TANEJA GROUP
TECHNOLOGY ANALYSTS

S O L U T I O N P R O F I L E

WAN File Tiering – High performance access for distributed data

August 2010

THE NET NET

"Set me free!" Too often today that is one of the clearest statements we hear from storage administrators - whether block or file - who feel tied to an infrastructure that seemingly restricts how they support the data and storage needs of their business. While the restrictions are multi-dimensional - sufferings revolve around capacity, performance, complexity, availability, adaptability, and more - one especially sticky problem is *access*. Moreover, such access-based problems underpin many of the other struggles. Digital data is undergoing explosive growth today in part because it is *important*. The problem is, the more important data is, the more people there are who will need to access it. Many years ago a nearly revolutionary technology came to market that changed how we access data: Network Attached Storage (NAS), with the promise of easy access over the network. But while NAS might have been revolutionary for shared access at one point in time, that point is past, and that historic revolution supported a different generation of compute - one where data and users were more often in close proximity to each other. In the age of the mobile workforce, a global marketplace, and follow the sun distributed workforces that use data 24x7, the workforce is now scattered, but we still require the same access to data. In an industry where the efficiencies inherent in consolidating data in one location has long been considered nirvana for IT shops everywhere (and reasonably so), accessing multi-terabyte data is now a problem.

For the most part, the vendors who are the leaders in shared storage have yet to respond. But that isn't the case everywhere. The founders behind Avere Systems set out to tackle all of the sticky problems surrounding NAS, and have found a particularly attractive problem to turn their attention to in access – especially when organizations need to access large data sets over distances. Avere Systems calls this WAN File Tiering, and in this solution profile, we'll discuss how it is equipping NAS users with a new set of data access capabilities, with compromises in neither their performance nor their quest for data consolidation. Key to the equation: Avere understands that access is about sufficiently performant access from *anywhere*.

Feeding the NAS beast

Over the past decade Network Attached Storage has rapidly become the lifeblood of mission critical work processes and data flows. Behind all manner of businesses, NAS may be found supporting key data. But such file storage over the network may be even more important in some business than others - specifically where NAS is the enabler behind enormous amounts of data storage and shared access; in such industries as oil and gas, media and entertainment, life

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-5040 Fax: 508-435-1530 www.tanejagroup.com

sciences, design and engineering, insurance and finance, and others. The examples are easy to identify - multi-terabyte seismic data files, high-definition video content, or the huge data sets behind microsecond sensitive automation or decision support applications like trading platforms.

Yet for such businesses, the same NAS-based digital data storage that may well be the foundation of their business too often also feels like a ball and chain shackled around their neck. The problem is that NAS within many businesses has become a beast that must constantly be fed more capacity and more performance in the form of more spindles and more NAS arrays. The constant feeding of this monster keeps the storage organization from investing their time and energy in better IT capabilities. Within such a paradigm, today's investments will never meet tomorrow's demands, and the cycle marches on.

Yet for the data driven business, today's problems are even bigger. The limitations that face an organization that is feeding this beast do not stop with perpetual upkeep. The organization that is tied to ever-growing file data on top of this NAS infrastructure often finds they become restricted in their *access and use of that data*.

Storage-inaccessible

It's ironic but the very NAS solutions that have long rested their differentiation on ease of access are now themselves *limiting access*. The problem is the digital age has changed the flow of data, and access is no longer about serving a local set of users and systems attached by a corporate network wire. Now, the user needing access to digital data can be anywhere. Moreover, in many cases the data a user needs access to is no longer made up of the run of the mill word processing documents and spreadsheets that have in the past been the focus of business productivity. Now all manners of business revolve around enormous sets of digital data, and many parts of that digital data may be made up of very large files or millions of small files.

Nearly every industry has its example, but some are more obvious than others – computer-generated imagery (CGI) is one example that is often turned to. Recently, in the course of investigating the issues confronting users, we talked with one CGI company that was experiencing just such a challenge. Core activities at this company revolved around computer generated digital content using work teams scattered across several corporate locations, yet only one location had the resources of a special purpose CG rendering farm. The challenge was two fold. First, their big movie projects are collaborative efforts between the locations, and moving large volumes of files between locations wastes time and money as idle humans wait for the data to arrive, over a period of time measured in hours or days. Second, even without a render farm, the users at other locations require high-speed performance from their storage in order to develop the wire frame models, textures, backgrounds, etc. that will ultimately be rendered into movie scenes. In turn the administrator is left with a single best practice that doesn't hardly look like it should be a best practice: Build out an expensive secondary high performance storage infrastructure in each location, along

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-5040 Fax: 508-435-1530 www.tanejagroup.com

with data protection. Then surround the organization with complex data transfer processes to make sure everyone knows who has the master copy when, and prepare the organization for idle time during transfers.

Better solutions look few and far between. WAN optimization solutions may accelerate data transmission *throughput* but simply cannot hide the *latency*, or waiting period, for each file as it transfers over the WAN. In addition, these solutions do little to address storage performance once data is at the remote destination. Cache solutions may run out of steam before caching a full set of data, especially with the project sizes involved in movie projects. And many vendors approach optimization with an office productivity application and CIFS focus that has yet to add NFS - the protocol behind large file, high performance storage needs - to their radar. Each of those solutions has its place in optimizing data transmission and remote access for typical business users, but the very large data set remains unaddressed. And large data sets sit behind an ever larger number of businesses today, crossing life sciences, oil and gas, media and entertainment, financials, and more.

NAS from storage industry innovators once helped roll the snowball of today's digital revolution off the precipice by providing shared access, but the innovators behind the shared access that has enabled render farms to create terabyte sized digital projects now has no answer for the demands of the next generation of the digital age - the distributed digital work process. In turn, for any business working with large data sets, the opportunity to utilize global workers, industry experts, business partners, and even digital distribution may be lost, or at best mired in a battle with the complexity of constant data movement.

Spotlight on Avere Systems: Data in the right place, at the right time

While established vendors have long proven reluctant to innovate in the face of this challenge, new vendors are recognizing there is opportunity while the giants rest on their laurels. Such vendors are making early use of the latest storage technology to leap well ahead of the capabilities of yesteryear's storage. But some also hope to depart from the traditional approach, and chase the problem not with an entirely different storage solution, nor an entirely additional layer of network or transmission technology. Instead they are coming at the problem by augmenting the capabilities of yesteryear's technology by helping your existing storage to hold up under varying performance requirements. It turns out that this performance can also be accessed from anywhere.

The leading vendor in this case is Avere Systems. Avere Systems set out to sever the performance constraints created by legacy systems and their deep integrations with capacity oriented rotational disk, but in that severing, Avere designed a system that can distribute the performance of a NAS infrastructure to wherever the users are, irrespective of where the master data may reside. Let's take a look at how Avere is using this approach to address the performance and access requirements of users today.

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-5040 Fax: 508-435-1530 www.tanejagroup.com

Performance Tiering – accelerating access

When applied to the problem of high performance access to stored data, Avere's FXT series appliances tier data from a traditional file storage infrastructure onto a collection of performance-optimized SSD and (optionally) SAS disk within an FXT appliance. Using Avere's Tiered File System (TFS), up to 25 FXT appliances can be clustered together as a single storage system in an N+1 high availability arrangement. In contrast to temporary caching, Avere policy-controlled algorithms can move active data to clustered FXT appliances that are transparently inserted into the network in front of existing NAS systems, and leverage FXT processing power, power-failure tolerant NVRAM, high-speed DRAM cache, up to 512GB of SSD, and optionally SAS disk to augment the performance of the NAS infrastructure. Avere offers a range of FXT appliances, from the entry point with the disk-based version, to the mid-range unit with larger drives, to the entirely SSD-based product. Moreover, these appliances can be mixed within a single cluster for an ideal balance of capacity and performance.

Leveraging the Avere policy engine, businesses can fine tune what data can be tiered, how long it is stored on the FXT, and when and why it is moved back to traditional NAS storage. Moreover, organizations can maintain the original data on the original NAS infrastructure, and set intervals at which Avere will automatically resynchronize ("write-back" in Avere terminology) an updated copy to the original NAS, so that the organization can continue to use their existing snapshot and data protection practices. FXT simply exclusively locks this data to protect against change while it is tiered, and it stays locked (but periodically updated) for as long as the FXT has a tiered copy of the data.

Taneja Group views this fundamental approach as a new innovation in "Performance Tiering" whereby data is not necessarily relocated, but copied and redistributed across performance augmenting storage systems that will extend the overall capabilities of the NAS infrastructure.

For the access conundrum, these capabilities are significant - it means that the file storage infrastructure can be grown in performance irrespective of capacity. Latency for access is massively shrunken, as an initial request for data can with the same request move a copy of data to the Avere FXT. Further reads and writes can then take place on the Avere FXT, including the creation of entirely new data. Only during scheduled intervals will the data be written-back to original storage - until then, performance is completely offloaded from the original storage system. Eventually, when one of the write-back actions determines the data is no longer being accessed, the data will be entirely removed from FXT, keeping the infrastructure performance optimized as usage patterns change.

Enter WAN File Tiering

Using the same technology foundation, it is easy to see that such performance-extending capabilities might have application to file storage needs outside the data center. Without a doubt, performance tiering could

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-5040 Fax: 508-435-1530 www.tanejagroup.com

extend the performance of the lower cost and performance storage systems that often equip the branch office, but the possibilities do not stop there.

With FXT in tow, organizations might do away with some additional storage systems altogether by extending access to their data center storage systems across far reaches of a moderate latency corporate network, while keeping data consolidated within the data center. To do so, an FXT cluster is simply deployed in the remote location - as data is requested, only the data that is needed is copied to the remote Avere FXT, and it is copied in the order it is requested. This optimizes the transfer of data, and makes it available for high performance shared access at the remote site as rapidly as possible - allowing clients to interact well before the entire data set is copied. Simultaneously, the organization can do this while providing high performance read access to the data as soon as the first data segments arrive, as if the entire file was already entirely stored within a solid state or disk based FXT appliance (writes today are executed in write-through mode, where the full write is carried back to the data center in order to guarantee data integrity).

Avere does this by employing the same, granular sub-file tiering that they use when deployed in the data center, and using that technology to move only a copy of the needed data across the WAN, while leaving the original data in place.

This isn't to be confused with WAN clustering of file systems - a technology that does not yet truly exist for the multi-site enterprise – Avere is employing a fundamentally more scalable and less constrained tiering approach. With tiering, files can be dynamically pulled from a data center to any remote location, enabling many remote sites to deploy FXT appliances for access to the same data center storage infrastructure.

Moreover, the Avere customer can elect their preference of two ways to implement Avere. One is a read-acceleration implementation, and one is a read/write acceleration implementation.

In the read-accleration implementation, a customer is typically focused on moving data to a new site where it is primarily *read* from high performance storage. For this, an unlimited number of remote Avere systems can be attached to a single file system back in the data center (or some other central location), and the remote Avere systems will be configured in write-through mode. Write-through can tier unlimited copies of data to multiple remote locations to accelerate the reading of data, but will only minimally improve the performance of writes across the WAN. True to the definition of write-through, write data will always be written-through to the original file system, with the latency of remote transmission, albeit with Avere optimization of any chattiness or redundant transmissions. Such write-through deployments are particularly useful when many users need access to the same master content, and where new versions or copies of the original data are created once the data is accessed – video content creation is one such example.

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-5040 Fax: 508-435-1530 www.tanejagroup.com

For customers desiring accelerated remote write performance as well, read/write acceleration attaches a single remote Avere system to a single mount point on a one-to-one basis, and configures the Avere system in write-back mode. This allows Avere to exclusively lock files, handle all remote writes, and periodically copy all changes to the original file system at an interval of the customer's choosing, while protecting and preserving data integrity even across temporary connection or power outages. Multiple remote locations can be served through multiple mount points, and customers may find they're able to share effectively single master data sets even with multiple mount points by performing automated copies or tiering back in the data center. Furthermore, since the remote site will not experience storage performance impacts when the remote data copy is written back to the original storage, data writes can be copied to the master mount point as often as the customer desires, with no impact upon remote storage performance. Storage administrators can then coordinate data protection practices – snapshots and backup – to protect the master mount point whenever writes are copied back from the Avere system. The net effect is a single set of management and data protection practices – back in the data center where such practices belong – that protect all the data tiered to a remote site, even when it is undergoing change at the remote site.

The net effect from either implementation approach is that remote access to file data is accelerated for any location needing access – even large files are accessible as soon as the initial sub-file segments start arriving – and continued access can scream along at a blistering speed. Moreover, the WAN transmissions to communicate with the data center storage are minimized once data is pulled across the wire.

The Avere approach to WAN File Tiering puts data into the "right place" at the "right time", and puts it behind the right performance at the right time, while keeping it in the right place for the long term too. Moreover, it lets businesses move the right data (with sub-file movement) and move it rapidly in response to changing demands. Once data is tiered out to a remote location, Avere FXT will work in the background to keep the data synchronized with any writes happening at the main data center, until policies determine that the data should no longer be tiered to the remote location.

Let's look at a few specific places where WAN File Tiering is bound to have a significant impact.

NFS High Performance

For the customer with high performance requirements, NFS performance in the data center can be a serious enough challenge, but performance at remote sites quickly begins to look like an insurmountable hurdle. It is not unusual for us to talk with customers with hundreds of NAS silos strung across huge geographies with enormous associated maintenance, operations, and routine replacement costs.

While some vendors have innovated around optimizing the performance of some remotely accessed data, their efforts to this day have remained focused on Microsoft

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-5040 Fax: 508-435-1530 www.tanejagroup.com

Windows and Office productivity applications, leaving customers with few mechanisms to accelerate NFS performance at a distance. Moreover, where they could accelerate performance for small data sets behind such applications, today's large data sets could never fit in their memory and capacity footprints. In turn, businesses requiring performance at a distance face the deployment of yet another storage system, with all of its requisite maintenance and management. And if it is for purposes of performance, that storage system is likely to be big, containing more capacity than needed in order to get the rotating spindles that drive performance that may in many cases be as much as 90% read transactions. Provisioning that remote storage system with all of those spindles carries with it a significant price tag beyond just purchase – in the form of management, power, cooling, and floor space.

Avere FXT is an entirely different solution – and one that is not a stand-alone, over-provisioned storage system that requires separate care and feeding. Instead, Avere's FXT is a small footprint, performance enhancing extension of the enterprise's consolidated data repository. And the extension of the consolidated repository can reach across data center boundaries to carry storage access right into the hands of remote offices.

WAN Movement of Large Data Sets

Similarly, the enterprise involved in the periodic movement of large data sets may well find that WAN File Tiering provides a solution where others fall short. With large data sets, the complexity of moving, tracking, and controlling data can be immense, not to mention the requirements of protecting and managing that data once a move is completed. For many businesses, data movement or transfer is a burdensome and manpower consuming practice, but is nonetheless necessary behind key business workflows.

Long ago, WAN optimization and file transfer vendors eagerly rushed to meet the needs of this market, and in part, they have helped - transfer vendors can reduce complexity, and optimization vendors can reduce the time and bandwidth required to move data. But today's businesses often revolve around 24x7 work processes, with bigger teams than ever before working among a globally distributed workforce. Movement solutions haven't focused holistically on providing "access" to the single master sets of data at the center of this global workforce - data that doesn't benefit from being scattered to every corner of the globe. Many solutions can move data a little faster or a little easier, but fail to focus upon accelerating the first possible moment of access, and are far from providing the integrated high performance storage necessary to make a first possible access moment useable. Such are the capabilities of Avere FXT, and for the large data set business, the FXT WAN File Tier can make the details of movement vanish, and turn enterprise storage into a movement-free experience in accessibility.

The 1-2 Punch

Truth be told, while we've separately discussed the benefits of both WAN data movement and NFS performance

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-5040 Fax: 508-435-1530 www.tanejagroup.com

acceleration behind remote-site Avere FXT appliances, many customers will simultaneously get both sets of benefits. Such was the case with the CGI customer mentioned earlier in this solution profile. That customer deployed Avere's FXT in a remote site to conquer both data movement challenges, and cost effectively provide enough NFS horsepower for remote users to carry out demanding movie-making tasks. Now, on demand, remote teams can begin accessing data, with no hassles of data transfer. Those same remote teams have all the horsepower needed to read data out into new movie projects, and they save on the significant CapEx and OpEx burdens of remotely deploying a standalone storage system to support this high bandwidth read IO. Meanwhile, the administrators are keeping data consolidated, where it is protected with a single set of practices, within a single infrastructure.

Taneja Group Opinion

Avere is a welcome breath of fresh air in an industry where the established vendors have seldom been able to address the differing demands for access and capacity. Moreover, they've recognized the access conundrum is about more than just severing the connection between performance and capacity, but also involves extending performance to anywhere. By wrapping a NAS infrastructure in a cluster of FXT appliances, a NAS infrastructure becomes focused on the task it does well - storing, managing, and protecting data. FXT in turn is applied to the task it does well - providing high performance access to that data, and making such performance scale. With FXT at remote locations, data stays optimized and consolidated, significantly reducing a key demand driving NAS sprawl.

For other administrators in businesses that work with large data sets – those data sets with a size untouchable by solutions of the past – Avere carries with it yet another one-two punch: not only does it enable new access to data, but it can significantly simplify the infrastructure and practices already in place. In this sense, the future might just demonstrate that Avere's performance and WAN tiering are changing the way storage is done, and this time the change will be about access.

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-5040 Fax: 508-435-1530 www.tanejagroup.com