

## **Intelligent File Virtualization and the FAN Paradigm**

September, 2006



Early this year, we broadly defined the term “File Area Networking” (FAN) for the IT community. Our goal was to provide the industry with a common reference architecture to describe the increasingly complex world of file management. Like the evolution of block IO that led to the SAN, we now see similar evolutionary steps transpiring with file data. And so, the FAN reference architecture becomes a means for end users and vendors alike to deal with the many different aspects of file management challenges by utilizing a common nomenclature. In our initial FAN research, we identified the emerging category of file virtualization or Network File Management (NFM) technologies as one of the key enabling components of any future FAN architecture. Within that category, we identify Acopia Networks as one of the leading examples of how intelligent file virtualization can help IT teams develop a proactive, strategic approach to their most daunting file challenges. We are very pleased to see that Acopia not only fundamentally grasps the FAN philosophy, but their ARX offering actually embodies the core capabilities on which a FAN can be architected. As such, we believe that Acopia and the ARX product family are extremely well positioned to take advantage of the changes coming in the file management realm. IT decision makers should familiarize themselves with this technology. Why so? Because the age of advanced, intelligent file networking is now upon us, and Acopia is not shy about staking their claim on its future.

### **What is a File Area Network?**

F-A-N: Yet another 3 letters for the IT lexicon. Since the IT world is already awash in an alphabet soup of acronyms, why add yet another to the mix? Do we *really* need to think about a File Area Network, or, FAN? The answer: *Absolutely*. This one matters. The criticality of this concept to the daily IT realm will be dramatic. The reason is simple: In all of our recent research on the topic of file management, we continue to find that IT teams are overwhelmed and confused by an entire range of file-related technologies now hitting the market.

Case in point: In a recent Taneja Group survey of 225 IT decision-makers, (June, 2006), we found that 62% of respondents now view file data growth and file management as either their “top” or “one of the top” IT priorities. Further, 57% of our respondents said that they are actively evaluating new and emerging technologies to help assist with these file challenges. However, in our follow-up discussions with IT teams, it has become very obvious that many managers are not comparing technological “apples-to-apples.” There are simply too many options, too many technologies, and too many overlapping approaches that make a coherent file management strategy difficult to achieve:

## T E C H N O L O G Y B R I E F

Virtualization, WAFS and WAN optimization, advanced replication, distributed or clustered file systems, document management software, classification and indexing software and “ILM” software, NAS clusters... the list is endless and growing. In the face of this complexity, we all need a common reference architecture that everyone – both IT teams and IT vendors - can point to and locate their technologies, challenges, and offerings without confusion. That file strategy framework is called a File Area Network (FAN).

A FAN is an architectural and methodological approach to enterprise-wide file management. It is an attempt to bring to the file realm the same manner of architectural commonality and nomenclature that was applied to block IO via the SAN over ten years ago. By agreeing to and implementing against a common architectural reference model, both the end user community and vendors alike can achieve much greater efficiency.

### Understanding the FAN

A FAN is ultimately about applying business level controls and intelligence to files; something not possible with block level data that is necessarily void of business or application level context. This key shift is achieved through a focus on network-centric file management rather than our collective historical reliance on physical devices. When users begin to think and deploy file management in this network-centric fashion,

the power of a FAN can become immediately evident:

**FAN Benefit:** A shift towards enterprise-wide, pervasive controls of all file information, and management of file attributes based on metadata and content values, regardless of platform.

**FAN Benefit:** The goal and ability to establish user file visibility and access rights based on business values (e.g. departments, projects, geographies) regardless of physical device residency, across the enterprise.

**FAN Benefit:** Establishing the expected best practice of non-disruptive, transparent movement of file information across all geographical boundaries.

**FAN Benefit:** The ability to create tailored file management software services that are deployed as true “services” to the entire infrastructure, not deployed in application-specific silos.

**FAN Benefit:** Finally establishing measurable ROI for file management due to consolidation of redundant file resources (e.g. de-duplication of redundant data, optimization of persistent capacities, appropriate data placement based on value, etc).

Certainly, those are all laudable goals for any infrastructure. However, it invites the question: What then are the network-centric technology elements that constitute a FAN? Below, we have identified 6 key elements in any FAN:

**T E C H N O L O G Y B R I E F**

**1. Storage Devices.** The foundational level atop which a FAN resides is the storage infrastructure. This can be either a SAN or a NAS environment. Ideally, this networked storage will be highly scalable and high performance. The only pre-requisite is that a FAN must leverage a networked storage environment to enable data and resource sharing, and the data to be managed is residing at the file level, meaning it has business and application context.

**2. File Serving Devices/Interfaces.** Either as a directly integrated part of the storage infrastructure (e.g. NAS), or as a gateway interface (e.g. SAN), all FAN must have devices capable of surfacing file level information in the form of standard protocols such as CIFS or NFS.

**3. Namespaces.** All FAN are premised on the existence of file systems with the ability to organize, present and store file content for their authorized end clients. This capability is referred to as the file system's "namespace". It is one of the central concepts around which the entire FAN must revolve. There are several kinds of namespaces possible in a FAN: Non-shared, shared, and Global. While beyond this scope of this profile, it should be understood that each kind of namespace enables varying levels of data sharing across physical file systems and platforms, with increasing levels of enterprise coverage. In no small way, the namespace is the heart of any FAN.

**4. File Management and Control Services.** The other central concept in the architecture of a FAN is the software intelligence that inter-operates with

namespaces to create new value across the entire enterprise. From a deployment perspective, these services might be integrated directly with file systems, or in networking devices, but they may also be stand-alone services. Examples include:

- migration
- replication
- classification
- load balancing
- data placement
- de-duplication
- access and controls
- retention (e.g. compliance)

To be effective, many of these services need to operate in real time, enforcing policies mapped to business requirements.

**5. End Clients.** All FANs have end client machines that access the namespaces created by file systems. Categorically speaking, these clients could be on literally any conceivable platform or kind of computing device.

**6. Connectivity.** There are many possible ways that a FAN connects its end clients to the namespaces. They are commonly connected across a standard local LAN, but they may simultaneously or alternatively leverage any manner of wide-area technologies, as well. More importantly, a FAN must truly function like a network wherein file requests can be received at any entry point in the FAN and be relayed to any other FAN device as necessary, just as an IP network operates today.

## **Key to FAN: Virtualizing Files**

One of the key technological advancements enabling the FAN to come into existence is the ability to abstract the file and its values out from the physical devices on which the file resides. This kind of file-level virtualization is essential because it ultimately allows the namespaces for multiple file systems to be adjusted based on business needs or enterprise-wide IT processes, not parochial device-level dependencies. Referred to variously as file virtualization or Network File Management (NFM), Taneja Group has been tracking this critical category since its infancy. Perhaps more than any other emerging technology category, we view this type of network-resident file abstraction as the linchpin to future file management strategies.

IT teams that are serious about investing in a file networking based approach need to seriously explore their options for abstracting – virtualizing - their file controls above their file serving devices. Without that critical move, it will be difficult to achieve the many economic and management gains a FAN can deliver. With that in mind, we will now turn our attention to Acopia Networks and explore that company's flagship ARX offering in the context of FAN. We believe that Acopia has done an excellent job foreseeing where the file networking space needed to evolve, and we view their ARX offering as perfectly compatible with the FAN approach.

## **Building a FAN with Acopia ARX**

Acopia Networks is a provider of intelligent file virtualization technologies targeted at

enterprise customers. For the purposes of this brief, we will broadly discuss the Acopia ARX family with an eye to its suitability for establishing an enterprise FAN.

Acopia's flagship offering family, ARX, is a comprehensive file control and delivery platform that inserts in-band between clients and enterprise NAS or file servers. From that position in the infrastructure, ARX is capable of providing a wide range of controls across the entire enterprise's file environment. In one integrated system, ARX offers three distinct and important functions necessary for building an effective FAN – virtualization, real-time policy enforcement, and sophisticated file routing capabilities.

It is important to note that ARX is not simply a stand-alone device. It plays a very real role as at the core of a FAN fabric as a router of enterprise-wide file information. At the most fundamental level, ARX should be understood as a means to elevate the entire range of file migration, replication, tiering, and load balancing functionality *above* heterogeneous physical file serving devices. This is a literal abstraction – or virtualization - of the file environment from the underlying physical infrastructure. ARX accomplishes this in a totally transparent and non-disruptive fashion to the end clients, and with total neutrality to the network protocols involved (e.g. NFS or CIFS.) It is also very important to note that ARX preserves the pre-existing file system lay-outs of the devices it virtualizes, but enables their namespaces to be accessed and controlled by both end users and administrators as if they were a single shared entity.

**T E C H N O L O G Y B R I E F**

Viewed from the perspective of building a FAN, it is clear that ARX has distinct advantages over historical approaches to file management. We believe that a file virtualization technology like Acopia's is essential to creating a powerful FAN. We can summarize these advantages as follows:

**Advantage: Heterogeneous support.**

A networked file management solution that can virtualize any NAS or file serving device while supporting both NFS and CIFS provides a massive efficiency gain. It means that an IT team could truly build a common management framework that is future-proofed against any architectural choices they make over time at the storage layer. That is a key requirement for a FAN and a powerful file virtualization technology such as Acopia ARX is at a distinct advantage due to its inherently heterogeneous architecture.

**Advantage: Advanced file controls**

By utilizing a file virtualization technology like ARX, it is much easier to establish a foundation for the range of advanced file level controls that are one of the key objectives in building a FAN. Because an in-band file virtualization platform is acting as a virtual mount for all file traffic, it means that it can exert very fine-grained, real-time policy-based controls across the entire infrastructure for file placement, file migration, file replication, and automated file movement or balancing. This is all achieved without requiring the IT team to deploy a new file system infrastructure, which is the only alternative for achieving this level of file control. Again, a significant advantage for ARX and file virtualization approaches.

**Advantage: Network Residency**

ARX is a truly network resident platform. This means that it does not touch the end client environment, nor does it touch the storage layer beneath it. As such, in a very real sense, we can consider ARX as a means of creating a "file fabric" that loosely parallels the block-level fabric first created by SAN switches in the prior decade. In the same way that SAN fabric switches enabled an explosion of flexibility and sophistication in the block IO world, a network-resident file virtualization tool becomes the heart of a FAN. It frees up the servers and their applications to deploy specific tools as required, regardless of platform. It also enables the storage layer to grow and expand in an open fashion, without tying file level controls to the fate of the array. In short, placing namespace creation, file level policy enforcement controls, and sophisticated file routing in the network creates a very powerful leverage point, and ARX is well positioned to provide that value to customers.

All of these advantages point towards the central role that file virtualization platforms will play in the future of File Area Networking. The gains that come from this network-centric architecture are difficult to argue with, and we already see enterprise customers agreeing with this assessment.

**Acopia's Key FAN Differentiators**

When we look specifically at the role Acopia ARX can play in building out a file management strategy, we see several distinct differentiators in the ARX design center that

**T E C H N O L O G Y B R I E F**

make it compelling. These are the factors that we believe will continue to help the company compete effectively in this extremely competitive emerging market.

**Differentiator: Enterprise-caliber**

The team at Acopia built the ARX solution with the very clear intention of being able to handle billions of files for high-end multi-petabyte NAS deployments. They also paid very close attention to the very complex availability, fail-over, and recovery requirements that any company developing an in-band platform must address when selling into today's Fortune 500 data center. ARX is designed to be a file virtualization and control platform for an *entire* enterprise, spanning not just all heterogeneous file capacities, but all geographies, as well. This is non-trivial engineering, and it adds up to a viable technology for even the most ambitious FAN deployments.

**Differentiator: Platform focused**

The design approach that Acopia took with ARX is decidedly platform-focused. The offering is intended to become an ongoing resource for conducting all manner of advanced and automated file management. ARX is *not* a point solution for migration, replication, or data placement. Rather it is a robust network-resident resource intended to enable and augment all manner of common file management tasks. ARX customers can significantly reduce the amount of management time spent dealing with individual file management operations because they are utilizing a common platform to achieve more with less. We believe this approach makes good sense, for the major reason that customers can turn

that converged functionality of ARX into a hard ROI story that justifies their FAN investment.

**Differentiator: High Performance**

One thing that stands out immediately with the ARX architecture is its ability to handle significant file throughput at over 2.5 Gigabytes per second, with support for billions of files. And by utilizing the offering's various file-level controls to optimize how data resides on a FAN, the platform can actually increase aggregate throughput. Why is this so important? IT teams should note that building a FAN also requires optimizing IO and throughput so that the file serving device layer can scale effectively to support the enterprise-wide file delivery mission. Additionally, ARX can bring even more high performance by enabling users to federate certain file workloads with dedicated high performance servers, thereby greatly accelerating NAS performance. These kinds of performance gains all improve the overall ROI of the FAN, lessening the need for incremental device-level investments. ARX should enjoy competitive advantage on this front if Acopia continues to invest aggressively in building multi gigabyte/second performance into their offering.

Because of these competitive differentiators, it is apparent to us that Acopia has what we might term "FAN principles" embedded into the ARX product. It is a platform very well suited for duty at the core of an enterprise FAN deployment.

## **Taneja Group Opinion**

File Area Networking is very much in its infancy as a discipline. In the same manner that SAN concepts and management evolved over several years, we expect to see FAN approaches and architectures evolve too, albeit with greater speed than in the block IO world due to the wide range of innovation now taking place.

Amongst the innovations that we deem central to the emerging FAN discipline is a category we term Network File Management (NFM), a core component of which is file virtualization. Without abstracting file control and delivery from the physical device layer, the FAN cannot happen. The namespaces of a FAN derive their management gains from that abstraction. Thus, we view file virtualization as the most elegant and forward looking way to achieve that abstraction: it is network resident,

scalable, heterogeneous, and adaptive over time.

Acopia and that company's ARX platform are excellent examples of how a FAN can be built around an intelligent file virtualization platform. We see Acopia's advantages in scalability, performance, and converged functionality as keys to its success in this emerging FAN market.

For IT decision makers who are serious about deploying a file management strategy that will scale and adapt across their entire enterprise for years to come, they need to evaluate Acopia ARX. We are quite confident that the FANs of the future will be using this class of virtualization technology aggressively.

---

*NOTICE: The information and product recommendations made by the TANEJA GROUP are based upon public information and sources and may also include personal opinions both of the TANEJA GROUP and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. The TANEJA GROUP, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors which may appear in this document.*