Technical Report

vSphere 5 on NetApp MetroCluster Solution

Santhosh Devaraju, Ashish Nainwal, NetApp
February 2013 | TR-4128

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Introduction

This technical report provides an insight into the NetApp® MetroCluster™ solution for vSphere® 5.0 or later. It provides advanced details for implementing NetApp's high-availability solution, MetroCluster, with VMware's industry-leading virtual infrastructure solution vSphere. Common production-specific use cases are included to provide a deep dive into the implementation procedures.

1.1 Intended Audience

This document is intended for:

- Customers and partners looking to implement a continuous-availability solution for their virtual infrastructure consisting of VMware® vSphere 5.0 or later and NetApp FAS storage.
- End users and managers interested in continuous-availability solutions in a production or dev/test environment.

1.2 Scope

The purpose of this report is to provide:

- An end-to-end architecture overview of the continuous-availability solution for virtual infrastructures
- A detailed design and implementation guide, and configuration best practices
- A high-level overview of architectural use cases with expected behaviors

The scope of this document is limited to the following.

- This report does not replace any official manuals or documents from NetApp and VMware on the products used in the solution or those from any other switch vendors referenced in this report.
- This report does not discuss any performance impact or analysis from an end-user perspective during a disaster.
- This report does not replace NetApp and VMware professional services documents or services.
- This report does not discuss a regional (long-distance) disaster recovery solution. If you are looking for a regional disaster recovery solution in addition to the high-availability option discussed in this paper, contact your NetApp representative for further assistance.

1.3 Assumptions and Prerequisites

This document assumes familiarity with the following:

- Basic knowledge of VMware's virtualization technologies and products:
  - VMware vCenter™ 5. 0 or later
  - VMware vSphere 5. 0 or later
- Basic knowledge of NetApp storage systems and the NetApp Data ONTAP® operating system

2 Background

2.1 Business Challenge

Ever-evolving business challenges and exponential growth put continuous pressure on availability and business continuity. As more business-critical solutions are hosted in virtualized data centers, there is increased emphasis on improving the robustness of the infrastructure. Such an improvement enables businesses to reap the economic and operational benefits of virtualization without compromising on availability or quality of service. Planning a robust high-availability infrastructure solution for virtual data center environments hosting mission-critical applications is of the utmost importance.

An efficient high-availability virtualized infrastructure is an easy-to-deploy and -manage solution that provides high standards of resiliency without compromising on either performance or cost effectiveness.

VMware, with its high-availability and fault-tolerant features, provides uniform failover protection against hardware and software failures within a virtualized IT environment.

NetApp MetroCluster is a highly cost-effective, synchronous replication solution for combining high availability and disaster recovery in a campus or metropolitan area to protect against both site disasters and hardware outages. NetApp MetroCluster provides automatic recovery for any single storage component failure and a highly efficient single-command recovery in case of major site disasters. It provides solutions with zero data loss and recovery within minutes rather than hours, hence, improved RPO and RTO.

Combining VMware high availability and fault tolerance with NetApp MetroCluster technologies offers a great value proposition for business-critical applications. The combination provides an architecturally simple and highly robust continuous-availability solution for both planned and unplanned downtimes in virtual data center environments.

Each of these solutions is discussed briefly in section 2.2, “Continuous Availability Solutions for vSphere Environments.”

2.2 Continuous Availability Solutions for vSphere Environments

The NetApp Unified Storage Architecture offers an agile and scalable storage platform. All NetApp storage systems use the Data ONTAP operating system to provide SAN (FC/FCoE, iSCSI) and NFS.

NetApp MetroCluster leverages NetApp’s HA (controller failover, a.k.a. CFO) functionality to automatically protect against controller failures. Additionally, MetroCluster layers local SyncMirror® technology, cluster failover on disaster (controller failover on demand, a.k.a. CFOD), hardware redundancy, and geographical separation to achieve extreme levels of availability. Local SyncMirror synchronously mirrors data across the two halves of the MetroCluster configuration by writing data to two plexes: the local plex (on the local shelf) actively serving data and the remote plex (on the remote shelf) normally not serving data. On local shelf failure, the remote shelf seamlessly takes over data-serving operations. No data loss occurs because of synchronous mirroring. Hardware redundancy is put in place for all MetroCluster components. Controllers, storage, cables, switches (used with fabric MetroCluster), and adapters are all redundant.

A VMware HA/DRS cluster is created across the two sites using ESXi™ 5.0 or 5.1 hosts and managed by vCenter Server 5.0 or 5.1. The vSphere management, vMotion®, and virtual machine networks are connected using a redundant network between the two sites, and the vCenter Server managing the HA/DRS cluster can connect to the ESXi hosts at both sites.

What Is vSphere Metro Storage Cluster?

vSphere Metro Storage Cluster (vMSC) is a new certified configuration for NetApp MetroCluster storage architectures. The vMSC configuration is designed to maintain data availability beyond a single physical or logical site. A storage device configured in the vMSC configuration is supported after successful vMSC certification. All supported storage devices are listed on the VMware Storage Compatibility Guide.

For more information on the design guidelines for vSphere Metro Storage Cluster, refer to the following documentation.

- vSphere 5.x Support with NetApp MetroCluster
- VMware vSphere Metro Storage Cluster - Case Study

Based on the distance considerations, NetApp MetroCluster for vSphere can be deployed in two different configurations:

- Stretch MetroCluster

- Fabric MetroCluster

Figure 1 illustrates high-level topology diagram of stretch MetroCluster.

Figure 1) High-level topology diagram of stretch MetroCluster.

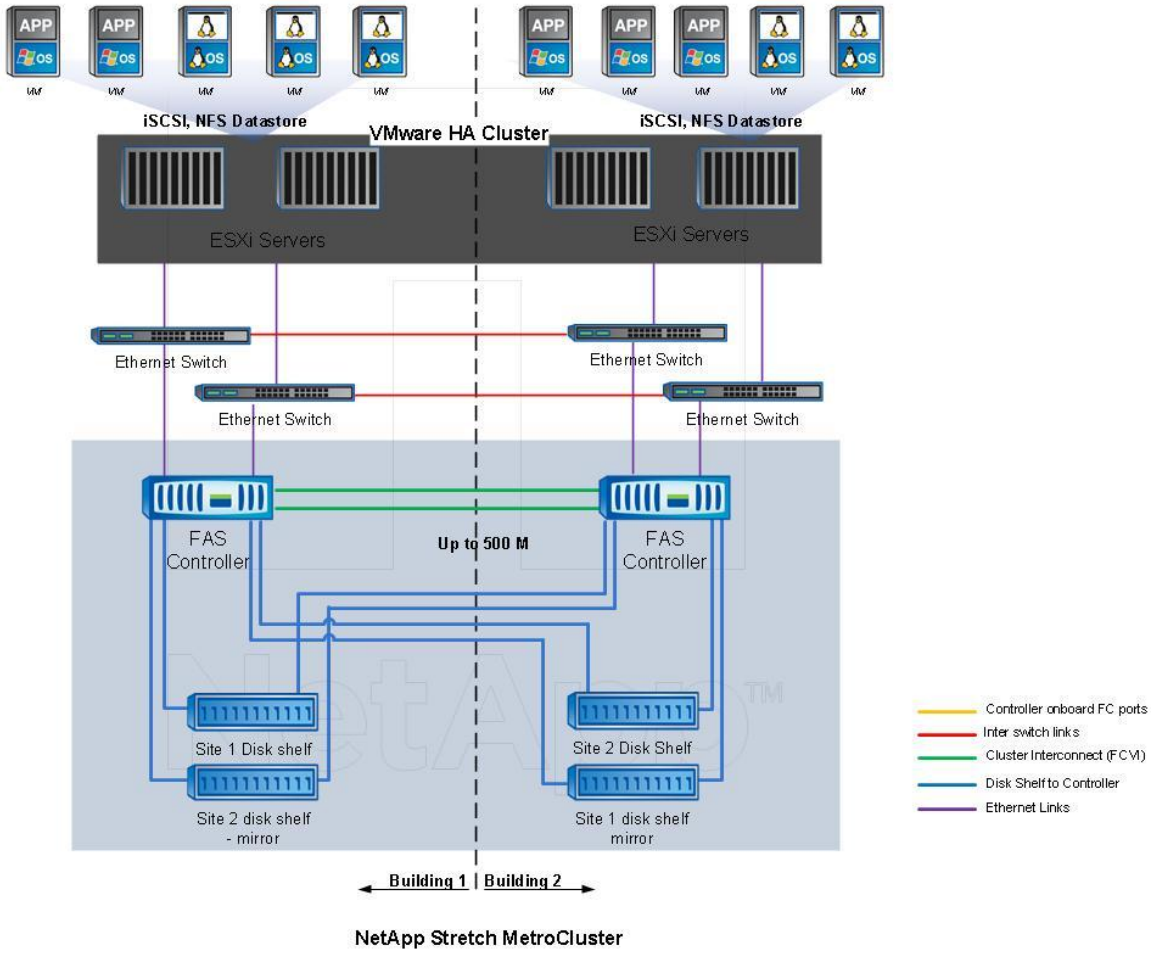

Figure 2 illustrates high-level topology diagram of fabric MetroCluster.

Figure 2) High-level topology diagram of fabric MetroCluster.

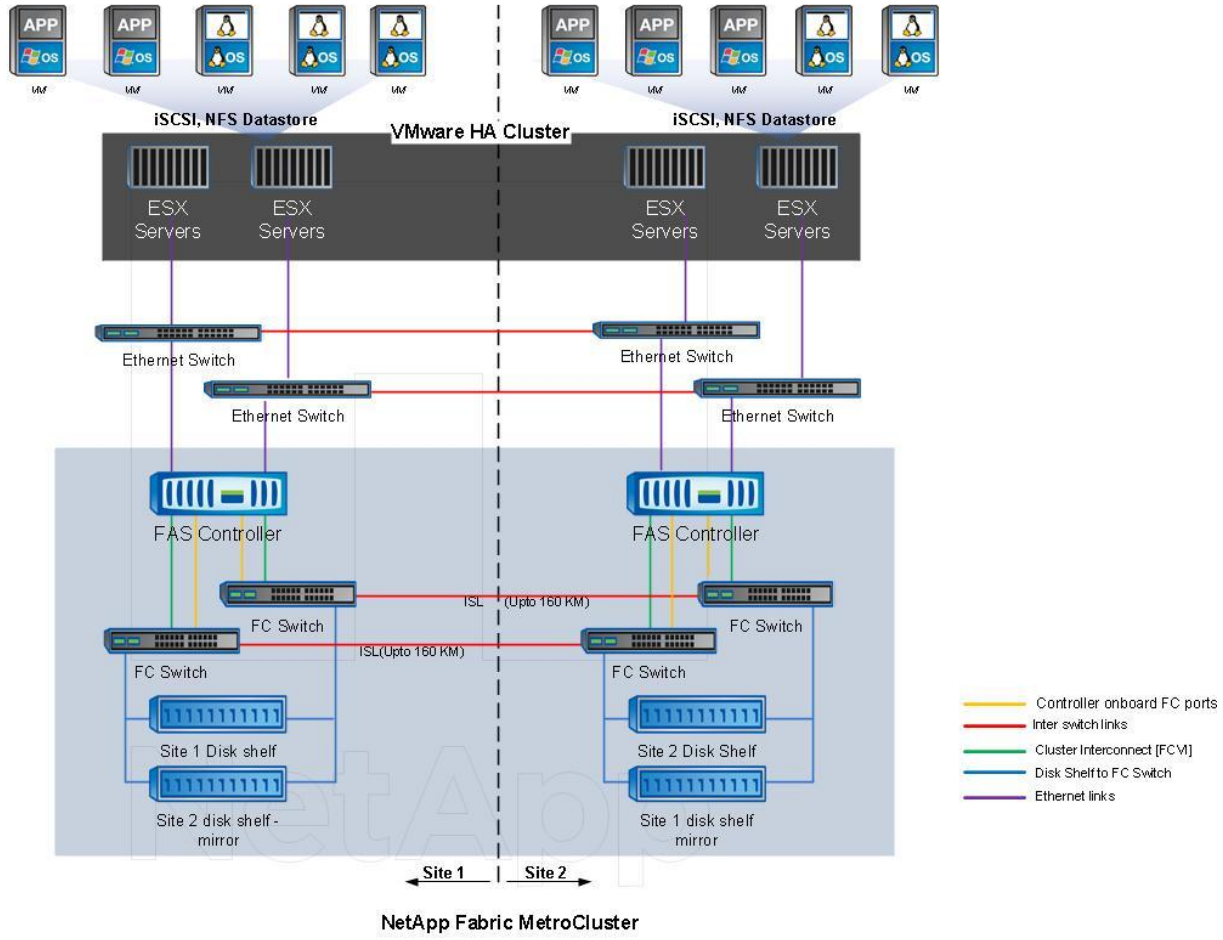

3 Introduction to NetApp MetroCluster

MetroCluster is an integrated, high-availability, business-continuance solution designed for large campuses and metropolitan areas. In the event of a disaster affecting a data center, MetroCluster enables the customer to quickly and easily fail over the mission-critical operation to another data center at a remote site. The operation is resumed without data loss and with minimal administrator intervention. The takeover process only takes minutes and imposes no additional disruption to users. The customer can achieve extremely fast recovery time compared to that of competing disaster recovery products, which often take hours or even days.

MetroCluster consists of the following software components.

- **Active-Active Controller.** Provides high-availability failover capability between the appliances at the local and remote sites.
- **SyncMirror.** Provides an up-to-date copy of data at the remote site. Data can be accessed only by the remote storage after failover.
- **Cluster Remote.** Provides a mechanism with which the administrator can declare a site disaster and initiate site failover to the remote site.

In addition, if the primary and remote sites are more than 500 meters apart, a pair of dedicated FC switches is required at each site to enable long-distance connectivity between the sites.

3.1 MetroCluster Components

A MetroCluster configuration includes the hardware and software components listed in Table 1 and Table 2.

Table 1) MetroCluster hardware components.

S.No.	Hardware Components
1.	Standard FAS HA pair of controllers running a compatible version of Data ONTAP.
2.	Four Fibre Channel (FC) switches with supported firmware supplied by NetApp, a pair at each location. (Fabric-attached MetroCluster configurations only.) Supported models might differ between fabrics, but must be the same within each fabric. These switches need to be dedicated to MetroCluster and cannot be shared by components other than MetroCluster. In addition, these switches are necessary; existing switch infrastructures at customer sites cannot be leveraged instead.
3.	Fibre Channel-Virtual interface (FC-VI) adapter for the cluster interconnect, except for the stretch MetroCluster configuration with an FC-VI cable distance of less than 30m. This is mandatory for FAS31xx, 32xx, and 62xx models.
4.	Copper/fiber converters for cluster interconnect. For FAS9xx, 30xx, and 60xx only.
5.	Associated cabling and SFP connectors.
6.	Minimum four FC initiator ports (storage adapters).
7.	Additional disk shelves to accommodate the mirrored data.
8.	If using SAS shelves, it requires FibreBridges 6500N, two per stack of SAS shelves.
9.	Fabric-attached MetroCluster configuration requires dedicated dark fiber links between the sites, and it might contain switch vendor–qualified xWDM.

Note: Refer to the NetApp MetroCluster Compatibility Matrix site to find out the latest supported models, specific firmware /fabric OS, and Data ONTAP version.

Table 2 lists the MetroCluster software components.

Table 2) MetroCluster software components.

S.No.	Software License to Be Enabled on Each Controller / Switch		
	For FAS Controller	For Brocade Switch	For Cisco Switch
1.	`cluster` - for controller failover functionality	Brocade extended distance license - for intersite distances >10km	Cisco® ENTERPRISE_PKG - to maximize buffer-to-buffer credits

2.	`syncmirror_local`- for synchronous mirroring across sites	Brocade ports-on-demand (POD) - to scale switch with additional ports	Cisco PORT_ACTIVATION_PKG - to scale switch with additional ports
3.	`cluster_remote` - for the site failover on disaster functionality		Cisco FM_SERVER_PKG – (optional) for enabling the use of the Cisco Fabric Manager GUI

3.2 Types of MetroCluster

MetroCluster can be configured in the following two ways.

Stretch MetroCluster (SMC) Topology

A stretch MetroCluster configuration is a direct connection of two or more mirrored HA pairs over limited distances (campus-sized areas). This configuration provides storage redundancy and disaster recovery across distances of up to 500 meters. Stretch MetroCluster provides data mirroring and the additional capability to initiate a failover if an entire site becomes lost or unavailable.

Similar to the mirrored active-active configurations, stretch MetroCluster contains two complete copies of the specified data volumes or file systems. These copies are called plexes, and every time Data ONTAP writes data to the disks, the copies are continually and synchronously updated.

Unlike mirrored active-active configurations, MetroCluster provides the capability to force a failover when an entire node (including the controllers and storage) is destroyed or unavailable.

Figure 3 illustrates the stretch MetroCluster topology.

Figure 3) Stretch MetroCluster topology.

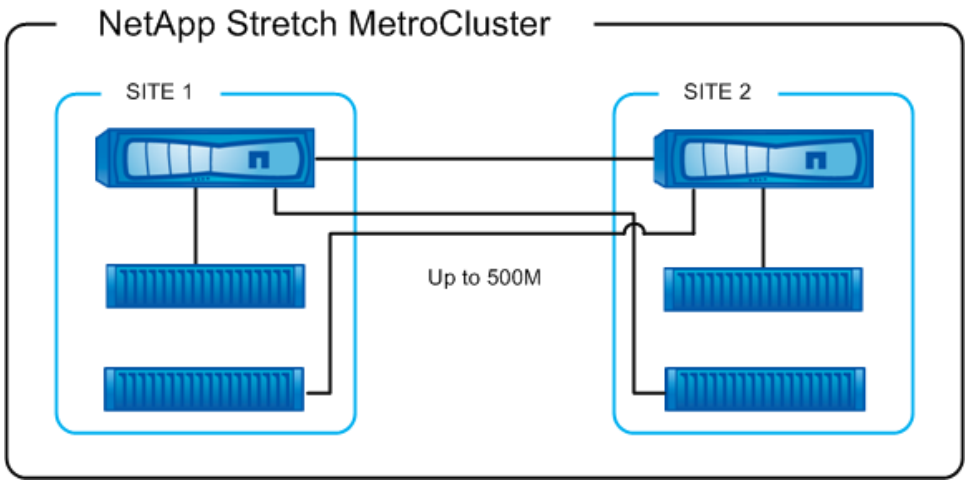

Note: Distance is the actual cable length connecting the primary and remote sites. Do not use surface distance, which is always significantly shorter than the actual cable distance.

Stretch MetroCluster Configuration

Figure 4 illustrates a typical stretch MetroCluster configuration with a FAS32XX single controller on each site with DS14 MK4FC shelves.

Figure 4) Stretch MetroCluster configuration.

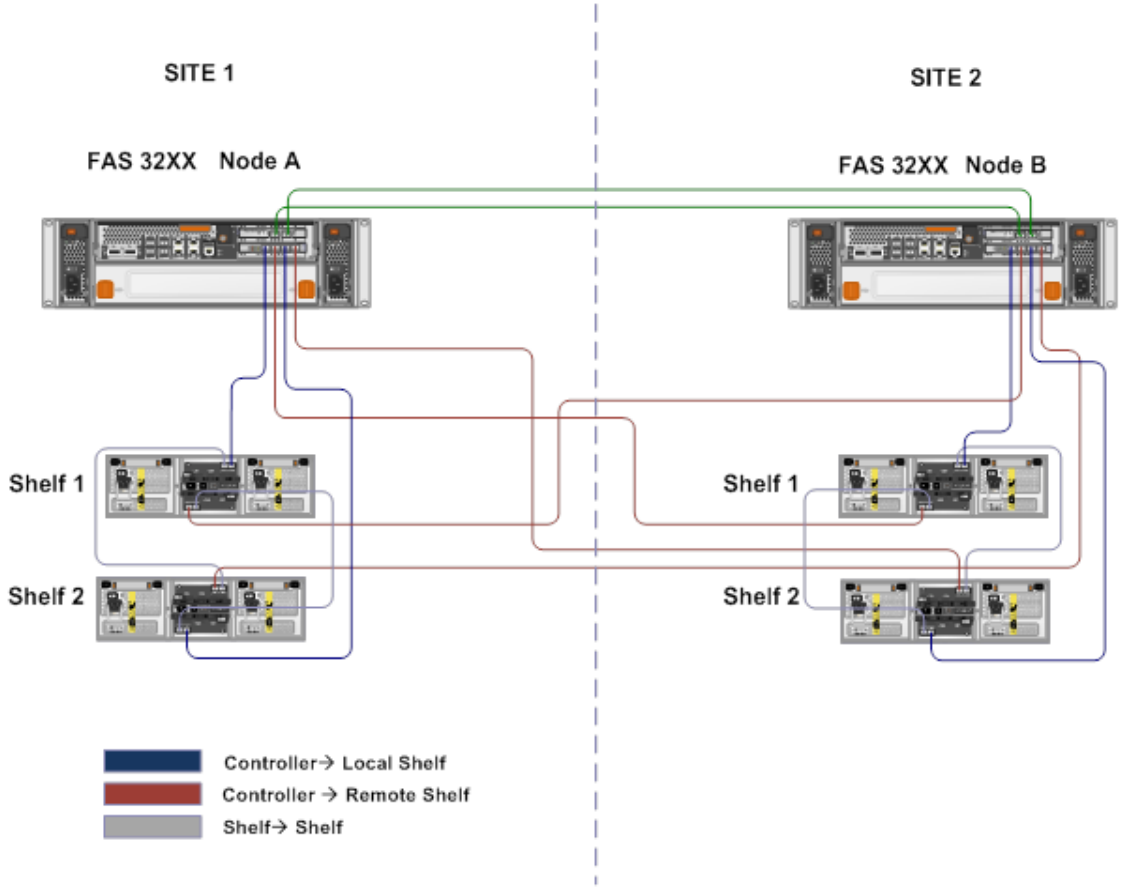

The reference stretch MetroCluster configuration includes:

- Standard HA pair of FAS32XX controllers, one in site 1 and the other in site 2 within a campus
- A pair of DS14 MK4 shelves in site 1 and site 2

Note: For detailed stretch MetroCluster configurations, refer to Chapter 6 Stretch MetroCluster Considerations.

Stretch MetroCluster Dual and Twin Configurations

Stretch MetroCluster supports controllers in a dual MetroCluster configuration with a single-head chassis connected to another single-head chassis as well as a twin MetroCluster configuration similar to a dual MetroCluster configuration with head A at the local site connected to head A at the remote site and head B at the local site connected to head B at the remote site. Figure 3 and Figure 4 illustrate dual and twin stretch MetroCluster configurations and topologies, respectively.

Figure 5 illustrates a dual MetroCluster configuration.

Figure 5) Dual MetroCluster configuration.

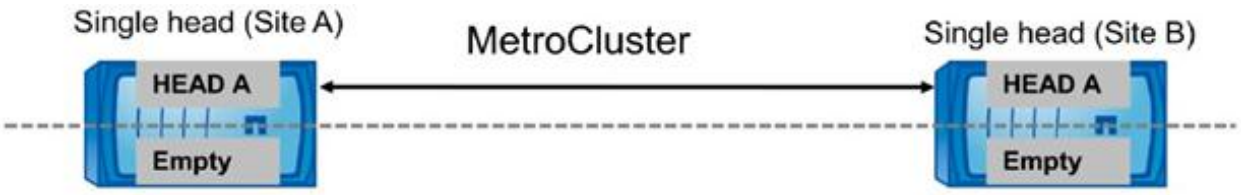

Figure 6 illustrates a twin MetroCluster configuration.

Figure 6) Twin MetroCluster configuration.

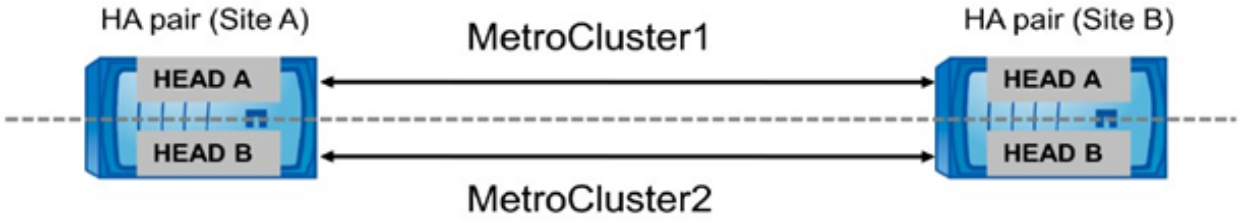

Fabric-Attached MetroCluster (FMC) Topology

A fabric-attached MetroCluster configuration is for distances greater than 500 meters (going up to 160 km) connecting the two nodes using four Brocade or Cisco Fibre Channel switches in a dual-fabric configuration for redundancy. A fabric-attached MetroCluster leverages SyncMirror to build a system that can continue to serve data even after complete loss of one of the nodes and the storage at that site. Data consistency is retained, even when the data is contained in more than one aggregate.

Figure 7 illustrates the fabric-attached MetroCluster topology.

Figure 7) Fabric-attached MetroCluster topology.

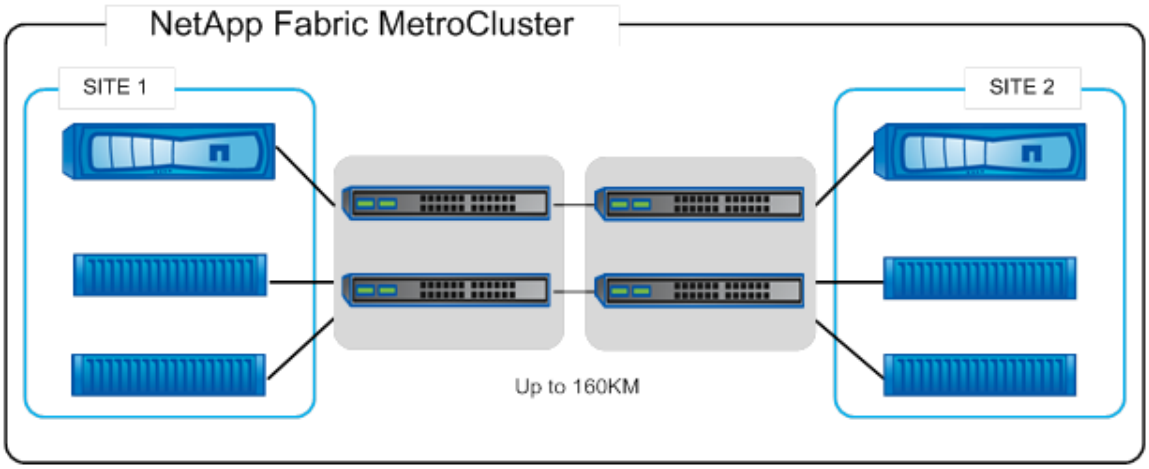

Note: Distance is the actual cable length connecting the primary and remote sites. Do not use surface distance, which is always significantly shorter than the actual cable distance.

Fabric-Attached MetroCluster Configuration

Figure 8 illustrates an example of a typical fabric-attached MetroCluster configuration with a FAS62XX single controller on each site with DS14 MK4 FC shelves and Brocade 5100 switches.

Figure 8) Fabric-attached MetroCluster configuration.

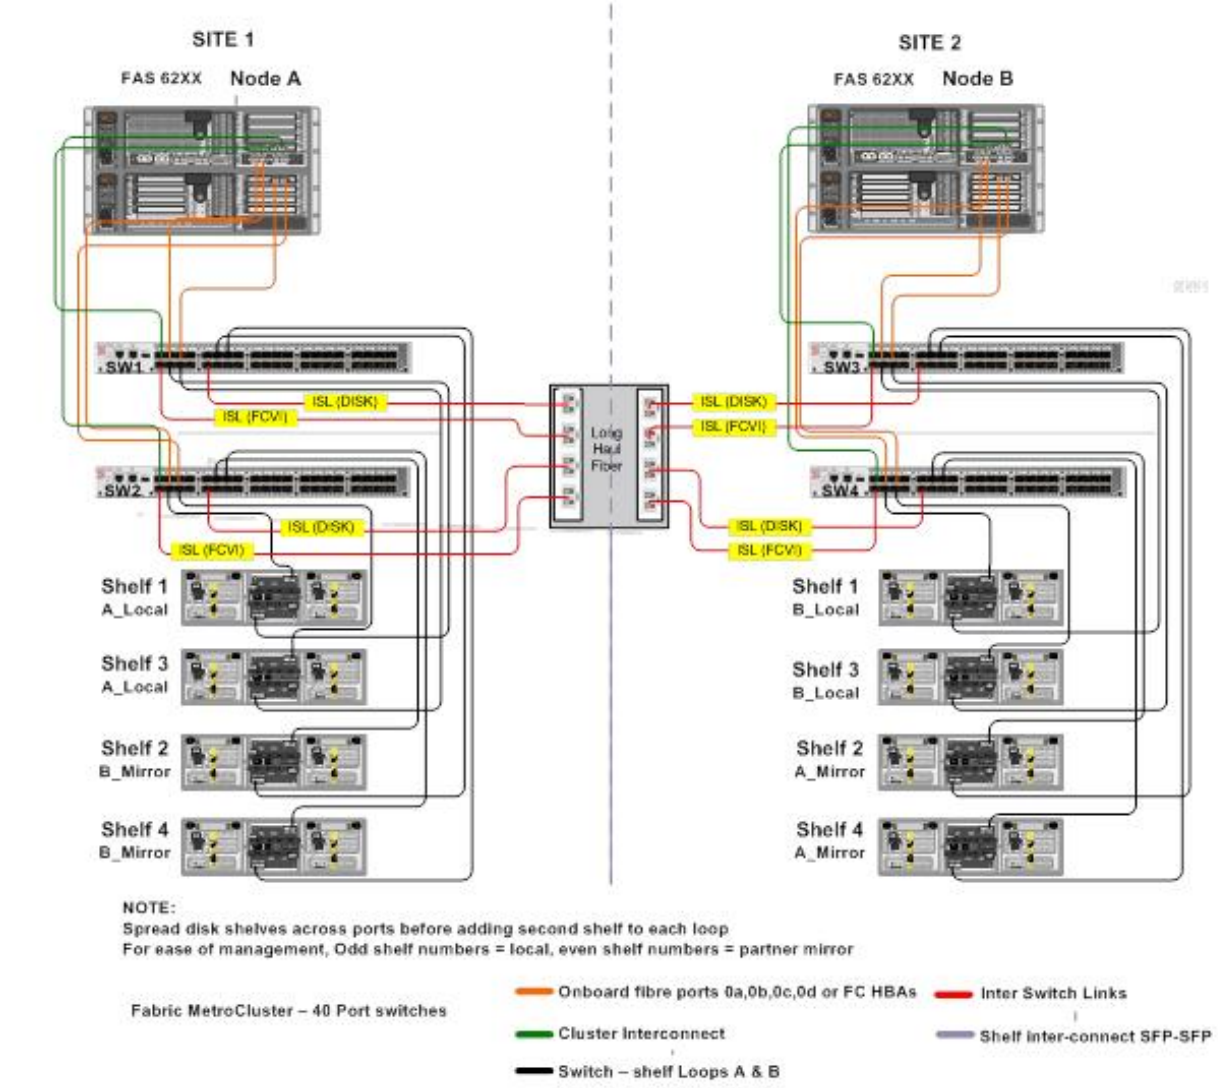

The reference fabric-attached MetroCluster configuration includes:

- Standard HA pair of FAS62XX controllers, one in site 1 and the other in site 2
- Two pairs of DS14 MK4 shelves in site 1 and site 2

4 vSphere Solution Overview

VMware vCenter is a centralized management tool for ESX® clusters that helps administrators perform core functions such as VM provisioning, vMotion operation, DRS, and so on. It also plays a vital role in VMware View™ vCloud® environments. The VMware virtual infrastructure should be designed considering service availability.

4.1 vCenter Availability

Availability designs are based on the uptime requirements of the service so that the guaranteed service level is achieved.

Option 1: Protecting a vCenter Virtual Machine in an HA Cluster

In this option, vCenter is installed on a virtual machine as part of an HA cluster. Upon host failure, the vCenter virtual machine is restarted on the surviving ESX host in the HA cluster. In this availability solution, the vCenter virtual machine services will be disrupted because the virtual machine will be restarted.

Figure 9 illustrates the VMware HA solution.

Figure 9) VMware HA solution.

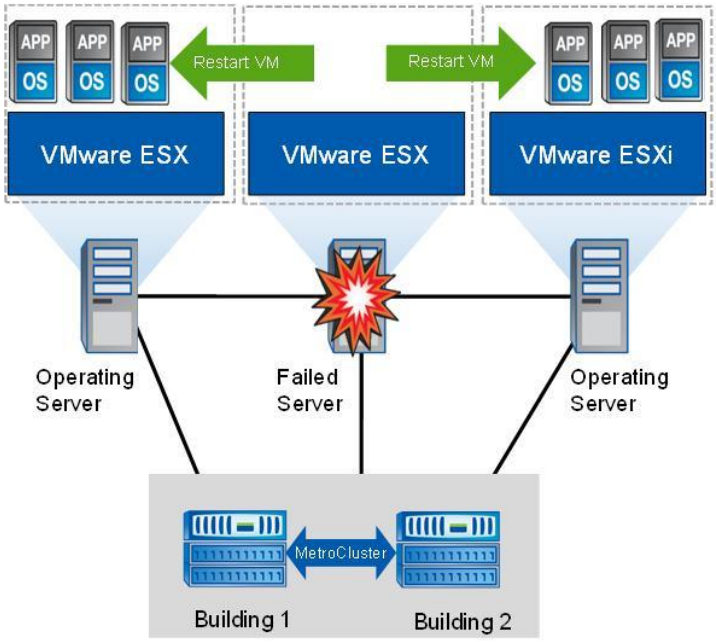

Option 2: Roll Out Your Own Clustering Configuration

Another way of designing the vCenter Server is to place it in a physical cluster using a third-party clustering solution at each site. If the storage that houses the vCenter cluster instance is at the failed site, it is necessary to perform the NetApp CFOD recovery.

For details on Microsoft Cluster Service (MSCS), refer to VMware KB.

For details on Oracle® RAC, refer to Oracle Databases on VMware RAC Deployment Guide.

Option 3: Protecting vCenter Using the VMware vCenter Heartbeat Solution

Because the VMware vCenter Server is used to manage many tier 1 applications, it renders itself as a tier 1 application. Therefore, it becomes very important for the VMware vCenter Server to be highly available. This is where VMware vCenter Server Heartbeat steps in.

VMware vCenter Server Heartbeat delivers high availability for the VMware vCenter management platform for consistent operation of the VMware vSphere environment. Architecturally, vCenter Server Heartbeat is implemented on the active-passive vCenter Server clones running on physical or virtual machines. In addition to server and network hardware, vCenter Server Heartbeat monitors the actual vCenter Server instance, its back-end database, and the underlying operating system. In case of failure, the passive node takes over and the vCenter Server Heartbeat software restarts the vCenter service. Failover can occur on both LANs and WANs.

To know more about the installation and configuration of VMware vCenter Server Heartbeat, refer to VMware vCenter Server Heartbeat.

- Protects VMware vCenter Server availability by monitoring all components of VMware vCenter Server, including VMware License Server and other plug-ins
- Minimizes downtime of critical functions such as VMware vMotion and VMware DRS
- Protects VMware vCenter Server performance, alerts, and events information, keeping it up to date even if the VMware vCenter Server experiences an outage
- Provides automatic failover and failback of VMware vCenter Server
- Enables administrators to schedule maintenance windows and maintains availability by initiating a manual switchover to the standby server
- Protects and recovers the VMware vCenter Server database
- Protects critical configuration, inventory, and other information stored in the VMware vCenter Server database, even if the database is installed on a separate server

Figure 10 illustrates the vCenter Server Heartbeat configuration used for this solution. The primary and the secondary vCenter Server VMs are deployed on separate ESX servers in separate sites and use separate datastores from the local NetApp storage controller in the MetroCluster configuration. The vCenter Server Heartbeat channel is configured over the LAN across the sites.

Figure 10) vCenter Heartbeat solution.

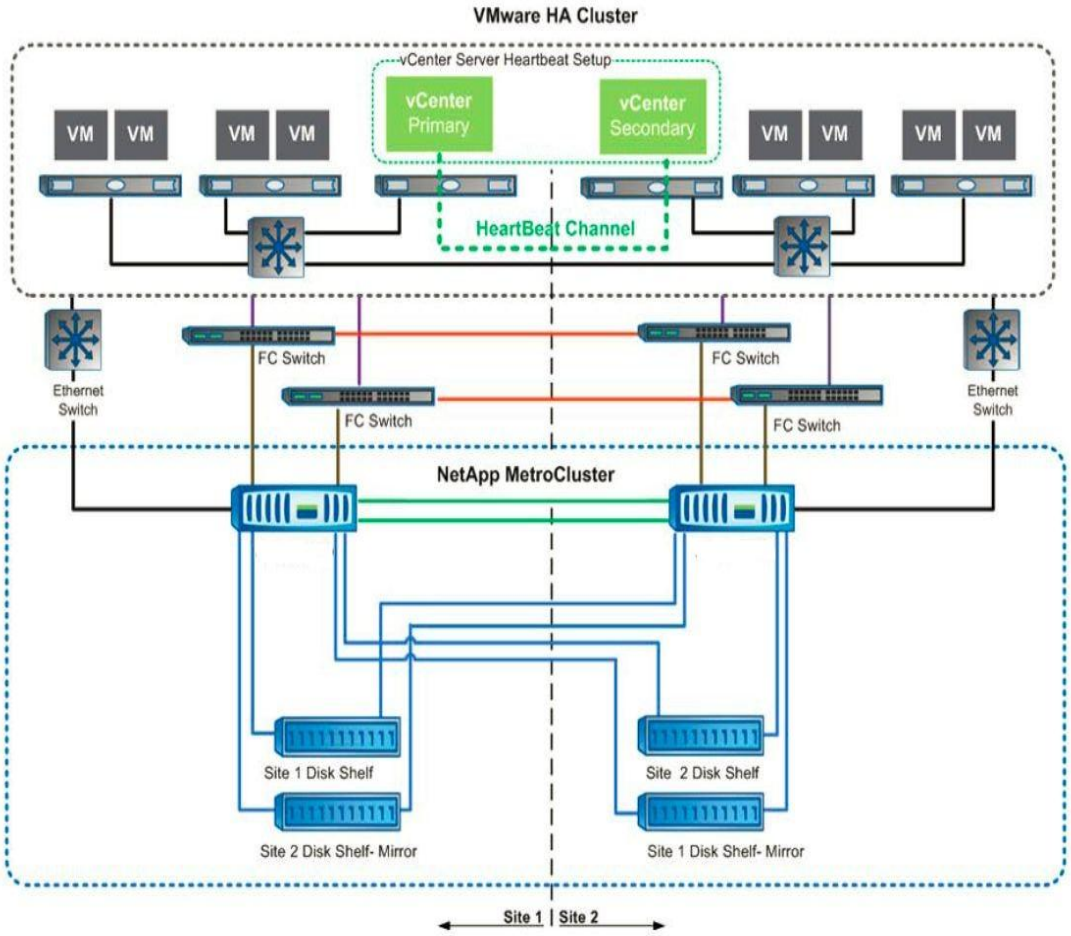

4.2 vSphere HA Implementation for NetApp MetroCluster

Overview

vSphere HA provides high availability for virtual machines by grouping the virtual machines and the hosts they reside on into a cluster. Hosts in the cluster are monitored and, in the event of a failure, the virtual machines on a failed host are restarted on surviving hosts.

In the NetApp MetroCluster solution for vSphere, the HA solution takes care of virtual machine failover in the case of compute failure in a site by restarting the virtual machines on the other surviving ESX hosts, which are located in the surviving site.

As soon as the HA cluster is created, all hosts in the cluster participate in election, and one of the hosts becomes a master. Each slave performs network heartbeat to the master, and the master in turn performs network heartbeat on all the slaves. It is the master's responsibility to determine if a slave has failed and to restart and place virtual machines on surviving hosts. Slave hosts monitor the state of their virtual machines and send updates to the master about any state changes. In addition, they monitor the health of the master by monitoring heartbeats. If the master becomes unavailable, the slaves initiate and participate in the election process.

HA architecture in vSphere 5.0 has changed substantially, and now there are different mechanisms to detect host failure or isolation; in earlier versions network heartbeat was the only mechanism available to detect failures. With network heartbeat, VMware reengineered HA with one more mechanism: datastore heartbeating. Datastore heartbeating is used by the master host in cases where it cannot communicate with slave hosts over the management network. With datastore heartbeating the master determines whether a slave host has failed, is in a network partition, or is in an isolated network.

For more information on vSphere HA, refer to the VMware document vSphere Availability.

Best Practice

In addition to using the network and heartbeat mechanism, NetApp recommends adding Isolation IP addresses in advanced settings of vSphere HA. These will enhance the reliability of isolation validation.

4.3 VMware DRS Implementation for NetApp MetroCluster

VMware DRS is a feature that aggregates the host resources in a cluster and is primarily used to load balance within a cluster in a virtual infrastructure. VMware DRS primarily calculates the CPU and memory resources to perform load balancing in a cluster. Many features are available within VMware DRS that can be leveraged in the NetApp MetroCluster environment.

Using VM–host affinity roles in VMware DRS, one can have a logical separation between site A and site B, so that the VM runs on the host at the same site as the array that is configured as the primary read/write controller for a given datastore. Also, VM–host affinity rules enable virtual machines to stay local to the storage, which in turn ascertains the virtual machine connection in case of network failures between the sites.

4.4 VMware Storage DRS Implementation with NetApp MetroCluster

The VMware Storage DRS feature enables aggregation of datastores into a single unit and balances virtual machine disks when storage I/O control thresholds are exceeded.

Storage I/O control is enabled by default on Storage DRS–enabled DRS clusters. Storage I/O control allows an administrator to control the amount of storage I/O that is allocated to virtual machines during periods of I/O congestion, which enables more important virtual machines to have preference over less important virtual machines for I/O resource allocation.

Storage DRS uses Storage vMotion to migrate the virtual machines to different datastores within a datastore cluster. In a NetApp MetroCluster environment, a virtual machine migration needs to be controlled within the datastores of that site. For example, virtual machine A running on the host at site A should ideally migrate within the datastores of the storage controller at site A. If it fails to do so, the virtual machine will continue to operate but with degraded performance, since the virtual disk read/write will be from a remote site B through intersite links.

NetApp recommends creating datastore clusters with respect to storage site affinity; that is, datastores with site affinity for site A should not be mixed with datastore clusters with datastores with site affinity for site B.

Best Practice

Whenever a virtual machine is newly provisioned or migrated using storage vMotion, NetApp highly recommends that all the VMware DRS rules specific to those virtual machines be manually updated accordingly. This will ascertain the virtual machine affinity at the site level for both host and datastore and thus reduce the network and storage overhead.

5 Design and Implementation Guidelines

5.1 NetApp Storage Configuration

Best Practice

Set the Data ONTAP configuration option `cf.takeover.change_fsid` to OFF. This option is supported on Data ONTAP version 7.2.4 and higher.

In the event of failure of the complete storage controller and/or all disk shelves (storage controller and associated local disk shelves), a manual failover of the MetroCluster system should be performed. If the `change_fsid` option is set to OFF on a NetApp FAS storage controller running Data ONTAP version 7.2.4 or higher, after performing a manual MetroCluster failover the UUIDs of the mirrored LUNs are retained and additional steps in the ESX Server side are not required to detect the VMFS volumes. After the VMFS volumes are detected, the VMs can be manually powered on.

On NetApp FAS storage controllers running Data ONTAP older than 7.2.4, after performing a manual MetroCluster failover the mirrored LUNs do not maintain the same LUN UUID as the original LUNs because this option is not available. When these LUNs house the VMFS-3 file system, the volumes are detected by ESX Server 3.x as being on Snapshot™ LUNs. Similarly, if a RAW LUN that is mapped as an RDM (Raw Device Mapping) is replicated or mirrored through MetroCluster, the metadata entry for the RDM must be recreated to map to the replicated or mirrored LUN. So that the ESX hosts have access to the VMFS volumes on the mirrored LUNs, see VMware KB 1001783.

Figure 11 illustrates how to set the `cf.takeover.change_fsid` configuration to Off.

Figure 11) Set the `cf.takeover.change_fsid` configuration to OFF.

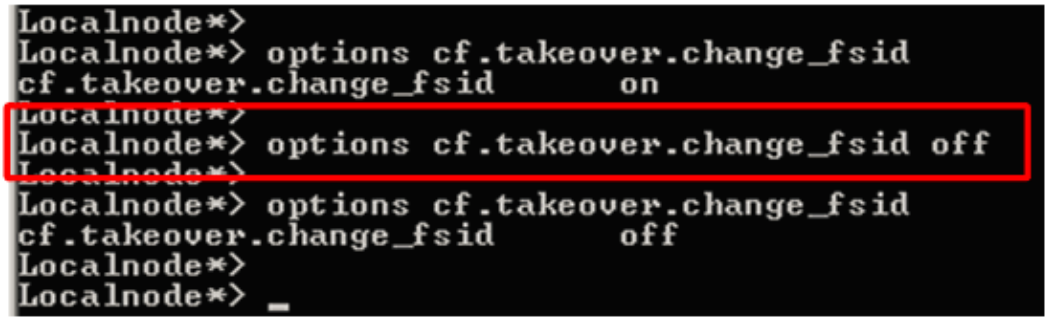

The FAS controllers should be licensed with the following features:

- cluster, cluster_remote, syncmirror_local
- iSCSI, FCP and/or NFS
- MetroCluster setup with software-based disk ownership for the NetApp FAS controllers with the Brocade switches is performed in accordance with the guidelines provided by:
  - High-Availability and MetroCluster Configuration Guide
  - MetroCluster Design and Implementation Guide
- In the FAS controllers on both sites, flexible volumes are created inside the same aggregate corresponding to two types of ESX datastores: VMFS (FC and iSCSI) and NFS.

Best Practice

Virtual Storage Console for VMware vSphere software is a single vCenter Server plug-in that provides end-to-end virtual machine lifecycle management for VMware environments running NetApp storage.

NetApp recommends using the NetApp Virtual Storage Console for vSphere to manage and provision the datastores.

Figure 12 illustrates the physical and logical storage configuration of NetApp FAS controllers in MetroCluster setup.

Figure 12) Physical and logical storage configuration of NetApp FAS controllers in MetroCluster setup.

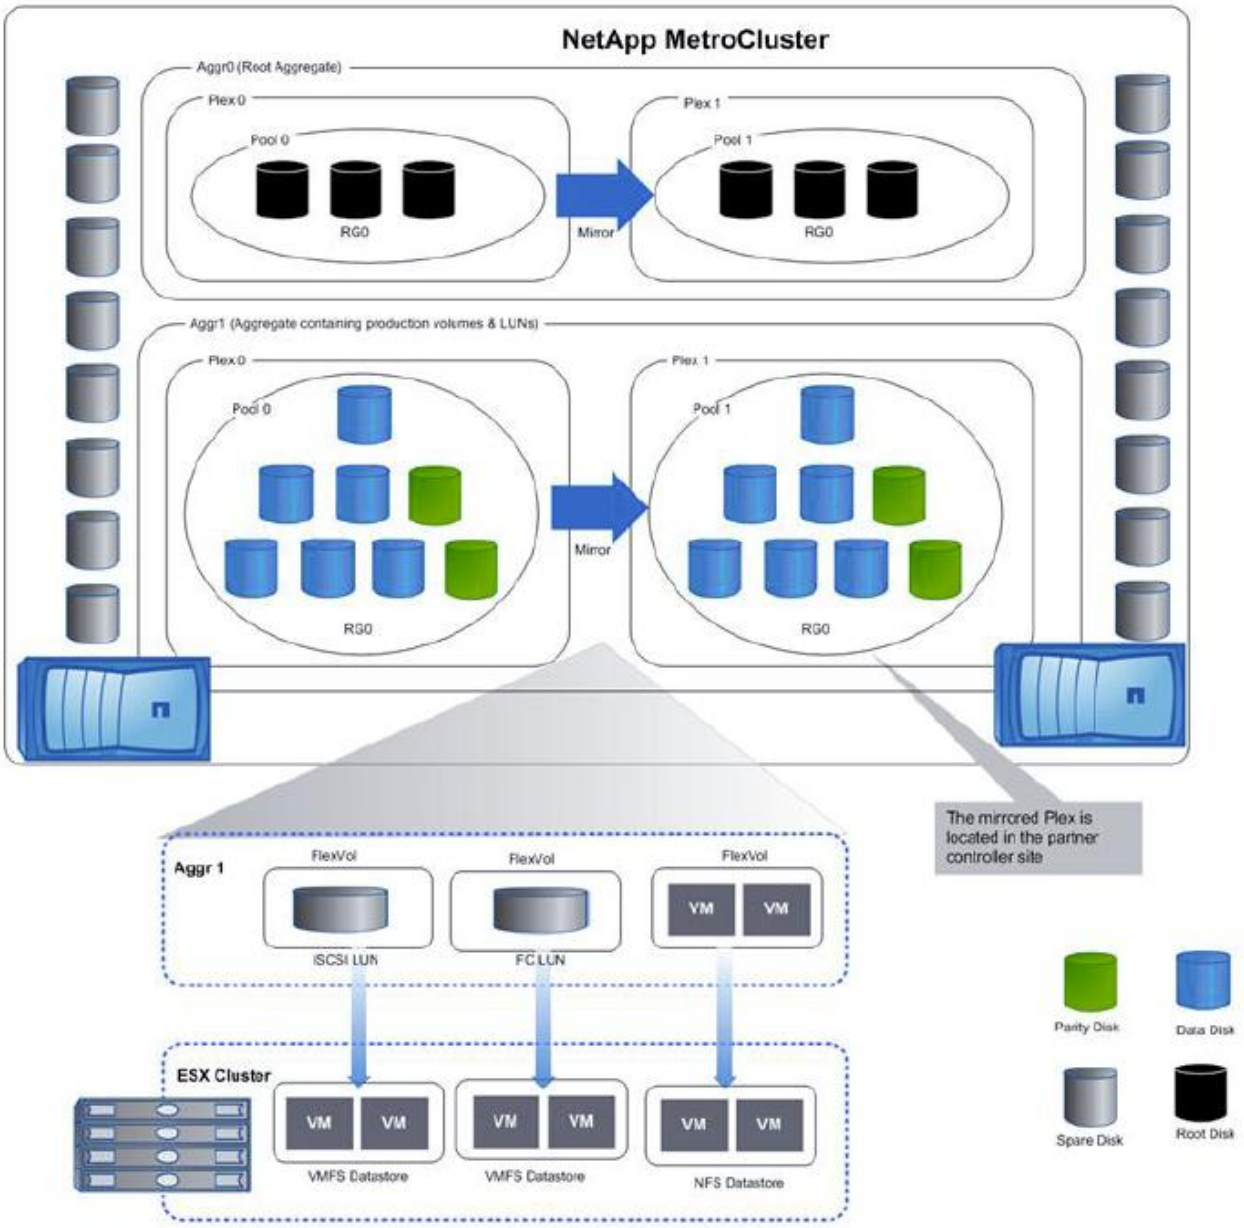

5.2 VMware vSphere Configuration

Configuring vSphere HA for NetApp MetroCluster

The following steps will help a user to configure HA settings in an existing cluster.

1. Connect to vCenter using the vSphere client.
2. Select Hosts and Clusters view.
3. Right-click the existing cluster and select “Edit Settings.”
4. Select the option “Turn On vSphere HA.”
5. On the left pane, select vSphere HA and verify that the option “Host Monitoring” is selected.

6. For the Admission Control option, select “Disable” since the solution’s main goal is maximum availability rather than performance in the case of host failure.

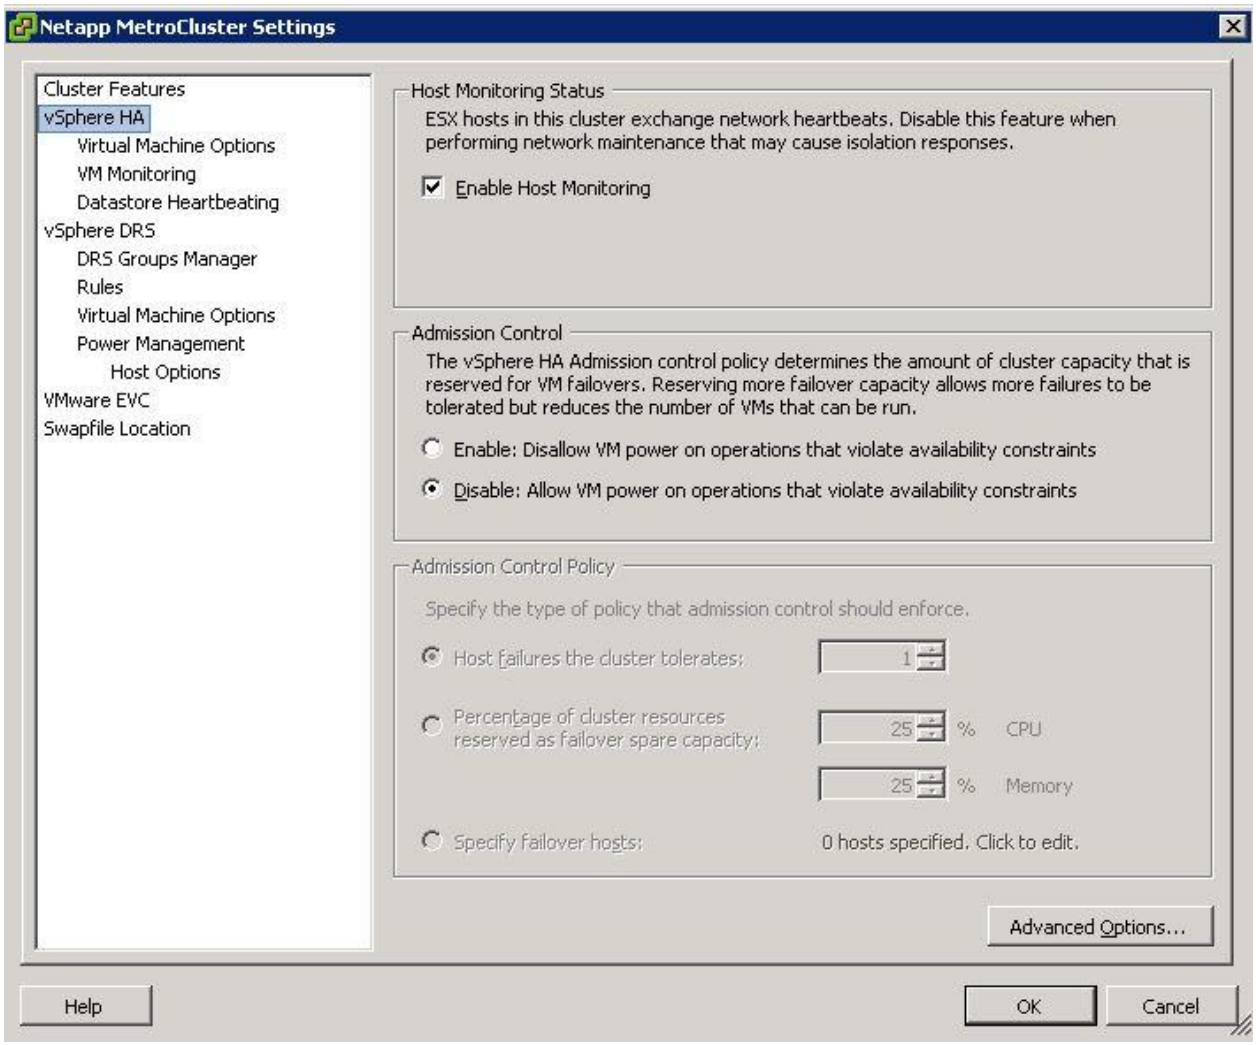

7. This step is optional. Select the “Advanced Options” settings to add the Isolation address by adding the parameter `das.isolationaddress` and the IP address next to it.

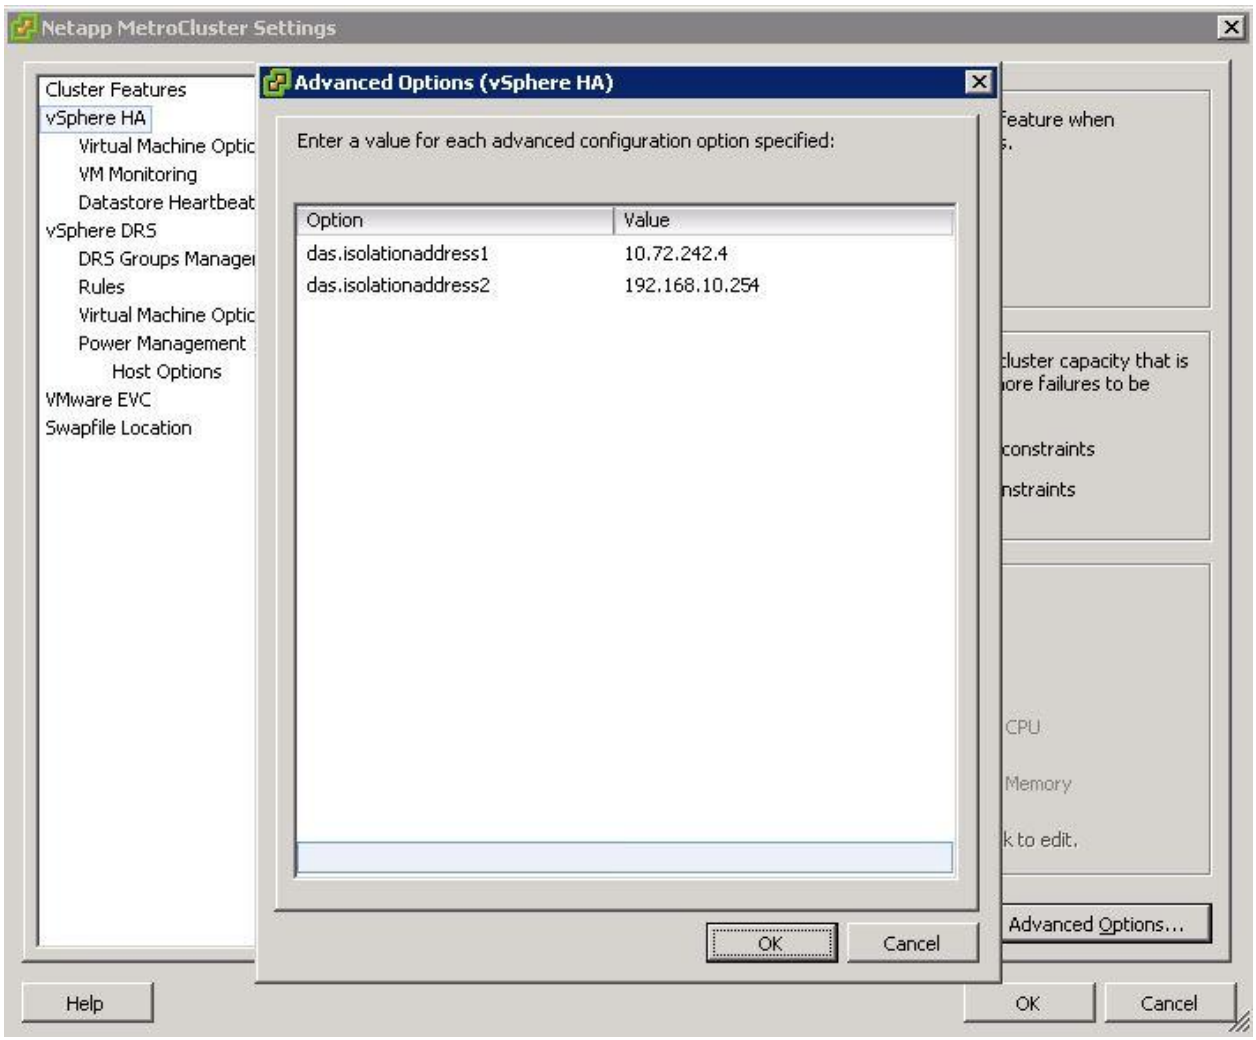

8. The VM restart priority can be customized per the requirement. By default, all of the virtual machines' restart priorities are set to "Medium."
9. For the Host Isolation response setting, NetApp recommends "Leave Powered On."

Best Practice

By setting the Host Isolation response to "Leave Powered On," you can avoid unnecessary downtimes for virtual machine restarts in situations in which the management network failure does not correspond to the virtual machine network.

Note: In iSCSI/NFS environments in which the management network correlates with the IP storage network, it is impossible for hosts to decide whether it is fully isolated. In these environments, it is better to change the setting to "Shutdown," which will gracefully shut down the VMs whenever there is an isolation response. This avoids split-brain scenarios too.

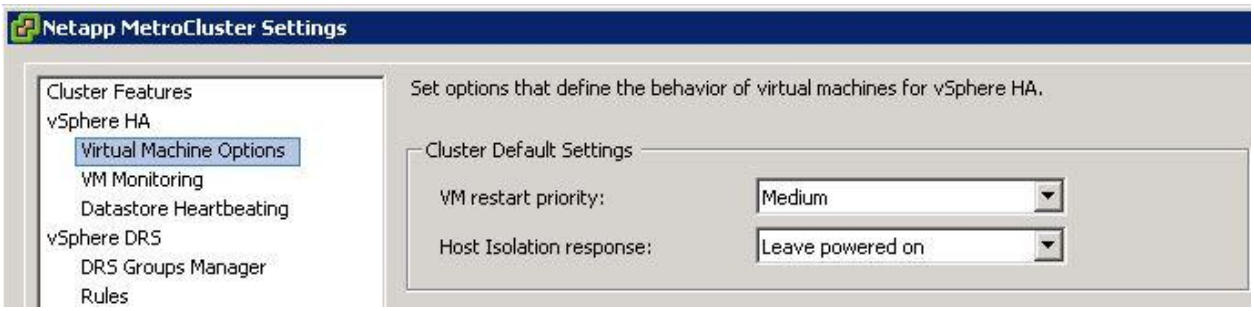

10. VM Monitoring settings can be customized based on the requirement.
11. In the Datastore Heartbeating section, select “Select any of the cluster datastores taking into account my preferences.”
12. Click OK.

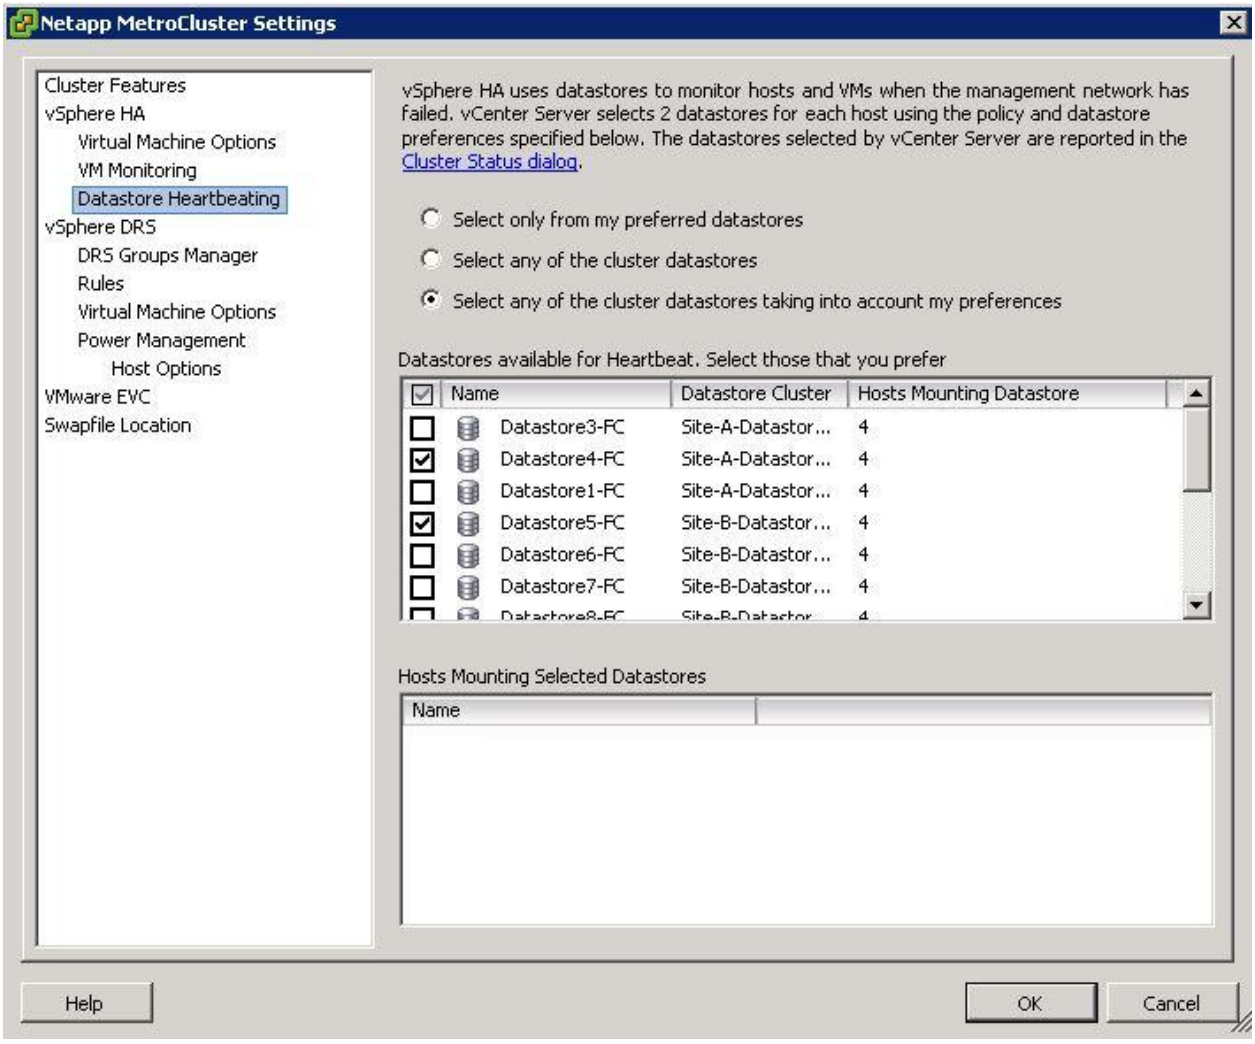

Configuring VMware DRS Groups in vSphere 5 for NetApp MetroCluster

The following steps help create DRS groups for VMs and hosts. This step is a prerequisite for configuring rules.

1. Connect to vCenter using the vSphere client.
2. In the vSphere client, right-click the cluster in the inventory and select Edit Settings.
3. Select the DRS Group Manager tab.
4. Create two virtual machine DRS groups, one for the local site and one for the remote site.
5. Add the VMs of the respective site to these groups.
6. Create two host DRS groups, one for the local site and one for the remote site.
7. Add the hosts to their respective host groups.

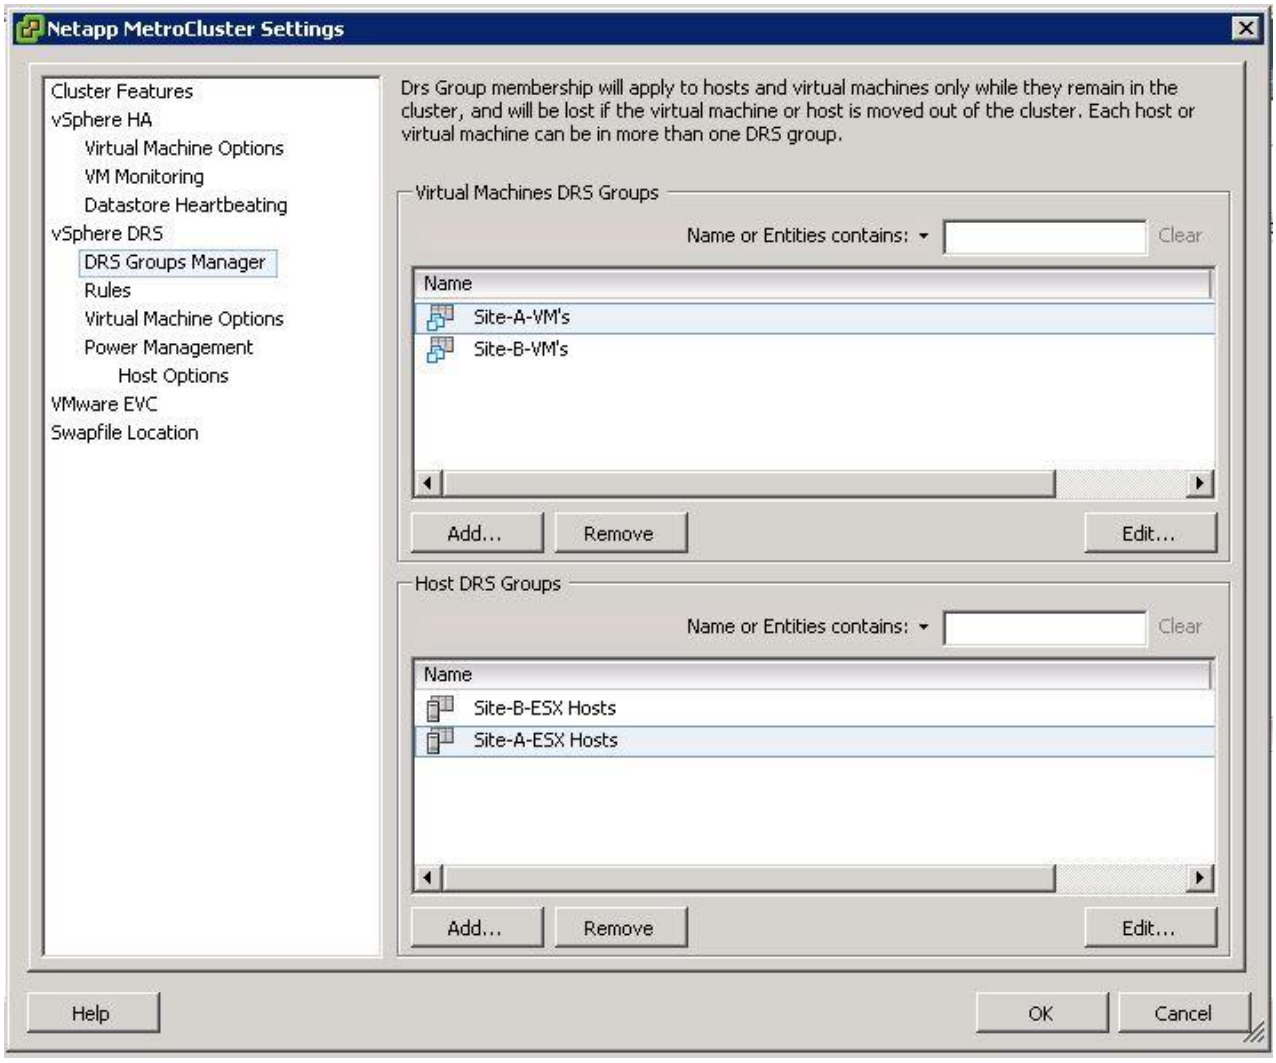

Configuring DRS Rules in vSphere 5 for NetApp MetroCluster

The following steps help create DRS rules specific to site A and site B.

1. In the vSphere client, right-click the cluster in the inventory and select Edit Settings.
2. Select the Rules tab.

3. Click Add.
4. In the Rule dialog box, type the name for the rule.
5. From the Type menu, select Virtual Machines to Hosts.
6. Select the virtual machine DRS group and the host DRS group to which the rule applies.
7. Verify that the VM and host DRS groups created for site A are selected for the site A rule.
8. Select the Specification "Should run on hosts in group." (Virtual machines in VM Group A should, but are not required to, run on hosts in Host Group A.)
9. Repeat steps 3 through 7 to add another rule for site B.
10. Click OK.

Best Practice

NetApp highly recommends the specification "Should run on hosts in group" rather than the specification "Must run on hosts in group." In the event of a site A host's failure, the VMs of site A need to be restarted on hosts at site B through vSphere HA, but the latter specification does not allow HA to restart VMs on site B because it's a hard rule. The former specification is a soft rule and will be violated in the event of HA, thus enabling availability rather than performance.

Note: You can create an event-based alarm that is triggered when a virtual machine violates a VM-Host affinity rule. In the vSphere Client, add a new alarm for the virtual machine and select "VM is violating VM-Host Affinity Rule" as the event trigger. For more information about creating and editing alarms, see the vSphere Monitoring and Performance documentation.

Figure 13 illustrates how to set DRS rules.

Figure 13) DRS rules.

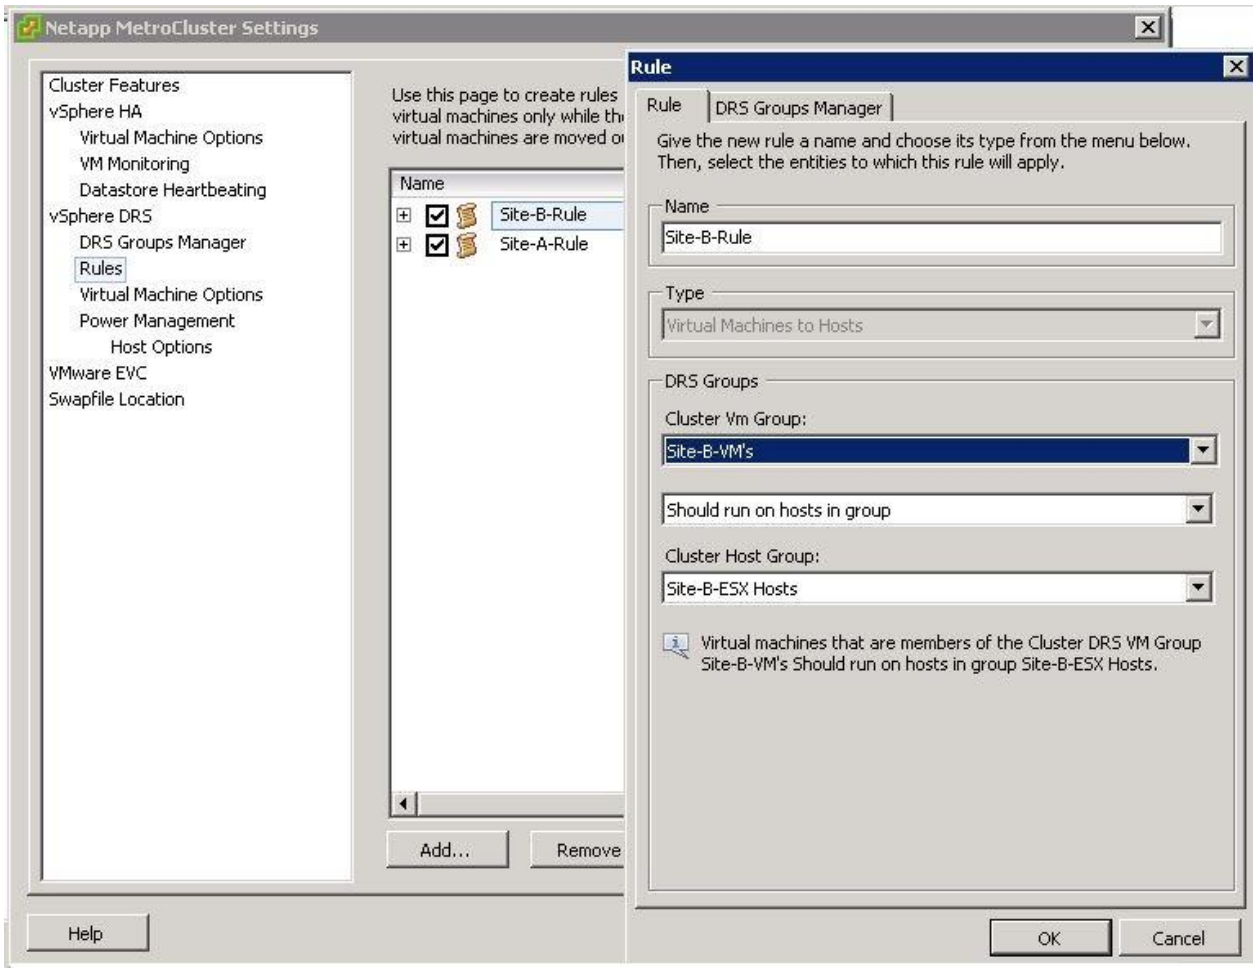

Creating a Datastore Cluster

The following steps will help you configure a datastore cluster.

1. Connect to vCenter using the vSphere client.
2. In the Datastores and Datastore Clusters view of the vSphere Client inventory, right-click the data center object and select New Datastore Cluster.
3. Name the datastore cluster and verify that the option “Turn on Storage DRS” is selected.

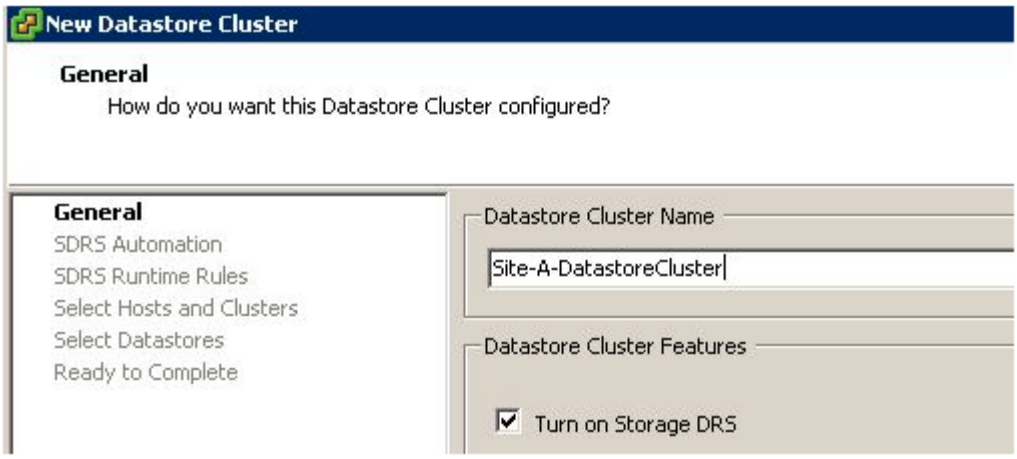

4. For the SDRS automation level, verify that “No Automation (Manual Mode)” is selected.

Best Practice

NetApp recommends that Storage DRS be configured in manual mode, so that the administrator gets to decide and control when migrations need to happen.

5. Verify that the setting “Enable I/O metric for SDRS recommendations” is selected; metric settings can be left with default values.

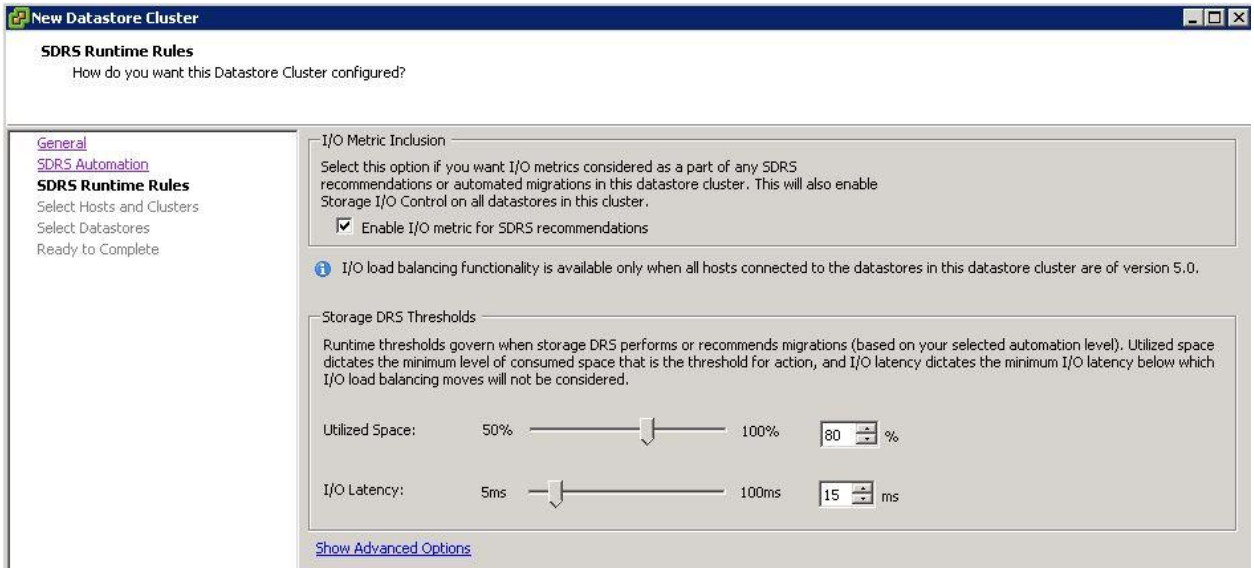

6. Select the appropriate host cluster.
7. In the Select Datastores screen, verify that only the datastores of site A are selected; thus they will be part of the site A datastore cluster.
8. Repeat steps 1 through 6 to create the site B datastore cluster and verify that only datastores of site B are selected.

6 Architecture Use Cases

6.1 Single Storage Path Failure

In this scenario, if components such as the HBA port, the network port, the front-end data switch port, or an FC/Ethernet cable fail, that particular path to the storage device is marked as dead by the ESX host. If several paths are configured for the storage device by providing resiliency at the HBA/network/switch port, ESX ideally performs a path switchover. During this period, virtual machines remain running without getting affected, because availability to the storage is taken care of by providing numerous paths to the storage device.

Note: There is no change in MetroCluster behavior in this scenario, and all the datastores continue to be intact from their respective sites.

Best Practice

In environments in which NFS/iSCSI volumes are used, NetApp recommends having at least two network uplinks configured for the NFS vmkernel port in the standard vSwitch and the same at the port group where the NFS vmkernel interface is mapped for the distributed vSwitch. NIC teaming can be configured in either active-active or active-standby.

Also, for iSCSI LUNs, multipathing needs to be configured by binding the vmkernel interfaces to the iSCSI network adapters. For more information, refer to the vSphere storage documentation.

Best Practice

In environments in which Fibre Channel LUNs are used, NetApp recommends having at least two HBAs, which guarantees resiliency at the HBA/port level. NetApp also recommends single initiator to single target zoning as the best practice to configure zoning.

The Virtual Storage Console (VSC) should be used to set multipathing policies because it sets policies for all new and existing NetApp storage devices.

6.2 Single ESX Host Failure

Figure 14 illustrates host failure.

Figure 14) Host failure.

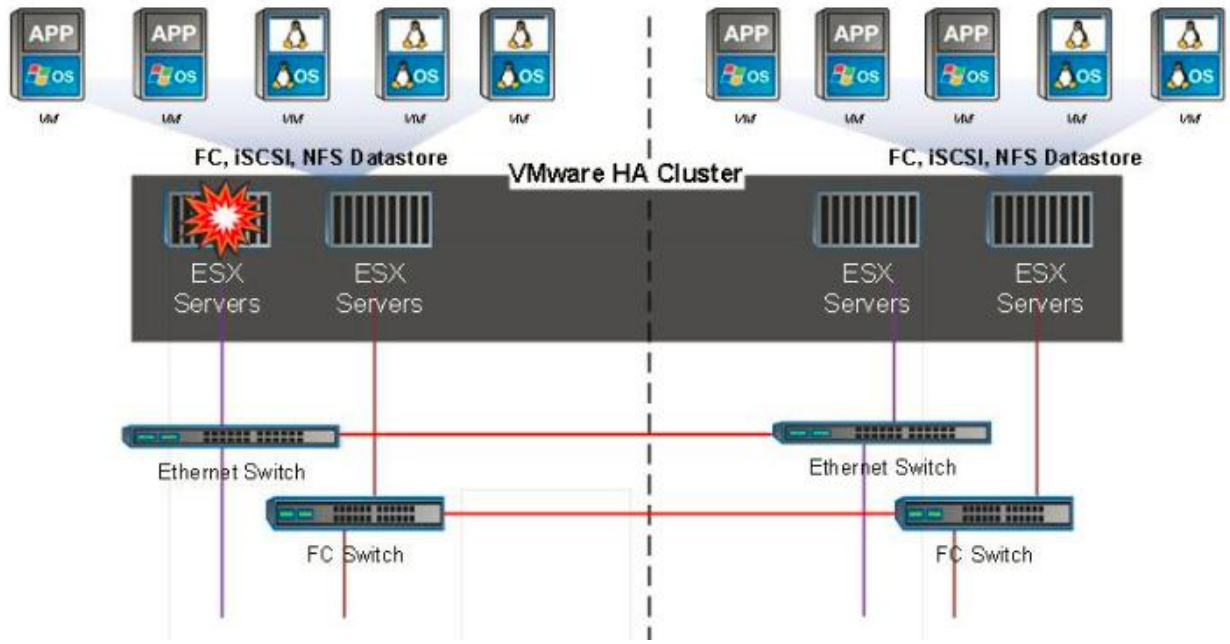

In this scenario, if there is an ESX host failure, the master node in the VMware HA cluster detects the host failure since it no longer receives network heartbeats. To determine whether the host is really down or only a network partition, the master node monitors the datastore heartbeats and, if they are absent, it performs a final check by pinging the management IP addresses of the failed host. If all these checks are negative, then the master node declares this host a failed host and all the virtual machines that were running on this failed host are rebooted on the surviving host in the cluster.

It is possible that the virtual machines will be started on the ESX hosts at the other site if there is a resource constraint in the local site. However, the defined DRS virtual machine-host affinity rules will correct if any rules are violated by migrating the virtual machines back to any surviving ESX hosts in the local site. In cases in which DRS is set to manual, NetApp recommends invoking the DRS and applying the recommendations to correct the virtual machine placement.

There is no change in the MetroCluster behavior in this scenario and all the datastores continue to be intact from their respective sites.

6.3 ESX Host Isolation

Figure 15 illustrates ESX host isolation.

Figure 15) ESX host isolation.

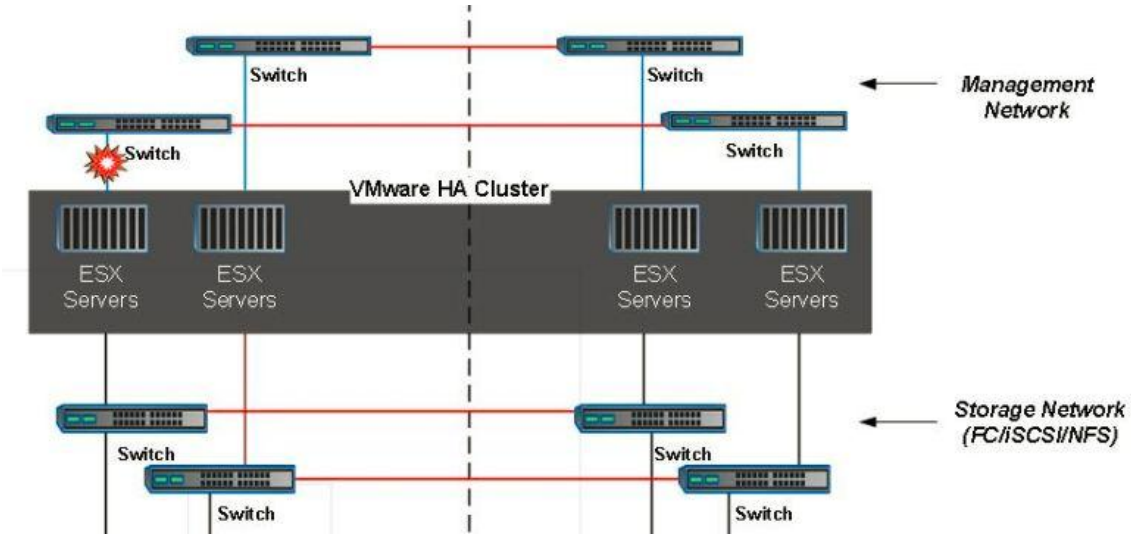

In this scenario, if the management network of the ESX host is down, the master node in the HA cluster will not receive any heartbeats, and thus this host becomes isolated in the network. To determine whether it has failed or is only isolated, the master node starts monitoring the datastore heartbeat. If it is present then the host is declared isolated by the master node. Depending on the isolation response configured, the host may choose to power off, shut down the virtual machines, or even leave the virtual machines powered on. The default interval for the isolation response is 30 seconds.

There is no change in the MetroCluster behavior in this scenario and all the datastores continue to be intact from their respective sites.

6.4 vCenter Server Failure

In this scenario, if vCenter service is unavailable, the only service affected will be the DRS feature, for which operations such as load balancing and placement of VMs as per the rules will be unavailable. The virtual machines will continue to run in their respective hosts unless there is a host failure, in which case HA will trigger the virtual machine restart on the surviving hosts. In a host failure scenario, the virtual machines may be restarted on hosts at the other site, and this might induce network latency because the virtual machines access the datastore remotely. After the vCenter service is up, the DRS host-VM affinity rule will correct VM placement issues.

Note: There is no change in the MetroCluster behavior in this scenario and all the datastores continue to be intact from their respective sites.

6.5 Interswitch Link Failure

Interswitch Link Failure at Management Network

Figure 16 illustrates the interswitch link failure at management network.

Figure 16) Interswitch link failure at management network.

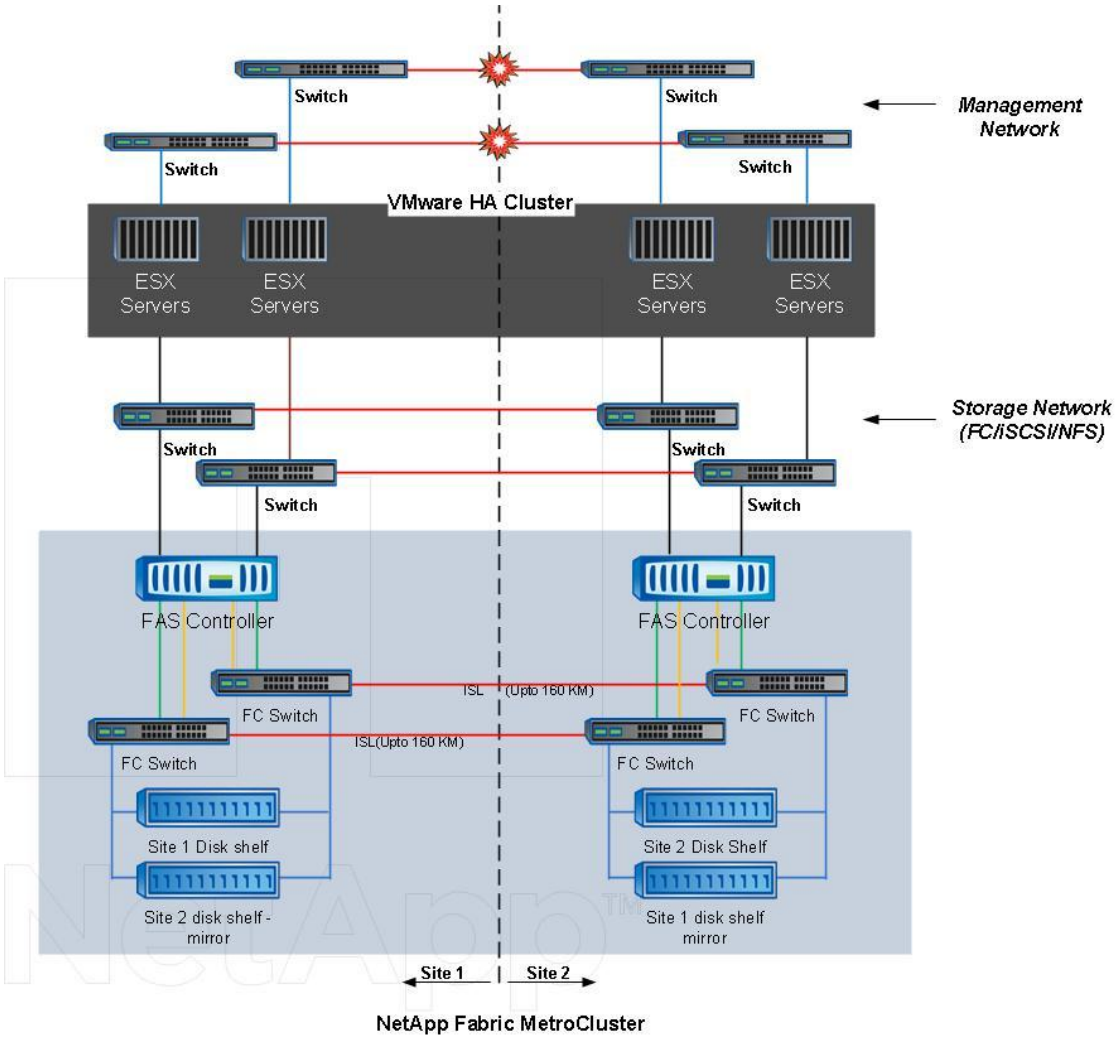

In this scenario, if the ISL links at the front-end host management network fail, the ESX hosts at site A will not be able to communicate with ESX hosts in the other site. This will lead to a network partition because ESX hosts at a particular site will be unable to send the network heartbeats to the master node in the HA cluster. As such, there will be two network segments because of partition and there will be a master node in each segment that will protect VMs from host failures within the particular site.

Note: During this period, the virtual machines remain running and there is no change in the MetroCluster behavior in this scenario. All the datastores continue to be intact from their respective sites.

Interswitch Link Failure at Storage Network

Figure 17 illustrates the interswitch link failure at storage network.

Figure 17) Interswitch link failure at storage network.

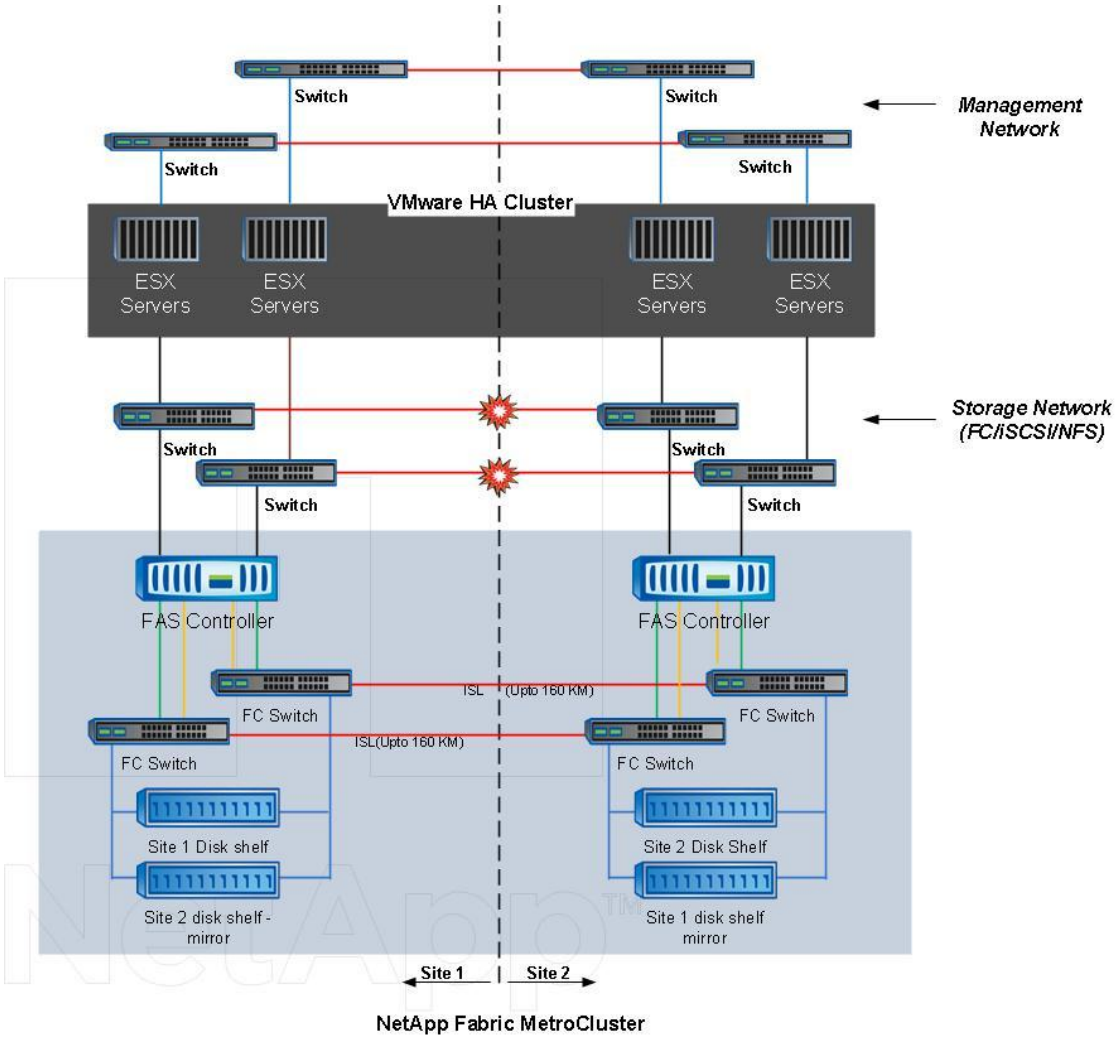

In this scenario, if the ISL links at the back-end storage network fail, the hosts at site A will lose access to the storage volumes/LUNs of storage array B at site B and vice versa. The VMware DRS rules are defined so that host-storage site affinity facilitates the virtual machines to run without impact within the site.

Note: During this period, the virtual machines remain running in their respective sites and there is no change in the MetroCluster behavior in this scenario. All the datastores continue to be intact from their respective sites.

If for some reason the affinity rule was violated, that is, virtual machine A, which was supposed to run from site A where its disks reside on local storage array, is running on a host at site B, the virtual machines disk will be remotely accessed via ISL links. Because of ISL link failure, VM A running at site B will not be able to write to its disks because the paths to the storage volume are down and that particular virtual machine is down. In these situations, VMware HA does not take any action since the hosts are actively sending heartbeats. Those virtual machines need to be manually powered off and powered on in their respective sites.

Interswitch Link Failure Between Controllers in NetApp Fabric MetroCluster

Figure 18 illustrates the interswitch link failure between controllers in NetApp fabric MetroCluster.

Figure 18) Interswitch link failure between controllers in NetApp fabric MetroCluster.

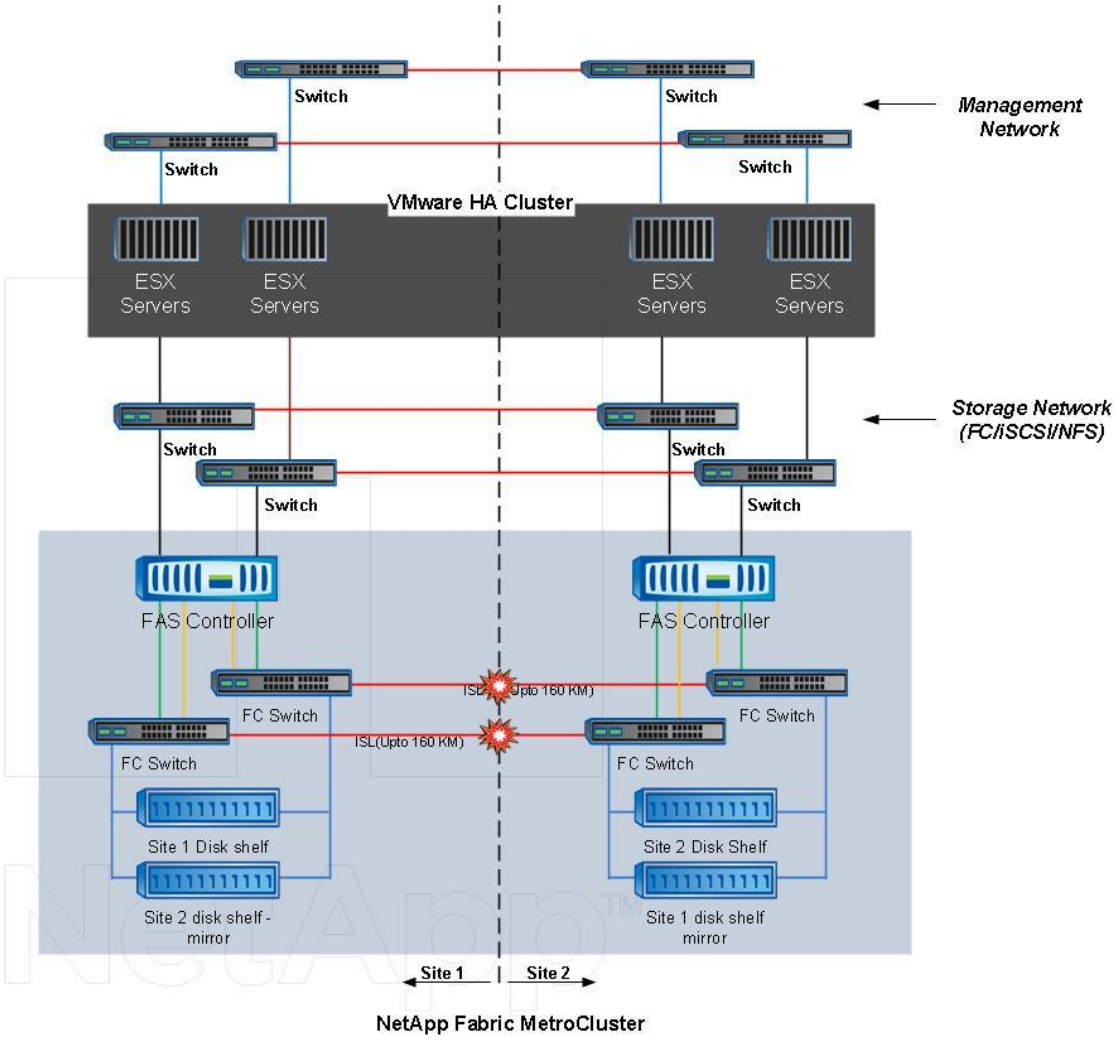

In this scenario, if the ISL links between the controllers in fabric MetroCluster fail, automatic controller failover is disabled and the storage volumes/LUNs of the respective controllers at each site are available for the hosts in their respective sites.

NetApp recommends checking the status of storage volumes/LUNs in the respective sites if the hosts at a site lost access to the local storage array due to rolling failures of back-end switches connecting the storage controller, and if the storage controller itself, "controller failover on demand," needs to be initiated from the surviving controller to restore the host access to storage volume/LUNs.

The following snippet was taken when the ISL links between the storage controllers were down; it has the output of command `cf status`.

```
sitea-3240> cf status
siteb-3240 may be down, takeover disabled because of reason (partner mailbox dis
ks not accessible or invalid)
sitea-3240 has disabled takeover by siteb-3240 (interconnect error)
VIA Interconnect is down (link 0 down, link 1 down).
The DR partner site might be dead.
 To take it over, power it down or isolate it as described in the Data Protectio
n Guide, and then use cf forcetakeover -d.
sitea-3240>
```

6.6 All Interswitch Failure or Complete Data Center Partition

Figure 19 illustrates all ISL failure.

Figure 19) All ISL failure.

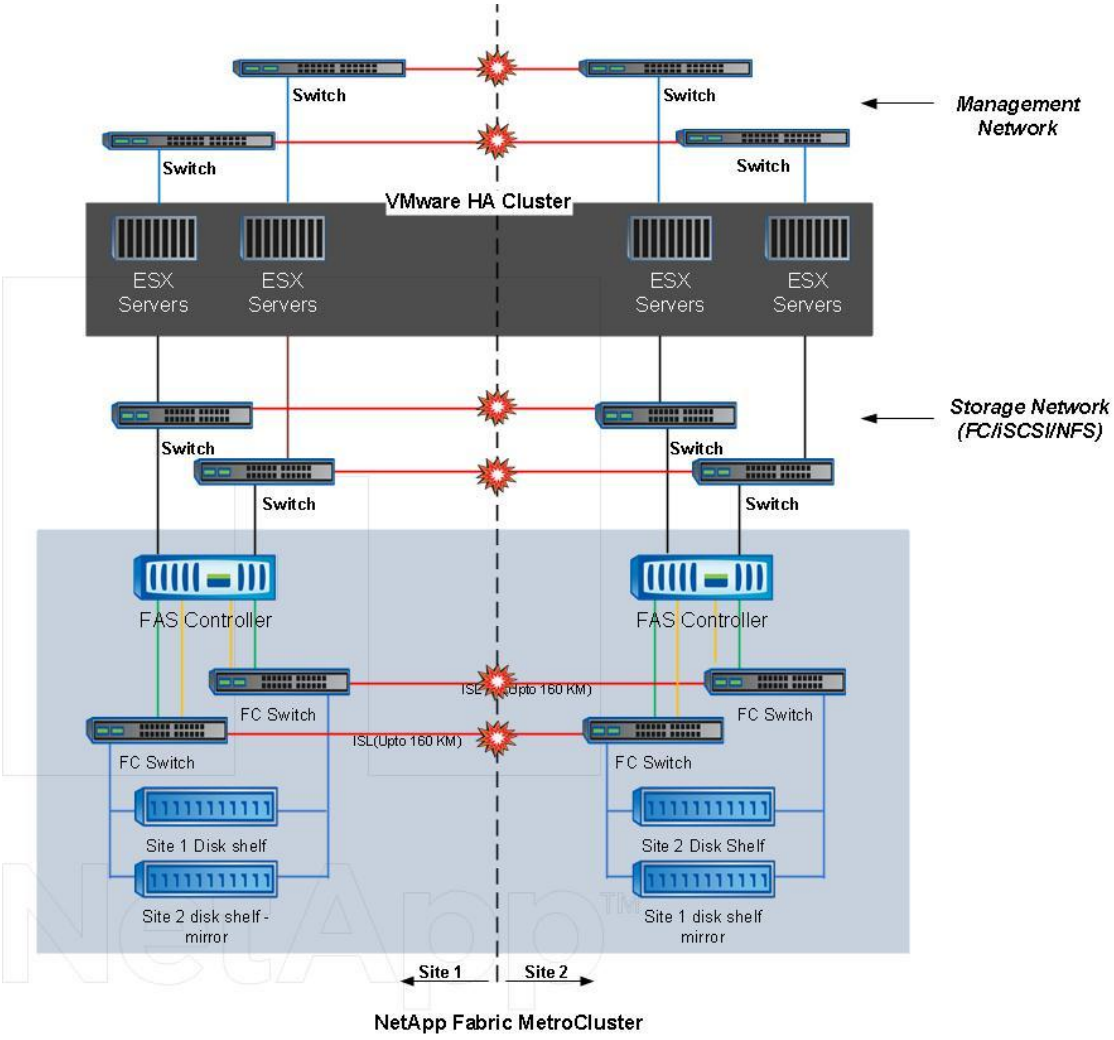

In this scenario, all the ISL links between the sites are down and both the sites are isolated from each other. As discussed in earlier scenarios, such as ISL failure at the management network and at the storage network, the virtual machines are not affected in complete ISL failure.

After ESX hosts are partitioned between sites, the vSphere HA agent will check for datastore heartbeats and, in each site, the local ESX hosts will be able to update the datastore heartbeats to their respective read-write volume/LUN. Hosts in site A will assume that the other ESX hosts at site B have failed because there are no network/datastore heartbeats. vSphere HA at site A will try to restart the virtual machines of site B, which will eventually fail because the datastores of site B will not be accessible due to storage ISL failure. A similar situation is repeated in site B.

NetApp recommends determining if any virtual machine has violated the DRS rules. The virtual machine running from a remote site will be down since it will not be able to access the datastore, and vSphere HA will restart that virtual machine on the local site. After the ISL links are back online, the virtual machine that was running in the remote site will be killed, since there cannot be two instances of virtual machines running with the same MAC addresses.

6.7 Storage Controller Failure

Figure 20 illustrates the storage controller failure.

Figure 20) Storage controller failure.

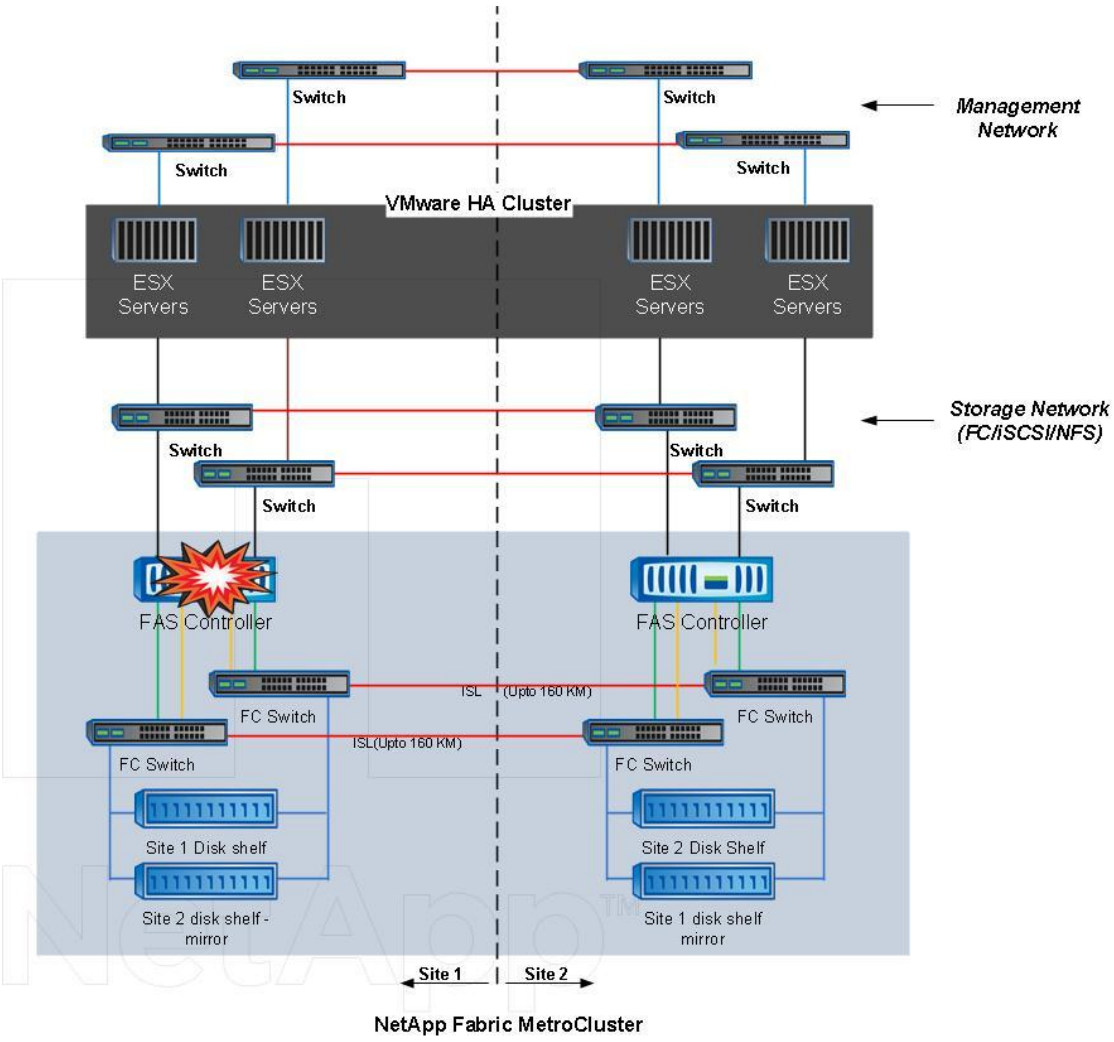

In this scenario, if the storage controller fails at one site, all the services of the storage controller at the failed site will be failed over to the storage controller at the surviving site. The surviving controller will start serving the data from the disk shelves at the failed site if it is only a controller failure, since disk shelves will be very much accessible to the surviving storage controller. During this process, ESX hosts at the failed site will lose access to the storage paths to the failed controller and initiate the path failover to the active paths of the surviving controller, which will be accessed through ISL links of the front-end storage switch.

Virtual machines will not be affected because the alternate paths to the same storage volume/LUNs will get active. In addition, vSphere HA will not take any action since the master node in the cluster will still be receiving the network heartbeats.

There may be situations in which there is a complete storage failure of the storage controller and all disk shelves at a particular site. During such situations, controller failover on demand needs to be initiated manually; this will fail over all the storage services of the failed controller to the surviving controller and also the plexes or mirror aggregates that were read-only will become read-write. It is a rare event in a

data center when both controller and disk shelves fail, but power loss to the storage rack or back-end switch failure at a particular site in fabric MetroCluster can lead to this happening.

After the failed storage controller is back online, controller failback needs to be initiated after all the mirror aggregates are synced. Also, in situations in which the automatic failover option is not enabled, manual failover needs to be initiated by using the CLI.

During this period, there is no impact on virtual machine I/O operations but performance is degraded since the data is being accessed from the remote storage controller through ISL links.

The following snippet was taken during storage controller failure at site A caused by a panic, but the disk shelves at site A were accessible to the storage controller at site B.

```
siteb-3240> Thu OctThu Oct 11 00:13:34 EST [siteb-3240:cf.fsm.firmwareStatus:inf
o 11 00:]: Failover monitor: partner Dumping sparecore
13:34 EST [sitThu Oct 11 00:13:34 EST [eb-3240siteb-3240:cf.fsm.takeover.panic:A
LERT]: Failover monitor: takeover attempted after partner panic
:cf.fsm.firmwaThu Oct 11 00:13:34 EST [sitereStatub-3240:cf.fsm.stateTransit:inf
o]: Failover monitor: UP --> TAKEOVER
```

During firmware upgrades in which the storage controller needs to be rebooted, downtime for the storage services is a challenge in tier 1 and 2 environments. In NetApp FAS storage in an HA pair, whenever rebooting the storage controller is required (it could be for several reasons, such as a firmware upgrade or a hardware refresh for the controller), the customer can perform a negotiated failover wherein the surviving storage controller takes over all the services of the controller that will be rebooted. After the controller is back online, failback is initiated from the surviving controller, which resumes services as before.

The following snippet was taken during a firmware upgrade process in which we were required to reboot the storage controller at site B and we initiated the storage takeover process from the storage controller at site A.

```
sitea-3240> cf takeover
cf: takeover initiated by operator
sitea-3240> Thu OctThu Oct 11 00:38:0 11 00:6 EST [sitea-3240:cf.misc.operatorTa
keover:notice]: Failover monitor: takeover initiated by operator
38:06 EST [sitea-3240Thu Oct 11 00:38:06 EST [site:cf.misa-3240:cf.fsm.nfo.accep
tTakeoverReq:notice]: Negotiated failover: accepting takeover request by partner
, reason: operator initiated cf takeover. Asking partner to shutdown gracefully;
 will takeover in at most 180 seconds.
```

6.8 Disk Shelf Failure

Figure 21 illustrates the disk shelf failure.

Figure 21) Disk shelf failure.

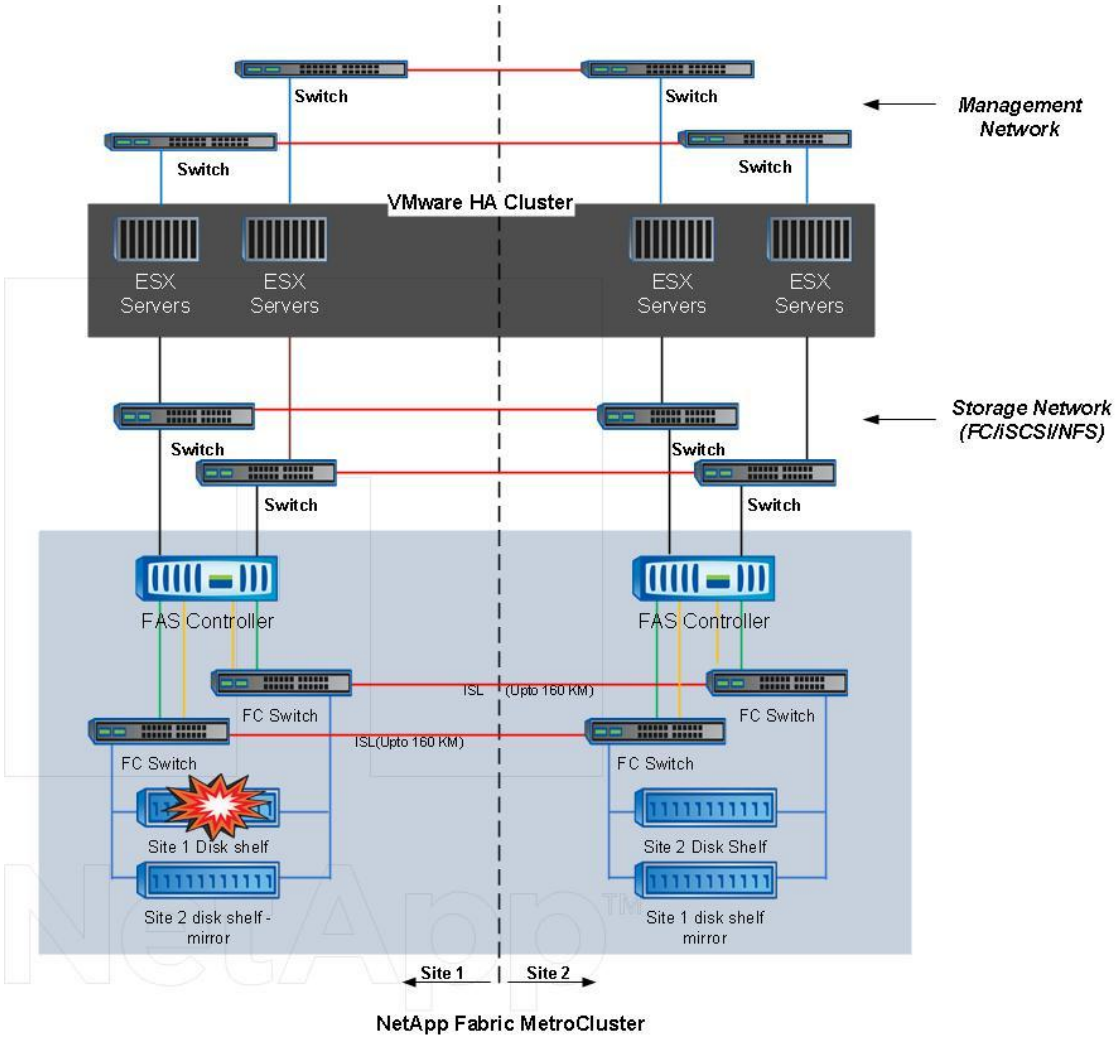

In this scenario, if there is a failure of the disk shelf at site A, which had volumes/LUNs that were in production at site A (a site 1 disk shelf failure is shown above), then the corresponding disk shelf at site B (that is, the site 1 disk shelf mirror shown in Figure 21.) where the mirror copy of the data is present will become read-write and will be available for the ESX hosts. In situations in which the disk shelves that have mirror copy fail, there is no change in MetroCluster behavior.

Note: During this period, there is no impact on the virtual machine I/O operations but there is degraded performance since the data is being accessed from the remote disk shelf through ISL links.

6.9 Complete Site Failure

Figure 22 illustrates the complete site failure.

Figure 22) Complete site failure.

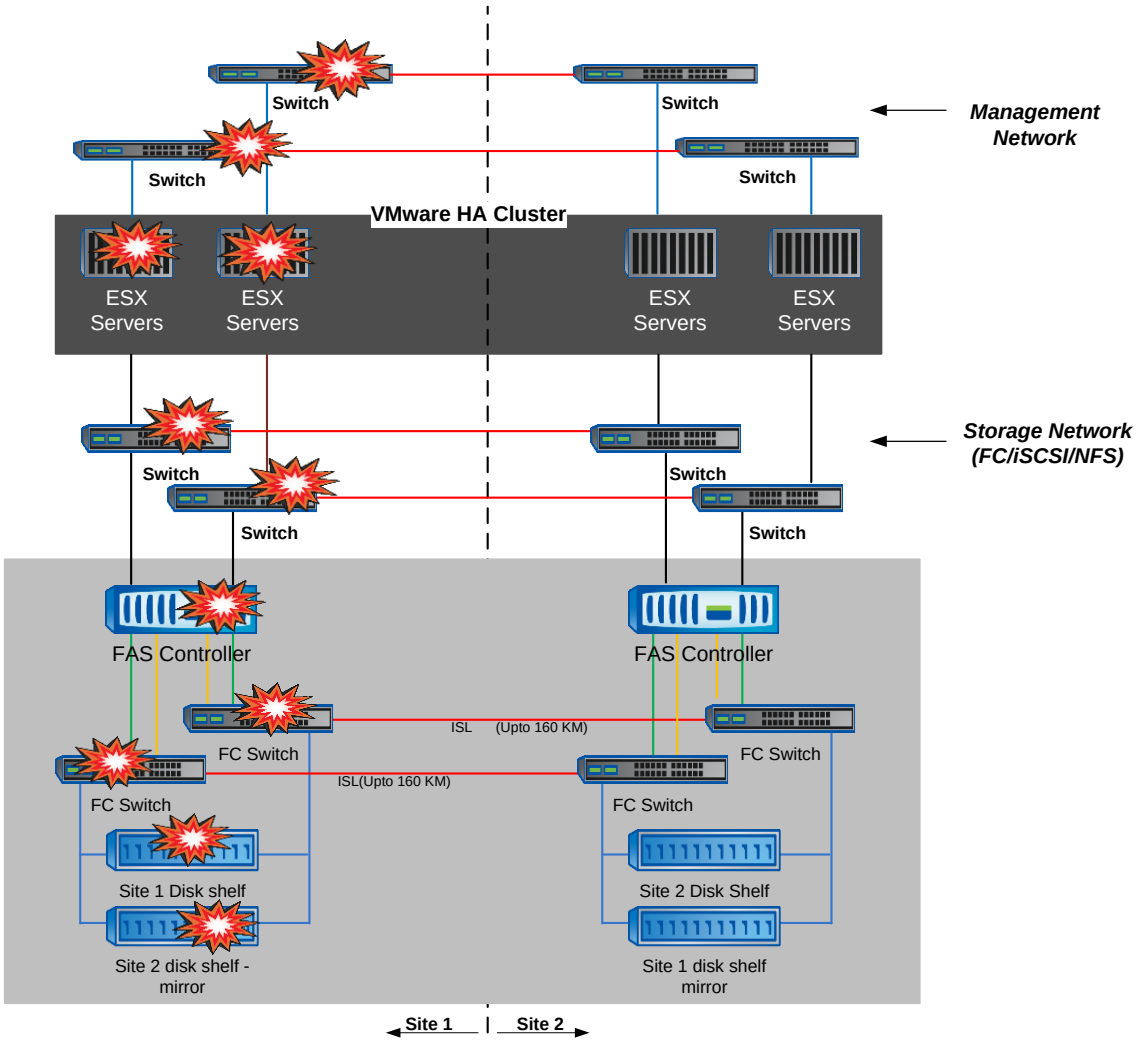

NetApp Fabric MetroCluster

In this scenario, if there is a complete site A failure, the ESX hosts at site B will not get the network heartbeat from the ESX hosts at site A, since they are down and the ESX hosts at site B will verify if datastore heartbeats are present. If datastore heartbeats are not present, the ESX hosts in site A will be declared failed and they will try to restart the site A virtual machines in site B. During this period, customers should initiate the controller failure on demand, and that will restore all the storage services of site A at site B. After the site A volumes/LUNs are available at site B, the hosts at site B will be able to access the virtual machines of site A and restart them successfully.

Best Practice

The controller failover on demand needs to be initiated within 30 minutes of the site failure because vSphere HA will stop trying to restart the virtual machines at a failed site after 30 minutes.

The following snippet was taken during complete site B failure. The storage controller at site A took over all the storage services of the failed controller by manually initiating controller failure on demand (CFOD) from the console of the storage controller at site A.

```
sitea-3240> cf forcetakeover -d
Following the command, mirrored volumes will be split and
Clients of the partner controller may experience data loss because of disaster.
Prior to issuing this command, the partner controller should be powered off.
If the partner controller is operational or if it becomes operational at any tim
e
while this controller is running in takeover mode, your filesystems may be destr
oyed.
Do you wish to continue [y/n] ?? y
cf: forcetakeover -d initiated by operator
sitea-3240> Thu OctThu Oct 11 05:09:32 EST [sitea- 11 05:3240:cf.misc.operatorDi
sasterTakeover:notice]: Failover monitor: forcetakeover -d initiated by operator

09:32 EST [sitThu Oct 11 05:09:32 EST [sitea-3240:cf.fsm.takeover.disaster:info]
ea-3240ver monitor: takeover attempted after 'cf forcetakeover -d' command.
:cf.misThu Oct 11 05:09:32c.opera EST [sitea-3240:cf.fsm.stateTransit:info]: Fai
lover monitor: UP --> TAKEOVER
```

7 Combination Tests (Failures That Affect Both Sites)

Table 3 lists the combination test 1: ESX host server and storage controller failure across sites.

Table 3) Combination test 1: ESX host server and storage controller failure across sites.

Tests Performed	1. Power off all ESX hosts at site 1. 2. Power off the storage controller at site 2.
Expected Results	• The controller at site 1 automatically takes over the powered-off controller. • All the VMs in site 1 are restarted on the surviving hosts in site 2.
Actual Results	Actual results were in line with the expected behavior, and the tests passed as expected.
MetroCluster Behavior	Partner controller in site 1 performs an automatic takeover.
VMware HA Behavior	VMs that were previously running in the failed hosts are automatically powered on in the surviving nodes.

Table 4 lists the combination test 2: Disk shelf failure in both sites.

Table 4) Combination test 2: Disk shelf failure in both sites.

Tests Performed	1. Power off disk pool 0 in site 1. 2. Power off disk pool 0 in site 2.
Expected Results	• VMs should not detect any changes and continue to operate normally. • VMs in DRS groups should not migrate.
Actual Results	Actual results were in line with the expected behavior, and the tests passed as expected.
MetroCluster Behavior	No MetroCluster event. Data is served from the mirrored copy.
VMware HA Behavior	No HA event.
Impact to Data Availability	None.

Table 5 lists the combination test 3: Controller and disk shelf failure.

Table 5) Combination test 3: Controller and disk shelf failure.

Tests Performed	1. Power off storage controller in site 1. 2. Power off disk pool 0 in site 2.
Expected Results	• VMs should not detect any changes and should continue to operate normally. • VMs in DRS groups should not migrate.
Actual Results	Actual results were in line with the expected behavior, and the tests passed as expected.
MetroCluster Behavior	Surviving storage controller performs automatic takeover.
VMware HA Behavior	No HA event.
Impact to Data Availability	None.

8 Conclusion

A plethora of high-availability solutions available in the market make it challenging to choose the optimum business continuity solution for the enterprise, one that can integrate with the business solution and meet the customer's business continuity standards. This paper tackles these challenges by presenting a consolidated solution for implementing the NetApp MetroCluster with VMware vSphere HA solution, a best-in-class business continuity enterprise solution.

This solution presents a unified approach to handling the critical business requirement for continuous service availability, and it provides the capability to meet varying customer requirements for RPOs and RTOs. This solution works across the domains of planned and unplanned outages and elaborates upon use case-based implementations covering all the disruption scenarios. A key benefit of this solution is the tight integration between VMware and NetApp's HA and FT features. This simplifies the implementation and maintenance of this disaster recovery solution without compromising on its agility, robustness, or cost effectiveness.

About the Authors

Ashish Nainwal is a Technical Marketing Engineer in NetApp Infrastructure and Cloud Enablement group, and he is focused on storage and business continuity solutions based on NetApp products.

Santhosh Devaraju is a Technical Marketing Engineer in NetApp Infrastructure and Cloud Enablement group, and he is focused on virtualization and business continuity solutions based on NetApp products.

Version History

Version	Date	Document Version History
Version 1.0	February 2013	Initial release

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.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®