White Paper

How Do I Get to Ethernet from Here?

Jason Blosil, Joshua Konkle, NetApp
April 2009 | WP-7071-0409

HOW TO MOVE TO A UNIFIED ETHERNET FABRIC

Consolidating all of your data networks onto a single Ethernet fabric is a great way to reduce cost and improve efficiency. Consolidation creates business and technical benefits, such as cost reductions associated with equipment purchases as well as reduced operational expenses associated with training, power and cooling, and overall network management. The benefits of data network consolidation are available to businesses of all sizes. This white paper is a primer for consolidating your data networks onto a single Ethernet fabric.

TABLE OF CONTENTS

1 INTRODUCTION

Napoleon Bonaparte commissioned the design and development of the Arc de Triomphe, which was completed in 1836, 30 years after being commissioned. The colossal monument was placed on a historical axis leading from the Louvre Museum to the outskirts of Paris, and now stands at the center of 12 avenues and streets. This location was aptly named the Place de l'Etoile, or Star Square, because so many avenues and streets converge on the site, and because avenues and streets were essential to communication and commerce in the 1800s. Today, no matter what your vantage point in central Paris, you can still see and find your way to the Arc de Triomphe. It is a powerful symbol of victory, communication, and progress.

In the 21st century, networks have become our digital avenues and streets. Networks are an essential part of any data center, large or small. Today, most data centers use one or more network types, which may include the Internet or other user data networks, storage, and voice networks, to transmit digital data. Each network uses a different technology to transmit data. Digital data is one of the core assets of almost any modern business; therefore, the data storage network is of particular focus because it manages the storing and retention of these assets.

For data center storage networks, the most common technologies in use today are Ethernet, Fibre Channel, and, to a lesser extent, Infiniband. Ethernet is the network technology of choice for the Internet and other user networks, and is rapidly gaining in popularity as a storage network. Fibre Channel is the dominant technology for storage networks, trailed by Ethernet, then Infiniband. It is common for data centers to include multiple single-purpose networks for data storage, with each storage network requiring separate capital and operational budgets. In good economic times, the capital and operating cost required to maintain separate networks may be justified to support optimizing the environment around other priorities, such as higher performance or additional support for new business applications. However, priorities shift in turbulent times. Though capital purchases, spare parts, power, and cooling remain a constant requirement, budgets tend to fall under increased scrutiny. Moreover, a stressed economy may drive increased data center consolidation because of corporate mergers, reductions in the work force, or outsourcing. Further, new data center projects may be postponed and existing data centers may be required to offer new services without the support of additional funding.

What if you could reduce both capital and operational expenses by consolidating onto a single storage network technology? The continual demands on IT to accomplish more with less money drive the need to reduce the number of different technologies in your storage network. By consolidating onto a single solution, those demands can be met by utilizing just one storage fabric, one designed for redundancy and high availability.

Advances in Ethernet as a storage fabric are positioning it as the best way to consolidate data center network technologies onto a unified fabric. The advances or enhancements to the Ethernet standard are collectively referred to as Data Center Bridging (DCB) (see WP-7052: "Converged Enhanced Ethernet—Good for iSCSI SANs"). DCB is commonly marketed as Converged Enhanced Ethernet (CEE) or Data Center Ethernet (DCE), and allows Ethernet to meet storage performance requirements that previously would have required Fibre Channel network technology. Moving a storage network onto an Ethernet fabric offers compelling advantages, including lower capital costs and improved operational efficiencies from simplified management. Just as the Arc de Triomphe represents a convergence point for avenues and streets, yesterday's pathways of commerce and communication, Ethernet is positioned today as a convergence point for modern pathways of commerce and communication in the form of digital storage networks.

This paper discusses the benefits of Ethernet as a unified fabric and also covers practical scenarios for data centers of all sizes that want a single data center storage network technology based on a unified Ethernet fabric.

2 CAPITAL AND OPERATIONAL DRIVERS FOR UNIFIED ETHERNET FABRIC

Storage administrators, IT directors, and CFOs should ask themselves this question:

> "If we reduce the number of storage network technologies from a few to one, will it produce meaningful results?"

The answer is, yes. Moving to a single storage network technology promises to reduce capital expenses and introduce operational efficiencies. A unified Ethernet fabric for storage immediately reduces the amount of capital equipment by eliminating the Host Bus Adapters (HBAs), cables, and switches required to support single-purpose

storage network technologies. Single-purpose storage network technologies, such as Fibre Channel and Infiniband, offer the advantages of dedicated road maps and a controlled, isolated network environment. As a result, the use cases are focused, which can facilitate high-speed communication and physical security, since the network is not shared with other types of data traffic. But, because of its specific functionality, the market for single-purpose networks is generally smaller and more specialized, resulting in higher deployment and management costs.

Data centers with many servers will commonly employ multiple network adapters in order to provide more bandwidth and redundant connectivity to user and storage networks. Deploying networks with multiple technologies increases cost and complexity. For example, blade server adoption drives the need for HBA and cable efficiency. Using a unified 10Gb Ethernet fabric, one can manage the fabric bandwidth to service both user and storage requirements. Further, one will see a continued reduction in adapters, cables, and switches as a result of deploying only one network technology.

Operational expenses are reduced because less hardware in the data center results in reduced power and cooling requirements. This also increases data center density, characterized by utilization of hardware per square foot, by eliminating the extra active hardware and spares needed to maintain two or three separate storage network types.

Moving to a unified Ethernet fabric doesn't remove the required skills of a storage team. A storage team's knowledge of SCSI, Fibre Channel, and storage fabric management is transferrable to the new unified Ethernet fabric. Over time, IT organizations can optimize their procurement and deployment around a unified Ethernet fabric instead of multiple single-purpose networks. Therefore, rather than spending time chasing down cable, adapter, and switch vendors, storage teams can spend that time focusing on storage provisioning, fabric design, and improved service levels. For example, data centers have servers that connect to user and storage networks. Each server requires input from multiple teams to address specific network technologies. Consolidating onto a unified Ethernet fabric reduces the operational overhead associated with procurement and installation. In addition, ongoing configuration and driver or firmware maintenance are simplified.

As an organization, you may be removing multiple single-purpose storage networks, but you aren't removing the logical storage fabric or a storage team's specialized knowledge. All the benefits of a unified Ethernet fabric combine to produce meaningful results, such as reduced administration, lower capital costs, and improved operational efficiency.

3 ETHERNET STORAGE BEFORE UNIFIED ETHERNET FABRIC

Ethernet has been used as a storage network technology for two decades. In 1989, the Internet Engineering Task Force (IETF) approved as a standard Network File System (NFS) RFC 1094. NFS is a specification for sharing a file system over network technologies. Over the years Ethernet became the de facto standard for file sharing networks. File sharing networks are similar to storage fabrics as they relate to the requirements for redundancy and uptime. However, access to file sharing networks expands beyond host servers to individual workstations, and, in some cases, redundancy between workstations may be necessary.

File sharing networks are comprised of multiple file sharing protocols, including Common Internet File Sharing (CIFS), Hyper Text Transfer Protocol (HTTP), FTP (File Transfer Protocol), NFS, and others. NFS and CIFS are the dominant file sharing protocols within an organization's Local, Metro, and Wide Area Networks (LAN, MAN, and WAN).

CIFS is primarily used in Microsoft® file sharing environments. For example, Windows® 3.11 for Workgroups was the first to introduce early CIFS functionality in the form of Server Message Blocks (SMBs). NFS is primarily used for UNIX® engineering applications and high-density server virtualization environments. High-density virtualization environments are those in which more than 100 server virtual machines are used by the virtual machine operating system, for example, VMware® ESX Server (see TR-3428: "NetApp and VMware Virtual Infrastructure 3 Storage Best Practices"). UNIX engineering applications, including products like IBM/Rational ClearCase (see TR-3677: "Data Protection: IBM Rational ClearCase Environments").

Both CIFS and NFS have significantly improved since their introduction in the early 1990s. NFS experienced a major upgrade in 2003 by adding security, high availability, internationalization, and performance (SHIP) features, all supported by Data ONTAP® 7.3 or higher. Microsoft delivered improvements to CIFS, or SMB, in the form of SMB 2.0 in Windows Server 2008. SMB 2.0 added high availability and performance features, all supported by Data ONTAP 7.3.1 or higher. Over time, as Ethernet networks increased from gigabit to 10-gigabit speeds and network-attached storage (NAS) appliances became faster and more highly available, these file protocols grew beyond file serving to

also support other applications. Today, many customers run their databases, virtualized servers, content management, and archival applications over NFS or CIFS.

Both protocols continue to add features that bring them in line with their respective functionality. In SMB 2.0, Microsoft added durable handles to support high availability. Durable handles allow a client to maintain access to a file opened on an SMB 2.0 server even during Ethernet network disturbances, a feature of NFS since it was introduced. In NFSv4, the standards group added support for Discretionary Access Control Lists (DACL). DACLs are lists of users who are authorized or unauthorized to read and write data. DACLs have been part of SMB and Microsoft environments since the introduction of NTFS. These added features, especially DACLs, support a unified Ethernet fabric. DACLs ensure that, as more systems converge onto Ethernet, security can be enforced between user networks and storage networks.

For decades, Ethernet has played a vital role in storing data for applications using a file format. However, some applications required storage that appeared like a local SCSI disk and offered low latency. Therefore, after the turn of the millennium, the iSCSI protocol was introduced to transmit block-level storage commands over Ethernet network technology. iSCSI is a network storage protocol that encapsulates SCSI commands into TCP/IP packets for transmission over a standard Ethernet network. iSCSI met the storage application requirement by presenting networked storage as a local disk with low latency. As is the case with other block storage protocols, such as Fibre Channel, early iSCSI deployments relied on dedicated HBAs at the host to transmit or “initiate” packets to the “target,” the network disk or logical unit (LUN). The iSCSI HBA managed the encapsulation of SCSI commands into TCP/IP packets and provided the additional benefits of being able to boot the host from a remote volume or to process security protocols, such as IPSec. Later advances by Microsoft and others introduced software drivers included with OS distributions that offered many of the same benefits only with the use of standard Ethernet adapters. The contribution of free software initiators greatly reduced the cost and complexity of deploying iSCSI storage networks and also increased the availability of iSCSI to a broader audience who previously might not have been able to afford the investment in a storage network.

iSCSI is easy to configure and leverages the skills found among administrators of the Internet and user networks, such as IP address configuration and IP routing. The lower entry cost and simple management have generated rapid adoption among small and medium-size enterprises looking to move away from direct-attached storage (DAS) to a shared storage network or storage area network (SAN). Introduced on 100Mb Ethernet, iSCSI gained popularity with the broad adoption of Gigabit Ethernet. iSCSI now supports 10-Gigabit Ethernet for high-bandwidth applications. Performance will continue to improve since iSCSI follows the Ethernet standards road map, which will include 40 gigabit and 100 gigabit in the near future. In just a few years, iSCSI storage has approached a US$1 billion market and remains the fastest growing type of storage network.

Still, some applications required more throughput or lower latency than iSCSI could offer using Gigabit Ethernet. Fibre Channel was the first to market supporting a storage fabric and still represents a market roughly three to four times that of iSCSI in terms of revenue. Since Fibre Channel is a dedicated storage network technology, its road map included higher bandwidth, moving from gigabit to 4-gigabit,and now 8-gigabit speeds, while iSCSI remained tethered to the Ethernet road map, which limited affordable bandwidth to a single gigabit. A large enterprise with many business applications might find that if even one critical application required Fibre Channel while the rest performed adequately over iSCSI, the company would elect to maintain separate storage networks.

For those organizations that require the benefits of Fibre Channel along with the efficiencies of a unified Ethernet fabric, a new standard is expected to be ratified in mid-2009 that will deliver the value of Ethernet to Fibre Channel users by enabling Fibre Channel traffic to run natively across an Ethernet network technology. This new standard is known as Fibre Channel over Ethernet (FCoE) (see FCIA_SNW_FCoE_WP_Final.pdf, published on www.fibrechannel.org).

FCoE will enable you to preserve your existing Fibre Channel storage investment while migrating to an all-Ethernet network technology. FCoE will be a key enabler of greater data center network efficiency by eventually removing the need for single-purpose storage network technology. FCoE offers a very attractive solution for data centers of all sizes, including those with large Fibre Channel investments.

As you consider adopting an FCoE storage network, recognize that FCoE requires switches and HBAs that support DCB. Though available today as a premium feature set, DCB functionality will likely become standard on most Ethernet switches in the future. The enhanced features will likely contribute to improved performance for iSCSI, NFS, and CIFS (see WP-7052: "Converged Enhanced Ethernet—Good for iSCSI SANs"). A small investment in DCB-enabled Ethernet equipment will offer compelling options in the future.

4 HOW DO I GET THERE FROM HERE?

NAS and iSCSI have been used in all sizes of data centers, from remote offices and small businesses to corporate data facilities. Moreover, higher performance requirements can be addressed with 10-Gigabit Ethernet for NAS, iSCSI, and now FCoE traffic with DCB functionality.

DCB allows you to consolidate bandwidth on 10-gigabit paths and share the bandwidth with a variety of traffic types. Further, DCB allows you to prioritize the most important data traffic for optimum performance with guaranteed bandwidth on a virtual LAN (VLAN). VLANs enable you to logically isolate an Ethernet network from another while sharing the same physical connection. VLANs can offer additional security as well as improve the ability to predict performance.

In order to determine which unified Ethernet fabric strategy is best for your organization, let's consider three scenarios:

- New data centers
- Existing data centers
- Small businesses and remote offices

SCENARIO I—NEW DATA CENTERS

For a new data center deployment, several options are available for both file and block storage access. In the case of file access, the choices are pretty clear: CIFS for Windows and NFS for all other operating systems, each without any special hardware or software requirements. For block access, the choice for a storage network protocol really depends upon the timing of the deployment, budget and performance requirements, and the available skill set.

If the data center will be deployed prior to the summer of 2009, then iSCSI is your only choice for block storage over a unified Ethernet fabric. iSCSI offers the flexibility of performance and the price points to address the majority of data center applications. If the data center is deployed in the second half of 2009, then FCoE becomes a viable option. Regardless, it may be wise to invest in DCB-enabled components if FCoE may be a requirement in the future, since the functional enhancements of DCB also extend to other Ethernet traffic.

If your new data center is optimized for budget over performance, then iSCSI is the right way to go for block storage traffic. iSCSI offers the greatest flexibility for balancing performance and overall cost in terms of equipment acquisition, personnel training, and overall manageability. iSCSI performs well on single Gigabit Ethernet and 10-Gigabit Ethernet with or without DCB. As a result, iSCSI can be deployed cost effectively according to performance and budget objectives.

If performance is the primary driver for the new data center, then FCoE should be considered. Fibre Channel, including FCoE, has a reputation for lower latency and improved performance, as well as a robust set of management and discovery tools to increase storage network productivity. Moving to FCoE requires the acquisition of DCB-enabled storage network adapters and network switches, which is currently a premium over standard Ethernet equipment, including 10-Gigabit Ethernet. Regardless of the protocol decision you make today, moving to a unified Ethernet fabric with DCB Ethernet will offer investment protection and flexibility as your business requirements evolve (see Figure 1).

The IT skills required to manage the new data center should also be considered when selecting network protocols. Though FCoE runs over Ethernet, the protocol is still Fibre Channel, and includes the same feature set and management paradigm. So, IT staff must be hired and trained for Fibre Channel in addition to Ethernet. As with NFS and CIFS, iSCSI runs over TCP/IP and therefore leverages a common skill set with the rest of the user network.

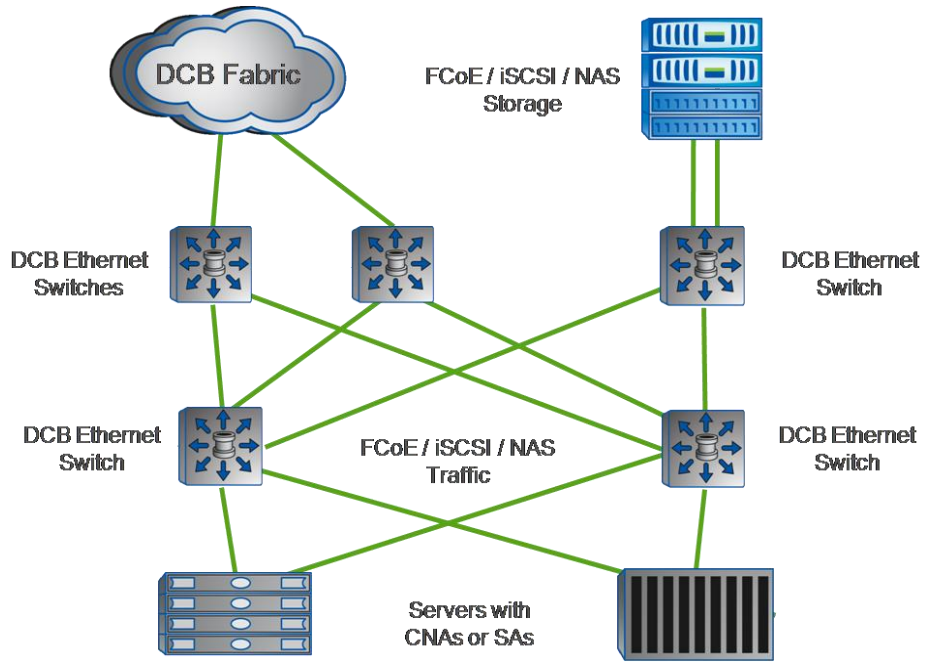

Figure 1) Example of a unified Ethernet fabric using DCB components with FCoE enabled. Switch topology represents a standard edge to core configuration.

SCENARIO II—EXISTING DATA CENTERS

Nearly all large data centers have a combination of storage network technologies that address specific business application needs using Ethernet and Fibre Channel SANs (see Figure 2). Ethernet has historically been used exclusively for the Internet or user networks and network-attached storage, such as NFS version 4 and SMB2.

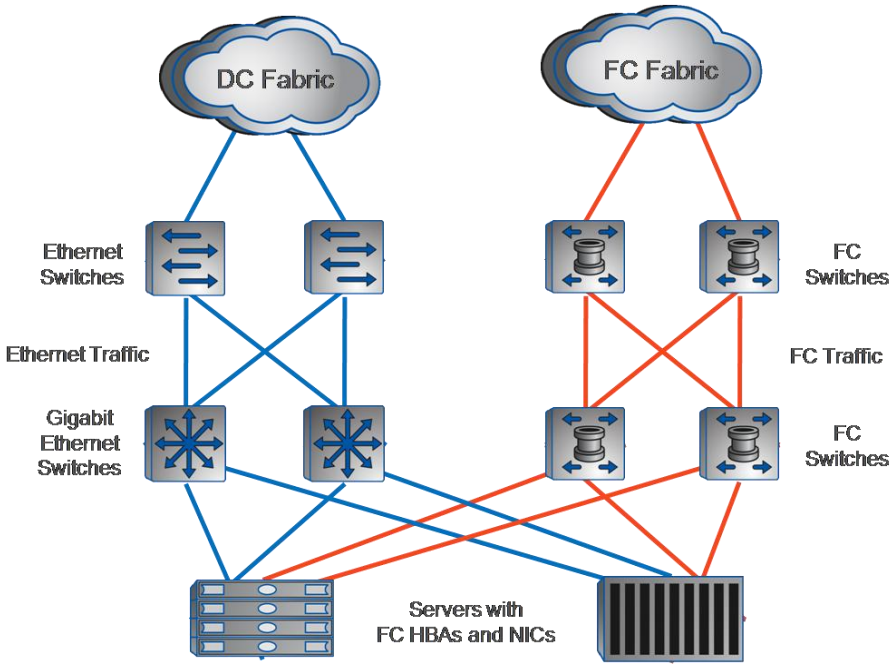

Figure 2) Typical redundant Ethernet and FC SAN infrastructure.

For highest performance, Fibre Channel is the predominant protocol. In order to preserve the investment made in Fibre Channel equipment and training, data centers with large Fibre Channel installations will likely transition to an FCoE network as a preferred method to converge onto a unified Ethernet fabric as a single storage network technology. Transitions may begin at the edge, where new servers and edge switches will be updated to support FCoE, in order to provide access to existing Fibre Channel storage. Switches such as Cisco's Nexus family support Ethernet, FCoE, and Fibre Channel protocols simultaneously to provide maximum flexibility during infrastructure and economic transitions to a unified Ethernet fabric (see Figure 3).

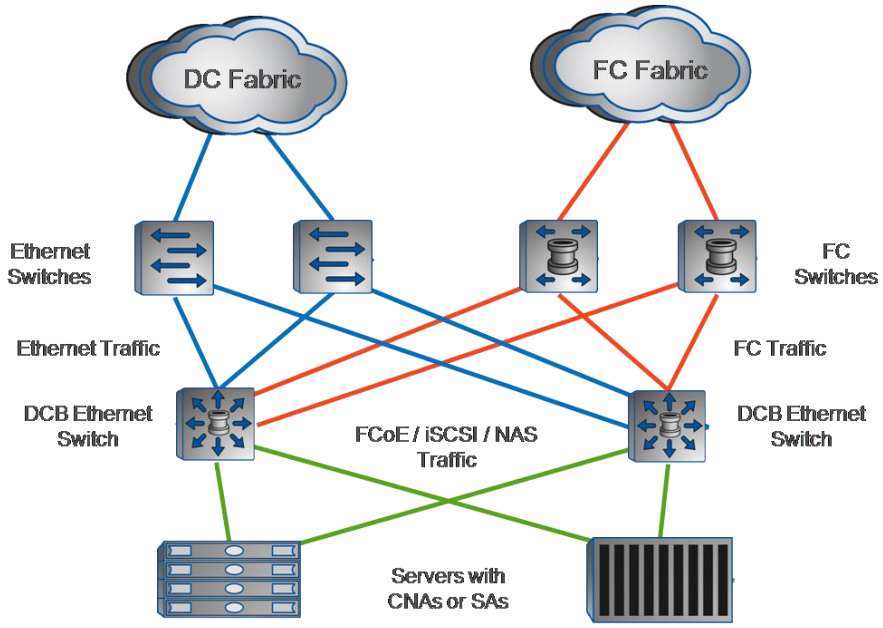

Figure 3) Servers with FCoE adapters attached to DCB edge switches with FCoE enabled simplify the network topology.

From the edge to the core, and eventually to the storage, FCoE will enable the transition to a unified Ethernet fabric while preserving your investment in Fibre Channel storage (see Figure 4).

The move to DCB-enabled switches to support FCoE does not preclude the use of iSCSI on the same network for business applications that may not require the improved performance associated with FCoE. iSCSI is increasingly being used for infrastructure servers and more broadly adopted in virtual server deployments. A unified Ethernet fabric offers flexibility that allows you to match your business applications with the proper storage technology without altering equipment purchase decisions.

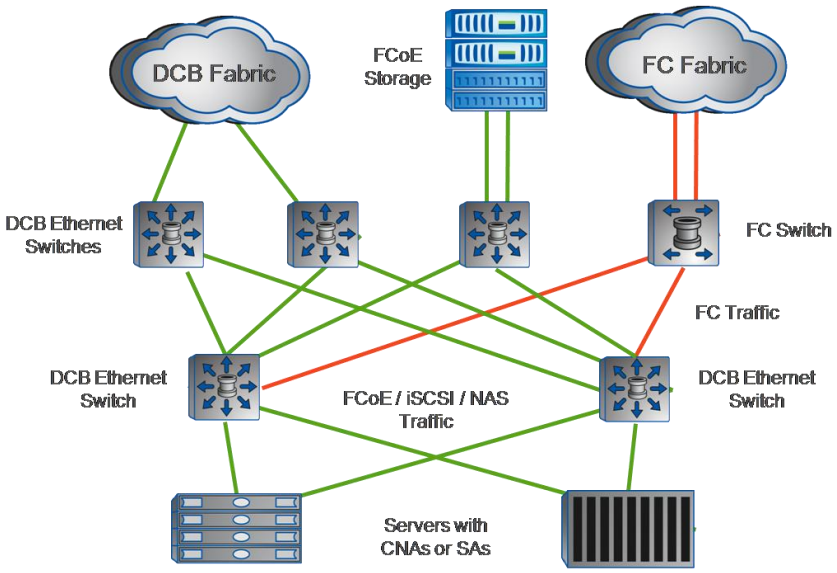

Figure 4) A fully converged network architecture with end-to-end FCoE plus FC investment protection.

IT business units may be engaged in data center consolidation for a number of reasons, such as to:

- Reduce costs by improving business application density using server virtualization.
- Reduce environmental costs by removing equipment.
- Improve manageability and reduce labor expenses.
- Improve service levels for customers with a unified Ethernet fabric.

The need to consolidate the data center exists today in order to experience the benefits of a unified Ethernet fabric through the use of iSCSI, FCoE, NFS, CIFS, and SMB2. The skill set to manage storage network technology is similar for the rest of the unified Ethernet fabric, such as user data and voice network technology.

If you currently have Fibre Channel storage and are looking to deploy major consolidation projects by the end of 2009 or thereafter, then you have a clear path to a unified Ethernet fabric with FCoE.

SCENARIO III—SMALL BUSINESSES AND REMOTE OFFICES

SMALL BUSINESSES

The choice between FCoE and iSCSI largely depends upon budget and business application requirements. Generally, the smaller the business, the higher amount of direct-attached storage, since DAS is cost effective and easy to deploy in very limited quantities, for example, a single server. As a result, the networking skill set of IT personnel tends to be optimized around Ethernet for user data and some NAS. In these cases, moving to shared storage and consolidating on Ethernet favors the use of iSCSI, NFS, CIFS, and SMB2. Ethernet can satisfy the performance requirements of the majority of business applications while also providing exceptional value and ease of transition from DAS.

Businesses on the larger end of small may often have deployed Fibre Channel networks already and will benefit from FCoE as a transition from single-purpose Fibre Channel storage network technologies to a unified Ethernet fabric. In some cases, iSCSI and NFS may replace Fibre Channel, if the performance profile is met with gigabit or 10-Gigabit Ethernet, with or without DCB.

REMOTE OFFICES

Ethernet can be managed over longer distances than a single-purpose Fibre Channel storage network technology. In order to operate efficiently, it is critical that medium-sized businesses manage the data for remote offices—management and reliability are paramount. Typically, remote office performance is generally less demanding, and single Gigabit Ethernet can satisfy the business application performance requirements, such as e-mail, operations software, and file sharing (NFS, CIFS, and SMB2).

A unified Ethernet fabric for remote offices delivers operational efficiency and necessary performance without compromising reliability or adding additional cost for single-purpose Fibre Channel network technology.

5 SCALE-OUT STORAGE FOR A UNIFIED ETHERNET FABRIC

In 2008, the Internet Engineering Task Force (IETF) approved NFSv4.1, an extension of NFSv4, to support parallel input and output (I/O) processing called parallel NFS (pNFS). pNFS supports parallel I/O access to data servers in a standard way so that clients and servers can be productized and supported across multiple brands, such as NetApp. pNFS is NOT limited to NFS for parallel I/O. pNFS supports three storage types:

- Blocks—The pNFS client uses the SCSI protocol, including traditional iSCSI, FC, and FCoE.
- Objects—The pNFS client uses the Object-based Storage Device (OSD) protocol. OSD moves the storage management to the storage subsystem, that is, the storage head/controller.
- Files—The pNFS client uses the NFSv4.1 protocol. This layout type supports all NFSv4 and v4.1 operations, including access control lists.

You will have a mixture of networks and storage requirements. Where you are in your application deployment, plus your budget and the size of your business, will influence which path you take in establishing your unified Ethernet fabric. Investing in Ethernet fabrics today will continue to yield positive results in terms of capital expense, operational efficiencies, and evolutionary changes.

6 CONCLUSION

As demands for performance, service, and capacity continue to increase and budgets continue to flatten or decrease, there is pressure to reduce capital expense and operational expense. Companies of all sizes are addressing budget constraints by consolidating. Ethernet presents a clear opportunity to consolidate data center networks onto a single network technology. Consolidating the network technologies for user, storage, and server applications reduces equipment requirements, resulting in lower capital expense and reduced total cost of ownership.

Fibre Channel over Ethernet will allow customers to migrate their FC network technology to an Ethernet network technology with minimal impact on staff training, tools acquisition, and so on. FCoE will complement other protocols, such as iSCSI, NFS, and CIFS, which are already available on Ethernet, to provide the greatest flexibility and cost efficiency for the data centers of today and tomorrow.

The question as to which file sharing or block storage protocol is right for you depends on your objectives. Ethernet addresses multiple goals, including performance, manageability, simplicity, and cost effectiveness. Ethernet offers more options to advance and optimize your data center without having to make network technology decisions to address single-purpose applications.

Just as the Arc de Triomphe is a point of communication convergence in the heart of Paris, Ethernet is a point of network convergence for data centers of all sizes. With Ethernet as your destination, there is a clear way to get there from wherever you are today.