

IP Storage Controllers for Next-generation Unified Network Storage Systems

A vision of the future of IP Storage Networking

Published: January 2007

Abstract:

Today, IP storage systems are built using off-the-shelf components. While this approach is functional, it is not at all optimal. We predict that the next-generation of intelligent unified IP storage systems will be based on purpose-built hardware and software designed from the ground up specifically for the unique requirements of IP storage. These systems will offer dramatically higher performance, enterprise-grade high availability features, lower power consumption, and better scalability, all in a smaller footprint and at a lower cost.

Nimbus Data Systems, a leading provider of IP storage systems and software, and AMCC, a leading provider of embedded PowerPC processors for storage and networking applications, have announced a strategic agreement to develop such purpose-built hardware/software platforms.

We call these optimized network storage platforms *IP Storage Controllers*.

Table of Contents

Table of Contents	2
About this White Paper.....	3
IP Storage: SAN and NAS.....	3
Classes of Storage: Primary and Nearline	4
The Future: Unified Storage on a Unified Network	4
Why IP Storage Controllers?.....	5
The Future of IP Storage Systems	6
Purpose-built Platforms Win.....	6
IP Storage Controllers v. Off-the-shelf Hardware	7
Superior Performance	8
Lower Power Consumption	8
Smaller Footprint	9
Enterprise-grade High Availability	9
Lower System Cost	9
Massive, Flexible Scalability	10
Easy-to-use, Multi-protocol Storage Software.....	10
Summary	12
More Information	12

About this White Paper

This white paper introduces the IP Storage Controller, a purpose-built hardware/software platform designed to power the next-generation of low-cost, low-power, high reliability, and highly adaptable IP storage systems. The intended audience is System Builders and OEMs looking to capitalize on the rapidly-growing market for intelligent unified IP storage systems. Storage technologists interested in the architecture and advantages of next-generation IP storage systems will also find this white paper relevant and interesting.

IP Storage: SAN and NAS

Storage area networks (SANs) based on IP have seen tremendous worldwide adoption that has especially accelerated since 2006. There are numerous reasons for the attraction to IP SANs:

1. **Simplicity:** Building networks using Ethernet is widely understood by IT managers, whereas Fibre Channel is more complex and less prevalent.
2. **Compatibility:** IP is more compatible and interoperable across different vendors of switches and adapters than Fibre Channel. All major operating systems, including Windows, Linux, Solaris, HP-UX, AIX, Netware, and MacOS, support iSCSI.
3. **Affordability:** Ethernet networking is significantly less expensive. NICs are approximately $\frac{1}{8}$ th the cost of Fibre Channel HBAs, and Gigabit Ethernet switch ports approximately 80% less expensive than Fibre Channel switch ports.
4. **Security:** iSCSI offers CHAP for server-to-storage authentication, IPsec for storage traffic encryption, and header and data digests for maximizing data integrity.
5. **High Availability:** iSCSI, by leveraging proven features of Ethernet, has built-in multi-pathing without requiring additional expensive server-based path management software.
6. **Management:** iSCSI storage systems are compatible with existing network management and monitoring software without requiring new Fibre Channel-specific management tools.
7. **Disaster Recovery:** IP storage systems are ideal for remote replication since they already offer IP connectivity. No costly Fibre Channel to IP routers are necessary.
8. **Application Support:** Microsoft offers broad iSCSI support in Windows and key applications like Exchange and SQL Server. VMware has integrated iSCSI support into its software to enable affordable virtual computing.

The statistics below further demonstrate the rapid market adoption of iSCSI and IP SANs:

- The iSCSI storage systems market is experiencing exponential growth, from \$216 million in 2003 to an expected \$4.7 billion by 2007. (Source: IDC)
- 42% of small and midsize businesses will implement IP SANs by the middle of 2006. (Source: Enterprise Strategy Group)
- 1.5 million servers will be connected to iSCSI storage by 2006. (Source: Gartner)

At the same time, the popularity of network-attached storage (NAS) continues to grow. With the ability to provide simultaneous read/write access to network shares across multiple operating systems, NAS offers an important feature that SANs lack, at the expense of performance and expandability. For creating sharable file storage, however, the simplicity of NAS is difficult to ignore, especially among small businesses. Since NAS systems are also based on IP, the emergence of IP SANs has spurred interest in “unified” IP storage systems that combine SAN and NAS into a single array based on a common and ubiquitous network: Ethernet.







Classes of Storage: Primary and Nearline

Just as SAN provides “block-level” storage ideal for server applications and NAS provides “file-level” storage best suited for multi-platform file sharing, Primary storage and Nearline storage represent distinct classes of storage suited for different purposes. Primary storage is typically defined as high-performance, high-reliability storage based on premium hard drives and fully redundant RAID levels like RAID 10. Nearline storage is typically defined as cost-effective, high density storage based on inexpensive hard drives and parity-based RAID such as RAID 5.

Primary storage, often featuring hard drives spinning at up to 15,000 rpm, is best suited for database and other server-based applications. The reasoning is logical: these hard drives are faster (twice as fast as most Serial ATA drives) and feature better MTBF (mean time between failure) ratings than lower cost Serial ATA drives. Nearline storage, on the other hand, offers the advantages of a significantly lower cost per GB and higher density, at the expense of performance and reliability (about 30% lower MTBF according to hard drive manufacturers).

The Future: Unified Storage on a Unified Network

With even small businesses now depending on a wide range of server applications for daily business operations, storage can no longer be viewed with a “one-size-fits-all” philosophy. Requirements for optimal storage vary as new applications are brought online. In some cases, SAN is the preferred solution, but just as frequently, NAS is the desired choice. Critical applications demand the reliability of Primary storage, but its high cost makes Nearline storage compelling for the majority of stored data. The table below shows how different applications match up in terms of SAN, NAS, Primary storage, and Nearline storage.

Applications and Their Ideal Storage Type		
Application	Type	Class of Storage
 Email servers	SAN	Primary
 Databases and ERP	SAN	Primary
 Web and content servers	NAS	Primary
 Network shares	NAS	Nearline
 Imaging and multimedia	SAN, NAS	Primary, Nearline
 Archives and disk-based backups	SAN, NAS	Nearline

Without versatile storage systems that are adaptable to the needs of different applications, IT departments will find themselves deploying “point products” that gradually increase the cost, complexity, and inefficiency of their storage infrastructure. Small and midsize businesses especially need versatile and easy-to-use storage systems since IT management resources and budgets are limited.

Why IP Storage Controllers?

IP Storage Controllers enable System Builders and OEMs to offer inexpensive, easy-to-use, and highly adaptable IP storage systems based on optimized hardware and powerful, multi-protocol software. As an integrated solution, developed from the ground up by Nimbus Data Systems, IP Storage Controllers provide a compelling alternative to single-purpose iSCSI arrays or NAS servers based on off-the-shelf server hardware and generic operating system software, including:

- Up to 4x greater performance
- Up to 80% lower power consumption
- Up to 80% smaller form factor
- 99.999% high availability
- Dramatically lower cost
- Support for concurrent SAN and NAS
- Powerful data protection features
- Easy-to-use administration software

The Future of IP Storage Systems

Today, IP storage systems are built using off-the-shelf components. In other words, IP storage system manufacturers build their products by assembling motherboards, processors, PCI cards, power supplies and other components that are generally available in the marketplace, usually into a rackmount enclosure with multiple bays for hard drives. These systems, when booted, usually run a generic operating system such as Microsoft Windows Storage Server or Linux. As a result, today's IP storage systems, from a hardware perspective, closely resemble the x86-based (Intel, AMD) servers, but with more bays to hold more hard drives and additional software features specific to the requirements of network storage.

While IP storage systems built this way are functional, they are by no means optimal. This reflects in their sub par ratings in the important criteria used in evaluating enterprise storage systems, such as:

- Performance
- Power consumption
- Footprint/density
- High-availability
- Expandability
- Cost
- Flexibility
- Software features

We predict that the next-generation of intelligent unified IP storage systems will be based on purpose-built hardware and software designed from the ground up specifically for the unique requirements of IP storage. IP storage systems based on this new combination of hardware and software will deliver dramatically higher performance, enterprise-grade high availability features, lower power consumption, all in a smaller footprint and at a lower cost.

On January 23, 2006, Nimbus Data Systems, a leading provider of IP storage systems and software, and AMCC, a leading provider of embedded PowerPC processors for storage and networking applications, announced a strategic agreement to develop such purpose-built hardware/software platforms, or IP Storage Controllers, combining processors from AMCC and other industry leaders with software and architecture design from Nimbus.

Purpose-built Platforms Win

When new technologies are introduced, it is common to use software, combined with off-the-shelf hardware, to bring functional versions of the new products to market. A good example of this was Sun Microsystems's early routers; these were based on Sun servers loaded with specialized networking software. It worked, but the solution was expensive, its features were limited, and performance was mediocre. When Cisco came to market with a family of multi-protocol routers with optimized hardware and software, it capitalized on purpose-built hardware and software to surpass the generic offering from Sun, and the rest is indeed history. This scenario is common:

as technologies mature, manufacturers look for ways to optimize price/performance, often by designing hardware/software solutions purpose-built for a specific application.

Not surprisingly, today's first generation IP storage systems are built using the functional method described earlier. Vendors taking this approach include startups and large companies alike, and all these products have numerous drawbacks. Performance is mediocre and often less than that of traditional direct-attached storage. Many IP storage systems lack dual-controller high availability. The scalability (i.e. maximum storage capacity) of these systems is very limited because off-the-shelf motherboards can only support a finite number of hard drives. Power consumption and cooling requirements are excessively high. Most importantly, the software is loosely coupled with the hardware, degrading its ease of use.

The best case against off-the-shelf hardware for next generation IP storage systems can be found in the enterprise storage industry itself. The fact is that no vendors of enterprise Fibre Channel storage systems use off-the-shelf motherboards as the architecture for their products. For example, Engenio, a leading provider of OEM storage systems to IBM, Sun, and SGI, designs its own series of Fibre Channel controllers to power its performance-leading storage arrays. Dot Hill acquired Chaparral Network Storage for its unique Fibre Channel controllers, which were based on custom hardware and software. EMC's Clariion and Symmetrix DMX series and Hitachi's enterprise storage arrays are also based on purpose-built hardware.

The evidence is clear: as the market continues to mature for IP storage, purpose-built hardware and embedded software will replace assembled off-the-shelf solutions.

IP Storage Controllers v. Off-the-shelf Hardware

System Builders and OEMs that select IP Storage Controllers can expect:

1. **Superior performance:** Enables greater storage consolidation in enterprises while meeting the needs of speed-sensitive applications.
2. **Dramatically reduced power consumption:** Increases reliability, cuts ongoing environmental costs, and reduces noise levels.
3. **Smaller footprint:** Increases storage density to reduce datacenter costs in rapidly-expanding storage environments.
4. **Enterprise-grade high availability:** Dual redundant controllers increase availability to 99.999%, reducing downtime to less than 6 minutes per year.
5. **Lower system cost:** Purpose-built hardware eliminates costly add-in cards, off-the-shelf processors, and other components while reducing power supply requirements.
6. **Exceptional scalability and versatility:** SAS technology with 100TB+ storage expansion and the convenience of Primary and Nearline storage in a single system.
7. **Integrated software:** Based on the same multi-protocol HALO software that powers current Nimbus IP storage systems, IP Storage Controllers provide an integrated, optimized hardware/software solution with enterprise-grade data protection features.

Now, let's outline how these advantages are achieved with IP Storage Controllers.

Superior Performance

In IP storage, data is encapsulated into IP packets prior to going “across the wire” and then reconstituted at the other end of the wire. Without additional processing to perform this function in real-time, latency is increased and I/O performance is reduced. As more servers are connected to the IP storage, the processing demands on the IP storage system increase in parallel and net performance declines further. As IP storage looks to move behind Gigabit Ethernet to 10 Gigabit Ethernet, the processing requirement grows ten-fold further.

Using the brute force of more processors in an off-the-shelf system will not solve the problem; in fact, it introduces new ones. Multi-core architectures cannot effectively divvy up the TCP/IP processing load, and generic processors are not optimized for handling TCP/IP processing. Adding cores adds cost, power consumption, heat, and size while also reducing reliability.

IP Storage Controllers incorporate an advanced protocol processor that provides wire-speed processing of TCP and iSCSI commands, enabling IP storage performance far greater than today's typical 100 MB/sec. In preliminary tests, these IP Storage Controllers achieved up to 4 times greater throughput than today's iSCSI storage systems. This will enable IP storage systems match or exceed Fibre Channel storage systems in both throughput and IO performance, while retaining existing advantages in cost, compatibility, and ease of use.

Lower Power Consumption

Energy efficiency has emerged as a top concern for IT datacenters. IDC found that only 56% of total IT costs are related to hardware and software; the remaining costs are ongoing operational expenses including power and cooling costs. With oil prices hitting record levels, businesses are looking to reduce power consumption while increasing both processing and storage capacity, a tough challenge.

Unfortunately, storage systems are significant consumers of energy in today's enterprise datacenter. IP storage systems often use dual Xeon or Opteron processors to handle the processing load of the TCP protocols. Each processor can consume up to 100 watts alone, while additional PCI cards used for internal storage connectivity or RAID functions can add an additional 30-100 watts. A motherboard and its related components can add an additional 200 watts, totaling approximately 430-500 watts, BEFORE the power consumption of hard drives is even considered. This in turn requires that storage systems be outfitted with high capacity power supplies that drive up cost, heat, and noise.

Not only does this level of power consumption increase electricity costs, it also decreases the reliability of the storage. More fans are needed to cool the system. Some server systems are even available with water cooling systems, evidence of just how hot today's general-purpose processors run. As more fans are added, more moving parts are introduced, decreasing overall system reliability. Cooling bills and maintenance costs escalate. Higher operating temperatures also cause systems to prematurely reach the end of their lives, resulting in greater hardware costs and decreased return on investment.

IP Storage Controllers, with a low-power, purpose-built design, use dramatically less power, drawing less than $\frac{1}{5}$ the power of traditional off-the-shelf based solutions. In turn, this reduces IT maintenance costs, increases system reliability, and extends the lifespan of the storage systems. Noise levels also drop dramatically as fewer fans are spinning at high RPM. The emergence of IP Storage Controllers will better enable OEMs and System Builders to meet the needs of tomorrow's datacenters for energy efficient, eco-friendly storage with maximum reliability.

Smaller Footprint

As the cost per square foot in datacenters rises, IT departments are looking for ways to fit more equipment in less space. IP Storage Controllers combine the functionality of several off-the-shelf components, such as the motherboard, processors, memory, network adapters, storage connectivity, and RAID, into one single board purpose-built for IP storage. As a result, OEMs and System Builders can develop enterprise IP storage systems, scalable to hundreds of terabytes, in a very small form factor. More details regarding the form factor of the IP Storage Controllers will be available in the coming months.

Enterprise-grade High Availability

A critical limitation of today's IP storage systems is their lack of a high-availability, dual-controller hardware design. With Fibre Channel enterprise storage systems, the storage controllers often take the form of "blades", a single board that plugs directly into the backplane of the storage system. Some manufacturers, such as Engenio, use a separate "processor enclosure" that houses the storage controllers. In both scenarios, if a storage controller fails, the storage system does not experience any downtime. All storage traffic automatically fails over to the second controller, and neither servers nor their business applications are disrupted.

With IP Storage Controllers, OEMs and System Builders will be able to develop IP storage systems with a high-availability, dual controller architecture. If an IP Storage Controller fails, the system can seamlessly failover to a second IP Storage Controller. Users will also have the flexibility to swap out the failed IP Storage Controller while the storage system is still operating, maximizing uptime. With these attributes, IP storage systems powered by IP Storage Controllers will be able to offer RAS (Reliability, Availability, Serviceability) features comparable to Fibre Channel systems.

Lower System Cost

Since IP Storage Controllers were designed from the ground up to meet the needs of IP storage systems, they contain all of the necessary components, such as the management processor, protocol processor, RAID, storage connectivity, and software on a single board. Using off-the-shelf components to provide comparable functionality would require one motherboard, multiple processors and PCI cards, and a multitude of internal cables. As a result, IP Storage Controllers

cost a fraction of off-the-shelf alternatives. More information regarding the cost advantages of the IP Storage Controllers is available upon request.

Massive, Flexible Scalability

IP Storage Controllers will leverage Serial-attached SCSI technology to deliver a backend network of hard drives with very high bandwidth, ultra-low latency, and exceptional expandability. OEMs will be able to build systems with over 150 TB of capacity using IP Storage Controllers. In addition, IP Storage Controllers provide the flexibility of supporting both Primary and Nearline storage on one system with simultaneous Serial ATA and SAS hard drive support. More information is available upon request.

Easy-to-use, Multi-protocol Storage Software

At the core of the upcoming IP Storage Controllers is powerful software called HALO. Developed by Nimbus Data Systems, HALO combines several technologies in one modular software stack:

- Multi-protocol support of iSCSI, NFS, and CIFS
- RAID 0, 1, 5, 6, 10 and 50
- Storage virtualization and dynamic volume management
- Volume access controls
- Point-in-time, mountable snapshots and snapclones with rollback
- Asynchronous replication and synchronous mirroring
- System management tools
- Capacity scalability beyond 150 TB
- Intuitive web interface (HALO Manager)

HALO is designed to be embedded on the IP Storage Controllers as flashable firmware. In addition to providing iSCSI target functionality, HALO provides RAID 0, 1, 5, 6, 10 and 50, dynamic volume management, and full system monitoring. Embedded virtualization technology enables storage capacity to be seamlessly pooled across multiple systems. Port aggregation, mirroring, failover, and replication are also integrated.

Using HALO Manager, users can deploy IP SAN and NAS in minutes. HALO Manager enables users to configure, partition, provision, and monitor storage systems from one intuitive and easy-to-use web interface. With HALO Manager, users can create partitions (or volumes) and assign them to specific servers based on stringent access privileges. Volumes can be reallocated, expanded, reduced, or deleted, working seamlessly with the server operating system to enable instant expansion of file systems without reformatting.

HALO Manager also features extensive monitoring and diagnostics to report on storage utilization, performance, and health status. Administrators are notified of events via email and/or SNMP based on severity and frequency. CHAP authentication provides the robust security to ensure that volumes are only accessible by authorized computers. Password-protected user accounts and SSL encryption further protect against unauthorized access.

HALO Manager features powerful yet easy-to-use data protection features, including snapshot and asynchronous replication. Using snapshot, users can make a point-in-time copy of data on a volume instantly without taking the storage system offline or performing a lengthy tape backup. In the event of data loss or corruption, the user can “roll-back” to the snapshot, allowing storage operations to immediately resume from that point-in-time. Asynchronous replication goes one step further, copying these snapshots to another IP storage system offsite up to several hundred miles away. In the event of a natural or other disaster at the primary site, the data will be securely stored offsite and available for rapid recovery.

Summary

As purpose-built hardware platforms with robust storage software, IP Storage Controllers offer numerous advantages over off-the-shelf alternatives:

- Up to 4x greater performance
- Up to 80% lower power consumption
- Up to 80% smaller form factor
- 99.999% high availability
- Lower cost

Combined with the enterprise-grade software features of Nimbus HALO, which includes iSCSI/CIFS/NFS support, RAID, volume management, snapshot, clustering, asynchronous replication, and web administration, IP Storage Controllers provide System Builders and OEMs with the best hardware/software platform for building next-generation IP storage systems.

More Information

OEMs and System Builders interested in learning more about IP Storage Controllers can visit our website at <http://www.nimbusdata.com>.

600 Townsend St., Ste 120e, San Francisco, CA 94103
www.nimbusdata.com toll-free (877) 6-NIMBUS international (415) 503-4371

