

# Comparison of Windows IaaS Environments

---

Comparison of Amazon Web Services, Expedient, Microsoft, and Rackspace Public Clouds

January 5, 2015



# © TABLE OF CONTENTS

---

<b>Executive Summary</b>	<b>2</b>
vCPU Performance Summary	2
Memory Performance Summary	2
Disk Performance Summary	2
Internal Network Performance Summary	3
<b>Introduction</b>	<b>3</b>
Why Does Performance Matter	3
Key Considerations	4
<b>Test Results</b>	<b>5</b>
vCPU Performance Results	5
Memory Performance Results	7
Disk Performance Results	9
Internal Network Performance Results	11
<b>Conclusion</b>	<b>13</b>
Further Study	13
About Cloud Spectator	14
<b>Methodology</b>	<b>15</b>
Test Setup	15
Test Selection	15
About the VMs	15
Testing Notes	16

## EXECUTIVE SUMMARY

---

Cloud Spectator was commissioned by Expedient to gauge the performance of virtual machines (VMs) on several public cloud providers. Amazon Web Services, Microsoft Azure and Rackspace were included in this study for comparison. Both performance and price-performance were examined to evaluate the value of each provider's VM. The purpose of the study was to understand the disparity of performance and value (defined as price-performance) between cloud providers for similarly sized VMs. The primary components of virtual servers, vCPU, memory, disk and internal network, were evaluated for each of the provider's VMs.

### **vCPU PERFORMANCE SUMMARY**

For this study, Cloud Spectator evaluated vCPU performance by benchmarking the VMs using PassMark, a suite of benchmark tests for the Windows operating system, over the course of a 7-day test period. Pricing was examined in conjunction with the performance tests.

#### **vCPU Performance Key Findings:**

- Rackspace exhibited 16-27% higher vCPU performance
- Amazon's vCPU performance had the lowest variability
- Expedient offered the highest price-performance value for vCPU

### **MEMORY PERFORMANCE SUMMARY**

Memory was evaluated using the PassMark benchmark suite over a 7-day test period. Pricing was examined in conjunction with the performance tests.

#### **Memory Performance Key Findings:**

- Amazon exhibited 16-48% higher memory performance
- Microsoft Azure's memory performance had the lowest variability
- Amazon offered the highest price-performance value for memory

### **STORAGE PERFORMANCE SUMMARY**

Storage was evaluated using the PassMark benchmark suite over a 7-day test period. Persistent storage (offered as "block storage") was used in all storage tests. Pricing was examined in conjunction with the performance tests.

#### **Storage Performance Key Findings:**

- Expedient exhibited 0.6x-32x higher storage performance
- Amazon's disk performance had the lowest variability
- Expedient offered the highest price-performance value for storage

## INTERNAL NETWORK PERFORMANCE SUMMARY

Internal network performance was measured as the throughput between VMs within the internal environments of the cloud provider (measured using iperf). Throughput was monitored over the course of a 7-day test period. Pricing was examined in conjunction with the performance tests.

### Network Performance Key Findings:

- Expedient exhibited 39-150% higher internal network performance
- Amazon 's internal network performance had the lowest variability
- Expedient offered the highest price-performance value for internal network

## INTRODUCTION

### Processor Performance Across the IaaS Industry




Figure 1: Variance in VM performance across the industry means that simply comparing features or pricing cannot make an apples-to-apples

many processor-intensive workloads. CSPs 2-5 exhibit a closer resemblance in processor performance, but also do not offer nearly as much processing power as CSP 1. Selecting the wrong provider to house an application can result in unnecessary overspending and/or application performance problems.

Cloud Spectator examines infrastructure performance on both a hardware and an application layer. By examining the underlying hardware (in most cases, the virtualized hardware), Cloud Spectator can provide an understanding of the theoretical maximum and sustained performance of each component that comprises the server. Afterwards, running application tests helps determine which hardware bottlenecks occur first on the server for different types of applications. This proactively addresses performance concerns in a test environment before it can affect a production environment. Table 1 below lists the 3 hardware components studied in this project, and each purpose as a function in the server.

### WHY DOES PERFORMANCE MATTER?

Cloud infrastructure performance is a key consideration that should not be mistakenly overlooked; differences in performance outputs of virtual machines across the industry make it hard to compare across IaaS providers in a standardized manner by simply examining features and/or pricing. Figure 1 on the left illustrates an example of the average processor performance from a sample of Cloud Service Providers (CSPs) as studied by Cloud Spectator. CSP 1 and CSP 6 (names removed) have a marked 3x difference in processor performance, giving CSP 1 a sizable advantage in

vCPU PERFORMANCE	MEMORY PERFORMANCE	STORAGE PERFORMANCE
<p>The performance of all applications is highly dependent on the vCPU. The vCPU is responsible for the processing and orchestration of all applications.</p>	<p>While memory performance is not considered one of the key bottlenecks in performance for many applications, a subset of applications—particularly HPC and in-memory databases—is highly dependent on large sustained memory bandwidth.</p>	<p>Because most applications and all data reside on the storage, having fast storage performance is a key consideration for smooth application performance in many cases.</p>

Table 1: A summary of each of the hardware components examined in this study and their functions within the server.

## PRICE-PERFORMANCE COMPARISONS

Cloud Spectator’s price-performance calculation, the CloudSpecs Score, provides information on how much performance is realized 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 value is scaled; e.g., a Cloud Service Provider (CSP) with a score of 100 gives 4x the value of a CSP with a score of 25. The CloudSpecs Scores in this report can only be compared with equivalent configurations; e.g., a 1vCPU VM on Provider A can only be compared to a 1vCPU VM on Provider B.

The calculation of the CloudSpecs Score:

1. provider\_value = [Provider Performance Score] / [Provider Cost]
2. best\_provider\_value = max{provider\_values}
3. CSP’s CloudSpecs Score = 100\*provider\_value / best\_provider\_value

## KEY CONSIDERATIONS

When examining the results of these tests, please keep the following in mind:

- Testing was conducted on one specific VM size for each provider. Different VM configurations may yield different comparative results between the providers. Amazon, Microsoft and Rackspace offer fixed VM configurations with different resource focuses, while Expedient offers independently scalable resources.
- Prices used in the price-performance comparisons are up to date effective December 21, 2014. Pricing may change for the specified VMs after the release of this report.
- The same VM was tested for each of the selected configurations on each of the providers. Users may experience different performance characteristics across different physical hosts. Factors such as user contention (“noisy neighbor”) or malfunctions of the physical hardware can cause less than optimal performance.
- The VMs selected were the base offerings across the providers; greater performance may be obtained on certain providers by paying for additional features/services.

## TEST RESULTS

### vCPU Performance Results

The VMs selected for each provider contained four virtual cores or vCPUs. The vCPUs were tested using the PassMark suite of benchmark tests for Windows operating systems. The tests included mathematical operations such as integer and floating point calculations, as well as common CPU tasks such as compression and encryption of files. The aggregate results of the individual CPU tests were compounded together for an overall CPU Mark score. The performance, determined by the overall CPU Mark score, was then matched to the price of the VM to calculate the CloudSpecs Score for price-performance comparisons.

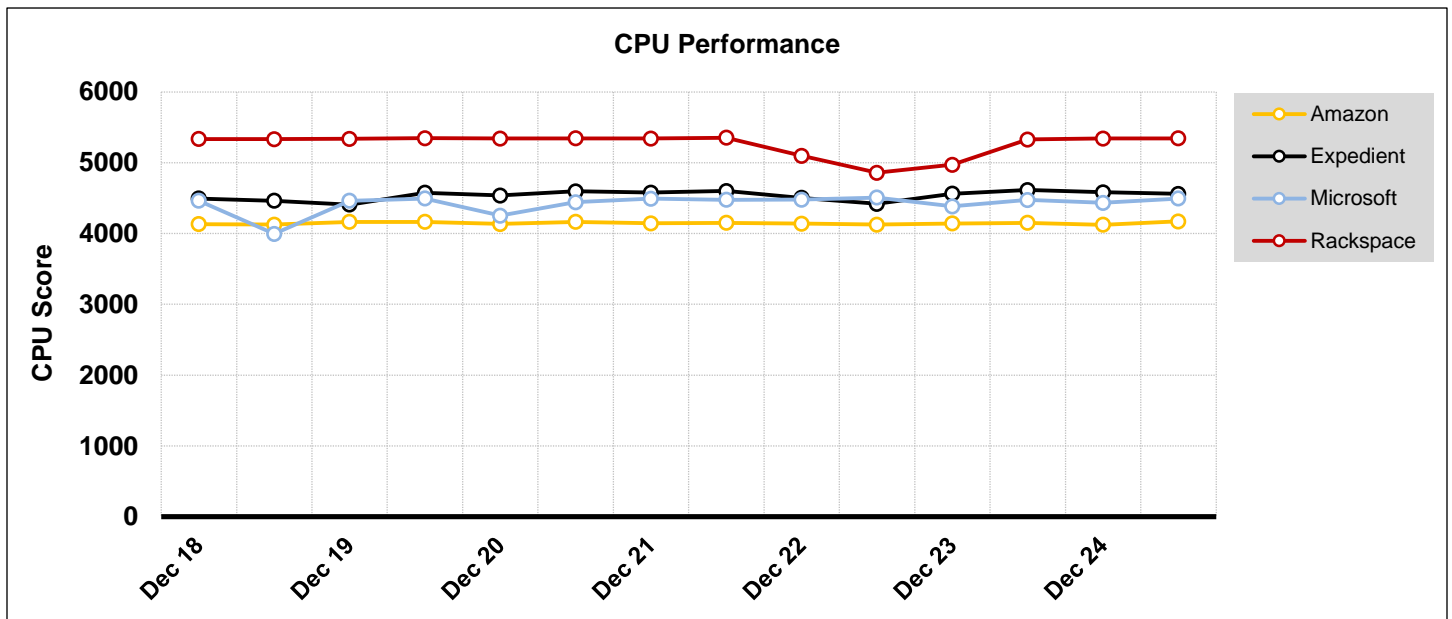


Chart 1: The chart displays the vCPU performance of the provider VMs over the entire 7 day testing period.

PROVIDER	AVERAGE	STDEV	CV	MAX	MIN
EXPEDIENT	4535	67	1%	4615	4409
AMAZON	4145	16	0%	4172	4124
MICROSOFT	4418	138	3%	4507	3996
RACKSPACE	5262	162	3%	5352	4859

### vCPU Performance (see chart above)

According to the CPU Score from the PassMark test results, Rackspace achieved between 16% to 27% higher vCPU performance than Amazon, Expedient and Microsoft. Performance variability, as determined by the coefficient of variation (CV), was stable for all four providers who exhibited CVs of 3% or lower.

The performance results were surprising because Expedient uses more advanced Intel Xeon E5-2695 v2 processors while both Amazon and Rackspace use the Intel Xeon E5-2680 v2 processors. Rackspace's performance superiority compared to the same processor utilized by Amazon and the higher-level processors on Expedient was noteworthy. Microsoft uses the Intel Xeon E5-2660 processors, which outperformed the newer processors utilized by Amazon as shown in the results. A possible rationale for the performance disparities given the aforementioned processor types was that contention for vCPU resources between users on each VM varied between provider to provider.

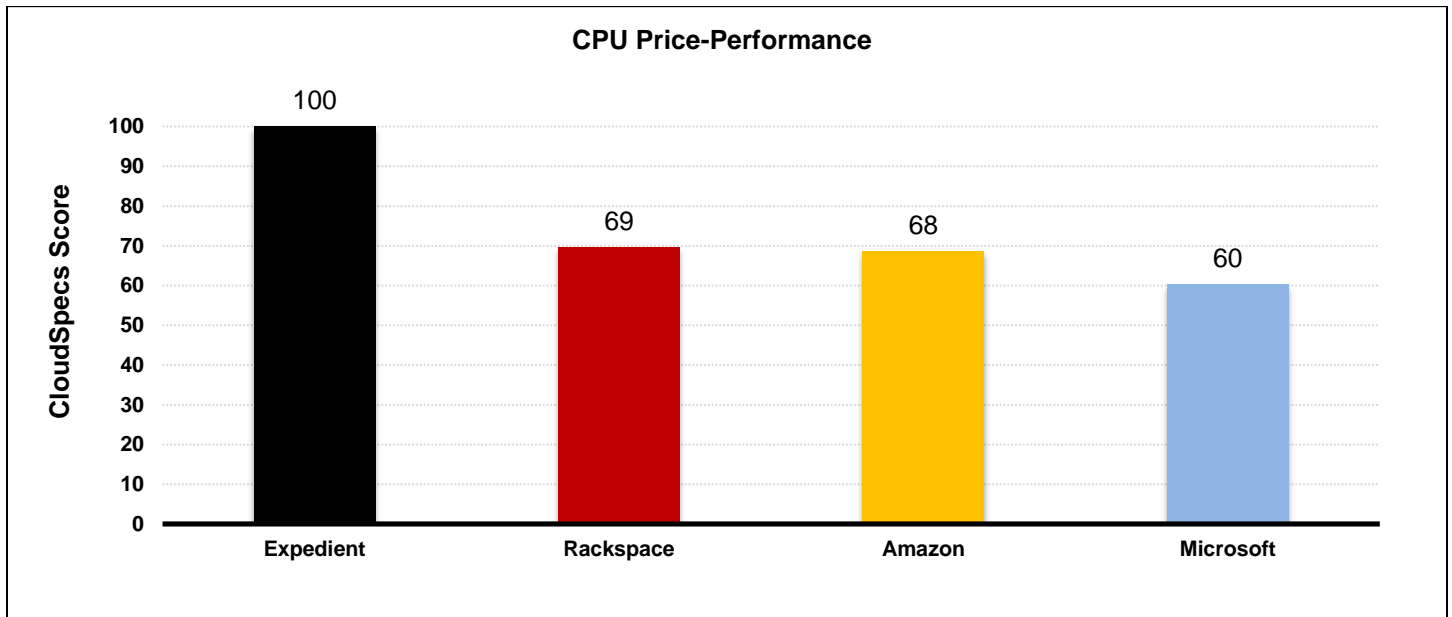


Chart 2: The chart displays the vCPU price-performance of the provider VMs over the entire 7 day testing period.

### vCPU Price-Performance (see chart above)

When considering the prices of the VMs along with performance scores in the calculated CloudSpecs Score, Expedient offered between 44% to 66% greater value than the three other providers included in this study. As evidenced by the graph above, Rackspace offered the second most value for vCPUs, trailed by Amazon then Microsoft. Despite Rackspace's higher vCPU performance (see *Chart 1*), Expedient's lower price in combination with its relative performance offered slightly greater value than Rackspace for vCPUs.

## Memory Performance Results

The VMs selected for each provider offered a relative 2:1 ratio of 2GB memory to each 1vCPU. The memory was tested using the PassMark suite of benchmark tests. The tests included common memory tasks such as reading and writing from the RAM. The aggregate results of the individual memory tests were compounded together for an overall Memory Mark score. The performance, determined by the Memory Mark score, was then matched to the price of the VM to calculate the CloudSpecs Score for price-performance comparisons.

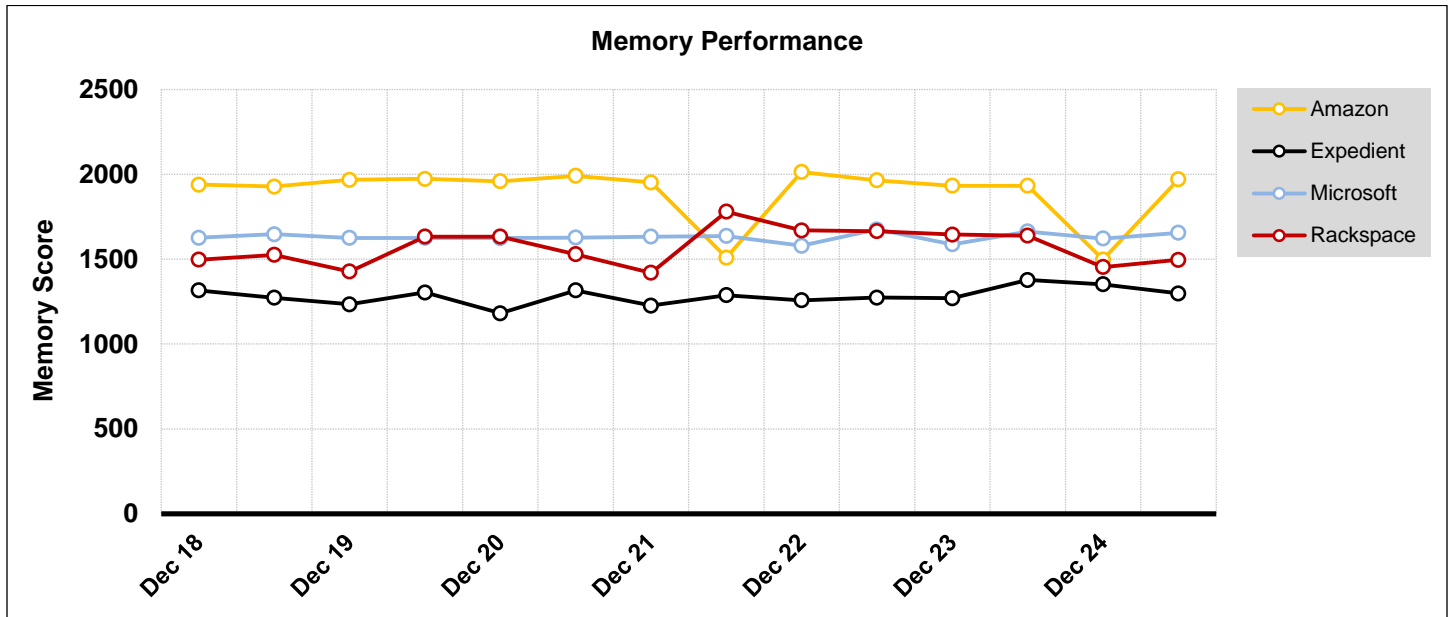


Chart 3: The chart displays the memory performance of the provider VMs over the entire 7 day testing period.

PROVIDER	AVERAGE	STDEV	CV	MAX	MIN
EXPEDIENT	1283	51	4%	1377	1181
AMAZON	1895	167	9%	2013	1498
MICROSOFT	1630	26	2%	1675	1579
RACKSPACE	1572	108	7%	1780	1420

### Memory Performance (see chart above)

According to the Memory Score from the PassMark test results, Amazon achieved between 16% to 48% higher memory performance than Expedient, Microsoft and Rackspace. Performance variability of memory was relatively stable for Expedient and Microsoft with CVs of 4% and 2% respectively. Amazon and Rackspace were slightly higher with CVs of 9% and 7% respectively.

A possible reason for the disparity in RAM performance between the providers could be due to the quantity of DIMM's populated in each physical server. Expedient populates 24 slots providing 384GB of RAM which changes the maximum frequency to 1066MHz.



Other service providers may only fill banks 1 and 2 for a maximum available amount of RAM of 256GB per server, changing the max frequency to 1866MHz.

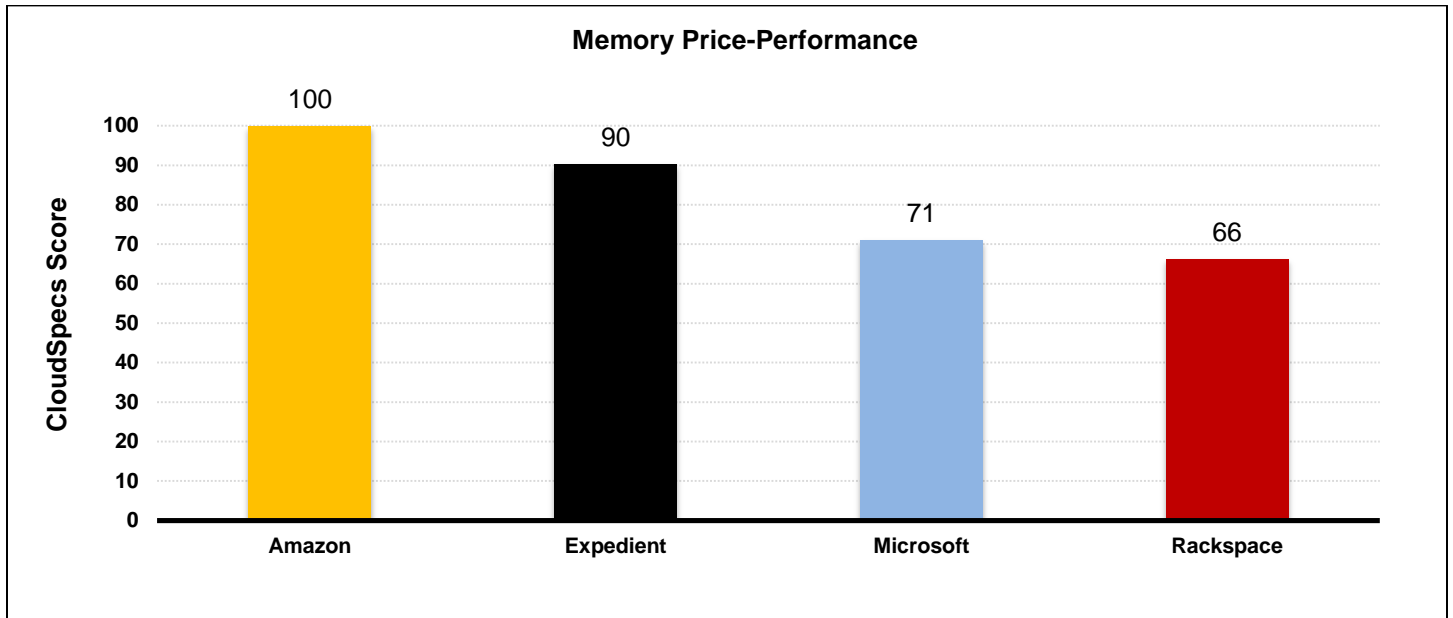


Chart 4: The chart displays the memory price-performance of the provider VMs over the entire 7 day testing period.

### Memory Price-Performance (see chart above)

The analysis of the results revealed the price-performance values for memory, wherein Amazon offered between 11% to 51% greater value than the three other providers included in this study. By offering the highest performance (see Chart 3), despite charging a higher price compared to Expedient, Amazon was able to offer a higher overall price-performance ratio, which resulted in the highest CloudSpecs Score. Nonetheless, Expedient and Amazon offer similar levels of value for memory. Microsoft and Rackspace's memory performance were offset by higher costs, resulting in relatively low price-performance values.

## Storage Performance Results

The storage performance for each provider was based upon testing of the persistent storage. Persistent storage can be included either as the primary storage or as an option in addition to ephemeral storage depending on the provider. Ephemeral storage puts the user at risk of losing data if the VM shuts off for any reason. Persistent storage maintains any stored data regardless of the life of any attached VMs. The storage was tested using the PassMark suite of benchmark tests. The tests measured the performance of sequential read, sequential write and random seek + R/W operations. The average results of the three storage tests were compounded together for an overall Disk Mark score. The performance, determined by the overall Disk Mark score, was then matched to the price of the VM to calculate the CloudSpecs Score for price-performance comparisons.




Chart 5: The chart displays the disk performance of the provider VMs over the entire 7 day testing period.

PROVIDER	AVERAGE	STDEV	CV	MAX	MIN
EXPEDIENT	8897	559	6%	9479	7506
AMAZON	916	27	3%	940	861
MICROSOFT	266	23	9%	297	208
RACKSPACE	5374	245	5%	5696	4927

### Storage Performance (see chart above)

The results of the storage tests show that Expedient achieved between 66% to 3239% higher storage performance than Amazon, Microsoft and Rackspace. Performance variability was stable for Amazon and Rackspace with CVs of 3% and 5% respectively. Microsoft's variability was slightly elevated with a CV of 9%.

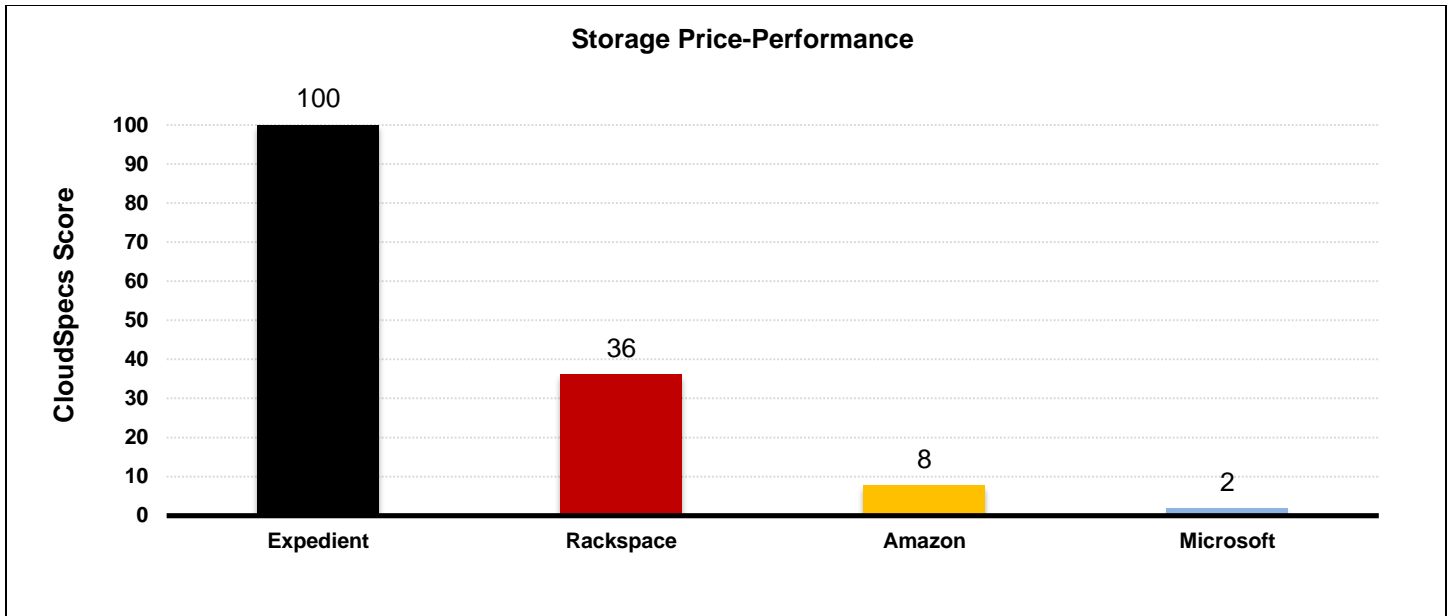


Chart 6: The chart displays the disk price-performance of the provider VMs over the entire 7 day testing period.

### Storage Price-Performance (see chart above)

As a result of testing the storage and examining prices of the VMs to find the price-performance values, the findings revealed that Expedient offered between 1.8x to 53x greater value than the three other providers included in this study. Rackspace offered the second highest value for storage, followed by Amazon then Microsoft.

## Internal Network Performance Results

The performance of the internal network on each of the providers was tested by examining the throughput between VMs within the same data center. The network was tested using Iperf which sends data bi-directionally between two servers for a set amount of time. The performance, determined by the network throughput, was then matched to the price of the VM to calculate the CloudSpecs Score for price-performance comparisons.




Chart 7: The chart displays the network performance of the provider VMs over the entire 7 day testing period.

PROVIDER	AVERAGE	STDEV	CV	MAX	MIN
EXPEDIENT	1468	182	12%	1616	887
AMAZON	1004	0	0%	1004	1003
MICROSOFT	1054	270	26%	1591	757
RACKSPACE	588	21	4%	607	531

### Internal Network Performance (see chart above)

The network throughput tests show that Expedient achieved between 39% to 150% higher network performance than Amazon, Microsoft and Rackspace. Performance variability of network was the most stable on Amazon with a CV of 0%. Expedient appears to be relatively as stable except for a drop in performance on December 23<sup>rd</sup>, possibly due to increased network contention due to the Christmas holiday. Expedient's 12% CV value is largely due to the dip in performance. Microsoft displayed variable performance throughout the testing period with a CV of 26%.




Chart 8: The chart displays the network price-performance of the provider VMs over the entire 7 day testing period.

### Internal Network Price-Performance (see chart above)

Taking into account the price of the VMs along with performance scores in the calculated CloudSpecs Score, Expedient offered between 1x to 3x greater value than the three other providers included in this study. As evidenced by the graph above, Amazon offered the second most value for network, trailed closely by Azure, offering 15% greater value than the latter provider. Although Microsoft performed 5% higher than Amazon (see *Chart 7*), the higher price offset the difference in performance and places Amazon slightly ahead for network price-performance value.

## CONCLUSION

---

The purpose of this study was to examine the differences in performance and value of the cloud offerings across various public cloud providers. Performance and price-performance of the servers were examined to determine the value offered, as determined by the price-performance ratios, albeit standardized into a general score for easier comparison (CloudSpecs Score). Performance must be examined first to determine the baseline amount of performance that will be offered by each provider. Pricing is then factored in to understand how much a user would need to be spent on each provider to obtain the given level of performance. The value of the cloud offerings is important for users to consider when determining the amount of resources needed for an application to run, and utilizing the optimal amount of spend for the selected resources. Consumers want to maximize their utility of the goods and services they purchase.

The study revealed that Expedient offered the highest overall value for the majority of services that were examined. Expedient exhibited the highest performance in storage and network tests, while being second highest in vCPU and fourth in the memory tests. Amazon offered the either the best or second best value for vCPU, memory and internal network while Rackspace was second best for storage value. Variability in performance was low for all providers for vCPU, with CVs under 4%. Variability for memory and storage varied between providers, but CVs remained under 9%. Microsoft exhibited significantly high performance variability for internal network, with a CV of 26%. Network variability was elevated for Expedient, largely due to an outlier in performance during a small window within the testing period.

Performance and value differences between providers were evidenced by this study. Providers such as Expedient have certain areas where they perform best, but also may offer the best value overall given the pricing of their services. At the same time, providers with more customizable offerings may allow you to purchase additional resources or performance services based on application needs.

### **FURTHER STUDY**

Further studies may include more VM sizes and varying configurations across the providers. Some providers throttle performance based on the amount of resources allocated. As a result, performance may scale differently on certain providers based on different vCPU to memory ratios or size of the disks.

Many providers allow you to pay for additional features and/or services that increase performance of the VMs. Users may be able to provision faster storage hardware, additional IOPS, larger network bandwidth, isolated physical environments, etc.

## **ABOUT CLOUD SPECTATOR**

Cloud Spectator is a cloud analyst agency focused on cloud Infrastructure-as-a-Service (IaaS) performance. The company actively monitors 20+ of the largest IaaS providers in the world, comparing VM performance (i.e., CPU, RAM, disk, internal network, and workloads) and pricing to achieve transparency in the cloud market. The company helps cloud providers understand their market position and helps businesses make intelligent decisions in selecting cloud providers and lowering total cost of ownership. The firm was founded in early 2011 and is located in Boston, MA.

**TESTING SETUP**

Cloud Spectator identified and provisioned equivalent sized VMs across each provider. The locations of the VMs were all in North America. The chosen benchmarking software was then installed on each VM and run twice a day for a period of 7 days.

**TESTS USED**

Test	Category	Description
PassMark	CPU	Nine separate CPU tests that are all aggregated into a final score. <ul style="list-style-type: none"> <li>Subtests include: Integer Math, Floating Point Math, Prime Numbers, Extended Instructions (SSE), Compression, Encryption, Physics, Sorting and Single Threaded.</li> </ul>
PassMark	Memory	Seven separate memory tests that are all aggregated into a final score. <ul style="list-style-type: none"> <li>Subtests include: Database Operations, Read Cached, Read Uncached, Write, Available RAM, Latency, and Threaded.</li> </ul>
PassMark	Disk	Average of Sequential Read, Sequential Write and Random Seek + R/W tests.
Iperf	Network	Transfers data bi-directionally between 2 nodes within 60 seconds, measured in megabits per second.

PassMark test details and benchmark software can be found at: <http://www.passmark.com/products/pt.htm>

**ABOUT THE VMs**

All VM sizes were selected by Cloud Spectator, standardized as closely as possible by vCPUs and memory. All VMs ran Windows Server 2012 R2 operating systems. Two VMs of identical configurations were provisioned on each provider, one VM is used for the CPU, memory and disk tests, while the second VM is used in testing the internal network performance.

SAMPLE SIZE	EXPEDIENT	AMAZON	MICROSOFT	RACKSPACE
OFFERING NAME	Level 15 Instance	m3.xlarge	D3	I/O1-15
vCPUs	4	4	4	4
MEMORY	15GB	15GB	14GB	15GB
BLOCK STORAGE	50GB	50GB	50GB	50GB
DATACENTER	Pittsburgh	US WEST (N. California)	EAST US	DFW
VM PRICE (\$/HR)	\$0.411	\$0.56	\$0.684	\$0.675*
STORAGE (\$/HR)	\$0.014	\$0.006	\$0.002	\$0.034

\*Rackspace charges a monthly Service Level fee on top of raw infrastructure costs. The lowest priced Service Level tier (Managed Infrastructure) was used in this study.



## TESTING NOTES

- Testing was conducted for a period of 7 days.
  - The testing spanned from 12/18/14 to 12/24/14.
  - Tests were conducted twice per day for each day of testing; once in the morning and once in the afternoon.
- The average performance values from the entire testing period were used in the price performance calculations as the [Provider Performance Score]
- The performance variability was measured using the coefficient of variation (standard deviation / average) of values from the entire testing period.
- Persistent storage (also known as block storage) was tested instead of ephemeral storage (also known as local storage) due to the greater dependability of the data volume as well as the common user need to retain data regardless of the on/off state of the VM.