Summer 2018

Cloud Performance Analysis Report

Prepared for NxtGen by Cloud Spectator LLC

Table of Contents

Executive Summary

NxtGen commissioned Cloud Spectator to evaluate the performance of virtual machines (VMs) on four different Cloud Service Providers (CSP), or providers: Amazon Web Services, Microsoft Azure, Google Compute Engine and NxtGen. Cloud Spectator tested the various VMs to evaluate the CPU performance, Random Access Memory (RAM), storage and internal network performance of each provider's VMs. The purpose of the study was to understand Cloud service performance among major Cloud providers with similarly-sized VMs using a standardized, repeatable testing methodology. Based on the analysis, NxtGen's VM performance was superior in nearly all measured VM performance dimensions. The primary drivers for NxtGen's strong performance were three factors:

1. NxtGen's competitive pricing model
2. NxtGen's excellent CPU and memory performance
3. NxtGen's strong IOPs results for both random reads and random writes.

Key findings and observations from this analysis are highlighted below, with more detailed analysis following in the body of the report.

Key Findings Observed in This Report

The following summary findings were gleaned from the testing performed by Cloud Spectator during this engagement:

vCPU and Memory Performance

Compute and memory performance were tested using the GeekBench4 test suite. The following highlights emerged from these tests:

- NxtGen VMs outperformed their rivals within each size category for multi-core testing, with some lower vCPU machines exceeding larger VMs from other providers.
- Azure Fs series VMs were the closest NxtGen competitor, achieving ~75% of NxtGen's equivalent scores.
- Other VMs, such as Azure's Dv3 series, Amazon's M5 and C5, as well as Google's n1-standard and highcpu offerings, trailed with their respective scores clustering near one another.

Storage Speed

Storage performance was tested using the FIO tool to perform a variety of tests. The 4K random read and random write results are summarized below.

- NxtGen VMs excelled at both read and write speed, although its 4K read throughput was 20x greater than its 4K write performance.
- For 4K random write performance, NxtGen was able to surpass other CSPs by a large margin, with their smallest VMs outperforming the largest VMs from Azure, AWS and Google.

Internal Network Throughput

Internal network throughput was evaluated using the IPerf test suite. The highlights of these tests are as follows:

- NxtGen achieved the highest overall throughput among all providers at 4 threads, pushing 17.7Gb/s, though performance dipped at 8 threads to just above 10Gb/s. This could be due to network saturation from other servers pushing large volumes of data through shared switching and routing infrastructure.
- AWS was extremely consistent, with throughput just below 10GB/s with thread count having no bearing on speed.
- Google and Azure both demonstrated a direct relationship between thread count and throughput. Google quadrupled Azure's performance, thereby achieving the highest eight thread value at ~14.4Gb/s.

Price-Performance

Price-Performance, or the value of a given Cloud service, derived by dividing performance by price (compute/memory and storage), is summarized below:

- NxtGen provided excellent price-performance value across the board, which is a result of its lower prices combined with higher performance.
- All NxtGen VM sizes were able to achieve the highest price-performance values, not only within their size category or class, but across all providers' respective VMs.
- The other VMs trailed NxtGen's performance, largely driven by the empirical fact that larger VMs tend to demonstrate lower price-performance value, with their VM class, e.g. standard or compute-optimized.

The details of the testing setup, design and methodology along with full results, are explained in the body of the report.

Introduction

NxtGen commissioned Cloud Spectator to assess the performance of virtual machines (VMs) from four different Cloud Service Providers (CSP) or providers: Amazon Web Services (AWS), Microsoft Azure, Google Compute Engine (GCE), and NxtGen. Cloud Spectator tested the various VMs from these providers to evaluate the CPU performance, RAM, storage and internal network of each provider's VMs. The purpose of the study was to understand the VM performance between Cloud providers with similarly-sized VMs using a standardized and repeatable testing methodology. Performance information for the specified 2, 4 and 8 vCPU VMs was gathered using Geekbench4, FIO and Iperf benchmarking tools for this analysis. Each VM type was provisioned with a duplicate VM to limit sampling error. Data was then collected during 100 test iterations.

This project focused on comparison of performance data for CPU, RAM storage and internal network throughput. The CPU-Memory composite and storage scores were evaluated on their own, and then were used to calculate the price-performance value for all providers' VM offerings. The Price-performance value for each VM was calculated by dividing performance averages by monthly cost in USD, with separate scoring performed for storage read and write. This simple price-performance formula allows the comparison of VMs offered by the respective Cloud Service Providers included in this analysis.

Using this proven Cloud sampling and testing methodology, Cloud Spectator evaluated the Cloud services based on price-performance calculations, while detailing specific strengths and weakness of each provider's VMs based on the objective performance results. Given the inherent variability of Cloud services, these methods are necessary to provide reliable and comparable analyses of Cloud-based infrastructure-as-a-Service (IaaS) services.

VM Specs and Selection Methodology

Virtual machines (VM) for this engagement focused on 2, 4 and 8 vCPU VMs. They were grouped and classified as small (2vCPU), medium (4vCPU) and large (8 vCPU) VM categories. All machines were deployed with a current release of Ubuntu 16.04 from the respective providers. Local Storage was used for NxtGen VMs, while persistent block storage was used for Azure, AWS and GCE. Two VM classes were tested in addition to the small, medium and large VM categories: standard and compute-optimized. Standard VMs are targeted for general purpose workloads, and typically are configured with a 1:4 RAM to vCPU ratio. Compute-optimized VMs, on the other hand, are typically configured with a 1:2 RAM to vCPU ratio.

The VMs selected for this engagement are listed in the tables below:

Provider	VM	VM Size Group	Class	vCPU	RAM (GB)	Disk (GB)	Storage Type	Location	Hourly Price	Monthly Price	Monthly Price INR
AWS	c5.large	Small	Compute	2	4	100	EBS SSD	Singapore	$0.114	$83.54	INR 5,569.33
Azure	Standard_F2s	Small	Compute	2	4	128	No Cache Premium LRS	Singapore	$0.123	$89.84	INR 5,989.33
Google	n1-highcpu-2 Intel Skylake	Small	Compute	2	4	100	Persistent SSD	East Asia	$0.099	$72.60	INR 4,840.13
NxtGen	Compute-Small	Small	Compute	2	4	100	Local	Bangalore	$0.052	$37.80	INR 2,520.00

Table 6.1 – Small Compute VMs

Provider	VM	VM Size Group	Class	vCPU	RAM (GB)	Disk (GB)	Storage Type	Location	Hourly Price	Monthly Price	Monthly Price INR
AWS	m5.large	Small	Standard	2	8	100	EBS SSD	Singapore	$0.136	$99.60	INR 6,640.00
Azure	Standard_D2s_v3	Small	Standard	2	8	128	No Cache Premium LRS	Singapore	$0.133	$97.14	INR 6,476.00
Google	n1-standard-2 Intel Skylake	Small	Standard	2	8	100	Persistent SSD	East Asia	$0.129	$94.36	INR 6,290.40
NxtGen	Standard-Small	Small	Standard	2	8	100	Local	Bangalore	$0.079	$57.60	INR 3,840.00

Table 6.2– Small Standard VMs

Provider	VM	VM Size Group	Class	vCPU	RAM (GB)	Disk (GB)	Storage Type	Location	VM Hourly Price	Monthly Price	Monthly Price INR
AWS	c5.xlarge	Medium	Compute	4	8	500	EBS SSD	Singapore	$0.278	$203.08	INR 13,538.67
Azure	Standard_F4s	Medium	Compute	4	8	512	No Cache Premium LRS	Singapore	$0.261	$190.39	INR 12,692.67
Google	n1-highcpu-4 Intel Skylake	Medium	Compute	4	8	500	Persistent SSD	East Asia	$0.235	$171.60	INR 11,440.27
NxtGen	Compute-Medium	Medium	Compute	4	8	500	Local	Bangalore	$0.141	$102.60	INR 6,840.00

Table 6.3 – Medium Compute VMs

Provider	VM	VM Size Group	Class	vCPU	RAM (GB)	Disk (GB)	Storage Type	Location	VM Hourly Price	Monthly Price	Monthly Price INR
AWS	m5.xlarge	Medium	Standard	4	16	500	EBS SSD	Singapore	$0.322	$235.20	INR 15,680.00
Azure	Standard_D4s_v3	Medium	Standard	4	16	512	No Cache Premium LRS	Singapore	$0.280	$204.26	INR 13,617.33
Google	n1-standard-4 Intel Skylake	Medium	Standard	4	16	500	Persistent SSD	East Asia	$0.295	$215.11	INR 14,340.80
NxtGen	Standard-Medium	Medium	Standard	4	16	500	Local	Bangalore	$0.195	$142.20	INR 9,480.00

Table 6.4 – Medium Standard VMs

Provider	VM	VM Size Group	Class	vCPU	RAM (GB)	Disk (GB)	Storage Type	Location	Hourly Price	Monthly Price	Monthly Price INR
AWS	c5.2xlarge	Large	Compute	8	16	1000	EBS SSD	Singapore	$0.556	$406.16	INR 27,077.33
Azure	Standard_F8s	Large	Compute	8	16	1024	No Cache Premium LRS	Singapore	$0.518	$378.22	INR 25,214.67
Google	n1-highcpu-8 Intel Skylake	Large	Compute	8	16	1000	Persistent SSD	East Asia	$0.470	$343.21	INR 22,880.53
NxtGen	Compute-Large	Large	Compute	8	16	1000	Local	Bangalore	$0.281	$205.20	INR 13,680.00

Table 6.5 – Large Compute VMs

Provider	VM	VM Size Group	Class	vCPU	RAM (GB)	Disk (GB)	Storage Type	Location	Hourly Price	Monthly Price	Monthly Price INR
AWS	m5.2xlarge	Large	Standard	8	32	1000	EBS SSD	Singapore	$0.644	$470.40	INR 31,360.00
Azure	Standard_D8s_v3	Large	Standard	8	32	1024	No Cache Premium LRS	Singapore	$0.556	$405.96	INR 27,064.00
Google	n1-standard-8 Intel Skylake	Large	Standard	8	32	1000	Persistent SSD	East Asia	$0.589	$430.22	INR 28,681.60
NxtGen	Standard-Large	Large	Standard	8	32	1000	Local	Bangalore	$0.390	$284.40	INR 18,960.00

Table 6.6 – Large Compute VMs

The test design and methodology used in this analysis are described in the following sections.

Test Design and Methodology

The test design and methodology are described below for each of the VM performance dimensions evaluated: CPU and RAM, storage and internal network throughput. Synthetic testing was performed on the selected VMs to enable objective comparisons of performance.

Synthetic Testing: CPU & Memory

CPU and memory testing was conducted with the Geekbench4 benchmarking suite, which allows modern testing scenarios such as floating-point computations, encryption and decryption, as well as image encoding, life-science algorithms and other use cases.

Synthetic Testing: Storage

Storage results were obtained using FIO (Flexible I/O tester) using 4KB blocks and threads corresponding to vCPU count. Several thousand 60-second random iterations were conducted to compensate for the high variability often seen when stressing storage volumes. Results were gathered and represented in IOPs (input/output operations per second).

Synthetic Testing: Internal Network

Throughput was evaluated using IPerf for each provider over their private networks only. IPerf works using a server node and client node. All transport was conducted using the TCP protocol, measuring both upload and download bandwidth. Threading for transport was determined based on vCPU count of each virtual instance.

Test Design Considerations:

Testing was conducted on specific VM types for each provider. Provider VM configurations may yield different results based on underlying infrastructure, virtualization technology and settings, and other factors. Furthermore, factors such as user contention or physical hardware malfunctions can also cause suboptimal performance. Cloud Spectator therefore provisioned multiple VMs with the same configuration to better sample the underlying hardware and enabling technology, as well as improve testing accuracy and limit the effects of underlying environmental variables.

The VMs selected for this engagement were generally-available specified offerings from the various providers. Greater performance can often be attained from these providers when additional features or support services are purchased. The selected VMs used in Cloud Spectator's testing do not leverage value-added services, which helps provide data and test results that are indicative of real-world customer choices from selected CSPs.

Error Minimizing Considerations

Duplicate VMs were deployed during testing to minimize sources of error prevalent in a Cloud hosting environment. The most notable challenge is the Noisy Neighbor Effect. Testing duplicate VMs mitigates most non-specific errors that could be attributed to a singular parent instance or storage volume. By minimizing possible sources of error, more accurate and precise performance samples can be collected during testing.

Performance Summary

In the sections below, the performance results from Cloud Spectator's testing is presented. Graphs of the performance results, along with interpretive narratives, are provided for all tests.

Price-Performance Ratio

Price-performance, also known as value, compares the performance of a Cloud service to the price of that service. Thus, price-performance offers a universal metric for comparing Cloud service value. Price-performance is calculated from the average Geekbench4 multi-core score divided by the monthly price in US Dollars (USD). A higher price-performance score indicates greater value for a given VM configuration.

Generally, smaller VMs achieve higher price-performance values than larger machines, as large VMs are typically used for specific use cases and have increased cost overhead. Thus, this evaluation was limited to small, medium and large VMs, while omitting extra-large VM configurations.

CPU Performance

The compute and memory performance for the specified VMs is summarized in the following sections.

CPU Performance for all VMs

The chart below depicts the CPU multi-core performance of all VMs tested for this study. As the vCPU count increases, VMs generally produce higher scores as depicted in the graph below. Compute-optimized class VMs are shown in lighter shades of each color, while standard class VMs are the darker shade of each color.

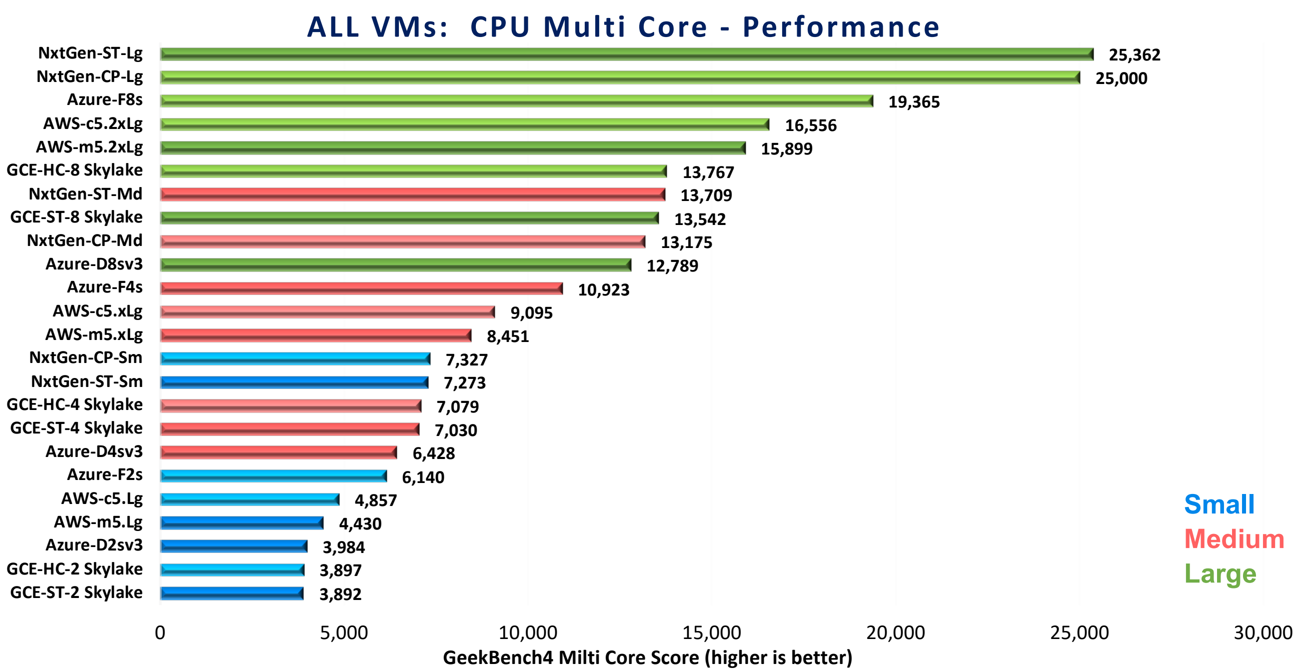

Figure 9.1 – All VMs CPU Multi-score

- NxtGen VMs, powered by current Intel Xeon Gold 6152 CPUs, posted impressive scores and outperformed all VM sizes and performance groupings (standard or compute-optimized).
- Microsoft Azure's compute-optimized Fs series machines generally follow NxtGen VMs in performance, while Azure's Dsv3 machines trail the AWS and GCE equivalents.
- Smaller NxtGen VMs outperformed VMs with double the vCPU count from competitors in many cases.

CPU Multi-core size breakdowns by VM size follow in the sections below.

CPU Multi-Core Performance for Small VMs

Below, the test results for the small VM category are presented, along with summaries of the findings. The data is broken by standard and compute-optimized small VMs.

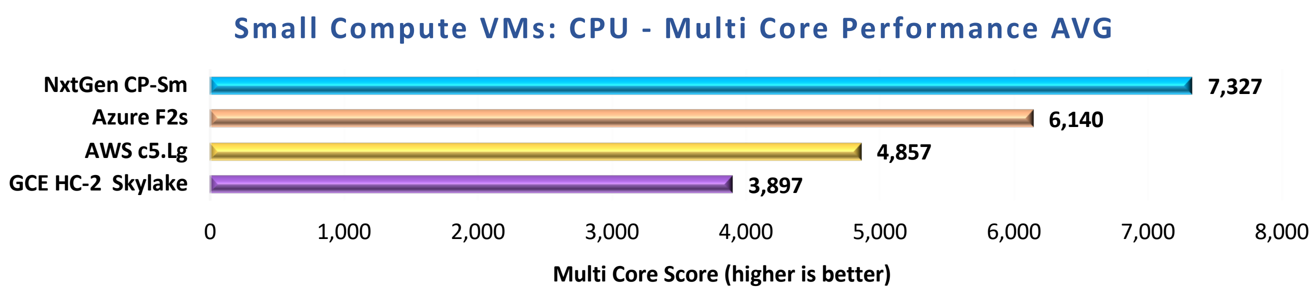

Figure 10.1 – Small Compute VMs CPU Multi-score

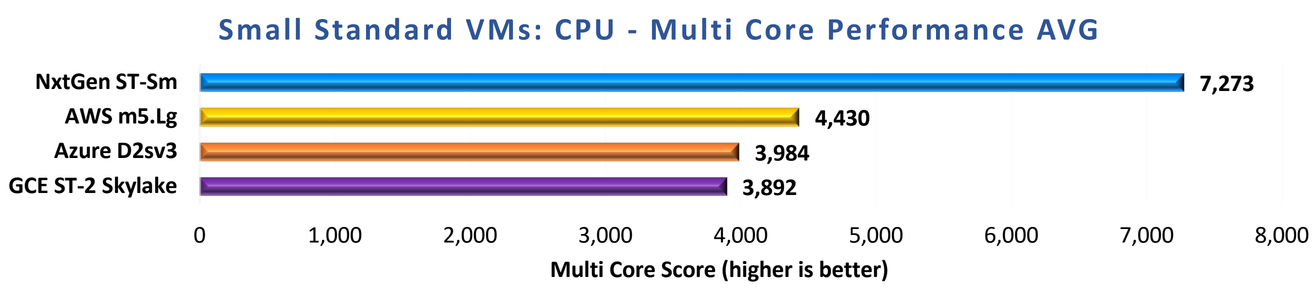

Figure 10.2 – Small Standard VMs CPU Multi-score

- NxtGen compute-optimized and standard VMs demonstrated similar performance, achieving the same CPU performance albeit with different RAM allocations. Both compute-optimized and standard varieties standout among all VMs.
- Azure's compute-optimized F2s VMs posted competitive scores, achieving 83% of NxtGen's compute-optimized small VMs. For standard small VMs, the Azure D2sv3 was able to achieve just over 50% of the CPU performance of NxtGen's standard medium VM offering.
- AWS performed better than GCE in both the compute-optimized and standard groups, though the c5.large showed ~10% greater performance than the M5.
- GCE 2vCPU VMs had similar scores in both the standard and compute-optimized categories.

The medium VM scores and analysis are presented in next section.

CPU Multi-Core Performance for Medium VMs

Medium VMs exhibited a similar performance distribution to the small VMs. The results and summary narratives are detailed below.

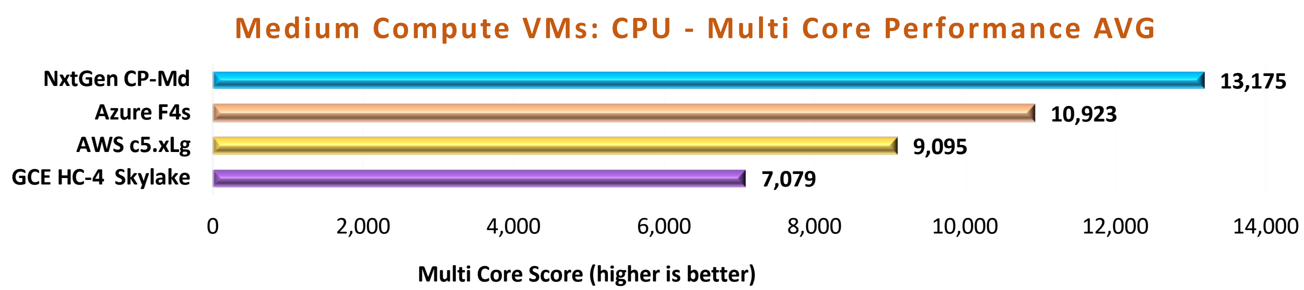

Figure 11.1 – Medium Compute VMs CPU Multi-score

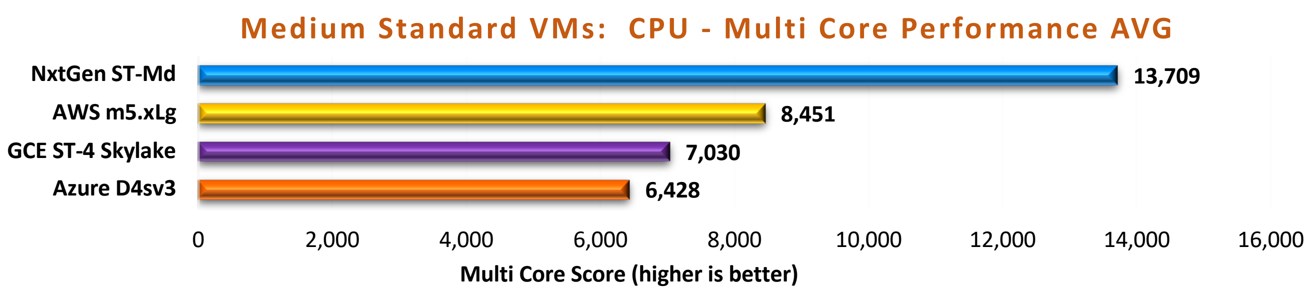

Figure 11.2 – Medium Standard VMs CPU Multi-score

- Both of NxtGen's medium compute-optimized and standard VMs stood out among rival offerings, showing markedly high CPU and memory performance results.
- In addition, Azure's F4s posted the second-highest compute-optimized class score, performing ~70% better than its standard D4sv3 offering.
- Amazon's c5.xlarge performed 8% better than its standard m5.xlarge counterpart, yet fell below Azure's F4s.
- Google's n1-standard-4 and n1-highcpu-4 VMs showed nearly identical CPU and Memory performance.

Large VM performance is reviewed in the next section.

CPU Multi-Core Performance for Large VMs

For large VM multi-core performance, the performance results were similar to the small and medium VMs. However, the large VM performance gaps were slightly more pronounced between NxtGen and the other providers' VMs, as illustrated below:

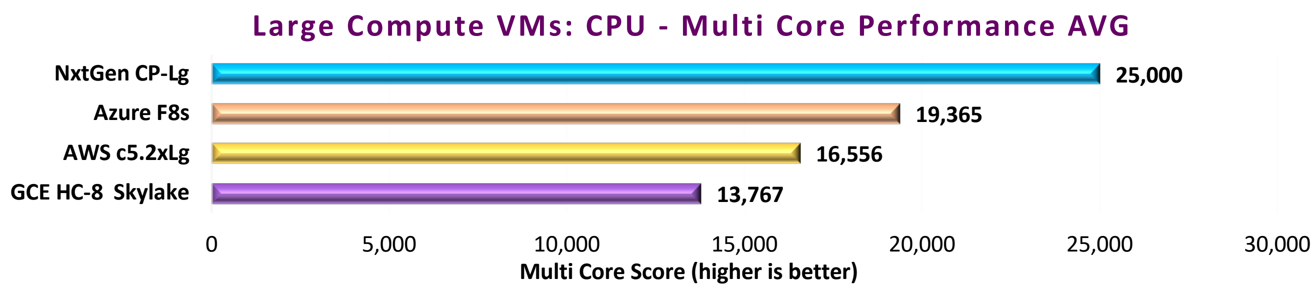

Figure 12.1 – Large Compute VMs CPU Multi-score

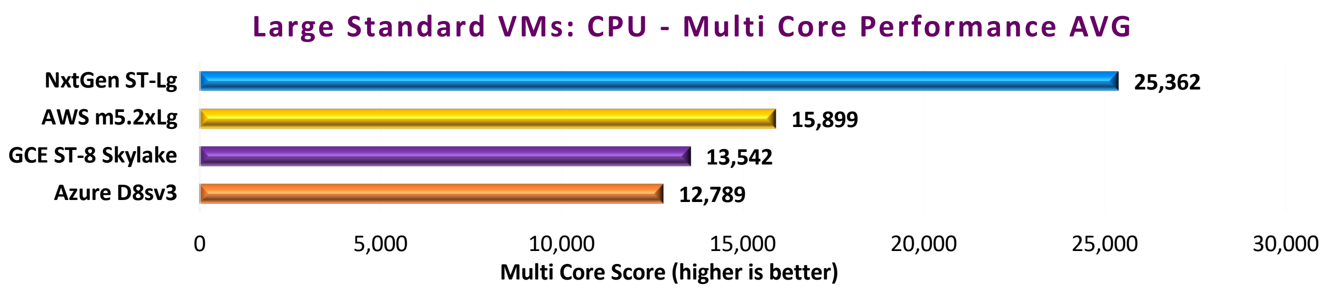

Figure 12.2 – Large Standard VMs CPU Multi-score

- NxtGen's compute-optimized large and standard large VMs scored similarly to one another, outperforming other compute-optimized and standard VMs in this size category.
- Azure's compute-optimized F8s VMs performed ~50% higher than its standard D8sv3 standard VM.
- Amazon's c5.2xlarge performed about 4% better than the AWS standard m5.2xlarge VM.
- Google's n1-standard-8 performs within 2% to its n1-highcpu-8 VM.

Next, the analysis will focus on Cloud storage performance.

Storage Performance

Storage performance results are summarized in the sections below. The testing methodology for storage ensured that all machines were tested for several thousand iterations using FIO with a block size of 4KB, Queue depth of 32 running at 1 thread per vCPU for random read and write.

Storage Performance for ALL VMs (READ)

Below, storage random read results for all providers and VM sizes are displayed. Summary details are provided below. The chart highlights compute-optimized class VMs in lighter shades of the respective colors than standard class VMs of the same size.

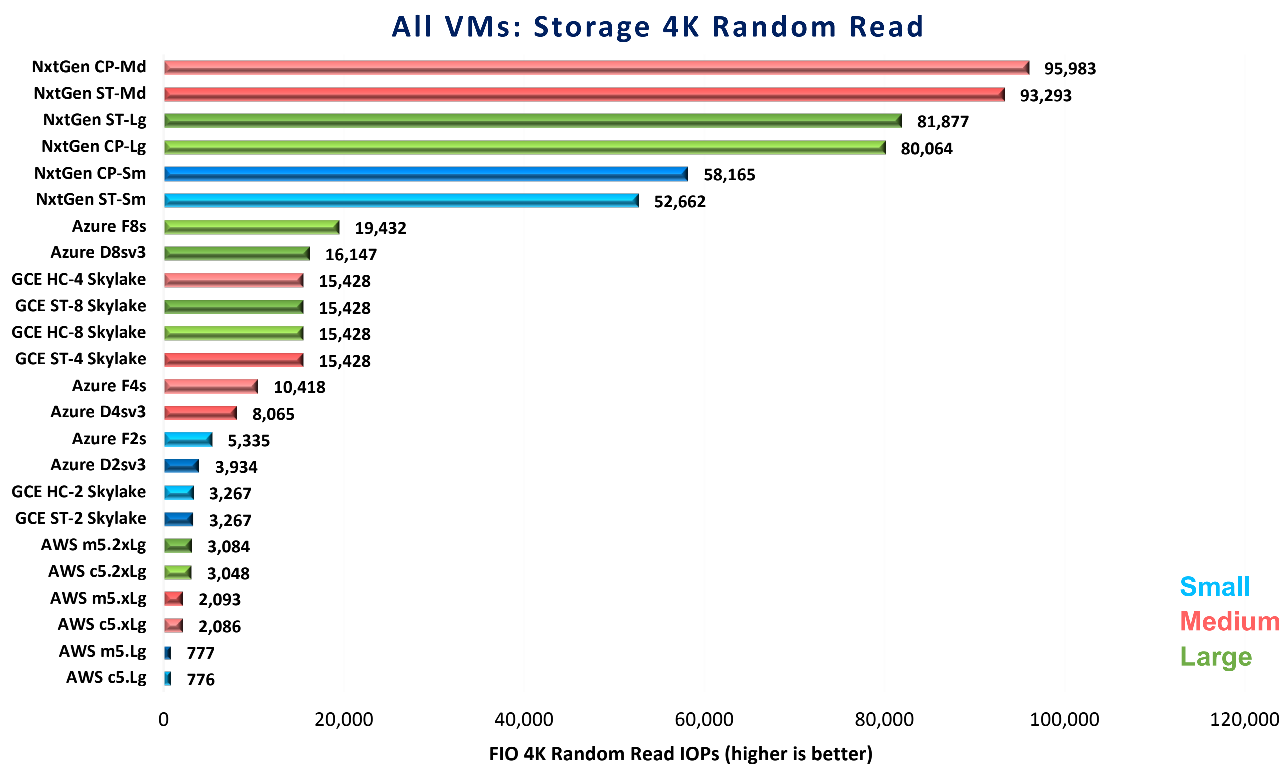

Figure 13.1 – ALL VMs 4K Read Overview

- NxtGen VMs performed extremely well against the other provider VMs, with its 4 vCPU VMs achieving the highest storage performance score based on optimized threading. In addition, even NxtGen's smallest VMs outperformed other provider VMs regardless of size, class or volume type.
- Azure posted reasonably competitive scores with their highly-redundant Premium LRS volumes. This was especially true for its largest VMs.
- Falling between Azure's 8 vCPU models and most Amazon VMs are the Google VMs. Google's compute-optimized and standard VMs, showed extreme consistency, with its 4 and 8 vCPU models obtaining the same average IOPs.
- AWS read speeds were significantly lower than the other VMs for all sizes and classes.

The test results illustrate that VM size, as described by vCPU count and RAM, plays less of a role in storage random read performance than CPU testing. Below, random write results are presented.

Storage Performance for ALL VMs (Write)

Unlike random read performance, random write performance requires a number of background tasks between operations. This includes data allocation and redundancy checking across the large storage arrays typically used in Cloud-base block storage. VM vCPU counts are less import factors in random write performance than vendor storage types. Below, storage write performance is depicted. The data compares all VM sizes and confirms that vCPU core count is a less relevant factor for storage performance in general than for CPU-Memory testing. Summary observations of these test results can be found below the chart.

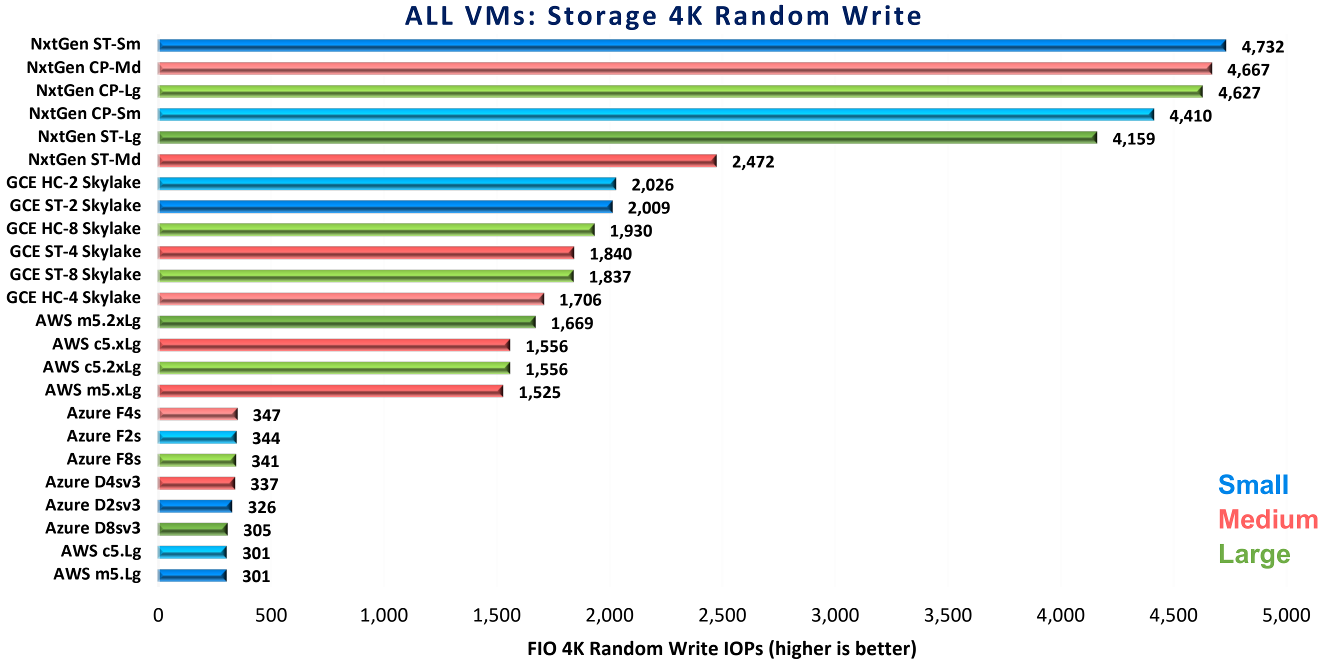

Figure 13.1 – ALL VMs 4k Write Overview

- All NxtGen VMs scored highest across the board, with its compute-optimized VMs outperforming its standard VMs by a modest margin.
- Google VMs achieved ~<50% of NxtGen's write throughput, with its smallest VMs performing slightly better than its larger counterparts.
- AWS VMs posted reasonable write scores -- roughly ~33% lower than read IOPs -- thereby showing consistency between read and write capabilities. There is no specific difference in average IOPs or EBS bursting between Amazon's 500GB and 1TB volumes, although a large difference is evident when compared to C5 and M5.large 100GB volumes.
- Azure's VMs achieved lower scores in exchange for extreme redundancy provided by their Premium LRS storage volumes. There is no correlation between machine size and speed, though scores show great consistency for all storage volumes and machines types.

Write performance is further broken down by server size in the following sections.

Storage Performance for Small VMs (Write)

The results of the five VM configurations in the small VM category are shown in the charts below, followed by summary analysis.

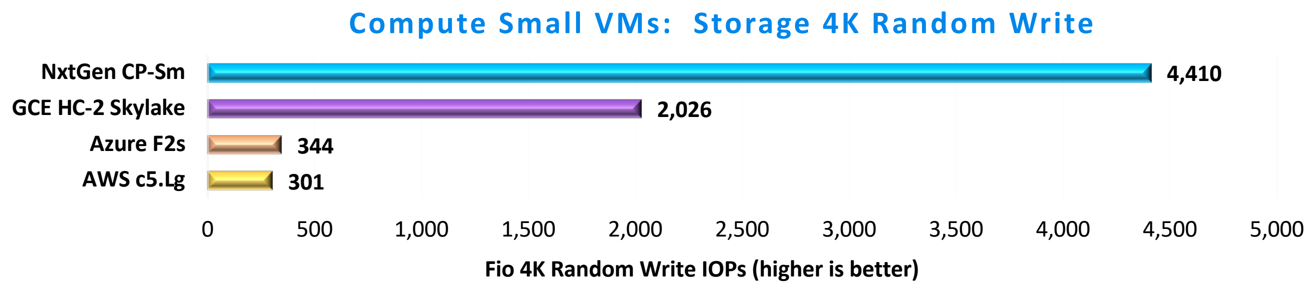

Figure 15.1 – Compute Small VMs 4K Write

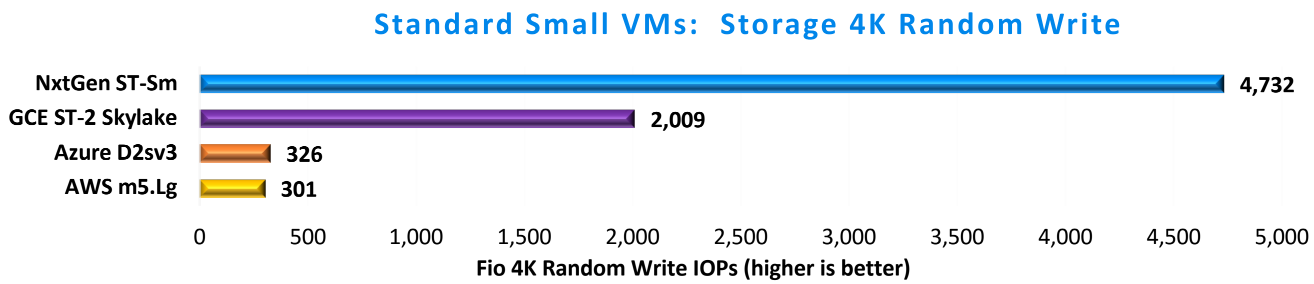

Figure 15.2 – Standard Small VMs 4K Write

- NxtGen's standard small VM was the highest-performing VM in the random write results.
- NxtGen's VMs roughly double the write speeds of Google's persistent SSD block volumes, and achieved over 10x the performance of Azure's Premium LRS and AWS' EBS volume types.
- Across all providers' offerings, neither compute-optimized nor standard class VMs demonstrated a significant difference in write performance.

NxtGen performed extremely well in the small VM category. In the following section, 4 vCPU medium VMs are compared.

Storage Performance for Medium VMs (Write)

As with the small VMs, NxtGen performed very well for random write speed. Amazon demonstrates substantial performance gains as server and volume sizes increase due to EBS burst. The Medium VM random write performance is summarized below.

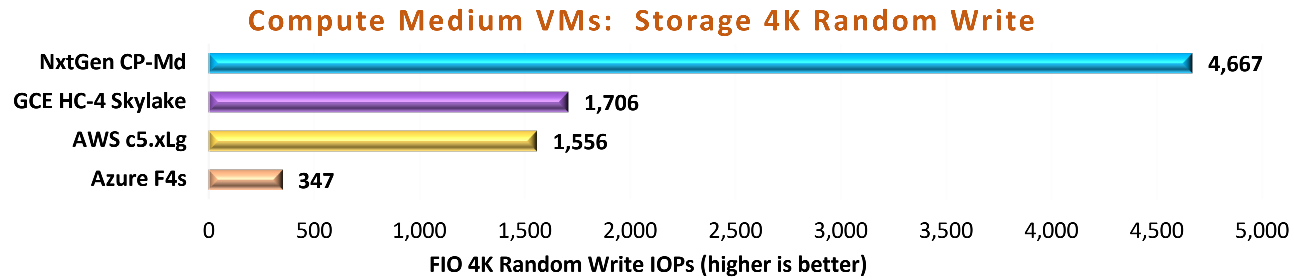

Figure 16.1 – Compute Medium VMs 4K Write

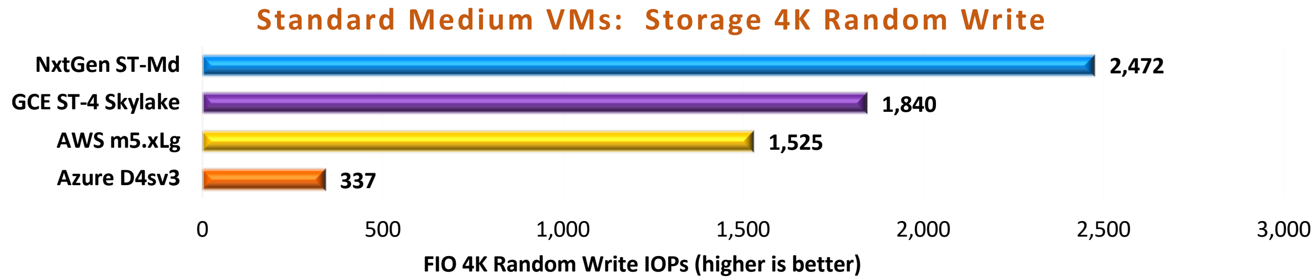

Figure 16.2 – Standard Medium VMs 4K Write

- There is a significant performance increase between NxtGen's compute-optimized and standard offerings. This trend is not observed in the small or large VM groups.
- NxtGen's compute-optimized medium VM handily outperformed AWS and GCE entries, with approximately 3x I/O's per second.
- Azure's F4s VM performed less optimally due the Premium LRS being a more read-oriented storage solution.
- The AWS VM write performance increase is pronounced due to the larger volume's increased EBS burst capability and base IOPs. Write IOPs of both c5 and m5.xlarge VMs jumped an astounding 500% over the smaller c5 and m5.large 100GB volumes.
- Google's medium VM's showed slightly lower write scores than their smaller counterparts, although its performance was similar between standard and compute-optimized VMs in this category.

NxtGen was the overall best performer for random writes in the medium VM category within both classes. Large VM random write test results are presented in the next section.

Storage Performance for Large VMs (Write)

Below, random write performance of large VM offerings is presented and discussed in the summary observations that follow.

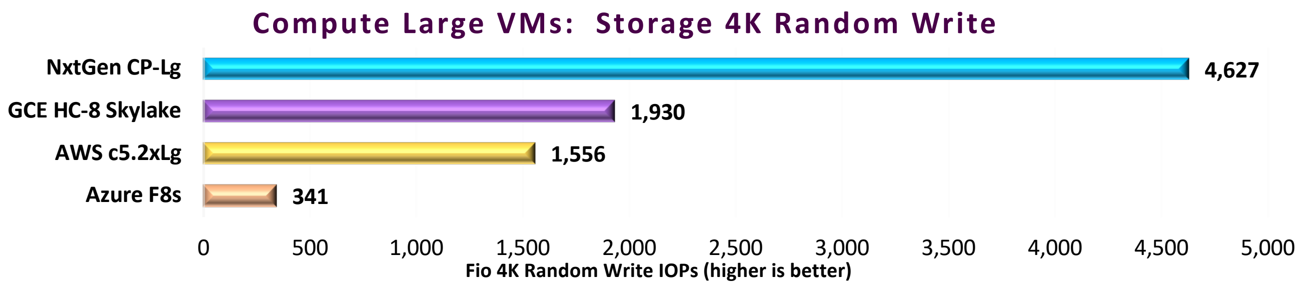

Figure 17.1 – Compute Large VMs 4K Write

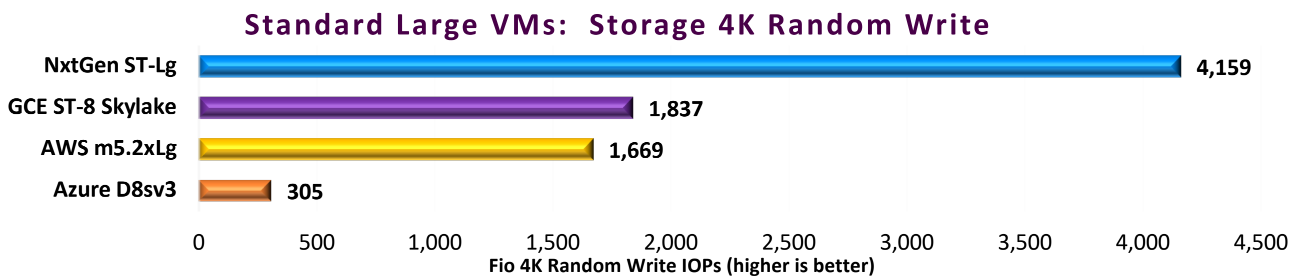

Figure 17.2 – Standard Large VMs 4K Write

- NxtGen's large VM displays 2-3x the write speed of GCE and AWS in both standard and compute-optimized classes.
- AWS burst effects are less pronounced between the medium VM 500GB and large 1TB volumes, and achieved roughly the same speed.
- GCE's persistent SSD block storage demonstrated consistent performance between medium and large classes, with a slight discernable difference between standard and compute-optimized VMs.
- Azure's 1TB Premium LRS performance, however, remains low and is on par with their other size offerings with machine type having little or no influence.

To summarize, NxtGen performed consistently well for random write performance, while also dominating the random read performance tests. As mentioned, Amazon's scores did increase with volume and VM size due to EBS burst and higher base IOPs, though write speed leveled off from 500GB to 1TB sizes.

Network Performance

Network performance data was collecting using the IPerf utility in a server-client scenario. Threads were configured to match the number of vCPUs of the serving VM, with all VMs existing on the same private network. These tests compared providers to one another rather than specific VMs. Scores in the chart below are composed of unweighted average upload and download speeds.

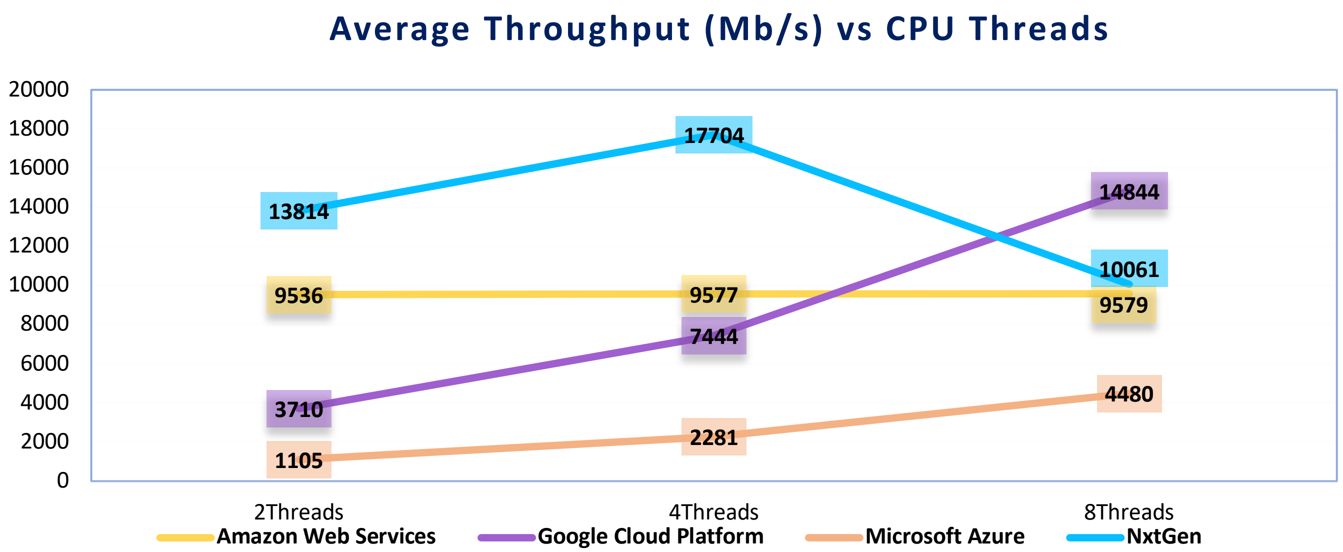

Figure 18.1 – Network Throughput

- NxtGen achieved highest network throughput, reaching an incredibly high 17.7Gb/s of throughput at 4 threads. This performance declines at 8 threads, in contrast to other providers that demonstrated stability or consistent improvement across thread count.
- Unlike Azure or Google, AWS' throughput does not scale with thread count. However, transfer speed is consistently fast, clocking in just below 10Gb/s.
- Azure implements internally-developed technologies, such as SONiC, on most networking devices within their data centers, which provides Azure with a universal network API for high-speed, dynamic routing to maintain optimal throughput. Azure's performance demonstrated a direct relationship in speed with threading from ~1Gb/s at one thread to about 4.5Gb/s at eight threads.
- Google, as with Azure, displays a relationship between threads and throughput, although to a greater degree, ranging from 3.7Gb/s throughput to ~15Gb/s. Google achieved the highest eight-thread throughput and the second-highest overall.

In the following section, CPU price-performance is analyzed.

VM CPU Performance Per USD (Price-Performance, or Value)

This section focuses on the compute and memory price-performance, or value. The values shown are linear and unweighted, using the multi-score performance scores (figures 9.1) and monthly price (figures 6.1 - 6.6). Higher scores indicate higher value for each USD spent vs competing VMs.

CPU Price-Performance ALL VMs

The chart below summarizes all VM sizes evaluated for price-performance. Breakouts by VM size category follow in the sections below. Compute-optimized class VMs are found in lighters shade of each color, while standard VMs are darker.

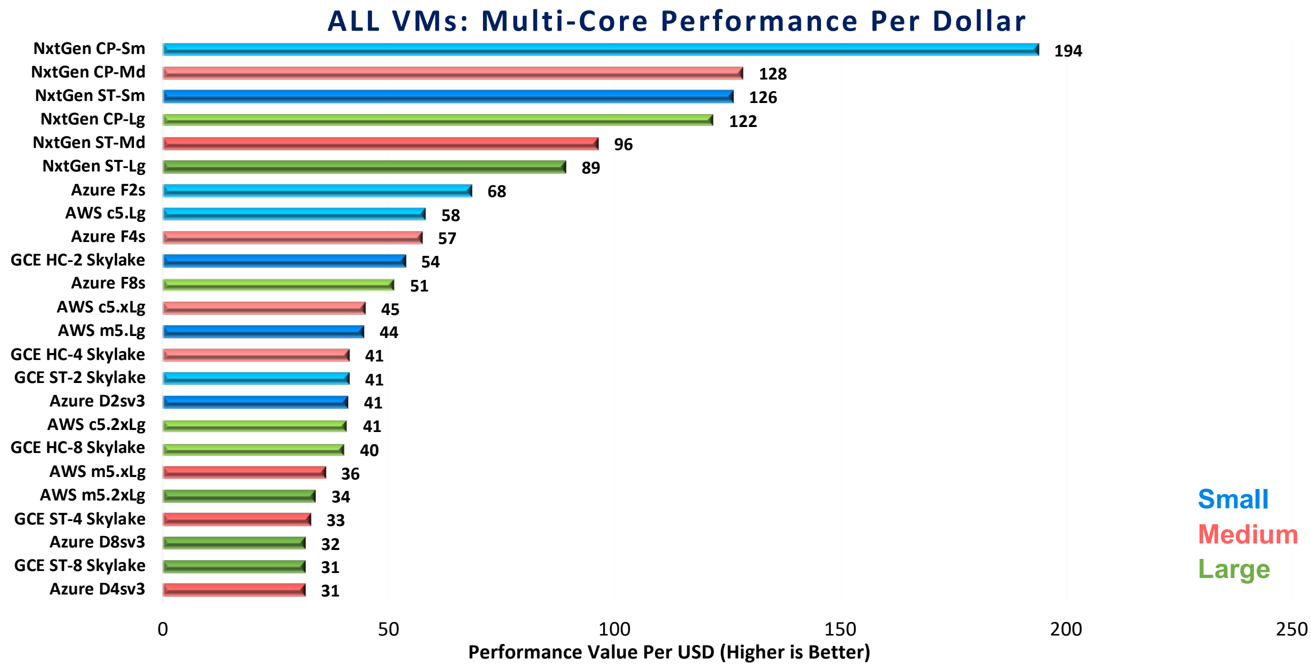

Figure 19.1 – ALL VMs Price-Performance

- NxtGen provides the highest price-performance ratios for compute-optimized and standard VM's.
- NxtGen's lower prices, coupled with higher CPU performance, make them prime general use-case choices where performance requirements may be variable or unknown.
- In general, as VM sizes increase, price-performance trends down. That said, multi-core machines may be required for compute-intensive use cases. Due to price of these machines, particularly Azure's with higher prices, use-case matching should be considered for optimal efficiency of funds spent.
- Overall, compute-optimized VM's showed greater value across all providers and sizes.

In summary, NxtGen VMs generally provide excellent value across all sizes and both groups, with smaller and compute-optimized class machines having an advantage over higher RAM standard or larger vCPU machines. In the following sections, CPU price-performance will be further detailed by class size.

CPU Price-Performance for Small VMs

Small VM price-performance results are displayed below. In addition, pricing is provided following the analysis.

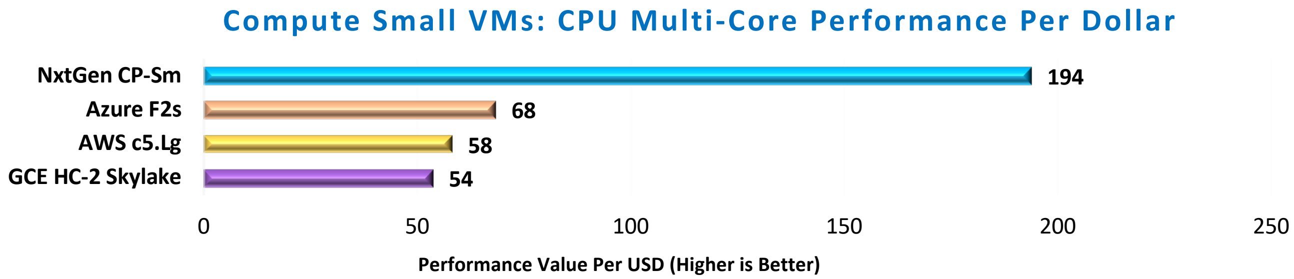

Figure 20.1 – Small Compute VMs CPU-Memory Price-Performance

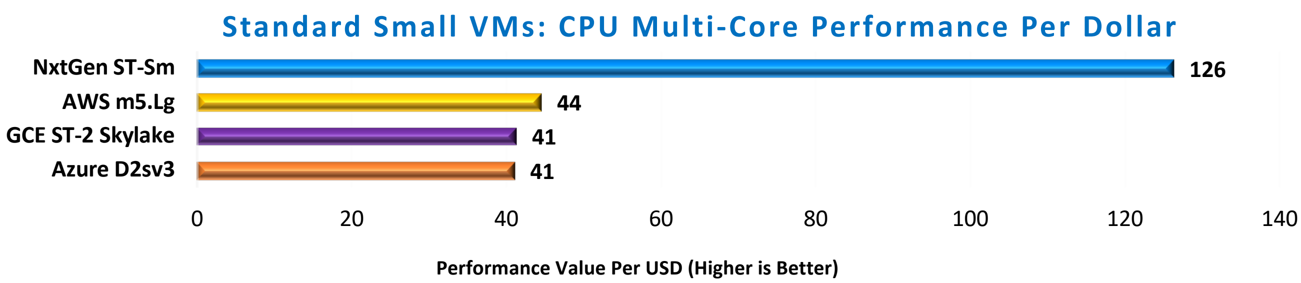

Figure 20.2 – Small Standard VMs CPU-Memory Price-Performance

- NxtGen provides ~3x computational price-performance value over other tested providers in the small VM category for both compute-optimized and standard VMs.
- Azure offers 17% greater price-performance ratio compared to AWS and GCE for the compute-optimized category.
- For the standard VM category, AWS offers less than a 10% cost benefit over Azure and GCE.

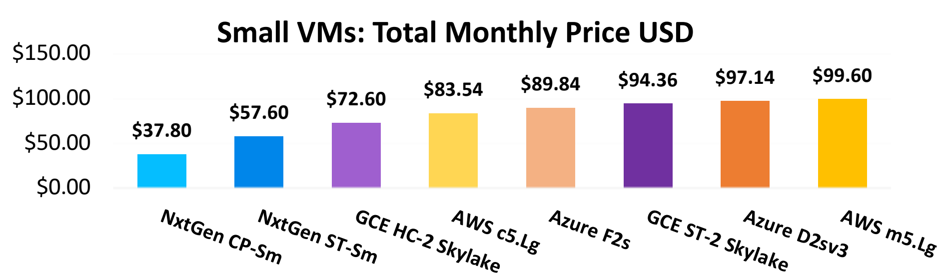

Figure 20.3 – Small VMs Total Monthly Price

NxtGen's small VM pricing and performance were found superior to comparable provider VM offerings, coming in at less than half the total monthly cost along with impressively high CPU multi-scores.

CPU Price-Performance for Medium VMs

Medium VM price-performance results are summarized in the charts below, followed by a summary analysis. In general, price-performance metrics are affected more by price differences than by performance alone, as illustrated in the small VM comparison. The summary observations for medium VM price-performance are provided below:

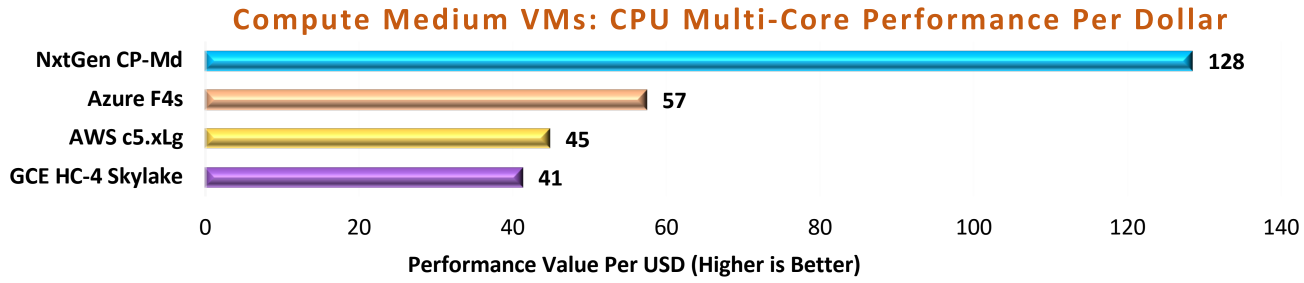

Figure 21.1 – Medium Compute VMs CPU-Memory Price-Performance

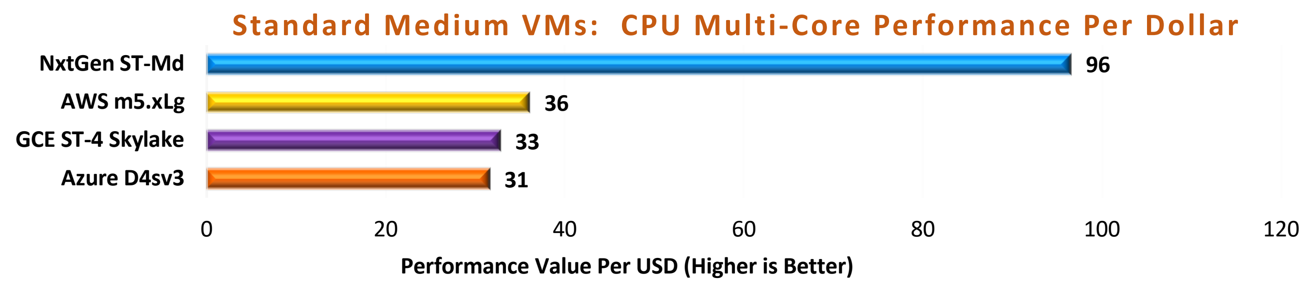

Figure 21.2 – Medium Standard VMs CPU-Memory Price-Performance

- The price-performance of the NxtGen compute-optimized and standard 4 vCPU models is 2x to 3x greater than Azure's, Amazon's or Google's equivalent VMs.
- As in the small VM compute-optimized category, Azure's medium Fs series VM achieved over 25% greater price-performance than AWS or GCE, but still lags behind NxtGen's offering.
- For the standard VM category, AWS offers an approximate 10% cost benefit over Azure, and a 15% price-performance advantage over GCE.

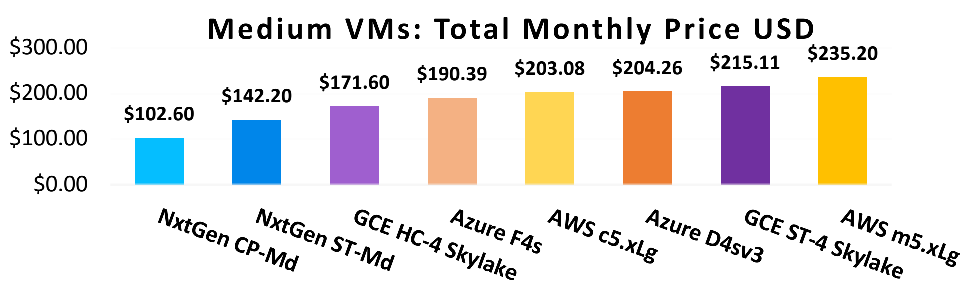

Figure 21.3 – Medium VMs Total Monthly Price

NxtGen's combination of low-cost and high performance for their medium VM's drive the greatest price-performance ratios when compared to AWS', Azure's and GCE's VMs. Moving forward, 8 vCPU standard and compute-optimized models are evaluated.

CPU Price-Performance for Large VMs

Large VMs are targeted for applications that may require more vCPUs, RAM or storage, such as database servers, moderately busy web applications or multi-role machines, all of which demand increased parallel processing. The 8 vCPU offerings from NxtGen, AWS, GCE and Azure all meet these needs. Their respective price-performance scores are compared below:

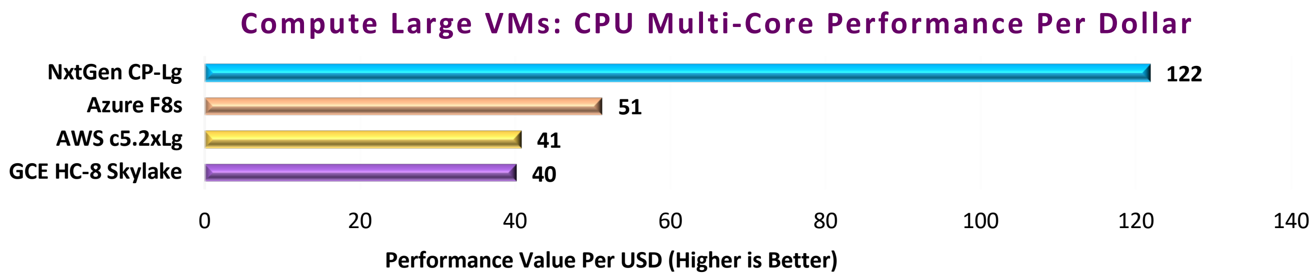

Figure 22.1 – Medium Compute VMs CPU-Memory Price-Performance

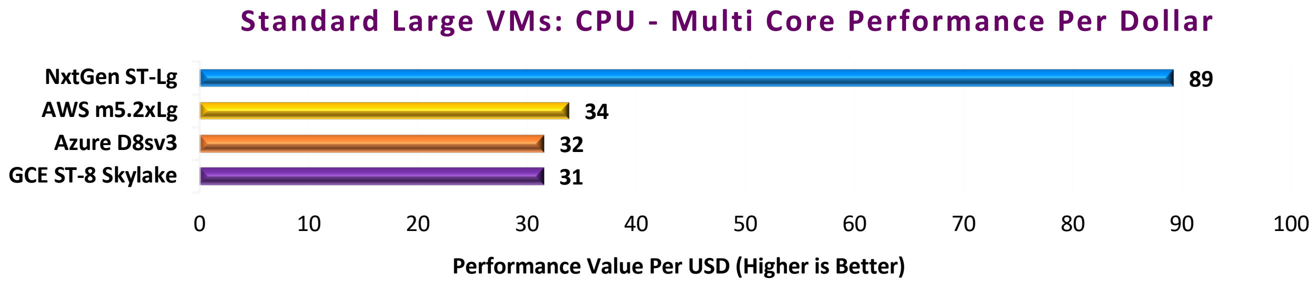

Figure 22.1 – Medium Compute VMs CPU-Memory Price-Performance

- The NxtGen 8 vCPU Machines in both compute-optimized and standard categories posted the highest price-performance scores. NxtGen's performance, in conjunction with their significantly lower costs compared to other VMs, makes them a compelling choice for use cases requiring large VMs.
- Azure's compute-optimized F8s VMs offer a price-performance ratio of 25% greater than AWS and GCE. Azure falls short of NxtGen's compute-optimized and standard VMs due primarily to its higher cost.
- AWS and GCE's compute-optimized VMs, as well as standard offerings from all other providers except NxtGen, achieve price-performance ratios within 10% of one another.

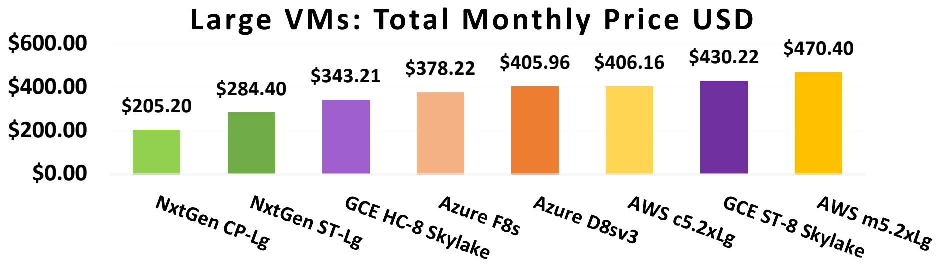

Figure 23.1 – Large VMs Total Monthly Price

CPU and memory performance have a major impact on the price-performance of a given VM, although storage performance must also be factored in. Storage performance can often bottleneck performance for various workloads based on random 4k read and write performance. Storage configurations not only influence speed, but they impact cost as well. In the following sections, storage price-performance is presented, along with summary observations.

Storage Performance Per USD (Price-Performance, or Value)

This section focuses on storage price-performance value. Storage performance can be a major bottleneck for certain applications, and often storage pricing can be a budgeting challenge for some enterprises. The values shown in the sections below are linear and unweighted, using the average performance scores for random read (Figure 13.1) and random write (Figure 14.1), as well as Monthly Prices (tables 6.1-6.6). Higher scores indicate a better price-performance value per USD spent compared to competing VMs in the same size category.

Price-Performance ALL VMs (Read)

In the chart below, price-performance is compared for the specified VMs based on random read speed. This is most useful when considering read-intensive use-cases.

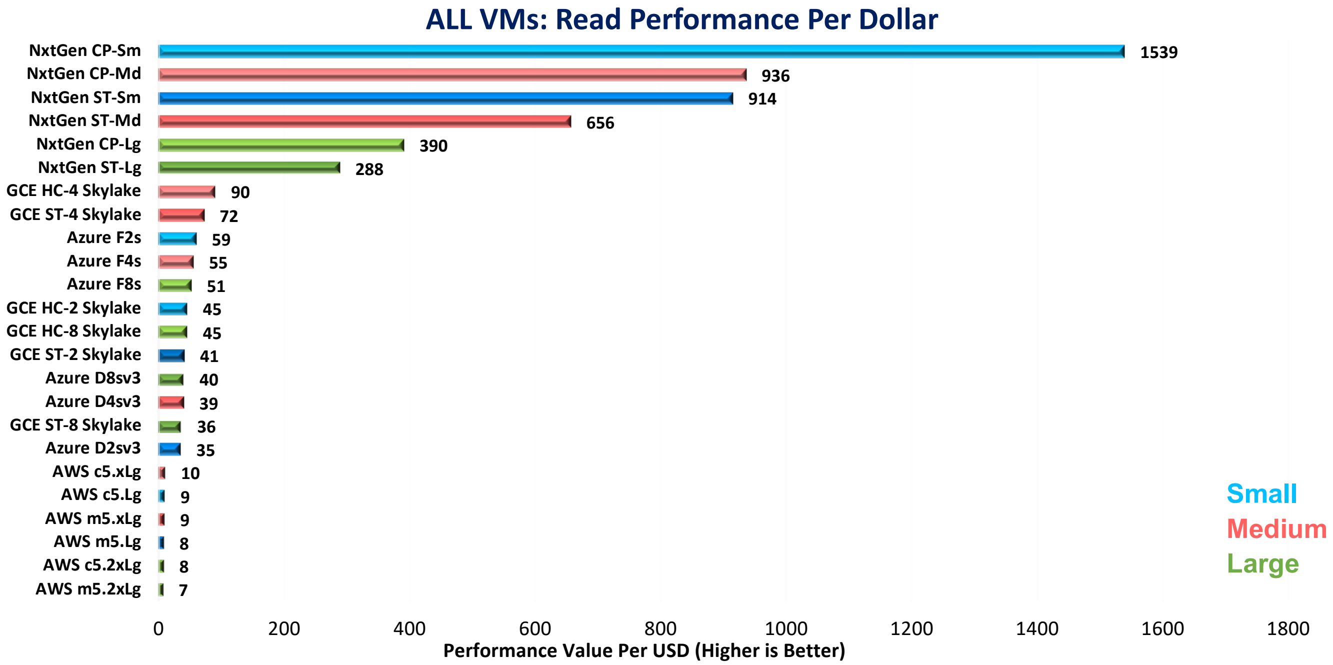

Figure 24.1 – ALL VMs Read Price-Performance

- NxtGen's VMs demonstrated the highest price-performance read values across the board, outperforming rival providers in every size and class.
- NxtGen took all top slots in price-performance value for two primary reasons:
  1. First, NxtGen's read performance was extremely strong, exceeding 90,000 IOPs in some cases, which was 4x greater than the nearest competing machines.
  2. Second, NxtGen offers a very competitive pricing model, with pricing set at one half or less of equivalent offerings from GCE, AWS or Azure.
- Azure and Google showed competitive read scores compared to one another, though NxtGen's offerings stood apart.
- AWS EBS read performance had some relation to volume size due to EBS burst, although overall price-performance in this category was lackluster compared to the other providers.

Overall, NxtGen's offerings were dominant in read speeds, as its local storage provides an edge over the highly distributed and redundant block storage offered by the other providers tested.

Price-Performance ALL VMs (Write)

Write speed, no matter the drive type or configuration, has always been a technology challenge. The following summary observations were gleaned from the analysis. Compute-optimized VMs are displayed in a lighter shade of each color, while the standard VMs are shown in the darker shade.

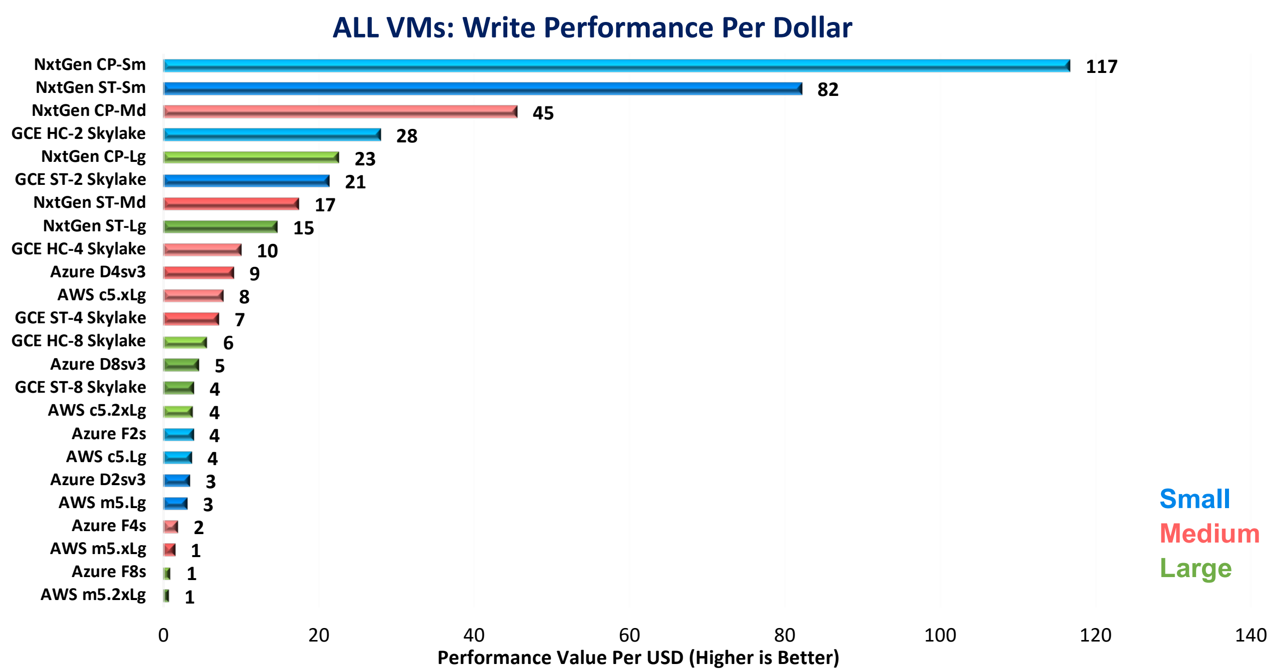

Figure 25.1 – ALL VMs Write Price-Performance

- NxtGen's compute-optimized small and standard small VMs delivered the highest price-performance value, followed by NxtGen's compute-optimized medium VM.
- NxtGen's high price-performance scores are due to a combination of high write speeds, which were often 200-300% higher than any alternative offering, along with substantially lower monthly pricing.
- In contrast to read testing, price-performance values for Azure and AWS write testing exhibited lower price-performance results.
- Google write price-performance was better than the AWS and Azure offerings, as compared to its read values.

In summary, NxtGen provides excellent storage price-performance for both read and write for all VM, configurations tested, regardless of machine classification or size.

Conclusion

For this engagement, Cloud Spectator tested Amazon Web Services, Google Compute Engine, Microsoft Azure and NxtGen VMs head-to-head across three VM size classes and two CPU:RAM categories: compute-optimized (1:2 vCPU to RAM ratio) and standard (1:4 vCPU to RAM ratio). The testing and data collection were performed by running exhaustive computational, storage and networking benchmarks on all VM configurations. From these results, the price-performance value for all VMs was determined. NxtGen provides far superior value in random read and write performance against all provider VMs tested, although NxtGen's write tests were not as decisive of an advantage.

The CPU and memory performance by all NxtGen VMs was impressive, given its use of Intel Xeon Gold 6152 CPU's. Only Azure's Fs series VMs were able to approach NxtGen's CPU performance, achieving roughly 80% in the compute-optimized class. However, NxtGen clearly stood apart in the standard VM category.

Google Compute Engine offerings were consistent in all areas. However, Google VMs were unable to achieve NxtGen's performance levels or compete with NxtGen's aggressive pricing model.

Based on Cloud Spectator's testing, Amazon and Azure's VM storage performance merits some additional analysis, as highlighted below:

1. AWS demonstrates size-dependent storage speed with its Elastic Block Storage "Burst." AWS small volumes perform at a fraction the speed of large volumes. This is due to larger volumes being allowed longer "burst" periods and higher minimum IO. Thus, AWS' storage speed is less reliable, particularly with smaller machines. If an enterprise is seeking larger volumes or machines from Amazon, this weakness can be compensated for.
2. Azure's Premium LRS, though highly redundant, is almost entirely geared toward read performance and had universally low write speeds. This performance attribute suggests that the tested Azure volume types are best suited for read-intensive use cases, as opposed to write-intensive use cases.

For network performance, NxtGen was able to achieve the highest network throughput, but its performance was not as dominant compared to GCE and AWS. In addition, NxtGen showed inconsistency across thread counts. Amazon was able to demonstrate a highly consistent performance for all thread counts at approximately 10Gb/s, while Azure and Google's performance increased.

The Cloud Spectator analysis revealed that NxtGen is highly competitive with major providers. NxtGen provides an aggressive pricing strategy, yet does not sacrifice high performance in CPU-memory testing or intensive read/write storage operations. Combined with excellent network throughput performance, NxtGen is positioned as a superior provider for generalized Cloud use-cases, and is well-suited for multi-role infrastructure with its VMs able to support workloads demanding high network throughput, such as clustered servers or load balanced web servers, regardless of budget.