White Paper

# Backup and Recovery Solutions with the MDS 9000 Family

## Purpose

The purpose of this whitepaper is to discuss backup and recovery architectures and solutions and outline the applicable features of the Cisco MDS 9000 Family of Multilayer Directors and Fabric Switches.

## Introduction

Global enterprises with mission critical data residing on their servers demand continuous availability for their applications. Applications such as supply chain management (SCM), enterprise resource planning (ERP), and customer relationship management (CRM) are creating voluminous amounts of data that must be protected at all costs. At a minimum, this data must be backed up to tape regularly as an insurance policy against a potential loss of data. However, growing data volumes requiring larger storage capacity, faster servers, and also require longer time windows for backup. One must also consider that data that takes several hours to backup will also take the same length of time for a full restore should the need arise. This restore time is often unacceptable as it translates into lost revenue due to extended downtime. Therefore, in many cases, tape backup is considered a minimum level of disaster recovery (DR) planning.

In order to ensure 99.999% uptime required by these enterprise applications, a storage design must incorporate additional high availability considerations at every level. A disaster recovery plan, imperative for all enterprises, must address this concern of potential extended outages and provide seamless failover to a secondary site during major outages. Corporations often utilize replication technology to remotely replicate a whole data center, in addition to tape backup, in their DR plan. Therefore, a recovery now can include a data center fail over to a live remote location in addition to a data restore from tapes. Disasters can be caused by a myriad of factors and are difficult to predict. Some key scenarios are listed below:

- Equipment failure
- Application failure
- Human error
- Natural and unnatural disasters

Each enterprise must prepare to recover from disasters and identify all critical data that must be preserved for continuous access. Business impact and risk analysis must be performed to identify locations, functions, or applications most critical to the enterprise. A remote data center, a mirrored image of the primary one, is used to provide full access after a major disaster. Many DR solutions involve keeping real-time replicated images of data in

conjunction with backing up to tape. Replication technologies also provide multiple options suitable for varying application requirements. Although replication can help you to recover from a catastrophic failure faster, it does have a limitation of replicating corrupt data along with valid data. Therefore, there will always be a need for tape backup as a way of archiving valid data. This paper primarily focuses on tape backup technologies, architectures, and options as a component of an overall DR plan.

## Tape Backup

In today's enterprise environment, most application servers are directly attached to a dedicated tape drive via a parallel SCSI connection. Dedicated resources are expensive to both deploy and maintain since the number of tape devices to be managed increases in direct proportion to application servers. However, directly-attached tape drives guarantee performance because that server is the only one who uses the drive. Cost considerations caused a migration to network backup models where tape drives were placed on a LAN and shared among multiple servers. A typical LAN based backup scenario is shown in the figure below. In this model both data and backup traffic traverse the same LAN. The networked model for backup optimizes tape utilization and enhances manageability but introduces concerns listed below.

Figure 1

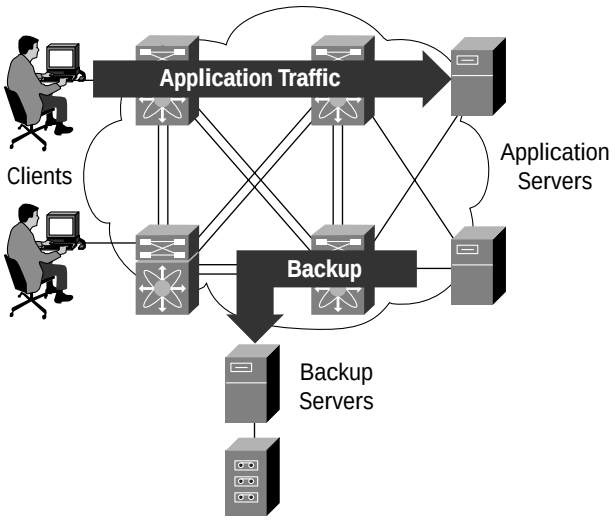


First, large volumes of data being backed up increases traffic on the LAN and may cause degradation in application performance. Backups are generally performed off-hours in order to minimize interruption to data traffic. Growing data volume leads to a longer backup window that can potentially extend into business hours. Globalization of enterprises continues to shrink available time windows for backup due to 24x7 uptime requirements. Secondly, sharing LANs for backup and application traffic may result in backup interruptions causing backup jobs to fail altogether. Thirdly, backup and data applications sharing the same LAN can often prove costly as a firmware upgrade or instability in one application environment can lead to an outage in the other as well. To alleviate these potential conflicts in a common LAN, administrators proposed a separation in application and backup domains. In newer implementations, customers are migrating towards LAN-free architectures to segregate backup traffic from applications as described below. A large number of customers have started to deploy dedicated storage networks for backup.

Current implementations of backup processes are manual, labor intensive, and inundated with problems caused by human errors. Backups that fail can often go undetected leading to potential data loss following a failure. Tapes must be manually inserted, rotated, and removed for off-site transportation to ensure recoverability. Due to high expenses associated with manual administration, a case for automation can be made where robots are used to improve tape management. Centralized backups group tapes together by pools whereby a number of backups can be multiplexed onto a tape. As a tape is filled the backup continues using another free tape within the pool which significantly enhances manageability. A reduction in the number of tape drives to be managed leads to cost reduction.

Backups require an increase in application server activity in order to fetch data from disk and write to tape. Application servers are usually busy processing large volumes of latency and performance sensitive data. The extra CPU cycles consumed for data movement and scheduling while conducting backups can often prove costly to the application itself. A server-free solution is targeted at eliminating the performance impact of a backup when performed by the application server itself. This architecture migrates the data-mover and connection-broker roles to a dedicated backup server. A robust backup solution must address the following concerns:

1. Optimize use of backup resources including tapes, drives, and operating time
2. Minimize impact to application traffic
3. Segregate backup domains to reduce CPU overhead on application servers

The following section provides details on most common implementations of backup solutions.

### LAN-Free Backup

LAN-ree backup enables each application server to move data directly to a tape device over the storage network without going through a dedicated backup server. Using a commonly known shared-storage option, each application server acts as a media server in that it moves backup data directly to tape. Each server then arbitrates for a tape drive and reserves that tape drive during the backup process. Application servers also can be configured using dedicated access to tape drives within the tape library unit instead of the shared option. The flow of data through the storage network allows for a reduction in LAN traffic. A tape may also be shared amongst applications with multiple backup streams multiplexed onto managed tape libraries and drives. The LAN may still be used to pass metadata, context tables that track location of changes in data, back and forth between the backup server and the client but actual backed-up data is passed over the storage network. A typical LAN-free implementation is depicted in the figure below.

LAN-free backup segregates data and backup domains but doesn't alleviate CPU load on the application servers that still must fetch backup data from disk. Server-less backup as described below addresses this concern.

Figure 2

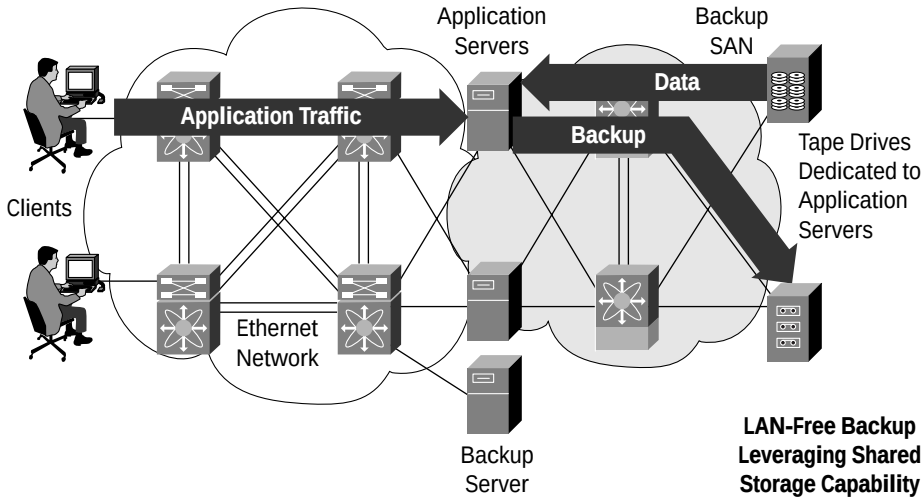


### Server-Less Backup

Server-less (or server-free) backup involved the task of moving backup data from a disk to a tape without the data traversing the application server itself. In server-less backup, a snapshot of the data to be backed up is taken with minimal or no disruption to applications on the server. The snapshot of the data is then moved intelligently from disk to tape without going through the server. This process alleviates additional CPU load on the application server as it is not in the backup path. The mechanism used to move data directly between disk and tape is the SCSI extended copy command. The component that performs the extended copy SCSI commands can be located in the SAN fabric or server software. Current implementations of server-less backup, such as Legato Celestra Power and Veritas NetBackup, manage the whole process via the backup server. A typical server-less backup scenario is depicted in the figure below. Most customers currently prefer to deploy separate SANs for backup due to a concern for segregated domains. Hence they may not deploy this solution as the disk subsystem and tape drive must coexist on the same SAN.

Figure 3

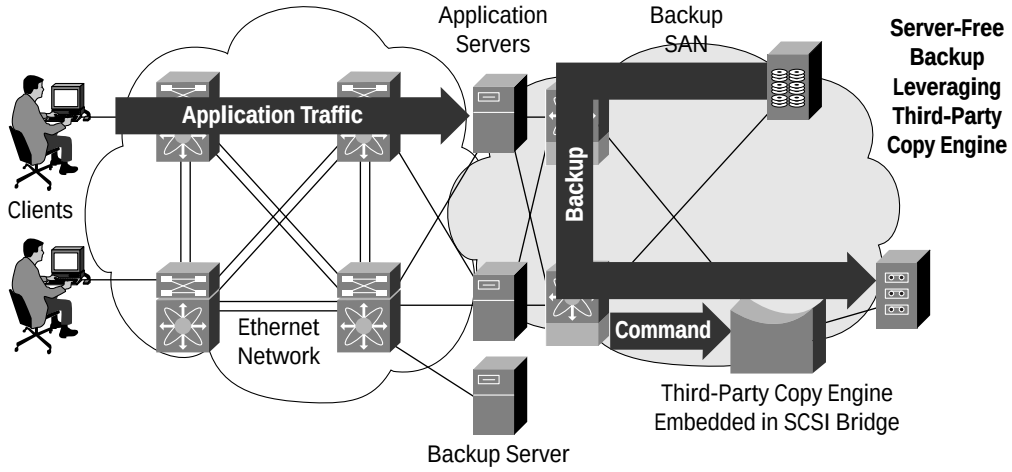


## Backup Industry and Product Overview

Many software vendors produce centralized backup software such as Veritas Netbackup, Veritas BackupExec, Legato Networker, and Computer Associates BrightStor ARCserve 2000 Advanced Edition. All these products follow a centralized mechanism whereby backups are performed via a dedicated backup server and directly attached archive device(s). A centralized backup scheme may also involve several different software and hardware modules. It is important to understand these various components of a backup solution and their role and impact on performance.

First, the central or *master server* controls the entire backup environment including indices, backup schedules, client group definitions, and hardware configurations. The *master server* is also responsible for logging problems with backups and reporting them to a system administrator. Sometimes this function is referred to as the *connection broker.*

The second type of server in a centralized backup system is a *media server* or storage node. The *media server* is attached to the type of storage medium for the backup, usually a tape device attached via a Fibre Channel or parallel SCSI connection. The *media server* is responsible for actually sending data to tape. This server takes its direction from the *master server* as to the files to backup and the particular media set to use. The *media server* may also be referred to as the *data mover.*

Third and most important is the backup client that actually refers to the server being backed up. Client software is installed to every system that requires backup services. Even the *master server* and *media server* usually have client software on them so that they may back themselves up.

## Major Advantages of SAN Deployment for Backup

As discussed above, enterprises have started to deploy dedicated Fibre Channel based storage networks to solve the performance bottleneck issues of a shared LAN. Additional advantages of storage networks are listed below:

### Data Availability

A storage network infrastructure provides multiple paths to storage subsystems, including disks and tapes, for both high availability and scalability. Customers can implement improved disaster recovery solutions, especially within an open systems environment. In legacy implementations, disaster recovery at a remote site is typically implemented by server-to-server communications over a LAN. In a networked model, data can also be mirrored between two storage subsystems freeing expensive server and LAN resources. Migrating backup traffic onto a storage network limits potential outages within each domain thereby protecting backups from any failure in LAN data traffic flow and vice versa.

### Lower TCO

Storage consolidation allows multiple servers to share the same storage devices and reduces the number of tape libraries required in the data center. Consolidation also makes it easier to reassign unused capacity amongst all servers to increase utilization and efficiency rates. Customers can gain additional savings by implementing an enterprise backup/recovery solution to reduce management and maintenance costs associated with server-attached tape drives. The cost of managing each individual component can be worsened by a high potential for human error resulting from manual processes. Mechanisms such as remote tape vaulting that deploy robots to eliminate manual transport of tapes between sites help reduce these errors. This also enhances reliability by eliminating damage caused by handling, potential for loss, and availability since the data that needs to be restored is never in transit. Storage networks enable pooling of backup resources that each server can draw upon without the limitations caused by a failure of an individual backup device. Lower total cost of ownership (TCO) is also delivered via enhanced scalability, availability, performance, and manageability of shared backup resources.

### Flexible Backup Options

Businesses are now open for longer hours and support a global community of customers which drives requirements for round-the-clock operation. The two backup options are available including hot and cold backup. Cold backup relates to the scenario where application data is tied up for the duration of the backup. However, a hot backup relates to the scenario where a system is performing this operation while applications are updating data.

Hot backup technologies such as copy-on-write and split-mirror snapshot use a mirrored image of the original data created at a particular instant to backup online without affecting the application. Both copy-on-write and split-mirror options copy data blocks to unused storage to create a point-in-time copy and are supported by most databases. The copy-on-write and split-mirror options manage the mapping process of physical data blocks and their correlation to a file system or database. The storage subsystem vendors that support these options include EMC Timefinder using the *Business Continuance Volume (BCV)*, HDS *ShadowImage*, and *FlashCopy* in the IBM Enterprise Storage Server (ESS) or the Modular Storage server (MSS).

Backup implementations strive to reduce the time required for backups to minimize disruption to user traffic (especially for cold backups) and increase performance. Hot backups utilize mirroring, whether it be local or remote, to ensure continuous availability for applications. A detailed discussion of remote mirroring technology is provided below.

### Remote Data Mirroring/Data Replication

Remote data mirroring, or *remote copy,* is the most commonly used mechanism for fast application and data recovery. This inherently implies that a mirrored volume within a disk subsystem is created to provide protection. The main volume being used by the application is considered the primary and the mirrored volume is considered as the secondary. The two main remote copy technologies that are implemented include:

- Host-based software for remote replication
- Storage controller-based hardware and firmware for remote replication

The most commonly known remote copy facilities available today include EMC's Symmetrix RemoteData Facility (SRDF). IBM eXtended Remote Copy (XRC), IBM's Peer-to-Peer Remote Copy (PPRC), HDS's TrueCopy, Compaq's Data Replication Manager (DRM), and Veritas Volume Replicator (VVR).

Both IBM XRC and Veritas VVR solutions are host based, software-assisted data mirroring facilities while PPRC, SRDF, TrueCopy, and DRM are hardware controller-based implementations of remote copy. It's important to point out that Veritas (VVR) provides remote data mirroring over an IP network instead of Fibre Channel or ESCON (Enterprise Serial CONnection).

## Cisco Products and Solutions for Disaster Recovery

As data availability becomes a key differentiator for enterprises, an increasing amount of resources are being spent in ensuring continuous operations. As explained above, dedicated networks are being provisioned to guarantee performance metrics as well as security for backup applications. Intelligent storage networks provide a new dimension to backup and recovery. In addition, remote data replication solutions offer higher degrees of availability and can be scaled to meet the requirements of the enterprise. Cisco Systems delivers state-of-the-art technology helping enterprises build end-to-end backup and recovery solutions along with disaster recovery solutions in a more scalable, secure and cost-effective fashion.

Figure 4

The Cisco MDS 9000 Family of Multilayer Directors and Fabric Switches. From left to right: MDS 9506 Director, MDS 9513 Director, MDS 9509 Director, and the MDS 9216 Fabric Switch

The Cisco MDS 9000 Family of Multilayer Directors and Fabric Switches targets enterprise and service provider storage network environments and delivers higher port density, high switching bandwidth, high performance, multiprotocol and greater reliability. The MDS 9000 Family products also targets heterogeneous storage area networks where the overall storage environment consists of multiple vendors' products. In those environments, the Cisco MDS 9000 Family products can serve as a centralized system to provide interconnection and advanced services.

The MDS 9000 Family consists of the MDS 9500 Series of Multilayer Directors and the MDS 9216 Multilayer fabric switch. The MDS 9000 Family products are modular systems that are optimized for high port density and performance for data center applications. For remote data centers used for backup and disaster recovery, the MDS 9216 Multilayer Fabric Switch provides a smaller product with all the same features and services as the MDS 9500 Family Directors.

In addition to the scale of these switching devices, the Cisco MDS 9000 Family of Multilayer Directors and Fabric Switches also provides an extensive list of features and services including Virtual SANs, advanced ISL link aggregation, LUN zoning, Call-Home, high availability, and hitless firmware upgrades. One of the issues plaguing most customers today is the lack of manageability and tools to allow for adequate support of a storage network. The Cisco MDS 9000 Family of products includes a robust embedded fabric manager application providing configuration, monitoring, and troubleshooting of storage networks.

### Virtual SANs (VSANs)

In many existing environments, backup solution design involved building a separate parallel network for backup traffic. From technical and operational points of view, this separate network approach provides a flexible, secure, and highly available backup solution albeit at a higher cost. The isolation of the tape storage network from the disk storage network eliminates an application-level impact due to failures such as a power reset of a tape library unit (TLU). The separate networks also protect against backup failures caused by a reset in a device on the application domain. It has thus become a de facto general practice to isolate tape devices from disk subsystems. With this isolation, a fabric configuration change remains local and nondisruptive fueling the continued deployment of segregated SANs.

While these separate storage networks help guarantee performance alleviate fabric wide disruptions, it is an expensive solution requiring separate switches and additional management complexity. In addition, this often leads to wasted ports that can also be costly to customers deploying these solutions. Cisco delivers advanced technology revolutionizing storage network deployments using a capability called a Virtual SAN or VSAN. VSANs provide a method of building separate virtually isolated fabric atop the same redundant physical infrastructure. VSANs therefore guarantee security and isolation of SAN domains as required in these designs. Cisco MDS 9000 Family of Multilayer Directors and Fabric Switches have the ability to create up to 4092 isolated VSAN topologies or layouts within the same physical infrastructure. This implementation is somewhat analogous to a VLAN in an Ethernet network. VSANs revolutionize SAN deployments by leveraging proven technology and ease of configuration of Ethernet networks to provide features like traffic isolation and security in a SAN environment. Some key features of VSANs are listed below:

- VSANs are isolated from each other using explicit tags placed on Fibre Channel frames to guarantee no leakage of frames from one VSAN to another.
- The implementation of VSANs allow for Domain IDs and FC_IDs to be reused within different VSANs
- Each VSAN has its own set of fabric services including zone server, name server, etc.
- Within each VSAN, a user can create zone sets as they would on a normal fabric
- A zone cannot span multiple VSANs.

VSANs apply proven technologies in Ethernet networking to the performance and isolation requirements of Fibre Channel storage networks to deliver best-in-class solutions. Zoning is very useful in restricting the access and traffic flow between devices within the fabric by securing access at the edge. VSANs offer a complimentary capability of isolating all services provided a fabric and "encapsulate" them within a VSAN therefore enabling a logically independent fabric.

Figure 5

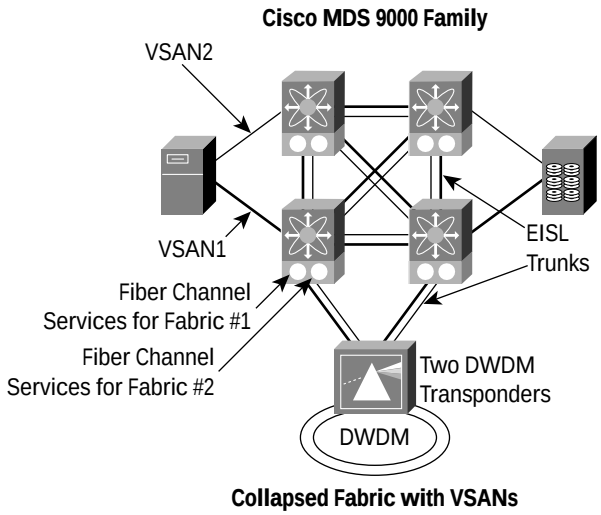


The storage network domains created by VSANs not only leap frog the existing security mechanisms but also deliver the ability to segregate even a single switch into multiple virtual environments. They provide complete separation among different VSANs and guarantee that a device outage or fabric instability is segregated to a single VSAN and doesn't cause a fabric-wide interruption. All of this is available without the need to implement expensive solutions requiring multiple physically-isolated fabric switches. VSANs also lower the storage network TCO by maximizing port utilization and thereby lowering the effective per-port cost. This type of implementation is quite efficient as multiple user communities can be grouped together thereby allowing a single fabric switch to service them using fewer overall devices. Cisco MDS 9000 Family delivers leading edge security and services while lowering capital and operational expenditures.

This architecture reduces the total number of SANs or fabrics deployed in the data center while maintaining status quo in terms of separating backup/recovery and remote data mirroring domains from the application SAN. Considering the data center real estate and storage consolidation model customers are adopting, VSANs enable a lower TCO solution for the enterprise.

### Multiprotocol Support

The MDS 9000 Family of Multilayer Directors and Fabric Switches integrate FCIP and iSCSI onto a single switching module for integrated Multiprotocol capability. The multiprotocol solution provides connectivity options over IP for applications such as disaster recovery. The integration of multiprotocol services along with high speed, high capacity Fibre Channel switching into a single platform allows customers to address their corporate-wide and departmental application requirements independently of each other however leveraging the same platforms. The integrated FCIP capability provides access to a remote data center over TCP/IP for seamless transition in case of a major failure. While

application requirements dictate the solutions that are deployed, FCIP is suitable for a large majority of remote data mirroring applications. FCIP is ideal for wide area connectivity between two data centers over a WAN link of varying speeds.

The integrated iSCSI capability is a low-cost host connectivity option leveraging an existing IP/Ethernet infrastructure. The iSCSI technology helps to optimize the use of expensive storage and tape subsystems by fanning the scope of the storage network to midrange and low-end servers in a cost-effective manner. This integration represents the industry's first Fibre Channel and IP storage integrated solution in a highly available director platform. Low-end servers with departmental or workgroup applications can easily be integrated into the same storage network fabric to provide easier access to corporate resources. Now workgroup applications are not tied to limited storage capacity or suffer from downtime caused during capacity upgrades. In addition, performance is not compromised, rather now delivered at a much lower cost. Furthermore, iSCSI provides centralized, shared tape option, LAN-free, and server-less capable backup and recovery solutions for departmental or workgroup application.

It is clear that Cisco's intelligent storage network feature set coupled with a focus on data center availability requirements delivers a breadth of solutions unmatched in the industry today. Cisco's products deliver enhanced management features to address enterprise concerns of reducing downtime caused by unplanned outages as well as planned upgrades. Technologies such as FCIP and iSCSI allow enterprises to reduce infrastructure costs while still providing integrated management across the two technologies.

Figure 6

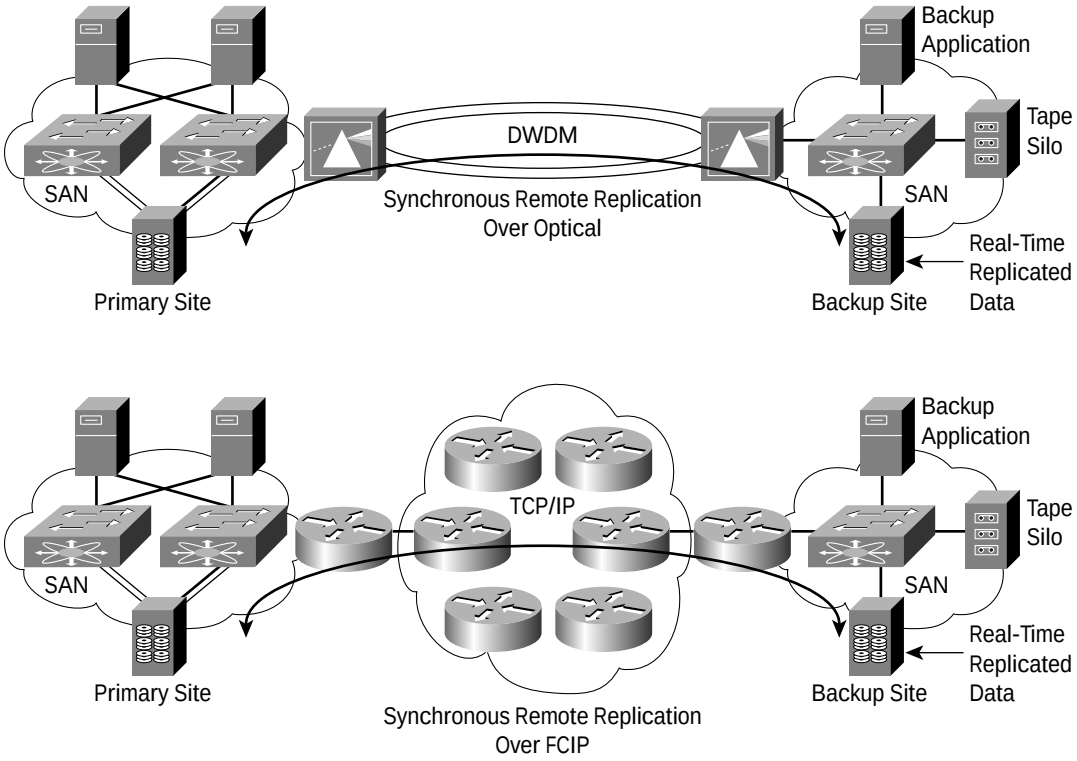

Synchronous mirroring applications for disaster recovery are also provided via Cisco's DWDM solutions including products such as the ONS 15540 DWDM switch and the ONS 15454 SONET switch platforms. DWDM is ideal for reliable metro area connectivity between two data centers while SONET provides high TDM bandwidth over longer distances. Both technologies provide excellent transport options for remotely replicated data over Fibre Channel or FCIP. More information on Cisco's DWDM solutions can be found at Cisco's website: http://www.cisco.com/warp/public/44/jump/optical_platforms.shtml

## Summary

Cisco MDS 9000 Family of Multilayer Directors and Fabric Switches delivers a best-in-class solution for Disaster Recovery, backup and restore applications. With a rich mix of technology, Cisco MDS 9000 Family is well positioned to address the needs of low-end, mid-range, and high-end enterprises. Cisco 9000 Family provides an end-to-end solution for customers deploying multiple applications in a heterogeneous transport environment. The MDS 9000 Family provides a seamless migration for customers to deploy new technologies but also co-exist with legacy environments. Given the focus on performance, port density, investment protection, and management, the Cisco 9000 Family enables large-scale deployments of highly effective SAN based solutions.


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: +65 317 7777
Fax: +65 317 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




LW3344 0802