EBOOK

The Cloud Migration Path for Financial Companies

NetApp

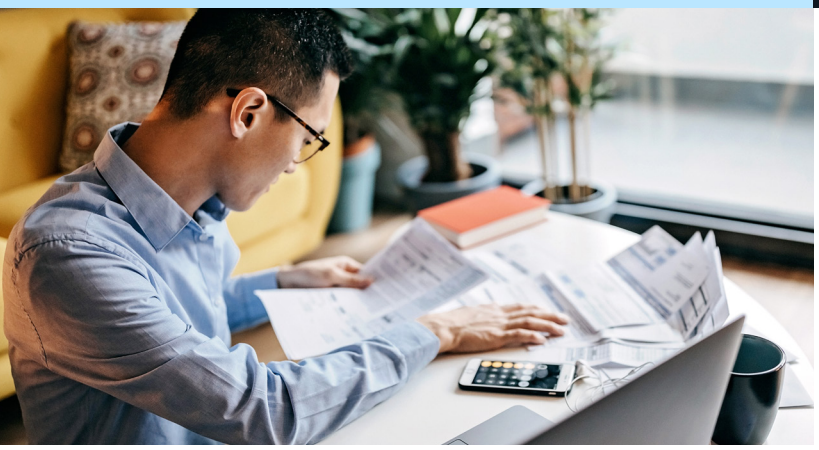

Executive Summary

When it comes to moving workloads into the cloud, Financial Service Industry (FSI) companies have a different starting point. They have a challenging set of constraints that frame every technology they adopt. These include regulations, data security, geographical storage requirements, governance/oversight, and more.

But there are also powerful motivators driving financial companies into the cloud, among them:

- Embrace tech like Artificial Intelligence and Data Science Machine Learning platforms (AI/DSML) to innovate.
- Scale operations fast to respond to market demands.
- Achieve profitable insights quickly.
- Increase workload resilience.

What does the migration path look like for your financial company? The five parts of this eBook will help you to navigate the options.

This includes insights into what industry leaders have already done:

1. **Approaching cloud adoption:** A checklist of concerns common to most financial companies, defined and illustrated, to help you understand your path into the cloud.
2. **Fundamentals of cloud migration:** You have your cloud adoption plan. Now, how do you get your data and workflows out there? Learn about the fundamentals and options to accomplish this in a secure and cost-effective manner.
3. **Benefits of a cloud data management solution:** There is so much innovation happening with cloud technology that once you get into the cloud, you'll need to adjust to the trends and innovations that keep developing. Crucial to your success is to avoid cloud lock-in, where your company becomes dependent on the offerings of one cloud platform provider. A cloud data management solution acts as a layer on top of the cloud platform, allowing your IT and application teams to use best-of-breed features in your solutions.
4. **Keeping Tabs on the Cloud:** Final thoughts on what it means to operate in the cloud, how the landscape changes quickly with time. Your company's staff must remain agile to monitor cloud tech and apply the best-fit services to your workloads.
5. **Case studies:** A selection of use cases for various types of financial companies. See how they approached their path to the cloud and the enabling technologies they used to get there.

This guidebook will help you understand the unique needs of your financial company so you can build a path to migrate workloads into the cloud. It will also provide some insight into how you can do that with the help of NetApp Cloud Volumes ONTAP, the data management platform for workloads in Azure, AWS, and Google Cloud.

Table of Contents

Introduction

Companies in the financial industry have typically built their IT infrastructure over generations of technology, acquisitions, and mergers. Regulations require that they store and secure their data in the geographical locations where the systems and services exist.

In many cases, these companies use disparate infrastructures that don't interact, creating data silos that prevent achieving the insights across the enterprise that can drive competitive innovation.

Financial companies, therefore, are very much on-premises IT operations. While companies in other industries have gone fully into the cloud, established financial companies cannot realistically do that to the same degree.

But there's "gold" in the data silos scattered across these enterprises. And there are two competitive threats applying pressure to adopt cloud tech to mine for that gold:

1. Early adopters that have explored and are growing proficient with cloud tech.
2. A new generation of small, agile challengers that provide a specific service to underserved markets and do it well.

PWC's Financial Services 2020 and Beyond report states that established financial companies risk losing 25% of their business to agile startups. Also, the early adopters have already unified their data in the cloud and are now achieving the insights that allow them to become more competitive.

Financial companies that wait to embrace cloud tech risk getting left behind. IDC's Driving Digital Transformation in Financial Services report labels early adopters as "thrivers" and those who have yet to adopt as "resistors." It describes the world of cloud tech as rapidly changing, where thrivers embrace a process of continuous evaluation, plugging in new technologies and casting others aside as services evolve. It's best to jump on this train as soon as possible.

If your company has been cautious about the approaching the cloud, this eBook will offer some important insights for you. It will help you assess cloud tech for financial companies so that you can create an adoption path that makes sense for your business model.

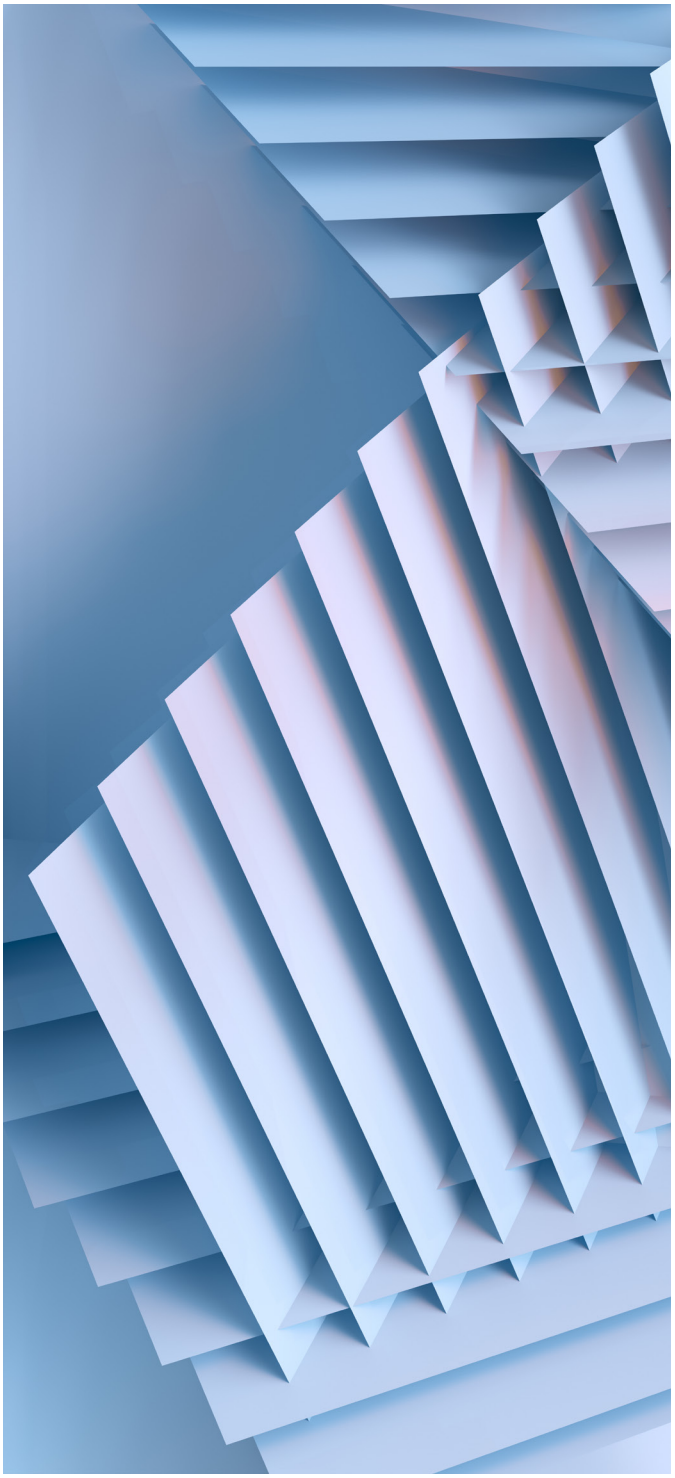

Part One

Approaching Cloud Adoption for Financial Companies

Before you can plan how to make cloud tech work for your company, you must first assess your business model and choose the workloads that make sense to move into the cloud. The plan should clearly identify what to move first and how the end goal will look.

You can achieve this by working through the following categories.

Benefits Motivating Cloud Adoption

What do you want to accomplish with cloud tech? To answer this, it helps to look at what the leaders are doing. According to the PWC report, the main driver is to leverage data with AI/ML tools to accomplish these goals:

- **Personalized offerings** to match customers' wants and needs. Financial companies must provide the web and mobile UX that users have become accustomed to from other industries.
- **Improved services** targeting primarily security and anti-fraud detection. Also, automating decisions for investing and underwriting frees skilled workers to contribute to risk management, product development and other expert advisory functions.
- **Enhanced risk analysis** by using AI to lower risks by monitoring increasing volumes of both structured data (databases) and unstructured data (social media, news). AI can assess risks in transactions, trading patterns, profitability, and more.
- **Understand and target the most profitable customers** by unifying data across the enterprise and using AI to derive segments that identify customers' needs. Marketers can then use these segments to create user experiences with personalized choices.

The ML platform market continues to innovate with a wide variety of offerings to use for implementing these innovations.

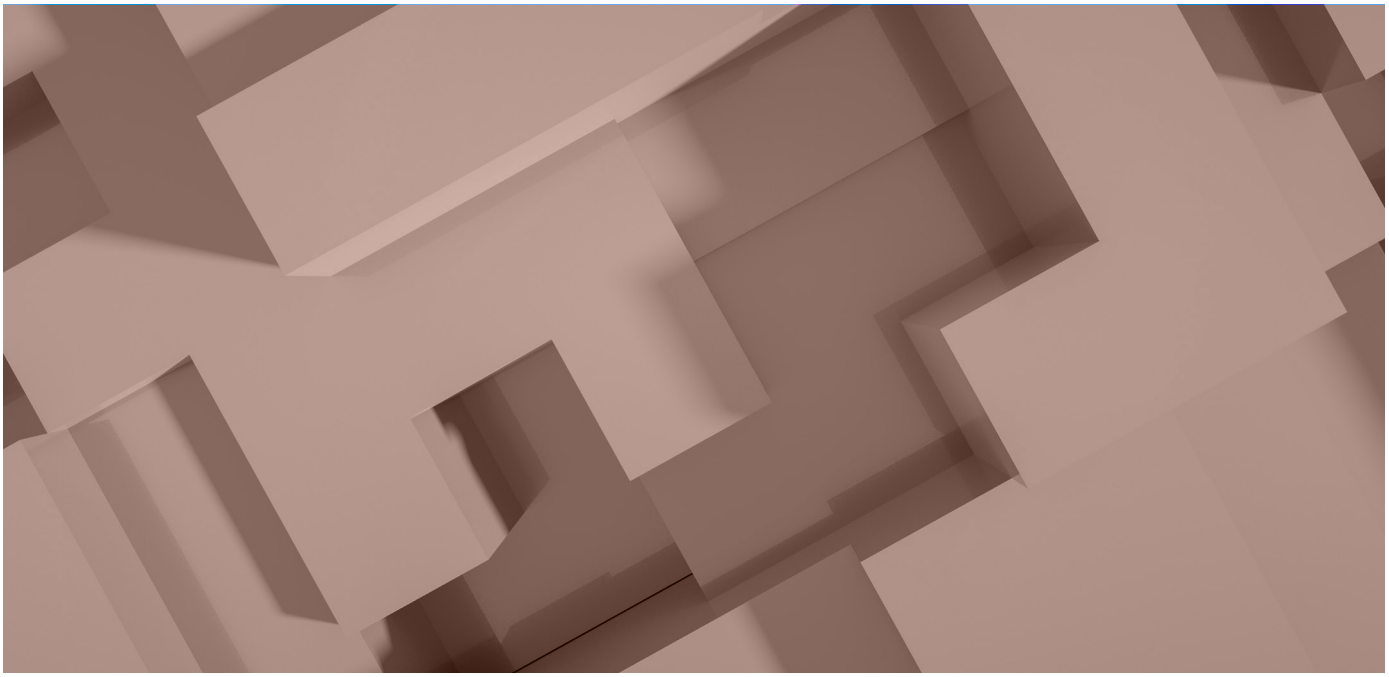

Using Cloud Tech to Derive Value

How does cloud tech enable financial companies to derive these benefits? A Deloitte on Cloud Banking: Financial Services and Banking of the Future cites these trends:

- **Data unification** using on-premises systems is complicated. Moving data to the cloud lets you share data between business units, enabling you to solve customer problems more quickly.
- **Connect data sets** to use with analytics to achieve insights for faster decisions. This becomes possible with unification. Once the data is all in one place, you can use cloud tools to discover and build relationships between the datasets.
- **Attract new talent** by providing access to modern cloud tech. Once you've invested in a pool of cloud tech savvy talent, they will keep the innovation and knowledge base growing as cloud tech continues to evolve.
- **Build resilient operations** unaffected by outages and physical disruption by making use of the cloud vendor's failover and load-balancing capabilities in the cloud.
- **Ability to respond to market shifts** by auto-scaling. One of the first big wins for the cloud was the ability to set metrics that trigger the cloud platform to scale up capacity to handle a bigger workload.

Cloud tech provides a scalable data laboratory with many tools. But financial organizations need to know that it's safe. Let's deal next with those concerns.

Using Cloud Tech to Address Performance and Governance Concerns

Financial companies have been slow to adopt new tech due to internal requirements and externally-imposed regulations. Cloud tech meets these needs and offers additional benefits.

- **Data visibility and governance** are mandatory requirements for financial companies. Regulators—now turning to technology to help solve their own problems—need access to financial institutions' systems to verify that all is well. Cloud tech helps financial companies comply.
- **Inconsistent data management** can arise once legacy, silo-based systems start to move components to the cloud. The tools you choose must impose the same data standards across all datasets. This makes unification easier to achieve.
- **Different products for different environments** were also typical for legacy systems, and one of the major barriers to data unification on-premises. The cloud vendor provides a common platform for all its tools.
- **Performance issues** can happen at many points. In legacy systems, performance bottlenecks were typically addressed by making new infrastructure expenditures, which was reactive, and sometimes costly. Cloud tech comes with standard monitoring tools that cover service availability, latency, throughput, application performance and more. Performance optimization is necessary to ensure peak performance.
- **Data security** is hard enough on-premises. The threat possibilities increase when you add the Internet of Things (IoT), mobile devices and third-party tools. It's natural to ask how secure your data is in the cloud. Cloud platforms help keep your data safe with built-in encryption and current firewall protections; however, financial companies need the highest levels of security and privacy, which may require solutions beyond those natively available from the cloud provider.

Cloud technology meets the concerns of financial companies with standardization and a service-level agreement to keep data safe. But financial companies have additional needs unique to the industry.

Addressing Security and Compliance Priorities

Financial data is among the most sensitive to manage. Financial companies must have solutions to secure the data and adhere to compliance regulations:

- **Data security** that includes backup, recovery, encryption and protection from external threats. Financial companies must have this on-prem and also in the cloud.
- **Ransomware prevention** is critical in the financial industry, as there are huge costs involved whether paying the ransom, violating compliance laws, or rebuilding systems from scratch if you're locked out.
- **Privacy of sensitive customer data** is now a mandate for all businesses that handle customer financial data.
- **Privacy of sensitive organizational data**. A similar set of concerns exist for a financial company's proprietary.
- **Compliance with industry regulations** such as PCI, Dodd-Frank, and GDPR bring with them their own record-keeping and enforcement requirements.

The cloud now offers numerous services that can help solve these problems. However, there are still many professionals in the financial industry who are reluctant to use these new services. This is because despite the standard functionality cloud-based services provide, they are usually very different from the existing on-prem services the financial companies have in use. That means replatforming or repurchasing to transition to these services can be complex. One advantage to being a resistor in this case is the ability to test these services with caution and adapt them to existing services. This way, when it is time to move to the cloud, the transition will be much easier.

Hybrid Cloud Architecture Goals

As mentioned in the introduction, established financial companies cannot fully move their data operations into the cloud. Most of your company's operations will remain on-premises. Most likely, workloads such as analytics and artificial intelligence/machine learning (AI/ML) and those that require unified data will be prioritized for migration to the cloud, either temporarily or permanently, while the rest of the deployment remains on-premises. This is what is known as a "hybrid cloud architecture."

According to the IDC study, the thrivers tend to do their most proprietary AI/ML processing in a private cloud that melds with the on-prem network, something like a VPN. Thrivers use the public cloud more sparingly, and employ Software as a Service (SaaS) to do the financial transactions there, since SaaS vendors already have security built in.

The benefit that comes with a hybrid private cloud implementation is data center experience that provides an end-to-end unified plane for the computation as well as the data. More specifically:

- **Single combined compute plane** and accompanying workload migration ability. This means the workloads happen in the tools and technologies provided by the cloud vendors, providing a comprehensive set of tools that work together.
- **A single control plane** and view of the environment. This means the tools provide a UX that lets you control the processing steps and environments acquired in the cloud, with full monitoring of performance and status.
- **End-to-end control over the IT environment**, including infrastructure, applications, and data. The hybrid private cloud connects all the pieces together so that the control plane manages the compute plane to carry out all the workloads.

The hybrid private cloud is the path to the cloud that's most compatible with the requirements and constraints of a financial company. This is where a cloud data management offering can become extremely useful.

Key Areas of a Cloud Data Management Layer

A cloud data management product provides a layer of abstraction over the cloud platform(s) that you'll use for the cloud implementation. Look for the following capabilities to start your path to cloud adoption:

- **File consolidation** means you can share on-prem files into the cloud and make them available via a global file cache for use in the cloud or in other on-prem locations.
- **Data analytics with cloud bursting** is how you will implement the workflow. First the data will get bursted into the cloud environment. Once there, analytics and AI can process it.
- **Provides data protection** by automatically backing up data stored in the cloud with protection against data loss, human errors, outages, cyber threats and natural disasters.
- **Supports DevOps** or the agile software development model, which rapidly brings services to the field. Your DevOps team uses a wide range of IT personnel (UX, UI, analytics, integration, big data, and AI/ML) who are trained to help accelerate the company's software development and release cycle.

These are the major features to keep in mind as you evaluate vendor offerings. Next, let's see how to move your workload into the cloud.

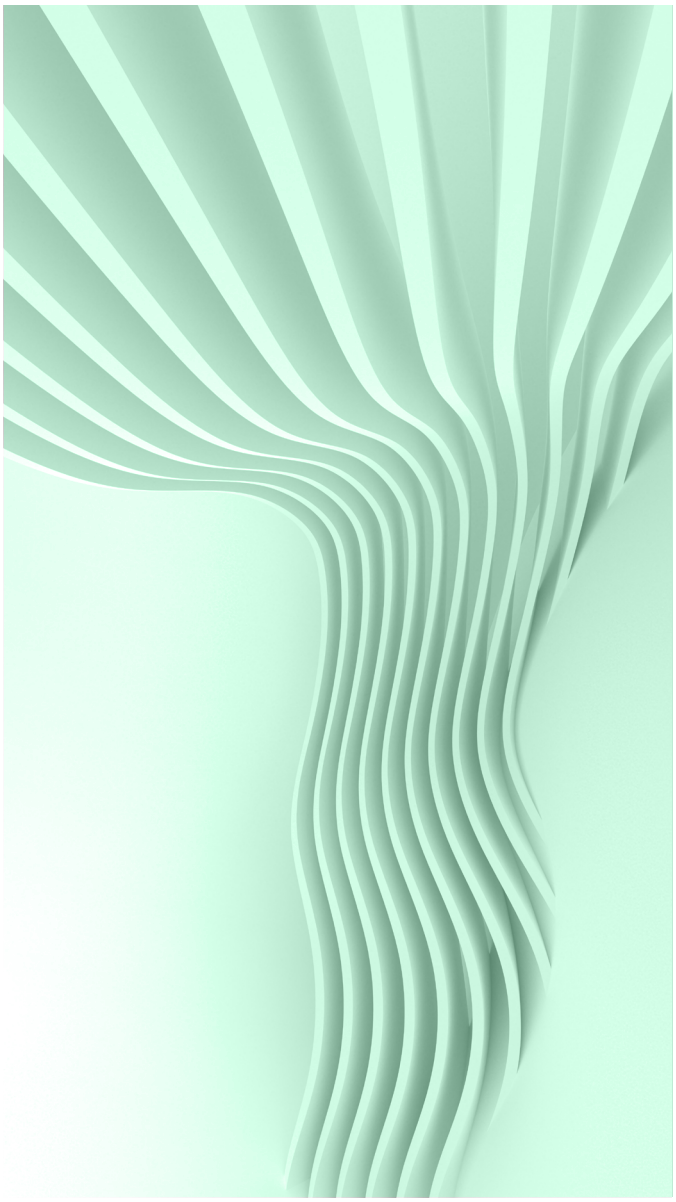

Part Two

Determining Workload 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.

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/s.

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 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.

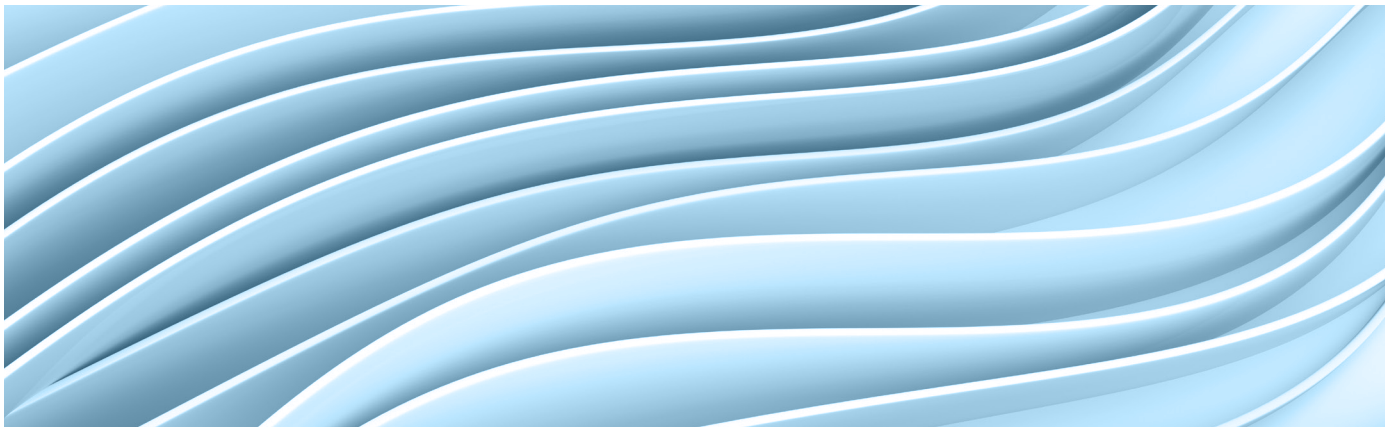

The service provider that you choose should also be consulted, because it can assist your move. Cloud providers have expert teams that are available to give advice and help you reorganize your architecture for cloud deployment, with security and compliance needs in mind. For existing NetApp storage system 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

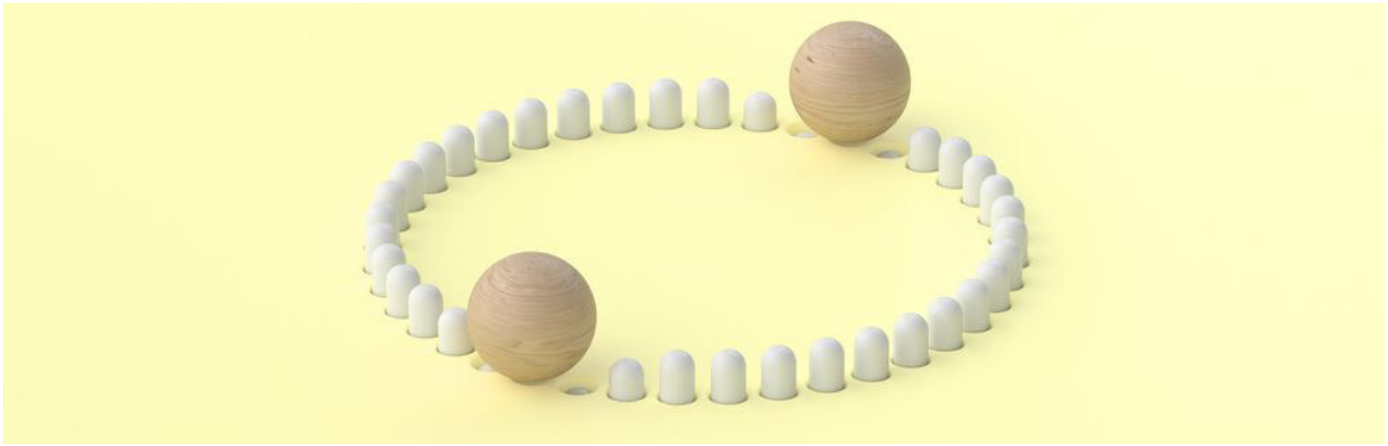

Finding the Right Migration Solutions

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

Part Three

Key Features of Cloud Deployment with Cloud Volumes ONTAP

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, as well as 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.

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.

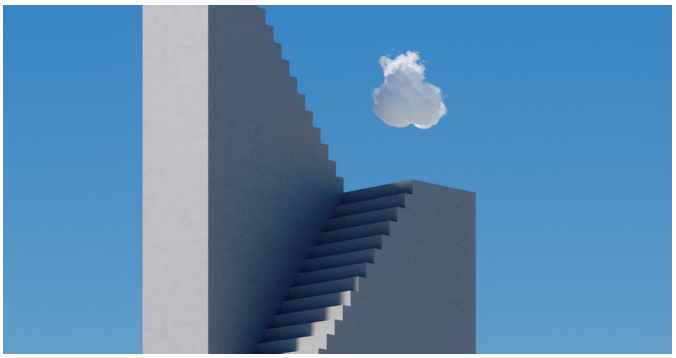

Cloud Volumes ONTAP has a suite of features that make deploying enterprise workloads using the cloud easy, cost efficient, and safe, including:

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 protection.
Cost-effective NetApp Snapshot™ copies and seamless disaster recovery capabilities keep your files safe from failures or data corruption.

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.

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.

Cloud Bursting.
In deployments that are primarily on-prem, cloud bursting allows companies to take advantage of the unlimited processing power available on-demand in the cloud. Cloud Volumes ONTAP includes FlexCache®, a unique caching technology that allows for a fast and cost-effective way to perform processing-intensive jobs on your on-premises data.

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.

High Performance.
Enhanced storage throughput due to intelligent optimization of costs/performance so that each block of data gets the performance it needs without overspending.

File Services and Caching.
File services with access to storage over NFS and SMB/ CIFS, allowing the same storage service to be configured and used for every use case. Cloud Volumes ONTAP unifies this data and NetApp Global File Cache ensures fast access to the consolidated data from remote branch offices.

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.

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.

Automation.
RESTful API calls allow developers to treat Infrastructure as Code (IAC), speeding up the dev/test workflow and cutting down TTM.

Part Four

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

Part Five

Financial Company Case Studies

Multinational Insurance and Financial Services Provider

This insurance and financial services provider serve over 26 million customers worldwide and has approximately one trillion Canadian Dollars in assets under management and administration.

They started their journey on Azure in 2017. As their cloud adoption grew, they Enterprise Technology Services team realized they would need a centralized platform that would allow them to attain better control over their cloud data and its associated costs. They found that solution in Cloud Volumes ONTAP.

The company moved hundreds of TBs from Azure Files to Cloud Volumes ONTAP which now serves as a consolidated data platform for multiple use cases, including Production Workloads, File Shares, DevOps, Archiving and Disaster Recovery.

Cloud Volumes ONTAP enabled this company to optimize their cloud costs in several ways:

- **Storage efficiencies:** Deduplication, compaction, compression, space-saving snapshots and clones reduce Azure storage and associated costs.
- **Thin provisioning:** Allows IT to service continuous requests for large storage capacities without actually acquiring and paying for the underlying cloud storage, only the capacity that's used.
- **Data tiering:** Optimized costs and performance by setting different policies, tiering log files, archives and disaster recovery environments to Azure Blob Storage to minimize costs.

American Multinational Investment Bank Hedge Fund

This hedge fund division belongs to a multinational investment bank and financial services company which identifies, models, and trades global financial markets by leveraging research, technology, and automation. The division uses cutting-edge technology to formulate risk models for their clients, powered by **Google Cloud**'s elastic grid computing.

The team wanted to find a way to perform its existing research analysis and simulation processing in Google Cloud. This required mirroring the primary on-premises data in Google Cloud. For elastic scaling and burst consumption they needed an NFS platform on Google Cloud that is scalable, easily managed, and integrates into their existing automation frameworks and processes.

Cloud Volumes ONTAP was the perfect solution to meet their needs and led to some major benefits:

- Quick set up using Cloud Manager's drag-and-drop capabilities.
- Seamless replication of data from on-premises to the cloud and back using SnapMirror.
- Integration with existing automation tools through the Cloud Manager API.
- Storage efficiencies that reduce cloud data storage costs of a massive 80 TB dataset.

Large Financial Institution

With origins as a government-sponsored institution, today this company is a fully privatized bank that owns nearly one-fifth of the student debt across the United States. Offering more than 500 savings plans and operating across several states, this massive financial institution has annual revenues of ~$1.6 billion and employs 1,500 people nationwide.

The company had made a strategic decision to transition workloads from their on-prem data center to the public cloud. However, they also needed to maintain a consistent Data Fabric across which data could flow seamlessly between their on-prem and AWS public cloud storage resources.

The first task was to replicate more than **50 TB** of production **WORM (Write Once, Ready Many)** data to Amazon S3 object storage from their NetApp on-prem storage arrays. They were delighted to see that they could accomplish this quickly and securely using the **Cloud Sync** automated data synchronization service.

Conclusion

Companies in the financial industry have a set of constraints that must be considered when adopting any new technology. This is true for migrating workloads into the cloud. Each such company must factor its unique business model when planning how to leverage the cloud.

The path takes advantage of the cloud tech by building a hybrid cloud approach, keeping data on-prem and geographically located, secure and available to oversight, and also bursting the data into the cloud so that it can be combined across silos to achieve the innovation and competitive insights enabled by analytics and machine learning.

Financial companies can build a seamless plane of control and processing—while avoiding cloud lock-in—by making use of a multi-vendor product like Cloud Volumes ONTAP that provides a level of abstraction over cloud vendors, while also offering powerful functionality of its own.

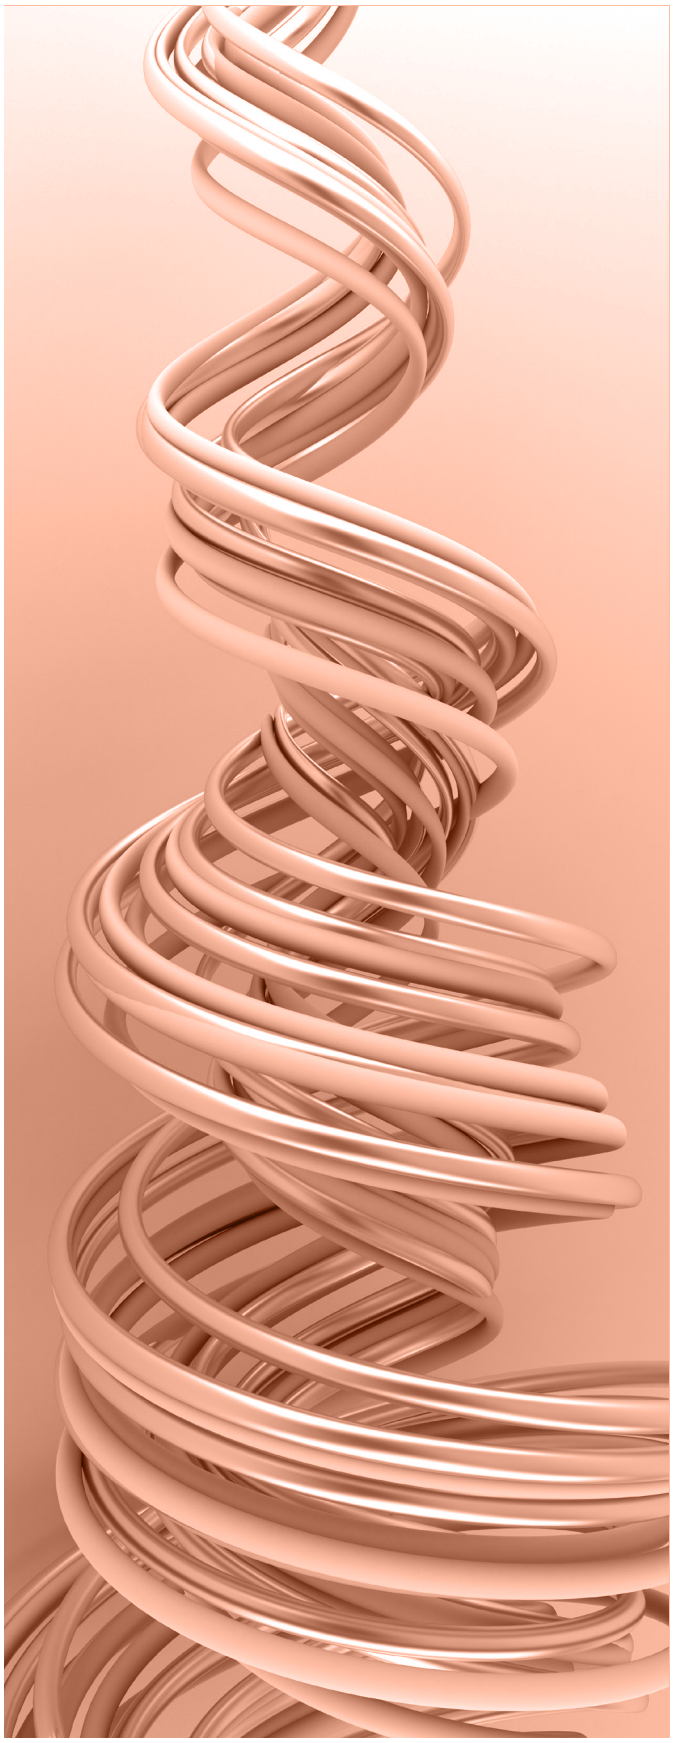

Start a free trial today with Cloud Volumes ONTAP on AWS, Google Cloud, or Azure.

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.

NA-000-1220