White Paper: **February 2023**

Comparing Services from the Big Three Cloud Providers

Table of contents

Introduction

Compiling a cloud services comparison is a daunting task in the rapidly evolving cloud environment. There are thousands of cloud services, dozens of cloud service providers, and numerous infrastructure-as-a-service (IaaS) providers offering pay-as-you-go pricing models—each one frequently changing and upgrading their portfolios.

This white paper provides a comparison of the cloud services from the top three IaaS providers—Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP)—because not only are these the IaaS providers most organizations are familiar with, they are also the IaaS providers whose services most organizations are likely to compare.

AWS, Azure and GCP all offer a very similar range of products and services. If one provider launches a new product ahead of its competitors, you can be almost certain the other two will soon release a comparable product or service, just with a different name. What generally differentiates similar products across different providers is how that given product interacts with other services within a specific provider's product portfolio.

This white paper provides an introduction to the range of services offered by the leading cloud service providers, information on regions and availability zones, a breakdown of cloud storage services, and more to help inform and equip you on your multi-cloud journey.

Cost and discounts

Cost comparisons between the big three cloud providers tend to be meaningful only for a limited time due to frequently changing prices, new product launches, and the increasing choice of discount programs. It's also the case that as cloud maturity increases, organizations change the ways in which they use the cloud, rendering previously useful cost comparisons irrelevant.

However, it is worth discussing the increasing choice of discount programs and the services they relate to. While this was once limited to AWS Standard Reserved Instances and Microsoft Enterprise Agreements, the days of this simplicity are far behind us.

Cloud service providers now offer various discount options for multiple services to help meet customer demands and stay ahead of the competition.

AWS was the first to offer a discount program in 2009 with Reserved Instances. Reserved Instances offer discounts on committed use that vary according to the length of the commitment and the manner of payment. Over time, Reserved Instances have become more flexible and convertible, and in 2019, AWS launched Savings Plans. The primary difference between the two programs is that Reserved Instances offer a discount against on-demand pricing depending on committed utilization, whereas Savings Plans offer a discount depending on committed spend.

Currently, the services covered by Reserved Instances include Amazon Elastic Compute Cloud (EC2), Amazon Relational Database Service (RDS), Amazon Redshift, Amazon ElastiCache, Amazon Elasticsearch, and Amazon DynamoDB, while Savings Plans include EC2, AWS Fargate, and AWS Lambda services. However, AWS is expected to extend these programs to cover more services in the future.

In 2017, Microsoft launched Azure Reserved Virtual Machine (VM) Instances, giving organizations that were not in a position to take advantage of Enterprise Agreements an opportunity to reduce costs by committing to a specified use over a one- or three-year term. Instance size flexibility and the Azure Hybrid Benefit program followed shortly after, and in 2020, Microsoft removed the requirement to pay all upfront for Reserved VM Instances, with no loss of discount.

Microsoft has since extended its no upfront committed use reservations program to include database services, analytics, object storage, and block storage. Although conditions apply to how these reservations can be used (for example, you have to reserve a minimum of 1TiB of premium SSD block storage to be eligible for a discount), the number of available discount programs is growing quickly.

It was also in 2017 that GCP launched its Committed Use Discount (CUD) program, a strong competitor to AWS Reserved Instances given that CUDs offered greater discounts upfront than AWS' no upfront payment option and did not require any upfront payment costs on their own. CUDs also applied to customizable VMs, giving customers greater flexibility.

While GCP's discount program applies to fewer services than AWS, Google Cloud also introduced sustained use discounts (SUDs) for most types of compute services. SUDs are applied automatically when services are running more than 25 percent of the month. Although the percentage savings are modest, it makes many of GCP's compute services more price-competitive than AWS' on-demand instances.

Table 1 provides an overview of the different committed use discount offerings available from AWS, Azure and GCP for their compute services.

Table 1: Comparison of committed use discount offerings[1]			
	AWS	Azure	Google Cloud
Commitment length	Reserved Instance: 1 or 3 years Savings Plans: 1 or 3 years	Reserved VM Instance: 1 or 3 years	Committed Use Discounts: 1 or 3 years
Cancellation options	Standard Reserved Instances: Available to see on Amazon Marketplace Convertible Reserved Instances and Savings Plans: No	Yes, with 12 percent fee on remainder value of Reserved VM Instance	No
Flexibility	Convertible Reserved Instance: Can be exchanged for Reserved Instances of equal or greater value Compute Savings Plans: Applied to any region or instance family	Can exchange for other Reserved Instances across any region and series	Discount is applied to all instances in the same region, but you can't change the size, family or region of the commitment after purchase
Payment options	All upfront (for highest cost savings) Partial upfront No upfront	All upfront Monthly (with no loss of discount)	Monthly
Potential cost savings compared to on demand	Up to 72 percent	Up to 72 percent (and up to 80 percent with Azure Hybrid Benefit)	Up to 70 percent

It's important to remember that because of the dynamic and ever-changing nature of the cloud, discount options offered by cloud service providers are constantly changing. Make sure you keep up to date with new cloud services and discounts, and integrate these into your greater cloud financial management practice.

1. This is a high-level comparison. For detailed specifications and pricing, please refer to each cloud provider's respective solutions pages.

Cloud compute services

Within cloud compute services, we find virtual machines (VMs/instances) configured for general purpose, memory optimization, compute optimization, and storage optimization.

General purpose

General purpose VMs provide balanced CPU-to-memory ratios and are ideal for testing and development, small to medium databases, and low- to medium-traffic web servers. This category of VM includes burstable VMs that run workloads using a fraction of the maximum available CPU capacity, and save excess capacity to cope with temporary increases in demand.

One distinguishing factor between the general purpose VMs offered by AWS, Azure and GCP is how frequently each provider introduces new generations of VMs. Typically, later generations have faster processors that improve performance and reduce latency. In some cases, they also enable customers to rightsize to a smaller VM size to reduce costs.

Memory optimized

Memory-optimized VMs deliver high memory-to-CPU ratios suitable for relational database servers, medium to large caches, and in-memory analytics. In addition to regular memory-optimized VMs, all three providers in our cloud service comparison offer super-memory-optimized VMs for large enterprises that provide more storage per vCPU.

Microsoft Azure also offers memory-optimized VMs with constrained vCPUs. These allow you to constrain the vCPU count to one-half or one-quarter of the original VM size to reduce the cost of software licensing while maintaining the same memory, storage and I/O bandwidth for database workloads, such as SQL Server or Oracle. You can replicate this feature on GCP with custom VM types.

Compute optimized

Compute-optimized VMs have a high CPU-to-memory ratio and are good for medium-traffic web servers, network appliances, batch processes, and application servers. Typical use cases include:

- Scientific modeling
- Distributed analytics
- Machine/deep learning inference
- Ad serving
- Highly scalable multiplayer gaming
- Video encoding

There are different sub-classes of high-performance compute VMs within the compute-optimized category, depending on whether you are running high-graphics, AI, or deep learning workloads, or need to support memory-intensive computational workloads. It is important to note that, currently, not every high-performance compute service is available in every region.

Compared to Azure and GCP, AWS currently offers significantly more options for compute-optimized instances in terms of sizes and specifications. However, at the top end of the spectrum, it is worth noting that the Azure ND 96asr A100 v4 VM (92 vCPUs) has five times as much memory as either AWS' p4d.24xlarge instance (96 vCPUs) or GCP's a2-megagpu-16g VM (96 vCPUs) for a lower on-demand price.

Storage optimized

Storage-optimized VMs provide high disk throughput and I/O, and are ideal for big data, SQL and NoSQL databases. AWS and Azure are fairly similar in their range of offerings of memory per vCPU. GCP doesn't offer a designated storage-optimized VM option. Instead, GCP allows you to add one of the following to an existing VM:

- Zonal standard (HDD) persistent disks are the slowest but cheapest, and better suited to data processing workloads that primarily use sequential I/O.
- Regional standard persistent disks are the same as zonal standard persistent disks but with synchronous replication across two zones in a region.
- Zonal balanced persistent disks are suitable for most general purpose applications at a price point between that of standard and SSD persistent disks.
- Regional balanced persistent disks are the same as above but with synchronous replication across two zones in a region.
- Zonal SSD persistent disks are faster and more suitable for enterprise applications and high-performance database workloads.
- Regional SSD persistent disks are suitable for workloads that may not have application-level replication.
- Local SSD disks (available with SCSI or NVMe protocols) have much higher throughput and lower latency.

GCP's approach probably gives more choice over how VMs are configured for storage, but there are a few trade-offs. For example, local SSD storage is not automatically replicated, and all data on the local SSD may be lost if the VM is terminated for any reason.

Container services and serverless computing

Managed container services

Containers are a hot topic in the rapidly evolving cloud world. In fact, based on our analysis of actual public cloud spend in 2020, investment in containers and serverless technologies increased by 38.7 percent and 13.5 percent, respectively, between January and September 2020.

While there are many ways to deploy containers in the cloud, our analysis shows customers prefer the managed container services offered by the big three cloud providers. AWS, Azure and GCP each offer a suite of products and services to help organizations build, deploy and manage containerized environments, many of which offer similar services and functionality.

Amazon Elastic Kubernetes Service (EKS) is a managed service to run Kubernetes on AWS without needing to install, operate and maintain your own Kubernetes control plane. Although the most complex of AWS' container services, EKS offers businesses greater container portability.

Amazon Elastic Container Service (ECS) is AWS' container orchestration service that supports Docker containers. ECS may be more natural for infrastructures built on EC2 as it provides built-in integration with other AWS services, but it's not as intuitive for developers more familiar with Kubernetes.

AWS also offers the serverless container service, Fargate, which is best used for event-driven containerized microservices. Fargate is equally as secure, reliable and scalable as ECS and EKS, but it reduces the operational overhead of scaling, patching, securing and managing servers.

Similar to EKS, Microsoft offers Azure Kubernetes Service (AKS), a highly available, secure and fully managed Kubernetes service. For Azure customers, AKS reduces the complexity of managing Kubernetes by offloading much of that responsibility to Azure.

Microsoft previously offered a standalone container management solution, Azure Container Service (ACS), but retired this in early 2021 due to the Docker and Apache Mesos services offered by ACS being absorbed into AKS.

For event-driven containerized workloads and those that don't require container orchestration, Microsoft provides Azure Container Instances (ACI). This service is the equivalent to Fargate for Azure customers.

Similar to Azure, Google Cloud Platform offers one managed container service, Google Kubernetes Engine (GKE), and one serverless container service, Google Cloud Run. GKE provides a fully managed Kubernetes environment for deploying, managing and scaling containerized applications using Google infrastructure, while Google Cloud Run is an event-driven service that removes the overhead of resource provisioning.

Serverless computing and function as a service

Serverless computing/function as a service (FaaS) eliminates the necessity to provision, manage or scale resources by allowing developers to upload code that performs a short-lived function when it's triggered by an event. As businesses only pay for the milliseconds when the function is executed, serverless computing can significantly reduce costs.

AWS pioneered the FaaS application model in 2014 with AWS Lambda, and Azure and Google followed suit with their own solutions: Azure Functions and Google Cloud Functions. While AWS Lambda had a head start and is still considered the leading solution on the market, Azure and Google are catching up in terms of performance and new features.

Table 2 has a quick guide of the different container and serverless computing services offered by AWS, Azure and GCP.

Table 2: Container and serverless computing services

	AWS	Azure	Google Cloud
Managed container services	Amazon ECS Amazon EKS	Azure Kubernetes Service	Google Kubernetes Engine
Serverless container services	AWS Fargate	Azure Container Instances	Google Cloud Run
Function as a service	AWS Lambda	Azure Functions	Google Cloud Functions

Database services

AWS, Azure and GCP all offer managed databases as a service (DBaaS) and give businesses the opportunity to deploy SQL and NoSQL databases on unmanaged VMs if a required feature is not supported by the managed database service.

While each cloud provider tends to prioritize their own cloud-specific database services—which typically integrate better with their other cloud-specific services—all three support database services, such as MongoDB or Cassandra, that are suitable for use in multi-cloud environments.

Important differences to know

AWS, Azure and GCP all offer solid services depending on your needs. However, each does specialize in certain areas. The decision about what service(s) to use should be based on an assessment of your current environment and needs.

For example, you might opt to use Azure database services if you are:

- Operating in a hybrid environment
- Migrating an existing database to the cloud
- Focusing on privacy and adhering to compliance

You might opt to use AWS database services if you are:

- Already using other AWS services
- Needing high performance and reliability
- Looking for the broadest range of options

You might opt to use Google Cloud database services if you are:

- Attaching a database to a microservices architecture
- Needing high performance for your workloads
- Looking for a user-friendly solution

One final consideration: Cost

One final consideration with regard to database services is cost. In most cases, there is little to choose between the big three cloud providers' on-demand pricing for database services, but the discounts available for committed use vary significantly.

For example, while it's possible to achieve discounts of up to 60 percent with all-upfront, three-year, standard Amazon RDS Reserved Instances, the average discount available with Azure database VM reservations is 36 percent (or up to 80 percent with Hybrid Benefit).

However, you can get discounts of up to 57 percent on the compute element of Google reserved VM instances if utilization of the instances is steady and predictable. If utilization is not steady and predictable, it will likely be more cost-effective to opt for the serverless pricing options available on Amazon Aurora, Azure SQL Database, or Google Cloud Spanner.

To help conduct an analysis based on your organization's requirements, Table 3 provides an overview of the various database services offered by AWS, Azure and GCP.

Table 3: Overview of database services			
	AWS	Azure	Google Cloud
Managed relational database as a service	Amazon RDS	Azure SQL Managed Instance	Cloud SQL
Managed relational DBaaS with serverless options	Amazon Aurora	Azure SQL Database	Cloud Spanner
NoSQL database as a service	Amazon DynamoDB	Azure Cosmos DB	Cloud Bigtable
In-memory database services	Amazon ElastiCache	Azure Cache for Redis	Memorystore
Document database services	Amazon DocumentDB	Azure Cosmos DB	Filestore
Data warehouse services	Amazon Redshift	Azure Synapse	BigQuery
Data analysis services	Amazon Athena	Azure Synapse	BigQuery
Ledger services	Amazon Quantum Ledger Database (QLDB)	Azure Blockchain Workbench	Cloud Spanner
Graph database services	Amazon Neptune	Neo4j (Azure partner)	Cloud Bigtable

Cloud storage services

When it comes to a comparison of cloud storage services, organizations have historically chosen to use the object storage facility offered by the service provider through which they provision VMs. Now, however, many organizations are opting for multi-cloud environments and have more options available to them, particularly with regard to infrequently accessed and archived data.

Amazon Simple Storage Service (S3) is undoubtedly the best known of all cloud storage services, but Microsoft and Google Cloud have equally reliable and robust services.

Explanation of cloud storage classes

Understanding the different storage classes, prices and levels of fault tolerance can be complicated. To conduct a like-for-like comparison of cloud storage services, it is necessary to understand what the different cloud storage classes are.

Here's a quick breakdown:

- Block storage – Units of storage attached to a VM, which can be either local or network-attached, and are treated as an independent disc drive
- Object storage – Units of storage for most types of data, which can be replicated across different regions and zones for durability and accessed via simple web service interfaces
- File storage – Systems that facilitate file shares in the cloud that allow servers and applications to access stored data through shared file systems
- Infrequent access storage – Used for storing backup data and disaster recovery data you might need in a hurry but are unlikely to access on a frequent basis
- Archive storage – Most often used for storing data for compliance purposes, and is intended for long-term data that can tolerate retrieval latency
- Physical bulk data transport – Used for physically moving large volumes of data from on-premises data centers to the cloud service providers' data centers

The way AWS, Azure and GCP name each of these storage classes varies. Table 4 provides an overview of how they compare.

Table 4: Comparison of storage classes

	AWS	Azure	Google Cloud
Block storage	Amazon Elastic Block Store (EBS)	Azure Disk Storage	Persistent Disk, Local SSD
Object storage	Amazon S3	Azure Blob Storage	Cloud Storage
File storage	Amazon Elastic File System (EFS)	Azure Files	Filestore
Infrequent access/ archive storage	Amazon S3 Glacier, S3 Infrequent Access	Azure Archive Storage, Azure Cool Blob Storage	Nearline Storage, Coldline Storage, Archive Storage
Hybrid storage	AWS Storage Gateway	Azure StorSimple	ClearSky
Bulk data transport	AWS Import/ Export Disk, AWS Snowball, AWS Snowmobile	Azure Import/ Export Service, Azure Data Box	Storage Transfer Service

When calculating the cost of any cloud storage service, remember to include the costs of PUT, POST, COPY and GET requests, and to account for minimum capacity charges or minimum duration charges. For example, AWS has a minimum 128KB capacity charge for its two Infrequent Access classes and a minimum of 90 days of storage for its Archive Storage class.

Cloud regions and availability zones

The number and locations of cloud regions and availability zones is an important consideration when selecting a cloud service provider. This is because the more extensive the network of data centers is, the less likelihood there is of latency. Also, extensive data center networks increase the options for replication and redundancy, and improve disaster tolerance in cases of outages.

As previously mentioned, cloud service prices are often subject to regional variations. However, possibly the most important reason for comparing network size is that the greater the number of regions and zones, the more likely it is a zone local to your business will support a full range of services.

Businesses operating in U.S. central zones are likely oblivious to how limited some services are outside primary zones, and we are not talking about the outermost reaches of Southeast Asia. For example, AWS' data centers in Ohio and northern California do not support AWS Cost Explorer. This means that if you want a better understanding of cost drivers in your AWS Cloud, you will have to migrate services to a different region or use a third-party solution.

Similarly, certain services (or elements of certain services) are not available in all Azure or GCP regions. For example, Azure doesn't supports Azure VMware Solution in the Central US, South Central US, West Central US, West US 2, East US 2, or Canada East regions. And customers looking to take advantage of the latest high-compute series of VMs will have to deploy their workloads in the East US, South Central US, or West US 2 region to be able to access this VM type.

Conclusion

Compute and storage services, and their local availability, will be the primary considerations for most organizations when comparing cloud service providers. Some may have other motives for conducting a cloud services comparison to take factors such as analytics, integrations, databases and DevOps tools into account.

Comparing cloud services can be rewarding in terms of cost reduction and enhanced performance for organizations already operating in a multi-cloud environment or those looking to expand their cloud strategy to a multi-cloud model. Those that select services from a variety of providers can create a custom multi-cloud environment that ensures their unique performance and financial requirements are achieved.

Make sure you have the necessary solutions in place to manage your multi-cloud environment by learning more about VMware Tanzu CloudHealth®.

Item No: 1703106aq-wp-comp-svcs-cld-pvdrs-a4 2/23