CloudHealth® by vmware®

Comparing Services For The Big Three Cloud Providers

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.

We have chosen to limit our cloud services comparison to the top three IaaS providers providing service in the Western Hemisphere—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.

This eBook 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 you on your multicloud journey.

1 PRODUCT

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

Take serverless Function-as-a-Service as an example. All three cloud service providers offer serverless computing as a product, but how serverless computing services differ across providers is a result of varying provider maturity and product portfolio capabilities, like the ability to pass data to backend services, perform calculations, transform data, store results, and quickly retrieve data.

AWS benefits from being the leader in serverless cloud computing due to the sheer size of its product portfolio and the range of products that can be integrated using Lambda. Azure Functions is a close second, followed by GCP Cloud Functions, which lags behind because of its recent launch and historic focus on building out machine learning and Kubernetes.

Use the chart below to see the major differences in product availability across the three providers (here's a hint—many are comparable in functionality with different names):

PRODUCT	aws	Microsoft Azure	Google Cloud Platform
Virtual Servers	Instances	VMs	VM Instances
Platform-as-a-Service	Elastic Beanstalk	Cloud Services	App Engine
Serverless Computing	Lambda	Azure Functions	Cloud Functions
Docker Management	ECS	Container Service	Container Engine
Kubernetes Management	EKS	Kubernetes Service	Kubernetes Engine
Object Storage	S3	Block Blob	Cloud Storage
Archive Storage	Glacier	Archive Storage	Coldline
File Storage	EFS	Azure Files	ZFS / Avere
Global Content Delivery	CloudFront	Delivery Network	Cloud CDN
Managed Data Warehouse	Redshift	SQL Warehouse	Big Query

2 COST

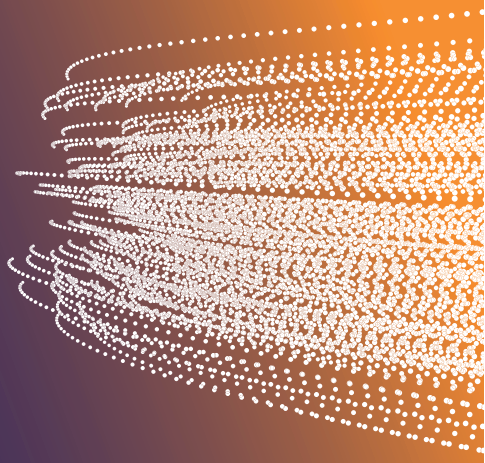

We don't often compare costs between AWS, Azure, and GCP. Each provider updates their pricing model multiple times each year in response to both market trends and business strategy, and to ensure they are staying competitive with one another. As one provider decreases cost or offers new discounts, other providers are forced to respond to stay financially competitive.

Generally speaking, AWS tends to be the most competitive cloud service provider on a CPU/Hour basis for most workloads, whereas GGP is the less expensive option for compute-intensive workloads. Azure comes into the equation if your business uses a range of Microsoft products that would qualify for a discount under an Enterprise Agreement. So if you're currently using AWS and are looking to expand to a second cloud, Azure might end up being the most cost-effective option after qualifying discounts, even though GCP's sticker price is lower. If you aren't eligible for some of Microsoft's enterprise discounts, GCP's Sustained Use Discounts, which is a discount based on the time you are using the resources, could be a big advantage.

The frequently-changing variety of discounts further complicates cost comparison between providers. For example, when AWS introduced one-year Convertible RIs, Microsoft launched Azure Reserved VM Instances with free exchanges or adjusted refunds. This was followed by GCP's announcement of Committed Use Discounts, which rivals AWS' reserved instances by offering discounts with no advanced payment.

3 CLOUD COMPUTE SERVICES

Within this category we find Virtual Machines (Instances/VMs) configured for General Purpose, Memory Optimization, Compute Optimization, and Storage Optimization. In addition to Virtual Machines, our comparison also encompasses Containers and Serverless Computing.

GENERAL PURPOSE

General Purpose Virtual Machines 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 Virtual Machine 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.

Both AWS’ EC2 Instances and Azure’s General Purpose VMs offer similar, and better, storage ratios than GCP’s standard General Purpose VMs. However, with GCP, you have the option of creating custom VM types to meet your specific compute requirements. If you’re currently an AWS user looking to adopt a multicloud model, you won’t gain any advantages by using Azure’s general purpose VMs because of how similar their storage services are to your current AWS offerings. Instead, look into GCP and see if it makes sense for you to take advantage of their custom VMs capabilities to ensure you’re meeting all of your environment’s unique compute needs.

3 CLOUD COMPUTE SERVICES

CONTINUED

MEMORY OPTIMIZED

Memory Optimized Virtual Machines deliver high memory-to-CPU ratios suitable for relational database servers, medium to large caches, and in-memory analytics. As well as “regular” Memory Optimized VMs, all three providers in our cloud service comparison offer super-Memory Optimized VMs for large enterprises - providing 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 in order 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 Virtual Machines 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 Modelling
- 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 not every High-Performance Compute service is available in every region at the present time.

3 CLOUD COMPUTE SERVICES

CONTINUED

STORAGE OPTIMIZED

Storage Optimized Virtual Machines 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 ranges offerings of memory per vCPU. GCP doesn't offer a designated Storage Optimized VM option, but instead allows you to add one of the following to an existing VM:

- Standard Persistent Discs, which are the slowest but cheapest.
- Regional Persistent Discs, which are equally as slow but with better redundancy.
- Standard SSD Discs, which are faster and necessary for high rates of random IOPS.
- Regional SSD Discs, suitable for workloads that may not have application-level replication.
- Local SSD Discs, which 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 tradeoffs. For example, local SSD storage is not automatically replicated and all data on the local SSD may be lost if the Virtual Machine terminated for any reason.

3 CLOUD COMPUTE SERVICES

CONTINUED

CONTAINER SERVICES

Containers are a hot topic for the rapidly-evolving cloud—so much so that Gartner has predicted "by 2020, more than 50 percent of global organizations will be running containerized applications in production, up from less than 20 percent today." All three cloud service providers offer some form of managed container service, but there are some differences in their level of service.

AWS is considered to be more secure, reliable, and scalable, but difficult to use for new developers unfamiliar with container services and Kubernetes. Azure is more intuitive for Windows developers, but doesn't support hybrid containers. GCP is usually first-to-market with new features, offers great integrations with other services, and is the least expensive of the three once discounts are taken into account. Plus, GCP developed Kubernetes, which ultimately gave them a head start in this space compared to the other two clouds.

Which provider is best for your container needs? It depends—if you're currently running on Azure but need to support hybrid containers, GCP might be the better option if you have a small budget or your developers are still learning Kubernetes. Already have a team of experts ready to rapidly expand your use of container services? AWS can provide the scalability needed to make it happen.

3 CLOUD COMPUTE SERVICES

CONTINUED

SERVERLESS COMPUTING & 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 is triggered by an event. As businesses only pay for the milliseconds when the function is executed, serverless computing can significantly reduce costs.

AWS leads the way in serverless computing, although Azure is catching up quickly. Google lags behind and is the most expensive of the three cloud service providers once the free tier of two million requests/400,000 GB of compute seconds per month is surpassed. Google also restricts projects to fewer than twenty triggers, and these have to be primed via its Cloud Pub/Sub service.

4 CLOUD STORAGE SERVICES

When it comes to a comparison of cloud storage services, organizations have historically chosen to use the storage facility offered by the service provider through which they provision Virtual Machines. Now that management of multicloud environments has become less complicated, organizations have more options available to them—particularly with regard to data that is infrequently accessed.

Undoubtedly AWS' Simple Storage Service (S3) is the best known of all cloud storage services. However, understanding the different storage classes, different prices, and different levels of fault tolerance can be complicated. Microsoft and Google Cloud have equally reliable and robust services, and especially with regard to frequently accessed data, just about beat AWS on price.

EXPLANATION OF CLOUD STORAGE CLASSES

In order 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 Virtual Machine. They 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 services interfaces.

FILE STORAGE: systems that facilitate file shares in the cloud that allow servers and applications to access stored data through shared file systems.

4 CLOUD STORAGE SERVICES

CONTINUED

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. This class of cloud storage is intended for long-term data that can tolerate retrieval latency.

HYBRID STORAGE: systems used for moving inactive data to the cloud while maintaining business-critical and sensitive data on-premises to reduce physical storage requirements.

PHYSICAL BULK DATA TRANSPORT: used for physically moving large volumes of data from on-premises data centers to the cloud service providers' data centers.

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 ninety-day minimum charge for its Archive Storage class.

5 CLOUD REGIONS & AVAILABILITY ZONES

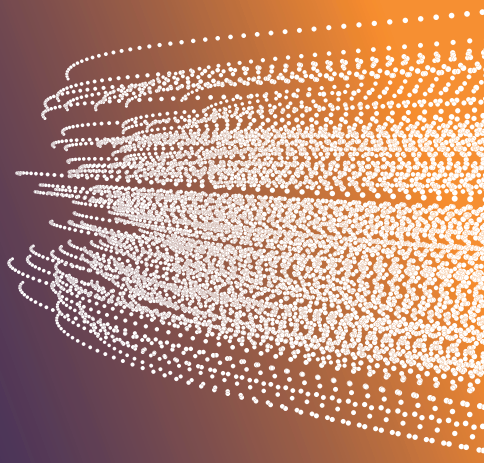

The number and locations of cloud regions and availability zones is an important consideration when selecting a cloud service provider—not only because the more extensive the network of data centers is, the less likelihood there is of latency; but also because extensive data center networks increase the options for replication and redundancy, and improve disaster tolerance in cases of outages.

It was mentioned earlier in our cloud services comparison that cloud service prices are often subject to regional variations, and this is certainly true for latency-tolerant services such as Archive Storage—with prices differing by as much as 50% depending on where archive data is stored. However, possibly the most important reason for comparing network size is, 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 South-East Asia. For example, AWS' data centers in Ohio and Northern California do not support Amazon Elastic Container Services for Kubernetes (EKS). If you want to use these services, you will have to run them via Northern Virginia or Oregon—which may also have implications for other cloud services you wish to take advantage of. Like AWS, Azure also offers their Kubernetes services, Azure Kubernetes Services (AKS), on both the east and west coast, while also providing Kubernetes support in Iowa with future availability in Texas. If you're looking to expand your container services, or any services for that matter, understanding how and where these services are supported is crucial to your multicloud strategy.

Each provider in our cloud services comparison publishes a webpage of services available per region, and these are well worth reviewing.

CONCLUSION

Compute services 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 in order to take factors such as analytics, networking, 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 multicloud environment or those looking to expand their cloud strategy to a multicloud model. Organizations that select services from one provider and different services from another provider can create a custom multicloud environment that ensures their unique performance and financial requirements are achieved.

Make sure you have the necessary solutions in place to effectively manage your multicloud environment by visiting **www.cloudhealthtech.com.**