TANEJA GROUP
TECHNOLOGY ANALYSTS

PRODUCT PROFILE

Intransa Opens its Platform

October 2005

THE NET NET

The battle for the SAN is on. FC SANs made serious inroads into the datacenter over the past five years but failed to gain much traction in the SMB space, primarily due to cost and management issues. Over the past two years, iSCSI, which was presented to the market as the simple and cost effective alternative to FC SAN, has been making accelerating headway into SMBs for applications below the mission critical tier. In this writeup we describe the three categories of iSCSI players in the market today and how each is impacting the market. We use Intransa, a new iSCSI player, as an example of a vendor that brings out the best in iSCSI. As we dig deeper we discover that much of the innovation brought about by Intransa and a few other players is really outside of iSCSI, per se, and the fact that they brought it within the context of iSCSI is more a matter of choice and less an artifact of iSCSI itself. In any case, it gives iSCSI a solid footing in the market and FC a run for the money.

The Battle for the SAN

The battle for the SAN continues. FC SANs have continued to penetrate large enterprises at a healthy clip over the last five years. We expect that 65% of all new storage shipments into a large enterprise are being implemented in a SAN. The remaining 35% is made up of 10% NAS and 25% DAS. We expect both the NAS and the SAN portions to grow by another 5% each and stabilize around these points. There will always be reasons to keep some systems DAS due to location, cost, simplicity or preference.

The story is entirely different in the SMB marketplace. DAS is still at 85% or more, with FC penetration at no more than 10%. iSCSI SANs are making some headway but SMBs are still largely virgin territory for SANs of either variety. It is within this environment that iSCSI SANs are entering the market.

The Current State of iSCSI

In spite of some vendor's claims to the contrary most iSCSI installations thus far have been in the SMB space and most run applications that certainly fall below the most mission critical. We have yet to find an iSCSI based trading application on the Wall Street or in London's financial district. Having said that, iSCSI solutions are making serious headway into running the second tier applications in the SMB space. In the larger enterprises we find most iSCSI SANs are built from what we call "excess capacity". Here is a classic example of what we mean. You are a NetApp customer and you have several filers in your data center and in your departments. Most of these filers have some spare capacity. On the server side you have several situations where the application is important but not mission critical and the performance needs are moderate. In this environment the IT administrator would install an iSCSI initiator (free) on the server and, given that NetApp ships the iSCSI target for no additional cost with each filer, configure the "excess" capacity for use by this application. This is what we call smart usage of storage. It is another matter that the customer paid higher cost per MB for this storage than for iSCSI storage available in the market or that when he needs more capacity he is now more likely to add storage to the same filer, even if the storage was required for the iSCSI application. If you replace NetApp's filer with EMC's DMX or a Clariion, the same logic applies. As a result of this phenomenon many established players are touting strong iSCSI revenues. A large % of these revenues are incidental, as seen above. There are certainly exceptions to the rule. Dell, for instance has been shipping iSCSI SANs very effectively using the AX100i platform from EMC. These are pure iSCSI SAN solutions with no legacy impact.

The newer crop of iSCSI vendors, including Intransa, EqualLogic and LeftHand Networks, present what we believe to be the most innovative iSCSI solutions in the market. These solutions have as much innovation outside iSCSI as within. They are designed to fully take advantage of the original iSCSI vision, which, at the risk of repetition, includes ease of installation and use, reuse of TCP/IP knowledge, and lower TCO and capital costs. The large majority of shipments from these vendors has been going into the SMB space, including departments or divisions of larger companies, and is being implemented as true iSCSI SANs from the ground up.

This is an extremely high growth area for the next three years, given that an iSCSI SAN can and will bring as many of the SAN benefits to the mid market as FC SANs did to the larger companies. By 2008 we expect that the DAS % in SMBs would be down from 85% today to 50%, with most of reduction being absorbed by iSCSI, which, we believe, will stand at 30% from a small number today.

Focus on Intransa

Intransa, a company headquartered in San Jose, CA, and with an R&D arm in Pune, India, has been shipping an iSCSI SAN solution since 2003. The company has since expanded its product line to include three performance levels and has just added a low end product line that we discuss in detail later. Unlike most US based startups that focus entirely on the US market initially and enter European and Pacific Asia regions in years three or later, Intransa, recognizing the attractiveness of a simple, low cost SAN solution to the cost conscious Asian customers, entered those markets from the get go. The implications of this move, which we discuss at length later are huge. Suffice to say that 60% of their revenues currently are from outside the US. By contrast we expect no more than 10% of the revenues from other pure play iSCSI players are from outside the US.

The company is adequately financed at $75M over four rounds. The lead investors include US Venture Partners. Menlo Ventures, Advanced Technology Ventures, and Sofinnova Ventures.

The Architecture

There are three basic architectures currently prevalent in the marketplace. The first, exemplified by the legacy players, such as Adaptec, EMC and HP, is to essentially take an FC-SATA storage product and replace the front end with an iSCSI target. All other features of the product are preserved. If there was array-based snapshot and mirroring available it is now available with an iSCSI front end. Period. No enhancements beyond that. The storage box, therefore, is as easy to use (or not) as it was with a FC front end.

The second architecture is similar to the first, except it is applied to the NAS storage system. The classic example here is that of NetApp. An iSCSI target is added to the front end but since a filesystem is inherently involved in dealing with all the traffic coming over the TCP/IP network (WAFL in the case of NetApp), the iSCSI block traffic at least touches the filesystem. These products minimize this touch and push the blocks as quickly as they can to the storage behind the filesystem. Depending on the design the impact on performance may be trivial to substantial. As in the first case, the basic characteristics of the storage system do not change. The same applications as before are available.

The third and the most interesting set of architectures are presented by the new crop of iSCSI players. Indeed there are differences in their approaches that are important to discuss but in general these players have taken a fresh look at iSCSI and what it could mean to an IT shop. They fundamentally rely on a nodal concept where the customer can start with a single node and grow by adding nodes. The real innovation is in the way the storage is virtualized, how it is provisioned, how it is managed on a day-to-day basis. The basic idea for all these vendors has been to aim for the ultimate goal of “set it and forget it” storage. While no one has reached that as yet the group has taken a huge step, relative to FC. The concepts of “external box” virtualization, online expansion, capacity efficient snapshots, grid-like non-disruptive scalability, high availability and no SPOF, are all at the heart of these systems.

But there are crucial differences, which we describe within the context of Intransa.

Intransa’s basic architecture starts with a Storage Controller that runs Linux operating system, implements the iSCSI target and presents one or more GE interfaces on the host side. The primary software that is at the heart of the Intransa system is StorControl. The controller has additional GE ports for inter-controller communication and Disk Enclosure (DE) connectivity, and for management. The company’s controller model is currently based on the X Series IBM X335. Each controller presents a total of four GE ports.

The complement of this storage controller is the Disk Enclosure (DE). It contains 16 slots in the current version to house hot pluggable SATA drives. There are four disk controller modules, each presenting two GE ports for connection to the Storage Controller(s). Each disk controller manages four disk drives and is hot pluggable. The disk controllers fundamentally manage the conversion of IP traffic to SATA and to direct the data to the right drives. They also manage the allocation of IP addresses to each disk drive. This aspect is unique to Intransa; each drive is addressable by a unique IP address. Noteworthy is the fact that the RAID functionality exists at the storage controller level and not at the DE level. Of course, there is redundant, hot pluggable power and fans to ensure no single point of failure.

The installation requires the connection of SC(s) and the DE(s) to the Ethernet switch and giving the system an IP address. The SC(s) speak to the DE(s) over TCP/IP using a low latency proprietary protocol, X-Block, developed by Intransa. The external communication to the hosts is all iSCSI. Once configured, the entire storage is automatically virtualized into a pool. Volumes can be created with RAID 0, 1, or 10 (RAID 5 has been added on newer low end models). Allocating volumes to servers and applications has been made almost trivial (5 clicks to provision a volume).

The key aspects of this architecture have to do with the fact that when the first DE is maxed out, a second DE can be added non-disruptively and the storage capacity is added to the pool automatically. Volumes can be extended non-disruptively. The same concept applies if an additional SC is added. The SCs are designed to perform load balancing between themselves automatically. The SC adds to the performance of the system and the DE adds to storage capacity. Both can be scaled independently of each other.

The other pure play iSCSI vendors also use the nodal concept, except they maintain a fixed processor-storage ratio. In other words, a node to them is a combination of storage controller and a certain maximum amount of storage. In the case of Left Hand Networks, the node is an x86 server and four SATA drives. Each node is identical to the next. EqualLogic's molecule is larger in that it has up to 14 disk drive slots along with the server but not all drives have to be present in the minimum configuration.

This primary architectural difference gives Intransa the ability to scale capacity independently of the processing power (very crucial in secondary disk applications, for instance). Adding compute capability prepares the system for IO intensive applications. Intransa's architectural difference also allows a third party storage box to be used instead of a DE and lends itself to a stronger OEM play.

The other key architectural difference for Intransa is the fact that each disk is IP addressable. In the case of EqualLogic and LeftHand Networks the compute engine does all the management and translation of iSCSI traffic to SATA. It also manages RAID and the allocation of volumes. Intransa's ability to "speak" to each drive will have future implications that are not visible today.

All vendors in this space have added various storage applications, such as snapshots, asynchronous replication directly from the array, mirroring (local and remote), etc. We will discuss Intransa's specific product line in more details.

In Intransa's architecture, a common pool is created, from which logical volumes may be synthesized. If a user desires to scale by adding a second Disk Enclosure, StorControl discovers the new disks and reports them. The user may then resize existing logical volumes to include the additional disks or create new volumes. The default is RAID 1 (mirror) and the system automatically creates the mirror when a volume is created or expanded.

The Disk Enclosures (DEs) are networked into a StorControl Cluster; StorControl is aware of their existence as networked disks. The DEs themselves are not intelligently aware of other DEs. When a node fails, the cluster automatically picks up the loads through failover. Logical volumes can be as large as the total population of disks (networked DEs) or any subsets.

Volumes are owned by one of the controllers in the cluster; but load balancing can shift the volumes. In contrast, the interesting challenges with the Equalogic architecture is that large logical volumes span controllers, with controllers being limited, at a moment in time, to reads and writes to physical volumes physically behind that controller. Where there are large logical volumes and a great deal of non-sequential I/O taking place, the Equalogic architecture must hop between controllers to hit the target.

Intransa sees all of disks in the networked DEs as being a single enclosure; reducing the need for controller hops, creating a likelihood of better performance.

With all of the variety of the product lines, Intransa has one controller model, based on the X Series IBM X335.

Intransa's Product Line

Until now Intransa's product line consisted of three products; the IntraStor IP 3000 (single SC), IntraStor IP 5500 (dual SC) and IntraStor IP 7500 (quad SC). The three products share common characteristics and run the same OS and applications. The difference lies in the number of Storage Controllers and Disk Enclosures. The IP 3000 has one SC and up to 2 DEs and currently delivers 12.8TB of capacity with 400GB 7.2K RPM SATA drives.

The IP 5500 comes standard with 2 SCs, with dynamic load balancing between them and up to 4 DEs can be attached on the backend for 25.6TB of capacity. The top end of the product line, the IP 7500 can be configured for a max of 4 SCs and up to 8 DEs for a maximum of 51.2 TB of capacity. The SCs are 1U and each contains 2GB of cache and come standard with 4 GE ports.

The application servers can be Windows (all popular versions are supported), Linux (RedHat or SuSE), Solaris, or HPUX. A cluster of SCs is called a Realm. Realm-to-Realm asynchronous replication, where the SCs can be local or remote, is supported. So are snapshots and local mirroring.

The list pricing of the product ranges from $30K to $238K, with the price per MB of storage 25% to 50% below comparable FC configurations.

Latest Additions: New Value Line

Intransa this month announced a new line of iSCSI products for the low end of the family.

The basis of these products is an iSCSI storage box from Rasilient, named IP 2000. This is a 3U box with an iSCSI interface to the outside via GE and 15 slots for SATA drives, for a max capacity of 6TB, using 400GB drives. RAID 0, 1, 10 and 5 are supported. There are two multifunction controllers with failover built in the IP2000. The design is cableless for simplicity and cost reduction.

The significance of the IP 2000 is that it provides an inexpensive way to implement an iSCSI SAN. The starting price for a 2 TB system is just under $20K. A 6TB system configured with 400GB drives lists for $5.7K per TB. But even more significant is the fact that one or two SCs can front end the IP2000 and create models IP2300 and 2500 respectively. These configurations use the IP2000 as a DE. The IP2300 places one SC in front of two or more IP2000s, and the IP2500 places two SCs in front. The complete functionality of Intransa's product line becomes available in these configurations. The IP2000 offers up to 6TB and supports Exchange, Backup and workgroup applications. The addition of the IP2300 or IP2500 enables much larger capacity configurations through clustering IP2000s and the creation of large logical volumes, higher performance offered through striping across IP2000's, and extends the reach into applications support for SQL, Oracle, remote replication, disaster recovery and fixed content archiving.

Along with IP2000 and the relationship with Rasilient, Intransa also announced iSCSI Storage Array Manager (iSAM), an extension to its StorControl software that provides support for third party iSCSI disk products, integrating them into an Intransa Realm as a DE and providing the same virtualization and data protection support as found in its own native DEs.

IP2000 is the first such use by Intransa but it signals their future intentions very clearly. StorControl, the essential Intransa software that runs across their product line, will now be able to manage "foreign" storage, using iSAM.

Relationship with H3C Announced

In addition to the announcing the new product line Intransa also announced that they have signed a development and marketing agreement with China's growing powerhouse, H3C, the Huawei and 3COM joint venture that has designs to compete with Cisco on a worldwide basis. According to the agreement H3C and Intransa will co-develop IP SAN solutions to address the expanding Chinese market.

Taneja Group Opinion

The time for iSCSi has finally arrived. After several false starts and hiccups the ecosystem is finally in place for the sales engines to turn. 2004 was approximately a $100M worldwide market, including the contributions from the legacy players, which in many ways distort the real iSCSI numbers. We expect this year to be at least double that number. A large majority of sales have come from the mid market where FC SANs had not yet taken hold and simplicity and cost effectiveness were key.

On the surface everything looks great for iSCSI. But as we look under the covers we discover many variations, several myths and many false expectations. In our view, the success of the so called iSCSI market has less to do with iSCSI, per se, and more to do with the simplicity and ease of use that several newcomers have brought to market. The fact they brought this market via iSCSI is simply a coincidence. The principals they apply as equally valid to FC. Except that the FC market, due to its ecosystem, has had difficulty implementing these features. As we interviewed early users of iSCSI systems we discovered that what they liked was the fact that storage and processing power could be kept in balance via the grid concept, the fact that all backend virtualization happened automatically, the fact that the time to production for a new iSCSI SAN could be two hours and not a week or more, that a volume could be allocated in minutes, expanded in seconds and that expansion was nothing more than adding another module. Note that these features have nothing to do with iSCSI but the fact that vendors first brought them out in conjunction with iSCSI makes them appear connected. They are not.

Regardless, we believe that Intransa and a few other players have changed the storage game and we are in the very early stage, in the scheme of things. To be sure, it seemed that Intransa had lost its step for a while. But now it is clear that they were focusing more on the China, India and the Far East markets and building relationships that are only now coming to light. The company is now committed, after much focus on overseas markets, to expand its North American presence and believes that the IP2000 coupled with some assertive marketing and sales focus will help it meet this objective. The fact that Intransa's architecture allows for the separation of front end storage controllers and the backend storage gives them the flexibility to work with other storage players much more easily than if the molecule consists of the combination of the two, as in the case of LeftHand Networks and EqualLogic. As we indicated before the real magic of any of these solutions is what they do with storage and not that they package the whole thing themselves. Intransa has shown the first card with the partnership with Rasilient. We expect much more, especially with the China connection.

We also like the fact that 60% of Intransa's revenue already comes from outside the US. The Far East markets will be crucial for volume sales. Many of those customers haven't gone to FC yet and they are ideal candidates to jump on the iSCSI bandwagon. These customers also like to buy from companies that have built relationships over time. Intransa already has an advantage here.

The legacy players stand to do well in iSCSI as well but for a totally different set of reasons. We think they will sell existing storage for use via iSCSI. Not to say they are not already selling iSCSI targets, such as EMC AX100i at an increasing pace. But frankly, these products do nothing more to simplify a SAN than their FC counterparts. They happen to communicate with the server via iSCSI rather than FC. Granted, there are savings involved (the switch, the NIC are practically free, even if the array costs the same as its FC counterpart) but, in our view, they miss what is really important about any SAN: simplicity of use and reduction in TCO.

Intransa has not done a great job of bringing out these differences to the market, especially in the US market. We believe they have architectural advantages that need to be marketed aggressively. We believe they need to not only differentiate themselves from the legacy players but also from the "new" players that are starting to get traction. A few from this pack will make it big, others will wither away as the big players get serious. The race for the finish line is on and Intransa has all the ingredients to win. All it needs now is the desire... to win. We think they do and they will.