CLOUD SPECTATOR

IaaS Generational Comparison

Amazon EC2: C3 and C4 Families

March 2015

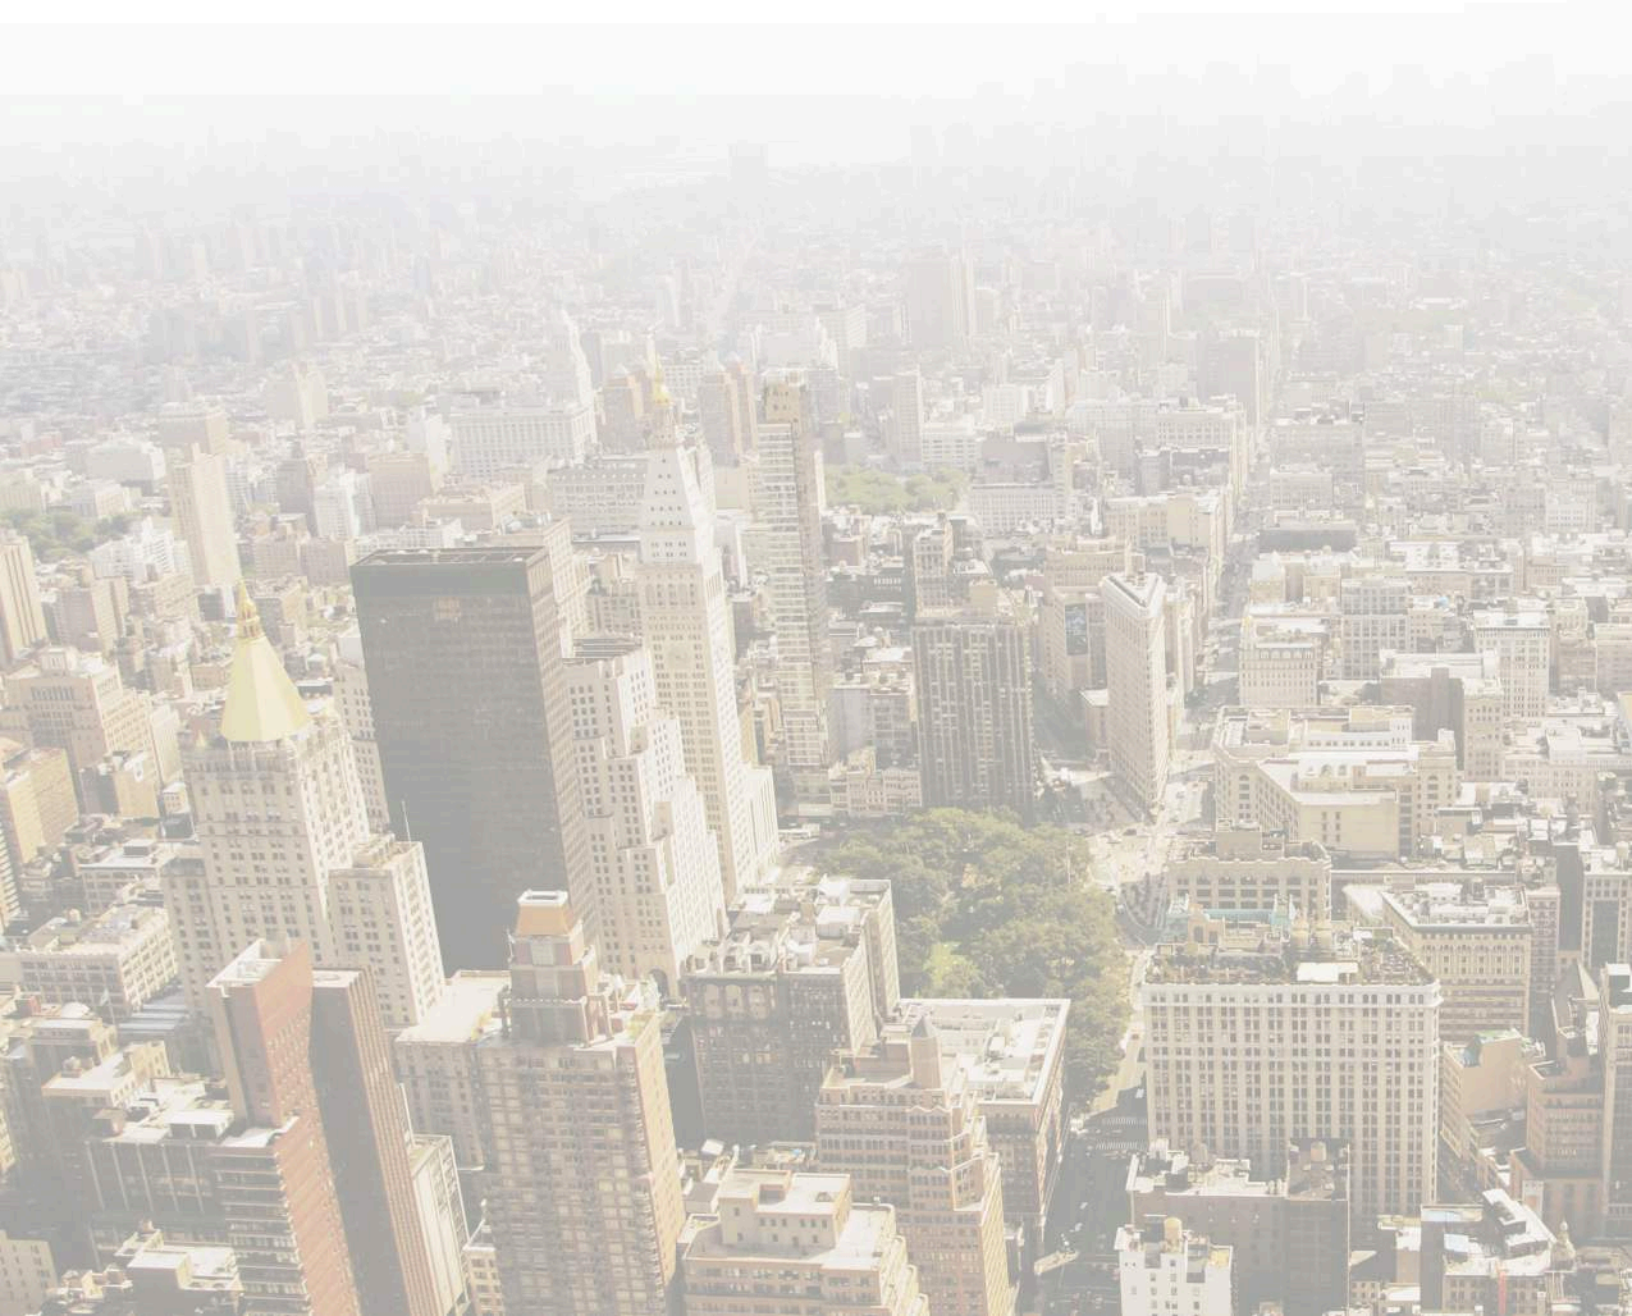

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Cloud Spectator examined Amazon EC2's C4 family, recently released on January 11th, and explored any potential performance and price-performance improvements over its previous generation, the C3 family. The results indicated that in general the C4 virtual machines (VMs) had 10%-20% higher vCPU performance and approximately 6 GB/s more memory throughput than the C3 VMs across different machine sizes. However, after factoring in the price increases, the price-performance values of the C4 VMs were on average the same as those of the C3 VMs. Both vCPU performance levels and network throughput displayed high stability over time and across all tested machines. The results indicated Amazon's effort to provide highly predictable performance outputs and to match its C4 family's price-performance with that of its earlier generation C3 family.

"Differences in performance outputs of VMs across IaaS providers can greatly impact quality of service as well as total cost of ownership."

Performance*	Key Findings		
vCPU	The C4 family's increase in integer performance over the C3 family ranged from 7% to 19% across different machine sizes.	The C4 family's increase in floating point performance over the C3 family ranged from 10% to 21% across different machine sizes.	Both families displayed high stability in vCPU performance over time, with the majority of virtual machine performance variability below 1%.
Memory	The C4 VMs exhibited an average memory throughput increase of approximately 6 GB/s on every machine size compared to the matching C3 VMs.	There are three tiers of maximum memory throughput. The xlarge, 2xlarge and 4xlarge share a tier. The large and 8xlarge have substantially different memory throughput than the other three.	Higher memory throughput fluctuations were observed on the C4 VMs than on the C3 VMs. The average performance variability of the C4 family was 9.5%, while that of the C3 family was 7.4%.
Internal Network	The C4 family displayed an average of 2% increase in public network throughput and an average of 3% increase in private network throughput from the C3 family.	Private network exhibited on average a 4.5% higher throughput than public network for all VM sizes except for the 8xlarge VM sizes, which doubled their throughput from public to private.	Little variability was observed for both families. The C4 family showed more stable network throughput than the C3 family, with its average performance variability being 0.2% compared to the C3's 0.7%.

**Storage Performance is not included in this report because there is no local storage on the C4 family VMs.*

Price-Performance	Key Findings		
Processing (vCPU + Memory)	On average, the C4 family displayed nearly identical price-performance value as the C3 family, with an average increase of 0.15%.	The C4's large and 8xlarge instances had 4% higher, while the C4's 2xlarge and 4xlarge instances had 3.5% lower price-performance than the C3 VMs on average. The xlarge instances from the two families were not distinguished in price-performance.	The C4 VM's processing performance was 10%-20% higher than that of the C3 VMs across different VM sizes.

Table 1: Key findings of vCPU, memory and network performance and processing price-performance.

INTRODUCTION

Performance and pricing are both key considerations in the public cloud industry. The price wars between some of the largest providers, such as Amazon EC2, Microsoft Azure and Google Compute, have pressured the industry to speed up innovation in order to stay competitive. Therefore, performance is quickly becoming an important focus as providers start introducing better performing offerings: Rackspace's Performance Servers (see http://bit.ly/rax_report for Cloud Spectator's generational performance report on Rackspace's cloud offerings), Microsoft Azure's D-Series (see http://bit.ly/ms_report for Cloud Spectator's generational comparison on Microsoft Azure's cloud offerings), and most recently, Amazon EC2's C4 family. This report specifically examines Amazon EC2's C4 family, and how it has improved over the C3 family from a performance and price-performance perspective.

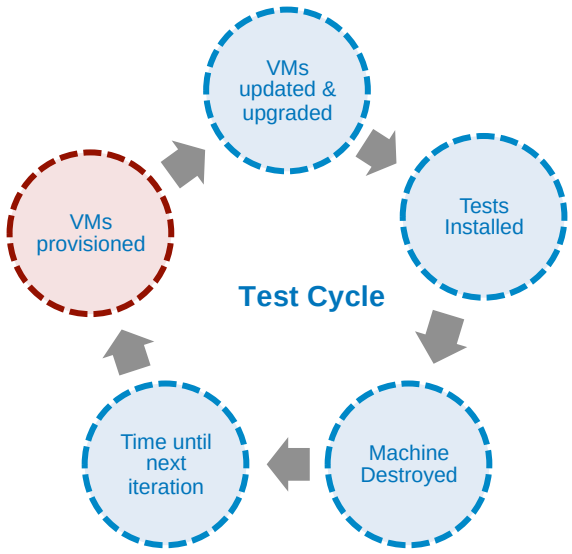

Cloud Spectator performed 10 iterations of vCPU, memory, and internal network tests on several C3 and C4 VMs of equivalent size (large – 2vCPU, xlarge – 4vCPU, 2xlarge – 8vCPU, 4xlarge – 16vCPU and 8xlarge – 32vCPU for C3 and 36 vCPU for C4; see the Appendix for more detailed information). Tests were conducted twice per day for five consecutive days to capture variation in performance across different physical hosts. Machines were terminated and recreated between each test session to increase likelihood of provisioning on different physical nodes, and therefore to provide better insight into variability that was caused by differences in physical hardware as well as potential "noisy neighbor" effects.

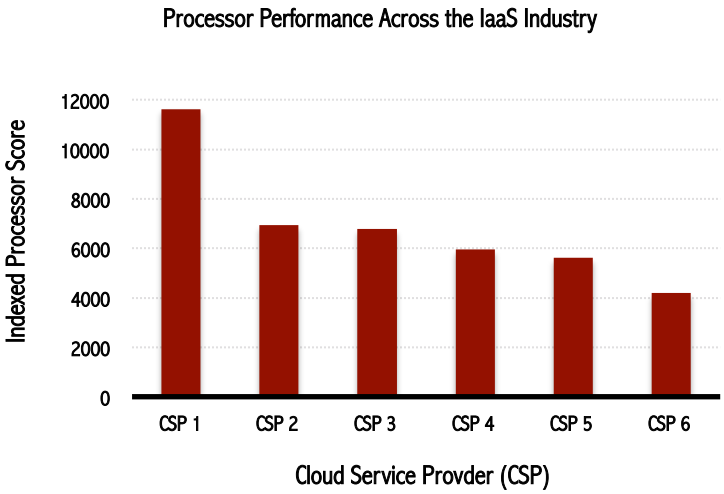

Figure 1: Variance in VM performance across the industry means that simply comparing features or pricing cannot make an apples-to-apples comparison.

WHY DOES PERFORMANCE MATTER?

Cloud infrastructure performance is a key consideration that should not be overlooked. Differences in performance outputs of VMs across IaaS providers can greatly impact quality of service as well as total cost of ownership. 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 notable advantage in 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 or machine size to house an application can result in unnecessary overspending or application performance problems.

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. Table 2 below lists the 3 virtual server components studied in this project.

vCPU	Memory	Internal Network
The performance of all applications is highly dependent on the vCPU. The vCPU is responsible for the processing and orchestration of all applications.	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.	In a cloud environment, internal network performance is critical. Distributed environments must maintain a strong network backbone for servers to communicate efficiently and synchronize data.

Table 2: A summary of each of the infrastructure components examined in this study and their functions within cloud environments.

KEY CONSIDERATIONS

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

- The C4 family was introduced three days prior to when Cloud Spectator began testing. The physical machines may not have been as saturated with users relative to the C3 VMs, which had been available for more than a year. Therefore, retesting the C4 family in a few months may result in different findings.
- The tests leveraged in this study were used to standardize and compare component performance. They may not be reflective of application performance, which varies on a case-by-case basis.
- Cloud Spectator tested each machine size multiple times, terminating and re-creating machines for each test iteration. This was done intentionally to sample different physical hosts.
- The instances were deployed in the same VPC to optimize networking. The instances were not deployed on dedicated hardware.
- The C4 and C3 price-performance comparison only took into account vCPU and memory, because of the C4's lack of local SSD storage.

PERFORMANCE FINDINGS

vCPU PERFORMANCE

Cloud Spectator used Geekbench 3 to test vCPU performance and memory bandwidth. Geekbench 3 is a modern stress-testing tool that simulates real-world scenarios/use cases for servers, desktops, and mobile devices. Such use cases include (but are not limited to) encryption, decryption, compression, and solving equations.

Amazon EC2's C3 family uses the Intel® Xeon® CPU E5-2680 V2 processor, while its C4 family uses Intel®Xeon® CPU E5-2666 V3 processor. The C4 VMs displayed higher performance than the equivalent sized C3 machines on both integer and floating point mathematics, which were measured as the two categories of vCPU performance.

The C4 family on average had 14% higher integer operation performance compared to the C3 family. The c4.8xlarge VMs exhibited the largest performance improvement with a 19% higher average integer performance score than the c3.8xlarge VMs (*See Figure 2 and Table 3*). Users with applications focusing on integer-intense operations such as AES encryption, JPEG compression and decompression, or PNG compression and decompression should note the integer performance difference between the two families.

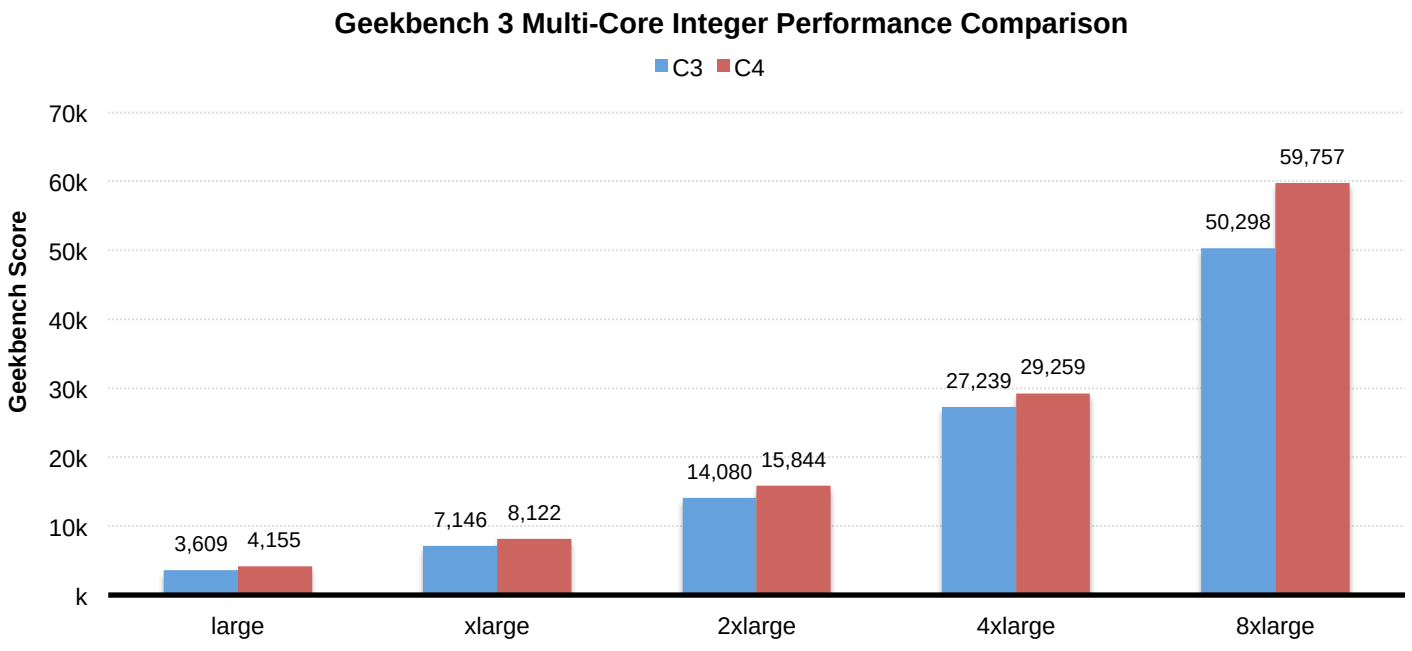

Figure 2: An overview of C3 and C4 VMs integer performance across different VM sizes. The Geekbench scores are average scores from 10 iterations.
*VM size information can be found in Appendix.

Average Geekbench 3 Multi-Core Integer Performance Scores

	large	xlarge	2xlarge	4xlarge	8xlarge
C3 Family	3,609	7,146	14,080	27,239	50,298
C4 Family	4,155	8,122	15,844	29,259	59,757
% Increase	15%	14%	13%	7%	19%

Table 3: A comparison of average integer performance across difference VM sizes. The scores are average scores from 10 iterations.

The C4 family on average had 13% higher floating point calculation performance than the C3 family. The c4.8xlarge VMs exhibited the largest performance improvement with a 21% higher average floating point performance score than the c3.8xlarge VMs (*See Figure 3 and Table 4*). The floating point performance differences are most relevant to users with applications focusing on floating point intensive operations such as options trading algorithms or image rendering.

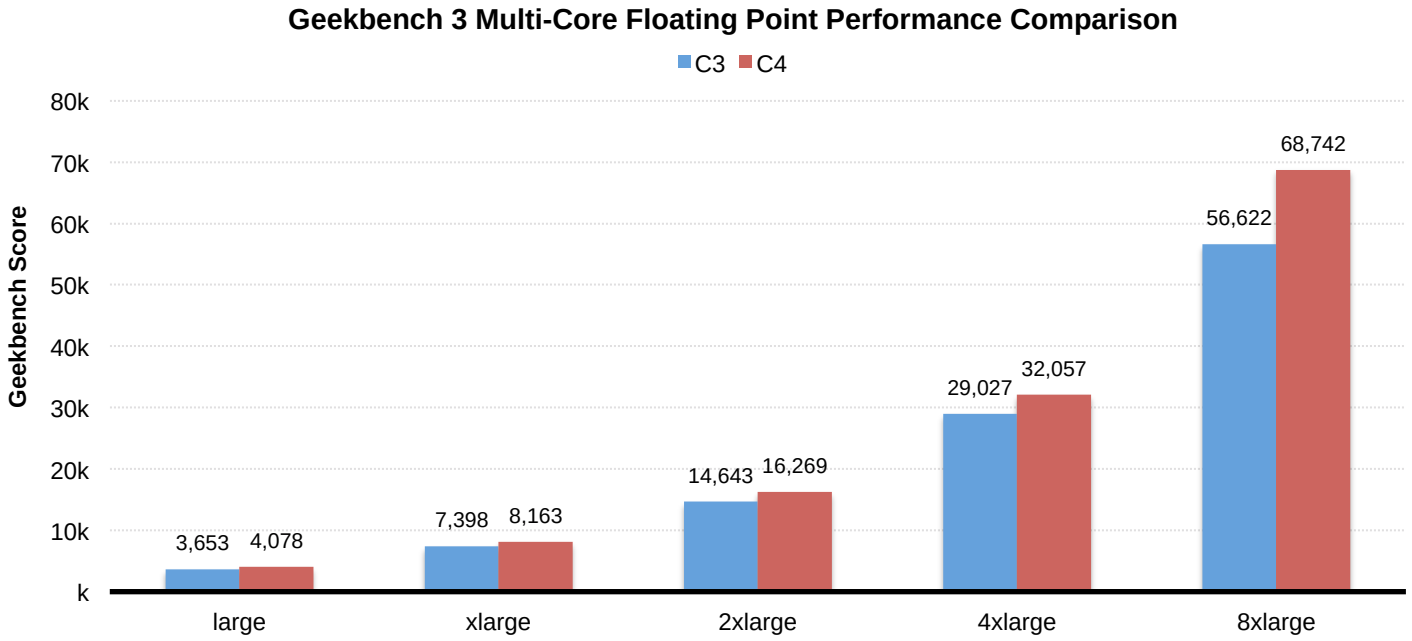

Figure 3: An overview of C3 and C4 VMs floating point performance across different VM sizes. The Geekbench scores are average scores from 10 iterations.

Average Geekbench 3 Multi-Core Floating Point Performance Scores

	large	xlarge	2xlarge	4xlarge	8xlarge
C3 Family	3,653	7,398	14,643	29,027	56,622
C4 Family	4,078	8,163	16,269	32,057	68,742
% Increase	12%	10%	11%	10%	21%

Table 4: A comparison of average floating point performance across difference VM sizes. The scores are average scores from 10 iterations.

To quantify the Performance Variability (*PV*) across different machines, Cloud Spectator calculated the coefficients of variance. This measure is defined as the standard deviation expressed as a percentage of the average:

*PV = (Standard Deviation / Average) * 100%*

Both the C3 and C4 instances exhibited a high degree of stability in the processing tests over the 10 test iterations, which likely took place on different physical hosts. The PVs for integer performance ranged from 0.3% and 3.4% (*See Figure 4 and Table 5*), and the *PV*s for floating point performance were below 1% for all machines (*See Figure 5 and Table 6*).

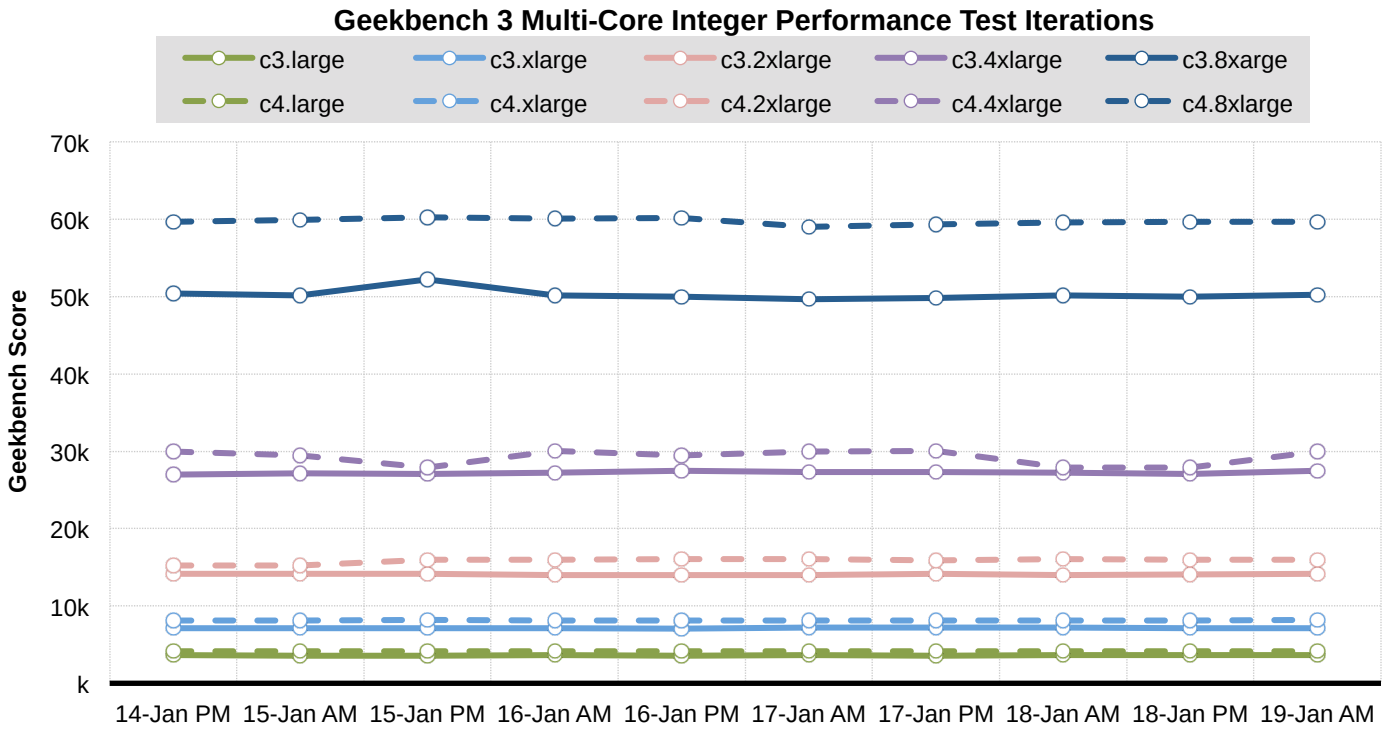

Figure 4: An overview of C3 and C4 VMs integer performance variability across different VM sizes over 10 iterations.

Multi-Core Integer Performance Variability (%)

	large	xlarge	2xlarge	4xlarge	8xlarge
C3 Family	0.9	0.5	0.5	0.6	1.4
C4 Family	0.3	0.6	2.1	3.4	0.6

Table 5: A comparison of integer performance variability across different VM sizes. The PVs are calculated from 10 iterations.

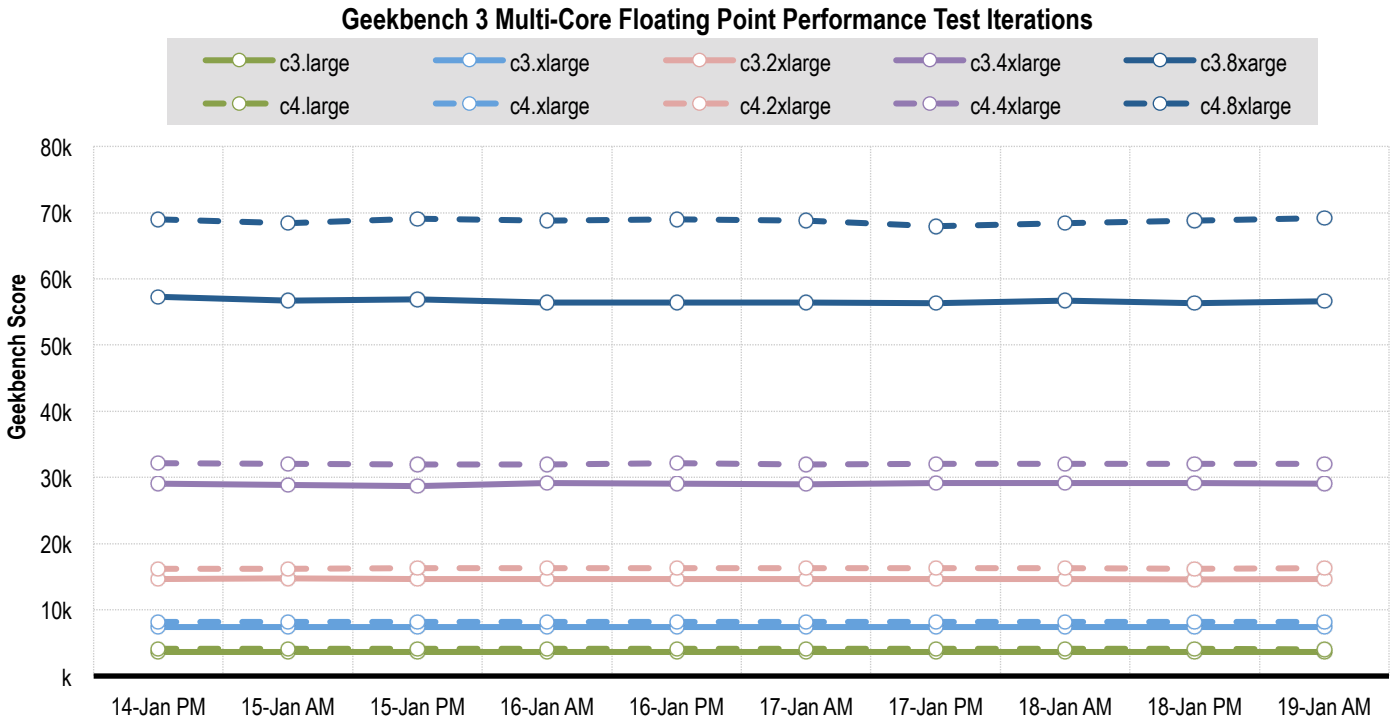

Figure 5: An overview of C3 and C4 VMs floating point performance variability across different VM sizes over 10 iterations.

Multi-Core Floating Point Performance Variability (%)

	large	xlarge	2xlarge	4xlarge	8xlarge
C3 Family	0.3	0.2	0.2	0.5	0.6
C4 Family	0.5	0.2	0.3	0.5	0.6

Table 6: A comparison of floating point performance variability across different VM sizes. The PVs are calculated from 10 iterations.

MEMORY PERFORMANCE

Cloud Spectator uses Geekbench 3's STREAM benchmark to test memory throughput performance. The STREAM benchmark is a simple synthetic benchmark program that measures sustainable memory bandwidth (in MB/s) and the corresponding computation rate for simple vector kernels.

Amazon EC2's C4 VMs showed an approximate 6 GB/s increase in memory throughput compared to their C3 counterparts across all tested VM sizes. The percentage increase was bigger for instances with lower memory throughput and smaller for instances with higher throughput as the absolute value of the increase remained relatively consistent across all VM sizes. (*See Figure 6 and Table 7*).

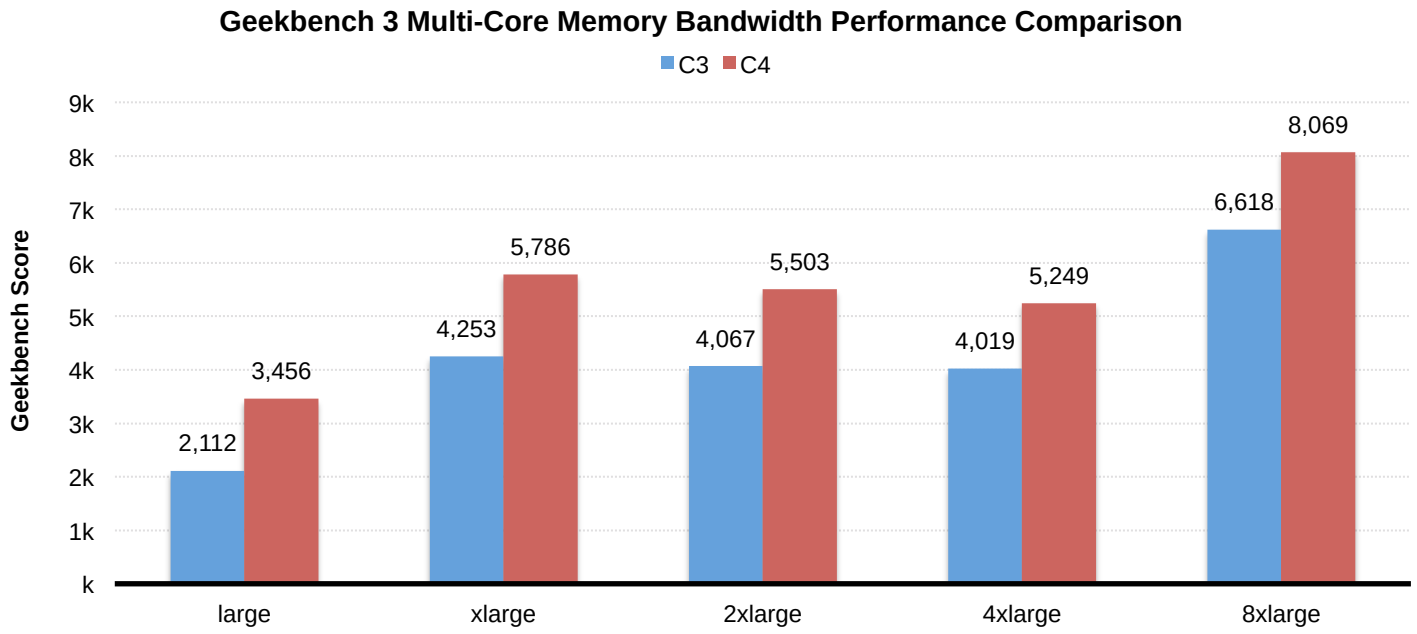

Figure 6: An overview of C3 and C4 VMs memory throughput across different VM sizes. The Geekbench scores are averaged over 10 iterations.

Average Geekbench 3 Memory Throughput Performance Scores

	large	xlarge	2xlarge	4xlarge	8xlarge
C3 Family	2,112	4,253	4,067	4,019	6,618
C4 Family	3,456	5,786	5,503	5,249	8,069
% Increase	64%	36%	35%	31%	22%

Table 7: A comparison of average memory throughput performance across different VM sizes. The scores are averaged over 10 iterations.

Depending on the VM size, the maximum memory throughput varies but can be grouped into three tiers: low (for large VMs), medium (for xlarge, 2xlarge and 4xlarge VMs) and high (for 8xlarge VMs). The memory throughput differences between the low and medium tiers and between the medium and high tiers for the C3 family were 9.1 GB/s and 10.9 GB/s respectively, while those for the C4 family were 5.9 GB/s and 10.2 GB/s respectively. The maximum memory throughput measurements for the C3 and C4 machines are displayed below (See Table 8).

Maximum Memory Throughput (GB/s)

	large	xlarge	2xlarge	4xlarge	8xlarge
C3 Family	9.7	18.8	18.7	18.7	29.7
C4 Family	20.0	25.8	25.7	25.9	36.1

Table 8: A comparison of maximum memory throughput for C3 and C4 family.

Both C3 and C4 family VMs showed variable memory throughput over several test iterations that likely took place on different physical hosts. The average performance variability of the C4 family is larger than that of the C3 family - 9.5% and 7.4 % respectively (*Figure 7 and Table 9*).

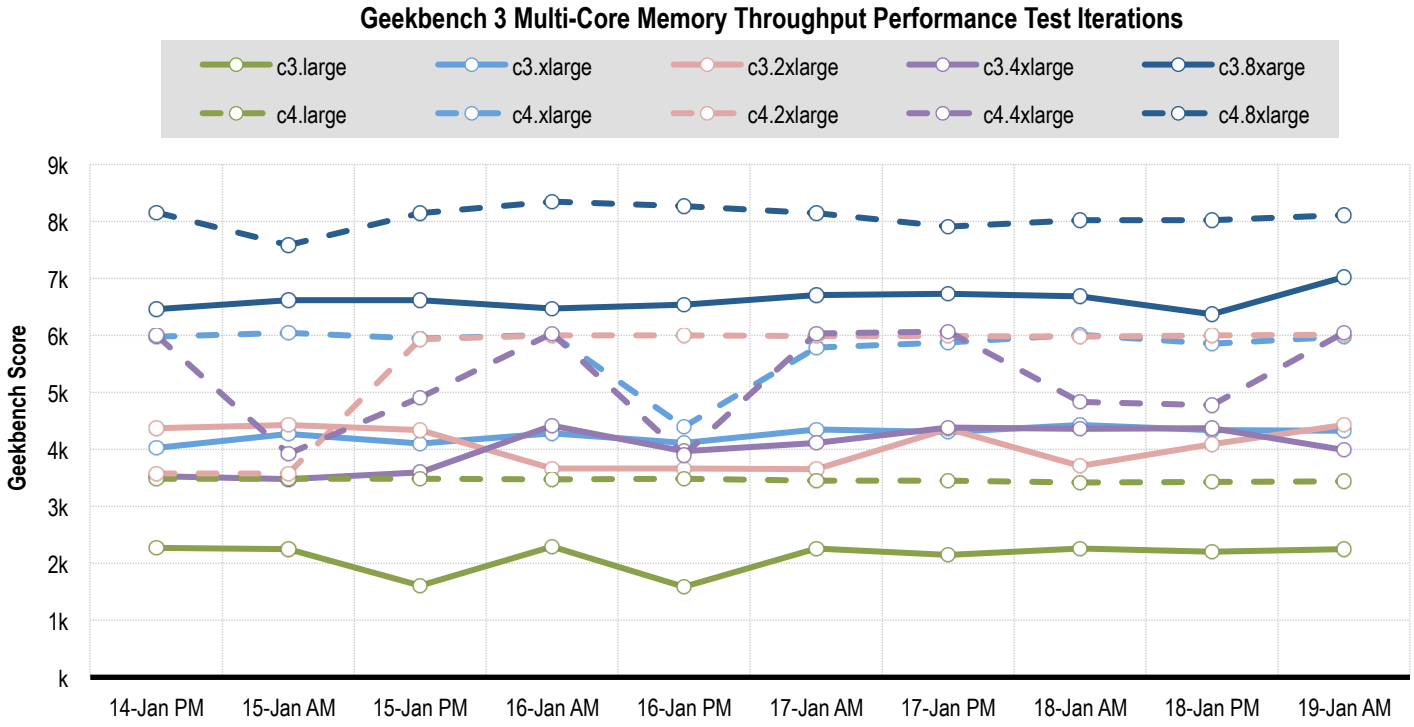

Figure 7: An overview of C3 and C4 VMs memory performance variability across different VM sizes over 10 iterations.

Memory Throughput Performance Variability (%)

	large	xlarge	2xlarge	4xlarge	8xlarge
C3 Family	13.0	3.0	8.7	9.3	2.8
C4 Family	0.7	8.6	18.5	17.0	2.6

Table 9: A comparison of memory performance variability across different VM sizes. The PVs are calculated from 10 iterations.

NETWORK PERFORMANCE

Internal network performance was evaluated by running tests measuring the throughput between machines using iperf (see Appendix for details). Network throughput is a critical component of public cloud systems, where distributed environments must maintain a strong network backbone for servers to communicate efficiently and to synchronize data. In-memory databases, distributed file systems, and high-performance application clusters require a large and fast connection across VMs within the cluster.

Internal network bandwidth increased with larger VM sizes for both families. Through running in a virtual private cloud (VPC) on Amazon EC2, the VMs were granted a more powerful network connection, which translated into more throughput, especially on the private IP addresses. For the same size VMs, the C4 family showed nearly identical public and private network throughput as the C3 family, with an average increase of 2% for public network throughput and 3% for private network throughput. Private network performance displayed an average of 4.5% increase compared to public network performance for VMs on both families except for the 8xlarge VMs. For the c3.8xlarge and c4.8xlarge instances, the private network throughput was more than twice a large as the public network throughput (*See Figure 8, Table 10 and Table 11*).

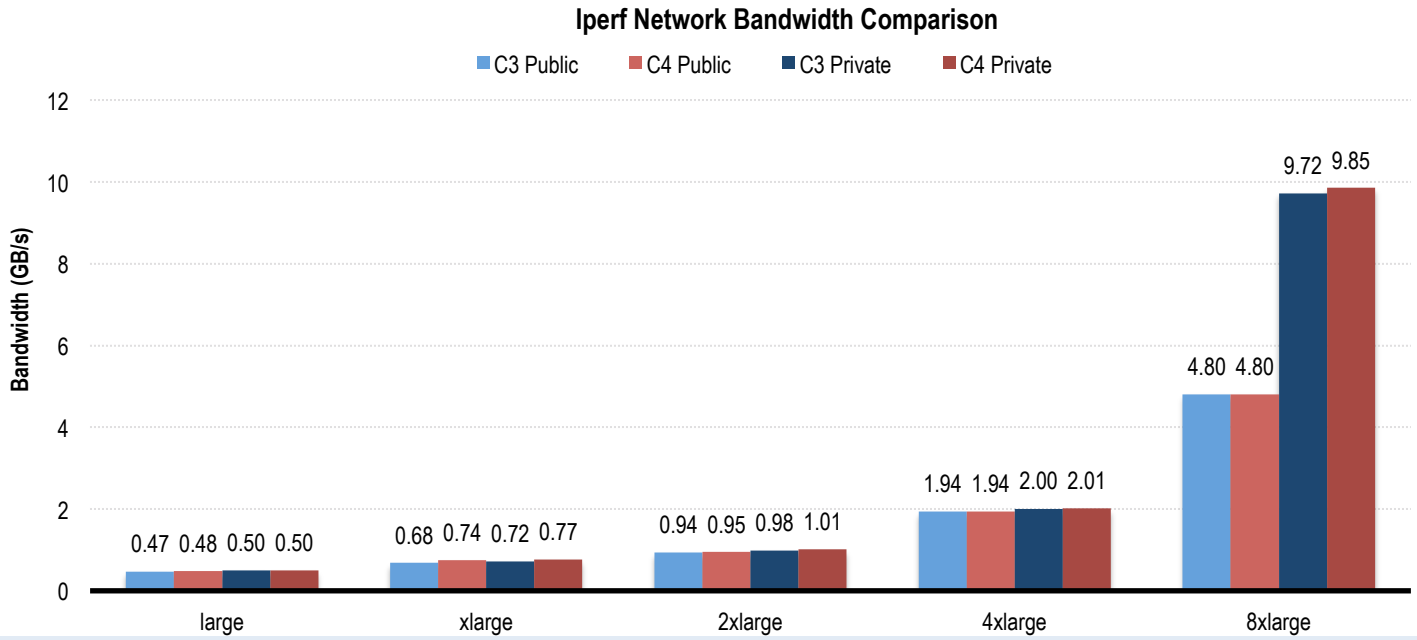

Figure 8: An overview of C3 and C4 VMs public and private network throughput across different VM sizes. The values represent averages of 10 iterations.

Average Public Network Throughput (GB/s)

	large	xlarge	2xlarge	4xlarge	8xlarge
C3 Family	0.47	0.68	0.94	1.94	4.80
C4 Family	0.48	0.74	0.95	1.94	4.80
% Increase	2%	9%	1%	0%	0%

Table 10: A comparison of average public network throughput across different VM sizes. The numbers represent averages over 10 iterations.

Average Private Network Throughput (GB/s)

	large	xlarge	2xlarge	4xlarge	8xlarge
C3 Family	0.50	0.72	0.98	2.00	9.72
C4 Family	0.50	0.77	1.01	2.01	9.85
% Increase	2%	8%	3%	1%	1%

Table 11: A comparison of average private network throughput across different VM sizes. The numbers represent averages over 10 iterations.

The internal network throughput for both the C3 and C4 families showed little variability. Most PVs for public and private network bandwidth were kept under 0.5% (*See Table 12 and Table 13*). The C4 instances were more stable than the C3 instances, with average PVs of 0.2% compared to the C3's 0.7%.

Public Network Throughput Performance Variability (%)

	large	xlarge	2xlarge	4xlarge	8xlarge
C3 Family	1.9	1.5	0.4	0.0	0.7
C4 Family	0.0	0.1	0.0	0.5	0.4

Table 12: A comparison of public network throughput PVs across different VM sizes. The PVs are calculated over 10 iterations.

Private Network Throughput Performance Variability (%)

	large	xlarge	2xlarge	4xlarge	8xlarge
C3 Family	0.3	0.2	1.3	0.2	0.7
C4 Family	0.0	0.1	0.0	0.3	0.1

Table 13: A comparison of private network throughput PVs across different VM sizes. The PVs are calculated over 10 iterations.

PRICE-PERFORMANCE COMPARISON

Cloud Spectator's price-performance calculation, the CloudSpecs Score™, provides information on how much performance the user receives for each unit of cost. For this report, Cloud Spectator focused on the C3 and C4 families' processing (vCPU and memory) price-performance.

Amazon's C4 family displayed on average a 15.6% improvement in processing performance, indicated by Geekbench 3's Total Score, as a result of higher integer, floating point and memory throughput performance. The largest performance enhancement appeared in its large and 8xlarge instances, which showed processing performance increases of 20% (*See Figure 9 and Table 14*).

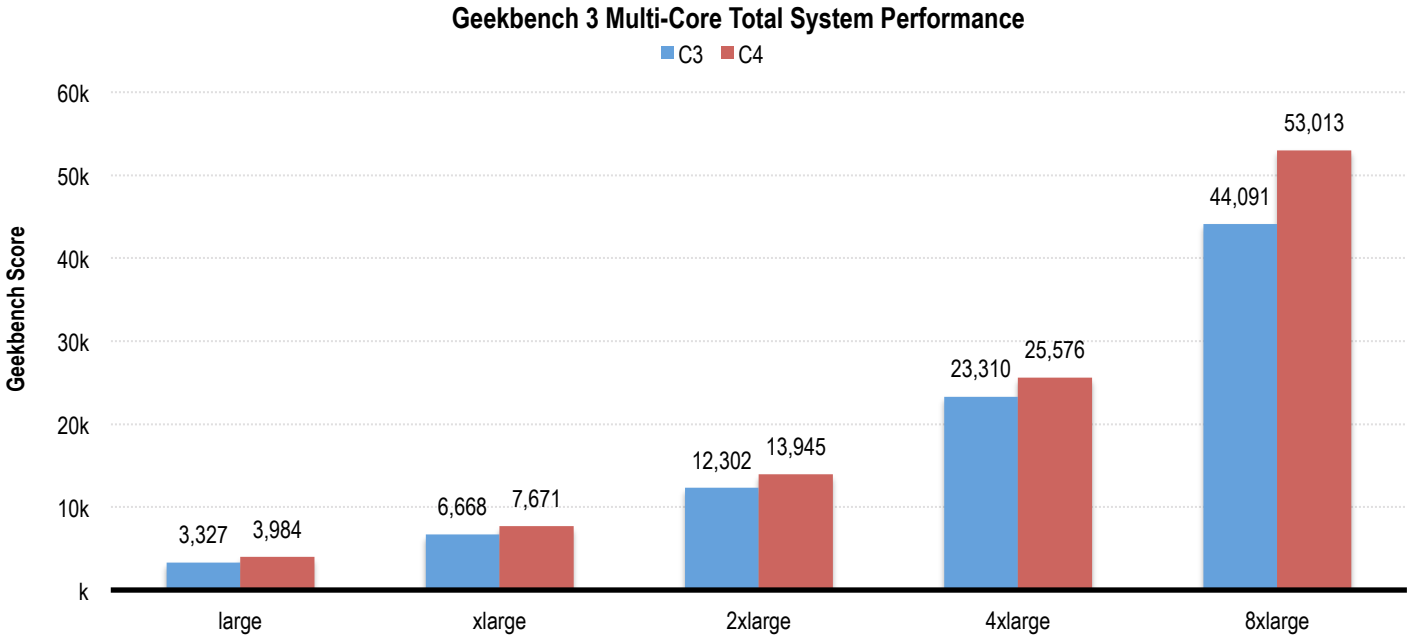

Figure 9: An overview of C3 and C4 VMs processing performance across different VM sizes. The Geekbench scores are averaged over 10 iterations.

Average Geekbench 3 Multi-Core System Performance Scores

	large	xlarge	2xlarge	4xlarge	8xlarge
C3 Family	3,327	6,668	12,302	23,310	44,091
C4 Family	3,984	7,671	13,945	25,576	53,013
% Increase	20%	15%	13%	10%	20%

Table 14: A comparison of average system performance across difference VM sizes. The numbers are averaged over 10 iterations.

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 = [provider performance score] / [provider cost]

best_provider_value = max{price-performance_values}

CloudSpecs Score™ = 100*provider_value / best_provider_value

The closer the score is to 100, the higher price-performance value it indicates. The score 100 represents the best-value provider among all in the comparison. The value is scaled; e.g., a Cloud Service Provider with a score of 100 gives 4x the value of a Cloud Service Provider with a score of 25. The CloudSpecs Scores™ in this report can only be compared with equivalent configurations; e.g., the 2 vCPU instance of the C3 family can only be compared to the 2 vCPU virtual machine of the C4 family.

The CloudSpecs Scores™ showed that the C4 family's processing price-performance value in general was not distinguished from the C3 family, despite its performance advantages. The C4 family had higher price-performance value for its large and 8xlarge VMs, but lower price-performance value for its 2xlarge and 4xlarge VMs. The C4 xlarge instances on average had an identical price-performance ratio to the C3 xlarge instances. The average CloudSpecs Score™ of the C4 family was 98.56, only a 0.15% increase from 98.41 of the C3 family (*See Figure 10 and table 15*).

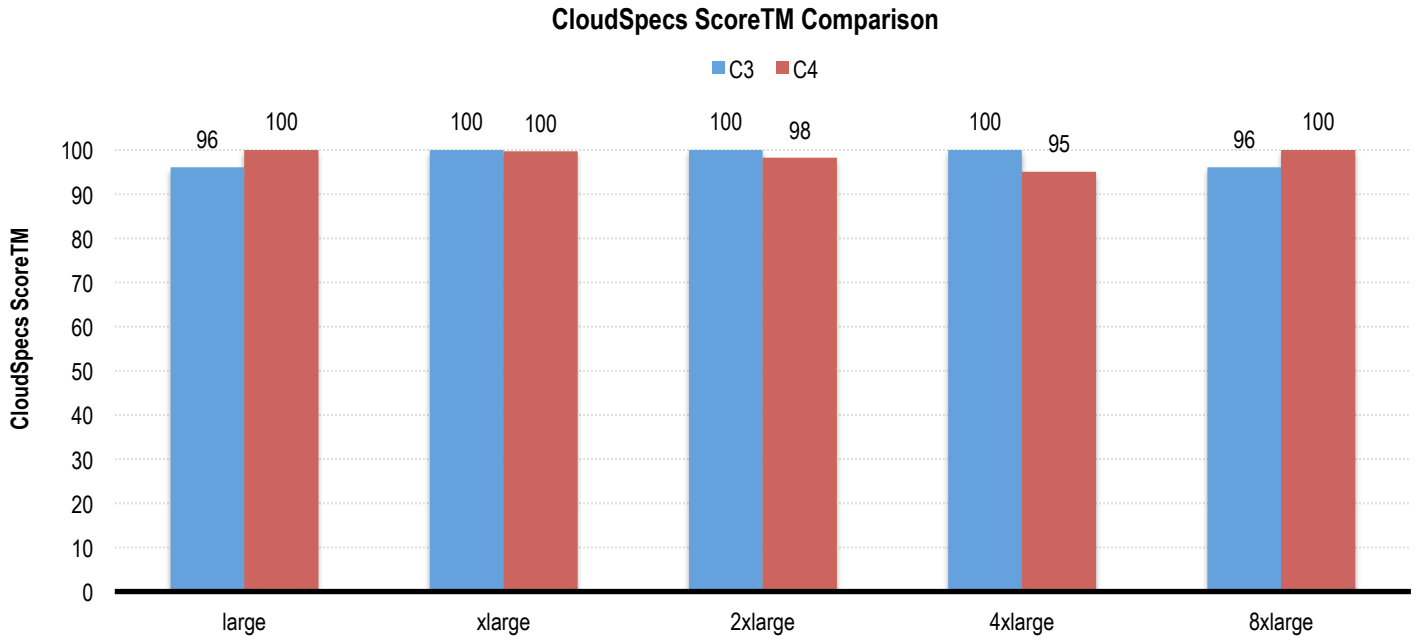

Figure 10: The CloudSpecs Scores™ of the C3 and C4 families.

Price-Performance Values

	large	xlarge	2xlarge	4xlarge	8xlarge
C3 Family	27,673	27,899	25,763	24,369	23,068
C4 Family	28,862	27,777	25,229	23,125	23,999
% Increase	4%	0%	-2%	-5%	4%

Table 15: A comparison of price-performance values across different VM sizes. The numbers are averaged over 10 iterations.

CONCLUSION

Amazon EC2's recently released C4 family presents a new generation of higher performing cloud compute instances. The C4 VMs displayed a 10%-20% improvement in vCPU (integer and floating point) performance and a 6 GB/s increase in memory throughput over the C3 VMs, which provide higher computing power for applications demanding such performance requirements. The tested VMs also displayed high stability, which is helpful for users whose applications depend on consistent performance. Rather than significantly boosting its price-performance for its new generational offering (as seen in Microsoft Azure's A and D series comparison: http://bit.ly/ms_report), Amazon EC2's C4 family preserves a close price-performance ratio with its C3 family.

Apart from the performance differences, the C4 family excludes local storage, so users must provision using EBS, AWS's scalable block storage. Thus, users considering high-performance environments such as distributed file systems or NoSQL database clusters looking for the superior performance of local SSDs may want to use the larger instances on the C3 family. By contrast, users looking to utilize block storage with a higher-performance equivalent of the current C3 generation should look to C4 instances. Block storage utilization is recommended to add a layer of resiliency to the instance; local storage, which may perform faster, can result in data being lost in the case of an accidental termination of the instance and is not scalable with increased data sizes.

Virtual machine performance is a very important consideration not only because users need their applications to be fast, but also because of the potential cost savings realized when buying less infrastructure on better-performing services. For example, a Microsoft D2 (2 vCPU) has a similar CPU performance level to an A3 (4 vCPUs), but the D2 can save a user $0.07 per hour per machine over the A3. In generational comparisons, usually a more significant improvement in price-performance can be seen, but the lack of a notable improvement Amazon EC2's new generation may be an indication that Amazon has made an effort to match its prices with the performance output of each machine.

FURTHER STUDY

This report provides an introductory analysis into Amazon Web Services' new C4 machines to create a foundation for further study. This report should only be used as a starting point to address further questions. Some additional studies should consider:

- Application performance, which may experience different bottlenecks for different requirements. For further interest in specific application performance, users should target application-focused benchmark tests.
- Comparing Amazon EC2's C4 family with recent offerings from other providers to achieve a better understanding of Amazon's position in the market from a performance and price-performance value standpoint.
- Testing Amazon EC2's local and EBS storage to gain a thorough understanding of storage performance.

APPENDIX

TESTING DESCRIPTIONS

Test	Category	Description
Geekbench 3	CPU	Separate CPU tests that are all aggregated into a final score. • Subtests include: Integer Math, Floating Point Math
Geekbench 3	Memory	Four separate memory tests that are all aggregated into a final score. • Subtests include: Add, Copy, Triad, Scale
Iperf	Network	Transfers data bi-directionally between 2 nodes within 60 seconds, measured in megabits per second.

Geekbench 3 test details and benchmark software can be found at: http://www.primatelabs.com/geekbench/

SERVER SETUPS

	Amazon EC2 C3 family	Amazon EC2 C4 family
2 core VM	**c3.large**	**c4.large**
Virtual Cores	2	2
Memory (in GB)	3.75	3.75
Price ($/hr)	$0.120	$0.138
4 core VM	**c3.xlarge**	**c4.xlarge**
Virtual Cores	4	4
Memory (in GB)	7.5	7.5
Price ($/hr)	$0.239	$0.276
8 core VM	**c3.2xlarge**	**c4.2xlarge**
Virtual Cores	8	8
Memory (in GB)	15	15
Price ($/hr)	$0.478	$0.552
16 core VM	**c3.4xlarge**	**c4.4xlarge**
Virtual Cores	16	16
Memory (in GB)	30	30
Price ($/hr)	$0.956	$1.104
32 (36) core VM	**c3.8xlarge**	**c4.4xlarge**
Virtual Cores	32	36
Memory (in GB)	60	60
Price ($/hr)	$1.912	$2.208

Operating System	Ubuntu 14.04	Ubuntu 14.04
Data Center Location	US Northern California	US Northern California

NOTES ON PROVISIONING

- All virtual machines were provisioned in the US Northern California data center.
- After each test iteration, the virtual machines were deleted; every test iteration started with a newly provisioned machine. This was done to enable the possibility of switching physical machines throughout the process of testing, which would increase the relevancy to any potential Amazon Web Services users and the effectiveness of the sample size.
- Once a new virtual machine was created for each new iteration of testing, “apt-get update” and “apt-get upgrade --yes” commands were issued, followed by a reboot of the system before installing any tests.
- All instances were previsioned in the same virtual private cloud (VPC) in order to enable private network testing.

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.

Cloud Spectator
800 Boylston Street, 16th Floor
Boston, MA 02199
www.cloudspectator.com
Phone: (USA) 1-617-300-0711
icontact@cloudspectator.com