Is vol rehost supported with MetroCluster?	55
Can I nondisruptively replace an ATTO 6500N with the ATTO 7500N?	55
Where can I find information about other maintenance operations with MetroCluster?	55
3.3 Transition from 7-Mode (Including Fabric and Stretch MetroCluster).....	55
How can 7-Mode (including FMC or SMC) customers transition to MetroCluster?	55
Can FMC switches and ISLs be shared while transitioning to MetroCluster 8.3.x or 9.0?	56
Can the 7-Mode Transition Tool (7MTT) be used to transition from 7-Mode to MetroCluster 8.3.x/9.0? ..	56
Is there a nondisruptive transition from 7-Mode (including FMC and SMC) to MetroCluster?	56
Can 7MTT be used to transition from single-node or HA pair 7-Mode systems to MetroCluster?	56
Can 7MTT transition 32-bit data from 7-Mode controllers to MetroCluster?	57
Can equipment that was previously running 7-Mode (FMC and SMC) be reused for MetroCluster 8.3.x/9.0?	57
3.4 Transition from Clustered Data ONTAP	57
How can existing clustered Data ONTAP customers transition to MetroCluster?	57
Is there a nondisruptive transition from clustered Data ONTAP to MetroCluster?	58
Can 7MTT be used to transition from clustered Data ONTAP to MetroCluster?	58
Can nodes that were previously running another Data ONTAP release be reused for MetroCluster 8.3.x?	58
4 Configuration	59
Does System Setup support MetroCluster?	59
How can MetroCluster be monitored?	59
Is there an easy way to configure the fabric switches?	59
How many ports for ISLs and stacks are included in the switch golden configuration files?	59
How is MetroCluster configured at initial setup?	59
How is the DR partner node chosen in an eight-node and four-node configuration?	60
What are the default passwords for hardware shipped from the factory?	60
What configurations are available for two-node stretch MetroCluster?	60
Can a MetroCluster configuration be returned to a configuration without MetroCluster?	60
4.1 Switchover and Switchback	61
How are switchover and switchback initiated?	61
When is automatic switchover initiated in the two-node, four-node, and eight-node MetroCluster configurations?	61
Does Automatic switchover require a Product Variance Request (PVR)?	61
What is the difference between negotiated switchover and switchover on disaster?	61
Is it possible to perform a switchover test without actually switching over any objects?	62
Is it possible to perform an automatic switchover on disaster?	62
5 Interoperability	63
What capabilities does OCUM provide for MetroCluster?	63
What events related to MetroCluster cause a health monitor alert?	63

Does System Manager support MetroCluster?	63
Is Tiebreaker available for MetroCluster?	63
Does switchover work if the cluster is not in quorum?	64
How do SnapMirror and SnapVault interoperate with MetroCluster?	64
Is SnapLock Enterprise supported in MetroCluster?	64
Is SnapLock Compliance supported in MetroCluster?	64
References.....	65
Contact Us	65

LIST OF TABLES

Table 1) MetroCluster 8.3.x and 9 features compared with Data ONTAP 7-mode SMC and FMC.	12
Table 2) Stretch MetroCluster maximum distance, 16Gbps.	13
Table 3) Maximum ISL distances between data centers in Data ONTAP 8.3.2 and ONTAP 9.0	15
Table 4) Disk shelf ID assignment.	30
Table 5) LUN assignment, four-node configuration.	35
Table 6) Sample valid Flash Pool configurations.	37
Table 7) Supported 8Gb Cisco SFPs and equivalent NetApp part numbers.	44
Table 8) Supported short-wave optics distances.	44
Table 9) Maximum supported SSD and HDD stack depths for the ATTO 7500N and 6500N.	47
Table 10) ATTO 6500N and 7500N factory default port configuration.	48

LIST OF FIGURES

Figure 1) FAS8020 configuration in clustered Data ONTAP 8.3.0.	18
Figure 2) FAS8020 configuration in clustered Data ONTAP 8.3.1.	18
Figure 3) FAS8020 configuration in clustered Data ONTAP 8.3.2.	18
Figure 4) Four-node MetroCluster configuration with FAS shelves.	22
Figure 5) Eight-node MetroCluster configuration.	22
Figure 6) Two-node fabric configuration.	23
Figure 7) Two-node FibreBridge stretch configuration.	24
Figure 8) Two-node optical SAS stretch configuration.	25
Figure 9) MetroCluster FCIP configuration in clustered Data ONTAP 8.3.2.	26
Figure 10) Example of switch port assignments in a MetroCluster configuration with array LUNs.	29
Figure 11) Pool-to-shelf assignment, four shelves per site.	31
Figure 12) Layout of pool-to-shelf assignment, two shelves per site.	32
Figure 13) Disk drive considerations.	34
Figure 14) Components for a four-node MetroCluster configuration.	50
Figure 15) Example HA and DR partners, four-node MetroCluster configuration.	51

1 Overview

This document provides answers to common questions about MetroCluster in Data ONTAP 8.3.x and ONTAP 9.

Unless otherwise stated, all references to “MetroCluster” in this document mean MetroCluster in Data ONTAP 8.3.x and ONTAP 9, which are releases for clustered Data ONTAP only. References to MetroCluster products for Data ONTAP operating in 7-Mode, including fabric MetroCluster and stretch MetroCluster, are specifically called out.

For more information about MetroCluster, see [TR-4375: MetroCluster for Data ONTAP 8.3 Overview and Best Practices](#).

1.1 MetroCluster Release Summary

What was released for MetroCluster in Data ONTAP 8.3.0?

Answer: Data ONTAP 8.3.0 includes MetroCluster in a four-node fabric configuration. Data ONTAP 8.3.0 is a release only for clustered Data ONTAP; therefore, MetroCluster in this version also supports only clustered Data ONTAP.

What was released for MetroCluster in Data ONTAP 8.3.1?

Answer: The major addition for the Data ONTAP 8.3.1 release is support of a two-node architecture (in both fabric and stretch configurations) for MetroCluster.

What was released for MetroCluster in Data ONTAP 8.3.2?

Answer: The Data ONTAP 8.3.2 release includes the following features:

- Increased Fibre channel (FC) ISL distance between the MetroCluster clusters
 - Up to 185 miles (300km) using the Brocade 6505 or 6510 switch
 - 185 miles (300km) is not supported using the Cisco FC switches
- Fibre Channel over IP (FCIP) Inter-Switch Link (ISL) connectivity between the MetroCluster clusters:
 - Uses the Cisco 9250i FCIP fabric switch
 - Supports up to 120 miles (200 kilometers) between the MetroCluster clusters
- Starting in Data ONTAP 8.3.2 GA, the ATTO 7500N FibreBridge is supported for all MetroCluster fabric configurations.
 - Supports dual-power supplies
 - Provides increased performance (compared to the ATTO 6500N)
- New functionality for shared FC-VI and FC connectivity on the same initiator card (X1132A-EN-R6-C):
 - NetApp FAS8020 and AFF8020 only

What was released for MetroCluster in ONTAP 9?

Answer: The ONTAP 9 release includes the following features:

- Support for eight-node fabric configuration
 - NAS only
- Support for unmirrored aggregates on the MetroCluster cluster
- Upgrade from two-node to four-node MetroCluster configurations
 - Two-node fabric configuration to four-node fabric configuration only
 - Upgrading from a stretch 2-node configuration to a 4-node requires no data migration. Upgrade from four-node to eight-node MetroCluster configurations

- Four-node fabric configuration to eight-node fabric configuration only
- Automatic unplanned switchover for four-node/eight-node architectures
- SnapLock Enterprise support
- SnapLock Compliance is not supported

Does the Data ONTAP 8.3.0 or ONTAP 9.0 release affect 7-Mode stretch MetroCluster or fabric MetroCluster?

Answer: No. Stretch MetroCluster (SMC) and fabric MetroCluster (FMC) remain at their current levels in Data ONTAP 8.2.x in 7-Mode and will continue to be supported for that version.

Is MetroCluster a chargeable feature or a separate product?

Answer: No. MetroCluster is included in Data ONTAP 8.3.x and ONTAP 9.0, and there are no separate chargeable features. MetroCluster does require additional hardware compared with a non MetroCluster configuration, as detailed in [section 2](#).

Which licenses are required for MetroCluster?

Answer: 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 9.0 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 [section 2.3](#).

What are the key benefits of MetroCluster over SMC and FMC?

Answer: MetroCluster provides all the clustered Data ONTAP benefits (such as nondisruptive operations [NDO]) and extends them with site-level recovery. The NDO capabilities of MetroCluster (including for planned and unplanned events, capacity balancing, and tech refresh) go well beyond those of a 7-Mode MetroCluster solution.

The MetroCluster eight-node and four-node configurations provide both local failover (HA) capability and sitewide switchover. Because SMC and FMC use a single-node cluster at each site, they cannot provide local failover. The MetroCluster eight-node and four-node configuration also provide greater performance and scalability than SMC and FMC, because it is a configuration of HA pairs in two clusters. Customers who consolidate SMC or FMC environments can leverage cost savings because the four and eight MetroCluster nodes share the same switch and ISLs. Configuration and management are also simplified because of the larger node configuration, compared with multiple SMCs or FMCs in 7-Mode.

MetroCluster also integrates with management tools such as NetApp OnCommand® Unified Manager (OCUM) and OnCommand Performance Manager (OPM) for monitoring and alerts. Unmirrored aggregates are supported by CLI at this time in ONTAP 9.0.

How does MetroCluster 8.3.x compare with SMC and FMC?

Answer: Table 1 compares MetroCluster 8.3.x and 9.0 with SMC and FMC.

Table 1) MetroCluster 8.3.x and 9 features compared with Data ONTAP 7-mode SMC and FMC.

Feature	MetroCluster 8.3.x	Data ONTAP 7-mode SMC and FMC
Configurations	8.3.0 four-node fabric configuration only. ATTO FibreBridges are required except for array LUNs with NetApp FlexArray® software or V-Series. FC switches and ISLs required. 8.3.1 adds two-node configuration in three options: fabric, stretch with optical SAS, and stretch with FibreBridges. ONTAP 9 adds support for eight-node fabric configurations. ATTO FibreBridges are required except for array LUNs with NetApp FlexArray® software or V-Series. FC switches and ISLs required.	Fabric and stretch configurations. Stretch optical cables are supported with SMC.
Local HA + remote switchover	Eight-node and four-node configuration: yes; most events need only local HA. Two-node configuration requires remote switchover; there is no local HA. Failure or panic of one node (if the storage is still available) triggers automatic switchover.	No, remote switchover is required for all events.
Unmirrored aggregate support	8.3.x No, all aggregates, including root aggregates, must be mirrored. Starting in ONTAP 9, data aggregates can be unmirrored or mirrored. At least one mirrored data aggregate per cluster must exist at each site. Prior to ONTAP 9, all aggregates, including root aggregates, must be mirrored.	Yes, unmirrored aggregates are supported.
Management tool support	Yes, including OCUM Config Advisor, OPM, and health monitors. Tiebreaker is available for monitoring; a Product Variance Request (PVR is required for automated switchover).	Little to none. Tiebreaker is available as PVR only.
Number of nodes	Eight nodes: four-node clusters at each site. Four nodes: two-node cluster at each site. Two nodes: single-node cluster at each site.	HA pair: one controller at each site.

Feature	MetroCluster 8.3.x	Data ONTAP 7-mode SMC and FMC
ISL encryption support	Yes, for the fabric configurations with Brocade 6510.	Yes, with Brocade 6510.

Are all clustered Data ONTAP features supported in MetroCluster?

Answer: MetroCluster 8.3.x and 9.0 supports all SAN and NAS protocol versions in clustered Data ONTAP for two- and four-node configurations. MetroCluster 9.0 supports eight-node NAS configurations only - but not eight node SAN configurations at this time. For information about management application support, see [section 5](#). MetroCluster supports all clustered Data ONTAP features and configurations with the following exceptions:

- NetApp Infinite Volume
- NetApp Storage Encryption (NSE) drive encryption

Note: Data ONTAP systems can attach to NetApp E-Series storage arrays with FDE drives. The root aggregate can reside on unencrypted drives. For the list of systems that support this configuration, see the Interoperability Matrix Tool (IMT).

- Disk partitioning on the root aggregate (FAS2x00 is not supported with MetroCluster)
- SSD partitioning for NetApp Flash Pool™ intelligent caching (ADP)

Starting in clustered Data ONTAP 8.3.1, NetApp All Flash FAS is supported without a PVR. For versions earlier than clustered Data ONTAP 8.3.1, a PVR is required to support All Flash FAS. ONTAP 9.0 does not require a PVR to support ALL Flash FAS. In 8.3.2, Individual SSD-only aggregates do not require a PVR, as long as the node also contains aggregates with an HDD.

Switchover is disruptive to all SMB applications and clients, including SMB 3.0. This applies to all MetroCluster architectures including 8.3.1 onwards and ONTAP 9.0.

What are the maximum distances for stretch MetroCluster configurations?

Answer: See Table 2 for the maximum distances for stretch MetroCluster configurations.

Table 2) Stretch MetroCluster maximum distance, 16Gbps.

Speed (Gbps)	Maximum Distance (m)			
	16Gbps SW SFP			16Gbps LW SFP
	OM2	OM3	OM3+/OM4	Single-Mode (SM) Fiber ¹
2	n/a	n/a	n/a	n/a
4	150	270	270	500
8	50	150	170	500
16	35	100	125	500

¹ Single-mode (SM) LW SFPs are only supported for FC-VI connectivity.

2 Requirements

What are the basic MetroCluster configurations?

Answer: All MetroCluster configurations consist of an independent cluster at each site. In Data ONTAP 8.3.0, the cluster consists of exactly two nodes per site or one HA pair: a total of four nodes. Starting in 8.3.1, the two-node configuration is also available, consisting of a single-node cluster per site. Starting in ONTAP 9, the eight-node configuration is available, comprising of a four-node cluster at each site.

FAS (including FlexArray) and V-Series controllers are supported; all the systems in a MetroCluster configuration must be exactly the same model. Specific supported items are documented in this section; however, always consult the IMT as the authoritative source for supported configurations. All hardware must be ordered and purchased from NetApp. The configuration is active-active so that each cluster can serve its own separate workloads while providing disaster recovery (DR) protection for the other.

In clustered Data ONTAP 8.3.0, a MetroCluster configuration is a fabric-attached configuration that consists of a clustered Data ONTAP two-node cluster at each of the two sites. The four nodes are collectively known as a *DR group*. Supported, dedicated FC switches and FibreBridges are also required at each site.

Clustered Data ONTAP 8.3.0 supports one DR group. Both the minimum and maximum supported configuration size is a two-node cluster at each of the two sites.

Starting in clustered Data ONTAP 8.3.1, in addition to the existing four-node fabric configuration, two-node configurations are available. There are three options for the two-node configuration:

- **Two-node fabric configuration.** This option is similar to the four-node configuration except that there is only one node per site. Supported, dedicated FC switches and FibreBridges are required at each site.
- **Two-node stretched configuration with stretched optical cables.** Storage is attached to SAS ports on both nodes by using SAS cables. The FC-VI connections are direct between the controllers.
- **Two-node stretched configuration with FibreBridges.** Storage is attached to FC ports on both nodes with FibreBridges. The FC-VI connections are direct between the controllers.

Starting in clustered Data ONTAP 9, in addition to the existing two-node and four-node fabric configurations, eight-node configurations (NAS only) are available.

The MetroCluster eight-node configuration consists of an ONTAP four-node cluster at each of the two sites. These are deployed as two separate DR groups. Supported and dedicated FC switches and FibreBridges are also required at each site.

In ONTAP 9.0, the MetroCluster two-node and - remains the same as described in the clustered DATA ONTAP 8.3.0, 8.3.1, and 8.3.2 descriptions. For ONTAP 9.0 AUSO is supported for four-node and eight node configurations, this provides automatic switchover when an HA pair goes down.

What is the maximum supported distance between sites in MetroCluster?

Answer: For fabric configurations in clustered Data ONTAP 8.3.2 and newer, the maximum distance is 185 miles (300 kilometers) when using the Brocade 6505 or 6510 FC switch. Other FC switches support a maximum distance up to 120 miles (200 kilometers). However, certain ISL speed and switch combinations support shorter distances. For the supported distances, see Table 3. For stretch direct-attached two-node configurations and bridge-attached two-node configurations using the ATTO 6500N, the maximum distance is 1,640 feet (500 meters). For stretch bridge-attached two-node configurations with the ATTO 7500N, the maximum distance is 885 feet (270 meters). To verify supported distances based on specific cabling configurations, see the IMT.

Table 3) Maximum ISL distances between data centers in Data ONTAP 8.3.2 and ONTAP 9.0

Switch Type	Protocol	ISL Speed (Gbps)	Distance	Max # of ISLs
Cisco 9148S	FC	4	80km	4
	FC	8	40km	4
	FC	16	20km	4
Cisco 9250i	FCIP	10	200km	1
Brocade 7840	FCIP	10	200km	1
Brocade 6505	FC	4	300km	4
	FC	8	300km	4
	FC	16	300km (200km)	2 (4)
Brocade 6510	FC	4	300km	4
	FC	8	300km	4
	FC	16	300km (200km)	2 (4)

What is the minimum supported eight-node MetroCluster configuration?

Answer: The minimum supported eight-node MetroCluster configuration consists of:

- **Eight FAS, AFF or V-Series controllers, with four-node clusters (HA pair) per site.** For supported controllers, see [section 2.1](#). All four controllers at in a DR group must be exactly the same model and be either all FAS, AFF, all FAS with FlexArray, or all V-Series controllers. In an eight-node configuration, the controllers do not have to be the same across DR groups. Example – 1 DR group can be AFF8080 while the other is FAS8080.
- **Local HA between the nodes (HA pair) at each site.** Local HA is used for NVRAM mirroring to the HA partner and enables local HA failover.
- **One dual-port 16Gb FC-VI card (X1928) per controller.** This connection provides NVRAM mirroring to the DR partner node. The 8Gb card (X1927) is not supported. The X1928 must be configured with 16Gb SFPs (the same as used for the 16Gb HBA X1143). The SFP part number is X6602A. If the cards are ordered as add-ons, the correct 16Gb SFP must be specified. See also the question [How can I verify that the correct SFP is used for the 16Gb FC-VI connection?](#) in [section 2.1](#).
- **Four FC initiator ports (8Gb or 16Gb) per controller, consisting of two ports per fabric.** The ports can use the onboard/UTA ports or FC HBA ports, or a combination of both. For supported HBAs, see the IMT:
 - Only 8Gb or 16Gb ports are supported; the onboard ports on FAS3200 are 4Gb and are not supported.
 - Two ports on each controller must connect to each switch at the local site. The two ports connected to each switch should belong to different ASICs on the HBA or to different port pairs for the onboard ports.
 - All four ports must operate at the same speed. Therefore, if a mix of 8Gb and 16Gb ports are used, all ports must run at the highest common speed (8Gb in this case).
 - Only four initiator ports are required and supported for FAS disk shelves.
 - Additional ports are required when mixing FAS and array LUN disk shelves.

- The Quote Tool selects X1143 as the default card for FC initiators with FAS8040 and higher. If you want to use the onboard UTA ports or 8Gb ports (X1132), you must select X1132 under the Adapter section.

Starting in Data ONTAP 8.3.1, **only** the FAS8020 can optionally be configured with two FC initiator ports per controller. This system has only one connection from each controller to each switch and can be applied to both 2-node, 4-node, and 8-node fabric MetroCluster configurations. All other platforms require four FC initiator ports. Figures 1 and 2 show the configuration for the FAS8020 in clustered Data ONTAP 8.3.1.

Starting in Data ONTAP 8.3.2, **only** the FAS8020 and AFF8020 can optionally be configured to use the X1132A-EN-R6-C initiator card for both FC-VI and FC connectivity. This feature allows the MetroCluster infrastructure to connect by using a single initiator card. The use of this configuration provides a free PCIe slot on the controller, which can be used for NetApp Flash Cache™ caching or additional client ports. Figure 3 shows the port configuration for the FAS8020 in clustered Data ONTAP 8.3.2.

Note: The X1132A-EN-R6-C initiator card for both FC-VI and FC connectivity must be installed in slot 1 of the FAS8020 controller.

- **Two or four cluster interconnect 10GbE ports per controller.** Use the cluster interconnect ports as NetApp recommends in the [Hardware Universe](#). Certain controller models default to four cluster interconnects but can optionally support two cluster interconnects. For information about the number of cluster interconnect ports required and recommended, see the SE technical presentation on the specific controller model that you have chosen.
- **Four cluster interconnect switches:** Four 10GbE cluster interconnect switches, two switches per site, are required.
- **Cluster management switches (customer supplied or ordered through NetApp).** The requirement is the same as for clustered Data ONTAP.
- **Two Ethernet ports (1Gb or 10Gb) per controller for the cluster peering network.** This network is used to replicate the cluster configuration between the sites by using intercluster logical interfaces (LIFs). The cluster peering network switches are customer supplied. The intercluster LIFs can be on ports, VLANs, or ifgroups as per the [Clustered Data ONTAP Network Management Guide](#).
- **Four FC switches, two switches per site.** The switches must be dedicated to MetroCluster. The same switches can be used to connect to NetApp disks and array LUNs. For supported switches, see [section 2.3](#).
- **Dedicated ISLs between the fabric switches.** The interswitch links must be dark fiber or xWDM as supported by the switch vendor. For more information, see [section 2.3](#).

Note: In the Data ONTAP 8.3.2 release, only one FCIP Ethernet port can be used for the ISL connection between the two FCIP switches.

- **Four FibreBridges, two per site.** Each storage stack requires two FibreBridges, and a minimum of one storage stack is required per site. Each additional storage stack requires two FibreBridges. Bridges are not required for MetroCluster configurations with array LUNs because the storage connects directly to the switches.
- **Minimum storage recommended for production:** NetApp recommends eight SAS shelves or 192 disks **per site** (sixteen shelves total) to allow per-shelf disk assignment, or equivalent storage array disks (array LUNs). For supported storage, see [section 2.2](#).

Minimum supported storage for FAS: In an eight-node configuration, four shelves or 96 disks **per site** (eight shelves total) are the minimum supported storage.

Note: 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 four shelves, or two shelf per HA pair, with 96 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.

Note: *With this configuration, shelves are shared between nodes. Therefore, disk assignment must be managed manually. This configuration may experience low capacity and performance due to limited disks. NetApp therefore does not recommend this configuration for production use without an understanding of the limitations.*

What is the minimum supported four-node MetroCluster configuration?

Answer: The minimum supported four-node MetroCluster configuration consists of:

- **Four FAS or V-Series controllers, with a two-node cluster (HA pair) per site.** For supported controllers, see [section 2.1](#). All four controllers in the DR group must be exactly the same model and be either all FAS, all FAS with FlexArray, or all V-Series controllers. (You cannot mix and match FAS, FlexArray, and V-Series between the two sites.)
- **Local HA between the nodes (HA pair) at each site.** Local HA is used for NVRAM mirroring to the HA partner and enables local HA failover.
- **One dual-port 16Gb FC-VI card (X1928) per controller.** This connection provides NVRAM mirroring to the DR partner node. The 8Gb card (X1927) is not supported. The X1928 must be configured with 16Gb SFPs (the same as used for the 16Gb HBA X1143). The SFP part number is X6602A. If the cards are ordered as add-ons, the correct 16Gb SFP must be specified. See also the question [How can I verify that the correct SFP is used for the 16Gb FC-VI connection?](#) in [section 2.1](#).
- **Four FC initiator ports (8Gb or 16Gb) per controller, consisting of two ports per fabric.** The ports can use the onboard/UTA ports or FC HBA ports, or a combination of both. For supported HBAs, see the IMT:
 - Only 8Gb or 16Gb ports are supported; the onboard ports on FAS3200 are 4Gb and are not supported.
 - Two ports on each controller must connect to each switch at the local site. The two ports connected to each switch should belong to different ASICs on the HBA or to different port pairs for the onboard ports.
 - All four ports must operate at the same speed. Therefore, if a mix of 8Gb and 16Gb ports are used, all ports must run at the highest common speed (8Gb in this case).
 - Only four initiator ports are required and supported for FAS disk shelves.
 - Additional ports are required when mixing FAS and array LUN disk shelves.
 - The Quote Tool selects X1143 as the default card for FC initiators with FAS8040 and higher. If you want to use the onboard UTA ports or 8Gb ports (X1132), you must select X1132 under the Adapter section.

Starting in Data ONTAP 8.3.1, **only** the FAS8020 can optionally be configured with two FC initiator ports per controller. This system has only one connection from each controller to each switch and can be applied to both 2-node and 4-node fabric MetroCluster configurations. All other platforms require four FC initiator ports.

Figure 1) FAS8020 configuration in clustered Data ONTAP 8.3.0.






Figure 2) FAS8020 configuration in clustered Data ONTAP 8.3.1.



Starting in Data ONTAP 8.3.2, **only** the FAS8020 and AFF8020 can optionally be configured to use the X1132A-EN-R6-C initiator card for both FC-VI and FC connectivity. This feature allows the MetroCluster infrastructure to connect by using a single initiator card. The use of this configuration provides a free PCIe slot on the controller, which can be used for NetApp Flash Cache™ caching or additional client ports. Figure 3 shows the port configuration for the FAS8020 in clustered Data ONTAP 8.3.2.

Figure 3) FAS8020 configuration in clustered Data ONTAP 8.3.2.



Note: The X1132A-EN-R6-C initiator card for both FC-VI and FC connectivity must be installed in slot 1 of the FAS8020 controller.

- **Two or four cluster interconnect 10GbE ports per controller.** Use the cluster interconnect ports as NetApp recommends in the [Hardware Universe](#). Certain controller models default to four cluster interconnects but can optionally support two cluster interconnects. For information about the number of cluster interconnect ports required and recommended, see the SE technical presentation on the specific controller model that you have chosen.

- **Optional:** An option is to use four 10GbE cluster interconnect switches, two switches per site. If a two-node switchless cluster is configured, these switches are not required.
- **Cluster management switches (customer supplied or ordered through NetApp).** The requirement is the same as for clustered Data ONTAP.
- **Two Ethernet ports (1Gb or 10Gb) per controller for the cluster peering network.** This network is used to replicate the cluster configuration between the sites by using intercluster logical interfaces (LIFs). The cluster peering network switches are customer supplied. The intercluster LIFs can be on ports, VLANs, or ifgroups as per the [Clustered Data ONTAP Network Management Guide](#).
- **Four FC switches, two switches per site.** The switches must be dedicated to MetroCluster. The same switches can be used to connect to NetApp disks and array LUNs. For supported switches, see [section 2.3](#).
- **Dedicated ISLs between the fabric switches.** The interswitch links must be dark fiber or xWDM as supported by the switch vendor. For more information, see [section 2.3](#).

Note: In the Data ONTAP 8.3.2 release, only one FCIP Ethernet port can be used for the ISL connection between the two FCIP switches.

- **Four FibreBridges, two per site.** Each storage stack requires two FibreBridges, and a minimum of one storage stack is required per site. Each additional storage stack requires two FibreBridges. Bridges are not required for MetroCluster configurations with array LUNs because the storage connects directly to the switches.
- **Minimum storage recommended for production:** NetApp recommends four SAS shelves or 96 disks **per site** (eight shelves total) to allow per-shelf disk assignment, or equivalent storage array disks (array LUNs). For supported storage, see [section 2.2](#).

Minimum supported storage for FAS: In a four-node configuration, two shelves or 48 disks **per site** (four shelves total) are the minimum supported storage. In a two-node configuration, one shelf or 24 disks **per site** (two shelves total) are the minimum supported storage. With this configuration, shelves are shared between nodes. Therefore, disk assignment must be managed manually. This configuration may experience low capacity and performance due to limited disks. NetApp therefore does not recommend this configuration for production use without an understanding of the limitations.

Note: 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.

What is the minimum supported two-node stretch MetroCluster configuration in Data ONTAP?

Answer: All two-node configurations must have the following items (depending on the specific topology that you choose, additional hardware is required as described in the following questions):

- **Two FAS or V-Series controllers, with a single-node cluster per site.** For supported controllers, see [section 2.1](#). Both controllers must be exactly the same model and be either AFF, all FAS, all FAS with FlexArray, or all V-Series controllers (you cannot mix and match FAS, FlexArray, and V-Series between the two sites). Because these are single-node clusters, no cluster interconnect ports or switches are required.
- **One dual-port 16Gb FC-VI card (X1928) per controller.** This connection provides NVRAM mirroring to the DR partner node. The 8Gb card (X1927) is not supported. The X1928 must be configured with 16Gb SFPs (the same as used for the 16Gb HBA X1143). The SFP part number is X6602A. If the cards are ordered as add-ons, the correct 16Gb SFP must be specified. See also question [How can I verify that the correct SFP is used for the 16Gb FC-VI connection?](#) in [section 2.1](#).

- **Cluster management switches (customer supplied or ordered through NetApp).** Same requirement as for clustered Data ONTAP.
- **Two Ethernet ports (1Gb or 10Gb) per controller for the cluster peering network.** This network is used to replicate the cluster configuration between the sites, using intercluster LIFs. The cluster peering network switches are customer supplied. The intercluster LIFs can be on ports, VLANs, or ifgroups as per the [Clustered Data ONTAP Network Management Guide](#).
- **Minimum storage recommended for production:** NetApp recommends two SAS shelves or 48 disks **per site** (four shelves total) to allow per-shelf disk assignment, or equivalent storage array disks (array LUNs). For supported storage, see [section 2.2](#).

Minimum supported storage for FAS: In a two-node configuration, one SAS shelf or 24 disks **per site** (two shelves total) are the minimum supported storage. With this configuration, shelves are shared between nodes. Therefore, disk assignment must be managed manually. This configuration may experience low capacity and performance due to limited disks. NetApp therefore does not recommend this configuration for production use without an understanding of the limitations.

Note: 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.

What is required for the two-node fabric MetroCluster configuration in clustered Data ONTAP?

Answer:

- **Four FC initiator ports (8Gb or 16Gb) per controller, consisting of two ports per fabric.** The ports can use the onboard/UTA ports or FC HBA ports, or a combination of both. For supported HBAs, see the IMT:
 - Only 8Gb or 16Gb ports are supported. The onboard ports on FAS3200 are 4Gb and are not supported.
 - Two ports on each controller must connect to each switch at the local site. The two ports connected to each switch should belong to different ASICs on the HBA or to different port pairs for the onboard ports.
 - All four ports must operate at the same speed. Therefore, if a mix of 8Gb and 16Gb ports are used, all ports must run at the highest common speed (8Gb in this case).
 - Only four initiator ports are required and supported for FAS disk shelves.
 - Additional ports are required when mixing FAS and array LUN disk shelves.
 - The Quote Tool selects X1143 as the default card for FC initiators with FAS8040 and higher. If you want to use the onboard UTA ports or 8Gb ports (X1132), you must select X1132 under the Adapter section.

Starting in Data ONTAP 8.3.1, **only** the FAS8020 can optionally be configured with two FC initiator ports per controller. This system has only one connection from each controller to each switch. All other platforms require four FC initiator ports.

Starting in Data ONTAP 8.3.2, **only** the FAS8020 and the AFF8020 can optionally be configured to use the X1132A-EN-R6-C initiator card for both FC-VI and FC connectivity. This feature allows the MetroCluster infrastructure to connect by using a single initiator card. The use of this configuration provides a free PCIe slot on the controller, which can be used for Flash Cache or additional client ports.

- **Four FC switches, two switches per site.** The switches must be dedicated to MetroCluster. The same switches can be used to connect to NetApp disks and array LUNs. For supported switches, see [section 2.3](#).

- **Dedicated ISLs between the fabric switches.** The ISLs must be dark fiber or xWDM as supported by the switch vendor. For more information, see [section 2.3](#).
- **Four FibreBridges, two per site.** Each storage stack requires two FibreBridges, and a minimum of one storage stack is required per site. Each additional storage stack requires two FibreBridges. Bridges are not required for MetroCluster configurations with array LUNs because the storage connects directly to the switches.

What is required for the two-node configuration with FibreBridges?

Answer:

- **Four FC initiator ports (8Gb or 16Gb) per controller, consisting of two ports per fabric.** The ports can use the onboard/UTA ports or FC HBA ports, or a combination of both. For supported HBAs, see the IMT:
 - Only 8Gb or 16Gb ports are supported. The onboard ports on FAS3200 are 4Gb and are not supported.
 - Two ports on each controller must connect to each switch at the local site. The two ports connected to each switch should belong to different ASICs on the HBA or to different port pairs for the onboard ports.
 - All four ports must operate at the same speed. Therefore, if a mix of 8Gb and 16Gb ports are used, all ports must run at the highest common speed (8Gb in this case).
 - Only four initiator ports are required and supported for FAS disk shelves.
 - Additional ports are required when mixing FAS and array LUN disk shelves.
 - The Quote Tool selects X1143 as the default card for FC initiators with FAS8040 and higher. If you want to use the onboard UTA ports or 8Gb ports (X1132), you must select X1132 under the Adapter section.
- **Four FibreBridges, two per site.** Each storage stack requires two FibreBridges, and a minimum of one storage stack is required per site. Each additional storage stack requires two FibreBridges. Bridges are not required for MetroCluster configurations with array LUNs because the storage connects directly to the switches.

What is required for the two-node configuration with optical SAS cables?

Answer:

- Four SAS ports per controller.
- Optical SAS or regular SAS cables. Regular SAS cables can be used for local storage connections; optical SAS should be used for remote storage connections.


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

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

What does the cabling and architecture look like in the four-node configuration?

Answer: Figure 4 shows a four-node MetroCluster configuration with NetApp FAS shelves. If additional storage is required, additional shelves and FibreBridges are also supported.


Figure 4) Four-node MetroCluster configuration with FAS shelves.



What does the cabling and architecture look like in the eight-node configuration?

Answer: Figure 5 shows an eight-node MetroCluster configuration with NetApp disk shelves. If more storage is required, additional shelves and FibreBridges are also supported.

Figure 5) Eight-node MetroCluster configuration.




What two-node configurations are available in Data ONTAP 8.3.1 or later?

Answer: There are three possibilities for a two-node configuration. As stated earlier, NetApp recommends two shelves as the minimum storage configuration per site. Figure 6 through Figure 8 show optional additional shelves.


- **Two-node fabric configuration.** This configuration is similar to the four-node architecture in Figure 4, except that there is only one node at each site, and there is no cluster interconnect.

Figure 6) Two-node fabric configuration.




- **Two-node FibreBridge stretch configuration.** Storage is attached by using FibreBridges directly to FC initiator ports on each node. The FC-VI connection is a direct node-to-node connection.

Figure 7) Two-node FibreBridge stretch configuration.



- **Two-node optical SAS stretch configuration.** Storage is attached by using SAS (normal or optical) cables to SAS ports on each node. Direct SAS connections or patch panels can be used.


Figure 8) Two-node optical SAS stretch configuration.



How does the MetroCluster FCIP configuration differ from the FC configuration?

Answer: The cabling configuration remains the same. The only difference is that the FC fibre switches are replaced with the FCIP switches and an IP ISL. Only the Cisco 9250i is supported for FCIP ISL connectivity with the Data ONTAP 8.3.2 release. The FCIP configuration requires a dedicated IP network for the IP ISL and may be directly attached or utilize intermediate switches. The latency across the IP ISL network must be less than 2ms. The MetroCluster FCIP configuration is shown in Figure 9.

Figure 9) MetroCluster FCIP configuration in clustered Data ONTAP 8.3.2 and ONTAP 9.0



What are the front-end network requirements for MetroCluster?

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

MetroCluster for clustered Data ONTAP does not implement the `/etc/mcsrc` file that is available in 7-Mode to provide a different configuration after switchover.

Do the IMT and Config Advisor support MetroCluster on clustered Data ONTAP 8.3.x/ ONTAP 9.0?

Answer: Yes. As always, the IMT is the authoritative reference for supported configurations. Config Advisor 4.x supports MetroCluster 8.3.x and has been enhanced with several configuration rules that are specific to MetroCluster. Config Advisor 4.x, together with the MetroCluster plug-in files and installation guide, are available on the Field Portal on the [Config Advisor](#) main page.

Does the NetApp AutoSupport® tool provide support for MetroCluster 8.3.x?

Answer: Yes, the NetApp My AutoSupport (ASUP™) tool supports MetroCluster 8.3.x.

How can I size a MetroCluster 8.3.x configuration?

Answer: SPM supports sizing for MetroCluster running on ONTAP 8.3.x and ONTAP 9. Individual workloads can be assigned to the cluster at either site, and a common FAS controller for all workloads is determined. SPM provides options to adjust the sizing based on the CPU impact in the event of an HA takeover. Sizing can also be adjusted to provide additional headroom to minimize the performance impact of a failover after takeover when one node is performing the work of all four original nodes.

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

Answer: Yes. The ISL can be shared with the two DR groups in the eight-node MetroCluster configuration. Note that ISL sharing is not supported on the two-node and four-node MetroCluster configurations.

2.1 Controllers

Which FAS controller models are supported?

Answer: **Always consult the IMT for supported controllers.** At the time of this writing, FAS3220, FAS3250, FAS6210, FAS6220, FAS6240, FAS6250, FAS6280, FAS6290, FAS8020, FAS8040, FAS8060, and FAS8080 EX controllers are being supported in 8.3.x and 9.0. Including the All-flash FAS versions of these controllers. The controllers supported are based on those currently available; FAS3240,

FAS3210, and FAS3270 are not supported. FAS2xxx controllers are also not supported; they lack sufficient ports and performance to run MetroCluster.

Note: All controllers across DR groups must be the same model.

V-Series and FAS controllers that run FlexArray virtualization software are also supported. Mixing FAS and V-Series 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 [Clustered Data ONTAP Storage Platform Mixing Rules](#).

Does the 8-node MetroCluster configuration support SAN protocols?

Answer: No. In the ONTAP 9.0 release only NAS is supported on the 8-node MetroCluster configurations. 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: No. In the ONTAP 9.0 release the 8-node MetroCluster configuration only supports NAS. The 2-node and 4-node MetroCluster configurations continue to support SAN configurations.

Are cluster interconnect switches required in each HA pair?

Answer: The cluster interconnect switches are optional in the four-node MetroCluster configurations. In the eight-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 are the supported configurations of the FAS controllers?

Answer:

All the Controller nodes in an individual four-node or two-node MetroCluster or a single DR group in an eight-node FMC have to be the same platform model

In four-node FMC both controllers in a CC (Ctrl, Ctrl) must belong to the same DR group

In a two-node FMC or SMC, only controllers in a CI (Ctrl, IOXM) and CB (Ctrl, Blank) configurations are supported. Controllers in a CC (Ctrl, Ctrl) are not supported for a two-node MetroCluster.

For eight-node FMC each of the two DR groups should have the same four controllers. But the controllers can be different between the two DR groups as long as all eight controllers are supported for MetroCluster. See IMT for controllers supported for MetroCluster.

Follow the general cDOT Storage Platform Mixing Rules to determine which two controller models can be in the four-node cDOT cluster at each site. See:

https://library.netapp.com/ecm/ecm_get_file/ECMP1644424

The following controllers are supported for SMC and FMC with the restrictions noted above:

Data ONTAP Release	8.3	8.3.1+	9.0
FAS2xxx	No	No	No

FAS3210, FAS3240, FAS3270	No	No	No
FAS60xx	No	No	No
FAS3220	CC, CI, CB	CC, CI, CB	CC, CI, CB
FAS3250	CI	CI	CI
FAS6210	CC, CB	CC, CB	CC, CB
FAS6240	CI	CI	CI
FAS6280	CI	CI	CI
FAS6220	CC, CB	CC, CB	CC, CB
FAS6250	CI	CI	CI
FAS6290	CI	CI	CI
FAS8020	CI	CI	CI
FAS8040	CC	CC	CC
FAS8060	CC	CC, CI	CC, CI
FAS8080	CI,	CI, CC	CI, CC

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 switchless or switched. Switched configurations can use any cluster interconnect switch supported in clustered Data ONTAP. The cluster interconnect switch must be the same at both sites. For the list of supported cluster interconnect switches, see the [Hardware Universe](#).

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 V-Series model, vendor array, disk type, and LUN size. Starting with Data ONTAP 8.3, FAS disk shelves are optional in a Data 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.

If the storage is purely array LUNs, does MetroCluster support combining the front-end and back-end fabric switches?


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

How many target and initiator ports are required in a MetroCluster configuration with array LUNs?

Answer: Each storage array requires exactly eight target ports, and each V-Series or FAS controller requires four FC initiator ports to attach the arrays, as shown in Figure 10. For each additional array, two more FC initiators ports would be required. Note that a maximum of 256 LUNs can be provisioned to each initiator port. If number of LUNs is greater than 256, an additional FC initiator and additional target ports are required. The two ports to each fabric should belong to different ASICs on the HBA (for

example, 0a and 0c to the first fabric and 0b and 0d to the second fabric). Note that the switch ports that are used in Figure 10 are a sample only, they do not represent the standard configuration.

Figure 10) Example of switch port assignments in a MetroCluster configuration with array LUNs.



If both FAS disk shelves (with FibreBridge) and array LUNs are used, separate initiator ports are required for each type of storage. This requirement means that two sets of four FC initiator ports are required on each FAS or V-Series controller, or eight ports total. For cabling and zoning information, see the product documentation. The same back-end FC switches are used for the two storage types, and separate ports are used for each storage connection type.

Can controller types be mixed in a MetroCluster configuration?

Answer: In the two-node and four-node configurations, all controllers in the MetroCluster DR group must be the identical FAS, AFF, or V-Series model (for example, four FAS8040s). In an eight-node configuration each node in a DR group must be identical. The two four node DR groups in an eight-node configuration can vary, e.g. DR group #1 can have AFF8080 while DR group #2 can have FAS8080. The NVRAM is mirrored both to the HA partner and cross-site to the DR partner node (see [section 3.1](#) for more information) and therefore the model strings must match exactly. AFF, FAS and V-Series nodes cannot be mixed in the DR group (for example, you cannot configure one cluster with a FAS3250 HA pair and the other cluster with a V3250 HA pair). 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 <https://kb.netapp.com/support/index?page=content&id=S%3A2022898>.

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 clustered Data ONTAP MetroCluster configurations.

2.2 Storage

Which storage shelves are supported?

Answer: **Always consult the IMT for supported controllers.** At the time of this writing, SAS shelves DS2246, DS4243, and DS4246 are being supported. DS4486 and DS14 shelves are not supported. SAS, SATA, and SSD drives are supported. All NetApp FAS disk shelves require a FibreBridge connection, except for the two-node MetroCluster. For information about the composition of the storage stacks, see [section 2.2](#).

Third-party arrays and E-Series storage arrays are supported with MetroCluster configurations, and FibreBridges are not used for array attachment. MetroCluster configurations with array LUNs must be identical at both sites.

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

Note: Mini-SAS cables do provide Glowworm support with some NetApp storage shelves. Check the IMT for updates to the supported disk shelves.

What are the disk requirements for NetApp SyncMirror with MetroCluster?

Answer: NetApp SyncMirror technology requires disks of the same architecture in both plexes; each aggregate must be mirrored to a disk of the same type. It is possible to disable RPM and size matching, but NetApp strongly recommends that you use an identical disk configuration at both sites.

What is the recommended way to assign the IDs to the shelves?

Answer: Shelf IDs should be unique within the DR group. Table 4 shows the numbering schema that NetApp recommends. For numbering in a two-node configuration, simply omit the A2 and B2 shelves and use the same A1 and B1 ranges. This approach allows a smooth conversion to a four-node or eight-node architecture.

Table 4) Disk shelf ID assignment.

Shelf ID Site A	Usage	Shelf ID Site B	Usage
Shelves 10 to 19	A1:Pool0	Shelves 20 to 29	A1:Pool1
Shelves 30 to 39	A2:Pool0	Shelves 40 to 49	A2:Pool1
Shelves 60 to 69	B1:Pool1	Shelves 50 to 59	B1:Pool0
Shelves 80 to 89	B2:Pool1	Shelves 70 to 79	B2:Pool0


How many disks and shelves are required?

Answer:

- Four-node minimum storage recommended for production:** NetApp recommends four SAS shelves or 96 disks, or the equivalent storage array disks (array LUNs), **per site** (eight shelves total). With four FAS shelves per site, each pool (local and remote for each node) has its own shelf, and the disks can be assigned on a per-shelf basis to each node. Future disk expansions should be planned so that the pool-shelf isolation is preserved.

Figure 11 shows the pool-to-shelf assignment with four shelves per site, using the shelf IDs that NetApp recommends. Separation of pools on a per-shelf basis is not enforced. Software disk ownership allows disks to be assigned any way across the shelves. In some configurations with SSD, sharing pools on the same shelf is the preferred solution for performance reasons, by positioning SSD-only shelves at the top or the bottom of the stack. Having multiple nodes owning disks on the same shelves does disable per-shelf disk assignment and also increases the scope of impact if a shelf fails.


Figure 11) Pool-to-shelf assignment, four shelves per site.



- Four-node minimum supported storage for NetApp shelves:** Two SAS shelves or 48 disks **per site** (four shelves total) are the minimum supported storage. In this configuration, each shelf contains disks from two different pools for different nodes. Because each shelf is shared between two nodes, manual disk ownership is required, rather than automatic per-shelf disk assignment. Any replaced failed disks must also be manually assigned to make them visible to the appropriate node. This requirement makes disk management more complex.

Figure 12 shows a layout of pool-to-shelf assignment with only two shelves per site. Note the limited number of disks available: 12 disks for each node. Each of the four MetroCluster nodes requires a dedicated root aggregate and a minimum of one mirrored data aggregate before running the `metrocluster configure` command. This layout has two or three drives used by the root aggregate (RAID 4 or NetApp RAID DP® technology) and one spare drive per pool. Therefore, either 7 + 2 or 6 + 2 drives remain for the data aggregate, which is very limited for capacity and performance. **NetApp does not recommend this configuration for production use.**

Figure 12) Layout of pool-to-shelf assignment, two shelves per site.



- **Two-node minimum storage recommended for production:** NetApp recommends two SAS shelves or 48 disks, or the equivalent storage array disks (array LUNs), **per site** (four shelves total). With two FAS shelves per site, each pool (local and remote for each node) has its own shelf, and the disks can be assigned on a per-shelf basis to each node. Future disk expansions should be planned so that the pool-shelf isolation is preserved.
- **Two-node minimum supported storage for NetApp shelves:** One SAS shelf or 24 disks **per site** (two shelves total) are the minimum supported storage. In this configuration, each shelf contains disks from two different pools for different nodes. Because each shelf is shared between two nodes, manual disk ownership is required, rather than automatic per-shelf disk assignment. Any replaced failed disks must also be manually assigned to make them visible to the appropriate node. This requirement makes disk management more complex.

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.

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 Data ONTAP configuration, having more drives in an aggregate is likely to give better performance, and aggregates with few drives do not perform optimally.

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.

Does MetroCluster support unmirrored aggregates?

Answer: Yes. Starting in ONTAP 9 MetroCluster supports both unmirrored and mirrored aggregates.

Can mirrored aggregates be converted to an unmirrored aggregate?

Answer: Yes. Starting in ONTAP 9 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.

What is the minimum number of aggregates required in MetroCluster?

Answer: Each node requires its own mirrored root aggregate.

In 8.3.0, the `metrocluster configure` command requires a minimum of two mirrored data aggregates on each cluster with at least 10GB free to initialize the environment. A best practice is to create an eligible data aggregate on each node, although it is possible to have the two data aggregates on just one node in each cluster. If additional empty data aggregates have been created, they must also be mirrored before running the `metrocluster configure` command, because MetroCluster does not support unmirrored aggregates. See the note in the following question for the PVR option to run MetroCluster 8.3.0 with only one data aggregate per cluster.

In clustered Data ONTAP 8.3.x, having two mirrored data aggregates continues to be the default and the configuration that NetApp recommends. For the two-node architecture, two data aggregates must be created on **each of the two single-node clusters** before running `metrocluster configure`. With an advanced mode option, any MetroCluster configuration can be set up with just one data aggregate per cluster.

Starting in ONTAP 9, MetroCluster supports both unmirrored and mirrored data aggregates. It should be noted that at least a single mirrored data aggregate per cluster is required to hold the SVM root and MDV volumes.

Why are two data aggregates on each cluster recommended to configure MetroCluster?

Answer: The `metrocluster configure` command creates two system metadata volumes (MDVs) on each cluster. These volumes are created at 10GB in size and are used to queue and log configuration replication services (CRS) transactions for configuration replication to the other cluster over the cluster peering network. The default MetroCluster configuration process requires that two data aggregates exist on each of the two clusters, with each aggregate preferably on a separate node, for resiliency reasons. If there are not two data aggregates on each of the two clusters with sufficient space to create the metadata volumes, the initial default configuration fails. These data aggregates do not need to be dedicated to the system volume; user volumes can be created within the aggregate. For more information about CRS, see [section 3.1](#).

Additional mirrored and unmirrored data aggregates can be created as required after MetroCluster is configured.

Note: In 8.3.0, a PVR is available for a post installation process to modify the configuration to contain only one data aggregate in each cluster. This single aggregate then contains both metadata volumes and all the data. This modification allows an active-passive configuration with only one larger data aggregate in each HA pair. Starting in 8.3.1, an advanced mode modifier on the `metrocluster configure` command allows initial setup with a single aggregate per cluster, and the PVR is no longer required. Although the best practice remains as two separate data aggregates, this configuration might be helpful in certain small configurations, particularly with FAS8020. For more information about layouts in small systems, see the presentation [Small Configurations for MetroCluster 8.3](#).


Does storage need to be added in fixed increments?

Answer: 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 is the maximum spindle count for MetroCluster?

The spindle count for all MetroCluster configurations is the same as for an HA pair of the same controller type, as documented in the [Hardware Universe](#). For example, consider a FAS8080 EX that supports 1,440 disks in an HA pair. In MetroCluster, up to 720 disks could be installed on each site, making a total of 1,440 disks for the DR group, allocated across all the local and remote plexes. For eight-node MetroCluster configurations, the above 1,440 disks are limits for one DR group. site – for the entire eight-node configuration, it is 2,880 disks. The disks can be assigned asymmetrically between the two clusters. For example, if one cluster needs 500 drives to support its workload, a maximum of 220 drives is available for the other cluster's workload. Figure 13 shows this disk assignment. Note that the **total number** of drives physically present at each site is the same because all data is mirrored. Figure 13 shows a four-node configuration. For a two-node configuration, the disk limit is the same: in this case, 720 drives across the two nodes.

Figure 13) Disk drive considerations.



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 Data 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.

How are LUNs assigned in a MetroCluster configuration with storage arrays?

Answer: The minimum required aggregate size for data and root aggregates is the same in MetroCluster as for clustered Data ONTAP. Therefore, sufficient LUNs must be created to satisfy this requirement, including NetApp Snapshot® reserve space. For more guidance, see the best practices in [TR-3461: V-Series Best Practices Guide](#).

Data and root aggregates are mirrored to both sites. At least one data aggregate per node is required; additional aggregates can be created as needed.

At least one spare per node should be available in each pool. The spare LUN is required in order to be able to write a core file in case of an ONTAP panic and block protocols are configured.

Pool0 and Pool1 for each node must contain the same number, type, and size of LUNs.

If only array LUNs are in the configuration (no FAS shelves), at installation, you must designate an additional LUN per node to be used for core dumps when they are generated. This LUN is not mirrored. The installation process prompts for the required size. For more information, see the [FlexArray Installation Guide](#).

Table 5 shows LUN assignment for a four-node configuration. For the two-node configuration, omit the rows for nodes A2 and B2.

Table 5) LUN assignment, four-node configuration.

Site A	Site B
Node A1 Pool0: • LUNs for root aggregate • LUNs for one or more data aggregates • Spare LUN (minimum 1) for core dumps (if there are no FAS shelves)	Node A1 Pool1: • LUNs for root aggregate • LUNs for one or more data aggregates • Spare LUN (minimum 1) for core dumps
Node A2 Pool0: • LUNs for root aggregate • LUNs for one or more data aggregates • Spare LUN (minimum 1) for core dumps (if there are no FAS shelves)	Node A2 Pool1: • LUNs for root aggregate • LUNs for one or more data aggregates • Spare LUN (minimum 1) for core dumps
Node B1 Pool1: • LUNs for root aggregate • LUNs for one or more data aggregates • Spare LUN (minimum 1) for core dumps	Node B1 Pool0: • LUNs for root aggregate • LUNs for one or more data aggregates

	• Spare LUN (minimum 1) for core dumps (if there are no FAS shelves)
Node B2 Pool1: • LUNs for root aggregate • LUNs for one or more data aggregates • Spare LUN (minimum 1) for core dumps	Node B2 Pool0: • LUNs for root aggregate • LUNs for one or more data aggregates • Spare LUN (minimum 1) for core dumps (if there are no FAS shelves)

Does MetroCluster support advanced disk partitioning for the root aggregate?

Answer: No. FAS2200 and FAS2500, which use root aggregate disk partitioning, are not supported on MetroCluster.

Does MetroCluster support SSD advanced disk partitioning for Flash Pool?

Answer: No. Although Flash Pool is supported on MetroCluster, SSD disk partitioning is not supported with SyncMirror, whether it is used in MetroCluster or as local SyncMirror.

Does SMC support SSD partitioning for Flash Pool?

Answer: No. 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 a MetroCluster All Flash FAS configuration require a PVR?

Answer: No. Starting in clustered Data ONTAP 8.3.1, a PVR is no longer required for All Flash FAS MetroCluster. All Flash FAS configurations other than MetroCluster now automatically use SSD partitioning for the root aggregate. However, All Flash FAS MetroCluster configurations do not support this partitioning and are installed with non-partitioned drives in the root aggregate. In 8.3.0, MetroCluster All Flash FAS configurations require a PVR.

What is the maximum number of SSDs/HDDs in a MetroCluster configuration?

Answer:

	MCC 2-node	MCC 4-node	MCC 8-node*
	Maximum HDDs	Maximum HDDs	Maximum HDDs
FAS8020	480	480	960*
FAS8040	720	720	1440*
FAS8060	1200	1200	2400*
FAS8080	1440	1440	2880*

*8-node may have different controllers per DR group, e.g. 4xFAS8080 and 4xFAS8020. The maximum for this controller mix would be 1440 (4xFAS8080) plus 480 (4xFAS8020) for total of 1920 drives.

Are SSD-only aggregates supported in a MetroCluster configuration without requiring a PVR?

No PVR is required

How much Flash Pool cache can be configured in a MetroCluster configuration?

Answer: Like any other mirrored aggregate, Flash Pool aggregates must be mirrored on both sites. In the event of a site switchover followed by HA takeover, one node services the workloads of all the nodes. Because of the considerations, Flash Pool cache sizes are halved in MetroCluster configurations compared with configurations other than MetroCluster. The Flash Pool size is also limited by the maximum number of SSDs supported. For the maximum HA pair Flash Pool sizes for each controller model, see the [Flash Pool Technical FAQ](#).

For example, consider a FAS8080 EX configuration: In Data ONTAP 8.3 without MetroCluster, the maximum cache size with Flash Pool is 144TiB for an HA pair. In MetroCluster with FAS8080 EX, the maximum cache size with Flash Pool is 72TiB across the four nodes in the DR group, with the equivalent 72TiB mirrored. The total amount of Flash Pool cache can be assigned in any way between the two clusters, as long as all of it is mirrored and the maximum cache size is not exceeded. The following Flash Pool configurations in Table 6 are valid, among many others.

Table 6) Sample valid Flash Pool configurations.

Sample Flash Pool Configuration 1: Even Distribution	
Site A	Site B
A local plexes: 36TiB	A remote plexes: 36TiB
B remote plexes: 36TiB	B local plexes: 36TiB
Sample Flash Pool Configuration 2: Different Amounts of Flash Pool Cache in the 2 Clusters	
Site A	Site B
A local plexes: 62TiB	A remote plexes: 62TiB
B remote plexes: 10TiB	B local plexes: 10TiB
Sample Flash Pool Configuration 3: Flash Pool Only on One Cluster	
Site A	Site B
A local plexes: 72TiB	A remote plexes: 72TiB

2.3 FC Switches and ISLs (Fabric Configurations Only)

Which 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. At the time of the writing, MetroCluster is supporting the following switches; however, you should always verify the current list on the IMT:

- Brocade 6505 and 6510 (16Gb switches, configured with 16Gb SFPs) are supported.
- Cisco MDS 9148 (8Gb switch, configured with 8Gb SFPs) is supported in Data ONTAP 8.3 only.
- Cisco MDS 9148S is supported in Data ONTAP 8.3.1 or later (16Gb switch configured with 16Gb SFPs).
- Cisco MDS 9250i is supported in Data ONTAP 8.3.2 or later for IP ISL connectivity.

Are 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.

Here 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 Data ONTAP 8.3.2. The Brocade 7840 is also supported.

Are any 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 pricelist but components/licenses need to be ordered individually.

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.

```
mcc4-siteB::> statistics show -object mcculp -instance Port_0_15_min -raw -node local

Object: mcculp
Instance: Port_0_15_min
Start-time: 6/4/2015 07:31:26
End-time: 6/4/2015 07:31:26
Node: mcc4-b2

  Counter                                     Value
  -----
  avg_latency                               55us
  instance_name                             Port_0_15_min
  max_latency                               1118us
  min_latency                               39us
  node_name                                 mcc4-b2
  process_name                             -
```

```
total_completions
7 entries were displayed.
```

145485

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

```
cisco9250i-fcs5# ips measure-rtt 192.168.0.10 interface IPStorage 1/1
Round trip time is 54 micro seconds (0.05 milli seconds)
```


For more information, see the [knowledge base article# 1015624](#).

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

How do I create a quote for an FCIP MetroCluster deployment?

Answer: The Brocade 7840 is in a different switch Quote Tool menu and not the same as the MetroCluster configuration of Cisco 9250i. This will be updated in a future release of Quote Tool. Follow these steps to create a quote for the Cisco FCIP switch in the Quote Tool:

1. Select **Configure a new cluster** and check the **Configure a new MetroCluster** option.



Quote: **13083510**

Configure a Cluster

What would you like to do?

- ☒ Configure a new cluster
- ☒ Configure a new MetroCluster
- ☐ Add to an existing cluster

2. Select the NetApp controller model and the number of nodes in the MetroCluster cluster.

Quote: 13083510

Configure a Cluster

- Cluster Advisor
- New HA Pairs
- * Add New MetroCluster System
 - Storage
 - Cabinets and Power
 - NICS and Adapters
- * Metro-Cluster Switch Services
 - Training
- * Cluster Software
- * Cluster Switches

Release & Product Info

How to Demos
Release Notes

System

Family: FAS
☐ SSD-Only

System: 8000 [cDOT Platform Mixing Rules](#)

Model: 8060

Controller Count: 4

Type: High Availability (Active Active) ?

System Options:

☐ FlexArray Virtualization (V-Series) ?
☐ IO Expansion Module

CNA OnBoard Port Pair 1 ?

Port: FC Target Type: Cable: Fiber Length(Meters): 5

CNA OnBoard Port Pair 2 ? ☒ Same as Port Pair #1

Port: FC Target Type: Cable: Fiber Length(Meters): 5 ☒ SFP +

☒ Metro Cluster

☒ Fabric with SAS to FC Bridge (Supports Short wave SFP only)
☒ Short wave SFP (for FC-VI)

Note: The Cisco MDS 9250i FCIP switch is only supported on fabric MetroCluster configurations.

3. Select the **Cisco** manufacturer option and the **Cisco MDS 9250i** model.

Quote: 13083510

Configure a Cluster

- Cluster Advisor
- New HA Pairs
- FAS 8060 System
 - Storage
 - Cabinets and Power
 - NICS and Adapters
 - Metro-Cluster Switch Services

Fabric Metro Cluster

 This section supports adding switches as part of the configured system. If a switch is not needed, skip this section.

Quantity: 4

MetroCluster Configuration

☐ Use Existing Switch

Switch

Manufacturer: Cisco Model: Cisco MDS9148S Please quote Cisco MetroCluster switches separately from Add Switch

4. Select the **Stand Alone Switch** option to choose the **MDS 9250i 12-port FCIP** type.

Switches

 This section supports adding switches as part of the configured system. If a switch is not needed, skip this section.

Switch Quantity: 1


Switch_Type: SAN or MetroCluster

Switch

☐ In-Cabinet Switch ?
☒ Stand Alone Switch

Manufacturer: Cisco Type: MDS9250i 12-Port FCIP w/3PSU for MC Ports: 12

5. Include the supported 10Gbps SFPs for FCIP configurations in the quote: Part number: X-SFP-10G-LR-R6.



Note: Refer to the [Cisco Storage Interoperability Matrix](#) for DWDM systems supported by the Cisco MDS 9250i FCIP ports.

Can I get more information about the Cisco MDS 9250i FCIP switch?

Answer: Yes. More information about the Cisco MDS 9250i FCIP switch is available in the [Cisco MDS 9250i Multiservice Fabric Switch Data Sheet Cisco MDS 9250i Multiservice Fabric Switch Data Sheet](#).

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

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

What are the steps to configure the Cisco MDS 9250i FCIP switch for MetroCluster?

Answer: The steps to configure the Cisco MDS 9250i FCIP switch are available in [knowledge base article# 1015624](#).

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.

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 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—including 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.

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
- Enterprise package feature for QoS capability; this feature must be selected through the Quote Tool

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

Can MetroCluster 8.3.x share the same switches and ISLs with an FMC during transition?

Answer: Yes. For information about Brocade 6510 only, see [section 3.3](#). Switch sharing is supported only during the transition from 7-Mode fabric MetroCluster to MetroCluster 8.3.x, and only one FMC and one MetroCluster 8.3.x can be attached at the same time for transition.

Is the 8Gb Brocade 6510 supported?

Answer: In FMC (7-Mode), the 6510 is sold with 8Gb SFPs; however, only 16Gb SFPs are supported with MetroCluster 8.3.x and ONTAP 9.0. All components in the 8.3.x and ONTAP 9.0 configurations (apart from ISLs) must connect only to 16Gb SFPs, including the FC initiators, FC-VI, and ATTO. Therefore, if Brocade 6510 switches that are currently used for FMC are reused in MetroCluster 8.3.x and ONTAP 9.0 or during transition, 16Gb SFPs are required for the MetroCluster 8.3.x and ONTAP 9.0 connections. For more information, see [section 3.3](#) on transition.

What are the requirements for the ISLs between switches?

Answer: ISLs must be dedicated to MetroCluster. Only ISLs over dark fiber (dedicated fiber cabling with “no protocol” on it) or xWDMs (as per the switch vendor’s interoperability matrix) are supported.

The minimum supported ISL speed is 4Gb; 8Gb and 16Gb are also supported. ISLs with a speed of 2Gb do not provide sufficient bandwidth for the number of nodes and could lead to performance issues.

The switch port and SFP must support the ISL speed and distance. SW SFPs that are supplied by NetApp support ISL distances of up to 15.5 miles (25 kilometers) with Brocade switches. For ISL distances and speeds that are supported with SFPs from NetApp, see the IMT.

For longer distances, LW SFPs or xWDM devices between the locations are required. For supported equipment details over longer distances, see the Brocade and Cisco interoperability matrixes. This issue is also addressed in the question “How do I check the required Cisco and Brocade LW SFP and xWDM support?” Any SFP that is supplied by a third party must be supported by the switch vendor and the supported switch firmware version.

One to four ISLs per fabric are supported. NetApp recommends two or more ISLs per fabric for all platforms except FAS8080 EX. For FAS8080 EX, NetApp recommends four ISLs per fabric for optimal performance.

In general, NetApp does not specify the requirements for the physical cable between sites. However, the link must have a dedicated lambda (wavelength), and the link must be supported by the switch vendor for the distance, switch type, and firmware version.

Here are a few additional considerations for ISL configuration:

- Metro-E (any non-FC native framing or signaling) is not supported. Metro-E is Ethernet framing or signaling either natively over a metro distance or through some TDM, MPLS, or xWDM technology.
- TDM (SONET/SDH) is supported as long as native FC frames are transmitted over the link without any conversion from FC to another protocol. TDMs might require changing signaling parameters on the switch, but those settings should come from the TDM or switch vendor.
- FCR (native FC routing) extensions are not supported for the back-end fabric.

Is trunking required for Brocade switches?

Answer: If two or more ISL connections are in the fabric, then trunking must be used when connecting the Brocade switches. 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).

How do I check the required Cisco and Brocade LW SFP and xWDM support?

Answer:

- **Brocade:** For distances of up to 15.5 miles (25 kilometers), SW SFPs that are supplied by NetApp can be used as shown in the IMT. Select the proposed switch, ISL distance, and speed. In the ISL Type column, the SFPs that NetApp supplies are listed if they are available.

For longer distances, either LW SFPs or xWDM equipment is required, and they should be selected from those listed in the Brocade interoperability matrix:

<http://www.brocade.com/downloads/documents/matrices/brocade-compatibility-matrix-fos-7x-mx.pdf>.

- **Cisco:** You can view the Cisco interoperability matrix for the MDS 9148, 9148S, and 9250i on the Cisco product support page. For example, the Cisco MDS 9148 fabric switch details can be found at <http://www.cisco.com/c/en/us/products/storage-networking/mds-9148-multilayer-fabric-switch/index.html>. The relevant row is for DS-C9148-K9. After you confirm that the SFPs DS-SFP-FC8G-XX are supported, see the table later in the document for 8Gb SFPs and verify the distance needed.

Table 7 shows the supported 8Gb Cisco SFPs and the equivalent NetApp part numbers at the time of the writing, but NetApp recommends that you always verify with the Cisco matrix. For example, at a 6-mile (10-kilometer) distance, NetApp part number X-DS-SFP-FC8G-LW-R6 is required. For longer

distances, xWDM equipment is generally required, and those items are also listed in a separate table in the Cisco matrix.

Table 7) Supported 8Gb Cisco SFPs and equivalent NetApp part numbers.

Cisco Part Number	NetApp Part Number	Wavelength	Maximum Distance
DS-SFP-FC8G-SW	X-DS-SFP-FC8G-SW-R5	850nm	-
DS-SFP-FC8G-LW	X-DS-SFP-FC8G-LW-R6	1310nm	6mi (10km)
DS-SFP-FC8G-ER	X-DS-SFP-FC8G-ER-R6	-	25mi (40km)

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 four). 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.

Can separate MetroCluster configurations share the same switches and ISLs?

Answer: No. Switch sharing is under consideration at this time, but currently each MetroCluster configuration requires separate switches and ISLs.

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

Answer: The maximum supported distance is 185 miles (300 kilometers), depending on the switch that you select.

What distances are supported using SW optics between the switch and the xWDM?

Answer: See Table 8 for the maximum distances supported using SW optics between the switch and the xWDM.

Table 8) Supported short-wave optics distances.

Speed (Gbps)	Maximum Distance (m)		
	16Gbps SW SFP		
	OM2	OM3	OM3+/OM4
2	n/a	n/a	n/a

Speed (Gbps)	Maximum Distance (m)		
	16Gbps SW SFP		
	OM2	OM3	OM3+/OM4
4	150	270	270
8	50	150	170
16	35	100	125

Note: Only multimode SFP+ cables are supported.

What distances are supported using LW optics between the switch and the xWDM?

Answer: MetroCluster supports the maximum distance defined by the SFPs specification. Check the IMT for more information.

Is traffic isolation used?

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

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

Answer: No. In-order delivery (IOD) is the default and the setting that NetApp recommends. OOD is not supported. IOD should be enabled, even if there is only one ISL per fabric.

Is ISL encryption supported?

Answer: Yes. MetroCluster supports ISL encryption with the Brocade 6510 switch only. No other encryption devices are supported anywhere else in the fabric, including in the ISLs across the WAN.

Is ISL compression supported?

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

Can the FC switch be shared with all four controllers in an eight-node MetroCluster configuration?

Answer: Yes. The FC switch can be shared with up to four controllers in the eight-node MetroCluster configuration. Be sure to verify you have licensed the required number of ports on the FC switch.

2.4 FibreBridges

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: Starting in Data ONTAP 8.3.2 GA, the ATTO 7500N FibreBridge is supported and recommended for attaching NetApp FAS disk shelves. The ATTO 7500N provides improved performance and lower latency than the ATTO 6500N FibreBridge and is ideal for AFF MetroCluster deployments. In addition, the ATTO 7500N has dual-power supplies for redundancy. A FibreBridge is not used for a MetroCluster configuration that uses solely array LUNs.

In Data ONTAP 8.3 and newer, the ATTO 6500N FibreBridge is available for attaching NetApp FAS shelves. A FibreBridge is not used for a MetroCluster configuration that uses solely array LUNs.

How many FibreBridges are required in a MetroCluster configuration?

Answer: Both the ATTO 6500N and the ATTO 7500N require two FibreBridges per disk stack. At least one stack is required at each site; therefore, a minimum of four FibreBridges is required. Each FibreBridge connects through one FC port to one switch and one SAS port to the SAS disk stack.

See Table 9 below for the SSD/HDD maximum stack depths.

How many shelves are supported in a FibreBridge-attached stack with MetroCluster?

Answer: Up to 10 shelves per stack are supported when only HDDs are used. Mixing disk types (for example, IOM3, IOM6) in the stack is supported as described the storage subsystem requirements.

If SSDs are present in the FibreBridge stack, the supported maximum stack depths are shown in Table 9. The total number of shelves includes SSD-only shelves, mixed SSD-HDD shelves, and HDD-only shelves. Up to 96 SSDs can be present in any single disk stack on the ATTO 7500N, and the SSDs can be distributed across any of the shelves. On the ATTO 6500N up to 48 SSDs can be present in a single disk stack. However, NetApp recommends that you not mix SSDs and HDDs in the same shelf, but instead configure SSD-only shelves.

Furthermore, for optimal performance, NetApp recommends that you place SSD-only shelves either in their own dedicated stack (no HDD shelves) or at the top or the bottom of a mixed SSD and HDD stack. In that way, the SSD shelves connect directly to the FibreBridge. In this instance, the SSD shelves likely contain pools from more than one node, so disk ownership must be manually assigned. For all-SSD aggregates, stacks should contain SSD only, with no more than two SSD shelves in the stack. Optimal performance is achieved with one SSD shelf per stack.

Table 9) Maximum supported SSD and HDD stack depths for the ATTO 7500N and 6500N.

Disk Configuration	Disk Type	Total shelves per FibreBridge pair					
		ATTO 6500N		ATTO 7500N			
All SSDs	SSD	2		4			
All HDDs	HDD	10		10			
Mixed SSD + HDDs	SSD	1	2	1	2	3	4
	HDD	6	2	9	8	7	6

Is SNMP supported for the FibreBridge?

Answer: No. SNMP is not supported on the ATTO 6500N or 7500N FibreBridge.

Is the ATTO 6500N still supported and available for MetroCluster deployments?

Answer: Yes. The ATTO 6500N is supported as an option for MetroCluster configurations and is available as an add-on.

What is the maximum number of disk shelves when using the ATTO 7500N?

Answer: The maximum number of disk shelves is the same for the ATTO 7500N and the ATTO 6500N. Both FibreBridges can support up to 10 shelves. In addition, the ATTO 7500N FibreBridges can support four stacks of disk shelves. The maximum number of shelves can vary depending on the disk configuration. See Table 9 for the maximum supported disk shelves.

Check the [Hardware Universe](#) and the [Clustered Data ONTAP Maximum Configuration Guides](#) for additional information.

What are the disk shelf limits in an All Flash FAS (AFF) configuration when using the ATTO 7500N?

Answer: In an AFF configuration a maximum of four SSD disk shelves is supported. This provides 96 SSDs disk drives on each pair of ATTO 7500N FibreBridges. See Table 9 for more information.

What are some of the key advantages of the ATTO 7500N compared to the ATTO 6500N?

Answer:

The key advantages of the ATTO 7500N include:

- Dual power supply
- Increased throughput performance
- Reduced latency performance
- 2 x 16G FC ports: both can be enabled
- 4 x 12G SAS ports: all four can be enabled
- Support for four SSD disk shelves with 96 SSDs connected to a pair of ATTO 7500Ns
- Support for four disk shelf stacks connected to a pair of ATTO 7500Ns

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.

What is the factory default configuration of the disk shelves between the four SAS ports?

Answer: By default, the disk shelves are equally distributed between the four SAS ports. In addition, both FC ports are connected. As a general best practice, SSD and HDD disk shelves should be connected to different SAS ports.

What is the factory default configuration of the ports on the ATTO 6500N and 7500N?

Answer: See Table 10 for the factory default port configuration.

Table 10) ATTO 6500N and 7500N factory default port configuration.

Port	ATTO 6500N	ATTO 7500N
Ethernet	Both enabled	Both enabled
FC	FC-1 and FC-2 enabled	FC-1 and FC-2 enabled
SAS	SAS-1 enabled, SAS-2 disabled	SAS-1/2/3/4 enabled

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 nondisruptively 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](#).

What firmware version is supported on ATTO 7500N?

Answer: Check the IMT to verify supported configurations, including the supported ATTO 7500N firmware versions.

What should I be aware of when deploying the ATTO 7500N?

Answer: When deploying the ATTO 7500N FibreBridges:

- Make sure to follow the cabling and zoning guidance in the [MetroCluster Installation and Configuration Guide](#)
- The management port is configured using the serial port.
Note: IPDHCP must be set to disabled.
- Health monitoring is not supported.
- SNMP is not supported.
- Read and understand the MetroCluster management, installation and service guides. For more information, view the [MetroCluster documentation for clustered Data ONTAP](#).

How do I configure the management port on the ATTO 7500N?

Answer: The management port is configured from the serial port. All Ethernet ports are enabled by default but at least one requires manual configuration for connectivity. The FibreBridge CLI commands are used to configure the management port. Configure IPDHCP to disabled. Use the following commands:

```
set IPDHCP [mp1 | mp2] disabled
set IPAddress [mp1 | mp2] <IP address>
set IPSubnetMask [mp1 | mp2] <IP subnet mask>
set IPGateway [mp1 | mp2] <IP gateway>
SaveConfiguration Restart
```

How do I upgrade the firmware on ATTO 7500N?

Answer: Instructions for upgrading the ATTO 7500N 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.

What distances are supported using SW optics with the ATTO 7500N?

Answer: See Table 8 for the maximum distances supported using SW optics with the ATTO 7500N.

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.

Is the four-port 8Gbps FC card (X1132A-EN-R6-C) supported with the ATTO 7500N?

Answer: Check the IMT to verify supported configurations, including support for FC cards.

What disk shelf models are supported with the ATTO 6500N and the ATTO 7500N?

Answer: NetApp DS2246, DS4246, and DS4243 disk shelves are supported. Check the IMT for updates to the supported disk shelves.

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.

2.5 SAN Host Support

Which SAN OSs are supported?

Answer: For SAN OS support, see the IMT. Configure the hosts according to standard NetApp best practices for clustered Data ONTAP.

What SAN host connectivity (front end) is required?

Answer: MetroCluster has full identity preservation on switchover and switchback. SVMs keep the same WWNN; FCP LIFs keep the same WWPNs; and iSCSI LIFs keep the same IP address and TPG tags. The only difference is that LUNs are presented on different ports on a different switch. Hosts do not need to rediscover the LUNs, which keep the same IDs. It is therefore required that SAN hosts have cross-site connectivity; that is, the SAN front-end fabric must span across both sites.

2.6 Configuring MetroCluster


Which components are required in a MetroCluster configuration?

Answer: Figure 14 shows the components for four-node MetroCluster. The two-node fabric configuration is similar, except that there is only one node per site.

When you configure SAS cables, the number of shelf-to-shelf cables required for a stack with X shelves is $(X-1) \times 2$. In the example given for the minimum recommended stack of four shelves, $(4-1) \times 2 = 6$ SAS cables per stack are required.

QuoteEdge is used for all MetroCluster configurations; for more guidance on using the Quote Tool for MetroCluster, see [MetroCluster 8.3.0: How to Quote](#) and the recorded deep-dive session about [quoting MetroCluster](#) on the Field Portal.

Figure 14) Components for a four-node MetroCluster configuration.




3 Architecture

3.1 Replication

In the eight-node and four-node configurations, every node in a MetroCluster configuration has a local HA partner in the same cluster and has a DR partner in the remote cluster.

As an example, consider a site A cluster with nodes A1 and A2 and a site B cluster with nodes B1 and B2. A1 and A2 are HA partners (Figure 15). Nodes A1 and B1 are DR partners. Nodes B1 and B2 are HA partners. Nodes A2 and B2 are DR partners.

Figure 15) Example HA and DR partners, four-node MetroCluster configuration.



For information about how the DR partner is chosen at installation, see [section 4](#). Manual intervention might be required to alter the configuration as shipped from manufacturing.

In the two-node configuration, there is only one node at each site and hence only DR NVRAM replication.

Which types of replication are used in MetroCluster?

Answer: Three types of replication are used in MetroCluster:

- **Storage.** All of a node's mirrored aggregates, including the root aggregate, are replicated to its DR partner through the ISLs. Each aggregate contains disks from two plexes. One plex is local, and the other plex is located on the DR partner. SyncMirror is used to keep the plexes in sync. The local node owns all disks in the aggregate during normal operation. Disk ownership is transferred to the DR partner during switchover.
- **NVRAM.** All NVRAM replication is over the FC-VI connection:
 - In the eight-node and four-node configurations, each node mirrors its NVRAM locally to its HA partner over the HA interconnect and remotely to its DR partner over FC-VI by using the ISLs.
 - In a two-node configuration, NVRAM is replicated between the two nodes. In the two-node fabric configuration, as in the four-node configuration, the FC-VI NVRAM replication traffic uses the ISLs. NVRAM is replicated between the two nodes. In the two-node stretch configurations, the replication is routed over the direct FC-VI connection, separate from the storage replication path.
- **Cluster configuration.** Each cluster replicates its configuration over the customer-supplied IP-based cluster peering network. The configuration mechanism is known as the *configuration replication service* (CRS). CRS sends over the configuration objects, such as SVMs, volumes, and LUNs, that must exist in the other cluster's replicated database (RDB) to allow a switchover. Throughput requirements for the cluster peering network are similar to those of SnapMirror. The cluster peering network should be configured with dual links for redundancy (two intercluster LIFs per node, each LIF to an independent network).

Must all aggregates be mirrored in a MetroCluster configuration?

Answer: Starting in ONTAP 9, all MetroCluster configurations support both mirrored and unmirrored aggregates. Only data contained in mirrored aggregates is replicated to the partner DR site. Root aggregates must be mirrored with SyncMirror.

Prior to ONTAP 9, all aggregates, including the root aggregate, must be mirrored with SyncMirror.

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.

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.

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

Is synchronous or semi-synchronous SnapMirror available in Data ONTAP 8.3.x?

Answer: No. MetroCluster uses RAID SyncMirror at the aggregate level. NetApp SnapMirror software in clustered Data ONTAP is at the volume level and operates in asynchronous mode only.

What is the read and write behavior with plexes in SyncMirror (including 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 nondisruptively, 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 eight-node and four-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.

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 two-node fabric MetroCluster, because storage and FC-VI replication occurs over the same ISLs, the behavior is the same as in four-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.

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.

Does epsilon play a role in MetroCluster availability?

Answer: No. MetroCluster 8.3.0 uses two-node clusters, and there is no epsilon in two-node clusters. In two-node, each cluster is a single-node cluster; therefore, there is also no epsilon node designation. For eight-node MetroCluster, epsilon does play a role in determining which sections of cluster continue serving data. It is similar to a 4-node cluster behavior.

How is the MetroCluster DR partner assigned in the 8-node MetroCluster configuration?

Answer: In the 8-node MetroCluster configuration the DR partner site will have two HA pairs. MetroCluster will match the lowest controller system id of the primary site with the lowest controller system id at the

partner site. All controllers in the DR group at each site must be the same platform. But the two DR groups may have different platforms, e.g. AFF8080 in DR group #1 and FAS8080 in DR group #2

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

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

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

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

3.2 Nondisruptive Operations

Does MetroCluster support clustered Data ONTAP nondisruptive upgrades (NDU)?

Answer: Yes. In the eight-node and four-node configurations, NDU is used in MetroCluster to upgrade the Data ONTAP version. This method uses takeover and giveback on the HA pairs at each site to upgrade Data ONTAP nondisruptively. For a major update, **the two sites must be updated in parallel** to preserve the DR protection. For minor release updates (for example, from 8.3.0 to 8.3.1), 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 Data ONTAP upgrades at each site.

In two-node configurations, negotiated switchover is used to update each cluster in turn.

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. 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 Data ONTAP release as the surviving nodes. After both clusters are running the same release, the switchback can proceed.

Is aggregate relocation (ARL) for nondisruptive controller upgrade supported with MetroCluster?

Answer: Yes. In the eight-node and four-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 nondisruptive 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 Clustered Data ONTAP*, which is available on the NetApp Support site.

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

Can a two-node configuration be converted to a four-node configuration?

Answer: Starting in ONTAP 9, two-node fabric and stretch configurations can be converted to four-node fabric configuration. The two-node configuration must be running ONTAP 9.0, as earlier versions are not supported for conversion. Prior to ONTAP 9, conversion from a two-node configuration to a four-node

configuration is not supported. NetApp strongly recommends that customers who need local HA functionality deploy the four-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.

Can a four-node configuration be converted to an eight-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.

Is vol move supported with MetroCluster?

Answer: Yes, vol move is supported in the eight-node and four-node configurations. As in standard clustered Data ONTAP, volumes move between aggregates in a cluster and stay within the same SVM. MetroCluster does not provide the capability to nondisruptively 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.

Can I nondisruptively 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](#).

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.

3.3 Transition from 7-Mode (Including Fabric and Stretch MetroCluster)

How can 7-Mode (including FMC or SMC) customers transition to MetroCluster?

Answer: All transitions from Data ONTAP operating in 7-Mode (including FMC and SMC, as well as configurations without MetroCluster) to MetroCluster 8.3.x require data migration; in-place transition is not available. NetApp anticipates that the majority of customers make the transition as part of a technology refresh cycle. The choice of method for data migration and transition is subject to the same considerations as transitioning from 7-Mode to clustered Data ONTAP.

Can FMC switches and ISLs be shared while transitioning to MetroCluster 8.3.x or 9.0?

Answer: Yes. During transition from a 7-Mode FMC to a fabric MetroCluster configuration, a MetroCluster 8.3.x configuration can be connected to the same Brocade 6510 switches that the 7-Mode FMC uses and can share the existing ISLs. Sufficient bandwidth must be available to support the resultant traffic, and the ISLs must be at a speed that MetroCluster 8.3.x supports. Documentation is available to describe the ports, zoning, and other steps that are required to set up this part of the process.

After the transition is complete, the 7-Mode FMC is disconnected from the fabric. The MetroCluster configuration can remain on its current switch ports. If you want to modify the configuration to the ports that were originally used by the 7-Mode FMC, a disruption is required.

Only Brocade 6510 switches are supported for fabric sharing during transition; for any other switch model used with FMC, separate supported switches and ISLs are required for the MetroCluster 8.3.x configuration. **Only one** 7-Mode FMC can share the switches with one MetroCluster 8.3.x configuration.

Note: In FMC running on Data ONTAP 7-Mode, 8Gb SFPs are used with the Brocade 6510 switch. MetroCluster 8.3.x supports only 16Gb SFPs for the 6510, and all components must connect to switch ports with 16Gb SFPs, including the FC initiators, FC-VI ports, and ATTO FibreBridges. The only exception to the 16Gb requirement is for the switch ports that attach to the ISLs. Therefore, the requisite number of 16Gb SFPs must be purchased to attach the MetroCluster 8.3.x components to a 6510. No extra license is required to enable support of 16Gb on the switch. Ports 23 and higher are used to attach the MetroCluster 8.3.x components, as documented in the [MetroCluster Installation and Configuration Guide](#). If these ports have not already been enabled, the requisite Ports on Demand (PoD) license is required to activate them.

Can the 7-Mode Transition Tool (7MTT) be used to transition from 7-Mode to MetroCluster 8.3.x/9.0?

Answer: Yes, 7MTT 2.0 or later supports transition from 7-Mode (including FMC and SMC) to MetroCluster 8.3.x. For 7-Mode MetroCluster migration to clustered Data ONTAP, you must use 7MTT's copy-based transition (CBT) method. Copy-free transition (CFT) in 7MTT is not supported for MetroCluster. Data is copied from the 7-Mode controllers to the MetroCluster 8.3.x destination cluster. Copying a volume into MetroCluster 8.3.x automatically syncs the data to the remote plex. The 7MTT prechecks advise that any volumes that are in unmirrored aggregates be transitioned to mirrored aggregates. The cutover requires a brief application outage. For more information, see the [7MTT Copy-based Transition Guide](#).

Application-, client-, and host-level methods, such as rsync or VMware Storage vMotion, can also be used. For more information, see [TR-4052: Successfully Transitioning to Clustered Data ONTAP](#) and [TR-4336: Enterprise Application Transition to Clustered Data ONTAP 8.3](#).

Is there a nondisruptive transition from 7-Mode (including FMC and SMC) to MetroCluster?

Answer: No. A disruption of applications is generally required when transitioning unless a nondisruptive application-specific method is used. The length of the disruption depends on the method that is used to migrate the data and how much manual reconfiguration is required. If you use 7MTT, the disruption time should be similar to the time required for transitioning to a clustered Data ONTAP environment without MetroCluster. For more information, see the 7MTT transition documentation.

Can 7MTT be used to transition from single-node or HA pair 7-Mode systems to MetroCluster?

Answer: Yes. 7MTT 2.0 or later supports transition from 7-Mode systems without MetroCluster to MetroCluster 8.3.x/9.0. The 7MTT prechecks advise that any volumes that are in unmirrored aggregates be transitioned to mirrored aggregates.

Can 7MTT transition 32-bit data from 7-Mode controllers to MetroCluster?

Answer: No. In any transition to Data ONTAP 8.3.x (MetroCluster or other than MetroCluster) by using 7MTT, only 64-bit data (including Snapshot copies) can be migrated. No 32-bit data can exist in clustered Data ONTAP 8.3.x. If 32-bit volumes or Snapshot data is detected on the source system, 7MTT prechecks block the transition. Before the transition, 32-bit aggregates must be converted to 64-bit (either by in-place upgrade in Data ONTAP 8.2.1 or, for earlier releases, by adding disks), and 32-bit Snapshot copies must be aged out or deleted. Before you can use 7MTT to transition to MetroCluster 8.3.x, the source 7-Mode controllers (whether MetroCluster or not) must be, at a minimum, the Data ONTAP 8.1.x release. For more information about 7MTT, see [TR-4052: Successfully Transitioning to Clustered Data ONTAP](#).

Can equipment that was previously running 7-Mode (FMC and SMC) be reused for MetroCluster 8.3.x/9.0?

Answer: Yes, controllers and storage can be reused. Verify that all hardware components (controllers, cards, storage, switches, ISL speed) are supported with the ONTAP version; additional or replacement parts might be required. In addition:

- Sufficient storage on each site must be available. Follow the recommended minimum disk shelves for the MetroCluster configuration.
- Follow the guidance for matching the controller models, as documented previously in this FAQ.
- Only the 16Gb FC-VI card (X1928) is supported with MetroCluster 8.3.x/9.0; the 8Gb FC-VI card (X1927) is not supported. If you order the FC-VI card as an add-on or quick shipment, you must select 16Gb SFPs (X6602A). These SFPs are automatically included in a MetroCluster quote.
- 16Gb SFPs are required for all ports on Brocade 6510 switches, apart from ISL ports. The Cisco 9148 switch is 8Gb speed only, so SFPs do not need to be replaced. The appropriate switch licenses must also be available.
- As always, check with the IMT to verify supported configurations, including ISL speeds. To determine which additional parts are required, NetApp recommends that you create a quotation for a new MetroCluster configuration and compare the features generated with the hardware available.

Data ONTAP licenses and entitlements must be converted to the correct clustered Data ONTAP configuration. To configure MetroCluster, all nodes must be wiped (zero all disks) and installed freshly with Data ONTAP 8.3.x/9.0, so all data must have been previously copied or transitioned.

Note: On Brocade 6505 switches, 16Gb SFPs are also required with MetroCluster 8.3.x/9.0. However, because this model is not supported with 7-Mode FMC, it is not a consideration for equipment reuse.

3.4 Transition from Clustered Data ONTAP

How can existing clustered Data ONTAP customers transition to MetroCluster?

Answer: Transitioning from clustered Data 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 clustered Data 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 Data ONTAP 8.3.x for MetroCluster.

SnapMirror can be used to transition data from a clustered Data 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 clustered Data ONTAP environment to MetroCluster.

Is there a nondisruptive transition from clustered Data ONTAP to MetroCluster?

Answer: No. Disruption of applications is required when transitioning unless an application-specific nondisruptive method is used. The length of the disruption depends on the method that is used to migrate the data and how much manual reconfiguration is required.

Can 7MTT be used to transition from clustered Data ONTAP to MetroCluster?

Answer: 7MTT does not support transition from clustered Data ONTAP to MetroCluster. It is for source systems that are running 7-Mode only.

Can nodes that were previously running another Data ONTAP release be reused for MetroCluster 8.3.x?

Answer: Yes, the nodes can be reused, as long as the hardware is supported with 8.3.x (including four identical controllers). In addition, all the required additional hardware for MetroCluster 8.3.x 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 Data ONTAP 8.3.x.

4 Configuration

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 6.3 and OnCommand Performance Manager 2.0 also support MetroCluster 8.3.x. This support represents a substantial functionality improvement over fabric and stretch MetroCluster, which did not have management tool support. OnCommand Balance 4.2 and later (for Data ONTAP 8.3.x) does not support MetroCluster configurations.

Is there an easy way to configure the fabric switches?

Answer: Golden configuration files (reference configuration files) for the Brocade 6505 and 6510 and the Cisco 9148, 9148S, and 9250i switches are available for download from the NetApp Support site. The files contain the required commands to automatically configure the switches.

How many ports for ISLs and stacks are included in the switch golden configuration files?

Answer: Four ISL ports are configured. For the 6510 and 9148/9148S/9250i, ports for four storage stacks are configured. For the 6505, two storage stack ports are configured. For the 6505, then, ports 6 and 7 are assigned for storage attachment, and ports 8, 9, 10, and 11 are assigned for ISLs. This setup then requires only one PoD license (12 ports). If more than two stacks and fewer than four ISLs are required, ISL ports can be reconfigured for storage. In that case, manual setup is required. If four ISLs and more than two stacks are required, an additional PoD license is necessary. For information about the options for the 6505, see the following knowledge base article

<https://kb.netapp.com/support/index?page=content&actp=LIST&id=1015278>.

How is MetroCluster configured at initial setup?

Answer: Setup and configuration of MetroCluster are covered in detail in the [MetroCluster Installation and Configuration Guide](#) in the NetApp [product documentation](#). The [MetroCluster Installation Express Guide](#) is also available, which is appropriate for factory-configured hardware with the default settings. As a rule, the clusters are shipped preconfigured from the factory, with the disks assigned equally across all nodes, Data ONTAP installed, switches and FibreBridges configured (as applicable), clusters peered, and MetroCluster configured. It is required only to change the IP addresses of the management LIFs to match the customer's network. If it is required or preferred to install from scratch, follow the process in the [MetroCluster Installation and Configuration Guide](#).

For a scratch installation, rack and cable the hardware, assign disk ownership, configure the switches and FibreBridges, install Data ONTAP, create the clusters. In eight-node and four-node configurations enable HA. Then execute the following steps.

Note: These steps are a summary only; follow closely the detailed process in the [MetroCluster Installation and Configuration Guide](#):

1. Create intercluster LIFs and define cluster peering relationships. This step is required for the cluster configuration replication.
2. Mirror the root aggregate on each node by using the `storage aggregate mirror` command.

3. If you use the default setup option, in an eight-node or four-node configuration, create a mirrored data aggregate on each node; in a two-node configuration, create two mirrored data aggregates on each node. If you use the advanced mode option to configure with a single data aggregate per cluster, only one mirrored data aggregate per node is required.
4. Run the `metrocluster configure` command on one of the clusters to create the MetroCluster relationship. This command configures both clusters for MetroCluster. Two MDVs are created on each cluster to store the configuration replication information. By default, each MDV is on a separate data aggregate; otherwise, if you use the advanced option, both MDVs for each cluster are created on the same data aggregate. The command also sets up and initializes the NVRAM and configuration replication relationships between the DR partners.
5. Verify the configuration with the `metrocluster check` command. Also run Config Advisor to check cabling and other components.
6. Create SVMs and populate them as desired.
7. Test switchover and switchback.

How is the DR partner node chosen in an eight-node and four-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 [MetroCluster Installation and Configuration Guide](#).

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.

4.1 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 two-node, four-node, and eight-node MetroCluster configurations?

Answer: In two-node configurations, hardware failure or panic of one of the nodes triggers an automatic switchover, similar to the behavior of FMC and SMC. In a four-node and eight-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 both two-node, four-node, and eight-node configurations, you must run the Tiebreaker software in automatic switchover mode in the event of a site disaster. See [section 5](#) for more information about the MetroCluster Tiebreaker software.

Does Automatic switchover require a Product Variance Request (PVR)?

No – Automatic switchover is now integrated into ONTAP 9.0 and doesn't need a PVR. However, this only handles controller failures while the storage is still accessible. If an Automatic Switchover is needed to handle Site disaster you need to deploy Metrocluster Tiebreaker Software in "observer-mode false".

What is the difference between negotiated switchover and switchover on disaster?

Answer: By default, the `metrocluster switchover` command runs as a negotiated switchover. In a negotiated switchover, both sites are active, and the switchover is a clean operation because both local and remote plexes are available. NVRAM can be flushed to disk on the source plexes before the switchover so that NetApp WAFL® (Write Anywhere File Layout) replay is not necessary on the remote plexes. Negotiated switchover is useful for test purposes or when it is known in advance that a site or all ISL links will become unavailable.

Switchover on disaster (forced switchover) is performed by specifying a flag on the switchover command. The nodes, switches, and storage on the site that is being switched over must be powered down or otherwise disconnected from the surviving site before the switchover is executed. The switchover proceeds without access to the source plexes. Switchover on disaster takes longer than a negotiated switchover because additional steps are required, including WAFL replay.

The time needed for detection of a disaster and for an administrator to follow the institutional procedure to declare and initiate a disaster must be factored in. These considerations are outside of the control of the switchover command and affects the ability to perform a switchover nondisruptively to applications. If Tiebreaker is used in automatic mode (with PVR approval), detection and switchover proceed without human intervention, and the switchover time is predictable to complete nondisruptively for applications.

Note: All switchovers are disruptive to all SMB/CIFS applications, including SMB 3.0.

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.

Is it possible to perform an automatic switchover on disaster?

Answer: Automatic switchover on disaster is now integrated within MetroCluster. However, it must have access to remote storage in order to perform a switchover. The MetroCluster configurations (and all similar solutions from other vendors) cannot correctly arbitrate a switchover in a scenario in which all links between the sites are disabled but both sites remain operational. Because neither site is able to contact the other, each site could attempt to switch over, resulting in a split-brain scenario in which both sites are active on the same plexes. For this reason, a third site is required to correctly determine whether one site is actually down or whether neither site can contact the other. The optional NetApp MetroCluster Tiebreaker software is available that runs on a third site. See [section 5](#) for more information about Tiebreaker software.

Switchover is automatic if all of the nodes in a cluster fail: for example, due to a hardware failure or panic. Because mailbox access on both nodes is required for an automatic switchover, the failed node's storage must remain available.

5 Interoperability

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.

NetApp recommends OnCommand Unified Manager 6.3 as the minimum version for monitoring and alerting of MetroCluster components.

What events related to MetroCluster cause a health monitor alert?

Answer: To display the details of available health monitor alerts, use the following commands:

```
system health alert definition show -subsystem MetroCluster
system health alert definition show -subsystem MetroCluster_Node
```

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.

Is Tiebreaker available for MetroCluster?

Answer: Yes. Tiebreaker capability is available for MetroCluster. To download the software plus the installation guide and information about the software prerequisites, go to the software downloads section of the NetApp Support site at <http://mysupport.netapp.com/NOW/cgi-bin/software/>, then select MetroCluster Tiebreaker. No additional Data ONTAP license is required for the Tiebreaker software.

The capability known as NetApp MetroCluster Tiebreaker software is a standalone application that runs on a Linux host. Tiebreaker software enables detection of split-brain situations in which each site is isolated from the other but is otherwise operational. Tiebreaker software is supported by NetApp.

Tiebreaker software operates in observer mode and can detect and send an alert if a disaster that requires switchover has occurred. The switchover then can be issued manually by the administrator. With a PVR, the Tiebreaker software can be configured for automatic switchover on disaster.

For ONTAP 9.0, automatic switchover with Tiebreaker no longer requires a FPVR. The functionality is included in ONTAP 9.0.

The switchover time within the client/host has RTO <= 120 seconds

Notes:

- (1) For CIFS, TBUSO is no different from USO. CIFS is disruptive across SO/SB. CIFS hosts re-connect post SO/SB.
- (2) Host I/O may time out during USO
- (3) ONTAP support not available
- (4) Testing required before PVR can be approved
- (5) All Tie Breaker configurations are PVR only
- (6) Intermittent failures with Solaris nSO and SB, where the zpool gets into degraded state after few iterations of nSO/SB. We had tested Solaris with only zpool config (Not SVM / metaset).
- (7) AIX testing was on bare metal and not VIOS

(8) Supported for non-MetroCluster configurations

Supported Operating systems with Tiebreaker:

AIX 7.1 TL4

ESX 5.5

RHEL6 U7

RHEL7 U2

Windows 2012 R2 (with Windows Failover Cluster – WFC)

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.

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.

Is SnapLock Enterprise supported in MetroCluster?

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

Is SnapLock Compliance supported in MetroCluster?

Answer: No. SnapLock Compliance is not supported.

References

- TR-4375: MetroCluster for clustered Data ONTAP
<http://www.netapp.com/us/media/tr-4375.pdf>
- MetroCluster for clustered Data ONTAP SE Technical Presentation
<https://fieldportal.netapp.com/content/203487?assetComponentId=203585>
- OnCommand Unified Manager
<https://fieldportal.netapp.com/content/206725?assetComponentId=206824>
- TR-3838: Storage Subsystem Configuration Guide
<https://fieldportal.netapp.com/content/190829?assetComponentId=190927>
- Hardware Universe
<https://hwu.netapp.com/Controller/Index>
- Clustered Data ONTAP Maximum Configuration Guide – MetroCluster Configurations
<https://fieldportal.netapp.com/content/203281?assetComponentId=203379>
- MetroCluster Service Guide
https://library.netapp.com/ecm/ecm_download_file/ECMP1650547
- MetroCluster Management and Disaster Recovery Guide
https://library.netapp.com/ecm/ecm_download_file/ECMP12458277
- MetroCluster Installation and Configuration Guide
https://library.netapp.com/ecm/ecm_download_file/ECMP12454947
- Tiebreaker Software Installation and Configuration Guide
https://library.netapp.com/ecm/ecm_download_file/ECMP12007400
- NetApp Interoperability Matrix
<https://mysupport.netapp.com/matrix/>
- Cisco Storage Interoperability Matrix
<http://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/interoperability/matrix/intmatrx/Matrix1.html>

Contact Us

Do you have a question that you think we should include in this FAQ?

Ask us at docfeedback@netapp.com.

Include MetroCluster 9 TECHNICAL FAQ in the subject line.

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 © 1994–2016 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, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Fitness, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANshare, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, SnapCopy, Snap Creator, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, WAFL, and other names are trademarks or registered trademarks of NetApp Inc., in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.