Xframe™ from S2io, Inc.

A White Paper from S2io, Inc.

Server Optimization

Maximizing Bandwidth in the Virtualized Data Center

August 2004

Executive Summary

Faced with explosive growth in application data, multiple silos of servers and storage arrays, and limited budgets and resources, the challenge for IT is to find cost-effective ways to transform rigid data centers into environments capable of adapting to change as it happens. By implementing a virtualization strategy via consolidated servers or by using existing assets, it is possible to pool data center resources, deploying them dynamically to address service-level demands, while enabling vital efficiencies such as:

- Lowering operating and administration costs
- Reducing application latencies
- Improving server back-up performance
- Optimizing bandwidth utilization
- Improving overall application performance

Key to reaching these goals is establishing a simplified infrastructure that reduces the complexity of data center management and adopts standards for, among other things, efficiently transferring data and applications between assets. With enterprise data doubling every six months, network bandwidth clearly plays an important role in creating an adaptive enterprise environment that extracts optimal utility from servers, storage arrays and other assets.

To achieve immediate bandwidth increases today — and even greater efficiencies in years to come — IT managers now have the opportunity to leverage a proven, IEEE standard technology whose existing infrastructure and vast ecosystem provides a low-risk, high-return approach to optimizing data center performance. By selecting 10 Gigabit Ethernet, IT managers can leverage the economies of scale available only with a truly ubiquitous technology whose protocols, management tools, and troubleshooting techniques already provide the foundation for 85 percent of networks worldwide.

Investing in 10 Gigabit Ethernet today is a cost-efficient and non-disruptive move toward outfitting enterprises for an eventual unified interconnect fabric, thus achieving a level of standardization and simplification crucial to implementing a truly adaptive enterprise.

Introduction: Solving today's data center problems

Organizations today are at a crossroads as they bridge the data center infrastructures they built in the 1990s with the dynamic service level demands they will face over the next decade. Most mid- to large-sized data centers are saddled with the byproducts of rapid growth experienced over the last 10 years. During the economic boom of the 1990s, enterprises moved mission critical applications — from Enterprise Resource Planning (ERP) and Customer Relationship Management (CRM) — to Internet-based paradigms, and ample budgets allowed data centers to spawn vast populations of servers, storage arrays and other assets.

The need for an adaptive enterprise

IT managers today face a long list of challenges that will only intensify in the decade to come.

- Data centers made up of disparate and inflexible silos, multiple OSs, and interconnect technologies
- A desire for end-to-end visibility of all assets
- Increasing demands on applications, servers, storage devices and networks, including enterprise data growing between 30 and 200 percent annually
- Increasing data back-up constraints
- Space and power limitations
- Growing expectations, but limited if any increases in budget
- Complicated network topologies increase security vulnerabilities

Further fueling the phenomenon of "server sprawl" was fierce vendor competition and technology advances that drove down server prices, making it all too easy for IT managers to absorb workload increases by adding servers. Often, with each application introduced to the data center, organizations deployed up to three or four servers — one to develop the application, another to conduct quality assurance testing, still another to test user acceptance, and finally another to host the production-ready application. Servers have been provisioned to handle peak workloads, even if only called on to operate at such levels once a quarter, or even once a year. Disk arrays and tape backups, too, were often dedicated to specific applications and their corresponding servers. Ownership issues abounded — and still do, in fact — leading to demands that under-utilized server banks remain the sole property of a particular line of business, such as sales, HR, or accounting. The result was the emergence of silos of information sitting on numerous disparate groups of servers, many with utilization rates of 30 percent or less, and patched together with multiple interconnect technologies and operating systems. This, too, has done little to help IT's nagging reputation among non-technical executives as a large corporate cost center.

While adding silos was an effective solution in the '90s, this trend led to disparate data center infrastructures that were not designed to address IT's major concerns today and for the long term. For instance, a key desire among IT managers — to have end-to-end visibility and manageability of all data center assets — is often impossible in today's enterprises. Complicated network and data center topologies also threaten an enterprise's ability to stay ahead of security risks, as more devices require more work to monitor, update, and repair them. The more complex the data center infrastructure, the more difficult and expensive it is to manage.

Making matters worse, organizations expect IT management to meet ever more demanding Service Level Agreements (SLAs) with budgets and resources that have difficulty keeping pace with the increase in workload. Today's applications generate more data than ever. From email to data warehousing, enterprise application data will increase at a compounded annual growth rate of 30 to 40 percent, according to a Salomon Smith Barney study[1], while some corporate estimates warn that enterprise data is growing so fast as to double every six months.

New generations of data-rich applications will put even more demands on data centers and the networks that serve them. Oracle Database10g, for instance, is specifically designed to handle databases of up to 8 exabytes in size in environments that depend on high-performance business intelligence services such as ETL (Extract, Transformation and Load), data warehousing, online analytic processing and data mining.

And as developing countries grow more dependent on the Internet for communications and business, that growth rate will intensify. Today, some 794.8 million people worldwide have Internet access, though that represents only 12.3 percent of the world's population. North American Internet users account for nearly 69 percent of the region's

population, indicating that the pace of growth here will slow in the coming years, easing somewhat from its 107 percent growth rate from 2000-2004. But some emerging economies, such as Asian giants China and India, will see phenomenal growth in coming years. With some 3.7 billion potential users and Internet use among less than 7 percent of the population, Asia will likely see its 115.5 percent growth rate only accelerate in the future. And Internet in the Middle East, with another 260 million potential users today and a historical four-year growth rate of 219.2 percent, still has only reached 6.5 percent of the region's populace. Factor in Latin America and the Caribbean (with penetration of just 9.4 percent) and Africa (at only 1.4 percent), and the coming onslaught of Internet traffic and data demands from around the world is clear.[2]

Yet IT services spending is expected to remain flat through 2007, according to Gartner/Dataquest,[3] while IDC expects IT spending overall to increase by only 5 percent annually.[4] So while IT demands are exploding, budgets are not.

In the face of these pressures, many enterprises are striving to optimize their data centers, so they can achieve economies and efficiencies that will achieve the following: lower operating and administration costs; reduce application latencies; improve server back-up performance; optimize bandwidth utilization; and improve overall application performance.

Yet they still have to retain the stability, security and vertical application relevance of today's data centers while overcoming the limitations and inefficiency of inflexible data silos that are difficult and costly to manage horizontally — all with little or no increase in budget or resources.

Enabling the adaptive enterprise

Senior management is also pressuring IT departments to view the data center not only from a technology or product perspective, but from an economic one as well. As a result, they are seeking cost-effective ways to establish a more agile, adaptive environment capable of meeting demand as workloads shift and users and applications are dynamically added or removed from the infrastructure.

For IT managers, the Holy Grail is adaptive or flexible computing, an approach to IT that pools and shares resources so utilization is optimized and supply automatically scales to meet demand. To evolve from vertical silos to horizontal pools of resources, enterprises need to *virtualize* their IT infrastructure. By partitioning applications on the

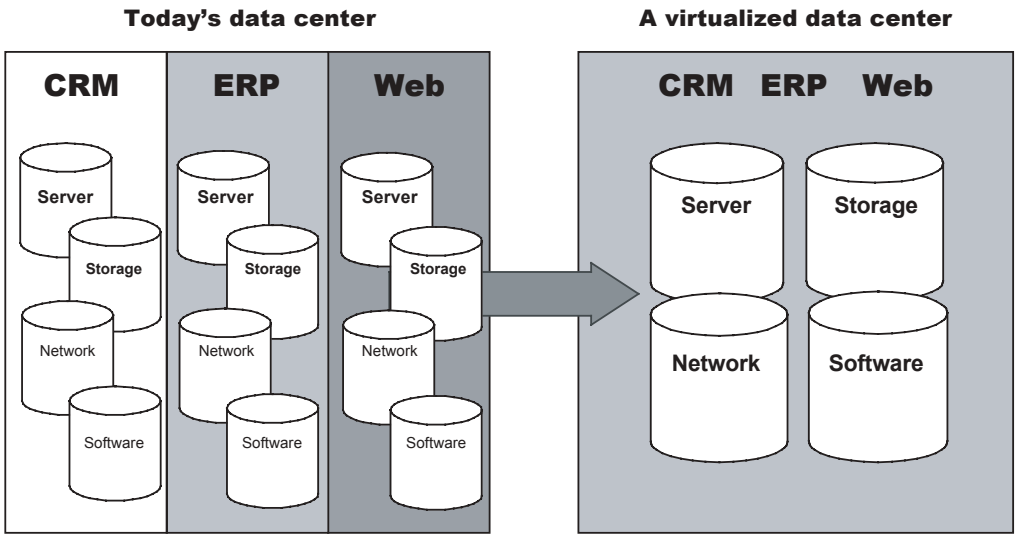

In the 1990s, enterprises expanded data centers by committing servers, storage arrays and other assets to specific applications, leading to disparate and inflexible silos of information.

By implementing a virtualized infrastructure, IT managers can pool assets to meet demand as it arises, while lowering TCO and simplifying server, storage and network administration.

Why virtualize?

- Lower TCO through optimal use of server and storage assets
- Reduced data center management and IT staffing costs
- Greater ability to adapt to changing service level demands
- Simplify interconnect fabric and server management
- Consolidation also leads to reductions in floor space, cabling, power and other physical plant requirements
- Free up expensive server I/O slots to make room for future expansion

fly to meet the demands of the business at hand, data center administrators can fully optimize their infrastructure. No longer are certain servers devoted to specific applications or tasks. In a virtualized environment, workloads shift from one server to the next, depending on the service level demands at the moment.

After weighing budgets against strategic goals, enterprises can implement virtualization in one of two ways: consolidation or virtualization of existing assets.

Consolidation: Organizations move away from several mid-sized systems to a few large multiprocessor servers, or many more low-cost servers. Consolidation might focus solely on data center assets, or it could involve re-hosting departmental servers in a more optimal data center environment. Consolidating from dozens to one or two large servers may be ideal for applications that have frequent spikes in demand. In fact, IDC reports that companies achieve seven-to-one management cost savings when they consolidate processes and servers.[5] A consolidated and virtualized large-server environment would allow for applications to be "resized" on the fly and reallocated as IT needs change, and can dramatically simplify an enterprise's security burdens with fewer devices to protect. In its annual survey of readers, Network World magazine found security to the No. 1 concern — just as it was the year prior.[6]

Virtualization of existing assets: For many — perhaps most — enterprises, this may present the most realistic option. For most data centers, servers of different sizes can accommodate a broad range of applications. With the right infrastructure, IT managers can move workloads between discrete machines, or re-host an application on a larger server during anticipated spikes, such as weekly payrolls runs, monthly sales reports, quarterly financial closes, or annual inventory reports. Such shifts can also occur several times a day, with workloads moving from smaller systems to a larger server to, for instance, run an ERP batch job at night. In the morning, the ERP application is replaced by an order entry application as Web traffic increases during business hours.

Essential to achieving a virtualized infrastructure is simplification. In fact, the mantra for many of today's IT managers is to "simplify, simplify, simplify." In doing so, they can dramatically cut ongoing administrative and operational costs, which make up to 90 percent of the Total Cost of Ownership (TCO) of a server.[7] One company, facing 35 percent annual sales growth and a major

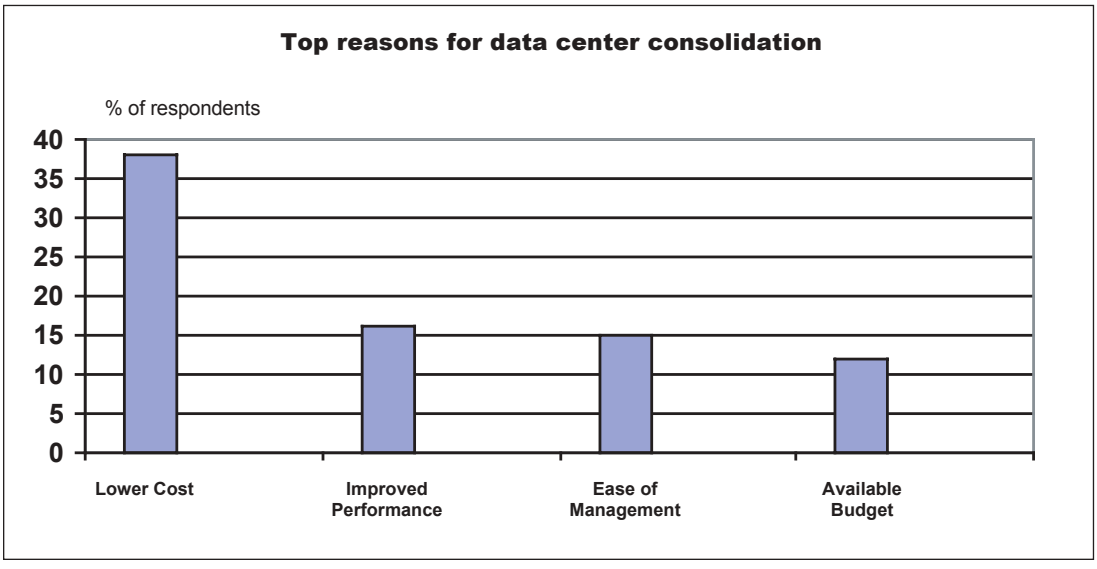

upgrade to all its Oracle enterprise applications, recently consolidated 11 underperforming servers down to one. The company created nine partitions that substantially improved fault tolerance and delivered instant capacity on demand. As a result, application performance increased by 250 percent, while the company trimmed floor space usage by 66 percent and power consumption by 90 percent. Even with equipment acquisition, technical service fees for conversion, and end-user training, the company dramatically cut overall data center TCO, saving $140,000 in annual management software costs alone.

Similar economies can be applied to storage environments. Case studies have illustrated that consolidated storage environments allow administrators to manage up to 20 times more capacity than they previously could using direct-attached storage configurations. Through virtualization, storage capacity can be dynamically reallocated without disruption, with capacity reassigned to a new server on demand. Adding storage capacity can also be achieved without costly downtime, typically allowing data centers to attain storage capacity utilization rates of 80 percent or better. By comparison, mature direct-attached storage configurations tend to suffer from utilization rates below 40 percent.[8]

Achieving an on-demand, adaptive enterprise does much more than make life easier for data center administrators and cut operational costs. It can help transform the image of IT itself, from a cost center to a strategic arm of the enterprise that is vital to customer satisfaction and employee productivity. Take financial institutions, for instance, where some 90 percent of IT budgets are spent on routine maintenance.[9] With an average cost-to-income ratio among U.S. banks of 55 to 60 percent,[10] trimming routine maintenance costs by even 20 percent can pull the cost-to-income ratio below the 50 percent mark.[11] For banks, which often spend between $1,500 to $2,500 to acquire the same customer more than once,[12] unifying data centers and cutting costs quickly become strategic measures that are meaningful to more executives throughout the enterprise.

IDC also has found that many organizations are not only consolidating large data centers, but are actively reducing the number of their data centers, consolidating them into a few strategic locations. In doing so, they can serve a growing and diverse user base — which includes employees, customers and partners — with a more efficiently managed network.[13]

A simpler, more standardized infrastructure also improves the organization's ability to remain agile in the face of constant and unpredictable change. With fewer types of assets to manage, enterprises are more able to operate in an ultimate state of fitness, with better and more efficient utilization not only of servers, but of floor space, cabling, storage arrays, and interconnect devices.

The growing need for bandwidth

Whether consolidating applications onto new servers or optimizing existing assets, organizations moving to a virtualized, adaptive enterprise can expect to see bandwidth requirements continue to soar. In a consolidation scenario, larger systems hosting multiple applications will receive, process and send exponentially greater amounts of data, thus putting pressure on bandwidth between servers and storage arrays. And when virtualizing existing assets, bandwidth arguably becomes even more crucial, particularly as servers dynamically exchange not only data, but the applications themselves. An increase in network bandwidth will be vital in aiding the flow of data between computing resources at a rate that minimizes latencies to produce response times acceptable to application end users.

Analysts at D.H. Brown Associates, Inc. saw this up close in their work with enterprises that had achieved up to 95 percent server consolidation. If the current fabric already experiences 70 percent consumption of available network bandwidth, they argue, the network cannot accommodate a consolidation, which "will require pushing a good deal more data traffic through it." And when planning for network bandwidth, they advise to plan ahead: "If an environment currently hosts 200 users, for example, and the applications are very interactive, plan bandwidth to accommodate 500 users. Such planning will save money later on, and will also avoid blame from people irritated by degrading performance as more users are added to the system."[14]

The current trend toward proliferating Gigabit Ethernet adapters on the desktop and on laptops — along with the growing practice of aggregating Ethernet and Gigabit Ethernet ports — reveal a thirst for more bandwidth even in data centers that have yet to implement a virtualized infrastructure. As time passes, pushing the limits of bandwidth is easier than many enterprises might assume. Just a hundred users with fast Ethernet connections — or only 10 Fast Ethernet users —can easily generate a Gigabit of backbone traffic. Gb Ethernet I/O ports have the potential to evolve a serious problem into a potentially severe one. Organizations that have begun to outfit hundreds or thousands of desktop and laptop system users with Fast Ethernet and Gb Ethernet I/O ports may soon find themselves with network utilization at intolerable levels. With ideal Ethernet network utilization rates most desirable at less than 30 percent, the combination of data-hungry applications fed by fast end-user I/O ports and consolidated application or database server environments driving multiple Gigabits of data per second can quickly bring network utilization to upwards of 80 percent, leading to dropped packets and a reduction in service quality and performance. Today's multi-processor servers equipped with PCI-X 1.0 or PCI-X 2.0 I/O ports, for instance, can quickly saturate even Gigabit Ethernet interconnects, leaving data centers at a potential I/0 deficit (see chart).

Today's servers quickly saturate even Gigabit Ethernet interconnects

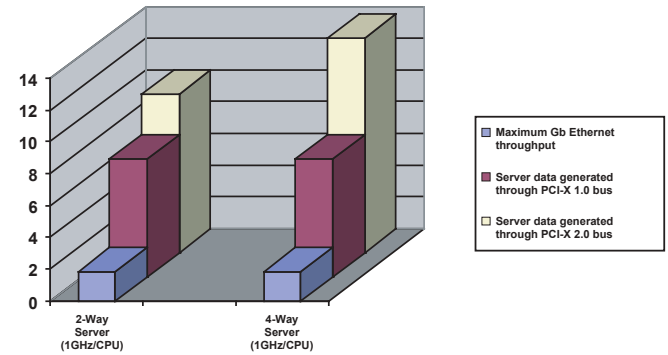

Estimates based on the following assumptions: Single Gb Ethernet adapter yielding maximum throughput of 1.8Gb/sec in full duplex mode; approx. PCI-X 1.0 throughput of 7.4 Gb/sec; PCI-X 2.0 processing power and assuming no significant server-side bottlenecks.

As data demands increase, with applications requiring streaming video and virtualized data centers dynamically partitioning and transferring data and applications, current network bandwidth problems will become more frequent and taxing. Michael Caton and Peter Coffee of eWeek chronicled the trend in July 2004: "Hand in hand with rich media collections," they wrote, "come vastly larger enterprise databases — rapidly passing the terabyte threshold with exabyte (megaterabyte) scale not far over the horizon."[15] Applications such as Oracle Database 10g, with an emphasis on grid and cluster computing, tend to prime users for big bandwidth solutions, creating an eventual upgrade need that many enterprises will find impossible to ignore.

The problem grows even more evident as servers scale upward in a consolidated IT environment: A single, two-way server can generate enough data to saturate a Gigabit Ethernet connection seven times over. Now imagine the bottlenecks faced by 16-, 32- and 64-way servers rapidly finding their way into newly consolidated data centers.

Choosing the right interconnect solution

For organizations pursuing a virtualization strategy, network bandwidth will play a decisive role in their ability to deliver reliable services in a fully integrated and adaptive environment. And those about to move toward a virtualized environment have the unique opportunity to standardize on an interconnect fabric that will serve them for years to come.

Among the high-bandwidth network options available, three have had widely varying levels of adoption in the marketplace: InfiniBand, FibreChannel, and Gigabit Ethernet.

InfiniBand: A low-latency, high-bandwidth solution, InfiniBand suffers from a relatively meager product ecosystem, which may pose threats to data centers that invest in InfiniBand as a long-term interconnect solution. Its relatively shallow market penetration may be a concern to IT managers who prefer a broad pool of InfiniBand-trained talent from which to hire network administrators. In addition, InfiniBand uses new types of costly, parallel copper or fiber cabling, and the lack of a broad competitive marketplace suggests prices will not lower anytime soon. While baseline InfiniBand products allow transmission of data to up to 10 meters, new products have extended InfiniBand multimode fiber networks to 150 and even 300 meters. Deployments of InfiniBand networks can

Trunking, TOE, or big bandwidth?

In response to increasing bandwidth limitations, some IT managers have "trunked up" multiple Gigabit Ethernet adapters. Such workarounds may appear cost-effective initially, but they can require extensive long-term management time and ownership cost. In fact, administrative expenses amount to as much as 90 percent of an asset's lifetime TCO. So the more cables, switches and adapter cards on a LAN, the more expensive it will be to operate until the day comes when even trunked up LANs must give way to a bandwidth upgrade.

By opting to upgrade bandwidth rather than investing in trunking or other workarounds, IT managers will solve today's network efficiency problems while seeing a dramatic reduction in the number of I/O cards, switch ports, cables and other assets that must be deployed and — more importantly — administered over time.

Implementing 10 GbE today will also prepare networks for TOE, RDMA and other enhancements that will become commonplace years from now, while protecting current server investments by freeing up server ports and preparing them for a unified interconnect environment.

take weeks, signaling potential disruptions to busy production environments. Based on semi-proprietary technology, InfiniBand lanes operate at 2.5 Gb/sec (2 Gb/sec with coding overhead), and can be linked to run as fast as 24 Gb/sec.

FibreChannel: Popular in high-performance computing environments such as research and science, as well as for use in Storage Area Networks (SANs), the semi-proprietary FibreChannel interconnect technology operates at a maximum of 10 Gb/sec, though the vast majority of installations still operate at 1 Gb/sec or 2 Gb/sec. However, FibreChannel's marketplace presence is kept relatively small because of its high cost, small ecosystem, long personnel training and potentially long and disruptive installation in production environments. For instance, upgrading from 1 Gb FibreChannel to 2 Gb or 4 Gb requires users to cut the switch cable and use a dongle to attach the 2 Gb or 4 Gb interface connector. Upgrading to 10 Gb FibreChannel requires significant retooling, with all new connectors and coding changes. For this reason, many believe that enterprises will hold onto their 4 Gb FibreChannel networks long after they have outlived their usefulness.

Gigabit Ethernet: Some 85 percent of the world's local area networks (LANs) are based on the global Ethernet standard, which has gradually matured over more than a quarter of a century to become the world's most ubiquitous interconnect technology. While Ethernet LANs certainly carry acquisition and maintenance costs, a sophisticated and robust ecosystem of Ethernet products exists to perpetually drive down the costs of acquiring, deploying and managing Ethernet network parts. Operating at a full 1 Gb/sec, Gigabit Ethernet uses the most affordable optical cabling and other accessories, and leverages the widespread knowledge of IT administrators well versed in Ethernet network administration. As a high-bandwidth interconnect, however, Gigabit Ethernet is becoming too slow to be useful to organizations looking to consolidate applications onto large multi-way servers, or for enterprises aiming to pool existing server and storage assets in a dynamic, transaction-intensive environment. The proliferation of Gigabit Ethernet ports on user desktops also means that a Gigabit Ethernet backbone must feed multiple clients with the same sized pipe. One option to partially overcome this bottleneck is full-protocol offload, though this can simply complicate administration and unnecessarily delay an inevitable network upgrade (see sidebar). In fact, the greatest value of Gigabit Ethernet is that it has prepared the market for an easy, low-risk, cost-effective transition to the next generation of Ethernet-based interconnection technology.

10 Gigabit Ethernet

For the vast majority of adaptive enterprises, the practical, long-term answer to achieving greater bandwidth is 10 Gigabit Ethernet (10 GbE), the latest enhancement to the ubiquitous standard which in its own right has earned IEEE ratification (IEEE 802.3ae).

With all the advantages of Ethernet, 10 GbE leverages the same broadly developed product ecosystem, including optical cabling, switches, routers, servers, computers, workstations and storage appliances. The result of a gradual but persistent evolution of Ethernet technology — from the original technology conceived by Bob Metcalfe to Fast Ethernet and 1 Gigabit Ethernet — 10 GbE retains the compatibility and ease of use for which Ethernet has become known.

The result is a solution that can be deployed today throughout the data center with minimal disruption. In a typical mid-size data center, implementing 10 GbE can usually be accomplished in a matter of hours, rather than the days or weeks required for InfiniBand and FibreChannel installations. This is largely due to 10 GbE's full compliance with exist-

ing Ethernet network infrastructures, which enables an efficient upgrade from Gigabit Ethernet deployments: 10 GbE uses the same low-cost multimode fiber cables, the same management tools, and the same operations procedures for troubleshooting and managing the network. Based on the results of an annual survey of Network World readers, that should be reassuring: The second most worrisome issue for survey respondents was "Acquiring skills in new and emerging technologies,"[16] something IT managers deploying 10 GbE need not worry about. Indeed, for most environments, upgrading from Gigabit Ethernet is a matter of swapping network adapters on servers and installing drivers, as 10 GbE is automatically compatible with major operating systems and thus requires no OS modifications. And because several 10 GbE switches are now available from major switch vendors, with competitive products continuing to enter the market, 10 GbE solutions will grow increasingly cost-effective. These competitive market forces ensure long-term viability for the 10 GbE space.

Creating a unified interconnect environment

A fundamental advantage of 10 Gigabit Ethernet is that it poises data centers for future networking developments that ultimately can create a unified interconnect fabric, leading to vastly more simplified administration and lower TCO.

10 Gigabit Ethernet will also benefit from emerging Ethernet technologies and enhancements such as TCP/IP Offload Engine (TOE) and Remote Direct Memory Access (RDMA) support will further optimize performance and will further pave the road for optimization and virtualization within data center, clustering and SAN environments. RDMA, for instance, will give Ethernet new low-latency advantages, while TOE will further optimize the efficiency of host systems and reduce CPU utilization well below that achieved in non-TOE environments.

Enhancements like these situate IT for "on-demand" or adaptive data centers, enabling easy administration of multiple smaller servers designed to balance loads evenly across the entire infrastructure.

The already growing popularity if 10 GbE offers early proof of the inevitable drop in price of adapters, switches and other 10 GbE components. Gartner/Dataquest, for instance, expects 10 GbE component volumes to triple year over year, from approximately 5,000 ports sold in 2004 to 185,000 in 2007. The resulting volumes and competitive marketplace will drive 10 GbE pricing down as quickly as previous generations of Ethernet components, with the average per-port pricing dropping to approximately $7,700 in the next three years.[17] InStat/MDR also has observed dramatic growth in 10 GbE port sales, marking fourth quarter 2003 sales growth of 300 percent over the third quarter.[18]

10 GbE has distinct advantages over Gigabit Ethernet. Clearly, there are the boosts in network throughput, server productivity and data center efficiencies, and the reductions in application latencies. Current leading 10 GbE adapters have achieved a throughput TCP rate of approximately 7.5 Gb per second, with today's PCI-X 133 bus proving to be a bottleneck to even faster throughput. With PCI-X 266 ready for deployment in early 2005, 10 GbE adapters will be able to generate line rate throughput. Next-generation enhancements including RDMA-enabled TCP/IP Offload Engines (TOEs) will scale to meet the needs of the most demanding end user. 10 GbE operates in full duplex mode, eliminating the need for collision detection protocols. It reduces TCO even further than Gigabit Ethernet, allowing IT managers to consolidate multiport trunks to attain vast reductions in the number of I/O cards, switch ports, cabling and accessory devices. As we have shown, this level of standardization and consolidation can dramatically reduce the lifetime TCO of network devices and lower data center management expenses. 10 GbE also protects current data center investments by integrating seamlessly with existing Ethernet environments and freeing up server ports for consolidation or future expansion.

The same flexible throughput that IT managers gain in server-to-server connections can benefit storage devices as well. While IDC estimates that 70 percent of all storage is direct-attached, less than half of it is utilized, according to Gartner/Dataquest.[19] In fact, smaller low-cost servers and even large multi-processor servers are now commonly installed without any local disk storage whatsoever. This puts even more importance on the ability of the data center's interconnect fabric to transfer data at extremely high rates. The greater the network bandwidth, the more efficiently the data center will be able to dynamically provision near-line and even offline storage. Whether engaged in real-time data access for transaction-based applications or daily backups of terabytes of enterprise data, companies today can ill afford to establish a virtualized server and storage environment with inadequate network band-

width. For instance, backing up 20 terabytes of data over Gigabit Ethernet takes more than 55 hours, but only 6 hours across a 10 GbE connection (see chart).

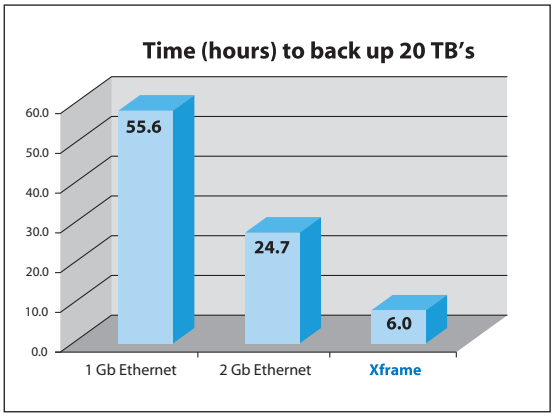

For enterprises seeking to lower the TCO of servers, storage devices and network components, 10 GbE promises even greater economies. For instance, reliance on a high-throughput Ethernet interconnect fabric provides the foundation for a more proven, reliable network. Congested networks with high utilization can discard packets when buffers are full, leading a rise in collision rates. While collisions are typical on an Ethernet network, throughput suffers when collision rates rise to 70 percent. As a result, networks can become unreliable, impacting server and application availability. And any downtime is expensive: In a survey of 94 U.S. IT professionals, TechWise Research, Inc. determined that the average cost per hour of cluster downtime is $145,000[20]. While clusters are not known to crash based on noisy Ethernet traffic, it is clear that a reliable network built to accommodate the growing demands of today's enterprise is far more cost-effective than attempting to make more from a slow, less efficient backplane.

A low-risk, high-reward approach

10 GbE defies conventional business wisdom that suggests substantial returns come only at great risk. Indeed, deploying 10 GbE to upgrade network bandwidth is a decidedly low-risk approach that lays the foundation for a highly optimized and virtualized data center. Because Ethernet is a proven, industry-standard technology, a 10 GbE network upgrade provides the greatest opportunity for a minimally disruptive deployment, while enabling long-term benefits such as lower application latencies, faster back-up rates, and improved application performance.

When upgrading their data centers to 10 GbE, however, IT managers should choose carefully among the network adapters available today. While cost-effective and non-disruptive deployment is crucial for production environments, the right solution should also enable enterprises to establish a dramatically simpler, unified interconnect fabric as enhancements to today's 10 GbE products come to market in the next 24 months.

In selecting from among currently available 10 GbE components, IT managers should demand a core set of capabilities and functionality from solution vendors. These include:

- **Carrier-class Reliability, Availability and Scalability (RAS).** 10 GbE adapters and switches should meet the highest levels of RAS. Ideally, they should provide full Single and Double Error Detection (SEC/DED) Error-Correcting Codes protection throughout the adapter to guard against internal errors and maximize availability.

- **Low-level compatibility with major operating systems.** IT managers should insist that hardware products and associated drivers are fully compatible with Microsoft Windows NT, Windows XP, and Windows 2000, as well as standard UNIX and Linux implementations. This spares enterprises from having to implement OS modifications when upgrading their networks — a common drawback to deployment of other network technologies.

- **Optimized for optical fiber-based networks.** Products designed to reduce CPU utilization, improve performance and reduce latency will feature such specialized technology as built-in EOI (electro-optical interface) transponders and ASICs, and standards-based MSAs (Multi-Source Agreements) for optical components.

- **True Quality of Service (QOS) support.** Products should provide full QOS level support to enable traffic to be consistently classified, prioritized, and queued at line rate. This is especially vital for environments in which multiple Gigabit Ethernet links are aggregated as part of a server consolidation effort.

- **Designed to extract superior performance from multi-way servers.** To take full advantage of the power of large, multi-processor servers, 10 GbE adapters should incorporate multiple unique interrupts, each of which can be associated with individual send and receive operations. Interrupt features supported should include interrupt-acknowledge (INTA), PCI 2.2 Message Signaled Interrupt (MSI), and next-generation MSI-X. IT managers should look for state machine designs, rather than a reliance on firmware to balance optimal performance with low power demands.

- **Stateless offload capabilities.** While full-protocol offload will be available for 10 GbE in the future, enterprises today should be able to achieve multiple simultaneous TCP offloads to maximize CPU utilization and ensure data integrity. To minimize overhead today, adapters should incorporate such features as TCP Large Send Offload that conform to Microsoft and UNIX specifications. For instance, adapters should feature the ability to perform UDP (User Datagram Protocol) checksum on packets that have already been IP-fragmented (UDP CoF). And look for Jumbo Frame support, important for back-end data center applications.

- **Single-mode and multi-mode fiber support.** Support for both single- and multi-mode fiber gives IT managers far greater throughput in short-reach environments (via multi-mode fiber) and maximum flexibility in determining the location of key server assets and stretching the LAN by up to 40 kilometers (via long-reach single-mode fiber).

- **Can co-exist with Gigabit Ethernet.** Adapters should be compliant with the PCI-X 1.0a bus specification and IEEE 802.3z/802.3ab Gigabit Ethernet standard.

- **Intent to support emerging standards.** Products sold today should be designed to easily comply with such emerging technologies as PCI-X 2.0, as well as the 10 GbE copper standard which, once ratified, will provide even more cabling options for data centers. Other networking technologies, including RJ-45 connectors, should also be supported once appropriate standards are ratified for 10 GbE.

- **Commitment for future product support of TOE and RDMA when ready.** Vendors must show a clear and viable product roadmap toward providing TOE and RDMA support, when the market moves to this architecture. These improvement should be made available via a simple upgrade and require no significant retooling.

Protect existing investments

As virtualization, consolidation and other server optimization trends exert ever-greater pressure on data center I/O fabric, temporary work-around solutions such as trunking Gigabit Ethernet ports can be difficult to manage and costly over the long run. And because of their expensive and proprietary nature, short-haul network technologies like InfiniBand and FibreChannel continue to be employed only for specialized and vertical applications.

To achieve immediate bandwidth increases today — and even greater efficiencies in years to come — IT managers should ask their vendors for an RFI regarding a proven, IEEE standard technology whose existing infrastructure and vast ecosystem provides a low-risk, high-return approach to optimizing data center performance. By selecting 10 GbE, IT managers can leverage the economies of scale available only with a truly ubiquitous technology whose protocols, management tools, and troubleshooting techniques already provide the foundation for 85 percent of networks worldwide.

Investing in 10 GbE today is a cost-efficient and non-disruptive move toward outfitting enterprises for an eventual unified interconnect fabric, thus achieving a level of standardization and simplification crucial to implementing a truly adaptive enterprise.

For more information on 10 Gigabit Ethernet, visit www.s2io.com.

Sources

[1] Salomon Smith Barney study cited in InformationWeek, "Every Little Gigabit Helps," Nov. 3, 2003.

[2] Internet World Stats, "Internet Usage Statistics," data updated July 26, 2004, http://www.internetworldstats.com/stats.htm, accessed July 29, 2004

[3] Gartner/Dataquest Global Industries Practice Report, January 7, 2004

[4] IDC report cited in "IDC: Tech's Fortunes are on the Turn," CNET, March 31, 2004

[5] IDC, Worldwide Server Consolidation Forecast, September 2002

[6] Network World, "What's Weighing on You?" July 21, 2003

[7] MetaGroup, Information Week TechWebCast: Server Consolidation Update, July 8, 2004

[8] "Storage Management: How to Simplify, Centralize and Optimize Your Environment, Datalink Corp., October 2003

[9] TowerGroup, "IT US Banking Company IT Spending in 2003: The 'Big Squeeze' on Operational Efficiency," February 2003

[10] Federal Deposit Insurance Corp., "FDIC Quarterly Banking Profile," 1997-2002

[11] CMP Executive Viewpoint, "Banking: IT consolidation lets banks optimize, innovate and thrive in an on-demand world," July 2003

[12] "Multiple Roads, One Destination: A Preview of How Banking Is Headed Toward On Demand," Vikram T.Lund, strategist, Global Banking Industry

[13] IDC, "IT Consolidation Solutions from Hewlett-Packard," 2002

[14] D.H. Brown Associates, Inc. "Enterprise Server Consolidation: Compaq Methodology, Issues and Best Practices," September 2001

[15] eWeek, "Oracle Database 10g Does More, Demands Less," July 6, 2004

[16] Network World, "What's Weighing on You?" July 21, 2003

[17] Network World, "10G: Worth your attention," August 18, 2003

[18] InStat/MDR, "4Q 2003 Ethernet Switch Market Analysis" (#IN0401115LS), March 2003

[19] IDC and Garter/Dataquest research cited in "Storage Consolidation: Reducing the Cost & Complexity of Storage," Adaptec, Inc., 2003

[20] TechWise Research, "Total Cost of Ownership for Entry-Level and Mid-Range Clusters," February 2004

Xframe™ from S2io, Inc.

www.s2io.com

S2io, Inc.
20230 Stevens Creek Blvd.
Suite C
Cupertino, CA 95014
Main Phone: 408-861-1250
Fax: 408-861-1258

S2io Technologies
349 Terry Fox Drive
Kanata, Ontario
K2K 2V6
Main Phone: 613-271-3730
Fax: 613-271-3758

Information: info@s2io.com
Sales contact: sales@s2io.com