




WHITE PAPER

Enhance DevOps in AWS Cloud with NetApp



Executive Summary	3
Introduction	3
What Is NetApp Cloud Volumes Service?	3
Custom NAS Solutions	3
AWS DevOps Services and Cloud Volumes Service	4
DevOps with Cloud Volumes Service	4
Conclusion	5

Executive Summary

Continuous integration (CI) and continuous delivery (CD) DevOps pipelines are making it easier and faster for companies to get their products to market using Amazon Web Services® (AWS). This white paper shows how AWS services for continuous integration (CI) and continuous delivery (CD), as well as infrastructure management, and application deployment such as AWS CloudFormation™, AWS CodePipeline™, and AWS OpsWorks™ can be integrated with NetApp® Cloud Volumes Service, NetApp's new fully-managed storage solution for deploying fast, scalable, and highly available shared file systems in AWS. It discusses all of the Cloud Volumes Service features for effective data management, such as NetApp Snapshot™ and cloning technologies, as well the RESTful APIs that allow DevOps engineers to integrate Cloud Volumes Service management with their workflows to cut time to market (TTM) to a minimum.

Introduction

DevOps engineers on AWS are able to automate the setup of test environments for complex application stacks composed of compute nodes, database services, file shares, and much more. Stateless services are easy to replicate, but working with large volumes of production data is significantly more complicated. For example, how can you efficiently make a test copy of a multi-terabyte file share to run a suite of integration tests? How can this be made a part of an automated CI process that might have to be executed several times a day?

NetApp [Cloud Volumes Service](#) for AWS is a fully-managed storage solution for deploying fast, scalable, and highly available shared file systems integrated to run with AWS compute services. Cloud Volumes Service for AWS also comes with a sophisticated set of storage management features, such as point-in-time Snapshot technology, volume cloning, and data synchronization, all of which can be accessed through the web-based user interface or from the RESTful APIs. This gives DevOps engineers the flexibility to create writable copies of production data for testing, in a way that can be integrated with existing automated processes.

This white paper examines the features of Cloud Volumes Service for AWS that can be used to enhance DevOps in an AWS environment. It considers how Cloud Volumes Service can be used with AWS services such as AWS CodePipeline, AWS CloudFormation, and AWS OpsWorks to create test environments that include cloned copies of production data.

What Is NetApp Cloud Volumes Service?

Cloud Volumes Service is a storage-as-a-service (SaaS) offering from NetApp that tightly integrates with AWS compute services. This means that end users can access, deploy, and manage the billing for Cloud Volumes Service for AWS just as they would for any other AWS cloud service, such as [Amazon S3](#) or [Amazon EBS](#). In the same way as any other fully-managed cloud service, end users are completely insulated from all management of the lower-level infrastructure that makes the

service work. Scalability, high availability, and performance are all managed automatically by Cloud Volumes Service. Users simply create a volume and start using it.

High-performance I/O is a hallmark of Cloud Volumes Service, which allows the IOPS of a file system to be adjusted based on its service level. Users can instantly change the capacity of a file system at any time, and they can change the service level on the fly as well. The available service levels are:

- **Standard.** Up to 1000 IOPS per TB (16k I/O) and 16MB/s of throughput per TB.
- **Premium.** Up to 4000 IOPS per TB (16k I/O) and 64MB/s of throughput per TB.
- **Extreme.** Up to 8000 IOPS per TB (16k I/O) and 128MB/s of throughput per TB.

Using a shared file system, such as NFS or SMB, allows hundreds or thousands of concurrent clients to access the same files at the same time. NetApp has decades of experience with building enterprise, on-premises NAS solutions to make Cloud Volumes Service a truly cloud-scale platform.

Custom NAS Solutions

NetApp Cloud Volumes Service use the advanced data protection features of NetApp [RAID-DP®](#). RAID-DP provides replication and allows all the data to be available to all availability zones, maximizing redundancy and availability with less cost and complexity normally associated with building a reliable NAS solution, and goes a long way toward safeguarding data without sacrificing performance.

Second, the balance between I/O performance and cost effectiveness needs to be managed at a much more granular level when building a custom NAS solution in the cloud. For example, to create a file system for capacity storage, the underlying storage would have to be created and scaled manually. If later it becomes necessary to move this storage to faster disks, provisioning the new storage, migrating the files, and decommissioning the old storage would also need to be performed manually. With Cloud Volumes Service, the size of a file system can be changed instantly at any time, simply by setting a new size value. The service level for each file system, which determines its performance, can also be changed on the fly.

Third, the complexity of building a solution that can support advanced features such as Snapshot copies, clones, data synchronization, and variable service levels (QoS) requires an organization to have a lot of in-house expertise. Using Cloud Volumes Service for AWS, a read-only, point-in-time copy of any file system can be created instantly, and writable clones of the file system can be spawned based on that copy. Using the built-in synchronization capabilities, data can also be brought in from on-premises systems or other cloud systems, and kept incrementally synchronized, without the need for custom scripts.

AWS DevOps Services and Cloud Volumes Service

This section looks at how Cloud Volumes Service for AWS can benefit DevOps workflows in the AWS ecosystem, including how it can be leveraged by AWS CloudFormation, AWS CodePipeline, and AWS OpsWorks.

One of the major benefits of using Cloud Volumes Service for AWS is the ability to create writable clones of existing storage without adversely affecting the source storage volume being cloned. This makes it very easy to set up development, test, and CI and CD environments that include access to an up-to-date copy of production data. A clone is based on a Snapshot copy, which determines the state of the data in the clone at the point in time at which it is created. To create a more up-to-date clone, users can simply use a more up-to-date Snapshot copy, which can also be created directly before cloning. A single Cloud Volumes Service storage volume can support many simultaneous clones, allowing DevOps engineers to provide cloned storage for multiple environments at the same time.

[AWS CloudFormation](#) gives DevOps engineers and cloud architects the power to deploy infrastructure by using code and to define templates to create an entire stack of AWS services. Cloud Volumes Service for AWS can be integrated into these templates by using a custom resource definition, allowing a Snapshot copy and clone to be created of existing storage and then connected to the new test servers deployed by the template.

AWS CloudFormation:

- Define infrastructure-as-code
- Deploy application services consistently
- Update deployments by using change sets

AWS CloudFormation templates can also be used with [AWS CodePipeline](#) for the implementation of a CI/CD process. This allows DevOps teams to manage test environments in the same way they would manage production deployments, where each test environment is effectively a Snapshot copy of the production environment at a point in time, complete with data. When testing has been completed, the test stack can be deleted and the clone storage freed up.

AWS CodePipeline:

- Automate continuous integration and continuous delivery
- Integrate with GitHub, Jenkins, TeamCity, and much more
- Deploy to AWS Elastic Beanstalk, Amazon ECS, AWS OpsWorks, and more

An alternative approach is to use [AWS OpsWorks](#), which can also be used to stand up an entire application stack based on a [Chef](#) cookbook recipe. A recipe can include calls to RESTful API services, making it possible to access Cloud Volumes Service for the purpose of creating storage clones. Chef provides strong support for AWS services and may be more suitable for users who have previously used it for deploying on-premises systems. AWS OpsWorks also offers support for [Puppet](#).

Without this type of instant mechanism for cloning storage volumes, getting a copy of production data into a test environment would be extremely cumbersome. Data would need to be copied from production storage, which could be terabytes in size. If test servers make changes to the data copied over from production while running test suites, the data must be refreshed before the tests can be repeated. This may be required many times throughout the course of a day, which adds significant load to production and is not a scalable solution.

AWS OpsWorks:

- Configuration management service for AWS DevOps
- Support for both Chef and Puppet
- AWS OpsWork stacks for managed application deployments

DevOps with Cloud Volumes Service

Cloud Volumes Service supports many different platforms, including file services, database systems, and analytics data repositories. Cloud Volumes Service can be used anywhere that there is a need for reliable, high-performance storage that can be integrated with DevOps processes. Cloud Volumes Service supports NFSv3, SMB with Active Directory integration, and (coming soon) NFSv4.1.

Using Snapshot technology, it is possible to instantly create a read-only, point-in-time copy of a Cloud Volumes Service file system. These copies can act as backups and offer protection against accidental data deletion or other data issues that require going back to a previous version. As mentioned earlier, these copies also form the basis for creating writable clones, which can be created instantly for any size of source data.

As well as providing access to all functionality through a web-based UI, Cloud Volumes Service also supports integration via REST. This makes it possible to use Cloud Volumes Service functionality, such as creating Snapshot copies and clones, creating and deleting volumes, changing service levels, and reverting to Snapshot copies, all with automated processes, which is how DevOps engineers typically prefer to work. RESTful APIs have become the standard mechanism for interacting with remote resources over a network, promoting a uniform protocol for accessing systems and building new processes.

Cloud Volumes Service test environments that are created with clones can be simply dropped at the end of a test cycle to free up the resources they were using. New Snapshot copies and clones can be created to capture an up-to-date copy of production storage volumes. In this way, getting the latest state of a file system becomes trivial, because all Snapshot and clone operations are instant, regardless of the source data size.

Scheduling enables Snapshot copies to be created on a regular basis, such as hourly, daily, weekly, etc. This is useful if it is necessary to create test environments based on the state of the system at specific times in the day. All file systems, Snapshot

copies, Snapshot schedules, and clones can be reviewed and managed through the Cloud Volumes Service UI, which provides convenient dashboards to simplify the administration of allocated storage volumes.

Cloud Volumes Service also provides a data synchronization service that enables data to be copied into or out of a Cloud Volumes Service file system and kept synchronized incrementally. This feature can be used to consolidate files from multiple locations, whether on the premises or from other cloud systems. If data is synchronized into Cloud Volumes Service from multiple data sources, creating a Snapshot copy or a clone of the destination file system captures the current state of all the files at the same time.

Incremental synchronization makes keeping files synchronized over time very efficient. After the initial baseline copy, only the files that have changed since the last synchronization are transferred to the destination. Without this feature, data would need to be fully synchronized every time changes were made to the source file system. Cloud Volumes Service can incrementally synchronize data from an external source or to an external destination. The source or destination can be on-premises systems or other systems in the cloud.

Conclusion

DevOps engineers are constantly required to create test environments that mirror production systems. Ideally, this would be done with an up-to-date copy of production data. However, achieving this result is complicated, because production datasets are usually large in size and test environments frequently need to be refreshed with the latest data. This data might be required for a file share, for user home directories, as storage for database systems, or for analytical processing.

NetApp Cloud Volumes Service for AWS is a high-performance, highly available platform for hosting NAS solutions in the cloud. With features such as NetApp Snapshot technology and

cloning, Cloud Volumes Service makes it very easy to create writable clones of production data. These clones can be made readily available to AWS services such as AWS CodePipeline, AWS CloudFormation, and AWS OpsWorks. Cloud Volumes Service supports access over REST, so DevOps engineers can quickly integrate the cloning of existing storage into their automated workflows and dramatically improve TTM.

To find out more about how you can speed up your AWS DevOps workflow and cut your TTM, [register for Cloud Volumes Service for AWS today](#).

About NetApp

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 like AWS, 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. #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.

WP-7279-1018