



iSCSI Whitepaper

Celeros
Bringing SANity to Storage Costs

iSCSI History

Servers typically talk to their Direct Attached Storage (DAS) through a *block* level IO interface called SCSI (Small Computer System Interface). As the need to realize the benefits of pooled, centrally managed storage increased, and servers needed to talk to storage systems that were not directly attached to them, alternative methods were deployed to establish this connection. The dominant technology that majority of the installation, mostly larger enterprises, deployed was Fiber Channel. Briefly, Fiber Channel is a *fabric*, or a high speed serial interconnect that connects servers and storage devices and carries the SCSI traffic between these devices. While it was faster than the other available technologies of the time, Fiber Channel proved to be expensive and complex. Since it is a fabric and not a network, it has limited address management and security, furthermore it has no native routing capability. Nevertheless it has become a de facto standard in majority of the enterprises.

Fiber Channels issues, especially the expense and complexity associated with its acquisition, deployment and ongoing management has become a major limiting factor in bringing the benefits of Storage Area Networks to small and medium businesses around the globe.

iSCSI, Knight In Shining Armor

Fiber Channel's initial claim to fame was its Gigabit speed. Another method that became attractive and almost a 'no brainer,' as Ethernet speed reached Gigabit and beyond and prices plummeted, was to attach servers and storage devices to Ethernet networks and wrap the SCSI traffic to and from these devices in the familiar and ubiquitous IP packets. In using such familiar technologies, Storage Area Network no longer needed to stay the domain of a few. As its reach through the IP Internet became global, its appeal to the enterprises, large and small, became universal. Internet Engineering Task Force (IETF), the de facto standards body of Internet ratified the *iSCSI* (Internet Small Computer System Interface) *protocol* as a proposed standard in early 2003.

iSCSI protocol allows a seamless connection between servers and consolidated storage systems. Its biggest benefits come from the fact that it does not rely on exclusive skill set required for operation and management of complex fabrics, and in that it benefits from the economies of scale that Ethernet and IP provide. Hence it has become the most affordable technology in acquisition and operation of SANs.

iSCSI, being a block level IO protocol, allows applications such as databases that require block level access, to use networked storage systems. NAS or Networked Attached Storage devices export a *file* level storage (such as NFS, CIFS) to servers and hosts. iSCSI based storage solutions are different but complementary to NAS. NAS devices can now use iSCSI to access the storage they export as a file system.

As the need for storage increases by leaps and bounds for businesses (more than 30-35% projected annually), the need for consolidation and centralization of storage assets also increase. Linux, UNIX, Novel, and Windows based servers running vital business applications such as customer transaction processing, Email, Customer relationship Management (CRM), Supply Chain Management (SCM), and databases large and small can benefit today from the affordable storage consolidation enabled by iSCSI. Backup, Archiving, and Disaster Recovery is now affordable and simplified.

iSCSI components:

Storage servers along with application servers and hosts that are clients and consumers of the storage subsystem are connected via Ethernet infrastructures interconnected by switches and routers.

Servers and hosts side iSCSI *drivers* that initiate SCSI commands are called *initiators*. These drivers either use standard Gigabit Ethernet (GbE) NICs or use special HBA (Host Bus Adaptors) that embed the iSCSI drivers to connect to the network. Some host resident drivers use special adapters to offload the TCP/IP stack processing from the main CPU, known as TOEs (TCP Offload Engines). Use of HBAs and TOEs help the performance of the client side of iSCSI and can offer some features not available to the host based software (such as remote boot and hardware based multipath support) but adds to the cost. They should be considered if the server workload and performance requirements demand it and added cost is acceptable. Rule of thumb is that on a server with a GbE card and heavy disk I/O load, 500MHz of CPU processing power is taken up by iSCSI protocol processing. Today's hosts with greater than 2GHz-processing power should be able to easily handle such load.

Major operating systems vendors, such as Microsoft, have aggressively supported iSCSI by providing free initiator drivers for their respective operating systems.

Storage servers such as *Celeros EzSAN XR11* and *XR23* Appliances that implement the native iSCSI commands are called *iSCSI Targets*. Connection between initiators and targets are authenticated using CHAP (Challenge/Response Protocol). Use of IPsec Digests and AntiReply adds another layer of security to *Celeros* implementation.

iSCSI vs. FCIP vs. iFCP

In existing Fiber Channel installations, iSCSI should not be looked at as a competitor, but more of a compliment. iSCSI with its cost advantages can add affordable storage to the existing pool. With its native IP reach, it can be added to remote offices and regional sites. To overcome its distance limitations, Fiber Channel uses FCIP (Fiber Channel on IP), which is a point-to-point tunneling protocol. While it uses IP as a transport, it only connects islands of FCP SANs.

Another method of connecting FC hosts and devices is through Gateways that connect to each other through the IP network using a protocol called iFCP (Internet Fiber Channel Protocol). It allows using the existing investment in the Fiber Channel technology. These and other methods are complimentary to iSCSI. They attempt to preserve the existing investment in Fiber Channel products.

Summary:

Storage Area Networks, by pooling and centralization of servers and hosts from their direct attached storage bring the following benefits:

- Improved storage utilization
- Reduced operational costs
- Improved manageability of resources
- Improved security and access to vital data
- Improved utilization of skilled resources

iSCSI use of standard technologies brings further benefits to users, namely

- Reduces acquisition costs
- Reduces administrative costs
- Reduces need for specialized skill sets
- Has a global reach by using IP and Internet
- Reduces TCO

For customers who do not have a SAN and those who do, using iSCSI for all the benefits outlined above is, well, a no brainer!.

About Celeros

At Celeros our mission is to make *reliable, high performance* storage solutions that are *easy* to operate and *affordable*. While we do not have any religion with respect to the type of technology that can help our customers, we are ardent believers in choosing appropriate technologies that cost effectively solve today's problems and scale to address tomorrow's needs.

To learn more, please visit www.celeros.com or email us at info@celeros.com