

The Last-Mile Solution Lies Within the First 50 Feet

InfiniBand: Breaking the Data Center Bottleneck

Abstract:

As companies work to meet the growing demands of the Internet economy, the computing model is evolving to become more distributed. The distributed computing model challenges the scalability, reliability, availability, and performance of computer and storage networks. Meeting these demands requires balanced system architecture, with equally strong performance in the memory, processor, and input/output (I/O) subsystems.

InfiniBand™ allows information technology professionals to achieve the balance necessary to deliver the advantages of high-speed computer networks to end users. InfiniBand is a high-performance, channel-based I/O protocol that delivers increased performance, scalability and reliability for servers and other network infrastructure equipment, including storage systems and network peripherals.

Under the standards developed by the InfiniBand Trade Association (IBTA, www.infinibandta.org), the InfiniBand I/O protocol provides a universal interconnect fabric that is packet-based, enabling higher performance, lower cost, greater efficiency, and increased reliability and scalability than is possible in today's networks. These point-to-point linking technology features will benefit end users by delivering faster, less-congested and more robust data transport through computer networks and the Internet.

Introduction:

When the UNIVAC I debuted fifty years ago, experts of the day thought the now-antiquated monstrosity would satisfy every potential business need – and predict an election or two – but did not think of the first modern computer as a solution to any specific problem.

Since then, the computer industry has gathered problems as fast as breakthroughs. Technological innovations spawned a labyrinth of unexpected collateral challenges, not least of which is the fact that simple handheld and desktop computers can now process far more data than they can efficiently transmit.

So it is with the Internet, this generation's solution to filling information gaps and finding instantaneous information at our fingertips. The more accessible the Internet becomes, the more problems it creates; the more applications we develop to fulfill its potential, the more the information clogs up, or gets lost in, the data pipes.

The best answer we currently have for feeding the ubiquitous Internet is the data center – a conglomeration of servers and peripherals that stores and processes vast quantities of data while consuming vast quantities of space and power. Data centers also generate a lot

of noise, heat and frustration. Some of those resulting problems, such as the data center's large footprint, have been addressed with more efficient versions, or "server farms," in which "thin client servers" with streamlined form factors are replacing cumbersome server boxes. These thin clients take up less space, generate less heat, and, individually, consume less power.

But as the demands for data and bandwidth increase with wider deployment of computers and the emergence of broadband, so do the sheer number of boxes in the data centers and server farms and the number of devices with which they must communicate.

Enter Infinite Bandwidth: InfiniBand

In a world where "cutting-edge" is overused and rarely delivered, everything about InfiniBand shouts it. Even its specification, which was not submitted to ANSI or any other traditional standards bodies for fear of hindering time to market, was created in a remarkably short time frame shortly after the competing Next-Generation I/O (NGIO) and FutureIO organizations merged to become the InfiniBand Trade Association. The standards development process was led by such industry heavy-hitters as Compaq, Dell, Hewlett-Packard, IBM, Intel, Microsoft and Sun, and now joined by over 200 firms, each weighing in with the benefits of a smooth process to develop and deploy a faster I/O technology.

Benefits You Can Really Plug Into

Data center players – computer, storage and networking companies – have been waiting for a technology to boost data transfers. InfiniBand 1.0 specifications allow scaling from 2.5 to 30 gigabits per second, which explains its descriptor of being the "next step" after such existing interconnection technologies as Fibre Channel, Gigabit, Ethernet, PCI and SCSI.

InfiniBand architecture boosts data transfers among servers, storage and peripherals by shifting I/O control from within processors to a point-to-point, full-duplex interconnect that is channel-based and built with a universal switched fabric. This enables higher performing, more reliable networks and data centers with critical scalability. There is also a lower cost of ownership associated with each of these features, since overall usability is increased and downtime is minimized.

Communicating by a packet-based network means that, rather than relying on a processor, as PCI does, InfiniBand employs intelligent I/O engines, referred to as channel adapters. These series of interconnected channels form a fabric, which physically combines connectors, cables and switches using one or more host channel adapters (HCAs).

Taken individually, think of an I/O fabric as a matrix – interwoven communications strands and links providing multiple, redundant patterns of data flow. The I/O fabric is controlled by switches, and -- with the connecting devices and peripheral devices -- makes up the high-speed network.

Next, InfiniBand switching technology links HCAs to target channel adapters (TCAs). The switch operates between the HCA, which generally resides near the processors, and the TCAs, which support the storage and peripheral devices, to manage and direct data packets. De-coupling the I/O subsystem from memory by utilizing channel-based, point-to-point connections rather than a shared bus, load-and-store configuration means the devices connected to the fabric don't have to share bandwidth, so they can communicate at maximum speed.

The technological innovations in the InfiniBand specifications deliver several critical improvements over today's interconnection technologies:

Performance – Today, no single I/O solution is optimal for every use. As much as these standards have improved and evolved over the years, none provide true interoperability or manageability, and therefore place tremendous burden on IT managers to understand, operate, and maintain multiple I/O systems. Likewise, none have been able to keep up with the tremendous increases in processing power.

In contrast, today's InfiniBand specifications support up to 30 gigabits per second of bandwidth, with the promise of even greater speeds to come. Lower latency, built-in security, and quality of service are additional performance and efficiency benefits.

Within the data center, InfiniBand will connect all the various components with a single, easy to use I/O technology that allows processors, storage, and peripherals to communicate at speeds that break through the bandwidth bottleneck.

Efficiency – The estimated cost of data centers averages about \$300 per square foot, forcing companies to address space requirements by wringing out as much revenue per square foot as possible. InfiniBand enables a denser computing environment, allowing for more profitable and efficient management of space and maintenance because of greater computing power per square inch.

Likewise, experienced personnel and available parts affect up-time and maintenance in data centers that have multiple, complex interconnection technologies. The InfiniBand fabric will coexist with, and eventually replace, existing I/O technology standards such as PCI, Fibre Channel and many proprietary standards as data center managers optimize the performance of each unique I/O demand.

And finally, as InfiniBand enables greater computing power in smaller space, it also enables greater computing power at lower rates of energy consumption.

Reliability – By working in addressable packets, data can be routed to any application once it hits the InfiniBand fabric, without the linear and dogged routing of existing protocols.

Robust performance is another critical aspect of reliability. At broadband speeds, data loss must be minimized. The open nature and ready availability of the IBTA's technical specifications contribute to innovation by being accessible for rapid development and deployment of software and hardware. As the industry migrates toward the InfiniBand solution, compliance and interoperability are critical at every stage of development, and InfiniBand's open architecture paves that way.

Scalability – Scalability is addressed both by InfiniBand's universal fabric and its physical modularity. As mentioned earlier, the InfiniBand I/O fabric easily handles workload increases that create latencies in shared-bus I/O architectures.

InfiniBand is a "hot plug" technology, allowing customers to pay only for what they need, instead of trying to anticipate future growth and buying excess capacity up front. Any resources, be they storage or processing components, can be added on demand by plugging them in, without taking down other systems. This makes it far easier to add, remove, maintain and upgrade an InfiniBand system than today's shared-bus I/O systems.

Current data center technology is slow, relatively unstable and highly demanding of IT professionals. In addition to taking up a lot of space, fixed performance levels, limited speed and finite scalability, Internet server farms are less efficient and less stable than they should be partly because I/O infrastructure has not kept up with the processing power it connects. Each server needs a complete I/O solution, increasing total cost of ownership. The multiple protocols in today's data centers make them difficult to manage and create unacceptable quality of service challenges for end users.

Conversely, the InfiniBand data center is dense, is easy to manage by hot-plugging new equipment, maintains a single network based on an open standard versus multiple proprietary ones, and allows for interoperability across vendors. It is scalable, with unlimited sever expansion, is less expensive since servers are sharing I/O systems, has higher performance to address bandwidth and latency issues, and is thus overall more reliable and more cost-efficient.

The Customer and End User

As end users have demanded more from their Internet, they have upgraded hardware and software systems to deliver rich, dynamic content – the highest performance processors, broadband modems, the fastest corporate networks and intranets. All this processing power has merely forced the bandwidth bottleneck backwards from the last mile to the data center – the first fifty feet of the Internet. In fact, these advances in technology

aggravate the challenges as more users become aware of them and develop expectations that exceed the bandwidth available even with their new tools.

High performance systems – such as those delivering high-definition video and other broadband services on demand – realize the potential of InfiniBand.

Potential Customers

So who would benefit from such a solution today?

Likely customers are financial centers and other data-intensive companies that must store and process vast quantities of data. Institutions with high broadband loads, such as universities and research and development organizations are also probable early adopters of InfiniBand-powered equipment.

Much of this equipment is already in development. QLogic was the first announced InfiniBand customer, and is developing next-generation InfiniBand-based SANblade channel adapter solutions using Banderacom's IBandit family of target channel adaptors.

Deployment: How Close Is It, and What Will Take Us There?

According to the research firm IDC, more than 4 million InfiniBand-enabled servers will ship in 2004. This will make up the bulk of the entire server market, and as deployment in servers, or hosts, increases, demand for the targets will naturally increase significantly. Once InfiniBand finds its place in server clustering, it will evolve to high-speed server I/O, then server-to-server and server-to-storage communications, and as the primary I/O subsystem in WANs, LANs, SANs and MANs within three years. Some predict migration of InfiniBand into high-end desktop computers by 2005.

InfiniBand evangelists believe that migration is inevitable, but there are still barriers to implementation – cost, availability and interoperability are examples. That said, most industry experts and analysts put early deployment in mid-2002 and full-scale deployment at 2003, as increased production lowers costs.

The InfiniBand Trade Association and its members are successfully addressing interoperability on two fronts: first, demonstrations at several networking and InfiniBand trade shows have tested InfiniBand's connections to PCI, Ethernet, Fibre Channel and TCP/IP; second, IBTA-sponsored working groups and plugfests, as well as the trade shows, have demonstrated interoperability among various InfiniBand hardware and software developers.

InfiniBand's end user benefits are tremendous, and will have an immediate impact. First generation InfiniBand fabrics will handle communications as fast as 2.5 gigabits/second between any two devices, with speeds expected to rise rapidly to 10 and 30 gigabits/second as the technology evolves over the next 12 to 18 months. For some

applications, InfiniBand is already ten times faster than the best possible connection currently available.

Technology is, by definition, disruptive to innovators and implementers. To be successful, though, technology must not be disruptive to the end user. By InfiniBand's nature, its deployment will take place in a transition, easing integration and minimizing the disruptive effect on the customer and the customer's customer. Ultimately, end users receive the greatest benefit through improved price and performance, greater flexibility and scalability, and more reliable and manageable data centers – though InfiniBand's most significant technological changes will be largely transparent to them.

Unlike many of today's solutions, InfiniBand isn't a technology looking for an application. It was envisioned, and is currently being developed, as the answer to a question critical to the success of the Internet economy – how to break the bandwidth bottleneck in the first fifty feet of the network infrastructure.