

DECEMBER 21, 2005



TWO PIONEER FIRMS OF SUPERCOMPUTING NOW PLAYING CATCH-UP

By Henry Newman

Another year has passed, so it's time once again to look at the predictions I made last year, and to look ahead at what may be in store for the storage industry in 2006. In so doing, we'll not only be judging my prognostication skills, but also hopefully gaining insight into where the storage market has been and where it's headed.

First, a look back at my predictions for 2005.

First Half 2005 Predictions

I said that 4Gb Fibre Channel HBA availability would become commonplace, but I was off by a little bit on that because widespread availability didn't occur until July or so. I also said that 4 Gb RAID's would become available, and thanks to SGI and Engenio, I was right.

I predicted that RAID vendors would announce bigger, faster midrange devices with 4Gb support, and in some cases support for up to 10Gb, which requires an infrastructure change. We'll thank Engenio again there for making me look smart.

Another prediction was that Fibre Channel switch density would increase in port density, as would the bandwidth that is required for 4Gb FC. I was wrong on this one, and, in fact, director-class switch port density did not increase unless high-density blades were used. These blades do not run at full rated 4Gb for all ports active at the same time.

PCI Express would become commonplace in the Linux cluster and on small Intel servers to replace PCI-X — PCI Express was available the first half of the year, but it was not commonplace.

Since we'd been stuck at 300GB Fibre Channel drives for a long time, I surmised that disk sizes would get bigger. HDS announced 400GB Fibre Channel drives, but when will we start to see deliveries in quantity and shipment from the RAID vendors? Except for SATA drives, this was not correct.

I also predicted increases in tape density for high-end tapes. Imation in May announced that they have built a plant that will create 1TB tapes, so this is a step in that direction, but will the performance increase at a commensurate rate? I was off by a couple of quarters on this one.

Storage management would continue to improve, I predicted, but storage management nirvana would not be reached. Nirvana to me will occur when you can manage the servers, file system, HBAs/iSCSI NICS, switches and RAID's for a whole enterprise — this does not even include a single console for storage and network management. No surprise there — we're still along way from nirvana.

I thought we were due for better support for host side failover for Linux and other operating systems. It got better, but not by much. I also said we'd get support for HSM from a number of shared file system vendors that currently do not support HSM, and a few vendors announced support for Linux and other clients.

Second Half 2005 Predictions

I said we'd see 10Gb Fibre Channel RAID devices, and two vendors made me right on that one. I also said we'd get closer to terabyte tapes, and we can thank StorageTek for that.

Vendor products based on the Object-based Storage Devices (OSD) standard would become available late in the year or early next year, I said, but so far Panasas remains the only such vendor.

We haven't seen 600GB ATA drives, so I was wrong there. I also suggested that we would get better management and Linux cluster tools, which we got, while nirvana remains but a dream.

Better HBA drivers and file system options for Linux would be coming, I said. A number of vendors have made drivers easier to use, and Lustre and Panasas have allowed large global namespaces. But I was wrong when I said better shared file system integration was coming for heterogeneous environments. I see some progress, but not enough.

Looking at the results, I'd say I batted just over .500 on the first half of the year, and .750 the second half. I did a far better job for the second half of the year, and I think the reason was that vendors had time to fix some of their mistakes and still get the products out.

Last Year's Predictions for 2006

We'll begin our look at the year ahead with some of the things I said might happen in 2006 in last year's piece.

For 1TB tapes with either LTO standards, IBM or StorageTek, I no longer think that 1TB tapes will happen before the end of 2006. There were too many delays this year.

I still think we'll see OSD shared file systems and object storage managers (RAID) integrated with HSM, with the announcement coming late next year and general availability in 2007.

PCI express in large servers with more than 24 CPUs is a big issue, since the memory backplane will need to be redesigned, but I still think we'll see them in the first half of 2006.

More usage of RDMA-TCP. This could change the nature of storage devices, since we might live in a TCP-IP world alone for storage if this takes off. I am sticking with this one for the first half of 2006.

Finally, and most importantly, we are finally going to get some real security for file systems, and storage in general. This might include encrypted file systems with real user interfaces and management that run at device rates. Still expected for the first half of 2006, but it might not contain per-user keys, just system-wide keys.

New Predictions for 2006

And now for my new predictions.

4Gb Fibre Channel director switches will become available in the first quarter of 2006. This technology generally follows edge switches, which have been out for a number of months.

8Gb Fibre Channel will not happen in 2006, and may not happen at all. We could just go straight from 4Gb to 10Gb, even though it will require an infrastructure change.

InfiniBand (IB) will continue strong growth in the storage market, replacing FC for some large clusters and other HPC environments. IB growth will accelerate, given that I/O bus bandwidth now allows IB to work at full rate with PCI Express.

Drive density will continue to slow. This is a trend that has been happening for a number of years and will continue.

Holographic technology might finally come to fruition by late in the year. There seems to be some excitement around this technology from a number of areas.

Predictions for 2007

Object-based storage will be available from the host to the storage device, including object managers that replace RAID controllers. This will be the first major change in storage technology end-to-end in many decades. This will affect areas such as HSM, write reconstruction and virtualization, to name just a few.

POSIX file system semantics will be changed to allow for things like lazy metadata updates for shared file systems. The standards are more than 25 years old at this point and need to be updated.

Metadata allocation and management for file systems will be addressed. Areas such as databases and other additions for file systems will address the problems of *ls -l* and *find* and the effect on file system performance.

Summing Up

I do not see lots of changes happening in 2006, given the recent availability of 4Gb Fibre Channel, but I expect OSD to make some big changes in the storage industry in 2007.

The biggest area of change that I think will happen in 2006 is the use of IB-based storage connectivity. This is already happening in large clusters, where they only want a single interconnect topology. Yes, you all remember about five years ago when IB was going to be the local connection fabric for everything, which of course did not happen. The reason the technology will fare better this time is that the PCI bus bandwidth even 12 months ago did not support IB performance, but today things are different.

While I don't see any breakthrough, must-have technologies for 2006, I think 2007 will be a significant year for a number of technologies. Storage technology needs some significant R&D dollars and some leadership compared to CPU technology (see Storage Headed for Trouble). OSD has the potential to manage data in a revolutionary way, but will need some revolutionary technology to address the physical layer. At least that's the way I see it.