

Is Your IT Infrastructure Ready for Tomorrow? — Preparing for a 10GbE Future

Analyst: David Reine

Management Summary

The success, in fact proliferation, of the automobile in the 20th Century has been credited at various times to the success of the internal combustion engine and a variety of technological innovations. No innovation, however, had as big an impact on the success of the automobile as Henry Ford's assembly line, which brought the price of the automobile down to a level that the middle class could afford. Unfortunately, breakdowns were frequent, fuel was difficult to obtain, and roads suitable for travelling were scarce. Until a network of gas stations was put in place, along with a suitable web of paved city roads and interstate highways, the intrepid driver would have to carry spare fuel, as well as spare parts, and know what to do with them, in order to be able to reach his destination. Without these extras, the automobile would not have achieved the rapid success that it did. In fact, the automobile infrastructure was as responsible, if not more so, for the success of the modern automobile as any single technical innovation. The future success of electric cars will depend upon the accessibility of recharging stations to supplement the gas pump at service stations around the world.

A similar infrastructure, an IT infrastructure will be the key to the success of the enterprise data center in the years to come. Because of the poor utilization of existing single-core microprocessor servers, the IT staff of the typical data center has embarked on a journey of consolidation and virtualization to reduce server sprawl, simplify data center complexity, and reduce the total cost of ownership (TCO) of the IT infrastructure. With the availability of the latest x86 microprocessor technology, e.g., the Intel *Xeon 5500* (née *Nehalem*), the data center has the quad-core and multi-threaded processing power to consolidate multiple mission- and business-critical applications on a single dual-socket server. **Unfortunately, consolidating multiple applications on a single platform can put an extremely high workload onto the I/O infrastructure of the server network.** With an existing architecture comprised of 1GbE ports either built into the motherboard or added through PCI-e adapters, an I/O bottleneck could quickly bring mission-critical applications to their knees, destroying whatever TCO gains were achieved through consolidation. **Today's IT staff must be concerned with the enterprise scalability of the I/O infrastructure, its bandwidth, its I/O virtualization capabilities, and its flexibility to support all of the servers and storage devices already deployed throughout the enterprise, as well as any new devices planned for the future.** This bulletin discusses the need for and the advantages of deploying a unified 10GbE for the IT infrastructure. If your data center is going to need higher throughput in the years to come, and most will, please read on to learn how **you can start to prepare now.**

Complexity of Enterprise Data Center

The CIO of every enterprise, large or small, is being overwhelmed by a variety of issues that have a serious impact on the bottom line of the enterprise. The complexity of the IT infrastructure, caused by server sprawl and uncontrolled growth in data storage, along with the under-utilization of these data

IN THIS ISSUE

- Complexity of Enterprise Data Center .. 1
- Innovative Throughput Technologies... 3
- Storage Convergence Options 4
- Conclusion 6

center resources, result in a tremendous waste of these resources and, ultimately, a significant increase in the TCO of data center operations.

Over the past decade, there has been a significant shift in the *use* of IT dollars. Ten years ago, over half of the IT budget was dedicated to server acquisition, with lesser amounts required for their management and administration, and even less for the energy required to power and cool them. Today, while the IT budget is being scrutinized extensively during an economic downturn, the percentage of that budget for server acquisition has fallen dramatically, down to about 30%. The most significant increase, in dollars and percentage, is allocated to server management and administration, now about 50% of the budget. In addition, for every dollar being spent on server hardware, another \$.50 is being spent on energy. With no end foreseen for the rise in the cost of energy, this portion of the budget will continue to rise. In addition, for some data centers in dense urban areas, there is a limit to the amount of energy available to them. For these, the next server deployed could force the enterprise to build a new data center at a cost of millions of dollars. Unfortunately, a significant portion of that energy is wasted, with as little as 15-to-20% of the process or being utilized.

In order to reduce TCO and create a more flexible infrastructure with superior performance, the IT staff must improve server utilization and reduce administrative overhead for mission- and business-critical applications. **Consolidation and virtualization are two of the strategies that many data centers are implementing in order to lower server count, improve server utilization, reduce energy demand, reclaim floor space, and redeploy IT resources to higher value projects, while maintaining, or improving, reliability.** Innovation in the form of multi-core, multi-threaded CPUs, such as Intel's *Xeon 5500*, enables the IT staff to deploy fewer physical servers throughout the enterprise. A virtualized environment provides a flexible server infrastructure to support multiple virtual machines that can provide the resources needed to deploy not only today's mission- and business-critical applications, but also to enable the data center staff to manage the risk associated with deploying tomorrow's applications in a changing economic environment. Virtualization provides a variety of benefits to the data center, such as reduced capital and operations expenses, and improved IT agility. **However, fully optimizing TCO benefits requires a holis-**


tic view of the data center, including creation of a shared storage area network (SAN), integrated management of physical and virtual assets, and advanced network storage, as fewer servers become responsible for processing more information.

When you change the paradigm of the data center from physical to virtual, you need an I/O infrastructure that can support the increased, mission-critical communications workload with wider I/O pipes that can maintain a higher throughput level not only between servers, but also from server to storage. You need a communications fabric to support a variety of applications for hundreds, or thousands, of users. **The enterprise needs to deploy an I/O infrastructure that will maximize enterprise ROI, reduce TCO, and lower risk.**

In order to reduce data center TCO, the IT staff must simplify the management of server and storage resources. They must streamline data center management with a single, unified network infrastructure. **A unified network can provide the enterprise with a cost-effective means to scale storage capacity and address the challenges of supporting high-bandwidth, mission-critical applications.**

One area where the IT staff *has been* optimizing data center operations for over two decades is the storage network, where direct-attached storage (DAS) has evolved into the file-based network-attached storage (NAS) and block-based storage area networks (SAN), designed to consolidate enterprise storage. SAN technology first appeared as the only viable means for a data center with multiple servers to access shared, centralized storage. The first open systems SANs employed a Fibre Channel (FC) architecture at 1 Gb/s as the transport mechanism to pass SCSI commands between the host and the storage target. More recently, many SMBs and enterprise departments have deployed iSCSI SANs, utilizing the integrated infrastructure of the 1Gb/s LAN to access centralized storage. Today, the majority of enterprise data centers have implemented an infrastructure consisting of a network of HBAs, switches, and storage arrays driving their FC SANs at speeds of 2Gb/s and 4Gb/s. Because FC did not lower the TCO to connect the many storage islands in the SAN down to Ethernet's level, many enterprises are looking at alternative technologies in order to lower the costs of data center operations and simplify the management of the data center.

Exhibit 1 — Comparing the Stacks



Innovative Throughput Technologies

The consolidation and simplification of the data center are based on two conflicting goals: simplification and optimization. Virtualization helps to optimize resource utilization, but it often creates additional complexity with the management of an increasing number of applications on a single server divided into multiple virtual machines. Today, the typical enterprise data center has an Ethernet LAN to connect servers to a SAN (typically FC or iSCSI) to connect those servers to centralized storage devices. Some enterprises may also deploy a third network, perhaps InfiniBand, to cluster servers within a high-performance computing environment. With three separate networks to manage, the IT staff must cope with added complexity and a redundant infrastructure. To reduce cost, the data center needs to use innovation to reduce infrastructure.

Fibre Channel

Fibre Channel has been a standard protocol since 1994. Its popularity as a storage consolidation tool took off, along with the dot.com bubble, in the late ‘90’s as the only method for a data center full of servers to access centralized storage. Initially introduced at 1Gb/s, FC has slowly evolved to 2 Gb/s, and, today, to 4 Gb/s. The specification for 8 Gb/s is complete, however adoption has lagged, as the 8Gb FC ecosystem has been slow to deploy. One significant reason is the cost, as 8Gb carries a significant cost premium above 4Gb, and also above 10 GbE, at a time when the enterprise does not want to spend one penny more than it has to. Unfortunately, a FC infrastructure is very expensive, requiring a dedicated fiber optic storage network, with expensive, storage specific FC switches and multiple FC HBAs installed in each server accessing the SAN. Fibre Channel and traditional networks have stacks of layers, where each layer in the stack represents a set of functionality. The FC stack consists of five layers, FC-0 through FC-4. (See Exhibit 1, above.)

In addition to acquisition costs, a FC SAN is very complex to manage, with administrative costs adding significantly to the TCO of the data center infrastructure. Furthermore, FC SANS are perceived to be very reliable, as they deliver fixed length frames, sequentially, with robust error checking. As a result of being on a self-contained network, FC SANs are not as susceptible to security risks, human error, or hardware failure. However, they are complex and have limited scalability in terms of device addressability. Many data centers currently are investigating other high-speed networking options to enable a simplification of data center I/O. One of these options is to move to *Fibre Channel on Ethernet (FCoE)*¹.

iSCSI

iSCSI², on the other hand, is a low-cost, IP-based storage-networking standard designed to link data storage facilities to internet-connected hosts. It was ratified by the Internet Engineering Task Force in 2003. iSCSI transmits SCSI commands over an IP network, using a traditional seven-layer OSI stack to define the physical and data link layers (as shown in Exhibit 1, above). It deploys quickly and flexibly and can be used to facilitate data transfers over a LAN, a WAN, or an intranet and can manage storage over long distances, in effect, creating a SAN using IP. Because of the higher cost and complexity of FC, iSCSI has become the low-cost alternative to FC. It does not require dedicated cabling – it can run over existing LAN switches and IP infrastructure. It runs on top of TCP/IP and was initially deployed over a 1GbE network, with 10GbE infrastructure now becoming available. An iSCSI SAN can be connected via embedded NICs with integrated software drivers to help lower TCO, or through dedicated iSCSI HBAs to accelerate performance.

¹ FCoE, still under development, will enable data centers that have already deployed FC to preserve their investment and take advantage of 10GbE. See further detail on p.4.
² Internet Small Computer System Interface.

It was initially deployed to connect SMBs and non-mission-critical servers to centralized Tier-2 storage, in storage consolidation and disaster recovery scenarios, but has now met acceptance throughout the enterprise with all of the leading storage vendors, such as Dell, EMC, HP, IBM, etc., offering iSCSI as a prime connectivity option. The extensibility of IP enables larger networks for iSCSI storage, with no physical limit to the distance between initiator and target, only the logical limitation of latency. **However, just because you can run iSCSI over your existing network does not mean that you should!**

The ubiquitous nature of the Ethernet requires iSCSI data to have stricter security measures than FC, which uses zoning services to enable security. **Reliability, availability, and security are critical in sharing your storage over your data network.** It is important to utilize the security features available within iSCSI, such as *access control lists*, *IP security (IPSec)* encryption and authentication protocol, *Challenge Handshake Authentication Protocol (CHAP)* to restrict access to targets, and the use of virtual private networks to segregate data. You must deploy an enterprise-class Gigabit switch, with redundancy, to ensure against storage corruption. In addition, in order to guard against a single point of failure, the data center should deploy multiple NICs, which will also boost bandwidth by load-balancing traffic across two or more links to the same, or multiple switches, using multi-path I/O. The IT staff also needs to ensure that each server manages the file systems on the logical drives that it accesses. iSCSI targets enable this security through an IP address or through an *iSCSI-qualified name (IQN)*. They also can configure the iSCSI targets with CHAP to password-protect sensitive volumes. The IT staff must be sure to take advantage of the security features available to protect their data.

Unlike FC, which requires support from a highly specialized administrative staff, iSCSI can take advantage of the knowledge base provided by existing IP network administrators to help lower TCO. Total costs are lowered even further by reduced controller and cable count, as a result of routing both IP and storage traffic through the same NIC or HBA. Because iSCSI rides over the Internet natively, there are fewer interoperability challenges. In addition, iSCSI arrays are also priced lower than their FC counterparts, making them especially attractive for the SMB or any cost-conscious enterprise department.

Storage Convergence Options

In order to reduce the TCO of the data center infrastructure, CIOs are looking at a variety of options to combine the storage and data networks into a single converged pipe. 10GbE provides the super-expressway to enable the increased throughput required to unify communications and enable network consistency. For the data center that has already deployed a FC infrastructure, one of those options is Fibre Channel over Ethernet (FCoE).

Fibre Channel over Ethernet

FCoE is an innovative technique to enable the enterprise data center to encapsulate FC packets over an Enhanced Ethernet³, creating a converged network with no performance degradation. Still in the early stages of development, the FCoE standard⁴ is not yet ratified, keeping FCoE out of the enterprise data center and off of critical server environments until the standard is passed. This could cause a delay in some virtualization projects where FC SANs have been deployed. When ratified, FCoE will enable the consolidation of FC storage and data networks, sharing IP and SAN traffic on the same pipe. FCoE provides the capability to carry the FC-2 layer over the Ethernet layer (as shown in Exhibit 1, on the previous page), enabling Ethernet to transmit the upper FC layers FC-3 and FC-4. This will enable the data center to reduce the count of under-utilized servers and the number of HBAs required, eliminate many cables and switches, reduce energy consumption, and improve reliability. It also will help to reduce network complexity.

Based on these TCO savings, this consolidation is generating interest. The new specification will enable the data center to leverage 10GbE networks while preserving the FC constructs, allowing a seamless integration with existing FC networks and management software. It will also enable the data center to leverage existing investment in FC resources.

Unfortunately for those who have already invested a significant amount of money in the existing FC infrastructure, it will cost even more to converge it with the data network. The data center will need to populate new FCoE infrastructure, namely, Enhanced Ethernet switches to provide 10Gb/s Enhanced Ethernet service, Converged

³ Enhanced Ethernet eliminates Layer 3 TCP/IP protocols in favor of native Layer 2 Ethernet.

⁴ FCoE is a proposed standard within the INCITS T11 FC-BB-5 committee.

Exhibit 2 — Benefit Comparison

Capability	FC	FCoE	iSCSI
Ethernet to FC Bridging	N/A	Easy	Hard
Management Services	FC	FC	iSNS
Runs on Ethernet	No	DCB	Yes
Routable(LANs, WANs)	No	No	Yes
Host Connectivity	Servers	Servers	Servers, Clients
Cost	High	High	Low

Network Adaptors⁵ (CNAs) to support both Ethernet and FC, and an FCoE forwarder to perform the stateless encapsulation/de-encapsulation function. Furthermore, FCoE will not work across a routed IP network.

As could be expected in this economic down spiral, there is resistance due to the anticipated high acquisition costs, as well as continued high costs for administration of the FC network. Some industry pundits, especially those who work in the FC industry, believe that FCoE will rule the data center. Others believe that this is unlikely due to the high costs involved, especially in view of the availability and rapid acceptance of iSCSI in the data center, specifically iSCSI over 10GbE. You should look for FCoE at the edge of the data center first, serving as a link to FC storage in the data center. Full adoption of FCoE is years away, at best.

iSCSI over 10 GbE

Today's medium-sized and large-enterprises are facing storage challenges, created by server virtualization, never experienced before. They face the requirement to redo the network for more storage capacity and better management at lower cost every day. Because multiple virtualized applications support heavy workloads on a single physical server, they are competing for throughput. This demand can easily exceed any single 1GbE connection on today's open system servers. **iSCSI over 10GbE provides them with the best opportunity for flexible, scalable storage with easier manageability and lower TCO than any other alternative.** Software iSCSI initiators are available on every x86 system to support LAN,

NAS, and SAN traffic. These initiators work with the NICs embedded on the motherboard of every rack, tower, and blade server.

For CIOs still deploying DAS or a 1Gb iSCSI SAN, or those with a FC SAN that are looking to reduce their overhead costs, deploying an iSCSI SAN over a 10GbE network will enable them to lower costs and improve performance. The 10GbE protocol, defined in 2002, is the fastest of the Ethernet standards, with a smooth evolution from 1GbE to 10GbE and a ten-fold improvement in throughput. 10GbE supports full duplex links, which can be connected by switches, but it does not support half-duplex operation or *CSMA/CD*⁶. Combined with a 10 GbE switch, iSCSI can deliver significant cost benefits to the enterprise data center, especially in the enterprise using unified storage. While client systems can access file data using CIFS and NFS, servers can access business- and mission-critical application data through the iSCSI protocol. In addition, pre-terminated iSCSI cables with *SFP+*⁷ connectors typically are less expensive than traditional FC cables, and often approach the same cost as 1 GbE options. In addition, dual-port SFP+ iSCSI adapters reduce per port energy consumption, which lowers the TCO. See Exhibit 2, above, for a comparison of the benefits of each I/O infrastructure.

iSCSI Misconceptions

There are several misconceptions about iSCSI, which has enjoyed wide deployment over the past few years, with the availability of iSCSI arrays from all of the leading vendors. First, it is believed that iSCSI cannot provide the performance necessary for mission-critical enterprise

⁵ A CNA is required to pass FC packets to the SAN and Ethernet packets to the LAN on the same adapter, with one or more Ethernet ports.

⁶ Carrier Sense Multiple Access with Collision Detection.

⁷ Small form-factor pluggable +.

applications. *This is wrong.* **With a 10GbE network, there is ample throughput.** Most of the bottlenecks actually occur in the actual reading and writing of the data. Secondly, iSCSI is more difficult to manage than FC. *This is wrong.* **In fact, management applications for iSCSI SANs are now integrated into the arrays as standard, wizard-based, or self-managing routines, removing much of the complexity of daily administration.**

iSCSI Advantages

Some of the more obvious TCO advantages of using iSCSI over 10GbE are:

- Support for 1GbE and native IP routing
- Support for standard and familiar Ethernet switches, routers, and cabling
- Highly scalable – extending the range of the SAN beyond the 10 km limit of FC to the “boundaries” of the Internet. This enables the enterprise to share data with branch offices and field offices around the world, as well as enabling a better disaster recovery solution using a pair of iSCSI SANs
- Because iSCSI packets are routable using standard network routers, it is easier to set up a disaster recovery site at a remote location/branch office than with FC, which requires converters, special cables, etc.
- Enables the data center to optimize storage for consolidation and protection for disaster recovery
- Possible consolidation of front- and back-end network staff
- With embedded NICs, iSCSI is much less expensive than FC

When the IT staff initiates the decision-making process for iSCSI vs. FC on 10 GbE, there are many factors that need to be included. Among them are performance, reliability, security, cost, scalability, flexibility, and ease of use. iSCSI meets or betters FCoE for each category.

Conclusion

If your enterprise data center is using FC today, and you are comfortable with the administrative complexity and cost, FCoE provides your enterprise with a viable solution for unified communication. However, this familiarity adds nothing to your data center environment except cost. If your data center does not have FC deployed, or the costs of FC are overrunning your budget, iSCSI over 10GbE may be the solution for you.

You need to be prepared for any unification project. Do your planning ahead of time to define

your application set, performance requirements, scalability needs, security levels, and availability. FC SANs are not intrinsically more secure than iSCSI SANs. It is simply that the complexity of FC that makes it virtually impossible to accidentally connect to one.

For many data center administrators, lowering costs and improving productivity are top priorities. To meet these challenges, every IT manager should give serious consideration to the economic and efficiency benefits of virtualizing their data centers with an iSCSI SAN infrastructure running over a 10GbE network.



About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

David Reine is Director, Enterprise Systems for The Clipper Group. Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. He joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- ***Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 123. (Please dial “123” when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, Clipper Notes, and “clipper.com” are trademarks of The Clipper Group, Inc., and the clipper ship drawings, *“Navigating Information Technology Horizons”*, and *“teraproductivity”* are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.