
Qumulo File Fabric extends high-performance file services to the cloud

SEPTEMBER 2017



The timing for Qumulo to extend its software-defined scalable file services to the cloud could not be better as public cloud utilization continues to grow at a phenomenal rate. Infrastructure spending on the public and private cloud is growing at double-digit rates while spending on traditional, non-cloud, IT infrastructure continues to decline and within a few short years will represent less than 50% of the entire infrastructure market. This trend is not surprising and has been widely predicted for several years. The surprising element now is how strong the momentum has become toward public cloud adoption, and the question is where the long-term equilibrium point will be between public clouds and on-premises infrastructure.

AWS was a pioneer in public cloud storage services when it introduced S3 (Simple Storage Service) over ten years ago. The approach of public cloud vendors has been to offer storage services at cut-rate pricing in what we call the “Hotel California” strategy – once they have your data, it can *never* leave. Recently, we have been hearing increased grumbling from customers that they are very concerned about losing the option to change infrastructure vendors and the resulting reduction in competition. In response to this, Taneja Group initiated multiple public and hybrid cloud research studies to gain insight on what storage services are needed across heterogenous cloud infrastructures. What we found is that IT practitioners are not only concerned about data security in the cloud; they are concerned about vendor lock-in created by the lack of data mobility between on-premises and public cloud infrastructures. Another surprising element we found is that IT practitioners predominately want file services across clouds and that object storage such as AWS S3 cannot meet their future cloud storage needs. This is actually not that surprising, as our research showed that many applications that businesses want to move to the cloud (to benefit from a highly dynamic compute environment) still rely on high-performance files access.

Enter Qumulo File Fabric (QF2). QF2 is a modern, highly scalable file storage system that runs in the data center and now in the public cloud. Unlike legacy scale-out NAS products, QF2 provides capacity for billions of files, closely matching the scale that could only previously be achieved with object storage solutions, but with the benefit of supporting file access protocols. Qumulo’s modern SDS, flash-first approach allows it to provide a very high-performance file storage system that can cover a wide variety of workloads. Its built-in, real-time analytics let administrators easily manage data no matter how large the footprint or where it is globally located. Continuous replication enables data to move where and when it’s required depending on business need. Qumulo refers to this unmatched file scalability and performance as universal-scale file storage.

Qumulo, founded in 2012, is rapidly growing its market presence and we recently validated their very high customer satisfaction and product capability through an extensive interview process with several customers. Qumulo recently extended their go-to-market ecosystem support through a partnership with Hewlett Packard Enterprise (HPE). Now with the launch of QF2 and support for AWS, we expect Qumulo to continue its rapid rise as a leading provider of file services with universal scale. They are also well positioned to capture a significant share of the emerging multi-cloud storage market. We found many companies still prefer file access and there are plenty of reasons why scalable file will continue to grow and compete effectively versus object storage centric architectures.

The emerging need for multi-cloud storage services

To gather data and develop insights regarding plans for public cloud use, Taneja Group initiated a primary research study earlier this year. We surveyed 400+ IT decision makers and practitioners representing a wide range of industries and business sizes to understand current and planned deployments of applications to the public cloud.

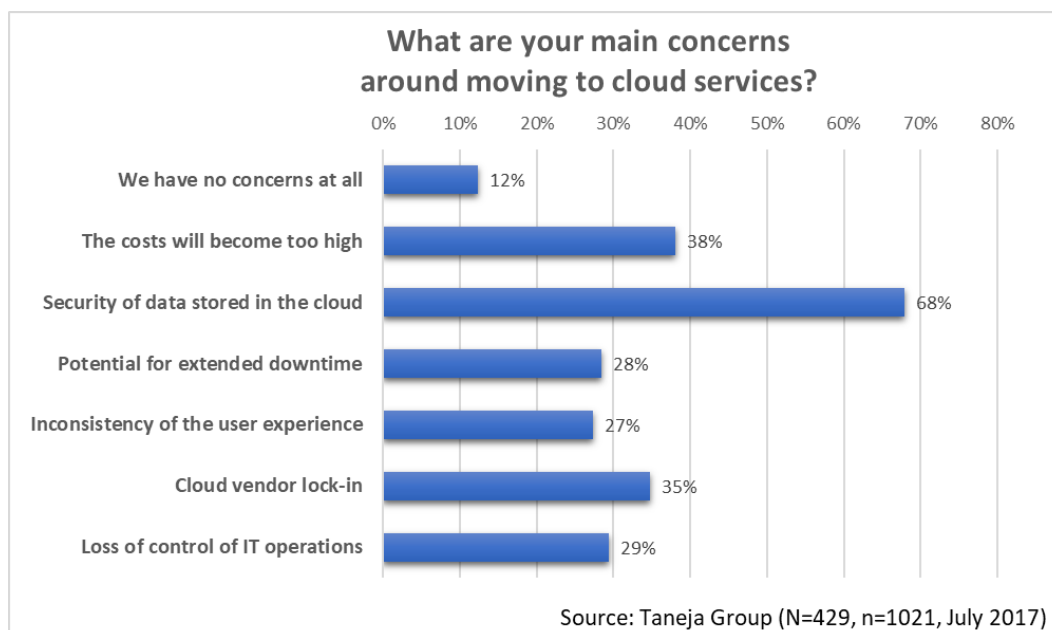
Specifically, we wanted to understand the need for an emerging set of storage products we call multi-cloud primary storage. These products provide their data services across more than one public cloud simultaneously.

Taneja Group Definition of Multi-cloud Primary Storage

Multi-cloud primary storage provides primary data services that can operate simultaneously across multiple heterogeneous cloud environments, and where the compute and applications can be collocated with the data services. The primary data storage services must support at least one large public cloud vendor (e.g. AWS, Azure, and Google). As an example, “multi-cloud” might include a private and public cloud, multiple public clouds, or any combination of these.

Obstacles to public cloud adoption

Increasing cloud costs and vendor lock-in are top concerns just behind the ever-present need for data security. One would argue increased costs and vendor lock-in concerns are related, as a lack of vendor choice often leads to higher prices. The graph below details these findings:



Multi-cloud storage solutions can help address vendor lock-in and data security concerns. Look for end-to-end security to be a major differentiator in the multi-cloud storage market going forward. Only 12% of respondents said they have no apprehensions about moving to cloud services, indicating that there is a sizeable opportunity for providers to offer products that relieve the concerns of the remainder of respondents to our survey.

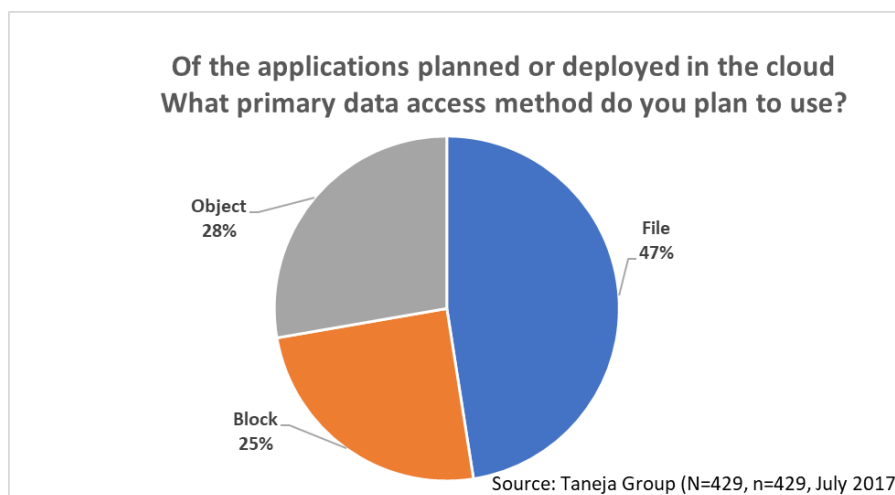
Multi-cloud benefits

The following are key benefits multi-cloud storage products and services promise to deliver:

- **Data portability between heterogeneous clouds.** The ability to shift data between public clouds as business needs require. This portability mitigates vendor lock-in and enables customers to more easily meet new regulatory compliance demands.
- **Ease of lifting and shifting applications between environments.** Certifying applications against a common set of data services makes shifting an application to a new environment easier to manage. Most multi-cloud storage enables application testing on-premises and in the cloud.
- **Control of end-to-end data security.** Multi-cloud storage should provide better visibility of data location and enable seamless security with access control and encryption across disparate cloud environments.
- **Enhanced data availability and disaster recovery options.** Although public clouds boast geographically dispersed availability zones, they often come with increased costs, and there are still occasional outages. Multi-cloud storage can allow for diversity of cloud providers for both continuous availability purposes and disaster recovery.
- **Consistent set of enterprise data services collocated with apps and compute resources.** Features such as replication, data locality, snapshots, and capacity optimization differ among cloud providers. Multi-cloud storage should offer a more consistent quality of service across clouds through a common feature set.
- **Bridging data services between private and public clouds.** Multi-cloud primary storage enables a more seamless transition for customers seeking hybrid cloud services.

File access still matters

The following figure shows the results when respondents were asked what data access method they prefer most for the applications they plan to move to the public cloud. File access led the way at 47%, followed by object at 28% and block at 25%.



Instead of redesigning legacy scale-out NAS products, most of the industry continues to focus on object storage for scale-out cloud workloads. Unfortunately, object-based storage has failed to provide the high-performance, enterprise-grade, POSIX-compliant file access that thousands of legacy applications require. It also fails to provide a performance level that can meet the requirements of many big data workloads like media and entertainment, life sciences and commercial HPC. Object storage companies are trying to address these issues by adding file

gateway accelerators in front of their object-based backend. But this approach adds another layer of complexity, leaving the door open for a new architecture that is focused on modernizing enterprise capable scalable file services.

Qumulo File Fabric is an ideal solution for very high-performance, large capacity multi-cloud storage workloads

Qumulo File Fabric (QF2) provides enterprise-grade, scalable file services running between public clouds and on-premises infrastructure. Qumulo universal-scale file systems are based on a distributed architecture where each node has both a software-defined storage (SDS) component and underlying direct attached storage running in the cloud or on-premises using commodity x86 hardware. QF2 has been developed from the ground up to handle both small file and large file workloads while scaling to tens of billions of files and hundreds of petabytes of capacity. Qumulo's flash-first hybrid architecture enables QF2 to provide low-latency write operations using standard SSD flash instead of the more typical battery-backed NVRAM approach. This keeps the system performant while running in low-cost compute environments. This flash-first hybrid architecture also allows QF2 to implement enhanced metadata aggregation techniques that enable real-time analytics on a scale not seen before. All metadata operations execute in the flash layer enabling additional performance benefits.

In the past year, we interviewed several Qumulo customers to validate their solutions in production. Now, with QF2 extending to public cloud, we see QF2 providing the following key benefits for multi-cloud storage workloads.

- **Billion-file scale:** With QF2, you can use any mix of large and small files and store as many files as you need. There is no practical limit with Qumulo's advanced file-system technology. Many Qumulo customers have data footprints in excess of a billion files.
- **Performance for mixed file sizes:** Qumulo QF2 has been implemented from the start to handle mixed file size workloads. Customers selected Qumulo over object-storage and legacy NAS vendors because of this capability. We found many of the Qumulo customers we talked to needed both small and large file size performance simultaneously.
- **Legacy application support:** Many businesses are not in a position to rewrite all their applications to change from a file access to an object API access method. Therefore, POSIX-compliant file access is a critical feature, which often rules out object storage alternatives. Performance is also an issue for many workloads, especially those that rely on high-performance file access.
- **Real-time analytics and control:** Qumulo provides real-time insights and control on a massive scale. Most scale-out NAS products today suffer from the metadata death march. Walking file trees just to see what and who owns a file is untenable at scale. Unlike legacy scale-out NAS vendors, Qumulo designed a flexible metadata approach into QF2 that can handle massive amounts of files in a distributed and non-disruptive fashion (Qumulo demonstrated 10 billion files using just their smallest footprint out of four appliances.) This capability enables administrators to pinpoint problems and instantly manage how storage is used. For example, with QF2, administrators can create directory-level quotas that take effect instantly, and snapshots that affect only a single subdirectory tree.
- **Ease-of-use:** The highly simplified interface does not require management specialists, which allows companies to use generalist IT resources for storage management. Even with large clusters, management of Qumulo systems remains a simple process. Customers and ecosystem partners get access to any feature or administrative task through self-describing REST APIs.
- **Modular scalability with subscription-based pricing:** Because Qumulo is based on scale-out SDS technology customers can take advantage of modular scalability, which lets them right-size

from the start and pay as they grow. QF2 is also available by subscription with a simple pricing structure. Subscription pricing is an important attribute and matches the public cloud pricing approach.

- **Cloud-based monitoring and trends:** QF2 includes proactive cloud-based monitoring that prevents problems before they happen. Access to historical trends helps lower costs and optimize workflows for best use of your storage investment. Cloud-based monitoring capability is a must for those businesses that might want to deploy Qumulo solely in the cloud.
- **Consistent global access to your data:** QF2 can be implemented across multiple operating environments that are connected using replication. Wherever you deploy QF2, it works the same way. Storage-as-a-service enables equally strong support for burst-to-cloud workloads, lift-and-shift application workloads, and cloud-native apps.
- **Continuous innovations through a cloud-like development approach:** Future Qumulo features are included in the QF2 all-inclusive subscription pricing model. Qumulo releases new fully-tested, production-ready features and performance improvements to customers frequently as part of their modern agile development approach.
- **No infrastructure provider lock-in:** Because QF2 is infrastructure independent, you have the flexibility to choose your operating environment. You can use standard, commodity hardware provided by Qumulo, hardware provided by partners such as HPE, or run in the public cloud with partners such as AWS. There's no infrastructure vendor lock-in, which lowers costs and creates operational flexibility.

Taneja Group Conclusion

The multi-cloud primary storage market is newly emerging and will evolve quickly in the next few years. The key question that will dictate the long-term success of this market is whether customers will value the business benefits that a consistent multi-cloud data service delivers. Qumulo customer examples are quickly emerging for why businesses might want to use multi-cloud primary storage technology to shift data from one cloud to another.

In one example, FuseFX, a visual effects company is using QF2 on-premises and in the cloud for video effects rendering. "We are building a fully orchestrated visual effects rendering solution that spans our on-premises data center and AWS," said Jason Fotter, Co-Founder and CTO at FuseFX. "We now have a QF2 cluster on AWS and in our data center creating a unified fabric that enables us to share file data between these two operating environments, maintain workflow consistency, and meet the high performance requirements for heavy compute workloads in the cloud."

In another example, the University of Washington is using the cloud to ramp GPU-based massively parallel computing to operate on their very large file-based data sets. QF2 enables them to move file-based data set to AWS and back to avoid buying expensive GPU infrastructure on-premises. Tyrone Grandison, CIO, Institute for Health Metrics and Evaluation at the University of Washington said "Ramping up GPU infrastructure on premises in our data center is far too expensive and complex. QF2 allows us to move file-based data sets to a QF2 cluster on AWS, complete our analysis, and move the artifact back to our on-premises QF2 storage cluster, saving us time and money. The flexibility for us to move our file-based data where we need it to be is something that nobody else in the market can provide at scale."

We expect the multi-cloud primary storage market to grow rapidly in the next couple of years as the public cloud providers just get bigger and bigger, leaving many businesses feeling like their data is being held hostage. We expect Qumulo to be on the forefront of this rapidly growing market category. Focusing on universal-scale file storage will give them a unique opportunity in a market where file access still matters a lot. Based on our review of the product and conversations with

Qumulo customers, we think they've succeeded in creating a product that will satisfy the needs of many multi-cloud primary storage users.

NOTICE: The information and product recommendations made by Taneja Group are based upon public information and sources and may also include personal opinions both of Taneja Group and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. Taneja Group, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors that may appear in this document.