

Understanding the **Alternatives** for Extending Storage-Area Networks

Storage capacity for the IT industry is growing 49 percent annually, according to a report by IDC covering the first quarter of 2003. Because stricter government regulations are demanding longer data retention and because greater amounts of data are being generated through business automation, this growth trend is likely to continue. To accommodate the vast amounts of data needed throughout enterprises, many IT organizations are switching from direct-attached storage to storage-area networks (SANs). In the latter model, storage resources are detached from the host servers and are shared among various servers in a SAN. The most popular protocols for SANs include Fibre Channel, Enterprise Systems Connection (ESCON), and IBM Fiber Connection (FICON).

With widespread power outages, potential threats of terrorism, and natural disasters still fresh in the minds of many IT planners, disaster-recovery applications that are designed to protect data centers from local catastrophes (such as geo-clustering and remote disk replication) have become hot topics. Because most of the SAN protocols were originally designed to function within a campus or, at most, within a metropolitan area, additional methods of transport are required to allow SANs to be truly geographically dispersed.

This paper explains the various options for SAN extension.

SAN Extension Deployment Considerations

Before selecting a SAN extension technology or solution, carefully consider the business drivers, technical system requirements, and infrastructure and service availability, outlined as follows:

Business Drivers

- *Regulations*—Recent events have driven regulatory authorities to mandate new standards for disaster recovery and business continuance across many sectors, including financial and banking, insurance, health care, and government. As an example, the Federal Reserve and Securities and Exchange Commission recently released a document titled “Interagency Paper on Sound Practices to Strengthen the Resilience of the U.S. Financial System,” which outlines objectives for rapid recovery and timely resumption of critical operations after a disaster. Similar regulations addressing specific requirements for health care, life sciences, and government have been issued or are under consideration.
- *Cost*—Factors include the cost of downtime (millions of dollars per hour for some institutions), more efficient use of storage resources, and reduced operational expenses.



Asynchronous Replication

Asynchronous replication is typically used for low-speed remote connections. For example, a data center in California might use asynchronous replication to a data center in Kansas. The asynchronous replication application sends input-output requests to the mirrored site as they are received without first waiting for acknowledgments from any previous input-output operation. It should be noted that the number of outstanding operations varies between applications. As such, the more outstanding operations that are allowed, the greater the length of latencies allowed to maintain the same link speed.

Table 1 provides a summary of synchronous and asynchronous applications from the leading storage vendors.

Table 1 Storage Vendor Applications

Storage Vendor	Synchronous Replication Application	Asynchronous Replication Application
EMC	SRDF Synchronous and Semi Synchronous mode	SRDF Asynchronous mode
IBM	PPRC	XRC
HP	Continuous Access-Enterprise Virtual Array (CA-EVA); originally DRM	CA-EVA
Hitachi Data Systems	TrueCopy	TrueCopy

- Protocols to be Transported

Fibre Channel

Fibre Channel is a layered network protocol suite developed by ANSI and typically used for networking between host servers and storage devices and between storage devices. Transfer speeds come in two rates: 1.0625 Gbps and 2.125 Gbps. With single-mode fiber connections, Fibre Channel has a maximum distance of about 6.2 miles (10 kilometers).

The primary problem with transparently extending Fibre Channel over long distances stems from its flow-control mechanism and its potential effect on an application's effective input-output performance. To ensure that input buffers do not get overrun and start dropping Fibre Channel frames, a system of buffer-to-buffer credits provides a throttling mechanism to the transmitting storage or host devices to slow down the flow of frames. The general principle is that one buffer-to-buffer credit is required for every 1.2 miles (2 km) to sustain 1 Gbps of bandwidth, and one buffer-to-buffer credit is required for every 0.6 miles (1 km) between two interfaces on a link for 2 Gbps. These numbers are derived using full-size Fibre Channel frames (2148 bytes); if using smaller frames, the number of buffer credits required significantly increases. Without SAN extension methods in place, a typical Fibre Channel fabric cannot exceed 6.2 miles (10 km). To achieve greater distances with Fibre Channel SAN extensions, SAN switches such as the Cisco[®] MDS 9000 Family are always used to provide additional inline buffer credits. These credits are required because most storage devices support very few credits (fewer than 10) of their own, thereby limiting the ability to directly extend a storage array.



ESCON

Enterprise Systems Connection (ESCON) is a 200-Mbps unidirectional serial bit transmission protocol used to dynamically connect IBM or IBM-compatible mainframes with their various control units. ESCON provides non-blocking access through either point-to-point connections or high-speed switches called ESCON directors. ESCON performance is seriously affected if the distance spanned is greater than 5 miles (8 km).

FICON

FICON (IBM Fiber Connection) is the next-generation bidirectional channel protocol used to connect mainframes directly with control units or ESCON aggregation switches (ESCON directors with a bridge card). FICON runs over Fibre Channel at a data rate of 1.062 Gbps. One of the main advantages of FICON is its performance stability over distances. FICON can reach a distance of 62 miles (100 km) before experiencing any significant drop in data throughput.

- *Present and future demand*—The type and quantity (density) of SAN extension protocols to be transported, as well as specific traffic patterns and restoration techniques need to be considered. The type and density requirements will help determine the technology options and specific products that should be implemented. Growth should be factored into the initial design to ensure a cost-effective upgrade path.
- *Distances*—Because of the strict latency requirements of SAN applications, especially those found in synchronous environments, performance could be severely affected by the type of SAN extension technology implemented. Table 2, near the end of this document, provides some guidance on distance restrictions and other considerations for each technology option.
- *Recovery objectives*—A business continuity strategy can be implemented to reduce an organization's annual downtime and to reduce the potential costs and intangible issues associated with downtime. Recovery with local or remote tape backup could require days to implement, whereas geographically dispersed clusters with synchronous mirroring can result in recovery times measured in minutes. Ultimately, the business risks and costs of each solution have to be weighed to determine the appropriate recovery objective for each enterprise.
- *Original storage manufacturer certifications*—Manufacturers such as IBM, EMC, HP, and Hitachi Data Systems require rigorous testing and associated certifications for SAN extension technologies and for specific vendor products. Implementing a network containing elements without the proper certification may result in limited support from the manufacturer in the event of network problems.

Infrastructure and Service Availability

When evaluating the various SAN extension alternatives, another variable is the availability of existing enterprise infrastructure or service provider solutions. The following sections describe technology alternatives that might be relevant.

In some instances, an enterprise might have existing equipment (Fibre Channel switches, channel extenders, IP routers, SONET/SDH) that could influence the final technology decision.

- *Dark fiber*—The enterprise builds a privately owned and managed Fibre Channel, coarse wavelength-division multiplexing (CWDM), dense wavelength-division multiplexing (DWDM), or SONET/SDH network.
- *Managed service offering*—The service provider builds, operates, and maintains a CWDM, DWDM, or SONET/SDH infrastructure dedicated to a specific customer.
- *Wavelength service*—The service provider leases a wavelength (or wavelengths) to a customer over a shared DWDM/CWDM network. Multiple customers share a common service provider infrastructure.



- *SONET/SDH service*—The service provider leases Synchronous Transport Signal level n (STS-n) and Synchronous Transport Module level n (STM-n) capacity to a customer over a shared SONET/SDH network. Depending upon the enterprise's requirements or existing equipment and service availability, the service provider might offer DS-3, Optical Carrier level n (OC-n)/STM-n, Gigabit Ethernet, or Fibre Channel over SONET/SDH client interfaces.
- *IP network-based service*—The enterprise might implement Fibre Channel Interface Protocol (FCIP) over a privately owned and operated IP network or over a service provider IP network. Quality-of-service (QoS) and service-level-agreement (SLA) requirements as well as support for various frame sizes (including jumbo frames) need to be considered, depending on the specific application requirements.

Overview of SAN Extension Technology Alternatives


Dark Fiber

In some instances, an enterprise might own or have access to dark fiber. If so, more options become available for extending the SAN across the metro and beyond. Depending upon deployment considerations, as outlined in the previous section, an enterprise might choose to implement one or more of the following solutions over dark fiber:

- A native Fibre Channel solution, extended directly over dark fiber.
- A native Fibre Channel solution, extended over CWDM. (See the following CWDM section for a more complete discussion.)
- A DWDM system capable of transporting several protocols over extended distances using a single fiber pair. (See the following DWDM section.)
- A SONET or SDH system, also capable of transporting several protocols over extended distances. (See the following SONET/SDH section.)

Because the last three options will be discussed in greater detail in following sections, the remaining portion of this section will take a closer look at the first option: native Fibre Channel over dark fiber. Figure 2 illustrates Cisco MDS 9000 multilayer fabric switches directly connected by diverse single-mode-fiber paths.

Figure 2
Native Fibre Channel over Dark Fiber





For more details about the Cisco MDS 9000 Family of products, visit the following URL:

<http://www.cisco.com/en/US/products/hw/ps4159/ps4358/index.html>

Technical Considerations

- Fibre Channel solution, 1 or 2 Gbps per fiber pair
- The underlying protocol (SAN or application) is responsible for failover protection
- Using short-wave (850-nanometer) Small Form-Factor Pluggable (SFP) optics in the Cisco MDS 9000, the reach is limited to about 984 feet (300 meters) over 62.5/125-micron multimode fiber
- Using short-wave (850-nm) SFPs, the reach is limited to about 1640 feet (500 meters) over 50/125-micron multimode fiber
- Using long-wave (1310-nm) SFPs, the reach is limited to about 6.2 miles (10 km) over 9/125-micron single-mode fiber
- Using CWDM (up to eight wavelengths) SFPs, the reach is limited to about 56 miles (90 km) over 9/125-micron single-mode fiber (distance assumes dark fiber with no CWDM filters)
- To cost-effectively accommodate future growth without increasing the number of fibers or replacing SFPs, consider using CWDM SFPs over dark fiber on day 1 (even if the extra reach is not required). This would allow the enterprise to easily scale bandwidth by adding passive Cisco CWDM filters in the future without requiring a replacement of the SFPs in the Cisco MDS 9000. The insertion loss of the CWDM filters needs to be accounted for in the initial design if this is a likely growth strategy.
- Cisco Virtual SAN (VSAN) technology in the Cisco MDS 9000 Family can be used to transport multiple SAN fabrics over the same fiber pair without actually merging the fabrics. As many as 1000 different virtual fabrics can be deployed and transported over a common fiber pair. VSANs provide an efficient mechanism to use the available transport bandwidth to extend multiple SANs for data replication services where each replication instance may only require a fraction of the available bandwidth on its own.

Note: Distance values are approximate. Please consult online documentation at the following URL for Cisco SFP interfaces, or contact your Cisco sales representative:

http://www.cisco.com/en/US/products/hw/modules/prod_module_category_home.html

Recommendations for Use

The solution illustrated in Figure 2 can offer several benefits, assuming that an enterprise has enough dark fiber to meet the capacity requirements and that the distance between locations is less than about 56 miles (90 km). Because this solution eliminates any sort of external transmission or media conversion from the network, this approach is among the simplest to implement and easiest to manage. This solution can offer low up-front costs, short deployment intervals (assuming the fiber is already in place or is being leased), and low operational expenses, primarily because of the reduced equipment requirements.

Coarse Wavelength Division Multiplexing

CWDM offers a convenient, cost-effective mechanism to increase network bandwidth while minimizing the number of fibers consumed between locations. By using the CWDM SFPs in the Cisco MDS 9000 Family (as discussed in the previous section) and by adding the Cisco passive CWDM filter solution, enterprises can easily and cost-effectively scale network bandwidth. The Cisco CWDM solution offers up to eight channels over a single fiber pair, each capable



of carrying 1- or 2-Gbps Fibre Channel. Cisco also offers a single-fiber, bidirectional CWDM filter that supports as many as four channels. In the event an enterprise is severely fiber-constrained or can only justify the expense of leasing a single fiber, the bidirectional filter offers the perfect solution.

It should be noted that the CWDM filter solution described in this section could also be used in conjunction with CWDM SFPs and gigabit interface converters (GBICs) available for the industry-leading Cisco Catalyst® switches. This allows an enterprise to freely mix Gigabit Ethernet data connectivity and SAN extension connectivity over the same fiber pair(s) for greater network optimization.

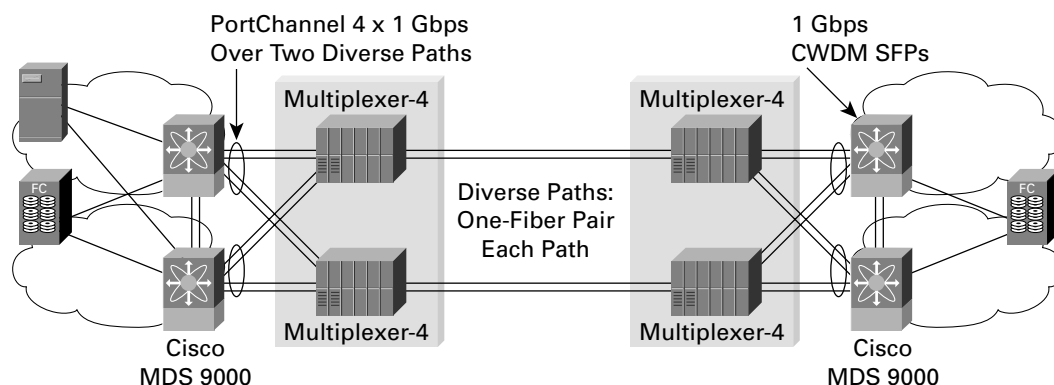
For more details about the Cisco CWDM solution, visit the following URL:

http://www.cisco.com/en/US/tech/tk713/tk772/tech_relevant_products.html

Figure 3 illustrates Cisco MDS 9000 multilayer fabric switches connected by CWDM SFPs and filters.

Figure 3

CWDM Filter Solution Using Cisco MDS 9000 Switches



Technical Considerations

- Fibre Channel solution, 1 or 2 Gbps per fiber pair.
- The application is responsible for failover protection.
- Using CWDM SFPs and filters, reach is limited by fiber attenuation and filter insertion losses. In Figure 3, the distance between locations would be limited to about 41 miles (66 km). This differs from the 56-mile (90-km) reach in the previous dark fiber solution because of the insertion loss of the 4-channel CWDM multiplexer/demultiplexer at each location.
- Currently limited to a maximum of eight channels over a single-fiber pair or four channels over a single fiber.
- CWDM channels cannot be amplified to extend distances.
- Cisco VSAN technology on the Cisco MDS 9000 Family can be used to transport multiple SAN fabrics over the same CWDM wavelength without actually merging the fabrics. As many as 1000 different virtual fabrics can be deployed and transported over a common wavelength. VSANs provide an efficient mechanism to use the available transport bandwidth to extend multiple SANs for data replication services where each replication instance may only require a fraction of the available wavelength bandwidth on its own.
- Cisco PortChanneling Inter-Switch Link (ISL) aggregation technology can be used to aggregate several extended Fibre Channel ISLs into a logical ISL for more efficient bandwidth use. Using the CWDM solution, as many as eight ISLs can be extended—each using a different wavelength—and aggregated to form a single 16-Gbps ISL.



Note: To determine more precise reach limitations for a CWDM solution, the exact fiber loss, network topology, and CWDM filter requirements would need to be determined. Please consult online documentation for the Cisco CWDM solution at the following URL or contact your local Cisco sales representative.

http://www.cisco.com/univercd/cc/td/doc/product/gbic_sfp/psv_doc/78_14167.htm

Recommendations for Use

The solution illustrated in Figure 3 is similar to the first scenario in Figure 2, with the exception that CWDM SFPs and filters have been added. This solution has many of the same advantages as described in the first scenario with the added benefit of reduced fiber requirements. Cisco offers flexible CWDM filter options to support 1-fiber or 2-fiber implementations, single channel add/drop, as well as support for various network topologies.

Dense Wavelength Division Multiplexing

Like CWDM, DWDM offers a cost-effective mechanism to increase network bandwidth while minimizing the number of fibers consumed between locations. However, in many instances the Fibre Channel-over-dark-fiber or Fibre Channel-over-CWDM solutions (as previously described) might not fully satisfy the SAN extension design requirements for carrying mission-critical data. DWDM offers a solution that provides high-density aggregation over a high-speed, low-latency network that can support important protocols such as Fibre Channel, FICON, ESCON, and the IBM Sysplex and Coupling Link protocols. In addition, DWDM offers the ability to support not only storage but also the network consolidation of multiple other protocols, including data (Gigabit Ethernet, 10 Gigabit Ethernet), voice, and video.

Cisco metro DWDM systems offer greater network capacity (32 DWDM channels for a maximum of 320 Gbps, compared with 8 CWDM channels for a maximum of 16 Gbps), high-density service aggregation, flexible transponder options allowing for service transparency across a wider range of client interfaces, comprehensive service protection options, and amplification as required to support extended distances. As examples of the service aggregation capabilities offered, Cisco ONS 15530 can aggregate as many as 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 as many as 256 Fibre Channel/FICON/Gigabit Ethernet interfaces or as many as 1280 ESCON channels per fiber pair, allowing customers to maximize the carrying capacity of their fibers. The use of Fibre Channel buffer-to-buffer credit mechanism in a DWDM deployment allows extension of storage applications across the metro DWDM network.

Figure 4 illustrates the Cisco metro DWDM options. For a more comprehensive view of the supported network topologies and service interface options, please ask your Cisco sales representative or consult Cisco online technical documentation at the following URLs:

<http://www.cisco.com/en/US/products/hw/optical/ps2011/index.html>

<http://www.cisco.com/en/US/products/hw/optical/ps2006/ps5320/index.html>



Figure 4
Cisco Metro DWDM Options

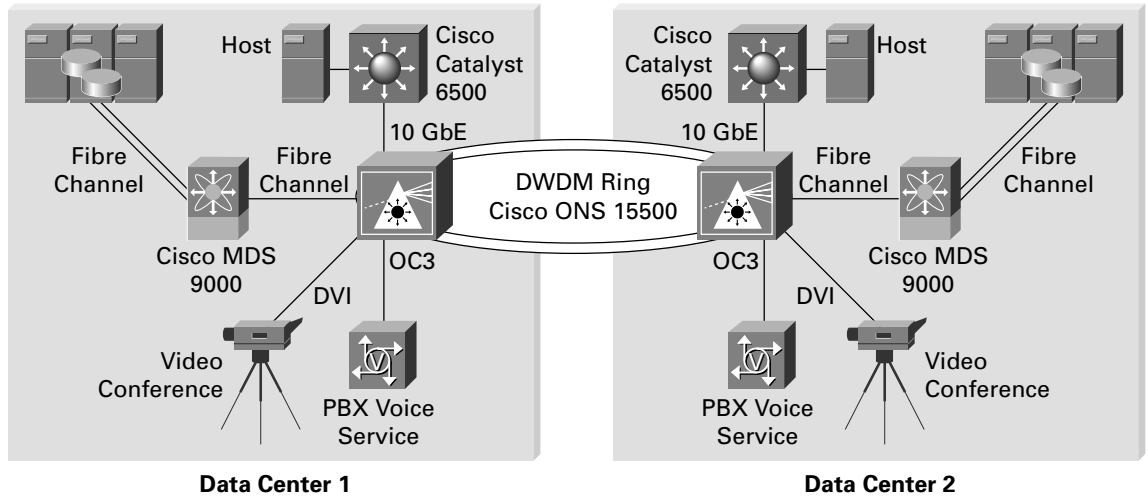
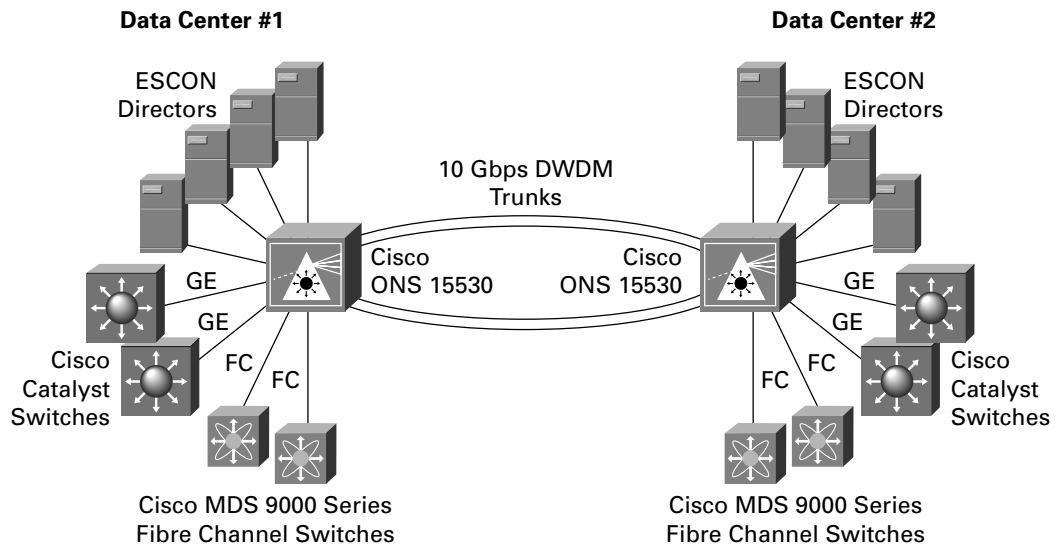


Figure 5 shows a Cisco ONS 15530 metro DWDM being used for high-density service aggregation of ESCON, Gigabit Ethernet, and Fibre Channel.

Figure 5
Cisco ONS 15530 Metro DWDM





Technical Considerations

The following technical considerations might result in DWDM being the required (or most appropriate) technology choice for SAN extension:

- The enterprise owns or has access to limited amounts of dark fiber with a need for more capacity than can be provided by CWDM.
- Requirement for a large number and wide variety of service types to be transported efficiently over a common transmission infrastructure. These services could include storage protocols like Fibre Channel, FICON, ESCON, data over Gigabit Ethernet, 10 Gigabit Ethernet, OC-n/STM-n; as well as voice or digital video. The Cisco ONS 15500 Series offers the industry's highest service density for business continuance solutions.
- High-speed, low-latency connections for synchronous storage applications from EMC, IBM, HP, Hitachi Data Systems, and others. The Cisco ONS 15500 Series is certified by all these major vendors for guaranteed interoperability.
- Extended distance support through a combination of optical amplification and protocol buffer-to-buffer credit mechanism as supported by the Cisco ONS 15500 Series.
- Protocol and optical performance monitoring and diagnostic capability is desirable from an operations perspective; these advantages are unique with the Cisco ONS 15500 Series.
- Network resiliency requirements of five nines or higher. The Cisco ONS 15500 Series offers a multitude of options for both wavelength and service protection: optical unidirectional path switched ring (UPSR), 1+1 optical trunk protection, 1+1 client protection, 1+1 wavelength protection, wavelength optical splitter protection, and even 1+1 service protection for some service types.
- Cisco VSAN technology in the Cisco MDS 9000 Family can be used to transport multiple SAN fabrics over the same DWDM wavelength without actually merging the fabrics. As many as 1000 different virtual fabrics can be deployed and transported over a common wavelength. VSANs provide an efficient mechanism to use the available transport bandwidth to extend multiple SANs for data replication services where each replication instance may only require a fraction of the available wavelength bandwidth on its own.
- Cisco PortChanneling ISL aggregation technology can be used to aggregate several extended Fibre Channel ISLs into a logical ISL for more efficient bandwidth use. Using the DWDM solution, as many as 16 ISLs (PortChannel maximum) can be extended—each using a different wavelength—and aggregated to form a single 32-Gbps ISL.

Recommendations for Use

Cisco metro DWDM systems are designed to carry mission-critical storage, data, and existing applications over a highly available, metro optical network. These platforms are certified or in the process of certification by leading system and storage vendors such as IBM for its GDPS applications, EMC for its SRDF applications, Hitachi Data Systems for its TrueCopy applications, Hewlett Packard for its Continuous Access applications, and others. Cisco metro DWDM systems offer the greatest degree of flexibility and investment protection when considering a wide array of SAN extension requirements.

SONET/SDH

SONET/SDH-based networks have been widely deployed in service provider and enterprise private networks alike since the mid-1980s. SONET/SDH technology is well-known and is trusted as a reliable and cost-effective means of transporting an assortment of electrical and optical client interface types. In recent years, additional capabilities such as the integration of Ethernet services and DWDM have further increased applicability for SONET/SDH network



elements. With a demonstrated ability to evolve to meet ever-changing network requirements and its widespread deployment, SONET/SDH remains a useful technology for delivering many types of services, including SAN extension (Figure 6).

Cisco has been an industry leader in introducing evolutionary capabilities into SONET/SDH networks. More than 33,000 Cisco ONS 15454 Multiservice Provisioning Platform (MSPP) systems have been deployed, including more than 45,000 Ethernet or Fast Ethernet ports and more than 4000 Gigabit Ethernet ports. Cisco Systems® plans to continue this multiservice-over-SONET/SDH (MSOS) leadership by introducing the SL Series Fibre Channel-over-SONET/SDH module. The SL Series module will allow Cisco to provide 100 percent coverage of the Fibre Channel transport space, providing end-to-end coverage of data center and enterprise storage needs for metropolitan and regional networks and beyond.

For more details about Cisco ONS 15454 MSPP platforms and multiservice-over-SONET/SDH strategy, visit the following URLs:


<http://www.cisco.com/en/US/products/hw/optical/ps2006/ps2010/index.html>

<http://www.cisco.com/en/US/products/hw/optical/ps2006/ps2008/index.html>

http://www.cisco.com/en/US/netsol/ns110/ns112/ns114/ns116/networking_solutions_white_paper09186a008011391e.shtml

Figure 6

Cisco Fibre Channel over SONET/SDH



Technical Considerations

The following list of technical issues should be studied when considering a Fibre Channel-over-SONET/SDH solution. (Options for using SONET/SDH in conjunction with Fibre Channel over IP [FCIP] will be discussed later in this document.)

- The enterprise owns or has access to dark fiber and needs to efficiently multiplex various service types over a common transmission infrastructure. Service interface needs may include: DS-1/E-1, DS-3/E-3/E-4, OC-n/STM-n, 10/100-Mbps Ethernet, and Gigabit Ethernet (including Fibre Channel for SAN extension).
- The enterprise selects a managed service from a local provider that offers a “direct” Fibre Channel-over-SONET/SDH solution. Technical considerations such as the mixture of service interface requirements, distances involved, application requirements, service endpoints, and the availability of existing SONET/SDH infrastructure would all be determining factors for the service provider.



- The enterprise (or service provider) desires an integrated solution to address the service interfaces as previously described for ease of installation and management.
- The Fibre Channel bandwidth requirements can be multiples of STS-n (SONET) STM-n (SDH) bandwidth, including DS-3 mapped for SONET up to full line rate 1- or 2-Gbps Fibre Channel. This granularity can offer the enterprise a balance between scalable bandwidth options and the cost of service. This option may be especially attractive to small and midsize enterprises that cannot justify the expense of a dedicated DWDM or wavelength service for their SAN extension.
- Distance requirements between storage locations and allowable delay for the application will affect the technology choice. For example, the maximum allowable delay (one way) for a synchronous application is about 500 microns. This delay is divided between fiber propagation (about 5 microns per 0.6 miles, or 1 km) and processing delays at the transmission nodes. Fibre Channel-over-SONET/SDH processing at end nodes typically adds 20–25 microns of processing delay, and an intermediate (through node) may introduce about 10 microns of delay. One of the advantages SONET/SDH can offer over higher-layer methods (such as FCIP) is delay predictability. Even so, worst-case delay must be considered, including the fiber and node processing delays under protection switch scenarios. Although the issue is beyond the scope of this paper, it should be mentioned that mechanisms such as large buffer-to-buffer credits (in the Fibre Channel switch) and spoofing of Fibre Channel acknowledgments by the transmission layer are both mechanisms to extend the effective throughput over longer distances.
- SONET/SDH protection mechanisms such as UPSR and subnetwork connection protection (SNCP), 2- and 4-fiber bidirectional line switched ring (BLSR) and multiplex section-shared protection ring (MS-SPR), 1+1 automatic protection switching/multiplex section protection (APS/MSP), or path-protected mesh networking (PPMN) might be necessary to achieve the required level of network resiliency.
- Cisco VSAN technology in the Cisco MDS 9000 Family can be used to transport multiple SAN fabrics over the same ISL extended over SONET/SDH without actually merging the fabrics. As many as 1000 different virtual fabrics can be deployed and transported over a common circuit. VSANs provide an efficient mechanism to use the available transport bandwidth to extend multiple SANs for data replication services where each replication instance might require only a fraction of the available circuit bandwidth on its own.
- Cisco PortChanneling ISL aggregation technology can be used to aggregate several extended Fibre Channel ISLs into a logical ISL for more efficient bandwidth use. Using the SONET/SDH solution, as many as 16 ISLs can be extended (PortChannel maximum)—each using a different SONET/SDH circuit—and aggregated to form a single 32-Gbps ISL.

Recommendations for Use

The extension of SANs over SONET/SDH allows both the service provider and the enterprise to make use of the enormous installed base of SONET and SDH networks. Additional advantages for the enterprise include cost-effectiveness through scalable SONET/SDH bandwidth options, suitability for SAN extension over greater distances, and guaranteed performance and quality of service. Additional Fibre Channel-over-SONET/SDH benefits for the service provider include the ability to gain incremental revenue from existing SONET/SDH infrastructure, rapid time to service, and integrated management and diagnostic capabilities.



Fibre Channel over IP

Fibre Channel over IP (FCIP) is a protocol specification developed by the Internet Engineering Task Force (IETF) that allows a device to transparently tunnel Fibre Channel frames over an IP network. An FCIP gateway or edge device attaches to a Fibre Channel switch and provides an interface to the IP network. At the remote SAN island, another FCIP device receives incoming FCIP traffic and places Fibre Channel frames back on to the SAN. FCIP devices provide Fibre Channel expansion port connectivity, creating a single Fibre Channel fabric.

One of the primary advantages of FCIP for remote connectivity is its ability to extend distances using TCP/IP. However, distance achieved at the expense of performance is an unacceptable trade-off for IT organizations that demand full utilization of expensive WAN bandwidth. IETF RFC 1323 adds TCP options for performance, including the ability to scale the standard TCP window size up to 1 gigabyte. As the TCP window size widens, the sustained bandwidth rate across a long haul (more latency) TCP connection increases. With the Cisco MDS 9000 IP Storage Services Module, the maximum TCP window size can be configured to be up to 32 MB. From early field trials, distances spanning more than 3600 miles were feasible for disk replication in asynchronous mode. Even greater transport distances are achievable. Theoretically, a 32-MB TCP window with a 1-Gbps bandwidth can be extended over 31,069 miles (50,000 km) with 256 milliseconds of latency.

Another advantage of FCIP is the ability to use existing infrastructures that provide IP services. For IT organizations that are deploying routers for IP transport between their primary data centers and their disaster-recovery sites, with QoS enabled, FCIP can be used right away for SAN extension applications. For larger IT organizations that have already invested in or are leasing SONET/SDH infrastructures, FCIP can provide the most flexibility in adding SAN extension services because no additional hardware is required.

For enterprises required to deploy SAN extension across various remote offices with the central office, a hub-and-spoke configuration of FCIP connections is also possible. In this manner, applications such as disk replication can be made between the disk arrays of each individual office and the central office's disk array but not necessarily between the individual offices' disk arrays themselves. With this scenario, the most cost-effective method of deployment would be to use FCIP along routers.

For FCIP over SONET, various configurations are possible. Figure 7 illustrates a basic configuration, in which the Gigabit Ethernet port of the Cisco MDS 9000 IP Storage Services Module is connected directly to the Gigabit Ethernet port of the Cisco ONS 15454. This scenario assumes a dedicated Gigabit Ethernet port is available on the Cisco ONS 15454. Another possible configuration is to include routers between the Cisco MDS 9000 and the Cisco ONS 15454 as shown in Figure 8. In this case, the Cisco ONS 15454 might not necessarily have a Gigabit Ethernet card, and thus a router is required to connect the Gigabit Ethernet connection of the Cisco MDS 9000 to the Cisco ONS 15454. With VPN Acceleration Modules on the routers, this scenario can also allow for data compression. In cases where encryption is a business requirement, Cisco Catalyst 6500 Series switches with VPN Services Modules (VPNSMs) can also be added to the configuration as shown in Figure 9. Likewise, if both encryption and compression were required, both routers and Cisco Catalyst 6500 Series with VPNSMs can be incorporated as shown in Figure 10.



Figure 7
Cisco FCIP over SONET with Direct Connection to the Cisco ONS 15454




Figure 8
Cisco FCIP over SONET with Routers for Compression

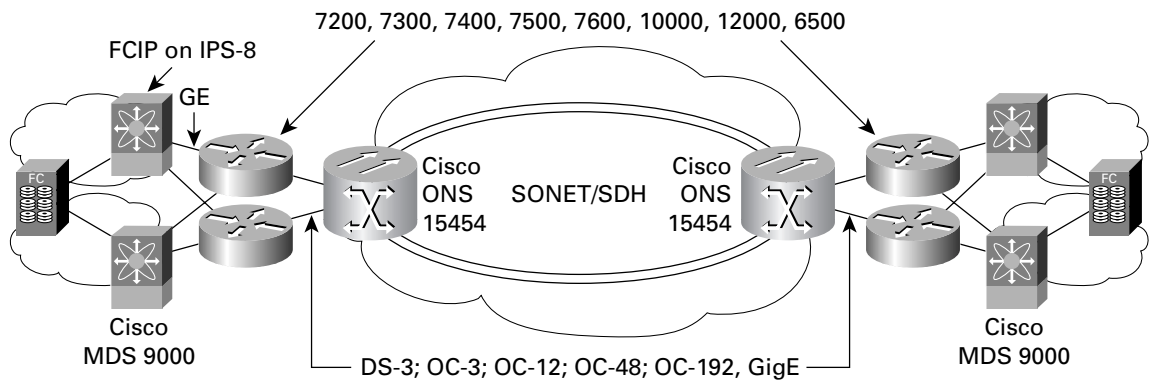


Figure 9
Cisco FCIP over SONET with Cisco Catalyst 6500 Series Switches for Encryption

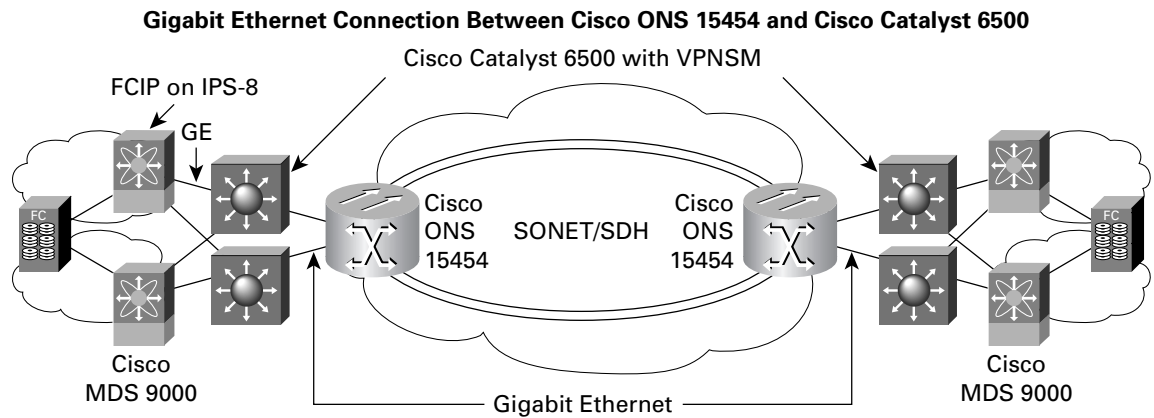
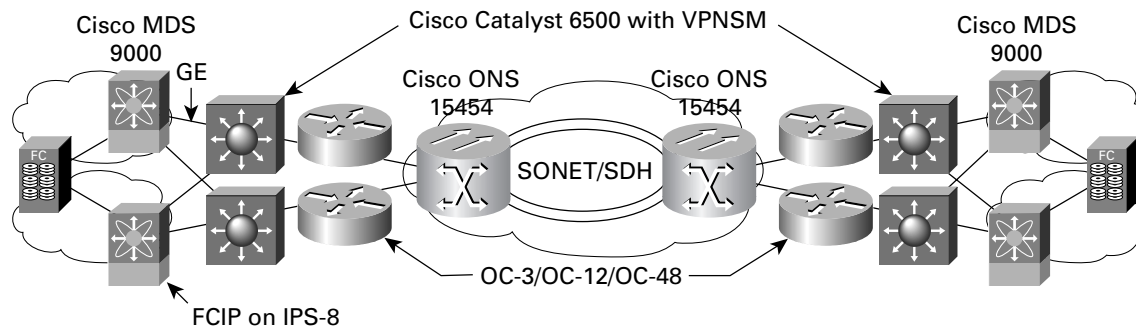




Figure 10

Cisco FCIP over SONET with Routers and Cisco Catalyst 6500 Series Switches for Encryption and Compression



Technical Considerations

The following technical issues should be addressed when considering Fibre Channel over IP:

- In cases where it is required to be able to switch between data and storage traffic in a limited capacity leased line, additional provisioning is not required. For enterprises that do not manage their own transport infrastructures, FCIP offers the most flexibility in handling both IP data and Fibre Channel storage traffic.
- In cases of considerable cost restrictions, FCIP would probably be the most cost-effective solution because it can make use of all existing IP infrastructures. In the simplest case, a pair of routers with a leased line between two locations is all that is required to begin SAN extension.
- In cases where compression and encryption is required, IP router-based compression and encryption can be used. Compression techniques are popular in ultralong storage extension networks to save bandwidth.
- For cases in which multiple remote offices are required to replicate data back to a central office, a hub-and-spoke configuration of FCIP connections using routers is perhaps the most cost-effective method of SAN extension for multiple site disk replications.
- For FCIP over SONET, in cases where it is cost-prohibitive to have a single leased OC-48 circuit dedicated to Fibre Channel over SONET, multiples of STS-n (SONET) STM-n (SDH) bandwidth can be used to address the Fibre Channel bandwidth requirements for the SAN extension.
- Cisco VSAN technology on the Cisco MDS 9000 Family can be used to transport multiple SAN fabrics over the same ISL extended over IP without actually merging the fabrics. As many as 1000 different virtual fabrics can be deployed and transported over a common circuit. Cisco VSANs provide an efficient mechanism to use the available transport bandwidth to extend multiple SANs for data replication services where each replication instance may only require a fraction of the available circuit bandwidth on its own.
- Cisco PortChanneling ISL aggregation technology can be used to aggregate several extended Fibre Channel ISLs into a logical ISL for more efficient bandwidth use. Using the FCIP solution, as many as 16 ISLs can be extended (PortChannel maximum)—each using a different IP circuit—and aggregated to form a single 16-Gbps ISL.

For more details about Cisco MDS 9000 platforms, visit the following URLs:

<http://www.cisco.com/en/US/products/hw/ps4159/ps4358/index.html>

http://www.cisco.com/warp/public/cc/pd/ps4159/ps4358/prodlit/90ip_ds.htm



Recommendations for Use

The extension of SANs using FCIP provides the most transparent method of SAN extension and allows both the service provider and the enterprise to make full use of the enormous installed base of IP infrastructure regardless of the actual transport. FCIP is ideally deployed in environments in which the flexibility of using both IP data and storage traffic is required. Also, in situations requiring encryption and compression, existing IP switching and routing infrastructures can be used without incurring additional costs.

Summary

The previous sections discussed SAN extension deployment considerations as well as various technology options. Table 2 summarizes each technology, the primary criteria for evaluating the technology options, and a list of Cisco SAN extension products.

Table 2 SAN Extension Options

	Fibre Channel over Dark Fiber	Fibre Channel over CWDM	DWDM	SONET/SDH	FCIP
SAN protocols supported	Fibre Channel	Fibre Channel	Fibre Channel, FICON, ESCON, IBM Sysplex Timer, IBM Coupling Facility	Fibre Channel, FICON	Fibre Channel over IP
SAN distances supported (Fibre Channel/FCIP only for comparative purposes)	56 miles (90 km)** **Assumes the use of CWDM SFPs, no filters	37–41 miles (60–66 km)** **Assumes point-to-point configuration	Up to 124 miles (200 km)* * Actual distances depend on the characteristics of fiber used	1740 miles (2800 km) with buffer credit support	Distance limitation dependent on the latency tolerance of the end application **Longest tested distance is 3604 miles (5800 km)
SAN bandwidth options (per fiber pair)	1-Gbps Fibre Channel (1.0625 Gbps), 2-Gbps Fibre Channel (2.125 Gbps)	1-Gbps Fibre Channel (1.0625 Gbps), 2-Gbps Fibre Channel (2.125 Gbps), up to 8 channels	Up to 256 Fibre Channel/FICON channels, up to 1280 ESCON channels, up to 32 channels at 10 Gbps	1-Gbps Fibre Channel (1.0625 Gbps), up to 32 channels with sub-rating, 2-Gbps Fibre Channel (2.125 Gbps), up to 16 channels with sub-rating	1-Gbps Fibre Channel (1.0625 Gbps)
Network protection options	Fabric Shortest Path First (FSPF), PortChannel, isolation with VSANs	FSPF, PortChannel, isolation with VSANs	Client, 1+1, y-cable, switch fabric protected, switch fabric protected trunk, and protection switch module, unprotected	UPSR/SNCP, 2F and 4F BLSR/MS-SPR, PPMN, 1+1 APS/MSP, unprotected	VRRP, redundant FCIP tunnels, FSPF, PortChannel, isolation with VSANs

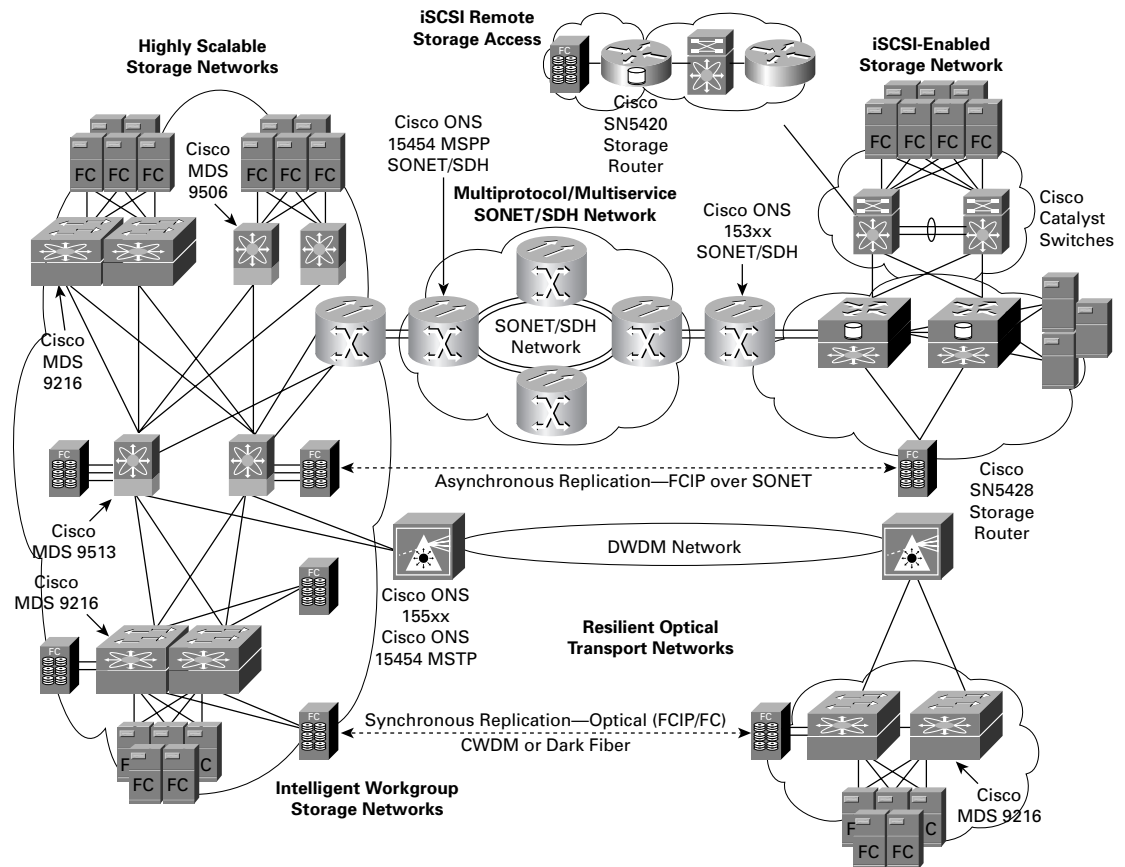


Table 2 SAN Extension Options (Continued)

	Fibre Channel over Dark Fiber	Fibre Channel over CWDM	DWDM	SONET/SDH	FCIP
Other protocols supported	—	CWDM filters also support Gigabit Ethernet	OC-3/12/48/192, STM-1/4/16/64, Gigabit Ethernet, 10 Gigabit Fast Ethernet, D1 Video	DS-1, DS-3, OC-3/12/48/192, E-1, E-3, E-4, STM-1E, STM-1/4/16/64, 10/100-Mbps Ethernet, Gigabit Ethernet	—
Cisco products	Cisco MDS 9000	Cisco MDS 9000, CWDM SFPs, and filters	Cisco ONS 15530, Cisco ONS 15540, Cisco ONS 15454 MSTP	Cisco ONS 15454 MSPP (SONET and SDH)	Cisco MDS 9000, Cisco 7200, 7400 series

Figure 11 illustrates the Cisco end-to-end solution set for storage networking and SAN extension.

Figure 11
Cisco SAN Solutions



Cisco Systems, Inc.

All contents are Copyright © 1992–2003 Cisco Systems, Inc. All rights reserved. Important Notices and Privacy Statement.



Regulatory requirements and business economics are making business continuance solutions mandatory for enterprises. Various methods and technology options are available for SAN extension, each with its own merits and limitations. Recognizing that no single solution is an appropriate fit for all applications, Cisco has made significant investments in both storage and optical technologies to meet customer needs for an end-to-end storage network that meets the most demanding application requirements.

The Cisco MDS 9000 Family provides a full line of products to meet requirements for storage networks of all sizes and architectures. To efficiently interconnect these SAN islands, Cisco offers a complete portfolio of optical transmission platforms that span CWDM, DWDM, and SONET/SDH technologies. The Cisco Catalyst 4000 Series, the Cisco ONS 15500 Series and the Cisco ONS 15454 each has its unique advantages in network diameter, service variety, service density, resiliency, performance monitoring, network management, and cost. Cisco FCIP options allow an enterprise to use its existing IP infrastructure for further network convergence, allowing for investment protection and streamlining of capital and operational expenses.

Cisco SAN extension solutions offer open interfaces, with certifications from all major storage vendors, including EMC, IBM, HP, and Hitachi Data Systems. Additionally, enterprises can integrate these solutions easily and quickly into their existing networks using the same Cisco management tools (CiscoWorks, CiscoView, Cisco Transport Manager, and Cisco IOS[®] command-line interface), further reducing training and IT costs.

References

Cisco storage networking products:

<http://www.cisco.com/en/US/products/hw/ps4159/index.html>

Cisco optical platforms:

<http://www.cisco.com/en/US/products/hw/optical/index.html>

Cisco optical network management software:

<http://www.cisco.com/en/US/products/sw/opticsw/index.html>

Cisco storage networking solutions:

http://www.cisco.com/en/US/netsol/ns110/ns258/net_solution_home.html

Cisco optical networking solutions:

http://www.cisco.com/en/US/netsol/ns110/ns112/net_solution_home.html

Cisco business continuous networking solutions:

http://www.cisco.com/en/US/netsol/ns110/ns112/ns113/ns283/networking_solutions_package.html

Cisco data center networking solutions:

http://www.cisco.com/en/US/netsol/ns110/ns53/net_solution_home.html



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. Catalyst, Cisco, Cisco IOS, Cisco Systems, and the Cisco Systems logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company.
(0304R) AW/LW5033 0903