



A SKY Computers White Paper

InfiniBand™ Gains Momentum

By: Steve Paavola, SKY Computers, Inc.

*SKY Computers, Inc.
27 Industrial Ave.
Chelmsford, MA 01824
978 250 1920
978 250 0036 FAX
www.skycomputers.com*

InfiniBand is a “clean-sheet” specification, developed from the requirement to handle very complex applications and ever larger databases. It provides high bandwidth, low latency, low processor overhead, simplified network management, reliability, availability and serviceability, all at a reasonable cost.

InfiniBand technology is ideal for high performance embedded computing (HPEC) applications because it provides the aggregate throughput needed by HPEC systems. In addition, as a complete fabric with supporting software, it enables high application availability (HAA). At full maturation, HAA will provide up time approaching 99.999%. Once operating systems, software toolkits and the InfiniBand standard come together, systems will be self-monitoring and self-healing. For example when there is a hardware, software, or user error, the system will be smart enough to re-route the processing around the error. No other interconnect fabric has the potential to provide this level of fail over response.

As InfiniBand enters mainstream computing, the rate of product announcements is increasing. Some current examples include:

- InfiniCon has announced availability of its I/O system
- InfiniSwitch has already shipped over 1,000 switch ports
- Mellanox has announced availability of its server blade platform
- VIEO has announced a system management blade

As InfiniBand continues to gain momentum this list will grow and change as member companies bring new products to market.

On the end user side, Sandia National labs is the first national lab to purchase a complete InfiniBand fabric. Sandia will use InfiniBand in development projects with Ohio State University, Ohio Supercomputer Center, MSTI and Argonne National Laboratory.

Cornell Theory Center, a high-performance research lab at Cornell University recently set up a numerical analysis project using industry-standard computing equipment. The aim of the project is to promote this technology for financial applications that traditionally rely on proprietary, high-priced systems for reliability. In order to achieve the required levels of reliability on industry standard equipment, the Center will use InfiniBand clustering, using InfiniSwitch as the interconnect technology.

InfiniBand Availability

Hundreds of engineers in the world's largest and most successful computer companies have been working on InfiniBand hardware and software for several years. InfiniBand fabric components are already in the second generation of products.

Second generation 8 ported 4X switch chips have been shipping for several months from IBM, Mellanox, and Redswitch. Also, second generation interface chips implementing 4X links with data rates of up to 10Gbits/s are now available from Banderacom, IBM, and Mellanox. These use four serial interfaces ganged together to provide four times the throughput of the first generation 1X links for an aggregate 2 GB/s throughput. This is eight times the bandwidth of Gigabit Ethernet, and four times the bandwidth of the latest Fibre Channel interfaces.

On the software side, the chip vendors such as IBM and Mellanox have always included software tools for their products. Now, Lane15 and VIEO are shipping sophisticated fabric management software to further extend the capabilities that InfiniBand offers. Lane15 has announced that they've already shipped 400 units of its fabric manager. VIEO has announced a complete suite of software tools that support the latest 4X technologies as well as a blade dedicated to supporting InfiniBand network management. As InfiniBand moves into the computing mainstream, more and more products are becoming available. Host Channel Adapter (HCA) boards are available in PCI and PMC form factors from Infiniswitch, JNI, Mellanox, and VMIC. InfiniBand switch boxes are available from FirstStar Networks, Infiniswitch, Mellanox, Paceline Systems, and Redswitch. End user systems are now available from Network Appliance, and Hitachi is working with Voltaire to produce solutions. Dell plans to integrate InfiniBand technology into its next generation PowerEdge "brick" servers with Intel's Xeon processors. Major computer companies such as HP, IBM, and Sun are expected to make announcements soon.

Embedded Applications

While InfiniBand was developed from the needs of the IT data center, those same needs exist for high performance embedded computing (HPEC) applications as well. Initially as HPEC vendors, like SKY Computers, considered new interconnect choices their selection criteria was based primarily on aggregate throughput performance. However, HPEC customers are now demanding:

- standard hardware and software solutions which allow easy migration of their application from platform to platform
- high system availability with a common uptime requirement of 99.999% (5 9's).
This translates to about 5 minutes per year that the application can be unavailable due to hardware, software, or user error

InfiniBand is a smart system architecture providing: system performance, system reliability, application resiliency, and high application availability. HAA addresses hardware, software, and application errors by providing a self-reliant, self-repairing interconnect for the HPEC system.

InfiniBand Availability

There is no doubt about InfiniBand's importance as a technology. In addition to the needed performance for both HPEC and IT, it is the only interconnect to provide a SMART architecture. As a SMART architecture InfiniBand is:

- Secure - provides guaranteed data communication integrity.
- Manageable – A series of hardware and software technologies provide system sensing, monitoring and management capabilities (proactive & reactive)
- Adaptable – System and application adapts “on-the-fly” to changes due to failures or addition of resources.
- Real-Time – The architecture and capabilities designed and tuned to signal & image processing.

InfiniBand has the backing of every major computer company. For example, Sun Microsystems, the leader in providing high performance engineering computers has

announced that ultimately all Sun servers using 4-to-8 or more processors will use InfiniBand at 10-Gbit/second and 30-Gbit/s (so-called 4x and 12x) speeds.

According to the Yankee Group, Boston, sales of InfiniBand-based switches and routers will grow from \$32 million this year to more than \$1.53 billion in 2006. Companies including Banderacom, InfiniCon Systems, Lane 15 Software, Paceline, Voltaire and IBM Microelectronics to name a few, are all continuing to move forward building silicon, systems and software. While there are reports of major vendors scaling back their plans for InfiniBand due to the sagging IT economy, the technology is being positioned into specialized areas-especially those that benefit from a low-latency, high-performance interconnect like the HPEC markets that SKY Computers serves.

Other technologies aren't standing still, but aren't yet competitive. For example, there is work underway to specify RDMA (Remote Direct Memory Access) for TCP/IP, which would enable Ethernet applications to reduce their processor overhead. This is one of the important capabilities already available with InfiniBand. The approved RDMA specification for TCP/IP is still nearly a year away making it two to three years behind InfiniBand. Integrating RDMA with TCP/IP is a difficult task.

In addition, there are some high speed chip-to-chip interconnects like HyperTransport, RapidIO, and PCI Express; and PCI extenders like StarFabric. Today, only HyperTransport and StarFabric are shipping. These fabrics are all designed to enhance or compete with PCI and don't address the system management and RAS, (reliability, availability, serviceability) requirements that InfiniBand provides.

InfiniBand technology is ideal for high performance embedded computing (HPEC) applications because it provides the aggregate throughput needed by HPEC systems. It

provides the high bandwidth, low latency, low processor overhead, simplified network management, reliability, availability and serviceability, required for high performance computing applications. No other interconnect fabric goes to this level. As a smart fabric with supporting software, it provides HPEC users with high application availability which ultimately translates to lower cost of ownership and quicker deployment of embedded systems. SKY Computers will join the ranks of the world's largest and most successful computer companies when it brings its InfiniBand-based smart systems architecture products to the market in early 2003. These systems will provide HPEC users with unprecedented application performance and reliability, all based on broadly accepted and adopted industry standards.