ActualTech Media

NetApp®

Take Amazon EBS from enterprise-grade to enterprise-great

Teren Bryson

CONTENTS

IN THIS PAPER

Block storage is a valuable cloud storage paradigm, necessary for many types of workloads. But cloud block storage offerings, although high performing, are often basic, and they lack features found in on-premises solutions.

Insufficient features across multiple cloud offerings further muddy the block storage waters, and the price-to-performance gap becomes untenable at scale. This brief explains how NetApp® Cloud Volumes ONTAP® brings feature parity to block storage in the cloud, reducing costs and optimizing usage.

Often, developers favor block storage for performance-intensive compute situations that require fast, efficient, and reliable data transportation. Examples include mission-critical Oracle, SQL Server, and SAP HANA databases; Microsoft Exchange; virtualization vendors like VMware; and anything that requires consistent IOPS. Although public cloud block storage offerings are high performing, they're general purpose and often lack the reliability and enterprise-grade features found in the on-premises data center. This shortcoming can cause impassable challenges to meeting hybrid and cloud mandates.

Insufficient features across leading public cloud offerings further muddy the block storage waters, and the price/performance gap becomes untenable at scale. When researching cloud storage solutions, organizations find that they can't get the same performance and reliability from the various public cloud offerings as they have in their own data center. Cloud storage is maturing all the time, but it still lags behind the on-premises enterprise-grade storage in terms of features and reliability.

Who needs block storage?

Block storage, sometimes referred to as block-level storage, is a technology that is used to store data files on Storage Area Networks (SANs) or cloud-based storage environments. Data is broken up into blocks that are then stored as separate pieces, each with a unique identifier. Block storage places those blocks of data wherever it's most efficient. Block storage also decouples data from user environments, allowing said data to be spread across multiple environments. This approach creates multiple paths to the data and allows users to retrieve it quickly.

Block storage is commonly used for enterprise-grade I/O-intensive applications such as databases, Hadoop, Java, PHP, .NET server-side processing applications, containerized applications, and media workflows in the cloud. It's also prominent in production-level, submillisecond, low-latency workloads, such as those in financial apps and high-performance computing (HPC), where shaving off every microsecond is important. All these workloads require consistent, predictable performance, which is where block storage can excel.

Limits of public cloud block storage

All three major cloud providers—Amazon, Google, and Microsoft—offer block storage, but their features vary, and consistency across the megaclouds is challenging at best. Moreover, no public cloud provider's block storage offers as robust a package as on-premises solutions have for years.

Storage networks in the cloud come with their own sets of challenges. These challenges are different from the ones that storage administrators have encountered in on-premises solutions, largely because features are typically tied to direct hardware rather than a more abstracted service.

Block storage in the cloud is effectively a giant, shared SAN, similar to an on-premises solution. It is, however, shared among users of the related service and is tightly coupled with the hardware underpinnings. This leads to a couple of challenges, including stuck or hung volumes.

> Storage networks in the cloud come with their own sets of challenges. These challenges are different from the ones that storage administrators have encountered in on-premises solutions, largely because features are typically tied to direct hardware rather than a more abstracted service.

Another performance bottleneck can result from noisy-neighbor shared networks connecting Amazon Elastic Block Store (Amazon EBS) and Amazon Elastic Compute Cloud (Amazon EC2) systems. Because multiple users share a storage pool, one user with busy traffic patterns can degrade performance for all users who are using the same hardware—but they won't know why. If you have degraded performance in block storage but no way to correlate it to a known problem, you have no way to solve the issue when it inevitably arises.

These block storage pain points can result in cumbersome and complex storage inefficiencies. To make matter worse, cloud infrastructure overspending can arise due to overprovisioning and scarce storage efficiency features. For example, if you don't have zero-impact snapshots or automatic tiering cold data off to cheap object storage, you'll end up overspending. Furthermore, DevOps principles can lead to quick setup or teardown of services, putting further strain on storage pools. Users often create "shadow IT," because it's quicker than using on-premises options or going through proper channels. Then storage sprawls across disparate management domains and platforms, causing costs to spiral out of control and storage tiers to become inefficient.

What you should expect from your block storage

On-premises block storage has traditionally offered many features to storage consumers and administrators. High-availability (HA) and application-integrated data protection with snapshots for DR, backup are must haves and not forgetting data deduplication, workload balancing, storage tiering, and an API-driven control platform are just a few such features. These capabilities have existed, in a mature form, in the data center for years.

A public cloud block storage offering ought to maintain these kinds of abilities. Moving block storage to the cloud shouldn't come with trade-offs and missing features. Low-latency connectivity, reliability, and consistency of I/O performance are all critical to database deployments such as Oracle, SQL Server, or SAP, and other mission-critical workloads such as Microsoft Exchange or SharePoint. Block storage has strong one-to-one mapping with the applications using it, so it offers highly predictable pricing and high-performance service. But when cloud block deployments don't maintain these capabilities, moving to the cloud involves a painful series of trade-offs.

Surprise billing also hurts deployment efficiencies in some block storage deployments. Cloud block storage should be no different, in actual use if not in technology, from on-premises solutions—it should offer the same predictable monetization of services. You shouldn't have to tolerate surprise bills when you move storage to the cloud.

Block storage for stateful containers

Stateful workloads in container deployments need enterprise-grade storage support. Thin provisioning, deduplication, snapshots, backup and restore, and nondisruptive volume expansion are necessary features for a strong containerized platform. Most cloud-based block storage can't do what's needed by the storage back ends of containerized workloads.

Another common problem with cloud-based block storage consists of mount times and stuck volumes. Spinning up stateful-storage–based containers can take as long as 2 minutes, drastically cutting the benefit of rapid cloud scaling—one of the primary motivations for moving workloads to the cloud.

Slow failover is also a challenge. Because the back end is tightly coupled to hardware, various problems might occur. For example, the storage resource can run out of mount points; existing mounts aren't released as you try to move resources; and APIs can fail. These problems make failover time stretch to several minutes. At this point, the storage is no longer highly available, so you aren't getting the value of persistent stateful storage.

Noisy neighbors are a problem for cloud-based block storage, too—more specifically, lack of IOPS. You can pay a lot for dedicated IOPS and largely avoid this problem, but why should you pay more for your block storage to actually work for you?

All of these challenges should be mitigated, and the features you want ought to work well in any cloud block storage solution. Sure, most public cloud providers offer block storage, but they don't include all these features. There isn't even feature parity or risk mitigation across platforms—something your applications expect, based on decades of running on-premises. Inconsistent feature sets across public cloud often lead to storage inefficiencies and unnecessary cloud waste.

No one-size-fits-all

No single public cloud file system fits all workloads. Applications require flexible and durable options so that you can optimize them to meet business needs. Ideally, you should be able to run applications on premises, colocated or hosted, connected to the public cloud with low-latency interconnects, on AWS or AWS Outposts.

> A cloud block storage solution should unify your on-premises and public cloud experiences by offering a consistent, enterprise-scale, elastic-cloud experience.

A cloud block storage solution should unify your on-premises and public cloud experiences by offering a consistent, enterprise-scale, elastic-cloud experience.

It should be offered as a flexible, pay-as-you-grow subscription-based service that enables movement of data between any physical location and AWS from a single pane of glass.

This is what you get with NetApp® solutions.

Get from enterprise-grade to enterprise-great on AWS

NetApp Cloud Volumes ONTAP® for AWS gives you feature parity with on-premises data center solutions. The nice part is ONTAP is ONTAP is ONTAP. From a hybrid AWS cloud to on-premises to AWS Outposts, you have the same storage efficiencies in a ubiquitous data fabric across any cloud. Not only do you get the same features your apps are currently using, but you also get that single-pane-of-glass interface for managing all of your block storage. (See **Figure 1** for an architecture overview.)

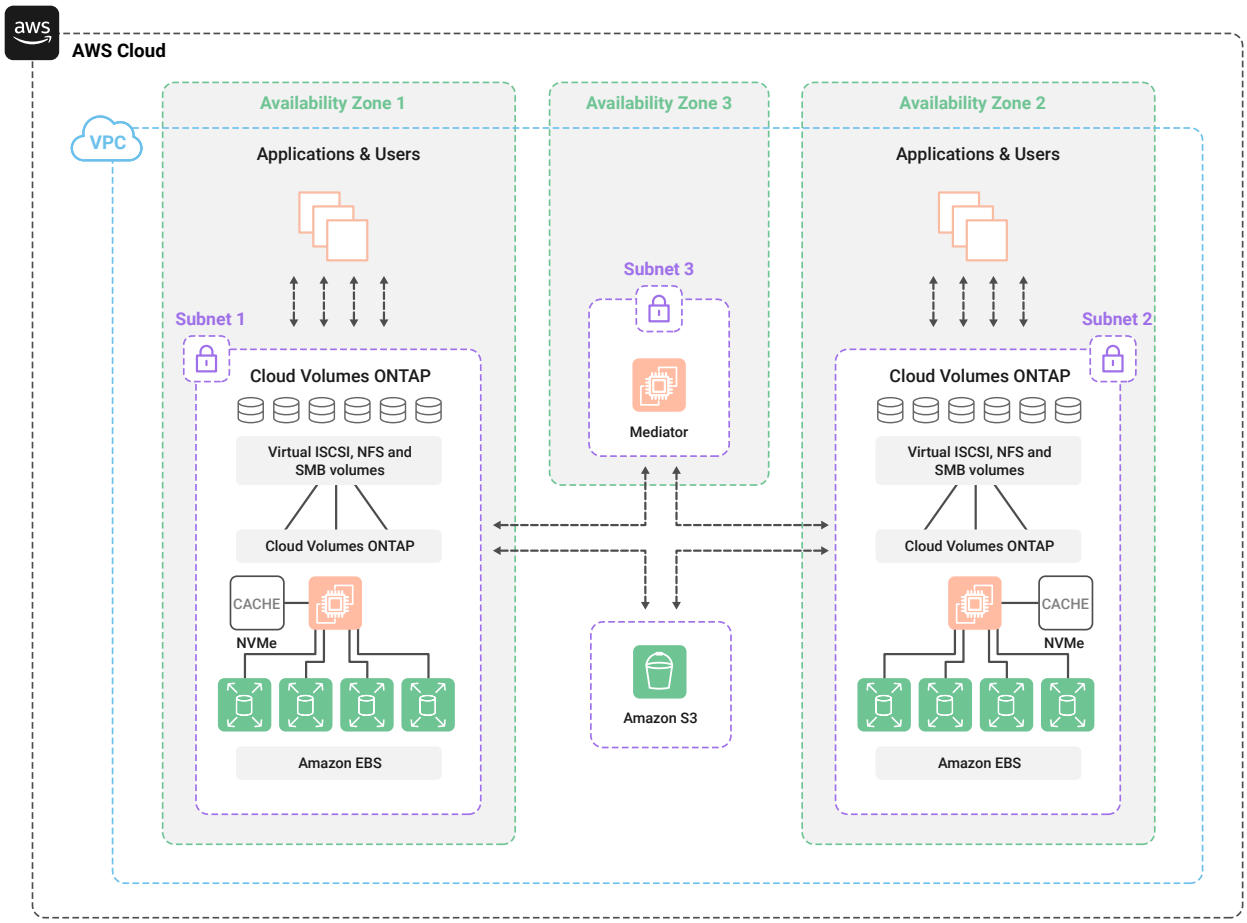

Figure 1: Cloud Volumes ONTAP architecture for AWS

NetApp has always been a leader in block storage, and analysts like Gartner and IDC recognize Cloud Volumes ONTAP as a proven, industry-leading storage solution. Paired with AWS, the acclaimed public cloud leader, NetApp technology provides enterprise-grade cloud storage to support you through the build, run, and optimize lifecycle—with storage cost and capacity efficiencies from data tiering and zero-impact NetApp Snapshot™ copies. Here are some of the things it can do for you:

- Connect data to the people who need it with persistent storage for NAS and SAN (iSCSI).
- Run and manage critical business applications in the cloud, with simplified data management and storage efficiencies to intelligently tier cold data to low-cost Amazon Simple Storage Service (Amazon S3) object storage.
- Guard your data like your business depends on it—achieve data protection from cross-zone high availability, cross region DR replication, and built-in backup to Amazon S3 with instant restores from point-in-time snapshots.

A complete approach

To succeed with application-driven systems, you need to understand how cloud infrastructure affects application resources. Cloud optimization can target both storage and compute infrastructures.

> By automating the allocation of resources, you prevent your CloudOps team from making error-prone manual changes.

Storage plays a pivotal role in supporting applications and in application-driven infrastructures, whereas complex workloads require visibility into your entire cloud infrastructure to maintain availability of resources, cost containment, and data protection. NetApp Cloud Manager offers a centralized control plane, making management of multiple cloud resources easier.

AUTOMATION

But those resources come in myriad types, configurations, and pricing models, making it challenging to ensure that they keep pace with your application's ever-changing needs. How can they keep up? Automation.

By automating the allocation of resources, you prevent your CloudOps team from making error-prone manual changes. And because your development teams can work faster without the constraint of insufficient infrastructure, the risk of shadow IT is reduced.

> To succeed with application-driven systems, you need to understand how cloud infrastructure affects application resources.

You can achieve AWS compute optimization and automation through Spot by NetApp. Spot can help you wave goodbye to cloud infrastructure complexity by automatically provisioning, scaling, and refreshing resources to avoid interruptions, meet demand, and minimize costs—all without requiring changes to applications. It's the only solution that continuously optimizes your cloud infrastructure's pricing, capacity, and availability with real-time analysis of reserved and on-demand compute resources.

SECURITY

But it's not enough to make apps more efficient—they also need to be kept safe. Data breaches occur every day around the world. This reality, combined with new regulations for data privacy and governance such as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA), have increased the need for privacy controls in storage infrastructures.

The NetApp Cloud Compliance service is an AI-based solution for your data privacy needs, offering a "set and forget" approach to protecting the data that matters most.

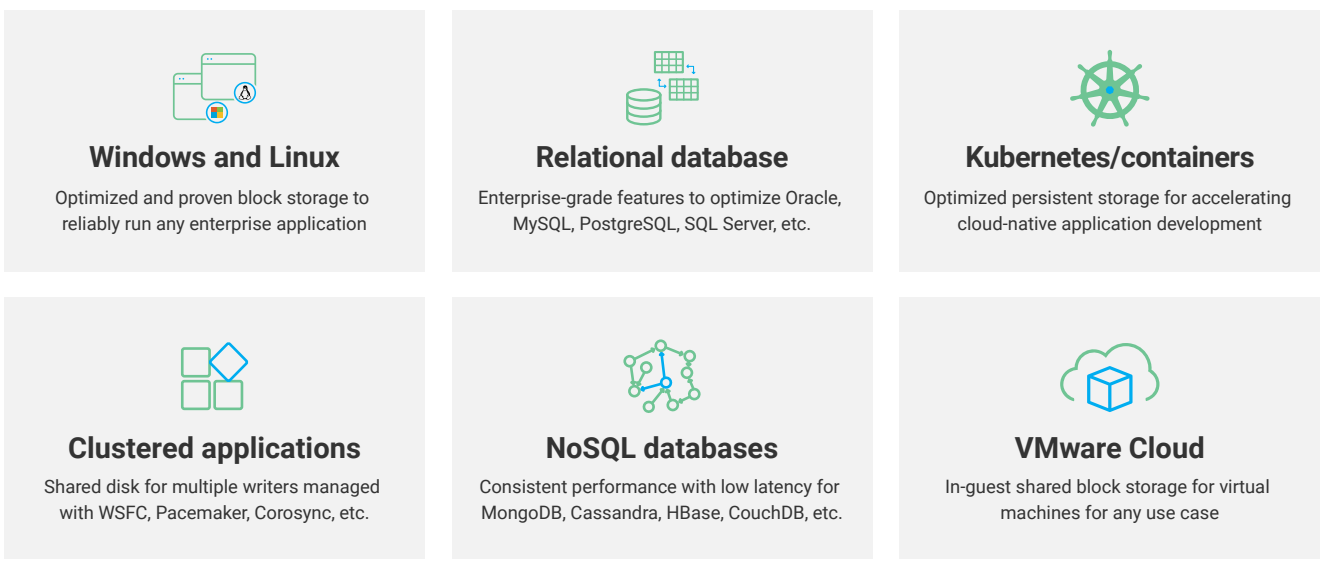

Figure 2: NetApp can help optimize a wide range of workloads

CONSOLIDATION

Managing islands of data in remote offices or across disparate geographic sites can be challenging. Excessive growth and duplication of data, inefficient use of server and storage resources, server sprawl, data integrity issues, and poor performance accessing data are all characteristics of an inefficient, unstructured storage system.

> Without visibility into all aspects of your cloud storage infrastructure and applications, you can't effectively or efficiently make business decisions about your data.

NetApp Global File Cache allows you to consolidate and centralize your data in the public cloud, taking advantage of the scalability and performance of cloud-based, enterprise-grade storage solutions.

MONITORING

Finally, without visibility into all aspects of your cloud storage infrastructure and applications, you can't effectively or efficiently make business decisions about your data. NetApp Cloud Insights allows you to monitor, troubleshoot, and thwart ransomware attacks in any cloud.

Use cases

Cloud block storage brings the enterprise benefits of on-premises SAN to AWS, letting you:

- Perform disaster recovery from an on-premises workload of Kubernetes cloud to AWS
- Unify your hybrid cloud into infrastructure as code with a single set of APIs
- Perform data deduplication to reduce storage and bring costs down
- Save money with cold data tiering and space-saving zero-impact Snapshot copies for backups to Amazon S3
- Unify siloed cloud storage with iSCSI block storage and multiprotocol file storage for NFS and SMB

Most large organizations have various workloads, as shown in **Figure 2**, that can take advantage of NetApp solutions.

Here's how NetApp can help your organization.

DATA PROTECTION

To be successful, business continuity and disaster recovery strategies need a secure, cloud-based 3-2-1 backup plan. Moving your applications to the cloud doesn't inherently ensure 100% uptime; site and service outages commonly affect application and data availability. But you can meet the 3-2-1 data backup rule with incremental, block-based, "forever" backups.

The right cloud storage includes built-in capabilities to keep your applications online and your data safe. NetApp provides several tools for safe, efficient backup and recovery, including:

- Built-in storage efficiencies with point-and-click simplicity for both cloud and on-premises backup and restore with NetApp Cloud Backup
- Protection of cloud-native apps like Salesforce and Microsoft 365 data with NetApp SaaS Backup
- Flexible, scalable disaster recovery with nondisruptive testing for your cloud and on-prem workloads with NetApp Cloud Volumes ONTAP for AWS
- Snapshot copies that provide near-instantaneous point-in-time backups of your data and don't consume additional resources or affect application performance
- NetApp Cloud Sync, which tiers Snapshot copies to cost-effective Amazon S3 storage, freeing up the Amazon EBS performance tier for "hot" data
- NetApp SnapMirror®, which offers data replication between on-premises systems and Cloud Volumes ONTAP
- NetApp SnapVault®, which creates copies of your source data with longer retention policies

DATABASES IN THE CLOUD

Databases require quick read/write operations, and volumes can be mounted directly to a cloud instance's file system, so there's only a minimal delay in writing and retrieving data. A database needs to be durable, available, and scalable enough that you can resize when you need to.

Cloud Volumes ONTAP for AWS delivers the high IOPS and sustained high throughput (with submillisecond latencies) required for your database performance. NetApp solutions also provide HA, performance durability, and 99.999% reliability for increased data protection. These features result in a robust database storage solution.

FILE SERVICES USING NFS AND SMB

Cloud Volumes ONTAP for AWS lets you quickly and easily create cloud-based NFS (Linux) and SMB (Windows) file shares. The same file system can be accessed concurrently by many different users and applications for read and write workloads.

DEVOPS: DEVELOPMENT AND TEST ENVIRONMENTS

Many DevOps shops rely solely on Amazon EBS and its various IOPS speeds, but those offerings can come with limitations. NetApp delivers a consistent, seamless DevOps experience across a hybrid multicloud environment that's designed to accelerate application development and streamline your pipeline. With NetApp DevOps solutions, operations teams deliver automated infrastructure with less engineering, and developers create in reliable and predictable environments with less friction.

> A database needs to be durable, available, and scalable enough that you can resize when you need to.

CONTAINERIZATION BLOCK STORAGE

Amazon Elastic Container Service (Amazon ECS) and Amazon Elastic Kubernetes Service (Amazon EKS) make delivering cloud-native applications on AWS easy. But some apps require persistent storage, something Kubernetes and containers don't provide natively. That's where NetApp comes in. You can integrate the management capabilities of Cloud Volumes ONTAP with NetApp FlexClone® technology so that any volume can be cloned instantly and consistently, regardless of source data size. NetApp also provides its NetApp Docker Volume Plugin as an open-source tool called NetApp Trident.

PRODUCTION CLOUD ENVIRONMENTS

Centrally managed, tiered storage can be connected to Amazon EC2 instances over iSCSI. Storage can also be served out to client users and hosts directly from Cloud Volumes ONTAP through either NFS or SMB.

How much can you save with NetApp?

Save up to 90% on infrastructure costs, or reduce storage by up to 90%.

With NetApp, **both** are achievable. The storage efficiencies of Cloud Volumes ONTAP enable you to save on standard AWS costs—up to thousands of dollars per month.

To find out how much you can save, NetApp has created a cost calculator that makes it easy.

What can you do with NetApp solutions?

NetApp Cloud Volumes ONTAP offers advanced block storage features to help you throughout the build, run, and optimize phases of deployment. Try it now for free, and see how much you could save. Visit NetApp Cloud Central to discover the broad range of cloud products, many of which include free trials.

About NetApp

In a world full of generalists, NetApp is a specialist. We're focused on one thing, helping your business get the most out of your data. NetApp brings the enterprise-grade data services you rely on into the cloud, and the simple flexibility of cloud into the data center. Our industry-leading solutions work across diverse customer environments and the world's biggest public clouds. As a cloud-led, data-centric software company, only NetApp can help build your unique data fabric, simplify and connect your cloud, and securely deliver the right data, services, and applications to the right people—anytime, anywhere.