

WHITEPAPER

Databases on Azure NetApp Files and Cloud Volumes Service




Table of contents

Executive summary	3
Database storage in the cloud	3
Database solutions with Cloud Volumes Service and Azure NetApp Files	4
Conclusion	5

Executive summary

Robust, high-performing, and scalable storage in the cloud is essential for deploying a database system, where the reliable storage and retrieval of data across hundreds or thousands of concurrent client connections is paramount. Large organizations rely on their databases to be the permanent system of record for all business transactions, which means that the data they contain must survive localized server and disk failures, as well as a site-wide failure necessitating disaster recovery. As data volumes grow, database administrators require the flexibility to quickly scale up both the size and performance of database storage volumes to meet demand.

NetApp® Cloud Volumes Service and Azure NetApp Files are fully managed, highly available, high-performance solutions for provisioning file services in AWS, Google Cloud Platform, and Azure, respectively. NetApp brings its industry-leading reputation in building data management systems to all the major cloud platforms, providing data protection, scalability, and advanced storage management. The Cloud Volumes Service is used like any other cloud marketplace service, allowing seamless administration of all cloud infrastructure while Azure NetApp Files is a first-party Azure service, delivered and supported by Microsoft built on NetApp Technology.

This white paper examines in depth the considerations for choosing a database cloud storage service, and shows how Azure NetApp Files and Cloud Volumes Service address these areas and offer solutions.

Database storage in the cloud

The primary goal of a database system is to promote efficient data manipulation while maintaining data integrity and security. This requires storage services to deliver on many fronts to be as effective as possible. The following list summarizes the most crucial features that the storage environment must support:

- **Performance.** Because almost all database operations involve reading or writing data, I/O performance determines the speed at which a database can operate. Ideally, the storage environment provides the ability to allocate storage pools of varying performance, with, for example, archive data on slower storage and live application data on faster storage. This gives database administrators the flexibility to match storage capacity to performance requirements in order to stay cost effective.
- **Reliability.** Because the storage environment is an integral part of the database platform, it must be available for sustained access by database servers without fail. If access to the storage is interrupted, database operations might come to a halt, potentially causing a major disruption for all dependent applications and systems.
- **Durability.** When a database system is ACID (atomicity, consistency, isolation, and durability) compliant, it means that the database system can guarantee that when a database transaction has been committed the data is durable; that is, it will survive a failure. Storage environments must ensure adequate data redundancy to protect against failures.
- **Security.** Organizations with strict requirements for data storage require features such as encrypted transport and data encryption at rest. These features are necessary to protect sensitive data, such as personal, financial, and healthcare information.
- **Disaster recovery.** Database storage that can withstand a severe outage without data loss (i.e., RPO=0) presents enormous value to database administrators setting up disaster recovery systems. Applications can continue operating after failover with minimal downtime, allowing businesses to remain operational in the worst-case scenario, and ideally with an RTO < 60 seconds.

There are other data management features that can significantly ease the burden of database management, especially when working with large volumes of data. In some cases, it may be possible to work around these issues, but having a data service that provides these capabilities can save considerable time and effort.

- **Snapshot copy and restore.** Leverage the capabilities of the storage infrastructure to take a NetApp Snapshot™ copy, which creates an instant, point-in-time copy of a storage volume. Being able to also instantly restore these copies makes for an unparalleled backup and restore solution. Snapshot copies help to prevent inadvertent data loss and data corruption.
- **Storage cloning.** A clone of a database system is often required to perform testing, such as for a database upgrade or an application deployment. Being able to clone existing storage volumes and quickly create temporary, writable, and up-to-date copies of large production databases, without any adverse effects to the live environment, is a big win for DevOps engineers and database administrators. Software developers with access to database test environments are able to accelerate the development and testing of new application features, ensuring faster time to market (TTM).
- **Migration services.** Storage environments that can be easily synchronized with on-premises data, or from other cloud environments, have an advantage by making it easier to migrate existing systems to the new platform.

Database solutions with Cloud Volumes Service and Azure NetApp Files

Azure NetApp Files and Cloud Volumes Service deliver flexible and ready-to-use file services in the cloud without administration overhead. Users can simply set up a new data volume based on their current capacity, performance, and protocol access requirements and immediately have access to a new file share. Azure NetApp Files and Cloud Volumes Service also offer a number of data management features that enable much greater flexibility in working with allocated volumes.

Database systems can realize the following benefits:

- **Support for all major database platforms.** Reliable database storage for all major database systems, including Oracle, PostgreSQL, MySQL, MongoDB, and Microsoft SQL Server (over SMB). Oracle Direct NFS is able to open multiple, parallel client sessions to NFS shares to further increase I/O performance and scalability.
- **Configurable performance.** Storage is allocated in accordance with the service level defined when a storage volume is created. The storage level can then be changed on demand to best suit the user's needs. This allows the performance of a volume to be controlled and storage pools of varying size and performance to be made available to a database system. Users can control cloud storage costs by allocating faster storage only where necessary.

Users can select one of three service levels:

Standard (up to 1000 IOPS/TB and 16MB throughput/TB), Premium (up to 4000 IOPS/TB and 64MB throughput/TB), and Extreme (up to 8000 IOPS/TB and 128MB throughput/TB).

- **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.

- **Robust data protection.** Azure NetApp Files and Cloud Volumes Service are cloud-native services for allocating storage in the cloud and therefore require no user management of underlying resources. Because volumes are highly available, customers can be sure that their data will be online when they need it. They can also use built-in replication capabilities to set up secondary regional copies of their data for enhanced protection. And users can be sure that their data is durable and always available.
- **Data encryption.** All data is encrypted at rest, and users can encrypt data both at rest and in transit by using a VPN. Users can encrypt data in transit from the database server to the storage volume. This provides transparent protection from malicious attempts to access the data.

- **Scalability.** Volumes can be expanded on the fly, whenever more storage is required, without compromising performance or data protection. They can also reduce the size of the volume as needed to use their resources better. Database administrators can easily allocate storage for new databases or existing databases that are growing in size, without the need to manage any of the underlying physical infrastructure. This substantially reduces the administrative overhead of reacting to a change in database storage requirements.

As well as providing a powerful solution for creating and managing cloud file service, Azure NetApp Files and Cloud Volumes Service offer an additional set of storage management features that simplify many common database administrative tasks. These features enable:

- **Snapshot copy and restore.** Instantly create a Snapshot copy of a storage volume of any size. The storage volume can then be instantly restored back to the point in time when the copy was created, whenever it is needed. To create a consistent Snapshot copy, NetApp recommends that users first quiesce their database system to ensure that in-flight I/O operations have been completed.
- **Storage cloning.** Azure NetApp Files and Cloud Volumes Service can quickly create writable copies of existing storage volumes. This is especially useful for database administrators who need the flexibility to rapidly create database test environments without the overhead of manually copying large volumes of data.
- **Replication and data synchronization.** Leverage NetApp replication and synchronization services that can read data from on-premises or cloud-based datasets and incrementally synchronize them with your cloud volume. This speeds up the process of synchronizing existing database systems to the cloud. Synchronization can also be performed in the opposite direction, out of the cloud volumes to other locations.

Using Azure NetApp Files and Cloud Volumes Service makes it easier to manage cloud storage deployments by providing the following tools and support to cloud architects and database administrators.

- **Cloud native.** Cloud Volumes Service is accessible to users just like any other cloud marketplace service provided by AWS or Google Cloud Platform. Azure NetApp Files is accessible just like any Microsoft first-party service in Azure. They provide consistency when adding the service to a new or existing cloud deployment.
- **Simple web interface.** Cloud Volumes Service for AWS and GCP provides an easy-to-use, web-based UI that allows users to manage storage volumes, create Snapshot copies and clones, and synchronize data with other locations. With Azure NetApp Files, users get all of these features through the native Azure interface itself.
- **Integration support.** All of the operations that can be performed through the UI will soon also be available by using the RESTful API. This makes it easy to integrate Azure NetApp Files and Cloud Volumes Service operations with other processes, such as creating consistent Snapshot copies or setting up and cloning database test environments.

Conclusion

Database systems are complex enterprise applications that depend heavily on the I/O systems they use. For the best results, storage services must combine performance, data protection, scalability, security, and flexibility into a single solution.

Azure NetApp Files and Cloud Volumes Service are designed to meet these demands, a step forward for cloud data management solutions. Because they are available through all major cloud vendors, users can employ Azure NetApp Files and Cloud Volumes Service just as they would any other cloud service from Azure, AWS, and Google Cloud Platform, respectively.

To get the performance and limitless storage that running your database in the cloud requires, [register for Azure NetApp Files or Cloud Volumes Service on AWS or GCP](#) today.





Copyright Information

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 © 2021 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).