Mellanox Technologies

White Paper

Comparative I/O Analysis

InfiniBand Compared with PCI-X, Fiber Channel, Gigabit Ethernet, Storage over IP, HyperTransport, and RapidIO

1.0 Overview

There has been much discussion on how InfiniBand™ complements, replaces or compares with a variety of I/O standards. This white paper provides a basic understanding of each I/O technology, how each compares with InfiniBand, and it's position in the market with reference to InfiniBand. A matrix of features incorporated within these technologies is included in this document.

2.0 Brief Overview of InfiniBand

InfiniBand is a point-to-point high-speed switch fabric interconnect architecture that features built-in Quality of Service (QoS), fault tolerance and scalability. The InfiniBand Architecture (IBA) Specification defines the interconnect (fabric) technology for interconnecting processor nodes and I/O nodes to form a System Area Network that is independent of the host operating system and processor platform. InfiniBand defines 1X (2.5 Gb/s) links, 4X (10Gb/s) links, and 12X (30Gb/s) links. It also uses IPv6 as its native network layer. Mellanox™ Technologies has been shipping 10Gb/s (or 4X links) InfiniBand capable silicon since early 2001.

This technology was created by the InfiniBand Trade Association (IBTA), which includes all of the major server vendors and Microsoft® to provide a logical successor to the shared PCI bus on servers. While creating a new I/O for servers, the IBTA determined they could also create a highly reliable fabric for data centers based on the same technology. Therefore, InfiniBand extends out of server motherboards over copper or fiber links as a new interconnect (called a "fabric") for data centers. InfiniBand is built from the ground up for Reliability, Availability, and Serviceability (RAS) for Internet and Enterprise data centers[1]

1. See Mellanox's white papers: "Introduction to InfiniBand" and "InfiniBand in the Internet Data Center", and for more information, visit http://www.mellanox.com/products.

3.0 InfiniBand and PCI-X

Peripheral Component Interconnect (PCI) is a bus standard designed to provide a low cost interface to allow the connection of peripherals into personal computers. Since its inception, PCI has scaled to PCI-X, which has bandwidth of up to 8.5 Gb/s (half duplex), but is implemented by many at 100 MHz (vs. 133 MHz) for 6.4 Gb/s. PCI has been very successful and is in virtually every personal computer and server shipped today. The industry does not expect PCI to exit the personal computer market any time soon, but PCI bandwidth capabilities are not able to keep up with heavy loads that servers place on it. Servers today host SCSI cards (soon Ultra320 SCSI), 1Gb Ethernet, Fiber Channel, Clustering cards and more. PCI cannot keep up with the I/O bandwidth required by these devices. So, as mentioned, the IBTA was formed to develop an I/O specification that can handle the bandwidth needs of today's servers. The IBTA specification was released in October, 2000, and subsequently, silicon was released by a number of established and start-up silicon companies.

The first implementations of InfiniBand will be attached through PCI and PCI-X cards. InfiniBand will leverage PCI for its host and target channel adapters to create the first InfiniBand fabric. Over the next year, InfiniBand will be supported on the host memory controller (or directly to the host processor) via newly architected I/O buses (e.g.: HyperTransport™, RapidIO™, etc.) that offer greater bandwidth than existing PC component interconnect buses. Buses like these will be required for servers and the embedded market, since most system controller I/O buses are limited to 10Gb/sec or less. Greater bandwidth is required as InfiniBand pushes past these limits. For example, Mellanox's currently shipping InfiniBridge MT21108 supports two 4X links that can create up to 40 Gb/s[1] of InfiniBand bandwidth; PCI limits excluded.

Mellanox has engineered a creative solution, called InfiniPCI™ Technology, to assist in the immediate migration of PCI out of the server and onto InfiniBand. InfiniPCI Technology creates a transparent, standard PCI to PCI bridge over InfiniBand. This allows existing PCI I/O cards in servers to be replaced by a single InfiniBand HCA, and the other I/O cards moved to an external or remote PCI chassis without any changes to the PCI hardware, device drivers, BIOS or Operating Systems.

4.0 InfiniBand and Fiber Channel

Fiber Channel was designed in the early 1990's as a bus standard designed to provide a high speed physical layer to meet the demands of storage I/O applications. It was released at 1Gb/s and is now shipping at 2Gb/s with future plans for faster technologies. The key benefits of a Fiber Channel Storage Area Network (SAN) are that it allows block transfers of SCSI data at 800Mb/s versus (at the time of FC introduction) 100 Mb/s Ethernet, and enables an independent SAN that relieves Ethernet network congestion by backing-up files with minimal performance impact on the network. SANs can also be scaled or maintained without impacting the operation of the data center.[2]

1. The 4x link bandwidth is 10 Gb/s but it is bi-directional or fully duplexed. Therefore, data can run in both directions at 10 Gb/s per link or 20 Gb/s, fully duplexed. Thus, dual 4x links can support up to 40 Gb/s of InfiniBand data).
2. For more information visit http://www.fiberchannel.com/

The InfiniBand architecture incorporates a super set of Fibre Channel mechanisms. Since the benefits offered by Fiber Channel data centers have been demonstrated in the market, this paper will briefly examine some of the key benefits that InfiniBand technology offers versus Fibre Channel. These include flexibility, Quality of Service (QoS), RAS, lower cost, and better performance.

4.1 Flexibility

InfiniBand has the flexibility to be used for multiple technologies that include server-to-server communication (IPC for clustering), switching, and storage as well as in-band management functions across the InfiniBand wire (unlike Fiber Channel that must use Ethernet). One key to this flexibility is, only Infiniband architecture has the necessary mechanisms to support the integration of the System Area Network by enabling “virtual fabrics” (through virtual lanes) to carry each type of traffic. This flexibility enables InfiniBand to be the “single” interconnect or fabric for the data center. Once achieved, this will greatly reduce the complexity of the data center, as only one set of wires will need to be managed and maintained.

4.2 Quality of Service

Infiniband implements key QoS features such as Virtual Lanes and end-to-end flow control enabling the implementation of robust fabric. These features ensure that should the network get congested, critical or priority functions will not get bogged down. Tasks such as back up will be assigned a low priority, and can be executed only once higher priority traffic had been handled.

4.3 Reliability, Availability, and Serviceability (RAS)

The InfiniBand Architecture has been developed from the ground up to provide a cost effective, high performance solution with RAS support for the Internet and Enterprise Data Center. The architecture supports many RAS features including Cyclical Redundancy Checks (CRCs), reliable transport, and failover. The InfiniBand link protocol incorporates multiple CRC fields, providing error detection capabilities, on both, a per-hop link level and an end-to-end basis. A second RAS feature is the InfiniBand transport support of reliability for both connection and datagram services. The InfiniBand architecture defines a failover mechanism, which allows a network to heal itself if individual links fail. InfiniBand subnets, which incorporate redundant connections, can detect link errors and migrate traffic from a failed link to a redundant link that is still functional. InfiniBand supports both managed and un-managed failover. Un-managed failover is essentially a hardware mechanism whereby traffic is switched to a pre-allocated redundant channel automatically, as errors are detected. Managed failover is a function of the subnet manager detecting errors in the fabric and reconfiguring the forwarding tables of individual switches.

4.4 Cost

InfiniBand will impact server costs through many different aspects. These include: infrastructure (cables), reduction in I/O devices, power, density, and performance.

Infrastructure: InfiniBand is the first technology to achieve 10Gb/sec (4x links) over serial copper connectors. Copper cables and connectors allow the lowest cost solution for connectors and cables for this very-high bandwidth. Copper transceivers generally cost hundreds of dollars less than fiber optic transceivers. Of course, InfiniBand supports fiber optical connections when needed for long connections such as building-to-building or site-to-site.

I/O Cost Savings: InfiniBand enables servers to connect to the data center with only one high performance I/O interface. By utilizing a fabric design, the current I/O devices or PCI cards in a server can be moved out of the physical server chassis to a remote location or eliminated all together. Clustering cards are eliminated, as the Virtual Interface Architecture is enabled in hardware. Ethernet NICs are not required in this environment as Ethernet is moved to the edge of the data center. Host Bus Adapters (HBAs) for Fiber Channel are eliminated as InfiniBand offers better bandwidth and can host native Storage Area Networks or use InfiniBand to Fiber Channel gateways to support existing SANs.

This disaggregation of I/O from the server enables configurations that require only a single PCI InfiniBand card that result in the ability to better utilize the 1U form factor. Also, with the advent of “server blades” even smaller form factors and greater density can be achieved in server racks. Server blades are a technology under development that utilize a CPU, system controller, memory and InfiniBand to create a fully functional server board in a form factor as small as 4 x 6 inches.

Power and density are improved with the reduction in I/O devices and the reduction in form factor. Mellanox’s new dual 4X port reference card provides 40 Gb/sec of bandwidth yet requires less than 10 Watts of power. This lower power and improved density allows a better return on the cost per square foot of a data center.(See Mellanox’s PPPV Metric for more information.)

4.5 Performance

Improved performance to 10Gb/sec will eliminate the need for additional servers to reach the same performance levels in clustering and 10Gb/sec will considerably enhance the bandwidth of a SAN (versus 1 or 2 Gb/sec Fiber Channel or Gigabit Ethernet). The Virtual Interface Architecture enables InfiniBand to RDMA data directly into memory, where it is needed most. InfiniBand transport mechanisms are implemented in hardware, so there is little need for host compute cycles, as compared to TCP.

IDC has recognized all of these benefit InfiniBand, and estimates InfiniBand capable servers will reach 80% of the server market in 2005 (May 2001). So, the adoption of InfiniBand is no longer a question and this has strong implications on the future of Fiber Channel.

Native InfiniBand-enabled servers will provide optimized service connections to SANs without the requirement for expensive Fiber Channel HBAs. As InfiniBand disaggregates the I/O from the server, it will be a natural migration to move storage onto the InfiniBand fabric. This means that InfiniBand will supersede Fiber Channel in SAN applications over time as Fiber Channel HBAs are eliminated from the servers. After all, why sustain a point-to-point connection, like Fiber Channel when one of the points is InfiniBand? InfiniBand is designed to support the requirements of SANs and will complement existing FC Storage Area Networks until they are replaced. An InfiniBand based SAN supports a super-set of features offered by legacy FC without compromising the operation of the network. During the transition period servers, gateways or hybrid switches will facilitate communication between InfiniBand fabrics and Fiber Channel SANs. A completely native InfiniBand Fabric will evolve over time allowing for greater simplicity: one fabric, based on one cable, single fabric management software, built-in in-band management and higher performance versus Fiber Channel.

5.0 InfiniBand and Ethernet

Ethernet was designed to enable a local area network for computers. Today, the desktop Ethernet is at 10/100 Mb/s and the backbone is quickly moving to 1Gb with 10Gb on the horizon for MANs. Ethernet is ubiquitous in computing today, and enjoys an extensive and well-developed infrastructure of switches, routers, software and more. Ethernet was designed for connecting disparate systems together into a network. To move data, TCP/IP is required to run over Ethernet and TCP is implemented in software. This is great for flexibility (modifications), but it requires extensive support from the host CPU to execute the stack. TCP also uses slow-start algorithms, unreliable service, re-order buffers, and out-of-order delivery, all of which add complexity and overhead for an uncontrolled environment like the Internet. Other tasks such as calculating the checksum and reassembling the original byte stream (sent out-of-order) add to the processing overhead. These complexities are not necessary for a controlled environment like a data center.

Currently 10 Gb/s Ethernet (10GE) deployment is targeting the MAN market were fiber is the standard interface. Once the standard is approved, sometime in mid-2002 the industry will begin to see this deployment. Since 10GE isn’t targeting desktops or the data center, it’s not currently looking at the lower cost copper connections. Also, since 10GE is natively more expensive, and targeting MAN installations, it will not enjoy the economics of scale of the PC market Ethernet has enjoyed in the past. 1GE is available as a NIC and can go into servers with copper connections, but this offers less bandwidth than the current rate of either Fiber Channel at 2Gb/sec or InfiniBand at 2.5 or 10 Gb/sec. Also, Ethernet still suffers from the need for a great deal of the host processor compute cycles. For storage applications, a single SCSI 3 device can generate up to 1.3 Gb/s (160 MB/s) of bandwidth, Ultra320 SCSI will generate twice that bandwidth and multiple SCSI devices coupled together generate even greater bandwidth. The implication is that Gigabit Ethernet does not keep up with the performance needs for storage and to get any bandwidth overhead 10GE would be required. Anyone running mission critical applications or on-line transactional processing would likely see Gigabit Ethernet as a bottleneck in their data center.

InfiniBand is designed to be a transport service independent of protocol. InfiniBand’s ability to use a single communication technology for storage, networks, audio/video, or to move raw data is superior to the common frame feature. It also implements the layer 4-transport service in hardware and does not require the overhead of network layer 3 support when operating within a single data center subnet (InfiniBand’s most typical configuration). And only InfiniBand supports the memory semantics required to closely couple multiple CPUs into a cluster. Both InfiniBand and Ethernet can be utilized for multiple technologies but only InfiniBand silicon is designed from the ground up to be the single fabric for the data center.

6.0 InfiniBand and Storage over IP

Storage over IP is a generic term that describes various protocols which operate over an IP network (typically implemented on 100Mb/s or Gigabit Ethernet) using SAN mechanisms (block level) rather than the file level used in Networked Attached Storage (NAS). There have been several proposals to accomplish this, which are: iSCSI (SCSI over IP), FCIP (FC tunneled into IP), and iFCP (maps FC ends devices to IP). For brevity, this discussion will focus on iSCSI as the means for storage over IP. iSCSI is an attempt to allow traditional IP networks to support SAN technology. The fundamental advantage claimed by proponents is the enormous installed base of IP routers and switching elements within the Internet. However, the specification for iSCSI is not expected to be completed until early 2002. And it is interesting to note that storage over IP advocates believe that their method will make storage management simpler, but it is difficult to comprehend how having three competing storage over IP standards will make life simpler.

As discussed in the previous section, Ethernet creates some serious bandwidth issues for Internet and Enterprise data centers that demand high bandwidth. Data centers require not only bandwidth, but security, CPU performance and RAS. Since Fibre Channel was created to eliminate congestion or bottlenecks on the network, and iSCSI (supposedly) puts that traffic back on the network, the concept of moving storage over ethernet comes across as a contradiction.

The proponents of iSCSI tout the following as their value proposition:

- Runs over existing equipment
- Allows re-integration of the data center network
- Uses well known management software
- Costs less than other storage technologies

This paper will examine each of iSCSI's claims:

1. Runs over existing equipment: This seems like an odd proposition. Since the original Fiber Channel SANs were created, in part, due to “too much” congestion on the network, it seems odd that anyone would want to recreate that problem by bringing storage back into the network. The only real way to overcome this problem is to either completely upgrade the existing network to the next higher Ethernet speed or create a separate Ethernet network. Of course, this means it can not run over existing equipment. It has already been mentioned that 1GE can't keep up with the bandwidth needs to SCSI, so the only option is adding an new 10 GE network for storage.

   If storage over IP is the choice, how does the storage convert it's block information onto Ethernet? To achieve this it will be necessary to add a gateway (FC to Ethernet) for existing SANs, or incorporate new Ethernet interfaces into storage devices. Therefore, new 10GE NICs will be needed in servers and new fiber cables will be needed as 10GE won't run on the existing copper (CAT5) cables. Also, since the TCP/IP stack runs on the host processor, new TCP (and even iSCSI) offload engines will be required to achieve any degree of performance for either Ethernet or the applications running on the server. It is also not unreasonable to expect that the increase in traffic may create congestion problems in Layer2 Ethernet switches that would probably require new Layer3 switches to resolve.

So much for the “claim” that iSCSI uses existing equipment. It seems like the data center manager will need expensive fiber cables, 10GE NICs in the servers, along with a TCP and/or iSCSI off-load engines, and either a gateway box or new storage interfaces to talk to the disks. So much for the iSCSI claims that it runs on existing equipment.

2. Allows re-integration of the data center network: IP is based on an unreliable connection that uses “best effort” to move your data from one point to another. In comparison, InfiniBand implements a reliable, in-order, connection transport service implemented in hardware. The lack of native QoS mechanisms within Ethernet should make everyone wonder if they would ever consider it as a replacement for Fiber Channel. But proponents of iSCSI state that QoS can be easily solved simply by over-provisioning the network, and recommend separate networks for both storage and communications. This creates the nasty liability of having two networks that are required to be kept separate. Invariably, the system admin will inadvertently connect the storage network to the communications network and they will become hopelessly intermeshed. Any credible argument of a unified data center network had better have a technology like InfiniBand to support multiple “virtual fabrics” to allow both storage and communication to co-exist peacefully[1]. It appears, to deploy an iSCSI solution, a totally new 10GE network is required and it needs to be dedicated to storage (shouldn’t carry the existing traffic). This looks more like a duplication effort rather than iSCSI’s claim of re-integration.
3. Uses well known management software: Ethernet is a well-known technology and does provide a multitude of management tools. But the reality is, the real challenge of managing storage networks has nothing to do with the basic management of networks and switches, but rather the management of storage virtualization. This challenge is not alleviated by using an IP network because virtualization, LUN management, security, and zoning are still required. The learning curve for system admins has been learning both Fiber Channel and storage management. It is important to understand that InfiniBand supports a super set of Fibre Channel mechanisms and even though InfiniBand management tools do need to be developed, this feature should allow a quick port of virtualization and other Fibre Channel tools to InfiniBand. iSCSI claims can confuse IT managers into thinking existing Ethernet tools can be used to manage storage, but this is not the case. InfiniBand’s architecture, again, offers a straight forward approach to address the issue by supporting Fibre Channel mechanisms.
4. Costs less than other storage technologies: As mentioned above, it looks like iSCSI not only requires a duplicate network but also new fiber cables, 10GE NICs, expensive TCP (and/or iSCSI) offload engines and either gateways or new storage devices. Implementing iSCSI isn’t a simple task that can generally be done on existing networks. iSCSI may claim they are a cheaper solution than other options, but considering all the costs, this is not an inexpensive solution. Also, it only services part of the data center when compared to InfiniBand.

Three recent articles discuss the target applications and limitations of iSCSI. First: “Commentary: Cast a skeptical eye on iSCSI” by the Meta Group (http://news.cnet.com/news/0-1003-201-5559972-0.html). And two articles from InfoSTOR assert that iSCSI is correctly positioned as “remote storage”. (See InfoSTOR: April 2001 Volume 5, No.4: “IBM readies iSCSI, NAS devices” (page 1) and “Emulex outlines iSCSI HBA plans”, (page 12)). These articles state that iSCSI is appropriately targeted in relatively low-end environments for remote backup.

1. See Mellanox’s white paper “Implementing Virtual Fabrics with InfiniBand” for more information

So, there may be a market for iSCSI (and/or the other storage over IP standards) as a way to provide remote storage or remote storage connections, but given the design of Ethernet (unreliable connections) and the performance demands of today's data centers, it doesn't seem likely that iSCSI could create a highly reliable and performance minded Storage Area Network. Or iSCSI doesn't seem to be the low cost and re-integrated solution that is so widely being discussed.

InfiniBand's architecture uses IP as its native networking layer. As 10Gb Ethernet is deployed in the MAN, it can be placed on the edge of an Infiniband data center and IP can be used for block transfers from InfiniBand data centers to remote locations. Or the InfiniBand architecture has the native capability to tunnel SCSI blocks over IP. Therefore, InfiniBand does include the key benefits that make storage over IP interesting, namely remote backup or remote storage.

7.0 InfiniBand and New Internal Buses: HyperTransport™, RapidIO™ and 3GIO

"HyperTransport™ technology is a new high speed, high performance point-to-point link for interconnecting integrated circuits on a motherboard. It can be significantly faster than a PCI bus for an equivalent number of pins. HyperTransport was previously code-named Lightning Data Transport, or LDT. HyperTransport technology was invented by AMD and perfected with the help of several partners throughout the industry. It is primarily targeted for the IT and Telecomm industries, but any application where high speed, low latency and scalability are necessary can potentially take advantage of HyperTransport technology[1]".

"The RapidIO™ architecture is an electronic data communications standard for interconnecting chips on a circuit board and circuit boards using a backplane…. An important bottleneck in networking and communications equipment is the speed at which the various components 'inside the box' communicate with each other. The RapidIO architecture eliminates this bottleneck. Current equipment is limited to hundreds of Mbits per second transfer rates using legacy bus technologies such as PCI. The new RapidIO interconnect increases this bandwidth significantly. Many believe that increases in bandwidth have already replaced increases in microprocessor performance as the key requirement for higher-performance Internet technology[2]."

There has been some discussion about Intel's next generation system bus interface called 3GIO or 3rd Generation IO. To date, Intel has released NO public information on how this new interface works or where it is targeted for. There is some speculation that 3GIO will only be used for PC's. Others think this might be a HyperTransport like technology. The only item confirmed is, Intel has stated they will disclose the details of 3GIO in late August 2001. Therefore, we can not comment on this technology except to say that it is late to market as compared to the other new buses that have already published specifications and seen new product announcements and releases.

We see HyperTransport and RapidIO as complementary technologies. Although InfiniBand technology can be utilized as a bus on a server logic board or backplane, its primary focus is to

1. From www.hypertransport.org:
2. From http://www.rapidio.org/tech.htm:

improve I/O bandwidth that move data from the system controller and memory to the edge of the data center and vice versa. HyperTransport and RapidIO's primary focus is to move data to the host processor. Although these two technologies can overlap on a server logic board or in backplanes, it's best to imagine an InfiniBand device that is Landed on the Motherboard (LOM) connected to the host processor or system logic via one of these new internal buses. In this application, the two technologies offer a great way to move unprecedented amounts of data from a System Area Network (based on InfiniBand fabric) to the host processor and vice versa.

It is important to note that InfiniBand is designed to support high bandwidth scalable backplane solutions for servers and communication equipment. InfiniBand uses an embedded clock in its signal that minimizes signal skew as a function of distance, unlike HyperTransport or Rapid IO that rely on exact routing layouts to keep accurate timing. This makes InfiniBand ideal for backplane designs. And the IBTA has already architected solutions that include connectors and physical designs that describe standard and tall InfiniBand modules (or blades) that can be either single or double wide.

8.0 Summary

Each of the technologies listed has a useful role to play in today's compute intensive world. PCI is a stellar way to add I/O functionally to personal computers, Fiber Channel is the only high-bandwidth SAN solution available today, Ethernet is the foundation of LAN communication, storage over IP has the promise to offer remote storage communication for disaster backup, and new internal buses offer greater bandwidth to the processor. But, none of these technologies offer a complete solution for all I/O and communication facets of the Internet and Enterprise data center. Only InfiniBand offers such a complete solution that provides a single unified fabric for the data center designed for the highest RAS available.

As servers mature and evolve into new form factors, it is clear that InfiniBand is the choice as the primary I/O interconnect on server boards and the single fabric that enables clustering, communication, storage and I/O for the Internet and Enterprise data centers of the future.

I/O Features Matrix

This document provides a feature comparison between a number of compute I/O technologies.

Feature	InfiniBand™	PCI-X	Fibre Channel	1Gb & 10Gb Ethernet	Hyper-Transport™	Rapid I/O
Bus/Link Bandwidth	2.5/10/30Gb/s[a]	8.51 Gb/s	1/2.1 Gb/s[b]	1 Gb, 10Gb	12.8, 25.6, 51.2 Gb/s[f]	16/32 Gb/s[c]
Bus/Link Bandwidth (Fully Duplexed)	5/20/60Gb/s[a]	Half-Duplex	2.1/4.2 Gb/s[b]	2 Gb, 20Gb	25.6, 51.2, 102 Gb/s[f]	32/64 Gb/s[c]
Pin Count	4/16/48[d]	90	4	4, Fiber	55,103,197[f]	40/76[c]
Maximum Signal Length	Km	Inches	Km	Km	Inches	Inches
Transport Media	PCB, Fiber and Copper Cable	PCB only	Copper and Fiber Cable	Copper and Fiber Cable	PCB only	PCB only
Simultaneous Peer to Peer communication	15 VLs + Mngt Lane			X		3 Transaction Flows
Native Hwd Transport Support with Memory Protection	X					
In-Band Management	X		Uses out-of-band mngt	Not native, can use IP		
RDMA Support	X					
Native Virtual Interface Support	X					
End-to-End Flow Control	X			X	X	X
Memory Partitioning[e]	X		X			
Quality of Service	X		X	Limited		X
Reliable	X		X		X[f]	X
Scalable	X		X	X	X	X
Maximum Packet Payload	4 KB	Not Packet Based	2 KB	1.5KB (10Gb no jumbo support)	64 bytes	256 bytes

a. The raw bandwidth of an InfiniBand 1x link is 2.5Gb/s (per link). Data bandwidth (due to 8b/10b encoding) is 2.0Gb/s for 1X, 8 Gb/s for 4X and 24Gb/s for 12x; twice that for full duplex or 4/16/48 Gb/s.

b. The bandwidth of 2Gb Fibre Channel is 2.1Gb/s but the actual raw bandwidth (due to 8b/10b encoding) is 20% lower or around 1.7Gb/s (twice that for full duplex).

c. Values are for 8 bit/16 bit data paths peak @ 1GHz operation. Speeds of 125, 250 & 500 MHz are supported.

d. The pin count for a 1x link is '4' pins up to '48' pins for a 12x link.

e. Memory partitioning enables multiple hosts to access storage endpoints in a controlled manner based on a key. Access to a particular endpoint is controlled by this key, so different hosts can have access to different elements in the network.

f. Based upon 8, 16, 32 bit HyperTransport (it can support 2 & 4 bit modes) with up to 800 MHz (DDR) operation (modes from 400 MHz can be supported). Error management features will be refined in future revisions of the specification.

9.0 About Mellanox

Mellanox is the leading supplier of InfiniBand semiconductors, providing Switches, Host Channel Adapters, and Target Channel Adapters to the server, communications, and data storage markets. In January 2001, Mellanox Technologies delivered the InfiniBridge™ MT21108, the first 1X/4X InfiniBand device to market, and is now shipping second generation InfiniScale silicon. The company has raised more than $33 million to date and has strong corporate and venture backing from Intel Capital, Raza Venture Management, Sequoia Capital, and US Venture Partners.

In May 2001, Mellanox was selected by the Red Herring Magazine as one of the 50 most important private companies in the world and to Computerworld Magazine Top 100 Emerging Companies for 2002. Mellanox currently has more than 200 employees in multiple sites worldwide. The company's business operations, sales, marketing, and customer support are headquartered in Santa Clara, CA; with the design, engineering, software, system validation, and quality and reliability operations based in Israel. For more information on Mellanox, visit www.mellanox.com.