



Breadth and Simplicity – NuView’s Solution Set For File Storage Virtualization

Analyst: Michael Fisch

Management Summary

The scope of what file storage virtualization can do for your enterprise may surprise you. When an enterprise reaches a certain point of scale, the number of file servers, NAS platforms, user desktops, and application servers and their interrelationships for file access become an intricate, unruly web. Management costs are high and utilization rates are low. File virtualization addresses this problem with a *global namespace*. This technology unifies and simplifies fragmented environments, delivering benefits in proportion to the environment’s size and complexity. This alone makes file virtualization a worthy investment, though it also does more.

It facilitates transparent data movement. This important activity is at the heart of storage and server consolidation, data distribution, data retention for regulatory compliance, centralized backup and restore, data protection, and business continuity. Furthermore, file virtualization facilitates the collection of enterprise-wide metadata for file classification. Both data movement and file classification are necessary components of information lifecycle management (ILM), which intelligently places files in different price-performance tiers of storage in order to balance service levels and cost. In short, the impact of file virtualization is broad and deep.

During the last five years, NuView has focused on delivering the promise of file storage virtualization by developing and enhancing a set of software solutions:

- **StorageX** – NuView’s main product, which includes a global namespace for heterogeneous *Windows* and *UNIX*, policy-based automation, data migration, differential replication, reporting, and automatic failover for disaster recovery.
- **File Lifecycle Manager** – A specific solution for automated file migration and ILM. It manages and migrates individual files through a tiered-storage infrastructure, from creation to deletion.
- **UNC Update** – A tool for revising UNC references in existing documents, which is useful when switching to a global namespace.

Enterprises face numerous file storage challenges, from high costs to rising service level expectations. NuView brings a solution set to the table with breadth and simplicity, in that it does not require a major infrastructure upgrade or reconfiguration. It just layers on top of an existing environment. Read on for details.

IN THIS ISSUE

➤ File Storage Challenges	2
➤ StorageX	2
➤ File Lifecycle Manager	4
➤ UNC Update	4
➤ Conclusion	4

File Storage Challenges

Enterprises are experiencing formidable challenges with their file storage infrastructure:

- **Data growth** – Continuous data growth (including unstructured data or files) is the backdrop on the stage of enterprise storage. All enterprises will need more capacity next year than they do today – the only question is how much more. Therefore, all must plan to add capacity periodically.
- **Wasted storage capacity** – As enterprises try to keep up with this data growth treadmill, most end up carrying too much capacity. Traditional NAS platforms and file servers scale in discrete units, so enterprises with requirements of any scale have multiple platforms. Each NAS platform likely contains multiple, discrete file systems (i.e., logical data containers) due to size limitations. Such a fragmented environment has poor utilization because capacity is not easily transferred between the storage “islands.” To make matters worse, enterprises tend to purchase capacity in large chunks due to the effort and disruption associated with each addition. This excess capacity ties up capital, consumes electricity, space, and administrative resources, and does not help further business objectives.
- **Labor-intensive management** – Storage hardware is not the only wasted resource in a fragmented environment. IT administrators also spend an overabundance of time and effort managing the storage islands: tracking resources, making changes, manual load balancing, upgrades, and consolidation. For each change or addition, administrators must update the share mappings or mount points on all affected clients. This is like completely rewiring a house every time a lamp is moved or new light bulb installed. There must be a better way. These excess IT staff hours have a major budget impact, over time.
- **File access disruption** – Another problem is downtime. During these changes and additions, clients may be denied access to data for periods of time that affect their productivity or, at minimum, are an inconvenience.
- **Over-reliance on top-tier storage** – There is an opportunity to lower file storage costs

yet more by implementing information life-cycle management (ILM)¹. Enterprises can classify files according to their service level requirements and migrate many (if not most) files from top-tier storage to low-cost media like SATA drives and tape. By streamlining the top tier, its performance also improves and activities like data replication, backup, and restore are faster.

- **Data retention and regulatory compliance** – Enterprises need to retain data for extended periods to meet regulatory, legal, and operational requirements.
- **Business continuity** – Enterprises also need more effective and attainable data protection and recovery solutions² to ensure business continuity.

File storage virtualization is a rising technology area that addresses these challenges through simplification and powerful management capabilities. The software vendor NuView has developed and enhanced a broad set of solutions in this area.

StorageX

NuView’s first and principal solution for file storage virtualization is *StorageX*. It inserts a software management layer on top of existing file servers and NAS platforms. In this way, it can unify, simplify, streamline, and fortify a distributed, heterogeneous file infrastructure. StorageX is a targeted solution, meaning that it does not involve a major IT renovation. No client installation is required. An enterprise can install it and begin to experience benefits within a day.

The foundational feature of StorageX is a *global namespace*, which operates in a manner similar to the World Wide Web. When surfing the Web, do you know where content comes from (physically)? Probably not. Web sites are hosted on servers, a multitude of which are located all around the world and connected to the Internet. You might download content

¹ See *Top 10 Things You Should Know About Information Lifecycle Management* in **The Clipper Group Explorer** dated May 10, 2004, at <http://www.clipper.com/research/TCG2041R.pdf>.

² See *Recovery in Perspective – Ensuring Access to Enterprise Data* in **The Clipper Group Explorer** dated January 28, 2005, at <http://www.clipper.com/research/TCG2005003.pdf>.

from many different states and countries in a single surfing session. The physical complexity of the practically infinite Internet is hidden from you. In other words, it is *virtualized*. You just point and click. This is what makes the Web so simple and powerful!

In a similar fashion, a global namespace makes an enterprise's file storage infrastructure easier to access and administer. It unites the many physical components and presents them to end users³ as a single, consistent "pool of files." StorageX uses the de facto standard *Windows DFS* and *UNIX automount* capabilities to create an enterprise-wide namespace for all Windows and UNIX clients.⁴ On the storage side, it uses the standard CIFS and NFS protocols to incorporate heterogeneous file servers and NAS platforms and offer centralized management. As such, it employs an open architecture. It is also out-of-band because it does not require an appliance in the data path or a proprietary file system.

This global namespace eliminates the direct physical dependencies that cause downtime and administrative work whenever the environment changes. Therefore, tasks like data migration, storage and server consolidation, and capacity additions become transparent to users and easier for administrators. It is feasible to add capacity in more precise and frequent allocations, instead of over-provisioning. These benefits improve the economics of storage as well as information access.

StorageX builds upon the global namespace with policy and data movement engines, which open up a number of possibilities. The data movement process itself is minimally disruptive, operates at the share, directory, or qtree level, does not require "desktop touches," and may be scheduled for off hours.

For instance, StorageX can monitor capacity utilization and migrate data automatically between storage platforms for load balancing purposes. If one platform becomes too full, it can move data to an underutilized platform. This raises utilization overall and avoids performance bottlenecks.

Server and storage consolidation projects

are a normal part of the life of the enterprise. The key to successful implementation of these projects is making them transparent to the end user, which StorageX can do. Moreover, its capability for policy-based data migration makes the process simpler, easier, and more cost-effective.

Furthermore, policy-based data movement enables ILM. If an enterprise has storage tiers⁵, such as a high-performance primary platform and a low-cost secondary platform, StorageX can classify files based on policies and move them to appropriate tiers. It uses file attributes (e.g., size, age, type, and directory) to make classification decisions that reflect business and operational requirements. Again, all of this is transparent to users. This is another way to lower costs and to increase performance of the primary tier. It is also useful for data retention and regulatory compliance. StorageX supports the NetApp *SnapLock* feature that ensures data is non-erasable and non-rewritable.

For business continuance, StorageX has a feature called *Byte-Level File Differential Replication (BFDR)* that replicates files at periodic intervals, including over long distances. BFDR only replicates changes after establishing the initial copy. This feature speeds up replication and consumes less network bandwidth, thus cutting costs. BFDR serves as a relatively inexpensive option for disaster recovery, offering a recovery point objective in between real-time mirroring and backup. It is also useful for file distribution to remote sites, consolidation from remote sites, and centralized backup.

Business continuity is not only about data replication, but also about the time it takes for the end users to connect to the replicated data. In this area, StorageX can automatically detect the failure of primary storage and transparently failover clients within minutes to a secondary, remote site.

Finally, StorageX generates reports on the file infrastructure. It scans and summarizes stored data, in terms of numbers of files, types,

³ Client PCs, workstations, and application servers

⁴ It technically creates two namespaces, one for Unix clients and another for Windows clients.

⁵ For details, see *Tiered Storage Classes Save Money – Getting The Most Out Of Your Storage Infrastructure* in **The Clipper Group Explorer** dated August 29, 2002, at <http://www.clipper.com/research/TCG2002030.pdf>.

capacity consumed, where located, and who owns them. This information is useful for data classification, capacity planning, and charge-back to lines of business.

StorageX supports Windows and Windows storage servers, UNIX storage targets running NFS v3, as well as NetApp and EMC NAS platforms and gateways. StorageX is also well integrated with NetApp's data management software offerings, such as *SnapShot*, *Snap-Mirror*, and *SnapLock*. NetApp resells StorageX under the name *VFM*.

File Lifecycle Manager

FLM is a specific solution for automated file migration. It complements StorageX by enabling more sophisticated ILM for files, though it does not require StorageX to run. It enables enterprises to take advantage of an unlimited number of disk-based tiers of storage.

FLM classifies, migrates, and restores individual files based on policy. Like StorageX, it uses file attributes and user-defined policies to make classification and migration decisions. Instead of moving entire directories and then reflecting changes in the global namespace, *FLM* migrates individual files and leaves a stub or pointer on the primary storage platform. The filename and metadata are still there, but the content (and bulk of the data) are not. When a client accesses a migrated file, *FLM* automatically restores it to primary storage. This process is completely transparent to end users. Therefore, like StorageX, it does not require any modification to the infrastructure or behavior changes on the part of the users. The ILM benefits are fundamentally the same – lower storage costs, faster backup and replication of the primary, data retention and archiving – but the result is more precise and nuanced.

FLM also does more for managing the whole data lifecycle from creation to deletion. It can block unwanted files (e.g., MP3 – sorry music fans) from being stored in the first place. Migration options include many-to-one and one-to-many. *FLM* can also delete files based on policy, when the enterprise no longer needs them for operational or regulatory purposes. Finally, it offers file reporting, capacity planning, and what-if analysis tools.

UNC Update

UNC Update takes care of a particular challenge that arises when switching to a global namespace or during the time of storage/server consolidation – references within existing documents that point to files with their physical paths. The UNC⁶ references, OLE⁷ links, and shortcuts in these documents become invalid. Therefore, *UNC Update* changes these references to reflect the new layout of the global namespace, preserving the internal consistency of an enterprise's file infrastructure.

Conclusion

File storage virtualization solves many problems and delivers many benefits. The first question to settle is whether your enterprise needs a global namespace. If the file storage infrastructure is of any size and complexity, the answer is likely *yes*. This compelling technology cuts costs, eases management, and improves information access and availability. From there, other questions to answer include whether you need enterprise-wide file reporting, transparent data migration, some form of ILM, data retention and archiving, and possibly replication and disaster recovery. Then, choose the best solution for your requirements.

In this process, be sure to consider NuView's solution set. **StorageX, FLM, and UNC Update are effective, targeted solutions for file storage virtualization and data management.** They are not obtrusive, over-engineered, or high-risk, but merely ride on top of what is already there. The price is such that a back-of-the-envelope ROI calculation should yield a solid multiplier. **So, if your enterprise needs a solution to shave costs out of the file infrastructure and improve service levels, StorageX may be a smart investment.**



⁶ Universal Naming Convention

⁷ Object Linking and Embedding

About The Clipper Group, Inc.

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

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

About the Author

Michael Fisch is Director of Storage and Networking for The Clipper Group. He brings over nine years of experience in the computer industry working in sales, market analysis and positioning, and engineering. Mr. Fisch worked at EMC Corporation as a marketing program manager focused on service providers and as a competitive market analyst. Before that, he worked in international channel development, manufacturing, and technical support at Extended Systems, Inc. Mr. Fisch earned an MBA from Babson College and a Bachelor's degree in electrical engineering from the University of Idaho.

- ***Reach Michael Fisch via e-mail at mike.fisch@clipper.com or at 781-235-0085 Ext. 25. (Please dial "1-25" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

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

Disclosure

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

Regarding the Information in this Issue

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