

## NETAPP SUBMISSION – July 16, 2004

### 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?

First of all, if this is the party line, we haven't been invited to the party, nor would we go if we were invited.

Our party line is that there's a place for FCP and a place for "native" iSCSI. That hybrid/bridge architectures only add to the complexity of designing, building and managing SANs.

Our customers tell us that the place for FCP is indeed in core fabrics and, more directly, datacenters, to connect relatively small numbers of larger servers to a high performance SAN. iSCSI fits where there's no existing FCP SAN, or where the cost of FCP interconnect, or complexity of design and management (and lack of a "local" skill-set), makes FCP impractical. In many ways, the performance issue is more perception than reality (see below) but at this point, most customers still see FCP as "higher performance".

Interestingly, the whole core vs. edge debate will begin again when 10Gig IP becomes pervasive.

#### 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?

One needs to recognize that the impact of transport speed differences increase proportionately with block size. Another impact point is that fibre channel protocol is a "slimmer" protocol and so utilizes less CPU cycles. The two need to be viewed in concert. Small block transfers which place the bulk of the burden on the protocol itself result in near equal performance. Our experience based on actual deployment examples highlight the fact that actual difference in performance between FC and iSCSI is often an academic concern. Many applications currently leveraging iSCSI themselves are the bottlenecks that gate the overall utilization of the storage pipe and so storage performance would be the same regardless of transport protocol.

We expect that initial 10Gbe deployments will roll out in a similar fashion as we've seen with faster FC protocols (1-2Gb, 2-4Gb). Due to the higher cost, these deployments will be limited initially to network interconnects and only those applications which can benefit from or require higher transport speeds such as those which utilize large block sizes (video, etc.) as well.

Another factor limiting 10GbE deployments will be the availability of 10GbE capable infrastructure - many organizations still have limited 1GbE capable infrastructure and are satisfied with iSCSI performance over 10/100 Mb infrastructure.

### 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?

As noted in Question 2, we've done lots of lab testing with hardware and software initiators with many applications. We find that the interconnect speed only matters when apps reach "line speed." Since the great majority of our production installs are for apps which \*never\* reach line speed (MSFT Exchange, SQL, SAP, Novell Netware) this is something a customer just will not see. Combine this with the fact that most of the servers running these apps are Intel based 2-ways, 4 ways and 8 ways and the issue of maxing out interconnect speed becomes even less relevant at the user level.

### 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?

We agree, but do note that FC does have provisions for in-band management and does provide a premise for more secure, less complex (i.e. the simplicity that comes with operating a single protocol) but integration with Enterprise network management applications seems to be a potential point of compromise.

## 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?

Enterprise customers ask us this question all the time. It points out one of the \*revealed truths\* of FCP SAN and IP SANS--that the choice of interconnect is often a financial/infrastructure decision, and this will dictate the protocol of the storage endpoints. The larger the organization, the more the issue of infrastructure change and ROI comes into play. Infrastructure related to both network hardware and organization processes and skill-sets.

In nirvana, if one \*transport\* were available that met all of an organization's needs at a given point in time that why would anyone want to try to manage two networks? For smaller organizations, or for the distributed parts of larger enterprises, it's much easier to get to nirvana. These places are also where you're less likely to find FCP SANS and also almost guaranteed to find a sys-admin or network admin that knows IP. For this reason, we've seen a lot of early success in deploying Unified IP Storage (CIFS or NFS plus iSCSI) in this segment of the market. Only one interconnect (and the skill to manage it) is necessary, and the user gets all the benefits of SAN and NAS without the cost of 2 networks.

By the way, we are at a loss to understand the recent announcement regarding FCP SANS for "small businesses." While it's possible to process an Internet order and then ship a configured FCP SAN from a factory to a user that fits that users' need at the time of shipment, what happens when the user's environment changes or grows? Do they need to go for FCP SAN training so they can learn about zoning, etc, or do they then need to call the vendors service department so they can send out the "guy in the white coat?" If the assumption is that these SANs will be self-serviced without training, the market for this class of FCP system might be limited to the same folks who build and fly ultra-light airplanes.

## 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 view iSCSI and CIFS / NFS (NAS) to FC gateways as being appropriate when re-purposing older FC storage that has reached the end of its economic life or which cannot meet the current performance requirements of the applications to which it is currently attached. It's also a good way to improve the utilization of an older storage system that was put in place with excess capacity, and has never reach, or never will reach a reasonable payback level. This combination can be use to meet the requirements of tier-2 applications. For tier-1 apps, customer are now purchasing storage solutions that provide native NAS and iSCSI capabilities to reduce the overhead impact and management complexity associated with protocol conversion.

## 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?

Yes.

FCP interoperability has gotten much better in the 8 years that have elapsed since vendors started marketing \*multi-vendor SANs \*. That said, the fact the IP SANS were interoperable out of the box has made it much easier to design reliable IP SANs that do useful things. It has possible to spend development dollars on features that solve customer problems, rather than constructing HUGE interoperability labs to allow the same engineers to get devices to talk to each other. We think that this accounts for the very rapid ramp we've seen in production deployments of iSCSI post standards "ratification" vs. the relatively long incubation period that FCP had after the standard was set.

## 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?

Our experience is that iSCSI deployments are not limited to the size of the company but instead by the current SAN deployments (See Q5), or other limitations including distance (i.e. outside the FC distance limit) and cost (i.e. is it sensible to spend more than the cost of the server itself to connect it to a storage network as observed in the Wintel / Lintel/Novell server base?). Many Fortune 500 companies have distributed offices which can definitely benefit from iSCSI as iSCSI will enable them to extend "IT literacy" to remote offices in more imaginative ways. An example would be to enable and monitor such mundane but essential tasks as local backup / restore, while also enabling disaster recovery-oriented data protection and regulatory compliance by enabling copies of distributed information to be cost-effectively gathered at a centralized facility for subsequent processing.

And if so:

- What do “the rest of us” require a SAN for? What is the killer application for iSCSI SANs?

First and foremost, it’s important to note that iSCSI enables a broader set of servers to cost-effectively enjoy the typical benefits of storage area networking i.e. easier management, higher performance, etc. That being stated, there a several killer apps" - remote backup of distributed servers and data; remote boot enabling rapid re-deployment of servers, and economical storage replication based disaster recovery.

- 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?

SAS doesn't leverage the ubiquity and distance which are the true enablers of broad-based deployments.

- 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?

Yes but this needs to be viewed as a specific market segment. Building larger arrays places more data at simultaneous risk to the universal single point of failure - the storage server OS.

- 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 again, this needs to be viewed as a specific market segment. Building larger arrays places more data at simultaneous risk to the universal single point of failure - the storage server OS. Don't put all your eggs in one basket!

## 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?

More than 90% of our production deployments are made with software initiators. There are a few reasons for this.

Software initiators for Windows and Netware and SuSE Linux are OS vendor provided and supported --and "FREE". So the ROI is infinite! They are easy to try and easy to deploy and many of the machines they get deployed on have plenty of CPU headroom.

Over time, as larger, more IO intensive applications are moved to IP SANs, and the servers they run on have less and less CPU headroom, we'll see more TOE/HBA deployments. Right now, in the early ramp for the technology, software initiators are free, they work, and they are directly supported by the OS vendors.

Also, in order to provide iSCSI net boot, the only option today is to use an HBA. Since this is one of the killer apps for IP SANs, early adopters who see the huge win that IP SAN net boot provides are also investing in hardware initiators.

## 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?

Let's break storage security down into its two major constituent parts - access security and data integrity security. There is a misinformed perception that FC SANs are more secure than IP SANs as breaching FC SANs will require an invasive "tap" that isn't likely to occur due to the inherent physical security of FC networks i.e. they're deployed within data centers and don't cross firewalls. This perception is void in bridged inter-data center networks which traverse LAN/MAN /WAN infrastructures (where the killer apps for iSCSI reside).

The IP-based management does expose FC products to DOS attacks, potentially adversely impacting availability.

One other soft aspect of security relates to simplicity. The more complex the requirements of implementing and managing security the higher the likelihood that it won't be implemented. As the number of servers scales with grid and blade environments, the more challenging FC zone management becomes. iSCSI on the other hand, leverages the simple, well understood and effective Ethernet and IP security mechanisms, augmented by mechanisms inherent in the protocol such as CHAP and digests, raising the likelihood that security will be implemented regardless of the size or competencies of the organization delivering storage services.

- 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?

We don't know.

- Why have no FC switch vendors implemented the FCP security standards from ANSI in their products?

It seems to us that this is another way for them to use proprietary features for competitive advantage? This is the opposite of IP network development philosophy.

## 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. We were the first vendor to qualify our target with their initiator, and we have about 20 storage server models and NearStore systems listed on the MSFT server catalog for iSCSI targets.

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.)

MSFT's initiator is fully compliant with the iSCSI spec. We are not aware of anything proprietary in their code.

## 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?

I haven't heard these claims, nor can I think of a single reason why they might be important. Our FAS systems use FCP drives no matter the protocol. Our NearStore systems use ATA drives regardless of protocol. There's no difference in how a subsystem presents itself to a host. Choice of drive protocol is simply a cost/performance/reliability/scalability tradeoff.

### **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?

Yes, but relative to DAS (non-networked storage) both do provide the benefits implied in Metcalfe's Law. Different types of customers see component cost savings as being more or less important to their buying criteria. For large sites, the issue is more to do with total costs than the costs of the switches or the NICS. In our experience, a large customer won't make a major architecture change unless half the cost savings is a result of decreased management costs. Why? Component costs are short term savings. Decreasing management costs pays back an annual annuity. Small and medium size businesses will see the component costs as the sole justification necessary to deploy IP SANs since management costs are a smaller piece of their total cost of ownership.

### **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?

See # 5, it's really about deployed infrastructure organization which have FC already deployed in the data center will likely continue to deploy FC while those who haven't will very likely undertake an iSCSI direction unless they have already procured FC expertise. iSCSI is an excellent fit for both intra- and inter-datacenter environments because of the likely presence of Ethernet and IP infrastructure as well as deployment and management expertise.



## 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?

Yes. iSNS, CHAP, teaming/trunking/link aggregation, QOS, SNMP-- and everything associated with IP that can be availed of by running SCSI CDB's across an IP network, is a major advantage for deploying IP SANS. More importantly, the fact that there are MILLIONS of people who know how these services work, vs.. HUNDREDS for FCP is the big win for IP SANS. Especially in places where there's no FCP knowledgebase. This is true in the remote office, regional office, and even has implications for adoption by geography. If you look at areas of the world where FCP never caught on due to the requirement for specialized expertise or relatively high cost (China, Eastern Europe, etc), this demographic factor is even more powerful. IP and knowledge about IP is pervasive and ubiquitous.

## 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?

This is an effort to create margin-segments based on reduced functionality deployments to protect the high-end margins. Nothing new in the industry - CPU vendors sub-optimize from their design point servers to enable low-end offerings; switch vendors and HBA vendors do the same as do software vendors -everybody does it.

What's most important is that this is a leading indicator that those vendors who stand the most to lose relative to the wide deployment of iSCSI solutions are seeking to preserve their high FC margins as long as possible - obviously they recognize the universal attraction of iSCSI and so are rightfully concerned about iSCSI!

## QUESTION 17

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

Not at the subsystem level--but the fact that IP Networks are ubiquitous provides a means to move data around an organization--or a country. We think IP SANS will enable the deployment of ILM architectures that FCP SANS can not. Being able to protect data, or effect compliance regulations on data that is remote from a datacenter, is one of the "v-word" advantages we see that IP SANS will offer.

## QUESTION 18

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

We have iSCSI enabled our entire product line (it's a natural progression for and very compatible with our IP-based NAS heritage). We offer Unified Storage solutions allowing customer to simultaneously deploy FC, iSCSI, CIFS and NFS protocols on any of our FAS200 Series, F800



Series, FAS900 Series and NearStore products. We have also enabled the majority of our existing install base to run iSCSI for free by offering them a free download of our iSCSI protocol license.

## **QUESTION 19**

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

In the case of NetApp, the solutions are the same (see question 18). In addition, the complete NetApp portfolio of data protection software (SnapRestore, SnapVault, SnapMirror, SnapDrive, SnapManager for Exchange/SQL/Domino) work with all of our storage protocols, both file based (CIFS and NFS) and block-based (iSCSI and FCP).

The cost difference between systems is simply a matter of the network infrastructure to which they are attached (FCP or IP networks and the choice of host initiator). Obviously, the larger the network the more significant the difference in infrastructure cost is in the total solution cost.

There's one exception to this. Part of our strategy to seed the market with iSCSI has been to offer a free-of-charge bundled iSCSI protocol license with every system we ship, and with every system in our installed base that is under service/support agreements. We have been running this promotion since we started shipping iSCSI protocol products in February 2003 (we announced the day after the protocol standard was finalized). This has played an important role in our position (IDC data) as number one market share in shipped iSCSI systems. It also accounts for the fact that we have the largest installed base of production iSCSI systems in the market. (IDC does not count the free licenses unless we confirm that the systems are indeed running iSCSI. We don't count the system as a production system unless we know that iSCSI protocol is running at loads that indicate production level usage. The latter stat is based on our phone-home support service).

## **QUESTION 20**

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

Absolutely! Let's look at backup and DR. Regarding backup, iSCSI enables the deployment of backup solutions across lower-cost IP infrastructure, providing for lower cost remote backup, which itself relieves remote offices of the requirement to IT staffing to manage local backup and restore. Regarding DR, iSCSI allows lower cost storage-based replication technologies as it removes the cost impediment associated with the requirements for FC to IP bridges or cards essential in FC-based storage replication deployments. Each of these impact points translates into more frequent and more cost effective and broader solution deployments.