

Product Brief

Avere: Delivering on the Promise of Tiered Storage?

Date: October, 2009 **Author:** Terri McClure, Senior Analyst

Abstract: *Storage has become increasingly complex, with multiple flavors of disk drives and storage networks available to meet multiple performance, availability, and price point requirements. But data performance and availability requirements change as data ages. Optimizing the storage environment by moving data to an appropriately priced storage tier as it ages is a great way to reduce cost and storage footprint, but lack of automated tools makes it a cumbersome and laborious process that few firms actually implement. Avere's new FXT Series addresses the automation issue and should help users finally realize the promise of a seamless, perpetually optimized file storage environment.*

Overview

The storage industry has an assortment of labels that describe the capability to align data storage performance, cost, and protection to the value of that data. The more valuable the data, the bigger and badder the system it's stored on—at a higher price. As data ages and access slows or stops, data is moved to bulk storage tiers. All in all, it is a beautiful, compelling story: This type of alignment can save a bundle in the data center as data is moved off of Tier 1 storage onto bulk storage tiers, prolonging the life of Tier 1 storage investments and reducing both CAPEX and OPEX. Queue the choir of angels; this is data center storage nirvana. Or is it?

Vendors have been selling this story for years; calling it storage tiering, data lifecycle management, information lifecycle management (ILM), and, way back when, it was called hierarchical storage management (HSM). In reality, users have not seen the promised benefits because:

- **Solutions lack standards.** Sure, vendors have tools to automate migration, but they typically require that all of the storage tiers come from them (or even that the data stay in the Tier 1 array on denser disk drives). So if you've deployed what you believe to be a best of breed architecture with vendor A for Tier 1 and vendor B for Tier 2, it won't work. You need to be homogenous, because vendor A's and B's products won't work together, nor will their engineering teams. That makes life that much harder.
- **Granularity is not properly addressed.** Almost everything traditional block storage vendors, and some NAS vendors, do is done at the LUN or volume (logical unit or a logical representation of a disk drive) level. "Hot" LUNs get moved to SSD or high performance disk. That's great, but if only a portion of the data on that LUN is hot—as is often the case, you are wasting expensive storage real estate on cold data. NAS vendors often have a file, file system, or directory level approach that again provides insufficient granularity to ensure storage resources are optimized. In both cases, for block and file data, if the container is not sufficiently granular, even automation won't help.
- **Data migration is not easy.** Copying or moving data from primary to secondary storage consumes lots of system overhead at either the storage or server processor, depending on the technology used. That means taking a performance hit for those really important, 24 x 7 x 365 high performance applications that are staying behind on Tier 1 storage. And all the cleanup work is not pretty: moving data between storage tiers means remapping data paths or mount points to point applications to the new data location. In a world where data locations are mapped on massive Excel spreadsheets, there is a lot of room for human error in the equation and organizations risk losing data rather than moving it to a long-term storage tier.

It is a long sad story that's prevented users from getting to the promised land of real storage optimization—and these challenges are exactly what Avere was formed to address. Avere's products can be dropped into an existing file storage environment to ensure data gets stored on the proper storage tier to meet performance and availability requirements. It does so automatically and at a very granular level. If Avere performs as promised, it could be the long awaited path to the promised land of storage optimization.

The Universal Data Lifecycle

There are four simple stages of life for any kind of information—regardless of its format or the application that generated it. To run a cost effective IT organization, data needs to be managed and stored according to what stage of life it is in.

1. **Dynamic Active Online Data:** Data is born of an application in some form and is constantly changing. This is the stage during which you should focus the most energy, money, and time. Data belongs on the fastest, most fault-tolerant, and most secure storage; this should be the most protected time of its life.
2. **Persistent Active Online Data:** Just about all data stops changing at some point and becomes persistent. This point varies by data type, application, and corporate and personal perception. Once data becomes persistent, it needs to be treated differently. It does not make sense to keep persistent data on expensive big-iron storage, backing it up at set intervals, if it doesn't change. At this stage, business requirements still necessitate that data remains readily available. If requirements do not dictate immediate access, it should move to stage three.
3. **Persistent Inactive Data:** Not only is this data unchanging, it is rarely accessed. However, when it actually is required, it needs to be delivered in a reasonable timeframe—typically within seconds or minutes. Consider this your online archive or “bulk” data store where most of your data lives and it should be on the most economical, greenest, and easiest to manage online or nearline infrastructure. At some point, either due to the age of the data or the expiration of regulatory requirements that would require speedy data retrieval, the economics of keeping this data online just don't make sense.
4. **Doomsday Vault:** Otherwise known as the persistent inactive offline data tier for disaster recovery. The doomsday archive is pretty self explanatory—this is the stuff you send offsite and pray you never need to use. Often overlooked, organizations need to have it just in case you experience a DR scenario that requires a full remote recovery.

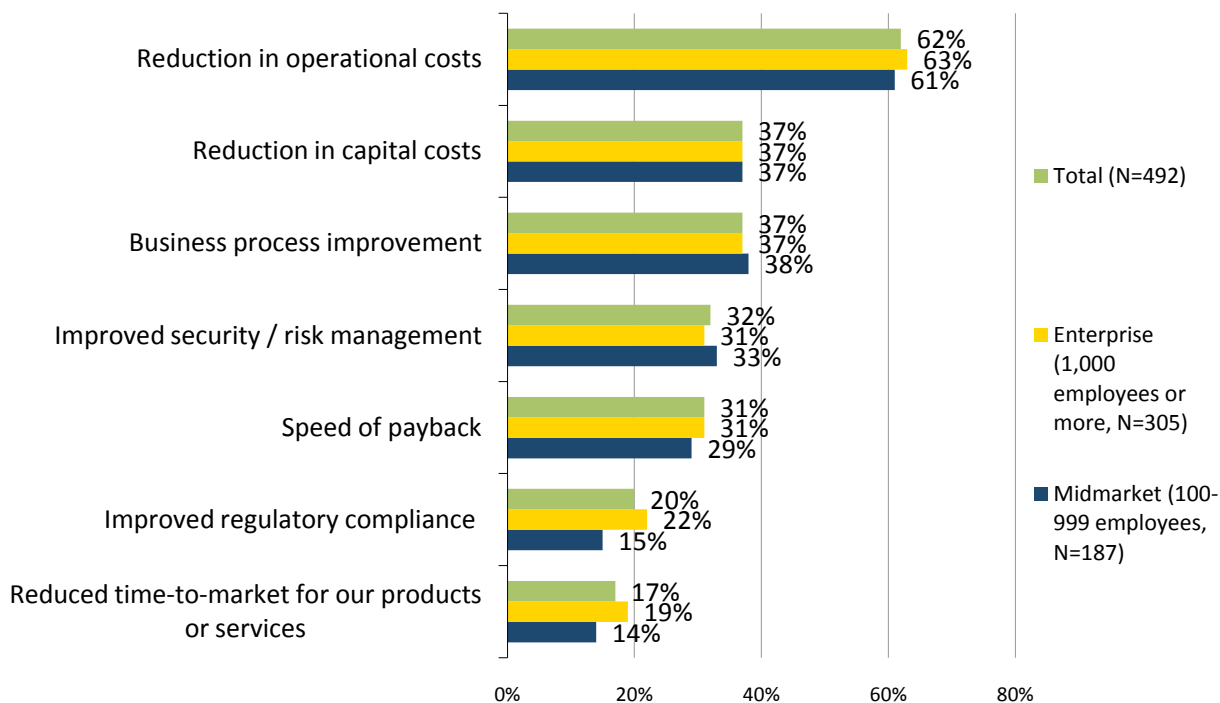
Each stage of the data lifecycle requires different infrastructure with different response times and different economics—and you need a simple way to migrate between storage layers. That is what has been missing. It would be even better if the solution would automatically migrate between tiers based on access patterns, with fine-grained control so that only truly active data blocks are on Tier 1 storage. There is considerable potential for reduction in both capital and operational expenditures with proper application storage alignment, thanks to putting less frequently accessed data on denser more cost-efficient storage media. But there is strong potential for even more long term operational cost reduction if users can eliminate the need to have storage administrators mapping, planning, and migrating the data—in a good sized data center, that would take a small army to accomplish efficiently.

Cost reduction is truly what it's all about when it comes to justifying IT spending. A big opportunity for savings comes from better data management. These strategies are now more important than ever as current economic pressures exacerbate the already-daunting challenge of efficiently managing data growth. In January and February of 2009, ESG conducted an in-depth survey of 492 senior IT professionals concerning their organizations' IT spending plans and priorities over the next 12-24 months. For both midmarket and enterprise IT professionals, reducing operational costs far outweighed any other criteria for justifying IT purchases for the foreseeable future (see Figure 1).¹

¹ Source: ESG Research Report, *2009 Data Center Spending Intentions Survey*, March 2009.

Figure 1. Most Important Considerations to Justify IT Spend

Which of the following considerations do you believe will be most important in justifying IT investments to your organization’s business management team over the next 12-24 months? (Percent of respondents, multiple responses accepted)



Source: Enterprise Strategy Group, 2009.

Automation and policy-based management are core requirements to effectively implementing the Universal Data Lifecycle, ultimately driving greater IT storage efficiency and reducing operational costs. Avere accomplished this through what it is calling its “demand-driven storage” architecture.

Avere FXT Demand-Driven Storage Enables the Universal Data Lifecycle

Avere’s new offering is an appliance that sits in the data path between NAS clients and mass storage arrays. There is no change required to the user’s environment; FXT supports standard file access protocols and clients (both NFS v3 and CIFS). Avere appliances contain either 1.2 TB (model FXT 2300) or 3.6 TB (model FXT 2500) of internal disk capacity and both models contain 64 GB of memory. They store data blocks that require high performance on either DRAM for highest IO performance or on 15K SAS disks for demanding performance requirements. Blocks of data that are less frequently accessed are dynamically moved to a back end, high density mass storage system of the user’s choice, allowing them to redeploy legacy NAS storage. FXT clusters start at two nodes, for availability, and can scale to twenty-five nodes with 90 TB of disk capacity and 1.6 TB DRAM per cluster. Avere claims that because of its scale-out architecture, which adds memory and processors as the system is scaled, it can attain performance of millions of operations per second and tens of GBs per second of throughput. FXT is a scale-out clustered platform that can independently scale performance and capacity while providing users with a single mount point to insulate them from underlying infrastructure changes.

Spindle and throughput bottlenecks are theoretically eliminated by storing hot data blocks in memory on a scale-out cluster of processors. This *significantly reduces the overall storage footprint* as any less active data is dynamically moved to high-density, more cost effective storage while maintaining application performance requirements. The big differentiator here is that Avere accomplishes all of this at the block level—a very fine grained unit of measure. Even systems that allow tiering within a frame only offer users the ability to move LUNs or volumes between physical disk types and onto higher capacity dense disk drives. The catch is that in a typical user environment, most of the capacity in the volume is static while about 20% is active. That means that over time, 80% of the blocks residing on the high speed

storage tier are inactive data. Avere has solved this problem. FXT offers the type of fine grained tiering that can provide true storage optimization. The best part is that this is an entirely automated process. There is no need to track all this movement through spreadsheets, do any manual migrations, or remap mount points. Avere actually delivers on what other vendors promise: the user sets policy and the system takes care of the rest. Avere delivers autonomic file storage.

Avere offers users a dashboard with a web-based management GUI that includes statistics reporting, graphical system performance monitoring, and all the usual e-mail and pager alerts, as well as SNMP discovery to plug into management frameworks. It supports existing backup and mirroring policies when dropped into a user environment, with a slider-based setting users can leverage to set recovery point objectives that trigger the system to copy data from the cluster onto the mass storage device so it can be backed up. This allows data to be backed up with existing data protection policies from the mass storage device without impacting user performance.

Avere's To-Do List

This is a new and different first generation product. That means it will take time to build end-user trust. IT has a lot of established best practices in place around storing and protecting data. Avere overcomes some of this by non-disruptively putting the scale-out appliance in the data path while preserving the ability to execute existing data backup and mirroring procedures; no need to change best practices. IT also has an inherent mistrust of true automation—users are control freaks, as well they should be when it comes to the company's data. No matter how much better or more efficient technology is, there needs to be a compelling event that jars and shakes up everything to create change. We could be at that point now, with the economy still creating budgetary pressure and server virtualization causing users to look at overall IT infrastructure optimization and virtualization. Avere's timing is good. But there are a few items that users will need to make them comfortable with the technology and that they expect from an enterprise storage platform.

- **Proof points.** Avere can't have enough users talking about their experiences with its products. The solution sounds almost too good to be true. Since users trust their peers' experiences and input much more than they believe vendor promises, Avere needs to get its users talking and connected with prospects. User proof points for new technology are the key to future growth, especially a fully automated technology that asks users for their trust that the system is making the right moves.
- **Support for multiple mass storage archive devices.** At launch, the Avere cluster only writes to a single mass storage archive device. If roughly 80% of the data in the data center is rarely (if ever) accessed, that can add up to a lot of capacity and could easily max out today's storage arrays even with high-density SATA drives. Support for multiple archive targets will be required as data growth accelerates due to corporate server virtualization and Web 2.0 initiatives. Down the road, extending to a cloud tier would make it even more compelling. It could be envisioned that one of the policies users set is whether or not the data can go offsite to create a policy-based hybrid cloud.
- **Support for NFS v4.** While NFS v4 has been slow to take off, it is starting to get some momentum thanks to better stability. Users want the NFS security enhancements incorporated in v4, but performance and stability issues with v4 have kept adoption low. Now that some of those issues seem to be worked out, expect to see users start to adopt v4. Avere will need support for v4 within the next twelve months or risk an adoption inhibitor.

The Bigger Truth

Storage has become incredibly complex! In the old days, it was disk drive or tape drive. We used to store data on disk drives and back it up to tape. Period. Now we have solid state, Fibre Channel, SAS, and SATA disks, all with different performance characteristics. The ability to move less active data off of expensive high-performance disk onto dense, slower, high-capacity drives gives users the potential for enormous cost savings while maintaining data accessibility never before available with a two tier disk-to-tape strategy. But the lack of tools to do this efficiently meant tiered storage (HSM, ILM, and so on) was only a pipe dream.

It is nice to see someone finally offer tiered storage at the block level to perpetually optimize the file storage environment. Even better, users don't need spreadsheets and analyzers to track data placement and do the data migrations between tiers themselves. Avere's value proposition is simple: it delivers storage that is cheaper, faster, and simpler to use. That's the same message NetApp used to make its name in the NAS space and it is a message that still resounds with users today. This is new and innovative technology that could shake things up.