

eBook

The NetApp Guide to Migrating Enterprise Workloads to the Cloud



Introduction: Migrating Enterprise Workloads to the Cloud

Your enterprise workloads are the backbone of your business operations. Whether they are production applications, customer-facing or back office systems, databases, big data, web applications, file services, or other custom-built applications, all of them must run uninterrupted, securely, and efficiently.

Enterprise workloads are the way that the world runs, and the cloud is making enterprise workloads run better, work harder, and cost less to maintain than ever before. As companies around the world pivot to cloud-enabled IT deployments, those that still have enterprise workloads running exclusively in on-premises data centers might need more guidance about the way to move forward. For a large enterprise with diverse customer bases and needs, totally abandoning an existing IT structure might seem too risky, no matter what kind of advantages the cloud gives the company over its competitors.

Running an enterprise workload in the cloud comes with challenges: from adapting to a cloud architecture and new application process management (APM) processes to renewed focuses on data integrity, business continuity, and

cost control, moving to the cloud is not just about solving problems, it's about learning to avoid new ones. To meet these challenges, the option preferred by most companies taking early steps into the cloud is to adopt a hybrid cloud architecture for their enterprise workloads.

This handbook acts as a guide for every step of the enterprise workload migration process to the public cloud using AWS, Azure, or Google Cloud, from the initial motivations behind the move to orchestrating the migration and making sure of its deployment. As challenges to moving an enterprise workload present themselves, we look at ways to handle these situations, including NetApp® technologies such as Cloud Volumes ONTAP® that provide solutions to support and simplify the migration process.

Table of Contents

Why Migrate an Enterprise Workload to the Cloud?	3
Determining Your Workloads' Migration Requirements	4
The Type of Migration: Choosing a Method	5
Testing the Workload	6
Finding the Right Migration Solutions	7
NetApp Solutions for Migrating Enterprise Workloads to the Cloud	8
Cloud Volumes ONTAP: Key Features	10
The Final Stage: Keeping Tabs on the Cloud	11
Cloud Migration Case Studies with Cloud Volumes ONTAP	12
The Cordant Group	12
McKesson	13
AdvancedMD	14
Summary	15



Why Should Enterprise Workloads Migrate to the Cloud?

The motivating factors behind a cloud migration are unique to every company, but behind most of them are several factors. There is the desire to keep pace with the technology that is being adopted by rival businesses, which allows them to scale larger while cutting costs at the same time. There is also the desire to increase business agility, cut time to market, and trade CAPEX spending for OPEX, all of which the cloud makes possible.

But that goal is at the end of a complex journey. To make sure that the migration will be a success, a careful planning process needs to be in place. Every aspect of the company's business needs and existing IT infrastructure has to be assessed so that the new cloud-based deployment can meet all of the enterprise's business needs and performance requirements. Risks have to be realistically considered and addressed.

Large enterprises might assume that the cloud won't be capable of absorbing their workloads, but the case histories show otherwise. One example of a company that made a major shift to the AWS cloud was [Philips](#). Known primarily for its lighting and appliance branches, Philips also does significant work in the healthcare industry. The switch of Philips to the cloud brought 15 PB of patient data into AWS and increased the company's application performance, scale, and reliability. Due to regulations in place to safeguard sensitive information—in this case, information such as medical records—migrations such as that of Philips can be particularly challenging.

[Geico](#) is an example of a major enterprise moving to Azure cloud services. By moving to a DevOps culture and shifting development to Azure, the insurance giant was able to increase its availability, cutting down the release time for new features and lowering operational costs throughout its business. But that success didn't happen overnight. Geico's initial strategy to move to the cloud was to shift to all-cloud deployment within one year. But after analysis within the company, Geico decided to gradually migrate workloads as part of a hybrid cloud architecture. This allowed them to validate each component that they moved to the cloud before attempting to migrate the next workload.

[PayPal](#) is a major online financial institution that billions of users worldwide rely on for electronic money transfers, banking, and other financial services. After breaking off from eBay, PayPal moved a number of its mission-critical applications that were on-prem to Google Cloud. But the approach to the cloud PayPal took was different: to gain flexibility, their deployment is hybrid as well as multicloud. This kind of architecture provides agility and expands the number of resources and services available to the company, but can add specific challenges when it comes to managing data across all those silos.

What these examples show is that even the biggest enterprises with complex customer requirements and needs can make the jump into the cloud—however, the migration method and business requirements will shape specific challenges they'll face, both getting to the cloud and maintaining their applications in it.

Cloud Migration Plan Challenges



Environment: Deciding which cloud (or clouds) will benefit your company the most



Complexity: Moving compute, network, and storage to a new environment



Validation: Quantifying and testing how systems will perform in the cloud



Heavy lifting: Finding the data sync and replication tools that will carry out the migration

Determining Your Workloads' Migration Requirements



No two companies have the same IT deployments, so no two companies migrate to the cloud the same way. Though every company can find increased scale, added agility, and reduced costs in the cloud, the parts of their business that are able to take advantage of those benefits differ. Before any sort of migration takes place, determining your company's requirements in the cloud has to take place.

What do these requirements look like? The initial planning stage of a migration has to identify the workloads that it is feasible to move to the cloud, because it might not be practical to move everything. The user base and the rate of usage for each workload that you plan to move need to be identified. Your customers or users are affected by the move if you don't plan to meet their usage needs during the shift, so you need to make sure the migration affects users as little as possible.

This is also the stage to determine network configurations. Your network might have interdependencies with the workloads you plan to move that might be affected by a transition to the cloud. Another major factor is costs. The three cost components of compute, storage, and networking must be calculated. However, because the storage component is the one that grows on a constant basis, minimizing storage costs should be a key consideration when planning the move.

Probably the most important factor to determine in this early stage is how available you need your workloads to be. Do you require your workloads to be highly available? Two factors come in when addressing your workloads' availability: one is your recovery time objective (RTO), and the other is your recovery point objective (RPO). These numbers represent the amount of time that it takes your business to recover from failure with acceptable losses and the point in time your business can operate without its data, respectively. For critical enterprise workloads, these numbers most likely need to approach zero. These numbers determine the shape that your high-availability, disaster recovery, and business continuity plans take, which most likely are also supported by the cloud.

Another important factor is your business protection requirements. Is protecting your workload data a key requirement? In disaster recovery (DR), secondary copies of data are crucial to making sure that your workloads can be restored in case there is ever a catastrophic event (such as a natural disaster, ransomware attack, or hardware failure). Your workload needs to be able to failover to a secondary site if and when such events occur and be able to fail back when the primary site is up and running again, keeping in mind all of your stated service-level agreements (SLAs).

Requirements that the business has in regard to meeting SLAs for users must also be considered here. There might also be compliance and regulation guidelines that your business is expected to follow, such as HIPAA in the health industry and FISMA at the U.S. federal government level.

In addition to the sensitive personal data that can be subject to some of the regulations mentioned above, there is also a chance that there is duplicate, stale, and unnecessary data in your on-prem repositories. Moving all of that data to the cloud will not only present higher migration costs, but also bring issues with it that will need to be addressed in the cloud, complicating your success in the new environment. Your migration planning should include a step for you to "clean up" your data sets through a systematic review of all the data you're considering to move to the cloud.

Key steps for determining the shape of an enterprise workload migration:

- Make an inventory of workloads you consider moving to the cloud.
- Identify usage base.
- Calculate compute, network, and storage costs.
- Determine security and recovery needs, including SLAs and RTO/RPO points.
- Research legal implications (that is, compliance).
- Choose your cloud provider.

The Type of Migration: Choosing a Method

After you have determined the needs of your enterprise workload in the cloud, it is time to determine the type of migration that best meets those requirements. Currently there are two main routes that your migration can follow: infrastructure as a service (IaaS) and platform as a service (PaaS). IaaS uses cloud-based virtual machines that customers can use and configure on their own as they see fit in terms of runtime, operating system, and middleware. With PaaS, those virtual machines are configured by the service provider itself and then offered to the customer.

AWS pioneered a large part of cloud culture, and they [famously identified six migration strategies for applications](#):

Retire will see the application and workload reach the end of its life

Retain keeps the application in the data center, with no move to the cloud (currently)

Rehost also known as the “lift and shift,” moves the application intact to run on cloud resources

Replatform the code that runs your workload is slightly modified to meet the cloud deployment requirements


Repurchase involves moving to a cloud-based SaaS product over the existing application

Refactor requires a complete code rebuild of the application to fundamentally take advantage of the cloud.

Retiring and retaining don’t get the application to the cloud, so we’ll focus on the four options that take existing workloads and can go between the IaaS and PaaS categories: rehost, replatform, refactor, and repurchase. With IaaS you can rehost (“lift and shift”) or replatform (“lift, tinker, and shift”), and with PaaS you can refactor or repurchase.

The PaaS options are different. With refactoring, your workload code is run on your cloud provider’s service. The drawback here is that you might lose some of the functionality that you once had, because the cloud provider infrastructure differs from your own. An additional drawback to a refactoring is that you need to recreate your APM processes. A refactoring means rewriting the code for your application from scratch, which aligns your workload most closely to the cloud provider’s services. However, a rebuild might also mean vendor lock-in with that provider. Repurchasing means scrapping the application you have previously worked with in favor of a provider’s SaaS offering. While this option may take less time than a refactor, it comes with the same potential loss of existing functionality, in addition to totally giving up control over the application itself.

Between these four options, the fastest way to get an enterprise workload into the cloud and running is to go with “lift and shift” rehosting. A refactoring rebuild is obviously the most costly, risky, and time-consuming form of migration.



Another step to take at this stage is to build the leadership team that is responsible for carrying out the migration. This team can be selected from cloud supporters within the organization, or it can be done with the help of a managed service provider. This team has to work closely with leaders in departments all around your company, from the IT department to marketing and sales teams, so it is important that the team includes point persons who can relay the migration plan's goals and needs in each field.

The service provider that you choose should also be consulted, because it can assist your move. Cloud providers have [expert teams](#) that are available to give advice and help you reorganize your architecture for cloud deployment, with security and compliance needs in mind. For existing NetApp storage system users, [turning to NetApp](#) at this time is highly advantageous. NetApp has cloud solutions such as Cloud Volumes ONTAP that work seamlessly with on-premises storage systems already in use at your data centers. Determining how to best transition those resources is an important decision that NetApp can help you make. If

you use other third-party solutions, look into their availability on your cloud provider's marketplace. They might have compatible cloud versions to use in your transition, but their use might also affect your existing agreements.

Key points to consider as you choose a migration method:

- Choosing the model that works best for your enterprise workload
- Cloud leadership team: finding the personnel who can best manage the transition
- Service provider migration resources and planning
- Planning how existing infrastructure and third-party solutions map out in the cloud

Testing the Workload


Testing how your workload runs in the cloud is the next important step in the migration. Here you should build a proof-of-concept model that allows you to see what the real costs of operating the cloud are and validate that the workload performance is acceptable in a test environment. This testing is to plan for the correct amount of services you require to run efficiently, also known as "tuning."

It is outside the scope of this handbook to detail the exact differences between all of the storage formats, compute types, databases, and networking services that the different cloud service providers make available, but it is important that you are aware of these differences. [Azure has a handy guide](#) to recognizing the differences between their services and those on AWS that gives some idea of what is available as you begin your research. Google Cloud has similar resources that compare their products and services, including [one for AWS](#) and [one for Azure](#). Understanding what your provider offers makes it clear whether or not your existing solution is integrable or upgradable in these regards.

Security is another concern to address during the testing process. It isn't easy for many enterprise companies to accept that their infrastructure partially or even entirely exists under the control of some other company. There should be no gap between the level of security that you currently use and the security you need to set up in the cloud. If anything, the migration should be a chance to increase your security levels by considering additional security tools such as [Amazon VPC](#), [Azure's network security groups](#), and [Google Virtual Private Cloud](#).

Key points at this stage:

- Figure out the total amount of storage and compute your workload needs
- Determine your expected cloud costs
- Put your workload through tests
- Set security guidelines and control parameters



Finding the Right Migration Solutions

This is where the heavy lifting gets done: the migration solutions necessary to bring an enterprise workload into the cloud need to be powerful. Deciding on the wrong solution can set back the migration and even lead to further headaches.

There are multiple ways to move data to the cloud. Some are native to the cloud service providers, others are open source, and there are also third-party vendor solutions and services. There are also solutions unique to specific use cases, such as databases and virtual machines. Migrating data files is one of the most critical and challenging moves that has to be orchestrated. When it comes to data, not only does the initial migration have to be considered, but also keeping that data up to date and in sync with sources on the premises and in backup locations. Time and costs are real factors to consider.

Database migrations, for example, can be done with the help of cloud-native services such as [AWS Database Migration Service \(DMS\)](#) for moving to Amazon RDS from MySQL, MS SQL, or PostgreSQL, and [Azure Database Migration Service](#), which can move SQL Server or on-premises Oracle databases to Azure's SQL options. Google offers similar services for SQL migration with its fully managed [Google Cloud SQL](#) and [Cloud Spanner](#) services

For the largest migrations—for data that exists on a scale that would take years or even decades to transmit electronically—AWS offers [AWS Snowball](#), [AWS Snowball Edge](#), and [AWS Snowmobile](#). Both Azure and Google Cloud Platform provide comparable offline bulk data transfer services, namely [Azure Data Box](#) and [Google Cloud Transfer Appliance](#).

As the actual migration takes place, it is important to make sure that there is no interruption of normal business for your workloads. Data must continue to be accessible to all customers, and updates to existing data must continue to take place as normal. The process should be carried out as quickly as possible, but can effectively be broken down into phases that see each workload element successfully migrated and validated before moving on to the next. You'll also need to find a way to synchronize changes that are made to the source data while the migration is ongoing. A good way to manage the task is to employ data management solutions that can be found on your respective public cloud's marketplace. Existing NetApp users benefit from having Cloud Volumes ONTAP, which extends enterprise-level on-premises storage into the cloud through NetApp SnapMirror® technology. In the next section we'll look in detail at what Cloud Volumes ONTAP and NetApp can offer enterprises migrating workloads to the cloud.

Key points to have in mind before migration day:

- Find the right solution to carry out the migration and provide support during the move
- Have plans in place for when the migration takes place, including contingencies for if things go wrong
- Test as you go to make sure that everything works

NetApp Solutions for Migrating Enterprise Workloads to the Cloud

Because migrating an enterprise workload to the cloud requires massive amounts of orchestration and support, many companies turn to solution providers such as NetApp to aid in the migration and for continued data management and support of their enterprise workloads.

[NetApp Cloud Volumes ONTAP](#) offers enterprise businesses a way to seamlessly transition their workloads into the cloud. Utilizing SnapMirror, Cloud Volumes ONTAP replicates files from on-premises NetApp storage systems and brings them into the cloud. Available on AWS, Azure, and Google Cloud Platform, Cloud Volumes ONTAP is an interface similar to the one that longtime NetApp storage system users are familiar with using, except now all of that functionality has been updated for performance in the cloud. With the ability to support SMB and NFS file shares, and iSCSI SAN storage, enterprise workloads can effectively leverage the cloud for all of their business demands.

[Azure NetApp Files \(ANF\)](#) is a fully managed Azure cloud-native file storage service that provides NAS volumes over NFS and SMB with all-flash performance. The service is integrated with Azure portal and accessed via NetApp Cloud Manager, REST API and Azure SDKs. Customers can seamlessly migrate and run applications in the cloud without worrying about procuring or managing storage infrastructure.

ANF simplifies storage management in Azure with NetApp out-of-the-box capabilities including file sharing, multiprotocol support, high availability, data protection, and more.

[NetApp@ Cloud Volumes Service for Google Cloud](#) (CVS for GCP) is a fully managed GCP cloud-native file storage service that provides NAS volumes over NFS and SMB with all-flash performance. The service is integrated with Google Cloud console and accessed via NetApp Cloud Manager and REST API.

CVS for GCP simplifies how you migrate and run enterprise workloads in Google Cloud with out-of-the-box capabilities such as file sharing, multiprotocol support, high availability, data protection, and more.

At the enterprise level, data management requires a seamless way to orchestrate the cloud environment from a single pane-of-glass, where resources can be easily launched and deployed with the click of a button.

For that, Cloud Volumes ONTAP comes with NetApp [Cloud Manager](#), the NetApp automation, orchestration, and management GUI. From Cloud Manager, tasks can easily be carried out through a drag-and-drop interface that connects, discovers, and manages resources throughout your deployment, both on premises and in the cloud. Scheduling, monitoring, and alert tools are all accessible through a single interface, so migration and maintaining a cloud deployment do not mean a difficult management process.



In a migration, it's crucial to understand which data is moving and where. [NetApp Cloud Data Sense](#) is a data governance and classification tool that can automatically discover, map, classify, and act on enterprise data, wherever it's stored—on-prem or in the cloud. With [Data Sense](#), users gain full visibility and control over data based on automated and AI driven analysis of metadata and data context.

For companies planning migrations, these capabilities are particularly useful in determining where private data you have is stored so you can take the proper precautions to protect it, in light of privacy regulations. It also makes it possible to pinpoint duplicate and stale data so you can carry out the kind of clean migration that will reduce your costs to migrate to and operate in the new environment.

[Sign up now to try Cloud Data Sense free for up to 1 TB of data.](#)

[Cloud Sync](#) is another NetApp solution for migrating data (from systems other than ONTAP) to the cloud. As a file transfer solution for companies migrating data to the cloud, Cloud Sync offers much more than open-source tools such as rclone and rsync, because it comes with a service's robust set of features. With automation for ongoing file transfers, parallel processing for the fastest transfer speeds, and data protection that never takes the data out of your security boundaries, Cloud Sync turns the movement of data into a task that companies can expect to be done affordably and quickly.

The [NetApp Data Migration Services for Cloud](#) helps customers to determine the best cloud transition strategy and it tailors and executes large-scale data migration plans to meet changing business needs.

Whether it's Hybrid Cloud, Lift-and-Shift, or Cloud First, the NetApp Data Migration Services for Cloud helps customers migrate from legacy, on-premises environments to public cloud-native storage with minimal business disruption and while controlling migration costs and sustaining data protection and regulatory compliance policies.

The packaged services include **Discovery, Planning, Data Movement, Testing/Validation, and Project Management.**




Cloud Volumes ONTAP: Key Features

Cloud Volumes ONTAP has a suite of features that make [deploying enterprise workloads](#) using the cloud easy, cost efficient, and safe, including:

-  **High Availability.** Two-node high availability with Cloud Volumes ONTAP makes sure that when your enterprise workload faces an outage, the redundant node takes over, keeping your RPO at zero and RTO below 60 seconds.
-  **Data Mobility.** [SnapMirror® data replication](#) seamlessly transfers data from on-premises or other clouds, avoiding vendor lock-in. SnapMirror also helps to set up and continuously synchronize DR data copies.
-  **Data Protection.** Cost-effective [NetApp Snapshot™](#) copies and seamless [disaster recovery capabilities](#) keep your files safe from failures or data corruption.
-  **Hybrid and Multicloud.** Cloud Manager gives users a single pane of glass to seamlessly manage, maintain, and monitor resources that span environments on AWS, Azure, Google Cloud, and the data center with drag-and-drop functionality or API calls.
-  **Automation.** [RESTful API calls](#) allow developers to treat Infrastructure as Code (IAC), speeding up the dev/test workflow and cutting down TTM.
-  **Storage Efficiencies.** Cloud Volumes ONTAP makes it possible to cut down on cloud storage costs through the use of [several storage efficiency features](#), including thin provisioning, data deduplication, compression, and zero-capacity data clones.
-  **Data Tiering.** Automatic tiering of “cold” data between highly performant disk storage to less-expensive object storage on Amazon S3, Azure Blob, or Google Cloud Storage saves costs and optimizes storage.
-  **Multiprotocol Access.** Access to storage over protocol of choice—NFS/SMB/ iSCSI—allows the same storage service to be configured and used for file share as well as block storage use cases.
-  **High Performance.** Enhanced storage throughput through features like WAFL and cloud data caching, best suited for performance intensive production workloads.
-  **Kubernetes and Containers Integration.** With NetApp Trident, Cloud Volumes ONTAP can be used to automatically provision persistent storage for stateful applications stored in Kubernetes and other containerized environments.
-  **Security and Safety.** Cloud Volumes ONTAP augments and integrates with the cloud provider storage features like encryption at-rest and in-transit, VNET integration for perimeter security, ransomware protection, and cloud WORM (write once, read many) storage
-  **Compliance.** The AI-driven Cloud Compliance add-on scans all your data in order to map, identify, and report on sensitive private data that could fall under regulatory scope of GDPR, CCPA, and other data privacy legislation.

The Final Stage: Keeping Tabs on the Cloud



The cloud is always changing. It's one of its biggest features, but also a concern for enterprise businesses that have to keep up. The upkeep of your cloud deployment requires constant monitoring and awareness of these changes. Make sure that you have round-the-clock support both for your resources with the cloud provider and with your own deployment.

Your SLAs with the cloud provider should be carefully monitored. SLAs are contract-bound expectations of service that cover everything from compliance to security and performance. Not every product from the same provider has the same SLA, and it is important to be aware of these. The SLA for Amazon EC2, for example, clearly says that every reasonable effort will be made by AWS to provide the service at a 99.95% availability. That in itself is not a guarantee, and that is where monitoring on your part has to take place.

Keeping tabs on the cloud means employing monitoring services. On AWS there are [AWS CloudTrail](#) and [Amazon CloudWatch](#), while Azure provides this service through [Azure Application Insights](#) and [Azure Monitor](#). Google provides [Cloud Monitoring](#), which aims to help users monitor their cloud deployments, as well as identify and diagnose issues. NetApp users can also turn to Cloud Insights to get in-depth analysis into the performance of their cloud services and applications, not just the storage level.

DR is another long-term planning goal for operating in the cloud. Before the advent of the cloud, disaster recovery

for enterprise workloads meant maintaining secondary and sometimes even tertiary physical backup sites to make sure of data safety and compliance. The cloud still provides that level of redundant protection, but in this case, you no longer have the outlays of real estate, maintenance, orchestration, security, and environmental controls involved with running the backup sites. Data stored for backup purposes in the cloud can be ready to use within seconds of a disaster scenario, making sure of business continuity when you would otherwise be in danger of missing your RPO and RTO. Cost control is the main concern when it comes to storing data for DR, and Cloud Volumes ONTAP is designed specifically to do that. With storage efficiencies; easy replication to your backup sites using SnapMirror; and tiering cold data to inexpensive object storage until it's needed, Cloud Volumes ONTAP can play a crucial role in maintaining a cost-effective DR plan.

Enterprise workloads have a particularly challenging set of requirements, most of which have to deal with their scale. The sheer amount of data that has to be stored can eat up an entire IT budget. Managing to keep these costs as low as possible is an ongoing task for any enterprise workload in the cloud. More than just business continuity, your industry might have compliance requirements and regulations that demand a level of redundant storage for certain types of data, no matter how much that costs your company. Once again, a solution such as Cloud Volumes ONTAP makes meeting these goals achievable and effective.

Points to keep in mind moving forward in the cloud:

- Be ready to adapt to constantly changing platform updates
- Provide vigilant system monitoring and upkeep
- Keep on top of SLAs to make sure you always get the performance you pay for and require
- Adjust solutions to keep storage and disaster recovery costs at a minimum

Cloud Migration Case Studies with Cloud Volumes ONTAP

Cloud Volumes ONTAP has helped thousands of companies in hundreds of industries take control of their cloud migration and deployment.



The Cordant Group

[The Cordant Group](#) is a UK company that started in security services and later expanded into facilities management as well as recruitment. With 190+ locations and 5,000+ clients throughout the UK—and growing—the company has an extensive IT infrastructure.

As part of a strategic decision to shift from a CAPEX to an OPEX model, the company had already migrated its infrastructure to a private isolated section on AWS. However, the company found itself restricted to the use of Windows DFS to manage its server resources. Another challenge was related to SQL databases that held business-critical information.

As a veteran NetApp customer, they were able to use SnapMirror technology to easily replicate data from its physical storage infrastructure to the AWS public cloud and reap the following benefits:

- NetApp Snapshots for quick and cost-effective backups in the cloud that complies with a rigorous data restore objective of less than an hour.
- Leverage [SnapManager® for SQL](#) for crash-consistent, application-aware backup copies with little or no impact on their storage resources or application performance.
- A robust and cost-effective cloud disaster recovery solution with automatic failover and failback capabilities when things go wrong.
- The ability to provision both NAS and SAN storage, with support and multi-protocol access for SMB/ CIFS, NFS, and iSCSI.
- Single-pane centralized management of their hybrid architecture via NetApp Cloud Manager.

In short, by migrating and then managing their data with Cloud Volumes ONTAP, the Cordant Group achieved the scalability they needed to grow their business, realized storage efficiencies and cost savings of 40-50%, gained the flexibility of a hybrid architecture, and could rest assured that their data was protected wherever it resided.

[Read the full case study here](#)



McKesson

McKesson is the largest healthcare company in the United States, distributing one-third of all medications across North America. The company needed a platform for developing, testing, and deploying their fulfillment and logistics applications at scale and at speed, but was feeling constrained by the managed private cloud solution it was using. They found the scale and advanced data management features they needed in Azure and NetApp.

The McKesson IT team adopted a data fabric strategy. The team uses NetApp AFF and NetApp E-Series flash storage in its data center while its public cloud is powered by Cloud Volumes ONTAP in Azure. Today, McKesson benefits from:

- Improved cloud performance at a much lower cost.
- More responsiveness, so its IT team can effectively address the needs of developers and business units.
- Disaster recovery as a service (DRaaS) by using Azure Site Recovery to replicate applications and virtual machines.
- Seamless integration across the hybrid deployment with SnapMirror and Cloud Volumes ONTAP replicating on-premises data to Azure.

[Read the full McKesson case study here](#)




AdvancedMD

[AdvancedMD](#) provides a SaaS medical office workflow solution that automates manual front office and billing processes, maintains always-accessible electronic medical and health records, and promotes patient engagement. With AdvancedMD deployed, physicians can focus on improving patient care, with the added bonus of increased revenue because their practice is running more efficiently.

When AdvancedMD was spun off from its parent company, ADP, they had to quickly move all their data out of the ADP data center. As ADP was an existing user of NetApp storage technology in-house, AdvancedMD decided to use Cloud Volumes ONTAP for AWS for moving to the cloud. The key benefits that they gained were:

- Easy lift and shift of their SQL database and their file shares, with no need to change their infrastructure architecture.
- High availability dual-node redundancy for AWS and for Azure to provide minimal downtime in failures or upgrades, a critical requirement when dealing with important healthcare data.
- Storage efficiencies, such as data deduplication, compression, and compaction as well as thin provisioning, automated data tiering, and instantaneous data cloning to writable volumes with zero impact on storage footprints.
- The intuitive Cloud Manager interface, with simple lift-and-shift drag-and-drop operations, as well as support for data management automation and orchestration.
- During the migration and afterwards: Enterprise-level data protection and encryption (especially important for sensitive healthcare data, HIPAA regulations) without sacrificing scalability or agility.

[Read the customer blog here](#)




Summary

This handbook provided a good idea of what can happen along the way to bringing an enterprise workload into the cloud, whether that's on AWS, Azure, or Google Cloud Platform. Hopefully with the answers you found here, you'll be able to make the next steps toward the cloud. For every step of the enterprise workload migration process, we've also shown how NetApp has a solution that is proven to help make the transition safe, cost-effective, and fast.

Migrations pose tough challenges that can be met with a careful strategy and the right resources at hand. With NetApp solutions such as Cloud Volumes ONTAP, Cloud Manager, Cloud Sync, and Cloud Insights, you'll find the service and solution you need to effectively make the transition of your enterprise workload to the cloud a success.

**Start a free
trial today with
Cloud Volumes
ONTAP**

Start now



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 © 1994–2022 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

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.

NA-000-1221