



IP Storage – an appealing solution?

Part 1

Which technologies will merge? What can be combined?

SAN projects don't always demand a clear decision between iSCSI and Fibre Channel; combinations are also possible. Despite the many differences, communication is possible. On the hardware side, a connection requires an IP storage switch or bridge, or a router.

The Internet Fibre Channel Protocol (iFCP) acts as a gateway between Fibre Channel and iSCSI because it provides NAT between FC and IP; that is, it maps FC addresses to IP addresses.

FCIP (Fibre Channel over IP) allows FC SANs to be connected over IP networks by tunneling. The FC packets are encapsulated in TCP/IP.

What are typical applications supported by iSCSI?

Primarily, business-critical enterprise applications at the department level, along with mission-critical applications of small and medium-sized businesses. The most widespread examples are mail servers, databases, and other applications that regularly access databases, such as eCommerce applications. iSCSI provides a powerful connection for server virtualization solutions.

Which customer problems does iSCSI solve?

iSCSI is mainly a good way to counter the uncontrolled proliferation of DAS servers and SAN islands in a company. The most important point in iSCSI's favor is the cost savings potential arising from centralized management. However, iSCSI also offers all the other advantages of a SAN (better storage utilization, smaller backup windows, cluster support, snapshots for SQL and Exchange servers, long ranges) and costs much less than Fibre Channel.

Is this new technology ready yet for the tough conditions of real-world customer use?

Absolutely; just think of consolidation, cost-efficient replication, easy administration, and the familiar TCP/IP environment. The decisive factor for its use in companies is probably the cost, along with the fact that differences in performance compared to FC are steadily decreasing.

What does an iSCSI solution consist of?

You need two iSCSI endpoints: One end is the iSCSI target, the other is the iSCSI initiator. Initiators are available for pretty much every major operating system currently in use. Alternatively, you can use an iSCSI HBA as an initiator. You also need Standard Cat5E or Cat6 Ethernet cabling, and Gigabit Ethernet network cards. Standard Ethernet switches are entirely sufficient.



Is it true that a lot of additional hardware is required to get iSCSI running? A cost factor!

In fact, just the opposite is true. With iSCSI installations you can use a lot of existing Ethernet infrastructure, such as standard network cards or existing server systems. iSCSI initiators are available for free, for example from Microsoft, or even as Open Source software in the form of Open iSCSI. Additionally required components include, for example, hard disk systems or tape drives, but those aren't needed for iSCSI itself, but for the intended functionality such as for backups. Compared to DAS systems or Fibre Channel SANs, iSCSI is a financially attractive solution.

What kinds of iSCSI products are currently on the market?

There is an abundance of products and solutions on the market that are based on different drafts of the iSCSI specification. The most important ones are iSCSI targets, such as hard disk arrays, storage routers, storage switches and iSCSI initiators integrated into network adapters, and specific host bus adapters.

Are pure IP SANs (iSCSI SANs) or FCIP (Fibre Channel over IP) already being implemented?

Yes, businesses today usually choose pure IP SANs with iSCSI in order to stay in a familiar, controllable environment. These SANs are often created in the course of consolidation projects, where DAS systems are replaced. FCIP is used in large companies, for example when FC SANs at different locations are connected via IP networks.

Which new features/characteristics will drive the adaptation of iSCSI this year?

For many businesses, virtualization is a very high priority; iSCSI provides a suitable and cost-effective solution for that. In addition, most companies are trying to optimize their backups, which also argues for greater iSCSI acceptance.

How does iSCSI relate to IP storage protocols such as iFCP (Internet Fibre Channel Protocol)?

Technically, iSCSI and iFCP are not very far apart. The principle of iSCSI is that SCSI commands are embedded in IP packets. iFCP works very similarly - at its core, SCSI commands are at work, which are transmitted over an IP network. The difference is at Layer Four of the protocol stack. With iFCP, this layer of the Fibre Channel stack is retained, allowing the use of management applications for Fibre Channel networks with iFCP.

iFCP and also FCIP are used to connect SANs to each other. In contrast to iFCP, FCIP creates a tunnel between FC SANs.

What kinds of storage data traffic are transmitted via iSCSI?

iSCSI offers block-based data transmission, for example for backups, databases, snapshots, remote boot, virtualization, disaster recovery, mirroring, and much more.

What factors are driving the transition from Direct Attached Storage systems to SANs? What is iSCSI's part in this?

Easy management, better utilization of storage capacity, high availability and better scalability are the main arguments against DAS systems. Financially, iSCSI SANs offer better TCO and faster ROI. DAS systems have usually developed over time; they are time-consuming to back up, complicated to manage, and not particularly reliable. High hardware and software costs are caused, for example, by a large number of local backups. Centralized backup solutions usually put too much pressure on the network and are not a particularly viable solution for 24x7 applications. iSCSI solves these problems and eliminates Single Points of Failure through features such as mirroring.

Fibre Channel is already designed to transmit SCSI commands via switched network fabrics.

So why should I bother with iSCSI?

If you prefer a unified network infrastructure and want to save money, there's no way around iSCSI. Another advantage is that it's easy to administrate with familiar management tools.

Is it true that IP SANs require you to take a very close look at the issue of "security"?

Of course, after all we're talking about storing valuable company data. On the one hand, data security is important; that is, ensuring backups, replication, snapshots etc... On the other hand, access to this data must also be strictly regulated. The necessary functionalities are, of course, available for iSCSI, for example IPsec, VPNs, CHAP, VLANs...





Part 2

What is an iSCSI initiator?

An initiator is an iSCSI client. Initiators typically have the same tasks as a SCSI host bus adapter. The difference is that iSCSI initiators don't connect SCSI devices (such as hard disks or tape changers) physically via cables. iSCSI initiators send commands over IP networks. There are two kinds of initiators: software and hardware iSCSI initiators.

What is the aim of iSCSI?

iSCSI is the industry standard for data transfer over Ethernet-based networks in IP-based storage networks. iSCSI systems use the existing Ethernet infrastructure to transfer SCSI commands over IP. Because IP-based networks are available everywhere, iSCSI can be used to transfer data in LANs or WANs, enabling location-independent data storage and recovery.

Which is preferable?

A hardware or a software initiator?

The advantage of a hardware initiator, or iSCSI HBA, is that the IP and iSCSI stack are offloaded at the hardware level. Therefore, no other hardware resources are required. In light of the capacities of modern hardware at affordable prices, however, it no longer makes much difference whether a hardware or a software initiator is used. A solution with a TCP Offload Engine (TOE) is therefore not necessarily recommended, as experience has shown that disadvantages, such as stack rewriting, may actually result in performance loss.

How do hardware and software initiators affect the performance of applications?

If an application accesses a large number of small files frequently (e.g. a heavily utilized OLTP database), this may result in the iSCSI initiator using up to 25% of CPU capacity. In such cases a hardware iSCSI initiator (or iSCSI HBA) would possibly become necessary, although a TNIC might still suffice.

How does iSCSI interoperability compare to Fibre Channel interoperability?

Both technologies are so mature that interoperability between most manufacturers has already been verified or certified.

What is the difference between iSCSI and Fibre Channel?

iSCSI is the transfer of SCSI commands over IP (Internet Protocol) networks, whereas Fibre Channel is the transfer of network and storage protocols (e.g. SCSI or IP) over fibre-optic or copper cables. Therefore iSCSI uses established network infrastructures to create an IP SAN, whereas with Fibre Channel, additional infrastructure must be purchased in order to create a SAN.

Will iSCSI replace Fibre Channel, in your opinion?

No, Fibre Channel offers a very solid basis for secure and high-performance storage solutions. iSCSI still suffers from the weakness that IP packets may be lost on the network and then have to be resent. If a network is overloaded, this may cause significant drops in performance.

When should I use iSCSI instead of Fibre Channel SAN?

You should perhaps be asking the opposite question: When should I use Fibre Channel instead of iSCSI? You should use Fibre Channel if the requirements of your IT systems cannot be covered by iSCSI, for example if you have very high data traffic. Due to the FC protocol, FC offers lower data packet loss than iSCSI, which uses TCP/IP as a transfer protocol.

How do the characteristics of iSCSI storage arrays compare to the characteristics of Fibre Channel arrays?

Apart from the connection type, the characteristics of iSCSI and Fibre storage arrays are comparable. What's interesting is to take a look at the features offered by the various manufacturers and array models: replication, backup/recovery, disaster recovery, redundancy, failover, price and performance.

How does iSCSI compare to Fibre Channel in terms of security?

Contrary to general opinion, FC SANs are no more secure than iSCSI SANs. Rather, FC installations have benefited indirectly from their complexity and their relative isolation in data centers.

When comparing the two, it's worth noting that iSCSI supports a pretty extensive list of security features, including Access Control Lists (ACLs), the IP Security Protocol (IPSec), the Challenge Handshake Authentication Protocol (CHAP) and the use of Virtual Private Networks (VPNs). Anytime an iSCSI SAN extends beyond the data center, the storage administrator must take the necessary steps to ensure secure access.

The trick when it comes to iSCSI security is not to find the right tools for the job, but to use the best methods and make the most of the tools that already exist.

At the moment, we're told, the cost of IP storage is far lower than that of Fibre Channel. Will this cost advantage continue to apply in the future?

From what we can expect, a TCP/IP infrastructure will always be more affordable than an infrastructure that involves significant additional costs for the Fibre Channel technology. It's unlikely that this will change.



Let's assume the difference in price diminishes or disappears. Would Fibre Channel then remain the preferred solution?

iSCSI solutions would still be interesting for small and medium-sized businesses, because implementing iSCSI solutions requires no specific knowledge, and in most cases no new hardware.

What about iSCSI backups?

iSCSI can be used to perform backups over the network. Typically, an iSCSI target is created, where a snapshot of the block device can be seen, and this block device is integrated into a backup server as write-protected. In this regard it should be noted that in order to perform the backup, the backup server must be able to read the iSCSI block device as a formatted disk with a file system (e.g. NTFS, XFS, EXT3).

How about IP storage and disaster recovery?

Disaster recovery: organizations mirror storage resources from one data center to another, remote data center, where they can serve as a "hot standby" in the event of a prolonged outage. Specifically, iSCSI SANs allow you to migrate entire disk arrays over a WAN, with minimal configuration changes, so that storage traffic can be routed in the same way as network traffic.

What's the difference between iSCSI software on a Gigabit Network Interface Card (NIC) and using iSCSI HBAs or TNICs?

iSCSI-based storage networks basically have three kinds of network interfaces:

- Simple network cards (software endpoints)

Simple NICs are the cheapest option for iSCSI; however, with these, the IP protocol is processed on the host CPU, which may cause high interrupt figures, and therefore additional CPU utilization, at higher I/O loads.

- TCP/IP Offload Engines (TOEs / TNICS) (software with hardware-supported endpoints)

In order to reduce the load on the host CPU and cut latency periods, TOE cards have been developed which "calculate" the TCP/IP protocol by means of hardware on the card. These cards allow all TCP/IP traffic to be offloaded from the host CPU; they also support other protocols besides iSCSI. Operating system support is, however, required for offloading normal TCP/IP traffic, and this is not always available.

- Lastly, iSCSI host bus adapters (hardware endpoints)

Offloading only the iSCSI protocol is less critical. In this scenario, TCP/IP and iSCSI are run as hardware implementations, keeping the load entirely off the CPU. Only these iSCSI HBAs are comparable to Fibre Channel HBAs in terms of CPU load and latency. iSCSI HBAs are also comparable to standard SCSI adapters in their functionality towards the server. HBA capacity should be taken into consideration when choosing an iSCSI HBA.

What does iSCSI API mean?

iSCSI API is really called iSCSI Management API, and is a standardized interface that applications use in order to manage iSCSI HBAs, regardless of manufacturer

What does IPsec mean?

IPsec is an abbreviation of Internet Protocol Security. It provides a security architecture for communication over IP computer networks. This protection consists of ensuring confidentiality, authenticity and integrity.



Janusz Bak

Chief Technical Officer bei Open-E

answers competently all questions concerning virtualization with iSCSI