TANEJA GROUP
TECHNOLOGY ANALYSTS

T e c h n o l o g y B r i e f

InfiniBand in the Enterprise:
Next Generation Data Center Computing

April 2008

There is a flurry of activity afoot in the industry today around the demands of the next generation data center. Without a doubt, the circumstances are challenging: space and power are at premium and only getting more so; there are thousands of different demands for I/O even in small data centers, with new demands popping up at frequencies that require every IT manager to have seer-like powers of prediction when designing the infrastructure; and the pace of change and reconfiguration around every server and storage array in the enterprise is phenomenal. In the midst of this, the typical storage manager finds they are managing a broader range of connectivity and protocols than ever before – spanning Ethernet, Fibre Channel, IP, iSCSI, NFS, and more. It is little wonder that all corners of the industry are grasping after new solutions and coming up with new standards for better I/O management with less complexity.

In the middle of this storm of activity, stands InfiniBand. Long looked to as the leader for high performance computing, InfiniBand is gaining increasing adoption in general purpose enterprise computing. Why? A few reasons include: cost effectiveness, simplicity, flexibility and most of all performance. At first glance, those characteristics might be easily dismissed as a set of buzzwords, but their combination in one I/O platform in fact brings the next generation of data center connectivity to bear today. Moreover, InfiniBand is leaps and bounds beyond the competition in each of those capabilities. In this Technology Brief, we'll look at the challenges surrounding enterprise I/O management, dive into key InfiniBand capabilities, and take a look at how InfiniBand is being adopted by enterprises today.

The Data Center Built of Sticks

Much like the story of the three little pigs, many current generation data centers are figuratively built of sticks. Those sticks are the I/O fabrics that connect systems together. Today's common fabrics are not up to the task of supporting an ever-changing and evolving infrastructure that is being bombarded by a deluge of requests from business users for more compute power and storage than ever before. In the face of these requests, the data center requires more ability to scale I/O bandwidth, manage I/O allocation, and consolidate I/O than ever before.

IT managers are searching for ways to obtain these capabilities in a data center fabric made out of heterogeneous strands of connectivity. The result? These managers must add more switches, more cables, more adapters, more protocols, and more network resident management and control devices in

order to effectively scale, adding to an ever increasing sprawl of expensive and entangled connectivity. In a data center faced with shrinking usable space, denser-than-ever devices, and increasingly precious power and cooling, each addition robs the enterprise of future flexibility. Worse yet each addition can be a step backwards because managers have little ability to better control or manage additional connectivity, much less granularly deploy it or guarantee levels of performance.

Like a house of sticks, the foundation is insufficient, and the structure is not sound. As managers pile on connectivity and reconfigure their infrastructures to support changing business needs, they throw more "sticks" on top of a jumble of Ethernet, Fibre Channel, devices, software, trunks and various protocols. In large data centers, the end result looks increasingly like it will be eventual collapse if the approach to connectivity does not change.

The increasing concentration of data center assets, rapidly evolving performance characteristics of systems, and constant application and infrastructure change are converging upon IT managers in a 1-2-3 punch that exponentially increases the challenges surrounding inflexibility, limited performance, and poor management within existing data center connectivity.

Data Center Evolution: Pressures in the New Age of Computing

The 1-2-3 punch from density, performance, and change is striking close to home for nearly ever data center manager today. Pressed by these forces and the limitations of existing I/O fabrics, the infrastructure within the data center has become inflexible and unmanageable, and is driving management overhead out of control.

These forces are in fact being driven by several significant trends:

Computing in the Data Center

High Performance Computing (HPC) has long been seen as a computing domain that is restricted to the ivory towers of academia. But in recent years, HPC has entered a new era – one that places it firmly in the center of the enterprise. Regardless of industry, today more enterprise IT managers than ever before own high performance computing solutions to address demanding business requirements for faster compute cycles. Taneja Group has observed that the rate at which the enterprise is deploying HPC-like solutions is rapidly increasing.

Clustering and distributed computing have always suffered in the hands of I/O fabrics that were restrictive in bandwidth, or introduced latency and processing overhead. For large amounts of data processing that such solutions are often purpose built for, limited bandwidth requires multiple connections. These multiple connections are technically challenging to aggregate together and create unnecessary cabling and cabling complexity. At the same time, traditional connectivity introduces latency in both the transmission media and host connectivity devices that can defeat many of the benefits of distributing processing tasks across multiple nodes. Host bus adapters within many fabrics can in fact triple or quadruple the transmission latency introduced by

media, while adding enormous host processing overhead.

For many enterprises processing significant data, or connecting many systems together, traditional connectivity just doesn't suffice, and the challenges of HPC have trickled down into these data centers. This is even getting to be the case for small computing clusters that are common to many applications.

Storage in the Data Center

Even without clustered and/or distributed computing systems that mimic HPC, enterprises everywhere are still being pressed to extract every ounce of available performance from their infrastructures. Enterprises often feel this pressure first around storage performance. While random disk I/O performance has often been blamed as a lagging technology that drags down other computing performance, the truth is most organizations that require more storage performance can harness it through over-provisioning disk spindles or turning to technologies like flash or dram-based solid state disk. Unfortunately, these organizations inevitably soon become constrained by limited I/O across their fabrics, or saturated system buses and poor adapter card implementations. The solution often involves the addition of more expensive HBAs, offload cards, and/or connectivity, or a nearly complete infrastructure replacement and migration in pursuit of the next generation of connectivity.

Virtualizing the Data Center

Without a doubt, server virtualization has promised to rescue the enterprise from an aging data center infrastructure. By consolidating multiple servers on single larger virtualization hosts, enterprises can reduce their hardware, connectivity, and multiple sources of I/O. Unfortunately enterprises that have turned to server virtualization often face their first rude awakening around I/O. The culprit is often the bottleneck created by the interface between storage adapters and the virtualization host. That bottleneck both constrains storage bandwidth, and creates significant CPU overhead and latency as the virtual host attempts to deliver storage I/O to running guests and processes. In fact, even as bandwidth grows, this overhead can easily overpower fast buses and processing resources, and in worst cases can easily consume 10's of times or more bandwidth on a system bus than on a connectivity fabric.

Of equal significance, server virtualization brings a new capability for rapid provisioning and change into the data center, while at the same time not fully isolating virtual servers from their dependencies upon static and complex I/O fabrics. The result is an infrastructure where some changes are simple (server provisioning and resource configuration) but where other changes (movement) can create significant overhead around managing the complex web of enterprise connectivity.

Managing the Data Center

Traditional connectivity in general has long fallen short in manageability, and the enterprise requires more manageability today than ever before. Lack of manageability is wreaking havoc on a data center infrastructure that is constantly

changing in the face of virtualization, normal equipment replacements, and evolving business needs, and the enterprise cannot efficiently harness the manpower necessary to control the chaos. Multiple types of connectivity, multiple protocols, and protocols that were designed for relatively static long-term infrastructure configurations can rapidly complicate the management of changing environments. When a single enterprise system is moved or replaced, whether virtual or physical, changes ensue. Ethernet connectivity, Fibre Channel connectivity, LUNs, VLANs, addressing, and any number of other configurations might need to be changed, in multiple locations and for multiple connections.

Moreover, as an infrastructure changes, all of the ensuing reconfiguration can often impact how fabrics perform. The enterprise has little ability to manage data inside of connections – as a result, many infrastructures have relied upon physical topologies to control access and bandwidth utilization. When physical cabling is changed, or new systems are brought on-line, the impact upon existing connectivity is often not understood, and performance, service-levels, and security can be thrown out the window.

Density in the Data Center

Finally, traditional connectivity is rapidly burying many datacenters beneath mountains of cabling. As form factors increase in density (e.g. blade servers), more cabling in smaller spaces not only make management and change nearly impossible, but consume unnecessary power while making the data center increasingly difficult to cool.

The Changing Landscape for Data Center 2.0

Taken together, it is clear that computing, storage, virtualization, manageability, and density are creating a unique set of demands for the next generation data center. A quick look around cutting edge data centers today suggests that savvy IT managers – the thought leaders in their professions – are fully aware of these challenges. Moreover, when they look around for an answer, they are finding that while there are emerging opportunities, there is only one technology that is well established and proven in the industry. That technology is InfiniBand.

In our research with end users and vendors alike around InfiniBand adoption, it is clear that InfiniBand is on an upswing. Revenues are rapidly growing for established InfiniBand vendors, such as Mellanox, Voltaire, and QLogic. Additionally, more appliances are emerging onto the market that use InfiniBand's ability to pool and granularly manage I/O resources, such as 3Leaf's V-8000 and Xsigo's I/O Director. In fact, the relative recent entry of 3Leaf and Xsigo into an IO virtualization market already established by Cisco and Voltaire emphasize how much opportunity there is to meet pressing data center needs with core InfiniBand capabilities. Similarly, InfiniBand native storage also indicates that these opportunities are beginning to capture the consideration of leading organizations.

So why is there such an uptick in InfiniBand adoption and why are data center thought leaders turning to InfiniBand to solve their challenges? To answer this question, and provide a better understanding of how InfiniBand fits into the enterprise, we'll next take a brief look at the key capabilities behind InfiniBand technology, exploring how it is differentiated from more complex traditional network and SAN connectivity, and then look at specific applications where InfiniBand can solve difficult challenges.

InfiniBand and Data Center 2.0

As we've discussed, the enterprise is dealing with foundation wrecking issues around data center demand - a deluge of new compute and storage systems in multiple physical and virtual formats; system performance that is scaling faster than connectivity can keep up; and more change than ever before. InfiniBand is finding adoption as a response to these issues because it can unify connectivity across storage and systems, rapidly scale and be managed with granularity and ease, and lower capital dollars spent on multiple fabrics. We'll next look at each of these capabilities.

Dense, Unified Connectivity

Nearly since its inception, InfiniBand has been a hotbed of deeply capable protocol heterogeneity, and has rapidly lent itself to providing highly efficient Remote Direct Memory Access (RDMA) to SCSI (and thereby Fibre Channel), IP, iSCSI, NFS, and various other protocols. RDMA provides efficient, direct data placement into host or client memory without processor overhead. Historically, RDMA has evolved for each of these protocols on InfiniBand first, as InfiniBand provides deep and consistent

InfiniBand hits its Stride

While InfiniBand fabrics often use gateway options that broker connectivity between InfiniBand and Fibre Channel systems, InfiniBand-native storage is gaining traction. The native InfiniBand storage market is small, consisting of a handful of vendors including DataDirect Networks, Engenio, and Xiranet, as well as such solutions as Terrascale and IBRIX, but our preliminary look into this space suggests that its growth trajectory may be on par with the fastest growing storage systems in the market, potentially approaching a CAGR of 60% - 70% when sampled over the past 3 years. With the current uptick in InfiniBand, multiple opportunities in financial and other compute intensive markets, we suspect this growth rate of InfiniBand-native storage systems will be sustainable in the next few years.

The factors driving this growth are demand for high capacity storage systems with high bandwidths, and ever-evolving compute intensive business needs. This combination requires high bandwidth, high performance, efficient management, and dense cabling in order to support a large number of systems in a small amount of space. In the past, many of these organizations have turned to InfiniBand to FC gateways for their storage needs, but as more InfiniBand storage choices become available, these same organizations are switching to native InfiniBand storage to increase their flexibility and reduce their headaches.

mechanisms for RDMA. While on the surface, this connectivity has created an InfiniBand unique set of terminology – SRP, IPoIB, iSER, NFSoRDMA – these technologies are in fact largely the same and actually simplified by consolidation onto a standard InfiniBand fabric that uses the same robust CX4 cabling commonly used with 10Gb Ethernet.

With direct memory access into a server, a CX4 cable and an InfiniBand Host Channel Adapter (HCA) become an extension of a server system bus, allowing system administrators to centralize heterogeneous physical IO connectivity in a device outside of a server. This is done through InfiniBand gateways and routers (or direct attachment to native InfiniBand storage) that in turn contain adapter cards for traditional fabrics such as Ethernet or Fibre Channel. The InfiniBand HCA inside a server then creates virtual NICs and virtual HBAs that appear as physical devices to the host operating system.

With this approach organizations are finding they can use a single fabric with fewer adapters to unify multiple connections to broadly heterogeneous systems, including physical and virtual machines. This unified fabric leads to rich services, including network gateway devices that can offload services such as NIC and HBA management that were previously host responsibilities. More importantly, sophisticated InfiniBand RDMA support makes some complexity unnecessary.

As an example, increasing Ethernet speeds often require proprietary 3rd party TCP offload processors, drivers, and cost. TCP offload has long been natively delivered by InfiniBand's RDMA-enabled protocols, and under a much broader set of conditions than Ethernet-based offload NICs and drivers can address. Moreover, InfiniBand's broad RDMA enablement of higher layer protocols is more effective than TCP offload alone – end user case studies have demonstrated the higher layer NFS/RDMA over IB protocol doubling the performance of the lower layer IPoIB offload effect.

Performant Connectivity

While capabilities such as TCP offload optimize CPU utilization, InfiniBand fabric is equally performant in several other dimensions. First, InfiniBand is the ultimate low latency fabric. In part this is due to broad RDMA support that avoids CPU and interrupt overhead, but InfiniBand has also been designed at every layer for low latency. In fact, even the poorest implementations demonstrate sub-millisecond latencies, and well-designed systems can yield end-to-end nanosecond latencies.

Equally important, InfiniBand's bandwidth is unmatched by other connectivity in the market, with single port speeds today ranging from 10Gbps to 40Gbps. InfiniBand utilizes many of the same encoding schemes as Fibre Channel for efficient, high bandwidth throughput. In testament to the efficiency of InfiniBand encoding, a number of organizations have turned to native InfiniBand over a WAN because it can more efficiently transfer data via products like Obsidian Research's Longbow WAN gateway. While InfiniBand can make very efficient use of native high-bandwidth port speeds, it is not restricted to just the data center.

Moreover, since InfiniBand's unified fabric provides deep visibility into all I/O and interactions on the fabric, InfiniBand switches, routers, and gateways automatically redistribute and load balance connectivity. This keeps a fabric performant, and scales a fabric in a grid-like fashion. Simply adding an additional switch can rebalance the I/O load for a complete fabric.

Flexible and Manageable Connectivity

With efficient, performant, high-bandwidth connectivity, InfiniBand is more than capable of consolidating all of the connectivity any system may require over single or redundant IB connections. This provides the basis for more flexible connectivity management, and allows IB switches, gateways, and routers to granularly manage I/O.

InfiniBand management encompasses both easy configuration and change, and comprehensive control over I/O. Since I/O connectivity is completely virtualized in an InfiniBand fabric, inflexible FC and Ethernet connections can be completely terminated in gateway devices, with full visibility to hosts. The gateway device can dynamically remap changes without reconfiguring an entire fabric. This means that a virtual HBA and virtual NIC in a virtual guest can move with a guest and maintain connectivity even if it moves to a new virtual host on a different InfiniBand switch.

Moreover, InfiniBand provides deep control over I/O. InfiniBand switches can pool multiple FC connections, and then use fine-grained quality of service and access control mechanisms to prioritize resources for different systems. Quality of service restrictions and/or guarantees can be configured for virtual I/O adapters, and follow a system anywhere it is moved. Using gateways, or native InfiniBand storage, InfiniBand can create a load balanced pool of QoS-managed connectivity that is more efficient than server-based multi-path approaches to aggregating connectivity and distributing load. The result is that a single high I/O and bandwidth connection to a server can provide better and more controlled performance than traditional storage fabrics.

While InfiniBand's ability to abstract I/O from servers creates a fabric where change can be easily and non-disruptively managed with less overhead than multiple fabrics of the past, the availability of the fabric is increased as well. InfiniBand virtualizes Ethernet and FC connectivity while pooling physical connections in a gateway device that isolates failure-prone devices from traditional fabrics. When these devices are placed in a gateway, InfiniBand's robust traffic management and redirection make failures in commodity NICs and HBAs non-disruptive by redirecting traffic across other interfaces while hardware is replaced and brought back on-line.

InfiniBand I/O virtualization vendors are commonly touting 50% or more decreases in management overhead when compared to traditional fabrics. We attribute these claims to significant reductions in time and effort previously associated with manual provisioning and changes. This number is likely even larger when considering bigger

picture time and effort associated with both indirect planning and change management processes, as well as the costs of downtime.

Cost Effective Connectivity

When considered alone, it appears as if 10G Ethernet may soon be at per port price parity with InfiniBand. But, in our view, that is only part of the story, as InfiniBand's capabilities rapidly recreate other total cost of ownership advantages.

First, many organizations have been ill at ease with sending their mission-critical storage traffic over an Ethernet fabric. These organizations may be considering 10G Ethernet in addition to expensive Fibre Channel SANs that can easily add $1500 or more cost per port of connectivity. InfiniBand's ability to disaggregate and abstract I/O connectivity can be leveraged as an FC and Ethernet gateway, consolidate and reduce the number of these ports required, and rapidly drive down costs. When assessed against Ethernet and FC connectivity, InfiniBand can rapidly become significantly cheaper, reflecting as little as 1/4th to 1/3rd the price of traditional multiple fabrics. This figure improves when total operational expenses are considered, including power and cooling. As an example, some InfiniBand vendors are claiming 70% improvements in power utilization and cooling over the multiple adapters, switches, and cables used in traditional fabrics.

Moreover, with performance and easy management for nearly any protocol, InfiniBand is a mature solution for completely managing a complex array of connectivity and systems. While FCoE is playing catchup with InfiniBand's basic capabilities, InfiniBand is poised to attack data center density with mature capabilities.

Finally, enterprises with InfiniBand can rapidly open the door for alternative, open storage solutions. Storage solutions such as IBRIX's Fusion, Isilon, Lustre, or several other scalable, modular software-based products that run on commodity hardware can deliver enormous performance at under $2/GB compared to traditional mid-range and enterprise systems that range between $10 to $30 per gigabyte of capacity. While InfiniBand isn't a prerequisite for these systems, InfiniBand can harness all of their performance potential, and make implementation and scaling easier.

InfiniBand Use Cases

For organizations willing to survey alternatives to traditional connectivity, InfiniBand's capabilities can rapidly create infrastructure agility alongside cost and management efficiencies. Some industries and users have been more pressed to consider alternatives than others. We'll look next at a few example enterprise-computing needs where InfiniBand is finding the most traction today.

The High Performance Computing Center You Don't Know You Have

Since the earlier days of academia-based HPC, similar demanding performance requirements have gradually crept into data centers across many industries. Most of these demands revolve around analytics and/or processing of large amounts of sequential digital data. Examples include

analytical processes in the oil and gas industry around seismic and survey data, complex computational processes common among Electronic Design Automation (EDA) tools in the semiconductor industry, digital content creation and editing in the video industry, video and instrumentation analysis in the defense industry, and historical data analysis in the financial industry.

The challenges that surface in these scenarios often revolve around limited I/O bandwidth to single large files. Such I/O patterns are extremely difficult to aggregate across multiple connections, but there often is also latency in other fabrics that can slow completion time for time sensitive calculations. For these types of demands, storage managers have turned to InfiniBand so they can stream huge volumes of sequential data from dense, InfiniBand-native disk platforms like DataDirect Network's S2A arrays, and perform random accesses against I/O optimized native InfiniBand SSD devices from Texas Memory Systems. In these areas, InfiniBand is simply unmatched for performance today and likely even tomorrow, as even emerging technologies such as 10gb Ethernet and FCoE will not match the bandwidth and latency characteristics of InfiniBand. For high bandwidth sequential needs, dense disk systems have been able to deliver astounding performance of multiple gigabytes per second, while SSD devices coupled with InfiniBand can deliver microsecond latencies that are impossible with other host bus adapter and fabric approaches. These demands are confronting financial institutions, aerospace companies, design teams in all manner of manufacturers, large

HPC Storage as a Service

Xasax normalizes exchange tick data from high speed data feeds in their 10 massive Points of Presence (POPs) scattered across the United States, and resells this normalized data to traders for up to the second analysis. More significantly, given the size of the data sets and the importance of latency to traders, Xasax sells storage and compute resources in their POPs in a Storage as a Service (SaaS) model. By buying into this service, traders can run trading and analysis software right next to high-speed real-time trading data.

In the financial world, where trading relies upon real-time data, every second counts. Sub-second delays in creating an execution order for a trade can easily turn a profit to a loss. To deliver the performance and responsiveness their customers require, Xasax turned to a combination of Voltaire's InfiniBand switches and FalconStor's IPStor software. Xasax uses FalconStor to virtualize extremely dense, industry standard servers into a pool of scalable, performant iSCSI storage. Voltaire switches then create a mesh of compute nodes that interconnect over InfiniBand, and access IPStor nodes via InfiniBand's RDMA-enabled iSER. The result is a datacenter at each POP filled with low latency compute resources.

Xasax uses their own InfiniBand network to cluster the computing resources necessary to normalize the high-speed data being fed to them from different financial exchanges. That clustered computing task can push nodes toward the 20Gbps mark, but in turn provides data to users at Xasax POPs in near real-time, with guaranteed compute performance.

and small oil and gas companies, video broadcasters, digital content creators, a long list of government agencies, numerous types of government contractors, and many other companies. As the enterprise evolves, data processing, data storage, and data access activities are exceeding the capabilities of traditional fabrics.

Today, storage and server administrators in even run of the mill enterprises are more often than not facing HPC-like computing needs, or will be facing one in the near future. In our view, once a system approaches 16 nodes and requires substantial I/O, connections to communication and storage fabrics make traditional connectivity impractical. Moreover, even many small 4 to 16 node clusters built for performance driven solutions such as OLTP can sacrifice much of their capability to latency and processing overhead introduced by poor connectivity. And as we have discussed, if a system requires much periodic change or management of connectivity or nodes, then other connectivity may become impractical well before this.

Server Virtualization

In many ways, InfiniBand is ideally aligned with the needs of server virtualization. Server virtualization enables the enterprise to rapidly deploy and easily change infrastructure, but these same capabilities are wreaking havoc on change control processes and forcing storage and server administrators to constantly react to utilization, performance, and configuration issues. While InfiniBand has long targeted I/O virtualization, new solutions specifically targeted at virtual servers are rapidly emerging, and are using InfiniBand as their foundation for granularly managing I/O. These solutions are using a software intelligence layer to automate fundamental InfiniBand capabilities such as QoS, virtual HBAs, and virtual NICs. The result is an overarching I/O management system that can control and automate many previously complex and disruptive aspects of provisioning and moving servers. Moreover, these management layers can provide total visibility while granularly reserving and allocating bandwidth for specific servers and applications to effectively put an end to hard-to-identify, run away resource consumption that can plague some virtual server environments.

Just as importantly, server virtualization with InfiniBand can also be highly performant. Like a mega-offload card, RDMA can bypass both the hypervisor and kernel, and reduce data copies, context switches, and virtual switches, thereby significantly reducing both latency and CPU utilization. Coupled with granular control over I/O load, this makes the most of precious host resources, and allows organizations to consolidate even more virtual guests on a single platform.

At Taneja Group, we recognize the obvious attractions of both SAN attachment and iSCSI (including InfiniBand's iSER) for server virtualization, and we firmly believe NFS has significant merit as well. In many ways, NFS offers more versatility for server virtualization, as it applies file level integrity guarantees to virtual images, and can make use of pervasive and robust management tools such as advanced file systems,

snapshots, thin provisioning, and various other technologies. Unfortunately, NFS has significant processing overhead for both the IP traffic that bears it and the NFS traffic itself. But when NFS is paired with InfiniBand, it can become much more performant and may make an ideal platform for many virtualization needs. More importantly, NFS can open the door to modularly scalable NAS storage systems such as IBRIX, Isilon, and Lustre. The cost and scalable capacity and performance of these systems are in many ways good fits for virtualization.

As the number of vendors with InfiniBand-based server virtualization I/O management solutions indicates, there is merit in the granular control InfiniBand can deliver to these environments. From our view, regardless of file storage format and protocol choices, organizations can significantly benefit from deploying their virtual server infrastructures in InfiniBand islands that then attach to traditional fabrics. This will allow these organizations to granularly control I/O for their virtual server infrastructures.

Digital Content

Another current niche where InfiniBand is finding rapid adoption is around digital content storage and creation. Digital content creation and/or access often require huge amounts of bandwidth to effectively stream video from disk storage. Whether rendering, creating, or editing digital content the bandwidth requirements can exponentially mutliply as users attempt to interact in real or near real time. Moreover, with huge amounts of content stored in these environments, cost can rapidly become a concern.

InfiniBand has easily stepped up to the plate for the most demanding applications in these environments. With its intrinsic QoS, sub-microsecond latencies, and ample bandwidth, InfiniBand can guarantee system responsiveness. Moreover, InfiniBand coupled with modular, scalable storage can build a truly adaptable storage platform that can match even unforeseen future demands.

Organizations considering storing and working with digital content of any type should consider the future demands upon their storage and data fabrics – InfiniBand might well be worth a look. Given the explosive data growth facing the enterprise today, and the fact that much of this data is rich content, even moderate needs that fit well within traditional fabric capabilities could grow to need substantial bandwidth in the near future. It is easy to imagine a digital content system built to edit and stream content to a few users that will grow to serve web, conference room, training, or other video to thousands of end points in the near future. Such traffic could quickly complicate creating and editing new content over an I/O constrained fabric. Rapidly changing patterns in the way rich content is used by the enterprise may require granular I/O management and QoS capabilities sooner rather than later.

Taneja Group Opinion

For a wide number of enterprises, the needs of the next generation Data Center are present and pressing today. While

InfiniBand isn't an across the board solution for all connectivity in the enterprise, selective use of InfiniBand can go a long ways in reinventing the enterprise infrastructure and creating a data center with second generation capabilities.

For the enterprise, we have long held that the time is right to start seriously considering domains of InfiniBand connectivity. While forklift upgrades of existing infrastructure are hardly justifiable or called for, for many new requirements, InfiniBand may fundamentally make sense. Implementation within subsets of servers or applications can still easily connect to existing FC infrastructure, but allow better connectivity and I/O management for groups of systems. The vendors currently storming onto the scene with InfiniBand-based IO management solutions are only strengthening our view.

For many organizations, the needs of Data Center 2.0 are here, and they are pressing. These users require more density, more performance, fewer airflow obstructing cables, more power efficiency, and easier management. Those capabilities are present and well tested in InfiniBand. Some users have in fact turned to InfiniBand to solve these issues, and are in fact finding good benefit in selective deployment throughout their organizations. The capabilities of InfiniBand will never replace traditional connectivity in the enterprise, but InfiniBand technology can interconnect with and leverage traditional connectivity to create a truly well managed and adaptable infrastructure foundation.