



## **Acopia ARX Takes File Virtualization Global**

Analyst: Michael Fisch

### **Management Summary**

No doubt your enterprise experiences its share of IT pains and challenges. Consider the following list:

- Administration of client mappings for network file systems
- Disruptive data and storage management tasks
- Poor storage utilization and access performance
- Moving data to/from remote sites
- Need to replicate data for resiliency and protection
- Fragmented, inefficient storage management
- Poor and/or unpredictable WAN performance
- Need to implement data (or information) lifecycle management (DLM)

If one or more of these strike a chord, you should have a look at the new Acopia *ARX* platform for file virtualization and data management. This solution actually addresses *all* of the problems, which is impressive considering there are products on the market for doing just one of them.

The ARX platform is an intelligent Gigabit Ethernet switch that resides (logically) between client servers/PCs and an enterprise's file servers and NAS platforms. It virtualizes heterogeneous file servers and presents them to clients as a single, global namespace. Clients mount the namespace and access it via standard CIFS and NFS protocols. The global namespace abstracts changes in the underlying physical storage, shares, and volumes, so clients do not need to be reconfigured. Furthermore, ARX can create an enterprise-wide file virtualization infrastructure by placing platforms in geographically-distributed sites and interconnecting them with WAN links. It can automatically move and replicate data among sites based on policy for purposes such as centralized backup, archiving, data distribution, and data protection. It can also migrate files among different storage tiers (again, based on policy) to perform DLM functions. And the list continues.

Business benefits of this solution include lower storage acquisition and operating costs, enhanced service levels, and broader, easier access to enterprise information. Read on for details about the ARX and the IT challenges it can address.

### **IN THIS ISSUE**

➤ <b>The Problems ARX Solves</b> .....	<b>2</b>
➤ <b>The Acopia ARX Platform</b> .....	<b>4</b>
➤ <b>Conclusion</b> .....	<b>4</b>

## The Problems ARX Solves

The Acopia *ARX* is a next-generation platform for enterprise-scale file virtualization. Its blend of capabilities is broad and unique, so it seems best to describe the *problems it solves* before detailing its *technical features*. This will provide some context and help you see where it might fit into an IT infrastructure. So, here is a list of the technical and business challenges that the ARX addresses:

### ***Administration of Client Mappings for Network File Systems***

Client mappings allow servers and PCs to read and write files on the network. Think of them as hard-wired connections that point to shares and volumes on file servers and NAS systems<sup>1</sup>. If storage never changed, mappings would be constant and never need to be updated. But this is like saying if grass did not grow, you would never need to mow it. The problem is that it *does*. Unstructured enterprise data, namely files, grow quickly. The only question is how fast, and administrators periodically need to add disk drives, file systems, and even file servers to keep up. They also move data around for load balancing and performance optimization purposes. Client mappings need to be updated when these tasks occur, and it is a time-consuming chore, especially when multiplied over the hundreds or thousands of clients.

The ARX solution to this problem is the equivalent of grass that never grows – a *global namespace*. It is a veil that masks the complexity of many, even heterogeneous file servers and presents them as an unchanging, monolithic pool. The result is that client mappings remain constant despite the data and storage turbulence underneath. It saves work and cost and makes storage more easily adaptable to changing business needs.

### ***Disruptive Data and Storage Management Tasks***

Adding capacity, removing obsolete file servers, and moving data between file servers are a normal part of managing a storage infrastructure. IT rationalization projects like server, storage, and data center consolidation also involve these activities. However, these tasks disrupt users and applications, causing data unavailability. Enterprises normally must choose between impacting user productivity or performing the tasks during off hours, which makes the jobs and lives of IT administrators more difficult (if this is an option in light of 7x24 operations of some businesses and tasks that are too lengthy for the off-hour window). However, ARX allows

moves, changes, and additions to occur behind the veil of the global namespace, so they are not disruptive.

### ***Poor Storage Utilization and Access Performance***

If you group a dozen glasses together and splash a gallon of water over them, each glass will probably have a different amount of water in it. Some may be full while others empty. A similar effect occurs when multiple file servers are storing and accumulating data. Uneven distribution of data lowers overall storage utilization and causes an enterprise to purchase more capacity than it really needs. If too many “hot” files reside on a single system, access performance can also suffer. ARX remedies these problems through non-disruptive data movement. It can automatically place files among multiple file servers to smooth out utilization. It monitors data access in real-time and can move data to a different platform or location to improve performance. Both functions help enterprises get more out of their storage assets.

### ***Moving Data to/from Remote Sites***

Enterprises with geographically-distributed locations are challenged to move data back and forth from the center (e.g., main office/data center) to the edge (e.g., remote and branch sites). There are a variety of reasons to do this:

- **Centralized backup and archiving** – Backup is necessary for data protection, and archiving is increasingly important for corporate governance and regulatory compliance. So, a distributed enterprise can do them the hard way or the easy way. Each site could perform these functions itself, but this involves a lot of redundant IT work and infrastructure and may be done poorly or haphazardly. A better way is to perform backup and archiving centrally by periodically moving new or updated files from the edge to the main data center. The result is greater consistency and economies of scale.
- **Data distribution** – Enterprises may need to periodically push content from the center to the edge. For instance, remote sites may need access to the latest marketing materials or software updates. Manual and ad hoc processes and physical CD or tape shipments are inefficient. FTP over the Internet may be insecure and slow. It is better to do it in an automated, secure fashion over a network.
- **Branch office reporting** – Enterprises may need to periodically collect reports on financials, sales, transactions, inventory, and so forth, from branch offices to be rolled up at headquarters into company reports.

<sup>1</sup> Both file servers and NAS appliances and gateways provide common file access and sharing over a network. While there are technical differences between them, this report will refer to them generically as *file servers*.

- **Improving local access performance** – When remote sites access files over a WAN, they must contend with delays and bandwidth limitations that can make it unacceptably slow. A solution is to keep file copies local, so users can access them over the LAN.

ARX can meet the challenge of data movement by interconnecting distributed platforms over a wide-area network (WAN). ARX platforms at each geographical location not only virtualize local file resources, but also connect to each other to form a dynamic, enterprise-wide infrastructure for file access and movement. Administrators can create policies for moving data between locations according to business and IT needs, and the ARX network executes them. The platforms can even monitor performance and usage in real-time and move data around to maintain acceptable access performance.

### ***Need to Replicate Data for Resiliency and Protection***

Related to data movement is replication, where data copies are kept in remote locations to protect against local loss, corruption, or disaster. Administrators can set up policies to periodically replicate files and directories, including incremental changes since the last copy. It gives enterprises a higher level of data protection and the accompanying peace of mind. Another purpose is to make point-in-time copies for non-disruptive backup operations, data warehouse loading, and testing, and development.

### ***Fragmented, Inefficient Storage Management***

The rule of thumb is that it is easier to manage one big thing than many smaller things. Many enterprises struggle with managing a plethora of file servers, whether collocated in one data center or distributed throughout an enterprise. ARX offers the holy grail of centralized management. Administrators can manage one global namespace and aspects of the underlying storage assets from one console. ARX provides leverage.

### ***Poor and/or Unpredictable WAN Performance***

Network congestion and failed connections can get in the way of data movement, replication, and remote access. However, interconnected ARX platforms are actually smart enough to compute alternative routing paths based on factors like throughput, latency, and packet loss. They help get the most out of the network and the data, where it needs to be in a timely fashion.

### ***Data Lifecycle Management (DLM)***

Enterprises increasingly want to exploit data lifecycle management (DLM)<sup>2</sup>. Similar to plants

and animals, data follows a cradle-to-grave path of existence where its value changes over time. The DLM process seeks to apply the optimal storage service level<sup>3</sup> to data at each point in its lifecycle. There can be multiple tiers of service – think in terms of “low, medium, and high” or “silver, gold, and platinum.” DLM strikes the optimal balance between meeting business requirements and minimizing storage costs. It ensures the right data has the right service level at the right time. Reasons to consider these solutions include business and regulatory requirements for data retention, improved performance of databases and file systems, lower storage acquisition costs, and speeding up data management operations like backup and replication.

ARX can enable a tiered storage environment<sup>4</sup> by:

- **Classifying data according to file attributes (e.g., modified date, directory), usage, and access performance.** It can apply user-defined policies to categorize the data based on these characteristics.
- **Virtualizing multiple storage tiers.** Since the ARX can include multiple, heterogeneous file servers in its global namespace, it can deliver multiple storage tiers to users and applications. For example, it might use high-end NAS platforms with Fibre Channel drives for high service levels and Windows-based NAS appliances with ATA drives for a lower and less-costly service level.
- **Moving files to the appropriate storage tier automatically and non-disruptively based on policies and classification.** As values change over time, ARX can dynamically respond and migrate data to ensure it resides in the most appropriate service-level tier.

In this way, ARX enables DLM for an enterprise-wide, virtualized file infrastructure.

### ***A Workhorse***

As you can see, the Acopia ARX is a workhouse that can simultaneously solve many IT and business challenges. If you require multiple capabilities listed above, and/or if you need them on a broad scale, then you are in the ARX’s sweet spot. The ARX can accomplish a lot in one fell swoop. But if the scope of requirements is limited, such as

---

**Group Explorer** dated May 11, 2004, available at <http://www.clipper.com/research/TCG2004041R.pdf>.

<sup>3</sup> Service levels describe the particular performance, availability, recoverability, and cost characteristics applied to a data set.

<sup>4</sup> 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>.

---

<sup>2</sup> See *Top 10 Things You Should Know About Information Lifecycle Management* in **The Clipper**

### Acopia ARX Series Specifications

	<i>ARX6000</i>	<i>ARX1000</i>
<b>Size</b>	13U modular chassis	2U fixed configuration
<b>Network connectivity</b>	24 x 1Gbps Ethernet	6 x 1Gbps Ethernet
<b>File system support</b>	NFS v2 & v3, CIFS	NFS v2 & v3, CIFS
<b>Data object capacity</b>	500 million	100 million
<b>Throughput</b>	10 Gbps read/write	3 Gbps read/write
<b>Simultaneous sessions</b>	20,000	5,000
<b>File placement rate</b>	60 million/hr	10 million/hr
<b>Average list price</b>	\$150,000	\$45,000

replicating data between two servers, the ARX could be more than you need. There are products on the market that can address any one of the above capabilities – the uniqueness of the ARX is that it can do all of them on a global-enterprise scale.

### The Acopia ARX Platform

The ARX platform is an intelligent, purpose-built switch. It subsumes multiple, heterogeneous file servers and NAS platforms, and then presents them to clients within a global namespace. It supports the standard NFS and CIFS protocols and the *Solaris*, *Linux*, and *Windows* operating systems on clients. It comes in a large and a small version: The ARX6000 is a modular chassis with up to 24 Gigabit Ethernet ports, support for 20,000 simultaneous sessions, and has been tested up to 500M objects. The ARX1000 is a fixed configuration with 6 Gigabit Ethernet ports, 5,000 simultaneous sessions, and tested to 100M objects<sup>5</sup>. See the sidebar above for technical specifications.

The platform is *in-band*, meaning that all I/Os pass through the ARX platform. This central position in the network confers several advantages:

- Encompasses conceivably all clients and file servers,
- No agents or changes on clients,
- Able to monitor and deliver intelligent data and network management in real-time, and
- Facilitates centralized management.

The network is the common ground between clients and file servers. Thus, it is a sensible place to put intelligence, especially for a solution of broad scope.<sup>6</sup> The downside is potentially being a single point of failure, but the ARX can be configured in clustered pairs for high availability.

The ARX ties together an existing storage

infrastructure and adds a layer of intelligence.<sup>7</sup> There is no need to sweep the floor of existing assets or to perform a major upgrade. The kind of incremental investment that ARX offers can be less disruptive and easier to justify financially to executives.

### Conclusion

While the Acopia ARX's set of capabilities is quite broad, its business benefits fall into four categories:

- **Lower storage acquisition costs** – Thanks to higher utilization and the use of tiered storage, enterprises can spend less on acquiring storage.
- **Lower operating costs** – Management is greatly simplified by a global namespace, centralized management, and policy-based automation. Simplicity saves time, and time is money. Operating costs are by far the largest component of storage total cost of ownership.
- **Enhanced service levels** – The capabilities of the ARX can improve performance, availability, recoverability, and cost.
- **Better information access** – The global namespace and wide-area connectivity of ARX help get the right information to the right people for getting the right work done. It breaks down information silos and improves accessibility. Even a file search is easier and more comprehensive. As a result, the business itself can be more productive and competitive.

If the ARX might address one or two (or more) of your enterprise's particular problems, take a closer look. You may find that it can solve those problems as well as provide advantages that you had not yet considered. **It is like getting many birds with one stone!**



<sup>5</sup> Both models could scale beyond the rated number of objects, but this has not been tested.

<sup>6</sup> See *Intelligent Storage Networking – Poised for Broad Adoption* in **The Clipper Group Explorer** dated October 16, 2003, available at <http://www.clipper.com/research/TCG2003054.pdf>.

<sup>7</sup> It also has an API (application programming interface) that can allow third-party software and policy managers to make inquiries and initiate tasks.

### ***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](http://www.clipper.com).***

### ***About the Author***

***Michael Fisch is Director of Storage and Networking for The Clipper Group.*** He brings over eight 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](mailto: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.