



# Unified Physical Infrastructure (UPI) Strategies for Storage Networking

**Using Centralized Architectures to Achieve  
Physically Consolidated Storage Environments**

WHITE PAPER

## Introduction

Businesses are taking a closer look at how they store data and how effectively that data can be accessed. Today's storage infrastructure satisfies both external business demands, such as 24-hour availability and compliance with government regulations, and internal needs, such as cost control, fast access, scalability and reduced power consumption.

Traditional Direct-Attach Storage (DAS) methods persist in many data centers due to low off-the-shelf hardware costs. However, enterprise applications are currently driving the adoption of more centralized networking environments such as Storage Area Networks (SAN) or Network Attached Storage (NAS) to realize the benefits offered by consolidating applications, hardware, and cabling infrastructure. These innovations in storage consolidation are driving toward faster transport speeds within the data center – 10 Gb/s and beyond – and require robust physical infrastructure technologies to support these advancements.

This paper examines the importance of implementing a solution based on Unified Physical Infrastructure (UPI) principles to optimize the performance of centralized storage architectures. After a review of the differences between DAS, NAS, and SAN methodologies, the paper surveys storage protocols (Fibre Channel, iSCSI, Fibre Channel over Ethernet) commonly deployed for each method. The paper then demonstrates the importance of deploying a UPI-based solution consisting of high-performance connectivity, cable management, and routing pathways to support these storage network architectures and manage risk within the data center.

## Is Your Data Center Ready for the Future of Storage?

Many data centers operate multiple parallel networks: one for IP networking, one for block I/O, and a third for server-to-server Inter-Process Communication (IPC) used by High-Performance Computing (HPC) applications. However, running several applications over multiple networks has a significant unfavorable impact on the company's bottom line. In response IT organizations are increasingly consolidating multiple applications onto a smaller number of more powerful servers to reduce their need for data center space and to lower operating costs.

The newest wave of consolidation leverages a unified high-speed fabric to extend network, storage, and IPC traffic over a single physical infrastructure and eliminate redundant adapters, cables, and ports. These innovative I/O consolidation techniques are made possible by leveraging new transport technologies such as Fibre Channel over Ethernet (FCoE) with emerging "data center backbone" class network switching platforms linked by 10, 40, or 100 Gb/s fiber and copper cabling systems.

However, a consolidated storage infrastructure can lead to I/O bottlenecks and added latency anywhere from individual disk mechanisms or arrays to LAN or SAN paths connecting the disk array and servers. Data center stakeholders and IT managers need to review their existing network, current facilities, and future build-out plans to determine how best to support the growing data storage and information exchange requirements placed upon the enterprise's IT networks and infrastructure. This process involves weighing the potential benefits that centralized storage architectures can offer.

## DAS, NAS, and SAN: What's the Difference?

Data center storage systems range from basic DAS farms to more complex NAS and SAN deployments. As businesses grow, storage systems evolve from scattered DAS islands to more centralized and/or tiered environments that store data in shared physical spaces under the control of a single network or file management system. Centralized storage infrastructures, often using high-performance arrays and storage networks, also serve to reduce backup windows and data restore times.

### Direct-Attached Storage

DAS refers to a storage space directly attached to a server or workstation. A typical DAS system is made of one or more enclosures holding storage devices such as hard disk drives, and one or more controllers. These systems are designed to extend storage capacity one server at a time while providing high data bandwidth and access rate.

The main protocols used in DAS are Small Computer Systems Interface (SCSI), serial-attached SCSI (SAS), and Fibre Channel, and the interface with the server or the workstation is made through a host bus adapter (HBA). While straightforward to design and deploy, DAS systems have some significant drawbacks:

- Each DAS drive must have a significant reserve capacity above the required storage space to ensure adequate room for data, resulting in a considerable amount of unused storage capacity.
- High-availability redundant paths are rarely deployed to all DAS servers, adding risk which can lead to server downtime resulting in lost productivity, and lost revenue.
- IT staff must deploy and administer scheduled backups on every server; periodically monitor, upgrade, and service the servers; and recover piecemeal from disasters that affect some or all of them.

### Network Attached Storage

A NAS unit is essentially a self-contained computer connected to a network, with the sole purpose of supplying file-based data storage services to other devices on the network. The operating system and other software running on the NAS unit manage data storage, file creation, and file access. NAS units usually do not have a keyboard or display, and are controlled and configured over the network, often by using a typical web browser.

NAS storage methods employ file-based protocols such as Network File System (NFS) or Server Message Block (SMB) and Common Internet File System (CIFS); all three protocols are designed for remote storage architectures, under which computers request a portion of an abstract file rather than a disk block. This makes NAS systems ideal for applications that are shared between multiple clients, such as Web content and e-mail storage. Also, since NAS systems are attached to an IP-based Local Area Network (LAN), they benefit from economies of scale achieved from using a standard IP infrastructure, products, and services across existing LAN and wide area network (WAN) infrastructures.

Several characteristics of NAS systems limit their use for general high-performance, enterprise-wide networked storage:

- Scalability can be difficult due to the need to replicate and manage multiple copies of data as the network grows and additional NAS appliances are deployed.

- NAS system performance, data availability, and associated application performance are dependent on IP LAN traffic and network congestion. The variability in performance is often too great for data-intensive enterprise-class transaction-based applications that require high availability with a high degree of predictability.
- The processing overhead associated with file-based data segmentation limits the speed and efficiency of large block data transfers. This may be unacceptable for data replication applications such as, off-site back-up, disaster recovery, and business continuity that require the quick and reliable exchange of extremely large amounts of data.

### Storage Area Networks

SAN systems connect storage devices, typically disk arrays, tape libraries and optical arrays, using high-speed networking architectures, allowing all clients and applications to access all storage disks. In this way SANs help to increase storage capacity by enabling multiple servers to share space on the storage disk arrays. This flexible disk sharing arrangement simplifies storage administration: cables and storage devices do not need to be physically moved from one server to another, and storage devices appear to the operating system as locally attached.

SANs are preferred in larger enterprise environments due to reliability of data transfer, and their ability to scale as business requirements evolve over time. The common application of a SAN is for the use of transaction-based data access, such as databases or high usage file servers that require low latency, high-speed I/O block-level access to the hard drives. To achieve this SANs often use a Fibre Channel (FC) fabric topology (see Figure 1), which is an infrastructure specially designed to provide faster and more reliable storage communications than NAS protocols, such as NFS or SMB/CIFS.

All major SAN equipment vendors also offer some form of Fibre Channel routing solution, and these bring substantial scalability benefits to the SAN architecture by allowing data to cross between different fabrics without merging them. These offerings use proprietary protocol elements, and the top-level architectures being promoted for each are radically different.

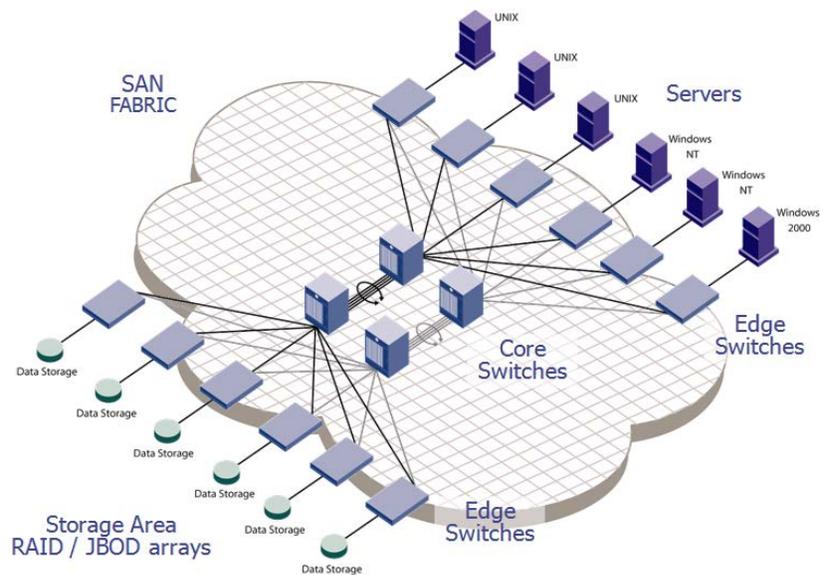


Figure 1. Example Fibre Channel fabric topology.

SANs are also preferred over DAS and NAS deployments to meet disaster recovery requirements in enterprise environments due to their ability to replicate data quickly between storage arrays and over long distances. In a disaster recovery plan, organizations mirror storage resources from one data center to another remote data center, where it can serve as a hot standby in the event of an outage. Traditional physical SCSI layers can support links only a few meters in length, which is not enough to ensure continued operation in a disaster. However, Fibre Channel links can be extended using the Fibre Channel over IP (FCIP) protocol to move FC-based storage traffic across thousands of kilometers via IP WANs, thereby enabling enterprises to deploy disaster recovery plans that meet continuity requirements.

Other SAN disaster recovery features include the ability to allow servers to boot from the SAN itself, which allows for a quick (<10 minutes) data duplication from a faulty server to a replacement server.

### IP-Based Protocols Support I/O, Storage Consolidation

A high-performance physical infrastructure is critical to support all storage goals and initiatives – from disaster recovery to virtualization – as any given physical link may carry mission-critical data at any given time. Under unified network fabrics, the physical infrastructure is no longer segmented between one cabling and connectivity system for transporting high-performance mission-critical information and another cabling system designed solely for lower-speed, non-priority data. The need for a high-performance physical layer is especially critical for virtual environments, where multiple virtual machines are run over the same physical link, demanding full channel bandwidth in order to maximize network efficiency for the highest possible port utilization.

Several techniques for merging I/O traffic (LAN, SAN, IPC) onto one consolidated physical infrastructure have been adopted for some data center applications with immediate benefits: simpler cabling infrastructure, higher density and space efficiency, lower power consumption, improved cooling airflow. To date, these I/O consolidation technologies address small niche applications such as InfiniBand for high performance clusters and grid computing or iSCSI for SAN device connectivity across a WAN network. As LAN speeds increase with Ethernet running at 10 Gb/s, FCoE and iSCSI will likely become credible alternatives to Fibre Channel.

#### **FCoE**

Fibre Channel over Ethernet (FCoE) is a standards-based encapsulation of Fibre Channel onto an Ethernet fabric. The FCoE specification transports native Fibre Channel frames over an Ethernet infrastructure, essentially treating FC as just another network protocol running on Ethernet alongside traditional IP traffic. These improvements allow existing Fibre Channel management modes to stay intact and enable the seamless integration of existing Fibre Channel networks with Ethernet-based management software.

Both Fibre Channel and FCoE require the underlying network fabric to be lossless. This requirement is fulfilled by using additions to the Ethernet standard such as Data Center Ethernet (DCE) or Converged Enhanced Ethernet (CEE) which provide the ability to pause communication traffic on a priority basis, enabling a lossless fabric implementation. Support of FCoE modules in new converged I/O switches that can be connected via Fibre Channel to existing storage networking switches ensures coexistence with prior and future generation FC switching and services modules. This extends network flexibility and protects existing Fibre Channel investments.

FCoE also enables consolidated I/O by using a converged network adapter (CNA), which appears as both a Fibre Channel HBA and an Ethernet network interface card (NIC) to the server, but appears as a single Ethernet NIC to the network. With separate LAN and SAN connections, a typical server may have five or six unique Ethernet and Fibre Channel connections with the associated copper and fiber optic cabling complexity. Using a CNA, the number of adapters and cables required can be reduced by 50% or more (see Table 1).

This reduction in cabling infrastructure can enable high equipment densities (such as dense rack-mounted and blade server areas) to support server virtualization initiatives, which often require many I/O connections per server. It can also result in improved data center airflow and cooling efficiencies due to the smaller size and number of cabling pathways associated with consolidated I/O networks.

**Table 1. Example Cabling Reductions Achieved Using FCoE Architecture**

Unique Connections (non-FCoE)	Consolidated LAN Connections	Consolidated SAN Connections	Reduced Number of Cables (FCoE)
4	From 2 to 1	From 2 to 1	50%
6	From 4 to 1	From 2 to 1	67%
8	From 6 to 1	From 2 to 1	75%

### iSCSI

The iSCSI protocol simply allows a host to negotiate and then exchange SCSI commands with a storage device using IP networks instead of Fibre Channel. By doing this, iSCSI takes a popular high-performance local storage bus and allows it to be run over wide-area networks, creating or extending a SAN. In particular, iSCSI SANs allow entire disk arrays to be migrated across a WAN with minimal configuration changes, in effect making storage "routable" in the same manner as network traffic.

Unlike some SAN protocols, iSCSI requires no dedicated cabling so it can be run over existing switching and IP infrastructure. As a result, iSCSI is used as a low-cost alternative to Fibre Channel which requires dedicated infrastructure. Since iSCSI runs on top of IP, latency is much higher than Fibre Channel which makes iSCSI unsuitable for enterprise-class transaction-based applications such as Enterprise Resource Planning (ERP) systems that have low latency and high bandwidth requirements.

### UPI Strategies Optimize the Storage Networking Physical Infrastructure

Whether deploying today's best-in-class storage networks or preparing for tomorrow's converged I/O fabrics, the need for a high-performance, high-reliability physical infrastructure becomes even more critical. The physical infrastructure must be able to grow as data storage and information exchange needs grow, must be agile to easily implement network reconfigurations as business needs change, and must ensure suitable bandwidth for proper application performance. All of these items must be addressed in a cost-effective manner within organizational budgetary constraints.

The storage network physical infrastructure is comprised of two key elements: cabling and connectivity to transport data signals; and physical infrastructure hardware to patch, house, and route storage network cabling. The former provides network performance and bandwidth while the latter ensures network flexibility and scalability; both contribute to overall network performance and reliability. While copper cabling and connectivity can be used for storage networking, such as the front-end of a NAS system, the vast majority of storage networks are implemented using a fiber optic infrastructure (see Figure 2).

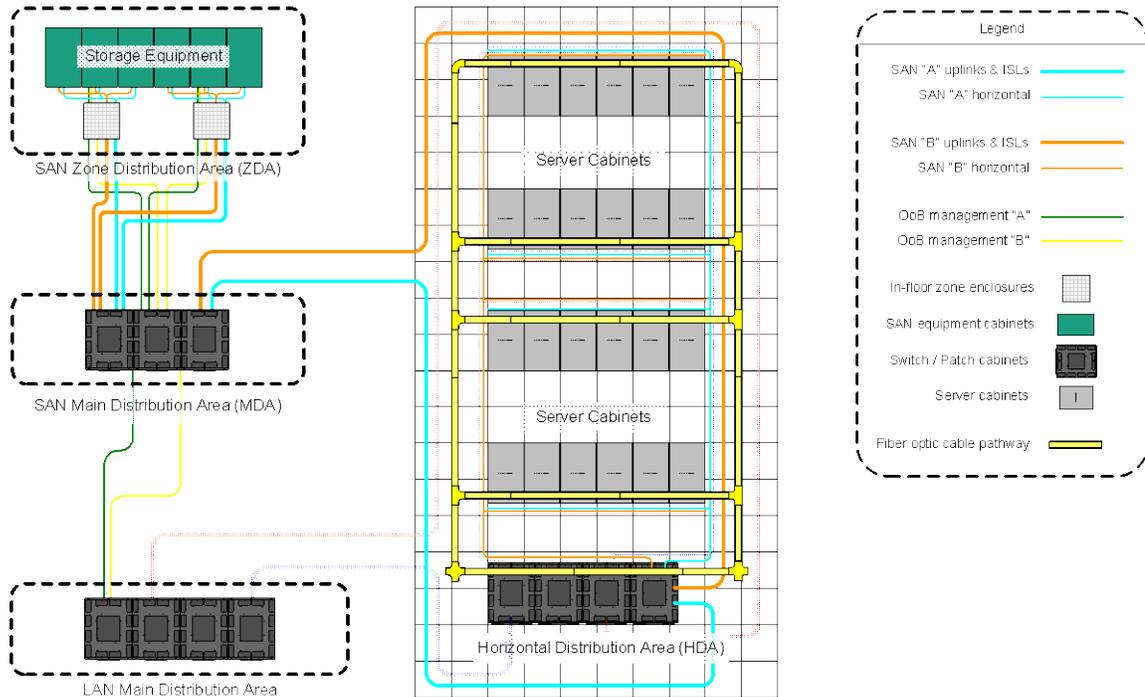


Figure 2. Centralized storage network based on modular layout and design.

Speed to deploy, increased installer productivity, efficient maintenance, and fast response to changes are challenges that all storage area network administrators face. The specific attributes required from the storage network physical infrastructure are:

- **Scalability:** The physical infrastructure and layout should easily scale as business information exchange and data storage needs grow without resorting to “rip and rebuild” methods.
- **Modularity:** Storage network gear should feature common storage gear form factors, accessories, mounting methods, and layout topologies should be specified so any component of the system can be easily replicated or utilized elsewhere in the system.
- **Density:** Escalating data center real estate costs translate to higher port counts and increased port and cabling densities throughout the storage area network, particularly in the core and edge switching areas.
- **Performance:** Ensuring high performance includes specifying high-bandwidth fiber optic media, low-loss connectivity, and efficient design topologies for cable management, distribution area layout, and pathway routing.
- **Agility:** The demands to provide high levels of availability have rapidly increased and “no downtime” is now the operating norm. The physical infrastructure must allow for rapid installations, permit easy and intuitive upgrades, and support rapid turn-around time for network reconfigurations.
- **Risk Mitigation, Redundancy and Disaster Recovery:** Storage networks, by nature, must be redundant in every aspect. This requires separate physical infrastructures for “A” and “B” fabrics with redundant switches, patch fields, pathways, and fiber optic connectivity to ensure uninterrupted operation. Furthermore, the fiber optic infrastructure must accommodate continuity provisions including off-site disaster recovery centers, hot-standby fabrics, and recovery from archival back-ups.

### Modular Design for Storage Networking Reliability, Scalability

To achieve a sturdy centralized storage architecture, *PANDUIT* recommends using a modular design that can scale from one switch with fewer than 100 host servers to multiple director-class core switches with several thousand host servers. These designs feature all of the following storage network elements and UPI-based best practices:

#### Design

- Divide SAN into functional areas (as shown in Figure 2) to enable manageable, scalable growth through modular network sections.
- At a minimum, create a Main Distribution Area (MDA, see Figure 3) for SAN Director switches, an Equipment Distribution Area (EDA) for servers, and a Zone Distribution Area (ZDA, see Figure 4) for storage devices (RAID arrays, tape libraries, etc.).
- Deploy separate SAN "A" and SAN "B" fabrics, to provide redundancy among all storage network connections.
- Specify sufficient media capacity for the initial commissioning and for future network growth.
- Defined upgrade path whether through capacity increases within existing enclosures and patch fields or via straightforward addition of more enclosures, patch panels, racks, and/or cabinets without costly physical migration, cut-over activities, or "rip-and-replace" schemes.

#### Deploy

- Employ cable management best practices, including physical protection, bend radius control, pathway sizing, and in-cabinet routing of mission critical fiber optic cabling and connectivity, to ensure maximum bandwidth and uncompromised uptime.
- Route cables along pathways in-cabinet, overhead, and under floor, that promotes proper airflow and hot aisle / cold aisle separation using thermal management best practices.
- Use high-density modular patch fields (see Figures 4 and 5), where all fiber optic adapter panels/cassettes, patch panels, and enclosures are 100% interchangeable, to speed new equipment connections or re-cabling of existing equipment.
- Logically arrange and label patch fields, where all ports on storage equipment, edge switches and servers are replicated, to ensure that patching connections and/or reconfigurations can be made using discrete short-run patch cords.
- Use physical network security devices to restrict unauthorized access to network ports, and to prevent inadvertent removal of critical patch cord connections.
- Provision for out-of-band management for equipment monitoring and control.



Figure 3. Two-cabinet SAN MDA with dual core switches and 4RU fiber distribution enclosures.



Figure 4. High-density fiber optic patch field for SAN cross-connect.



Figure 5. Typical SAN storage equipment ZDA housed in 2 ft x 2 ft in-floor enclosure.

### Matching the Right Fiber Cabling to your Storage Network Application

Several fiber grades find use in data center installations, each with a different reach and bandwidth capability (as summarized in Table 3). With the rise in popularity of 8 Gb/s Fibre Channel transceivers and the recent introduction of 16 Gb/s devices, any high-performance storage network deployment should consist of OM3, OM4, or Panduit's Signature Core™ Fiber Optic Cabling System.

It is possible to transmit high-speed 8 Gb/s and 16 Gb/s Fibre Channel signals short distances over OM2 50/125 µm fiber. However, data transmission rates will continue to increase, with 16 Gb/s and 32 Gb/s native Fibre Channel and 40 Gb/s and 100 Gb/s Ethernet. It is extremely unlikely that legacy OM1 and OM2 grade multimode optical fiber will support those data rates even over very short distances.

In addition to choosing the proper fiber grade to ensure bandwidth over the required distance, it is equally important to select a fiber optic cabling system that can support the various physical and electro-optical needs of high-speed storage networking. A fully-modular, pre-terminated fiber optic cabling system provides the best combination of performance, reliability, and network agility. This pre-terminated fiber optic cabling system should feature low-loss connectivity – patch cords, cassettes, interconnect cables, trunk cable assemblies – to allow for maximum signal integrity. Low-loss connectivity also enables heightened network flexibility by building more headroom into the channel for patch fields and cross-connects and thereby increasing the number of physical infrastructure topologies available.

Finally, high-precision testing is recommended to guarantee end-to-end signal integrity across the storage area physical infrastructure. Differential Mode Delay (DMD) testing is used to certify bandwidth performance of OM3 and OM4 multimode fiber, and Bit Error Rate Testing (BERT) is the definitive measurement of fiber optic network performance as defined by IEEE 802.3ae standards.

**Table 3. Fiber Cabling Characteristics for 8 Gb/s Fibre Channel and 10 Gb/s Ethernet**

Fiber Type	Core/ Cladding (µm)	minEMBc <sup>[1]</sup> (MHz•km)	Max. Reach at 8 GbFC (m) <sup>[2,3]</sup>	Max. Reach at 10 GbE (m) <sup>[2]</sup>
OM1	62.5/125	200	21	33
OM2	50/125	500	50	82
OM3	50/125	2000	150	300
OM4 <sup>[4]</sup>	50/125	4700	125	Not defined
OS1	9/125	Not defined	10,000	10,000

[1] Minimum Calculated Effective Modal Bandwidth.

[2] Data based on use with 850 nm VCSEL-based serial transceivers for multi-mode fiber and 1310nm long wavelength serial LR transceivers for single-mode fiber.

[3] Per INCITS working draft proposal FC-PI-4.

[4] OM4 fiber grade has been agreed to by the ISO/IEC JTC 1/SC 25/WG 3 as of February 2008 and has been proposed to the governing ISO/IEC/IEEE standards bodies for inclusion in the ISO/IEC/IEEE 802.3 and INCITS specifications.

## Conclusion

An increasing number of enterprise applications are requiring the speed and power of centralized storage environments to consolidate IT assets, handle dramatic storage growth, support virtualized environments, and meet backup and disaster recovery requirements. Although multiple protocols and architectures are available to achieve a centralized storage infrastructure, a SAN is the preferred architecture for enterprises that require high speed, availability, and scalability, with Fibre Channel or FCoE the current SAN protocols of choice due to its speed, reliability, and flexibility to meet various application and networking requirements.

Data center stakeholders need to specify and implement a robust storage network physical infrastructure that can reliably support current transport speeds and business requirements, and which provides a migration path toward future speeds and technologies. For example, next-generation storage networks will increasingly rely on converged protocols such as FCoE over fiber optic media, and many will leverage consolidated I/O and virtualization technologies to achieve networking and application efficiencies. Modular storage network designs based on UPI principles enable more effective and efficient management of storage application, hardware, and infrastructure assets to support uptime goals, maintain business continuity, and scale with business needs. By mapping centralized storage onto a robust physical infrastructure, organizations can mitigate risk across the network to build a smarter, unified business foundation.

## About PANDUIT

*PANDUIT* is a world-class developer and provider of leading-edge solutions that help customers optimize the physical infrastructure through simplification, increased agility and operational efficiency. *PANDUIT's* Unified Physical Infrastructure (UPI) based solutions give Enterprises the capabilities to connect, manage and automate communications, computing, power, control and security systems for a smarter, unified business foundation. *PANDUIT* provides flexible, end-to-end solutions tailored by application and industry to drive performance, operational and financial advantages. *PANDUIT's* global manufacturing, logistics, and e-commerce capabilities along with a global network of distribution partners help customers reduce supply chain risk. Strong technology relationships with industry leading systems vendors and an engaged partner ecosystem of consultants, integrators and contractors together with its global staff and unmatched service and support make *PANDUIT* a valuable and trusted partner.

[www.panduit.com](http://www.panduit.com) · [cs@panduit.com](mailto:cs@panduit.com) · 800-777-3300