EMC AND CISCO: BUILDING DISASTER RECOVERY AND BUSINESS CONTINUANCE SOLUTIONS

This document discusses important business continuance challenges, Cisco Systems® and EMC offerings to meet them, and the real-world solutions that can you can implement using Cisco® and EMC products. It includes the following sections:

- Business Continuance Overview
- Cisco and EMC: Partners in Business Continuance
- Building Business Continuance Solutions with EMC and Cisco Products
- Summary
- Additional References

BUSINESS CONTINUANCE OVERVIEW

This section describes some of the basic concepts and the issues that are important when considering a business continuance solution. It includes the following topics:

- Planning for Business Continuance
- Recovery Point and Recovery Time Objectives
- Disaster Radius
- Storage-Area Network (SAN) Extension Technologies for Business Continuance
- Replication Technologies for Business Continuance

Planning for Business Continuance

Business continuance ensures availability of the functions that are critical to the day-to-day business environment in the event of a natural disaster, hardware failure, or application failure. Whereas disaster recovery is a reactive plan to resume IT operations after a failure, business continuance is a proactive plan to maintain business operations despite planned or unplanned outages. Planning for business continuance requires carefully analyzing network requirements and typically requires developing some or all of the following documentation:

- *Business continuance plan*—Defines the resources, actions, tasks, and data required to manage the business recovery process in the event of a business interruption.
- *Data (or application) classification study*—Determines the value of each application to the business and assigns a service level to each, such as mission critical, business critical and critical. Applications that supply data to critical applications must also be identified.
- *Business impact analysis (or assessment)*—Defines the effect on the organization when an application becomes unavailable. A business impact assessment identifies the risks and potential costs associated with application downtime, data loss, and disrupted user access.

In addition, you may need to perform risk and threat assessment, backup and recovery planning, and server/site redundancy planning. Perform periodic audits of your business continuance plans to test organizational readiness.

Recovery Point and Recovery Time Objectives

The recovery time objective (RTO) is the maximum allowable downtime after an outage for recovering systems, applications, and functions (see Figure 1). RTO provides the basis for developing cost-effective recovery strategies and for determining when and how to implement these recovery strategies during a disaster situation.

Figure 1
RTO and Recovery Point Objective

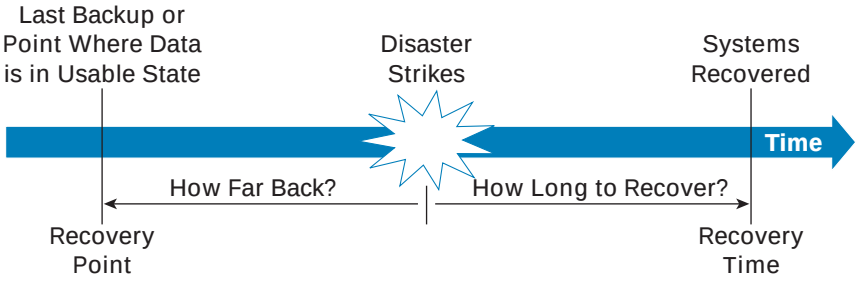

The recovery point defines how current or fresh the data is after a disaster. The recovery point objective (RPO) is the earlier point in time to which systems and data must be recovered after an outage. RPO defines the maximum amount of data that your organization is willing to sacrifice after a disaster. A zero RPO business continuance solution can survive a disaster without any loss of data.

Together, RTO and RPO provide a measurable target for a business continuance and disaster recovery solution to achieve. Improving RTO and RPO requires increasing your investment in networking and storage technologies and processes. Also, the physical distance between your data centers and how well your applications tolerate network latency affect how close you can get to zero RPO. That is why you should limit your RTO and RPO to whatever levels your organization can tolerate.

Disaster Radius

The probable distribution of a disaster, called the disaster or threat radius, affects the business continuance solution. The probability and extent of damage from earthquakes, floods, fires, hurricanes, cyclones, or terrorist threats varies according to the region in which the data center physically resides. To be effective, the backup site must not be within the disaster radius (see Figure 2).

Figure 2
Disaster Radius of a Data Center

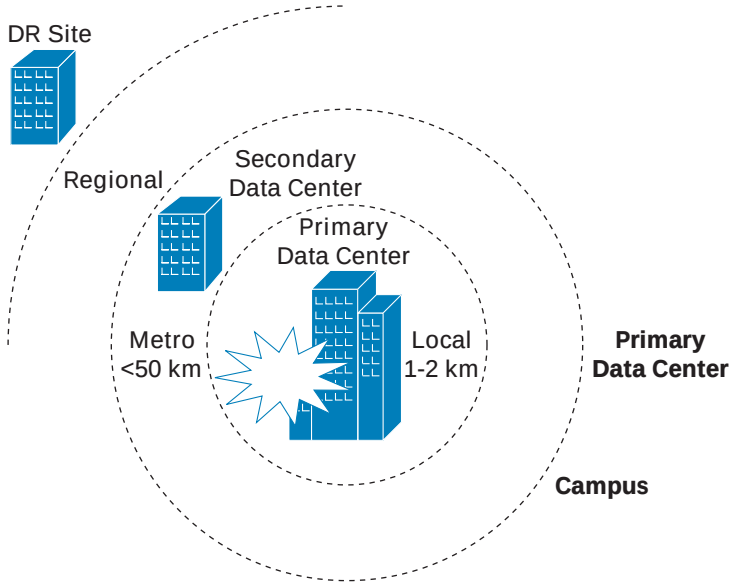

Defining the disaster radius may be more complicated than identifying a limited geographic region. For example, an earthquake might destroy both primary and secondary data centers if a major fault line connects them even though they are geographically separated. Many enterprises adopt a multihop strategy to be safe, using two data centers separated by metro distances and a third site located out of the region.

SAN Extension Technologies for Business Continuance

Modern data centers are built around SANs that provide secure, reliable, and scalable centralized storage facilities. Interconnecting or extending the SAN between the primary and the secondary data center ensures uninterrupted access to data independent of any disaster that occurs at either data center. The technical challenge is to provide distance extension over the distances involved in the most cost-effective and reliable way.

Cisco provides products and technologies that ensure high availability, scalability, and security within the data network and allows distance extension over regional, metro, and long-distance optical and IP networks. EMC builds on this network infrastructure with products that ensure that your organization can achieve the RTO and RPO required by your enterprise applications.

As shown in Figure 3, you can use optical and IP networks to extend a SAN over distance. Optical networks can be based on dark fiber or a managed fiber or wavelength service. Five basic technologies are available for building the network required for distance extension:

- Coarse wavelength-division multiplexing (CWDM)
- Dense wavelength-division multiplexing (DWDM)
- SONET/SDH
- Fibre Channel over IP (FCIP)
- EMC Symmetrix Remote Data Facility (SRDF) over IP with the EMC Gigabit Ethernet Director

Figure 3
Distance Extension Technology

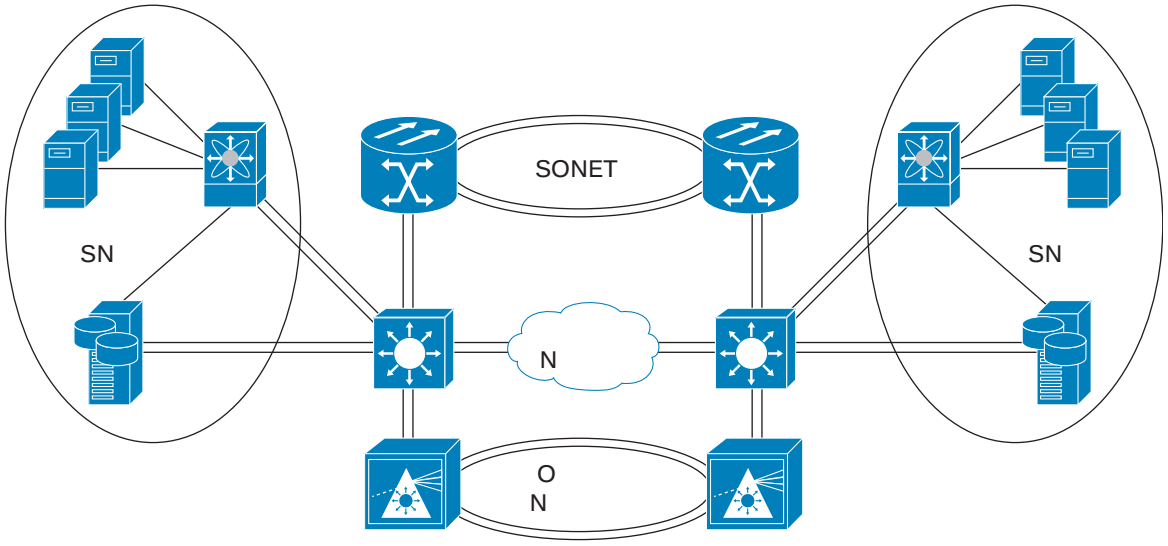

CWDM or DWDM technology works with either dark fiber or managed service. SONET/SDH uses time-division multiplexing (TDM) to enable the familiar T1 and higher leased-line services. Multiservice SONET/SDH solutions also support data, Fibre Channel, and FICON services. FCIP is an open standard for carrying blocks of data transparently over an IP network, which allows distance extension over LAN and even WAN networks.

EMC offers SRDF links with Fibre Channel, Gigabit Ethernet, and Enterprise Systems Connection (ESCON). ESCON is a legacy interface designed by IBM for mainframe connectivity. Similar in function to Fibre Channel and Gigabit Ethernet, ESCON links can be extended over CWDM, DWDM, or SONET/SDH links. The Gigabit Ethernet Director for the EMC Symmetrix provides native connection of Symmetrix storage into the IP network for replication over a WAN.

Cisco provides a full range of products for distance extension using various combinations of these technologies, as described in the section “Using Cisco Transport Technologies for SAN Extension” on page 9. EMC hardware and software products provide efficient and reliable methods of protecting and restoring your data. EMC products are described in the section, “Using EMC Products for Business Continuance and Disaster Recovery” on page 6.

Replication Technologies for Business Continuance

Different backup and replication technologies are required to address different levels of RTO and RPO required by the different application classes in an Enterprise data center:

- *Consolidated backup*—Consolidated tape or disk media backup is the least expensive solution and has the slowest recovery times. It is well suited for addressing limited disruptions, such as data corruption.
- *High availability storage networks*—A SAN solution can overcome local server failures by providing access to a standby or clustered server system to ensure continuous operation.
- *Remote point-in-time update replication*—Point-in-time update replication copies the changes made to data in another building or city. Changes can be replicated at scheduled times during the day or whenever changes occur. With no latency limitations to overcome, this technology accommodates any distance requirements. It offers faster recovery times than tape backup, but it cannot guarantee zero RPO.
- *Asynchronous replication*—Asynchronous replication has a significantly lower RPO than point-in-time update replication and reduces bandwidth costs compared to synchronous replication. Asynchronous replication allows the primary and remote copies to be out of synchronization by a range of seconds to minutes.
- *Synchronous disk replication*—Synchronous replication is suitable for applications that require the fastest recovery with zero RPO. All disk writes are synchronously copied to a remote site across a high-performance network before a transaction is acknowledged, eliminating any transaction loss. This technology is sensitive to network latency, which limits the practical distance between sites.

CISCO AND EMC: PARTNERS IN BUSINESS CONTINUANCE

This section describes the benefits of integrating Cisco and EMC products to create a business continuance solution and provides a brief summary of the products and technologies available from each company. It includes the following topics:

- Business Continuance Services from Cisco and EMC
- Using EMC Replication Products
- Using Cisco Transport Technologies for SAN Extension

Business Continuance Services from Cisco and EMC

Cisco and EMC offer services that help you find the best solution for your business continuance requirements. Cisco and EMC carefully consider each situation to develop the best service solution, and collaborate at every level to optimize all aspects of the business continuance solution, through each of the following phases:

- Planning
- Design
- Implementation
- Operation

Planning

During the planning phase, Cisco and EMC use assessments to generate a thorough understanding of your current environment and to gather the requirements for business continuance. These requirements are translated into a plan that maps to the overall architecture and identifies existing gaps. Potential obstacles are identified immediately with recommendations for overcoming these obstacles. The result is a thorough planning document that serves as the guideline for the design phase.

Design

Cisco and EMC each bring to the table the expertise required for designing the business continuance solution based upon the requirements identified during the planning phase. The entire solution is designed in a coherent way that maximizes efficiency and availability. Cisco and EMC carefully review each decision during the design process so that every component performs in an optimal way. The design process includes frequent reviews to guarantee a design that meets your business continuance requirements as well as any other related objectives.

Implementation

A business continuance solution typically is implemented in an operational environment where unscheduled downtime has a disastrous effect. Often, migration of existing applications or hardware is required as part of the implementation. In such circumstances, it is important to carefully plan the implementation. Cisco and EMC jointly develop the implementation plan and work together to execute it. This collaboration significantly reduces risk and ensures successful and rapid deployment.

Operation and Optimization

New applications and other developments have an impact on the business continuance requirements, which continue to evolve. Cisco and EMC provide ongoing services to solve problems in a fast and efficient way, to monitor changing requirements, and to update the business continuance requirements as required. New technologies and features are evaluated on an ongoing basis to validate their utility for achieving these requirements. These services help to maintain an optimal business continuance solution at all times and to quickly resolve any problems that might occur.

Using EMC Products for Business Continuance and Disaster Recovery

This section describes the replication products and technologies available from EMC that are useful in creating a business continuance solution. It includes the following topics:

- SRDF/Synchronous
- SRDF/Asynchronous
- SRDF/Automate Replication

Refer to http://www.emc.com for more information about EMC management and monitoring applications.

SRDF/Synchronous

EMC SRDF/Synchronous (SRDF/S) software offers a host-independent, mirrored data-storage solution that replicates production or source data to a remote site (see Figure 4).

Figure 4
SRDF/S

You can use SRDF/S for a range of applications, from disaster recovery to data movement and migration. If source data storage becomes unavailable during disaster recovery, SRDF/S enables fast switchover to the recovery data so that critical information is restored within minutes. Business operations and related applications resume full functions with minimal interruption.

Symmetrix production and recovery systems can be adjacent or many miles apart. In either case, the same enterprise-level information protection capabilities are offered. After a disaster, SRDF/S can resynchronize data to either the production site or to the recovery site, depending on your preference. For a complete description of SRDF/S features and functions refer to the EMC SRDF/S Production Description Guide, at the following Website:

http://www.emc.com/products/product_pdfs/pdg/srdf_pdg.pdf

SRDF/S enables business continuance by allowing data at a recovery site to be easily available for use as the primary data on demand. EMC business continuance solutions create an environment where planned workload balancing or data center outages easily can be accommodated without disrupting application performance.

SRDF/S operates in the following modes:

- Synchronous mode
- Semisynchronous mode
- Adaptive copy mode

The appropriate mode of operation depends on a range of factors including latency toleration, SRDF/S usage (disaster recovery versus data replication), performance requirements, and RPO. Each situation is different and often requires assessment by EMC to determine the appropriate mode of operation.

Synchronous Mode

Synchronous mode is appropriate for an implementation that requires data at the recovery site to be synchronized continuously to the source. In a disaster, the recovery data is accurate up to the last transaction and immediately is available for use. Synchronous mode mirrors data between the source Symmetrix and the target Symmetrix system. In this mode, data is written first to the source system and then to the target system. The source must receive an acknowledgment before a subsequent host I/O operation is allowed.

Semisynchronous Mode

Semisynchronous operation masks any performance impact that results from I/O propagation delays. Semisynchronous mode writes data to the source system, completes the I/O, and then synchronizes the data with the target system. Because the I/O is completed prior to synchronizing data with the target system, this mode provides an added performance advantage. A second write is not accepted on a volume until the target has been synchronized, so this is not asynchronous.

By presenting an earlier I/O complete signal to the host/server, the host/server continues processing reads to the logical volume. Most applications do a read before they do a write. This implementation successfully masks the I/O elongation that is created by the propagation delay to the target system in a pure synchronous mode.

Adaptive Copy Mode

The Adaptive Copy mode is useful for data center migrations or for migrating large blocks of data from one place to another. When choosing this mode of synchronization, the user should be aware that data is not transferred in a serialized manner. Therefore, Adaptive Copy mode is not a prime choice for real-time recovery.

In Adaptive Copy mode, the user specifies that the primary mode of operation for the source volume be suspended. Data is transferred from the source volumes to the remote volumes without waiting for an acknowledgment, and this allows a lag between source and target volumes of the specified number of I/Os. This allows SRDF/S to manage the Adaptive Copy volumes at a lower priority, enabling applications using that volume to avoid propagation delays while data is transferred to the remote site.

SRDF/Asynchronous (SRDF/A)

SRDF/A is an innovative approach to asynchronous replication that delivers a consistent and restartable remote copy of production data over extended distances and with minimal impact on host applications. Its unique Delta Set architecture provides efficient utilization of bandwidth requirements while minimizing data exposure without sacrificing performance. Delta Sets are cache-resident collections of dependent writes that have occurred within a specific period of time. Delta Sets also provide effective utilization of bandwidth between sites in two ways:

- Locality-of-reference
- Sizing communication links to average write workloads

With traditional asynchronous solutions, every write must be time stamped, sequenced, and sent to the remote side. Every write is treated as a discrete write and must be moved over the communication link. Through locality-of-reference, Delta Sets allow data to be rewritten, so only the final set of updates is sent to the remote site.

Another bandwidth savings provided by SRDF/A is the ability to size communication links to average write workloads rather than to peak workloads. Peak workloads might occur only for a brief period while batch updates are occurring or some other period of time when processing activity is high. Because Delta Sets collect the new writes in cache while experiencing locality-of-reference benefits, SRDF/A efficiently rides through periods of increased write activity.

SRDF/S and TimeFinder/Mirror

Business continuity[1] solutions must be robust, flexible, and scalable to meet mixed service-level requirements. Deploying a single solution to solve all your business continuity challenges is not an option as more and more data must be kept online. This means that multiple applications must have access to the same data at the same time.

Fortunately, you can deploy EMC business continuity solutions with any combination of products to solve your mixed service level requirements, knowing that the products work well together. Your investment in an EMC business continuity solution protects data in the event of a disaster and can help drive operational efficiencies.

SRDF can be used with either of the following EMC Symmetrix local replication technologies:

- Full-volume TimeFinder/Mirror Business Continuance Volumes (BCV)
- Space saving TimeFinder/Snap snapshot software

Either option allows you to perform parallel processing activities at both the source and target locations without temporarily suspending the remote mirroring operation and exporting your data.

When a TimeFinder/Mirror BCV or TimeFinder/Snap snapshot is paired with a secondary device, you can restore data from the BCV or snapshot to the secondary device without propagating the data across the SRDF/S links to the primary device. Ensure that the secondary device is suspended on the SRDF/S link. Use the appropriate TimeFinder/Mirror or TimeFinder/Snap software restore command options to restrict operation to the desired secondary volumes. However, if you want the restored data automatically to move across the SRDF/S links to the primary device, make sure that the secondary device is *active* (not suspended) on the SRDF/S links.

USING CISCO TRANSPORT TECHNOLOGIES FOR SAN EXTENSION

This section describes the different technologies and products available from Cisco that provide the network infrastructure for distance extension and for supporting EMC business continuance applications. It includes the following topics:

- Overview
- Cisco CWDM Solutions
- Cisco DWDM Solutions
- Cisco SONET/SDH Solutions
- Cisco FCIP Solutions

1. Business continuity is the term that EMC uses to refer to what Cisco calls business continuance.

Overview

Figure 5 illustrates various methods of interconnecting data centers and some of the hardware products from Cisco that allow cost-effective distance extension for business continuance.

Figure 5
Business Continuance Products from Cisco and EMC

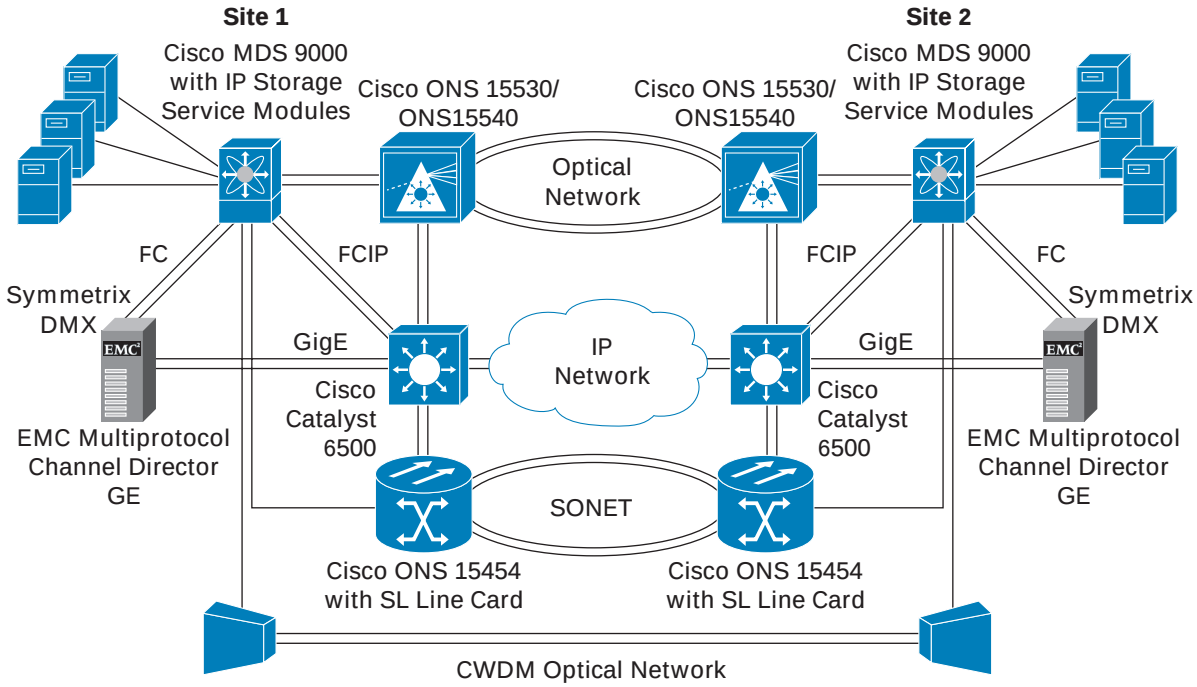

Cisco provides a full range of products for creating a resilient, fault-tolerant, and scalable infrastructure using any combination of distance extension technologies, including the following technologies:

- *Cisco Catalyst® 6500 Series*—Multilayer modular switch providing LAN/metropolitan-area network/WAN convergence
- *Cisco ONS 15530*—Provides distance extension supporting Fibre Channel, FICON, ESCON, or FCIP over a DWDM network
- *Cisco ONS 15540 ESP Extended Service Platform*—Provides distance extension supporting Fibre Channel, FICON, ESCON, or FCIP over a DWDM network
- *Cisco ONS 15454 SONET Multiservice Provisioning Platform (MSPP)*—Provides distance extension supporting Fibre Channel, FICON, or FCIP over a SONET/SDH network
- *Cisco MDS 9000 with IP Storage Services Module*—Provides distance extension using FCIP over an IP network

Note: The Cisco MDS 9000 family of products is cobranded within the EMC Connectrix product line.

The following paragraphs describe each solution.

Cisco CWDM Solutions

CWDM is a convenient and cost-effective solution for extending Fibre Channel (and Gigabit Ethernet) links over data center, campus, and metropolitan access networks. CWDM only incurs latency due to the speed of light through fiber, so it is ideally suited to the low-latency requirements of SRDF/S.

CWDM allows up to eight 1 Gbps or 2 Gbps channels to share a pair of fibers or up to four channels to share a single fiber. Each channel uses a Small Form-Factor Pluggable (SFP) tuned to a different wavelength to send and receive over the shared fiber. These channels are networked with a variety of wavelength-specific optical add/drop multiplexers (OADMs) to enable an assortment of ring or point-to-point topologies. The CWDM wavelengths are not amplifiable and are limited in distance according to the number of joins and drops. A Cisco CWDM SFP has a 28 dB power budget at 2 Gbps, allowing it to reach up to ~90 km in a point-point topology or around 40 km in a ring topology.

Protection against failure is provided by Cisco MDS 9000 Series multilayer switches or directors at the end points. Diverse optical paths must be employed to ensure high availability. PortChannels are recommended to maintain fabric availability through path failures. The Cisco CWDM solution consists of two component sets:

- CWDM SFP transceivers (eight)
- CWDM OADMs

Cisco DWDM Solutions

DWDM is the most common technology used for synchronous replication between metro data centers. As with other pure Fibre Channel over optical solutions, DWDM links only incur latency from the speed of light through fiber.

Cisco DWDM systems enable high density, 32-channel (lambda) aggregation between data centers up to 320 km apart. Each channel can operate up to 10 Gbps for a total of 320 Gbps over a single fiber pair. In contrast with the passive multiplexing of CWDM, DWDM platforms, such as the Cisco ONS 15530, Cisco ONS 15540, and Cisco ONS 15454, offer intelligent protection schemes to guard against failures in the fiber plant.

Additional advantages include a comprehensive selection of DWDM transmission elements enabling support for amplification, multiple network architectures, node configurations, in-service scalability, and easy-to-use network management capabilities for point-and-click provisioning.

In an IBM mainframe environment, the Cisco ONS 15530 and Cisco ONS 15540 can aggregate ESCON, IBM Fiber Connection (FICON), plus Sysplex and Coupling link protocols. This is in addition to Fibre Channel, Gigabit Ethernet, and 10 Gigabit Ethernet, making them ideally suited for metro data center deployments and SRDF synchronous replication business continuity purposes.

The Cisco ONS 15530 can aggregate up to 8 Fibre Channel or FICON services into a single wavelength or as many as 40 ESCON services into a single wavelength. Multiple Cisco ONS 15530 platforms can be deployed to support up to 256 Fibre Channel/FICON/Gigabit Ethernet interfaces or as many as 1280 ESCON interfaces per fiber pair. The Cisco ONS 15530 and Cisco ONS 15540 platforms offer the following wavelength and protection options for achieving network availability requirements of five nines or even higher:

- Optical unidirectional-path switched ring
- 1+1 optical trunk protection
- 1+1 wavelength protection
- Wavelength optical splitter protection

- 1+1 service protection
- 1+1 client protection

The Cisco ONS 15454 MSTP solution provides intelligent DWDM functionality for metropolitan as well as regional networks. This solution delivers a comprehensive suite of transparent wavelength service interfaces, including Fibre Channel, ESCON, FICON, Gigabit Ethernet, 10 Gigabit Ethernet (LAN/WAN), and support for SONET/SDH interfaces (OC-3/STM-1 to OC-192/STM-64). By supporting distances of up to 600 km, the ONS 15454 provides an ideal solution for both secondary data centers as well as out-of-region backup sites located outside of the disaster radius.

With the Cisco MDS 9000, PortChannels of up to 16 Fibre Channel Inter-Switch Link (ISL) or Ethernet ISL (EISL) links can be used with any of these protection mechanisms for enhanced resilience.

Cisco SONET/SDH Solutions

SONET/SDH remains the predominant basis for most service provider networks worldwide. As with other optical solutions, Fibre Channel over SONET/SDH is suitable for SRDF synchronous replication deployments, subject to application performance constraints as the distance increases.

SONET is the North American standard, while SDH is the standard elsewhere in the world. SONET/SDH can transport a number of optical and electrical WAN interface types, such as T1, DS3, STS-3, and so forth in North America. In the rest of the world, slightly different protocols, such as E1, E3, STM-1, and STM-4, are used. SONET/SDH is also deployed in some enterprise networks, often in addition to DWDM. The Cisco ONS 15454 platform offers SONET/SDH client connection options, in addition to Gigabit Ethernet and Fibre Channel.

Fibre Channel over SONET/SDH is provided by the SL line card for the Cisco ONS 15454. The Cisco ONS 15454 offers several SONET/SDH protection options in addition to client-level protection available through Fibre Channel switches. PortChannels are recommended with virtual SAN (VSAN) trunking where multiple links are used.

Cisco FCIP Solutions

FCIP links do not use the Fibre Channel Buffer to Buffer Credit (BB_Credit) mechanism used with Fibre Channel over optical networks. FCIP links are therefore not constrained by the distance limitations of optical networks. Instead, FCIP uses TCP for flow control. The 32 MB maximum TCP window size per FCIP link enables the Cisco MDS 9000 FCIP to maintain Gigabit throughput rates over intercontinental distances.

FCIP is implemented on the MDS 9000 IP Storage Services Module for extending Fibre Channel fabrics over an IP network. FCIP is used where optical networks are unavailable or where the distance between data centers is so great that BB_Credit limitations would cripple throughput. Each Cisco MDS 9000 IP Storage Services Module has 8 Gigabit Ethernet interfaces, each supporting up to 3 FCIP point-to-point links for a total of 24 FCIP links per card.

Up to 16 FCIP links can be PortChanneled for greater resilience and throughput, and VSAN trunks allow for SAN scalability. The Cisco MDS 9000 with the IP Storage Services Module sustains wire rate performance with very low delays through the switch and can be used for SRDF synchronous replication over low-latency Metro Ethernet or optical networks.

The TCP implementation on the Cisco MDS 9000 IP Services Module is optimized for carrying storage traffic. The FCIP TCP implementation can avoid the low throughput, slow start behavior of normal TCP implementations, and it recovers more quickly from packet loss. The Cisco MDS 9000 FCIP implementation employs variable, per flow traffic shaping that yields high instantaneous throughput while minimizing the possibility of overruns or congesting on downstream routers and switches.

BUILDING BUSINESS CONTINUANCE SOLUTIONS WITH CISCO AND EMC

This section describes some examples of proven solutions using EMC and Cisco products for campus, metro, and wide-area environments, using both optical and IP technologies. It includes the following topics:

- Campus, Metro, and Wide-Area Solutions
- Campus and Metro Optical: EMC SRDF/S with Cisco MDS 9000 and Cisco ONS 15530/Cisco ONS 15540 DWDM Metro Optical
- Campus and Metro: EMC SRDF/S with Cisco MDS 9000 and CWDM
- Campus and Metro: EMC SRDF/S over FCIP with Cisco MDS 9000 and IPS Module
- Wide Area: SRDF/A and SRDF/AR over FCIP with Cisco MDS 9000 and IPS Module
- Wide Area: SRDF over IP with the Multiprotocol Channel Director Using GigE

Campus, Metro, and Wide-Area Solutions

The best business continuance solution for your organization is determined by the distance between your data centers, by the RTO and RPO, and by how well your enterprise applications tolerate network latency. Shorter distances reduce network latency, which allows using synchronous disk replication and data center mirroring. Business continuance solutions can be divided into two general categories, based on the distances between the local and remote data center:

- Campus and Metro: Distances up to 200 km
- WAN: Long distance over 200 km

Based on these categories, the best use of EMC SRDF products within different scenarios is as follows:

- SRDF for campus/metro distances when zero data exposure is required
- SRDF/A = campus/metro or WAN distances when zero host impact is required but some data exposure can be tolerated
- SRDF/AR option for metro or WAN distances when synchronous regional protection is required in addition to long distance

The following sections describe end-to-end solutions using Cisco and EMC products for various distance requirements.

Campus/Metro: EMC SRDF/S with Cisco MDS 9000 and Cisco ONS 15530, Cisco ONS 15540, Cisco ONS 15454, and DWDM Metro Optical

Metro optical networks, such as DWDM, provide the lowest latency network options for SRDF synchronous data replication. The Cisco ONS 15540, Cisco ONS 15530, and Cisco ONS 15454 Multiservice Transport Platform (MSTP) platforms each can multiplex up to 32 protected channels (or wavelengths) per fiber pair. Each channel can be Gigabit Ethernet, SONET/SDH, ATM, Fibre Channel, FICON, or ESCON. The Cisco ONS 15530 can aggregate up to 10 ESCON links into a single lambda. The Cisco ONS 15454 additionally can transport SONET/SDH from OC-3/STM1 to OC-192/STM64.

In the solution shown in Figure 6, two replication VSANs are defined between the Symmetrix storage arrays in each data center for carrying the SRDF replication traffic.

Figure 6
Synchronous Replication Over Cisco ONS 155x0 DWDM Metro Optical

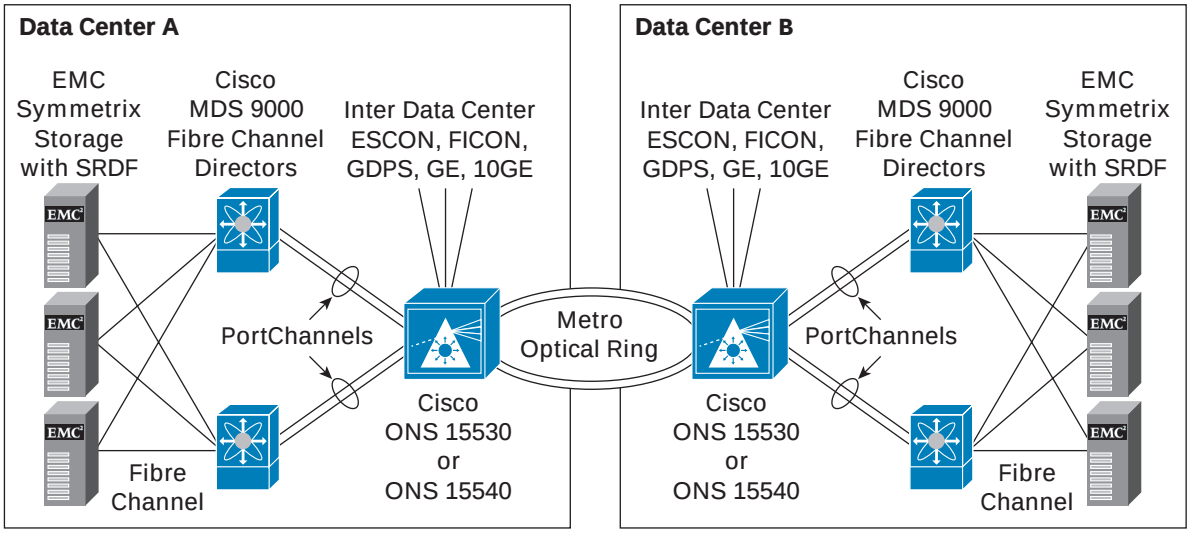

In this example, a PortChannel is configured for each replication VSAN, with each link in the PortChannel routed over alternate paths in the DWDM ring. This is augmented with facility protection (also known as splitter-based protection) on the DWDM platforms. Facility protection switches traffic to an alternative protected fiber upon detection of a fiber cut or signal degradation in the DWDM network. The combination of PortChannels and DWDM facility protection means the bandwidth and connectivity of the SRDF replication network is completely protected against any single fiber cut.

VSANs enable enterprises to maintain separation between SAN fabrics without additional fabric switch hardware. In this example, additional Symmetrix storage arrays can be added on existing or new SAN fabrics without changes to the Cisco ONS 15530/Cisco ONS 15540 or Cisco MDS 9000 infrastructure. Additional links can be added to the Cisco MDS 9000 PortChannel, and lambdas can be added to the Cisco ONS 15530/Cisco ONS 15540 platforms as throughput requirements increase.

The Cisco MDS 9000 16-port Fibre Channel line card supports up to 255 BB_credits per port, allowing DWDM metro optical fabric extension up to 250 km at 2 Gbps or up to 500 km at 1 Gbps without BB_Credit-imposed performance degradation.

Campus and Metro: EMC SRDF/S with the EMC/Cisco MDS 9000 and CWDM Transport

CWDM provides a convenient low-cost alternative to DWDM for Gigabit Ethernet and Fibre Channel connectivity over a data center, campus, or metro area for SRDF synchronous replication. CWDM allows multiplexing up to eight Fibre Channel or Gigabit Ethernet links at one or two Gbps over a fiber pair. In the solution shown in Figure 7, two replication VSANs are defined between the Symmetrix storage arrays in each data center for carrying the SRDF replication traffic.

Figure 7
SRDF Synchronous Replication over a PortChanneled Cisco MDS 9000 Based Point-to-Point CWDM Network

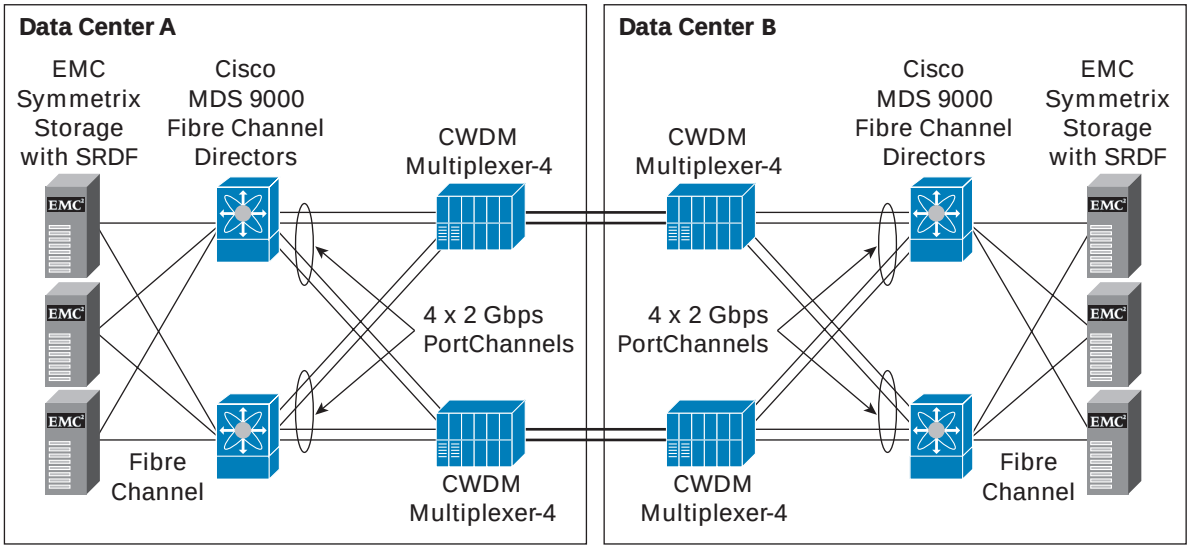

A PortChannel of four 2 Gbps Fibre Channel links is configured for each replication VSAN, with two links from each PortChannel routed through a different CWDM multiplexer and over a different fiber path between data centers.

Because CWDM OADMs are passive, network protection from fiber cuts and failures for each VSAN is provided by the diversely routed PortChannels on the Cisco MDS 9000 fabric switches. PortChanneled ISLs or EISLs are split over alternative optical paths, mitigating the effect of a fiber cut.

Using multiplexer-8 CWDM multiplexers in place of the multiplexer-4s shown in this example doubles the available Fibre Channel bandwidth to eight 2 Gbps channels per fiber pair. Alternatively, the additional CWDM channels can be used for Gigabit Ethernet connectivity between Cisco Catalyst family switches in the data centers.

The CWDM SFPs have an optical power budget of 28 dB at 2 Gbps and are capable of traversing a distance of around 90 km in the topology shown in this example.

Campus and Metro: EMC SRDF/S over FCIP with the Cisco MDS 9000 and IP Storage Services Module

FCIP allows you to use existing IP network infrastructure for distance extension. FCIP may be the only solution available if service providers do not offer native Fibre Channel transport service or when dark fiber for DWDM or CWDM is unavailable, too costly, or the distance is beyond metro optical range.

FCIP can be used over any distance for SRDF/S, SRDF/A, and SRDF/AR replication applications and the optional SRDF/AR configuration. The underlying TCP/IP network lets FCIP use a wealth of IP services such as compression, QoS, IP Security (IPSec), and a shared enterprise IP network.

An example dual data center deployment using FCIP with the IP Services Module of the Cisco MDS 9000 is shown in Figure 8.

Figure 8
Replication Using FCIP Transport over SONET/SDH, Metro Ethernet, or Transparent LAN Service

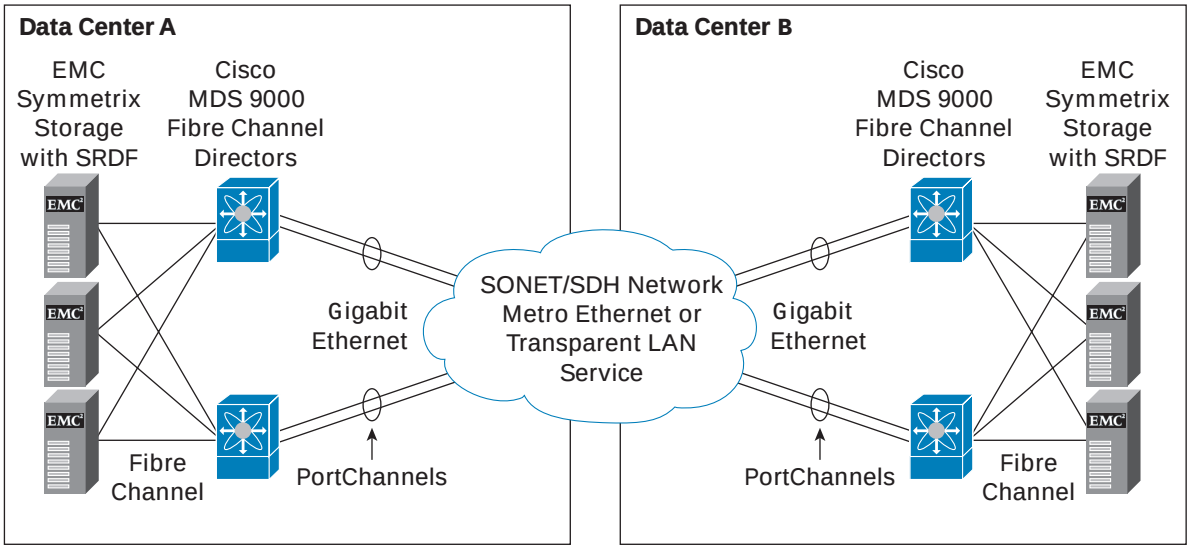

This example uses a Gigabit Ethernet service from a SONET/SDH or Metro Ethernet service provider. The topology is similar to that of a high availability optical solution with split PortChannels routed over dual interdata center links. In this design, a single link failure or Ethernet switch failure does not disrupt either of the replication VSANs other than by diminishing the available bandwidth.

In this application, use quality of service (QoS) options to provide low latency queuing and ensure optimum service for storage traffic. You must be careful to ensure that queuing and congestion in the IP network is kept to a minimum because synchronous replication is sensitive to network latency.

The Cisco MDS 9000 Write Acceleration feature lowers network latency by reducing the usual two round trips per write to one. A Fibre Channel Shortest Path First preferred path topology is required with Write Acceleration to ensure that all traffic within a particular VSAN traverses the same bidirectional path.

Wide Area: SRDF/A and SRDF/AR over FCIP with Cisco MDS 9000 and IP Storage Services Module

Synchronous replication is sensitive to latency and is harder to implement as the distance increases. In contrast, asynchronous replication is not sensitive to latency and can be deployed at almost any distance, as long as adequate network bandwidth is available.

Some of the possible WAN deployment options are shown in Figure 9. In this dual data center deployment, the Fibre Channel SANs are linked using FCIP on the Cisco MDS 9000 IP Services Module.

Figure 9
Replication Using FCIP Transport Over WAN Service

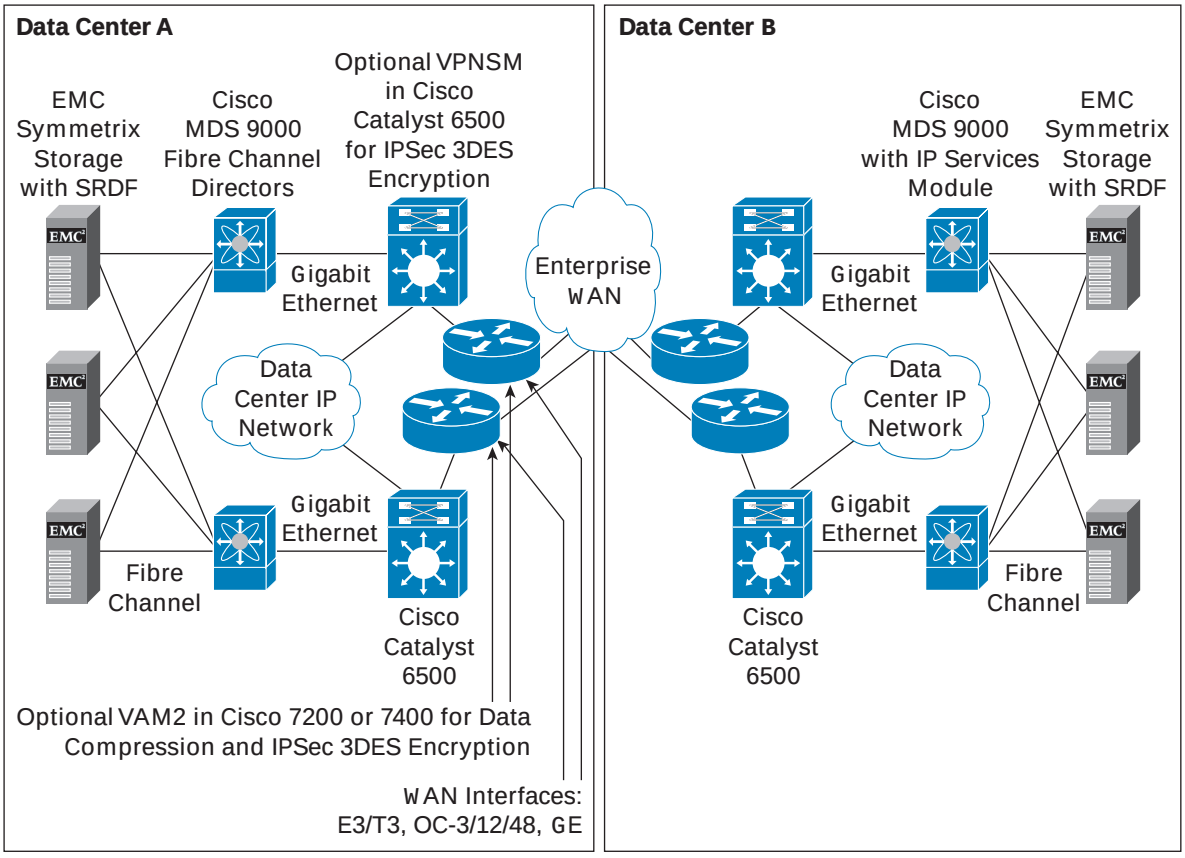

In this example, the Gigabit Ethernet ports on the IP Services Module are connected to a pair of Cisco Catalyst 6500 Series switches and then to a wide area router, such as the Cisco 7200 Series Router. The Cisco Catalyst 6500 Series switches interconnect to other data center LANs for carrying the shared WAN links with the FCIP traffic. These switches optionally can be equipped with VPN service modules to provide encryption of FCIP storage traffic at rates over 1 Gbps.

Compression offers a way to maximize throughput of WAN links. Compression can be configured on each Cisco MDS 9000 FCIP link in either high-throughput or high-compression mode. This can be combined with the Cisco Catalyst 6500 Series VPN Service Module for a compression/encryption solution that you can use with SRDF/A or SRDF/AR replication traffic over FCIP. The SA-VAM2 Port Adapter for the Cisco 7200 Series Router provides a single-chassis solution for compression and encryption of traffic over OC-3/STM-1 or DS3/E3 links.

Wide Area: EMC SRDF Family over IP with the EMC Gigabit Ethernet Director

The Gigabit Ethernet Director for the EMC Symmetrix provides native connection of Symmetrix storage into the IP network for replication over a WAN. This eliminates the need for any protocol conversion or bridging hardware. Gigabit Ethernet directors are available with Symmetrix 8000 and Symmetrix DMX models. The Symmetrix DMX also has on-board compression.

An SRDF deployment using the Gigabit Ethernet directors is shown in Figure 10. In this implementation, two Cisco Catalyst 6500 Series switches aggregate the storage and application data traffic for transport over a pair of Cisco routers.

Figure 10
Gigabit Ethernet Director for EMC Symmetrix

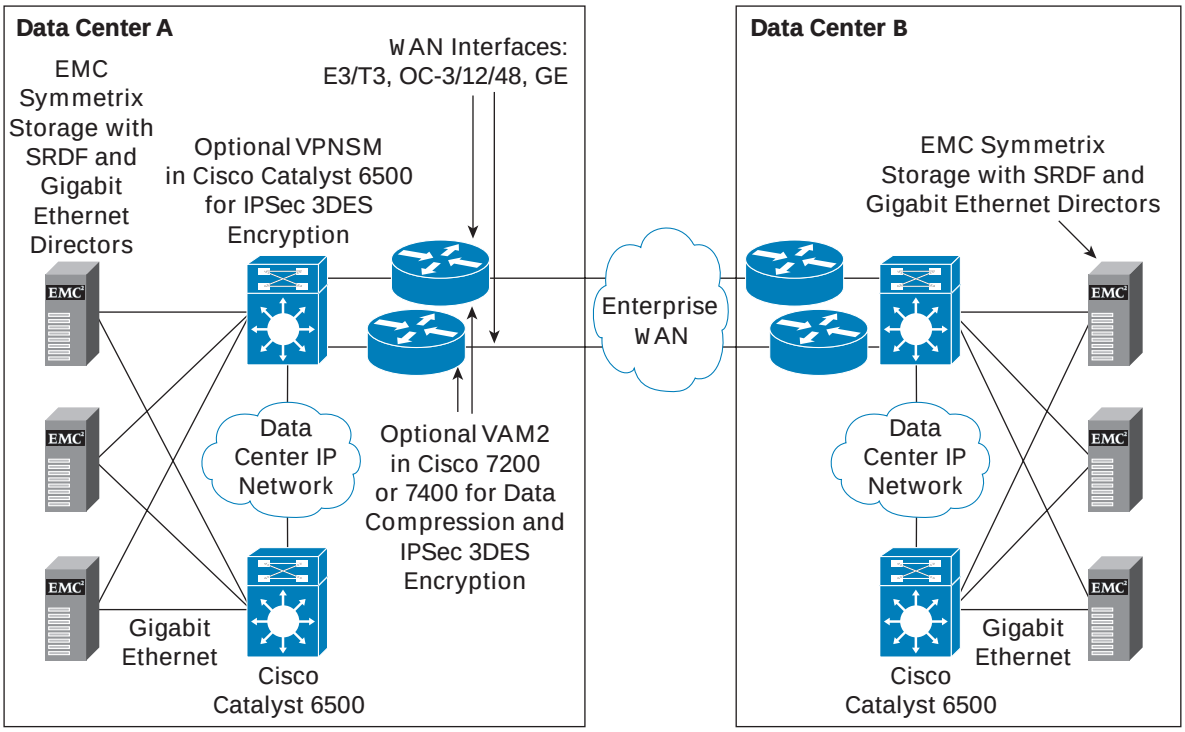

If data privacy is required, VPN Service Modules (VPNSM) can be used to create a Triple Data Encryption Standard IPSec encrypted tunnel between the data centers. Optionally, with Symmetrix 8000 or mixed Symmetrix 8000 and Symmetrix DMX configurations where the on-board compression cannot be used, you can obtain external compression with the VPN Adapter Module (VAM-2) on a Cisco 7200 or 7400 router at the data center edge.

SUMMARY

Together, Cisco and EMC provide a full range of optical and IP distance extension solutions over a variety of media for the transport of EMC SRDF/S and SRDF/A replication data streams. To ensure the success of any distance extension deployment, carefully consider the following:

- Enterprise RPO and RTO
- Applications that are most essential to the enterprise
- Possible threats around each data center
- Maximum tolerable disk I/O service time (latency) for the major enterprise applications

Analysis of this information, with special consideration of the trade offs between application performance, geographical separation, and RPO/RTO, will provide a solid basis for determining the appropriate replication mode and the distance extension solution for your enterprise.

ADDITIONAL REFERENCES

EMC Cisco MDS FCIP and Small Computer System Interface over IP solutions white paper. At time of publication, this paper is available only to EMC employees; it soon will be available on Powerlink.

"Wavelength-Division Multiplexing in EMC Environments"

SRDF/S Product Description Guide:
http://www.emc.com/products/product_pdfs/pdg/srdf_pdg.pdf

Understanding Alternatives for Extending Storage Networks:

http://www.cisco.com/en/US/products/hw/optical/ps2011/products_white_paper09186a008020805e.shtml

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.
168 Robinson Road
#28-01 Capital Tower
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 Website at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus
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

Mi/LW6058 04/04