Networks to Converge at 10 Gigabits:

SANs to Seamlessly Connect with LANs and MANs

AVAGO TECHNOLOGIES

White Paper

Today, general networking, storage networking and inter-processor communications have each developed different standards. As we move toward a 10 gigabit-per-second world, total cost of ownership and interoperability will drive the industry toward a single networking interconnect standard.

Networking Today

In general networking, including all Internet applications—e-mail, web services, downloads and multimedia, including voice and streaming video—interconnect technology is IP over Ethernet at 10 Mb/s, 100 Mb/s and 1 Gb/s. Although IP over Ethernet scales across the LAN, MAN and WAN, and dominates the industry in port count, interoperability and vendor support, the IP over Ethernet protocol stack (usually TCP/IP) is executed in the host CPU. This makes this architecture unsuitable for latency-sensitive applications such as data base storage and inter-processor communications.

Storage networking incorporates two applications. Network Attached Storage (NAS), which uses IP over Ethernet to transport data in *file* formats between storage servers and their clients, and Storage Area Networks (SANs), which transport *blocks* of data over Fibre Channel. Fibre Channel is the performance leader today at 1 Gb/s and 2 Gb/s link speeds, and offers excellent latency characteristics resulting from a fully offloaded protocol stack. This is one reason why Fibre Channel based SANs are often applied in performance sensitive applications while Ethernet-based NAS is used where cost and ease of use are more important.

There are also inter-processor communications (IPC) networks, which are used for server clustering. Although generally limited to high availability (HA) clusters, Ethernet is the most widely used IPC interconnect technology today. High performance parallel processing clusters tend to use a proprietary interconnect designed for very low latency. IPC is one of the most important applications targeted by InfiniBand. With an architecture optimized for low latency and high bandwidth, InfiniBand appears ideally suited for this application.

Starting down the road to 10 Gb/s

The Internet's insatiable appetite for performance will continue to drive Ethernet development to faster link rates at a quicker pace than Fibre Channel. Where 1-Gigabit Ethernet development leveraged mature Fibre Channel technology, Fibre Channel development will now leverage 10 Gb/s Ethernet standards. InfiniBand will be introduced at 2.5 Gb/s but will quickly move up to 10 Gb/s. See Figure 1.

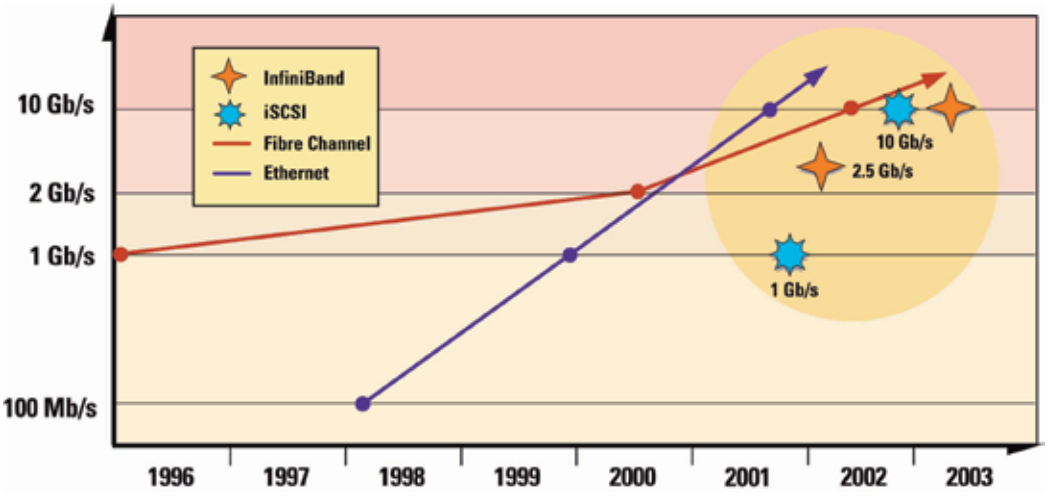

Figure 1: 10 Gb/s Ethernet products will be available first, with 10 Gb Fibre Channel, iSCSI and InfiniBand all expected to begin shipping in the late 2002, early 2003 timeframe. Note, though, that the PCI-X host interface will limit Ethernet and Fibre Channel to sustained throughputs of 5 Gb/s to 8 Gb/s.

For general networking, there are no serious challengers to Ethernet. CPU utilization will improve dramatically with the application of TCP Offload Engine (TOE) technology, which moves protocol stack execution from the host processor to the I/O card.

The choice of interconnect technology for interprocessor communications also seems to be fairly clear. Designed from the ground up for IPC and destined to be a standard I/O port on many chipsets, InfiniBand is well positioned to dominate large segments of this application. Ethernet with VI (virtual interface)/TCP/IP in offloaded hardware at 10 Gb/s will have excellent latency and throughput characteristics and could challenge InfiniBand for certain segments of the IPC application market.

The storage networking application is much less clear-cut than either of the others. Today, Fibre Channel dominates high performance storage networking applications because of the low latency, high bandwidth characteristics of the fabric. Even as the Internet Engineering Task Force enables block storage over TCP/IP through its iSCSI standards effort, Ethernet must have similar performance characteristics to Fibre Channel if it is to seriously compete for the storage networking application.

Today several companies are working to offload the iSCSI/TCP/IP protocol stack from the host processor. This will result in CPU utilization numbers similar to Fibre Channel. As these designs evolve, applying hardware acceleration to the protocol stack, latency will no longer be an issue at equal link rates. This technology will also dramatically reduce the latency of NAS, IPC and general networking applications. The offloaded iSCSI/TCP/IP design, when interfaced to the host system through an "IP Network Offload Engine" (iNOE) architecture allows networking traffic of all types to be offloaded and accelerated through a single I/O card. In the near term, Fibre Channel at 2 Gb/s will remain unchallenged in the data center while iSCSI enters the storage networking market through applications less sensitive to performance and high availability qualities. As all link speeds move up to 10 Gb/s, Ethernet will become a very serious contender for the entire storage networking market. In this same time frame NAS will begin to seriously challenge block storage for data base applications. See Figure 2.

The vision of InfiniBand is to become the single data center network. Storage would be connected directly to an IB switch with an Ethernet connection at the edge of the data center for LAN and WAN connectivity. Although very promising, IB is *all new* and is not likely to be a serious challenge to Fibre Channel and Ethernet in the near term.

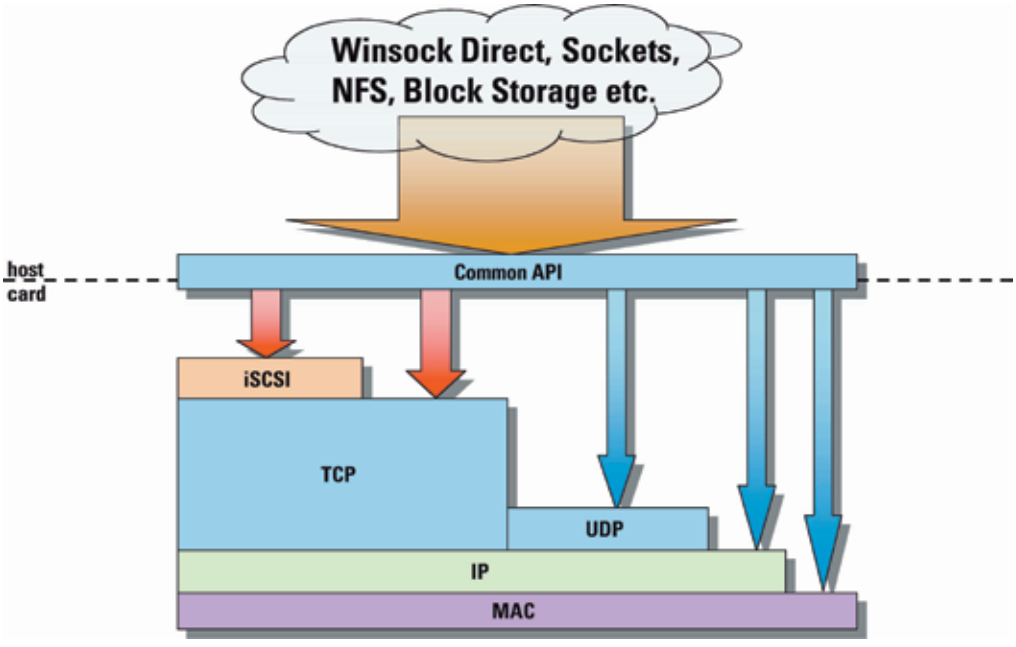

Figure 2: The IP Networking Offload Engine (iNOE) provides a fully offloaded, hardware accelerated multi-function architecture with simultaneous user interface for all popular networking protocols. This single solution will support high performance storage networking (both block and file storage), IPC and general networking.

Storage networking cross-talk

Given the Fibre Channel installed base and with FC over Ethernet and iSCSI contending for the storage networking market, the need for multi-protocol routers is clear. Several companies have announced bridges and routers, some including very powerful storage virtualization capabilities. Others are working to connect SAN islands through the use of a Fibre Channel over IP (FCIP) "tunneling" protocol. These developments, along with a strong commitment to interoperability by the iSCSI community, suggest that interconnecting networks will be easily accomplished.

When 10 Gigabit isn't really 10 Gigabit

First generation 10 Gb/s FC and Ethernet protocol chips, NICs (network interface cards) and HBAs (host bus adapters) will support a 10 Gb/s link speed, but will be limited to 5 Gb/s to 8 Gb/s by the PCI-X host interface. To achieve a sustained 10 Gb/s throughput a new interconnect standard is required. InfiniBand, PCI-X double data rate, Rapid I/O, Hypertransport, 3G I/O and others are contending for acceptance as the 10 Gb/s interconnect standard.

Peering in to the crystal ball

It appears that the *best* candidate for one networking technology providing transport for all applications is the IP Networking Offload Engine (iNOE). Using IP over Ethernet as a foundation, this architecture is a fully offloaded, hardware accelerated multi-function design providing simultaneous user interface for all popular networking protocols including TCP, IP, UDP, MAC, iSCSI and VI, as well as support for QOS. This single solution will support high performance storage networking (both block and file storage), IPC and general networking.

The world will not switch to a single interconnect technology for all applications very quickly. The application of iNOE and similar architectures to general networking is straightforward, will occur very quickly and will be unchallenged. IPC applications will have a choice between InfiniBand and Ethernet. If IB achieves the stated design objectives, iNOE will be applied to the portion of the market that is more sensitive to total cost of ownership and ease of use. Storage networking applications will continue to be dominated by Fibre Channel for several more years while iSCSI solutions mature. By the time that iSCSI becomes the preferred storage networking solution, the FC installed base will be very large, helping to drive the market for multi-protocol routers.

For product information and a complete list of distributors, please go to our web site: **www.avagotech.com**

5988-4858EN - October 25, 2007

AVAGO TECHNOLOGIES