IDC PERSPECTIVE

Growth of File Storage Services in the Public Cloud

Deepak Mohan Amita Potnis Andrew Smith

EXECUTIVE SNAPSHOT

FIGURE 1

Executive Snapshot: Growth of File Storage Services in the Public Cloud

Recent years have seen a surge of demand and offerings delivering shared file storage in the public cloud environment. This includes marketplace offerings like SoftNAS and Elastifile, and newer services like Amazon EFS and the recent Azure NetApp Files. The growth of file services is a reflection of new needs of applications moving to public cloud and the growth of applications with storage requirements that prioritize performance and on-premises consistency over scale.

Key Takeaways

- File services in public cloud are growing and will soon become an integral part of all major public clouds.
- Common reasons customers use file services on public cloud are consistency with on-premises environment for hybrid operation or lift and shift migration and performance and low-latency operation for dynamic data sets such as application state information and configuration information.
- File services offer on-premises consistency and latency performance relevant for both existing applications moving to public cloud and new applications being developed on public cloud.
- Customers have options: overlay, native, and partner delivered native file services, in public cloud today.

Recommended Actions

- Evaluate storage requirements for applications targeted for public cloud — and build a good understanding of prioritization of scale requirements, performance requirements, and on-premises consistency requirements for different storage needs.
- Experiment and pilot projects to build internal familiarity with the different file storage options available in public cloud today. The low entry cost makes this possible without any up-front investment.
- Use multiple services during the pilot and early phases of the application deployment to allow selection of the optimum choice(s) for the storage need. The standard network file storage protocols (NFS, SMB) available across the different file services limit the level of lock-in, especially early in the life cycle.

Source: IDC, 2018

June 2018, IDC #US44002318

SITUATION OVERVIEW

Public cloud storage introduced an alternative that addressed several key challenges associated with storage infrastructure operations — the need for specialized administration, the long lead times for capacity expansion, the risk associated with investing too much or too little into new technologies, and the constant changes in facilities requirements. The practically unlimited scalability and flexibility allowed customers to avoid investment risk, offload facilities overhead, and greatly simplify their storage capacity usage and planning.

Object-based storage architecture lends itself well to this model and was the de facto shared storage mechanism in public cloud platforms, pioneered and established by early offerings like the Amazon Simple Storage Service (Amazon S3). This also matched up well with the early set of applications and use cases moving into public cloud platforms — typically comprising a mix of emerging web-scale applications, test and development use cases, and unstructured backup and archival use cases. File-based shared storage did not have a major presence in the early growth period of public cloud storage services.

The expansion in types of applications being deployed on public cloud has changed this status quo, and the past four years have seen a fresh focus on shared file storage services within public cloud — from both vendors and buyers. File-based storage has reemerged as an important part of the public cloud world, particularly for applications and use cases where:

- Customers are moving an enterprise application from on-premises to public cloud and prefer to avoid investment in refactoring or re-architecting the application to use object storage.
- Consistency with on-premises is important, and the application uses a file-based storage back end to share data with other applications or processes on-premises.
- The application does not have high scalability or volume requirements but does have strong performance and latency requirements.

This IDC Perspective is a brief overview of shared storage evolution, the current trend, and the types of public cloud-based file storage services available to customers.

Enterprise Storage Systems Evolution — From DAS to Object Storage

The earliest prevalent storage architecture used for enterprise storage needs was direct-attached storage (DAS) introduced in the 1960s. DAS essentially consisted of hard disk drives directly attached to a server through SCSI, SATA, Fibre Channel, or other direct connectivity. This storage capacity and data can be made available over the network to other servers on the network, through the server, enabling a shared storage functionality from the host server. DAS continues to be a popular storage architecture even today for local storage and for small-scale use cases.

The 1990s to the Early 2000s — File-Based Shared Storage Takes Center Stage

While simple, DAS architectures have intrinsic constraints around scalability and shared usage. These led to the emergence and growth of storage area networks (SAN) and network-attached storage (NAS) architectures. A SAN is a dedicated network for storage, typically over Fibre Channel, that connects a shared pool of storage to multiple servers on this dedicated network. The storage capacity appears to servers as locally attached volumes and is accessed at the block storage level. When being used in a shared storage use case, this includes a file system layer on top, which provides a consistent namespace and enables file-sharing across all the clients on the network.

A NAS is a dedicated storage device that sits on the local network, providing file-level access to storage capacity for servers on the network. With the success of early network file protocols like Novell NetWare's NCP and Sun Microsystem's NFS, followed by positive adoption of file-based storage servers brought to market by Sun, IBM, NetApp, and EMC (now Dell EMC), file-based storage evolved to become a prevalent mechanism for shared storage at enterprises.

File-based shared storage systems, typically NAS but sometimes also built on underlying SAN architectures, have since grown to become one of the most common shared storage frameworks in the enterprise IT environment.

The Late 2000s and the Early 2010s — Growth of Object-Based Storage

Since its introduction in the mid-1990s, object-based storage (or object storage) has been seen as an answer to the growing scale-out needs of modern applications. Object storage offered end customers a flat address space, a key-value based referencing, and RESTful API-based access. The major factors that enabled the growth of object storage were:

- **Ability to support massive scale-out requirements:** The underlying flat architecture of object-based storage makes it much better suited than hierarchical file systems for highly scalable infrastructure needs. This also allows it to easily support large volume storage requirement going to petabytes of capacity.
- **Low management overheard:** Owing to its simplicity, object storage solutions typically result in lower management overhead and easier scaling operations compared with file storage solutions. IDC's *File- and Object-Based Storage Survey Findings, 2017: Adoption and Workloads Trends — Part 1* (IDC #US43630018, March 2018) survey results indicate that over 50% of the respondents looking to add object-based storage to their environment were doing so because of management complexity with NAS solutions.
- **Exponential growth in popularity due to public cloud:** With the introduction of public cloud-based object storage services — such as Amazon S3, Azure Blob Storage, and Google Cloud Storage — object storage popularity grew rapidly both among the enterprise and the developer community. This also acted to fuel growth in on-premises object storage platforms.

In the public cloud environment — object storage was seen as the ideal complement to the scalability and ease of use value proposition of public cloud. Object storage was able to truly deliver the scalability and ease of use needed for the early applications moving to public cloud. Enterprise applications were not a primary target during the early days of public cloud. A large portion of this growth came from applications with less stringent hardware performance requirements — scale-out applications that were designed to deal with hardware failures, test, and development use cases that were not business critical to the enterprise and backup and archival use cases.

As public cloud usage expanded into more use cases, there was also the expectation that software architectures would evolve in the coming years and that software applications moving to the cloud would be re-architected to make them more cloud native in their functionality, aligning with the initial wave of applications on public cloud.

The Shifting Profile of Public Cloud Applications

The early 2010s saw a gradual rise in acceptance of public cloud-based infrastructure services for enterprise IT use cases. This resulted in gradual growth in the number of enterprise IT applications running on public cloud. Initially, these started with backup and archival use cases, further facilitated by cloud storage gateways like StorSimple (now Microsoft) and the AWS Storage Gateway. This was

soon followed by other use cases and workloads, including a growing number of production applications. Enterprise workload movement to public cloud has accelerated since 2015, and IDC's 2016 *Amazon AWS Infrastructure as a Service (IaaS) Survey* showed that over 90% of enterprise IT organizations using AWS were using it for production IT applications. New types of applications and use case scenarios that were increasing using public cloud included:

- **Enterprise IT application "lift and shift"** – where the IT organization moves entire applications into a public cloud IaaS environment – has been an increasingly common starting point among enterprises, as they look to jump-start their public cloud infrastructure usage. One of the enablers for smooth lift and shift workload movement is the ability to use similar underlying services in the public cloud environment as within the on-premises operating environment.
- **Hybrid applications** are partly a by-product of growing enterprise adoption. A majority of enterprise IT organizations anticipate that they will be using a mix of public cloud and on-premises resources for the foreseeable future. Hybrid applications are applications running in the public cloud or on-premises, which are architected in such a manner as to allow easy interoperability across public and private cloud resources. These could be for movement or for interactions with another application on-premises. Consistency across the environments makes hybrid operations easier, and an IDC survey of enterprise IT organizations revealed that consistency between public cloud and on-premises environments was the top priority for enterprises when evaluating future infrastructure investments.
- **High-performance applications** – applications with high-performance requirements – require low latency and strong consistency for their shared storage needs. The growth in acceptance of public cloud infrastructure has seen a growing number of speed and performance-sensitive applications move to public cloud. These include business-critical applications moving from on-premises into public cloud, as well as new applications being developed on public cloud.

Demand for File-Based Shared Storage Within Public Cloud

With the growth of these new types of applications on public cloud, there has also been a shift in storage requirements from the public cloud environment. Key among this is the growing preference to use file-based shared storage for certain use cases within the application. Factors driving the growth in interest in file-based shared storage in public cloud are:

- For several enterprise applications moving to public cloud, it helps to have consistency with existing on-premises environments – either to enable an easy lift and shift of applications into public cloud or to enable a hybrid infrastructure where both public cloud and on-premises offer similar environments to the application. IDC's *File- and Object-Based Storage Survey Findings, 2017: Adoption and Workloads Trends – Part 1* (IDC #US43630018, March 2018) survey results show that 79.3% of respondents use file-based storage for unstructured data associated with mission-critical workloads. As more enterprise applications are deployed to the cloud, a file-based storage that can offer similar functionality as the on-premises environment can help reduce challenges to movement.
- In reality, not all applications require the massive scalability offered by object-based storage. IT organizations and application owners have developed a good understanding of the benefits, challenges, and effort involved in re-architecting applications to use object storage.
- Even when considering new applications, there are readily available development libraries and SDKs to develop client software to access shared file storage services. These intrinsically handle underlying errors and common failure issues associated with these services. Unless the effort justifies the benefit, shared file storage offers a simpler path to completion.

- Performance and latency-sensitive applications prefer the low-latency access and hierarchical framework in file-based storage services. These include applications with highly dynamic data sets that need to be stored and retrieved at a high frequency, such as stateless application architectures that store state data into shared storage for rapid failover or for timely event triggered execution.

Types of Shared File Storage Services Available on Public Cloud Today

In response to the growing demand for file-based storage services in the public cloud, a number of solutions have been introduced to the market – both by public cloud providers and by other technology vendors. IDC views the shared file storage services available in public cloud today as being grouped under three types of offerings – overlay file storage services, partner-delivered native file storage services, and native file storage services.

Overlay File Storage Services

Overlay file storage services were one of the early shared file storage options in public cloud. This typically consists of an OS image that embeds the functionality of a file server and presents a network file sharing interface (such as NFS or SMB) to clients that can access it over the network. In the back end, the overlay service uses underlying native block and/or object storage service available natively within the public cloud environment. Overlay services may include high-availability configurations, as well as support for multiprotocol access to files stored on the service.

Overlay file storage services are typically available to customer through the major hyperscalers' marketplace or downloadable directly from the third-party provider's website. These are available as open source, licensed pay as you go, or perpetual licensing models and may also include offerings that are formally endorsed and supported by the public cloud service provider.

Popular examples of overlay file storage services include SoftNAS, NetApp Cloud Volumes ONTAP, GlusterFS, Elastifile, and Avere Systems.

Native File Storage Services

As the name suggests, native file storage services are file storage services that are available natively within the public cloud service provider's ecosystem and delivered directly by the provider. These are delivered, billed, and supported just like other services from the public cloud provider.

One of the key benefits of native file storage services is that that they are deeply integrated infrastructure up with the other resources used for the provider. This minimizes integration and connectivity issues and ensures a low-latency, high-throughput connectivity between the file storage service and other resources.

Examples of native file storage services in public cloud include Amazon Elastic File System, Azure Files, and IBM Cloud File Storage.

Partner-Delivered Native File Storage Services

Partner-delivered native file storage services are a recent model in the market, where the file storage service is available to customers from within the environment of the public cloud ecosystem – just like other native services from the public cloud service provider. However, the file storage service is openly based on technology delivered by a partner provider. The billing and support may, in theory, come either directly from the public cloud provider or from the partner technology provider.

Partner-delivered native file storage services bring together a number of benefits to end customers – the ease of access and use by virtue of being within the public cloud providers core workflows and the specialized technology expertise of the technology provider. The pioneer example of this new partner-enabled file storage service is the Azure NetApp Files, delivered by NetApp and available within the Azure public cloud ecosystem.

IDC expects file storage services available within the major public cloud ecosystems to grow in the coming years and become an important component of the storage portfolio at all major public cloud providers. In acknowledgement of the importance of these offerings, the portion of public cloud IaaS storage revenue attributable to native file storage services is now tracked in IDC's Worldwide Storage Software and Cloud Services QView.

ADVICE FOR THE TECHNOLOGY BUYER

The growth of file storage service options available in public cloud bodes well for public cloud customers, as it increases the options available for shared storage on the public cloud. This applies to both existing applications moving to the cloud and new applications being developed targeting a public cloud environment. For existing applications moving to public cloud, this reduces the pressure on application architects and owners to re-architect and rewrite the application to use object storage on the cloud. This also allows for interoperability with other on-premises applications, enabling easier hybrid operations. For new applications, this increases the storage choices available and more easily allows architectures that prioritize latency performance or hybrid operation needs. This is particularly useful for new applications built on state machine architectures, such as those built on container or serverless compute, and use shared persistent storage to maintain rapidly changing state data.

Along with adoption of file storage services in the public cloud, it would also help consider other data management capabilities that may be required alongside – such as replication, backup and recovery, multiprotocol access, and tiering. Early consideration and incorporation of these capabilities would help build a future resilient data platform and would be beneficial as use cases scale or evolve in the long run.

File storage services can be expected to become first-class services within all major public cloud environments, and customers should leverage the growing file storage options where applicable. This can be done by evaluating and building a good understanding of the storage priorities such as scaling requirements, latency performance requirements, and hybrid operation requirements. The low cost of entry on public cloud makes it feasible to quickly pilot, build internal familiarity, and understand the benefits of each of the different file storage options available with public cloud. The availability of standard network file storage formats (NFS, SMB) across multiple file storage services also reduces the effort needed to switch between file storage services, particularly early in the life cycle when data volumes are still not a barrier. So organizations should also invest in pilot projects to quickly start using currently available file services in the cloud, thus being better prepared to select from and build on these services when executing on production deployments that require file storage services.

LEARN MORE

Related Research

- *Key Highlights from NetApp Analyst Day 2018* (IDC #US43815718, May 2018)

- *Shifts in Storage Pricing Models: Can OEMs Transform Hardware into Services* (IDC #US43181817, February 2018)
- IDC's Worldwide Storage Software and Cloud Services QView

Synopsis

This IDC Perspective provides an overview of the growth of file services in the public cloud environment and the types of file storage services that are available to public cloud customers.

"The past few years have seen strong growth in both enterprise IT applications and new performance sensitive applications being deployed by customers on public cloud. Storage requirements for these applications often prioritize factors like on-premises consistency and latency performance, over the massive scalability available with object storage," said Deepak Mohan, research director, Public Cloud Infrastructure as a Service. "The growth of file storage services available in public cloud is a response to this customer need, and shared file storage will soon become an integral part of all major public cloud environments."

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com