# *Unleashing the SAN in Your IP Network™*

***January 2004***

## Introduction

Computer storage holds the lifeblood of today's economy. From movies, music, books, and classroom materials to corporate, personal, and financial data of all sorts — everything is being stored digitally. For this reason, the information and the knowledge derived from today's Internet-reliant world have become the core elements by which our society increases its productivity.

Digital data is growing at unprecedented rates. According to International Data Corporation (IDC), the demand for enterprise storage capacity is growing 75% every year. IDC projects that total storage terabytes will grow from 1.9 million in 2003 to 3.3 million in 2004.

This deluge of digital data is driving the need for reliable, high-performance storage that is easier to implement and manage. These requirements favor network-based, centrally managed storage over alternative storage solutions and infrastructures. Internet Protocol (IP) has become firmly established as the predominant general-purpose networking protocol. The IP storage concept naturally followed as a way to meet growing networked storage requirements.

StoneFly Networks' complete line of Storage Concentrators™ provide IP Storage Area Network (SAN) infrastructures that are highly scalable, intelligent, reliable, manageable, and cost effective. Storage Concentrators are IP-based storage provisioning appliances that enable mid-sized to large enterprises to manage and optimize their storage assets in real-time. As a result, they provide a rapid return on investment by delivering more efficient storage usage, improved administrative productivity, and reduced network complexities.

IP storage protocols are enabling SANs to be built and interconnected using prevalent IP-enabled technologies such as Ethernet. From reading this paper, you will learn how to drive down your Total Cost of Ownership (TCO) by unleashing the SAN in your IP network.

## The Roots of Storage Networking

Storage networking was developed to provide a high-speed network for storing and retrieving data. With storage networking, a dedicated high-speed network allows files and data to transfer between storage devices and client machines directly, bypassing the traditional server bottlenecks and network control.

In this way, increased flexibility and performance are achieved by separating control from data. Storage management processes can operate in the background and in a continuous mode. By not conflicting with network traffic, network performance is maintained and even enhanced. The following section traces the roots of storage networking, starting from its Small Computer Systems Interface (SCSI) origins and progressing to the IP-based storage solutions of today.

### *SCSI*

SCSI is an intelligent, parallel Input/Output (I/O) bus on which various peripheral devices and controllers can exchange information. For over two decades, SCSI has been the dominant protocol for transporting block-level data among servers. In fact, the SCSI standard was one of the key factors in the development of open systems, providing a low-cost, high-performance storage interconnect.

Because of its longevity in the marketplace, the parallel SCSI interface enjoys a depth and breadth of products that exceeds that of any other I/O interface. It comes as no surprise, then, that a survey by Gartner/Dataquest found over 100 million SCSI-based disks, tapes, and other storage devices being directly connected to individual systems.

However, in a world inundated with multi-terabytes of data, it is clear that SCSI's inherent limitations cannot keep pace with the demands of distributed networked system environments. These limitations include:

- **Performance:** Most SCSI implementations today provide a single path with throughput from 40 to 80 MB/s.
- **Configuration:** SCSI ties storage directly to a single server. Even newer SCSI implementations permit only 15 devices to be connected to each bus.
- **Distance:** Common SCSI configurations limit the distance between storage elements and servers to 6 meters. Newer SCSI technologies address this limitation, but still restrict users to 25 meters.
- **Availability:** Because SCSI environments tie storage to a single server, all access to associated storage is lost if that server fails. Even the newest SCSI implementations support a maximum of two initiators, limiting cluster configurations and shared-storage resources.
- **Bus:** Some SCSI configurations support multiple buses. However, space limitations make it impractical to maintain all of the storage inside one box. Therefore, administrators in a SCSI environment may need to add servers and incur the associated expenses of maintenance and management, to increase storage capacity.
- **Costs:** In addition to the expenses associated with managing the servers added simply to increase storage capacity, the close coupling of server and storage prevents effective sharing of tape and disk resources.

Beyond these limitations, the SCSI paradigm of directly attaching a server to a storage device creates islands of computing that must be configured and managed separately. Efforts to overcome these limitations with Local Area Networks (LANs) have added extra layers of cost and complexity. This leads to duplication of hardware, software, and data-management resources. It also adds to the total cost of computing, while restricting a company's ability to meet the growing data demands of its business.

## Fibre Channel Storage Area Networks

To address the limitations inherent in SCSI, storage buyers adopted SANs as a solution. The Storage Networking Industry Association (SNIA) defines a SAN as "a network whose primary purpose is the transfer of data between computer systems and storage elements and among storage elements."

Currently, Fibre Channel (FC) is the prevalent technology used to transport data between computer systems and storage devices. This is primarily due to FC's high performance, connectivity, and availability to support block-oriented storage protocols.

A FC SAN overcomes the limitations found in SCSI by creating a switched infrastructure for storage. This infrastructure achieves its high performance by relegating most of the protocol processing to hardware.

Key characteristics of a FC SAN include:

- Support for as many as 16 million SCSI devices in an extended SAN.
- Support for operating distances up to 10 km.
- Support for 2 Gbps speed

## Network Attached Storage (NAS)

To address the need for shared file-level data access, the user community adopted the concept of Network Attached Storage (NAS). NAS allows scalable storage devices equipped with a built-in network interface to connect directly to an existing LAN. By providing a direct connection to the network, NAS facilitate easy installation and maintenance. In addition, as the number of users grows and free space becomes low, additional NAS devices can attach to the LAN as necessary.

NAS devices support local file systems such as the Microsoft Windows' NT File System (NTFS) and Unix File System (UFS). This support enables multiple NAS devices to access local file systems remotely using file-interchange protocols such as the Common Internet File System (CIFS) for Microsoft Windows and the Network File System (NFS) for Unix. The file system accesses data residing on disks that are either internal to or external from the NAS device using SCSI or FC interconnections.

## iSCSI

In a world in which IP dominates LANs, Metropolitan-Area Networks (MANs), and Wide-Area Networks (WANs), and data storage requirements increase dramatically, it seems inevitable that these two forces would converge. That convergence has been achieved with the iSCSI protocol.

iSCSI combines two widely used protocols from the storage and networking arenas:

- From the storage arena, iSCSI uses the SCSI command set.
- From the networking arena, iSCSI uses IP and Ethernet, protocols that are the cornerstone of most corporate networks.

Together, this technological synergy enables block-level storage data to be carried over TCP/IP networks.

iSCSI encapsulates SCSI packets in TCP for reliable transport and routing using IP. As a result, standard SCSI commands can pass between host systems (or "initiators") and storage devices ("targets") over a standard Ethernet/IP TCP/IP network infrastructure, rather than via SCSI cabling or FC connections.

## Technology Comparison

The following table provides a side-by-side comparison of SCSI, Fibre Channel, and iSCSI technologies.

| Specification | SCSI | Fibre Channel | iSCSI |
|---|---|---|---|
| **Installation:** | Mature, proven technology that IT professionals are comfortable installing, using, and maintaining. | Fairly new technology. Most IT professionals have limited knowledge about FC, which can make FC SAN installations difficult. | Because an iSCSI SAN leverages the existing SCSI and Ethernet technologies, installing an iSCSI SAN ranges from moderate to easy. |
| **Maintenance and Ease of Use:** | SCSI requires a certain amount of expertise for configuring RAID drives and troubleshooting. | After installation, using and maintaining the FC SAN require a moderate amount of expertise. | Easy, since iSCSI uses the mature and proven Ethernet technology to transport data. |

## The Shortcomings of First-Generation SANs

To date, SAN deployments have been driven by an insatiable demand for storage and the user benefits delivered by networked storage. While FC SAN and NAS solutions have been designed to address these issues — as well as SCSI's inherent limitations — they provide shortcomings of their own. These shortcomings are providing the impetus for IP-based SANs.

### *FC SAN Limitations*

The following limitations are associated with FC SANs:

- **Expensive, Difficult, and Time Consuming to Deploy**
  Adding a FC SAN requires a second network based on a new, unfamiliar technology. Deploying a second network using a new, different type of technology requires a significant overhaul for existing computing environments. It is also time consuming and adds significant cost for separate network management systems, training and education for IT personnel, and separate

sourcing of new networking equipment.

- **Interoperability Issues & High Implementation Costs**
  While standards exist for Fibre Channel technology, the standards are interpreted differently by each vendor, precluding interoperability among vendor products.

- **Lack of Management Standards**
  This void has prompted the proliferation of vendor-specific management tools. Some vendors have invested millions of dollars in interoperability testing centers to overcome these deficiencies by certifying their products with very specific combinations of hardware and software. Others have formed alliances to ensure interoperability between their SAN solutions.

- **No Help from Emerging Standards**
  Emerging SAN IP standards, such as the Internet Fibre Channel Protocol (iFCP) and Fibre Channel over IP (FCIP), address only FC SAN-to-FC SAN interconnections over an IP network. These standards ignore the multitude of deployed SCSI devices and the block-level data residing on them.

### *NAS Limitations*

The following limitations are associated with NAS:

- **Inability to Access Block-Level Data**
  NAS does not support direct access to block-level data residing on SCSI-attached devices or FC SANs. Rather, such access is only available by routing the connection through a file system. Some newer NAS devices may allow block access by bypassing the file system, but these devices are still optimized for file I/O and not block I/0.

- **Too Slow to Support Raw I/O**
  An application running on a client that accesses storage on a NAS device must go down the entire seven layers of the networking protocol, across the LAN wire, up the seven layers, fetch the data, and transmit back the data. The steps required to support raw I/O prevent NAS from being an effective solution with data-intensive server applications, such as direct connections to database and Exchange server applications. In fact, some high-performance database systems actually perform direct, raw I/O using SCSI commands to avoid the inefficiency of the server's operating system (OS) file system.

## Factors Driving IP Storage

The shortcomings of FC SAN and NAS solutions, coupled with the growth of SANs and the pervasiveness of IP, provide IP-based storage solutions with a decisive advantage over other storage solutions.

IDC predicts that IP SANs will show explosive growth over the next few years. IDC forecasts that the market for iSCSI-based disk arrays will grow from $216 million in 2003 to $4.9 billion in 2007, a four-year CAGR (compound annual growth rate) in excess of 118%. iSCSI-based switches used in IP SAN implementations will grow from some 18,000 total ports in 2003 to 6.94 million ports in 2007, a four-year CAGR of more than 342%.

The concept of using a single networking technology for both LAN and SAN is compelling. Organizations no longer need to maintain equipment, technical staff, and expertise for both IP and FC technologies. In addition, many smaller companies can take advantage of IP-based storage because IP is a familiar technology that is less expensive to adopt, deploy, and maintain than FC.

The following list identifies additional key advantages of IP-based storage:

- Enables organizations to “universalize” storage networking by getting both message/file and storage I/O onto an Ethernet/IP network and ultimately “converging” networking and storage architectures.

- Provides a natural fit within the prevalent IP/Ethernet infrastructures deployed in today’s business environments.

- Enables block storage over widely deployed IP-based networks, allowing easy access to storage over long distances.

- Provides a familiar networking technology and management, with virtually no additional training and support required.

- Empowers organizations to transition from 1 Gigabit Ethernet to 10 Gbps and beyond, thereby protecting investments with standards-based performance upgrades.

- Reduces TCO by delivering Ethernet economics to storage networking.

### Storage Provisioning

Combining the functionality of an iSCSI router and bridge with the power of a storage provisioning engine, Storage Concentrators are installed as a centralized appliance at the core of the IP network for logical volume management of cost effective IP SANs. While Fibre Channel and server based volume management can be costly and complex, StoneFly's Storage Concentrators allow simple, intelligent, and affordable SANs by utilizing the existing IP network infrastructure. Storage Concentrators require no specialized expertise to install and are fully interoperable with leading initiators and storage arrays.

An IP SAN utilizing a Storage Concentrator can use storage provisioning to create a logical volume from one or more physical volumes, and then improve performance or data reliability by aggregating multiple disks. Regardless of the logical volume's physical composition, the host computer sees only an available data container with a defined capacity; the Storage Concentrator handles all other technicalities.

## The iSCSI Framework

The iSCSI framework connects “initiator” and “target” devices over IP.

- iSCSI initiators are devices such as file servers that exchange block-level data with iSCSI target devices and encapsulate SCSI commands into TCP/IP for routing over an IP network.

- iSCSI targets are devices that receive iSCSI commands and exchange data across the IP network and/or through internetworking equipment (such as storage routers and switches). An example of an iSCSI target is a Storage Concentrator. Others include disk arrays, RAID devices, and tape libraries.

Figure 1 shows an example of an IP SAN based on iSCSI. iSCSI establishes communication sessions between initiators on the hosts and target storage devices, and provides ways for them to authenticate one another. Thereafter, iSCSI uses standard SCSI commands to facilitate the exchange of data between initiators and target devices. In addition, TCP/IP provides recovery in case the connection fails.

Because iSCSI is a native IP-based protocol, data can be transported over the IP infrastructure rather than through FC or SCSI cabling. In this way IP SANs enable data transfers using readily available infrastructures in Local-, Metropolitan-, and Wide-Area Network topologies. The Storage Concentrator from StoneFly Networks is installed at the core of the network as a FailOver Cluster, acting simultaneously as an iSCSI router, bridge and volume manager for optimization of the attached RAID storage.

Other configuration examples are described below under "Implementing an IP SAN".

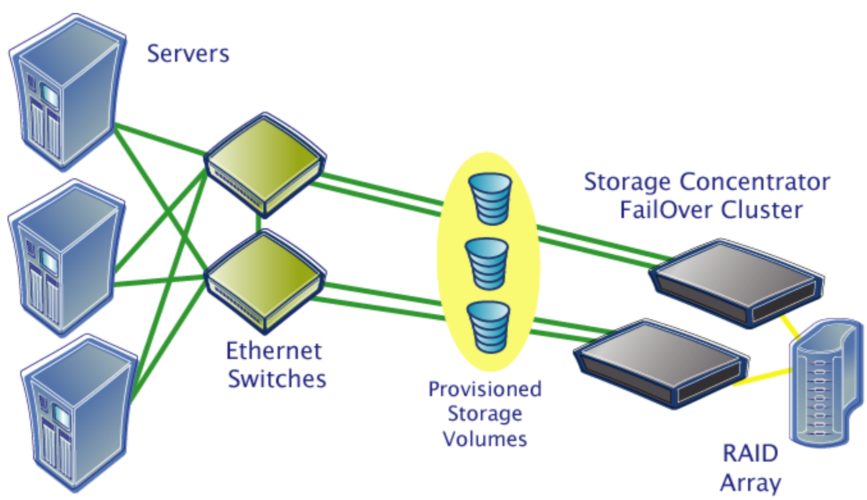


**Figure 1. Example of an IP SAN**

## How iSCSI Works

The following steps describe how iSCSI works.

1. An application or user issues a request for data, a file, or an application.
2. The operating system generates the SCSI commands and data request. The command and request are encapsulated and a packet header is added. The packets are transmitted over an Ethernet connection.
3. At the receiving end, the packet is disassembled or decrypted if it was originally encrypted, separating the SCSI commands and data.

4. The Storage Concentrator acts as the iSCSI target and receives the iSCSI commands. It interprets the iSCSI headers and instructions and uses its storage provisioning engine to locate the data requested by the server.

5. The Storage Concentrator then sends the SCSI commands and data to the storage device. Similarly, data is returned in data packets to the Storage Concentrator, which serves them up to the initiating server in response to the request using the iSCSI protocol.

## Implementing an IP SAN

The previous section described how IP SANs connect over a WAN using standard Ethernet equipment. IP SANs can also connect to FC SANs using an IP storage router or switch that handles the conversion between the FC protocol and iSCSI. This arrangement allows organizations to extend the reach of an IP SAN by bridging IP SANs to FC SANs. Figure 2 shows an example where a Storage Concentrator can share storage on a Fibre Channel SAN with IP-connected computers. The Storage Concentrator FailOver Cluster serves up volumes from the storage on the Fibre Channel SAN to the servers on the IP network. In some cases, configurations such as this can be less costly by using the Ethernet infrastructure for departmental servers but still taking advantage of centralized data stores.

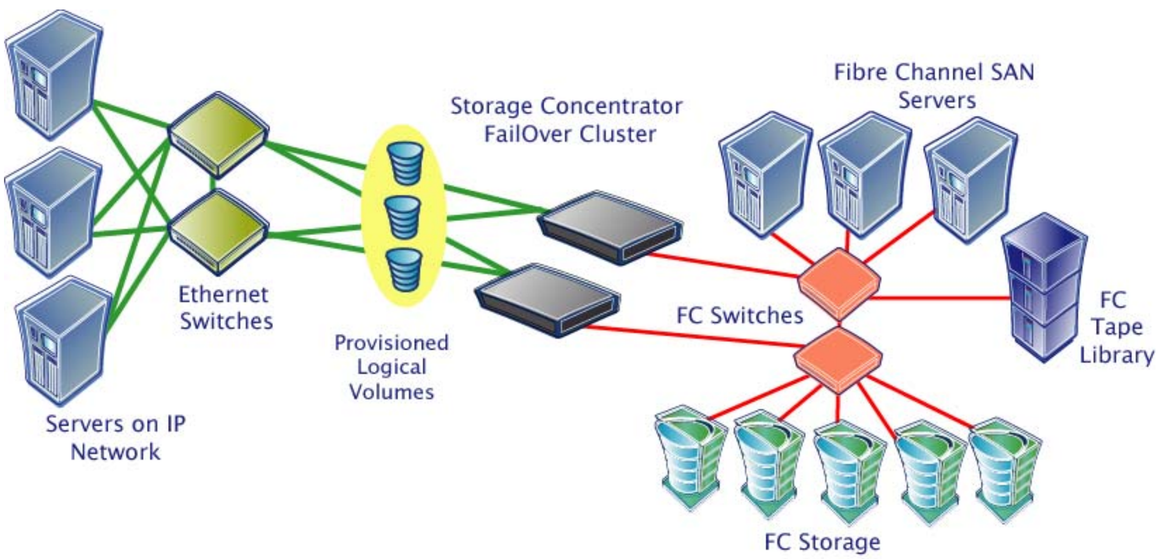


**Figure 2. Providing FC SAN Connectivity to IP-Connected Computers**

IP networks can also be deployed directly in SANs. This implementation allows organizations to maximize IP and Ethernet networking capabilities, while making use of FC and SCSI interfaces on servers or storage devices. Figure 3 shows a configuration where all inter-switch links are built with standard IP and Gigabit Ethernet connected to the enterprise Fibre Channel SAN. This configuration mirrors the networking implementation model deployed in the enterprise backbone. By using this common infrastructure, and Replicator software from StoneFly Networks, deployment time and costs are significantly reduced.

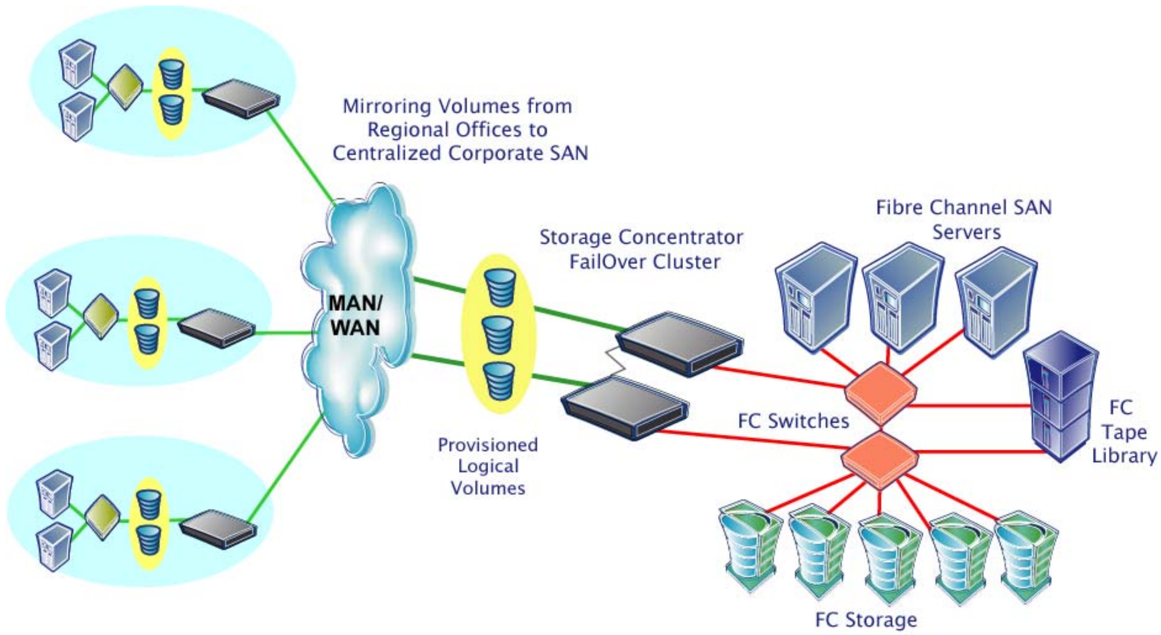


**Figure 3. Remote IP SANs Connected with Standard IP to Corporate Fibre Channel Network**

## IP Storage Applications

The fundamental customer requirements that gave rise to SAN technology have not changed. What has changed is the ability to fulfill those requirements efficiently and cost-effectively by leveraging mainstream storage and IP networking technologies. The following sections describe applications that can benefit from IP storage.

### *Server Clustering*

The 7x24 operating characteristics faced by E-commerce and Internet-based businesses means that loss of data access is lost business revenue. For these demanding environments, high availability of data requires server clustering. With server clustering, an individual server failure enables other servers to assume the tasks of the failed device and provide continuous access to its data.

IP SAN technology is an enabler for high-availability server clustering. In addition, IP SANs are inherently scalable, enabling large server clusters to be deployed over time, using Gigabit Ethernet backbones and link aggregation to maximize cluster performance.

### *Storage Consolidation*

Many enterprises suffer from spontaneous, unplanned storage growth. This dilemma has forced administrators into:

- Managing dispersed storage resources.

- Tracking multiple versions of applications, data sets, operating systems, and service packs.

- Accommodating tape backup for hundreds, even thousands, of distributed storage devices.

More storage requires more administrators, which translates into higher cost-per-megabyte of managed data. Replacing dispersed storage with larger centralized storage arrays can reduce maintenance time and costs, but requires shared access by servers.

Since the rationale driving storage consolidation is the reduction in TCO, consolidation goes hand in hand with IP SANs.

- An IP SAN provides shared storage access and enables storage consolidation over conventional Gigabit Ethernet and IP networks, reducing both storage and network administrative costs.

- An IP SAN further reduces administrative overhead by utilizing IP network personnel to deploy and manage the SAN infrastructure.

- An IP SAN is also highly scalable and can easily accommodate growth in server population and the installation of additional large arrays.

In addition, Gigabit Ethernet and IP services such as subnetting, Virtual LANs (VLANs), traffic prioritization, and link aggregation can be leveraged to ensure security and QoS.

### *Local Tape Backup*

Tape backup is a universal problem, with data storage increasing from gigabytes to terabytes of information. For LAN-based backup processes, the avalanche of tape backup data streams can swamp LAN resources.

An IP SAN overcomes LAN-based backup issues by creating a separate gigabit network that accommodates storage traffic, including tape backup. This removes potential bottlenecks and provides a robust, high-performance infrastructure that can support multiple tape backup streams concurrently.

### *Scaling a Data Center IP SAN*

One of the shortcomings of first generation SANs was their inability to build enterprise-class configurations to support hundreds or thousands of servers and storage devices. IP networks, by contrast, are designed to support thousands (or, for the Internet, millions) of devices and assume a much higher level of interoperability.

By using 10 Gigabit Ethernet interfaces, large IP SANs can be constructed so that bandwidth bottlenecks do not occur as the SAN population increases. With proper bandwidth allocation, a wide variety of SAN applications can be supported concurrently, further maximizing resource utilization and reducing overall costs.

For organizations with more demanding requirements, director-class Gigabit Ethernet switches and 10 gigabit or trunked interswitch links can be used to deliver true enterprise-class SANs.

### *Remote Backup/Vaulting*

Like local backup, remote tape backup presents significant challenges for organizations. Remote backup may be required for disaster recovery or streamlining backup processes over multiple sites.

FC was originally designed for high-performance channel applications within the circumference of a data center (500 meters using multi-mode fiber cabling and 10 kilometers using single-mode cabling). This design did not accommodate distances extending hundreds or thousands of kilometers, which might be typical of disaster-recovery applications. Therefore, remote connectivity for FC SANs is inherently problematic, since it pushes the architecture beyond its original purpose and scope.

By contrast, IP SANs have no inherent distance limitations. As a result, they can accommodate wire-speed area connectivity for storage traffic at distances exceeding thousands of kilometers. Rather than restricting remote backup distances or manually transporting tape cartridges off-site, companies can now protect their data over significant distances, well beyond the perimeter of a potential disaster area.

### *Remote Data Replication*

Remote data replication protects mission-critical data from unplanned events that can interrupt business operations. Remote data replication is a requirement for business continuance and avoids the arduous process of restoring data from tape should a disaster occur.

Remote data replication can be deployed via disk mirroring, from one site to another, or mutual mirroring between two sites. A primary data center, for example, can create a remote mirror to a secondary data center. In the event of failure at the primary data center, user requests are directed to the secondary site. Alternately, two active data centers can mirror each other to ensure data integrity.

Remote data replication has been a problem for FC SANs. While FC extensions permit data replication over short distances, the distances are insufficient for most disaster-recovery strategies.

An IP SAN solution, on the other hand, provides an effective business-continuance solution by enabling data replication from the primary data center to a remote location hundreds or thousands of miles away. This scenario safeguards the replicated data against any disaster at the primary site. Easy to use mirroring software, such as StoneFly Replicator, provides local and remote, synchronous and asynchronous mirroring as shown in Figure 4.

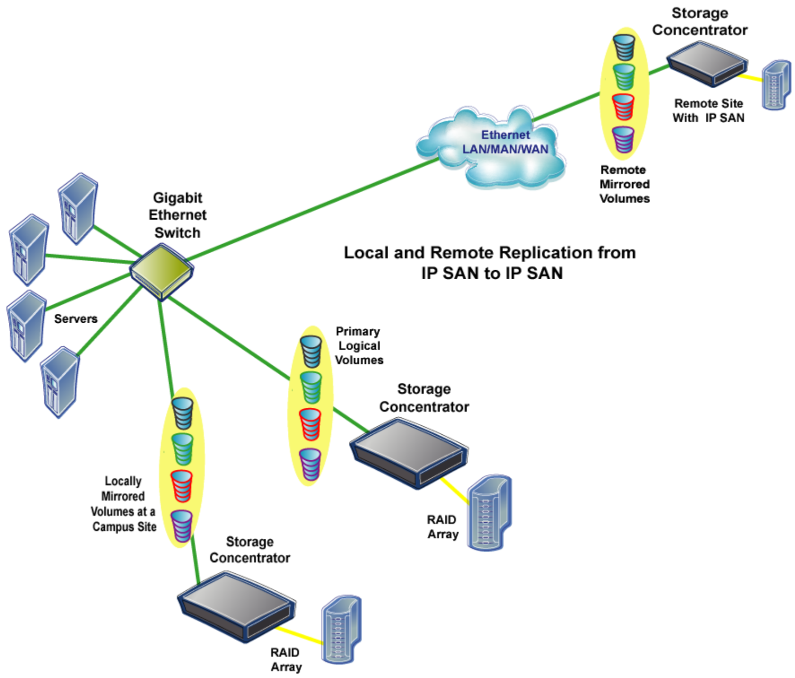


***Figure 4: Local and Remote Data Replication***

***Business Continuity***

Disaster-recovery strategies are designed to provide business continuance in the event of a catastrophe. Simple FC extension or tunneling (FCIP) cannot provide a robust, stable infrastructure for disaster recovery.

An IP SAN, however, offers high-speed data-transfer performance over thousands of miles, allowing companies to design multi-site, highly stable disaster-recovery configurations to facilitate business continuance regardless of geographical location. As a result, data centers can use gigabit and multi-gigabit links spanning enormous distances to safeguard storage data in secure locations.

## IP Storage: Implementation and Vision

By combining best-of-breed approaches from the storage and networking industries, Storage Concentrators offer significant advantages, including:

- Unified management and proven, readily available, high-speed networking.
- Familiar and more readily supported IP network equipment.
- Enhanced Gigabit Ethernet and security options.
- Ability to leverage existing infrastructure technology, then build upon it to develop an integrated, native IP storage network that works in local-, metropolitan-, and wide-area environments.
- Unique scalability from small to large enterprise-class SAN installations.
- The ability to connect FC SANs to standard IP and Gigabit Ethernet SANs.

Integrating SANs with mainstream IP networking places manageable shared storage solutions within the reach of every enterprise, large or small. By removing barriers to distance, device support, and protocols, the convergence of networking and storage provides a firm foundation for constructing IP storage networks today that will accommodate the evolving storage needs of tomorrow.

## Six Frequently Asked Questions About IP Storage

The following are six frequently asked questions about IP storage.

***Q*** *Which technology is best to use to implement a SAN - IP storage or FC?*

***A*** Companies with larger budgets and a willingness to learn FC technology may enjoy the benefits of Fibre Channel. However, companies more concerned with TCO and leveraging their network infrastructure should adopt IP storage.

***Q*** *What is the difference between IP storage and iSCSI?*

***A*** IP storage is a technology that encompasses block-storage transfer over an IP network. iSCSI is a protocol standard for encapsulating SCSI in TCP/IP for transfer over an IP network.

***Q*** *Why is iSCSI so important?*

***A*** iSCSI represent a change in direction for storage attachment to servers and SANs that changes the assumptions associated with deployment of storage by allowing the deployment of storage-area networking over an IP network.

The installation of a SAN forces companies to consider not only infrastructure but also technical abilities of the staff required to support a new storage model. Most storage installed today is directly attached to the servers that use the resource. iSCSI, by comparison, allows a company's existing Ethernet network to transfer SCSI commands, and data, with total location independence.

Installing a shared storage system also requires IT staff to learn new concepts such as fabrics, Logical Unit (LU) masking and mapping, and resource allocation. iSCSI, however, enables storage networking using familiar technology and commodity components. customers.

***Q*** *Why is iSCSI getting so much attention?*

***A*** One reason why iSCSI is garnering so much attention is because it represents a change in direction, deployment, and assumptions involving storage. Another reason can be attributed to the speed improvements that have been made, and that are planned, for network signaling. The increases from 10/100 Mbps to Gigabit Ethernet, and eventually 10 Gigabit Ethernet and beyond, make the Ethernet physical connection as fast as FC.

**Q** *Will IP storage mean the death of FC?*

**A** Absolutely Not. Most industry analysts believe that FC will coexist with iSCSI. Fibre Channel is very efficient for large-scale deployments and are operating comfortably in many Fortune 500 companies. iSCSI is ideal for mid-tier or departmental workgroups in large companies and for smaller companies.

**Q** *Does sharing storage traffic and data traffic within a common IP network lead to congestion and bottlenecks?*

**A** While combining messaging and storage traffic on a single network is possible, a more pragmatic implementation is to segment the IP network infrastructure (by using VLANs or subnetting, for instance) and move storage and data traffic using different paths. This approach protects a company's investment in IP networking while maximizing the efficiencies of both types of traffic over a common infrastructure.