

NOVEMBER 10, 2005



## HIGH-SPEED LINKS HEAD FOR MAINSTREAM

By Mary Jander

Suppliers are working hard to bring high-speed interconnection technologies from supercomputing sites to the more mundane realm of corporate data centers. And if consumers aren't ready for the migration -- and it seems quite a few are not -- they'll have to find ways to fight off a major pitchfest.

The Supercomputing 2005 conference next week in Seattle, for instance, will be the venue for a flurry of announcements about interconnecting data center systems and storage. A key aim will be to make technologies accessible to ordinary users as well as those employing high-performance computing (HPC).

**There will be data center clustering wares from Microsoft, as well as Isilon and Panasas; InfiniBand storage systems from DataDirect Networks; and the unveiling of an interoperability partnership between Foundry and Myricom. A multivendor "cross-continental InfiniBand cluster" demonstration from the OpenIB Alliance will be featured. (See OpenIB Showcases InfiniBand.) And at least one startup, Obsidian Research, will claim a breakthrough in extending InfiniBand technology to support data replication on WANs.**

The common denominator is bandwidth. For months now, suppliers have trumpeted InfiniBand, 10-Gbit/s Ethernet, and iWarp as faster ways to link systems and storage in consolidated data centers. What's more, as Fibre Channel vendors argue over whether to move ahead from 4-Gbit/s, InfiniBand is already hitting 20-Gbit/s in new Double Data Rate (DDR) equipment -- boosting its rep as a Fibre Channel replacement in transaction-heavy networks. (See Report Hedges High-Speed Bets, High-Speed Links Favor Ethernet, and Ohio Opts for iWarp.)

But lots of users aren't hearing their marching orders. InfiniBand is widely adopted for clustering in labs and is deployed inside NAS boxes, which continue to support faster upgrades, as next week's announcements will show. (See InfiniBand Gets Second Looks.) But InfiniBand doesn't support the distance extension that could make it a more mainstream technology.

At the same time, 10-Gbit/s Ethernet is too chatty to handle hefty data center links without the use of extra TOE (TCP/IP offload engine) cards, which tend to spark design debates. IWarp, which uses RDMA (remote direct memory access) over IP to reduce CPU processing on data center connections, is still the focus of intensive supplier-driven API development. Warning: Software interoperability issues loom.

Users, caught in the confusion and clamor, are opting to buy high-speed links only if they really need to -- and if they do buy, they're buying cautiously. The results aren't necessarily in keeping with some vendors' hopes.

Earlier this week, for instance, Huntsville, Alabama-based Colsa Corp., which provides high-performance systems and services to a range of U.S. Government agencies, revealed that it has chosen Myrinet, a longstanding and proprietary HPC interconnect, over either InfiniBand or Gigabit Ethernet to connect a massive cluster of Apple Xserve machines. (See Colsa Upgrades Apple Cluster.)

"Ease of operation and management was far better on Myrinet than anything else we evaluated," says Mike Whitlock, program director of Colsa's hypersonic missile technology team. He refuses to say what other vendors he considered before opting for Myricom's Myrinet-2000 technology.

Notably, Whitlock merely doubled his interconnect rate when he replaced Gigabit Ethernet switches with 2-Gbit/s Myrinet gear. For other users, too, there simply isn't a compelling requirement to move to 10- or 20-Gbit/s, especially if it means deploying all-new technology or risking interoperability woes.

"No, we haven't looked at higher-speed interconnects," says Steve Forstner, IT manager of systems engineering for the City of Richmond, Va. Indeed, the team just did a major virtualization upgrade to improve utilization of existing storage capacity. (See City of Richmond.) "The connection would be faster than we could make use of," Forstner says.

Some of the more aggressive interconnect players no doubt wish consumers like Whitlock and Forstner would remain silent. But it can't be denied that many users hesitate to buy some of the new faster links, either because they don't need them yet or don't trust them.

On the other hand, there *are* satisfied customers in the HPC space who are ready to testify to the virtues of InfiniBand and other fast solutions. One of these, an engineer at a national lab in the U.S. who asked not to be named, says he's used InfiniBand for years successfully in research clusters. Further, his group uses 10-Gbit/s Ethernet to link clusters over the WAN. He says he hasn't experienced interoperability problems yet and he's satisfied that vendors are working in plugfests to iron out any potential pitfalls.

Vendors at next week's show and elsewhere will work hard to make this kind of HPC user experience resonate with corporate customers. Whether they succeed or not will help determine the future of data center and storage networking.