International Journal of Computer Trends and Technology
ISSN: 2231–2803 / https://doi.org/10.14445/22312803/IJCTT-V72I11P127
Volume 72 Issue 11, 248-255, November 2024

Review Article

Cost-Performance Evaluation of General Compute Instances: AWS, Azure, GCP, and OCI

Jay Tharwani[1], Arnab Purkayastha[2]

[1]*Independent Researcher, Charlotte, NC, USA.*
[2]*Western New England University, Springfield, MA, USA.*

[1]*Corresponding Author : jtharwan@alumni.uncc.edu*

Received: 13 October 2024 Revised: 14 November 2024 Accepted: 26 November 2024 Published: 30 November 2024

Abstract *- Cloud computing has become the cornerstone of modern IT infrastructure, offering a wide range of general-purpose instances optimized for diverse workloads. This paper compares the cost and performance of general-purpose compute instances across four major cloud providers: AWS, Azure, Google Cloud Platform (GCP), and Oracle Cloud Infrastructure (OCI). Using standardized configurations of 4 vCPUs and 16 GiB of RAM, the study evaluates instances based on processor architecture (Intel, AMD, ARM), pricing models, and performance benchmarks. Key findings reveal that ARM-based instances deliver superior price-performance ratios for cost-sensitive workloads, while Intel-based instances excel in enterprise-grade applications requiring versatility and reliability. The results aim to guide organizations in selecting the most cost-effective and performance-efficient cloud resources for their specific needs.*

Keywords *- Cloud computing, Cost-performance analysis, General purpose instances, Processor architectures, Virtual machines.*

1. Introduction

Gone are the days when every business had its own data centre for IT infrastructure. Businesses are widely adopting the cloud for their business-critical workloads. Cloud computing has revolutionized modern IT infrastructure, enabling organizations to deploy, manage, and scale workloads with unprecedented flexibility and efficiency. Among the diverse computing options cloud providers offer, general-purpose instances stand out for their balanced configuration of computing, memory, and networking resources. These instances are widely used for a variety of workloads, including web servers, relational databases, application hosting, and development environments. This paper focuses on evaluating general-purpose compute instances across four leading cloud platforms: Amazon Web Services (AWS), Microsoft Azure, Google Cloud Platform (GCP), and Oracle Cloud Infrastructure (OCI). Each provider offers instance families powered by different processor architectures—Intel, AMD, and ARM—catering to diverse workload requirements and cost considerations. As the cloud landscape evolves, making informed decisions about instance selection becomes critical, particularly for organizations optimizing their infrastructure for performance and cost. The objective of this study is to conduct a comprehensive comparison of these general-purpose instances by analyzing their:

- Cost efficiency: Comparing hourly pricing and 1-year commitment discounts.
- Performance metrics: Using standardized benchmarks
- Processor architectures: Comparing cost per performance between manufacturers and instruction set architectures

This analysis leverages data from official documentation, pricing calculators, and performance benchmarks. The findings aim to guide organizations in selecting the most suitable general-purpose instances, balancing cost-effectiveness with workload requirements. This paper seeks to contribute to better decision-making in selecting cost-efficient general-purpose cloud instances by providing actionable insights.

2. Architecture Overview: x86 (CISC) vs. ARM (RISC)

2.1. Overview of CISC and RISC Architectures

2.2.1. CISC (Complex Instruction Set Computing)

CISC architectures, such as x86, are designed to execute complex instructions in fewer lines of code. They offer a rich set of instructions and can directly perform high-level tasks. This design simplifies programming at the cost of increased hardware complexity and power consumption.

2.2.2. RISC (Reduced Instruction Set Computing)

RISC architectures, like ARM, focus on executing simple instructions that take a uniform amount of time. The simplicity enables faster execution of individual instructions and lower power consumption, making RISC processors more efficient for specific workloads [21, 22].

2.2 Advantages of Each Architecture

See Table 1 for the advantages of each architecture.

Table 1. Advantages of each architecture

Feature	CISC (x86)	RISC (ARM)
Instruction Set	Rich and complex, requiring fewer lines of code.	Simplified, with fixed instruction lengths.
Power Consumption	Higher due to complexity.	Lower, making it energy-efficient.
Performance per Watt	Moderate, better for single-threaded, high-performance tasks.	Excellent, ideal for multi-threaded workloads.
Programming Simplicity	Simplifies high-level programming tasks.	Requires optimized software for peak performance.
Hardware Complexity	Higher, leading to increased power and thermal output.	Lower, enabling lightweight and efficient designs.
Cost	Typically, it is more expensive.	Generally cheaper due to simpler manufacturing.

2.3. Scenarios Where Each Architecture Excels

2.3.1. CISC (x86)

- High-Performance Applications

Suitable for compute-intensive tasks such as database management, large-scale data analytics, and virtualization, where high single-threaded performance is critical.

- Legacy Software Support

Offers excellent backward compatibility, making it a preferred choice for applications dependent on older software [22, 24].

2.3.2. RISC (ARM)

- Energy-Efficient Workloads

Ideal for scenarios where power efficiency and cost savings are paramount, such as web hosting, content delivery networks, and microservices.

- Cloud-Native and Parallel Workloads

Excels in distributed systems and containerized environments due to its ability to handle multi-threaded workloads efficiently [22, 24].

2.4 Applicability in Data Centers

2.4.1. CISC (x86) in Data Centers

- Advantages

Dominates traditional data centers, particularly for workloads requiring consistent high performance and compatibility with legacy software.

- Use Cases

Enterprise applications, virtual machines, high-frequency trading systems, and HPC (High-Performance Computing) [25, 26].

2.4.2. RISC (ARM) in Data Centers

- Adantages

Emerging as a viable alternative for energy-efficient cloud deployments. ARM processors like AWS Graviton3 provide excellent performance per watt, reducing operational costs.

- Use Cases

Cloud-native applications, serverless computing, web hosting, and workloads with predictable patterns [25, 26].

2.5. The Data Center Trade-Off

The choice between x86 (CISC) and ARM (RISC) in data centers depends on specific workload requirements:

- If performance and compatibility are critical, x86 is the go-to architecture.
- ARM is a compelling choice for cost-sensitive, energy- and scalable cloud environments.

3. Methodology

3.1. Cloud Service Provider Selection Criteria

The cloud service providers selected were based on cloud market share, with Amazon web services leading the approximately 300-billion-dollar[16] market at 36%.

Microsoft Azure is shortly after them at 23%, and Google Cloud is shortly after at 7%. Oracle Cloud Infrastructure was chosen as the 4th cloud provider as it is quickly gaining much momentum [17][18][19].

3.2. Instance Selected

3.2.1. AWS

- M6i: Custom Intel Xeon Platinum[10] 8375C (Ice Lake), x86, 4 GiB per vCPU. Use Case: Web servers, application servers, small-medium databases.
- M6a**:** AMD EPYC 7R13 (Milan), x86, 4 GiB per vCPU. Use Case: Cost-sensitive general-purpose workloads, scalable applications.
- M7g**:** AWS Graviton3 [6](ARM-based, Neoverse V1 cores), 4 GiB per vCPU. Use Case: Cloud-native apps, gaming servers, caching fleets, microservices.

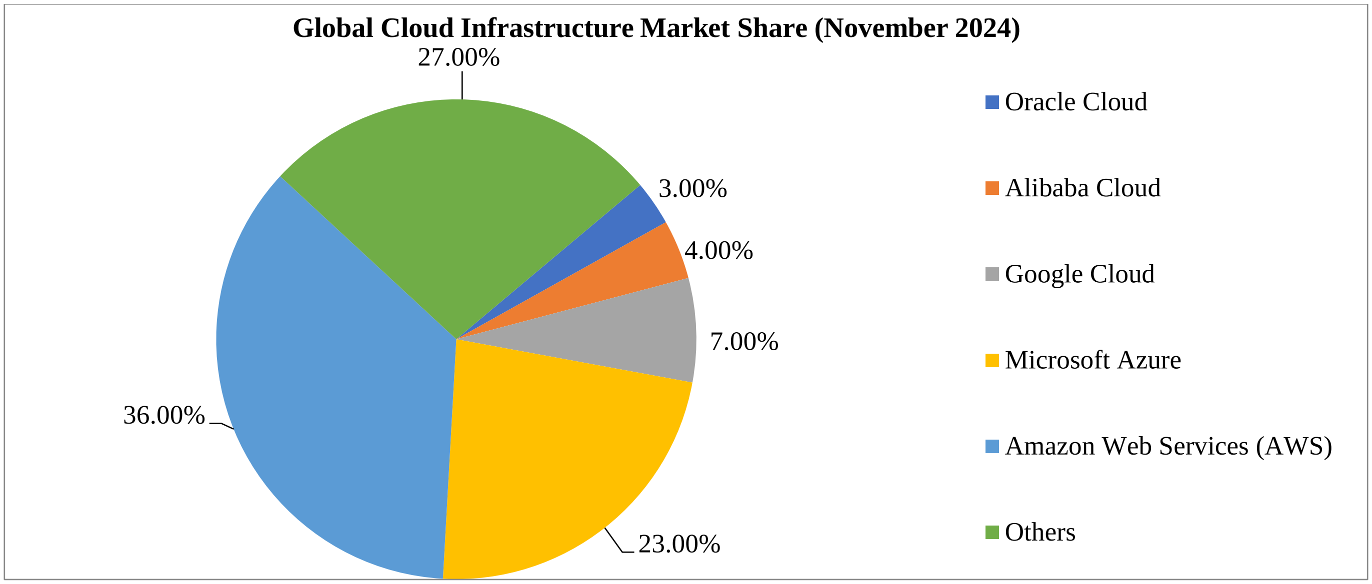

Fig. 1 Global cloud infrastructure market share

3.2.2. Azure

- Dv5: Intel Xeon Platinum 8370C (Ice Lake), x86, 4 GiB per vCPU. Use Case: Enterprise applications, relational databases, web services.[7]
- Dasv5: AMD EPYC 7763 (Milan), x86, 4 GiB per vCPU. Use Case: Cost-efficient general-purpose tasks, database servers.[7]
- Dpsv5: Ampere Altra (80 cores @ 3.0 GHz), ARM, 4 GiB per vCPU. Use Case: Scale-out workloads, open-source databases, and modern cloud-native applications.

3.2.3. Google Cloud Platform (GCP)

- N2: Intel Xeon Platinum 8273CL[13] (Cascade Lake), x86, 4 GiB per vCPU. Use Case: Web and application servers, enterprise apps, analytics.
- N2D: AMD EPYC 7B12(Rome)[13] ,x86, 4 GiB per vCPU. Use Case: Cost-optimized general-purpose workloads scalable applications.[8]
- Tau T2A: Ampere Altra Q64[13], ARM, 4 GiB per vCPU. Use Case: Cost-sensitive apps, containerized microservices, development/test environments.[8]

3.2.4. Oracle Cloud Infrastructure (OCI)

- VM.Standard3.Flex: Intel Xeon Platinum 8358 (Ice Lake), x86, 4 GiB per OCPU. Use Case: Enterprise applications, dynamic web servers, small-medium databases.[9]
- VM.Standard.E4.Flex: AMD EPYC 7742 (Milan), x86, 8 GiB per OCPU. Use Case: Cost-optimized workloads, scalable general-purpose applications.[9]
- VM.Standard.A1.Flex: Ampere Altra Q80-30, ARM, configurable memory (up to 64 GiB per core). Use Case: Cloud-native apps, portable applications, web hosting services.[9]

3.3. Instance Selection Criteria

Among the large selection of instance types between all cloud providers, the selected instances are the most popular and can efficiently run general-purpose workloads. These instances represent general-purpose workloads because:

- Balanced Resources: They provide a consistent compute-to-memory ratio (e.g., 4 GiB per vCPU), ideal for web servers, databases, and application servers.
- Versatility: Suitable for various workloads, from development to enterprise applications.
- Popularity: Were actively promoted by cloud providers as default general-purpose choices in the past (e.g., AWS M6i, Azure Dv5, GCP N2, OCI VM.Standard3).
- Diverse Architectures: Include Intel, AMD, and ARM options, catering to both traditional and modern workloads.
- Cost-Effectiveness: Lower costs than specialized instances while maintaining performance for common use cases.
- Global Availability: Widely available across regions, ensuring consistent performance and pricing.

All instances amongst the 4 cloud providers were chosen to have 4 Virtual processors and 16GiB of RAM to ensure a fair performance comparison.

3.4. Sources of Data

The data was collected from official cloud provider documentation, pricing calculators, and performance benchmarks. For instance, details and specifications referred to publicly available resources on AWS, Azure, Google Cloud, and Oracle Cloud Infrastructure websites. Performance metrics were obtained from an independent benchmark, Geekbench6. Pricing information was retrieved using cloud provider calculators and websites. The instances are also standardized (no. of vCPUs and RAM) for a fair comparison.[1][2][3][4]

3.5. Benchmark Selection

The Geekbench6 benchmarking suite was run across all instances to measure the performance of the respective instances. Geekbench6[15] is a cross-platform benchmark that measures a processor's single-core and multi-core performance by running some tests that many modern-day applications might use and finally gives a Single-Core Score and a Multi-Core score depending on how well the tests performed. Tests include File Compression, Navigation, HTML5 Browser, PDF Renderer, Photo Library, Clang, Text Processing, Asset Compression, Object Detection, Background Blur, Horizon Detection, Object Remover, HDR, Photo Filter, Ray Tracer, Structure from Motion.

3.6. Abbreviations and Calculation Overview

The paper uses the following abbreviations in the data tables that follow this section to be space efficient.

3.6.1. CSP

This was the cloud service provider against which the instance was chosen. AWS is Amazon web services, Azure is Microsoft's Azure, GCP is Google cloud platform, and OCI is Oracle Cloud Infrastructure.

3.6.2. Type

This is the type of instance series of the chosen cloud provider. The instance series are typically identified by a certain type of physical processor on which the instance runs, including its manufacturer, instruction set architecture and other parameters. Cloud providers also classify their instance series on a parameter they are optimizing, e.g., memory optimized, and many of these series offerings exist. However, for this paper, each cloud provider has chosen a general compute series. In addition, the chosen have 4 Virtual Processors and 16(GiB) of RAM.

3.6.3. Net(Gbps)

This is the max network bandwidth offered by the instance in the chosen series measured in Gigabit per second.

3.6.4. $/hr

Dollar per hour: This is the fee the cloud service provider charges per hour, in USD, when using their instance. Typically, the CSPs list this data as on-demand usage on their websites

3.6.5. $/hr 1 year

Estimated dollar per hour charges when a customer commits to using the instance for a year.

3.6.6. SCP

A single core performance score was given by Geekbench when geekbench6 was run in this instance.

3.6.7. MCP

A multi-core performance score was given by Geekbench when geekbench6 was run on this instance.

3.6.8. MCP/$/hr

The on-demand pricing ($/hr) was divided by the multicore performance score to get on-demand pricing per 1 multicore performance point per hour.

3.6.9. MCP/$/hr 1yr

The yearlong committed pricing was divided by the multicore performance score to get a yearlong commitment pricing per 1 multicore performance point per hour.

3.7. One Year Commitment Pricing

This paper has used the pricing websites of various cloud providers and pricing calculators to estimate the per hour pricing of the cloud infrastructure when committed for a year. Typically, cloud providers offer a discount when committed to their infrastructure for a year. However, OCI was an outlier in this case, as they run on a universal credits model and the yearly pricing is unavailable on their website. They mentioned contacting the sales department and offering an estimated discount of 20% compared to their on-demand pricing. Due to a lack of data, this paper has kept the on-demand and one-year commitment pricing the same for OCI. However, customers may be able to get a 20% when contacting their sales team.

4. Results and Discussion

4.1. Intel-Based Instances

Table 2 shows the performance, price and price per performance scores for general purpose Intel instances [1][2][3]**Error! Reference source not found.**

Key insights

- OCI leads the price per performance score with an on demand MCP/$/hr score of 0.0000320 and a 1-year MCP/$/hr score of 0.0000320.
- If taking OCI price as a reference, the on demand MCP/$/hr score of AWS is 81.88% higher, Azure is 70.31% higher, and GCP is 155% higher.
- The 1-year commitment price is much more comparable with OCI. AWS is 12.50% higher, Azure is 16.88% higher, and GCP is 60.63% higher in price per performance point.
- Overall, single core and multicore performance is highest in the Azure Dv5 series. However, the performance might not be a comparable metric due to some minor changes in the underlying hardware configuration offered by each Cloud service provider.

4.2. AMD-Based Instances

Table 3 shows the performance, price and price per performance scores for general purpose AMD instances [1][2][3][4]**Error! Reference source not found.**

Table 2. Intel based instances

Cloud	Type	Net(Gbps)	$/hr	$/hr 1 yr	SCP	MCP	MCP/$/hr	MCP/$/hr 1yr
AWS	M6i	Up to 12	$0.192	0.1185	1538	3297	0.0000582	0.0000360
Azure	Dv5	Up to 10	$0.192	0.1317	1661	3524	0.0000545	0.0000374
GCP	N2	Up to 10	$0.2187	0.1378	1205	2682	0.0000816	0.0000514
OCI	VM.Standard3.Flex	Up to 10	$0.104	0.104	1543	3254	0.0000320	0.0000320

Table 3. AMD based instances

CSP	Type	Net(Gbps)	$/hr	$/hr 1 yr	SCP	MCP	MCP/$/hr	MCP/$/hr 1 yr
AWS	M6a	Up to 12	$0.1728	$0.12695	1611	3641	0.0000475	0.0000349
Azure	Dasv5	Up to 10	$0.1720	$0.1175	1606	3718	0.0000463	0.0000316
GCP	N2D	Up to 10	$0.19032	$0.1199	1561	3634	0.0000524	0.0000330
OCI	VM.Standard.E4.Flex	Up to 10	$0.074	$0.074	1559	3588	0.0000206	0.0000206

Table 4. ARM based instances

Cloud	Type	Net(Gbps)	$/hr	$/hr 1 yr	SCP	MCP	MCP/$/hr	MCP/$/hr 1 yr
AWS	M7g	Up to 12	0.1632	$0.1199(-27%)	1462	4796	0.0000340	0.0000250
Azure	Dpsv5	Up to 10	$0.1540	$0.1058	1108	3721	0.0000414	0.0000284
GCP	Tau T2A	Up to 10	$0.154	$0.09702	1125	3785	0.0000407	0.0000256
OCI	VM.Standard.A1.Flex	Up to 10	$0.064	$0.064	1122	3747	0.0000171	0.0000171

Key insights

- OCI again leads the price for performance score with an on demand MCP/$/hr score of 0.0000206 and a 1-year MCP/$/hr score of 0.0000206.
- Taking OCI price as a reference, the on demand MCP/$/hr score of AWS is 130.58% higher, Azure is 124.76% higher, and GCP is 154.37% higher.
- The 1-year commitment price is better. AWS is 69.32% higher, Azure 53.40% higher, and GCP 60.19% in price per performance point.
- Overall single core and multicore performance are comparable and highest in AWS & Azure m6a & Dasv5 series. However, the performance might not be a comparable metric due to some minor changes in the underlying hardware configuration offered by each cloud service provider.

4.3. ARM-Based Instances

Table 4 shows the performance, price and price per performance scores for general purpose ARM instances[1][2][3][4]**Error! Reference source not found.**

Key insights

- ARM architecture provides the lowest price per performance across all cloud providers. This is also in line with the energy savings of the ARM architecture.
- OCI yet again leads the price per performance score with an on demand MCP/$/hr score of 0.0000171 and a 1-year MCP/$/hr score of 0.0000171.
- If taking OCI price as a reference, the on demand MCP/$/hr score of AWS is 98.83% higher, Azure is 142.11% higher, and GCP is 138.01% higher.
- The 1-year commitment price is much more comparable with OCI. AWS is 46.20% higher, Azure 66.08% higher, and GCP 49.71% higher in price per performance points
- Overall, single core and multicore performance is significantly higher in the AWS m7g series. However, the performance might not be a comparable metric due to some minor changes in the underlying hardware configuration offered by each Cloud service provider.
- AWS has its ARM architecture CPU, while the other cloud providers provide instances on Ampere Altara CPUs.

4.4. Discussion

The data reveals key trends across Intel, AMD, and ARM-based general-purpose instances:

4.4.1. Performance Efficiency

- Intel (x86): AWS M6i leads with an MCP Score of 3297, but OCI VM.Standard3.Flex offers competitive performance (3254) at the lowest cost.
- AMD (x86): All providers deliver near-parity performance, but OCI VM.Standard.E4.Flex stands out with the lowest price-per-performance ratio.
- ARM: AWS M7g (4796) and GCP Tau T2A (3785) showcase ARM's growing relevance in price-performance with OCI VM.Standard.A1 offers cost-efficient options.

4.4.2. Cost Insights

- ARM instances consistently provide the best price-performance ratio, which is ideal for cost-sensitive workloads.

- OCI offers the lowest hourly rates across all architectures, making it appealing for budget-conscious deployments.

4.4.3. Network Bandwidth

- Intel instances on AWS benefit from slightly higher bandwidth (12 Gbps), but parity exists across most providers and architectures.

4.4.4. Use Cases

- Intel: Best for enterprise-grade applications and legacy software.
- AMD: Cost-efficient for analytics and databases.
- ARM: Optimal for cloud-native, containerized, and energy-efficient workloads.

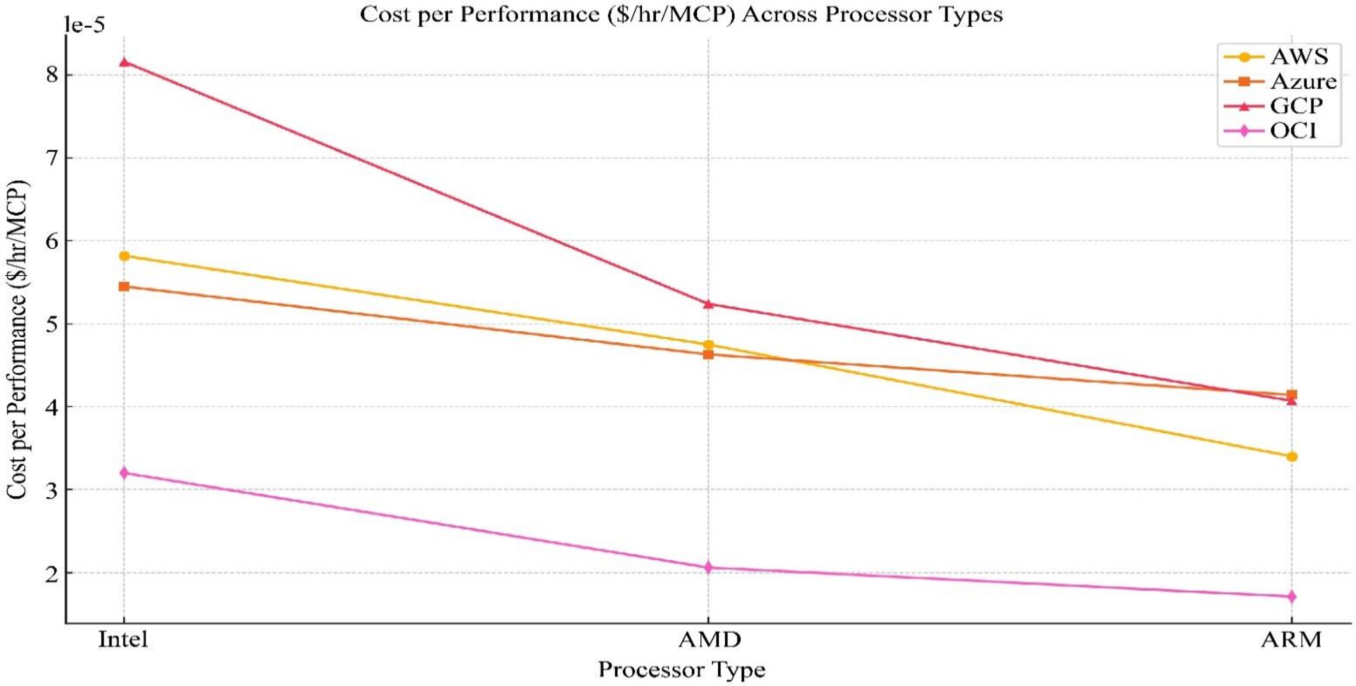

Fig. 2 Cost per performance

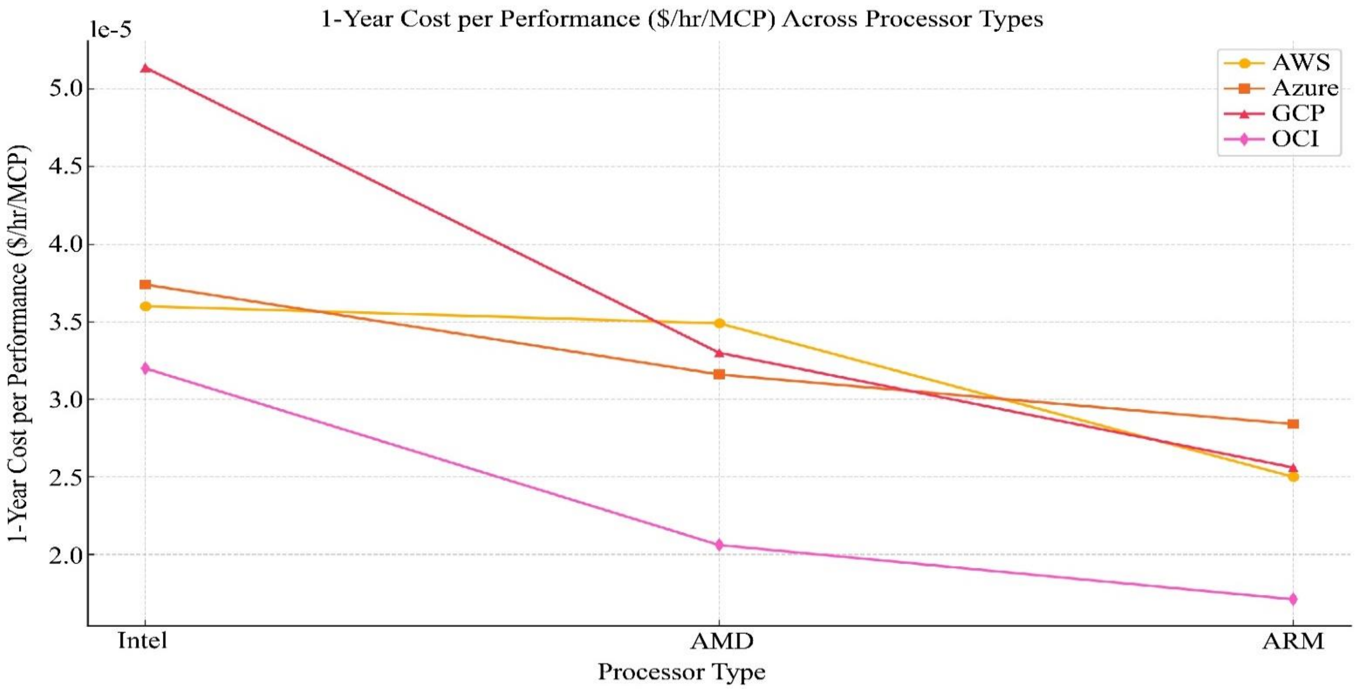

Fig. 3 1-Year cost per performance

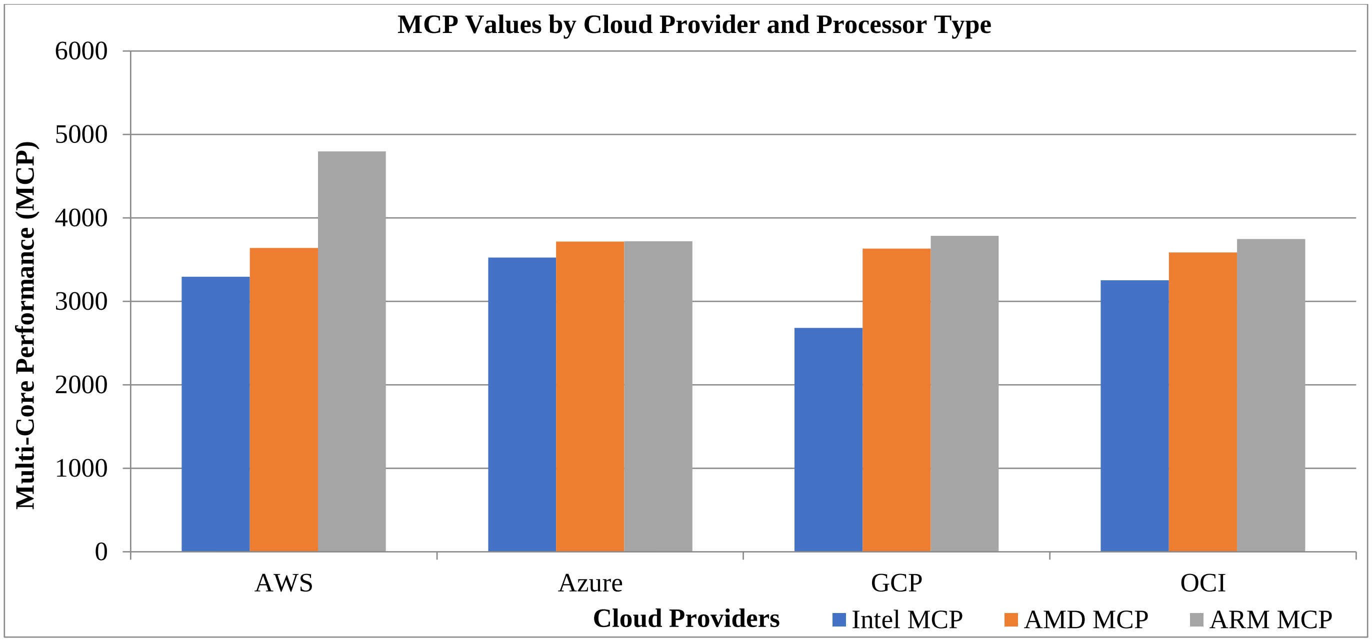

Fig. 4 MCP values by cloud provider and processor type

5. Future work

This paper measures comparable general purposed compute instances across large-cap cloud providers, namely AWS, Azure, GCP and an emerging cloud provider, OCI. The paper measures money spent per CPU performance point per hour. However, the view for this paper is only limited to CPU performance. Many real-world applications might rely not only on CPU performance but also on GPU performance. There are different instance types for cloud providers providing GPU support. Also, many workloads run on a cluster of virtual nodes where network bandwidth is key. Hence, more work can be done here to identify major systems of key importance for different applications, assigning a score to the different systems. Examples include GPU performance, memory performance and storage performance, having a mean average for these systems, assigning an overall score and calculating the per hour money spent for a unit of performance score. This can help enterprises to know which cloud provider is the best choice for their money. Even after we measure the entire system's performance, as the earlier paragraph describes, some cloud providers may still be better than others for some enterprises because of the services they provide. E.g., Oracle database services can be best provided by OCI. Hence, more work can be done to segregate service use cases and measure the dollar spent per performance unit.

6. Conclusion

The comparative analysis of Intel, AMD, and ARM-based general-purpose instances across AWS, Azure, GCP, and OCI highlights the evolving landscape of cloud computing. ARM-based architectures, such as AWS Graviton3 and GCP Tau T2A, deliver exceptional price-performance ratios, making them a compelling choice for cost-sensitive, cloud-native workloads. AMD-based instances offer cost-efficiency for general-purpose tasks. Intel-based instances remain preferred for enterprise-grade applications requiring high single-threaded performance and compatibility with legacy software. Oracle Cloud Infrastructure (OCI) emerges as the cost leader across all architectures, particularly for AMD and ARM instances, offering an attractive proposition for budget-conscious organizations. AWS maintains a balance of performance and availability, while Azure and GCP cater to diverse workloads with competitive options. The findings underscore the importance of aligning cloud instance selection with workload requirements. Organizations should prioritize ARM for scalability and efficiency, AMD for cost-optimized workloads, and Intel for performance-critical applications. By understanding these trade-offs, enterprises can optimize performance and cost, ensuring their cloud infrastructure effectively supports evolving business needs.

References

[1] Amazon Web Services, AWS Pricing. [Online]. Available: https://aws.amazon.com/pricing/
[2] Microsoft Azure, Azure Pricing Calculator. [Online]. Available: https://azure.microsoft.com/en-us/pricing/
[3] Google Cloud Platform, Google Cloud Pricing. [Online]. Available: https://cloud.google.com/pricing/
[4] Oracle Cloud Infrastructure, OCI Pricing and Services. [Online]. Available: https://www.oracle.com/cloud/pricing/
[5] Geekbench, Processor Performance Benchmarks. [Online]. Available: https://browser.geekbench.com/

[6] A. Jassy, Introducing AWS Graviton3 Processors, AWS Blog. [Online]. Available: https://aws.amazon.com/blogs/aws/introducing-aws-graviton3-processors/

[7] Microsoft Azure, Dv5 Sizes Series, 2024. [Online]. Available: https://learn.microsoft.com/en-us/azure/virtual-machines/dv5-dsv5-series

[8] Google Cloud, Machine Families Resource and Comparison Guide. [Online]. Available: https://cloud.google.com/compute/docs/machine-types

[9] Oracle Cloud Infrastructure Documentation, Compute Shapes. [Online]. Available: https://docs.oracle.com/en-us/iaas/Content/Compute/References/computeshapes.htm

[10] Phoronix, Amazon EC2 M6i Performance for Intel Ice Lake in the Cloud Benchmarks, 2021. [Online]. Available: https://www.phoronix.com/review/ec2-m6i-icelake

[11] Amazon EC2 m6a.Large. [Online]. Available: https://browser.geekbench.com/v5/cpu/21610849

[12] Microsoft Azure, Dpsv5 Sizes Series. [Online]. Available: https://docs.microsoft.com/azure/virtual-machines/dpsv5-dpdsv5-series

[13] Google Cloud, CPU Platforms. [Online]. Available: https://cloud.google.com/compute/docs/cpu-platforms

[14] OCI, Compute Shapes. [Online]. Available: https://docs.oracle.com/en-us/iaas/Content/Compute/References/computeshapes.htm#vm-standard

[15] Geekbench6. [Online]. Available: Geekbench 6 - Cross-Platform Benchmark

[16] Cloud Market Share. [Online]. Available" https://siliconangle.com/2024/11/16/cloud-market-share-shows-vendors-eyeing-1t-opportunity

[17] Oracle Announces Fiscal 2025 First Quarter Financial Results. [Online]. Available" https://investor.oracle.com/investor-news/news-details/2024/Oracle-Announces-Fiscal-2025-First-Quarter-Financial-Results/default.aspx

[18] Future Looks Good for Oracle on Back of Cloud and AI: Times. [Online]. Available: https://www.thetimes.com/business-money/markets/article/future-looks-good-for-oracle-on-back-of-cloud-and-ai-8q78nvdtt

[19] Synergy, Cloud Market Growth Stays Strong in Q2 While Amazon, Google and Oracle Nudge Higher, 2024. [Online]. Available: https://www.srgresearch.com/articles/cloud-market-growth-stays-strong-in-q2-while-amazon-google-and-oracle-nudge-higher

[20] Xinghan Chen et al., "X86 vs. ARM64: An Investigation of Factors Influencing Serverless Performance," *Proceedings of the 9th International Workshop on Serverless Computing*, pp. 7-12, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[21] S. Vignesh Bharadwaj, and Chetan Kumar Vudadha, "Evaluation of x86 and ARM Architectures using Compute-Intensive Workloads," *2022 IEEE International Symposium on Smart Electronic Systems (iSES)*, Warangal, India, pp. 586-589, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[22] Vasco Nuno Caio dos Santos, "*RISC vs. CISC – The Post-RISC*," University of Minho, pp. 105-112, 2024. [Google Scholar] [Publisher Link]

[23] Sebastian Kmiec et al., "A Comparison of ARM Against x86 for Distributed Machine Learning Workloads," *Performance Evaluation and Benchmarking for the Analytics Era*, pp. 164-184, 2018. [CrossRef] [Google Scholar] [Publisher Link]

[24] On Power Efficiency: x86 vs. ARM, pp. 1-28, 2022. [Online]. Available: https://indico.cern.ch/event/1128343/contributions/4787174/attachments/2412950/4129612/PowA_GridPP47.pdf

[25] Rafael Vidal Aroca, and Luiz Marcos Garcia Gonçalves, "Towards Green Data Centers: A Comparison of x86 and ARM Architectures Power Efficiency," *Journal of Parallel and Distributed Computing*, vol. 72, no. 12, pp. 1770-1780, 2012. [CrossRef] [Google Scholar] [Publisher Link]

[26] Emily Blem, Jaikrishnan Menon, and Karthikeyan Sankaralingam, "Power Struggles: Revisiting the RISC vs. CISC Debate on Contemporary ARM and x86 Architectures," *2013 IEEE 19th International Symposium on High Performance Computer Architecture (HPCA)*, 2013. [CrossRef] [Google Scholar] [Publisher Link]