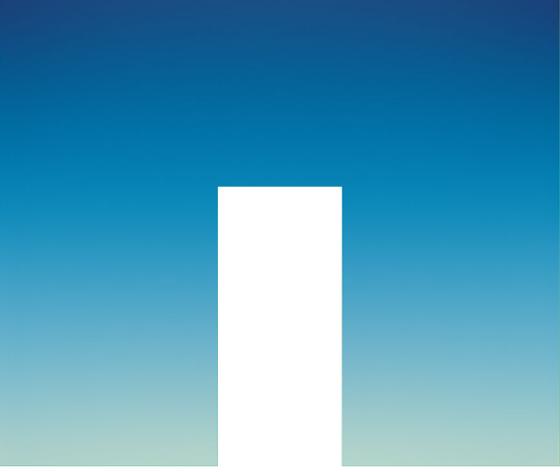


WHITE PAPER

# What Is Cloud Volumes Service? NetApp's Latest Cloud Solution



Executive Summary	3
Cloud Volumes Service and Azure NetApp Files: An Introduction	3
Use Cases for Cloud Volumes Service and Azure NetApp Files	3
File Services	3
Benefits: File Services	4
Databases	4
Benefits: Databases	4
Analytics	4
Benefits: Analytics	4
DevOps	4
Benefits: DevOps	4
Getting Started: Using Cloud Volumes Service	5
Registration	5
Create a 100TB Volume	5
Mounting and Functionality	5
Data Transfer	5
Upcoming Features	5
Conclusion	6

## Executive Summary

Cloud Volumes Service is NetApp's new, fully managed hyperscale cloud service for enterprise users seeking high performance and reliability.

Cloud storage comes in many forms, from object storage to block-level devices. However, many organizations simply require a ready-to-use, highly available, high-performance shared file service where scalability is not an issue. This is the solution that NetApp's [Cloud Volumes Service](#) offers as a native cloud marketplace service in AWS and Google Cloud Platform, and as a first-party service called [Azure NetApp Files](#), delivered supported by Microsoft, built on NetApp technology in Microsoft Azure.

Being able to provision NFS and SMB volumes on demand helps to simplify the deployment of cloud-based file services, database environments, and analytics clusters, to name a few use cases. [Cloud Volumes Service](#) and Azure NetApp Files come with many features to help ease the management of large volumes of data, such as support for NetApp Snapshot™ copies and cloning, which DevOps engineers can use to rapidly create writable clones of production data. Using synchronization features, users can consolidate data in the cloud from multiple sources, either on premises or from other systems in the cloud.

This white paper examines the capabilities of Cloud Volumes Service and Azure NetApp Files in detail and looks at the considerable advantages of using them in a number of use cases, including file services, database, analytics, and DevOps. It also walks you through the easily configurable initial steps to setting up Cloud Volumes Service for an example deployment on AWS.

## Cloud Volumes Service and Azure NetApp Files: An Introduction

Cloud Volumes Service and Azure NetApp Files are native cloud services, meaning that they are available to end users in the same way as any other cloud storage service, such as Amazon S3 or Azure Disk Storage. This removes any requirement to install or set up the service, enables seamless administration of all cloud services, and unifies the billing process for cloud deployments.

Using Cloud Volumes Service and Azure NetApp Files, highly available NFS and SMB file shares can be created instantly, with the ability to rapidly scale capacity up or down at any time. Performance is configured by selecting a service level for each storage volume, which is used to determine the level of IOPS and throughput the storage volume is able to deliver. Service levels help to control costs by giving users the option to use high-performance I/O where it's really needed, and a more cost-effective solution for colder data, such as file backups or archived information.

Snapshot technology is used to instantly create point-in-time copies of the data in a Cloud Volume. With an upcoming

version of the UI and the release of the RESTful API, users will be able to use Snapshot copies to instantly restore the active file environment to the snapshot. Clone volumes, a tool of choice for DevOps engineers, are writable copies of the data at the point in time at which the copy was created. They can be created quickly and used to test data changes or software releases before deployment to production systems.

Cloud Volumes Service and Azure NetApp Files also come with synchronization services that can efficiently copy data in to and out of a Cloud Volume, and ensure that data is kept synchronized through incremental updates. This reduces the complexity of deploying new cloud-based systems and enables centralized data repositories to be created in the cloud with data from various source systems. Synchronization can be scheduled to occur as often as required, such as once per hour or once per day.

Cloud Volumes Service operations on [AWS](#) and [GCP](#) are performed through a modern, web-based UI that is both fully featured and easy to use, while users of Azure NetApp Files can work directly with the native Azure UI. The Cloud Volumes UI dashboard in AWS and GCP and the Azure NetApp Files dashboard in the native Azure UI can streamline the process of creating, managing, and reviewing the file environments and Snapshot copies. The UI can also be used to create clone volumes and to manage synchronization relationships. Cloud Volumes Service and Azure NetApp Files will also soon be accessible through RESTful API calls, which will provide programmatic access to users who need to integrate operations, such as Snapshot creation or cloning, with an automated process.

## Use Cases for Cloud Volumes Service and Azure NetApp Files

This section describes a number of the ways in which Cloud Volumes Service and Azure NetApp Files expedite and support the deployment of various cloud-based systems through rapid delivery of cloud-based shared file service and a rich set of storage management features. The subsections focus on the following use cases: file services, databases, analytics, and DevOps.

### File Services

Cloud Volumes Service and Azure NetApp Files are highly available and enormously scalable platforms for creating cloud-based NFS and SMB file systems. NetApp's long and extensive experience in delivering enterprise, on-premises NAS solutions means that Cloud Volumes Service and Azure NetApp Files come with a complete range of supporting features, such as read-only and read-write client access control, connections over NFSv3, NFSv4 (coming soon), Active Directory integration for SMB file systems, and much more.

Synchronization services speed up and simplify the process of feeding new cloud-based shared file systems with data from existing systems. After the initial baseline copy, only the changes made to the source data are transferred to the destination, which makes it very efficient to keep data

synchronized over time. Data can also be synchronized from Cloud Volumes and Azure NetApp Files out to other systems.

Snapshots are used to create instant backups. They can be accessed through regular file-browsing tools for a read-only view of the file environment at the time the snapshot was created. This can be vital for accessing previous versions of files on a shared file environment. Snapshots can also be scheduled to occur on a regular basis through the Cloud Volumes Service UI or in Azure NetApp files in the native Azure UI.

#### **Benefits: File Services**

- Supports both NFS and SMB file shares
- Shared file access across Windows and Unix/Linux operating systems
- Complete integration with file directory metadata
- Business continuity
- Data protection
- Simple cloud deployment
- Secure multicloud and hybrid data mobility
- Integration and automation enabled via RESTful APIs (coming soon)

#### **Databases**

Database platforms require reliable, scalable, and high-performance storage environments to deliver low-latency and high-throughput data services to client applications. Cloud Volumes Service and Azure NetApp Files deliver a storage solution that meets these requirements by using NFS and SMB, which are supported by all major commercial and open source database systems, such as Oracle, Microsoft SQL Server, PostgreSQL, MySQL, and MongoDB.

Cloud Volumes Service supports different levels of performance for each file environment. Database administrators can allocate individual storage pools for hot and cold data, giving them fine-grained control over the use of high-performance storage or more cost-effective capacity storage. Because Cloud Volumes Service and Azure NetApp Files ensure that file environments are always available and resilient against system failures, it's easy to set up highly available database services in the cloud.

Snapshot copies enable instant backups of any database, eliminating the need to create a full copy of the data. In order to implement application-consistent Snapshot copies, the upcoming RESTful API will make it possible to integrate with processes that quiesce the database. Writable clones of production databases can quickly be created for testing purposes, such as to test a database upgrade or a new data transformation routine.

#### **Benefits: Databases**

- Clone technology
- Guaranteed protection
- Business continuity
- Data protection
- Simple cloud deployment

- Secure multicloud and hybrid data mobility
- Integration and automation enabled via RESTful APIs (coming soon)

#### **Analytics**

Cloud Volumes Service and Azure NetApp Files can be used to create data lakes in the cloud, synchronize data with on-premises systems or other data sources in the cloud, and provide direct access to the data from cloud analytics services, such as Azure HDInsight using [NetApp In-Place Analytics](#).

Incremental synchronization of data from repositories across an organization reduces the time and effort required to make data available for cloud-based analytical processing. The results of the analysis can also be synchronized out of Azure NetApp Files or Cloud Volumes Service back to other systems, whether on the premises or in the cloud. Targeting file environments directly from analytics services, such as Azure HDInsight, means that copy operations are not necessary, and allows a single repository to serve data for both regular use and analytics.

Snapshot and cloning technologies can be used to create temporary working copies of analytics data stored in Cloud Volumes or Azure NetApp Files, for use by data engineers to test data enrichment or transformation processes. Writable clones are created very quickly and can simply be dropped when testing is complete to free up the resources they were using.

#### **Benefits: Analytics**

- Data consolidation
- Integration with public cloud analytics
- Secure multicloud data mobility
- Predictable performance
- Clone technology

#### **DevOps**

DevOps engineers are tasked with the automated set up of test and staging environments for use by software developers or to support continuous integration and continuous delivery processes. It's difficult to set up these test environments with an up-to-date copy of production data, because production datasets are usually large in size and copies of the data may need to be refreshed repeatedly in order to rerun test suites.

Cloud Volumes Service uses NetApp cloning technology to rapidly create writable clones of existing data volumes. Data in the clone can then be added to, updated, or deleted without affecting the primary storage. When testing is complete, the clone is dropped to free up the resources it was using. Cloud Volumes Service and Azure NetApp Files support the use of multiple clones of the same source file environment concurrently.

#### **Benefits: DevOps**

- Fast and efficient cloning
- Faster TTM with and dev/test processes

## Getting Started: Using Cloud Volumes Service

Using Cloud Volumes Service and Azure NetApp Files is easy to understand. This section walks through the basic steps to begin using Cloud Volumes Service on AWS, for example.

### Registration

The initial step to begin using Cloud Volumes Service is to register on AWS. Either click the sign-up link or visit the AWS Marketplace. After signing up, the next step is to visit [NetApp Cloud Central](#), which is the basic launchpad for all the cloud services provided by NetApp. Existing NetApp users can sign in with their SSO, and new users can begin a registration here. After signing up, users will be ready to start using Cloud Volumes Service.

### Create a 100TB Volume

One of the most powerful features of Cloud Volumes Service is the speed and ease of volume creation. To create a new volume, follow these steps.

1. Select Create New Volume in the Cloud Volumes Service interface.

2. Enter the following information:

- **Name.** You can name the volume anything you choose. As an operational best practice, NetApp recommends keeping the name simple and easy to remember.
- **Region.** Select the region in which to create the new volume. Currently Cloud Volumes Service is available in the us-east-1 and us-west-1 regions. More regions are expected to be added in the coming year.
- **Time Zone.** Select a time zone for the volume.
- **Generated Volume Path.** Define the generated volume path. This path name can be automatically generated by Cloud Volumes Service with names that are based on the same format that is used in Docker (randomized adjectival words plus notable names from the computing world, such as “fastasticgates”). This convention keeps volume paths easy to remember and share, as opposed to long number-and-letter combinations.
- **Volumes.** Create the volume itself. Volumes can be instantly created on demand, or they can be based on Snapshot copies that the user already owns.
- **Service Levels.** Select a Cloud Volumes Service service level. There are three choices: Standard, Premium, and Extreme. The difference between these service levels is the level of performance (IOPS and latency) that is available and its corresponding price. Select the performance level that best fulfills your deployment requirements.
  - **Standard service level.** Up to 1000 IOPS per TB (16k I/O) and 16MB of throughput per TB.
  - **Premium service level.** Up to 4000 IOPS per TB (16k I/O) and 64MB of throughput per TB.
  - **Extreme service level.** Up to 8000 IOPS per TB (16k I/O) and 128MB of throughput per TB.

- **Quota.** Specify the desired volume size. 100TB is the current size limit that Cloud Volumes Service offers, although size increases up to the multipetabyte level are planned for the future.
- **Export Policies.** Set the volume's export policy for accessing the volume. This is useful in giving access to specific IP addresses, read-write or read-only permissions, and protocol to expose. SMB, NFSv3, and (coming soon) NFSv4, are all options. Multiprotocol selection is also on the way, which will allow both NFS and SMB to be served at once, if necessary.
- **Snapshot Policy.** Define the Snapshot policy. The policy can be set to an automatic creation schedule based on a day of the week and time. The number of Snapshot copies to create and store is also configurable here.

3. Click Create Volume. Within seconds, a volume as large as 100TB will be ready for use on AWS.

### Mounting and Functionality

Once the new volume is created, you can test its functionality by mounting it to an Amazon EC2 instance. The Cloud Volumes Service interface provides clickable mount instructions, a resource for carrying out this procedure. (The instructions vary depending on the protocol you have chosen.)

### Data Transfer

Cloud Volumes Service uses NetApp [Cloud Sync](#) to establish sync relationships between the Cloud Volumes Service volume and any other storage device or devices in use. Target sources can be on the premises, on AWS, on other Cloud Volumes, or even on a different CSP storage system. Sync relationships can be one-time-only transfers or used on a continuing sync basis.

### Upcoming Features

Cloud Volumes Service and Azure NetApp Files will also be programmable through the RESTful API (to be released soon), which will make it easy to automate any Cloud Volumes Service function.

## Conclusion

Cloud Volumes Service and Azure NetApp Files are flexible, scalable, highly available, high-performance cloud-native services for deploying shared file service in AWS, Google Cloud Platform, and Azure, respectively. As discussed earlier in this paper, Cloud Volumes Service and Azure NetApp Files can be used to build out cloud-based file services, database systems, data lakes for analytics, and much more.

Along with enterprise-grade NFS and SMB file services, Cloud Volumes Service and Azure NetApp Files also provide advanced capabilities for managing storage volumes, such as Snapshot copies, cloning, and incremental data synchronization. These features help organizations to quickly bring new cloud systems online and reduce the complexity of many operations, such as DevOps environment setup.

Using Cloud Volumes Service and Azure NetApp Files, storage capacity can be increased from a few gigabytes to many terabytes within seconds. The performance of the volumes is determined by the service level the end user chooses, which allows a balance between performance and cost.

Visit us [online](#) to find out more about Cloud Volumes or start using it today by [signing up on AWS or GCP](#), or for [Azure NetApp Files](#).

NetApp is the data authority for hybrid cloud. We provide a full range of hybrid cloud data services that simplify management of applications and data across cloud and on-premises environments to accelerate digital transformation. Together with our partners, we empower global organizations to unleash the full potential of their data to expand customer touchpoints, foster greater innovation, and optimize their operations. For more information, visit [www.netapp.com](http://www.netapp.com). #DataDriven

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

### Copyright Information

Copyright © 2018 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

### Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.