White Paper

# Interconnecting Fibre Channel SANs Over Optical and IP Infrastructures

*An Evolutionary Approach to Extending Storage Over IP Metropolitan and Wide-Area Networks*

# Table of Contents

## Abstract

Interconnecting Storage Area Networks (SANs) over Optical and Internet Protocol (IP) infrastructures enables cost-effective, high-performance remote storage services. As storage-consolidation, disaster-recovery, and business-continuance applications take advantage of SAN infrastructures, Fibre Channel SAN deployment accelerates and interconnection demand grows. As a result, the requirement for ubiquitous, cost-effective, high-performance network interfaces becomes paramount for delivering mission-critical services. This creates new opportunities and delivers increased benefits to enterprise, service-provider, and carrier markets.

Though transporting Fibre Channel traffic across dedicated fiber networks is possible, remote access to terabyte-storage resources has been cost prohibitive. Network Attached Storage (NAS) provides a solution for transporting file system data over long distances; however, NAS has significant performance limitations when transporting large blocks of storage data. Since NAS is a file-based storage-networking service, it lacks the scalability to meet the performance requirements of today's high-bandwidth remote-storage applications. Requirements for wire-speed throughput, improved cost efficiencies, and aggregation of data and storage services drive the need for better solutions to remote storage applications.

Gigabit Ethernet can extend the SAN infrastructure across Metropolitan-Area Networks (MANs) and Wide-Area Networks (WANs). It also addresses performance issues by providing higher bandwidth at greater efficiencies and lower costs. Most telecommunications carriers offer Gigabit speeds through SONET/SDH infrastructures, with OC-48 (2.5Gbps) at most central offices and OC-192 (10Gbps) throughout the core of their fiber-based telecommunications infrastructures. This provides opportunities to transport Gigabit Ethernet links over carrier infrastructures that are more cost effective than dedicated fiber networks, but still too costly for most small-to-medium enterprises.

Next-generation carrier networks offer leased fiber and managed IP services that operate at gigabit speeds — often at lower cost points — and provide optical packet-switched infrastructures, with the inherent reliability of a mesh-network architecture. These alternatives empower users to leverage the existing infrastructure for remote access to centralized storage resources, with instant improvements in performance and significant cost reductions over dedicated fiber links.

The evolution from Gigabit Ethernet to 10 Gigabit Ethernet addresses the need of connecting storage to dark-fiber networks and SONET-optical infrastructures. The adoption rate of this technology promises to be an evolutionary rather than a revolutionary process, offering several key advantages with each advance in technology. The initial market opportunity for each of these phases is significant, with SAN-extension services representing well over a multi-billion market opportunity.

## Introduction

SANs are high-speed networks that enable direct connection between heterogeneous storage devices and servers. The high-speed direct connectivity that SANs provide offers:

- Greater flexibility
- Improved data management
- Increased system throughput
- Simplified and reduced costs of storage administration

Conventional enterprise networks can neither protect nor prioritize high-bandwidth storage traffic from existing data traffic. Moreover, storage traffic would choke most Local Area Networks (LANs).

One of the major benefits of a SAN is its ability to provide an independent storage network within the data center, thereby sidestepping traffic and congestion on a company's LAN. IP-SANs extend the benefits of the SAN beyond the data center, enabling the use of cost-effective remote storage applications for Enterprise and Storage Service Provider deployments.

## SAN Benefits

The SAN architecture was created to address issues associated with scalability, reliability, and performance within a data center. By externalizing the storage traffic from the LAN, the SAN removes many bottlenecks associated with directly attached storage architectures. Key differentiators between SAN and other storage network architectures are:

- Fast data access
- Scalability
- Higher availability and reliability
- Improved storage resource sharing
- Extensive management capabilities

Figure 1 illustrates a typical SAN architecture.

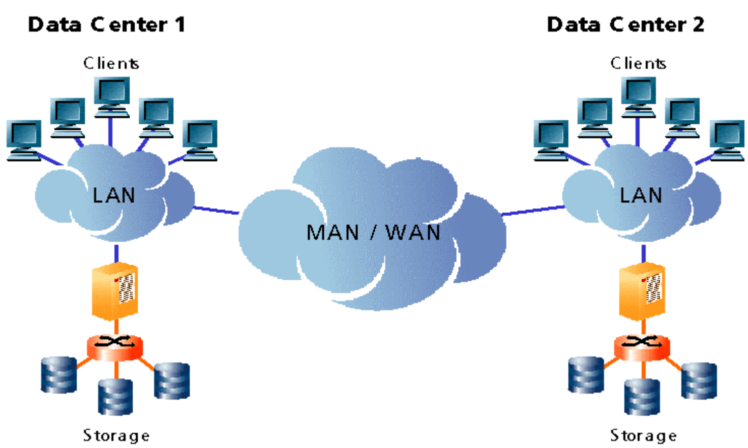


**Figure 1: Typical SAN Architecture**

A SAN solution is ideal for handling a corporation's multi-terabyte storage-networking requirements. Fibre Channel SANs have proven scalability to meet the high-end requirements of an Enterprise Data Center. Moreover, initial SAN installations typically range from 100 to 500 GB, although a SAN can be a good business decision for companies with as few as two Microsoft Windows NT servers.

SAN application server hardware can be upgraded without taking the shared storage resource out of service. Moreover, storage resources can be upgraded without halting the application. These advantages offer convenient management and reliability, while decoupling stored data from application hardware.

In addition, Fibre Channel devices — the core technology behind SAN architecture — are often dual-port solutions. Dual-port capability offers multiple redundant data connections, which provide excellent disaster protection for the network.

Furthermore, today's comprehensive management tools allow IT departments to manage capacity within SANs far better than IT departments can with direct-attached storage. With the introduction of newer virtual-storage technologies, SAN administrator can dynamically assign, reassign, and allocate storage resources in real time.

Additionally, a storage strategy that includes SANs is scalable, reliable, and gives IT departments the power to service both internal and external customers. As more SANs are deployed, the ability to interconnect them across long distances extends the capabilities of the storage infrastructure, enabling remote-storage applications and opening new market opportunities.

## SAN Limitations

While SANs are based on a high-speed, highly configurable architecture, they suffer from a few limitations.

Fibre Channel standards have been available for some time, but their implementation has been subject to interpretation. This interpretation has led to issues with interoperability, forcing many customers to turn to costly turnkey solutions for their SAN requirements. To remove the stigma of incompatibility, the Fibre Channel Industry Association (FCIA) has embarked upon the *SANMark* initiative, a certification program for Fibre Channel products.

Fibre Channel SANs are limited to distances of 10km (6.2mi) or less. While this is not an issue for SANs that are confined to individual data centers, it is a limiting factor for business-continuance and disaster-recovery applications. Remote-site data replication can lag significantly behind the primary storage — especially with NAS or when using magnetic tapes, which are often stored in vaults offsite using a transport method humorously referred to as Chevy Truck Access Method (CTAM).

Early SAN adopters have primarily been large organizations, whose business requirements center around geographic diversity and the need for heterogeneous connectivity. Fibre Channel SANs offer a top-tier solution for these campus and data-center deployments; however, very few solutions are available for connectivity outside the data center. And those solutions that do exist are expensive, software-based, and relatively slow.

Without the proper management tools, extending SANs over long distances can tax an Information Technology (IT) department significantly. An optical transport mechanism, such as DWDM, does not provide the necessary performance, configuration, and statistical data required to manage these remote data services. Without the ability to extend the management intelligence within a data center, an IT manager has no visibility of the network. Consequently, a simple outage can result in hours of expensive downtime.

## Market Drivers

The demand for solutions to extend SANs beyond the data center is building. In fact, rapid consolidation in almost every industry has created a need for connecting geographically distant stored data. Connecting these “Islands of Storage” is now a top priority for IT departments throughout the world.

Large enterprise and Internet data centers currently are extending SANs using dedicated fiber or Dense Wave Division Multiplexing (DWDM) links. While this may be an option for larger, more capitalized organizations or financial institutions, it is cost prohibitive for the majority of enterprises due to the expenses required to implement and deploy these types of dedicated storage infrastructures. To address this need, new, cost-effective solutions are quickly emerging into the marketplace.

The requirement for high-availability data in the e-business/e-commerce arena has spawned widely accepted solutions such as data and site mirroring. Current methodologies, however, are not suited to Metropolitan- or Wide-Area applications.

Backing up data to tape remains a critical component of a disaster recovery plan, despite data/site mirroring and other tools. The backup windows in effective disaster-recovery strategies should decrease as the volume of data increases. Full functionality despite limitations becomes critical as these windows of opportunity narrow.

Two key concerns of network design are Return on Investment (ROI) and Total Cost of Ownership (TCO). Network designers must be familiar with all the expenditures associated with managing a network. Adding complexity to an already complicated, understaffed operation does not serve the needs of the Chief Information Officer (CIO) in particular, nor the objectives of the company as a whole. The use of existing infrastructures is an important consideration in determining requirements and the way the operations are managed.

## Applications

SAN architecture allows a rapid deployment of resources without interrupting mission-critical applications — a requirement for emerging Storage Service Providers. Any current operation can benefit from a high-availability, fault-tolerant, scalable architecture. Mission-critical applications can include:

- On-Line Transaction Processing (OLTP)
- Data Replication and Mirroring
- Disaster Recovery
- Data Warehousing
- Email
- Database

Each of these applications must be available 24x7x365 and be geographically expandable, adding value to remote locations and/or removing untimely, batch-based reporting.

As the technology and architecture behind SANs develop into high-speed interconnections over IP, data inter-linkages between remote sites become important tools for enabling IT departments to provide enhanced storage services. The following sections describe a few examples of how high-speed connectivity over IP & SONET networks can provide a valuable technical solution and expand the ROI of a storage-network design.

### *Outsourced Storage Services*

Managed Service Providers (MSPs) and Storage Service Providers (SSPs) are a breed of outsourced service companies that can profit greatly from SAN-over-IP services. Current deployments of storage services are either in a co-located facility or in a building located near the data center with dedicated fiber to their customers. Wire-speed IP-based storage connectivity allows these companies to provide services, while maximizing their own ROI, by locating their equipment in the most cost-effective location.

SSPs purchase and procure equipment within Internet Data Centers (IDCs) in metropolitan areas to house the storage, servers, and networking equipment that make up their collective service offering. These co-located facilities provide hosting services for e-commerce and enterprise customers, with extensive fiber infrastructures connecting the IDC to the rest of the world.

An essential element of an SSP's differentiation is its Service Level Agreements (SLAs). SLAs provide guarantees of availability, response time, throughput, security, and other key attributes of an effective service. To measure the worth of an SLA, a suite of statistical tools are needed that can determine the service delivery based on measured, statistically tracked, and verified criteria. Without this reporting, the value of a connectivity solution and the SLA that stands behind it are subject to interpretation.

### *The Enterprise*

As consolidation of many industries continues, the need for accessing remote storage as part of the ongoing business continues to grow. These industries require global 24x7x365 service, mainly because outages that result in service interruption are staggering to e-commerce and business applications.

Disaster recovery has also become a real-time network design effort. Backup windows are minimized and the requirement for remote storage of mission-critical data is now more time-sensitive than ever. An organization's ability to access its archived data instantly is crucial to business continuance and, hence, the storage-network design.

### *E-commerce*

If past experience is any indication, the explosion of e-commerce will continue. During the 1999 Christmas season, online shoppers spent over $10.5 billion. It is estimated that almost $20 billion was spent online over the 2000 Christmas season. As electronic commerce providers of business-to-business (B2B) and business-to-consumer (B2C) services grow exponentially and create extremely competitive markets, the demands for uptime become the life-blood of the organization.

Mirrored servers around the globe are a standard method of providing instant access to companies' products. These mirrored sites must be kept in constant synchronization. Reports of downtime at major sites now reflect directly to their stock price, market capitalization, and earnings per share.

High-speed, IP-based data connectivity provides a competitive edge in service levels to customers. The network infrastructures to support these networks are changing dramatically. Therefore, it is necessary to understand these changes when examining all of the elements in the solution.

## Metropolitan- and Wide-Area Network Infrastructure

To understand where networks are going, it is important to understand infrastructure origins. Telecommunications carriers built networks to support telephony services over very long distances. These Time Division Multiplexing (TDM) network architectures, shown Figure 2, were built mainly around a 3-tier Synchronous Optical Network (SONET) ring topology. The local ring that connects to Central Offices (COs) and Internet Points of Presences (POPs) operates at speeds of 622 Mbps (OC-12). The regional ring connects the local rings and operates at 2.488 Gbps (OC-48); this ring then connects to the long-haul networks for transport around the globe.

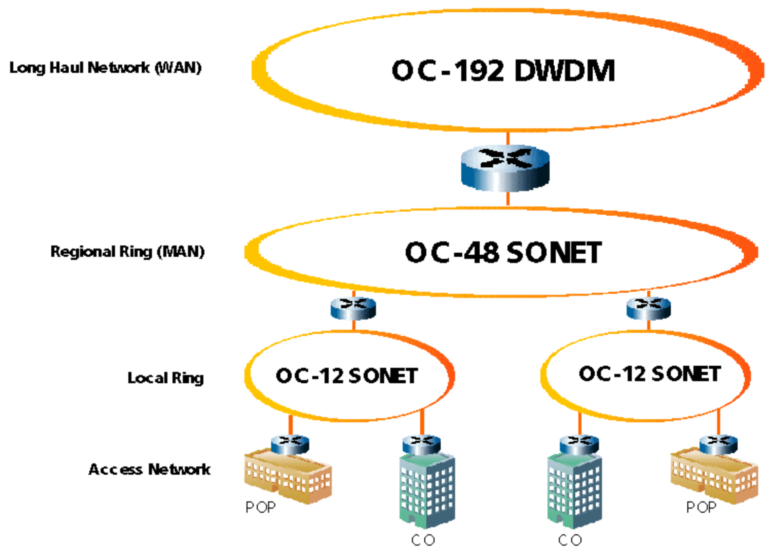


**Figure 2: Typical Telecommunications (SONET) Infrastructure**

With the explosive growth in the Internet, carriers have been upgrading long-haul networks (as shown in Figure 2) with Dense Wave Division Multiplexing (DWDM) technology to support 10Gbps (OC-192). These networks are moving toward 40Gbps (OC-768) to keep pace with the increased traffic demands, with the push for bigger pipes occurring in both the regional and local rings.

As bandwidth demand increases, a new generation of carriers has entered the market, building optical networks that are IP-based and optimized for packet-based traffic. By using optical technology throughout the network, these next-generation IP networks can deliver higher bandwidth into the long-haul (WAN), regional (MAN), and local-access networks. This results in a mesh architecture (see Figure 3) that offers optical networking from the core to the access network. This arrangement, in turn, enables all traffic to be carried in IP packets, and allows the entire network to be easily configured and reconfigured.

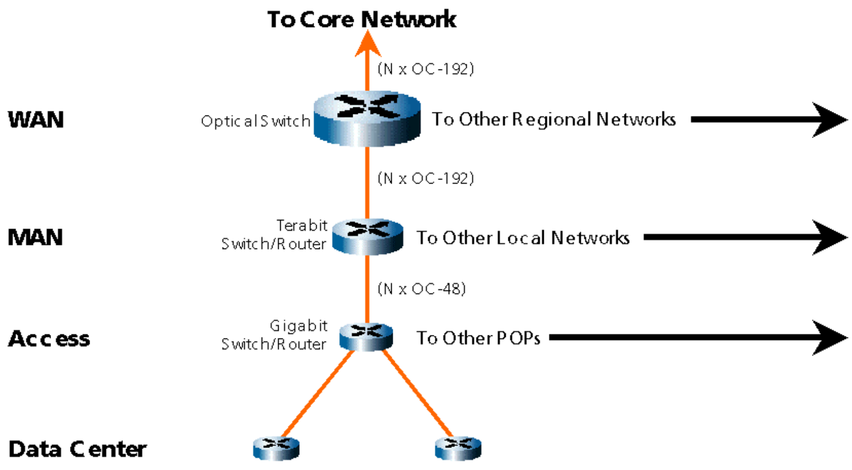


**Figure 3: Optical IP Network Architecture**

Feeding into these infrastructures are the enterprise connections and Internet Data Centers (IDCs), which are either self-contained or housed within co-location facilities. As the demand for bandwidth grows and the amount of data stored at the data center increases, there becomes a need for faster connections (e.g., OC-192) closer to the network edge and data centers.

A number of vendors are introducing10 Gbps Ethernet (10GbE) products to market, even before the IEEE802.3ae standard is officially adopted, in an effort to provide a satisfactory transport facility into the larger infrastructure. As these high-bandwidth connections become more ubiquitous, they will push from the core of the network to the edge. OC-192 interfaces are even surfacing within the data centers themselves, with 10GbE interfaces expected to follow soon. SANs within the data center will also push the need for additional bandwidth, especially as the requirement for inter-SAN connections increases.

## *Transporting SANs over IP Networks*

The growing need for storage data that is permeating the business community, coupled with the available bandwidth afforded by IP- and SONET-based networks, are making SAN extensions an increasingly attractive option. With SAN extensions, end users can connect data centers at opposite ends of a city. The challenge is to do so at full-wire speed, with the same reliability and availability as the storage traffic within each data center.

Issues such as network latency, packet jitter, and guaranteed delivery over MAN and WAN infrastructures become critical issues for the storage-network designer. In SONET networks, these issues are very tightly controlled by the TDM nature of the network. IP networks do not deliver the same Quality of Service (QoS) as SONET. However, by using mechanisms such as Diffserv, VLAN Tagging, and UDP port prioritization, traffic can be prioritized for transport across the IP network to deliver a QoS approaching that of SONET networks.

Interconnecting these remote storage islands across the IP infrastructure requires a gateway to transport Fibre Channel data across an IP network. This IP-SAN gateway technology is responsible for enabling the efficient transfer of Fibre Channel traffic over the IP network (as shown in Figure 4), managing flow control and jitter, and providing additional robustness for end-to-end packet latency. Encapsulating Fibre Channel traffic into the Gigabit Ethernet frames and shaping the traffic to ensure consistent, ordered packet delivery ensure that the network remains fully utilized even over Metropolitan and Wide Area Networks.

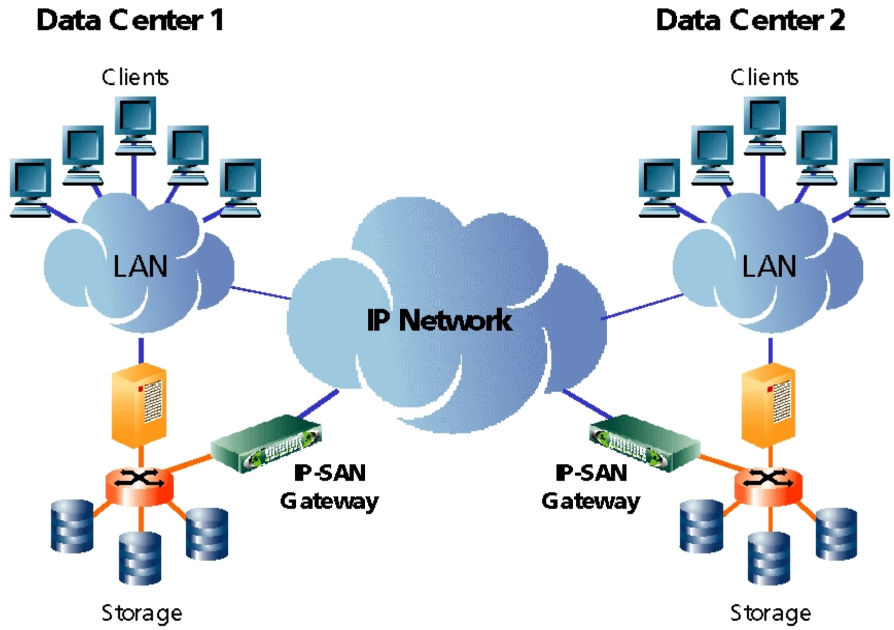


**Figure 4: FC SANs Extended over an IP Network**

SAN-to-SAN traffic now has the performance, reliability, and scalability of inter-SAN traffic, extending the benefits of a SAN across the data center boundaries. The following sections explore different deployment scenarios and the inherent benefits of IP-SAN technology to each.

## *IP-SANs and SONET*

Transporting SANs over SONET offers inherent benefits to an existing storage application. By grooming the Fibre Channel traffic, then encapsulating it into Gigabit Ethernet ports, the traffic can be aggregated by an edge router to a SONET network. In Figure 5, two ports from a Fibre Channel SAN switch are brought into the IP-SAN gateway, which then transports the encapsulated storage traffic into a switch or router. This configuration aggregates the two Gigabit Ethernet channels to the OC-48 Add Drop Multiplexer (ADM), then transports the data to the remote data center on the other side of the SONET network.

The IP-SAN gateway performs traffic shaping, prioritization, and flow control at each end point, ensuring that the data sent into the SONET network is tightly controlled. The data center can utilize the existing SONET infrastructure to transport storage traffic, leveraging the investment made on its SONET and IP equipment.

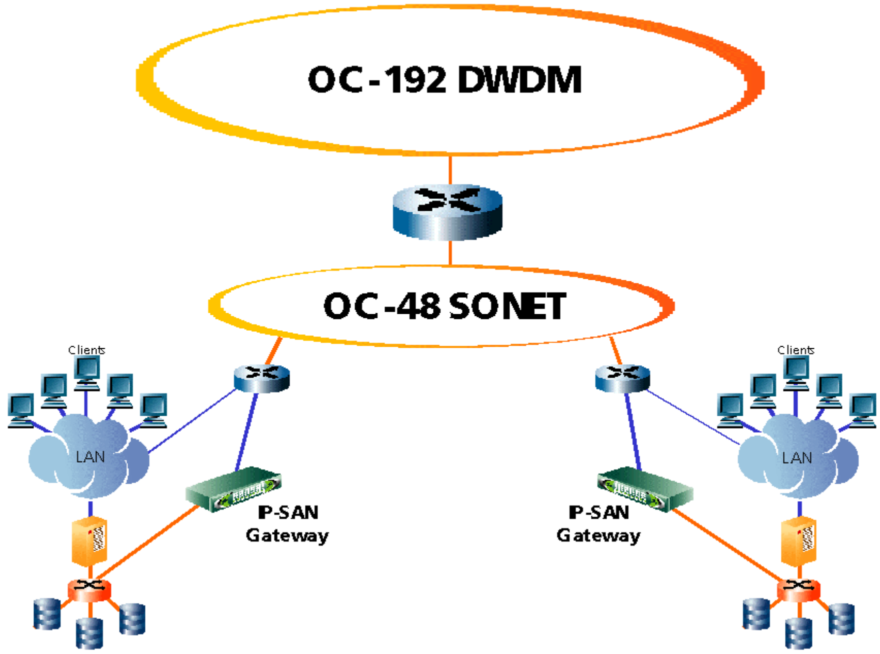


**Figure 5: IP-SANs Extended over a SONET Infrastructure**

Redundancy can be achieved by setting up dual paths from the Fibre Channel fabric to individual ports on the IP-SAN gateway. This arrangement enables the traffic to be switched within the Fibre Channel fabric to alternate routes across the network.

Aggregation to higher-speed connections is also possible, especially considering the momentum behind OC-192 ports moving from the core to the network edge. With a 10Gbps OC-192 port, up to 10 individual Fibre Channel ports can be aggregated to an OC-192 link at the edge of the data center and transported optically across the SONET infrastructure at 10Gbps.

## IP-SANs and Optical Networking

The mesh architecture of an all-optical network offers ease of provisioning and dynamic wavelength mapping to services. This is critical, especially considering the speed at which these services need to be provisioned and the number of wavelengths that will be driven over the network.

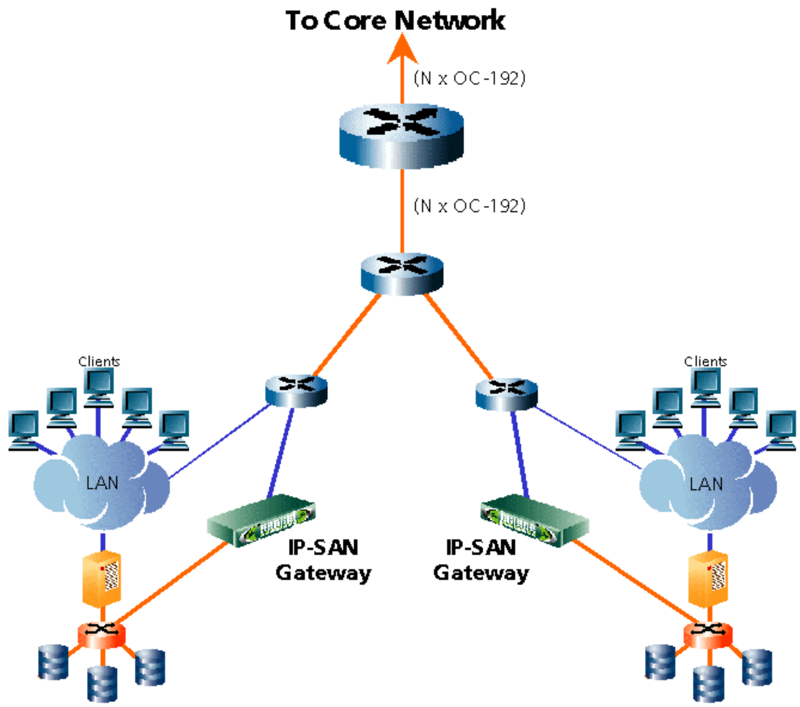


**Figure 6: SANs over an Optical IP Infrastructure**

Existing optical networking standards permit uninterrupted wire-speed transmission distances in excess of 400km (248 mi) using DWDM technology. Recently announced products can increase this distance to 2,000km (1240 mi), thereby eliminating the cost and latency associated with repeaters and amplifiers. This new technology is critical in providing seamless access across distant data centers.

Current buffer-to-buffer credit schemes in Fibre Channel products generally do not take into account network latencies. Because Fibre Channel was originally designed as a data-center technology, the current inability of Fibre Channel products to "fill the pipe" over a MAN or WAN exacerbates any network equipment-based latency. These shortfalls can be accommodated by intelligent implementation of buffer management in the IP-SAN gateway to ensure that the links across the optical network remain fully utilized, as shown in Figure 6.

Management of the optical network becomes a major factor when considering total cost of ownership. Typical optical networks offer physical layer management capability, with no visibility to the protocols operating on the optical network. The IP-SAN gateway provides the inherent management capability to provide performance and statistical information about the status of the link and can even be mapped to Service Level agreements. In this way, a service provider can manage the entire network effectively from a central Network Operations Center (NOC), determine root cause of a possible network-performance issue, and take corrective action ***before*** the issue interrupts service.

## Conclusion

A careful analysis of network requirements and service availability should be performed before linking distant data storage sites. Return on Investment, including an analysis of Total Cost of Ownership, is one of a number of criteria to consider when evaluating appropriate network configurations and deployment models.

Under most circumstances, solutions that leverage the current infrastructure are more attractive to network designers because cost effective management of new network interconnects are generally a driving factor for a strong ROI. IP-SAN connectivity provides a seamless methodology for connecting remote storage without disrupting an organization's infrastructure.

While standards-based approaches are attractive solutions for remote-storage connectivity, it is important to consider the costs associated with waiting for standards to be developed and losing a competitive advantage. IP-SAN gateway products with hardware-based designs offer excellent wire-speed solutions. Software-based network processor solutions, by contrast, cannot offer gigabit throughputs and fall short of meeting the network requirements for the Enterprise or Service Provider data center.

Managing Fibre Channel's buffer-to-buffer credit schemes is the most important part of ensuring reliable flow control and data delivery. The network designer must be cognizant of the distances involved and the latencies imposed as part of the existing infrastructure. Without a statistical analysis of the connection(s) as an integral part of a network's architecture, management tools are ineffective and SLAs are highly difficult to enforce. High availability and data-transmission reliability are cornerstones of a SAN extension; consequently, redundancy of key components and data paths are required to support fault-tolerant configurations.

## SAN Valley Systems SL1000 IP-SAN Gateway

The SL1000 IP-SAN Gateway extends SANs over IP networks with wire speed throughput, enabling cost-effective deployment of remote-storage applications. The SL1000 connects Fibre Channel SAN switches and storage resources across networks with wire-speed connectivity over MANs and WANs. The SL1000 provides the high-speed, high-availability connectivity essential for centralizing storage, and supporting backup and storage-mirroring applications in real time. The SL1000 connects storage islands over IP networks with over 8 Gbps of total chassis throughput, providing the fastest most scalable solution available today. With independent, dedicated, hardware-based Fibre Channel (FC)-to-Gigabit Ethernet (GbE) encapsulation engines, the SL1000 provides superior channel performance, scalability, and throughput without microprocessor interrupts. Support for standards-based QoS ensures that storage traffic is always prioritized for delivery across the IP network. The SL1000 also maps Fibre Channel buffer-to-buffer flow control to the Gigabit Ethernet channel, maximizing data throughput and efficiency.