

WHITE PAPER

The Economic Impact of File Virtualization: Reducing Storage Costs and Improving Management Efficiency

Sponsored by: Acopia Networks Inc.

Richard L. Villars
May 2007

IDC OPINION

Enterprises are adopting new strategies for building and managing file-based storage solutions in response to the high growth in file-based information and the critical need to more cost-effectively manage these information assets. In a series of interviews with U.S. and European enterprises, IDC found that companies are increasing their file-based storage by 50–120% a year and place a high priority on boosting the efficiency and reliability of their management processes for file-based information.

Network-based file virtualization is a key technology for more effectively managing file-based information. By using the virtualization and data management functions in Acopia Network's ARX products, interviewed companies, from a cross-section of industries, stated that they could:

- Rapidly and non-disruptively migrate data from file servers to more scalable NAS systems (cutting migration times by up to 90%)
- Automate movement of data to lower-cost tiers of storage based on preset policies (reducing spending on disk capacity by 50–80%)
- Reduce resources set aside for backup (cutting tape consumption and backup times by 50–80%)
- Share critical business files across applications, departments, and geographies (boosting end-user productivity and reducing development times)

BUSINESS PRIORITIES IN A SHIFTING WORLD

Companies rely on a growing set of applications to compete in today's changing business environment. They are:

- Consolidating the collection, storage, and analysis of information about products, customers, and transactions
- Relying on email, ecommerce systems, and Web sites to communicate and conduct business with customers and business partners
- Digitizing records, images, and other fixed content to deliver new services, boost efficiency, and comply with evolving government regulations

This expansion in the range of applications that companies are developing and deploying has a significant effect on how they organize and store information.

When Content Becomes King: Enterprises' Evolving Information Requirements

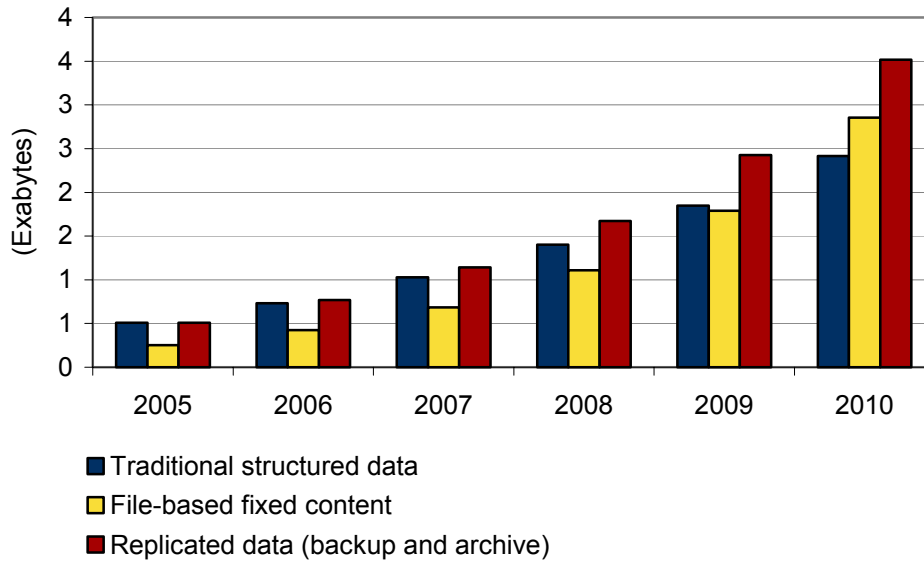
Historically, mission-critical, transaction-intensive applications built upon a block-based storage model consumed most enterprise storage assets. Over the past three years, however, the type of applications deployed by companies changed significantly. A new set of applications (e.g., Web servers, email, electronic records, and digital media) that rely upon a more file-centric storage infrastructure are accounting for an ever-greater portion of companies' corporate data.

Over the past six months, IDC spoke with a number of senior IT executives at major telecommunications and financial services companies where file-based storage grew from 5% of their installed storage capacity to over 25%, at a time when overall capacity was increasing anywhere from 50% to 70% annually. When asked to identify which of these application types would be the biggest drivers of storage capacity expansion in the coming years, file-based applications such as email, digital content, and long-term archives were the most widely cited.

To better gauge the pace of this transition, IDC developed an enterprise disk storage consumption model that tracks deployment of new storage capacity (See Figure 1). Across all regions, enterprises will consume as much disk storage for file-based content as for traditional block-based structured data in 2009. Add in backup, and archiving and managing file-based information throughout the data's life cycle becomes critical.

FIGURE 1

The Content Explosion



Source: IDC Enterprise Disk Storage Consumption Model, 2007

The Economic Perils of Managing File-Based Information

The rapid growth in file-based information is exposing the weaknesses of many existing IT management practices for file-based storage. Many companies continue to deploy large numbers of dispersed, underutilized file servers in support of different workgroups and geographically remote sites. Others have consolidated file-based storage on larger NAS systems but are still struggling to deal with rapid data growth and rapid technology changes.

The explosion in file-based information has a direct economic impact on enterprises in a number of ways.

- ☒ The use of dispersed file servers leads to severe underutilization of storage capacity, often at levels below 15%. In addition, enterprises require an increasing allocation of resources to make unneeded copies of older, unchanging files, often 70–80% of all files being backed up. Rapidly growing and dispersed file system environments also create management complexity and inhibit employee collaboration by preventing the sharing of information across groups or geographies.
- ☒ Efforts to migrate and consolidate data from even a limited number of file servers onto larger, more centralized NAS systems can often take six months to a year. In addition, ongoing integration and tuning of independent file structures as data grows and organizations change lead to severe and ongoing disruptions, driving down employee productivity and increasing the burden on IT staff.

- ☒ The long-term retention of files creates a growing performance and cost mismatch in terms of NAS systems deployed and backup policies followed. IT executives now acknowledge that they can't limit long-term retention of seldom-used files (due to user resistance or regulatory mandates).

Companies need to manage files throughout their life cycle, from creation to archive to possible destruction. The inability to automatically, but non-disruptively, migrate older data to lower-cost storage leads to excess spending on disk systems that may cost three to four times as much to acquire and administer.

In addition to these economic disruptions, the inability to separate different types of older, unaltered files forces companies to continue expending excessive resources on tape libraries and tape media used to back up the same unchanging data on a weekly or even daily basis. This reduces companies' ability to meet or extend their business continuity goals.

The Next Generation of File-Based Storage

Given the high growth in file-based information and the critical need to more reliably and cost effectively manage these information assets, enterprises are adopting new strategies for building and managing file-based storage solutions. These next-generation solutions must meet the following requirements:

- ☒ Allow enterprises to deploy and consistently manage across a wide range of disk storage tiers with different performance, capacity, availability, and cost characteristics, without requiring wholesale displacement of existing assets
- ☒ Provide a common set of scalable and highly available global namespace, data migration, data life-cycle management, and data protection services across these multiple storage tiers

The remainder of this white paper will look at one of the key technologies, network-based file virtualization, that enterprises are already using to meet these requirements and more effectively manage their file-based information assets. It will highlight comments from a number of enterprises that are using file virtualization solutions from Acopia Networks to address three key business requirements:

- ☒ Accelerating the migration and consolidation of dispersed corporate information onto more manageable and centralized NAS systems while minimizing business disruptions and reducing the IT management time needed for planning and reconfiguration
- ☒ Reducing the cost of doing business while boosting business integrity by making it easy to automatically and non-disruptively move information between tiers of storage based on preset policies and reducing the time and resources required to back up information
- ☒ Enabling more effective use of business information through consolidation and simplified access across applications, departments, and geographic boundaries

ENSURING BUSINESS GROWTH WITH FILE VIRTUALIZATION

Virtualization is currently a widely discussed trend throughout the IT industry. For old datacenter managers it means portioning resources in mainframe and Unix environments. For many IT managers, today, it means consolidating many x86-based applications onto a single server.

Virtualization is an equally important trend in the storage market; however, the focus is almost exactly the opposite. Rather than consolidating multiple servers onto one piece of hardware, network-based storage virtualization is most often about allowing storage managers to manage multiple storage systems (e.g., SAN-attached arrays or NAS systems) as a common pool of capacity. Such systems allow IT managers to more easily and non-disruptively migrate information between storage, boost the utilization of installed capacity, implement tiered storage policies to manage costs, and scale data protection processes.

Acopia's Intelligent File Virtualization

Acopia Networks Inc., with its ARX Series of intelligent file virtualization devices, is a leading provider of network-based file virtualization. The key features of the ARX Series include:

- ☒ Virtualizing all of the file-based storage infrastructure that is already in place, with a common global namespace irrespective of file system, platform, vendor, or protocol (NFS, CIFS, multiprotocol)
- ☒ Accelerating information migration and consolidation for dispersed file servers or legacy NAS systems through the use of data migration services
- ☒ Enabling adoption of tiered storage architectures that reduce spending on NAS systems and backup systems through use of intelligent and automated data lifecycle management and replication services
- ☒ Boosting the aggregate performance and utilization of virtualized devices through intelligent load balancing and enhancing the availability of all information on a growing range of virtualized storage systems

Acopia has a wide range of customers in industries, including finance, manufacturing, technology, media and entertainment, telecommunications, life sciences, energy, education and government. Its systems are deployed in the United States, Europe, and Asia.

As part of IDC's efforts to assess the economic impact of the explosion in file-based information on enterprises' IT environments, we conducted a series of interviews with a number of Acopia's current customers (see Table 1). The amount of file-based capacity installed at these companies ranged from 5TB to over 600TB, and all were experiencing growth rates in their file-based storage of 40–120% a year. They also placed a very high priority on boosting the efficiency and reliability of their management processes for file-based information.

The goal of these interviews with senior IT executives was to gain qualitative and quantitative information about how the deployment of file virtualization helped companies better meet business objectives and manage growth in IT expenditures.

TABLE 1

Adoption of File Virtualization: Company Profiles and Interview Results

Company Description	Current Environment	Challenges	File Virtualization Use Case
US-based resort and media company	<ul style="list-style-type: none"> • 3 NAS systems (6TB) • 1 NAS system for backup/archive (40TB) • Capacity demands increasing >50% a year 	<ul style="list-style-type: none"> • Support rapid growth in digital media archives • Reduce cost and improve reliability of data protection • Reduce administration of user namespace 	Using ARX for >2 years <ul style="list-style-type: none"> • Dynamically tier older data sets and media archives on lower-cost storage • Reduce IT staff overhead with global file sharing
Global financial services company	<ul style="list-style-type: none"> • Hundreds of file servers in offices around the world • 30 NAS systems in regional datacenters • 600TB of useable file-based capacity growing 100% a year 	<ul style="list-style-type: none"> • Consolidate data from dispersed file servers onto centralized NAS • Enable long-term retention of files on disk vs. tape • Coordinate application development between U.S., Europe, and Asia teams 	Using ARX for >2 years <ul style="list-style-type: none"> • Reduce time, cost, and disruptions of migrating from file servers to NAS • Provide global mount points for data sets in United States, Europe, and Asia
Global financial services company	<ul style="list-style-type: none"> • Multiple NAS systems with both high performance and high capacity (SATA disks) • 150TB of file-based data growing 100% a year 	<ul style="list-style-type: none"> • Excessive time and cost associated with migration of files from legacy NAS system onto NAS systems from another vendor • Reduce cost of storing and protecting older files and Exchange PSTs 	Using ARX for >1 year <ul style="list-style-type: none"> • Enable resumption of planned replacement of legacy NAS system through use of data migration services • Implement high-capacity and low-cost storage tier for older files and PSTs
Global media company	<ul style="list-style-type: none"> • Multiple NAS systems with both high performance and high capacity (SATA disks) • 150TB of file-based data and currently growing 1TB per week 	<ul style="list-style-type: none"> • Continually rebalancing data and namespaces on multiple NAS systems as new files are added to media distribution system • Reducing complexity for business partners linking to fast growing media library 	Using ARX for >2 years <ul style="list-style-type: none"> • Accelerate migration from legacy NAS systems through use of data migration services • Automate rebalancing of files on NAS systems • Eliminate disruptions in business partner applications with global namespace
European-based international design company	<ul style="list-style-type: none"> • Multiple NAS systems with both high performance and high capacity (SATA disks) • 30TB of file-based data growing 120% per year 	<ul style="list-style-type: none"> • Managing escalating storage costs as employees create, use, and archive more image-intensive files • Continually rebalancing data and namespaces on multiple NAS systems as projects end and teams move onto different projects 	Using ARX for >2 years <ul style="list-style-type: none"> • Dynamically tier storage of archived project files on lower cost storage • Boost access times for shared reference data • Automate rebalancing of files on NAS systems
European operations for global trading company	<ul style="list-style-type: none"> • 80 files servers and small NAS systems spread across 50 offices in multiple countries • 5TB of file-based data growing 30% per year 	<ul style="list-style-type: none"> • Assuring protection of information currently on widely dispersed servers at sites with little or no IT staff • Managing escalating costs associated with information recovery for growing number of ediscovery and audit requests 	Using ARX for <1 year <ul style="list-style-type: none"> • Non-disruptively migrate to consolidated NAS systems • Automate disk-to-disk back up for stored files • Enable tiered storage to reduce impact of ediscovery and audit requests

Source: IDC, 2007

Interview Findings

For all of the interviewed companies, finding ways to more efficiently and better manage the explosion in file-based information were high priorities. The three overarching themes (see Table 2) related to use of Acopia's file virtualization solution include:

- ☒ Enabling rapid, non-disruptive migrations of data from widely dispersed, hard-to-manage, and expensive-to-expand departmental file services to more scalable and cost-effective NAS systems (cutting migration times by up to 90% and eliminating staff time needed for migration planning and implementation)
- ☒ Enabling automated and non-disruptive movement of data based on preset policies as part of a dynamic tiered storage implementation to take advantage of low-cost disk storage and reduce resources set aside for backup (reducing spending on disk capacity by 50–80% and cutting tape consumption and backup times by 50–80%)
- ☒ Enabling the sharing of critical business files across departments, applications and geographies while eliminating disruptions due to capacity management, data administration and ineffective version control (boosting end-user and application productivity, and reducing development times through more effective sharing of reference data and more timely collaboration across geographies)

Interview Methodology

In the process of researching this white paper, IDC conducted interviews with six Acopia customers. The interviews were conducted to help IDC develop and articulate the file-based management rationale and represent the business benefits of Acopia's approach and solution. The interviewed companies were chosen by Acopia.

TABLE 2

The Business Value of Network-Based File Virtualization

Use Case	Key Findings	Interview Highlights
<p>File migration and consolidation: Consolidating dispersed file servers or legacy NAS systems onto a smaller number of NAS systems to boost utilization, enable greater data sharing, and improve backup/data retention processes</p>	<p>Challenges: Migration times for even a limited number of file servers can take six months to a year to plan and complete, consuming large amounts of IT management resources and planned downtime hours while severely disrupting end user productivity.</p> <p>Impact of file virtualization: Companies were able to reduce migration times (including both planning and downtime) to less than one month while dramatically reducing disruptions for end users.</p>	<p>"In one European office, it took almost a year to migrate 10 file servers, including the dedication of one FTE for the entire time. More recently, we did a similar migration [in the United States] and it took us a month, including the planning and the coordination. The actual migration [with Acopia] took us two weekends." (Global financial services company)</p> <p>"We were adding one terabyte of capacity a week for rich content and had a dedicated person doing the juggling of files and file directories across multiple NAS systems. Partners had to stop using the system when we made changes. With Acopia's ARX, we automated the redistribution of files and migrated to next-generation NAS systems without disrupting our partners' applications." (Global media company)</p>
<p>Dynamic tiered storage: Deploy a solution that allows the movement of old, infrequently accessed data, and back up copies of data to NAS systems optimized for capacity and low cost</p>	<p>Challenges: End users refuse to eliminate old files from networked drives and email systems and refuse to manually move data to separate archival drives, so in many case 60–80% of stored data is infrequently accessed. Backing up infrequently used data increases backup times and leads to massive duplication of data on tapes.</p> <p>Impact of file virtualization: Companies were able to reduce spending of disk storage hardware by 40– 60%, improve recovery times for files and email from days to hours, and reduce the time and costs of backup to tape by 50% to 80%.</p>	<p>"Every week, the help desk would get panic calls saying I can't save my budget spread sheet. What is going on? It turned out that less than a 30% of the files were used within the last year. We used Acopia's file virtualization and automated data migration facilities to transparently move all the older files to NAS with SATA storage (75% less expensive) and reduced backup for those older files from once a week to once a month." (U.S.-based Resort and Media Company)</p> <p>"We have a large number of unchanging reference files as well as a large archive of older project documents that we need to archive. Managing the reallocation of this fixed data across file servers and NAS systems consumed 100% of the time of one of my three storage administrators. We used Acopia's ARX to filter files away from our primary storage and onto our second tier (80% less expensive) with absolutely no impact on end users access to their data. We also limited 'backup to tape' to recently changed and modified data." (European-based international design company)</p>
<p>Global file sharing: Enable reliable, easy to manage, and secure sharing of files across diverse departments and geographies to boost collaboration and minimize version control issues</p>	<p>Challenges: Each department or office has its own file namespace so sharing either requires opening up access to all potential users (posing major security risks) or the continual replication of files over long distances.</p> <p>Impact of file virtualization: Companies improved productivity by allowing employees in multiple workgroups and remote offices to collaborate more effectively through common access to share files.</p>	<p>"One of our business groups is rapidly expanding its presence and development efforts. We could not meet our development plans without delivering a global view of the application development environment across four locations: London, Hong Kong, Tokyo, and New York. We were looking at Acopia for a file server consolidation project and determined that its networked-based approach would be the least intrusive while helping cut application development times by 50%." (Global financial services company)</p>

Source: IDC, 2007

Challenges / Opportunities

While Acopia currently offers a broad portfolio of data management and replication capabilities that complement its global namespace and file virtualization functions, the company must continue to expand the range of heterogeneous storage and application sets that it supports. In particular, Acopia must continue to enhance the ability of its solutions in the area of automated data life-cycle management and integrated access control for security. This effort includes expanding partnerships with content management, archiving, and data classification solution providers.

Finally, Acopia must further educate its customers on the benefits of setting up an environment that takes advantage of the automated data migration and data lifecycle management capabilities for improved disaster recovery and compliance. This coordinated storage architecture will help meet the rapidly expanding and increasingly demanding file-based storage requirements.

FINAL THOUGHTS: ENSURING ECONOMIC GAINS WITH EFFECTIVE IMPLEMENTATION

As comments from the interviews with Acopia's customers make clear, enterprises struggling to better manage their file-based information assets can achieve significant economic benefits from deploying a network-based file virtualization solution. They can:

- Accelerate migrations of data from file servers to more scalable and cost-effective NAS systems, while reducing disruptions for end-users
- Automate movement of data to lower-cost storage tiers, lowering spending on disk capacity
- Reduce resources set aside for backup, by cutting backup times and tape consumption
- Reduce the operational overhead associated with manual data management, capacity allocation, and data protection processes
- Boost the sharing of critical business files across applications, departments, and geographies

The key to attaining these benefits, of course, is effective implementation of the solution, including the adjustment of related processes (e.g., backup processes). A well-thought-out implementation (e.g., file server migration) also makes it easier to expand use of the solution into new user bases and for additional services (e.g., tiered storage or data replication). Finally, a sound implementation makes it easier for an enterprise to react to changing business conditions and new application needs.

When evaluating solutions such as Acopia's ARX products, IT managers must place equal focus on the ability of the solution provider and its business partners to design and deploy an effective and adaptable solution.

The companies that IDC interviewed stated that the ARX solution was easy to deploy and extend as they activated additional capabilities for data consolidation, data protection, and dynamic data life-cycle management. Those that acquired the solution through an Acopia business partner also stated that the partner provided a support infrastructure as well as a set of complementary implementation and planning services that allowed them to meet current and future needs of the business.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2007 IDC. Reproduction without written permission is completely forbidden.