

iSCSI Design Using the MDS 9000 Family of Multilayer Switches

Introduction

As enterprises migrate from DAS to SAN environments, and the need to consolidate enterprise storage resources increases, there is high demand for extending the consolidation effort to mid-range and low end application servers. In addition, the need to extend the reaches of a consolidated SAN over metro and wide area networks becomes a necessity. Cisco's MDS 9000 Family of Multilayer Directors and Fabric Switches provide enterprises with the ability to build large-scale Fibre Channel SANs and extend these SANs to mid-range servers and metro and wide area networks. By using FCIP and iSCSI protocols, enterprises can now leverage Ethernet and IP technologies to further extend their storage environment and continue to realize the cost savings derived from storage consolidation. Using the 16-port non-blocking Fibre Channel (FC) switching module or the 32-port shared-bandwidth Fibre Channel switching module, enterprises can attach their storage devices, tape libraries and host bus adapters to build up to 224 ports into a single switch. With the addition of the IP Services switching module providing 8 Gigabit Ethernet ports for iSCSI and FCIP services, enterprises can extend the SAN to other low to mid range servers with iSCSI or connect SAN islands over IP via the FCIP protocol. Using all of the options available in the Cisco MDS 9000 Family of switches, large-scale, high-port density SANs become reality. Customers may use their existing IP infrastructure along with their in-house IP expertise to optimize enterprise storage consolidation. Management of the enterprise SAN is also made simpler with the extensive multiprotocol management features of the Cisco MDS 9000 Family.

This design guide will focus on the aspects of extending the SAN utilizing the iSCSI protocol within the Cisco MDS 9000 IP Services switching module. Design considerations and typical implementations will be discussed to guide end users on how to implement an iSCSI solution in the enterprise with Cisco's MDS 9000 IP Services switching module. This paper will not discuss configuration of applications servers pertaining to the MDS implementation of iSCSI and is out of scope of this paper. For specific application notes

for the MDS implementation of iSCSI, please refer to the Cisco Connection Online website at:

<http://www.cisco.com/go/storagenetworking>.

iSCSI Basics

The iSCSI protocol is designed to carry the SCSI protocol using TCP/IP. Conceptually, iSCSI+TCP+IP provides a similar transport model to serial Fibre Channel Protocol (FCP) which also transports SCSI. The basic idea of iSCSI is to leverage an investment in existing IP networks to build and extend



Storage Area Networks (SANs). This is accomplished by using the TCP/IP protocol to transport SCSI commands, data, and status between hosts or initiators and storage devices or targets such as storage subsystems and tape devices.

Traditionally SANs have required a separate dedicated infrastructure to interconnect hosts and storage systems. The primary transport protocol for this interconnection has been Fibre Channel (FC). Fibre Channel networks provide primarily a serial transport for the SCSI protocol. In addition, IP data transport networks have been built to support the front-end and back-end of IP application servers and their associated storage.

Unlike IP, Fibre Channel cannot be easily transported over lower bandwidth long distance WAN networks in its native form and therefore requires special gateway hardware and protocols. The use of iSCSI over IP networks does not necessarily replace a FC network but rather provides a transport for IP attached hosts to access Fibre Channel based targets.

IP network infrastructures provide major advantages for interconnection of servers to block-oriented storage devices. Primarily, IP storage networks offer major cost benefits as Ethernet and its associated devices are significantly less expensive than the Fibre Channel equivalents. In addition, IP networks provide enhanced security, scalability, interoperability, and network management over a traditional Fibre Channel network.

IP network advantages include:

- General availability of network protocols and middleware for the management, security, and quality of service (QoS)
- Applying skills developed in the design and management of IP networks to IP storage area networks. Trained and experienced IP networking staffs are available to install and operate these networks
- Economies achieved from using a standard IP infrastructure, products, and service across the organization
- iSCSI is compatible with existing IP LAN and WAN infrastructures
- Distance is only limited to application performance requirement, not by the IP protocol

Value of iSCSI

By building on existing IP networks, users are able to connect hosts to storage facilities without additional host adapters. In addition, iSCSI SANs offer better utilization of storage network resources and eliminate the need for separate parallel WAN and MAN infrastructures. Since iSCSI uses TCP/IP as its transport for SCSI, data can be passed over existing IP based host connections commonly via Ethernet. Additional value can be realized by being able to better utilize existing FC back-end storage resources. Since hosts can utilize their existing IP/Ethernet network connections to access storage elements, storage consolidation efforts can now be extended to the mid-range server class at a relatively lower cost while improving the utilization and scalability of existing storage devices.

iSCSI Standards Track

The iSCSI standard is one of several protocols continually developed and delivered by the IP Storage (IPS) working group in the IETF. The IP Storage working group continues to work on new services including enhanced security services, directory services, and diskless client boot services. In addition, because iSCSI mainly uses Ethernet, interoperability of the transport protocol is well established in the networking industry. This fact removes one major hurdle that Fibre Channel still suffers from even today.




iSCSI Terminology and Protocol

The iSCSI standard uses the concept of a *Network Entity* which represents a device or gateway attached to an IP network. This *Network Entity* must contain one or more *Network Portals* providing the actual connection to the IP network. An *iSCSI Node* contained within a *Network Entity* can utilize any of the *Network Portals* to access the IP network. The *iSCSI Node* is an iSCSI initiator or target identified by its iSCSI Name within a *Network Entity*. For iSCSI, the SCSI device is the component within an iSCSI Node that provide the SCSI functionality. There is exactly one SCSI Device within an *iSCSI Node*.

A *Network Portal* is essentially the component within the *Network Entity* responsible for implementing the TCP/IP protocol stack. Relative to the initiator, the *Network Portal* is identified solely by its IP address. For an iSCSI target, its IP address and its TCP listening port identify the *Network Portal*. For iSCSI communications, a connection is established between an initiator *Network Portal* and a target *Network Portal*. A group of TCP connections between an initiator *iSCSI Node* and a target *iSCSI Node* make up an *iSCSI Session*. This is analogous to but not equal to the SCSI I_T Nexus.


Figure 1
iSCSI Client/Server Architecture



The iSCSI protocol is a mapping of the SCSI Initiator and Target (Remote Procedure Call, Reference SCSI Architecture Model, SAM) model to the TCP/IP protocol. The iSCSI protocol provides its own conceptual layer independent of the SCSI CDB information it carries. In this fashion SCSI commands are transported by iSCSI requests and SCSI response and status are handled by iSCSI responses. Also, iSCSI protocol tasks are carried by this same iSCSI request and response mechanism.



Figure 2
iSCSI Protocol Model



Just as with the SCSI protocol, iSCSI employs the concepts of an *initiator*, *target*, and communication messages called *protocol data units* (PDU). Likewise, the iSCSI transfer direction is defined relative to the *initiator*. As a means to improve performance, iSCSI allows a “phase-collapse” enabling a SCSI command or response and its associated data to be sent in a single iSCSI PDU.

Cisco MDS 9000 Family IPS Implementation of iSCSI

iSCSI Naming and Addressing

An *iSCSI Node Name* is location-independent in that it does not contain an IP address, a globally unique address, or a permanent identifier for an iSCSI initiator or iSCSI target node. This makes it reachable via multiple network interface or network portals. There are two types of naming conventions based on the iSCSI standard: *iSCSI Qualified Name* (*iqn*) and the *EUI format*. The Cisco MDS 9000 Family with the IP Storage switching module implements both types of the naming formats. However, the most commonly used naming method is the *iqn* naming format.

An *EUI* name comprises an *eui*, extended unique identifier, followed by a unique 64-character string. The 64-character string is the same name used in a Fibre Channel Worldwide Name (WWN). An example of this format is: *eui.02004567A425678D*.

An *IQN* name comprises an *iqn* key word followed by a qualified domain name. An example of this format is: *iqn.5886.com.acm.diskarrays-sn-a8675309*.

Management or support tools use the iSCSI address format to identify an iSCSI node. An iSCSI address ties the node name to the network address where it can be accessed. An example of an iSCSI address is:

iSCSI://172.16.1.1:3260/eui.02004567A425678D or *iSCSI://172.16.1.1:3260/iqn.com.acme.diskarrays.jbod1*




VLANs

On the MDS IPS module, Virtual LANs (VLANs) are supported. Virtual LANs (VLANs) create multiple virtual layer 2 networks over a single physical LAN. VLANs provide traffic isolation, security, and broadcast control. Each Gigabit Ethernet port can be configured as a trunking port and uses the IEEE 802.1Q standard tagging protocol for VLAN encapsulation.

iSCSI Access Methods

The iSCSI access method for the Cisco MDS 9000 iSCSI implementation is for iSCSI initiators to communicate with Fibre Channel targets. This is the first implemented mode. The reverse of this mode will be included as a future software feature.

Figure 3
iSCSI Access Method



To understand this access method, it is important that the concept of an *FV_Port* be introduced. The *FV_Port* is a logical port created by the IP Storage switching module for the purpose of forwarding frames between the Gigabit Ethernet and the Fibre Channel devices. Just as each physical FC port on the Cisco MDS 9000 Family negotiates to become an F_Port, FL_Port, E_Port or TE_Port and able to forward FC frames based on the hardware index assigned to this port, each of the Ethernet ports on the IP Storage switching modules require a similar index.

iSCSI initiator to access FC target

There are 4 basic steps required for an iSCSI initiator to be able to access FC targets through the MDS 9000 Family switch. A sample step-by-step configuration is shown in appendix A.

1. Configure the MDS 9000 IP Storage switching module for iSCSI access
2. Configure the iSCSI initiator node name or IP address and add it into a valid VSAN
3. Create iSCSI targets and map them to FC targets
4. Configure a FC zone containing the iSCSI initiator and FC target(s)

Configuring MDS 9000 IP Storage Switching Module for iSCSI

The first step is to configure the IP address for iSCSI clients to access. One can configure the Gigabit Ethernet ports with different parameters, such as MTU size, authentication mode etc. Once the Gigabit Ethernet ports have been configured, one will then need to enable each required port specifically as an iSCSI port. Since within the MDS 9000 IP Storage switching module the Gigabit Ethernet ports can support both iSCSI and FCIP simultaneously, it is necessary to enable each required Gigabit Ethernet port to specifically run iSCSI.



Configuring an iSCSI Initiator, IP Address, and VSAN

Depending on the iSCSI driver, one can configure a unique iSCSI initiator node name. If one does not statically assign one, the driver will automatically create a unique iSCSI node name. If the node name is dynamically created, the iSCSI initiator must login at least once to the MDS 9000 IP Storage switching module to allow recognition of the assigned node name. This node name is required so it can be added into the proper VSAN and zoned accordingly. In the MDS 9000 IP storage implementation, iSCSI initiators are allowed to span across multiple VSANs thus being able to access any FC targets on any VSAN.

The MDS 9000 IP Storage switching module iSCSI implementation also allows for zoning by IP address. Prior to configuring any zoning, adding the initiator's IP address into the specific VSAN is required. As with iSCSI initiators spanning multiple VSANs, the IP Address can span across multiple VSANs as well.

Creation of iSCSI Targets and FC Targets

The iSCSI initiator does not directly attach to Fibre Channel targets. An iSCSI initiator only connects to iSCSI *virtual* targets created as a representation of one or more Fibre Channel targets. To enable this function, the MDS IP Storage switching module must perform the conversion of Fiber Channel target(s) into iSCSI target(s) by advertising all available Fibre Channel targets to the iSCSI initiator in the IQN-format. The IP Storage module does this by pre-pending Fibre Channel WWNs with the desired *iqn* string. The Fibre Channel WWN of a target is learned by the IP Storage switching modules through a basic Fibre Channel name server query. These iSCSI targets are then made available to the iSCSI initiator when a *SendTargets* iSCSI command is received by the MDS 9000 IP Storage switching module from an iSCSI initiator.

There are two modes of operation to create Fibre Channel targets which can be exported as iSCSI targets. Creation of iSCSI targets can be done dynamically, the preferred method, or configured statically through the creation of virtual iSCSI targets. Essentially, a virtual target is defined manually through the process of target and LUN mapping from Fibre Channel to iSCSI. By creating virtual targets, an explicit target name is given to the initiators which they can use to access specific Fibre Channel target and specific LUN(s).

For the FC target devices in the SAN, an IP Storage switching module portrays an iSCSI initiator as an N_Port device in the SAN with its own FC_ID assigned by the SAN and an associated pWWN.

To represent FC target in iSCSI, each IP Storage module Gigabit Ethernet port advertises an iSCSI target as “iqn.xxx” with its own portal group tag (PGT). The group tag is unique within the physical switch.

Zoning of iSCSI Initiators or IP Addresses with FC Targets

By utilizing zoning capabilities within the fabric, iSCSI initiator node names and/or IP addresses can be added to a zone like any other Fibre Channel entity connected to the Fibre Channel fabric. This implementation provides extreme flexibility, especially in multi-pathing environments. The Fibre Channel standard allows the zoning of a *symbolic node-name*, which represents iSCSI initiators or IP addresses. Like any Fibre Channel initiator, iSCSI initiators can be in multiple zones.



Access Control

Access control in a traditional Fibre Channel SAN is achieved by implementing zoning services. With the introduction of VSANs in the Cisco MDS 9000 Family, both VSANs and zoning are used for access control. VSANs are used to divide the physical Fibre Channel SAN into logical fabrics. This functionality is very analogous to the role provided by VLANs in an Ethernet environment. Zoning services provide the ability to restrict communication between various endpoints within a VSAN. Each VSAN has its own set of zoning services.

Fibre Channel or iSCSI initiators only access Fibre Channel or iSCSI targets that are in the same zone and within the same VSAN. With the MDS implementation of iSCSI, an iSCSI initiator is not limited to any particular VSAN. Instead, an iSCSI initiator can be configured to be included in any VSAN of choice. This flexibility allows the iSCSI initiator to access any Fibre Channel device on any VSAN of the network if configured to do so.

Besides the normal access control, iSCSI also implements IP-based authentication mechanisms to restrict access to any targets. The authentication procedure occurs at the iSCSI login stage. The authentication algorithm implemented by the Cisco MDS 9000 Family of switches is the common Challenge Handshake Authentication Protocol (CHAP). Authentication can also be disabled if desired although not recommended. Other authentication algorithms such as SRP, Public Key method (SPKM-1 or 2) can also be used by iSCSI and will be implemented in future software releases.

iSCSI LUN Mapping

The Cisco MDS 9000 implementation of iSCSI supports advanced LUN mapping functionality to increase the availability of the physical disk and provide a high level of flexibility. The following are the methods of LUN mapping available:

- Map LUNs of different FC targets to one iSCSI virtual target (supported in future release)
- Map subsets of LUNs of one FC target to multiple iSCSI virtual targets

Many storage arrays support capabilities enabling many LUNs to be visible from one Fibre Channel target port. Having the capability of LUN masking/mapping of a Fibre Channel target to multiple logical iSCSI Virtual Target(s) provides flexibility to the IT administrator. This flexibility enables the logical division of the expensive disk array resources with huge volumes into multiple iSCSI targets which can be used by different iSCSI user groups. Previously, this was only accomplished through LUN masking and mapping on a disk array controller. However, with the Cisco MDS 9000 IP Storage switching module, this functionality can be achieved in the network. This feature also provides added security in term of access control. If an iSCSI host is not specifically allowed to access the logical iSCSI LUNs determined through the authentication process, access is denied.

iSCSI High Availability

The Cisco MDS 9000 iSCSI implementation supports iSCSI redundancy capabilities to increased high availability. These redundancy capabilities include EtherChannel and the Virtual Router Redundancy Protocol (VRRP). EtherChannel allows the bundling of multiple physical Ethernet links into a single higher bandwidth logical link. At initial release, EtherChannel only supports two contiguous links in an EtherChannel bundle which are required to be on the same IP Storage switching module. Full support of the 802.3ad port aggregation standard will be provided in a future software release. VRRP allows for the creation of a virtual IP Address (layer 3) and a virtual MAC address (layer 2) pair to be shared across multiple Ethernet gateway ports. The Cisco MDS 9000 Family iSCSI



implementation supports VRRP across multiple ports on the same or different physical MDS 9000 switches or IP Storage switching modules. If the VRRP function is invoked due to a gateway failure, TCP session(s) information is not synchronized which requires iSCSI initiators to re-establish a connection to the standby switch or gateway port.

Securely Integrating an iSCSI Host into a Fibre Channel SAN

The Cisco MDS 9000 Family of switches, with their industry-leading availability, scalability, security and high performance architecture also enable the extension of SANs to the IP world with the availability of the IP Storage switching module. Fibre Channel storage connected to a fabric based on the MDS 9000 Family can be extended to mid-range servers that do not have Fibre Channel Host Bus Adapters (HBA) through the use of the iSCSI protocol. Servers with a 10/100Mbps or Gigabit Ethernet NIC, or for higher performance requirements using a TCP Offload Engine (TOE) NIC card can now access Fibre Channel storage. Combined with the support of FCIP in the IP Storage switching module, the Cisco MDS 9000 family is a truly industry-leading integrated multi-protocol switching platform.

Fibre Channel security mechanisms such as VSANs and zoning inherent in the MDS 9000 Family are augmented with the use of added security capabilities provided by iSCSI and its associated services. iSCSI additional security services such as iSCSI initiator authentication through CHAP extends SAN security measures to securely incorporate iSCSI hosts. The flexibility of creating iSCSI “virtual-targets” provides LUN-level granularity in assigning Fibre Channel storage to iSCSI initiators. This capability is especially useful in scenarios where many iSCSI initiators with low I/O requirements need access to storage through a single Fibre Channel storage array interface.

Using the iSCSI protocol as a transport for the block-oriented SCSI protocol, many low to mid-range servers can now be incorporated into the SAN and centrally managed. Today, many such servers use Direct Attach Storage (DAS) and are difficult to scale properly and don't fully utilize their storage resources. For example, Server-A and Server-B may both have 100GB of direct attach storage. However, Server-A may only utilize 30% of it's storage and Server-B is at 90%. With DAS, one cannot easily migrate the under-utilized storage on Server-A to Server-B where it is needed. A Fibre Channel SAN would be an obvious solution to facilitate sharing of the storage resources, however many enterprises do not opt for a SAN due to the excessive port costs often prohibitive to such low and mid-range servers. Also, the typical I/O requirement for such servers is low, between 5MBps – 30MBps, and doesn't justify the migration to Fibre Channel networks. Now with the iSCSI protocol and Cisco's MDS 9000 iSCSI implementation, one can enable these types of servers to join the SAN easily and in a more cost effective manner. With the bandwidth provided by a Gigabit Ethernet link along with the often lower I/O requirement of iSCSI servers, one may be able to connect many iSCSI servers to a single Gigabit Ethernet port. With the 8 Gigabit Ethernet ports provided by the IP Storage switching module, scaling iSCSI clients is made even easier. Utilizing servers network interface card (NIC), either 10/100Mbps or Gigabit Ethernet, and iSCSI drivers provided by Cisco and Microsoft for the Windows platform, such servers can fully realize the benefits of a SAN.

With the addition of iSCSI to the IP stack within an iSCSI initiator, the iSCSI client's CPU will need to do additional processing to transmit and receive iSCSI packets and maintain iSCSI sessions. Therefore, iSCSI may potentially increase the overall CPU utilization of the system. To assist the system with this additional processing, some traditional HBA and network vendors have built iSCSI host bus adapters known as TCP Offload Engines (TOE Cards). Most vendors provide their own iSCSI drivers for their TOE cards for different platforms. Some vendors provide total offload capability of the iSCSI stack from the host CPU and others simply provide the offload of the TCP stack only.



iSCSI Performance Benchmarking

The performance of the IP Storage switching module for the Cisco MDS 9000 Family was measured using a well known tool, IOMeter. The purpose of this section is to illustrate the impact of different I/O patterns on the performance of iSCSI on the IP Storage services module. The various benchmark tests utilize different I/O patterns with different block sizes and different percentages of reads and writes.

Test Configuration

The following section outlines the test configuration used to collect the results outlined in this paper.


Server: Windows Dell 1650 with Embedded GE NIC, 1.13 GHz CPU, 2GB RAM, Windows 2000 Server SP3. Server: Cisco's iSCSI driver version 3.1.1 and a Qlogic 2300 Fibre Channel host bus adapter was used for baseline. A third party TOE card vendor was used that did TCP offload not full iSCSI offload.

Storage: Xyratex 2Gig RAID Controller Storage with 8 73GB 10K RPM drives

Switch: Cisco MDS9216 with an IP Storage switching module running version 1.1.(1)

The Xyratex storage array was connected to the MDS 9000 Family switch and the servers were connected to the MDS 9000 Family switch using a QLogic 2300 host bus adapter configured for 1Gbps operation. The LUNs on the Xyratex array were created as RAID 0 LUNs spread over 8 independent disks. The test was conducted on the disks with the NTFS file systems for Windows.

Figure 4
iSCSI Test Scenario




I/O Size Number of Threads: 4KB, 16KB, 64KB, 128KB, 512KB

Test Results

Detail test results are located in Appendix B.




Figure 5
IOPs Comparison—100% Reads – 100% Sequential



The number of I/Os per second in the different tests shows that as block sizes increase, the gap between the number of I/Os in the test scenarios decreases. Since iSCSI adds additional overhead to the CPU, the smaller the block size, the more CPU resources are required thereby explaining the I/O gap between FC and iSCSI.


Figure 6
IOPs Comparison—100% Writes – 100% Sequential



The write performance as shown by this diagram indicates all three test scenarios are quite comparable. It should be noted that with the smaller number of drives used in this test, there weren't enough spindles to saturate the FC HBA or the iSCSI TOE card from a CPU perspective. More spindles will support more I/O and consume more of the unused CPU.




Figure 7
Throughput Comparison—100% Reads – 100% Sequential



Looking at the diagram, iSCSI performs equally if not better on reads with larger block sizes. The throughput is affected with smaller block sizes in the different tests because of the higher CPU utilization needed for iSCSI.

Figure 8
Throughput Comparison—100% Writes – 100% Sequential



In this diagram, writes throughput shows iSCSI can perform equally if not better than Fibre Channel. With the smaller block size, throughput can be negatively affected due to the small number of drives and their inherent I/O processing capabilities. If more drives are added to the scenario on the back-end, performance will even further increase.



Figure 9
CPU Comparison—100% Reads – 100% Sequential





Figure 10
CPU Comparison—100% Writes – 100% Sequential



In both of the diagrams above, since iSCSI increases overhead on the CPU, the diagram shows the difference on CPU utilization between the tests. With TCP Offload Engines to alleviate CPU utilization, this CPU overhead is reduced. TOE card vendors that perform full iSCSI offload, the CPU utilization would decrease even further.




Conclusion

Enterprise environments now have the ability to create large Fibre Channel SANs with the MDS 9000 Family. However, utilizing the MDS 9000 Family IP Storage switching module, highly available and scalable multi-protocol SANs that support FCIP and iSCSI can be deployed. The Cisco MDS 9000 Family delivers a multi-protocol SAN enterprise solution providing high availability, scalability, and easier manageability for the Enterprise. With the capability of extending the SAN to low and mid-range servers, storage managers can now fully utilize the benefits of the SAN throughout their application environments and to all application servers. The ability to incorporate low and mid-range application servers into a centralized SAN utilizing an existing IP infrastructure provides a complete overall storage solution for the enterprise and an excellent return on investment.

Appendix A

Below is a sample configuration involving a basic iSCSI initiator connection to a Fibre Channel target. Using the following diagram, directions are provided on how to configure iSCSI on the MDS 9000 Family IP Storage switching module. With this basic configuration, all the initiators and storage ports are in VSAN 1, which is the default VSAN.

Figure 11
iSCSI Sample Configuration



The following steps are required in order for the above server to access the Fibre Channel storage. Prior to configuring iSCSI, the Fibre Channel storage must be connected on the MDS on module fc1/1 and enabled.

1. Configuration of the IP Storage switching module Gigabit Ethernet port for iSCSI access in VLAN 5:

```
interface GigabitEthernet2/1.5
ip address 10.10.11.30 255.255.255.0
no shutdown
```

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

2. In this section, zoning is performed by IP address. Therefore, the iSCSI initiator's IP address must be added into VSAN 1 where the storage resides:

```
iscsi initiator name 10.10.11.230
vsan 1
```

3. In this section, the dynamic creation of FC targets into iSCSI targets is enabled. Also, CHAP authentication is enabled. Here is the output of the configuration:

```
iscsi authentication chap
iscsi import target fc
username cisco password 7 fewhg1xnkfy1sewsm1 iscsi
```



4. With the above steps completed, one now needs to zone the iSCSI initiator's IP Address and the Fibre Channel storage into a zone. Here the configuration:

```
zoneset name ZS1 vsan 1
member Path1
zoneset activate name ZS1 vsan 1
zone name Path1 vsan 1
member pwn 21:00:00:04:cf:e6:e1:5f
member symbolic-nodename 10.10.11.230
```

5. Since the iSCSI initiator's IP Address is in a different subnet then the IP Storage switching module Gigabit Ethernet address, one needs to create a static route for the initiator to talk to the MDS 9000 IP Storage switching module.

The following is the configuration:

```
ip route 10.10.11.0 255.255.255.0 10.10.1.2
```

Appendix B

The following charts contain the actual performance results gathered from the successive tests run against the test infrastructure.

100% Reads - 100% Sequential

IOPS	FC	GE	TOE
4KB	22517.75	11275.21	13815.29
16KB	6076.81	5809.4	6900.96
64KB	1555.13	1410.71	1407.68
128KB	784.87	699.31	709.1
512KB	196.33	165.58	187.49

100% Writes - 100% Sequential

IOPS	FC	GE	TOE
4KB	9568.51	9253.31	9332.11
16KB	5954.32	6655.51	6304.96
64KB	1490.47	1718.75	1763.09
128KB	760.38	828.39	853.25
512KB	190.69	204.66	206.27



100% Reads - 100% Sequential

Throughput	FC	GE	TOE
4KB	87.96	44.04	53.97
16KB	94.95	90.77	107.83
64KB	97.2	88.17	87.98
128KB	98.11	87.41	88.64
512KB	98.15	82.79	93.74

100% Writes - 100% Sequential

Throughput	FC	GE	TOE
4KB	37.35	36.15	36.45
16KB	93.02	103.99	98.52
64KB	93.15	107.42	110.19
128KB	95.05	103.55	106.66
512KB	95.33	102.33	103.13

100% Reads - 100% Sequential

CPU	FC	GE	TOE
4KB	57.32	99.56	69.28
16KB	19.55	99.39	45.53
64KB	8.21	86.17	10.41
128KB	5.54	85.32	11.12
512KB	3.88	83.28	8.23

100% Writes - 100% Sequential

CPU	FC	GE	TOE
4KB	22.21	83.71	68.99
16KB	16.32	92.43	43.09
64KB	6.13	53.95	15.78
128KB	4.13	41.64	5.16
512KB	3.77	39.05	4.74



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 and VCO 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/LWX4449 0403