ESG Lab White Paper

Price Comparison: Google Cloud Platform vs. Amazon Web Services

By Aviv Kaufmann and Kerry Dolan, ESG Lab Analysts

June 2015

This ESG White Paper was commissioned by Google and is distributed under license from ESG.

Contents

ESG Lab Reports

The goal of ESG Lab reports is to educate IT professionals about data center technology products for companies of all types and sizes. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments. This ESG Lab White Paper was sponsored by Google.

Executive Summary

Cloud computing has established a strong foothold in organizations across the globe, and it's easy to see why. Virtualization enables organizations to use remote cloud computing services to augment—and even replace—onsite infrastructure. This saves on equipment, management, and data center floor space, and enables organizations to gain greater business agility and flexibility.

While cloud computing features vary among providers, the cost is always a key factor. ESG Lab was asked to compare Google Cloud Platform and Amazon Web Services pricing structures, and to validate cost differentials based on both companies' pricing calculators. We validated cost savings for various virtual machine (VM) instance types, but based our analyses on equivalent VMs with comparable 2.5-2.6 GHz Intel processors. Additional networking, middleware, and database costs were not included in this evaluation.

ESG Lab evaluated the online, public pricing calculators for Google Cloud Platform and Amazon Web Service EC2 and used the information to create a comprehensive, three-year cost model with a goal of comparing the many options between the two pricing models. Our analyses found consistent price advantages for Google. Google's on-demand, real-time pricing is simple and straightforward, with discounts for sustained usage, and a stated commitment to pass future price reductions on to customers. Amazon's pricing structure is complex, and provides substantial discounts only with long-term commitments and up-front payments. As a result, AWS pricing can negatively impact the agility and flexibility benefits for which organizations choose the cloud.

ESG Lab modeled pricing for multiple scenarios, from a single instance VM to various sizes of production applications, and included both "always-on" static workloads and dynamic workloads consisting of peak demand windows. Our modeling determined that realistic pricing for every scenario resulted in a Google price advantage, regardless of which AWS pricing model was selected (see Figure 1).

FIGURE 1. Google Cloud Platform Pricing Advantage vs. Amazon Web Services

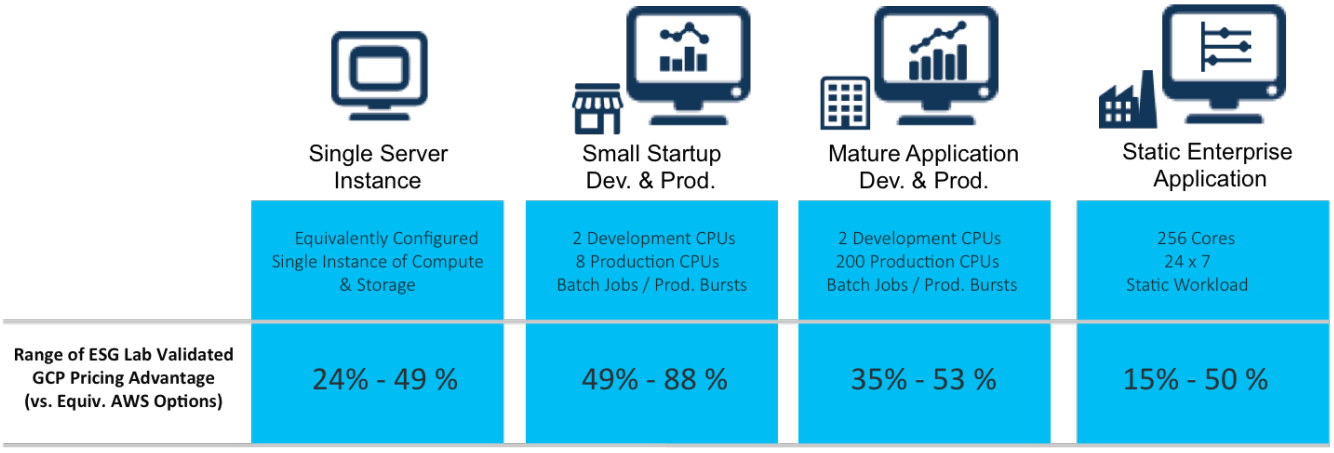

Source: Enterprise Strategy Group, 2015.

Why Cloud Computing

It wasn't long ago that pundits were still arguing over how to define "cloud computing." Today, the message from organizations around the globe is loud and clear: Publicly available cloud services, such as infrastructure-as-a-service (IaaS), are mainstream, essential parts of the IT armory. They tried it, they liked it, and they want more. Every company has its own reasons for using cloud computing, but they boil down to three critical values: Cloud computing can *reduce costs* while enabling *greater business agility* and *flexibility.*

- **Reduced costs**. Cloud computing enables organizations to pay only for what they use. Instead of standing up dedicated infrastructure to run each application, you spin up virtual machines on infrastructure owned and managed by a provider. You pay for the time you use the VMs, instead of setting up servers on your own infrastructure. You save on power, cooling, and floor space; you save on management since you don't have to install, operate, and troubleshoot it yourself. And you're not depreciating the equipment—someone else is. The ability to start small and grow organically as your business requires it, instead of having to guess at what you'll need next week, next month, and next year, lets you match your costs with actual usage. In addition, your computing costs in the cloud are usually operational expenses paid monthly rather than hefty up-front capital expenses.

- **Greater business agility**. Agility is really about responsiveness. Cloud computing lets you respond quickly to business opportunities and threats. What if your product team suddenly figures out how to make the ultimate widget? Or a competitor suddenly starts gaining on your market share? You can scale quickly in the cloud, adding VMs to cover spikes in production and ramp up sales. With physical infrastructure, scaling is often a lengthy process that starts with requisition, justification to senior management, and purchase, followed by waiting for delivery, and then managing deployment, testing, re-configuration, and, finally, production. Equally important, in the cloud you can scale back down when a utilization spike has passed. With strictly physical infrastructure, you've made an investment that likely sits idle waiting for another spike.

- **Flexibility**. Flexibility gives you choice. With the cloud, you can instantiate or destroy VM instances as you need to, move workloads around, and change your mind and revert—without wasting already purchased resources. You can move, resize, consolidate, and make choices to optimize any business metric.

For these reasons, cloud computing gives organizations both freedom and control for the workloads running on it.

Strong Cloud Foothold in the Market

ESG demand-side research with enterprise and midmarket organizations makes the status of cloud computing—and customers' reasons for choosing it—very clear. Our research shows that cloud computing has been moving up the IT priority list, and in 2015, 66% of survey respondents reported that they expect to increase spending on cloud computing services this year.[1] Its rise appears to be tied directly to its cost-saving impact. Since the global financial crisis of 2008, ESG has provided the same list of nine cost containment methods as responses in our research survey, and over that span, the use of cloud computing services has moved from the least commonly selected technique in 2009 to the second most selected in 2015. Other cost containment methods that trailed cloud computing include purchasing new technologies with improved ROI; postponing projects; and changes to headcount/hiring.

Another statistic relates to the adoption of infrastructure-as-a-service. Our research revealed that 40% of respondents currently use IaaS, with another 29% planning to do so. This means that more than two out of three organizations are willing to pay another IT organization in order to gain increased agility and flexibility along with expected cost savings.

[1] Source: ESG Research Report, *2015 IT Spending Intentions Survey*, February 2015. All ESG research references in this white paper have been taken from this research report, unless otherwise noted.

Google Cloud Platform and Amazon Web Services

If you're looking for cloud computing/IaaS services, you won't find two more renowned and respected providers than Google and Amazon. Both have a proven track record of delivering world-class, online solutions that are very well-known—Google Apps and search, Amazon e-commerce, and more. Both organizations have leveraged their advanced infrastructures to deliver public computing services, with Google Cloud Platform (GCP) and Amazon Web Services (AWS). A full description of the features of these services is beyond the scope of this paper, but at a high level, they offer remote computing services via virtual machine instances in various sizes and configurations. And they each have advantages: Google Cloud Platform tends to have an advantage in storage and network performance, while Amazon Web Services has an advantage in cloud features and points of presence around the globe.

As for any solution, cost is a key differentiator. ESG Lab evaluated the pricing structures for GCP and AWS using publicly available information, and our goal in this paper is to provide a high-level overview based on our own modeling and analyses. We have divided our analyses into three areas: cost-containment, agility, and flexibility. Our methodology included comparisons of effective monthly, yearly, and three-year costs to ensure an "apples to apples" comparison. In this paper, we have selected several representative pieces of our model to illustrate the most important differences between pricing concepts, including comparisons using a single VM instance, a startup application, and a mature application.

Cost Containment

It's surely obvious that organizations strive to keep costs down, but success and growth also depend on investing in the future of an organization. The challenge is to invest wisely—to use your money for the greatest value. Productivity is paramount in business today, as organizations continually ask employees to do more with less. Optimizing your financial resources is essential to getting the most out of every corporate dollar/euro/yuan. As ESG research demonstrates, more organizations are turning to cloud computing as a key strategy for cost reduction.

Cost starts with price, but involves so much more. For cloud computing services, it's important to understand how vendors determine prices, what factors are taken into consideration, and whether you pay once, over time, or a combination of both. Some of these are decided by the vendor, while others are customer options. Other cost factors may be less well-known at the time of purchase and include how much you actually utilize the resources, what you are able to do with them, and what the opportunity costs are of using your money one way versus another. We'll cover all of that in this paper.

Let's start with the simplest cost comparison between Google Cloud Platform and Amazon Web Services. Google's cost structure is simple and straightforward. AWS' is not.

Google Cloud Platform: Simple Pricing with Automatic Reductions

A key feature of GCP pricing is that there is only one method—you pay monthly for on-demand usage of virtual machine instances. The minimum is ten minutes, and usage is rounded up to the nearest minute. In addition, Google has publicly committed to passing along to customers any future price reductions Google achieves through technology-driven advancements in density, scale, power, and cooling. Just a look at its blog will reveal the history of price reductions. What Google provides is on-demand, real-time pricing.

In addition, Google Sustained Usage discounts are designed so that the more you use, the greater the discount. Google combines all your VMs into usage units that maximize the discount possibility—if you shut down a VM for part of the month, another VM's usage fills the gap to achieve sustained usage. Here's a simple example: You have four VMs, two running from June 1-15 and then shutting down, and another two running from June 15-30. From a pricing perspective, Google automatically converts those four VMs running for 50% of the month into two VMs running for 100% of the month, so you can get the maximum Sustained Usage discount and the lowest price.

Amazon Web Services: Complex Pricing Options

In contrast, AWS offers several pricing options. It offers on-demand usage pricing that is rounded up to the nearest hour, and this is Amazon's most expensive pricing. Another option is called a Reserved Instance, where you make a commitment for either one year or three years for a specific VM instance. These come with payment options: to pay in full up front, to pay a portion of the cost up front, or to pay nothing up front. AWS' lowest prices come from paying all of the cost up front. Any discounts are tied to up-front payments, which can create a tradeoff in agility and flexibility—but more on that later.

Figure 2 provides a simplified look at AWS pricing options for a single VM over three years. It should be immediately apparent that it is complex to parse, just to figure out what you would pay. Based on what you want, you must weigh the time you expect to use the resources with the contract length and level of up-front payment. In addition, this does not take into account your future needs, potential price reductions, or the cost of capital (we'll get to those, too).

FIGURE 2. Simplified AWS Pricing Options for Single Instance over Three Years

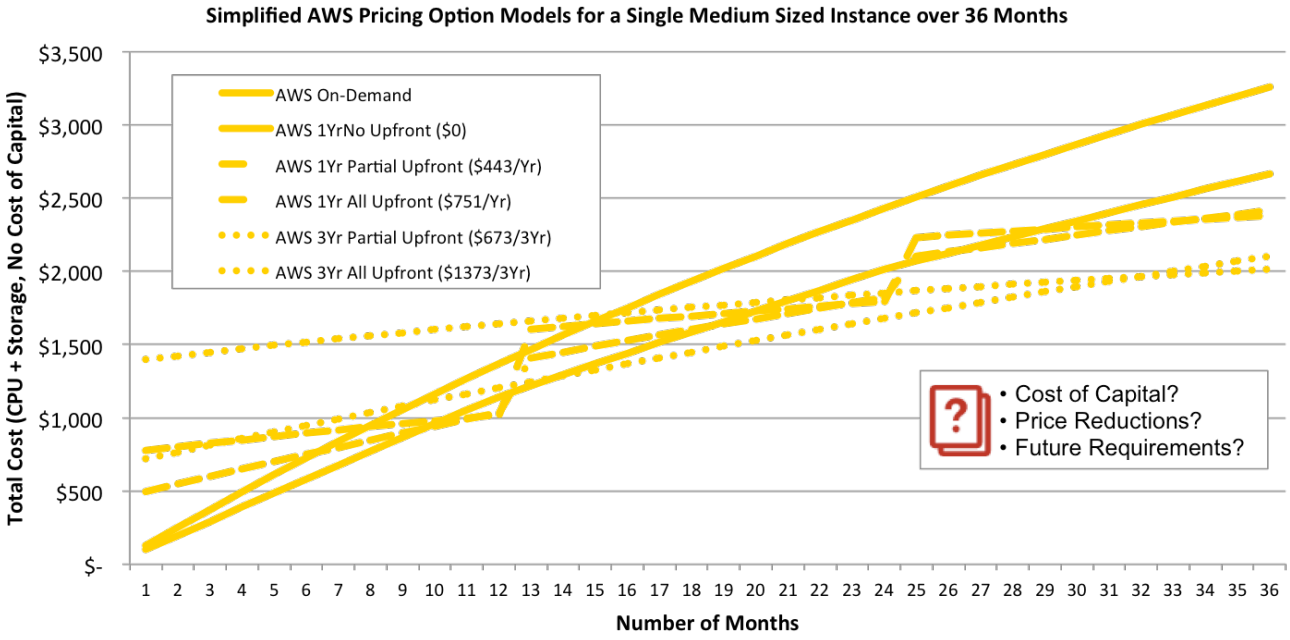

Source: Enterprise Strategy Group, 2015.

Moore's Law

The opportunity for on-demand, real-time pricing ensures that customers will gain from industry realities, and Moore's Law is an important one. The essence of Moore's Law—and evidence over the past 40 years—is that innovations in semiconductor technologies result in a doubling of density every two years. This results in computing power price reductions of 25% annually, driven by reductions in data center floor space, power, and cooling. Where that fits in your cost calculation is this: On-demand, real-time pricing will be able to reflect those reductions, while predetermined, locked-in prices paid up front will not.

Simple Cost Comparison: GCP On-demand vs. AWS On-demand

With these concepts in mind, ESG Lab compared the cumulative expected costs over a three-year period for an equivalently configured, single instance cloud server using both GCP and AWS on-demand pricing. The Google and Amazon online pricing calculators were used to generate the cost of a single, medium-sized instance comprised of two vCPUs, 7.5GB of memory, 32GB of SSD storage, 128GB of magnetic storage, and 128GB of snapshot storage space. A 2% monthly price reduction was included in the model based on an expected 25% annual cost reduction, as historically seen by cloud providers. In reality, cost reductions are unpredictable and would happen only a few times per year, but for simplicity of modeling we averaged the expected annual reduction across the entire year. As shown in Figure 3, the three-year cost of operating the on-demand GCP instance at 100% utilization (24 x 7) was 49% lower than the cost of operating the AWS on-demand instance. ESG Lab also validated that the GCP on-demand pricing advantage was relatively consistent across other comparably sized instances as well as between the number of cores per instance (see Appendix, Table 3).

FIGURE 3. Cost Comparison: Single Instance Deployment

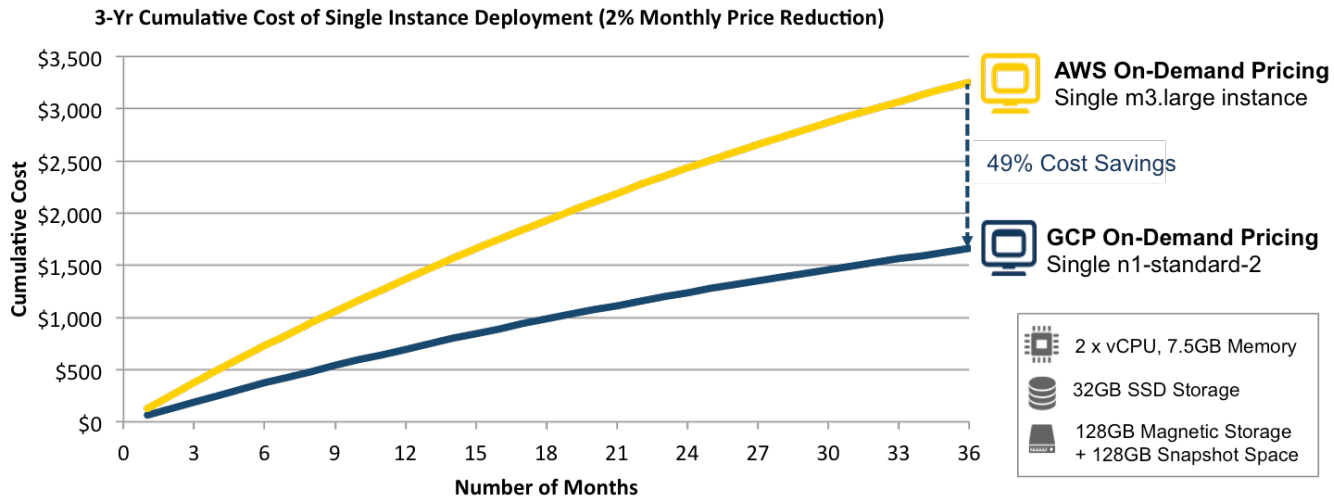

Source: Enterprise Strategy Group, 2015.

Why This Matters

So, why does this matter to you? Google offers a very simple, on-demand, real-time pricing structure, and works to automatically save you money. Customers know that Google will automatically discount prices based on sustained usage and will pass on any price reductions. AWS pricing is complex, and ESG Lab analysis comparing GCP and AWS on-demand pricing demonstrated a 49% cost advantage for equivalent VMs. Google delivers lower costs and less pricing complexity.

Agility

Another reason cloud computing is advantageous is that it enables organizations to be more agile, more responsive to changes in conditions as they arise. Organizations that can respond quickly—to competitive threats and opportunities, internal company vagaries, industry-wide changes such as innovations or supply problems, and the like—have a better chance at success.

Cloud computing enables organizations to grow and shrink on-demand to meet business needs. Sometimes that means adding/removing VM instances, turning resources up in one location or down in another, or efficiently handling peaks and troughs. Many organizations need additional computing power at fiscal milestones such as quarter- and year-ends; development/test organizations may have spikes during peak hours when they need more VMs to manage a new product feature; and sales organizations may need more resources to respond to a successful marketing campaign. Being able to move resources around is also important. If you have cloud instances in London and San Francisco and one location takes off while the other stalls, the ability to move your compute resources to where you want maximum performance enables business agility. These spikes could certainly be handled by adding more resources, but that's not efficient, just more expensive. If you can shift resources around, you are both agile *and* cost-efficient.

Google Cloud Platform's on-demand, real-time prices are in sync with this type of agility. There are no time commitments; you pay more when you use more, and pay less when you don't. In addition, your Google instances are not restricted by region—you can move them as needed. You can increase or decrease your resources and continue to pay only for what you're using.

With AWS's pricing model, agility costs you money. If you want to scale up and down, move resources, and pay on demand, you're going to pay the highest price. For any pricing discount, you must select Reserved Instances—these are specifically-sized instances that you contract for a period of time.

- In terms of pricing, you are locked in for one to three years. You can save some money by paying more up front, but that has its own opportunity cost.
- Reserved Instances cannot be resized, so if you need more compute power, you would need to add instances. This is where it gets to be like your smartphone plan: You start with a two-year contract, and after a year, add your child's phone to it (with another two-year contract); now you have overlapping contracts that keep you locked in over time.
- Reserved Instances are restricted regionally as well, so you are unable the move resources if you need to—you would have to add them.
- Similarly, you cannot shrink your resources because they are committed, so if you need fewer VMs, you still pay for what you contracted. Amazon does have an option where you can try to sell your instances in an online marketplace, usually at a discount. This burden is entirely on you, and Amazon charges a 12% fee. So you may or may not recoup your money.

With AWS, the best price comes with the three-year commitment where you pay it all up front. But isn't a long-term commitment the antithesis of agility?

Cost of Capital

Another important point about up-front payments and long-term commitments that many companies neglect to consider is the cost of capital. Companies have to choose what they spend their cash on, so there is a cost of capital for each company individually. It varies based on how companies are financed and it changes over time, but in this paper, we use a typical average of 8%. So if you want AWS' lowest price, assume it will cost you an additional 8% on whatever you spend. Anything you purchase up front comes with the risk of missing future price reductions, as well as the opportunity cost of other ways the company could spend that money. When it comes to cloud computing services, paying up front essentially takes flexible OpEx spending and transforms it into financially inflexible CapEx.

Agility Comparison: Startup Application, Production Environment

ESG Lab validated the three-year cumulative costs that could be expected for a typical startup application production environment. We modeled deployment of compute instances with GCP on-demand pricing versus an equivalently configured solution using AWS three-year, all-up-front Reserved Instance pricing (the lowest-priced three-year option for AWS). The startup environment consisted of eight medium-sized instances consisting of a single vCPU, 3.75GB of memory, and 4GB of SSD (n1-standard-1 for GCP and m3.medium for AWS). Reserved Instance pricing was amortized over time at an estimated 8% cost of capital.

As Figure 4 shows, AWS on-demand pricing remains the most expensive. At 100% utilization (24 x 7 operation for all eight servers), AWS three-year, all-up-front Reserved Instances (amortized) appear to be 16% less expensive than GCP. However, this data point represents an unrealistic "best case" scenario for AWS pricing. In reality, GCP on-demand pricing would take advantage of reductions in price over the three-year period. In addition, the startup production environment was sized for "peak" production times (such as the 12 busiest hours of the day), whereas in reality, customers can reduce instances when they are not required, saving money during these "troughs."

When ESG Lab factored in an expected 25% annual price reduction and typical case 50% utilization, GCP on-demand pricing produced a three-year expected cost that was 44% lower than the best case AWS Reserved Instance pricing. The blue dotted line in Figure 4 represents the benefits of agility delivered by Google's pricing model (see Figure 4).

FIGURE 4. Cost Comparison: Startup Application, Production Environment

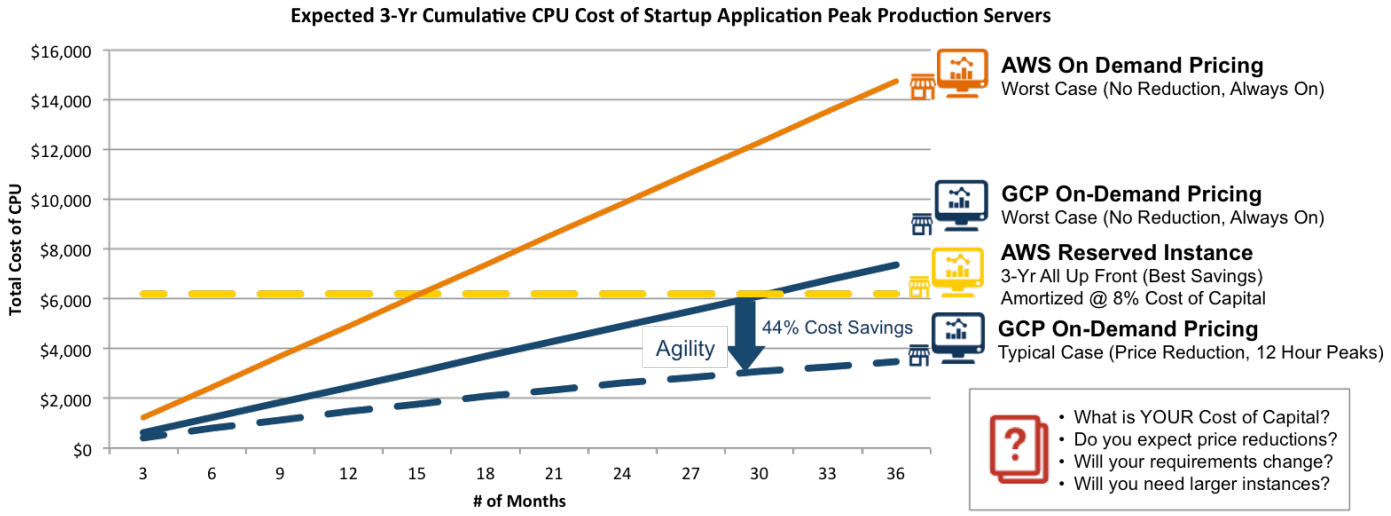

Source: Enterprise Strategy Group, 2015.

Why This Matters

Organizations go to the cloud to be more agile and responsive to business needs, as well as to reduce costs. So if your costs are locked in and/or paid up front, then the agility you need (to scale up and down, and move workloads to serve your business needs) is only attainable at a steep price. Google's on-demand, real-time pricing lets you be responsive without adding cost, while AWS' multiyear lock-in pricing negatively impacts agility and the ability to take advantage of price drops. ESG Lab analysis shows that GCP's three-year realistic cost for a startup application was 44% lower than AWS' three-year, all-up-front Reserved Instance pricing.

Flexibility

Organizations also go to the cloud for flexibility, which is really about choice: for example, choosing to offload workloads to the cloud for operational benefit, to select locations that will serve customers best, or to optimize a particular metric in a specific circumstance, without wasting resources.

Google's on-demand, real-time prices are flexible without wasting resources. In addition, Google's Sustained Usage discount is delivered to you automatically so your flexibility may result in a discount. If you underestimated the computing resources that you would need a year from now, you can add VMs and simply pay the price at that time, gaining whatever cost advantages are available then. You can also choose your instance size and region, and re-choose without penalty at any time. With AWS' pricing model, if you want flexibility, you pay a steep price because on-demand prices are its most expensive. If you want better prices, long-term lock-in and up-front payments are your only option, and you cannot resize instances or move them to other locations.

Another important consideration is utilization. When you pay up front (where AWS offers its best prices), you are paying for 100% utilization. You get no benefit from shutting down VMs. So if your compute resources aren't running 24 x 7, you're not really getting what you paid for. This is like paying your home electric bill up front every month, as if you had all the lights on all the time—you would get no benefit from turning off the lights.

Think of a typical development/test scenario. Developers want to turn on numerous VM instances, run batch jobs for a couple of minutes or hours, and shut the VMs down again—and they might do this several times a day. If you paid up front, your costs don't go down even though your resource usage does. Google's on-demand pricing starts with ten minutes and rounds to the nearest minute, so you only pay for the actual time you are using the resource. In contrast, a 15-minute batch job with AWS on-demand pricing is billed for four times as many minutes. (It should be noted that Reserved Instance pricing for the batch processing portions of the ESG Lab model was so high that our models did not consider it as an option in our analysis, opting instead for on-demand instances for this portion.)

Flexibility Comparison: Utilization of Mature Application, Production Environment

To better understand the Google Sustained Usage discount, ESG Lab again modeled the expected three-year cost, this time of a mature production environment consisting of peak workloads that require 200 medium-sized instances (n1-standard-1 for GCP and m3.medium for AWS) across a range of expected utilizations. As shown in Figure 5, GCP Sustained Usage provides deeper discounts at higher utilizations, providing up to 50% lower three-year costs when compared with AWS on-demand pricing. When compared with the best case AWS price (for three-year, all-up-front Reserved Instances), GCP on-demand instances provided the flexibility to only pay for what is used during lower instance utilization, providing a three-year cost that was up to 83% lower than AWS Reserved Instance pricing (see left side of Figure 5).

It should also be noted that the simplified model chosen to illustrate the concept did not take into consideration the flexibility to take advantage of price reductions or the chance to resize instances, both of which would be expected to lower the overall cost of GCP on-demand pricing well below AWS Reserved Instance pricing.

FIGURE 5. Cost Comparison: Utilization of Mature Application, Production Environment

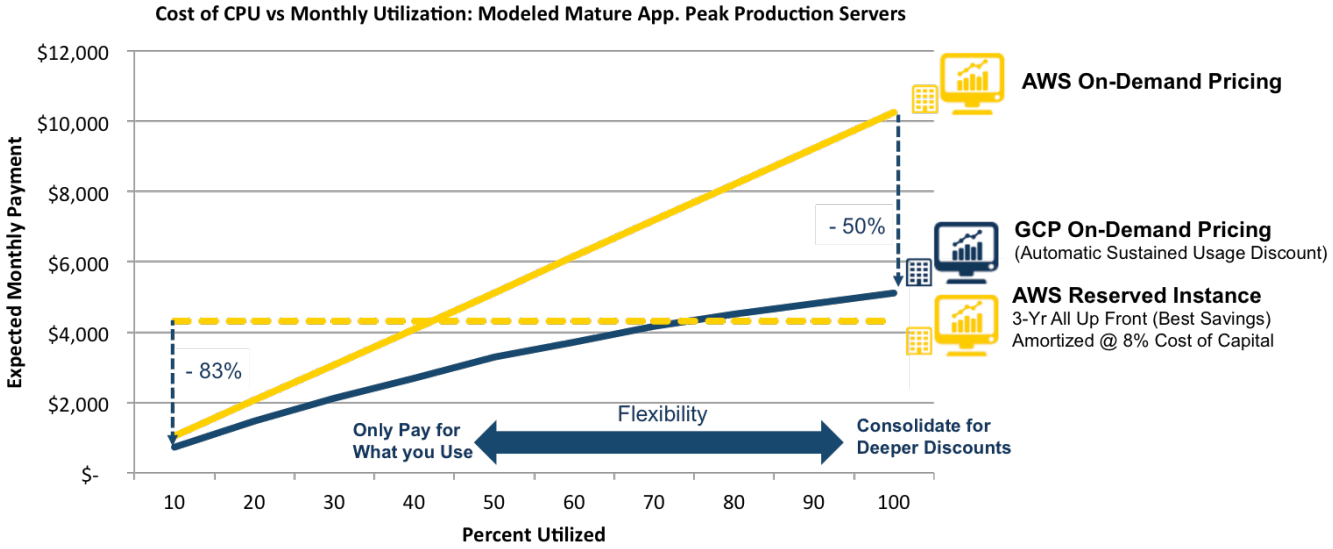

Source: Enterprise Strategy Group, 2015.

Why This Matters

With AWS, having to lock in VM instances and pricing for one to three years eliminates the flexibility to move instances, resize them, and change your mind. It also results in paying for 100% utilization of your resources even though you're not using them 100% of the time. With Google, changing your mind doesn't cost you a penalty, and the more you use, the greater your discount. ESG Lab validated that for a mature application in production, GCP's realistic, cumulative three-year cost was 35%-53% lower than the various AWS pricing options.[2]

[2] Note: Detailed analysis resulting in the reported range is not shown in the illustrative example in Figure 5. Please refer to Table 1.

The Bigger Truth

When you go to the cloud for computing resources, it's important to remember *why* you are going there: to get lower costs, greater business agility, and more flexibility. Virtualizing workloads and consolidating them on fewer servers eliminates the need for dedicated hardware for every application, saving on both infrastructure and management costs. Leveraging cloud provider infrastructure-as-a-service extends virtualization benefits by freeing up organizations from having to deal with infrastructure at all, and enabling them to pay only for the units of infrastructure time they need.

With this in mind, what sense does it make to buy infrastructure-as-a-service, and have its pricing scheme restrict your cost benefits, your business agility, and your flexibility? Does it make sense to purchase infrastructure-as-a-service, but be committed to specific instances over time that cannot be resized, cannot be moved, and are locked into prices and up-front payments? With those restrictions, your cloud deployment basically eliminates all the reasons you went to the cloud in the first place.

Comparing pricing schemes over time can be a complex and difficult task—so difficult, in fact, that Google has developed a head-to-head TCO calculator to help you out.[3] After validating key pricing concepts, ESG Lab used its own three-year pricing models to validate the costs and associated savings output of the Google TCO calculator for a given set of inputs and for several pricing scenarios. Although the models were created by different teams using separate sets of assumptions and methodologies to calculate the end results, the models arrived at very similar pricing advantages for each scenario (within 15% in nearly every case).

ESG Lab's comprehensive, three-year pricing analyses and modeling validate that Google has a price advantage any way you slice it, after looking at various scenarios of application characteristics and AWS pricing options (see Table 1). Google Cloud Platform's on-demand, real-time pricing structure matches the cloud paradigm, delivering the agility and flexibility customers seek, automatically, without penalty.

TABLE 1. Results of ESG Lab Three-year Comprehensive Cost Analyses: Google Pricing Advantage vs. AWS Options

Model Scenario	GCP vs. AWS On-demand	GCP vs. AWS 1-Yr Partial Up-front / All Up-front	GCP vs. AWS 3-Yr Partial Up-front / All Up-front
Single Instance	-49%	-33% /-32%	-24% / -24%
Startup Application	-49%	-50% / -50%	-60% / -88%
Mature Application	-42%	-37% / -47%	-35% / -53%
Static Enterprise Workload	-50%	-28% / -27%	-16% / -15%

Source: Enterprise Strategy Group, 2015.

It is important to note that ESG Lab's analyses were focused strictly on price. We have not compared the two solutions in terms of features, partnerships, or points of presence, where AWS would likely have an advantage. We expect that Google will catch up on those fronts, but it's not there yet. Nor have we had the chance to include price/performance, another metric that would be useful to customers. But what really matters is how any cloud computing solution works in your environment, so ESG recommends that you evaluate all aspects of any cloud solution before making a decision.

Why does Google insist on being cheaper? Because it's the new kid on the block, and its pricing advantage should result in gaining market share. That's competition at its best, working for the customer. So if you are looking at cloud computing, do your homework about features, benefits, and costs; remember that these are not discrete silos, but actually work together to deliver a complete picture of value. And if you are looking for cost-effective, agile, and flexible cloud IaaS, ESG Lab suggests that you make sure the pricing model helps deliver on these values.

[3] Amazon also has a calculator that compares to your current infrastructure.

Appendix

TABLE 2. Price Configurations Used in ESG Lab's Three-year Pricing Analyses

Single VM	Static Enterprise Application
GCP: 1 x n1-standard-2 (2 CPU Core, 7.5GB Mem. Per Instance) 32GB SSD Storage, 128GB Persistent Storage 128GB Snapshot Space	**GCP:** 256 x n1-standard-1 (1 CPU Core, 3.7GB Mem. Per Instance) All Instances: 24 x 7 Operation (100% Utilization)
AWS: m3.large 2 CPU Cores, 7.5GB Mem, 32GB SSD 128GB Persistent Storage (50IOPS) 128GB Snapshot Space	**AWS:** 256 x m3.medium 1 CPU Core, 3.7GB Mem, 4GB SSD All Instances: 24 x 7 Operation (100% Utilization)
Small Startup Application	
GCP (Development): 40 x n1-standard-1 (1 CPU Core, 3.7GB Mem. Per Instance) Each Instance: 15 minutes daily (1% Utilization)	**AWS (Development):** 40 x m3.medium 1 CPU Core, 3.7GB Mem, 4GB SSD All Instances: 1 Hr Daily (1Hr/Min) (4.2% Utilization) (On-demand Billing For All Pricing Cases)
GCP (Production Peaks): 8 x n1-standard-1 (1 CPU Core, 3.7GB Mem. Per Instance) Each Instance: 12 Hours Daily (50% Utilization)	**AWS (Development):** 8 x m3.medium 1 CPU Core, 3.7GB Mem, 4GB SSD Each Instance: 12 Hours Daily (50% Utilization)
GCP (Production Trough): 1 x n1-standard-1 (1 CPU Core, 3.7GB Mem. Per Instance) 24 x 7 Operation (100% Utilization)	**AWS (Development):** 1 x m3.medium 1 CPU Core, 3.7GB Mem, 4GB SSD 24 x 7 Operation (100% Utilization)
Mature Application	
GCP (Development): 80 x n1-standard-1 (1 CPU Core, 3.7GB Mem. Per Instance) Each Instance: 90 minutes daily (6.25% Utilization)	**AWS (Development):** 80 x m3.medium 1 CPU Core, 3.7GB Mem, 4GB SSD All Instances: 2 Hr Daily (1Hr/Min) (8.3% Utilization) (On-demand Billing For All Pricing Cases)
GCP (Production Peaks): 200 x n1-standard-1 (1 CPU Core, 3.7GB Mem. Per Instance) Each Instance: 12 Hours Daily (50% Utilization)	**AWS (Development):** 200 x m3.medium 1 CPU Core, 3.7GB Mem, 4GB SSD Each Instance: 12 Hours Daily (50% Utilization)
GCP (Production Trough): 33 x n1-standard-1 (1 CPU Core, 3.7GB Mem. Per Instance) 24 x 7 Operation (100% Utilization)	**AWS (Development):** 33 x m3.medium 1 CPU Core, 3.7GB Mem, 4GB SSD 24 x 7 Operation (100% Utilization)

Source: Enterprise Strategy Group, 2015.

Key ESG Lab Three-year Pricing Model Assumptions:

- Yearly price drop due to Moore's Law = 25%
  - 2% price drop factored in to all on-demand pricing (AWS and GCP)
  - 25% yearly price drop factored in every 12 months for 1-year RI pricing (all and partial up-front)
  - No price drop factored in for 3-year Reserved Instances
- Annual cost of capital = 8%
  - 1-year Reserved Instance up-front costs are financed at 8% cost of capital for 12-month period
  - 3-year Reserved Instance up-front costs are financed at 8% cost of capital for 36-month period
  - Pricing was also analyzed as flat payment with no cost of capital for some illustrative purposes (graphs)

- Cumulative costs were calculated on a monthly basis
- Total cost was analyzed over 1st year, 2nd year, 3rd year, and 3-year total
  - 3-year totals are reported in Table 1
  - 1st year totals were compared with Google's TCO tool results
- ESG analyses shown in the paper are based on the most comparable CPU offerings between GCP and AWS
  - GCP standard instances = 2.6GHz Intel Xeon E5 Sandy Bridge / 2.5Ghz Intel Xeon E5 v2 (Ivy Bridge)
  - AWS m3 instances = 2.6GHz Intel Xeon E5 Sandy Bridge / 2.6Ghz Intel Xeon E5 v2 (Ivy Bridge)

TABLE 3. ESG Lab On-demand Pricing Analysis of Instance Size and Core Count

						Price / Hour (100% Util)	
	GCP	AWS	GCP Memory (AWS)	# of Cores	GCP On Demand Pricing Advantage	GCP Price	AWS Price
Standard Server	n1-standard-1	m3.medium	3.75	1	**50%**	$0.04	$0.07
	n1-standard-2	m3.large	7.5	2	**50%**	$0.07	$0.14
	n1-standard-4	m3.xlarge	15	4	**50%**	$0.14	$0.28
	n1-standard-8	m3.2xlarge	30	8	**50%**	$0.28	$0.56
High Memory	n1-highmem-2	m2.xlarge	13 (17)	2	**64%**	$0.09	$0.25
	n1-highmem-4	m2.2xlarge	26 (34.2)	4	**64%**	$0.18	$0.49
	n1-highmem-8	m2.4xlarge	52 (68.4)	8	**64%**	$0.35	$0.98
High CPU (AWS w/ SSD)	n1-highcpu-2	c3.large	1.8 (3.75)	2	**50%**	$0.05	$0.11
	n1-highcpu-4	c3.xlarge	3.6 (7.5)	4	**50%**	$0.11	$0.21
	n1-highcpu-8	c3.2xlarge	7.2 (15)	8	**50%**	$0.21	$0.42
	n1-highcpu-16	c3.4xlarge	14.4 (30)	16	**50%**	$0.42	$0.84
High CPU (No SSD)	n1-highcpu-2	c4.large	1.8 (3.75)	2	**54%**	$0.05	$0.12
	n1-highcpu-4	c4.xlarge	3.6 (7.5)	4	**54%**	$0.11	$0.23
	n1-highcpu-8	c4.2xlarge	7.2 (15)	8	**54%**	$0.21	$0.46
	n1-highcpu-16	c4.4xlarge	14.4 (30)	16	**54%**	$0.42	$0.93

Key Takeaways: Google Pricing Advantage is Consistent Across # of Cores
Google Pricing Advantage is relatively consistent across Instance Size and Type (50% to 64% Lower)

Prices as of 6-10-15

Source: Enterprise Strategy Group, 2015.

ESG Lab's analyses included storage pricing for the single instance case only, as the amount and type of storage can vary drastically between organizations and applications. ESG Lab performed a cursory analysis of various storage offerings and pricing and the preliminary results appear to give GCP a large advantage in both price and performance (see Table 4).

TABLE 4. ESG Lab Quick Comparison of Storage Pricing Models

		Monthly Cost For Storage						
		GCP	AWS	GCP Savings	AWS (w/ 50 IOPS)	Diff GCP/AWS	AWS (w/100 IOPS)	Diff GCP/AWS
Persistent Disk	1GB	$0.04	$0.05	20%	$6.64	99%	$13.23	100%
	10GB	$0.40	$0.50	20%	$7.09	94%	$13.68	97%
	100GB	$4.00	$5.00	20%	$11.59	65%	$18.18	78%
	1000GB	$40.00	$50.00	20%	$56.59	29%	$63.18	37%
Snapshot Space	1GB	$0.03	$0.10	70%	N/A		N/A	
	10GB	$0.26	$0.96	73%	N/A		N/A	
	100GB	$2.60	$9.50	73%	N/A		N/A	
	1000GB	$26.00	$95.00	73%	N/A		N/A	
Persistent SSD (General Purpose)	1GB	$0.17	$0.10	-70%	AWS Limited to 3 IOPS with 1GB, GCP to 30 IOPS			
	10GB	$1.70	$1.00	-70%	AWS Limited to 30 IO/Sec with 1GB, GCP to 300 IOPS			
	100GB	$17.00	$10.00	-70%	AWS Limited to 300 IO/Sec with 1GB, GCP to 3,000			
	1000GB	$170.00	$100.00	-70%	AWS Limited to 3000 IO/Sec with 1GB, GCP to 10-15K IOPS			
					No Charge for above IOPS for either vendor			
Local SSD (@ 375GB)	0 IOPS	$82.12	$46.88	-75%	GCP: Local SSD available in SCSI or Higher Performing NVMe			
	10K IOPS	$82.12	$696.88	88%	AWS: Offers finer granularity of volume size (GCP must be multiple of 375GB)			
	20K IOPS	$82.12	$1,346.88	94%	GCP Price includes IOPS (70K to 100K for SCSI, 90K to 170K for NVMe)			
	> 20K IOPS	$82.12	N/A		Price config shown normalized to 375GB and AWS Max of 50K IOPS			
					GCP Price calculated pricing with instance then removing cost of instance alone			
Note AWS Charges per IO (Per volume, per month), GCP does not; Prices effective as of June 12, 2015								

Source: Enterprise Strategy Group, 2015.

ESG Lab's analysis did not include any networking costs, as network traffic can vary significantly between organizations and applications. A quick comparison of networking costs shows that both GCP and AWS can demonstrate a pricing advantage depending on mix of traffic, infrastructure location, and IP configuration.

TABLE 5. ESG Lab Quick Comparison of Network Pricing Models

			Rate in Cents/GB		
	From	To	GCP	AWS	Difference
Data Ingress	Internet	Instance	0.00	0.00	0.00
(Data In to Instance)	Another GCP/AWS Region	Instance	0.00	0.00	0.00
(ex - Download, DB Update In)	AWS/GCP Product	Instance	0.00	0.00	0.00
	Anywhere	Instance (Private IP)	0.00	0.00	0.00
	Anywhere	Instance (Public/Elasitc IP)	0.00	0.01	-0.01
	AWS Product in other Zone	Instance	0.00	0.01	-0.01
Data Egress	Instance	Google or AWS Product (Repective)	0.00	0.00	0.00
(Data Out from Instance)	Instance	Google or AWS Service in Same Region	0.00	0.00	0.00
(ex - Upload, DB Read Out)	Instance (Private IP)	Google or AWS Service in Same Zone	0.00	0.00	0.00
	Instance (Public/Elasitc IP)	Google or AWS Service in Same Zone	0.00	0.01	-0.01
	Instance (Private IP)	Same Zone	0.00	0.00	0.00
	Instance (Public/Elasitc IP)	Same Zone	0.00	0.01	-0.01
	Instance (Private IP)	Zone in same Region	0.01	0.00	0.01
	Instance (Public/Elastic IP)	Zone in same Region	0.01	0.01	0.00
	Instance	Another US region*	0.01	0.02	-0.01
	Instance	Worldwide Other	0.8-0.12	.05-.09	0.03
	Instance	Australia	0.15-0.19	.05-.09	0.10
	Instance	China	0.20-0.23	.05-.09	0.14 to 0.15
Notes:	AWS Has 10 Regions, GCP Has 3				
	Google Egress Pricing automatically discounted atfter 1TB, 10TB of data. AWS Tiers at 10/40/100/350 and higher.				
	Google Products = YouTube, Google Drive, Maps etc.., AWS Products = S3, Glacier, DynamoDB etc...				
	* = Promotional pricing rate subject to change				
	First 1GB Egress Free on AWS				

Source: Enterprise Strategy Group, 2015.

Enterprise Strategy Group is an integrated IT research, analysis, and strategy firm that is world renowned for providing actionable insight and intelligence to the global IT community.

 www.esg-global.com

 contact@esg-global.com

 P. 508.482.0188