



Cloud Bursting for High-Performance, Scalable NAS

Key Benefits

- Flexibility to deploy storage and compute on premises or in the cloud
- Complete NAS feature set including, NFS and SMB/CIFS protocols, cloud snapshots, AES-256 encryption, and compression
- Clustering provides extreme performance and capacity scaling
- Active/active failover ensures data is highly available
- WAN caching hides the latency to remote data
- Global namespace integrates public cloud storage and NAS into a single logical pool
- FlashMove® software provides transparent data mobility between on-prem and in-cloud storage
- FlashMirror® software enables disaster recovery by replicating data between on-prem and in-cloud storage
- Simplified management provided by powerful GUI and API
- Cost savings of 70% or more over traditional NAS

On-demand cloud-computing infrastructure and the availability of near-infinite processing capabilities make the cloud attractive for enterprises seeking relief from the runaway costs of growing and maintaining private data centers. Compute clouds can serve as permanently provisioned IT infrastructure or provide overflow burst-compute resources to temporarily add horsepower for applications like 3D-rendering, video transcoding, genomic sequencing, financial analytics, and big data.

Utilizing the cloud to supplement on-premises infrastructure can help enterprises meet short-term peaks or unanticipated demand, accessing nearly instantaneous compute capacity without planning requirements for rack space, capital budget, or additional staffing. But for time- and performance-sensitive big-data projects, the job of presenting datasets that are maintained on premises to the compute cloud application instances can be a deal-killer. How do you non-disruptively move datasets back and forth across a private data center and multiple public cloud service providers? Or avoid network bottlenecks and high latency that could critically degrade application performance? What about file interface incompatibilities? And, how can you keep datasets available to both cloud and on-premises applications?

The Avere Hybrid Cloud NAS architecture with FXT Edge filer provides a high-performance, scalable NAS solution that eliminates the current obstacles enterprises face when extending compute-intensive applications to run in the cloud. The Avere Virtual FXT (vFXT) solution provides best-in-class NAS functionality (including NFS and SMB access) and blazing-fast I/O capabilities that distinguish the Avere Edge-Core architecture. The Avere Virtual FXT delivers performance and scalability to meet hefty compute demands, along with the simple installation and zero hardware deployment that enterprises require to take full advantage of cloud bursting.

Challenge: Use the Cloud, Lose the Latency

The cloud offers push-button deployment of pay-as-you-go compute solutions for organizations that need more compute power but lack the physical space, capital budget, and staffing resources to continuously expand in-house IT infrastructure just to keep up. Public cloud service providers simplify the process of moving even the most demanding applications onto their cloud platforms. But the more challenging aspect can be feeding the necessary datasets to the applications now running on cloud compute infrastructure. To make bursting applications to the cloud effective, existing datasets living in a private data center must be made available to the cloud compute without negatively impacting performance or limiting access to applications still consuming the data on premises.

Avere FXT Series

Purpose-Built for the Cloud

- Flexibility to deploy public cloud storage and NAS
- Scale application performance on premises and in the cloud
- AES-256 encryption, FIPS 140-2 Level 1 compliance keeps data secure
- Compression for storage efficiency
- Cloud snapshots for data protection

Best-In-Class NAS Features

- NFS and SMB protocol support
- Clustering scales performance to hundreds of GB/s throughput and over ten million IOPS
- Active/active failover ensures HA
- Dynamic tiering hides WAN latency

Simplified Management

- Global namespace integrates public cloud storage and NAS
- Transparent data mobility and replication across heterogeneous storage
- GUI provides rich, historical statistics & graphical monitoring
- SNMP & KMIP support, XML-RPC interface, email alerts

Cloud Support

- Microsoft Azure Cloud Compute and Storage
- Amazon Elastic Compute Cloud and Amazon Simple Storage Service
- Google Compute Engine and Cloud Storage (Multi-Regional, Regional, and Nearline)

Traditional storage vendors offer less-than-ideal options, as illustrated in Figures 1 and 2. In the first scenario, cloud application instances access their data via a virtual private network (VPN) to the enterprise's on-premises NAS. In this case, network latency and throughput issues promptly make the solution fail to scale for heavy I/O workloads.

Other options, as depicted in Figure 2, deploy a virtual NAS in the cloud. But this self-contained NAS silo in the cloud will affect the ability of on-premises applications to access the data due to network throughput and latency limitations. A possible solution is to replicate, move, and synchronize data between the cloud and your data center on-premises NAS before and after bursting your application into the cloud. But administrative overhead and revision management quickly scales out of control. In this bring-your-own-data scenario, data is not readily available for simultaneous reading and writing to both on-premises and cloud-based applications, and data management can be a very complex exercise in these cases.

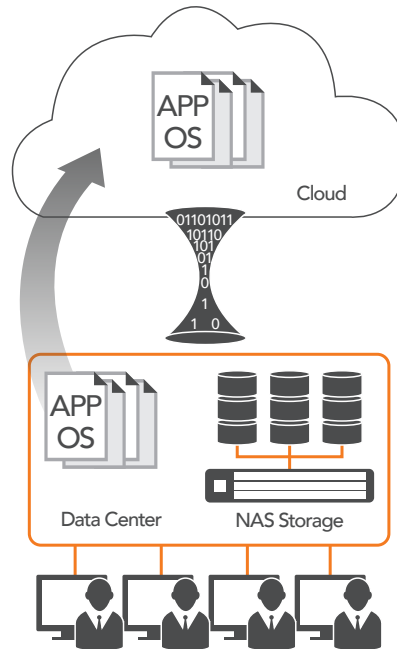


Figure 1




Figure 2

For burst compute efforts to be feasible and economical, organizations need a complementary storage solution running in the cloud that provides seamless access to on-premises NAS while minimizing latency and eliminating throughput bottlenecks. The solution should provide dataset access to multiple cloud regions, ensure scalability for growing workloads, and enable fast storage-resource set-up and tear-down for maximum agility and efficiency.

The Solution: Avere Virtual FXT to Run Big Data Applications in Cloud Compute

Avere Virtual FXT solutions enable the Avere Edge-Core architecture to extend into cloud infrastructure while resolving otherwise deal-breaking performance, scalability, access, and management challenges. The Avere Virtual FXT operates in the compute cloud, presenting your on-premises data storage resources to I/O-hungry big-data applications that require mounted NAS filesystems.

The Avere Virtual FXT solution defeats the challenge of trying to place physical hardware in a public cloud data center. Add a Virtual FXT node and scale the Edge filer cluster from three to 24 nodes to handle the I/O storm from large-scale cloud compute instances running your application.

Avere Virtual FXT solutions with Avere FlashCloud™ technology can also eliminate object and file-based interface incompatibilities, making it easy to present object-based, REST API resources to applications relying on NAS file interfaces.

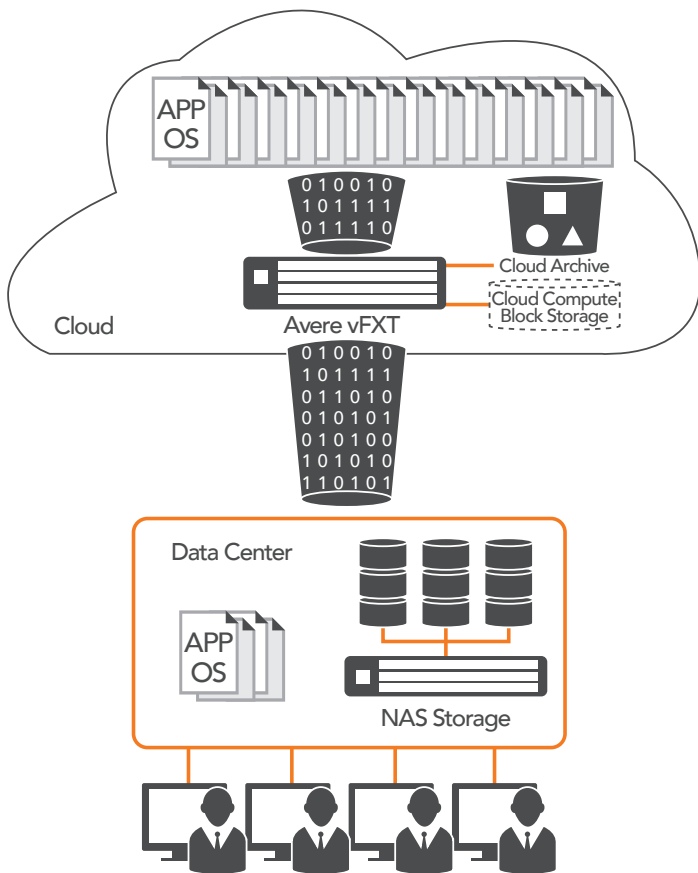


Figure 3: Avere vFXT eliminates WAN latency, delivers scalable performance, simplifies management with GNS, and provides an on-ramp to cloud storage.

Benefits

Performance for Big Data

The Avere Virtual FXT solution effectively creates a performance storage service tier in the compute cloud, allowing applications access to the data they need at multi-gigabyte/sec transfer rates and hundreds of thousands of I/O operations per second. Avere vFXT nodes can be dynamically added to or removed from the Edge filer cluster to ensure you can non-disruptively manage storage performance resources as application workloads expand.

The Avere Virtual FXT eliminates three big technical impediments of cloud computing: NAS versus REST-API incompatibilities, the bottlenecks of network connections to cloud service providers, and the data management synchronization complexities of shifting data from on premises to the cloud and back. The Avere vFXT provides a cloud-based front-end to your on-premises NAS—which can be your single, consolidated data repository—to provide data access from both on-premises and cloud-based applications. For maximum flexibility, the cloud-based vFXT can link directly to data center core filers or to existing Avere FXT Edge filer clusters elsewhere in your organization.

The Virtual FXT solution also delivers best-in-class NAS functionality, including NFSv3 and SMB2 access, high availability failover, and snapshot support for data stored in the cloud. Built-in global namespace (GNS) functionality hides the inevitable complexity of sprawling storage environments from users and applications to provide seamless access to enterprise-wide NAS data. Run your applications on compute cloud infrastructure without changes to client code or embedded pathnames and without complex data management processes.

Time and Cost Savings

Enterprises with data centers plagued by both limited space/power resources and inefficient use of existing compute resources can leverage the Avere Virtual FXT solution to take immediate advantage of the elasticity offered by pay-as-you-go compute cloud resources that scale to your most dynamic needs. The fully virtualized vFXT solution eliminates traditional data center real estate, power, and cooling expenses and allows you to tap into the unlimited compute power the cloud offers. Easy to set up, manage, and tear down, the Avere Virtual FXT solution also reduces IT administrative costs.