

Exclusive: Intransa realizes dreams of 10Gb iSCSI **StorStac PCU100 combines superior performance, scalability, and manageability**

By Logan G. Harbaugh May 7, 2007, *InfoWorld Magazine*

Ever since the announcement of 10Gb Ethernet, many storage managers have looked forward to 10Gb iSCSI. The hope was that for the first time, iSCSI would actually have not only a cost advantage over FC (Fibre Channel), but a performance advantage as well.

Intransa's StorStac PCU100 is one of the first 10Gb iSCSI systems available. It won't be formally released until later this month, but InfoWorld was given an exclusive opportunity to review it. Based on my testing, I can say that storage managers will be pleased to know that their hopes for 10Gb iSCSI have been realized.

Intransa's system boasts some unique features, different from most iSCSI systems, such as its basic architecture. Unlike other iSCSI and FC storage systems, the StorStac connects all drives via Ethernet. Each drive has its own IP address and every two drives share a GigE (Gigabit Ethernet) connection, providing each drive with more bandwidth than it needs.

High-speed storage

My test system had two storage controllers and two servers, connected to a 10GigE switch, and four storage units connected to 2GigE switches with 10GigE uplink ports. The storage controllers were connected to the switches via the uplink ports. The two servers ran Windows 2003 and RedHat Enterprise Linux 4.0 and were equipped with Neterion 10Gb Ethernet NICs.

The PCU100's performance is all that one might hope for. In testing, I saw sustained throughputs of 700Mbps (5.6Gbps), which is better than 4Gb FC performance. With two controllers, I saw throughputs in the range of 830Mbps to 900Mbps.

IOPS will not be as high with iSCSI as with high-performance FC arrays, due in part to the differences in latency between the two technologies. However, according to Intransa, targeted applications for its product typically use streaming content, such as video processing, video on demand, medical imaging, video surveillance, and disk-to-disk backup and recovery, where IOPS performance is less relevant.

Intransa uses the same drives and enclosures for all three of its systems, which run at 1Gbps, 2Gbps, and 10Gbps. The differences between the three are in the head units rather than in the drive subsystems. This approach ensures that your investment in storage is largely retained should you upgrade from the lower speeds to 10Gb.

Superior scaling

The StorStac system offers excellent scalability, not only in allowing existing users to easily and simply upgrade to a higher-speed system, but also in allowing for as many as four storage controllers and 15 storage shelves. You can even add iSCSI storage from other vendors to the PCU.

The four controllers are configured as a cluster. Bringing an additional controller online is pretty straightforward: The system's cluster management software will detect it once you configure the IP address, and it will ask if you want to add the new controller to the cluster. Adding additional disk shelves is just as easy: Connect Ethernet and power cables, power on, and the controller auto-discovers the drive shelf. In about two minutes, the additional drives are available.

Notably, the clusters provide automatic failover. If one controller fails, the other controllers continue to make the storage available to the same hosts. In testing, I ran IOMeter against a partition on the StorStac, then disabled one of the two controllers. There was a short (15 second) delay during switchover, but the IOMeter job continued to run without interruption or error. Controllers don't load balance automatically; with one LUN (logical unit number), there's no gain in performance if you've got two controllers. Multiple LUNs can be spread across two or more controllers, providing the same 500Mbps to 700Mbps performance to each, and file systems can stripe across multiple LUNs for doubled performance.

The controllers are all managed through a single virtual IP address. Conveniently, all the storage is addressed through that IP address, regardless of how many actual controllers there are. You can set up more than one IP for segregation of storage, but it isn't necessary.

There's also a single management interface for all controllers, offering excellent, easy-to-use monitoring of all hardware as well as configuration and management of RAID levels and other software functions. Each partition can have different RAID levels, numbers of drives, and stripe sizes, and these settings can be saved as a policy and applied to new partitions as well. You can also easily expand LUNs in size as desired -- without downtime.

Current supported RAID levels are 0, 1, 10, and JBOD (just a bunch of disks), with RAID 6 coming soon. The RAID 10 that Intransa uses can handle multiple drive failures, as long as two drives in the same pair don't fail simultaneously. If a

drive fails, an unused drive is automatically added to the defined partition, and rebuilding takes place automatically. Rebuilding occurs at about 50Mbps -- even while the array is in use.

Adding hosts that are allowed to use the controller is also simple. The controller will register any host attempting to connect to storage. Through the management interface, an admin can then grant access, with one click, for any LUN defined on the system.

Finely tuned features

The system's replication features are what you'd expect, including cloning (mirroring), snapshots, and asynchronous replication. The latter, of course, doesn't require a WAN gateway, since it's already occurring via TCP/IP. Intransa also adds some OS-specific and extremely useful functionality. On Windows, the Intransa system integrates with Microsoft Volume Shadow Copy for snapshots through the OS. On Linux, iSCSI devices (/dev/sda, /dev/sdb, etc.) usually are not persistent -- if you have more than one iSCSI LUN attached and reboot the Linux server, the devices may change positions and thus designations. Intransa drivers rename the devices using the LUN name from the controller to prevent this.

The Intransa StorStac PCU100 offers serious speed at accessible pricing. Intransa does not sell directly, but through individual partners who set the pricing. According to Intransa, a basic 10Gb system with one controller and one 4TB drive shelf should start at about \$75,000; about \$130,000 with 24TB. A dual-controller system with 24TB should cost about \$180,000.

This is a fair chunk of change, but not more than a 4Gb FC solution; the cost of host adapters and switches for 10Gb Ethernet should continue to fall more rapidly than 4Gb FC. The Intransa StorStac PCU100 delivers the performance level you'd expect from high-end storage and ease-of-use is great. Expanding the system with more controllers or storage is a snap, plus you can combine 1Gb, 2Gb, and 10Gb controllers with the same disks. Although the PCU100 is not the ideal solution for databases and other transactional systems that need high numbers of IOPS, it's an excellent solution for streaming data at very high speeds or aggregating a whole lot of smaller streams.

Intrinsa StorStac PCU100

Intrinsa, intrinsa.com

The logo for InfoWorld, with 'Info' in red and 'World' in black.

Excellent 9.0

criteria	score	weight
Performance	9.0	30.0%
Management	9.0	20.0%
Reliability	9.0	20.0%
Scalability	9.0	20.0%
Value	9.0	10.0%

Cost:

Pricing has been announced starting at \$75,000, although that may change at the time of launch.

Bottom Line:

The Intrinsa PCU100 offers superb 10Gbit iSCSI performance, scalability, and manageability in a relatively inexpensive package. It uses a unique architecture that connects each drive via TCP/IP, rather than relying on an internal bus. The system is oriented toward streaming data rather than transactional, however, so applications such as databases will not see as great a benefit. The system delivers excellent redundancy with multiple controllers and drive shelves, as well as an easy setup and a great management interface.