



iSCSI Questionnaire

Coming on the heels of the release of last year's SCSI over IP (iSCSI) standard from the Internet Engineering Task Force, 2004 is seeing a modest increase in the number of iSCSI products entering the market. Questions abound over the fit for this protocol and for storage topologies based on it.

The following questions are intended to identify your company's view of the market for iSCSI technology.

QUESTION 1

In its early development years, iSCSI had several prominent champions within the vendor community, including IBM and Cisco Systems. The early position of iSCSI advocates was that it would replace Fibre Channel as an interconnect for building storage area networks. With the delays in standards development, the party line seemed to change: FC would be used to build "core" fabrics, while iSCSI would be used to connect outlying servers to FC fabrics.

What is your position on the technical fit for the burgeoning technology?

While we see iSCSI being used to connect "stranded servers" on the IP network, we still find interest in using iSCSI in SMB storage. The reason is the desire for Ethernet solutions running on the existing Ethernet network, along with the perception that if it's on Ethernet it will interoperate and be easy to configure (rightly or wrongly, this is the perception)

One thing is still inhibiting adoption of iSCSI SANs – the lack of bundled solutions that include the storage, storage management and either HBAs or software.

QUESTION 2

As an IP-based protocol, iSCSI is limited in terms of speeds to available bandwidth less overhead, which is generally interpreted to mean that the technology is capable of delivering roughly 75 percent of the rated speed of the

TCP/IP network pipe in Mb/s or Gb/s. FC advocates have leveraged this as a major differentiator between FCP and iSCSI solutions.

How meaningful is this speed difference today? How meaningful will it be next year with the introduction of 10 GB/s IP nets?

Certainly today with FC running at 2G and iSCSI at 1G there is a meaningful speed difference, and this is about to become even more significant with FC going to 4G later this year. And true, as iSCSI runs over TCP/IP, there is more protocol overhead than FCP. However, the fact that iSCSI runs over TCP/IP allows the traffic to traverse long distances and over many types of networks. You have to start with what problem you are solving – if connecting remote servers over distance, the limit of the WAN connection may make the comparison moot – FC couldn't be used anyhow.

In SMB SANs, iSCSI speed may not be an issue either. One of the applications cited as appropriate for iSCSI is Microsoft Exchange. In a small network the slower speed of iSCSI may not be an issue. But without a doubt, FC costs are going down, bundled solutions ensure interoperability, and the performance of FC can't be beat – it's a well proven technology for high performance storage networking.

10G Ethernet pricing is coming down after shipping for two years already. By 2006, 10G switch ports for both fiber and copper (CX4) will be dramatically cheaper than today. FC's next step is to 8G in late '06 or '07, not 10G. So by 2006 iSCSI on 10G should have a better story to tell than FC.

QUESTION 3

Related to the above, how important is interconnect speed to applications? Haven't we made do with much slower storage interconnects in the recent past?

For remote servers, the WAN link, often a T1, will be the bottleneck. Software iSCSI will have plenty of performance to operate over WAN links. In SMB, many applications will not be as demanding as in the large enterprise. Each application will have to be assessed as to its speed requirements. Interactive DataBase apps do require low latency due to large numbers of interactions. Email can make do with higher latency as milliseconds of extra time to deliver messages is usually not an issue. Trading floor archiving requires high throughput to complete before next day trading resumes.

QUESTION 4

Both FC fabrics and iSCSI SANs utilize IP-based applications for management. In the case of iSCSI, management (or control path) is handled in the same network pipe as data and SCSI command traffic. In FCP, the control path and data path use different wires.

From the standpoint of scaling, simplified infrastructure, and design elegance, iSCSI would seem to have the advantage over Fibre Channel's "dual network" design. What do you think?

Even in many IP only networks, there is often a dual network, with management traffic using out of band Ethernet management ports even though in band management is possible. Servers will continue to ship with a management port as it is unknown what kind of storage attach HBA will be added later on, so this back door will exist for some time. IT folks like redundancy and back door fail back scenarios. So the philosophy of each IT group will weigh in both ways here. A wash I believe.

QUESTION 5

Both iSCSI and Fibre Channel use a serialization of SCSI, a channel protocol for storage I/O. The key technical difference is the transport used by each interconnect (TCP for iSCSI, FCP for FC fabrics).

If the two are more similar than dissimilar, why should a company field separate channel interconnect rather than use existing investments in networks to interconnect storage and servers?

We do not agree that they are similar. Flow control and handling of dropped frames differ considerably, both operate on different layers etc.

FC gained acceptance for storage traffic due to it's low latency, and was ahead of Ethernet in providing 1G link speeds. FC continues to outpace Ethernet with 2G and 4G link speeds. With bounded distances and topologies, FC gets away without the complication and workload of a TCP type layer. This has been the reason. Things could now change with Ethernet at 10G now and FC going to 8G, but not for a few years. Plus iSCSI over RDMA (iSER) with 10G gives iSCSI FC latency and performance.

QUESTION 6

FC SANs are increasingly seen behind NAS heads, which are said to act as gateways to SANs and provide hosting for SAN management utilities.

Taking this design choice to the next level, what is your opinion about using NAS gateways to support both NFS/CIFS and iSCSI on the front end in order to aggregate storage traffic?

We already see this happening. Network Appliance has offered an iSCSI option on their NFS/CIFS filers for over one year now. The Quantum SNAP NAS box also offers an iSCSI option. For both of these we still see the backend being FC for SAN connection, but we could see iSCSI on the back end soon. For the SMB,

this makes very good sense to have a single appliance aggregating storage traffic – keeps things simple.

QUESTION 7

iSCSI standards do not seem to have been “held hostage” to proprietary vendor interests the way that FCP standards have been at ANSI (it is an established fact that vendors can develop FC switches that fully comply with ANSI standards, yet fail to be compatible with one another).

From the consumer’s perspective, do you feel it’s smarter to go with iSCSI-based technologies because of product interoperability?

Switches from different vendors can usually establish links with each other, it’s the value added features at the protocol level that causes trouble.

The FC world has mitigated this issue with the extensive interoperability testing and “recipes” that are offered. iSCSI devices will have to undergo similar interoperability testing and recipes. It would seem that trying to mix different vendors switches in the same network won’t make the effort easier than sticking with a single vendor, even if there were no compatibility issues. Familiarity with a vendor’s UI and operational issues would dictate sticking with one vendor for the fabric. So I would answer no to the above.

QUESTION 8

At one point, vendors touted iSCSI as the foundational technology for building “SANs for the rest of us” – that is, companies that are not necessarily Fortune 500 status.

Do you embrace this view? **Partially – there are several options for SMB storage.**
And if so:

- What do “the rest of us” require a SAN for? What is the killer application for iSCSI SANs? **SANS are for low latency block access. For many SMBs, MS Exchange may be the only application requiring block access. Many SMBs don’t have a block storage problem at all – NAS is really all they need.**
- What is the advantage of iSCSI over burgeoning protocols for large-scale device interconnection like Serial Attached SCSI (SAS), which, with expanders, offers connectivity for up to 16,000 nodes? **This is somewhat mixing apples and oranges, as iSCSI is a transport protocol, and SAS is a drive technology. But having said that, it is possible with SAS expanders to build a SAN out of SAS technology. The links would run at 3G vs. iSCSI at 1G today. The iSCSI advantage would come if distance were a consideration – iSCSI can go anywhere over the Internet. SAS will be limited to local topologies only.**

- With burgeoning drive capacity improvements, already at 200 GB for SATA and SCSI, can arrays be built with adequate capacity to meet the needs of SMBs without resorting to SANs? *As you make drives larger in capacity you can then create even larger DAS islands. Yes, it could buy you time in deploying a SAN – but in the meantime you lose the SAN advantage of data sharing, ease of backup, ease of drive replacement.*
- With removable/exchangeable disk/tape hybrids, such as Spectra Logic's RXT platforms, can SMBs achieve capacity scaling requirements without deploying SANs at all? *Yes – but how inconvenient do you want to make DAS solutions to avoid SANs?*

QUESTION 9

What has happened to TCP Offload Engine (TOE) technology, once touted as a prerequisite for iSCSI SANs? Was it simply hype intended to keep Host Bus Adapter vendors from losing market share to vendors of simple NIC cards in an iSCSI world? Or, has TOE development proved more daunting than originally thought? Why aren't we hearing more about TOE?

Nothing happened to TOE technology - as far as I know, all the iSCSI HBAs in the market contain TOE engines. As TOE engines have become required for iSCSI offload HBAs if they are to provide the low latency, speed and reliability required of block storage. I believe vendors have stopped touting them – they have become a “Jacks or better” requirement. One other thought is a card with only TOE functionality has not proved popular, so maybe the TOE hype has died down.

QUESTION 10

FC fabric advocates claim that FC fabrics are more secure than iSCSI SANs. What do you think?

- How is an FC fabric any more secure than an IP-based iSCSI SAN if it uses an out-of-band, IP-based, connection for fabric management? *FC fabrics are perceived to be more secure because they are isolated from the data network. iSCSI SANs can be built the same way – with separate infrastructure from the data network. The IP-based fabric management is out of band from the data, so should not put data at risk.*
- How can FC advocates justify the claim that FCP remains a mystery to hackers, but also argue that the protocol is becoming more familiar and less of a training hurdle for customers? *What's becoming more familiar is set up of FC infrastructure, not the protocol itself. Again – proper data center firewalls and other protection can make either FC or iSCSI SANs secure.*
- Why have no FC switch vendors implemented the FCP security standards from ANSI in their products? *I think due to the perception that FC*

infrastructure is isolated from the data network, and hence, protected from security risks.

QUESTION 11

Microsoft's iSCSI initiator seems to be winning mindshare among vendors (Cisco recently opted to use the Microsoft initiator in place of its own in Windows shops).

Do you support the Microsoft iSCSI initiator with your products? **Yes** Does a target device also need to utilize Microsoft target definitions to work with a Microsoft initiator? (Microsoft says it does, some target vendors say it doesn't.) **No it doesn't. We have built Targets that do not support Microsoft Target definitions and they work with the Microsoft Initiator just fine.**

QUESTION 12

Some vendors seem to be suggesting that Fibre Channel is superior to iSCSI because of its end-to-end support of "native Fibre Channel drives."

Is there such a thing as a "native Fibre Channel drive" or are we really talking about SCSI drives with integral Fibre Channel to SCSI bridges in the electronics of the controller or disk? **All FC drives are "native FC drives." But that does not give FC an advantage end to end. Both FC and iSCSI transport SCSI commands and data.**

QUESTION 13

Fibre Channel fabrics do not seem to respond to Metcalfe's Law of networks, which states that the value of a network should increase and cost per node should decrease as more nodes are deployed. Fibre Channel fabrics seem, in fact, to become more difficult to manage as they scale (in many cases eliminating many of the value gains promised by vendors) and, in general, remain the most expensive platform for data storage. FC fabric per port costs have been extremely slow to decline.

By contrast, per port costs of GigE switches and GigE NICs have fallen dramatically in only a two to three year time frame. 10GbE is expected to follow this pattern as well.

From a cost standpoint, does iSCSI have a better story to tell than Fibre Channel to price-sensitive consumers?

If we look at today's situation, one has to run dual port 1G iSCSI to match 2G FC speeds. When you do the math, it's hard to see an advantage for iSCSI. There are low cost FC switches appearing now that are on parity with 2 Ethernet switch ports. One thing to point out here is that there are different classes of switches. To really compare to FC Director class switches, you would have to compare to Ethernet high switches supporting traffic shaping, MPLS, priority queueing etc.

This may change down the road, as 10G Ethernet pricing is coming down after shipping for two years already. By 2006, 10G switch ports for both fiber and copper (CX4) will be dramatically cheaper than today. FC's next step is to 8G in late '06 or '07, not 10G. So by 2006 iSCSI on 10G should have a better story to tell than FC.

QUESTION 14

The industry has given mixed messages about the fit for iSCSI: Is it a data center technology because that is where the big switches are located, or is it an "edge technology" because workgroups and departments do not require the speeds and feeds of data centers? What is your take?

I think it is both – but the time frames are different. In the short term iSCSI is more effective as an edge technology due to its ability to provide connectivity to a FC SAN, and ability to run as software Initiator at the edge, making it low cost. Not until 10G Ethernet comes down in cost by 2006, will iSCSI have a good fit in the data center. Today, data center choices are made by performance, and FC wins now.

There is another aspect of this that bears discussion here. Blade servers and compute clusters today have multiple fabrics – Ethernet and FC to handle block and file level traffic. At 10G, there could be a single Ethernet fabric supporting both traffic types – with HBAs supporting Block (iSCSI) and File level traffic, sometimes referred to as a Unified Wire HBA. This will make clusters much cheaper and less complex. This is a trick that FC will not be able to do – FC doesn't support file level traffic. In the long run Ethernet will displace FC when 10G Ethernet reaches cost parity with 8G FC and RDMA/iSCSI solutions are available.

QUESTION 15

With Simple Network Management Protocol (SNMP), Dynamic Host Communications Protocol (DHCP), and other established protocols in the IP world, it would seem that iSCSI will hit the ground running with services that were missing altogether from FCP. Is this an advantage in your opinion?

No. The problem to be solved in SANs is discovery of Targets. SNMP and DHCP do not solve discovery. New protocols such as iSNS are required for discovery or SMI. Both iSCSI and FC have the same issue here.

QUESTION 16

Some vendors are "dumbing down" their Fibre Channel products to facilitate their deployment in SMBs. Is this your strategy and what do you see as the benefits and drawbacks of such an effort?

No, this is not our strategy. I only see drawbacks in doing this. SANs only grow over time, so having a full featured, full performance HBA for the SMB will allow this growth easily. I am also skeptical about the new SMB offerings as I notice the same ASIC is on these SMB HBAs as on the enterprise versions – with dumbing down of the product – exactly where are the cost savings coming from?

QUESTION 17

Does iSCSI offer anything that FC fabrics do not to facilitate storage virtualization?

No. Virtualization concatenates multiple targets into a virtual target – neither transport offers an advantage over the other here.

QUESTION 18

Describe the products that your company is developing that support iSCSI.

LSI Logic does not see much of a market for 1G iSCSI HBAs – software provides adequate performance and CPU load based on testing of the Microsoft Initiator and Linux iSCSI Initiators in our lab.

We are aiming for the 10G space instead – supporting multiprotocols – RDMA, TOE and iSCSI. This is the earlier referred to Unified Wire HBA. We are building two versions of HBAs with the technology – one with 4 X 1G ports and another with a single 10G port.

QUESTION 19

Compare key pricing and capability differences for your iSCSI solutions versus comparable FC solutions.

We have not priced either HBA yet – but some general comments can be made. For a given link speed, a FC and an iSCSI HBA have the same optics transceiver at the same cost. This is a big component of the BOM of the HBA. The other big item is the ASIC itself. What's interesting here is the FC chip can be thought of as sort of equivalent to an Ethernet MAC in terms of functionality and the networking layers supported. On the other hand, iSCSI chips have to have a TCP/IP offload engine, and the HBA has to have more buffer memory than FC HBAs do. This is due to the need of iSCSI to deal with out of order packets and buffer and re-ordering. So, fundamentally iSCSI HBAs have more complexity and cost in them than FC HBAs do. Market pressure and expectations that Ethernet solutions must be cheap has caused margins on iSCSI HBAs to be lower than FC. The real cost savings for iSCSI comes at the switch port end of the connection, where Ethernet switches are much lower cost than FC switches, as long as you compare like levels of switches.

Soon, iSCSI may not be burdened with the buffering problem. iSCSI over RDMA (iSER) will allow the direct placement of data without buffering, similar to the way FC operates. But for now, a lot of buffer memory is required on iSCSI HBAs.

QUESTION 20

Does iSCSI contribute to data protection in a networked storage world? If so, what?

Yes. iSCSI has IPsec and VPNs. FC relies on physical network isolation.