NetApp®

Database Acceleration for AWS

NetApp Cloud Volumes Service Speeds Up Database Applications

Cloud Volumes Service for AWS: Database Use Cases

Cloud Volumes Service delivers great performance and availability for database applications that demand high transaction rates, high throughput, and low latency:

- **OLTP**
- **Analytics**
- **Data warehouses**
- **Development and test**
- **Copy data management**

WHY RUN YOUR DATABASE IN AWS?

Enterprise IT teams are increasingly turning to the cloud. Many organizations are mandating that all new applications be hosted in the cloud, adopting a cloud-first strategy. Some are closing data centers and undertaking large-scale migrations, while others move to cloud infrastructure during technology refresh cycles or to improve scaling.

Databases can pose significant challenges during cloud migration, but there are many advantages to moving databases to the cloud:

- Eliminate upfront capital expenditures and infrastructure refreshes every 3-5 years
- Take advantage of cloud elasticity to scale up or down in response to changing load
- Reduce the time spent deploying and managing infrastructure

Moving databases and associated applications to AWS—and running new projects in AWS from the start—is a proven approach that, done correctly, can make life simpler for IT, business, and development teams.

Moving a database to the cloud requires careful consideration and planning. Enterprise databases demand high service levels for performance, availability, durability, security, and data management.

Database and Cloud: Common Questions

A number of questions come up over and over when NetApp talks to customers who want to run databases in the cloud:

- *How can I move my database quickly and without loss of productivity?*
- *As my database grows, will performance scale with it?*
- *What if I make a mistake? How hard will it be to fix it?*

NetApp Cloud Volumes Service for AWS is the answer. It can help you move legacy database applications without refactoring and scale performance on the fly without locking you into a set level of performance.

With Cloud Volumes service you can move databases and applications to the cloud months—or even years—sooner than with other methods.

Amazon Web Services (AWS) offers a range of database-as-a-service (DBaaS) offerings that make it easy to get a database application up and running in the cloud. Services such as Amazon Aurora, Amazon RDS, and others address a range of database needs, but they aren't ideally suited for every enterprise use case:

- **Lack of control.** With DBaaS, you may not have full administrative control over your database.
- **Size limitations.** DBaaS offerings may have limits on allowable storage capacity.
- **Data protection and DR.** Database services may not meet your recovery point objective (RPO) or recovery time objective (RTO).
- **Storage performance.** With database services, performance is often tied to how much capacity you provision. You may have to over-provision capacity to get the performance you need from your database or opt for more expensive provisioned performance storage. Provisioning the performance you need for peak periods can increase your costs significantly.
- **Copying.** The instant copy capabilities of enterprise storage systems are in wide use by database teams as an efficient way to support testing, analytics, and a variety of other use cases. Because DBaaS offerings may rely on physical copies or limit the number of copies you can create, it can disrupt these workflows.
- **Lift and shift.** In many instances, you may need to "lift and shift" your current database applications to the cloud with a minimum of operational changes. This may not be possible with DBaaS offerings.
- **Hybrid and multicloud.** Many enterprises have databases that need to interoperate with services running on-premises or in another cloud. Choosing DBaaS may make this impossible. A corollary is that you may want to keep your options open so you can easily move a database between providers, either to get the best price or as an extra level of disaster recovery.

For these reasons, an infrastructure-as-a-service approach—combining AWS EC2 compute instances with the right cloud storage and your preferred database software—is often a better choice than DBaaS for enterprises. NetApp Cloud Volumes Service for AWS is the ideal cloud storage option to address database challenges, making it simple and affordable for your organization to run high-performance database workloads using a convenient service in the AWS cloud. With Cloud Volumes Service, you can lift and shift existing databases or spin up new instances without having to worry about performance or capacity limitations, and without sacrificing enterprise capabilities you rely on.

This eBook explores the capabilities of Cloud Volumes Service for AWS for supporting database needs.

NetApp®

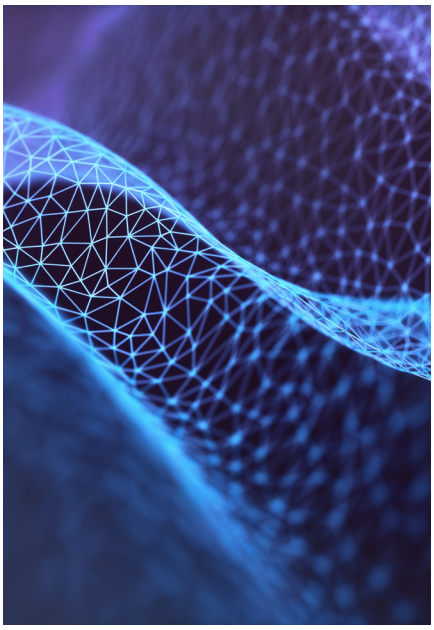

Introducing NetApp Cloud Volumes Service for AWS

If your organization is planning to run popular databases such as Oracle, Microsoft SQL Server, MySQL, PostgreSQL, and MongoDB in the cloud, NetApp Cloud Volumes Service for AWS is the ideal solution to meet your data storage needs. Cloud Volumes Service is a fully managed cloud service that is able to deliver extreme performance to meet the needs of even the most demanding databases.

Cloud Volumes Service for AWS offers managed file services for NFS and SMB protocols with dual-protocol access when needed, providing a number of advantages for database operations versus other cloud storage solutions:

- **High performance.** High throughput and low latency access meets the demands of the largest databases.
- **Low cost.** Cloud Volumes Service for AWS can be significantly less expensive to deploy and manage. Because you can adjust performance levels on the fly, it's possible to optimize performance for peak loads without overpaying during off-peak periods.
- **Fully managed.** NetApp configures and manages your infrastructure, so you don't have to.
- **Easy to deploy.** You can deploy a 100TB cloud volume in seconds and start deploying database applications immediately.
- **Enterprise features.** Cloud Volumes Service provides enterprise features like Snapshots, backup, and synchronization that facilitate data management for database environments and associated development efforts.
- **Guaranteed SLAs.** NetApp service level agreements (SLAs) for performance, availability, and durability allow you to be certain your cloud operations satisfy workload needs.

Why Choose NetApp for Cloud Database Needs?

If you're unfamiliar with NetApp, you may ask why you should trust us to deliver data services for databases running in the cloud.

Founded in 1992, NetApp pioneered the idea of network attached storage (NAS) and was the first to offer high-performance NAS appliances that greatly simplified the configuration and management of file services. NetApp has supported high-performance database environments for decades, including full support for Oracle Direct NFS (dNFS) since it was introduced in Oracle 11g.

NetApp was instrumental in developing many of the storage efficiency and data protection technologies we take for granted today, including Snapshots, cloning, and much more.

Recognizing early that the public cloud would reshape the way that enterprises approach IT, NetApp created the Data Fabric, an architecture and set of data services that provide consistent capabilities spanning on-premises and cloud environments.

NetApp Cloud Volumes Service for AWS is a key element of the Data Fabric vision. NetApp understands database performance and data management, so you don't have to; you can focus instead on innovation that moves your business forward.

NetApp®

Reasons to Choose Cloud Volumes Service for AWS

Fully managed. Eliminate the challenge of configuring and managing high-performance cloud storage

Integration. Full compatibility with popular databases, including Oracle and Microsoft SQL Server, as well as popular services such as Active Directory

Data Protection. Efficient Snapshots prevent data loss or corruption

Easy Migration. Move data to/from on-premises or other repositories via Cloud Sync

High availability. Ensure data remains available with guaranteed SLAs

Automation. Schedule tasks easily or take advantage of full REST APIs

Cloud Volumes Service for AWS: The NetApp Difference

The advanced data management that NetApp Cloud Volumes Service for AWS delivers for database workloads allows enterprises to move databases into the cloud with confidence. Cloud Volumes Service matches your need for performance and scale with the flexibility to automatically grow or shrink capacity and performance as necessary, significantly reducing your spending on storage capacity.

Most organizations need multiple copies of data for testing and development. Creating and refreshing those copies is a time-consuming and tedious process. Cloud Volumes Service for AWS allows you to fast copy and backup database files, drastically improving the process of copying, backing up, and reverting for a quicker time to market.

With NetApp Cloud Volumes Service, your organization can get access to the performance and capacity necessary for a wide range of database services with just a few mouse clicks, making you instantly productive.

NetApp Cloud Volumes for AWS Specifications	
Performance Tiers	Standard | 16MB / second throughput per allocated TB Premium | 64MB / second throughput per allocated TB Extreme | 128MB / second throughput per allocated TB
Advanced Data Management	• Cloud Sync • Snapshots • Rapid Copies • Rapid Revert • Backup • Encryption
Protocol Support	• SMB: 2.1, 3.0, 3.1.1 • NFS: NFSv3

NetApp Cloud Volumes Service for AWS is fully managed and optimized to deliver the performance, availability, scalability, security, and data management features you need. Choose the performance level you require today and switch performance levels on the fly with no penalties.

NetApp®

Cloud Volumes Service for AWS: Cost Savings

Many database workloads have fairly predictable peak times where they need maximum performance—but less performance the rest of the time. Examples include online sales or Uber and Lyft peak periods.

Suppose you need 30TB with Extreme performance 20% of the time. If you configured a Cloud Volume for peak performance and leave it set there, that costs about $9,000 per month.

However, the same Cloud Volume configured for Extreme performance 20% of the time and Standard performance 80% of the time, would cost about $3,600 per month, a savings of $6,400.

It's easy to switch performance levels automatically if you know when your peak periods are, and you can change performance levels manually on the fly for unexpected spikes. Cloud Volumes performance changes instantly without the need for data movement.

CLOUD VOLUMES SERVICE FOR AWS: OPTIMIZED PERFORMANCE AND COST

IT teams moving databases to the cloud for the first time are frequently surprised how much it can cost to deliver the necessary level of performance. With Cloud Volumes Service for AWS, performance concerns stop being an issue. You choose the performance tier you need up front and switch between three different performance levels dynamically, without the need to move data or reconfigure everything from scratch. This capability enables you to easily optimize your performance and cost, switching to high performance levels only when needed, which can yield substantial savings (see sidebar).

Guaranteed IOPS and Bandwidth

Cloud Volumes Service for AWS provides three different performance levels, allowing you to choose the right performance for each volume to meet the needs of your database applications. Assuming an I/O size of 4KB, Cloud Volumes Service delivers the performance levels (IOPS and bandwidth) shown below:

NetApp Cloud Volumes for AWS Performance Tiers		
	IOPS per TB	**Throughput per TB**
Standard	4,000	16MB/sec
Premium	16,000	64MB/sec
Extreme	32,000	128MB/sec

In NetApp internal testing, we've found that Cloud Volumes Service delivers up to 20x the performance of build-your-own file services, and NetApp has carefully benchmarked the performance of Cloud Volumes Service for a variety of databases. These benchmarks can help you make sure the service will meet your performance needs prior to provisioning.

NetApp®

Understanding Storage Performance

Different database applications generate different I/O workloads. It's useful to think about the I/O characteristics of the workload or workloads you need to support in terms of three metrics.

IOPS

Transaction-oriented applications such as online transaction processing (OLTP) generate small, random reads and writes. This type of storage performance is measured in I/O Operations Per Second or IOPS.

Throughput

Applications such as data warehouses and analytics rely on sequential access to data in large blocks, resulting in a workload that's dramatically different from transaction-oriented applications. Throughput is a measure of the amount of data that can be moved in or out of storage and is usually reported in MB/sec or GB/sec.

Latency

Latency is a measure of how long it takes to satisfy an I/O. It is an important metric for time-critical applications such as real-time trading and OLTP, which are extremely latency sensitive. Interactive users notice and react to differences in latency.

Oracle IOPS with Cloud Volumes Service

NetApp tested the IOPS performance of Oracle running against a single Cloud Volume using the SLOB2 benchmark, the de facto standard for Oracle performance testing.

As shown in the following figure, we first compared a single Oracle instance running on two different AWS compute instances: c5.9xlarge and c5.18xlarge with a read workload.

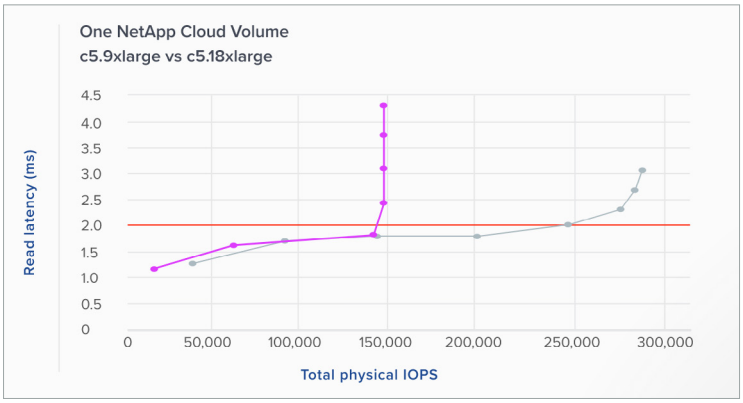

Oracle is able to drive 250,000 file system IOPS in 2 milliseconds (ms) using the c5.18xlarge instance and a single Cloud Volume, or 144,000 file system operations at below 2ms using the c5.9xlarge instance. (Latency may vary by AWS region. These results are for us-east-1.) A native storage configuration with the same compute instances yielded maximum performance of about 80,000 IOPS (not shown); in that test, maximum performance was clearly I/O-limited.

This benchmark not only demonstrates that Cloud Volumes Service is capable of delivering 250,000 IOPS to a single Oracle instance—a very high level of performance—but that your database is much less likely to be artificially I/O-limited when using Cloud Volumes Service. In many cases, you end up over-configuring cores and memory to compensate for I/O limitations. Depending on your database, this can significantly increase database licensing costs; Cloud Volumes Service may reduce your spending on database licensing.

Our second test looked at the performance of one or two Cloud Volumes in conjunction with the c5.18xlarge compute instance. The SLOB2 workload was varied from read-only to mixed read/write to write-only. As the figure shows, a second Cloud Volume can deliver additional low-latency IOPS, especially for a mixed workload.

Sizing A Cloud Volumes Deployment

Accurate sizing to deliver the necessary bandwidth for your database can be critical when it comes to delivering the fastest results and optimizing costs.

This video series will help you size a workload for Cloud Volumes Service to get the most from the service.

- **Part 1. Architecture**
- **Part 2. Performance**
- **Part 3. Oracle on AWS**

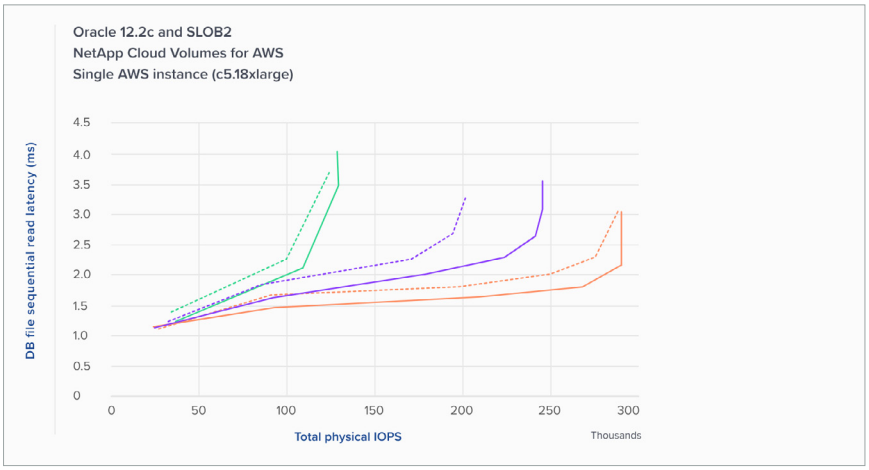

To see all the latest performance benchmarks for Cloud Volumes Service for AWS, visit: cloud.netapp.com/cloud-volumes-service/aws-benchmarks

Cloud Volumes Service for AWS: Scalability

Scaling is another significant challenge for databases in the cloud. In addition to providing extreme levels of I/O performance, NetApp Cloud Volumes Service for AWS is highly scalable and the only file system with guaranteed SLAs[1]. The ability to scale up automatically without needing to tune infrastructure is a major advantage of Cloud Volumes Service for AWS.

If you exceed the purchased capacity in a Cloud Volume, you can continue to write data to the volume beyond the allocated capacity, however the performance of the volume will not increase until the allocated capacity is increased.

Cloud Volumes Service for AWS: Availability and Data Protection

Availability and data protection are vital for business-critical database workloads. Cloud Volumes Service is designed to deliver 99.99% data availability using built-in HA features including network failover and advanced data protection. Cloud Volumes Service for AWS is immune to Availability Zone failures and provides consistent performance in each region. Redundancy is implemented at every level. NetApp technologies protect against a variety of disk errors, improving data durability and integrity beyond that of native AWS services.

Cloud Volumes Services for AWS provides instant protection for database files of any size through the use of space-efficient Snapshots. A Snapshot copy takes only moments to create, and Snapshots remain completely stable and incur no performance overhead. You can store up to 255 Snapshot copies per volume. Snapshots can be run manually or on a set schedule.

NetApp Cloud Backup Service complements Cloud Volumes Service for AWS with backup to Amazon S3 object storage for long-term archive and recovery, reducing backup costs.

NetApp®

Cloud Volumes Service for AWS: Enterprise Data Management

Data management is a challenge in any environment. Cloud Volumes Service for AWS gives you the capabilities you need to take the pain out of moving, organizing, protecting, and accessing data in your database environment.

Instant data copies. A range of database use cases from dev/test to analytics require copies of existing database instances. Cloud Volumes Service Instant Copies solve this challenge. You can create new copies of a database from Snapshots of a volume. Copying and Snapshots are essential tools for admins, DBAs, and busy development and test teams.

Data security. Every organization is concerned with data security, particularly if you're in a regulated industry such as financial services, healthcare, or government. With Cloud Volumes Services, data at rest is encrypted by default, so your data is always protected. Data encryption over the wire is available for users of the SMB3 protocol.

Cloud Volumes Service for AWS: Migrating Data

Cloud Volumes Service for AWS capabilities take the pain out of moving data from existing databases running on-premises to the cloud. You can move data efficiently to the cloud using either the native log shipping capabilities of your database or NetApp Cloud Sync.

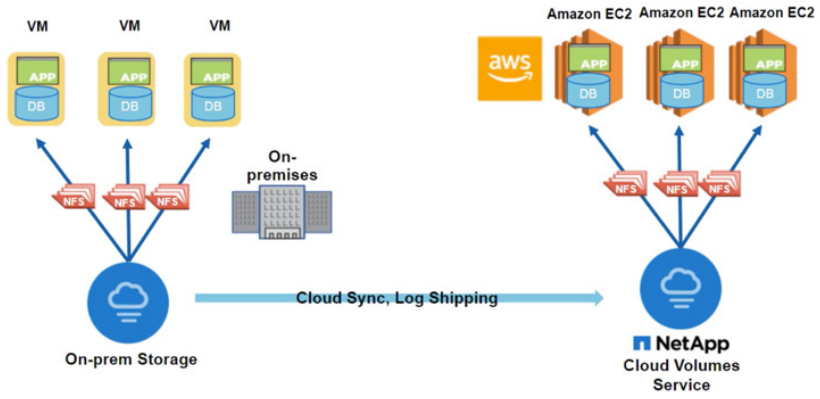

Cloud Sync is NetApp's service for rapid and secure data synchronization. Whether you need to transfer database files between on-premises repositories and Cloud Volumes Service or between Cloud Volumes in different regions, Cloud Sync does the job quickly and securely. Cloud Sync is integrated and included for free with Cloud Volumes Service for AWS.

NetApp®

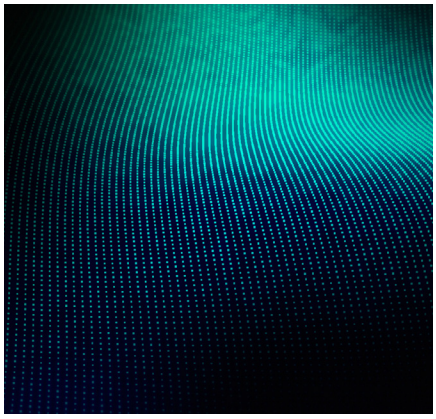

Getting Started with Cloud Volumes Service

Cloud Volumes Service for AWS is an enterprise-grade solution for deploying high-performance file sharing to support a wide range of cloud uses cases, including genomics, electronic design automation (EDA), seismic data processing, media and entertainment, high-performance databases, and large software builds.

This fully-managed solution lets you easily deploy scalable file shares, migrate data as needed, and get to work without delays. Flexible service levels allow you to optimize performance and cost to meet your application needs without breaking your budget.

Visit the Cloud Volumes Service page on NetApp Cloud Central for the latest information. There, you can request a demo with a NetApp Cloud expert or sign up for the Cloud Volumes Service. You can also go directly to the AWS Marketplace to get started.

[1] Guaranteed SLAs so in the event any of the "Service" does not meet the Service Levels for performance and/or availability, you will be eligible to receive a "Service Credit" from NetApp for the future use of the Service.

NetApp®