

White Paper

Evaluating High-Performance File Systems for Cloud

A V  R E

By Jeff Tabor
Sr. Director of Product Management



April 2017

1.0 Determine the Best Path to Cloud

Perhaps elements of the following scenario sound familiar: one large hedge fund, appreciating the tremendous scale and potential economies of the cloud, made the decision to run more of its financial simulations in a public compute cloud. Hoping to minimize changes to existing file-based applications, the firm's first step was to explore options for making their data accessible to the cloud. Quickly discounting cloud-native file systems because of both expense and technology limitations, the firm selected a popular open source solution that touted scalability to hundreds of terabytes. Next, the firm devoted considerable time and technical resources tuning the file system to deliver sufficient performance to support thousands of compute cores. After months of effort and expense, the firm regrettably abandoned the project because of the complexity of scaling to and beyond a 10,000 CPU compute cluster.

Unfortunately for many organizations ready to embrace the cloud for select applications or even to adopt a cloud-first policy, the challenges associated with accessing existing file-based workloads can stop the process cold. However, viable solutions do exist. This paper reviews best-fit scenarios for three different types of file system options: simple cloud-provider file services such as Amazon Elastic File System (EFS) file storage; file system solutions like GlusterFS; and the Avere vFXT enterprise-grade clustered file system.

The discussion centers on the performance, manageability, and cost benefits of using Avere clustered file system solutions for data-heavy applications such as those typically found in genomics, animation and visual effects, and the real-world financial modeling scenario described above...basically any use case that requires analysis by large numbers of compute nodes. In these settings, Avere solutions offer distinct read/write throughput and cost advantages over Amazon EFS. Likewise, Avere provides a high-performance alternative to GlusterFS, in many cases offering equivalent or better throughput with the deployment and management simplicity.

This information is meant to help enterprises evaluating file systems for implementation into hybrid cloud IT operations avoid costly missteps and accelerate the path to cloud. The aforementioned hedge fund manager, for example, ultimately achieved success with the deployment of an Avere Virtual FXT (vFXT) Edge file system solution that delivered the ideal balance of price/performance and manageability. Within just 24 hours of the Avere implementation, fund analysts were running applications on more than 40,000 cloud cores—an outcome the firm had hoped to achieve months and thousands of dollars in operating costs sooner.

2.0 Consider Performance

The Avere vFXT system is a multiprotocol, highly scalable hybrid cloud file system that enables sharing and management of cloud compute, cloud storage, and on-premises resources. Avere vFXTs provide our FlashCloud technology and supports NFS-only, SMB-only, and multi-protocol NFS/SMB access. The Avere vFXT solution gives file-oriented applications running in the compute cloud seamless access to cloud and on-premises data sets.

The Avere vFXT serves as both a file system and a read/write cache that automatically places working data sets close to the compute cluster for optimal performance. Avere vFXT uses clustering to scale performance to hundreds of thousands of operations per second and throughput beyond tens of gigabytes per second.

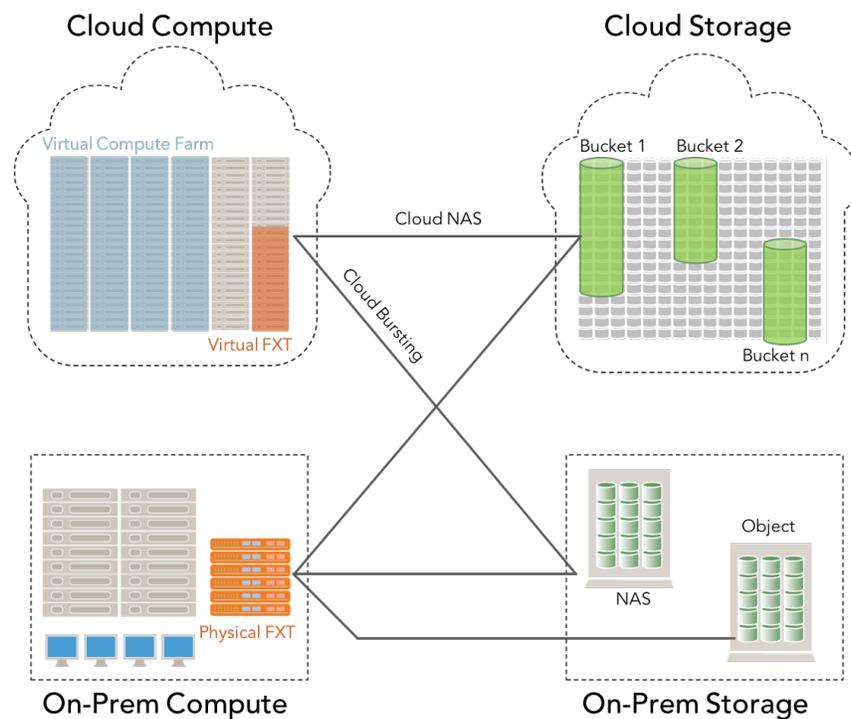


Figure 1: Virtual FXT File System accelerates performance of data stored on premises or in the public cloud

Evaluating suitability of cloud file systems requires taking into account expected workloads and data set sizes. Amazon EFS file storage for the Amazon Web Services (AWS) Cloud, for example, provides an easy-to-use solution with a simple interface that

allows fast creation and configuration of file systems. Ideally suited to “data light” applications, Amazon EFS offers an affordable solution at low capacities.

One issue with Amazon EFS, however, is that performance is extremely dependent on stored data set size. Throughput suffers on workloads accessing small data sets of one terabyte or less in size, and rates drop on random input/output (I/O) operations at any size data set. In both of these environments, an Avere vFXT can deliver significantly higher performance and throughput with a proven rate of more than one gigabyte per second for a one terabyte data set.

The Avere vFXT solution also provides more cost-effective scaling. In a 12-node cluster, for example, the Avere vFXT can achieve rates of up to 10 gigabytes per second of read throughput. The same 12-node cluster can deliver performance scaling to as many as 400,000 read operations per second. Amazon EFS does offer good performance at high capacity, but the speed comes at a very high price—currently \$30,000 per month for 100TB.¹

Options for more “data heavy” applications would include high-performance distributed or parallel file systems like GlusterFS. GlusterFS is available as free, open source software, and Red Hat offers a commercial version called Red Hat Gluster Storage. The GlusterFS network file system enables performance and capacity scaling for data- and bandwidth-intensive tasks. GlusterFS is well suited for HPC and big data workloads.

These high-end, parallel file systems offer tremendous flexibility and scale and have been used successfully in some of the world’s largest compute environments, national laboratories, and universities. But implementation and management complexities tend to make the systems impractical for most enterprise environments.

There are also, perhaps surprisingly, issues with performance. GlusterFS, for example, requires client code (which also adds to complexity) for optimal performance. GlusterFS recommends NFSv3 for small-file workloads, but performance is mediocre. In comparisons of random I/O performance, Avere vFXT can deliver throughput rates up to eight times higher.

¹ For current pricing, see <https://aws.amazon.com/efs/pricing/>

3.0 Weigh the Costs

As evidenced in the hedge fund scenario, free is not really free. To understand the true costs of any solution, enterprises must consider product and management costs in light of expected workloads and use cases. In terms of initial purchase prices, for example, Amazon EFS offers inexpensive options for low capacities, but, as previously mentioned, EFS is very low performance for data sets of one terabyte or less in size. EFS provides better performance for larger data sets on the order of one hundred terabytes in size, but remember this performance will come at a high cost.

Avere vFXT performance and pricing, however, allows starting small and achieving the performance needed for small data sets with the ability to grow to large data sets without “breaking the bank.” Avere uses caching to accomplish this, storing all the data on cost-effective S3-based cloud storage at under \$0.03 per gigabyte per month² and caching the 5-10% of the data that is active onto more expensive SSD/Flash-based block storage for low latency access to the compute nodes where the application is running. Caching data that is stored in S3-based cloud storage is not an available option with either Amazon EFS or GlusterFS. Without the ability to utilize lower-cost capacity, enterprises can expect to pay from three to six times more for block-level storage.

Support costs also vary widely. Although open source file system software sounds enticing, the costs of implementation and on-going management can make this a much higher-priced option in the long run. Enterprises that opt to use GlusterFS but lack in-house expertise, for example, should consider purchasing a commercially supported version. Even in this case, though, enterprise users need to weigh the costs of configuration, scaling, data access management, and performance tuning. Harder to quantify but obviously critically important is the net gain or loss of business opportunity—a file system implementation that enables nearly instant and seamless access to cloud compute can mean a much faster time to revenue.

² For current pricing, see <https://aws.amazon.com/s3/pricing/>

4.0 Factor in Complexity

Simplicity is not always good, and complexity is not always bad. An easy-to-use file system that has a high support overhead or doesn't offer required functionality—like snapshots or mirroring for data protection and recovery—may turn out to be a poor choice in a business-application environment. Similarly, a more complex file system that can be tuned to deliver high-speed access to large files may be the perfect solution for a scientific research environment. The key to making the right file system selection is to consider ease-of-use and functionality in the context of enterprise application requirements, resources, and expertise.

In regards to the Avere solution, consider that vFXT Edge filers offer the performance and versatility of high-end parallel file systems with the added benefit of management simplicity. Avere vFXT solutions provide storage access via standard NAS protocols and eliminate the need for rewriting existing applications for either access or performance tuning.

Avere solutions also were built to scale and can grow from 10 to 100s of thousands of clients. Each node added to a cluster contributes more CPUs and DRAM for performance, as well as more SSD capacity to support larger working sets and higher cache hit rates. Using Avere solutions, enterprises can easily take advantage of more cloud cores without making changes to existing workflow processes and without the complex configuration and setup associated with specialized cloud file services or systems. The hedge fund customer described earlier in this paper, for example, now uses an Avere vFXT cluster to run many times more cloud cores than was possible in its earlier GlusterFS implementation.

Another benefit of the Avere solution is the ability to run analyses on model data stored on a NAS array at the firm's on-premises data center via cloud bursting. When risk analysis jobs run in the compute cloud, the firm's Avere vFXT cluster populates its cache with the latest quant model data from the NAS system.

Avere support for cloud bursting enables enterprise clients to access on-premises data without having to migrate large data sets. In contrast to GlusterFS, Avere also provides technology to simplify data movement and mirroring. Avere FlashMove software enables simple, non-disruptive movement of live data, and Avere FlashMirror software enables replication for data protection and recovery.

5.0 Find the Balance

No single cloud file system perfectly addresses the requirements of every application workload. But with the range of options available, enterprises can expect to find solutions that effectively pave the way to cloud benefits. The table below offers a brief comparison of the key functionality available with the cloud file systems discussed in this paper. While each solution has its application-environment sweet spot, the Avere vFXT solution strikes an impressive balance of performance, cost, and management benefits to be a compelling option for enterprises working towards cloud readiness.

Feature	Avere vFXT	Amazon EFS	RedHat Gluster
NFS	yes	yes	yes
SMB	yes	no	yes
Multi-protocol (NFS & SMB)	yes	no	yes
Cloud storage support (S3)	yes	no	no
On-prem storage support	yes	no	no
Cloud compute support	AWS and GCP	AWS only	GCP only
Performance scaling	yes	yes	yes (proprietary client code)
Tiering (SSD + S3)	yes	no	no
Snapshots	yes	no	yes
Global namespace	yes	no	yes
Non-disruptive data migration	yes	no	no
Mirroring/DR	yes	no	yes
Encryption	yes	no	yes
Compression	yes	no	no
Analytics	yes	no	no

Table 1: Feature comparison of file systems available for the public cloud

For more information on using Avere solutions to take advantage of cloud infrastructure, visit <http://www.averesystems.com>.

Sources and Resources

Avere NAS Filer Systems for Cloud Enablement & Performance

<http://www.averesystems.com/products/products-overview>

Amazon Elastic File System

<https://aws.amazon.com/efs/>

File Servers on Google Compute Engine

<https://cloud.google.com/solutions/filers-on-compute-engine>

GlusterFS

<https://www.gluster.org/>

Red Hat Gluster Storage

<https://www.redhat.com/en/technologies/storage/gluster>

About Avere Systems

Avere helps enterprise IT organizations enable innovation with high-performance data storage access, and the flexibility to compute and store data where necessary to match business demands. Customers enjoy easy reach to cloud-based resources, without sacrificing the consistency, availability or security of enterprise data. A private company based in Pittsburgh, Pennsylvania, Avere is led by industry experts to support the demanding, mission-critical hybrid cloud systems of many of the world's most recognized [companies and organizations](#). Learn more at www.averesystems.com.

All trademarks, trade names and service marks referenced herein belong to their respective companies.

© 2017 Avere Systems, Inc. All rights reserved.