



MetroCluster IP

Solution Architecture and Design



Abstract

Lenovo® MetroCluster™ is a continuously available storage solution for Lenovo ONTAP® running on Hybrid and AFA systems. MetroCluster IP is the latest evolution that uses an Ethernet-based back-end storage fabric. MetroCluster IP provides a highly redundant configuration to meet the needs of the most critical business applications. Because MetroCluster IP is included in ONTAP, it does not require a separate license, and it provides NAS and SAN connectivity for clients and servers that use ONTAP storage.

© Copyright Lenovo .

LIMITED AND RESTRICTED RIGHTS NOTICE: If data or software is delivered pursuant to a General Services Administration (GSA) contract, use, reproduction, or disclosure is subject to restrictions set forth in Contract No. GS-35F-05925

TABLE OF CONTENTS

1	MetroCluster Overview	6
1.1	Continuous Availability Solution Overview	6
1.2	MetroCluster IP Compared to MetroCluster FC	7
2	MetroCluster IP Architecture	8
2.1	Disaster Recovery Group	9
2.2	Replication in MetroCluster IP	9
2.3	Network	11
2.4	Storage	13
3	Solution Design	14
3.1	Confirming Support	14
3.2	Sizing a Solution	14
3.3	Hardware Components	16
3.4	MetroCluster IP Switches	16
3.5	Network Adapters	17
3.6	Network Configuration	17
3.7	AFA DM7000F and DM7000H MetroCluster IP Node to 100Gb Switch Connections	18
3.8	AFA DM5000F and DM5000H MetroCluster IP Node to 25Gb Switch Connections	20
3.9	Cluster Fabric ISL	20
3.10	MetroCluster IP ISL Link Design	20
3.11	Shared Layer 2 MetroCluster IP ISL	25
3.12	Using Compliant Switches – MetroCluster IP with Existing Switches	25
3.13	Active-Passive Configurations	25
4	Operation and Administration	27
4.1	Switchover and Takeover	27
4.2	ONTAP Mediator Software	27
4.3	TieBreaker Software	28
5	Interoperability	31
5.1	SnapMirror	31
5.2	Lenovo FlexGroup Volumes	31
5.3	Lenovo FlexCache	31
5.4	Lenovo FabricPool	31
5.5	SVM Mirror and SVM Disaster Recovery	31
	Where to Find Additional Information	33

Contacting Support	33
Notices	34
Trademarks	35

LIST OF TABLES

Table 1) MetroCluster IP switch models.....	16
Table 2) MetroCluster IP network adapters.....	17
Table 3) AFA DM7000F cluster network node to switch cables.....	19
Table 4) AFA DM7000F MetroCluster IP network node to switch cables.....	19
Table 5) ISL maximum characteristics.....	20
Table 6) 40Gb 3m to 5m distance between switches (approximate cable length).....	21
Table 7) 100Gb 3m to 5m distance between switches (approximate cable length).....	21
Table 8) Short-range optical module for 40GbE switch.....	22
Table 9) Short-range optical module for 100GbE switch.....	22
Table 10) 40Gb and 100Gb optical cables.....	22

LIST OF FIGURES

Figure 1) MetroCluster IP and VMware vSphere Metro Storage Cluster.....	7
Figure 2) MetroCluster IP architecture.....	8
Figure 3) Storage and server.....	8
Figure 4) MetroCluster HA and disaster recovery.....	9
Figure 5) Mirroring write data blocks.....	11
Figure 6) MetroCluster IP combined fabric.....	12
Figure 7) AFA DM7000F one site network example.....	16
Figure 8) DM5000F one-site network.....	17
Figure 9) MetroCluster IP 100Gb switch port assignments.....	18
Figure 10) MetroCluster IP 40Gb switch port assignments.....	18
Figure 11) MetroCluster IP 10/25Gb switch port assignments.....	18
Figure 12) AFA DM7000F node to switch network connections.....	20
Figure 13) Site A passive DWDM example using 10Gb optical modules and QSA.....	23
Figure 14) ISL with 10Gb port adapter.....	24
Figure 15) Active-passive cluster or site.....	26
Figure 16) Active-passive HA.....	26
Figure 17) MetroCluster Mediator site.....	28
Figure 18) MetroCluster tiebreaker site.....	29
Figure 19) Tiebreaker site link failure.....	29
Figure 20) Tiebreaker site failure.....	30
Figure 21) SVM disaster recovery.....	32

1 MetroCluster Overview

Lenovo MetroCluster configurations are used by thousands of enterprises worldwide for high availability (HA), zero data loss, and nondisruptive operations both within and beyond the data center. MetroCluster is a free feature of ONTAP software that synchronously mirrors data and configuration between two ONTAP clusters in separate locations or failure domains.

MetroCluster provides continuously available storage for applications by automatically handling two objectives:

- Zero recovery point objective (RPO) by synchronously mirroring data written to the cluster
- Near zero recovery time objective (RTO) by mirroring configuration and automating access to data at the second site

MetroCluster provides simplicity with automatic mirroring of data and configuration between the two independent clusters located in the two sites. As storage is provisioned within one cluster, it is automatically mirrored to the second cluster at the second site. Lenovo SyncMirror® provides a complete copy of all data with a zero RPO, this means that workloads from one site could switch over at any time to the opposite site and continue serving data without data loss.

MetroCluster handles the switchover process of providing access to NAS and SAN-provisioned data at the second site. The design of MetroCluster as a validated solution contains sizing and configuration that allows a switchover to be performed within the protocol timeout periods or sooner (typically less than 120 seconds). This results in a near zero RPO with the recovery for storage occurring within the storage protocol timeout periods. Applications can continue accessing data without incurring failures.

MetroCluster is available in several variations defined by the back-end storage fabric. There are two main types of storage fabric for MetroCluster: FC and Ethernet. The Ethernet storage fabric is referred to as MetroCluster IP.

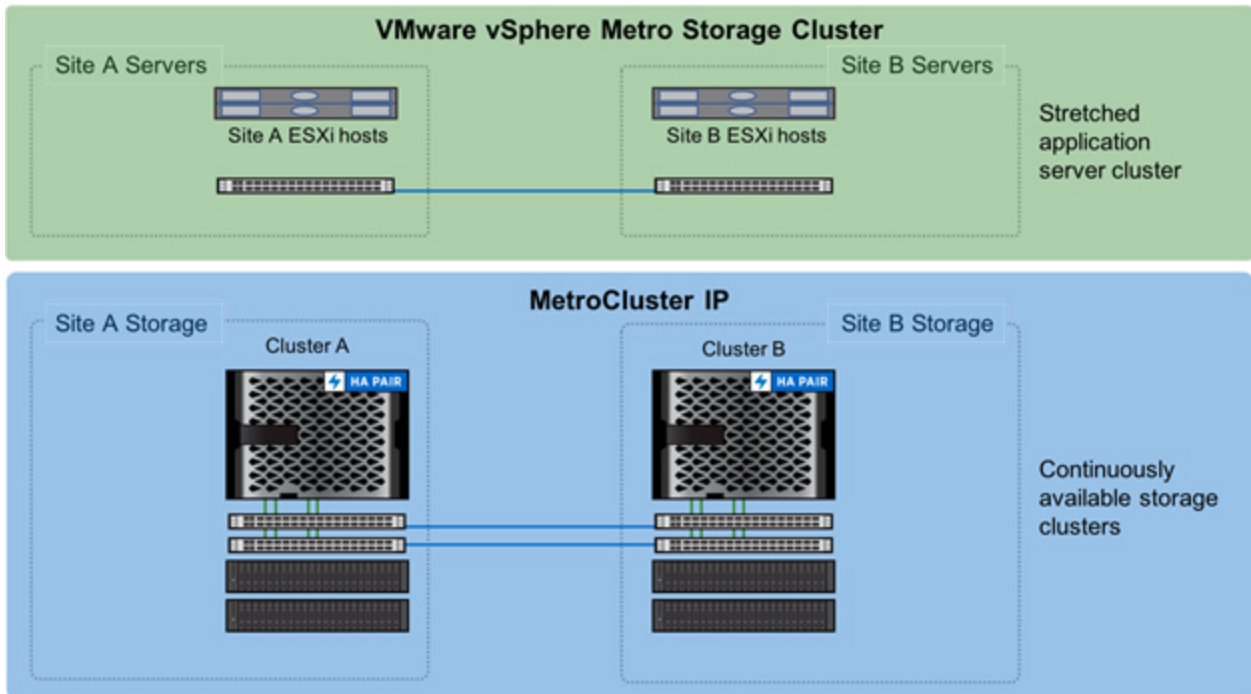
With the evolution of Ethernet storage and the adoption of high-performance Ethernet, MetroCluster IP offers a compelling solution for continuously available storage with a lower component count, lower cost, and increased functionality.

1.1 Continuous Availability Solution Overview

MetroCluster fulfills the need for continuously available storage. When combined with similar application availability products, the complete solution provides a highly resilient architecture that can continue operating even in the event of a site-wide disaster.

One example is using MetroCluster IP with VMware vSphere Metro Storage Cluster (vMSC). Combining the two products creates a highly resilient virtualized infrastructure that addresses the needs of business-critical applications. MetroCluster IP provides storage availability and vMSC provides a cross-site compute cluster that is available to operate even in the event of a complete site outage.

Figure 1) MetroCluster IP and VMware vSphere Metro Storage Cluster.



Similar multisite application solutions are available for databases and other applications that work well with MetroCluster.

1.2 MetroCluster IP Compared to MetroCluster FC


The following features outline the differences between MetroCluster IP and FC:

- MetroCluster IP uses an Ethernet back-end storage fabric rather than an FC back-end storage fabric, eliminating the need for dedicated FC switches.
- MetroCluster IP collapses the intercluster switches for both local and remote replication, eliminating the need for FC switches.
- MetroCluster IP does not require SAS bridges.
- MetroCluster IP replicates NVRAM with iWARP by using the same back-end Ethernet ports as the storage network.
- MetroCluster IP accesses remote disks using iSCSI protocol with the remote disaster recovery node acting as the iSCSI target, supporting flash systems with integrated storage.

2 MetroCluster IP Architecture

MetroCluster IP uses an Ethernet storage fabric. The MetroCluster storage fabric, also referred to as the back-end storage fabric, is used solely by ONTAP. It is a separate dedicated network for ONTAP cluster interconnect, MetroCluster SyncMirror, and MetroCluster NVRAM mirror communications.

Figure 2) MetroCluster IP architecture.




MetroCluster IP hardware summary:

- One HA pair controller per site
- Two high-speed Ethernet switches per site:
 - Collapsed intracluster and intercluster switches for local and remote replication
- Disk shelves or internal storage

MetroCluster extends the availability of ONTAP by mirroring data between two independent ONTAP clusters. Each cluster is in a site or failure domain and leverages the standard HA features on the Hybrid or AFA systems. MetroCluster provides the capability to mirror both data and configuration between the two ONTAP clusters. MetroCluster includes validated system parameters and limits designed to provide failover from one site to the other within standard timeout periods for storage protocols.

MetroCluster features and hardware are a certified subset of the typical ONTAP Hybrid and AFA systems.

Figure 3) Storage and server



MetroCluster has an architecture that can be broken down logically into several functional areas or components. Understanding how these components, such as replication, operate is important to build a well architected solution and to administer the solution.

Following are the main MetroCluster components:

- Disaster recovery group
- Replication
- Network
- Storage

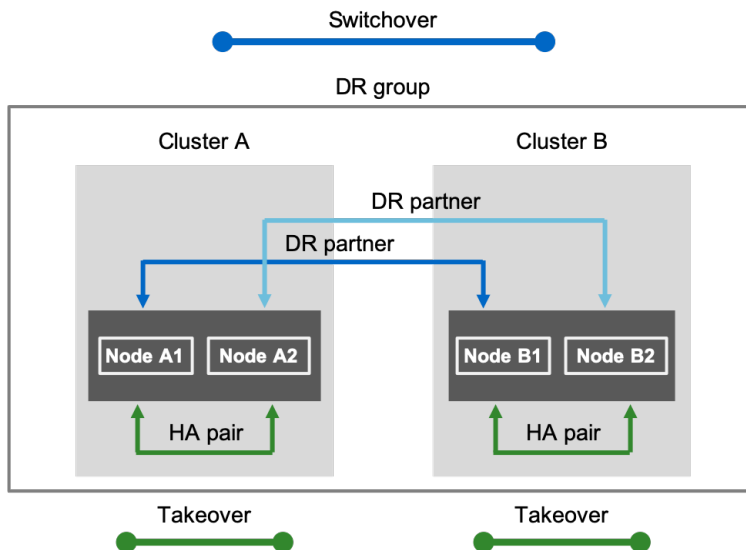
2.1 Disaster Recovery Group

MetroCluster IP uses the concept of disaster recovery for group and partners to determine the relationship for failover and switchover. The two clusters (site A and site B) are configured together as a disaster recovery group. Within the group, the nodes are associated as disaster recovery partners.

The HA relationship is the same as in a standard cluster. HA protects against single controller faults and performs failover locally. HA is also leveraged for nondisruptive ONTAP updates. For site-wide faults, the disaster recovery relationship is used to switch from site A to site B, which is referred to as switchover.

The disaster recovery partner relationship is configured in the initial MetroCluster setup and does not change. There is one command that assigns one node from cluster A and one node from cluster B as partners. The remaining nodes are automatically assigned to complete the disaster recovery group configuration.

Figure 4) MetroCluster HA and disaster recovery.



2.2 Replication in MetroCluster IP

MetroCluster IP leverages direct attached storage, which eliminates the need for external SAS bridges to connect disks to the storage fabric. Each node in the disaster recovery group acts as a storage proxy or iSCSI target that exports its disks to the other nodes in the group. iSCSI (SCSI over TCP/IP) is the storage transport protocol for the IP fabric that allows the iSCSI initiator and targets to communicate over a TCP/IP fabric. Each node in the disaster recovery group accesses its remote storage through an iSCSI initiator that establishes an iSCSI session with a remote disaster recovery partner iSCSI target.

The use of iSCSI and direct-attached storage also enables the use of systems that have internal disks. iSCSI allows the nodes to provide the disaster recovery partner node access to internal storage in addition to storage devices located in external disk shelves.

MetroCluster has three planes of replication:

1. Configuration replication
2. NVRAM replication
3. Storage replication

Configuration Replication Service

Configuration is replicated using the configuration replication service (CRS). CRS replicates the configuration synchronously from the local node to the disaster recovery partner in the partner cluster. This replication is carried out over the cluster peering network. The peering network is a customer-supplied IP network with intercluster LIFs. The peering network for MetroCluster is the same as a regular ONTAP cluster, such as the peering network used for ONTAP SnapMirror®. It can also be the same front-end network that hosts use to access storage.

Note: The cluster peering network is typically the front-end or host-side network. This traffic does not use the MetroCluster IP back-end storage fabric network.

The information, referred to as objects, that are replicated includes the cluster configuration and the storage virtual machine (SVM) configuration. CRS replicates configuration objects between the clusters including:

- SVMs, LIFs, volumes, aggregates, and LUNs
- Protocol objects such as CIFS, NFS, and SAN

New objects are transferred to the remote cluster as standby objects. Object updates are propagated as they occur. If there is an interruption in the cluster peering network that affects CRS, replication catches up automatically once the connection is reestablished.

CRS requires a small volume on a data aggregate to store metadata referred to as the metadata volume (MDV). Each cluster requires one MDV. Typically, this is only a concern when an active-passive site configuration is being planned. In this case, a small volume in a data aggregate must be created on the cluster in the passive site.

NVRAM Replication

NVRAM replication mirrors the local node NVRAM to the NVRAM for the remote disaster recovery node. MetroCluster IP uses iWARP to replicate NVRAM over a TCP/IP connection. The iWARP protocol is offloaded in hardware with RMDA-capable network adapters to make sure that latency is not affected by the IP stack.

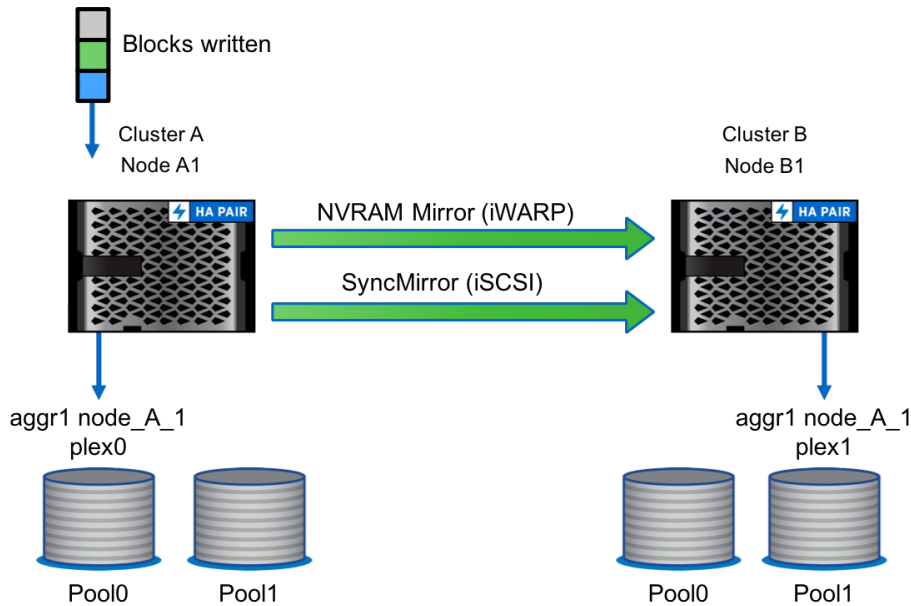
Storage Replication

Storage replication mirrors the local and remote back-end disks using RAID SyncMirror (RSM). MetroCluster IP presents the back-end storage as logically shared by making each node in a disaster recovery group serve as a remote iSCSI target. For a node to access its remote back-end disks, it goes through its remote disaster recovery partner node to access the remote disks that are served through an iSCSI target.

Figure 5 illustrates the MetroCluster IP planes of replication for NVRAM and storage. NodeB1 exports its locally attached disks to remote partner nodes in the disaster recovery group through an iSCSI target. NodeA1 pool0 disks are locally attached to NodeA1, while pool1 remote disks are exported through the iSCSI target hosted by B1. The aggregate `aggr1 node_A_1 local plex 0` consists of locally attached

disks from pool0. The aggregate `aggr1 node_A_1 remote plex 1` consists of disks directly attached to B1 and exported to A1 through the iSCSI target hosted in B1.

Figure 5) Mirroring write data blocks.



Blocks are written to both sides of the cluster with both NVRAM (or NVMEM) and SyncMirror. SyncMirror writes data to two plexes for each mirrored aggregate, one local plex and one remote plex. SyncMirror writes occur in the RAID layer, which means that any storage efficiencies such as deduplication and compression reduce the data written by the SyncMirror operations.

Blocks read are obtained from the local storage and do not affect performance or utilization of the Inter-Switch Links (ISLs) for read operations.

Encryption for Replication

MetroCluster does not provide a mechanism to encrypt data being sent between the sites. There are currently two options for ensuring that site-to-site data is encrypted. Lenovo recommends using a DWDM device to encrypt all data going across the ISLs. However, not all implementations use a DWDM.

Alternatively, you can use host-side encryption of the data. The disadvantage is that this negates any storage efficiencies that ONTAP normally provides.

While it is possible to use Lenovo Volume Encryption to encrypt data written to a volume, any writes are also sent with NVRAM replication including unencrypted block data written by the host.

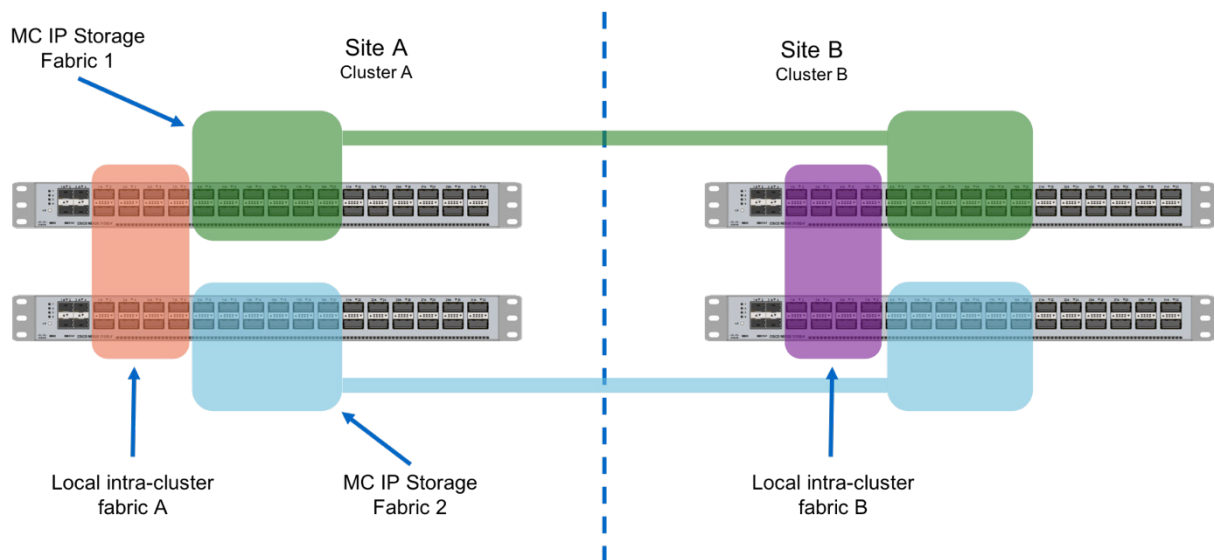
2.3 Network

There are two independent storage fabrics for MetroCluster:

- MetroCluster IP network
- Cluster interconnect

Each network is dedicated to certain functions. There are specific virtual LANs (VLAN) that map to each of the networks to create separate data link layers, or layer 2 in the OSI standard.

Figure 6) MetroCluster IP combined fabric.



Lenovo provides a standard switch configuration known as a reference configuration file (RCF). RCF files are specific to the switch model and are available on the Lenovo support site in the [Drivers&Software downloads section](#). The RCF files must be used to make any configuration changes to the switches. The RCF file is a bundle of four individual configuration files – one per switch.

The RCF designates the VLAN and the channel group identifiers (IDs). These are used only within the back-end storage switches. There are specific requirements for IP addresses. See the [MetroCluster IP Installation and Configuration guide](#) for a worksheet and description of the requirements.

Cluster Network

The ONTAP cluster network is a local only network. The cluster network does not connect between sites. All cluster traffic is local to the site where the nodes are located.

On the MetroCluster IP switch, there are ports dedicated to the ONTAP cluster network. There are two ISL ports that connect the two site switches together to pass cluster traffic forming a local VLAN that spans the switches. This provides redundancy for the local cluster network.

For nodes that connect at a nonnative port speed, such as the Lenovo AFA DM7000F all-flash storage system, there are designated ONTAP cluster network node breakout ports.

The AFA DM5000F and DM5000H use cluster ports that share the same traffic with MetroCluster IP networking. The combined ports use VLANs to separate the MetroCluster IP traffic.

If you order MetroCluster IP without switches because you already have existing compliant switches, the cluster network ports are cross-connected to create a switchless cluster network. This is a new feature for ONTAP 9.7. For more information, see section 3.12.

MetroCluster IP Network

Each site has two independent MetroCluster IP storage fabrics. Each fabric is connected to a similar remote fabric. The two local MetroCluster IP fabrics do not connect to each other. This is different from the cluster network.

On the MetroCluster IP switch, there are several ports reserved for MetroCluster IP node connectivity. Only two of the ports are currently used for a standard four node MetroCluster IP configuration that includes two nodes per site. Each of the two nodes connects to each switch.

For nodes that use the native port speed of the 100Gb switch, ports 9 and 10 are used. These ports are connected to VLAN 10 that is used for the MetroCluster IP storage fabric. For nodes that connect at a nonnative port speed, such as the AFA DM7000F, there are specific MetroCluster IP node breakout ports designated. For example, with an AFA DM7000F using the 100Gb switch, the node ports are 25 and 26 on the switch. These ports are configured to use the break-out cables and offer a 25Gb port speed.

Node connections to the switch are provided by different Ethernet interfaces. The standard cluster network ports are used, typically onboard ports or the same cluster network interfaces available for standard ONTAP HA configurations.

For the MetroCluster IP network, the Ethernet adapter is a specialized card that is optimized for the internet-wide area remote direct memory access (RDMA) protocol (iWARP). In addition, this includes both a TCP offload engine (TOE) and an iSCSI offload capability. iWARP provides RDMA over high-speed Ethernet. The adapter is used for both iWARP and iSCSI traffic for local and remote replication. Each node has two iWARP/iSCSI adapter ports installed and each port goes to a separate switch. Each switch forms a separate fabric that is not connected locally as shown in Figure 6.

For the Lenovo AFA DM5000F and DM5000H, software iWARP is used. These platforms have a fixed number of network ports. To further enhance the use of network ports for front-end, host-side data access, the two onboard 10GbE ports, e0a and e0b, that are typically reserved as cluster interfaces are combined. This feature allows cluster traffic and MetroCluster IP traffic to share the same ports, and the remaining four network ports can be used for host-side data access.

Note: iWARP is a standards-based protocol described in [IETF RFC 5040 – A Remote Direct Memory Access Protocol Specification](#).

2.4 Storage

Storage for MetroCluster IP is not directly shared across the two sites. It is not necessary to configure unique shelf IDs across the sites. The storage at each site is only directly accessible by the local HA pair. The remote storage is made available by the local nodes using iSCSI as described in the Storage Replication section.

SyncMirror

SyncMirror, or RAID SyncMirror (RSM) is the technology used in MetroCluster to mirror the aggregates between the sites. It enables two plexes to be configured in each aggregate, referred to as pool0 and pool1. Pool0 contains the local storage for a node and pool1 contains the remote mirror copy.

ADP

As of ONTAP 9.4, MetroCluster IP supports Advanced Disk Partitioning (ADPv2) on AFF systems. ADP allows a disk to be partitioned into one root and two data partitions. This allows a more granular disk allocation for better use of capacity especially for creating more efficient root aggregates. With ADP, one node owns a whole disk. Disks must be assigned to each node in the HA pair so that a node has capacity for its root aggregates.

Disks are automatically assigned for configurations that use one, two, four or multiples of four shelves per site. Other multiples require manual disk assignment as described in the installation guide.

See the [MetroCluster IP Installation and Configuration Guide](#) for more information about the ADP configuration. The Lenovo Storage Sizing Tool provides sizing for systems that support ADP.

3 Solution Design

Proper design of a solution is the key to addressing performance, capacity, and resiliency requirements. The overall steps for designing a solution include checking for supported hosts and platform configurations as well as sizing to meet capacity and performance needs. The following issues must be considered:

- Ensuring support for hosts and protocols
- Sizing of a solution for performance
- Sizing of a solution for capacity: active-active, active-passive configurations for capacity
- Reviewing systems limits
- Sizing ISLs between sites
- Cabling requirements

3.1 Confirming Support

Review the [Lenovo Interoperability Matrix](#) to verify that the host-side protocol and operating system versions are supported in the same way as any ONTAP design. Check any alerts noted in the results pages to see if they apply to MetroCluster.

The [Lenovo Press](#) lists system specifications and supported limits. Starting with ONTAP 9.6, the Lenovo Press also contains interoperability information for ONTAP 9.6 and later.

3.2 Sizing a Solution

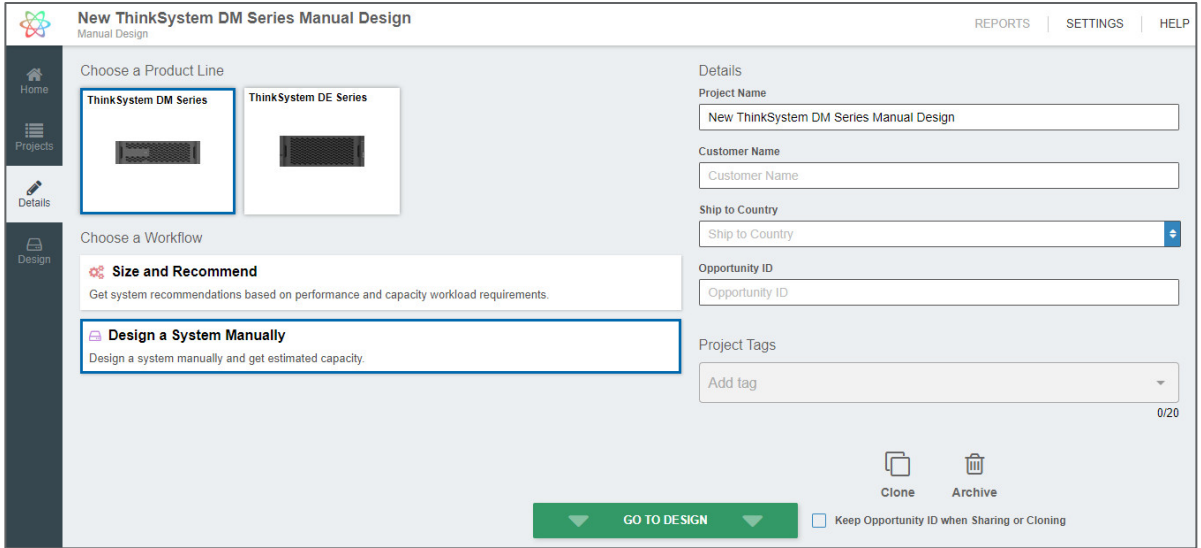
A solution can be sized to meet specific storage capacity or performance requirements. MetroCluster sizing is similar to sizing an HA pair with respect to capacity. With MetroCluster, the storage devices are double the capacity used for an HA pair to provide the mirror copy of the data at the opposite site.

With respect to sizing the performance, the ISLs are a factor that can be accounted for using the ISL sizing spreadsheet.

Lenovo Storage Sizing Tool

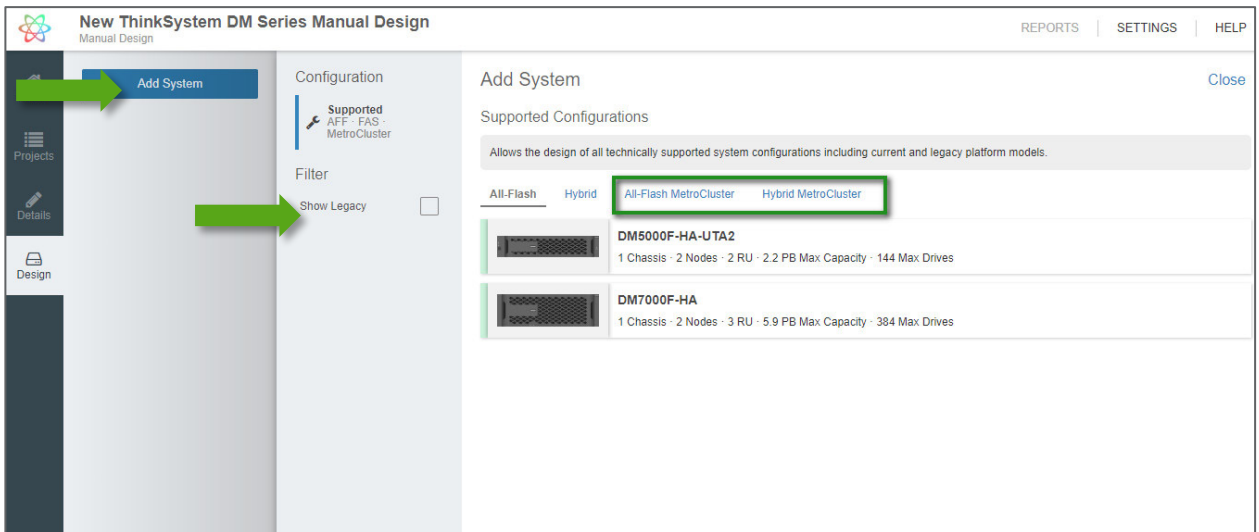
The Lenovo Storage Sizing Tool provides sizing for MetroCluster using the ThinkSystem DM Series manual design capability as follows:

1. Log in to the Lenovo storage sizing portal: <https://storagesizingtool.lenovo.com> (Lenovo login required).
2. Select ThinkSystem DM Series, and then click the Design a System Manually button.



3. Enter details about the project, then click the Go to Design button.
4. In the design page, click the Add System button. Choose Supported from the configuration types.
5. In the supported section, choose All-Flash MetroCluster or Hybrid MetroCluster.
6. Select the desired platform configuration with MetroCluster IP.

Note: The filter defaults to 4-Node MetroCluster IP. If it has been changed, it is necessary to select these filter settings.



7. After selecting the platforms, it is possible to select and add storage. If ADP is available, it is automatically configured for a symmetrical layout.

Inter-Switch Link Sizing

The size and number of ISLs can be determined using a sizing spreadsheet. The spreadsheet allows you to determine maximum performance based on the number and characteristics of the links.

It is also possible to evaluate inter-switch link throughput using ONTAP when validation of link bandwidth is required.

3.3 Hardware Components

Hardware components include storage platforms, controllers, shelves, and storage fabric/cluster switches and cabling. These are all detailed in the [Interoperability Matrix](#) and the [Lenovo Press](#).

For details about the hardware components, see the [MetroCluster IP Installation and Configuration Guide](#).

3.4 MetroCluster IP Switches

The MetroCluster IP switches are platform dependent. Each deployment requires four switches, two per site for redundancy. It is not possible to mix the switch models in a single MetroCluster IP deployment.

See the [Interoperability Matrix](#) on the Lenovo Support site and the [Lenovo Press](#) for information about supported switch models for a specific platform and version of ONTAP.

For port speeds lower than the native port speed, breakout cables are used. For example, when using 10Gb ISL links, breakout cables are used to connect the optical modules.

Table 1 describes the currently available switch models. See [Lenovo Press](#) for specifications.

Table 1) MetroCluster IP switch models.

Lenovo PN	Model Description	Native Port Speed
X190001	Cisco Nexus 3132Q-V	40Gb
X190100	Cisco Nexus 3232C	100Gb
X190005	BES-53248	10/25Gb

Figure 7) AFA DM7000F one site network example.





Figure 8) DM5000F one-site network.



3.5 Network Adapters

MetroCluster IP uses platform dependent specialized network adapters. The network adapters offer high-speed Ethernet and the ability to offload iWARP operations. The adapters have dual ports to attach to the two separate stretched layer 2 Ethernet fabrics.

Table 2 describes the currently available network adapters for specific platforms. See the [Lenovo Press](#) for current details for the specific platform and ONTAP version in use.

Only one MetroCluster IP network adapter is supported per node and must be installed in a specific slot. The MetroCluster IP network adapter provides the MetroCluster node to switch connection that is used for storage and NVRAM replication.

There is a separate network adapter or network ports that are used for the cluster interconnect.

The DM5000F and DM5000H use software iWARP combined on the cluster interfaces. This allows the DM5000F and DM5000H to share traffic on the onboard e0a and e0b interfaces. This feature reduces the port count required for backend storage and provides the maximum number of ports for host-side data interfaces.

The A320 uses onboard ports for both cluster and MetroCluster IP networks.

Table 2) MetroCluster IP network adapters.

Lenovo PN	Description	Platform Supported	Slot
X1116A	MetroCluster IP Dual Port 25GbE Network Adapter	AFA DM7000F / DM7000H	1

3.6 Network Configuration

Network configuration requires that VLANs and IP addresses do not overlap with other networks. VLANs are for the private back-end fabric and are assigned automatically in the switch RCF files. Prior to ONTAP 9.6, RCF files were created by switch model, platform, and ONTAP release. Starting with ONTAP 9.6, there is a new RCF File Generator utility that is described on the RCF downloads page on the [Lenovo Support Site](#) downloads section.

In a typical switch configuration, most of the ports are open for future expansion. The switches create two redundant fabrics. Each node has a cluster connection and a MetroCluster IP node connection.

There are several choices for cabling depending on the distance of the nodes to the switches. If possible, the optimal solution is to locate the nodes and the switches in the same equipment rack. This allows the use of copper twinax cables rather than requiring optical cabling.

The switches are capable of native port speed, either 10/25Gbps, 40Gbps or 100Gbps depending on the model. The native ports are also capable of operating in break-out mode. In this mode, the port is divided into four separate lanes that are used as individual interfaces.

When operating in break-out mode, a port on the 40Gb switch is able to operate as four 10Gbps interfaces that are also dependent on the cable or optical module. Not all optical modules support break-out mode.

The 100Gb switch supports break-out mode as well. When operating at native port speed, a single physical interface operates as four 25Gbps interfaces. Specific cables and optical modules support operating in break-out mode. The 100Gb switch is able to operate ports at 40Gbps for compatibility with 40Gb cables and optics. The 100Gb switch also supports break-out mode when operating ports at 40Gbps. This provides four separate 10Gbps interfaces for each physical port.

The RCF files preconfigure the break-out ports for the specific speeds required for each platform.

Figure 9) MetroCluster IP 100Gb switch port assignments.




Figure 10) MetroCluster IP 40Gb switch port assignments.





Figure 11) MetroCluster IP 10/25Gb switch port assignments.



3.7 AFA DM7000F and DM7000H MetroCluster IP Node to 100Gb Switch Connections

The AFA DM7000F and DM7000H platforms have 10Gb and 25Gb network ports. The ports are connected to the switch to specific ports that are configured in break-out mode in the RCF. The two types of connections, cluster and MetroCluster IP, use different node network ports and cables. The cluster

network uses the onboard 10GbE interfaces and the MetroCluster IP uses 25GbE interfaces on the MetroCluster IP network adapter. This requires the use of two different cable types that correspond to the 10Gb cluster ports and the 25Gb MetroCluster IP ports respectively.

Note: This section describes the cabling required for the Cisco 3232C 100Gb switch.

The cluster network for the AFA DM7000F and DM7000H uses the onboard 10GbE ports. The node's cluster ports connect to the switch using a break-out cable. The break-out cable and RCF configuration settings convert some of the 100Gb switch ports to four 10Gbps links. A single link is used per cable.

Supported Cluster Node to Switch Cable for Use with the AFA DM7000F or DM7000H

Four cables are required per site, and each cable connects one 10GbE port on the node to one port on the switch that is configured for break-out mode.

Table 3) AFA DM7000F cluster network node to switch cables.

Cable PN	Description	Quantity		
		Node	Site	Total
X66120-3	Cable, Cu, 40GbE, QSFP+/4xSFP+, 3m	2	4	8

The MetroCluster IP node connection uses a different break-out cable connected to the MetroCluster IP Ethernet adapter at 25Gb. This cable and corresponding RCF configuration settings converts a switch port to four 25Gbps links. A single link is used per cable.


Supported MetroCluster IP Node to Switch Cable for Use with the AFA DM7000F or DM7000H

Four cables are required per site, and each cable connects one 25GbE port on the node to one port on the switch.

Table 4) AFA DM7000F MetroCluster IP network node to switch cables.

Cable PN	Description	Quantity		
		Node	Site	Total
X-QSFP-4SFP25G-CU2M	Cable, Cu, 25GbE, QSFP+/4xSFP+, 2m	2	4	8

Figure 12) AFA DM7000F node to switch network connections.



3.8 AFA DM5000F and DM5000H MetroCluster IP Node to 25Gb Switch Connections

The AFA DM5000F and DM5000H can use the Broadcom BES-53248 (Lenovo P/N: X190005) 10/25Gb switch for a more economical solution. You can also use the 40Gb and 100Gb Cisco switches with the AFA DM5000F and DM5000H platforms.

3.9 Cluster Fabric ISL

The cluster network fabric is local to a site. Each switch has dedicated ports to the local cluster interswitch link. The cluster ISL uses native port-speed copper cables.

3.10 MetroCluster IP ISL Link Design

ISLs are site dependent and can be the more complex part of architecting a MetroCluster IP solution. This is because of the significant variety of site-to-site links and distances between sites.

Designing an ISL configuration includes doing sizing calculations for the links to determine how many links are required to meet a certain performance from the storage platform. There is an existing sizing tool for performing this calculation on the [Lenovo Storage Sizing tool](#) (login required). The result from the tool provides the number and size of links mapped to the performance of the platform.

The next steps are to determine the components required for each link and the distances that are supported. The distance possible for a link depends on several factors. The ONTAP software version has a maximum supported distance for the ISL link shown in Table 5.

Table 5) ISL maximum characteristics.

ONTAP Release	Maximum Distance	Maximum Round-Trip Latency
9.5 and later	700km	10ms
9.3, 9.4	100km	1ms

For ONTAP 9.5 and later, the site-to-site connection should have a maximum path round trip time (RTT) of 7ms, which is approximately 700km if the latency is purely from distance. In addition to 7ms, a jitter of

up to an additional 3ms is supported, providing a maximum latency in cases when the network is glitchy of up to 10ms, as is noted in the preceding table.

There are several other factors that determine the possible distance for the configuration. Optical modules and matching optical cable configurations provide various maximum supported distances. Multimode optics and cables provide shorter distances at a lower cost. For longer distances, long-range optics and single-mode fiber are required.

For maximum distances, you should consult with the telecommunications provider. Long-distance links use specific telecommunications equipment that can include amplification of the signal to extend the range beyond the optical module's capabilities.

For configurations that are within a single data center, separated by racks in separate availability zones, it might be possible to use standard Ethernet cabling. This greatly simplifies the design and can provide ISLs at native switch port speeds of either 25Gb, 40Gb or 100Gb depending on the switch and modules.

Factors to consider for MetroCluster IP ISL:

- Storage performance requirements
- Existing customer site-to-site network capabilities
- Direct fiber availability
- Multiplexing devices
- Distance: rack to rack, campus, or site distance
- Existing fiber infrastructure (cabling, connectors, patch panel connectors)

Rack-to-Rack: Short Distances

For MetroCluster IP configurations that are within a data center and have close proximity between the sites, it might be possible to use copper cabling for the ISL links. This is often the case with laboratory or test configurations. Table 6 and Table 7 show the part numbers for the switch to switch copper cables for both the 40Gb and 100Gb switches.

Table 6) 40Gb 3m to 5m distance between switches (approximate cable length).

Distance	Cable PN	Description
Less than 1m	X66100-3	Cable, Copper, QSFP+-QSFP+, 40GbE, 1m
Less than 3m	X66100-3	Cable, Copper, QSFP+-QSFP+, 40GbE, 3m
Less than 5m	X66100-5	Cable, Copper, QSFP+-QSFP+, 40GbE, 5m

Table 7) 100Gb 3m to 5m distance between switches (approximate cable length)

Distance	Cable PN	Description
Less than 1m	X66211A-1	Cable, Copper, QSFP28-QSFP28, 100GbE, 1m
Less than 2m	X66211A-2	Cable, Copper, QSFP28-QSFP28, 100GbE, 2m
Less than 5m	X66211A -5	Cable, Copper, QSFP28-QSFP28, 100GbE, 5m

Similarly, the use of optical cabling between racks is possible. This enables a simple ISL configuration when the distances are within the specification of the optical modules.

Table 8) Short-range optical module for 40GbE switch.

Distance	Module PN	Description
Up to 400M on OM4	X65401	XCVR, QSFP+, Optical, 40GbE, Shortwave

Table 9) Short-range optical module for 100GbE switch.

Distance	Module PN	Description
Up to 100M on OM4	X65405	XCVR, QSFP28, Optical, 100GbE, Shortwave

Table 10) 40Gb and 100Gb optical cables.

Length	Module PN	Description
2M	X66200-2	Cable, Optical, OM4, MPO/MPO Type B
5M	X66200-5	Cable, Optical, OM4, MPO/MPO Type B
15M	X66200-15	Cable, Optical, OM4, MPO/MPO Type B
30M	X66200-30	Cable, Optical, OM4, MPO/MPO Type B

Campus Links

Campus links that use direct fiber connections between short distances can be similar to using rack-to-rack ISLs. One potential difference is the use of long-range optics and single-mode cabling to achieve longer distances compared to multimode cabling and short-range optics.

Currently, Lenovo does not offer long-range optical modules for either the 40GbE or 100GbE switches. For designing links that require long-range optics, see the [Cisco support matrix](#) for the specific switch model and the Cisco optical module datasheets to determine distance and connection specifics.

Dedicated Fiber Links

Dedicated fiber links are more common for campus networks connecting buildings located in close proximity. With dedicated fiber links you might want to multiplex signals from many fiber connections onto fewer fiber links. Doing this can maximize utilization and reduce the required number of fibers between the site. Multiplexing of optical signals is called wavelength division multiplexing (WDM) and it is available in two types, coarse wavelength division multiplexers (CWDM) and dense wavelength division multiplexers (DWDM).

CWDM is able to multiplex a smaller number of wavelengths compared to DWDM.

CWDMs are commonly passive devices that optically multiplex and demultiplex the light from the optical modules into a single signal that can be transmitted across a single fiber pair. The optical modules are wavelength specific, sometimes referred to as channel. To multiplex two different fiber signals each source signal must be generated from an optical module that uses a different wavelength. CWDM optical modules are available from Cisco and support eight different wavelengths. This enables the multiplexing of eight fiber signal links onto a single fiber link. The CWDM multiplexer is passive and only contains optics that multiplex and demultiplex the signals. This typically is a lower cost for multiplexing devices and associated optical modules compared to DWDM.

DWDMs use a similar method for merging signals as CWDM devices. The primary difference is that the optical modules are more precise in the signals they generate, enabling a narrower spectral width for a narrow signal and less spacing between the signals. This enables a higher number of signals to be

combined for transmission on a site-to-site fiber link. DWDM devices can be active or passive. Passive devices use the same approach as CWDM where the optical modules transmit a specific wavelength or channel that is merged in the DWDM device to produce a single signal. This signal is transmitted on the longer distance fiber cable between the sites.

DWDM devices are also available as active devices. In this case, the signals between the switch and the DWDM device use standard optics and rely on the DWDM device to produce a signal at the wavelength that can be merged onto the site-to-site fiber link.

For distance, the optical modules provide specifics on the allowable distance and the link characteristics required to meet the specifications. Transmitting at longer distances might require a signal amplifier. There are several types of amplifiers that apply for DWDM such as an optical amplifier. Lenovo recommends consulting with a telecommunications specialist to help design the optimal configuration.

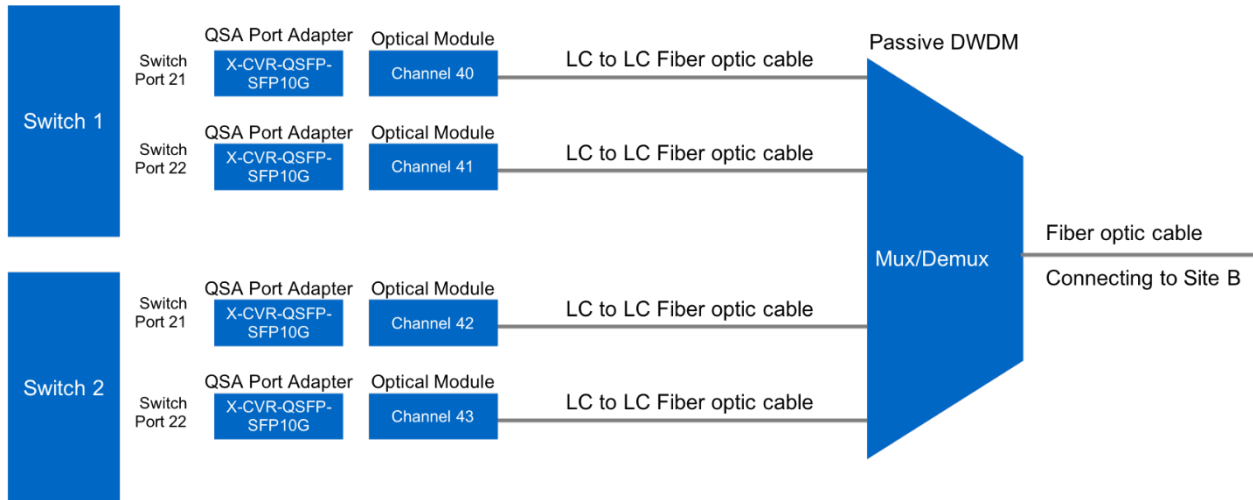
Cisco provides 10Gb SFP+ modules that can be used for coarse or dense wavelength multiplexing. xWDM enables multiple optical signals to be combined or multiplexed on a single fiber pair between sites, then demultiplexed before routing the signal to a switch or device.

Figure 14 is an example of using a passive DWDM and shows a possible mapping for the optical modules on specific channels. All the optical modules for each site use a unique channel. Match the opposite site B optic, same channel.

Example using DWDM modules

- Site A switch 1 port 21 – Site B switch 1 port 21 using two optical modules on channel 40
- Site A switch 2 port 21 – Site B switch 2 port 21 using two optical modules on channel 41
- Site A switch 1 port 22 – Site B switch 1 port 22 using two optical modules on channel 42
- Site A switch 2 port 22 – Site B switch 2 port 22 using two optical modules on channel 43

Figure 13) Site A passive DWDM example using 10Gb optical modules and QSA.



The first module in the example is Cisco part number DWDM-SFP10G-45.32, that is a 10GBASE-DWDM SFP+ module operating on the 1545.32nm wavelength (100-GHz ITU grid) which is ITU channel 40. To complete the configuration in this site, three more modules must be supplied each corresponding to channels 41, 42 and 43. Site B then contains the exact same configuration of optical modules, port adapters, and passive DWDM.


Intercity Links

For the longest distance links, active DWDM or Telco circuits are often used. Connection from the switch to most telecommunications or active DWDM devices is done with the same optical modules that would be used in a data center or rack-to-rack configuration. See section “Rack-to-Rack: Short Distances” for rack-to-rack cabling and modules.

Note: Depending on the equipment, it is also possible that an active DWDM can provide encryption of the ISL traffic.

Lenovo recommends that you consult with a telecommunications specialist to help design the optimal configuration for intercity links.

Figure 14) ISL with 10Gb port adapter.



For more information about cabling and optical modules, see the MetroCluster IP switch technical references in [Where to Find Additional Information](#).

Cisco DWDM optical modules are available in three different configurations:

- Cisco DWDM-SFP10G-XX.XX modules
 - DWDM fixed module supports 40 non-tunable ITU 100GHz wavelengths
- Cisco DWDM-SFP10G-C module
 - DWDM tunable module supports 96 tunable ITU 50GHz wavelengths
 - up to 80km
- Cisco DWDM-SFP10G-C-S module
 - tunable transceiver modules are Ethernet only
 - DWDM tunable module supports 96 tunable ITU 50GHz wavelengths
 - up to 70 km

For more information about supported optical modules and PNs to order from a Cisco partner, see the section “[Where to Find Additional Information](#).”

3.11 Shared Layer 2 MetroCluster IP ISL

ONTAP 9.6 adds support for shared layer-2 MetroCluster IP ISL connections. The MetroCluster IP switches are a required part of the solution. However, having a shared ISL means the switches can connect to an existing site-to-site network. The site-to-site network is a stretched layer-2 network. There are specific requirements for the site-to-site network configuration including VLAN and quality of service settings.

Details for configuration of a shared inter-site link are located in the [MetroCluster IP Installation and Configuration Guide](#).

For the ONTAP 9.6 RC1 release, the 10/25Gb BES-53248 switch is not supported for shared layer-2 ISL. In a future release shared layer 2 will be supported with the 10/25Gb switch.

For the AFA DM5000F and DM5000H, the VLAN IDs are currently fixed at VLAN 10 and VLAN 20. It is possible to perform shared layer 2 with these platforms using the Cisco 40Gb or 100Gb switch. However, the fixed VLAN IDs must be available for exclusive use by MetroCluster IP. It also means that multiple AFA DM5000F and DM5000H configurations of MetroCluster IP are not able to share the layer-2 ISL because the VLAN IDs must be dedicated for each MetroCluster IP configuration. In a future release, the VLANs will be configurable.

3.12 Using Compliant Switches – MetroCluster IP with Existing Switches

ONTAP 9.7 adds support for MetroCluster IP without Lenovo validated switches on some platforms. This allows the use of existing switches that can be used for more than just the MetroCluster IP backend storage fabric. The solution uses a switchless cluster configuration at each site (two-node switchless cluster). In this configuration, the cluster network interfaces are cross connected. The MetroCluster IP interfaces are plugged into existing switches that are compliant with MetroCluster specifications.

3.13 Active-Passive Configurations

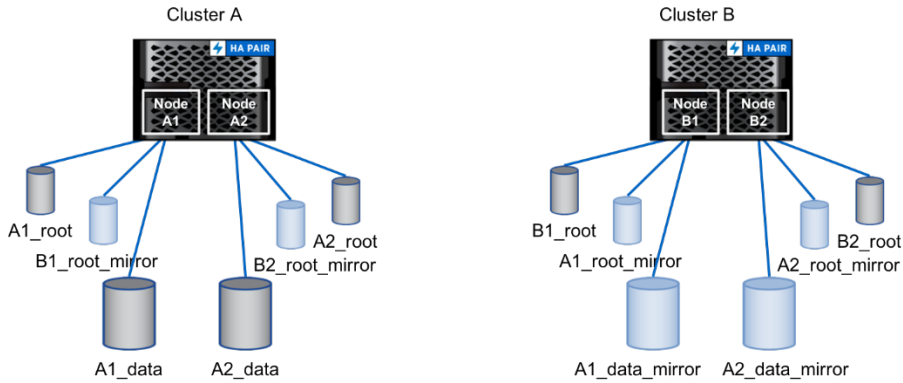
There are two types of active passive configurations:

- Active-passive clusters
- Active-passive HA in a cluster

An active-passive cluster or site is when one site is used for production workloads, the active site, and the second site has minimal capacity that is used for failover. This enables smaller storage configurations where the active site contains the active storage and workloads. It is also cost efficient because you do not need disk capacity when the second site is used only for operations in a switch-over or site failure situation.

For an active-passive cluster configuration, one cluster has all pool0 disks and the other cluster has all pool1 disks. For an active-passive cluster or site a small data aggregate must be created to host the MDV for MetroCluster. Except for root volumes and a small data volume for the MDV, the passive site only contains mirror copies of the data.

Figure 15) Active-passive cluster or site.




An active-passive HA configuration is when the storage is allocated to one of nodes in an HA pair. This configuration is typically done to maximize capacity in smaller configurations.

Note: For AFF, you must have storage distributed equally between the nodes.

In this example, each node owns disks with root volumes. The local active data volume is hosted node 1 and the remote mirror copy is hosted on node 2.

Figure 16) Active-passive HA.



For more information about active-passive HA pairs, see [High-Availability Configuration Guide](#) in the ONTAP 9 Documentation Center.

4 Operation and Administration

Operation and administration for MetroCluster includes checking or validating MetroCluster health and monitoring. For most operations the ONTAP documentation provides the steps to administer storage that includes MetroCluster. For specifics about MetroCluster feature management and operations see the [MetroCluster Management and Disaster Recovery Guide](#).

4.1 Switchover and Takeover

With MetroCluster, there are two types of terms to describe when a controller is hosting data and interfaces for a failed node. The two types are used to describe local level recovery or cluster (site) level recovery. Local recovery is referred to as takeover and giveback. Cluster or site level recovery provided by MetroCluster is referred to as switchover and switchback.

Local failures and nondisruptive operations, such as ONTAP upgrades, are handled by the HA partner. MetroCluster uses the standard ONTAP terminology for HA operations.

HA: local level recovery

- Takeover
- Giveback

For more about HA, see the [High-Availability Configuration Guide](#).

MetroCluster: cluster-level recovery

- Switchover
 - Negotiated switchover (NSO)
 - Unplanned switchover (USO)
- Switchback

For more information about switchover and switchback, see the [MetroCluster Management and Disaster Recovery Guide](#).

Starting with ONTAP 9.6, Storage Manager provides a graphical interface to manage planned and unplanned switchover.

Automatic Healing

Automatic healing is a feature added in ONTAP 9.5 and enhanced in 9.6 that allows MetroCluster IP systems a simpler process for negotiated switchover and switchback. This is especially helpful for performing disaster recovery testing as it simplifies the process to switchback to the original site.

For more information about automatic healing, see the section for [Automatic Healing](#) in the MetroCluster Management and Disaster Recover Guide.

4.2 ONTAP Mediator Software

ONTAP 9.7 includes the release of a new ONTAP Mediator software solution for MetroCluster IP. The software resides in a third failure zone and enables MetroCluster IP to perform automated unplanned switchover (AUSO). Additional functionality includes the disabling of AUSO in the event that the two sites have a failure with mirroring data between them. This prevents automatic switchover if the inter-site links are down, allowing an administrator to decide if it is appropriate to switch over manually.

The new mediator is configured from one of the MetroCluster IP nodes at one site. ONTAP automatically performs the configuration for all of the nodes and for the second cluster. The initial release requires, MetroCluster IP running ONTAP 9.7 or later and the mediator software, release 1.0 or later.

New commands are provided in ONTAP for the configuration of the mediator.

```
metrocluster configuration-settings mediator add -mediator-address <mediator-ip>
metrocluster configuration-settings mediator remove
```

When adding a new mediator, the command prompts for valid mediator credentials. The credentials are set during the installation. They can also be changed by logging into the mediator system and using the following commands:

- Change the mediator account name:

```
/opt/netapp/lib/ontap_mediator/tools/mediator_change_user
```

- Change the mediator password:

```
/opt/netapp/lib/ontap_mediator/tools/mediator_change_password
```

- Check the mediator status:

```
systemctl status ontap_mediator
```

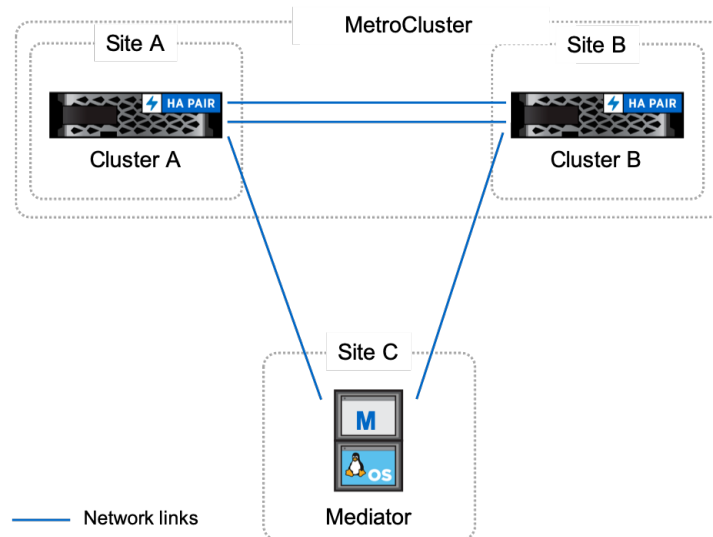
- Run the log collection tool:

```
/opt/netapp/lib/ontap_mediator/tools/mediator_generate_support_bundle
```

For more information about the requirements for ONTAP Mediator and details about failures, see the [MetroCluster IP Installation and Configuration Guide](#).

Note: Managing the same MetroCluster configuration with both Tiebreaker and ONTAP Mediator is not supported. Only one of the products can be used to manage a MetroCluster configuration.

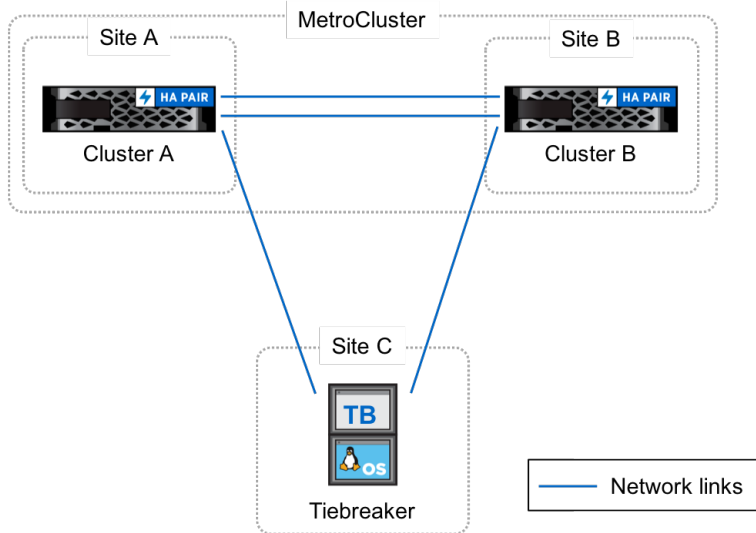
Figure 17) MetroCluster Mediator site.



4.3 TieBreaker Software

MetroCluster TieBreaker management pack provides monitoring of MetroCluster systems and provides the ability to detect site disasters and ISL failures. Tiebreaker software is installed on a Linux host, typically on a virtual machine, that is located in a third failure domain separate from the failure domains of either cluster in the MetroCluster solution.

Figure 18) MetroCluster tiebreaker site.



Tiebreaker software monitors each controller by establishing redundant connections through multiple paths to a node management LIF and to the cluster management LIF.

Note: Managing the same MetroCluster configuration with both Tiebreaker and ONTAP Mediator is not supported. Only one of the products can be used to manage a MetroCluster configuration.

Site Failure Symptoms

During a site failure, when one cluster is unreachable from the Tiebreaker software and the other cluster is reachable, the cluster that is reachable must also indicate it has lost communication with the partner cluster before Tiebreaker software triggers an alert. If the two clusters can still communicate, Tiebreaker identifies the loss of connectivity in the network between the Tiebreaker software and the cluster that is not reachable.

Figure 19) Tiebreaker site link failure.


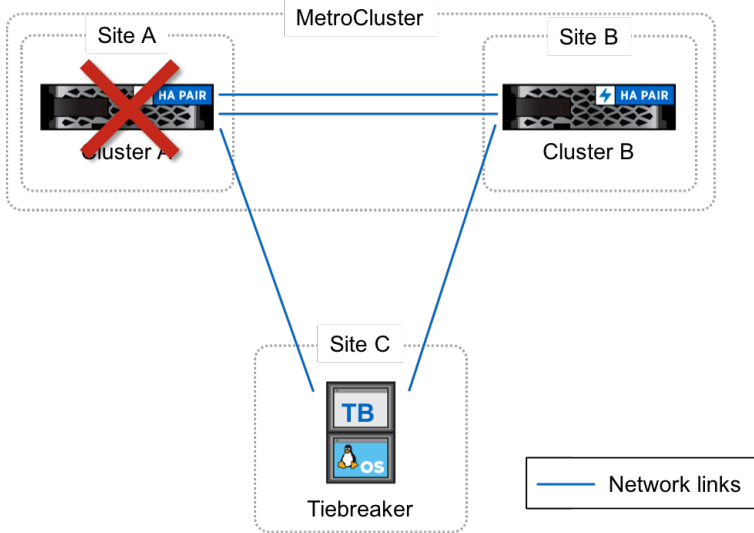


Figure 20) Tiebreaker site failure.



For more information, see the [Tiebreaker Software Installation and Configuration Guide](#).

5 Interoperability

MetroCluster supports the most common ONTAP features in the initial releases. However, some ONTAP features, such as FabricPool, are not currently supported by MetroCluster. You can see the ONTAP documentation of the feature for guidance on support for MetroCluster. For example, the [Considerations and requirements for using FabricPool](#) section in the ONTAP Documentation Center.

If information is not available in the documentation center or the [Interoperability Matrix](#), see the current interoperability section for more information and support of specific ONTAP features with MetroCluster in the *Technical FAQ: MetroCluster IP* on [Lenovo Support Website](#).

5.1 SnapMirror

SnapMirror asynchronous is supported with MetroCluster and allows additional protection and backup of data by asynchronously replicating volumes to a third cluster. MetroCluster systems are configured as a source for SnapMirror replication relationships.

5.2 Lenovo FlexGroup Volumes

Lenovo FlexGroup volumes are supported with MetroCluster starting with ONTAP 9.6. A FlexGroup volume is a scale-out NAS container that provides high performance along with automatic load distribution and scalability. For more information, see the [Scalability and Performance Using FlexGroup Volumes Power Guide](#) in the ONTAP documentation.

5.3 Lenovo FlexCache

Lenovo FlexCache® technology is supported with MetroCluster IP starting with ONTAP 9.7. See the following resources for more information.

- [FlexCache Volumes for Faster Data Access Power Guide](#)

5.4 Lenovo FabricPool

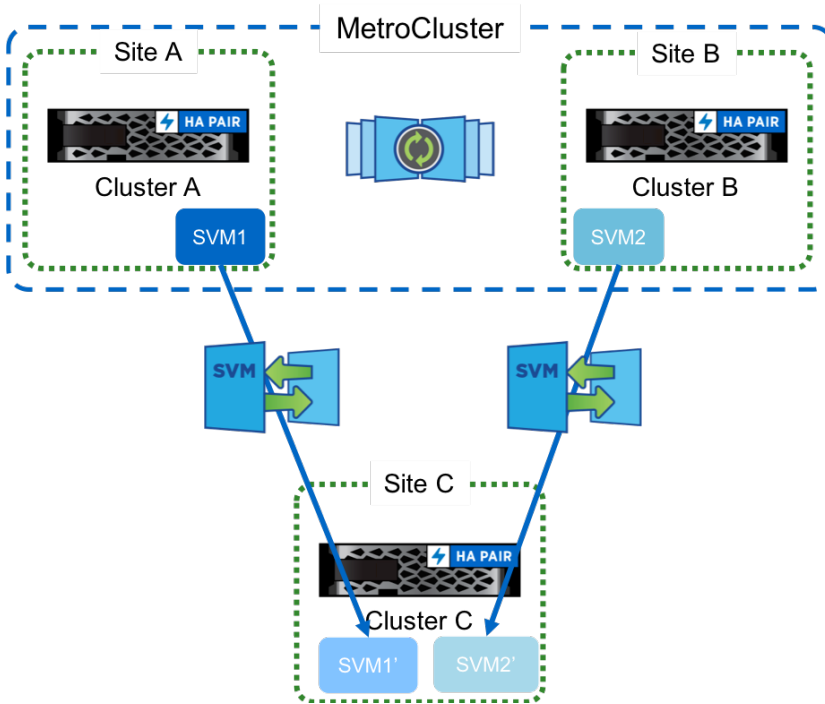
Lenovo FabricPool is supported with MetroCluster IP starting with ONTAP 9.7. See the following resources for more information:

- [Managing storage tiers by using FabricPool](#)

5.5 SVM Mirror and SVM Disaster Recovery

SVM mirror, also referred to as SVM disaster recovery is supported with MetroCluster IP starting in ONTAP 9.5. SVM mirror uses asynchronous replication to replicate data associated with a specific SVM, including data volumes and configuration.

Figure 21) SVM disaster recovery.



Where to Find Additional Information

To learn more about the information described in this document, see the following documents and/or websites:

- MetroCluster IP Installation and Configuration Guide
http://thinksystem.lenovofiles.com/help/topic/MCC-MetroCluster_IP_Installation_and_Configuration_Guide/overview.html
- MetroCluster Management and Disaster Recovery Guide
http://thinksystem.lenovofiles.com/help/topic/MCC-MetroCluster_Management_and_Disaster_Recovery_Guide/overview.html
- ONTAP 9 Documentation Center
<http://thinksystem.lenovofiles.com/help/index.jsp>
- Lenovo Interoperability Matrix
Go to [Lenovo Support Website](#), and select your product and machine type. On the product page, click the **Knowledge Base & Guides** tab, and look for the *Interoperability Matrix* document for your product.
- MetroCluster Resources page
https://thinksystem.lenovofiles.com/help/nav/7_1_0_6
- VMware vSphere 5.x and 6.x support with Lenovo MetroCluster (2031038)
<https://kb.vmware.com/s/article/2031038>

For more information about supported optical modules and part numbers to order from a Cisco partner, see the Cisco optical module support matrix for the specific model of switch.

- Cisco module and switch support matrix
<https://tmgmatrix.cisco.com/home>
- Cisco CWDM SFP 10 Gigabit Ethernet Solution Datasheet
<https://www.cisco.com/c/en/us/products/collateral/interfaces-modules/transceiver-modules/datasheet-c78-734047.html>
- Cisco 10GBASE SFP+ Modules Datasheet
https://www.cisco.com/c/en/us/products/collateral/interfaces-modules/transceiver-modules/data_sheet_c78-455693.html
- Cisco 10GBASE Dense Wavelength-Division Multiplexing SFP+ Modules Datasheet
https://www.cisco.com/c/en/us/products/collateral/interfaces-modules/dwdm-transceiver-modules/data_sheet_c78-711186.html

Contacting Support

You can contact Support to obtain help for your issue.

You can receive hardware service through a Lenovo Authorized Service Provider. To locate a service provider authorized by Lenovo to provide warranty service, go to <https://datacentersupport.lenovo.com/serviceprovider> and use filter searching for different countries. For Lenovo support telephone numbers, see <https://datacentersupport.lenovo.com/supportphonenumber> for your region support details.

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area.

Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document is not an offer and does not provide a license under any patents or patent applications. You can send inquiries in writing to the following:

*Lenovo (United
States), Inc. 8001
Development Drive*

*Morrisville, NC 27560
U.S.A.*

Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT,

MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Trademarks

LENOVO, LENOVO logo, and THINKSYSTEM are trademarks of Lenovo. All other trademarks are the property of their respective owners. © 2019 Lenovo.

