Why The Best Way to Reduce Compute Cloud Spend is OCI

September 2024

Written by: Marc Staimer, President Dragon Slayer Consulting, Sr. Contributor to theCUBE research

Introduction

There continues to be a severe market backlash against the regularly surprisingly high cloud compute costs of the big three cloud service providers (CSP) – Amazon Web Services (AWS), Microsoft Azure (Azure), and Google Cloud Platform (GCP). The backlash has taken the form of "optimizing cloud spend." Optimizing cloud spend is actually a euphenism for cutting cloud spend. And it's a major priority of enterprises, midtier, and SMBs alike.

The problem becomes how? Some customers have looked at how they can reduce their cloud spend within the structure of their current CSP. Others have started repatriating applications and data to their premises. Repatriation significantly increases costs temporarily because of CSP egress fees – data movement out of the public cloud – and more often than not in the long term as well depending on circumstances. The most intriguing way CSP customers are reducing their cloud spend is by moving it to the Oracle Cloud Infrastructure (OCI).

This research investigates the pros and cons of each of the different methods for reducing compute cloud spend, and makes recommendations based on the findings. Spoiler alert: the title provides a hint at the results.

Table of Contents

Premise

Customers want to scale back their cloud spend. They can always just scale back their cloud compute use, but that tends to be more difficult than it seems. Applications are more likely to use more cloud resources over time, not less. There is always some waste, such as forgotten devops instances racking up fees even when they're no longer being actively used. Or applications that have fallen out of use but are still active in the cloud service provider (CSP). This waste can be eliminated if it's tracked down. These are ultimately a small part of the cloud spend and should always be part of disciplined cloud compute usage.

CSP Configuration Options

Real, noteworthy cloud spend reductions requires much more. CSPs know this and it's why they provide a vast number of fixed shape compute configuration options with different processors, cores, and memory at vastly different price points – AWS by itself offers 709 different instance configurations as of this writing with Azure offering 647 instance cofigurations. Oracle Cloud Infrastructure (OCI) offers markedly simpler compute sizing with flexible instances. Google Cloud (GCP) also offers fixed instances and flexible instances – flexible instances cost less and have finer granularity.

AWS, Azure, and GCP fixed compute instances are optimized one of four ways:

1. Compute optimized
2. Memory optimized
3. General purpose, a.k.a. balanced compute and memory
4. Accelerated computing

Every fixed shape instance configuration has its own unique price point. There is a fixed ratio of compute to memory in each case. The underlying premise being that cloud spend will be reduced with less waste through the correct alignment of workload to instance configuration. Customers select the size they determine or believe is the best fit for their application workload. Should they find they need a bit more than the size they selected in a fixed shape instance, the next step up is 2X their current shape. Vice versa, going down a step reduces the compute vCPUs[1] and memory by 50%.

Because OCI allows the customer to choose the CPU type – AMD, Intel, or Ampere ARM, the number of vCPUs, and the memory amount – GB increments – to be sized independently of one another. This provides extremely fine granularity, simpler workload correlation, and much more adaptability scaling up or down. Customers can make any shape optimized as compute, memory, or balanced, and can modify as demanded. GCP also has flexible instances, albeit for only a select few CPUs and at a higer cost. This research examines how proper configuration alignment reduces cloud spend.

Note: Previous theCUBEresearch – formerly known as Wikibon – found *OCI flexible instances* were conclusively proven to be considerably lower cost, more flexible, easier to align with workload requirements, and substantially more cost/performance effective than AWS, Azure, and GCP fixed instances or GCP flexible instances. This research extends the previous research.

Discounting

A significant market perception about cloud spend reduction is that it can be reduced through aggressive CSP discounting practices. All CSPs offer a variety of discount programs. Discounting is usually tied to spending commitments over a specific time period, or risk[2] from infrastructure, performance, or availability that can vanish or become unavailable when needed with little to no notice.

Discount programs are supposed to reduce cloud spend. In most cases they do. However, under specific circumstances, they can also increase it. For example 1 year or 3 year savings plans can appear to save the

[1] OCI does not use vCPUs in its configurations or pricing. It uses OCPUs. The OCPU is 1 core whereas the vCPU is ½ core. 1 OCPU = 2 vCPUs.

[2] Risk as defined as the infrastructure resources being cut off with little to no notice as with spot pricing, or infrastructure resources not being available for bursting, or not getting the performance expected because of noisy neighbors.

customer costs or cloud spend. But, these plans assume the customer will be using 100% those contracted instances 100% of the time, all the time, for the entire committed time period. That's simply unrealistic. Many of the CSP discount programs are additionally complicated and confusing. This research examines the veracity of discount programs and it's actual impact cloud spend reduction.

Multi-region Uplift

A more difficult compute cloud spend reduction problem is multi-region uplift costs. AWS, Azure, and GCP all have fees for different regions. They typically charge more for the exact same cloud services depending on the country, sovereignty, or government region. Most services tend to start in the USA in the Virginia East region. They frequently do not offer all their cloud and compute services in every region. The result is applications developed in Virginia East are much costlier when deployed in other regions. The exception is Mumbai India. Whereas the uplift ranges from 7% to 41% over Virginia East for the vast majority of regions, Mumbai India can be substantially less with as much as a 35% downlift

OCI uniquely does not charge any uplift, anywhere. Their price is the exact same for compute, memory, and storage, everywhere.

Burstables

Burstables are promoted as a way to reduce cloud spend. Burstable instances are really only for applications that use a small fraction of the fixed shape's resources the vast majority of the time. They can intermittently and infrequently burst up to full capacity for a brief amount of time.

Burstable instances are bought at a specific percentage of the fixed instance shape's total vCPU resouces. Applications using a burstable instance can use the entire instance's vCPU resources, but only for a very limited time. Each CSP uses a different methodology for measuring the time that bursting is available.

Typically, customers have to decide how much of the total vCPU processing they will commit to financially. Let's say the customer anticipates using approximately 20% on average of the vCPUs in the shape over the coming year. They commit to that 20% for a full year. For all the time their application uses less than the 20% level, they will receive credits up to a maximum number. The credits are then used to pay for those times when the application bursts above the 20% level. That maximum credit number noticeably limits the amount of time an application can burst above the commitment level. And there is no guarantee the resources will available.

Keep in mind that AWS, Azure, and GCP generally oversubscribe their cloud compute infrastructure. These CSPs are operating under the assumption that most customer workloads will under-utilize resources most of the time. They are mostly correct. Statistics show all of the instances in a physical host are unlikely to need their full contracted resources simultaneously. There are several fallacies with those oversubscription assumptions. The first is that it fails to take into account "Murphy's Law" – whatever can go wrong, will go wrong, and typically at the worst possible time. No customer wants to find out they do not have the resources they need when they absolutely need them. That unfortunately happens.

OCI alone does not oversubscribe their fixed or flexible instance shapes. It means customers get the resources they're paying for with enterprise-grade SLAs.

AWS, Azure, and GCP limit their burstable instances to typically older, lower-performing processors. In contrast, OCI burstables are not relegated to older processors and are available on the latest CPUs. AWS burstable policy is a good example of the older processor policies. It uses the Graviton2 ARM, 1st Gen AMD EPYC processors, and Intel Xeon Scalable (Skylake, Cascade Lake) processors. Those are all older or just less performing processors. They do not use the latest AMD 4th generation EPYC processors or even the 3rd generation. Nor do they allow burstable instances on Intel's latest generation of processors. Both Azure and GCP are similarly restrictive. OCI burstable instances are available on the latest processors from AMD and Intel.

Putting that in perspective, AMD 4th Gen EPYC processors are up to **8X faster** than AMD 1st Gen EPYC processors, up to **12X** the number of **cores**, and up to **6X memory**. OCI AMD 4th Gen EPYC processors burstable instance shape

commitments start at 2 vCPUs[3] and scale to 160 vCPUs. AWS tops out at 8 vCPUs, Azure at 20 vCPUs, and GCP at 2 vCPUs in their burstable older processor shapes.

Burstable shapes on the older processors from the other CSPs are seriously problematic for mission-critical and business-critical applications. These applications absolutely demand performance for low application response times, better end-user experiences, higher employee morale, better quality work, faster-time-to-actionable-insights, faster-time-to-market, and faster-time-to-revenues.

What this means is that customers should only use AWS, Azure, and GCP burstable instances for applications that can tolerate the lower performance without guarantees of additional resources. OCI offers better performing and more reliable burstable flexibility.

Instance Spot Pricing

Spot pricing offers excess instance resources at extremely low pricing. The fixed instances are the same as on-demand with the same services and integration capabilities. Unfortunately, there's a catch in that the CSP can cancel the spot instance at any time, under any circumstances, with very little notice – as long as 2 minutes or as little as 30 seconds.

Spot instances need to be carefully utilized, if at all. Only application workloads that can tolerate sudden termination make any sense at all. Some batch processes may be okay. Spot pricing looks better on paper than in reality because of the high risk of the resources being pulled with little to no notice.

Storage Cost

Storage cost is rarely discussed as an opportunity to reduce cloud spend and for good reason. Data has become the life blood of every organization. Data consumes storage. In fact, storage is the only cloud technology that is both used and consumed. New data cannot be stored in the same storage location as data that's already stored there until the original data is erased. Data is rarely erased. Therefore, as storage is consumed its cost consistently trends upwards. Storage is the tail that wags the budget dog.

Keep in mind there are a wide variety of cloud storage types including block, file, and object, as well as a plethora of different storage media from NVMe SSDs to HDDs. And each cloud storage type has its own variations based on performance and price. The broad choice of cloud storage options is so great that this entire document could be entirely on cloud storage. However, since the focus of this document is on reducing cloud spend, the storage comparison research is limited to the balanced and high performance block storage – the storage most utilized for mission-critical and business-critical applications – in as close to an apples-to-apples comparison[4] as possible.

Data Egress Cost

Data egress fees are the cloud service provider charges or costs for moving, copying, or just reading data out from the cloud storage where it resides into a different public or private cloud. These data egress costs are commonly referred to as bandwidth or data transfer fees. They are completely separate from cloud compute and cloud storage fees. Data egress fees are tied to the CSP's network costs which is one reason why it substantially varies from region-to-region. OCI is egress fees also change per region, but a much lower variance.

Another data egress cost factor is whether the data egress is through the public cloud network gateway or through a VPN. The former is less expensive and faster.

High data egress costs have been known to hold an organization's data hostage within the CSP because of prohibitive costs. It's why examining the details of each CSP's data egress fees is an essential to keeping cloud spend down.

Reducing Cloud Spend Focus

The objective of this research is to take a hard-nosed look at how AWS, Azure, GCP, and OCI compare on cloud spend. More specifically, how they compare at reducing compute, block storage, and data egress cloud spend.

---

[3] 2 vCPUs = 1 core = 1 OCPU

[4] OCI offers external NVMe high performance block storage with similar latencies, higher IOPS, higher throughput, and the ability to scale performance and capacity independently. AWS, Azure, and GCP do not offer anything similar making comparisons apples to oranges.

Instance Comparison Cost

With the many hundreds of potential and different fixed and flexible instance shapes, it is imperative to be as fair as possible and level the playing field when comparing pricing. That requires any comparisons be normalized as much as possible. It starts with the assumptions made about compute and memory configurations.

Assumptions

Normalization requires this research look at the most common shapes in cloud compute. Conducting nonbiased evaluations necessitates equal assumptions for all of the CSPs for on-demand, in addition to 1-year commitment, and 3-year commitment discounts.

As previously noted, both OCI and GCP provide much more finely granular configurations than AWS or Azure via their flexible instances – enabling independent sizing of compute and memory. But since AWS and Azure do not, all comparisons will be for the exact same configurations[5] in processor, vCPUs and memory.

Compute and Memory Configuration Assumptions

CPU Type	vCPUs	Mem	OS	AWS	Azure	GCP	OCI
Gen4 AMD EPYC	4	16 GB	Linux - Free Distribution	m7a.xlarge	D4as v6	c3d-standard-4	VM.Standard.E5.Flex
Gen3 AMD EPYC				m6a.xlarge	D4as v5	n2d-standard-4	VM.Standard.E4.Flex
Intel Xeon Ice Lake				m6i.xlarge	D4s v5	n2-standard-4	VM.Standard3.Flex

Be aware that not all CPU types are available in all areas by all cloud service providers. AWS, Azure, and GCP generally make available new services such as the new Gen 4 AMD EPYC processor, into a single region – typically the lowest cost region such as Virginia, USA. Then over a period of time, they'll make it available piecemeal into other regions.

OCI has a different approach. When OCI announces a new service they make it available in all or just about all regions. When it's not available immediately in a given region, OCI makes every effort to get it to that region as soon as possible.

Pricing Assumptions

Type	Hrs Commit	Locations									
On-demand	730/mo	Virginia, US	California, US	Brazil	London, UK	Frankfurt, Germany	Mumbai, India	Singapore	Sydney, Australia	Tokyo, Japan	US Gov't VA
1 yr. Commit	8,760/yr										
3 yr. Commit	26,280/3 yrs										

AWS, Azure, and Google pricing varies by region. The lowest cost region or baseline region is primarily Virginia, USA. Most other regions have an uplift cost or tax on the baseline regions pricing. OCI is once again different. OCI charges the same rate worldwide regardless of region.

On-Demand Pricing Comparison

On-demand pricing is based on actual usage, not projected or guaranteed committed usage. It's a pay for use based on billing increments. It's extremely rare for an application workload to be using resouces 24 hours a day, every day of the month and year. Consider that Linux on-demand minimum billing increments is universally 1 minute from OCI, AWS, Azure, and GCP. The additional follow on base increments for OCI, AWS, and GCP are 1 second, and for Azure 1 minute. That means the on-demand costs are likely to be variable and much less than 730 hours per month. The assumption is 730 hours per month to provide a maximum baseline.

Windows on-demand billing increments are slightly different[6]. The comparison is on lower cost Linux instances.

[5] Flexible instances are candidly much more cost effective and reduce cloud spend considerably as shown in Appendix A.

[6] The minimum Windows billing increment is 1 minute for AWS, Azure, and GCP, but more finely granular at 1 second for OCI. The following base increments are 1 second for OCI and GCP. AWS is also 1 second for Windows, Windows with SQL Enterprise, Windows w/SQL Standard, Windows w/SQL Web. Each has a different charge.

On-Demand Instance Pricing Comparison Gen 4 AMD EPYC Processors

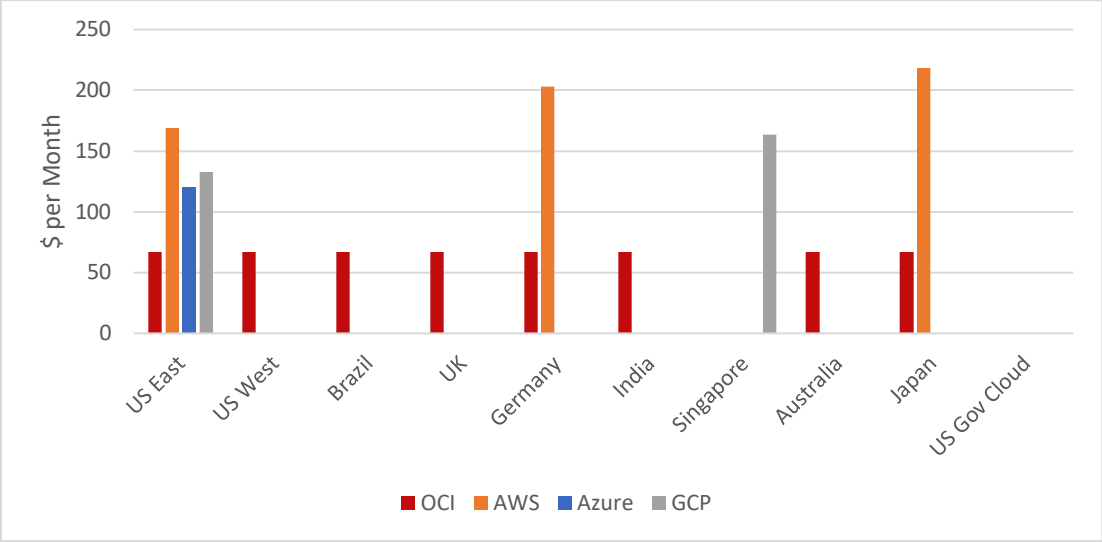

Chart 1: On-Demand Instance Pricing Comparison Gen 4 AMD EPYC Processors – Lower is Better

The chart above and detail below clearly shows OCI has much lower on-demand pricing than AWS, Azure, or GCP. AWS is ~ **2.52X**, Azure is ~ **1.79X**, and GCP is ~ **1.97X** – US East.

	OCI			AWS		Azure		GCP[7]	
	$/hr		$/mo	$/hr	$/mo	$/hr	$/mo	$/hr	$/mo
US East	vCPU	0.0150	**67.16**	0.02318	**169.24**	0.1650	**120.45**	0.1816	**132.57**
	GB	0.002							
US West	vCPU	0.0150	**67.16**						
	GB	0.002							
Brazil	vCPU	0.0150	**67.16**						
	GB	0.002							
UK	vCPU	0.0150	**67.16**						
	GB	0.002							
Germany	vCPU	0.0150	**67.16**	0.27772	**202.74**				
	GB	0.002							
India	vCPU	0.0150	**67.16**						
	GB	0.002							
Singapore	vCPU	0.0150						0.22403	**163.54**
	GB	0.002							
Australia	vCPU	0.0150	**67.16**						
	GB	0.002							
Japan	vCPU	0.0150	**67.16**						
	GB	0.002							
US Gov Cloud	vCPU	0.0150							
	GB	0.002							

Table 1: On-Demand Instance Pricing Comparison Gen 4 AMD EPYC Processors – Lower is Better

Additionally, OCI provides the Gen 4 AMD EPYC processor instances in most regions whereas AWS and GCP at the time or this research – are limited to 2 regions, and Azure is in preview in only a single region (US East). OCI treats all of its regions equally. They make a concerted effort to roll out new services to all regions concurrently. When OCI can't, they make every effort to bring orphaned regions up to speed as quickly as possible.

What's clearly obvious is that OCI does not uplift its instance fees to any region. They provide the same price regardless of region. AWS, Azure, and GCP do not have the same policy that can cause significant price jumps in many regions.

[7] Google commonly provides a 20% discount for on-demand pricing

But then, what about on-demand pricing for slightly older AMD processors or Intel processors? A look at the same configurations shows a similar gap. For Gen3 AMD EPYC processor: AWS is ~ **2.34X**, Azure is ~ **2.32X**, and GCP is ~ **2.06X** more costly than OCI. For Intel Xeon Ice Lake processor: AWS is ~ **1.85X**, Azure is ~ **1.85X**, and GCP is ~ **1.49X** more costly than OCI – based on US East. The gap grows noticably bigger in other regions with the exception of Mumbai India, where the gap is somewhat less.

On-Demand Instance Pricing Comparison Gen 3 AMD EPYC Processors

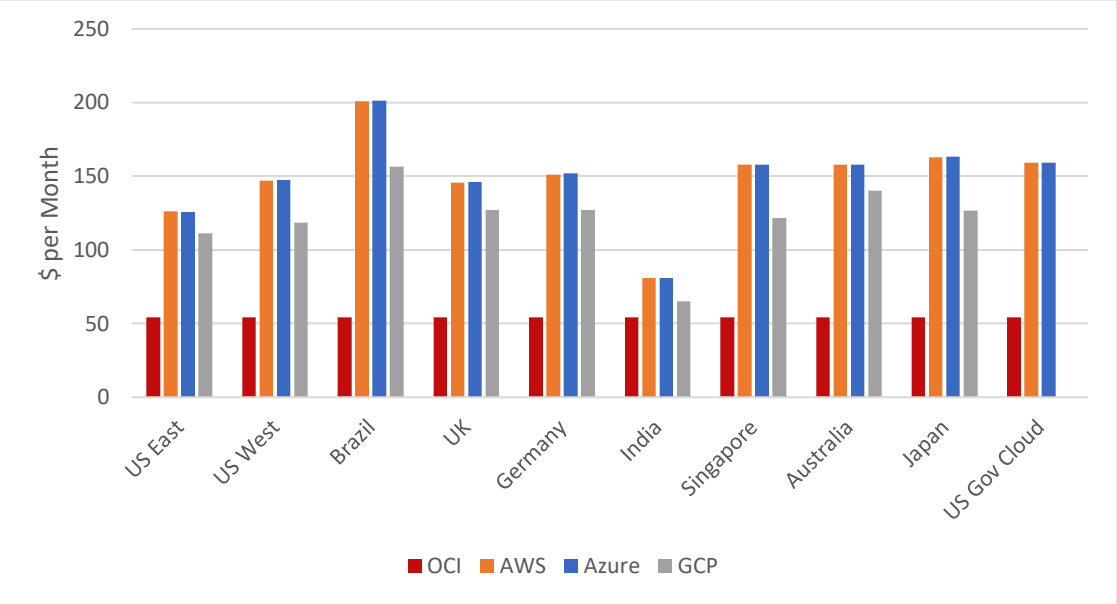

Chart 2: On-Demand Instance Pricing Comparison Gen 3 AMD EPYC Processors – Lower is Better

	OCI			AWS		Azure		GCP	
	$/hr		$/mo	$/hr	$/mo	$/hr	$/mo	$/hr	$/mo
US East	vCPU	0.0125	**54.02**	0.1728	**126.14**	0.1720	**125.56**	0.1903	**111.15**
	GB	0.0015							
US West	vCPU	0.0125	**54.02**	0.2016	**147.17**	0.2020	**147.46**	0.2030	**118.53**
	GB	0.0015							
Brazil	vCPU	0.0125	**54.02**	0.2754	**201.04**	0.2760	**201.48**	0.2682	**156.65**
	GB	0.0015							
UK	vCPU	0.0125	**54.02**	0.1998	**145.85**	0.2000	**146.00**	0.2177	**127.14**
	GB	0.0015							
Germany	vCPU	0.0125	**54.02**	0.2070	**151.11**	0.2080	**151.84**	0.2177	**127.14**
	GB	0.0015							
India	vCPU	0.0125	**54.02**	0.1111	**81.10**	0.1110	**81.03**	0.1115	**65.13**
	GB	0.0015							
Singapore	vCPU	0.0125	**54.02**	0.2160	**157.68**	0.2160	**157.68**	0.2085	**121.75**
	GB	0.0015							
Australia	vCPU	0.0125	**54.02**	0.2160	**157.68**	0.2160	**157.68**	0.2398	**140.03**
	GB	0.0015							
Japan	vCPU	0.0125	**54.02**	0.2232	**162.94**	0.2240	**163.52**	0.2168	**126.61**
	GB	0.0015							
US Gov Cloud	vCPU	0.0125	**54.02**	0.2180	**159.14**	0.2180	**159.14**		
	GB	0.0015							

Table 2: On-Demand Instance Pricing Comparison Gen3 AMD EPYC Processors – Lower is Better

On-Demand Instance Pricing Comparison Intel Xeon Ice Lake Processors

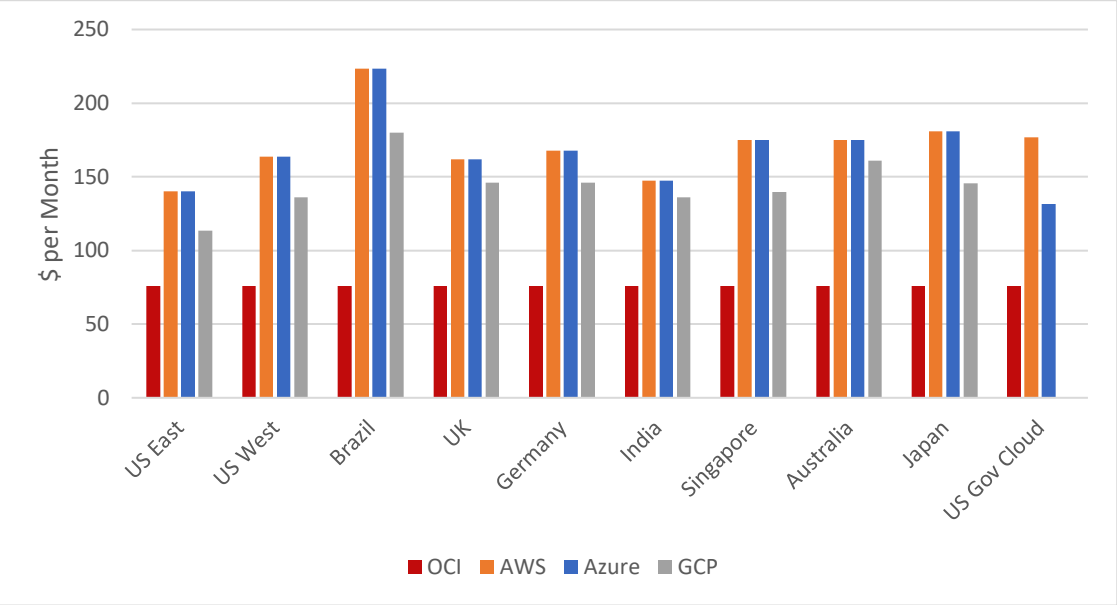

Chart 3: On-Demand Instance Pricing Comparison Intel Xeon Ice Lake Processors – Lower is Better

	OCI			AWS		Azure		GCP	
	$/hr		$/mo	$/hr	$/mo	$/hr	$/mo	$/hr	$/mo
US East	vCPU	0.02	**75.92**	0.1920	**140.16**	0.1920	**140.16**	0.1942	**113.41**
	GB	0.0015							
US West	vCPU	0.02	**75.92**	0.2240	**163.52**	0.2240	**163.52**	0.2333	**136.25**
	GB	0.0015							
Brazil	vCPU	0.02	**75.92**	0.3060	**223.38**	0.3060	**223.38**	0.3083	**180.05**
	GB	0.0015							
UK	vCPU	0.02	**75.92**	0.2220	**162.06**	0.2220	**162.06**	0.2502	**146.120**
	GB	0.0015							
Germany	vCPU	0.02	**75.92**	0.2300	**167.90**	0.2300	**167.90**	0.2502	**146.120**
	GB	0.0015							
India	vCPU	0.02	**75.92**	0.2020	**147.46**	0.2020	**147.46**	0.2333	**136.250**
	GB	0.0015							
Singapore	vCPU	0.02	**75.92**	0.2400	**175.20**	0.2400	**175.20**	0.2396	**139.93**
	GB	0.0015							
Australia	vCPU	0.02	**75.92**	0.2400	**175.20**	0.2400	**175.20**	0.2756	**160.95**
	GB	0.0015							
Japan	vCPU	0.02	**75.92**	0.2480	**181.04**	0.2480	**181.04**	0.2492	**145.53**
	GB	0.0015							
US Gov Cloud	vCPU	0.02	**75.92**	0.2420	**176.66**	0.1800	**131.40**		
	GB	0.0015							

Table 3: On-Demand Instance Pricing Comparison Intel Xeon Ice Lake Processors – Lower is Better

The comparisons make it crystal clear, based on the published on-demand instance pricing, in identical configurations, OCI will reduce cloud spend by 50% or more – with the exception of Gen3 AMD EPYC processors in Mumbai, India, where OCI still reduces cloud spend by 33% or more.

Next step is to examine instance price comparisons for 1-year and 3-year commitments, a.k.a. savings plans.

1-Year and 3-Year Savings Plan Pricing Comparison

AWS, Azure, and GCP all offer significant discounts for both 1-year and 3-year commitment savings plans. The discounts can be substantial. AWS discounts are ~ 27%, Azure ~ 32%, and GCP ~ 28% for the 1-year commitment. For the 3-year commitment those discounts go up to AWS ~ 49%, Azure ~ 62%, and GCP ~ 46%. Seems like a pretty good deal. But those commitment discounts require an obligation to pay for every minute of every hour for the

entire 1 or 3-year commitment whether it's being used or not. And realistically, most workloads do not come anywhere near that level of utilization.

OCI does not offer 1 or 3-year commitments or commitment discounts because its highly flexible instances and very low, on-demand pricing don't need additional commitment discounts. This allows users to pay only for what they use, which is generally a lot less than every hour of every day for 1-year (8,760 hrs) or 3-years (26,280 hrs). But in order to get some semblance of an apples-to-apples comparison, this research assumes that OCI instances will in fact be used every second, minute, hour, of every day for 1 or 3-years. This is extremely unlikely and will grossly overstate what customers actually pay for OCI flexible instances while skewing the results towards AWS, Azure, and GCP.

However as the following results show, OCI costs are still substantially lower even with the unfair skew. The charts and tables below show both the 1-year and 3-year commitment savings plans of AWS, Azure, and GCP versus the on-demand flexible instance pricing of OCI.

Remember, that at the end of the day, what really matters is the customer's net cost, not the discount. The discount is only a methodology that gets them to a net price.

1-Year Commitment Savings Plan Instance Pricing Comparison Gen 4 AMD EPYC Processors

To keep things simple, the same series of charts and tables as the on-demand pricing are used.

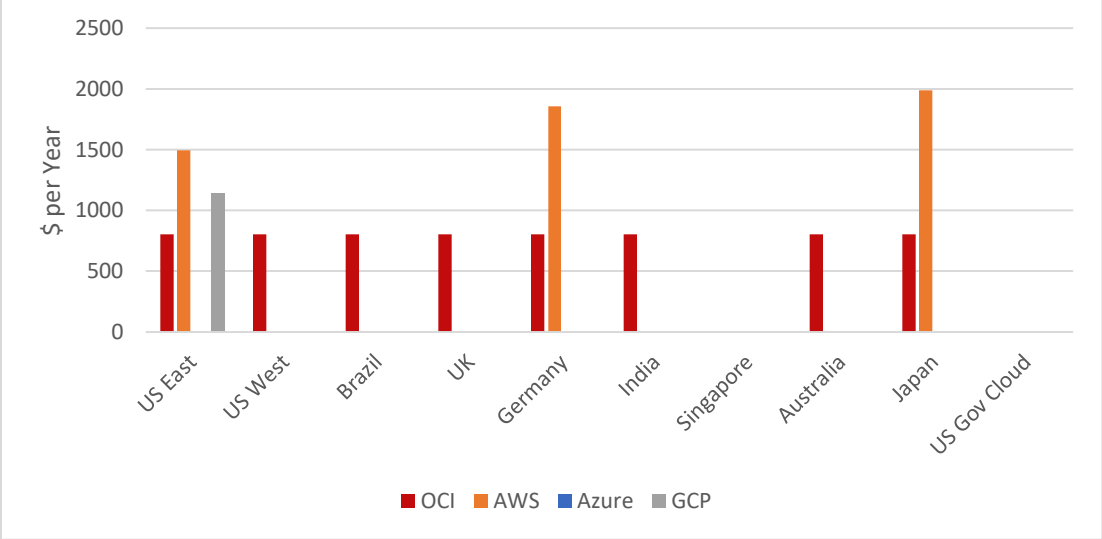

Chart 4: 1-Year Commitment Savings Plan Pricing Comparison Gen 4 AMD EPYC Processors – Lower is Better

The chart and table makes it clearly evident once again for US East – the other regions are mostly not supported yet – that OCI's flexible instances deliver much lower cost to the customer across the board in every region without having to make a commitment versus customers that make the 1-year commitment savings plans of AWS, Azure[8] and GCP. OCI Flexible instances on-demand pricing for the exact same number of hours (8,760) are ~ **46%** less than AWS and ~ **30%** less than GCP for US East – the only place GCP currently offers Gen4 AMD EPYC Processors. The OCI to AWS difference grows bigger in the other two regions AWS offers Gen4 AMD EPYC Processors.

[8] Azure currently has no savings plan for Gen 4 AMD EPYC Processors.

	OCI			AWS		Azure		GCP	
	$/hr		$/yr	$/hr	$/yr	$/hr	$/yr	$/hr	$/yr
US East	vCPU	0.0150	805.92	0.1703	1,492.09			0.130752	1,145.39
	GB	0.002							
US West	vCPU	0.0150	805.92						
	GB	0.002							
Brazil	vCPU	0.0150	805.92						
	GB	0.002							
UK	vCPU	0.0150	805.92						
	GB	0.002							
Germany	vCPU	0.0150	805.92	0.2118	1,855.06				
	GB	0.002							
India	vCPU	0.0150	805.92						
	GB	0.002							
Singapore	vCPU	0.0150							
	GB	0.002							
Australia	vCPU	0.0150	805.92						
	GB	0.002							
Japan	vCPU	0.0150	805.92	0.2268	1,986.94				
	GB	0.002							
US Gov Cloud	vCPU	0.0150							
	GB	0.002							

Table 4: 1-Year Commitment Instance Savings Plan Pricing Comparison Gen 4 AMD EPYC Processors – Lower is Better

Although the 1-year commitment instance savings plan narrowed the gap, surely the 3-year commitment instance savings plan discounts should cause AWS, Azure, and GCP to match or come in lower. Well, yes and no.

3-Year Commitment Savings Plan Instance Pricing Comparison Gen 4 AMD EPYC Processors

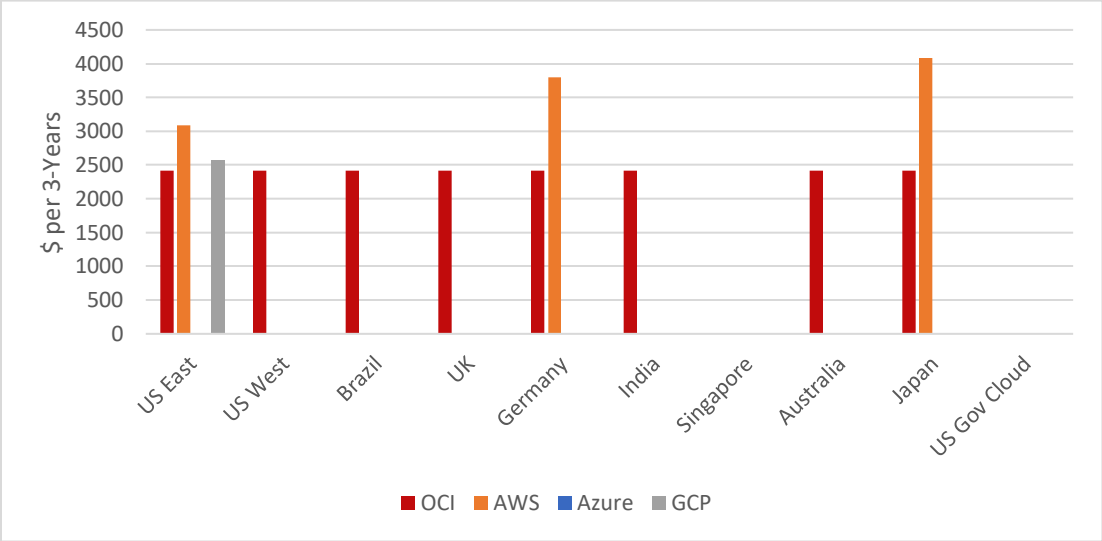

Chart 5: 3-Year Commitment Savings Plan Pricing Comparison Gen 4 AMD EPYC Processors – Lower is Better

As the chart and table show for US East, AWS comes in ~ **28%** more than OCI and GCP comes in at ~ **7%** more. It's important to note, OCI flexible instances will be much lower cost than these comparison. The finer granularity results in considerably greater savings.

	OCI			AWS		Azure		GCP	
	$/hr		$/3 yr	$/hr	$/3 yr	$/hr	$/3 yr	$/hr	$/3 yr
US East	vCPU	0.0150	2,417.76	0.1176	3,090.00			0.098064	2,577.12
	GB	0.002							
US West	vCPU	0.0150	2,417.76						
	GB	0.002							
Brazil	vCPU	0.0150	2,417.76						
	GB	0.002							
UK	vCPU	0.0150	2,417.76						
	GB	0.002							
Germany	vCPU	0.0150	2,417.76	0.1445	3,796.15				
	GB	0.002							
India	vCPU	0.0150	2,417.76						
	GB	0.002							
Singapore	vCPU	0.0150							
	GB	0.002							
Australia	vCPU	0.0150	2,417.76						
	GB	0.002							
Japan	vCPU	0.0150	2,417.76	0.1556	4,088.64				
	GB	0.002							
US Gov Cloud	vCPU	0.0150							
	GB	0.002							

Table 5: 3-Year Commitment Instance Savings Plan Pricing Comparison Gen 4 AMD EPYC Processors – Lower is Better

Results are similar with the older Gen 3 AMD EPYC processors.

1-Year Commitment Instance Savings Plan Pricing Comparison Gen 3 AMD EPYC Processors

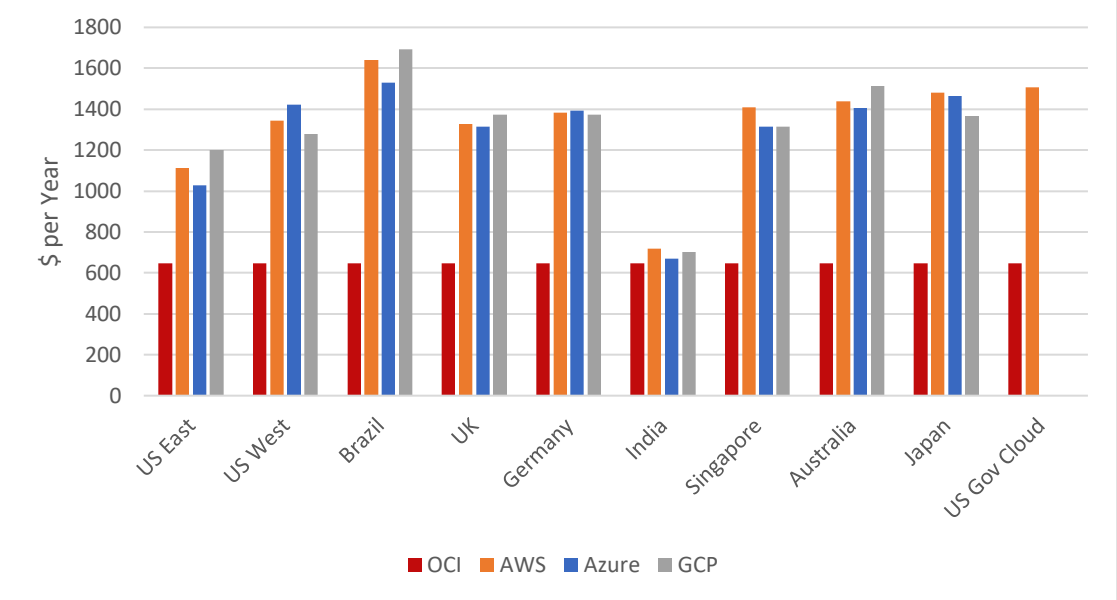

Chart 6: 1-Year Commitment Savings Plan Pricing Comparison Gen 3 AMD EPYC Processors – Lower is Better

This chart for Gen 3 AMD EPYC processors shows that once again, even with a 1-year commitment savings plan, OCI ranges from ~ **62%** less to ~ **46%** less across regions. The only region where AWS, Azure, and GCP get close is Mumbai in India. Here OCI is only ~ **3%** to ~ **10%** less. This again assumes the application workload is absurdly using 100% of the OCI flexible instance resources every minute of every day.

	OCI			AWS		Azure		GCP	
	$/hr		$/yr	$/hr	$/yr	$/hr	$/yr	$/hr	$/yr
US East	vCPU	0.0125							
	GB	0.0015	648.24	0.12695	1,112.08	0.1175	1,029.30	0.137030	1,200.39
US West	vCPU	0.0125							
	GB	0.0015	648.24	0.1534	1,343.52	0.1625	1,423.50	0.146137	1,280.16
Brazil	vCPU	0.0125							
	GB	0.0015	648.24	0.1871	1,639.05	0.1748	1,531.25	0.193136	1,691.87
UK	vCPU	0.0125							
	GB	0.0015	648.24	0.1517	1,329.07	0.1502	1,315.75	0.156750	1,373.13
Germany	vCPU	0.0125							
	GB	0.0015	648.24	0.1578	1,382.68	0.1590	1,392.84	0.156750	1,373.13
India	vCPU	0.0125							
	GB	0.0015	648.24	0.0820	718.32	0.0765	670.14	0.080300	703.43
Singapore	vCPU	0.0125							
	GB	0.0015	648.24	0.1608	1,408.96	0.1501	1,314.88	0.150097	1,314.85
Australia	vCPU	0.0125							
	GB	0.0015	648.24	0.1642	1,438.30	0.1605	1,405.98	0.172642	1,512.34
Japan	vCPU	0.0125							
	GB	0.0015	648.24	0.1691	1,480.97	0.1673	1,465.55	0.156090	1,367.35
US Gov Cloud	vCPU	0.0125							
	GB	0.0015	648.24	0.1720	1,506.72				

Table 6: 1-Year Commitment Instance Savings Plan Pricing Comparison Gen 3 AMD EPYC Processors – Lower is Better

The 3-year commitment comparison shows similar results in that OCI flexible instances cost less in every region.

3-Year Commitment Instance Savings Plan Pricing Comparison Gen 3 AMD EPYC Processors

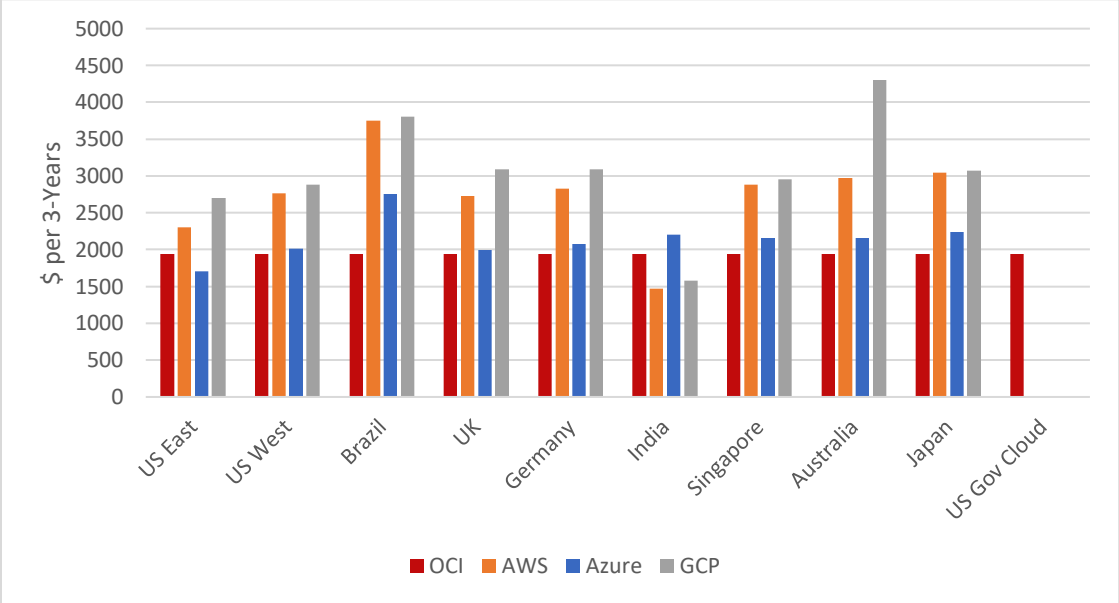

Chart 7: 3-Year Commitment Savings Plan Pricing Comparison Gen 3 AMD EPYC Processors – Lower is Better

	OCI			AWS		Azure		GCP	
	$/hr		$/3 yr	$/hr	$/3 yr	$/hr	$/3 yr	$/hr	$/3 yr
US East	vCPU	0.0125							
	GB	0.0015	1,944.72	0.0876	2,303.18	0.0654	1,718.71	0.102773	2,700.87
US West	vCPU	0.0125							
	GB	0.0015	1,944.72	0.1052	2,763.87	0.0767	2,015.68	0.109603	2,880.36
Brazil	vCPU	0.0125							
	GB	0.0015	1,944.72	0.1429	3,754.89	0.1049	2,756.77	0.144852	3,806.70
UK	vCPU	0.0125							
	GB	0.0015	1,944.72	0.1039	2,731.54	0.0760	1,997.28	0.117562	3,089.54
Germany	vCPU	0.0125							
	GB	0.0015	1,944.72	0.1077	2,829.57	0.0790	2,076.12	0.117562	3,089.54
India	vCPU	0.0125							
	GB	0.0015	1,944.72	0.0558	1,467.48	0.0837	2,199.64	0.060225	1,582.72
Singapore	vCPU	0.0125							
	GB	0.0015	1,944.72	0.1095	2,878.71	0.0821	2,157.59	0.112573	2,958.41
Australia	vCPU	0.0125							
	GB	0.0015	1,944.72	0.1131	2,972.01	0.0821	2,157.59	0.129481	3,402.77
Japan	vCPU	0.0125							
	GB	0.0015	1,944.72	0.1160	3,047.43	0.0851	2,236.43	0.117068	3,076.54
US Gov Cloud	vCPU	0.0125							
	GB	0.0015	1,944.72						

Table 7: 3-Year Commitment Instance Savings Plan Pricing Comparison Gen 3 AMD EPYC Processors – Lower is Better

OCI costs less than AWS and GCP in every region but Mumbai India and less than Azure in all regions but US East.

1-Year Commitment Instance Savings Plan Pricing Comparison Intel Xeon Ice Lake Processors

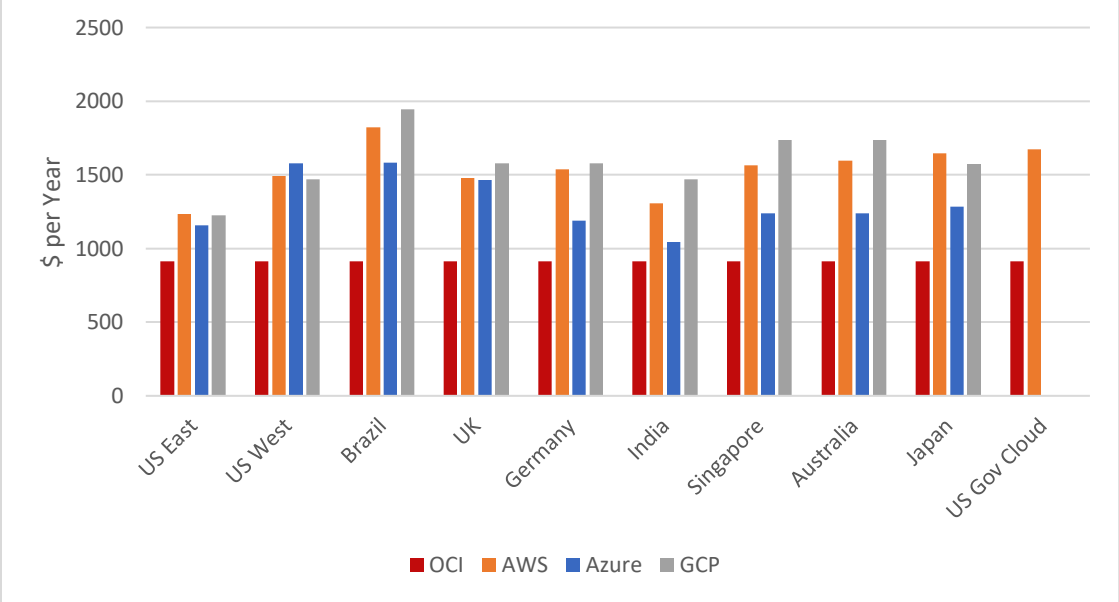

Chart 8: 1-Year Commitment Savings Plan Pricing Comparison Intel Xeon Ice Lake Processors – Lower is Better

As both the chart and table demonstrate, OCI on-demand flexible instance pricing for Intel Xeon Ice Lake processors delivers lower costs over a year than AWS, Azure, or GCP in all regions based on their 1-year commitment savings plan pricing.

	OCI			AWS		Azure		GCP	
	$/hr		$/yr	$/hr	$/yr	$/hr	$/yr	$/hr	$/yr
US East	vCPU	0.02	911.04	0.14110	1,235.69	0.1320	1,156.32	0.139824	1,224.86
	GB	0.0015							
US West	vCPU	0.02	911.04	0.1704	1,492.79	0.1800	1,576.80	0.167976	1,471.47
	GB	0.0015							
Brazil	vCPU	0.02	911.04	0.2079	1,821.20	0.1806	1,582.06	0.221976	1,944.51
	GB	0.0015							
UK	vCPU	0.02	911.04	0.1686	1,476.76	0.1670	1,462.92	0.180144	1,578.06
	GB	0.0015							
Germany	vCPU	0.02	911.04	0.1754	1,536.24	0.1357	1,188.73	0.180144	1,578.06
	GB	0.0015							
India	vCPU	0.02	911.04	0.1491	1,306.12	0.1192	1,044.19	0.167976	1,471.47
	GB	0.0015							
Singapore	vCPU	0.02	911.04	0.1787	1,565.50	0.1416	1,240.42	0.172512	1,511.21
	GB	0.0015							
Australia	vCPU	0.02	911.04	0.1824	1,598.09	0.1416	1,240.42	0.198432	1,738.26
	GB	0.0015							
Japan	vCPU	0.02	911.04	0.1879	1,645.57	0.1464	1,282.46	0.179424	1,571.75
	GB	0.0015							
US Gov Cloud	vCPU	0.02	911.04	0.1913	1,675.05				
	GB	0.0015							

Table 8: 1-Year Commitment Savings Plan Pricing Comparison Intel Xeon Ice Lake Processors – Lower is Better

3-Year Commitment Instance Savings Plan Pricing Comparison Intel Xeon Ice Lake Processors

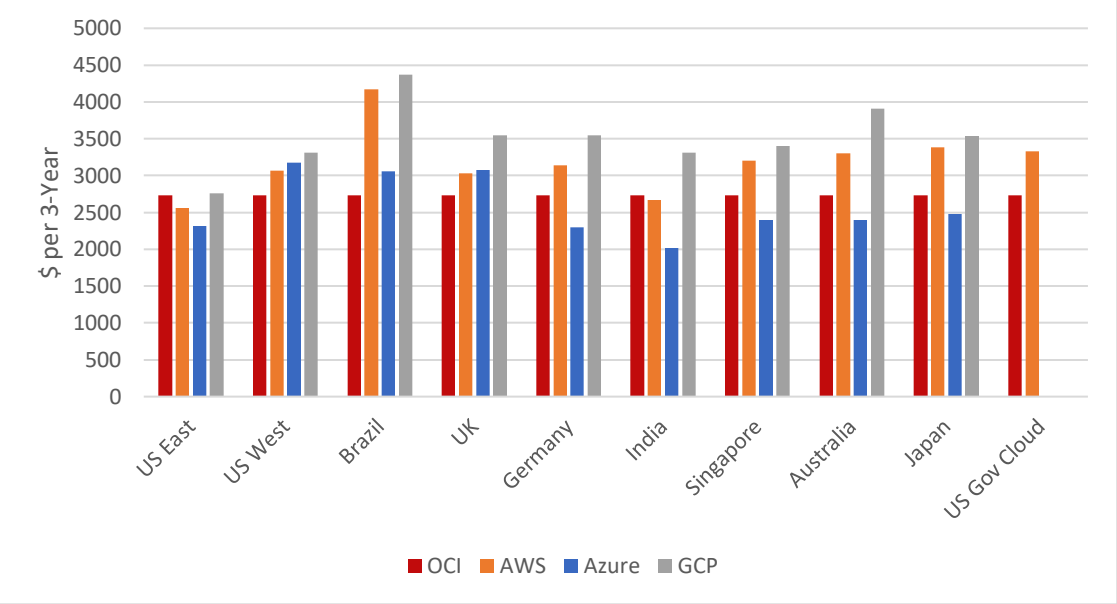

Chart 9: 3-Year Commitment Savings Plan Pricing Comparison Intel Xeon Ice Lake Processors – Lower is Better

The 3-year commitment and instance savings plans pricing chart and table for Intel Xeon Ice Lake Processors demonstrate that OCI is comparable in pricing to AWS, Azure, and GCP depending on the region. As previously discussed, real-world use and OCI flexible instances will usually be much lower cost.

	OCI			AWS		Azure		GCP	
	$/hr		$/3 yr	$/hr	$/3 yr	$/hr	$/3 yr	$/hr	$/3 yr
US East	vCPU	0.02	2733.12	0.0974	2,558.88	0.0880	2,312.64	0.104868	2,755.93
	GB	0.0015							
US West	vCPU	0.02	2733.12	0.1169	3,071.08	0.1190	3,172.32	0.125982	3,310.81
	GB	0.0015							
Brazil	vCPU	0.02	2733.12	0.1588	4,172.21	0.1163	3,056.36	0.166482	4,375.15
	GB	0.0015							
UK	vCPU	0.02	2733.12	0.1155	3,035.08	0.1170	3,074.76	0.135108	3,550.64
	GB	0.0015							
Germany	vCPU	0.02	2733.12	0.1196	3,143.88	0.0874	2,296.87	0.135108	3,550.64
	GB	0.0015							
India	vCPU	0.02	2733.12	0.1015	2,667.95	0.0767	2,015.68	0.125982	3,310.81
	GB	0.0015							
Singapore	vCPU	0.02	2733.12	0.1217	3,198.80	0.0912	2,396.74	0.129384	3,400.21
	GB	0.0015							
Australia	vCPU	0.02	2733.12	0.1257	3,302.08	0.0912	2,396.74	0.148824	3,911.09
	GB	0.0015							
Japan	vCPU	0.02	2733.12	0.1288	3,385.92	0.0943	2,478.20	0.134568	3,536.45
	GB	0.0015							
US Gov Cloud	vCPU	0.02	2733.12	0.1226	3,328.10				
	GB	0.0015							

Table 9: 3-Year Commitment Savings Plan Pricing Comparison Intel Xeon Ice Lake Processors – Lower is Better

IT organizations need to be aware about a commitment savings plan "gotcha" that can raise costs quite a bit. If the fixed size instances from AWS or Azure, and the flexible instances from GCP need to change to a smaller size. The cost does not decrease. The commitment does not change. The CSPs will usually let the customer increase their shape and commitment, just not decrease the commitment. OCI Flexible instances empower the customer to decrease or increase their shape with no financial penalties.

Yeah, But What About…?

The charts and tables confirm OCI flexible instances on-demand pricing definitively reduce cloud spend when compared to AWS, Azure, and GCP, even when compared against their 1-year and 3-year commitment savings plans. But there are some who will believe the comparisons are somehow rigged in some way to favor OCI. In reality, the comparison is heavily skewed in favor of AWS, Azure, and GCP specifically for the 1-year and 3-year commitment savings plans. This means OCI reduces cloud spend more than expressed in the charts and tables.

Then there are the doubters who say, "yeah, but what about?" What about other spot pricing, burstables, other configurations, and other discounts? Research reveals that these really do not move the needle. OCI costs are consistently and considerably much lower.

Spot Pricing

Spot pricing is incredibly cheap. AWS and Azure discount as much as 90% off of on-demand and GCP discounts as much as 91%. OCI discounts a flat 50%, which can still be less expensive. OCI Gen 4 and Gen 3 AMD processor spot pricing is generally lower than AWS and GCP while being higher than Asure. OCI Intel Xeon Ice Lake spot pricing is generally lower than AWS but higher than Azure and GCP. The reason this is qualified is because AWS changes its spot pricing every 5 minutes, Azure daily, and GCP roughly every 30 days. So their spot pricing discount changes constantly.

Because spot instances can be preempted anytime with minimal notice – 2 minutes from OCI and AWS, 30 seconds from Azure and GCP – for any reason whatsoever, spot instances cannot rationally be used for mission-critical or even business-critical applications. That means very few applications can use spot instances. They'll have at most a nominal impact on decreasing cloud spend.

Burstable Pricing

Burstables are genuinely a step up from spot pricing. But there is no guarantee that the instance's total resources above the customer's commitment level will actually be available, or whether the customer has enough credits to use those resources, or the customer needs more bursting time then they're allowed by the CSP. In other words,

the customer cannot count on resources above their commitment level because those resources are way oversubscribed; require the earn credits for time using fewer resources than their commitment level; and burstables have a hard limit of how much continuous bursting a given workload can run above its commitment.

As previously discussed, AWS, Azure, and GCP provide burstable pricing only on older and slower processors. OCI enables burstables on all processors.

Burstables are appropriate for less demanding workloads such as websites. They are not necessarily appropriate for ecommerce, any mission-critical, or business-critical applications, which can and do spike. If the burstable resource is not available or there are not enough crdits for what's needed, that performance dip has an enormous oversized negative impact on productivity, quality, time-to-actionable insights, time-to-unique-revenues, and time-to-profits.

Effective burstable use greatly depends on the demand profile, commitment level, configuration, processor type, memory, and CSP oversubscription. Customers must profile all their applications response time requirements and whether they are a fit for burstables. If they are a fit, can they tolerate older and slower processors or times when the resources they require will not be there.

Burstables can reduce some cloud spend when properly utilized. OCI has the most options, best performance, and lowest cost.

Other Configurations

The configuration used in the previous examples was not cherry picked. It was simply a very common configuration. A look at a configuration range in USA Virginia East regions, for Gen 3 and Gen 4 AMD EPYC processor hardware shows it was not a fluke. From 1 OCPU (2 vCPUs) through 32 OCPUs = 64vCPUs, using the same ratio of vCPUs to memory (GiB), standard on-demand pricing, running instances full for all 5 weekdays, and just 20 hours per each weekend day, OCI's cost remains consistently lower by substantial amounts – increasing as fixed shapes jump.

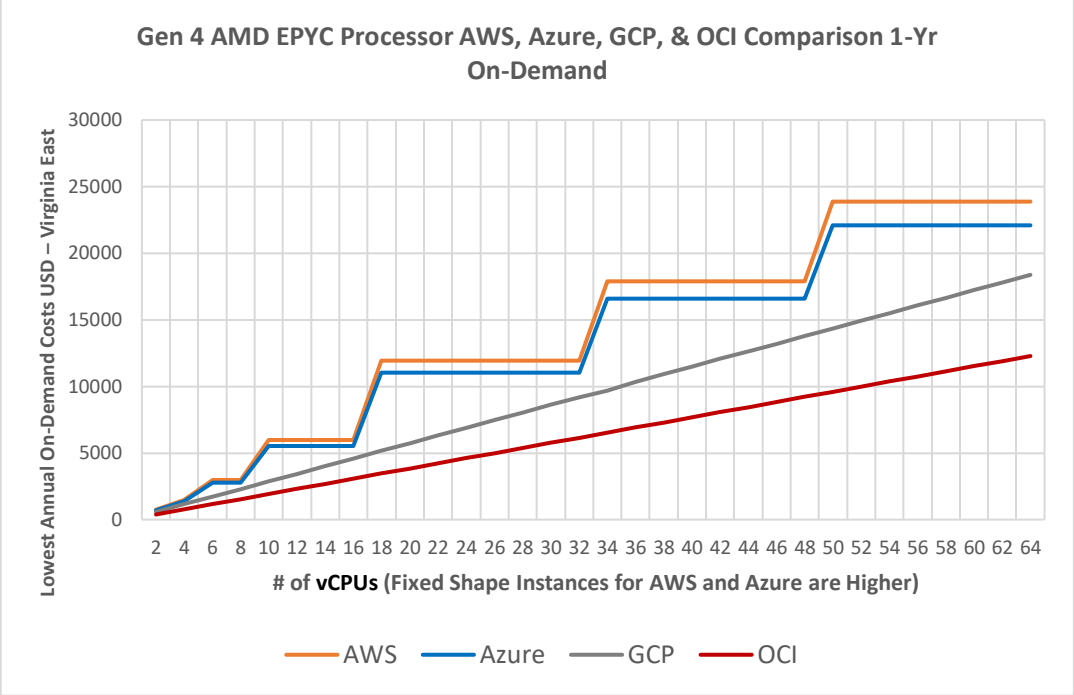

Chart 10: On-Demand Different Shape Instance Pricing Comparison Gen 4 AMD EPYC Processors – Lower is Better

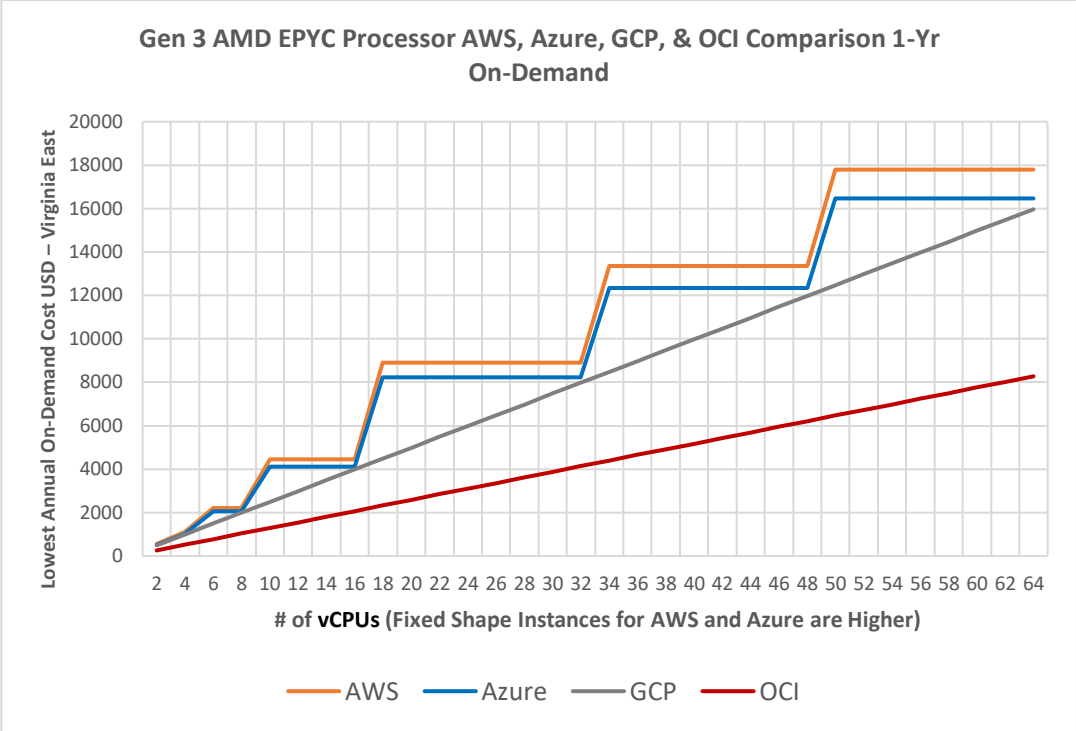

Chart 11: On-Demand Different Shape Instance Pricing Comparison Gen 3 AMD EPYC Processors – Lower is Better

As the scaling of the instances occur, the same OCI much-lower-cost ratios show up. There are also similar results for 1-year and 3-year commitments as well as Intel Xeon Ice Lake processors. Fixed configuration, coarse granularity doubles in size each time. The finer granularity of OCI and GCP grow at a steady rate. Most application workloads don't need to double the vCPUs and GiB when they max out the current shape, greatly increasing the cost gap.

Other CSP Discounting

Other discounting is complicated and comes in two primary flavors – account wide and usage based.

Discounts		OCI	AWS	Azure	GCP	
Account-wide discount						
		Universal Credits	Enterprise Discount Program (EDP)	Enterprise Agreement (EA) a.k.a. Azure Consumption Commitments	Committed Use Discount (CUD)	
Commitment Length		At least 1 yr	3 yrs or more	1-3 years or more	*20%* on-demand, *28%* 1-yr, *46%* 3-yr	
Usage based discount						
Typical committed usage discount	$.0 - $.1M/yr	*0%*	*Up to 7%*	*0%*	Sustained Use Discount (SUD) applies when there's no other discount & use > 25% of billing month. Discount increases incrementally up to 30% net discount off VM instances that run the entire month.	0-25%/mo usage
	$.1M - $.5M/yr	*Up to 5%*				*No Discount*
	$0.5-1M	*Up to 10%*		*Up to 5%*		25-50%/mo usage
	$1-5M	*Up to 15%*	*Up to 10%*	*Up to 10%*		*20% incremtal disc.*
	$5-10M	*Up to 20%*	*Up to 18%*	*Up to 20%*		50-75%/mo usage
	$10-30M	*Up to 25%*				*40% incremtal disc.*
	$30M+	*Up to 25%*	*Up to 21%*	*Varies*		75%+/mo usage
						60% incremtal disc.
Committed usage rollover		No Unused usage expires at the end of the term.	No Consumption measured yearly, & not over contract life.	No Any commitment added during the contract period is coterminus.	N/A	
Overconsumption pricing		Same discount	Same discount	Same discount	N/A	
("billing overage")		Once total amount of credits is exhausted, pricing remains at discounted price for contract term, but billed monthly	Once total amount of credits is exhausted, pricing remains at discounted price for contract term, but billed monthly	Price is protected during term of the EA, whether credits are available or not		

Table 10: CSP Discounting

Discounting Notes: GCP's CUD is included in the on-demand, 1-year commitment savings plan, and 3-year commitment savings plan pricing used in the price comparisons.

OCI Universal Credits are a highly flexible discount consumption model for all OCI infrastructure services. Customers purchase a prepaid amount of Universal Credits annually and the prepaid amount is draws down based on actual usage. Universal Credits may be applied to eligible services at any time in any OCI region – with the exception of Oracle EU Sovereign Cloud regions – at any time. The customer's OCI service usage is debited from their total contractual amount of Universal Credits. By definition, the Universal Credits must be consumed by the end of the contract term are forfeited.

A careful examination of the other CSP discounts reveal that they nominally affect the gap between AWS, Azure, GCP, and the much lower cost OCI. Discounts tend to be comparable and competitive. Some are simpler than others such as the OCI Universal Credits. But in the end, discounts are a means to a net price or cost. It is the net that matters, not the discount.

Getting the most out of any CSP discounts is essential to reducing cloud spend. However, some discounts come with caveats, as previously discussed, that can increase cloud spend based-on customer changes. OCI overall is still the best way to reduce cloud spend.

Block Data Storage Comparison Cost

As discussed earlier in this research, the block data storage comparison cost analysis is limited to balanced and high-performance block data storage. As noted, these storage types are the most common for mission-critical and business-critical applications.

As with the compute comparison, this research takes great pains to make sure any storage cost comparisons are as close as possible. There are several different balanced and high-performance block data storage options and pricing from AWS, Azure, and GCP. This research specifically chose the lowest cost option from each that provided the capacity and performance equivalence between them. When more than one option met the criteria, the lower cost option was chosen. OCI's choices are much simpler and cost substantially lower.

There are major differences in how each CSP prices out their storage for both balanced and high-performance. CSP Pricing is primarily based on capacity, IOPS, throughput, and region. Some CSP block storage have a fixed plus variable fees. Others are purely variable. Fixed commonly includes a preset amount of IOPS and throughput, whereas variable is priced per IOPS and throughput above and beyond what's included in the fixed.

Like their compute services, AWS, Azure, and GCP have different block storage fees in different regions. OCI does not. Therefore, this comparison wil look at 5 regions – Virginia USA, São Paulo Brazil, London UK, Tokyo Japan, and US Government, VA, where offered.

Based on these factors, here are the storage assumptions for both Balanced and High-Performance block storage..

Assumptions for Balanced Block Data Storage

Capacity	100 GB	500 GB	1000 GB	10000 GB
Minimum IOPS	6,000	25,000		
IOPS block size	4 KB			
Minimum Throughput	48 MB/s	240 MB/s		
Throughput block size	1 MB			
Timeframe	730 hrs			

Table 11: Common Assumptions for AWS, Azure, GCP, and OCI Balanced Block Data Storage

Other assumptions such as the service selected are vendor specific.

	AWS	Azure	GCP	OCI
Service	gp3	Premium SSD v2, LRS	Hyperdisk Balanced Generally - Extreme PD, zonal, LRS for 1 TB	Block

Table 12: CSP Block Balanced Performance Storage Services Compared

A quirk in GCP's pricing has the Extreme PD, zonal LRS lower cost than Hyperdisk Balanced at 1,000 GB (1 TB). All the other GCP configurations come in cheaper on Hyperdisk Balanced.

Assumptions for High-Performance Block Data Storage

Capacity	100 GB	500 GB	1000 GB	10000 GB
Minimum IOPS	7,500	37,500	75,000	375,000
IOPS block size	4 KB			
Minimum Throughput	59 MB/s	293 MB/s		

Throughput block size	1 MB
Timeframe	730 hrs

Table 13: Common Assumptions for AWS, Azure, GCP, and OCI High-Performance Block Data Storage

Other assumptions such as the service selected are vendor specific.

	AWS	Azure	GCP	OCI
Service	io2	Premium SSD v2, LRS Generally - Ultra Disk, LRS for 10 TB	Hyperdisk Balanced Generally - Hyperdisk Extreme for 10 TB	Block

Table 14: CSP Block High-Performance Storage Services Compared

Quirks in Azure and GCP pricing at 10,000 GB (10 TB) made the Azure's Ultra Disk LRS and GCP Hyperdisk Extreme lower priced than Azure Premium SSD v2, LRS and GCP Hyperdisk Balanced.

Balanced Block Data Storage Pricing Comparison

Balanced Performance	USA East VA		Sao Paulo Brazil		London UK		Tokyo Japan		US Gov't	
OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
100 GB	$ 4.25	0%	$ 4.25	0%	$ 4.25	0%	$ 4.25	0%	$ 4.25	0%
500GB	$ 21.25	0%	$ 21.25	0%	$ 21.25	0%	$ 21.25	0%	$ 21.25	0%
1,000 GB	$ 42.50	0%	$ 42.50	0%	$ 42.50	0%	$ 42.50	0%	$ 42.50	0%
10,000 GB	$ 425.00	0%	$ 425.00	0%	$ 425.00	0%	$ 425.00	0%	$ 425.00	0%
AWS	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
100 GB	$ 23.00	541%	$ 43.70	1028%	$ 26.68	628%	$ 27.60	649%	$ 27.60	649%
500GB	$ 154.60	728%	$ 293.74	1382%	$ 179.29	844%	$ 185.52	873%	$ 185.52	873%
1,000 GB	$ 194.60	458%	$ 369.74	870%	$ 225.69	531%	$ 233.52	549%	$ 233.52	549%
10,000 GB	$ 914.60	215%	$1,737.74	409%	$1,060.89	250%	$1,097.52	258%	$1,097.52	258%
Azure	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
100 GB	$ 23.70	558%	$ 43.70	1028%	$ 27.00	635%	$ 27.40	645%	$ 27.40	645%
500GB	$ 159.62	751%	$ 293.74	1382%	$ 181.59	855%	$ 183.94	866%	$ 183.94	866%
1,000 GB	$ 200.12	471%	$ 369.74	870%	$ 228.09	537%	$ 232.44	547%	$ 232.44	547%
10,000 GB	$ 929.12	219%	$1,737.74	409%	$1,065.09	251%	$1,105.44	260%	$1,105.44	260%
GCP	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
100 GB	$ 23.00	541%	$ 36.70	864%	$ 27.10	638%	$ 28.20	664%	Not Available	
500GB	$ 154.00	725%	$ 245.90	1157%	$ 182.10	857%	$ 188.10	885%		
1,000 GB	$ 187.00	440%	$ 255.00	600%	$ 204.00	480%	$ 221.00	520%		
10,000 GB	$ 914.00	215%	$1,452.40	342%	$1,046.60	246%	$1,157.10	272%		

Table 15: CSP Balanced Block Data Storage Pricing Comparison

As evident in table 15 above, the balanced block data storage pricing of AWS, Azure, and GCP are pretty comparable. But when compared to OCI pricing they are all considerably more expensive. This is true at all capacities in all regions. Chart 12 below makes it obvious at every capacity, OCI is much lower cost.

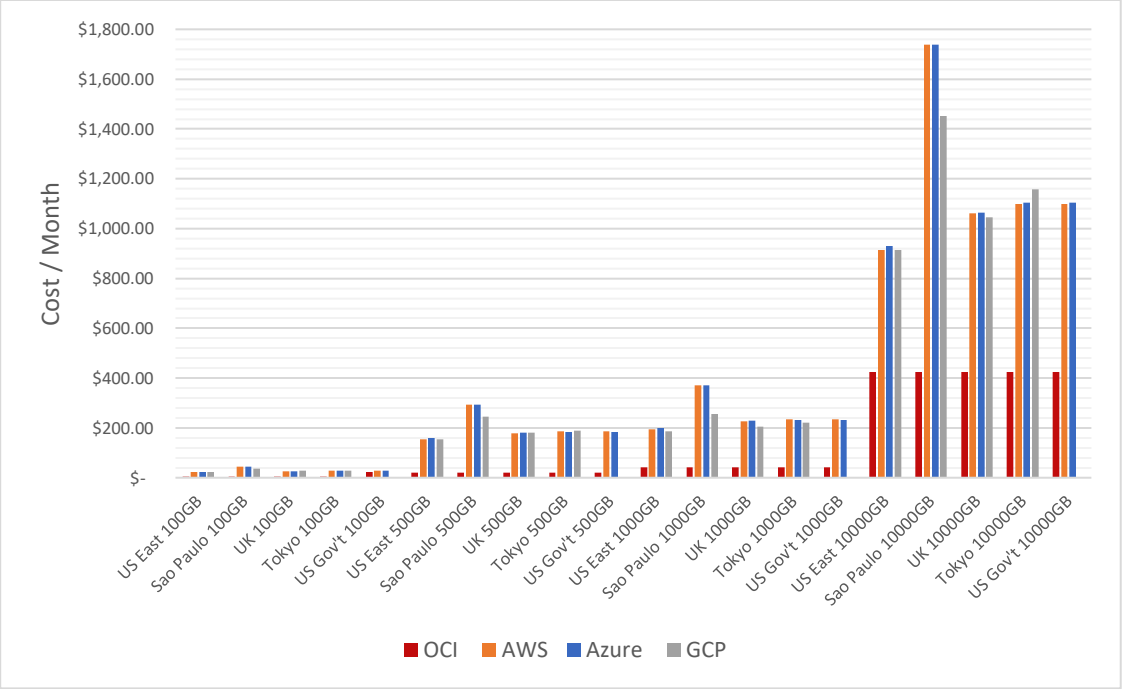

Chart 12: Pricing Comparison Block Data Storage Balanced Performance

High-Performance Block Data Storage Pricing Comparison

High Performance	USA East VA		Sao Paulo Brazil		London UK		Tokyo Japan	
OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
100 GB	$ 5.95	0%	$ 5.95	0%	$ 5.95	0%	$ 5.95	0%
500GB	$ 29.75	0%	$ 29.75	0%	$ 29.75	0%	$ 29.75	0%
1,000 GB	$ 59.50	0%	$ 59.50	0%	$ 59.50	0%	$ 59.50	0%
10,000 GB	$ 297.50	0%	$ 297.50	0%	$ 297.50	0%	$ 297.50	0%
AWS	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
100 GB	$ 500.00	11765%	Not Available		$ 570.00	9580%	$ 570.00	9580%
500GB	$ 2,395.50	11273%			$ 2,796.00	9398%	$ 2,796.00	9398%
1,000 GB	$ 4,029.00	9480%			$ 4,680.00	7866%	$ 4,680.00	7866%
10,000 GB	$ 10,321.00	2428%			$11,957.00	4019%	$28,086.52	9441%
Azure	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
100 GB	$ 31.50	741%	$ 57.95	1364%	$ 35.85	844%	$ 36.25	853%
500GB	$ 224.62	1057%	$ 412.49	1941%	$ 255.34	1202%	$ 257.69	1213%
1,000 GB	$ 460.12	1083%	$ 844.74	1988%	$ 523.09	1231%	$ 527.44	1241%
10,000 GB	$ 19,316.05	4545%	$38,622.48	9088%	$24,135.92	5679%	$28,086.52	6609%
GCP	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
100 GB	$ 37.20	875%	$ 48.70	1146%	$ 36.10	849%	$ 37.20	875%
500GB	$ 265.80	1251%	$ 349.29	1644%	$ 259.54	1221%	$ 265.80	1251%
1,000 GB	$ 541.80	1275%	$ 712.79	1677%	$ 530.04	1247%	$ 541.80	1275%
10,000 GB	$ 11,825.00	2782%	$18,845.00	4434%	$13,315.00	3133%	$15,150.00	3565%

Table 16: CSP High-Performance Block Data Storage Pricing Comparison

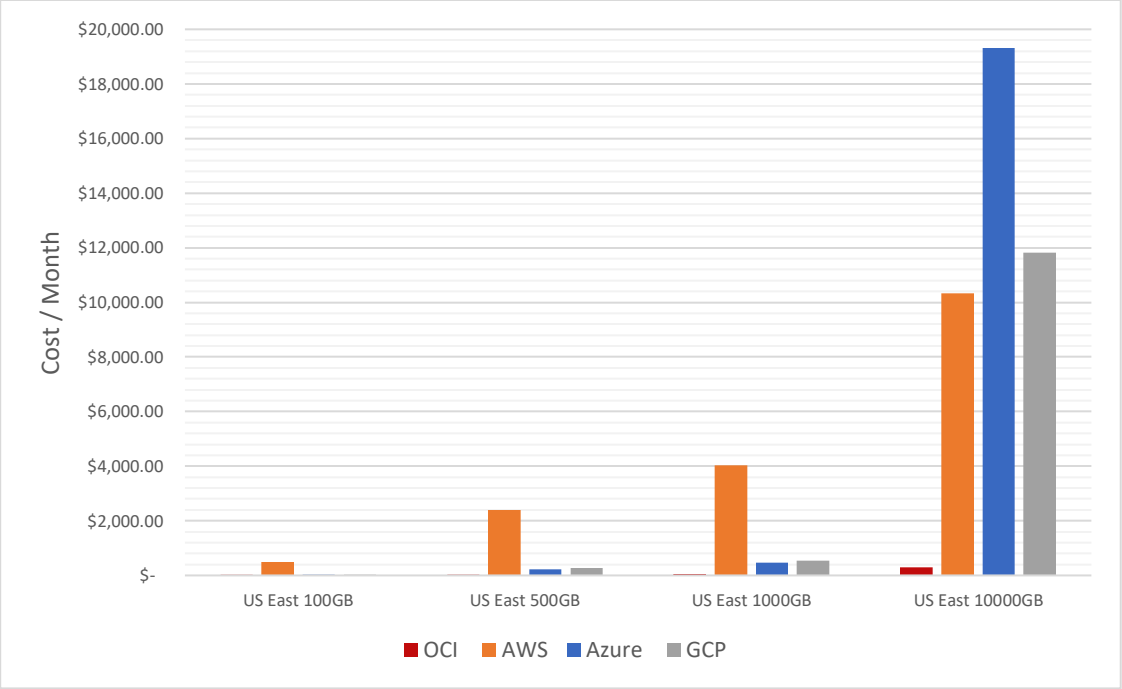

Chart 13: Pricing Comparison Block Data Storage High-Performance For US East (Lower is better)[9]

Clearly OCI block data storage pricing is dramatically lower than AWS, Azure, and GCP. In fact, it is so much lower when comparing high-performance block data storage pricing that it is difficult to see OCI on the chart.

Egress Comparison Cost

Egress costs are a way of life for cloud storage users. If the data is stored in AWS, Azure, GCP, or OCI, it's going to cost money to move it out of that cloud. A 2021 IDC survey found that 99% of cloud storage users have incurred egress fees averaging 6% of their cloud storage costs. Another poll from Global Market Intelligence found that 34% of enterprises using cloud storage had been affected by egress fees. This caused many of them to repatriate data on-premises or shift to service providers with lower egress costs.

The question becomes how costly are those egress fees can be. AWS, Azure, GCP, and OCI have different fees and rules that affect the cost. For example, GCP will not charge any egress fees if the customer is completely terminating their business with GCP and moving their data to another cloud or on-premises. OCI does not charge any egress fees for the first 10 TB per month of data being moved out using the public gateway.

A look at the these top 4 cloud service providers demonstrates a huge difference in egress cloud spin.

Egress Public Gateway Assumptions

Capacities measured	1 TB		10 TB		50 TB		100 TB		500 TB	
Regions data moves from	USA	South America	Europe	Middle East	Africa	APAC India	APAC Singapore	APAC Japan	Oceania (Australia)	US Gov't
Multiple locations/region	Cheapest Location Chosen									
Multiple tiers	Standard Tier is Chosen									
Timeframe	730 hrs									

Table 17: Egress Across Public Gateway Assumptions

The following Egress Public gatway pricing comparisons are sampled from 10 different regions worldwide.

[9] Only 1 region (US East) was used in the chart to make it readable. All other regions will show relatively similar ratios

Egress Pubic GW	From USA		From South America		From Europe		From Middle East		From Africa	
OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
1 TB	$ -	0%	$ -	0%	$ -	0%	$ -	0%	$ -	0%
10 TB	$ -	0%	$ -	0%	$ -	0%	$ -	0%	$ -	0%
50 TB	$ 340	0%	$ 1,000	0%	$ 340	0%	$ 2,000	0%	$ 2,000	0%
100 TB	$ 765	0%	$ 2,250	0%	$ 765	0%	$ 4,500	0%	$ 4,500	0%
500 TB	$ 4,165	0%	$ 12,250	0%	$ 4,165	0%	$ 24,500	0%	$ 24,500	0%
AWS	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
1 TB	$ 81	Infinite	$ 135	Infinite	$ 81	Infinite	$ 99	Infinite	$ 139	Infinite
10 TB	$ 891	Infinite	$ 1,485	Infinite	$ 891	Infinite	$ 1,089	Infinite	$ 1,525	Infinite
50 TB	$ 4,291	1262%	$ 7,005	701%	$ 4,291	1262%	$ 4,489	224%	$ 7,405	370%
100 TB	$ 7,791	1018%	$ 13,305	591%	$ 7,791	1018%	$ 8,339	185%	$ 13,705	305%
500 TB	$ 28,791	691%	$ 59,505	486%	$ 28,791	691%	$ 31,439	128%	$ 59,205	242%
Azure	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
1 TB	$ 72	Infinite	$ 108	Infinite	$ 72	Infinite	$ 99	Infinite	$ 99	Infinite
10 TB	$ 792	Infinite	$ 1,188	Infinite	$ 792	Infinite	$ 1,089	Infinite	$ 1,089	Infinite
50 TB	$ 3,392	998%	$ 4,588	459%	$ 3,392	998%	$ 4,089	204%	$ 4,089	204%
100 TB	$ 6,392	836%	$ 8,588	382%	$ 6,392	836%	$ 7,589	169%	$ 7,589	169%
500 TB	$ 23,392	562%	$ 38,838	317%	$ 23,392	562%	$ 32,089	131%	$ 32,089	131%
GCP	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
1 TB	$ 68	Infinite	$ 96	Infinite	$ 68	Infinite	$ 88	Infinite	$ 88	Infinite
10 TB	$ 833	Infinite	$ 1,176	Infinite	$ 833	Infinite	$ 1,078	Infinite	$ 1,078	Infinite
50 TB	$ 3,433	1010%	$ 4,576	458%	$ 3,433	1010%	$ 4,078	204%	$ 4,678	234%
100 TB	$ 6,683	874%	$ 8,826	392%	$ 6,683	874%	$ 7,828	174%	$ 9,178	204%
500 TB	$ 25,683	617%	$ 41,076	335%	$ 25,683	617%	$ 36,078	147%	$ 34,678	142%

Table 18: Egress Across Public Gateway Cost Comparison over 5 Regions

Egress Pubic GW	From Asia Pacific (India)		From Asia Pacific (Singapore)		From Oceania (Australia)		From Asia Pacific (Japan)		From US Government	
OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
1 TB	$ -	0%	$ -	0%	$ -	0%	$ -	0%	$ -	0%
10 TB	$ -	0%	$ -	0%	$ -	0%	$ -	0%	$ -	0%
50 TB	$ 1,000	0%	$ 1,000	0%	$ 1,000	0%	$ 1,000	0%	$ 1,000	0%
100 TB	$ 2,250	0%	$ 2,250	0%	$ 2,250	0%	$ 2,250	0%	$ 2,250	0%
500 TB	$ 12,250	0%	$ 12,250	0%	$ 12,250	0%	$ 12,250	0%	$ 12,250	0%
AWS	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
1 TB	$ 99	Infinite	$ 108	Infinite	$ 103	Infinite	$ 103	Infinite	$ 140	Infinite
10 TB	$ 1,089	Infinite	$ 1,188	Infinite	$ 1,129	Infinite	$ 1,129	Infinite	$ 1,535	Infinite
50 TB	$ 4,689	469%	$ 4,588	459%	$ 5,049	505%	$ 4,689	469%	$ 6,135	613%
100 TB	$ 8,689	386%	$ 8,688	386%	$ 9,749	433%	$ 8,989	399%	$ 10,635	473%
500 TB	$ 40,689	332%	$ 40,788	333%	$ 46,649	381%	$ 42,689	348%	$ 37,885	309%
Azure	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
1 TB	$ 99	Infinite	$ 99	Infinite	$ 99	Infinite	$ 99	Infinite	$ 72	Infinite
10 TB	$ 1,089	Infinite	$ 1,089	Infinite	$ 1,089	Infinite	$ 1,089	Infinite	$ 792	Infinite
50 TB	$ 4,289	429%	$ 4,089	409%	$ 4,089	409%	$ 4,089	409%	$ 3,392	339%
100 TB	$ 7,789	346%	$ 7,589	337%	$ 7,589	337%	$ 7,589	337%	$ 6,392	284%
500 TB	$ 32,289	264%	$ 32,089	262%	$ 32,089	262%	$ 32,089	262%	$ 23,392	191%
GCP	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
1 TB	$ 88	Infinite	$ 88	Infinite	$ 96	Infinite	$ 88	Infinite	Not Available	
10 TB	$ 1,078	Infinite	$ 1,078	Infinite	$ 1,176	Infinite	$ 1,078	Infinite		
50 TB	$ 4,078	408%	$ 4,078	408%	$ 4,576	458%	$ 4,078	408%		
100 TB	$ 7,828	348%	$ 7,828	348%	$ 8,826	392%	$ 7,828	348%		
500 TB	$ 36,078	295%	$ 36,078	295%	$ 41,076	335%	$ 36,078	295%		

Table 19: Egress Across Public Gateway Cost Comparison over and Additional 5 Regions

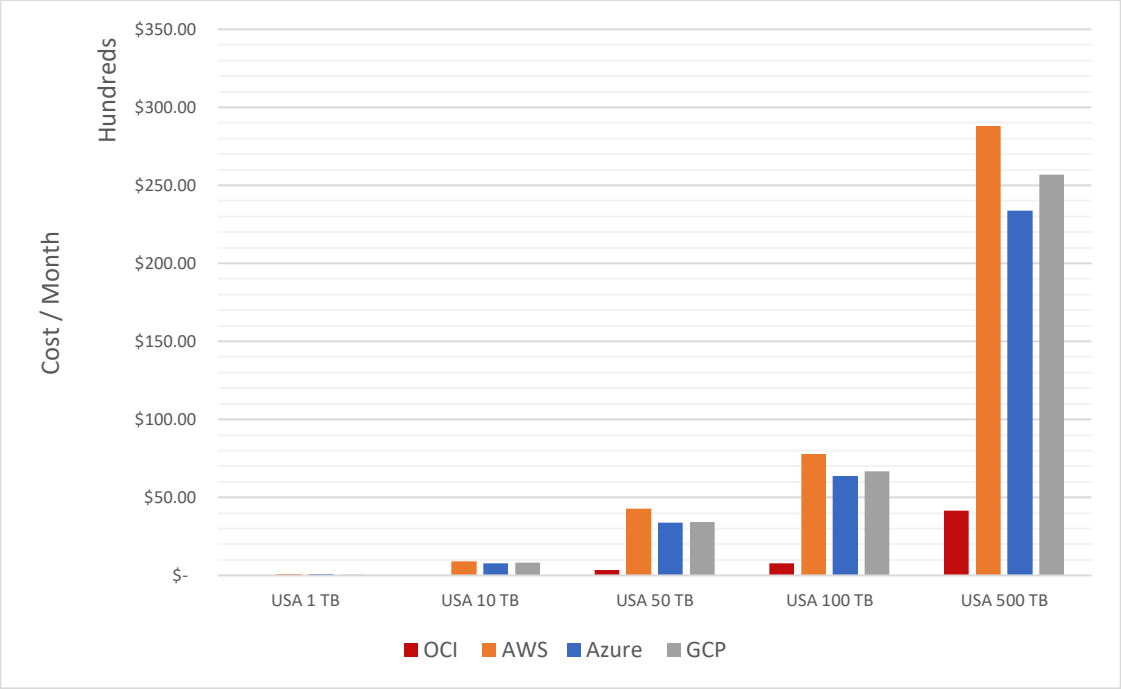

Chart 14: Pricing Comparison Egress Across Public Gateway For From USA (Lower is better)

The Egress pricing through the public gateways demonstrates once again the significantly lower OCI cost when compared with AWS, Azure, and GCP. The cost gap between AWS, Azure, GCP and OCI is in a word…huge.

Egress Private Line Assumptions

Max data moved/mo	3285 TB					
Percentage moved	10%	15%	20%	30%	50%	75%
Multiple locations/region	Cheapest Location Chosen					
Multiple tiers	Standard Tier is Chosen					
Timeframe	730 hrs					
Dedicated Line Bandwidth	10Gbps					
AWS Metered	Port Rate & Transfer					
Azure Metered & Unlimited	Port Rate & Transfer or unlimited					
GCP Metered	Port Rate & Transfer					
OCI Metered	Port Rate only					

Table 20: Egress Across Private Line Assumptions

The cost comparisons don't include cost of the private line the customer needs to purchase from their local Telco.

Egress Private Line	From USA		From South America		From Europe		From Middle East		From Africa	
OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
10%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%
15%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%
20%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%
30%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%
50%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%
75%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%
AWS	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
10%	$ 8,213	882%	$ 50,918	5471%	$ 10,906	1172%	$ 37,778	4059%	$ 37,778	4059%
15%	$ 11,498	1235%	$ 75,555	8118%	$ 15,538	1669%	$ 55,845	6000%	$ 55,845	6000%
20%	$ 14,783	1588%	$ 100,193	10765%	$ 20,170	2167%	$ 73,913	7941%	$ 73,913	7941%
30%	$ 21,353	2294%	$ 149,468	16059%	$ 29,434	3162%	$ 110,048	11824%	$ 110,048	11824%
50%	$ 34,493	3706%	$ 248,018	26647%	$ 47,961	5153%	$ 182,318	19588%	$ 182,318	19588%
75%	$ 50,918	5471%	$ 371,205	39882%	$ 71,120	7641%	$ 272,655	29294%	$ 272,655	29294%
Azure	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
10%	$ 11,613	1248%	$ 49,390	5306%	$ 11,613	1248%	$ 49,390	5306%	$ 49,390	5306%
15%	$ 15,719	1689%	$ 72,385	7777%	$ 15,719	1689%	$ 72,385	7777%	$ 72,385	7777%
20%	$ 19,825	2130%	$ 82,000	8810%	$ 19,825	2130%	$ 82,000	8810%	$ 82,000	8810%
30%	$ 28,038	3012%	$ 82,000	8810%	$ 28,038	3012%	$ 82,000	8810%	$ 82,000	8810%
50%	$ 44,463	4777%	$ 82,000	8810%	$ 44,463	4777%	$ 82,000	8810%	$ 82,000	8810%
75%	$ 51,300	5512%	$ 82,000	8810%	$ 51,300	5512%	$ 82,000	8810%	$ 82,000	8810%
GCP	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
10%	$ 8,269	888%	$ 37,834	4065%	$ 8,269	888%	$ 34,549	3712%	$ 37,834	4065%
15%	$ 11,554	1241%	$ 55,902	6006%	$ 11,554	1241%	$ 50,974	5477%	$ 55,902	6006%
20%	$ 14,839	1594%	$ 73,969	7947%	$ 14,839	1594%	$ 67,399	7241%	$ 73,969	7947%
30%	$ 21,409	2300%	$ 110,104	11830%	$ 21,409	2300%	$ 100,249	10771%	$ 110,104	11830%
50%	$ 34,549	3712%	$ 182,374	19594%	$ 34,549	3712%	$ 165,949	17830%	$ 182,374	19594%
75%	$ 50,974	5477%	$ 272,712	29300%	$ 50,974	5477%	$ 248,074	26653%	$ 272,712	29300%

Table 21: Egress Across Private Line Cost Comparison over 5 Regions

Egress Private Line	From Asia Pacific (India)		From Asia Pacific (Singapore)		From Oceania (Australia)		From Asia Pacific (Japan)		From US Government	
OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
10%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%
15%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%
20%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%
30%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%
20%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%
30%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%	$ 931	0%
AWS	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
10%	$ 29,565	3176%	$ 31,208	3353%	$ 44,348	4765%	$ 31,129	3344%	$ 8,213	882%
15%	$ 43,526	4676%	$ 45,990	4941%	$ 65,700	7059%	$ 45,911	4933%	$ 11,498	1235%
20%	$ 57,488	6176%	$ 60,773	6529%	$ 87,053	9353%	$ 60,694	6521%	$ 14,783	1588%
30%	$ 85,410	9176%	$ 90,338	9706%	$ 129,758	13941%	$ 90,259	9697%	$ 21,353	2294%
50%	$ 141,255	15176%	$ 149,468	16059%	$ 215,168	23118%	$ 149,389	16050%	$ 34,493	3706%
75%	$ 211,061	22676%	$ 223,380	24000%	$ 321,930	34588%	$ 223,301	23992%	$ 50,918	5471%
Azure	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
10%	$ 19,825	2130%	$ 19,825	2130%	$ 19,825	2130%	$ 19,825	2130%	$ 12,463	1339%
15%	$ 28,038	3012%	$ 28,038	3012%	$ 28,038	3012%	$ 28,038	3012%	$ 16,569	1780%
20%	$ 36,250	3895%	$ 36,250	3895%	$ 36,250	3895%	$ 36,250	3895%	$ 20,675	2221%
30%	$ 52,675	5659%	$ 52,675	5659%	$ 52,675	5659%	$ 52,675	5659%	$ 28,888	3104%
50%	$ 82,000	8810%	$ 82,000	8810%	$ 82,000	8810%	$ 82,000	8810%	$ 45,313	4868%
75%	$ 82,000	8810%	$ 82,000	8810%	$ 82,000	8810%	$ 82,000	8810%	$ 64,125	6890%
GCP	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI	Price/mo	% > OCI
10%	$ 16,482	1771%	$ 17,599	1891%	$ 15,496	1665%	$ 15,496	1665%	Not Available	
15%	$ 23,873	2565%	$ 25,549	2745%	$ 22,395	2406%	$ 22,395	2406%		
20%	$ 31,264	3359%	$ 33,498	3599%	$ 29,293	3147%	$ 29,293	3147%		
30%	$ 46,047	4947%	$ 49,398	5307%	$ 43,090	4630%	$ 43,090	4630%		
50%	$ 75,612	8124%	$ 81,196	8724%	$ 70,684	7594%	$ 70,684	7594%		
75%	$ 112,568	12094%	$ 120,945	12994%	$ 105,177	11300%	$ 105,177	11300%		

Table 22: Egress Across Private Line Cost Comparison over and Additional 5 Regions

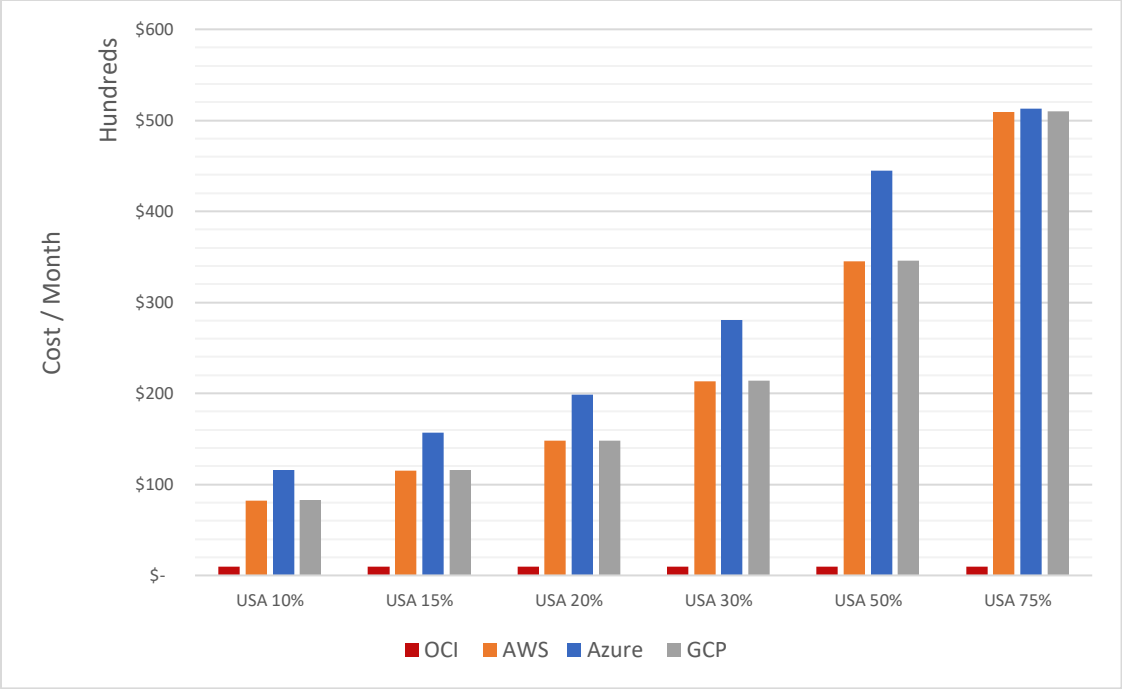

Chart 15: Pricing Comparison Egress Across Private Line For From USA (Lower is better)

The Egress pricing through a private line demonstrates a considerably bigger gap between OCI and AWS, Azure, or GCP. OCI is hands down much lower cost.

Conclusion

The cost numbers for compute, storage, and egress make it crystal clear. OCI significantly and dramatically reduces customer cloud spend. And that's before customers take into account database or SaaS services which are much lower cost on OCI than what AWS, Azure, or GCP offer. This is why OCI customers who have moved from the other clouds commonly save more than 50% on their cloud spend. Many save a heck of a lot more.

Any organization looking to reduce their cloud spend without compromising on performance, flexibility, reliability, and security should take a very long look at OCI.

For Information Go To

OCI

OCI Cloud Economics

OCI Pricing

