

iSCSI High Availability Design using the MDS 9000 Family of Multilayer Switches

Introduction

In enterprise environments, many of the benefits of migrating to a SAN-based model for storage deployment and management are being realized. The ability to share storage resources across different applications and servers significantly reduces costs associated with management and utilization rates of the consolidated environment. At the same time, this trend to migrate mission critical data onto shared storage resources accessible via multiple servers also mandates uninterrupted storage access. With the availability of IP SAN attachment capabilities in the MDS Family using the IP Services module, application deployments must continue to offer the utmost in high availability and resiliency. All enterprises can benefit from the cost and complexity savings offered by iSCSI solutions as it represents another choice for cost of connectivity to the shared storage resources. However, any loss of data can be potentially damaging and costly for businesses. As such, businesses are now demanding 99.999% availability for their data, regardless of whether it is access via Fibre Channel or IP (iSCSI).

High availability has long been a mandatory storage design constraint for mission-critical applications in such industries as finance, manufacturing, and telecommunications. Traditionally, high availability meant fault tolerant systems only. This implied a hardware system configuration that allowed protected system components to fail-over to redundant components or set of components so that downtime is minimized. High availability must address the following three concerns:

- Provide continuing service or minimize downtime for planned maintenance
- Provide HA recovery for unexpected hardware and software failures to entities such as server, network switch, storage subsystem, operating system, application, and so on

- Provide HA recovery in case of disaster involving fire, earthquake, flood, or even human errors

A typical fault tolerant system consists of three layers:

1. Protection for the server with clustering and alternate pathing techniques that allow applications to migrate to functioning server platform or switch path in the event of server or path failure
2. Redundant networking for reliable connections in case of hardware failures, also includes intelligent data and storage management systems
3. Protection for storage with RAID (redundant array of inexpensive disks) technology, cache mirroring, data mirroring and/or remote replication



A highly available SAN can be attained by building redundancy into various layers of the network. Each component in the network design serves as a building block towards attaining this goal. One basic form of high availability is to utilize redundant hardware components within the fabric. While many director class SAN fabric switches tout 99.999% uptime, high availability is a concern on an architectural level and doesn't just stop with the individual platform. Dual-fabric SAN solutions can be deployed to protect against outages caused by single points of failure.

On the server side, a high availability technology called clustering allows redundancy at the application level where multiple servers can be configured in the same cluster to provide backup during node failures. Clustering was not configured for this paper and was considered out of scope for this paper.

While both networking and system HA features can be configured separately, clustering may use multi-pathing to manage the underlying redundant network connections thereby tying the two HA systems together. Clustering provides a server or application level fail-over capability. Thus a standby server in a cluster can be configured for each application on an as-needed basis. A server in a cluster can also be located remotely to protect against disaster at a particular location or building. If a disaster were to occur and disrupts local operations, the entire server cluster can fail-over to a remote facility. Many SAN designs implement tiered and distributed architectures to provide maximum protection. Many techniques including clustering and mirroring provide the necessary intelligence to enable these solutions.

There are several software packages that provide a feature called Multi-Pathing. Typically these packages provide two services: during normal operation mode the data traffic is load balanced across the available interfaces and paths. In the event of a path failure, the software reroutes traffic to the remaining available paths. There are also many situations where more than two paths may exist, however this paper only addresses the scenario involving dual paths. The Multi-Pathing feature is used for redundancy in terms of access to storage devices where a SAN provides multiple logical paths to storage. Redundancy in a path can be constructed via multiple host I/O controllers, switches, cables, and ports on storage devices. Due to the cost impact of redundancy, and stringent network requirements, administrators may choose to implement redundancy at only one component level. At each individual component level there must be robust management and monitoring techniques built into the network so the switchover can occur with minimal downtime.

In a typical multi-pathing implementation, each path may traverse multiple fabrics to complete the connection between initiator and target. Failure anywhere in a chosen path can cause a fail-over event to occur. Thus multi-pathing software must provide proactive monitoring and fast fail-over should an existing utilized path fail.

During a failure event, it is important for either the network recovery mechanisms to maintain access to all devices (targets, LUNs) or the multi-pathing implementation to recognize and recover from any loss in connectivity. The key objective for redundancy design is to maintain access at the application layers and minimize any outages. A combination of multi-pathing software and iSCSI network redundancy is used in the scenarios presented in this paper to ensure true application layer protection.

There are multiple levels of considerations for high availability design including both the network and the systems using the network. For instance, redundant connections throughout the network provide maximum protection against failures and user errors that may potentially cause network outages. In addition, redundant SAN hardware solutions must also provide a built-in mechanism for the recovery from any failed component or service. Security is also an area that must be considered as a breach could heavily impact storage availability. Sharing expensive resources such as storage mandates tight security measures to protect against unauthorized access. Mechanisms such



as zoning, initiator authentication, and access lists ensure that only a designated initiator-target pair is allowed to communicate. This not only maintains security but also protects against data corruption due to inadvertent access by unauthorized servers thereby minimizing downtime.

The 'back-end' of the infrastructure must consist of highly available, scalable information storage subsystems providing scalable capacity, data integrity services, and data replication functionality. Robust software and firmware enable continuous access to all data regardless of application. No single point of failure within the storage subsystem architecture contributes to system high availability. Disk array RAID Levels such as 1, 5 or 0+1/1+0 should be used to guarantee data redundancy to protect vital information from physical disk failures.

Scenarios

This white paper provides an overview of the combination of factors defining high-availability IP storage area networks. It describes methodologies that assist users in architecting and deploying an appropriate level of availability across the application infrastructure. This paper also describes how Cisco solutions assist organizations of any size in achieving application high-availability computing goals.

The Cisco's MDS 9000 Family of Multilayer Switches provides an optional IP Services module providing integrated iSCSI and FCIP services. The IP Services module provides services for building redundancy and enhancing system availability for IP storage networks. The key areas of focus for this paper include:

- Design and high availability configuration using iSCSI and the Cisco IP Services module
- A case study involving an iSCSI SAN design without Multi-Pathing software
- A case study involving an iSCSI SAN design in conjunction with Multi-Pathing software

This paper focuses on the SAN configuration and options. Host clustering was not configured as it was outside the scope of this paper. The storage subsystem implementation and configuration is also not included in this discussion.

Hardware/Software Environment

The following table outlines the hardware and software environment that was used for all tests.

Table 1

Hardware	Versions and Details
Dell PowerEdge 1650 Server	2 Gig of RAM 18GB local disk Windows 2000 Server with SP3 Cisco iSCSI drivers for W2K version 3.1.1 Veritas Volume Manager 3.1 for Windows
Integrated Intel GE NIC	Intel ProSet Utility (Adapter Teaming)
Catalyst 3550	IOS (C3550-I5Q3L2-M)12.1(8) EA1c
Cisco MDS 9216 Fabric Switch	Firmware 1.1(1)
Xyratex RAID	2GB Controller with 8 LUNs




Test Topology 1

In this topology, multiple Cisco MDS 9000 Family features are used pertaining to iSCSI. One of the features is the Virtual Redundancy Router Protocol (VRRP). VRRP provides redundant router gateway services whereby should a Gigabit Ethernet port on the Cisco MDS 9000 IP storage module fail, another Gigabit Ethernet port on a redundant IP Storage module resumes the iSCSI service and its attributes and continues to provide access for affected iSCSI sessions. Another feature provided by the Cisco MDS 9000 Family IP storage module is EtherChannel. EtherChannel enables the combining of two Gigabit Ethernet links into one higher bandwidth logical link. At initial release, EtherChannel is limited to two contiguous Gigabit Ethernet ports on the same IP storage module to protect against link failures. However, in this test scenario, EtherChannel was not used.

The other Cisco MDS 9000 IP storage feature that is used is PWWN aliasing. PWWN aliasing provides recovery capability on the Fibre Channel end of the solution. Using PWWN aliasing, fail-over capability is provided to a redundant Fibre Channel port in the event of a failure on the active Fibre Channel port connecting to the actual physical storage target. The concept is that both Fibre Channel ports have access to the same LUNs and provide redundant paths to the physical storage residing on the Fibre Channel SAN. However, not all storage arrays provide the required active-active LUN capability across multiple storage subsystem interfaces. Please consult your storage vendor for details on this feature.

For some environments, multi-pathing software on the host is too costly to implement for all servers. However, deploying “adapter teaming”, a NIC-based feature on Intel-based cards, and combining this with MDS 9000 features listed above, customers can still deploy a highly available architecture for their low to mid end servers. The “adapter teaming” feature allows two Ethernet network interfaces to act as one logical interface, using only one IP address and one MAC address. The “adapter teaming” feature can run in two modes, active/active or active/passive. Please contact the chosen NIC vendor for further details.

Figure 1
Test Scenario #1 and Failure Scenario Results





Setup and Configuration

The setup and configuration detail for each individual component is listed in Appendix 2.

Results

The following table outlines the results based on a series of staged failures within the SAN. An analysis of each failure scenario is provided in the 'notes' section. As you can see, two scenarios resulted in a completely non-disruptive recovery of I/O from the host perspective.

Table 2


Staged Failure	Active Paths Before Failure	Fail-Over Time	Active Paths After Failure	Recovery Time	Notes
Host N1 NIC failure	2	0s	1	0s	Teaming of Host NICs allows transparent access to storage in case of NIC failure
IP Storage module port failure	2	~20s	2	Manual Switch-back	VRRP fails over to standby IPS GE port and I/O is held for ~20 seconds before recovery is complete
Storage port failure	2	0s	2	0s	pWWN Aliasing of storage port allows recovery to standby storage port transparently
Single MDS switch failure	2	~20s	2	Manual Switch-back	VRRP and pWWN Aliasing will allow access to storage in case MDS switch has failed or IPS Line Card has failed



Test Topology 2

As shown in Figure 2, a simple configuration with dual paths was setup between the server and the storage target. This topology is the same as topology 1 from a physical connectivity standpoint. However, in this topology, 'adapter teaming' was not used for redundancy on the host component and pWWN Aliasing, EtherChannel, and VRRP were also not used. The pWWN Aliasing, EtherChannel and VRRP features can be used for this topology, however they were not required due to the use of multi-pathing software. Multi-Pathing software was utilized to provide HA services for host connectivity and storage access. Using multi-pathing software, all paths to the storage ports can fully be utilized using a load-balancing capability that is included with some software solutions.

Figure 2
Test Scenario #2 and Failure Scenario Results



Setup and Configuration

The setup and configuration detail for each individual component is listed in Appendix 3.



Results

The following table outlines the results based on a series of staged failures within the SAN. An analysis of each failure scenario is provided in the 'notes' section. As you can see, all scenarios leveraging the multi-pathing software resulted in a completely non-disruptive recovery of I/O from the host perspective.

Table 3


Staged Failure	Active Paths Before Failure	Fail-Over Time	Active Paths After Failure	Recovery Time	Notes
Host N1 NIC failure	4	0s	2	0s	With N1 failure, there still exists 2 paths to storage but is coming from N2 to one single IP-P2 GE port to both storage ports.
IPS IP-P1 failure	4	0s	2	0s	With IP-P1 GE port failure, there still exists 2 paths to storage but is coming from N2 to one single IP-P2 GE port to both storage ports.
Storage PWWN1 port failure	4	0s	2	0s	With PWWN1 Storage port failure; there still exist 2 paths coming from both NICs to both IPS GE ports but to one PWWN2 storage port.
Single MDS switch failure	4	0s	1	0s	There is only one path to storage in this scenario b/c both the IPS GE port and Storage port PWWN1 is gone, which leaves only N2 to IP-P2 GE port and PWWN2 storage port.



Test Topology 3

In this topology, the two MDS 9216 switches were not linked to create one fabric. Instead, the two MDS 9216 switches were deployed as two separate fabrics allowing for high availability in the event of a fabric disruption. Note: one could also implement this fabric isolation using the Virtual SANs capability of the MDS 9000 Family of switches. As in topology 2, this topology does not utilize the VRRP, EtherChannel and PWWN aliasing features. Host-based multipathing software was used in this test scenario.

Figure 3
Test Scenario #3 and Failure Scenario Results



Setup and Configuration

The setup and configuration detail for each individual component is listed in Appendix 4.



Results

The following table outlines the results based on a series of staged failures within the SAN. An analysis of each failure scenario is provided in the 'notes' section. As you can see, all scenarios resulted in a completely non-disruptive recovery of I/O from the host perspective.

Table 4

Staged Failure	Active Paths Before Failure	Fail-Over Time	Active Paths After Failure	Recovery Time	Notes
Host N1 NIC failure	2	0s	1	0s	Second Path allowed for continue access to storage
IPS IP-P1 failure	2	0s	1	0s	Second Path allowed for continue access to storage.
Storage PWWN1 port failure	2	0s	1	0s	Second Path allowed for continue access to storage
Single MDS switch failure	2	0s	1	0s	Second Path allowed for continue access to storage

Conclusion

It is quite clear that high availability is a critical design characteristic and needs to be an end-to-end solution. High availability starts with building products that incorporate a high availability feature set for a robust and reliable network foundation. In a network consisting of multiple components like servers, switches, fabrics, and storage devices, redundancy doesn't always guarantee high availability. Network, storage, and server administrators must carefully design the network (LAN and SAN) architecture to consider the impact of failover on applications at all infrastructure points. Cisco MDS 9000 Family provides a highly resilient networking solution with fast recovery capability to both preserve storage transactions transparently and minimize impact at the application layer.

It is shown through the documented test cases that the HA features of Cisco MDS 9000 Family with the IP Services module provides a high availability IP SAN infrastructure that compliments host-based HA applications such as multi-pathing. Given the mission critical nature of many enterprise applications, availability is not only determined by uptime, but also by the ability to deliver consistent performance. The Cisco MDS 9000 Family of switches along with the IP Storage module providing wire-rate iSCSI performance provides the required infrastructure to build multiprotocol highly resilient SAN infrastructures.

References

1. Marc Farley, *Building Storage Networks* (Network Professionals Library)



Appendix 1: Hardware Component Information

Server Setup

In this setup, a SQL server was implemented having two Gigabit Ethernet connections. Each Gigabit Ethernet connection was connected to a separate Catalyst 3500 switch. Veritas Volume Manager for Windows software, version 3.1, was used to control pathing in Topology 2 and Topology 3. For Topology 1, 'Adapter Teaming' technology was utilized from Intel. A simple SQL script was created to validate data inputted into the SQL database stored on the iSCSI target while an injection of errors in the environment was simulated.

Catalyst 3500

In this case, a redundant Ethernet switch configuration was built using dual Cisco Catalyst 3550 switches. Since the switches were used primarily for bridging VLANs, any recent Catalyst software version could be used.

MDS 9216 Configuration


Xyratex Fibre Channel array storage ports were connected on separate MDS 9216 switches for high availability. Each of the MDS IP Storage modules' Gigabit Ethernet ports were connected to each separate Catalyst 3550 switch. The creation of a Fibre Channel port-channel between the two MDS9126 switches for Topology 1 and Topology 2 was done to create one single fabric and one single VSAN. For Topology 3, there were no port-channels created because of the tested logical and physical segregation of fabrics.

Xyratex Storage

A Xyratex RAID controller was used leveraging two Fibre Channel ports for redundancy. Each of the Xyratex FC ports was exporting the exact same LUNs for all topologies in the test scenarios. Please consult your storage vendor for supported configuration for multi-pathing solutions.



Appendix 2: Setup and Configuration for Topology 1



Please refer to the Cisco MDS 9000 Configuration Guide for specific MDS 9000 configuration command details.

Server Setup

Topology 1 required the use of ‘Adapter Teaming’ on the installed Intel NICs. The creation of the adapter team utilized the “Adaptive Load Balancing” mode which allowed for the use of both NICs. Using Cisco’s iSCSI driver version 3.1.1, authentication through CHAP was used. The configuration of the MDS 9000 switch required the enabling of CHAP for server access to storage devices.

Catalyst 3500

A basic configuration of the Catalyst switch was used and trunking of links between the two Ethernet switches was configured.

MDS 9XXX Configuration

The following MDS features were used in this Topology:

- Port Channel between the two MDS switches
- VRRP—Virtual Router Redundancy Protocol
- pWWN Aliasing in conjunction with the creation of a “Virtual-Target”

Below are the steps that were completed to build this topology on the MDS 9216 switches.

MDS 9216-1 (Switch1)

1. Creation of a Port-Channel between the 2 MDS switches was done. Here is the output of this configuration:

```
vsan 1 interface port-channel 1
interface fc1/7
channel-group 1
no shutdown

interface fc1/8
channel-group 1
no shutdown
```



2. Configuration for iSCSI was completed. Below are the necessary steps to complete this task:

- a. First step was to configure the IP Storage module Gigabit Ethernet port. Since we also used VRRP, configuration of VRRP is also done at the interface level. Here is the output of the configuration:

```
interface GigabitEthernet2/1
ip address 10.10.10.3 255.255.255.0
no shutdown
vrrp 1
priority 100
address 10.10.10.1
no shutdown

interface iscsi2/1
mode store-and-forward
no shutdown
```

- b. Other iSCSI parameters were needed to ensure that the iSCSI initiator was able to talk to the MDS. The parameters required for this environment were “import of fc targets”, putting the iSCSI initiator’s IP Address into proper VSAN, and authentication. Here is the output:

```
iscsi authentication none
iscsi import target fc
iscsi initiator name 10.10.10.230
vsan 1
```

- c. Using the PWWN Aliasing feature, we needed to create a “Virtual-Target” on the switch to achieve this. The creation of this “Virtual-Target” was also necessary on the secondary MDS switch because of the VRRP feature we used in this topology. Below is the output of the “Virtual-Target” with PWWN Aliasing:

```
iscsi virtual-target name JBOD1
pwwn 22:00:00:04:cf:e6:e1:5f secondary-pwwn 21:00:00:04:cf:e6:e1:5f

initiator 10.10.10.230 permit
```

3. Creation of the ZoneSet and Zone for the iSCSI initiator and Fibre Channel target was needed to complete the setup of the MDS switch. Below is the output of the zoning configuration:

```
zone name Path1 vsan 1
member pwn 21:00:00:04:cf:e6:e1:5f
member symbolic-nodename 10.10.10.230
member pwn 22:00:00:04:cf:e6:e1:5f

zoneset name ZS1 vsan 1
member Path1

zoneset activate name ZS1 vsan 1
```



MDS 9216-2 (Switch 2):

Below was the configuration needed for Topology 1 on the second MDS 9216 switch.

1. Creation of a Port-Channel between the 2 MDS switches was done. Here is the output of this configuration:

```
vsan 1 interface port-channel 1
```

```
interface fc1/7
channel-group 1
no shutdown
```

```
interface fc1/8
channel-group 1
no shutdown
```

2. Configuration for iSCSI was done next. Below are the steps necessary to complete this task:

- a. First is to configure the IP Storage module Gigabit Ethernet port. Since we were also using VRRP, configuration of VRRP was also done at the interface level. Here is the output of the configuration:

```
interface GigabitEthernet2/5
ip address 10.10.10.2 255.255.255.0
vrrp 1
address 10.10.10.1
no shutdown
```

```
interface iscsi2/5
mode store-and-forward
no shutdown
```

- b. Other iSCSI parameters were needed to ensure that the iSCSI initiator was able to talk to the MDS 9216 switches. The parameters required in this scenario were “import of fc targets”, putting iSCSI initiator’s IP Address into proper VSAN, and authentication. Here is the output:

```
iscsi authentication none
iscsi import target fc
iscsi initiator name 10.10.10.230
vsan 1
```

- c. Using the PWWN Aliasing feature, we needed to create a “Virtual-Target” on the switch:

```
iscsi virtual-target name JBOD1
pwwn 22:00:00:04:cf:e6:e1:5f secondary-pwwn 21:00:00:04:cf:e6:e1:5f
```

```
initiator 10.10.10.230 permit
```


3. Creation of the ZoneSet and Zone for the iSCSI initiator and Fibre Channel target was NOT needed because the activation of the ZoneSet is propagated through the fabric from the primary MDS 9216 switch (switch-1). Nothing needed to be configured on this switch pertaining to zoning.

Xyratex Storage

The verification of access to the same storage (LUNs) through two paths was required.



Appendix 3: Setup and Configuration for Topology 2



Please refer to the MDS Configuration Guide for configuration commands.

Server Setup

Veritas Volume Manager 3.1 for Windows was installed. The Dynamic Multipathing (DMP) feature was verified to be configured properly and operational. Please consult Veritas Volume Manager Documentation pertaining to DMP. There was a registry setting that also needed to be changed to allow a quicker detection of a storage access failure. This entry, if not already added, defaults to the time-out of 90 seconds. Using the default value, the detection of a storage failure can be lengthy which affects the ability of the application to recover. To add/edit this entry, please run “regedt32” and follow these steps:

Go to “HKEY_LOCAL_MACHINE”

Go to “SYSTEM”

Go to “CurrentControlSet”

Go to “Services”

Go to “iscsi”

Go to “Parameters” and see if you have the entry “DisconnectTimeOutValue”

If entry exists, then make necessary changes

If not, then please add this new value.

Click on “Edit” -> “Add Value”

Type in the “Value Name section” -> DisconnectTimeOutValue”

Set “Data Type” to REG_DWORD”

Enter in a value. Default is set to Hex so if you want set the value to 10 seconds, then type in the letter “A” for Hex. Otherwise select “Decimal” and type in the “10”.

Catalyst 3500

A basic configuration of the Catalyst switch was used. However, no trunking of links between the two Ethernet switches was configured or necessary in this topology.



MDS 9216 Configuration

The MDS 9000 PortChanneling feature used between the two MDS 9216 switches provided the capability to zone members on different switches. This allowed for the creation of four paths to the storage from the host, thus increasing availability without utilizing more ports.

Below are the configurations needed to build the environment. The PortChanneling configuration was the same in Topology 1. Please refer to Appendix 2 for the PortChanneling configuration.

MDS 9216-1 (Switch1)

The iSCSI configuration was completed first. Below were the steps necessary to complete this task:

First step was to configure the IP Storage module Gigabit Ethernet port. Here is the output of the configuration:

```
interface GigabitEthernet2/1
ip address 10.10.10.1 255.255.255.0
no shutdown

interface iscsi2/1
mode store-and-forward
no shutdown
```

Other iSCSI parameters were needed to ensure the iSCSI initiator was able to talk to the MDS 9216 switch. The required parameters in this environment were “import of fc targets”, putting iSCSI initiator’s IP Address into proper VSAN, and authentication. Here is the output:

```
iscsi authentication none
iscsi import target fc
iscsi initiator name 10.10.10.230
vsan 1
```

The creation of the ZoneSet and Zone for the iSCSI initiator and Fibre Channel target was needed to complete the setup of the MDS switch. Below is the output of the zoning configuration:

```
zone name Path1 vsan 1
  member pwwn 21:00:00:04:cf:e6:e1:5f
  member symbolic-nodename 10.10.10.230

zone name Path2 vsan 1
  member symbolic-nodename 10.10.10.230
  member pwwn 22:00:00:04:cf:e6:e1:5f

zone name Path3 vsan 1
  member pwwn 21:00:00:04:cf:e6:e1:5f
  member symbolic-nodename 10.10.11.230

zone name Path4 vsan 1
  member pwwn 22:00:00:04:cf:e6:e1:5f
  member symbolic-nodename 10.10.11.230
```



```
zoneset name ZS1 vsan 1
  member Path1
  member Path2
  member Path3
  member Path4
```

MDS 9216-2 (Switch 2)

Below is the relevant configuration needed for Topology 2 on this MDS switch. Again, the PortChannel configuration was the same in this topology as in Topology 1. Please refer to Appendix 2 for configuration.

Configuration steps for iSCSI were completed next. Below were the steps necessary to complete this task:

First was to configure the IP Storage module Gigabit Ethernet port. Here is the output of the configuration:

```
interface GigabitEthernet2/5
ip address 10.10.11.1 255.255.255.0
no shutdown

interface iscsi2/5
mode store-and-forward
no shutdown
```

Other iSCSI parameters were needed to ensure that the iSCSI initiator was able to talk to the MDS 9216 switch. The required parameters in this environment were “import of fc targets”, putting iSCSI initiator’s IP Address into proper VSAN, and authentication. Here is the output:

```
iscsi authentication none
iscsi import target fc
iscsi initiator name 10.10.11.230
vsan 1
```


Creation of the ZoneSet and Zone for the iSCSI initiator and Fibre Channel target was NOT needed because the activation of the ZoneSet is propagated through the fabric from the primary MDS 9216 switch (switch-1). Nothing needed to be configured on this switch pertaining to zoning.

Xyratex Storage

The verification of access to the same storage (LUNs) through two paths was required.



Appendix 4: Setup and Configuration for Topology 3



Please refer to the MDS Configuration Guide for configuration commands.

Server Setup

The server configuration did not change from Topology 2. Please refer to Topology 2 Appendix 3 for configuration details.

Catalyst 3500

A basic configuration of the Catalyst switch was used with no trunking of links between the two Ethernet switches.

MDS 9216 Configuration

In this topology, logical and physical separation of the fabric was configured. The following was the configuration necessary to build this environment:

MDS 9216-1 (Switch 1)

1. The configuration for iSCSI did not change from Topology 2 to Topology 3. Here are the relevant configuration details:

- a. First step was to configure the IP Storage module Gigabit Ethernet port. Here was the output of the configuration:

```
interface GigabitEthernet2/1
ip address 10.10.10.1 255.255.255.0
no shutdown

interface iscsi2/1
mode store-and-forward
no shutdown
```

- b. Other iSCSI parameters were needed to ensure that the iSCSI initiator was able to talk to the MDS 9216 switch. The required parameters in this environment were “import of fc targets”, putting iSCSI initiator’s IP Address into proper VSAN, and authentication. Here is the output:

```
iscsi authentication none
iscsi import target fc
iscsi initiator name 10.10.10.230
vsan 1
```

2. Creation of the ZoneSet and Zone for the iSCSI initiator and Fibre Channel target was needed to complete the setup of the MDS switch. Below is the output of the zoning configuration:

```
zone name Path1 vsan 1
  member symbolic-nodename 10.10.10.230
  member pwwn 22:00:00:04:cf:e6:e1:5f
zoneset name ZS1 vsan 1
  member Path1
```

MDS 9216-2 (Switch 2)

1. The configuration for iSCSI was completed first. Below are the steps necessary to complete this task:

- a. First step was to configure the IP Storage module Gigabit Ethernet port. Here was the output of the configuration:

```
interface GigabitEthernet2/5
ip address 10.10.11.1 255.255.255.0
no shutdown

interface iscsi2/5
mode store-and-forward
no shutdown
```

- b. Other iSCSI parameters were needed to ensure that the iSCSI initiator was able to talk to the MDS 9216 switch. The required parameters in this environment were “import of fc targets”, putting iSCSI initiator’s IP Address into proper VSAN, and authentication. Here is the output:

```
iscsi authentication none
iscsi import target fc
iscsi initiator name 10.10.11.230
  vsan 1
```

2. Since this switch was its own fabric, creation of the zoning information was necessary for the iSCSI initiator to access the secondary path. Below is the configuration in this environment for this switch:

```
zone name Path2 vsan 1
  member symbolic-nodename 10.10.11.230
  member pwwn 21:00:00:04:cf:e6:e1:5f

zoneset name ZS1 vsan 1
  member Path2
```

Xyratex Storage

The verification of access to the same storage (LUNs) through two paths was required.



Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the

Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland
Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland
Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

All contents are Copyright © 1992–2003 Cisco Systems, Inc. All rights reserved. Cisco, Cisco IOS, Cisco Systems, and the Cisco Systems logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company.
(0303R) EW/LW4449 0403