Intransa: 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?

We believe that iSCSI is getting to be suitable not only for the edge, but also for the core of the SAN fabric and this will become more prevalent as 10GE takes a stronger hold. One look at the progress made by IP in terms of technological advances and higher penetration over the last several years as compared to FC provides us with a good sense of this direction.

We would also like to make a note that Intransa was one of the early champions of iSCSI. In fact Intransa led the efforts to organize the IETF BoF session which led to the establishment of the IPS Working Group.

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?

Contrary to what FC advocates may claim, iSCSI is capable of far greater than 75% of link bandwidth; in fact, just as with FC, iSCSI can achieve wire speed. This has been demonstrated by Intransa's IP5000. There is nothing inherent in the protocol which limits its throughput. As a matter of fact, in the wide area networks where latency is an important aspect that impacts throughput, iSCSI has the advantage over FC by providing immediate data transfer for write commands, which saves a round-trip per write request compared to 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?

The importance of interconnect speeds is certainly increasing over time, considering that the CPU's are getting faster on both servers and clients, servers are typically serving more clients, disks are getting faster, and multiple disks in arrays are connected on the same interconnects. As we discussed in the previous question, iSCSI and FC provide comparable performance within the same data center, and iSCSI has the performance advantage over FC for interconnecting multiple data centers. Therefore, iSCSI is extremely well positioned to handle the challenge of increasing interconnect speeds.

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?

iSCSI does indeed have an advantage over FC's "dual network" design that is explained as part of question number 18 below.

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?

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?

Using NAS gateways to support both NFS/CIFS and iSCSI on the front end is certainly possible, but with some caveats. First of all, native iSCSI SANs such as Intransa's IP5000 use a cluster of controllers to provide active-active high availability, load balancing, and scalability of throughput. NAS heads do not provide the same high availability and performance.

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. Even though iSCSI has been a standard for much less time than FC, already the iSCSI products are more interoperable with each other. Due to its open and participatory nature, IETF has a very good track record of successful, interoperable deployment of its standards.

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? Yes

And if so:

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

  There are many killer applications for iSCSI SANs in the SMB market, including e-mail, data warehousing, databases, software development repositories, etc.

- 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?

  Unlike SAS, iSCSI is a true network protocol, and is not limited to short distances. Furthermore, being a network protocol, it makes it much easier to provide controllers that add value, such as storage virtualization, load balancing, high availability, etc.

- 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?

  The issue is that IP SANs would address is not so much one of capacity for SMBs. IP SANs address issues such as simplified storage

management, sharing of storage resources by multiple servers, high availability, serverless backups, etc.

- With removable/exchangeable disk/tape hybrids, such as Spectra Logic's RXT platforms, can SMBs achieve capacity scaling requirements without deploying SANs at all?

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?

TOE development has indeed proved more daunting than originally thought. Moreover, as computers are getting faster, the need for TOEs is getting less. However, we expect TOEs to have more of an impact during the first few years of widespread 10GE deployment (atleast for a while).

QUESTION 10

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

iSCSI is just as secure as FC. One of the security "features" of FC is that by being a different network, it is physically isolated from the IP network. However, that is not an inherent advantage of FC. In fact, Intransa recommends the iSCSI network between the servers and the Intransa IP5000 to be physically isolated in order to provide the same level of physical security as FC. However, with IP SANs, such physical separation is not a necessity, it's a choice that the customers make depending on their environment and their security needs.

- 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?

If an FC fabric is using an out-of-band, IP-based connection for fabric management, then their management security is also equivalent to the IP SANs

- 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?

These are indeed contradictory positions by FC advocates.

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

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? 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.)

Yes, Intransa fully interoperates with the Microsoft iSCSI initiator. We have been working very closely with Microsoft ever since the early days of iSCSI. We also wrote an Internet Draft jointly with Microsoft and Intel, titled "iSCSI Implementation Guidelines for Fault Tolerance and Load Balancing using Temporary Redirection".

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?

As already answered in the question itself, there is no such thing as a "Native FC Drive", and they are really SCSI drives with integral FC-SCSI bridges. Note that by the same definition Intransa provides disk drives with

integral IP bridges, and each disk drive in the Intransa IP SANs is IP addressable.

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. The reasons for the cost difference are very well explained in the question itself. Please also see the TCO analysis slide between Intransa/NAS/FC (attached).

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?

iSCSI is both an edge technology and a data center technology. Indeed, in our case we are using iSCSI in both environments. At Intransa we run our entire operations on IP5000 systems with 5 9's availability.

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?

Many of the problems that were solved years ago in the IP world are recently being rediscovered in the FC world. SNMP and DHCP are excellent

examples provided in the question. Other examples are DNS and SLP for discovery of hosts and services, RADIUS, Active Directory, or Kerberos for authentication infrastructure, IPSEC, SSL, and SSH for data and management security. There is even a recent effort to define routing for Fibre Channel - something that has been a non-issue in the IP world for decades.

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?

Our strategy is not to provide a "dumbed down" product. To the contrary, Intransa offers intelligent storage management that simplifies the users' day-to-day tasks through automated policies for volume configuration, automatic load balancing and failure detection/recovery, and a wizard-driven GUI. Underneath that automation lies an intelligent, full-featured, highly flexible IP SAN.

QUESTION 17

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

Unlike FC, iSCSI is not geographically limited. Therefore iSCSI storage virtualization is not limited to only one data center, but could be geographically distributed. This allows applications such as remote replication, content distribution and delivery, etc.

QUESTION 18

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

- **Architecture: IntraStor™ Architecture**

The Intransa architecture is based on providing virtually limitless, incremental scalability by using commodity low cost, high performance ATA drives and Internet Protocol over ubiquitous Ethernet cabling.

This unique and only architecture in the industry today separates the storage controllers (and along with them the control function) from the storage devices (disks) by a standard GigE switching fabric, allowing all disks on the network to be virtualized and treated as a single pool of storage. This can dramatically improve the overall utilization of storage devices. Further, this architecture allows for capacity and performance to be scaled independently. Initial system capacity can be expanded by simply plugging an additional disk enclosure into the network; the added drives identify themselves and automatically become part of the free storage pool. Performance can be scaled by plugging additional controllers into the network; the new controllers identify themselves and join the cluster after which the system automatically rebalances itself and transfers control of virtual volumes to the new members.

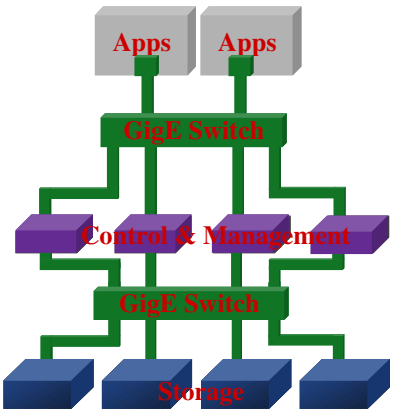

Product Overview

An Intransa IP SAN consists of three main components supplied by us—storage controllers, disk enclosure and software—plus Gigabit Ethernet switches, uninterruptible power supplies and iSCSI initiators provided by other vendors. This separation of functions allows intelligence to be added where it can be meaningful and expansion in multiple dimensions based on implementation needs.

The SC5100 storage controllers are powerful dedicated appliances which provide visibility for the storage to networks. Each multiprocessor based

module contains 2GB of cache and four Gigabit Ethernet copper network ports. The SC5100 employs dual hot-swappable power supplies and redundant cooling fans for a highly available hardware design. Processing throughput can be increased by adding additional SC units.

The DE5200 disk enclosures are robust capacity units containing 2TB and 4TB (eight or 16 250GB 7200 RPM disk drives, respectively) each with its own system-assigned IP address in a 3U form factor. High availability is ensured by including dual Ethernet ports for each set of four drives.

The StorControl™ software runs on the storage controllers and handles all the core tasks such as virtualization, data mirroring, snapshots and host access control, with remote replication a major feature of the next software release. StorControl™ software allows customers to manage, monitor and administer the IP SAN system. Customers can add, delete, resize, and mirror volumes, redistribute resources, assign hosts read/write access, and monitor or manage performance and health at multiple levels – all from a single Java™-based GUI.

The IP SAN leverages existing networking equipment to create the storage infrastructure. Reduce cost by connecting hosts to IP SAN with low cost NICs, software iSCSI drivers, and GigE switches which eliminates the need for expensive Fibre Channel HBAs and Switches. For additional performance, low cost iSCSI Host Bus Adapters are available but not required.

QUESTION 19

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

Pls find attached a slide that compares Intransa implementation against both FC and NAS

QUESTION 20

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

As discussed before, iSCSI allows remote replication natively, which facilitates disaster recovery immensely. Furthermore, both for remote

replication and for local backups, IP SANs can easily integrate with host agents such as Microsoft's VSS/VDS to provide application-level consistency.