


McDATA's Multi-Protocol SAN Technology*



EXECUTIVE SUMMARY

The maturity and mission-critical deployment of Fibre Channel (FC) in storage area networks (SANs) creates a unique class of multi-terabit networks with demanding throughput, latency, scalability, robustness, and availability requirements. This paper reviews the state of and critical system-level requirements for SANs. It describes how Internet SCSI (iSCSI), FC over IP (FCIP), and Internet FC Protocol (iFCP) integrate with FC SANs and discusses associated benefits and challenges. Finally, the paper examines case studies in performance and protocol tuning in high-speed, long-delay networks, which are increasingly critical for FC-to-IP integration opportunities and challenges.



*Originally published in the Proceedings of MSST2004, NASA/IEEE Conference on Mass Storage Systems and Technologies as Fibre Channel and IP SAN Integration.

INTRODUCTION

Information technology (IT) is a key driver and challenge for businesses, government, and research/development centers. Data centers are a critical asset and provide the infrastructure that houses information processing, storage, and communication resources. Corporations are under tremendous pressure to manage return on investment, massive growth in information processing and storage needs at a global scale, management, performance, availability, and scalability requirements, and the IT infrastructure. To add to the challenges, there are many new technology and deployment decisions that have significant implications in terms of value and impact to the data center.

SANs are a critical part of the data center, and are based on high speed, high bandwidth, low latency, and low error rate interconnects for scaling application, database, file, and storage services. FC is the key technology and standard that drive rapid growth of SAN deployment. The development of global and distributed file systems, content-addressable storage, object-oriented storage, cluster and blade servers, and utility computing is driving more integrated IP and FC network usage. The evolution of the data center and new information and computing trends drives the data center toward a more dynamic resource and performance provisioning and management model, which demands more efficient and scalable computing, information storage, and networking. In addition, business and operational requirements in the data center drive the scaling and evolution of larger SANs encompassing metropolitan and wide-area distances, high security and availability, and multi-protocol networks. In the face of these trends, ease of use, configuration, and management of the SAN is even more important.

This paper reviews important requirements and deployment examples. It describes emerging IP SAN technologies and how these technologies interface and integrate with FC. It also examines several protocol and design considerations, system-level behaviors, and areas that need further research and enhancement. This paper leverages the efforts of many engineers, architects, and researchers from the industry. The paper uses their findings and recommendations, and tries to relate them to SAN applications.

THE FC SAN TODAY

FC SAN Overview

FC technology^[1] and product deployment has evolved from 1 gigabit per second (Gbps) to 2 Gbps links, and there is development to introduce 4 Gbps and 10 Gbps links. An FC network or fabric is a multi-terabit, low-latency switching network, mainly used to interconnect servers to storage. Although a FC fabric is designed to support any-to-any connectivity, the actual use tends to be some-to-some. Each server talks to a few storage devices or each storage device talks to a few servers, with occasional traffic for backup or other purposes involving devices shared by many sets of storage and servers. Deployment of mid-range to high-end FC fabrics is

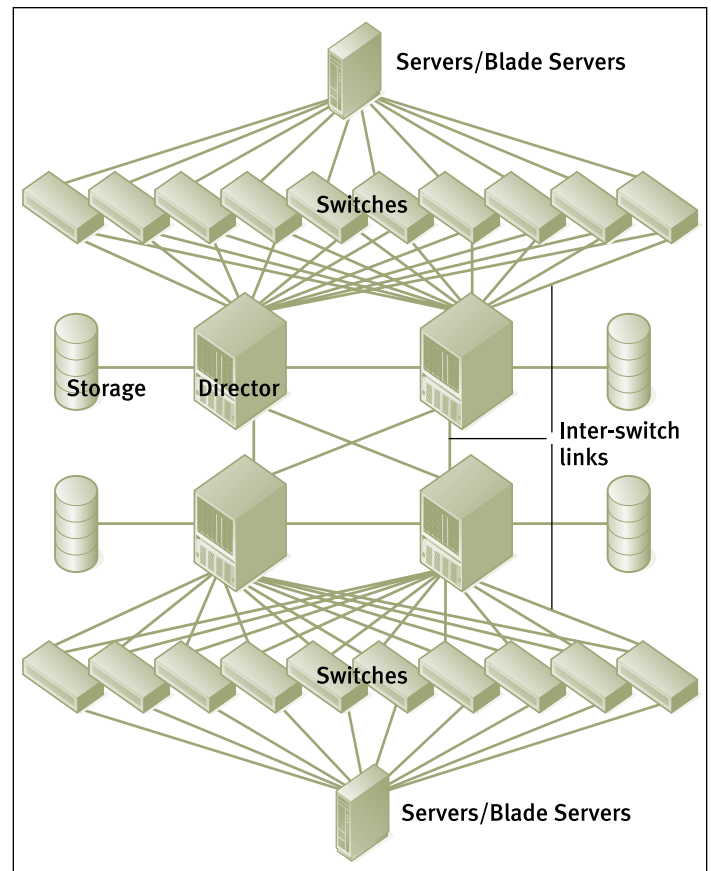


Figure 1 Example of a Large FC Fabric



based on FC directors^[2], which are high-availability switches with high-aggregate switching bandwidth and high port density. For the edge part of a large or small fabric, smaller and lower-cost FC switches are typically used. Directors and switches use one or more interswitch links (ISLs) to connect and form a larger fabric. It is common to deploy one or more isolated FC fabrics, called SAN islands. SANs are also extended to campus, metropolitan, and wide-area distances using T1/T3, ATM, IP, SONET, dark fiber, and DWDM technologies.

Figure 1 shows an example of a large (approximately 1000 node) fabric, with directors and switches configured to provide high availability and high-aggregate bandwidth. Servers are typically aggregated at the edge of the fabric, and storage arrays are typically configured near the core of the fabric. It is typical to over-subscribe server link bandwidth in comparison to storage array link bandwidth (more servers with respect to a given storage array). A network of directors forms the core (or backbone) of the fabric.

For a fabric to be operational, there is a fabric initialization, involving all switches, directors, and devices. Initialization steps include parameter exchanges, Principal Switch selection, address assignment, path computation, and zone merge operations. As part of the path computation, directors and switches in a fabric run Fabric Shortest Path First (FSPF) routing protocol to build the forwarding data base. FSPF is a link state protocol that computes shortest path routes for frame forwarding. Within a FC fabric, there are name services and state change notification protocol services for resource discovery, configuration, and change management. FC zoning is an overlay network mechanism to limit the visibility and connectivity of servers to storage devices. A device can be in one or more zones, thereby enabling the sharing of servers (or clusters of servers) and storage resources. When an ISL changes state, all these protocols normally run, and when an end device comes up or goes down, name services and state change notification services run. These services consume more and more resources as the fabric size grows.

Traffic Patterns in Fibre Channel

Most FC traffic uses the SCSI-FCP protocol^[4] on top of FC Class 3 (unacknowledged datagram) service. SCSI-FCP is a request-response protocol that provides frame sequencing within transactions provided by lower-layer FC protocols. On frame loss or error, the protocol performs a transaction-level time out and retransmission. No retransmission of individual frames is supported. Time-out values are typically pre-configured and not based on actual round trip delay. The performance of SCSI-FCP is therefore sensitive to frame loss or frame level errors. Table 1 shows example read and write transactions and protocols frames.

TRANSACTION	PROTOCOL DIRECTION	FRAME TYPE	TYPICAL FRAME LENGTH
Read	Server to Storage	FCP_CMD (Read)	68 Bytes
	Storage to Server	FCP_XFER_RDY	48 Bytes
	Storage to Server	FCP_DATA (one or more)	Up to 2084 Bytes
	Storage to Server	FCP_RSP	64 Bytes
Write	Server to Storage	FCP_CMD (Write)	68 Bytes
	Storage to Server	FCP_XFER_RDY	48 Bytes
	Server to Storage	FCP_DATA (one or more)	Up to 2084 Bytes
	Storage to Server	FCP_RSP	64 Bytes

Table 1 Example SCSI-FCP Read and Write Protocol Frames

As bandwidth and delay product increases, it is critical to understand performance tuning and error recovery mechanisms. For configurations with long delay, it is important to consider the way data is moved (write or read). As shown in Table 1, the write transaction has one additional round trip delay more than the read transaction. Therefore, the read operation is faster when network delay is long.

Critical Factors in SAN Deployment

SAN deployments today range from small fabrics with less than 100 devices to large fabrics with several thousand devices. The following are factors critical to SAN design and deployment:

- High availability: The impact of down-time and lost of information to business is severe. High availability requirements are quantified to vary from several 9's, to

99.999%, to no down time. Most highly available fabrics are based on dual-rail redundancy and highly available directors, switches, and gateways. Servers and storage devices may have redundant paths through one fabric or through separate redundant fabrics with no shared single point of failure. Directors and some switches are designed with high-availability features, including fully redundant and hot swappable field-replaceable units (FRUs) and hot software download and activation, meaning that operation may continue through a software upgrade.

- **Robustness and stability:** Some FC servers, associated host bus adapters (HBAs) and storage devices are extremely sensitive to frame loss and frame out of order delivery. Error recovery in the SCSI-FCP protocol is based on command and transaction level time-out and retry. Therefore, SCSI-FCP expects very low frame loss rate, since frame loss has significant performance impact. The design of SANs has to account for the following factors:
 - It is important to limit and reduce FC fabric size in terms of number of switching nodes. The goal is to limit the frequency of fabric initialization, FSPF route computation, and traffic for state notification and name services.
 - It is critical to ensure there is adequate aggregate bandwidth (fabric-wide and for individual links), to avoid severe and prolonged congestion. FC fabrics use a link-level, credit-based flow control, which is useful for handling short-term, bursty congestion. In FC, it is not common to use active queue management techniques (e.g., based on random early detection) to minimize queue build up. It is typical for a FC switch to discard frames that have been queued for a pre-determined time (e.g., 0.5 to 1.0 second), as part of the stale frame discard policy. As the deployment of multi-speed (1 Gbps, 2 Gbps, 4 Gbps, and 10 Gbps) ramps up, the design of the network and switching architecture becomes more challenging. As the size of network grows, comprehensive congestion management mechanisms become more critical and current link-level flow control may no longer be adequate.
- **Performance:** Most FC switches and directors specify best-case frame latency to be less than a few microseconds. But latency grows with loading and can result in effective bandwidth to be significantly less than nominal bandwidth. Measured frame latency at 70% link utilization^[3] showed it was 5.2 to 6.5 microseconds for one vendor's product and 2.6 to 222.6 microseconds for another vendor's product. The lesson is that not all switches are designed equal. Switching architecture

issues like head-of-line blocking and internal resource bandwidth (throughput or frame rate) limitations impact throughput, latency, and congestion loss, especially at higher offered load.

- **Distance extension:** Requirements for disaster recovery and business continuance (file/data mirroring, replication, and backup) are driving the deployment of SAN extension to deliver better performance and availability. In addition to robustness, stability, and performance considerations, it is important to understand the configurations, products, and protocols and system tuning parameters with respect to distance extension technology. We examine this topic later.
- **Scaling the SAN:** A large number of FC fabrics deployed today are small islands of fabrics that are not inter-networked into a large and connected SAN. Reasons for deploying isolated islands include early adopters learning new technology, difficulty and lack of confidence in management and operational stability of a large fabric, and insufficient business and operational drivers (for connecting islands of FC fabrics). However, there are many benefits of interconnecting FC islands. Resource sharing (such as tape library for backup) and the ability to dynamically provision and allocate resource are some of the benefits. When scaling an FC SAN, it is important to maintain performance and availability properties. Since a FC fabric is similar to an IP layer 2 switching network, it is important to constrain the number of switches in a fabric so the resulting fabric is stable and robust. When interconnecting FC fabrics, it is critical to consider isolating FC fabric local initialization and services, while allowing servers and storage devices to be interconnected regardless of locality. This is an area of further research and standardization work, and currently ANSI T11 has a fabric extension study group addressing these topics.

FC & IP INTEGRATION & CHALLENGES

IP SAN Developments

The emergence of iSCSI, FCIP, and iFCP standards^[5, 6, 7, 8, 9] enables IP technology to enhance the deployment and benefits of SANs. FCIP and iFCP protocols use a common framing and encapsulation design. We examine the applicability, design, and limitations of these technologies in the following sections. These protocols leverage the matured IPsec standard and technology to enable security (including authentication, integrity, and privacy). As part of the protocol suite, Internet Storage Name Service (iSNS)^[10] provides a method to manage



and configure names, registry, discovery, and zones for multi-protocol SANs. The use of Service Location Protocols (SLP)^[11] to discover services and resources is another critical part of the standard.

iSCSI

iSCSI is a SCSI over TCP transport protocol used between a SCSI initiator and a SCSI target for storage-block level transport of SCSI commands and payloads. iSCSI protocol uses TCP/IP and IPSec as its network transport and security protocols. It has many features designed to leverage standard TCP/IP protocols to block storage needs. These features include the use of multiple TCP connections (for a given session), cyclic redundancy check (CRC) digests, out of order data placement, and TCP connection failure recovery options. iSCSI design and analysis have been presented in several papers^[12, 13, 14, 15, 16].

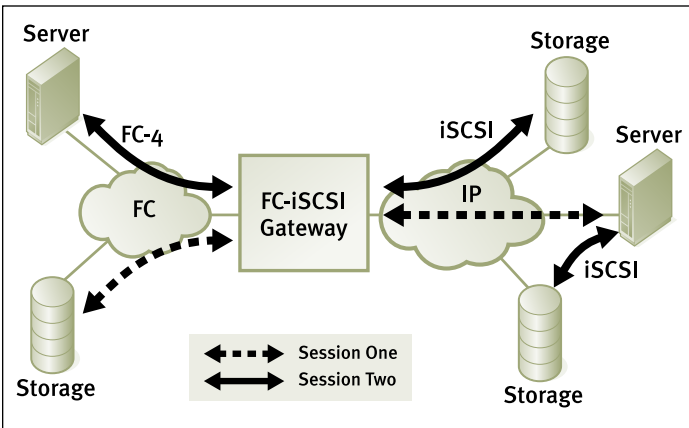


Figure 2 FC-iSCSI Gateway

In Figure 2, the FC-iSCSI gateway provides the internetworking of iSCSI devices with FC devices, while communicating with each of the networks appropriately. While the FC SAN and the IP SAN are operating independently, the gateway maps selected iSCSI devices into the FC SAN and selected FC devices into the IP SAN. When a FC server creates a SCSI-FCP session to a storage device, the gateway intercepts the request and acts as a proxy for the storage device. On the IP side, the gateway acts as a proxy initiator (for the server), and creates an iSCSI session for the storage device. The gateway maintains and manages the state of the gateway portion of supported

sessions. For an IP-based server creating an iSCSI session to a FC storage device, the gateway performs similar roles as proxy target on iSCSI session and proxy initiator for the SCSI-FCP session.

An iSCSI gateway performs several important functions, including FCP and iSCSI session-level protocol translations, command and payload forwarding, error checking, and command/session-level error propagation. A gateway has to manage device discovery and registry (on the IP side with an iSNS server, and on the FC side with FC name services), authentication of FC and IP devices, and mapping of device names to local addresses, etc. It is important that a gateway is as transparent as possible to the servers and storage devices using the gateway, while maintaining high data integrity. It should have very low latency and sufficient bandwidth to forward commands and payloads, and support a sufficiently large number of sessions to enable storage consolidation (a high end storage array on the FC side shared by a large number of IP based servers). Management of the multi-protocol SAN is a critical part of the deployment success.

FCIP

FCIP is a tunneling protocol that transports all FC ISL traffic. Similarly, FCIP uses TCP/IP as the transport protocol and IPSec for security. A FCIP link tunnels all ISL traffic between a pair of FC switches, and may have one or more TCP connections between a pair of IP nodes for the tunnel end points. From the FC fabric view, an FCIP link is an ISL transporting all FC control and data frames between switches, with the IP network and protocols invisible. One can configure one or more ISLs (using FCIP links) between FC switches using FCIP links. Figure 3 shows an example of FCIP links being used as ISLs between FC switches A and B.

A key advantage of the FCIP tunnel approach is transparency to a fabric, as existing fabric tools and services are used. Once a FCIP link is configured, existing fabric operations and management continue. Similarly, fabric initialization, FSPF routing protocol, and name/state change services run transparently over FCIP links. However, since FC fabric-level control protocols run over the FCIP tunnel, IP and TCP connection failures can disrupt the FC fabrics on

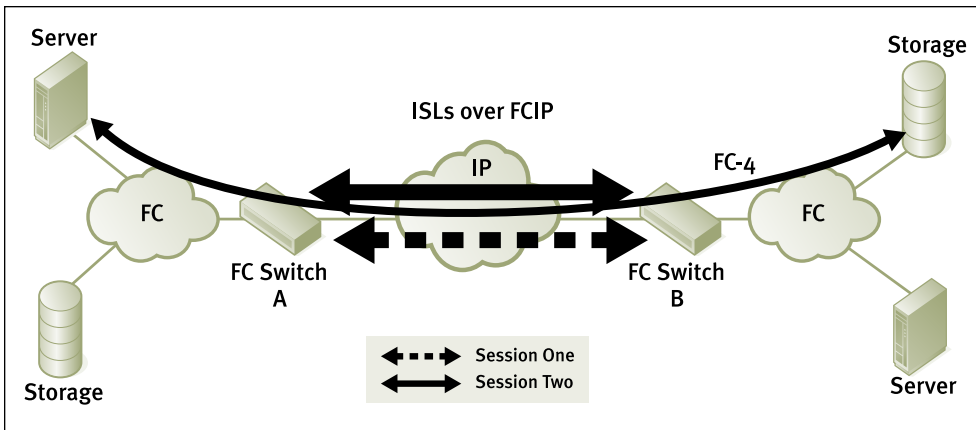


Figure 3 FC-FCIP Tunnel

both sides. Given the speed and bandwidth differences between FC and a typical IP network used to interconnect remote SANs, the design and management of congestion and over-load conditions is important to understand.

For the FCIP tunnel, a simple FIFO (first in first out) frame forwarding queue design can result in head-of-line blocking of fabric initialization protocol frames when the tunnel is congested, or the TCP connection is in slow-start recovery mode. Another case to consider is when a SCSI-FCP transaction time out occurs, the entire transaction (such as 1 MB block) might be retransmitted over an FCIP link that is experiencing congestion. In addition, there might be multiple application streams using the same FCIP link, and there is no mechanism to help reduce or avoid network congestion. These are possible scenarios of overload and congestion that can result in performance and stability issues that impact the entire fabric. For a medium to large fabric, these are critical issues for concern. Most FCIP deployments are based on small fabrics, where there are a small number of devices

and switches at each end of the FCIP link, and these issues are less critical.

iFCP

iFCP technology is a gateway-to-gateway protocol for providing FC device-to-FC device communication over TCP/IP. For each pair of FC devices, there is an iFCP session created between a pair of gateways supporting the devices. An iFCP session uses a TCP connection for transport and IPSec for security, and manages FC frame transport, data integrity, address translation, and session management for a pair of FC devices. Since

an iFCP gateway handles the communications between a pair of FC devices, it only transports device-to-device frames over the session, and, hence, the FC fabrics across the session are fully isolated and independent. This is a major difference between iFCP and FCIP, in that FCIP builds an extended fabric, tunneled over IP.

In contrast to an FC-iSCSI gateway, an iFCP gateway transports FC device-to-device frames over TCP, and in most cases original FC frames, including the original CRC and frame delimiters, are

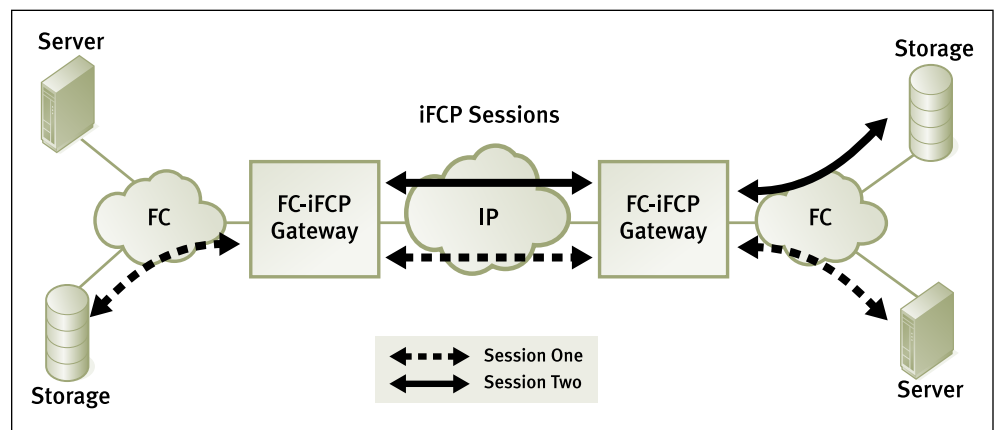


Figure 4 FC-iFCP Gateway



transported. An FC-iSCSI gateway terminates and translates SCSI-FCP protocol from the FC side and similarly terminates and translates iSCSI protocol from the IP side.

The iFCP draft standard specifies an address-translation mode as well as an address-transparent mode, depending on whether the FC addresses of devices are translated or not. An FC device exchanges login and protocol parameters with another FC device using FC link service protocol frames as part of the session creation and parameter exchange protocol. In address-translation mode, a gateway intercepts these device-to-device link service protocol frames and translates device addresses embedded in the frames. It regenerates frame CRCs when the original frame content is changed, which imposes extra overhead on the iFCP gateway. Address translation is a particularly useful feature when interconnecting FC fabrics. It enables installation of a gateway between existing fabrics without requiring fabric address changes. A gateway manages the session state, addresses translation and mapping, and provides proxy functions for remote devices. In addition, a gateway performs security functions (like authentication of devices), and works with an iSNS server for registry and discovery functions.

The configuration and management of an iFCP gateway is more involved than for an FCIP gateway, as each device-device session has to be set up. Also, an iFCP gateway has more device proxy-related states to manage. As the number of device-to-device sessions increases, an iFCP gateway design becomes more complex and may result in performance and stability issues. However, one can use admission control techniques to limit the number iFCP sessions allowed for a gateway. Since an iFCP gateway is managing device-to-device communications, it can enforce some degree of flow control by pacing command forwarding at the time of congestion. The iFCP specification allows an optional unbounded connection feature, which sets up and uses a pool of backup TCP connections for fast-session fail-over support. This assists a gateway in providing faster connection fail-over.

TCP/IP & Transport Protocol Discussions

Some classes of applications have different requirements for transport services and protocols. For example, applications that prefer timeliness in delivery over reliable data delivery (such as RealAudio, Voice over IP) prefer a different transport service and protocol design^[17] than that of TCP. Also, for applications that prefer a different type of fault tolerance, reliability, and a non-byte stream-oriented transport service, a different type of transport protocol might be needed (such as Stream Control Transmission Protocol (SCTP)^[18]). These are examples of new transport protocol research and standard development activities. TCP protocol is undergoing many enhancements to improve performance under different operating conditions, and these enhancements include High Performance Extensions (TCP Window Scaling Option, Round-Trip Time Measurements, Protect Against Wrapped Sequence Numbers)^[19], Selective Ack Option^[20, 21], Explicit Congestion Notification^[22, 23], Eifel Detection Algorithm^[24], and High Speed TCP (HSTCP)^[25].

As part of the design considerations for an IP SAN, the design and tuning of TCP for the SAN is critical. For iSCSI servers and storage devices, the design and tuning of protocol off-load, zero-copy, interrupt coalescing, and buffer-MTU-MSS tuning are critical (MTU is the maximum transfer unit, MSS is the maximum segment size). For iSCSI, FCIP, and iFCP gateway design, buffer-MTU-MSS tuning is very critical and several of the aforementioned TCP enhancements are important considerations for scaling the IP SAN for 1 to 10 Gbps speeds. For long and fast network (LFN), HSTCP enhancement is an important design. Multiple TCP connections for iSCSI, FCIP link, and unbounded iFCP connections are critical considerations for load balancing and high availability.

In addition to the IP based transport, there are developments for operating Gigabit Ethernet and FC protocol directly over SONET-based transports for Generic Frame Protocol ITU-T G.7041 standards^[26].

SOME CASE STUDIES

Experiment of 10 Gbps Transcontinental Data Access

As part of the Supercomputing Conference 2002 demonstration of SAN extension over a multi-gigabit transcontinental network,^[27]

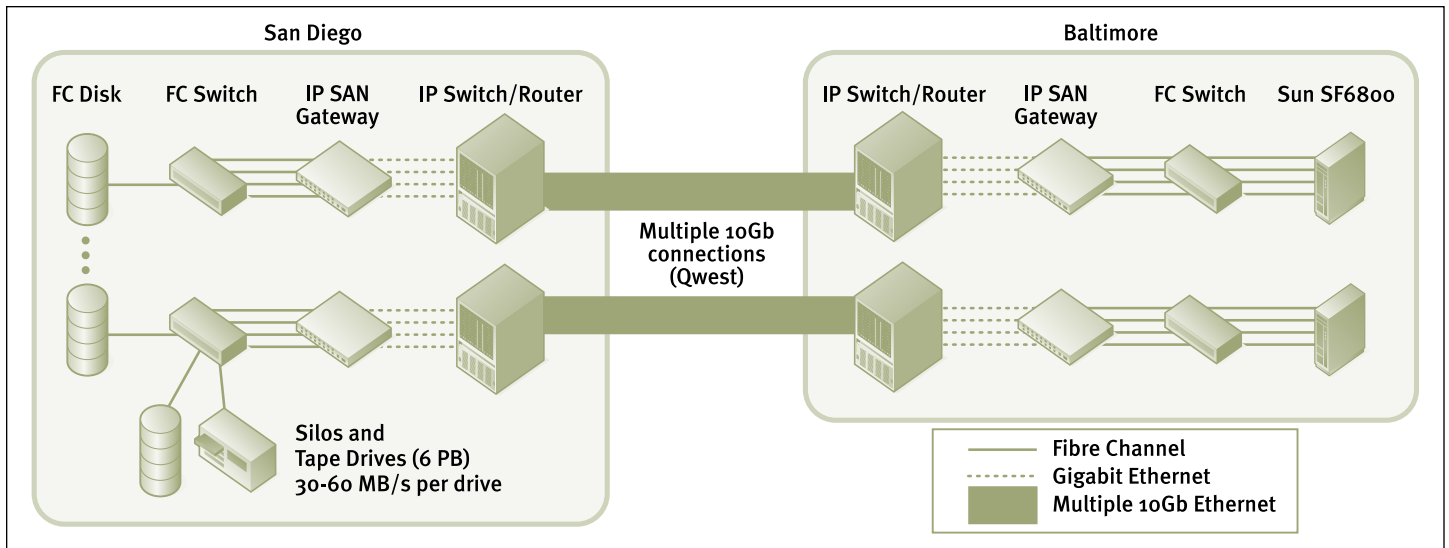


Figure 5 Schematic of Data Access for the SC'02 demonstration

test results of an FC SAN interconnected with iFCP gateways over a 10 Gbps link from San Diego to Baltimore were presented. Figure 5 shows the configurations used for the experiment.

The Supercomputing '02 experiment proves the operation of a network running FC over IP network, using iFCP gateways, between the San Diego Supercomputer Center (SDSC) and the SDSC booth in Baltimore. The experiment demonstrates that FC traffic, using iFCP gateways,

runs over a 10 Gbps link in excess of 2,600 miles, with a round-trip latency of 70 to 90 milliseconds. Aggregate throughput was relatively constant at 717 MB/s, and read performance was slightly better than write performance. In addition to the IP/iFCP based demo^[27], there was another experiment of FC traffic over FCIP using a 10 Gbps SONET link, configured between the Pittsburgh Supercomputing Center (PSC) booth and the SDSC booth at the SC'02 show.

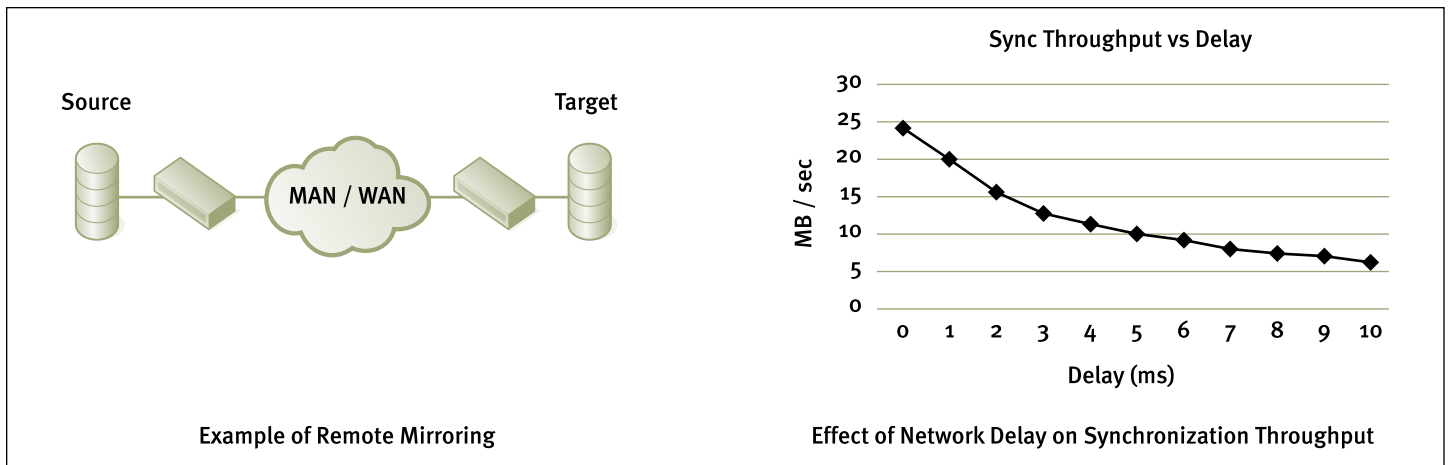


Figure 6 Remote Mirroring – Throughput vs Delay



Remote Mirroring

When performing remote mirroring of logical units (LUNs), remote copy operations must synchronize data copied to each of the LUNs to ensure data coherency within the mirror group. The effective throughput of the remote mirroring of 12 LUNs was shown to drop from 25 MBps to about 5 MBps as the round trip delay increases from 0 to 10 ms, as shown in Figure 6^[28]. It is important to configure and tune file and block size, MTU, MSS, and synchronization rate. In addition, the use of compression to reduce the amount of data transfer is important.

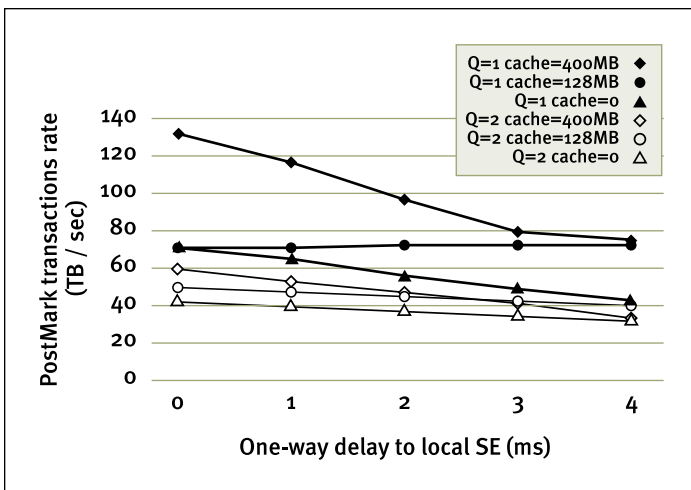


Figure 7 Delay and Cache Effect of PostMark Experiment

Delay and Cache Effect on I/O Workload

PostMark was used to test the delay and cache sensitivity of the I/O workload of a large email server^[29]. Figure 7 shows the PostMark transactions rate of I/O from a FreeBSD host to a storage element (SE) with varying delays (to the SE) and cache sizes in FreeBSD VM cache. The transaction rate declines as the delay is increased, and with larger cache sizes the transaction rate increases. Application performance sensitivity with respect to delay and error recovery is an area that needs further research and understanding.

Long Fast Network Experiment

In another case^[30], the University of Tokyo conducted an experiment using ‘iperf’ running TCP between Maryland and Tokyo, traversing the Abilene and APAN networks. The result was surprising in that Fast Ethernet is sometimes faster than Gigabit Ethernet on LFN. The main cause of the throughput degradation with Gigabit Ethernet LFN tests was congestion overflow of an intermediate router, resulting in cranking of TCP time out, slow start and congestion control mechanisms. Transmission rate control is important to mitigate the overflow in the bottleneck’s buffer in addition to the window size control. Therefore, transmit rate or bandwidth limiting is another important mechanism that avoids or mitigates the impact of congestion overflow in an intermediate network.

Fast Write™

We examine a method to improve the write performance over a long delay network. As shown in Table 1, a SCSI write transaction incurs two round trip delays for a data block. The maximum block

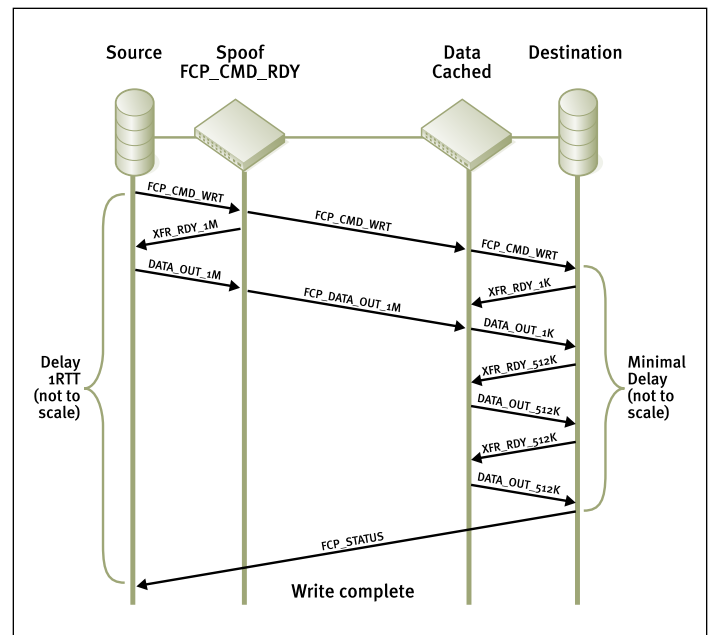


Figure 8 Fast Write Example

size is determined by the target (storage) device's buffer capacity and is specified by the target in the XFR_RDY message. For example, writing one MB of data using 64 KB blocks takes 16 transactions, which is 32 round trips plus data transfer time. Fast Write^[31] is a way to minimize round-trip delay overhead and accelerate SCSI write performance leveraging a gateway's buffer capacity. The XFR_RDY is spoofed by the gateway on the initiator (server) side of the network, and the data is buffered by the gateway on the target side of the network until the target sends its own XFR_RDY. In addition, the use of TCP protocol with selective retransmission (on error) provides better frame loss recovery than retransmitting the entire block on timeout (as in the SCSI-FCP case). With Fast Write, the number of round trip involved for a 1 MB transfer is reduced to two round trips.

Figure 8 shows an example of Fast Write, where a 1 MB transfer is negotiated between the source and the left-hand gateway as well as between the two gateways. For the write operations between the final gateway and destination, the maximum block size is specified by the destination. Most of the round-trips required are over the SAN between the right-hand gateway and destination. Therefore, the SAN should have higher bandwidth, lower latency, and lower error rate than the WAN connection between gateways. Fast Write is an innovative method of using standard protocols to leverage the capability of a gateway and leverage TCP protocol benefits over WAN.

SUMMARY

We present several of the critical requirements and best practices for FC SAN deployments. We examine IP SAN technologies and protocols, and show that FC and IP integration works well - integrated SANs are a critical part of today's data center. We explore how several new high-speed protocol extensions work, and areas that need further research and development. The deployment of high-speed and long-distance networks for data centers (while providing good performance and reliability) is becoming very important and has potential value as well as challenges.



REFERENCES

1. Latest FC standards and drafts can be found at www.t11.org.
2. H. Yang, "Performance Considerations for Large-Scale SANs", McDATA Corporation, 4 McDATA Parkway, Broomfield, CO 80021, December 2000. www.mcdata.com.
3. Randy Birdsall, "Competitive Performance Validation Test", Miercom Labs, Princeton Junction, NJ, June 2002.
4. Stephen Trevitt, "Traffic Patterns in Fibre Channel", McDATA Corporation, May 2002. Available at www.mcdata.com.
5. RFC 3643 - Fibre Channel (FC) Frame Encapsulation. www.ietf.org.
6. Draft-ietf-ips-iscsi-20.txt, January 19, 2003. www.ietf.org.
7. Draft-ietf-ips-fcovertcpip-12.txt, August 2002. www.ietf.org.
8. Draft-ietf-ips-ifcp-14.txt, December 2002. www.ietf.org.
9. T. Clark, "IP SANs A Guide to iSCSI, iFCP, and FCIP Protocols for Storage Area Networks", Addison-Wesley, ISBN 0-201-75277-8, October 2002.
10. Draft-ietf-ips-isns-21.txt, October 2003. www.ietf.org.
11. Draft-ietf-ips-iscsi-slp-06.txt, December 2003. www.ietf.org.
12. K. Meth, J. Satran, "Design of the iSCSI Protocol", IEEE/NASA MSST2003 Twentieth IEEE/Eleventh NASA Goddard Conference on Mass Storage Systems & Technologies, April 7-10 2003, storageconference.org/2003.
13. H. Thompson, C. Tilmes, R. Cavey, B. Fink, P. Lang, B. Kobler, "Considerations and Performance Evaluations of Shared Storage Area Networks At NASA Goddard Space Flight Center", IEEE/NASA MSST2003 Twentieth IEEE/Eleventh NASA Goddard Conference on Mass Storage Systems & Technologies, April 7-10 2003, storageconference.org/2003.
14. K. Voruganti, P. Sarkar, "An Analysis of Three Gigabit Networking Protocols for Storage Area Networks", 20th IEEE International Performance, Computing, and Communications Conference, April 2001.
15. S. Aiken, D. Grunwald, A. Pleszkun, J. Willeke, "A Performance Analysis of the iSCSI Protocol", IEEE/NASA MSST2003 Twentieth IEEE/Eleventh NASA Goddard Conference on Mass Storage Systems & Technologies, April 7-10 2003, storageconference.org/2003.
16. P. Sarkar, K. Voruganti, "IP Storage: The Challenge Ahead", 10th Goddard Conference on Mass Storage Systems and Technologies/ 19th IEEE Symposium on Mass Storage Systems, April 15-18, storageconference.org/2002.
17. Draft-ietf-dccp-problem-00.txt, October 23, 2003. www.ietf.org.
18. RFC 2960 – Stream Control Transmission Protocol, October 2000. www.ietf.org.
19. RFC 1323 – TCP Extension for High Performance. www.ietf.org.
20. RFC 2883 – An Extension to the Selective Acknowledgement (SACK) Option for TCP, July 2000. www.ietf.org.
21. RFC 3517 – A Conservative Selective Acknowledgement (SACK) – based Loss Recovery Algorithm for TCP, April 2003. www.ietf.org.
22. RFC 3168 - The Addition of Explicit Congestion Notification (ECN) to IP, September 2001. www.ietf.org.
23. S. Floyd, "TCP and Explicit Congestion Notification", Lawrence Berkeley Laboratory, One Cyclotron Road, Berkeley, CA 94704, ACM Computer Communication Review, V. 24 N.5, October 1994, p. 10-23. www.icir.org/floyd/papers.html.
24. RFC 3522 – The Eifel Detection Algorithm for TCP, April 2003. www.ietf.org.
25. RFC 3649 – HighSpeed TCP for Large Congestion Windows, December 2003. www.ietf.org.
26. ITU-T G.7041 Standards, reference www.itu.int/ITU-T/.
27. P. Andrews, T. Sherwin, B. Banister, "A Centralized Access Model for Grid Computing", IEEE/NASA MSST2003 Twentieth IEEE/Eleventh NASA Goddard Conference on Mass Storage Systems & Technologies, April 7-10 2003, storageconference.org/2003.
28. EMC Corporation, "MirrorView Using Fibre Channel over IP", December 30, 2002.
29. E. Gabber, J. Fellin, M. Flaster, F. Gu, B. Hillyer, W.T. Ng, B. Ozden, E. Shriver, "StarFish: highly-available block storage", Proceedings of the FREENIX track of the 2003 USENIX Annual Technical Conference, San Antonio, Tx, June 9-14.
30. M. Nakamura, M. Inaba, K. Hiraki, "Fast Ethernet Is Sometimes Faster Than Gigabit Ethernet on LFN – Observations Of Congestion Control Of TCP Streams", University of Tokyo, 7-3-1, Hongo, Bunkyo-ku, Tokyo, Japan, www.data-reservoir.adm.s.u-tokyo.ac.jp/paper/pdcs2003.pdf.
31. McDATA Corporation, "Maximizing Utilization of WAN Links with Nishan Fast Write", McDATA San Jose Development Center, 3850 North First Street, San Jose, Ca 95134. 2002.

