



Industry Trends and Technology Perspective Solutions Brief

IBRIX Fusion™ Scalable Storage for Unstructured Data

Enabling Scalable Storage Performance, Availability and Capacity for unstructured data and Web 2.0 applications

By Greg Schulz

Founder and Senior Analyst, the StorageIO Group



January 24th, 2008

Compliments of IBRIX

The growth of structured and unstructured data continues at an explosive rate in most environments resulting in a constantly expanding data footprint requiring data and storage management resources. Similarly, the relative ease of use of NFS and Windows CIFS file sharing based storage, also known as Network Attached Storage (NAS), has led to a proliferation of NAS and Windows file servers (similar to how the ease of use of personal computers (PCs) has resulted in desktop and server sprawl) and a resulting urgency to consolidate to maximize resource usage and IT management tasks. This Industry trends and perspectives solution brief looks at the capabilities of IBRIX Fusion™ data and as a solution to meet these and other challenges for structured and unstructured data.



Introduction

The popularity of rich media and Internet-based applications has resulted in explosive growth of unstructured file data requiring new and more scalable storage solutions. Applications including video pre- and post-production processing; animation rendering; on-demand video and audio; social networking websites and the ability to easily digitally create information from cell phones, PDAs and other sources have increased burdens from a storage performance and capacity standpoint. Unstructured data includes spreadsheets, PowerPoint™, slide decks, PDF and word documents, web pages, video and audio JPEG, MP3 and MP4 files among others.

The diversity of rich media and Internet applications ranges from many small files with various access patterns to more traditional large video stream access. Consequently, storage systems, in order to scale with stability to support Internet and Web 2.0 applications, will need to support variable performance characteristics from small random access of meta-data or individual files to larger streaming video sequences. For some environments, the StorageIO Group finds data growth rates in the low double digits while we see other environments in high double to triple digit growth as the result of more data being generated, more copies of data being made and more data stored for longer periods of time.

Background and Issues

While structured data in the form of databases continues to grow, for most environments and applications the high growth area and expanding data footprint along with associated performance bottlenecks is centered on semi-structured email data and unstructured file data. Unstructured data has varying I/O characteristics that change over time. For example, starting out with a lot of activity, then going idle and then extensive reads as in the case of a video or audio file becoming known and popular on a media, entertainment, social networking or company sponsored website. Another example is when a development or research project is completed, the data or intellectual property (IP) is archived or migrated to lower cost, lower performance bulk storage until needed for subsequent research or sequel projects.

Faced with rising energy costs, limits on reliable electrical power, proliferation of file and storage servers, constraints on floor space, even with lower costs for higher capacity magnetic hard disk drives (HDD), simply throwing more hardware at the expanding data footprint of unstructured data is no longer a viable option for many environments. As with larger data center environments, Internet companies regardless of size typically do not have the luxury of switching out technology while leveraging existing hardware technologies; often working with hardware that might otherwise be considered obsolete and dispensable along with leveraging commodity hardware to stay within cost budgets.

Common challenges and activities for organizations faced with managing data growth include:

- Server and storage consolidation, including Network Attached Storage (NAS) file servers without introducing bottlenecks or excessive energy consumption while enabling modular growth
- Need for cost effective bulk storage and near-line for occasional access of unstructured data such as for media and entertainment or social networking websites or archival of project data
- Boost performance by eliminating bottlenecks to improve productivity or support more growth; for example, speed up sequential processing of call detail records and subsequent small random I/O database updates to expedite telecom billing and other similar transaction processing applications
- Support on-demand growth while improving energy efficiency and reducing IT management costs



An expanding data footprint means more data storage is needed with different tiers or levels of service characteristics to meet different application service level objectives. Similar to how server clustering or computer grids are used across a wide spectrum of applications and business environments to support diverse application processing needs, IT data centers can now effectively manage the growth in structured and unstructured data in a cost-effective way without compromising application service levels.

Applications and situations that benefit from ultra-scalable and flexible clustered file server storage solutions enabling dynamic data movement across different tiers of existing or third party hardware with support for both small random I/O operations and large sequential stream data access include:

- Unstructured data files including spread sheets, PDFs, slide decks and other common documents
- Email systems including supporting Microsoft Exchange personal (.PST) files stored on file servers
- Users' home directories and storage for commonly used unstructured files and data objects
- Web-based managed service providers for on-line data storage and data protection or backup
- Rich media content data delivery and hosting and social networking Internet sites
- Media and entertainment content creation including animation rendering and post-processing
- High performance databases such as Oracle with Network File System (NFS) direct I/O (DIO)
- Financial services and telecommunications, transportation, logistics and manufacturing
- Project-oriented software and technology development, simulation and energy exploration
- Low cost high performance cache for transient or temporary along with look-up or reference data
- Enhance real-time performance applications including fraud detection and electronic surveillance
- Life sciences, chemical research, software development and computer aided design (CAD)

The bandwidth intensive and parallel access performance characteristics associated with clustered storage are generally known; what is not so commonly known is the break-through to support small and random I/O operations per second (IOPS) associated with database, email, general purpose file serving, home directories as well as meta-data lookup. Clustered storage is no longer exclusive to the confines of high performance compute (HPC) sequential and parallel access scientific workloads or ultra large environments. Small sized files and I/O (read or write), including meta-data information, are also being supported by a new generation of multi-purpose, flexible clustered storage solution that can be tailored to support different applications workloads and characteristics similarities.

Meet IBRIX Fusion™

IBRIX Fusion™ is an example of a new breed of modular clustered file and data storage servers that support variable performance requirements of both unstructured data, normally placed on NAS or Windows file servers, along with structured high performance transaction processing databases, traditionally hosted on large expensive monolithic enterprise systems.

The Many Faces of Clustered Storage Myths and Realities

Myth: Clustered storage is only for HPC

Reality: Some vendors only support HPC

Myth: Only for parallel or sequential access

Reality: Some solutions also support OLTP

Myth: Only for ultra large environments

Reality: Simplified growth for small sites

Myth: It's only for performance applications

Reality: Effective for NAS consolidation



Founded in 2000 and privately held, Boston, MA-based IBRIX (www.ibrix.com) has developed the IBRIX Fusion™ file serving solution enabling scalability and modular growth using industry standard hardware components, including servers and storage, eliminating performance bottlenecks and management complexities. Unlike other clustered storage systems, whose capabilities are tied to supporting specific application workloads and/or hardware platforms, IBRIX is a different type of open and heterogeneous clustered file server that supports scaling in terms of variable performance for both small random and large sequential concurrent or parallel access of data. IBRIX Fusion™ file server gateways, unlike many early generations of clustered file servers, is not tied to any particular vendor's servers, networks or storage devices as evidenced by relationships with leading Tier-1 vendors.

For example, IBRIX customers can deploy the IBRIX Fusion™ file serving solution software on any industry standard x86 servers, transforming them into powerful and intelligent file servers using either dedicated direct attached storage (DAS), external SAS or Fibre Channel and iSCSI storage area networks (SAN) storage devices or any combination. Another example of the IBRIX Fusion™ flexibility is the ability to configure different file server gateway nodes for different tiers of access and performance using variable memory sizes, I/O adapters and compute power with different tiers and types of disk storage to address common management issues associated with data growth and increased data retention needs.

What differentiates IBRIX from other NAS and Clustered Storage Solutions?

IBRIX Fusion is a software-based solution comprised of a scalable integrated volume manager, high availability and data integrity features, including transparent automatic failover, along with a patented parallel segmented¹ file system designed for the ultra scaling needs of clustered storage with the flexibility to support NFS, CIFS and optional parallel access by application servers.

For applications that leverage shared and clustered NAS file systems for caching and concurrent access by many users, access to metadata information can become a severe bottleneck or single point of failure even as additional clustered storage nodes are accessed. For example, by supporting a truly distributed architecture, IBRIX Fusion™ enables metadata access along with general file access from multiple nodes concurrently, eliminating a common point of congestion found in early generation clustered file servers. Consequently, scalable and flexible clustered file serving solutions need to scale with stability. Scaling with stability means that as additional IBRIX Fusion™ file server gateway nodes are added, application performance does not degrade nor does application availability or data integrity suffer even when faced with accessing metadata across different nodes concurrently.

Scalability with Stability
✓ Boost performance without downtime
✓ Increase capacity without complexity
✓ Performance not impacted by growth
✓ Enterprise class data protection features
✓ Support large and small I/O workloads
✓ No single points of failure or congestion
✓ Transparent storage load balancing

IBRIX Fusion™-enabled scalable and flexible clustered NAS storage

Clustered storage solutions, like those enabled by IBRIX Fusion™, go beyond meeting the basic requirements of supporting large sequential parallel or concurrent file access along with support for small random IOPS typical of on-line environments. Scalable and flexible clustered file servers that leverage commonly deployed servers; networking and storage technologies are well suited for new and emerging applications, enabling growth without complexity or the downtime of traditional storage solutions.

¹ Refer to the IBRIX document “IBRIX File System: Scaling Bandwidth and Throughput in File Systems”

IBRIX Fusion™ File gateway nodes, front-end access and back-end storage can be tailored using varying amounts of processing, memory, I/O bandwidth and different types and tiers of disk drives to meet specific access and application performance requirements. For example, if a particular file or volume becomes performance bound, additional memory can be added to an IBRIX Fusion™ file access gateway, more access gateways added, faster disk drives installed or varying combinations to meet particular needs. For environments that are already undergoing rapid growth, costly downtime and disruptions can be reduced using the IBRIX Fusion™ Rebalancing feature; this facilitates transparent load-balancing and redistribution of data across new storage to improve utilization as well as performance.

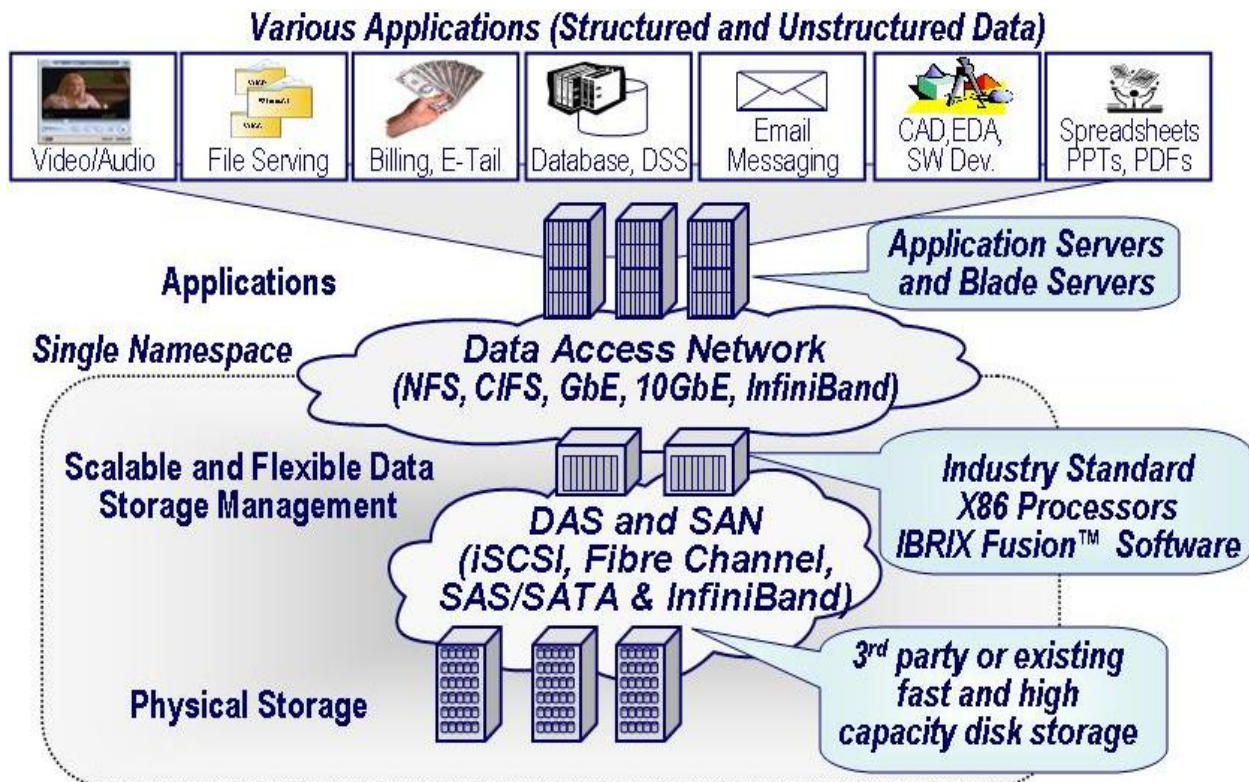


Figure-1: IBRIX Fusion™ Scalable and Flexible Clustered NAS and Storage Management

IBRIX Fusion™ components (Figure-1) deployed on industry standard x86 based processes include IBRIX Fusion™ patented segmented file system™ with distributed metadata and multi-petabyte name space to support ultra-scalability using standard NFS and CIFS access or optional parallel access from client application servers. FileRAID provides file level data protection with Fusion Snap enabling on-line backups and rapid file recovery in the event of file cooption or accidental deletion.

To eliminate bottlenecks, separate I/O paths are used in the IBRIX Fusion™ architecture including those for data access networks (for application server access), metadata (communication between IBRIX Fusion™ file server gateway nodes) and physical storage access networks using DAS or SAN or various combinations. Network interfaces include Gb Ethernet (GbE), 10GbE and InfiniBand while data access by application servers include GbE and 10GbE LAN, InfiniBand, TCP/IP, NFS and CIFS. The flexibility of IBRIX Fusion™ enables various topologies and combinations of I/O or networking interconnects to meet different application performance and service level requirements.



StorageIO Group Perspective

More nodes, ports, memory and disks do not guarantee more performance for applications; it depends on how those resources are deployed and how the storage management software enables those resources to avoid bottlenecks. For some clustered NAS and storage systems, more nodes are required to compensate for overhead or performance congestion of processing diverse application workload and performance characteristics. Other items and features to look for and consider include support for industry standard interfaces, protocols and technologies.

Scalable & Flexible NAS Storage Clusters

- ✓ Lower storage costs via modular growth
- ✓ Support server & storage consolidation
- ✓ Leverage existing or 3rd party hardware
- ✓ Enable application tuned performance
- ✓ Optimize energy effectiveness
- ✓ Inherent availability and data integrity
- ✓ Simplify data and storage management

Additional items to consider for clustered NAS storage solutions include:

- ✓ Do the data integrity features align to meet or exceed disaster recovery and data protection needs?
- ✓ Is there support for large file systems supporting many small files along with support for large files?
- ✓ What is performance for read/write of small random IOPS and bandwidth for large sequential I/O?
- ✓ How is performance enabled across different application workloads in the same cluster instance?
- ✓ Are I/O requests, including metadata lookup, funneled through a single node resulting in bottlenecks?
- ✓ How disruptive to applications and time-consuming is adding new or replacing existing storage?
- ✓ Is proprietary hardware required or can you tailor preferred resources of to application needs?
- ✓ What features exist to reduce storage management costs and complexity such as data load balancing?

Conclusion and Summary

Moving forward, The StorageIO Group sees vendors of scalable and flexible clustered file server storage like IBRIX having the potential of leveraging the inherent processing capabilities of constantly improving underlying hardware platforms to deliver new and robust value-added features for data and storage management. StorageIO sees and believes that clustered storage- including clustered NAS file serving for concurrent access of data using NFS and CIFS by application servers or optional parallel access- is no longer exclusive to niche applications or scientific and HPC environments. Organizations of all sizes can benefit from ultra-scalable flexible clustered NAS storage that supports different application performance needs form small random I/O, metadata lookup and large stream sequential I/O with scales with stability to grow with your business and application needs.

About the author

Greg Schulz is founder and senior analyst of the StorageIO Group and author of the books *The Green and Virtual Data Center (CRC)* and *Resilient Storage Networks — Designing Flexible Scalable Data Infrastructures* (Elsevier Digital Press).

All trademarks are the property of their respective companies and owners. The StorageIO Group makes no expressed or implied warranties in this document relating to the use or operation of the products and techniques described herein. The StorageIO Group in no event shall be liable for any indirect, consequential, special, incidental or other damages arising out of or associated with any aspect of this document, its use, reliance upon the information, recommendations, or inadvertent errors contained herein. Information, opinions and recommendations made by the StorageIO Group are based upon public information believed to be accurate, reliable, and subject to change. This Industry Trends and Perspective Solution Brief is compliments of IBRIX.