



Technical Research Study



Going Cloud-Native: Tips for Choosing the Right Cloud-Based Storage Solution for Your Business

Prowess Consulting looked at two competitors' offerings to see which one provided a more seamless, scalable, and performant cloud-native NAS solution.

Executive Summary

Businesses are storing some of their resources in the public cloud for its scalability, flexibility, security, and manageability. While network-attached storage (NAS) is a familiar, trusted storage technology, public cloud-based versions are still emerging, so their feature sets can vary widely from one offering to the next.

We started our research by looking at the challenges that businesses face when migrating their data into the cloud. These issues led to the question, "What storage features do I need to be successful in my cloud journey?"

We established a list of must-have features that we felt addressed this question. This list helped us uncover two standouts in the cloud-native NAS market: Dell™ APEX File Storage for AWS and NetApp® Cloud Volumes ONTAP® high availability (HA) configuration.

We also did a side-by-side comparison to identify any differences between the two solutions. Based on its superior migration ease, bandwidth, and scalability, we concluded that Dell APEX File Storage for AWS comes out a winner over NetApp Cloud Volumes ONTAP HA configuration.

Why So Many Companies Are Moving Their Data into the Cloud:

“Clearly, hybrid cloud is the new normal.”¹

Many organizations are storing some of their data in the cloud these days, driving phenomenal growth in hybrid and multicloud computing environments. And in some cases this requires a data migration. According to the financial advisory firm Deloitte, the cloud computing market has exhibited triple-digit growth in the last few years.¹ This expansion to hybrid environments is being fueled by organizations seeking improvements across a variety of business objectives, such as employee productivity, market agility, resource allocation, and business continuity.

Ranked by IT decision makers (ITDMs) near the top of the many services moving into the cloud is network storage. Just as you would purchase software as a service (SaaS), you can also purchase infrastructure as a service (IaaS). According to a 2022 survey conducted by Foundry, ITDMs listed IaaS as one of their top-five priorities in their cloud-planning strategies.² However, nearly all respondents also reported running into obstacles during the execution of their plans. Nearly one-quarter of these challenges were associated with the high cost of data migration, and another one-third were attributed to a shortage of cloud-management skills.²


Intrigued by the potential of combining high-throughput NAS with the availability and flexibility of the public cloud, Prowess Consulting decided to dig deeper into this new usage for a trusted storage technology. Our report examines the cloud-native storage offerings that are currently available, and we take a look at their benefits and pitfalls.

What We Wanted from an Enterprise Cloud Storage Solution

Unlike traditional NAS deployed on premises, cloud-native NAS is hosted in a distributed cloud environment. Cloud-native NAS facilitates your ability to scale storage capacity as needed, making it ideal for ever-shifting storage demands, such as an unexpected influx of large data volumes. Using NAS as a cloud-based IaaS lets you pay for storage as you grow.



Data mobility



Seamless integration



Consistent user experience



High-performance cloud instances



Multiple file protocols



Scalable clusters

Figure 1 | To help you experience a successful journey into the cloud, we recommend looking for cloud-storage solutions with these essential features

Cloud-based data storage is a distributed, highly available, and fault-tolerant environment that includes built-in redundancy and failover capabilities. Authorized users can easily access cloud resources and applications from anywhere, using a browser. You can also use cloud-native NAS as part of a hybrid environment that keeps certain data on premises for data sovereignty compliance.³

When using on-premises, hybrid, and multicloud storage from the same vendor, a consistent user experience helps simplify your journey into the cloud. For example, having the same user location information (ULI), command-line interface (CLI), and application programming interface (API) for cloud and on-premises environments helps streamline management. These common interfaces allow IT staff to manage your cloud storage without additional training, because they are using the same familiar controls that they already use on premises.

Another must-have feature is an “easy button” for data mobility. In other words, a tool that smoothly transfers your data while maintaining its accuracy and integrity. If you are seeking ways to lower the total cost of ownership (TCO) for data storage, you want data mobility that preserves existing data services so that applications are immediately available after the move into the public cloud. Native data replication, which helps eliminate errors introduced during migration, can significantly speed up the time-consuming chore of revalidating your migrated data in order to make it usable.⁴

We wanted a high level of performance from our cloud-native storage. The right file structure is critical for collecting and storing bulky, unstructured datasets, such as those processed by artificial intelligence (AI), big data analytics, streaming content services, high-throughput imaging, real-time language processing, write-intensive content creation, and other modern workloads. We prefer a scale-out architecture because it is less likely to experience latency bottlenecks than scale-up storage.⁵

A common question that arises when you move data into the public cloud is, “What about security?” Fortunately, most public cloud storage providers provide enhanced data and system security as part of their service. That said, data protection is a security feature you will want from a storage solution. Thanks to their greater redundancy and higher failure tolerance, scale-out architectures often offer stronger data protection than scale-up architectures.

Many storage vendors offer additional valuable features. To help ensure success in your public cloud migration, we recommend choosing vendors with a long history of storage technology expertise, including 24/7 tech support and regular rollouts of new features and upgrades. If you are processing high volumes of write data, which wears out drives more quickly than read operations, then you want public cloud servers that are configured with high-reliability, long-endurance drives. While these features may not be as critical as seamless integration, easy migration, scalability, high performance, and a consistent user experience, they are certainly worth considering when planning your cloud migration.

Our List of Essential Features

For this study, we focused on NAS designed for enterprises and small to medium-sized businesses (SMBs). This meant limiting our selections to high-performance storage that is capable of handling massive unstructured datasets and heavy analytics workloads. We also passed over single-user or personal cloud solutions (PCS), which lack the capacity, security, scalability, performance, and management required by most businesses.

For a more uniform side-by-side comparison, we wanted to limit our selections to the same public cloud-computing platform. We also wanted our findings to be useful to as many cloud users as possible. We chose Amazon Web Services® (AWS®) which represents one-third of all public cloud infrastructure providers.⁶

Within this business-class target, we looked for solutions that met our list of must-have features:

- Easy data mobility that includes native data replication within the same storage architectures
- Seamless integration with hybrid and multicloud environments
- A consistent user experience across on-premises, hybrid, and cloud environments
- Support for multiple file protocols
- Scalable server clusters
- High performance in the cloud

We determined that there are only a few NAS offerings that could fully support public cloud and hybrid environments. As expected with any emerging technology, many offerings marketed as public cloud storage solutions lacked one or more of our essential criteria. For example, we removed from consideration solutions that do not include integrated administrative controls and that instead rely on service provider management.

After sifting through the available choices, we found two public cloud-based NAS solutions that measured up on our list of must-have enterprise-grade features.

A Side-by-Side Comparison: Dell™ APEX File Storage for AWS and NetApp® Cloud Volume ONTAP® high availability (HA) configuration

Dell APEX File Storage for AWS is a suite of high-performance, cloud-based NAS storage solutions. The scale-out architecture supports accelerated file workloads, large-scale data workloads, and high-volume digital archives.

NetApp Cloud Volumes ONTAP high availability (HA) configuration offers subscription-based enterprise NAS solutions for hybrid and multicloud environments. This high-performance, high-availability storage is built on a scale-up architecture.

Table 1 | Cloud-native NAS features for the two competitors' offerings

	Dell™ APEX File Storage for AWS	NetApp® Cloud Volumes ONTAP® high availability (HA) configuration
Hybrid/multicloud support	Yes	Yes
Data mobility with native replication	Yes, within the same architecture	Yes, within the same architecture
NAS architecture	Scale-out	Scale-up
Maximum nodes per server	6 nodes/cluster	2 nodes/cluster
Policy-based load balancing	Yes, Dell SmartConnect	No
Maximum write performance	Up to 3,930 MB/s ⁷	Up to 897 MB/s ⁸

In our side-by-side comparison of these two public cloud-native storage solutions, we found a number of similarities. Both solutions offer seamless integration with existing infrastructures and support for hybrid and multicloud environments. Being fully public cloud-compliant, they both deliver single sign-on (SSO) access, data protection, and storage redundancy.

Easy-Button Data Mobility

Migrating data, applications, and systems into the public cloud is a multi-step process that requires careful planning and plenty of manual intervention. The best data-mobility tool helps streamline the process, while still maintaining the integrity and accuracy of all your data, including user permissions, access control lists (ACLs), hashing algorithms, and other metadata.

Data validation is a time-consuming but necessary chore that helps ensure that your data is as intact and error-free as possible before moving it into the public cloud. After migration, you typically must re-validate the data to check for and repair corrupted data and errors introduced during the transfer process.⁹ The last thing you want to discover is that the files you moved into the public cloud are no longer accessible because their access rights were broken or lost during the transfer.

Here is an example of how migration that uses native data replication helps minimize application downtime for your users. Home directories are repositories containing a default directory for individual user sign-ins, in addition to their personal files, directories, and programs. These repositories can be massive if you have lots of users. Migrating this data into the public cloud is potentially a massive task involving months of planning, validation, migration, and re-validation.

An “easy button” migration tool with native data replication lets you continue to use data services that you’ve already invested in by ensuring that your data and metadata stay intact and errorfree during the transfer.

Both Dell APEX File Storage for AWS and NetApp Cloud Volumes ONTAP HA configuration provide native data replication for data being moved within the same on-premises or public cloud architecture. This capability can reduce your post-migration re-validation time to nearly zero.

Where the Dell APEX File Storage for AWS comes out ahead is in its high-performance data transfer, which can slash data-transfer times dramatically. For equivalent volumes of data, the Dell solution’s higher write-bandwidth helps migrate your data quickly.^{7,8}

Scale-Out Architecture

One of the major benefits of cloud-based storage is easy scalability. NetApp Cloud Volumes ONTAP HA configuration uses a highly available scale-up architecture; however, it has a maximum limit of two storage nodes per server cluster. A two-node cluster also lowers the ceiling of your storage resiliency, which has the potential to disastrously impact data availability. Should one node fail, you will most likely notice a performance drop as all the cluster traffic gets redirected to the one remaining node. And should both nodes fail, the cluster shuts down and you could lose access to your data.

The Dell APEX File Storage for AWS scale-out architecture supports up to six nodes per cluster. If a smaller server footprint is important, you can start with four nodes. As your storage demands increase, you can quickly scale out storage capacity by adding new nodes to your cluster, up to six. If one node fails, the Dell cluster uses policy-based load balancing to distribute workloads across the remaining five nodes in the cluster, ensuring data availability.

The Dell cluster protects your data from multiple node failures, something the two-node NetApp cluster cannot do. Dell APEX File Storage for AWS clusters are protected at a +2N level, meaning the cluster is still available if one or two nodes fail. You may see a performance drop; however, and most importantly, your data stays accessible because the cluster is still active.

More nodes per cluster also means more bandwidth for write-intensive applications and high-volume datasets. For example, creative professionals working in media and entertainment (M&E) rely on demanding write-intensive workloads for high-definition content creation. If you are using generative design in the public cloud to create and edit backdrops for film and video, you will want plenty of bandwidth for your animation and rendering workflow.¹⁰ Streaming video transcoding, which is a write-intensive operation performed in real time, is particularly sensitive to traffic bottlenecks.

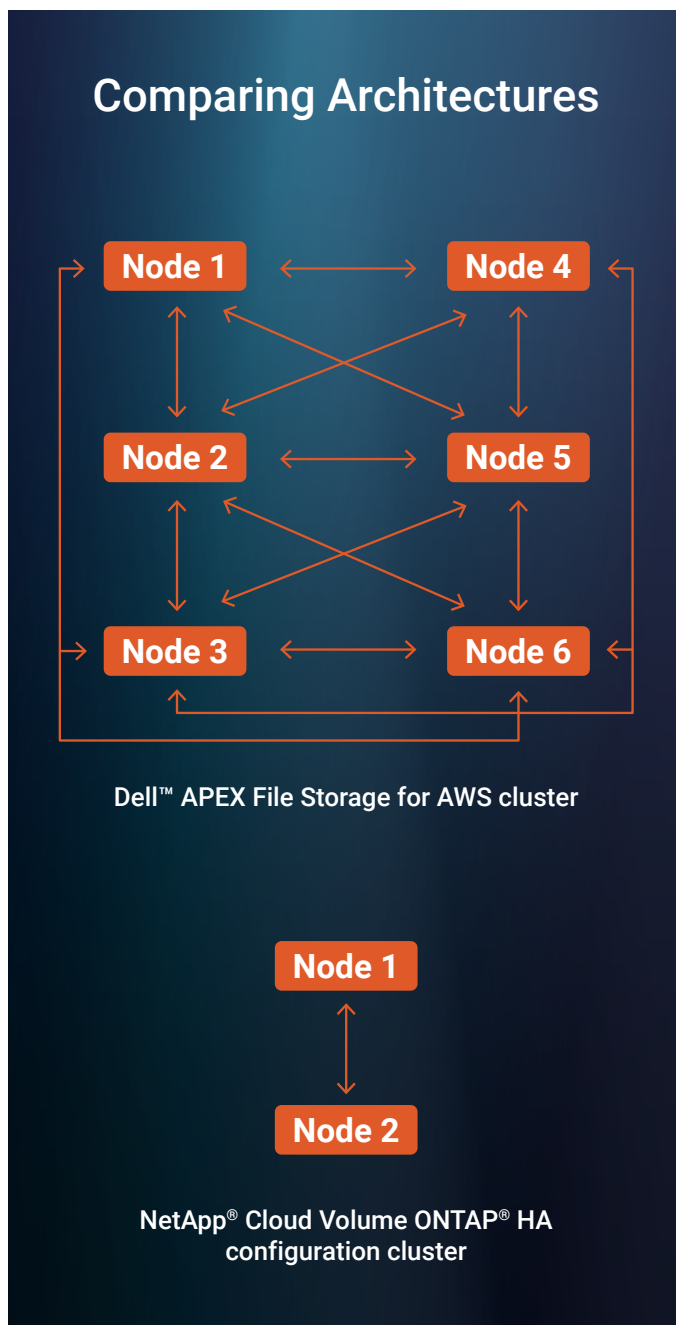


Figure 2 | The Dell™ APEX File Storage for AWS cluster supports a maximum of six storage nodes, whereas the NetApp® Cloud Volumes ONTAP® HA configuration cluster only supports up to two storage nodes^{7,8}

High-Performance Storage for Demanding Workflows

The best cloud storage solutions not only integrate seamlessly and migrate easily, but also perform flawlessly. Your public cloud-native experience should look and feel very similar to workloads running on premises. To get the best performance for big data analytics, media streaming, and AI workflows, your public cloud storage needs to support high volumes of unstructured data. In addition to supporting high data throughputs, your public cloud storage should support multiple file protocols, such as network file system (NFS), server message block (SMB), and Amazon Simple Storage Service (Amazon S3) object protocols.

The NetApp solution falls short in the performance category when compared to the Dell solution. Figure 3 shows that Dell APEX File Storage for AWS delivers 4.3x higher bandwidth per server cluster than NetApp Cloud Volumes ONTAP HA configuration.

With a scale-up architecture limited to a maximum of two nodes per pair (cluster), NetApp Cloud Volumes ONTAP HA configuration storage quickly reaches its bandwidth limit when working with high data volumes. If you need higher bandwidth, you need to switch to a larger instance size. However, opening a new instance means starting your workflow over from the beginning. Once you have reached the maximum instance size, the only way to increase performance is by deploying another cluster.

The Dell APEX File Storage for AWS scale-out architecture allows you to add more nodes—up to six nodes in a cluster—for four times more available bandwidth before you max out the cluster. This higher bandwidth makes the Dell solution a better choice for supporting the automatic scaling capabilities of Amazon S3 storage, which has an upper limit of thousands of responses per second.¹¹

The many advantages of the Dell solution are underscored by the technology consultant Gartner, which named Dell™ PowerScale™ storage software as an industry leader in object and file storage for the last seven years.¹²




Figure 3 | Dell™ APEX File Storage for AWS delivers up to 4.3x higher cluster bandwidth (3,930 MB/s normalized to 4.3) than NetApp® Cloud Volumes ONTAP® HA configuration (897 MB/s normalized to 1.0) for 100 percent sequential write workloads^{7,8}

Conclusion

To help ensure that your migration into the public cloud is a painless experience, we recommend choosing a public cloud storage solution with the following capabilities: easy data mobility with native data replication; seamless on-premises/cloud integration; simplified and consistent management; high performance; multiprotocol data support; and, of course, high scalability.

If you are planning to extend your storage into a public cloud, get a head start on your journey by using the findings in this report. We uncovered two public cloud-native NAS solutions with clear advantages over other public cloud-storage offerings. The NetApp Cloud Volumes ONTAP HA configuration solution checks off nearly all the boxes in our must-have checklist. However, with its faster migration, more nodes per cluster, and higher bandwidth, the Dell APEX File Storage for AWS storage solution comes out ahead in the side-by-side comparison.

Learn more

Learn more about [Dell APEX File Storage for AWS](#) for cloud environments.

Learn more about [NetApp Cloud Volumes ONTAP HA configuration](#) cloud storage infrastructures.

See more research reports by [Prowess Consulting](#).

¹ Deloitte. "The cloud migration forecast: Cloudy with a chance of clouds." Accessed April 2023.

www2.deloitte.com/xe/en/insights/industry/technology/technology-media-and-telecom-predictions/2021/cloud-migration-trends-and-forecast.html

² Foundry. "Cloud Computing Study: Executive summary outlining the 2022 research findings." Accessed April 2023.

https://1624046.fs1.hubspotusercontent-na1.nethubfs/1624046/FINAL_R-ES_CloudComputing22_040122.pdf

³ Imperva. "What is Data Localization?" Accessed April 2023. www.imperva.com/learn/data-security/data-localization.

⁴ Built In. "What Is Data Validation?" March 2023. <https://builtin.com/data-science/data-validation>.

⁵ ServerWatch. "Scale Up vs Scale Out: Data Center Infrastructure." March 2022. www.serverwatch.com/storage/scale-up-vs-scale-out/.

⁶ Synergy Research Group. "Huge Cloud Market Still Growing at 34% Per Year; Amazon, Microsoft & Google Now Account for 65% of the Total." April 2022.

www.srgresearch.com/articles/huge-cloud-market-is-still-growing-at-34-per-year-amazon-microsoft-and-google-now-account-for-65-of-all-cloud-revenues

⁷ Dell Technologies. "Introduction to APEX File Storage for AWS." May 2023.

<https://infohub.delltechnologies.com/t/introduction-to-apex-file-storage-for-aws>

⁸ NetApp. "Technical Report: Performance Characterization of NetApp Cloud Volumes ONTAP for Amazon Web Services." May 2020.

www.netapp.com/pdf.html?item=/media/9088-tr4383.pdf

⁹ FirstEigen. "How to Ensure Data Integrity During Cloud Migrations." Accessed April 2023. <https://firsteigen.com/blog/how-to-ensure-data-integrity-during-cloud-migrations/>

¹⁰ Forbes. "Cuebric: Generative AI Comes To Hollywood." March 2023. www.forbes.com/sites/tomdavenport/2023/03/13/cuebric-generative-ai-comes-to-hollywood/?sh=231c3263174b

¹¹ Amazon Web Services. "Best practices design patterns: optimizing Amazon S3 performance."

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/optimizing-performance.html>

¹² Dell Technologies. "Distributed File Systems and Object Storage Leadership." November 2022.

www.dell.com/en-us/blog/distributed-file-systems-and-object-storage-leadership/



The analysis in this document was done by Prowess Consulting and commissioned by Dell Technologies. Prowess Consulting and the Prowess logo are trademarks of Prowess Consulting, LLC. Copyright © 2023 Prowess Consulting, LLC. All rights reserved. Other trademarks are the property of their respective owners.