

Modern File Services Enabling Enterprise Cloud Transformation

by DCIG Lead Analyst, Storage, Ken Clipperton

File-based workloads are at the heart of innovation and of collaborative workflows.



COMPANY

Nasuni Corporation
One Marina Park Drive
6th Floor
Boston, MA 02210
(857) 444-8500

www.nasuni.com

INDUSTRY

Information Technology

NEW DIMENSIONS OF SCALABILITY

- Volume
- Variety
- Velocity
- Veracity

DRIVERS FOR CLOUD TRANSFORMATION

- Access cloud-native workloads
- New enterprise applications
- Scalable storage and compute
- Multi-location access (ROBO)
- Distributed workforce (SOHO)
- Disaster Recovery/Business Continuity
- Simplified management

MULTIPLE APPROACHES TO HYBRID CLOUD FILE SERVICES

- Colocation facilities near public cloud providers
- Traditional filers in public cloud data centers
- Cloud-native file services
- Edge caching


File-based workloads are at the heart of innovation and of collaborative workflows in many enterprises. Enterprises increasingly store many petabytes of unstructured file data. They frequently need to process that file data in the cloud and to access and collaborate on that data from many different locations.

Public cloud providers offer some file services of their own yet have welcomed other file service providers into their clouds. They do this because the third-party file services offer additional capabilities enterprises require in order to move workloads to the cloud.

New Workloads are Driving Demand

Many enterprises are now running workloads that until recently were the domain of academic and scientific high-performance computing (HPC) environments. Many of these workloads rely on file-based storage. These include video content creation and management, surveillance, AI/ML, autonomous vehicle technology, radiology, oil and gas exploration, and more.

New Dimensions of Scalability for Enterprise File Services



As a result of these new workloads, enterprises are capturing file data from more sources and in a greater variety of file sizes than in the past, creating new dimensions of scalability for enterprise file services. These include:

- Volume (scale of data, stored capacity)
- Variety (capturing, storing, and processing data from many more sources with varied characteristics)
- Velocity (ingest rate, analysis of data, overall bandwidth, and bandwidth per process)
- Veracity (accuracy, but also concerns such as “Is the file I am using the most recent version?”)

Capacity. A growing list of enterprises are storing nearly an exabyte of data. Many organizations are deploying large numbers of surveillance cameras and higher-resolution cameras. As a result, it is now common for a business to store more than a petabyte of video surveillance data alone.

On a recent DCIG call, a client described how they are now using artificial intelligence and machine learning to significantly enhance service to customers. He also

mentioned that his next meeting was with his CFO to discuss buying more storage.

File counts. A programmer at another enterprise recently described a current project to present telecommunications services bills to consumers on demand via the web rather than mailing the invoices. What he had proposed as an exploratory proof-of-concept had somehow become one of the top priorities for the business, with executive team visibility. The project required the business to store millions of telecommunications services bills each month as PDFs and make them accessible via the web.

File size. Some applications, including geoseismic analysis, create individual files larger than 20 terabytes in size and data sets of more than 200 terabytes. As enterprises deploy these new workloads and expanded usage scenarios, many are discovering that their legacy filers no longer meet all their needs.

“Enterprises are capturing file data from more sources and in a greater variety of file sizes than in the past, creating new dimensions of scalability for enterprise file services.”

Drivers of Cloud Transformation

For a variety of reasons, businesses want to process file data in the cloud or access and collaborate on that data from many different locations. Public cloud providers offer some file services of their own, but have welcomed other file service providers into their clouds because those providers help move more workloads into their clouds.

The drivers of cloud transformation include:

- Access cloud-native workloads
- New enterprise applications
- Scalable storage and compute
- Multi-location access (ROBO)
- Distributed workforce (SOHO)
- Disaster Recovery/Business Continuity
- Simplified management

The value of being able to collaborate on corporate files from anywhere is evident to everyone who suddenly

found themselves working from home due to the COVID-19 pandemic; or relocating to avoid the fires, floods and other natural disasters.

Multiple Reasons Enterprises Need Modern File Storage

Workload migration to the cloud

- Artificial intelligence and machine learning (AI/ML)
- File sizes, file counts, and file volumes that exceed the limits of legacy infrastructure
- Distributed operations
- Integrating new data streams into business processes
- The need to simplify infrastructure management and drive out costs while providing services to a distributed workforce

Multiple Approaches to Hybrid Cloud File Services

- Deploy or rent capacity on traditional filers in colocation facilities near the public cloud providers to provide hybrid/multi-cloud.
- Deploy traditional filers in public cloud data centers
- Cloud-native file services running directly on public cloud infrastructure
- Edge caching in front of traditional filers or cloud-native file services

Hybrid Cloud File Services Providers

Nasuni's cloud-native global file system, UniFS, unifies the capabilities of enterprise NAS, backup solutions, and disaster recovery infrastructure. It uses object storage as the backing store. It uses caching appliances (physical or virtual) to provide secure file-based access to an enterprise's global file repository at local network speeds.

Qumulo's software-defined hybrid file system provides high-performance file services for trillions of files in multi-petabyte data sets. This is true on-premises, natively in the public cloud, and as a system that scales across on-premises and public cloud environments. It provides the scalable data services and real-time visibility customers need to manage the ever-increasing volume, velocity and variety of data in the enterprise.

CTERA is another cloud-native solution that provides a global file system and caching at the edge. CTERA's focus is on transitioning file workloads to cloud without sacrificing security, performance, and control.

NetApp's leaders recognized that enterprise requirements were changing. To meet these emerging requirements, NetApp introduced the data fabric. Ironically, NetApp is pivoting to become a hybrid and multi-cloud data management services company rather than a seller of networked storage appliances.

NetApp recently acquired Talon FAST and now markets the solution as the NetApp Global File Cache. This product layers the Talon FAST environment on top of an existing NetApp cloud service (Azure NetApp Files,

Cloud Volumes ONTAP, Cloud Volumes Service). NetApp deploys its solution as caching software installed on physical Windows servers or Windows Server virtual machines at edge locations. Real-time central file locking is provided through a single ONTAP instance in the cloud.

Dell EMC **Isilon** (now PowerScale) has a huge installed base in the enterprise. Dell EMC's approach to Isilon hybrid cloud is to place Isilon hardware in colocation facilities near the major public cloud facilities and manage the Isilon clusters for its customers. Isilon can also tier cold data to the cloud via CloudPools, an extension of its SmartPools data tiering framework.

“Nasuni's cloud-native global file system, UniFS, unifies the capabilities of enterprise NAS, backup solutions, and disaster recovery infrastructure.”

Existing HPC solution providers are entering into the enterprise marketplace now. For example, **DDN** acquired Tintri and IntelliFlash. In doing so, DDN gained storage analytics and proactive support capabilities that enterprises expect from their infrastructure providers. DDN will likely integrate these capabilities with their HPC products, much as HPE did with the InfoSight technology it acquired with Nimble Storage.

Understanding Requirements is Foundational

Enterprises are embracing digital transformation. Many digital transformation initiatives are pushing enterprises to consider new file service infrastructures to create new value and carry them into the future. If the COVID-19 pandemic has taught us anything, it is that we do not know what the future will require of us.

Nevertheless, we must act in the present in the face of uncertainty. Speed matters. So does flexibility and optionality. If the incumbent filer solution meets your needs and offers a cloud option, then the fastest path may be a hybrid-cloud implementation with the current provider.

If individual file sizes exceed 16TB, then you will certainly need to look at new solutions architected for the cloud such as Nasuni and Qumulo. If your challenge is providing fast access and file-based collaboration to distributed locations, then a global file system with caching at edge locations may be the best fit. Depending on overall capacity requirements and the number of locations you must support, your short list may include CTERA, Nasuni, Panzura, Qumulo, or NetApp ONTAP and its Global File Cache.

Many solution providers have stepped up to help enterprises create new value from file-based workflows. If you approach the refresh of your file infrastructure with a good understanding of your requirements, you are sure to find a solution that will meet your needs and enable your enterprise cloud transformation. ■

About DCIG

DCIG, the Data Center Intelligence Group, empowers the information technology industry with actionable analysis. DCIG provides informed third-party analysis of various cloud, data protection, and data storage technologies. Learn more at www.dcig.com.



DCIG, LLC // 7511 MADISON STREET // OMAHA NE 68127 // 844.324.4552

dcig.com

© 2020 DCIG, LLC. All rights reserved. Other trademarks appearing in this document are the property of their respective owners. This DCIG report is a product of DCIG, LLC. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. Product information was compiled from both publicly available and vendor-provided resources. While DCIG has attempted to verify that product information is correct and complete, feature support can change and is subject to interpretation. All features represent the opinion of DCIG. No negative inferences should be drawn against any product or vendor not included in this report. DCIG cannot be held responsible for any errors that may appear. This report was commissioned by Acronis.

This DCIG Executive White Paper is licensed to Nasuni Corporation with perpetual, unlimited, global distribution rights.

October 2020 2