



Mosaic IT Directors Series

iSCSI Basics: A Practical Introduction

iSCSI Basics: A Practical Introduction

Table of Contents

Overview	3
Internet SCSI	3
SCSI	4
Fibre Channel/SAN	4
NAS	4
How iSCSI Works	5
iSCSI in the Enterprise	6
Where iSCSI Fits In The Enterprise	7
Why iSCSI?	8
About Mosaic Technology	9

iSCSI Basics: A Practical Introduction

Overview

Over the last two decades, computers evolved from simple productivity tools into the core infrastructure of the entire global economy. That rise to prominence is reflected in the increasing focus of IT in all business processes and in evolving oversight regulatory mandates – which require stringent controls around data handling and access.

Business information and data repositories have become indispensable to the success of any business. Today, workers require information access in order to do their jobs effectively. As a result, the need to provide access to a dramatically increasing information/data set is an essential element of corporate planning and long-term success.

Storage sits at the center of information access. It has gotten faster, smaller, and -- on a gigabyte basis – cheaper. It has also evolved into a major management challenge in terms of building effective and efficient (cost and performance) storage systems.

-

In response to this challenge Technology evolved from SCSI to NASs and SANs. Now iSCSI offers an additional storage option. One that promises simplified management and implementation at reduced costs.

This guide is an introduction to iSCSI as a technology and an enterprise level storage solution.

Internet SCSI

Internet Protocol (IP) is the most widely used standard anywhere. The technology is well understood. It's easy to implement and is affordable. Most corporate data traffic uses a common IP network...except for storage data. Access to high performance storage data traditionally requires direct-attached devices or a Fibre Channel (FC) storage network.

Internet SCSI (iSCSI) transports traditional high performance “block-based” storage data over a common IP network. Which means - it can be used in remote mirroring, remote backup and similar applications since an IP network has no distance limitations.

As iSCSI begins to achieve widespread market adoption, barriers to implementing and managing networked storage can be removed by incorporating IP networking into a storage network.

SCSI

Most servers access storage devices through the Small Computer Systems Interface (SCSI) standard, moving “blocks” of data among computer systems. But its limitations became clear as demand for storage capacity grew.

SCSI’s built-in limitations on distance, number of devices supported and exclusive ownership of a server to its respective SCSI storage device prohibited the creation and sharing of a common pool of storage.

Fibre Channel/SAN

Creating a shared storage environment is best done with networked storage. Fibre Channel (FC) is the basis for building a Storage Area Network (SAN) and popular because it offers:

- high performance,
- support for longer distances than SCSI,
- ability to transport block-level data, and
- capability to scale to large network configurations.

Although no longer proprietary, FC based storage networks do have some baggage:

- complexity,
- cost
- difficult to install and manage.

NAS

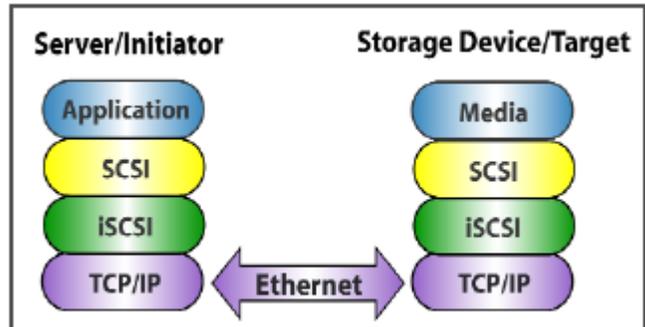
Another method of networking storage is to attach the storage device directly to an existing Ethernet network in a Network Attached Storage (NAS) configuration. NAS is typically a simpler and lower cost than a SAN. It leverages the benefits of using an IP network, such as support for significantly longer distances. However, NAS is slower than FC SAN. Data transport is managed at file level – not the block level. This significantly slows data access in transactional or database applications. Issues of scalability, centralization and management of a single pool of storage increase as the number of NAS systems deployed in a site increases.

A new solution was needed that offered the low cost and ease of deployment of NAS systems and the functionality of SAN. That’s where iSCSI plays.

How iSCSI Works

Internet SCSI (iSCSI) enables server host applications to perform traditional “block-level” transactions over a common IP network. It’s built on two of the most commonly understood protocols: SCSI and Ethernet, the dominant standards for storage and networking.

Using an ordinary IP network, iSCSI transports block-level data between an iSCSI initiator on a server and an iSCSI target on a storage device.



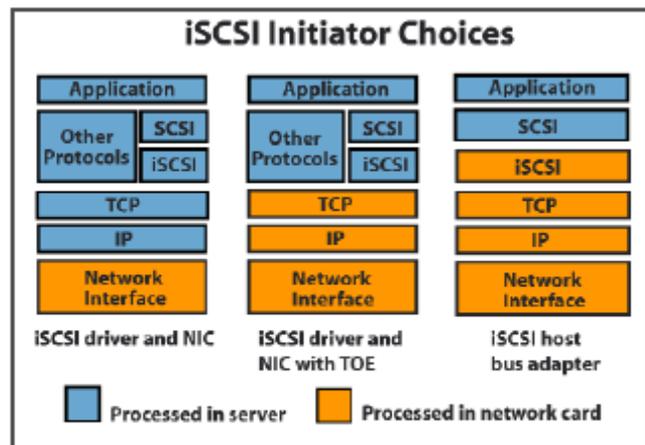
Communication between computing and storage platforms with iSCSI

The iSCSI protocol encapsulates SCSI commands and assembles the data in packets for the TCP/IP layer. Packets are sent over the network using a point-to-point connection. Upon arrival, the protocol translates data back to SCSI. Security is provided through iSCSI authentication and virtual private networks (VPNs), as needed.

When an iSCSI initiator connects to an iSCSI target, the operating system sees the storage as a local SCSI device that can be formatted as usual. The process is transparent to applications, file systems, and operating systems.

By consolidating storage with an iSCSI SAN, different platforms can share the same storage, greatly improving utilization and efficiency. Multi-protocol switches let iSCSI and Fibre Channel SANs co-exist

To access iSCSI storage, a server needs an iSCSI initiator connected to a network. An initiator can be an iSCSI driver with a standard network card, or a card with a TCP offload engine (TOE) to reduce CPU utilization. HBAs are available that offload both TCP and iSCSI.



iSCSI processing can be off-loaded to dedicated intelligent HBAs.

On the target side, storage devices also implement the iSCSI protocol stack.

With the advent of Gigabit Ethernet, iSCSI can deliver performance approaching a Fibre Channel SAN. Recent advances in 10GB Ethernet promise to make iSCSI a high performance solution capable of meeting and or exceeding other storage solutions performance levels.

The primary benefits of iSCSI leveraging IP are clear:

- no distance limitation
- lower cost
- well understood technology
- easy to administer and manage
- high availability (multiple connection paths)
- better performance than NAS (block level access instead of file level)
- maximum utilization of resources (share disk and tape devices across a heterogeneous environment).

FC SANs may be faster and have greater adoption and more robust management tools than iSCSI offers today. But, IP SANs are simpler and more affordable to implement than traditional FC SANs. And, with the 10 GbE, iSCSI on the horizon, IP SANs will benefit from the faster transfer rates. But iSCSI need not replace FC to fit in the enterprise.

iSCSI in the Enterprise

It would be prohibitively expensive to maintain all company data on instantly available and immediately recoverable storage systems. Enterprises routinely move less critical data to less available storage or archive it offline. But business and regulatory compliance requirements are beginning to demand rapid access to some of these archived records, which are no longer on immediately available storage.

Availability of data -- once seen as a business differentiator by providing better customer service or a greater ability to execute -- has moved into "must have" territory. On-demand data access is now a standard requirement of doing business. Add regulatory requirements mandating that large amounts of information be readily available and storage issues are compounded.

These business and compliance demands force companies to move data onto higher performance systems. Yet, redundant systems, backup and restore policies, snapshot or mirrored copies, and resources managing data are wasted if less critical data is moved to these systems. Moving volumes of data onto higher value storage is expensive and results in wasted resources.

Data classification seeks to overcome this potential weakness by aligning information with business needs. Categorizing data based on business needs and using the resulting classifications as a roadmap for retaining and storing information is a fundamental underpinning of information lifecycle management. Data classification aligns business requirements to infrastructure, so that infrastructure service delivery properly supports data storage and management.

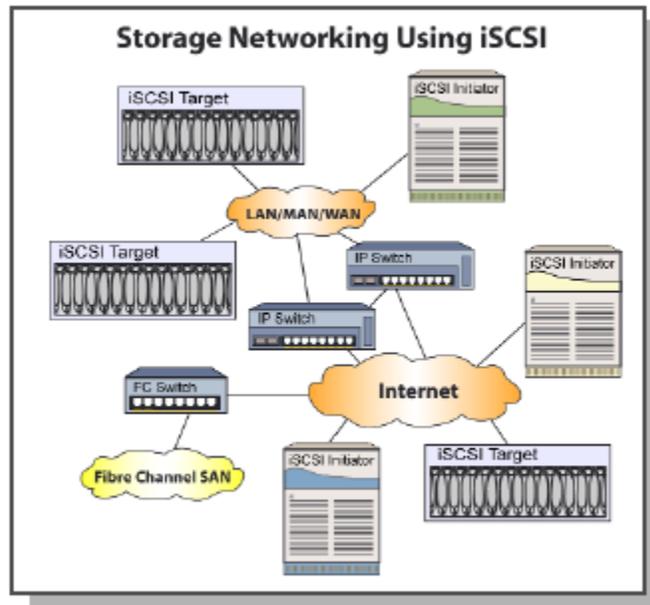
Data of more value should be supported with a higher service level of protection, such as mirroring, automated snapshots and remote replication for recovery. But the cost for this level of service is high. Less critical data can be archived to slower media, such as tape, at a much lower cost. With a hierarchy of storage devices in place, storage administrators can match the level of service available to the requirements of the business application.

Most enterprises are segmenting their server and storage systems into three or four categories, roughly; business-critical, business-essential, business-utility applications, and non-essential. Of course, these categories are arbitrary and which applications are critical, essential and utility varies by the enterprise.

Where iSCSI Fits In The Enterprise

On a price per port basis, connecting a server with Fibre Channel is high and limits the number of servers that attach to a SAN. For most IT budgets, storage and associated costs eats up one-third to one-half of total IT spending. A tiered-storage approach (matching relative value of data to associated storage cost) can substantially reduce costs.

Connecting high-availability servers can cost considerably less with iSCSI. This lets an enterprise increase the number of servers attached to a SAN. iSCSI is appealing for companies with remote sites. It brings stranded servers into the storage network, and consolidates backups. In a tiered-access strategy, iSCSI is a complementary technology to Fibre Channel.



iSCSI integrates into virtually any environment

iSCSI gives networked storage access to servers that might never have been connected, if it were not for the significantly lower cost of connecting to a SAN.

So, on one hand, an FC SAN is well suited for the high performance and functionality needed for business critical data, and can justify the expense to connect and maintain a select number of servers to a high performance storage network.

On the other hand, iSCSI is a good fit for all types of data with greater economies of scale, ease of use and the ability to bring more servers into the SAN than would otherwise be possible.

For the enterprise, it is no longer a question of either FC or iSCSI. They are enabling technologies that work hand-in-hand, as complementary technologies that fit together, in a tiered enterprise.

Why iSCSI?

The price of storage hardware has been decreasing. However, the unchecked demand for data and the associated increase in administrative costs work against all other factors in reducing the total cost of ownership.

The compound annual growth rate for storage is 68%. Administrative costs run seven times the cost of hardware. Until now midrange businesses could choose inefficient direct-attached storage (DAS) or invest in complex and often proprietary technologies. There's no guarantee that hardware will work together, and storage management remains cumbersome. iSCSI creates a standard for networked storage that brings the benefits of consolidated storage to a broad range of businesses.

Developed by the Internet Engineering Task Force (IETF) as a response to the need for interoperability in networked storage, iSCSI lets businesses leverage existing skills and network infrastructure to create IP-based SANs that deliver the performance of Fibre Channel, but at a fraction of the cost.

Benefits of iSCSI

- Makes consolidated storage possible for a wide range of businesses.
- Enables cost-effective, scalable, secure, and highly-available SANs.
- Leverages existing management skills and network infrastructure.
- Delivers performance comparable to Fibre Channel.
- Provides interoperability using industry standards.
- Implemented by the top system, storage, and network providers.

With iSCSI, businesses can get a handle on storage administration expenses without retrofitting their existing network infrastructure or investing in hardware that quickly becomes obsolete. Enabling low cost, interoperable, and high performance SANs, iSCSI is about to revolutionize the world of networked storage.

About Mosaic Technology

For over 10 years Mosaic Technology has been bringing an objective and unbiased view of technology to our customers. In the areas of storage, backup, data management, and server technology we offer a cross section of leading technologies. We know that one size does not fit all. We have long term relationships with companies such as EMC, Quantum, and Hitachi that enable us to give our customers solutions based on their specific needs rather than trying to force fit a product.

Headquartered in Salem, NH Mosaic has offices in Bellevue, WA, New York, and Connecticut. For more information call 603-898-5966. Or visit our website at www.mosaictec.com.