

# Dispelling the FUD surrounding iSCSI

A Cisco executive counters the fear, uncertainty, and doubt generated by opponents of the iSCSI specification, which will enable IP-based SANs.

■ BY CLINT JURGENS

When a technology as promising as IP storage networking emerges—and threatens established technologies—fear, uncertainty, and doubt (FUD) usually follow, spread mostly by potential competitors.



CLINT JURGENS  
Cisco Systems

Such has been the case with the introduction of the Internet SCSI (iSCSI) protocol. iSCSI, with Fibre Channel over IP (FCIP) and Internet Fibre Channel Protocol (iFCP), make up a suite of proposed storage networking protocols that operates over TCP/IP networks. Despite the best efforts of backers, much of the early buzz about iSCSI has attempted to raise doubts about whether it is both powerful and reliable enough to perform enterprise-class storage networking applications. With so much contradictory information floating around, skepticism and confusion abound.

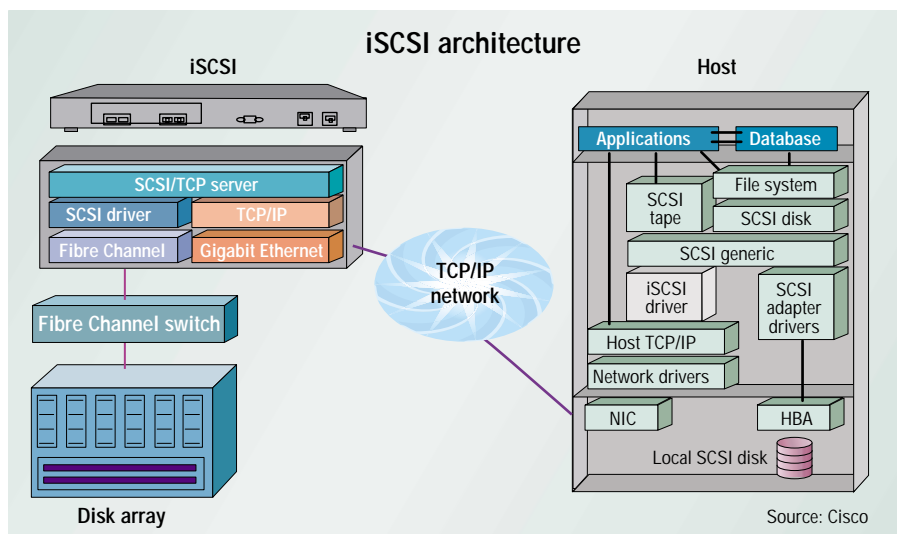
iSCSI is a proposed standard jointly authored by Cisco and IBM and submitted in February 2000 to the Internet Engineering Task Force (IETF) for ratification. Today, more

than 250 companies have committed to developing solutions using iSCSI. Several companies such as Adaptec, Alacritech, Emulex, Cisco, IBM, and Intel have announced iSCSI-based products, some shipping now (in “pre-iSCSI” implementations) and others by year-end.

iSCSI encapsulates block-level SCSI data in a TCP/IP frame, thus allowing servers to access storage resources over an existing IP infrastructure. It requires two components: an initiator (e.g., an iSCSI driver embedded in a server) and a target (e.g., a gateway device that routes data between servers and the storage at the back-end).

Common storage networking applications for iSCSI will include

- **Storage consolidation**—The ability for servers running multiple network operating systems to share a disk storage subsystem over an iSCSI-enabled IP network.



Source: Cisco

- **Remote IP access to storage**—Servers located in remote locations can have access to centralized storage (located in a storage network at headquarters, for example) using iSCSI. This is an ideal application for a small to medium-sized branch office that might not have the benefit of a SAN administrator onsite. Using iSCSI allows the local network administrator to deploy, configure, and manage iSCSI devices like any other network devices.
- **Business continuance**—IT operations can use a TCP/IP network to remotely mirror servers and storage. With iSCSI, these devices logically appear as local devices but provide insurance against a local catastrophe shutting down operations.

As for what iSCSI *isn't*, the rest of this article will focus on dispelling the common misperceptions.

**FUD:** *iSCSI takes up a great deal of CPU processing power in servers, which hinders performance.*

**FACT:** It is true that it takes significant processing cycles to process the TCP/IP protocol stack at gigabit speeds. This is true for all traffic, not just iSCSI. It is also true that this challenge is being addressed by a variety of networking and storage vendors. New adapters that have special engines to process the protocol stack and iSCSI are due by year-end. The result will be hardware acceleration in Gigabit Ethernet adapters that consume no more CPU cycles than Fibre Channel.

**FUD:** *iSCSI is designed to work over IP networks, which introduces latency issues, especially over longer distances.*

**FACT:** The laws of physics work the same for both Fibre Channel and iSCSI. If an application needs to provide storage access over long distances, then the latency will be the same. For longer distances, Fibre Channel encapsulates frames and sends them via IP. The number of router hops is identical for both Fibre Channel and iSCSI. What is different is the additional buffering (latency) to encapsulate and recover Fibre Channel frames. Some Fibre Channel

---

## *If the whole world has access to internal networks via TCP/IP, then there is a security risk.*

---

encapsulation schemes use UDP/IP. While this may be faster, it is an unreliable connectionless service that can add delays for Fibre Channel to recover due to frame loss and out-of-sequence delivery.

**FUD:** *Unlike Fibre Channel, iSCSI has no reliable flow control method, leading to congestion and dropped frames.*

**FACT:** Both Fibre Channel and iSCSI over Ethernet use two forms of flow control: link layer and end-to-end. Link-layer flow control is between two adjacent nodes in the network, and end-to-end flow control is between the sending and receiving nodes in the network.

At the link layer, Fibre Channel uses a buffer-to-buffer flow control system. As Fibre Channel frames are received, they are buffered. A primitive signal called Receiver Ready, or R\_Rdy, is sent from the receiver to the transmitter to signal the ability to receive another frame. This process ensures the receiving node is able to buffer a frame before the transmitter sends it, thus preventing over-run at the receiver.

Full-duplex Ethernet provides a similar mechanism of link-layer flow control. Each node has the ability to signal XON and XOFF, or transmitter on and transmitter off. As long as the receiving node is able to move the received frames through its buffers without congestion, there is no signal sent to the transmitting node. If the receiving buffer reaches a high-water mark, then a transmitter off (XOFF) signal is sent, stopping the transmission temporarily. As soon as the receiving buffer begins to empty and is able to receive more frames, the transmitter is turned on (XON).

Both the Fibre Channel and full-duplex Ethernet methods of link-layer flow control are implemented in hardware. Therefore, data is sent as fast as the receiving node is able to handle it without congestion or dropped frames.

The second method of flow control is

end-to-end, which only works in Fibre Channel for acknowledged connectionless service (Class 2). (*Note:* This also works for Class 1, but no one uses Class 1. For storage, only Class 2 and Class 3 are used.) A significant amount of Fibre Channel traffic is sent using unacknowledged connectionless service, or Class 3. This is similar to UDP or datagram traffic. There is no assurance that the information is received. Class 2 can operate in different modes, acknowledge each frame, acknowledge “*n*” frames, or acknowledge a complete sequence. Most implementations of Fibre Channel send a single acknowledgment for a complete sequence. The whole block of data is sent and if it arrives without error, then it is acknowledged.

iSCSI uses TCP, which provides an acknowledged connection-oriented service. TCP works by starting slowly and gradually building up speed until data is lost. Then TCP recovers and backs off to a speed that accommodates the capacity of the end-to-end link. The mechanism that TCP uses to control the flow of data is a sliding window, which starts out as a small number of frames. This quantity of *n* frames is sent by the sending end and acknowledged by the receiving end. Once the sending end receives the acknowledge signal, a second batch of *n* frames is sent. This window of *n* frames is increased as long as no congestion occurs. Therefore, TCP uses a flow-control mechanism that tunes itself to maximize throughput according to the capacity of the network. iSCSI keeps TCP connections open and begins the next iSCSI sequence of operations using the last negotiated window size.

**FUD:** *Using iSCSI introduces security risks inherent to IP networks that are not prevalent in Fibre Channel networks in the data center.*

**FACT:** If the whole world has access to

internal networks via TCP/IP, then there is a security risk. However, this problem is not limited to iSCSI and has been addressed in depth by the networking community. Private networks, encryption, and strong authentication are some of the ways iSCSI and other TCP/IP-based applications provide security. The bottom line is that there are ways to prevent nodes on a TCP/IP network from gaining unauthorized access to iSCSI.

With Fibre Channel, there is only switch zoning. Today, Fibre Channel does not provide encryption, and a rogue node may bypass the zoning to gain unauthorized access to data. If there is no switch, and Fibre Channel is operating on a loop, then there is no mechanism to prevent unauthorized access to data on the loop.

**FUD:** *iSCSI introduces more traffic over corporate LANs, which are already overburdened with networking traffic.*

**FACT:** The intent of SCSI over TCP/IP is to run storage traffic on high-performance switched Ethernet networks. It may be desirable to run storage traffic on an Ethernet network physically or logically separate from the primary LAN network to meet performance and service-level requirements. For example, virtual LAN (VLAN) technology would ensure separate bandwidth for the two networks, while still running over the same physical network infrastructure. Just as Fibre Channel is a

private network, most iSCSI storage networks will be private networks. iSCSI provides a single type of network for training and provisioning, with the additional benefit that the data network is available for use if required.

**FUD:** *Transporting iSCSI data over IP networks is unreliable because of IP's tendency to drop and disorder packets.*

**FACT:** Both Gigabit Ethernet and Fibre Channel use the same physical layer. In fact, most physical layer 8B/10B encoding chips and transceivers are designed to work in either application. Therefore, both operate on equally reliable physical links.

iSCSI uses the reliable connection-oriented service provided by TCP. While it is true that IP is a connectionless service that can disorder packets, TCP fixes the packet ordering and iSCSI only sees in-order packets that are confirmed to be sent correctly from the sending node to the receiving node.

This is not true for Fibre Channel. Much of Fibre Channel storage networking uses Class 3 service or unacknowledged connectionless service. This is satisfactory for loop operation, because the loop establishes a virtual point-to-point link between two nodes, and there is no way for the packets to get lost or congested. However, switches that use Class 3 have no way to guarantee delivery. There is risk of out-of-order delivery and packet loss via congestion. The result is that TCP/IP can be more reliable than Fibre

Channel Class 3 operation.

**FUD:** *Ethernet's small frames will hinder performance of transporting iSCSI data.*

**FACT:** Technology has moved beyond interrupting the host to handle every frame that arrives on the network. ASICs have the intelligence in hardware and firmware to move data with only a single interrupt to discover where to store the data. The speed at which the data is moved to the application memory space is paced by the speed of the network. Internal architectures do not slow down the network or hinder the movement of iSCSI data.

**FUD:** *Standard applications will not run unless modified for iSCSI.*

**FACT:** One of the fundamental benefits of iSCSI is that applications do in fact operate without modification. The iSCSI driver, when installed in a server, simply intercepts the SCSI command data block and creates an iSCSI protocol data unit that is sent via a standard network interface card. The iSCSI driver is a layer below the SCSI driver and does not impact applications. With time and factual information, iSCSI should live up to its promise as the technology that changes the face of storage networking. □

---

**Clint Jurgens** is business development manager in Cisco Systems' Storage Router Business Unit, based in Maple Grove, MN.