# 2018 TOP HYPERSCALE CLOUD PROVIDERS
# AMAZON - AZURE - GOOGLE

**Price-Performance Analysis of North American Hyperscale IaaS Providers**

CLOUD SPECTATOR

# TABLE OF CONTENTS

# EXECUTIVE SUMMARY

This report examines the results of a study measuring and comparing the performance and price-performance value of three hyperscale Cloud Service Providers (CSPs) within the North American region. While the CSPs included in the study did not have to be headquartered in North America, they must have at least one data center located within the North American continent (see Methodology page 11).

The CSPs included in this report are Amazon AWS, Google Compute Engine and Microsoft Azure.

The performance results are separated into two categories: VM Performance and Block Storage Performance. VM Performance tests the CPU and memory of the virtual machine. This performance data is aggregated into one score that includes both CPU and memory. Block storage is evaluated using two different tests as detailed in the methodology section (page 11).

# PRICE-PERFORMANCE KEY FINDINGS

*The chart below displays the overall price-performance scores of the providers included in this report. Price and performance of the VM and storage are incorporated into the values.*

Chart 2.1: Overall CloudSpecs Ranking

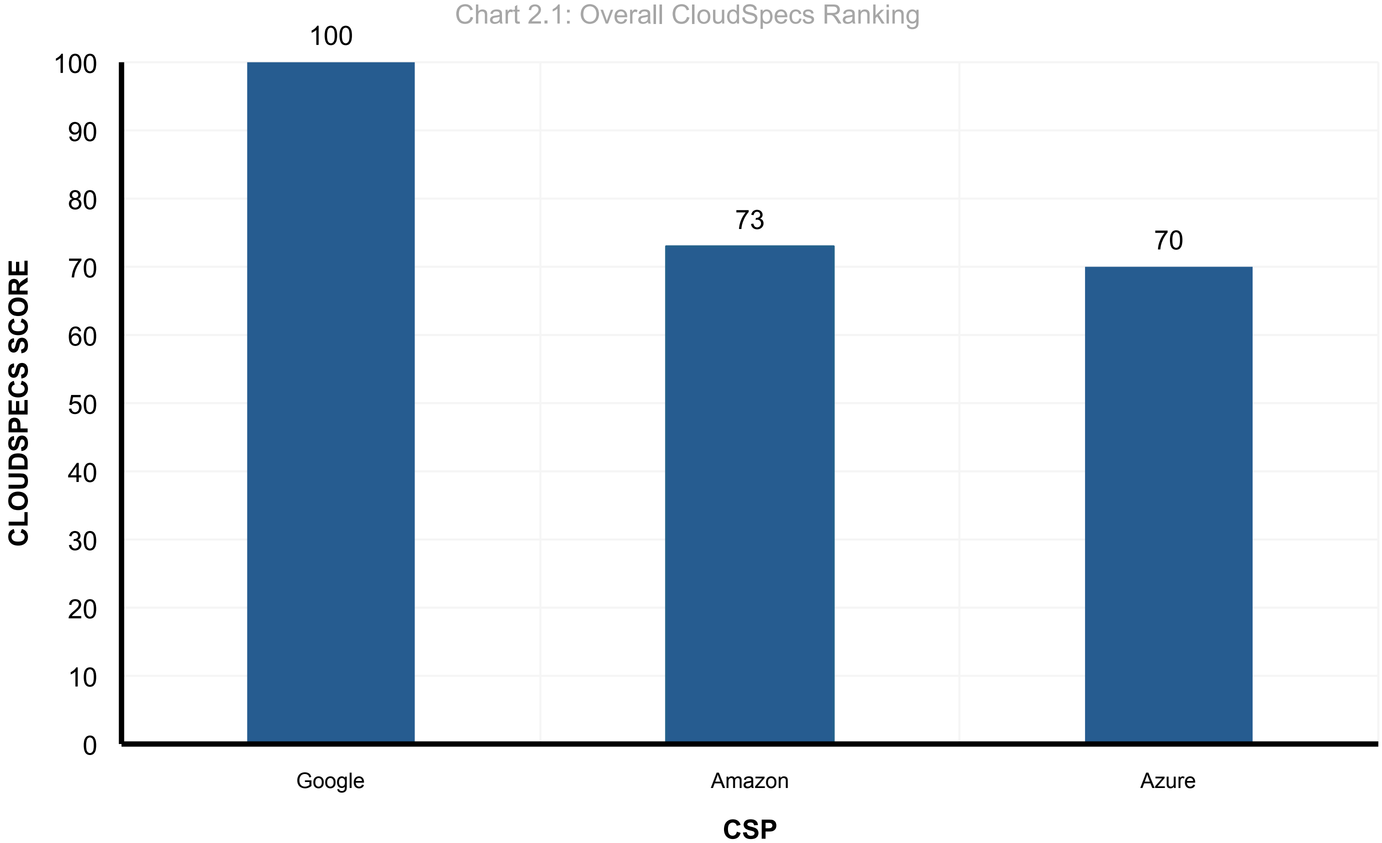


- Value, defined as the ratio of price and performance (see Methodology page 16) varies by 1.4x across the compared IaaS providers.
- Google achieves the highest CloudSpecs Score™ in the Hyperscale cloud IaaS providers ranking. This is due to strong storage performance and the most inexpensive packaged pricing found in the study.

# VM PERFORMANCE KEY FINDINGS

The chart below displays the median performance and performance variability captured across the VM testing on the providers included in the study.

Chart 2.2: VM Performance and Variability Over 24 Hours

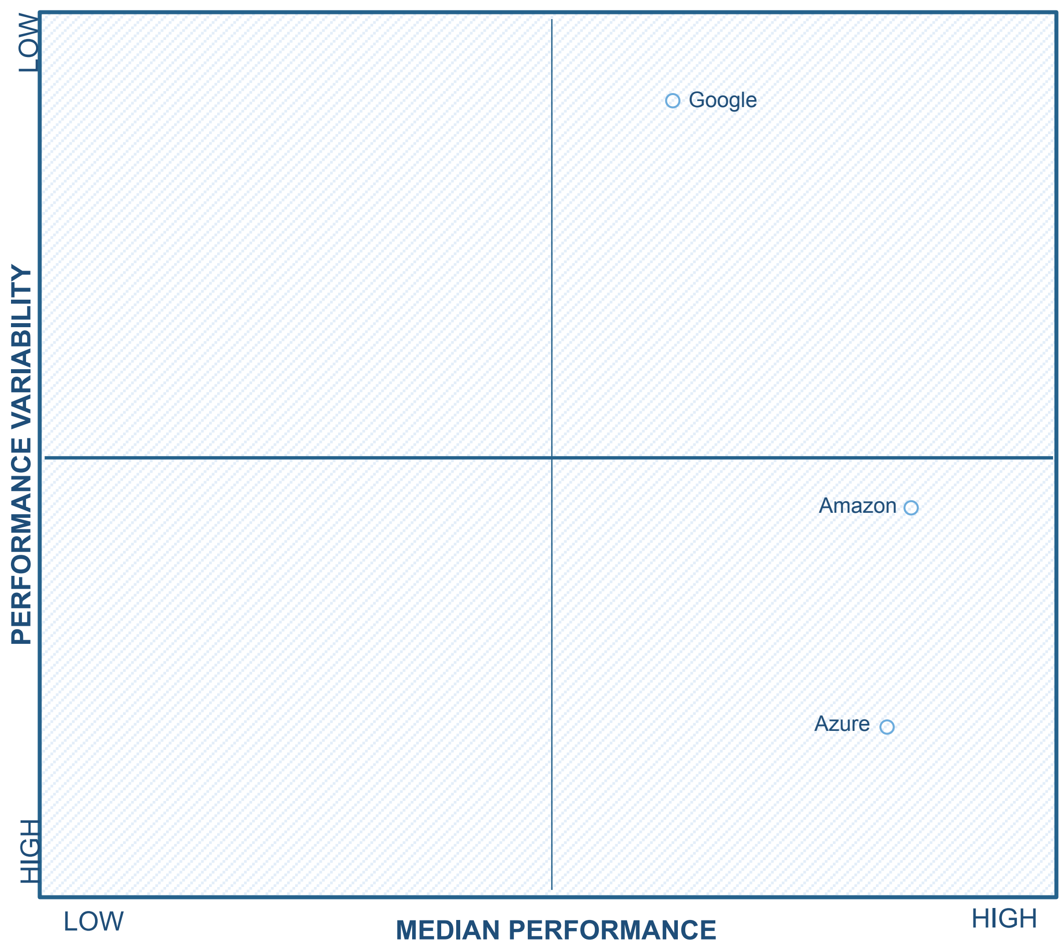


- Providers exhibited a difference of approximately 1.2x in VM (CPU & memory) performance, emphasizing the need for performance testing to understand value.
- While Amazon and Azure displayed relatively similar performance within the right-most grids, the variability is shown to be largely scattered between the three providers.
- Google Compute Engine showed the least performance variability in the 24-hour testing period.

# BLOCK STORAGE PERFORMANCE KEY FINDINGS

The chart below displays the median performance and performance variability captured across the Type 1 & 2 storage testing on the providers included in the study.

Chart 2.3: Block Storage Performance and Variability Over 24 Hours

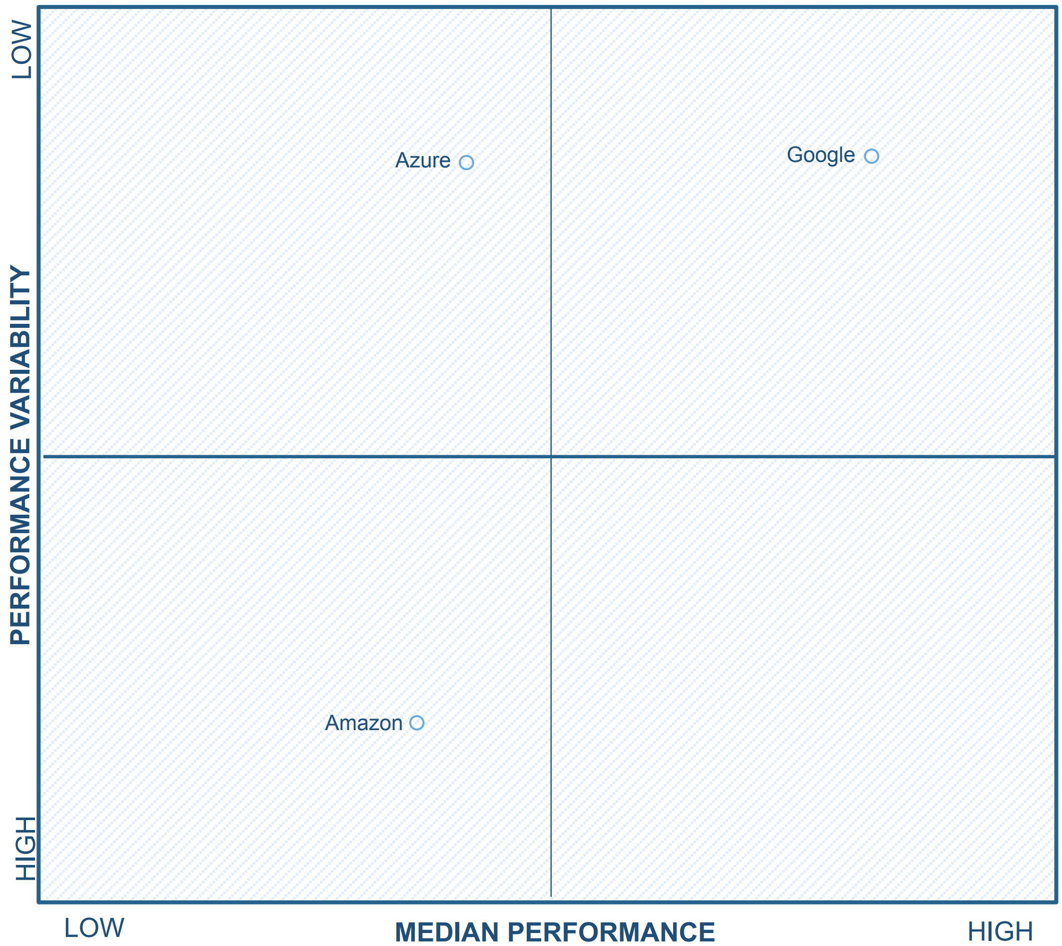


*Amazon's disk variability is artificially high due to an initial burst period that lasted until the volume ran out of I/O credits.

- Relative performance of all VM sizes on Type 1 and Type 2 varied up to 9.8x between the providers.
- Google tested highest in median disk performance across both Type 1 and Type 2 scenarios.
- Azure and Google showed the low performance variability over the testing period. Amazon AWS demonstrated controlled performance throttling on disk IOPS as well. The level of throttling on AWS EBS disks is determined by the size of the disk. The high amount of variability on AWS is due to a burst function that is built in to EBS and is not representative of an unstable environment.

Table 2.1 below lists the indexed performance scores and variability percentages by CSP. These numbers are used in generating Charts 2.2 and 2.3.

Table 2.1: Performance and Variability of CSPs Over 24 Hours

| | VM | | Block Disk | |
|---|---|---|---|---|
| | **Performance Index** | **Variability** | **Performance Index** | **Variability** |
| Amazon* | 100 | 0% | 45 | 38% |
| Azure | 98 | 1% | 50 | 0% |
| Google | 84 | 0% | 98 | 0% |

*Amazon's disk variability is artificially high due to an initial burst period that lasted until the volume ran out of I/O credits.

The Performance Index is calculated by indexing the individual performance scores achieved by each VM category (categorized as Small, Medium, Large and Extra Large; see Methodology for more information) on a scale of 0-100 with 100 as the highest possible score. An average across all VM categories is calculated to represent the Performance Index for each provider.

Variability is calculated as the average coefficient of variation (CV), which is the standard deviation expressed as a percentage of the mean performance for the VM categories of each CSP. Higher CV correlates to more fluctuation in performance (i.e., higher performance variability) over the test period.

# INTRODUCTION

Public cloud service providers (CSP) purport to offer instantaneous, scalable virtual infrastructure with utility-style pay as you go billing. Unbeknowst to many is that there is a wide variance in CSP cloud performance. While the public cloud IaaS industry offers these benefits, a lack of understanding of price and performance data can lead to businesses overspending in order to obtain the necessary performance requirements for their applications.

Cloud Spectator tested three of the largest, most well-known public cloud providers with data centers in North America. This report measures and ranks CSPs using a comprehensive performance and price-performance methodology designed by Cloud Spectator specifically for the purpose of measuring cloud services. This report examines the performance of vCPU, memory and block storage, as well as the overall value (the CloudSpecs™ Score) as defined by the relationship between price and performance.

In conjunction with an appropriate cloud vendor selection process, this report will assist purchasing decisions by providing performance and price-performance from a holistic and vendor-neutral perspective. The report is specifically designed to educate IT leadership on the variation in performance and price-performance value across public cloud providers. Performance is a critical and often overlooked component when making a cloud purchasing decision, and can have substantial impact on annual operating costs.

## WHY IS THIS INFORMATION NECESSARY?

A lack of transparency in the public cloud IaaS marketplace for cloud services performance often leads to misinformation or false assumptions. Users and potential users may be led to view cloud computing as a commodity, differentiated mostly by variety of services. In reality, cloud performance is impacted by a variety of factors from provider to provider, involving everything from the physical hardware (e.g., Intel or AMD, SSD or spinning disk), to the cost of the virtualized resources. By evaluating cloud services based on performance rather than solely on price or VM configurations, users are able to maximize value in the cloud.

## MISCONCEPTIONS ABOUT PERFORMANCE IN CLOUD

A number of common misconceptions about performance of cloud-based servers continue to exist. A few of the top misconceptions revolve around similarity of performance between CSPs, price tied to performance, and resource contention.

**1. VM performance is the same from CSP to CSP.**

While CSPs often use the same terms to label cloud resources (i.e., vCPUs, RAM or memory, and block storage), differences in the underlying hardware, architecture, and performance tuning lead to entirely different results from the same terms such as vCPUs. For example, on VM performance alone (the virtual processor and memory), the three IaaS providers in this report exhibited differences of up to 1.2x. With block storage performance, differences exceeded 9x.

**2. There is no correlation between price and performance.**

When it comes to additional services such as support, security, geographical location, and managed services on CSPs, you get what you pay for. However, with respect to performance, this study found no correlation between price and performance. The study demonstrated the best-value CSPs in this report (defined as the ratio of price and performance as ranked by the CloudSpecs Score™) offer virtualized resources at the lowest prices. Similarly-sized VMs within the three IaaS providers tested displayed a spectrum of prices with up to a 1.2x difference between the least and most expensive CSPs.

3. **Resource contention, known colloquially as the Noisy Neighbor Effect, is not a concern with most providers.**

   A public cloud environment offers multi-tenant physical hosts, which means a business may share the same physical resources with different users on the same hardware. This means that another user with-resource-greedy applications could potentially effect the performance of other VMs on the host machine. While resource contention has been addressed by many of the largest providers in an attempt to stabilize VM performance, the block storage offerings still exhibit high levels of performance fluctuations, which may be related to other activity on the same physical host as Cloud Spectator's test VMs. The fluctuation in performance evidenced in some CSPs can significantly affect hosted applications within those environments.

   **3a. If Noisy Neighbor is a concern, then performance is too unpredictable.**

   In public cloud environments, some providers, especially major ones such as Microsoft Azure and Amazon Web Services, use performance throttling to deliver a consistent user experience regardless of the actual user load on the physical machine. This means that, while performance may be lower for the VM, the user will not see much change over time. See ***Performance by VM Size*** to view the performance variation of different CSPs over the 24-hour period of the study.

# METHODOLOGY

The Cloud Spectator team designed this methodology to measure the performance of various public cloud infrastructure services. These results will provide a general insight into the public cloud industry.  However, businesses have varying needs when defining performance requirements, so they should apply testing methodologies relevant to their business and technical use cases to yield more relevant results.

## THE CRITERIA

In order to be considered and tested as one of the IaaS providers in this report, a CSP must meet the following criteria as part of its IaaS offering:

1. North American Data Center:  The CSP must have at least one data center located within the North American continent. The CSP does not need to be headquartered in North America.
2. Self-Sign-up: A user must be able to sign up for a CSP's services online, rather than reaching out to a sales representative. Contact forms that request users to message the CSP for signup are not considered self sign-up.
3. Self-Service: A user must be able to log into a portal that allows the user to provision, manage, and terminate virtual machines and other cloud-related services.
4. Hourly Billing Intervals: the CSP must provide billing by the hour or less. Some CSPs offer billing by the minute.
5. Block Storage Offerings:  Only providers with persistent block storage offerings are included in this study. Cloud Spectator measured disk performance by running performance tests on block storage.

# THE SETUP

The team set up anonymous accounts on all cloud service providers. No CSPs provided the team an account to provision virtual machines. For all VMs, Ubuntu 16.04 images were operating systems of choice. In cases when Ubuntu 16.04 was not available, version 14.04 was used, if no Ubuntu images were available, Debian was used. Virtual machines were tested according to four separate categories: Small, Medium, Large and Extra-Large. Each category contained a prerequisite allocation of VM resources, as shown in the table below:

Table 4A: VM Sizing

| SIZE | vCPU COUNT | RAM COUNT (GB) | DISK SIZE (GB) |
|---|---|---|---|
| Small | 2 | 4 | 100 |
| Medium | 4 | 8 | 150 |
| Large | 8 | 16 | 200 |
| Extra Large | 16 | 32 | 500 |

CSPs were segmented into two categories: (1) Packaged Offering CSPs and (2) Customizable Offering CSPs. Packaged Offering CSPs include providers such as Amazon AWS and Microsoft Azure, which deliver VMs based on pre-packaged sizes.  For example, a customer can purchase an instance size of c4.xlarge on AWS. Customizable Offering CSPs allow users to define custom VM sizes by setting resources such as vCPUs, RAM, and disk space. Only block storage was tested for disk because of its durability and persistence. This meant that CSPs that only offer local storage were not included in the report. A single block storage size was paired and tested with each VM size. For other CSPs such as Amazon AWS, which offers local and persistent block storage, the local storage was not measured and did not affect the performance or price-performance ranking of the CSP beyond the potential effect on pricing if local storage is included in packaged VM prices. **Please see the Appendix for a specific list showing what was tested on which providers.**

For Packaged Offering CSPs, the team selected VMs that most closely corresponded to the four categories of sizes. For Customizable Offering CSPs, the team provisioned servers designed to the exact requirements of the four categories of sizes when possible.

For each CSP, the team provisioned three copies of VMs for each size; i.e., three Small, three Medium, three Large, and three Extra-Large VMs. All VMs were provisioned and tested simultaneously for 24 hours. This means that, for each CSP, twelve various-sized VMs were running from the account on the corresponding provider for 24 hours.

Please note that some CSPs do not offer any VMs with resource allocations that would qualify for the Extra Large size.

## SIMULTANEOUS TESTING OVER TIME

Three resources were examined to compare performance: vCPU, memory, and storage. Performance tests were run in a continuous, iterative sequence according to the following order: vCPU tests and memory tests followed by block storage tests. Each complete sequence of testing comprised a single cycle, and cycles were repeated without pause for the duration of 24 hours. Different providers completed varying numbers of cycles within the 24-hour time limit, with the number of cycles completed being impacted by the performance levels of the resources tested (higher performance allows each test to be completed faster), which allows the performance data to be compared.

Testing over several iterations impacted the ranking of performance for CSPs. In an uncontrollable multi-tenant environment, VM performance can be affected by issues that arise with neighboring VMs. While these issues may be mitigated with resource planning as a responsibility of the CSP, sometimes performance levels cannot be guaranteed or sustained in the public cloud.  Therefore, measuring to examine sustainable performance is just as important on a public cloud as measuring to examine achieved performance. This is why the Cloud Spectator team chose to test over a period of 24 hours.

Three VMs of each category size were tested in parallel. Single-VM performance may not necessarily be reflective of the potential performance a CSP's VMs can achieve if the provisioned VM is faulty for any number of reasons. Measuring more than a singular VM of each size mitigates the possibility that the performance may be an unusual outlier due to a VM provisioning issue, so results are a more accurate reflection of a VM type's potential performance.

At other times, the physical host itself may experience issues, affecting all VMs residing on it. By provisioning all VMs simultaneously, Cloud Spectator may increase the possibility of measuring on multiple physical hosts with different users and resource contention issues, which would be more representative of a VM size's performance. While all of these processes are implemented to increase the accuracy of the measurements, it should be noted that these practices cannot guarantee 100% accuracy. Even by provisioning three of the same VMs of each category, the VMs still have the possibility of residing on the same physical host, depending on the provider's capacity.

## DATA COLLECTION

Testing was conducted at the end of Q4 2017 and early Q1 2018. The rankings were produced based on the CloudSpecs Score™, which is a price-performance ratio of the cost and median performance output of the VM. Each VM size category received a VM CloudSpecs Score™ and a block storage CloudSpecs Score™, which were averaged to calculate a CloudSpecs Score™ for the VM. The CSPs with the highest average CloudSpecs Scores™ across all VMs were then ranked. All three tested CSPs were ranked according to price-performance.

# TESTING USED

The following sections lists the tools and parameters used for the evaluation of the providers in this study.  Table 4B lists the testing tools used in this study.

Table 4B: Testing Tools

| TEST | TOOL | TASKS |
|---|---|---|
| vCPU Testing | Geekbench 3 | Integer and Floating Point |
| Memory | Geekbench 3 (using STREAM) | Reads and writes |
| Block Disk | Fio | Reads and writes |

**vCPU and Memory**

vCPU performance was measured with integer and floating point tasks from the Geekbench 3 benchmark suite. The Geekbench 3 benchmark suite was also used in collecting memory bandwidth data, which was used to measure the performance of the system memory (RAM).

**Storage**

Storage performance was measured using fio. Two storage scenarios were run to capture performance data: Type 1 and Type 2. In both scenarios, random read & write IOPS were recorded as the indicator of performance over a test period of 60 seconds. Type 1 used a large file size with a small block size, while Type 2 used a small file size with a large block size. The total number of files used in testing varied with the category of VM.

Table 4C: Testing Specifics

| CATEGORY | TYPE 1 | TYPE 2 |
|---|---|---|
| Block Size | 4KB | 128KB |
| File Size | 5GB | 128MB |

Table 4D: Total Files Used in Block Storage Testing

| SIZE | TYPE 1 | TYPE 2 |
|---|---|---|
| Small | 1 | 2 |
| Medium | 2 | 4 |
| Large | 4 | 8 |
| Extra Large | 8 | 16 |

In both testing scenarios, the number of parallel jobs run were set equal to the number of virtual processors in the VM. Each test scenario was run for 12 hours, for a total of 24 hours.

Table 4E: Type 1 Scenario

| SCENARIO | BLOCK SIZE | FILE SIZE |
|---|---|---|
| Type 1 | 4KB | 5GB |
| Type 2 | 128KB | 128MB |

# RANKING CALCULATION

The rankings of the three CSPs were determined by calculating the median performance of both vCPU-memory and storage with the monthly cost corresponding to each VM size for two price-performance scores per VM size (one for vCPU-memory and one for storage). The resulting ratios were normalized in relation to the highest-value provider for each resource, which receives a score of 100. Then the two price-performance scores for each VM size were averaged together to get one score per VM size. The providers were then ordered based on their value across all each VM size, and then their scores were averaged for all VM sizes to come up with a final score. The providers that sustained higher ratios across all VM categories ranked highly.

# PRICE-PERFORMANCE VALUE (THE CLOUDSPECS SCORE)

Cloud Spectator's price-performance calculation, the CloudSpecs Score™, provides information on how much performance the user receives for each unit of cost. The CloudSpecs Score™ is an indexed, comparable score ranging from 0-100 indicative of value based on a combination of cost and performance.

The calculation of the CloudSpecs Score™ is:

*price-performance_value = [VM performance score] / [VM cost]*

*best_VM_value = max{price-performance_values}*

*CloudSpecs Score™ = 100*price-performance_value / best_VM_value*

**CPU and Memory**

Cloud Spectator used the median Geekbench 3 performance scores as the [VM performance score] to calculate each machine's CPU and memory CloudSpecs Score™.

**Block Storage**

For both storage scenarios, median random r/w IOPS are used as the [VM performance score] to calculate each machine's Type 1 and Type 2 storage CloudSpecs Score™. Type 1 and Type 2 scores were averaged to calculate a single block storage CloudSpecs Score™.

**Overall**

Overall storage CloudSpecs Score™ was calculated by averaging block storage and vCPU-memory price-performance scores together so that they have equal weight for each VM size. Then, all resulting VM size scores were averaged together.

1. For block storage performance, the normalized sequential and random CloudSpecs Scores™ were averaged together.
2. The Type 1 and Type 2 CloudSpecs Scores™ were averaged together to create a single storage CloudSpecs™ score per VM size.
3. The overall CPU, memory and storage CloudSpecs Scores™ were calculated by averaging the CPU and memory CloudSpecs Score™ and overall storage CloudSpecs Score™ for each VM size.
4. All VM size scores were then averaged for each provider and normalized to get the final scores on the scale from 1 to 100.

# VARIABILITY

Variability is calculated by taking the coefficient of variation (CV) of each VM size's individual performance data points. The CVs are averaged for all VM sizes per CSP. The coefficient of variation is the standard deviation expressed as a percentage of the mean.

# CONSIDERATIONS

There are a few considerations to take into account when evaluating the providers under the methodology used in this study. The limitations of the methodology and shifting nature of the pricing and services offered by the providers introduce a number of variables that have to be taken into account.

**Limitations within the Methodology**

The IaaS industry lacks a standard methodology for evaluating CSPs. While the most effective methodology for measuring value of a CSP varies among use cases, the methodology developed by Cloud Spectator for this study was designed to capture performance statistics based on synthetic performance uniquely adopted for cloud infrastructure, which requires steps including extended testing over a period of time and running multiple VMs in parallel.

Furthermore, the synthetic testing conducted in this study is for measurement of maximum sustainable performance over a period of 24 hours, and is not representative of any specific workload. Therefore, the results are used for comparison purposes only, and cannot be applied to predict application performance. For example, on AWS, the gp2 block volumes demonstrated periods of burst that were limited due to the continuous bursting over the 24-hour period, but the gp2 would not be a recommended option for applications that demand the continuously high IOPS performance.

**VM Sizes**

The performance data in this report only applies to the tested VM and block sizes. Larger VMs may yield better results with both VM scores and block storage scores. Larger block sizes may also yield better block storage performance scores.

**Pricing Calculations and Discounts**

In this report, Cloud Spectator used monthly pricing to calculate the cost of VMs on providers. Some providers offer sustained-use discounts based on a monthly interval, while others discount for monthly commitments. Certain providers offer similar discounts on an annual basis or longer. These longer-term discounts were not factored into the analysis. Where available, monthly discounting was factored into the pricing calculations.  Therefore, for longer or shorter time commitments, the rankings may change.

**The Ranking System**

The three providers included in this report were ranked based on a calculation that considers both performance and cost of the environments. The performance results of vCPU, memory, and block storage are all included in the calculation. While some providers may exhibit high performance on vCPU, memory, and/or block storage, those CSPs may not necessarily rank highly depending on the cost of their environments as compared to their competitors.

**Additional Features and Costs**

Only the VM, block storage, and costs of those two components were examined in this study. Additional features, such as support costs (where applicable), public and private networks, traffic, and other services that may increase the overall cost of a CSP's offering, were not examined in the report. Depending on the types of use cases, the features not examined may impact the overall rankings. However, if a user can select between different base infrastructure options that have a difference in performance (e.g. SSD vs. magnetic storage), the options yielding higher performance outputs were chosen with exception to the pay-to-scale IOPS option, such as Provisioned IOPS. The amount and cost of the pay-to-scale IOPS option can affect a provider's ranking.

## DATA CENTER LOCATIONS

All VMs were provisioned in the North American data centers of each CSP. Specific locations, as described by each CSP, are listed in Table 4E.

Table 4F: Data Center Locations

| Provider | Data Center Location |
|---|---|
| Amazon | US East (N. Virginia) |
| Azure | US East |
| Google | Northern Virginia |

# PRICE-PERFORMANCE VALUE

This section examines the price-performance value (i.e., the CloudSpecs Score$^{TM}$) of the three IaaS providers, which is used in determining each CSP's ranking in this report. The CloudSpecs Score$^{TM}$ is calculated as the ratio between the price, defined as the monthly cost of the VM and block storage, and median performance of the VM and block storage. For more information on the calculation of the CloudSpecs Score$^{TM}$, please see the Methodology.

Google Compute achieves the highest CloudSpecs Score$^{TM}$ in this study (a CloudSpecs Score of 100).

# OVERALL CLOUDSPECS RANKING

The ranking of the Top three CSPs based on CloudSpecs Score™ is displayed in Chart 6A.

Value based on price-performance in this study is ranked in relation to the highest-value CSP, Google. A difference in value of over 1.4x exists between Google, the highest-ranked CSP, and Azure, the lowest-ranked CSP in the Top 3.

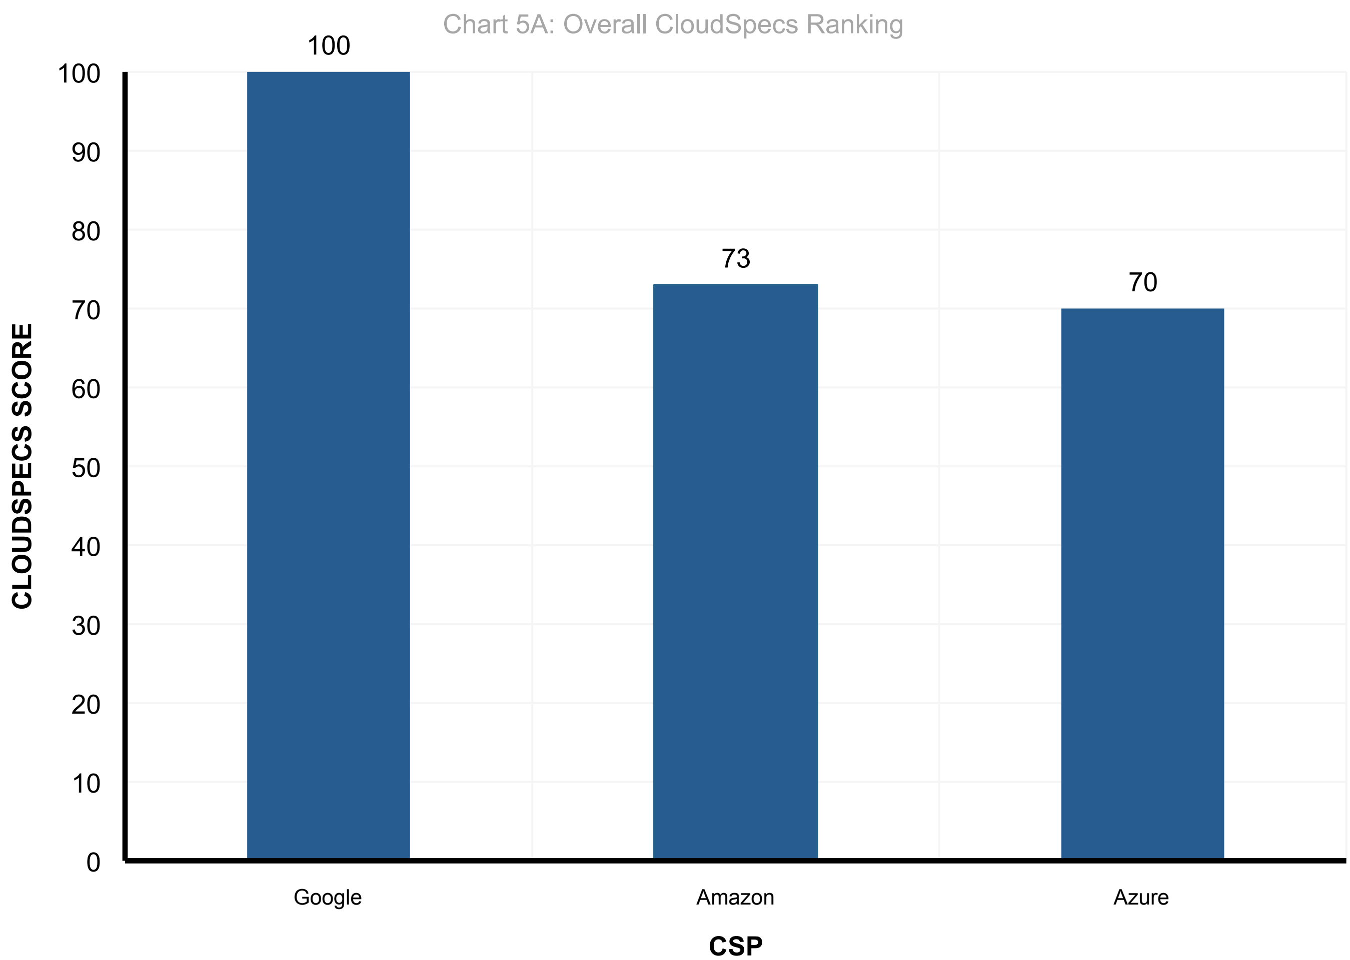


The following sections (vCPU and Memory Value and Block Storage Value) illustrate the individual Value scores segregated by section, which are the scores used to calculate the overall CloudSpecs ranking. Google achieves the second highest CloudSpecs ranking in the vCPU and Memory Value category and highest in the Block Storage Value category.

## VCPU AND MEMORY VALUE

*The chart below displays the overall price-performance values of the providers' VMs*

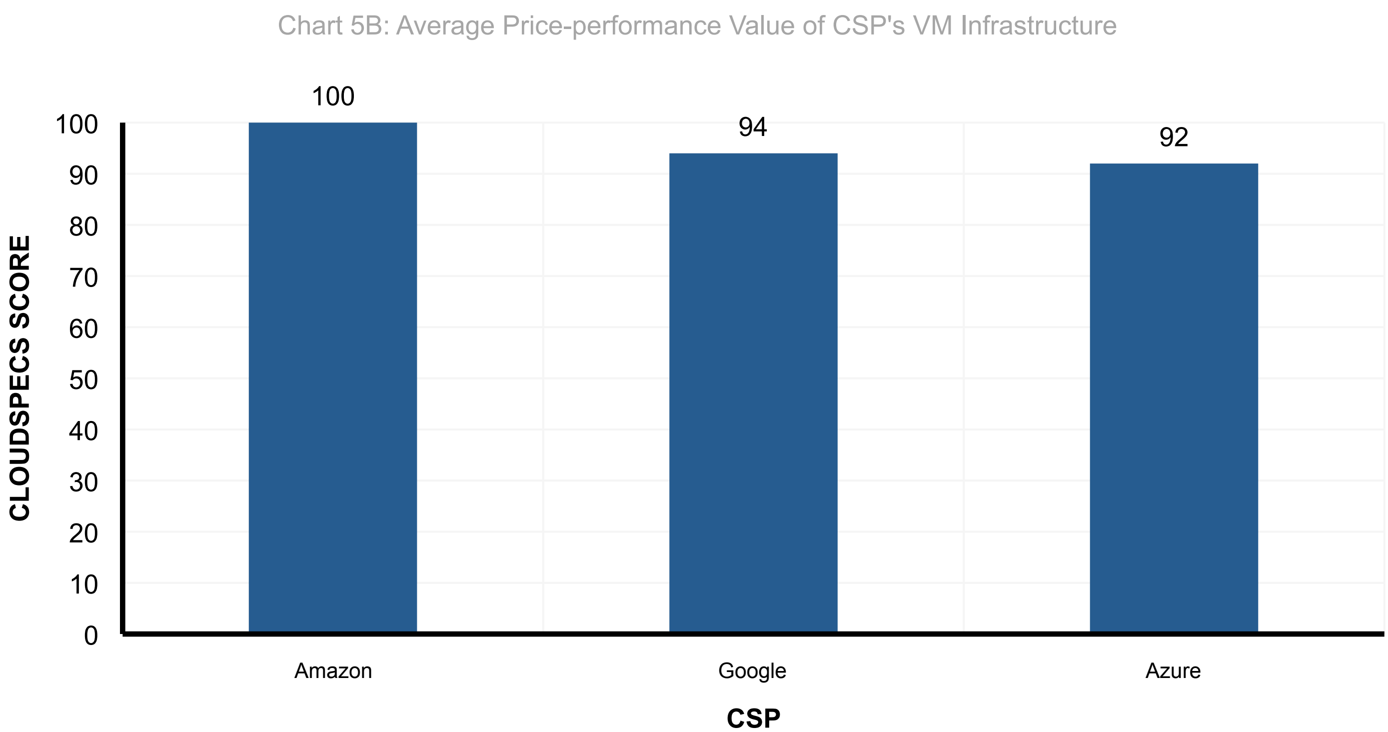


## BLOCK STORAGE VALUE

*The chart below displays the overall price-performance values of the providers' storage system*

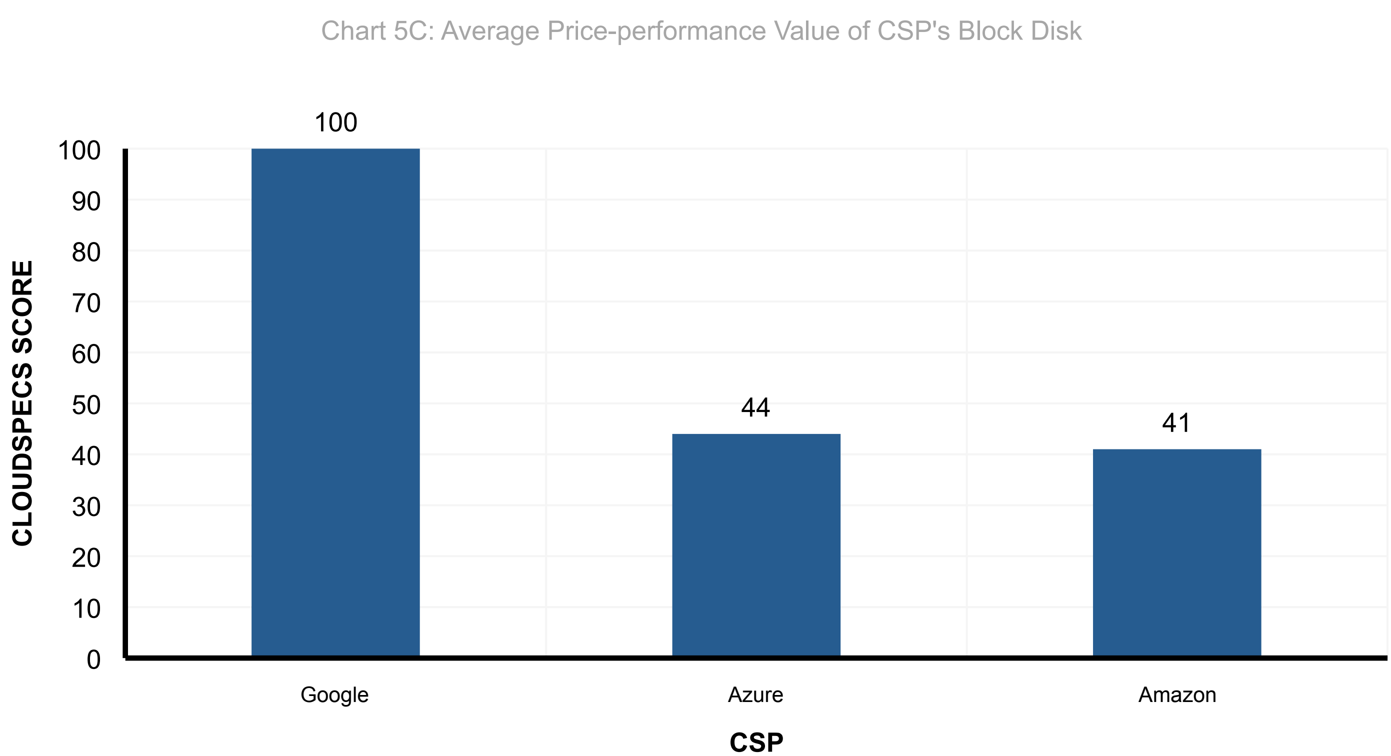

# PERFORMANCE

This section examines the performance of the three IaaS providers ranked in this report. This section does not use the CloudSpecs Score™, which is used to rank providers and can be found in the Price-Performance section of the report.

The period of 24-hour testing across three parallel machines for each category of VMs demonstrated much higher overall stability of performance in the vCPU and memory components for all providers, as compared to block storage performance over the same period. Performance differences are more noticeable as VMs scale up in size (e.g., the XL size), although a noticeable difference exists in the small VM category as well.

For detailed information on performance scores by VM size, see Performance by VM Size on page 32.

# VCPU AND MEMORY PERFORMANCE

Performance differences between CSPs on vCPU and memory remained relatively consistent across the VM sizes:

- All four VM categories exhibit a difference of 1.2x between the highest and lowest-performing CSP VMs.

*The chart below displays the median performance scores exhibited by the VMs in this study.*

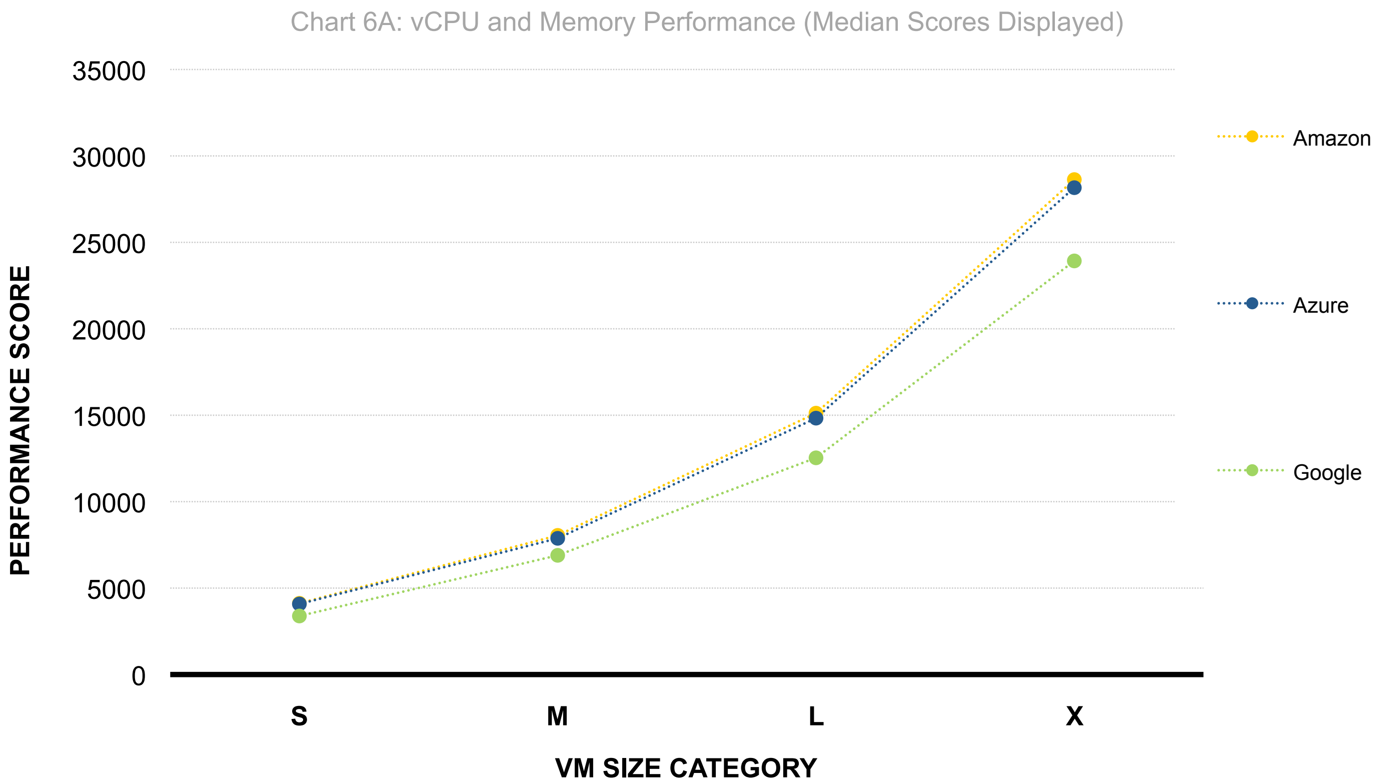

# BLOCK STORAGE PERFORMANCE

The performance of read and write in both Type 1 and Type 2 disk scenarios are shown on the following page. Detailed results can be found in the Performance by VM Size on page 32. Each varying disk size corresponded with a VM category (see Methodology page 11). More information on the two scenarios can be found in the Methodology (page 11).

- Block storage is not created equally across CSPs in regards to hardware, architecture, or performance. A difference of more than 9x can exist between highest and lowest-performing block storage offerings across CSPs.
- Reads and writes were targeted to the storage disks themselves and made to avoid hitting the cache. In testing the storage, the number of threads matched the number of vCPU's available in the VM. It is possible that as the disks became oversaturated as the thread count increased, which resulted in the lower performance as the VM size and disk size increased.
- While Amazon AWS's Small, Medium, and Large VMs show performance fluctuation, the variance is controlled. On AWS, block storage is allocated a limit of burst-performance time; the limit is dependent on the size of the block storage volume—the larger the volume, the longer the limit for burst performance. After the burst time limit expires, performance is throttled also based on the size of the volume

*The charts on the following page display the median performance scores achieved by each of the providers for Type1 and Type2 Read/Write operations.*

Chart 6B: Scenario Type 1 - Read Performance (Median Scores Displayed)

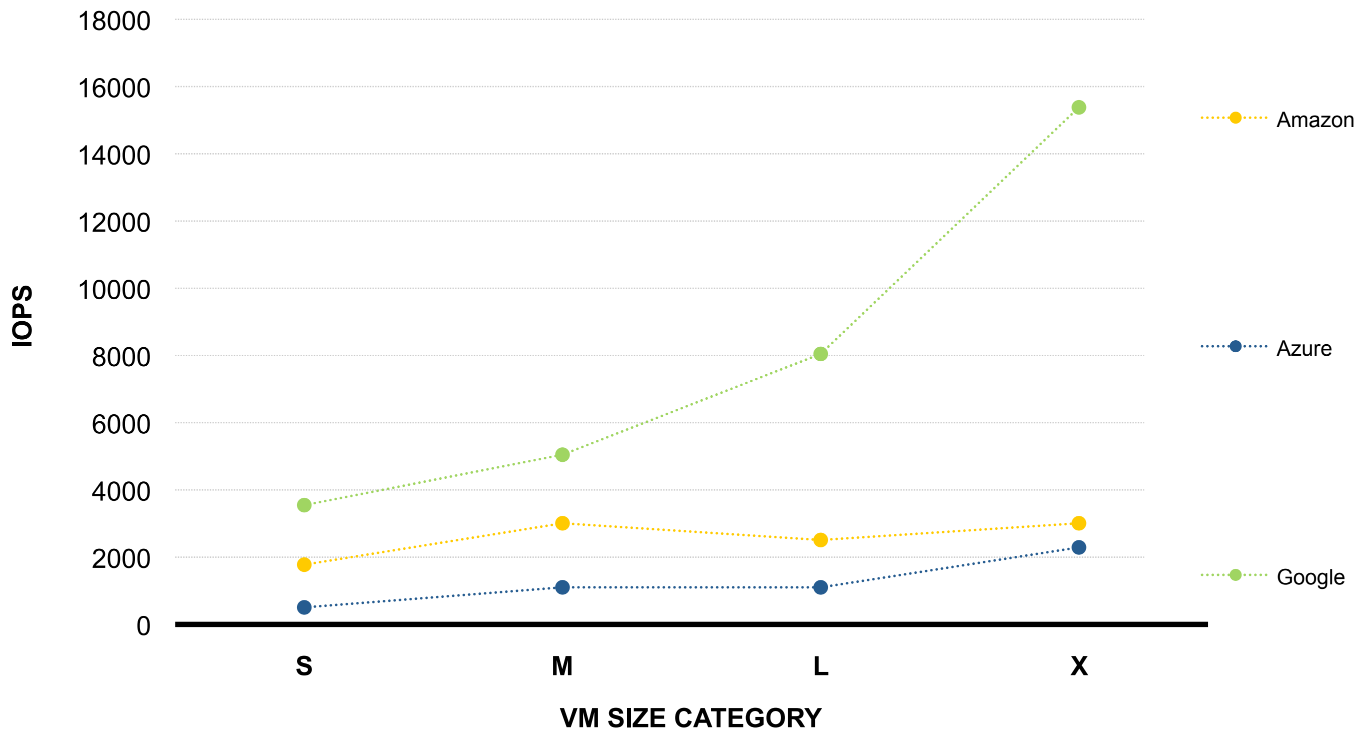


Chart 6C: Scenario Type 1 - Write Performance (Median Scores Displayed)

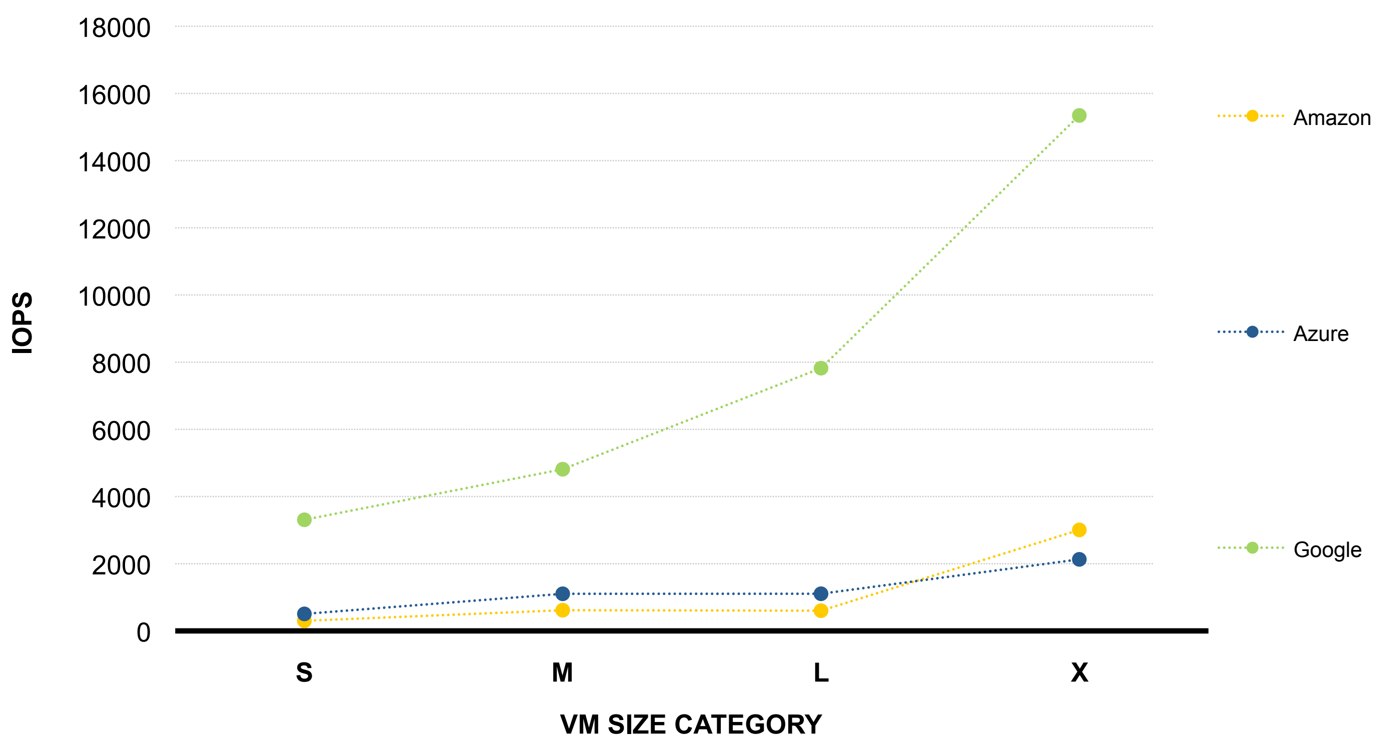

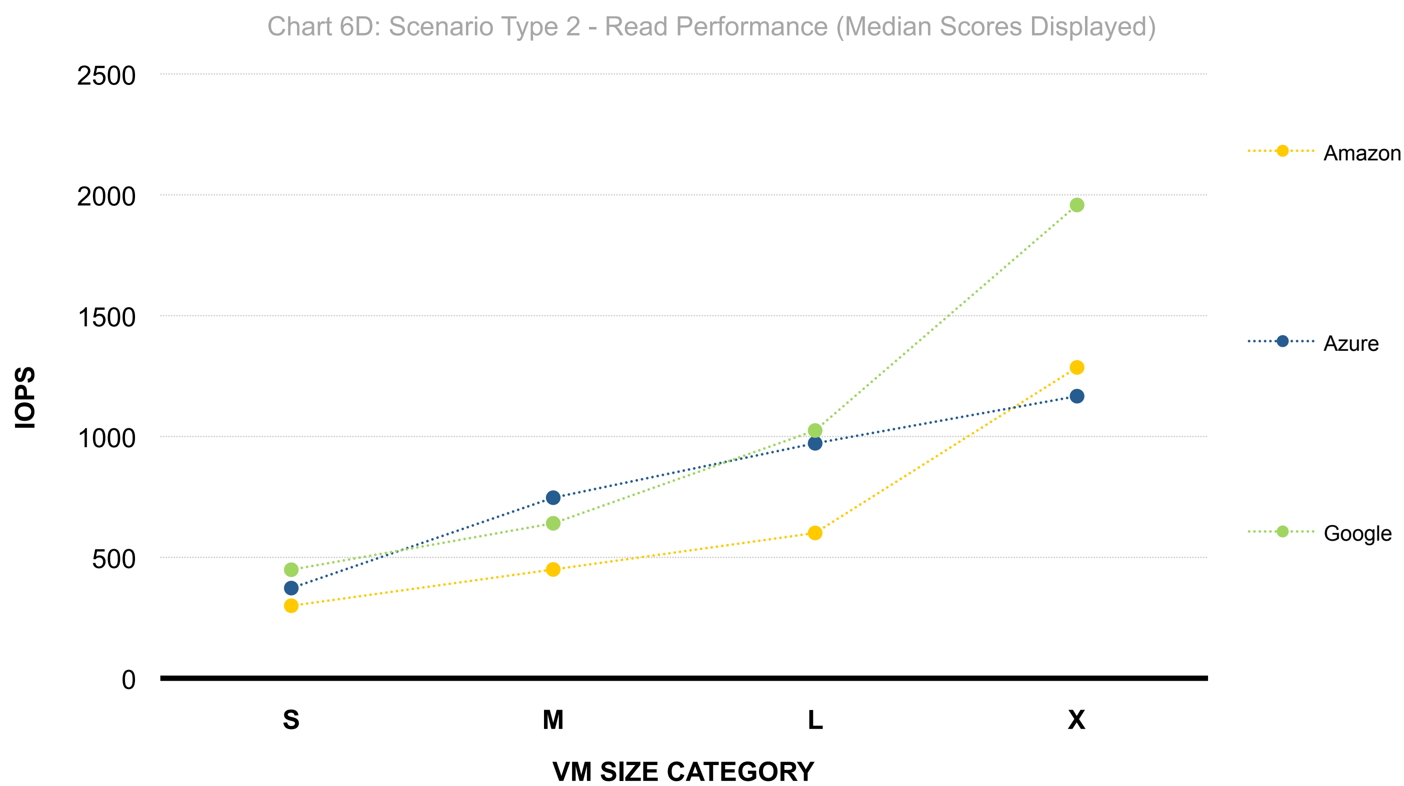
Chart 6D: Scenario Type 2 - Read Performance (Median Scores Displayed)
IOPS
2500
2000
1500
1000
500
0
S
M
L
X
VM SIZE CATEGORY
Amazon
Azure
Google

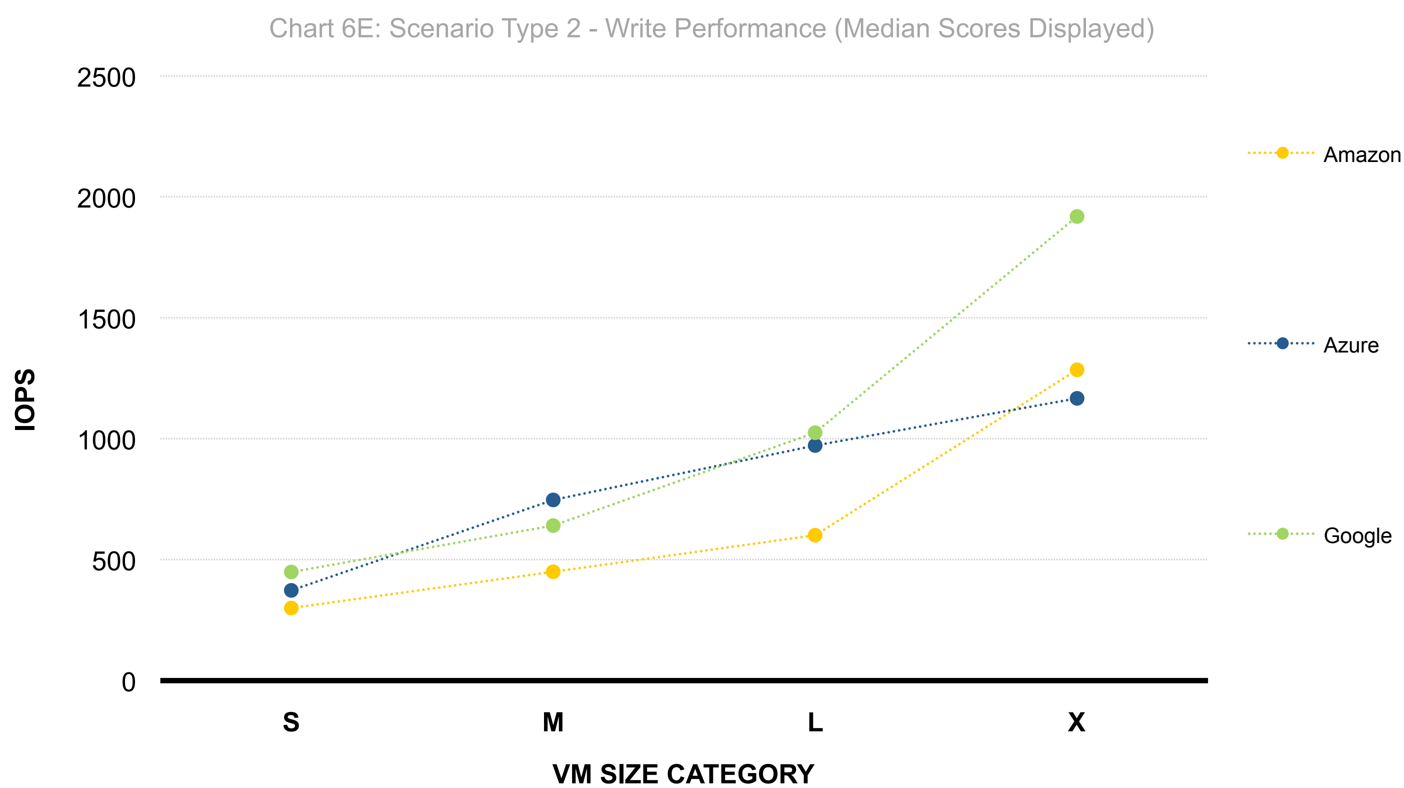
Chart 6E: Scenario Type 2 - Write Performance (Median Scores Displayed)
IOPS
2500
2000
1500
1000
500
0
S
M
L
X
VM SIZE CATEGORY
Amazon
Azure
Google

# PRICING

This section outlines the cost of the VMs and block storage for each size examined in the study across all CSPs. Additional services, unless required are not included in the final cost of the VMs. Only the cost of the VM and tested block storage were factored into the final cost. Please keep in mind that some providers may charge for add-on services such as support, while other providers include it into the cost of the VMs.

# OVERALL PRICING

The final monthly cost of each VM category for each CSP is calculated as the cost of the VM and the cost of the attached block storage. Google maintained top ranks as the least-expensive provider in the price ranking.

*The chart below displays the monthly costs of the providers' VMs and storage of the resources examined in the study.*

Chart 7A: Monthly Cost of VMs Across CSPs

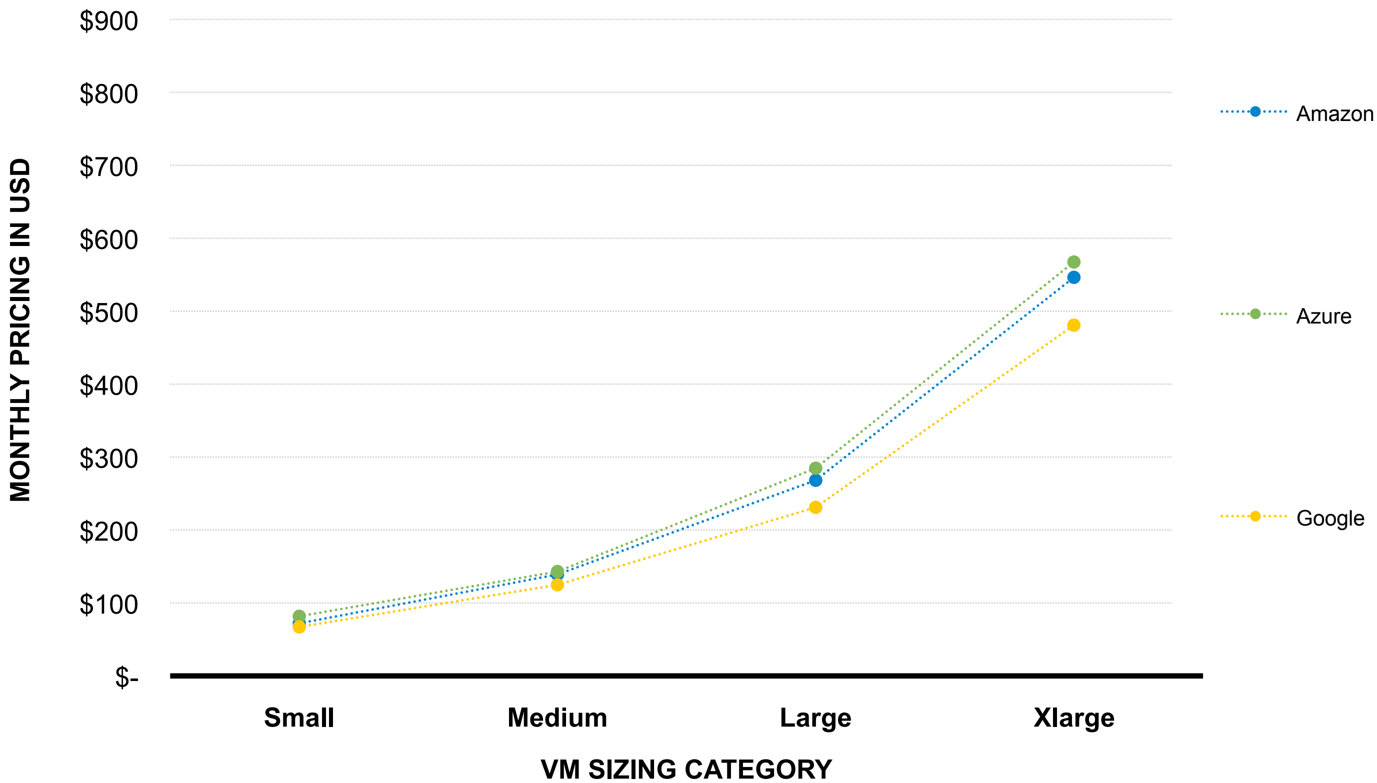


Table 7A: Monthly Cost of VMs Across CSPs

| | Small | Medium | Large | Extra Large |
|---|---|---|---|---|
| Amazon | $ 72.05 | $ 139.10 | $ 268.20 | $ 546.40 |
| Azure | $ 81.76 | $ 143.08 | $ 284.76 | $ 567.43 |
| Google | $ 67.12 | $ 124.89 | $ 231.08 | $ 480.86 |

# PRICING BY VM CATEGORY

From small to large-sized VMs, the provider rankings remained the same from the least to most expensive provider. Google offers the least-expensive VM, followed by Amazon and then Azure.

Chart 7B: Monthly Cost of Small VMs

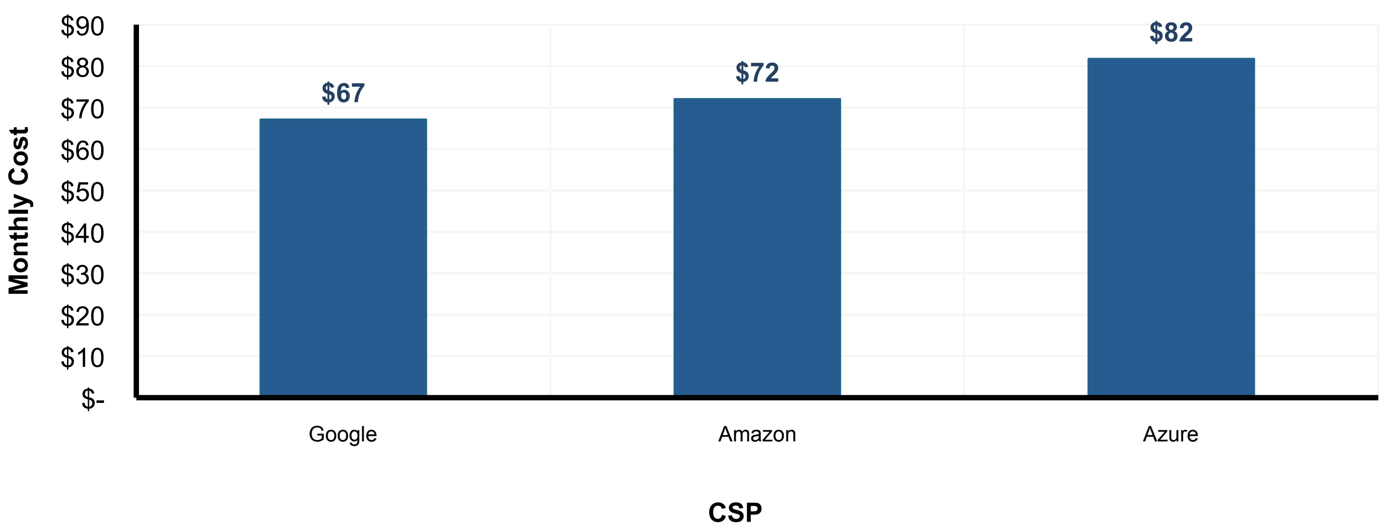


Chart 7C: Monthly Cost of Medium VMs

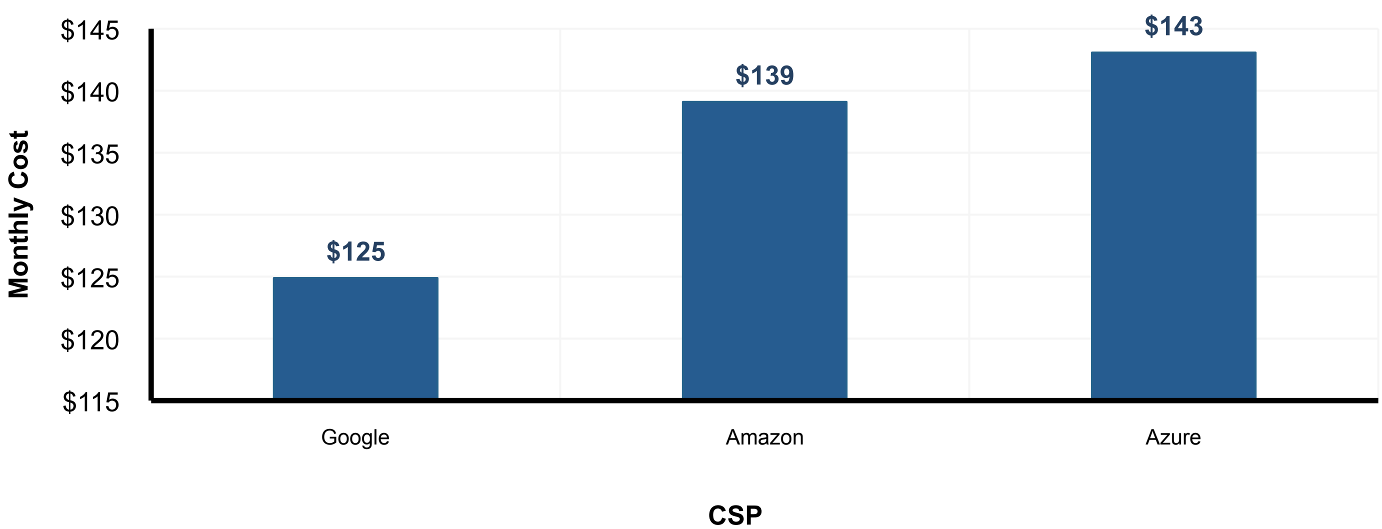

Chart 7D: Monthly Cost of Large VMs

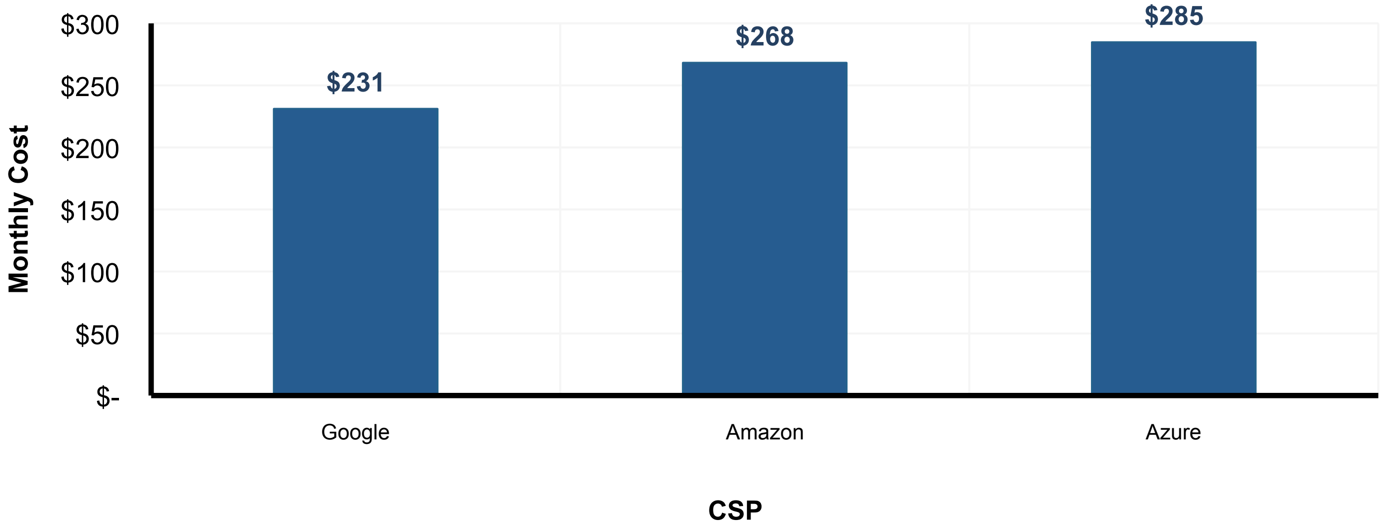


Chart 7E: Monthly Cost of Extra Large VMs

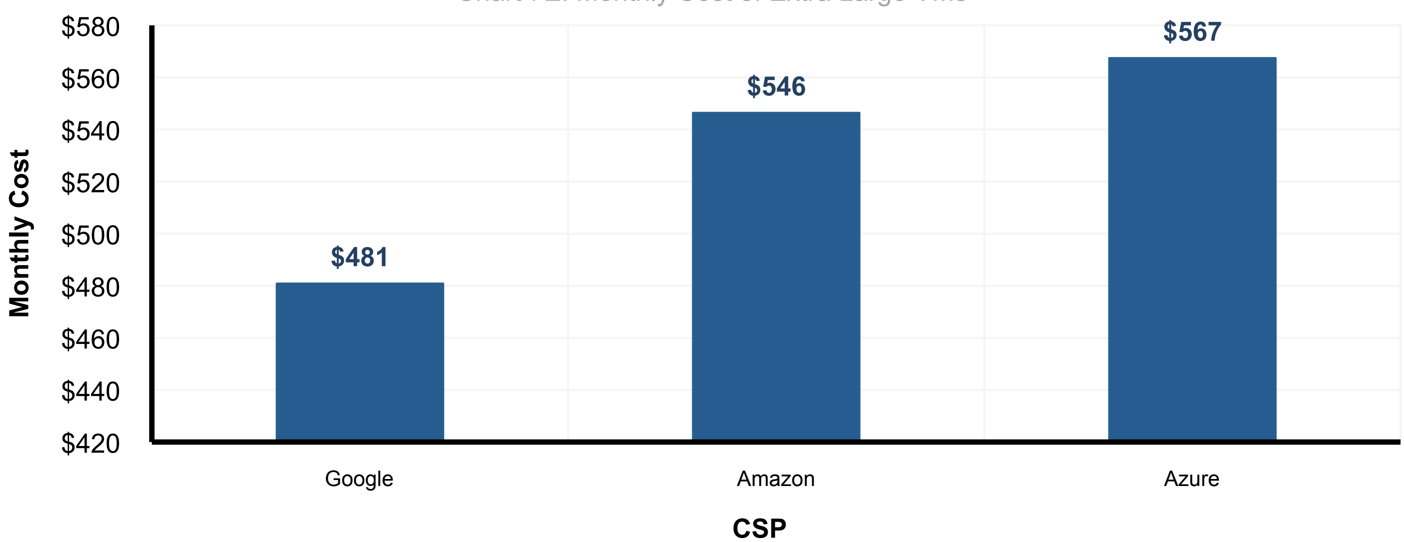

# PERFORMANCE BY VM SIZE

## UNDERSTANDING THE CHARTS

VM performance is illustrated using percentile scores retrieved from all data points collected. 5th percentile and 95th percentile scores are often used instead of minimum and maximum scores in order to exclude potential outliers. Median scores are used instead of mean to avoid values being skewed by outliers. The information has been integrated into percentile graphs and value tables designed to visualize performance variation captured while testing over time. An example of the performance percentile graph along with a corresponding value table is displayed in Chart 8:

Chart 8: Sample Chart Guide

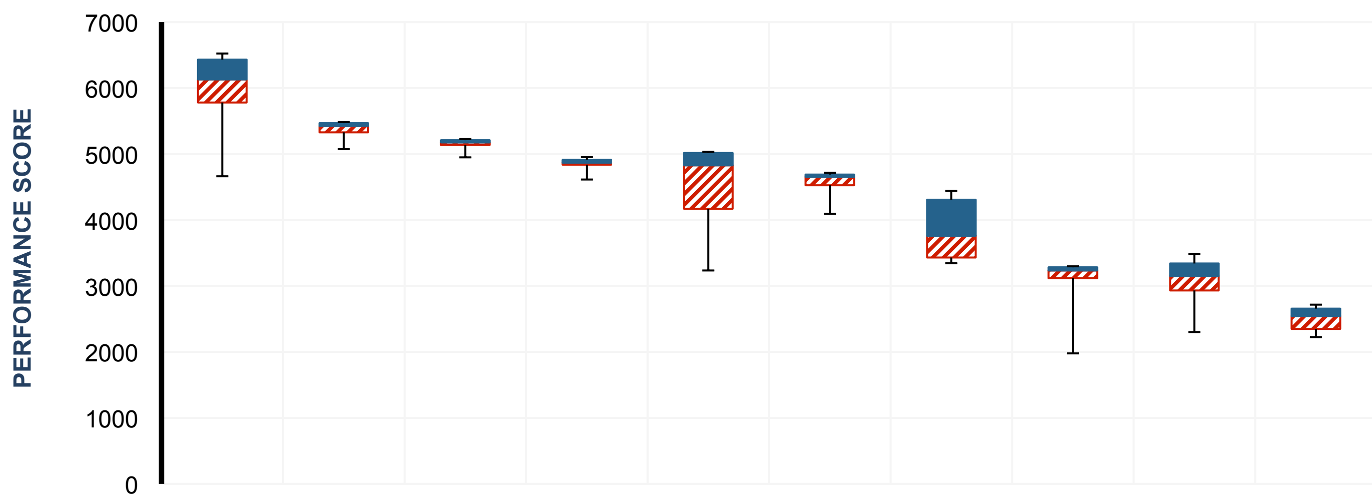


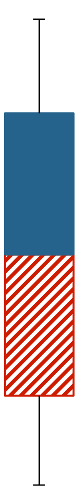

**Maximum (MAX)**: The highest performance score(s) achieved on the VM over the course of the

**95TH Percentile (95TH)**: 95% of all scores on the VM achieved this performance score or

**Median (MED)**: The number separating the higher and lower half of scores. If the median is closer to the 95TH percentile, then more high performance scores were observed than low performance scores (and vice versa).

**5TH Percentile (5TH)**: 5% of all scores on the VM achieved this performance score or

**Minimum (MIN)**: The lowest performance score(s) achieved on the VM over the course of the

# SMALL VMs

Chart 8A.1: VM Performance (Small VMs)

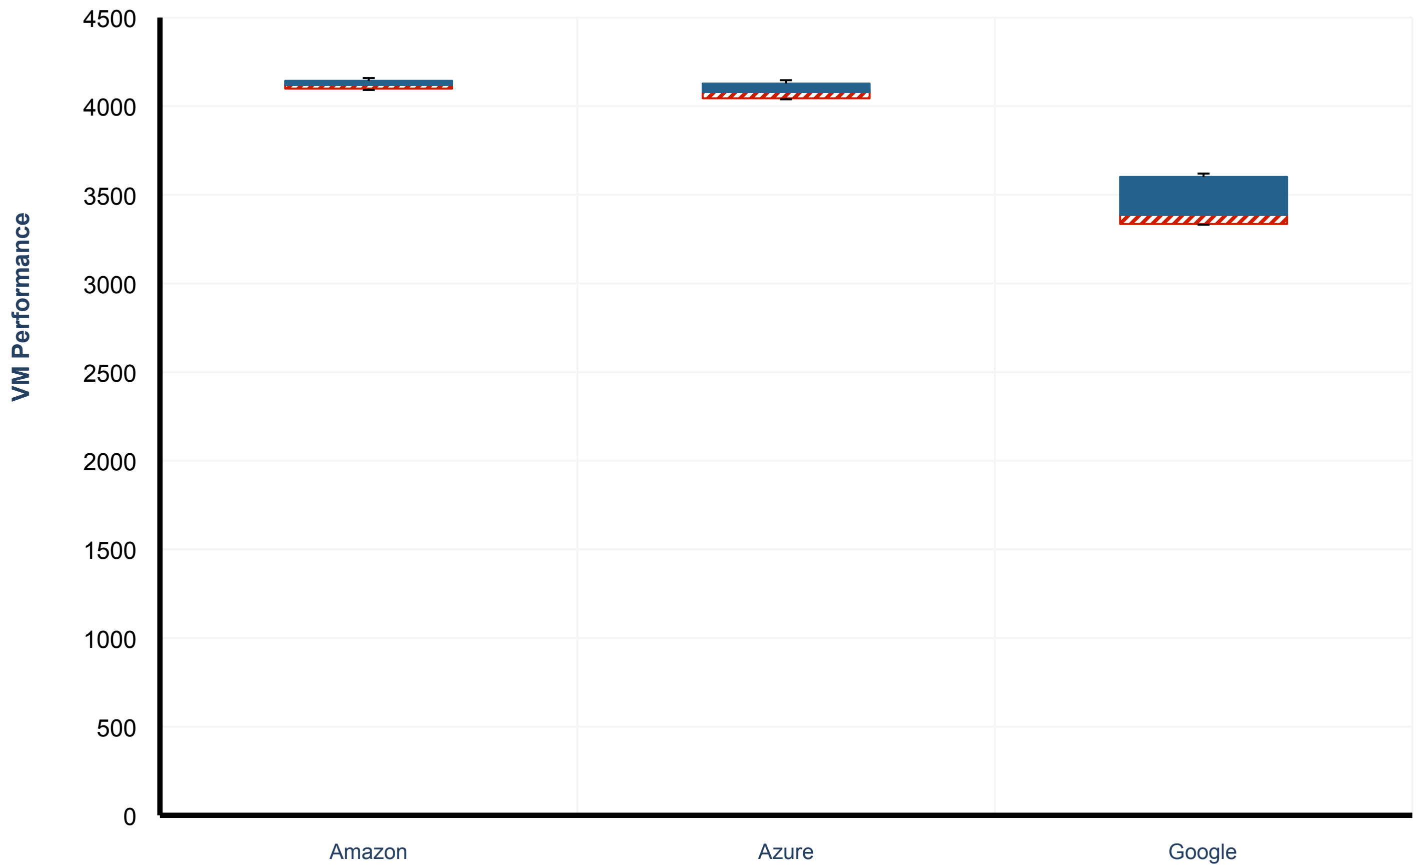


Table 8A.1: VM Performance (Small VMs)

| PROVIDER | MIN | 5TH | MEDIAN | 95TH | MAX | STDEV | CV |
|---|---|---|---|---|---|---|---|
| Amazon | 4091 | 4099 | 4120 | 4143 | 4159 | 13 | 0% |
| Azure | 4039 | 4045 | 4080 | 4128 | 4147 | 28 | 1% |
| Google | 3332 | 3336 | 3387 | 3601 | 3620 | 108 | 3% |

Chart 8A.2: Read Block Disk Performance Type 1 (Small VM)

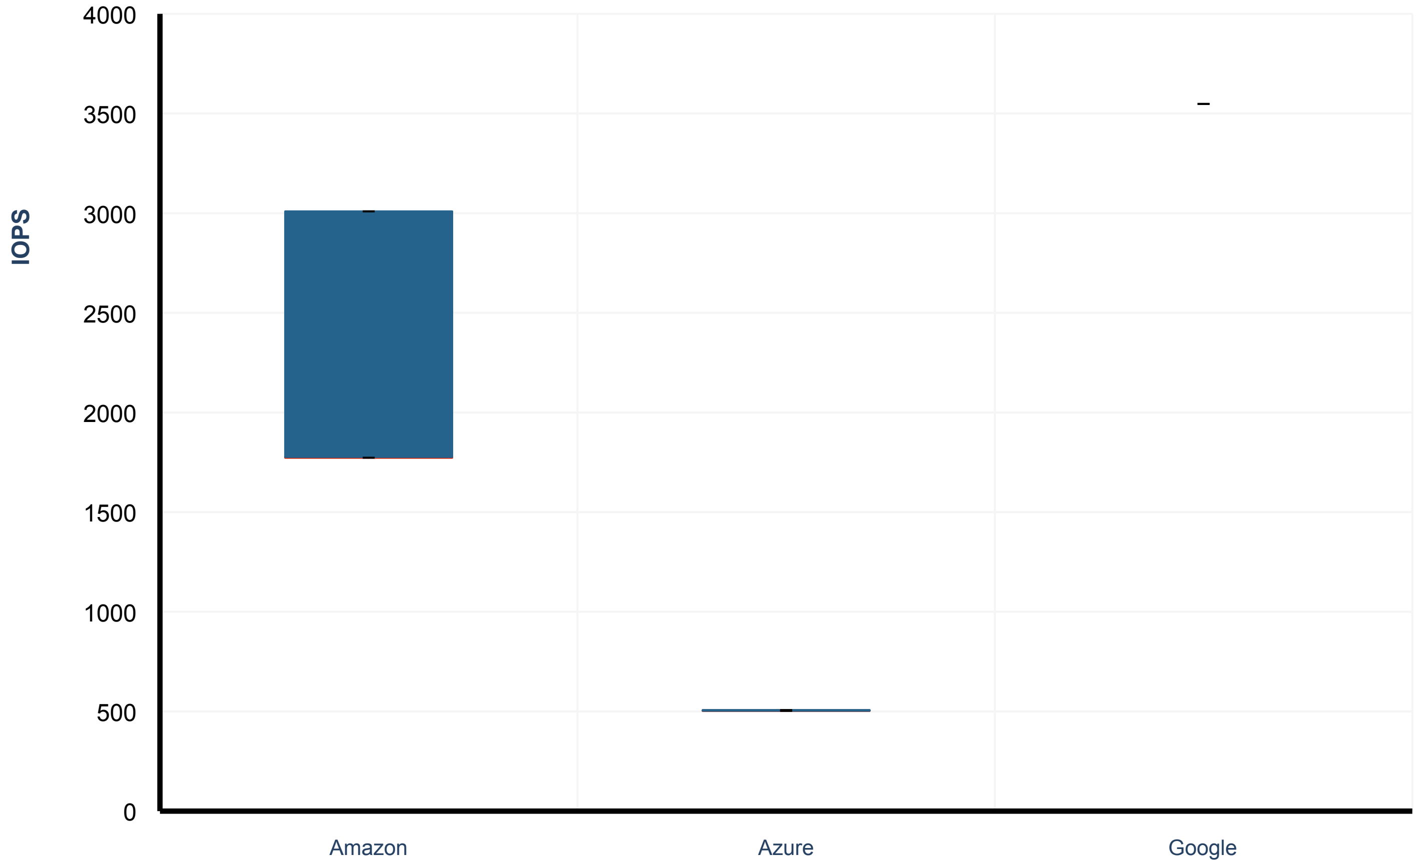


Table 8A.2: Read Block Disk Performance Type 1 (Small VMs)

| PROVIDER | MIN | 5TH | MEDIAN | 95TH | MAX | STDEV | CV |
|---|---|---|---|---|---|---|---|
| Amazon | 1773 | 1773 | 1777 | 3009 | 3009 | 411 | 21% |
| Azure | 503 | 505 | 506 | 507 | 507 | 1 | 0% |
| Google | 3548 | 3548 | 3548 | 3548 | 3548 | 0 | 0% |

Chart 8A.3: Write Block Disk Performance Type 1 (Small VM)

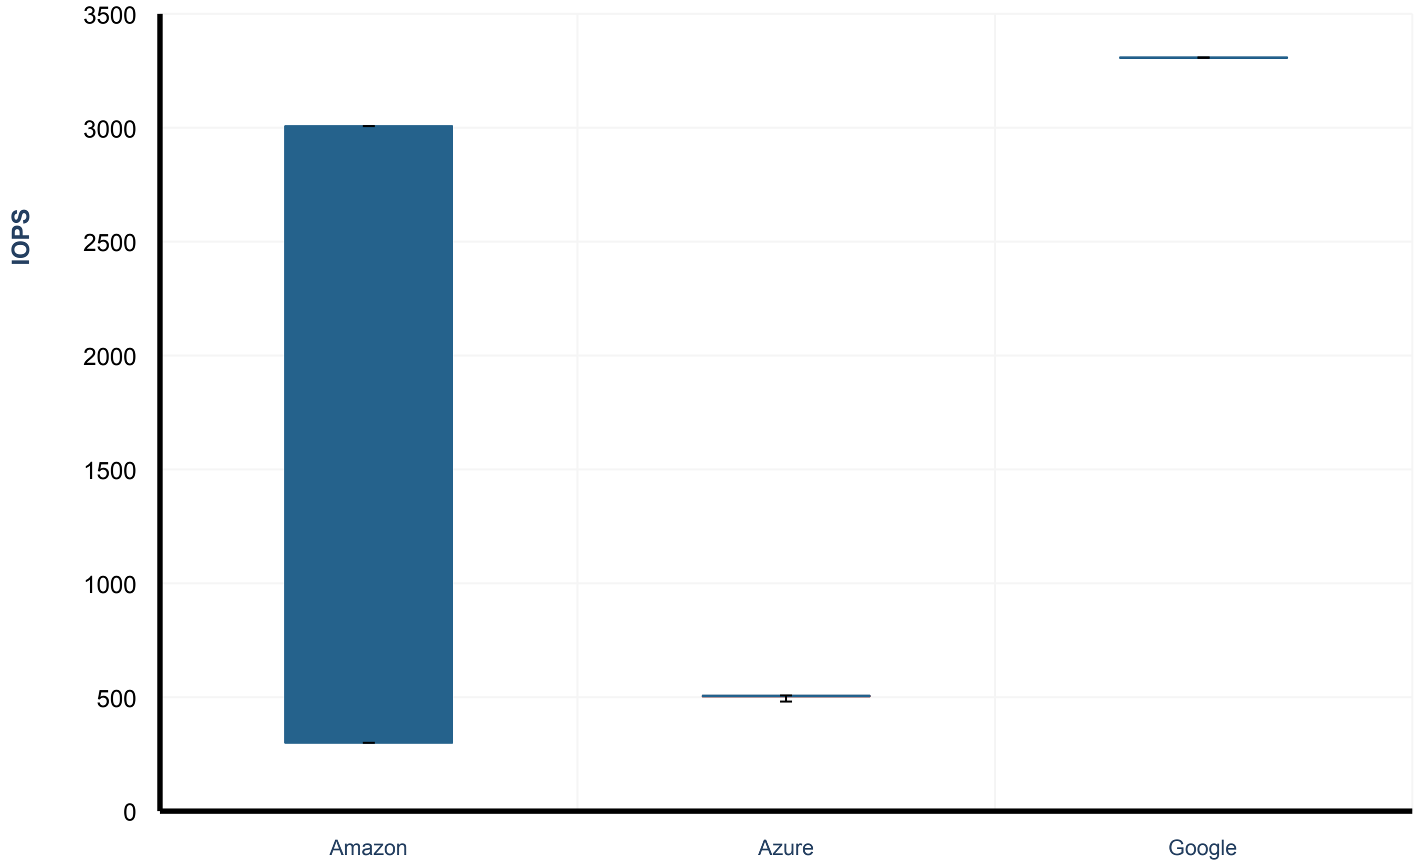


Table 8A.3: Write Block Disk Performance Type 1 (Small VMs)

| PROVIDER | MIN | 5TH | MEDIAN | 95TH | MAX | STDEV | CV |
|---|---|---|---|---|---|---|---|
| Amazon | 300 | 300 | 300 | 3007 | 3007 | 902 | 141% |
| Azure | 481 | 505 | 506 | 507 | 508 | 4 | 1% |
| Google | 3307 | 3308 | 3308 | 3309 | 3309 | 0 | 0% |

Chart 8A.4: Read Block Disk Performance Type 2 (Small VM)

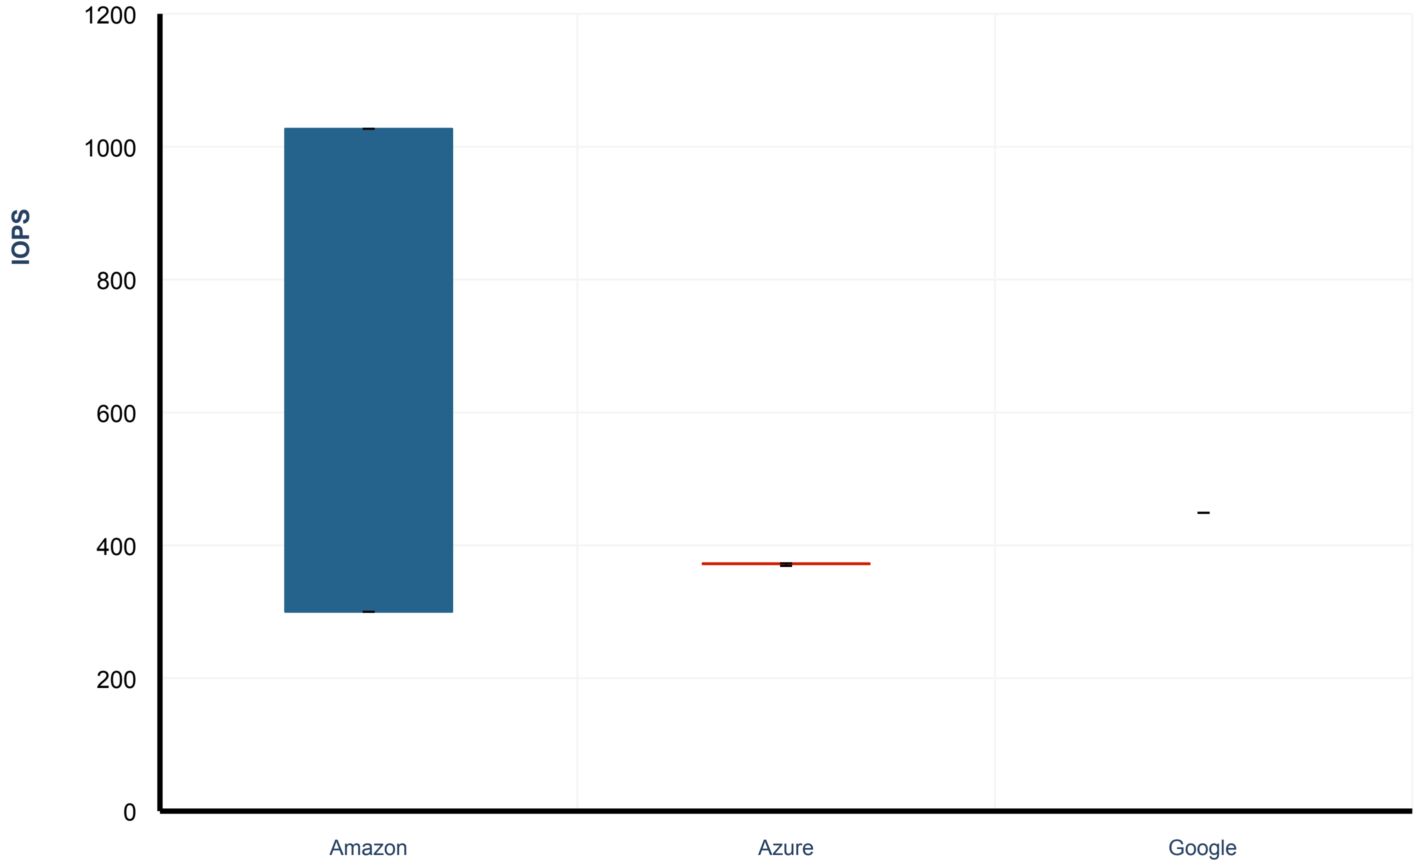


Table 8A.4: Read Block Disk Performance Type 2 (Small VMs)

| PROVIDER | MIN | 5TH | MEDIAN | 95TH | MAX | STDEV | CV |
|---|---|---|---|---|---|---|---|
| Amazon | 300 | 300 | 300 | 1027 | 1027 | 242 | 62% |
| Azure | 369 | 372 | 373 | 373 | 373 | 1 | 0% |
| Google | 449 | 449 | 449 | 449 | 449 | 0 | 0% |

Chart 8A.5: Write Block Disk Performance Type 2 (Small VM)

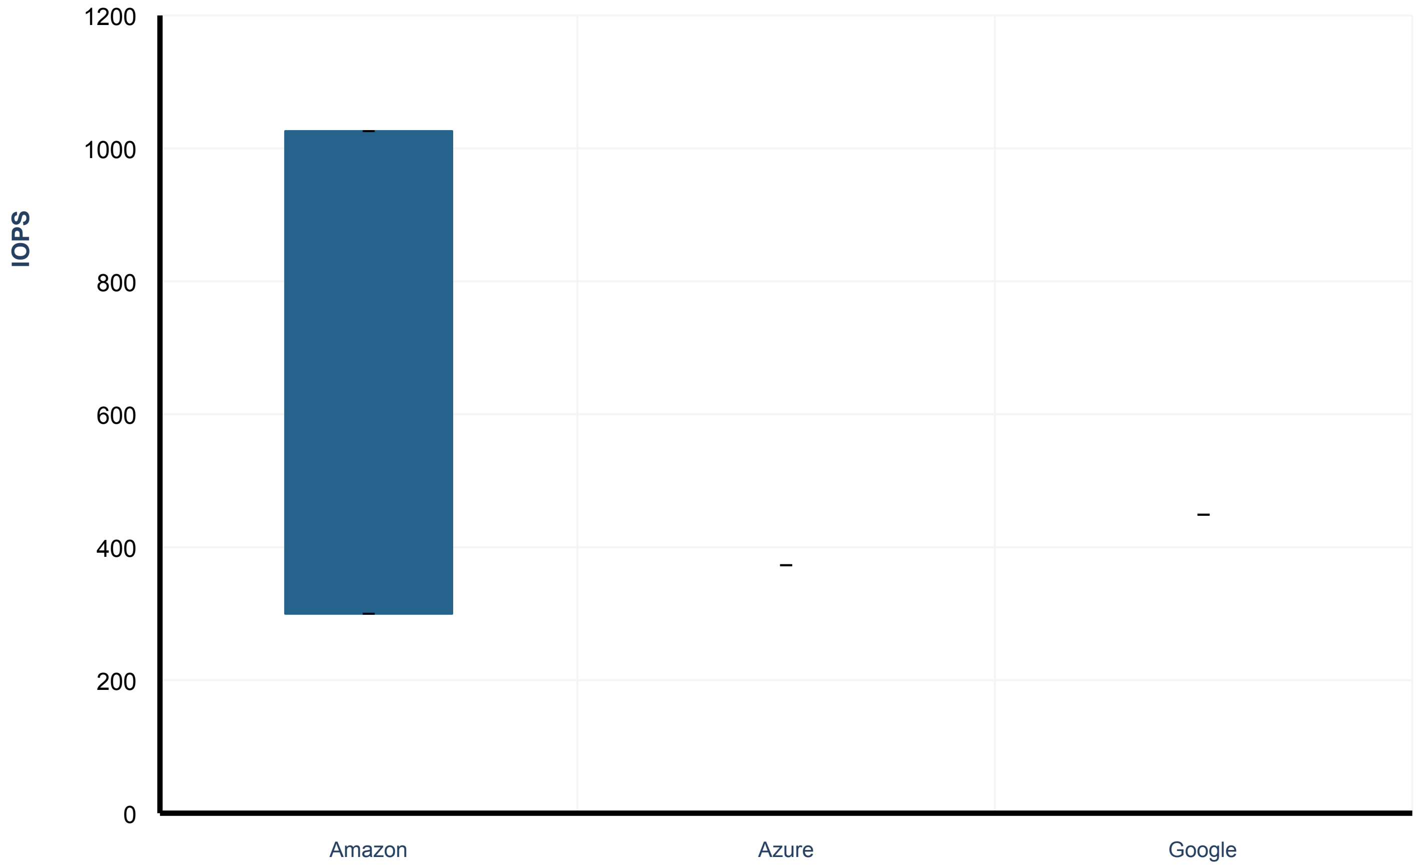


Table 8A.5: Write Block Disk Performance Type 2 (Small VMs)

| PROVIDER | MIN | 5TH | MEDIAN | 95TH | MAX | STDEV | CV |
|---|---|---|---|---|---|---|---|
| Amazon | 300 | 300 | 300 | 1026 | 1026 | 242 | 62% |
| Azure | 373 | 373 | 373 | 373 | 373 | 0 | 0% |
| Google | 449 | 449 | 449 | 449 | 449 | 0 | 0% |

# MEDIUM VMs

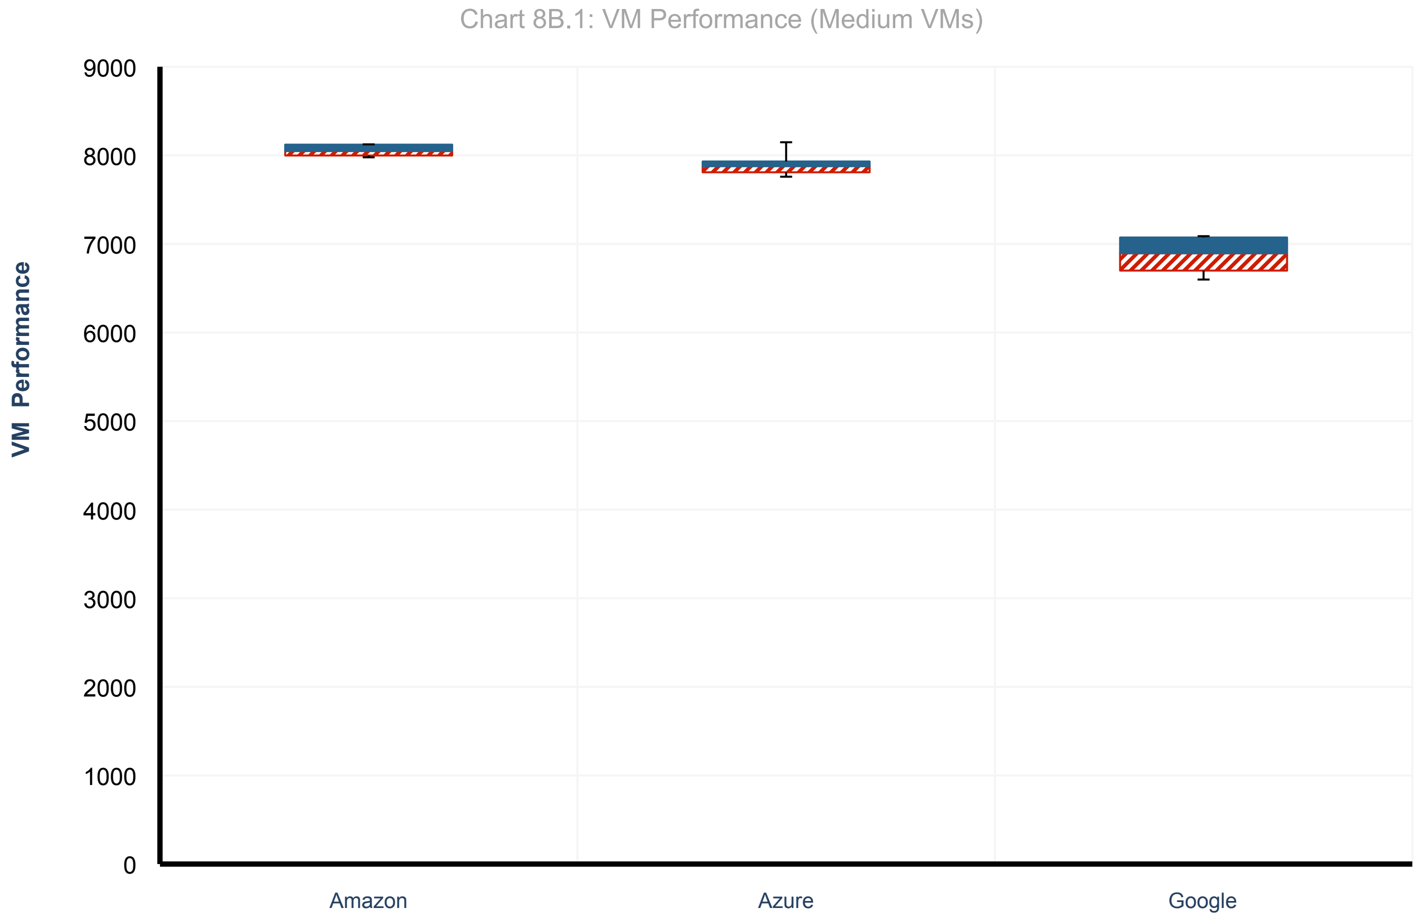


Table 8B.1: VM Performance (Medium VMs)

| PROVIDER | MIN | 5TH | MEDIAN | 95TH | MAX | STDEV | CV |
|---|---|---|---|---|---|---|---|
| Amazon | 7978 | 7998 | 8049 | 8121 | 8124 | 42 | 1% |
| Azure | 7759 | 7809 | 7878 | 7931 | 8148 | 67 | 1% |
| Google | 6598 | 6699 | 6895 | 7073 | 7088 | 114 | 2% |

Chart 8B.2: Read Block Disk Performance Type 1 (Medium VM)

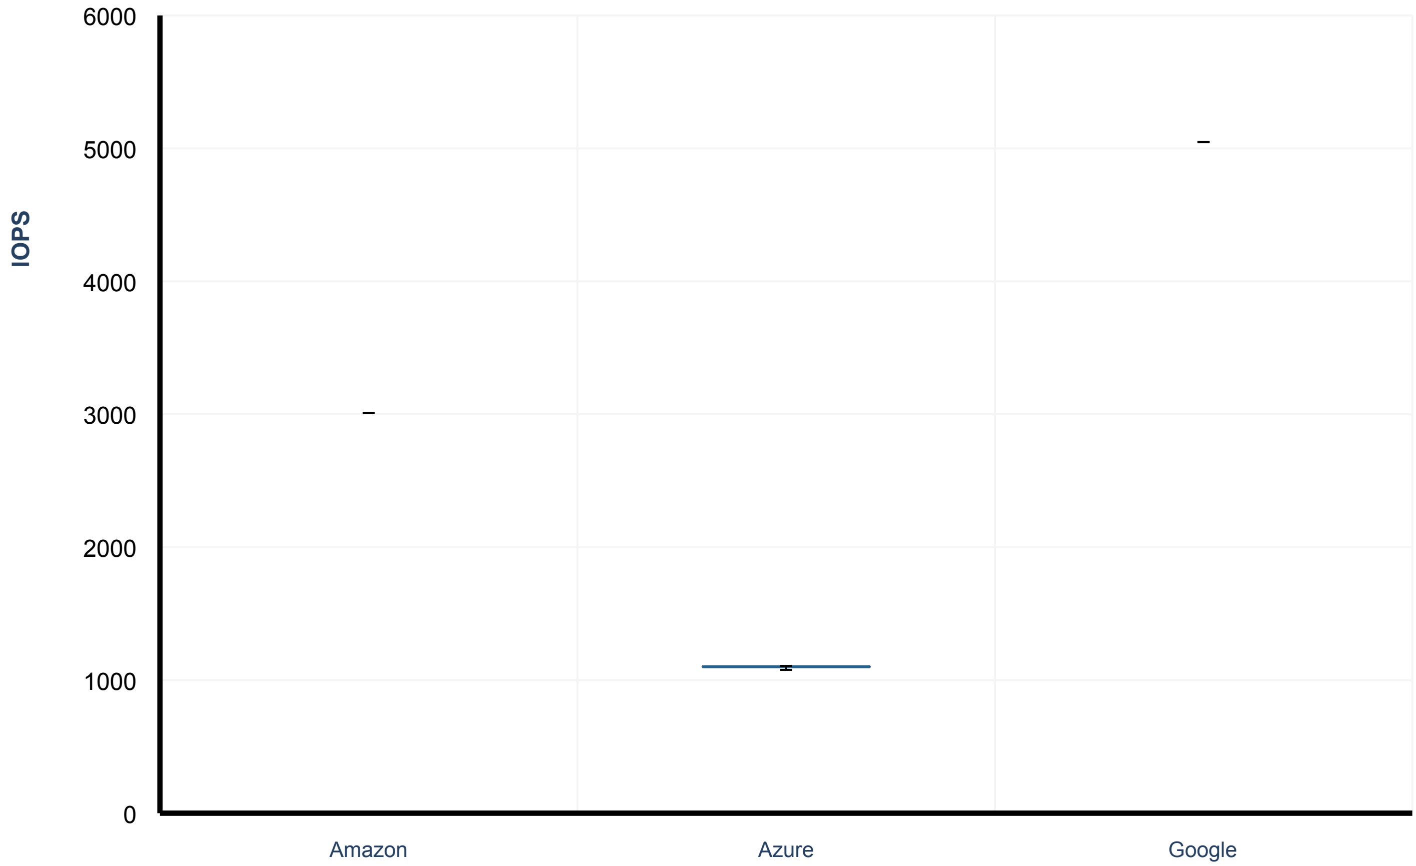


Table 8B.2: Read Block Disk Performance Type 1 (Medium VMs)

| PROVIDER | MIN | 5TH | MEDIAN | 95TH | MAX | STDEV | CV |
|---|---|---|---|---|---|---|---|
| Amazon | 3009 | 3009 | 3009 | 3009 | 3009 | 0 | 0% |
| Azure | 1078 | 1099 | 1103 | 1105 | 1108 | 5 | 0% |
| Google | 5047 | 5048 | 5048 | 5048 | 5048 | 0 | 0% |

Chart 8B.3: Write Block Disk Performance Type 1 (Medium VM)

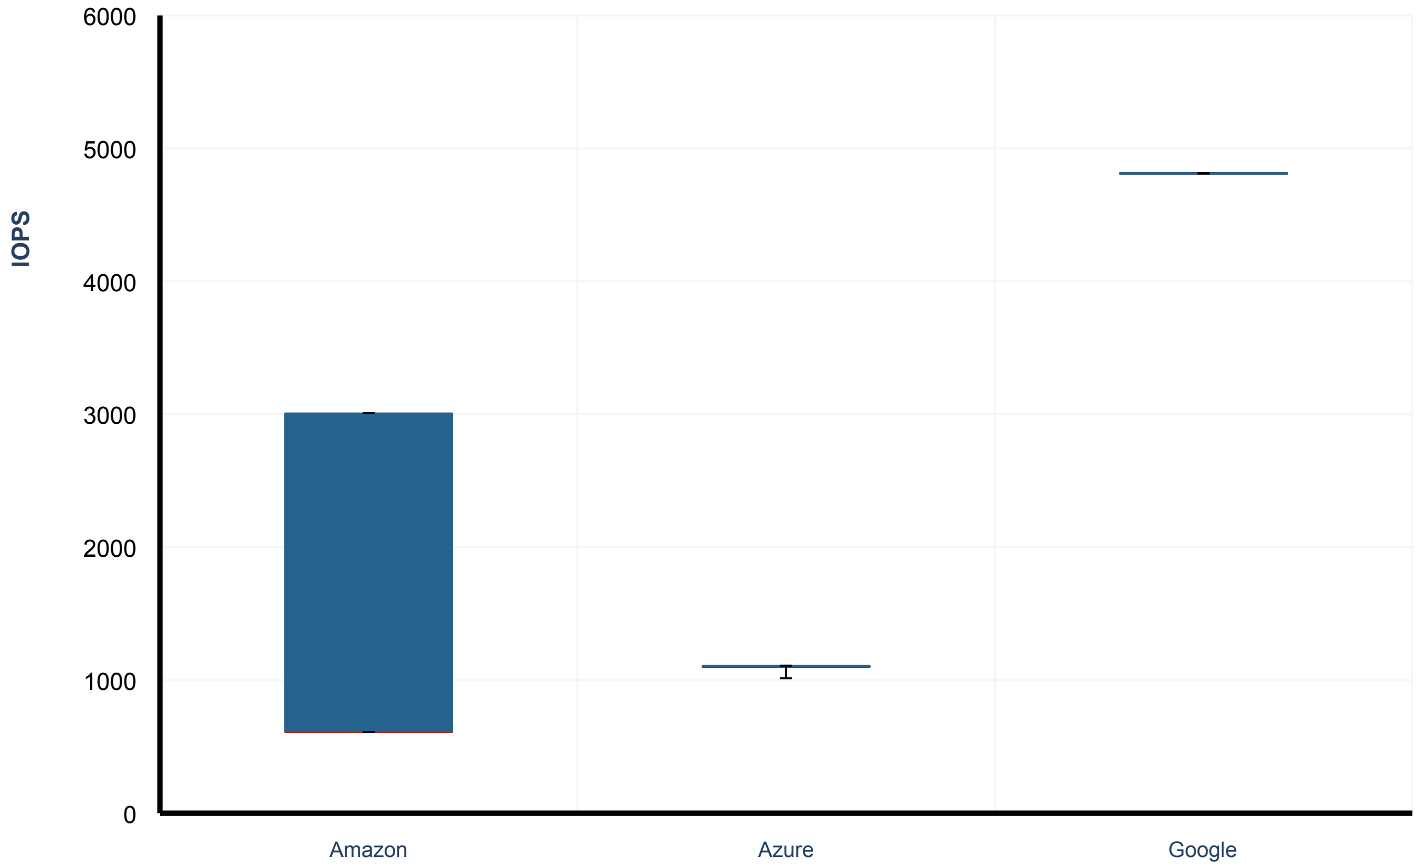


Table 8B.3: Write Block Disk Performance Type 1 (Medium VMs)

| PROVIDER | MIN | 5TH | MEDIAN | 95TH | MAX | STDEV | CV |
|---|---|---|---|---|---|---|---|
| Amazon | 611 | 612 | 616 | 3006 | 3009 | 903 | 85% |
| Azure | 1015 | 1100 | 1105 | 1109 | 1109 | 17 | 2% |
| Google | 4811 | 4812 | 4813 | 4814 | 4814 | 1 | 0% |

Chart 8B.4: Read Block Disk Performance Type 2 (Medium VM)

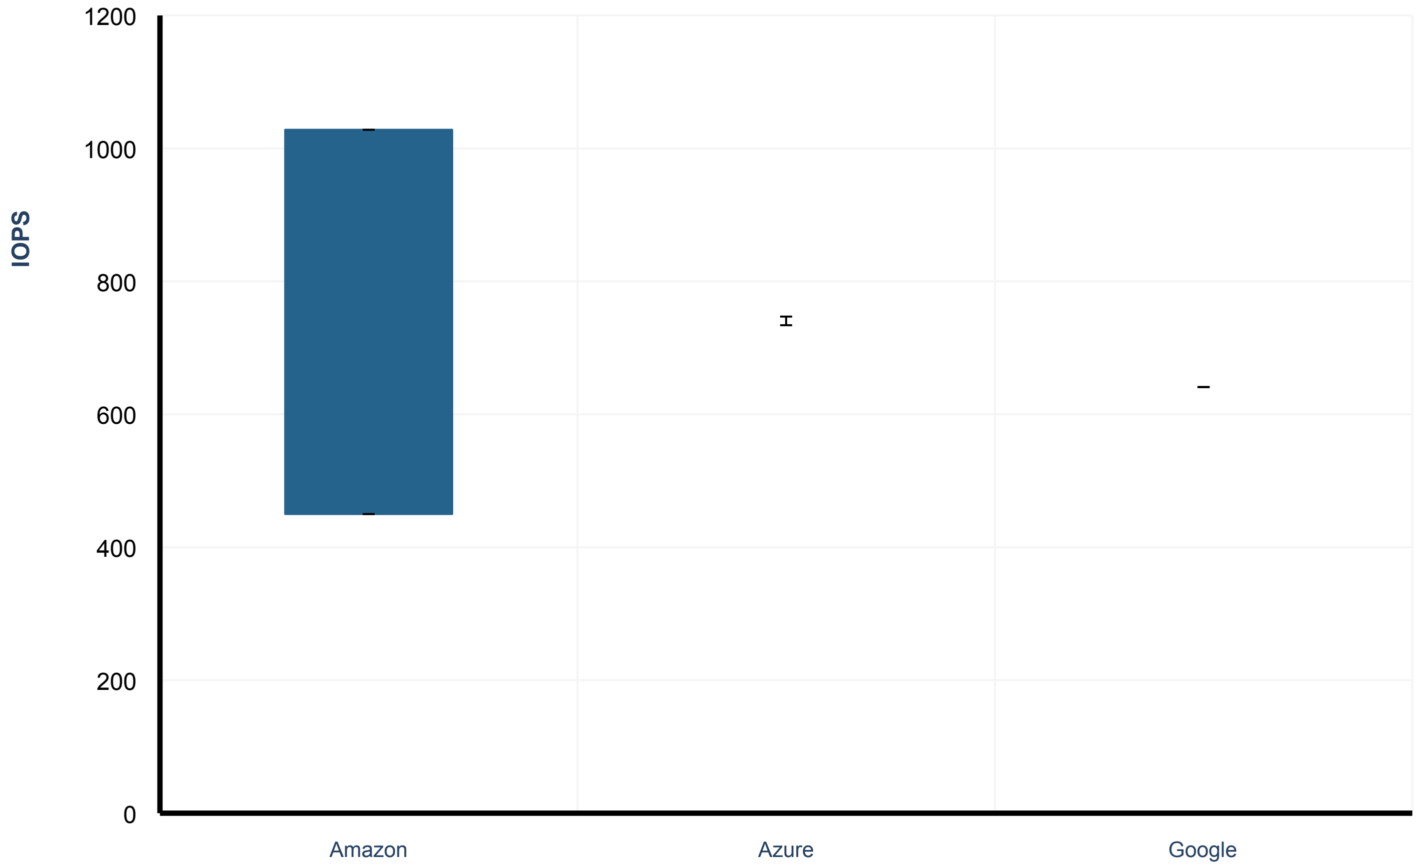


Table 8B.4: Read Block Disk Performance Type 2 (Medium VMs)

| PROVIDER | MIN | 5TH | MEDIAN | 95TH | MAX | STDEV | CV |
|---|---|---|---|---|---|---|---|
| Amazon | 450 | 450 | 450 | 1028 | 1028 | 217 | 40% |
| Azure | 734 | 747 | 747 | 747 | 747 | 2 | 0% |
| Google | 641 | 641 | 641 | 641 | 641 | 0 | 0% |

Chart 8B.5: Write Block Disk Performance Type 2 (Medium VM)

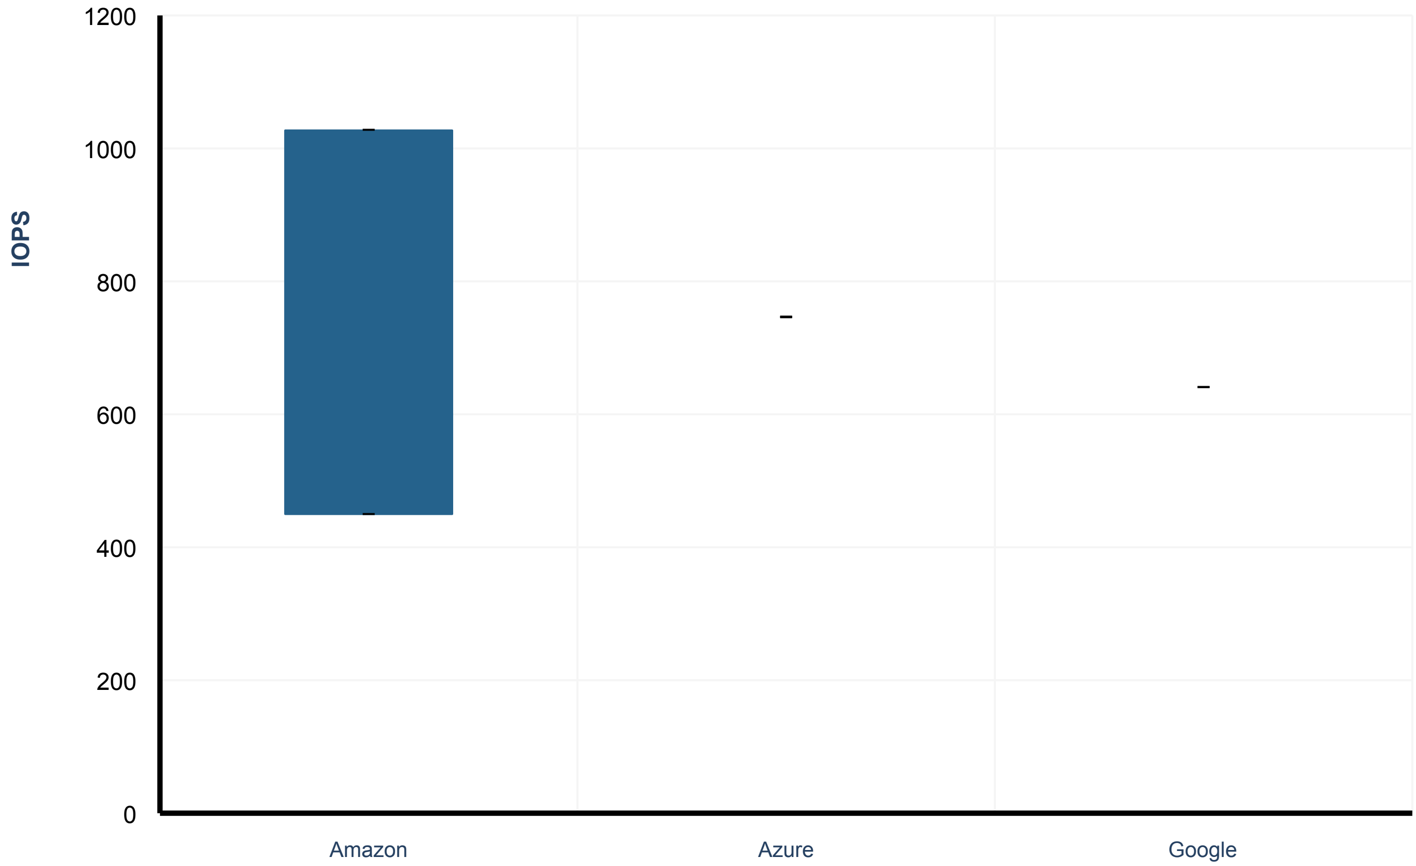


Table 8B.5: Write Block Disk Performance Type 2 (Medium VMs)

| PROVIDER | MIN | 5TH | MEDIAN | 95TH | MAX | STDEV | CV |
|---|---|---|---|---|---|---|---|
| Amazon | 450 | 450 | 450 | 1027 | 1028 | 216 | 40% |
| Azure | 746 | 747 | 747 | 747 | 747 | 0 | 0% |
| Google | 641 | 641 | 641 | 641 | 641 | 0 | 0% |

# LARGE VMs

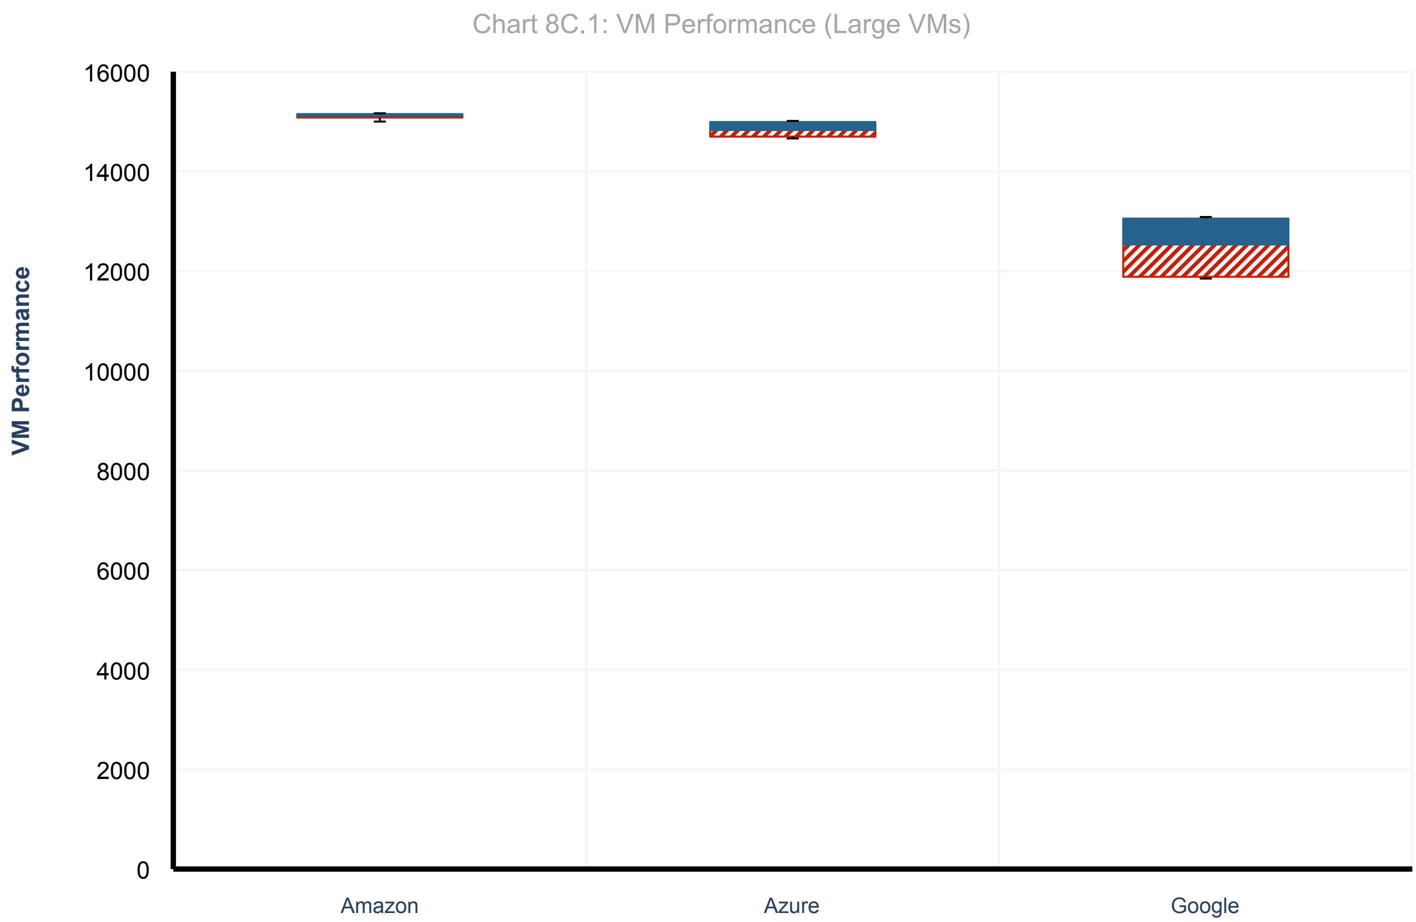


Table 8C.1: VM Performance (Large VMs)

| PROVIDER | MIN | 5TH | MEDIAN | 95TH | MAX | STDEV | CV |
|---|---|---|---|---|---|---|---|
| Amazon | 15000 | 15083 | 15119 | 15150 | 15167 | 24 | 0% |
| Azure | 14662 | 14699 | 14831 | 14991 | 15013 | 99 | 1% |
| Google | 11850 | 11886 | 12536 | 13053 | 13085 | 436 | 3% |

Chart 8C.2: Read Block Disk Performance Type 1 (Large VM)

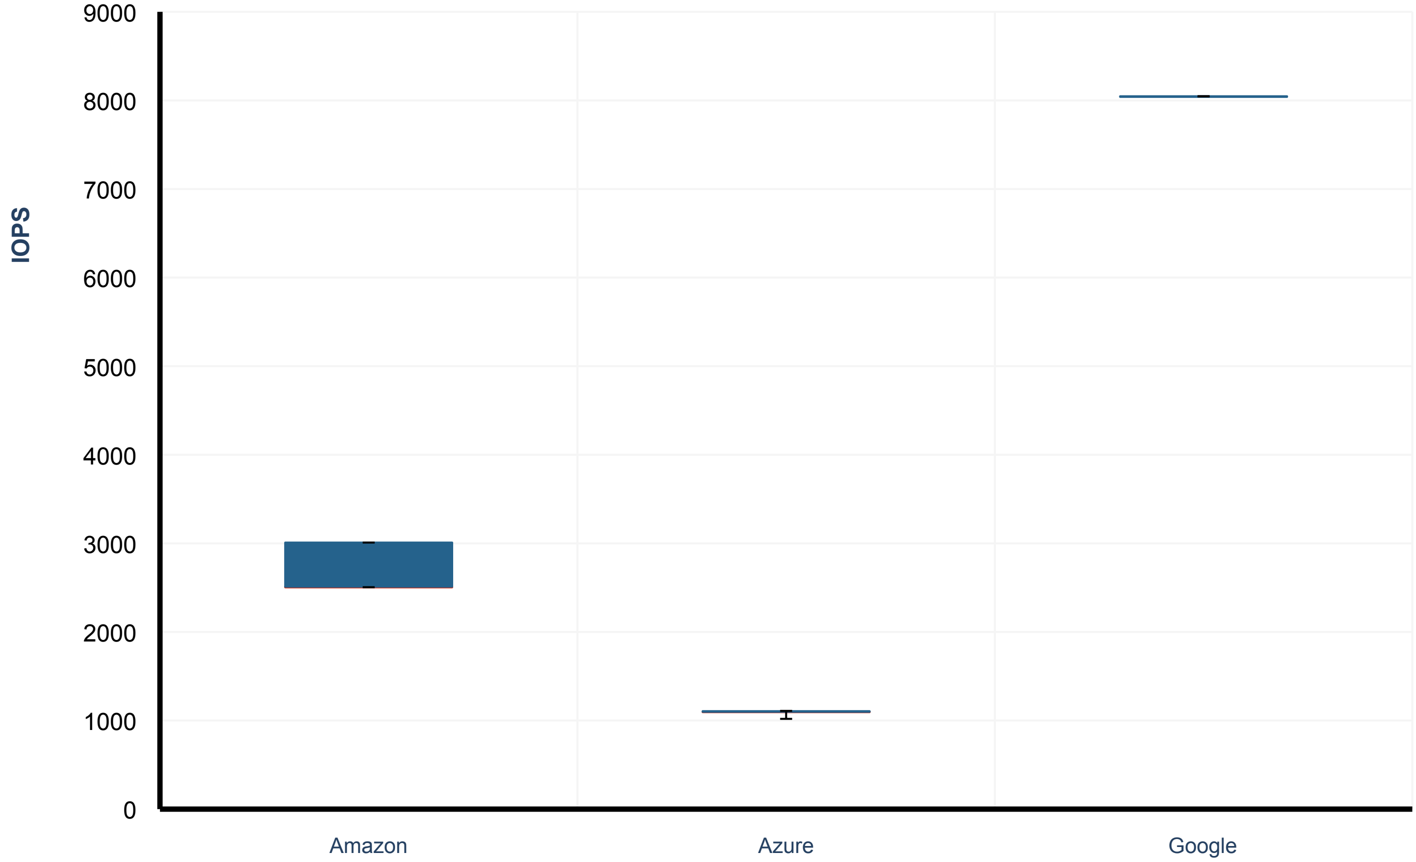


Table 8C.2: Read Block Disk Performance Type 1 (Large VMs)

| PROVIDER | MIN | 5TH | MEDIAN | 95TH | MAX | STDEV | CV |
|---|---|---|---|---|---|---|---|
| Amazon | 2505 | 2508 | 2511 | 3009 | 3009 | 168 | 7% |
| Azure | 1019 | 1098 | 1102 | 1106 | 1109 | 12 | 1% |
| Google | 8046 | 8046 | 8046 | 8047 | 8047 | 0 | 0% |

Chart 8C.3: Write Block Disk Performance Type 1 (Large VM)

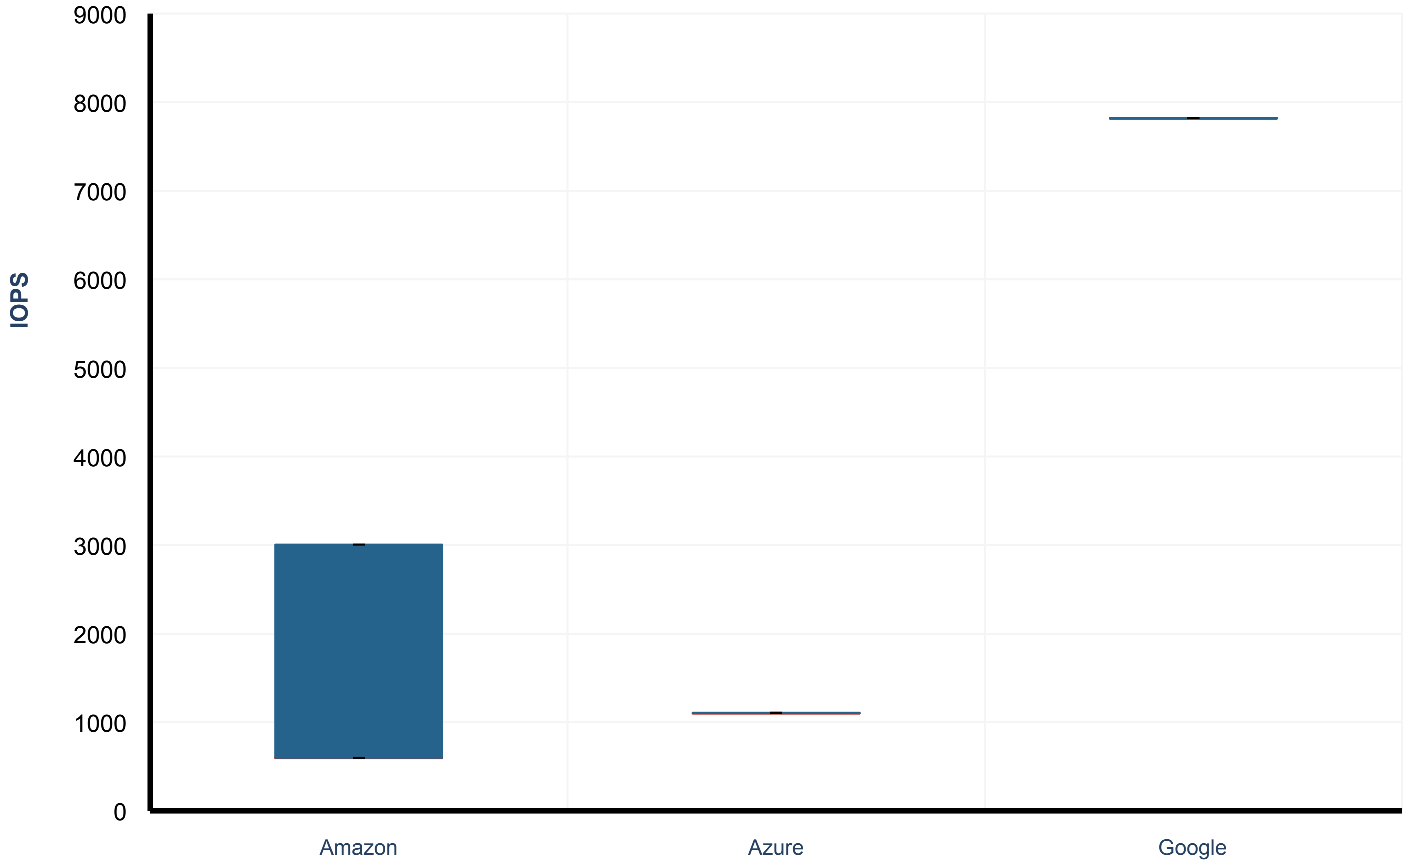


Table 8C.3: Write Block Disk Performance Type 1 (Large VMs)

| PROVIDER | MIN | 5TH | MEDIAN | 95TH | MAX | STDEV | CV |
|---|---|---|---|---|---|---|---|
| Amazon | 600 | 600 | 601 | 3006 | 3006 | 811 | 89% |
| Azure | 1102 | 1103 | 1106 | 1109 | 1109 | 2 | 0% |
| Google | 7820 | 7820 | 7821 | 7823 | 7824 | 1 | 0% |

Chart 8C.4: Read Block Disk Performance Type 2 (Large VM)

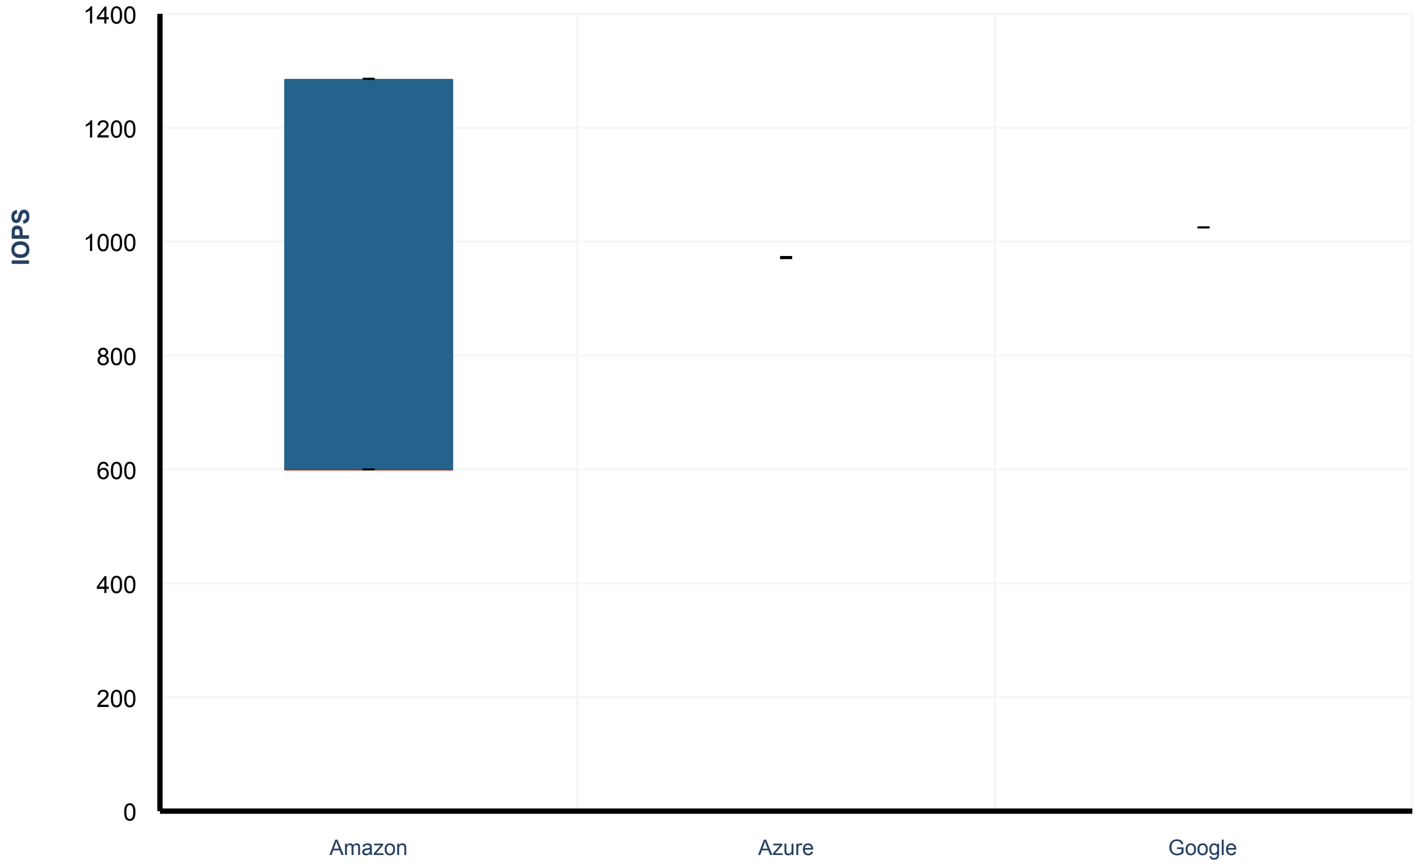


Table 8C.4: Read Block Disk Performance Type 2 (Large VMs)

| PROVIDER | MIN | 5TH | MEDIAN | 95TH | MAX | STDEV | CV |
|---|---|---|---|---|---|---|---|
| Amazon | 600 | 600 | 601 | 1284 | 1286 | 210 | 31% |
| Azure | 971 | 972 | 972 | 972 | 973 | 0 | 0% |
| Google | 1025 | 1025 | 1025 | 1025 | 1025 | 0 | 0% |

Chart 8C.5: Write Block Disk Performance Type 2 (Large VM)

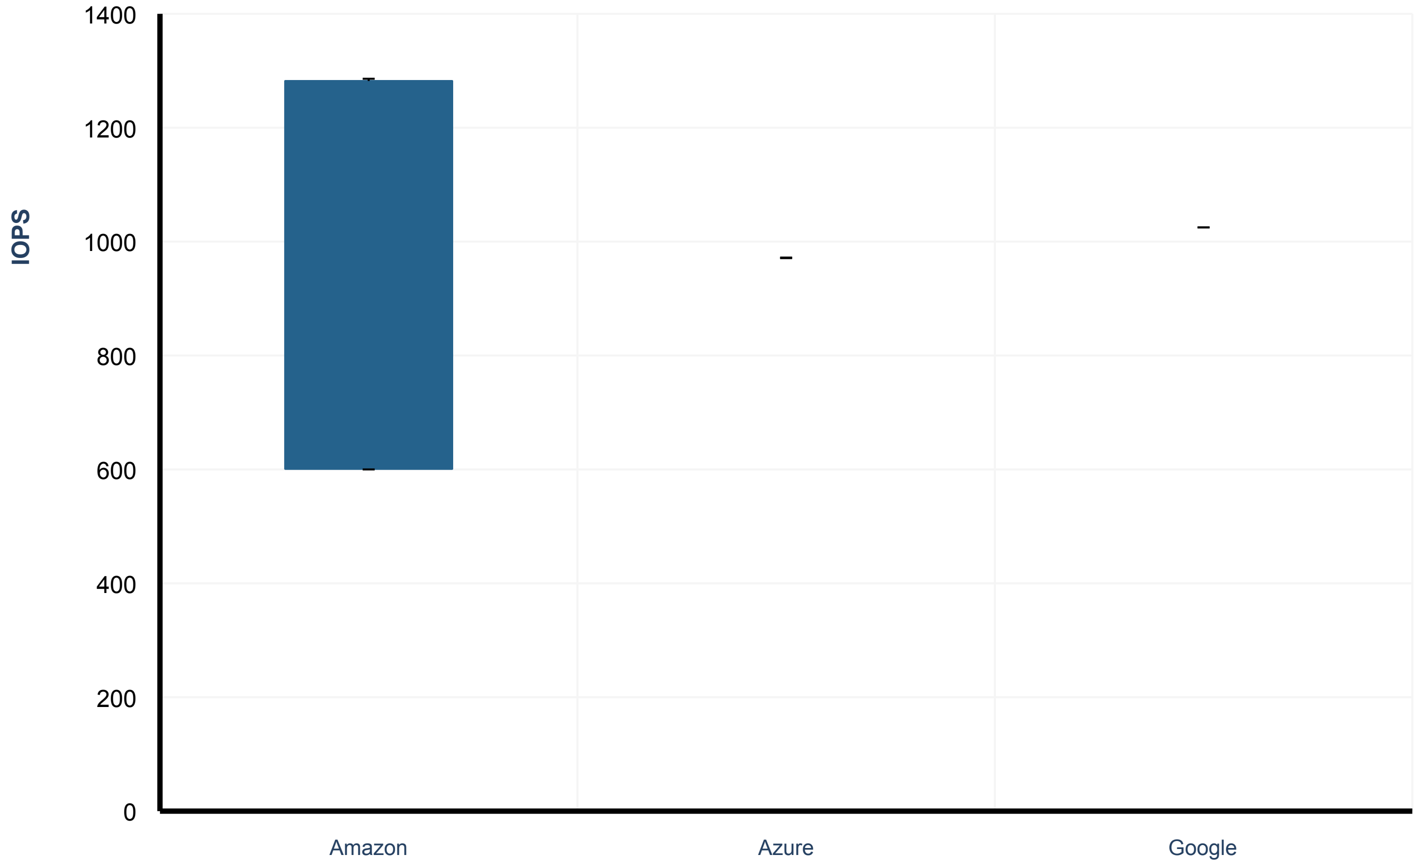


Table 8C.5: Write Block Disk Performance Type 2 (Large VMs)

| PROVIDER | MIN | 5TH | MEDIAN | 95TH | MAX | STDEV | CV |
|---|---|---|---|---|---|---|---|
| Amazon | 600 | 601 | 601 | 1282 | 1286 | 203 | 30% |
| Azure | 971 | 972 | 972 | 972 | 972 | 0 | 0% |
| Google | 1025 | 1025 | 1025 | 1025 | 1025 | 0 | 0% |

# EXTRA LARGE VMs

Chart 8D.1: VM Performance (Extra Large VMs)

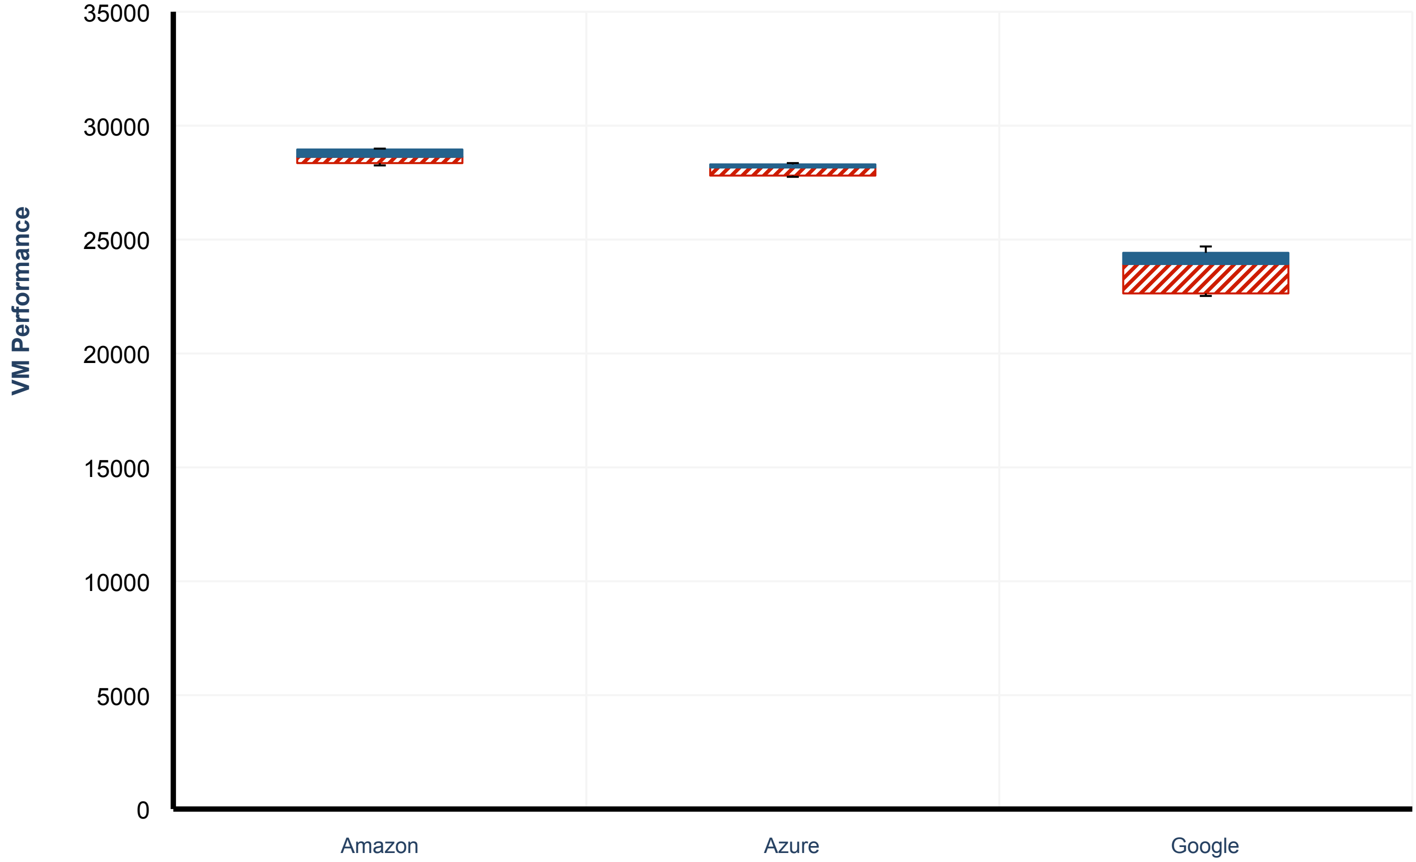


Table 8D.1: VM Performance (Extra Large VMs)

| PROVIDER | MIN | 5TH | MEDIAN | 95TH | MAX | STDEV | CV |
|---|---|---|---|---|---|---|---|
| Amazon | 28254 | 28356 | 28631 | 28954 | 28992 | 239 | 1% |
| Azure | 27753 | 27808 | 28163 | 28301 | 28357 | 176 | 1% |
| Google | 22525 | 22636 | 23926 | 24417 | 24697 | 697 | 3% |

Chart 8D.2: Read Block Disk Performance Type 1 (Extra Large VM)

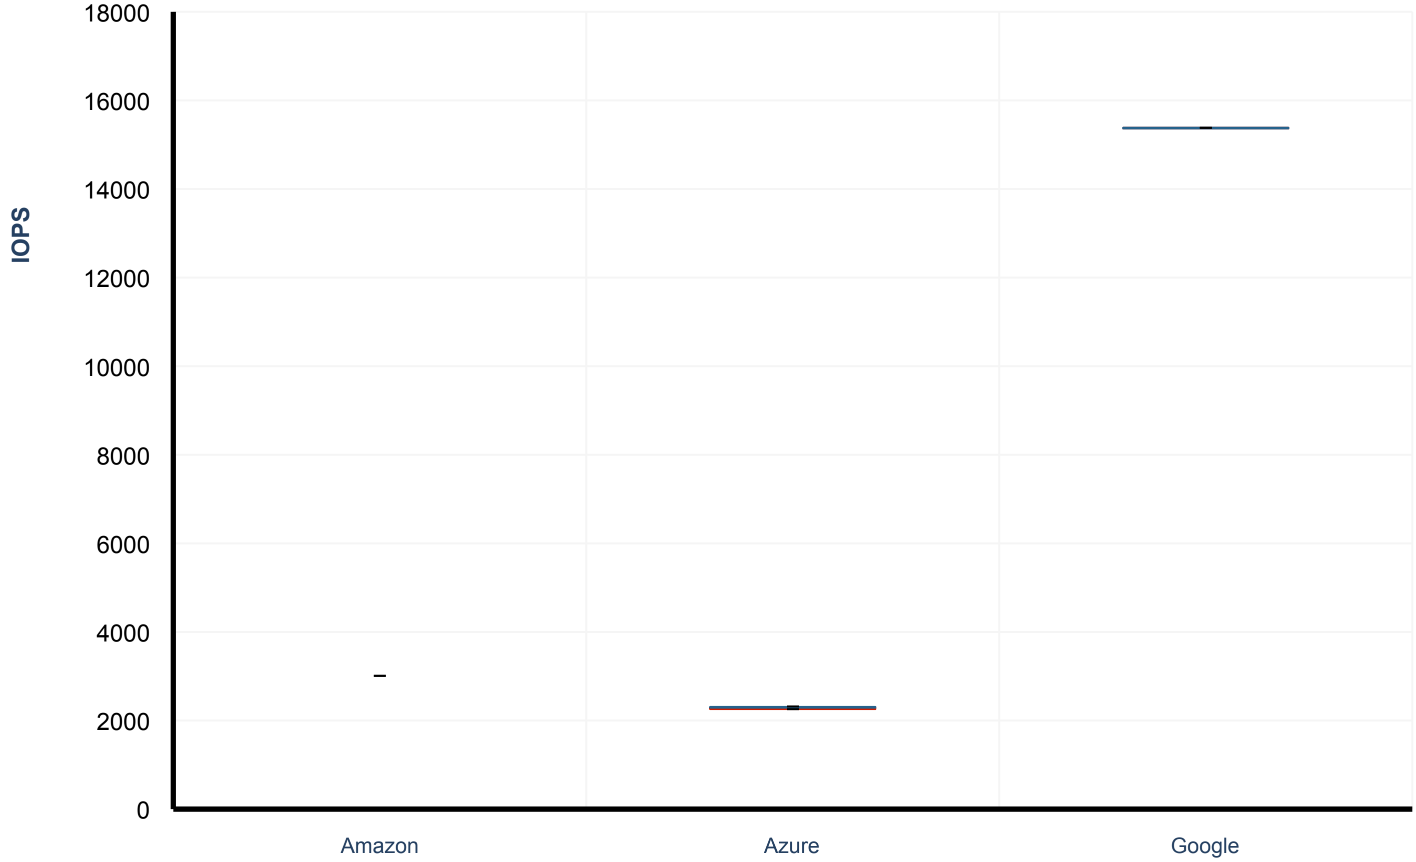


Table 8D.2: Read Block Disk Performance Type 1 (Extra Large VMs)

| PROVIDER | MIN | 5TH | MEDIAN | 95TH | MAX | STDEV | CV |
|---|---|---|---|---|---|---|---|
| Amazon | 3009 | 3009 | 3009 | 3009 | 3009 | 0 | 0% |
| Azure | 2257 | 2260 | 2291 | 2307 | 2318 | 16 | 1% |
| Google | 15380 | 15381 | 15382 | 15383 | 15383 | 1 | 0% |

Chart 8D.3: Write Block Disk Performance Type 1 (Extra Large VM)

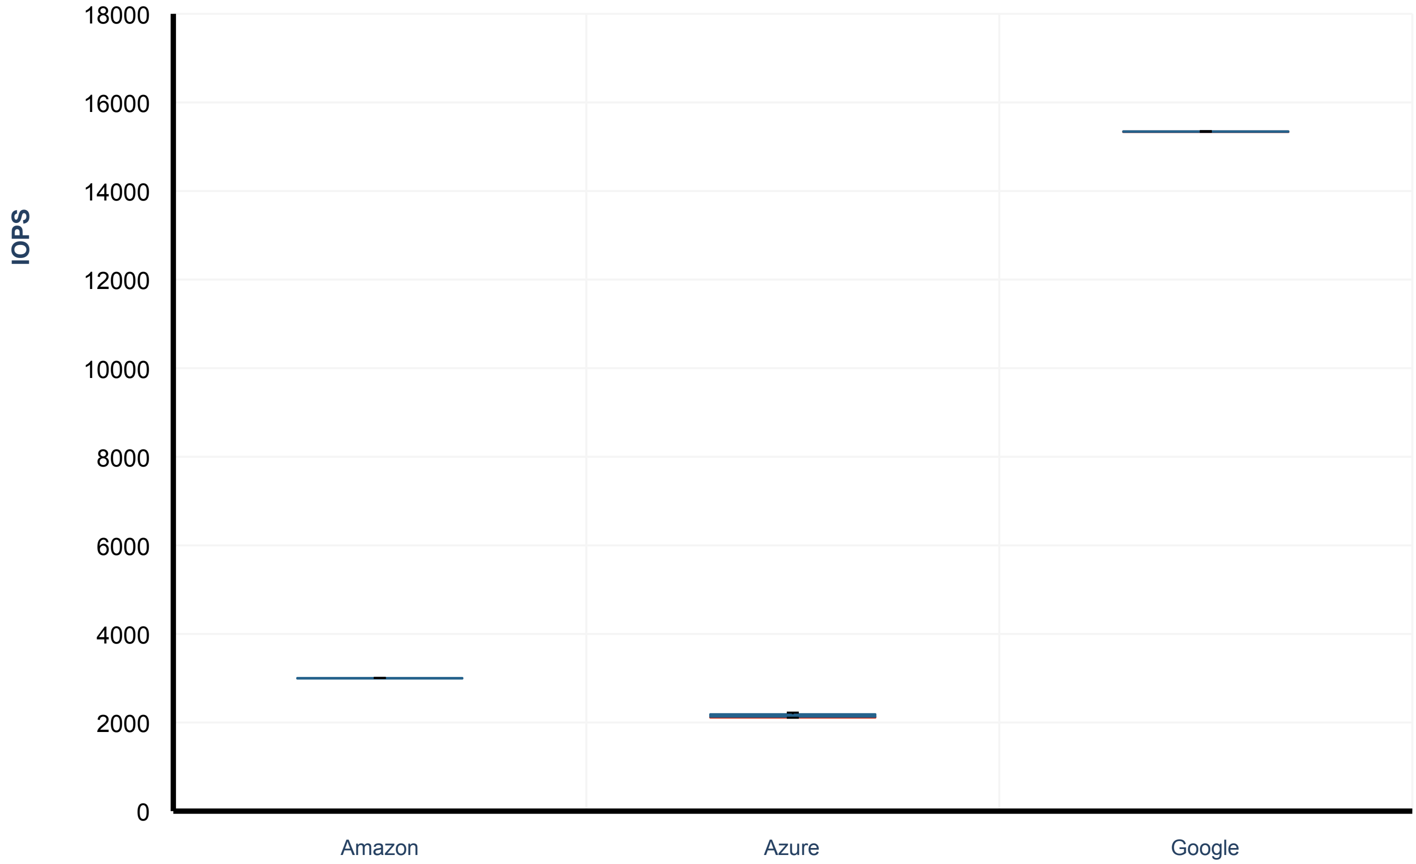


Table 8D.3: Write Block Disk Performance Type 1 (Extra Large VMs)

| PROVIDER | MIN | 5TH | MEDIAN | 95TH | MAX | STDEV | CV |
|---|---|---|---|---|---|---|---|
| Amazon | 3005 | 3006 | 3006 | 3007 | 3009 | 1 | 0% |
| Azure | 2108 | 2113 | 2130 | 2189 | 2224 | 26 | 1% |
| Google | 15337 | 15343 | 15345 | 15349 | 15352 | 2 | 0% |

Chart 8D.4: Read Block Disk Performance Type 2 (Extra Large VM)

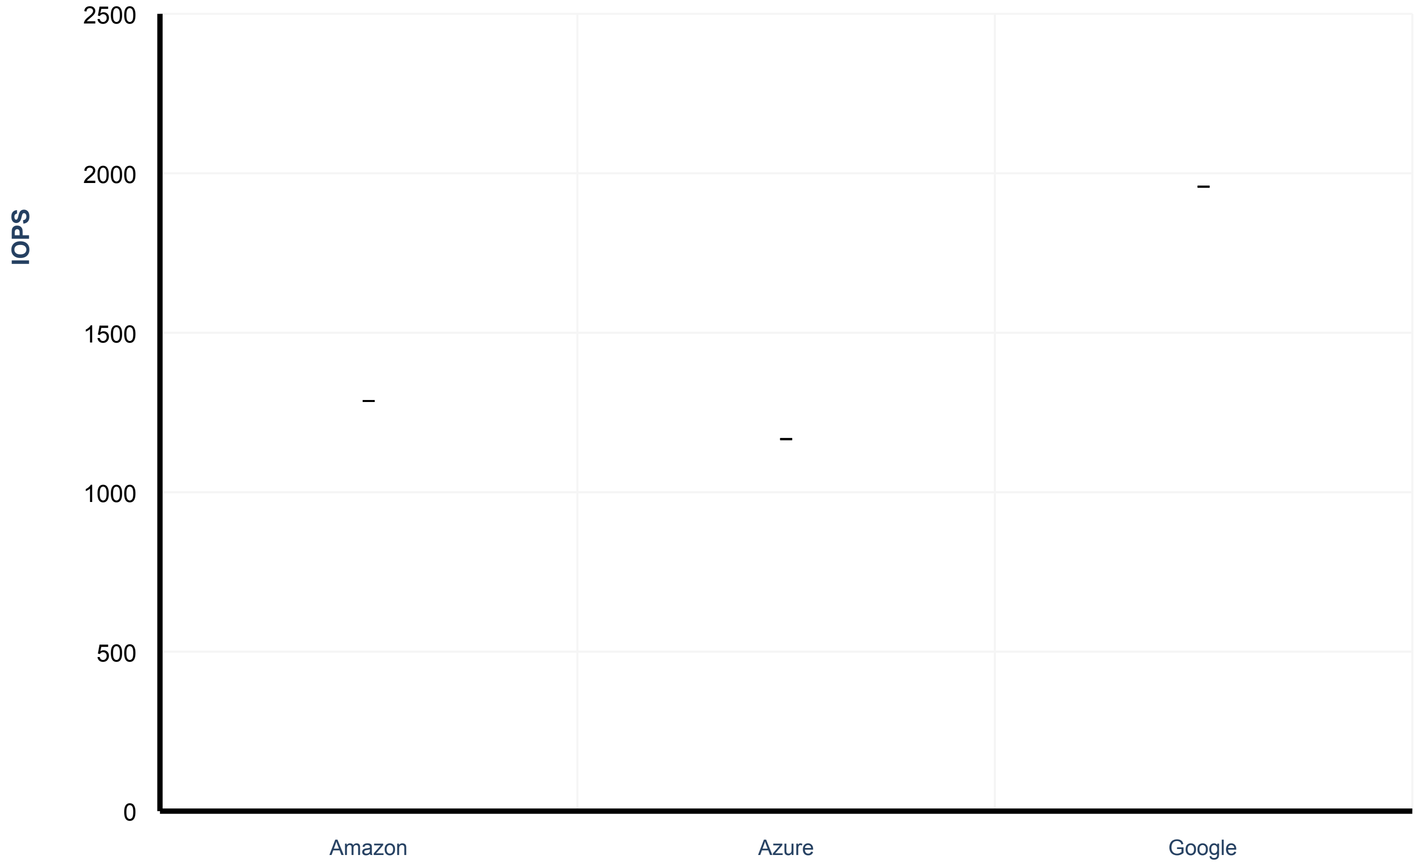


Table 8D.4: Read Block Disk Performance Type 2 (Extra Large VMs)

| PROVIDER | MIN | 5TH | MEDIAN | 95TH | MAX | STDEV | CV |
|---|---|---|---|---|---|---|---|
| Amazon | 1286 | 1286 | 1286 | 1286 | 1286 | 0 | 0% |
| Azure | 1166 | 1167 | 1167 | 1167 | 1167 | 0 | 0% |
| Google | 1958 | 1958 | 1958 | 1958 | 1958 | 0 | 0% |

Chart 8D.5: Write Block Disk Performance Type 2 (Extra Large VM)

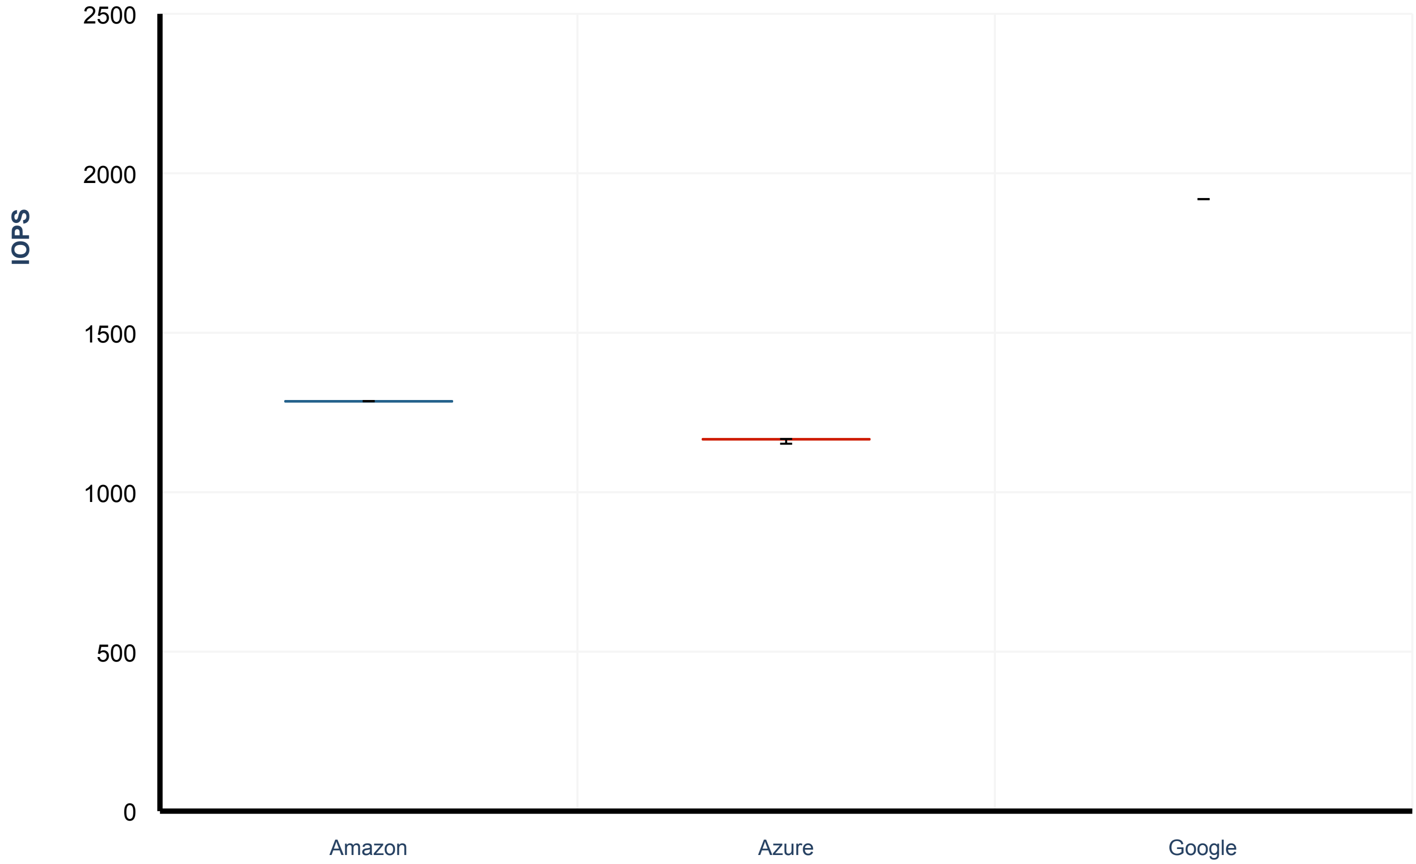


Table 8D.5: Write Block Disk Performance Type 2 (Extra Large VMs)

| PROVIDER | MIN | 5TH | MEDIAN | 95TH | MAX | STDEV | CV |
|---|---|---|---|---|---|---|---|
| Amazon | 1285 | 1285 | 1285 | 1286 | 1286 | 0 | 0% |
| Azure | 1152 | 1167 | 1167 | 1167 | 1167 | 2 | 0% |
| Google | 1919 | 1919 | 1919 | 1919 | 1919 | 0 | 0% |

# ABOUT CLOUD SPECTATOR

Cloud Spectator is a data-driven cloud consultancy specializing in cloud price-performance analysis and cloud consulting.

Cloud Spectator actively monitors many of the largest cloud Infrastructure as a Service (IaaS) providers in the world to evaluate and compare Cloud service performance (i.e., CPU, RAM, disk, internal network, external network and workloads) and pricing to achieve transparency in the cloud market.

Cloud Spectator provides full spectrum cloud consulting services including strategy and planning, architecture and technology selection, deployment and implementation, as well as cloud migration services. In addition, Cloud Spectator also helps cloud providers understand their market position within a competitive landscape.

The firm was founded in early 2011 and is located in Boston, MA.

For questions about this report, to request a custom report, or if you have general inquiries about our products and services, please contact Cloud Spectator (www.cloudspectator.com) at +1 (617) 300-0711 or contact@cloudspectator.com.

CLOUD SPECTATOR

www.cloudspectator.com


## Tested VM & Storage Configurations

| VM Size | Provider | Instance | Storage Type | vCPU | RAM | Storage | Data Center |
|---|---|---|---|---|---|---|---|
| Small | Amazon | c5.large | SSD EBS | 2 | 3.75 | 100 | US East (N. Virginia) |
| | Azure | F2s v2 | Premium Storage P10 (128GB) | 2 | 4 | 128 | US East |
| | Google | Customized | SSD Persistent Disk | 2 | 4 | 100 | US East |
| Medium | Amazon | c5.xlarge | SSD EBS | 4 | 7.5 | 150 | US East (N. Virginia) |
| | Azure | F4s v2 | Premium Storage P10 (128GB) | 4 | 8 | 128 | US East |
| | Google | Customized | SSD Persistent Disk | 4 | 8 | 150 | US East |
| Large | Amazon | c5.2xlarge | SSD EBS | 8 | 15 | 200 | US East (N. Virginia) |
| | Azure | F8s v2 | Premium Storage P15 (256GB) | 8 | 16 | 128 | US East |
| | Google | Customized | SSD Persistent Disk | 8 | 16 | 200 | US East |
| Extra Large | Amazon | c5.4xlarge | SSD EBS | 16 | 30 | 500 | US East (N. Virginia) |
| | Azure | F16s v2 | Premium Storage P20 (512GB) | 16 | 32 | 512 | US East |
| | Google | Customized | SSD Persistent Disk | 16 | 32 | 500 | US East |