

# A Comparative Analysis of Cloud Computing Services: AWS, Azure, and GCP

Ahmad A. A. Alkhatib, Ameen Shaheen Rand N Albustanji,

[Ahmad.Alkhatib@zuj.edu.iq](mailto:Ahmad.Alkhatib@zuj.edu.iq) [AmeenShaheen@zuj.edu.iq](mailto:AmeenShaheen@zuj.edu.iq) [Randalbustanji@gmail.com](mailto:Randalbustanji@gmail.com)

**Abstract-** This paper presents a comprehensive comparative analysis of three leading cloud computing platforms: Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP). By applying five different benchmarks, including latency, framework, RAM, GPU, and CPU (single and multi-core), the paper conducts an in-depth exploration and comparison of the virtual machines supplied by each platform. Additionally, it reviews features, services, and key pricing models, providing valuable insights for organizations seeking to choose the most suitable cloud provider for their specific needs. The analysis covers various dimensions, such as regions, availability zones, computing, storage, database, networking, security, machine learning, IoT, serverless computing, containers, big data, analytics, application development, content delivery networks (CDNs), hybrid/multi-cloud support, and more. By evaluating the strengths and weaknesses of AWS, Azure, and GCP across these dimensions, the paper offers a detailed guide to aid decision-making processes in selecting a cloud platform. The findings highlight the unique offerings of each platform, enabling organizations to leverage cloud technologies effectively. In conclusion, this review paper serves as a valuable resource for companies and IT professionals navigating the increasingly complex cloud computing landscape. By shedding light on the nuances of AWS, Azure, and Google Cloud, the paper empowers organizations to make informed decisions and maximize the potential of cloud services to meet their specific business requirements.

**Keywords:** Amazon Web service, Cloud Computing, Google Cloud Platform, Microsoft Azure, Virtual Servers.

## 1. Introduction

In recent years, the realm of cloud computing has undergone unprecedented growth and transformation, reshaping the landscape of modern technology infrastructure[1]. This burgeoning field has not only revolutionized the way organizations manage their data and applications but has also spurred intense competition among industry players vying for market dominance. As businesses increasingly migrate their operations to the cloud[2], understanding the intricacies of cloud computing services, including pricing models and competitive strategies, has become paramount.

We examine a deep comparison of these platforms through two main approaches. First, we use five Benchmarks (latency, frame, RAM, GPU, and CPU - single and multi-core) to assess the performance of each provider's virtual machines. Secondly, we conduct a broad feature analysis that includes different services, pricing models, and functions across different dimensions. These dimensions include areas, availability, and basic functions such as storage, database management, security, and advanced services such as machine learning and big data analytics. We also explore support for app development, content delivery networks (CDNs), hybrid/multi-cloud environments, and more.

Informed decision-making in cloud service selection necessitates a nuanced understanding of the unique value propositions and trade-offs associated with each platform. Leveraging insights gleaned from scholarly research, industry reports, and empirical studies, this review paper seeks to distill the complexities of cloud computing into actionable intelligence. By critically evaluating the strengths and limitations of AWS, Azure, and GCP across various domains, this paper equips organizations with the knowledge needed to make informed decisions aligned with their strategic objectives. Ultimately, the goal is to empower businesses and IT professionals to harness the full potential of cloud technologies in driving innovation and achieving competitive advantage in the digital era. Moreover, as cloud computing continues to evolve, the need for scalability, flexibility, and robust security features has become increasingly crucial for organizations across various industries. Each of the leading cloud platforms—AWS, Azure, and GCP—offers a distinct combination of tools and services tailored to address these needs. However, the challenge lies in selecting the platform that best aligns with an organization's specific requirements, whether it be for general computing purposes, data-intensive tasks, or specialized applications like machine learning and IoT. This paper delves into these aspects,

providing a detailed comparison that goes beyond surface-level features to uncover the strategic advantages and potential drawbacks of each platform.

In an era where digital transformation is a key driver of business success, the choice of a cloud provider can significantly impact an organization's operational efficiency, cost structure, and innovation capability. By offering a thorough analysis of AWS, Azure, and GCP, this paper aims to bridge the knowledge gap for decision-makers and IT professionals. It underscores the importance of aligning cloud service choices with long-term business goals, ensuring that the selected platform not only meets current needs but also supports future growth and technological advancements. Through this comprehensive examination, the paper aspires to contribute to more informed, strategic decisions in the rapidly evolving cloud computing landscape.

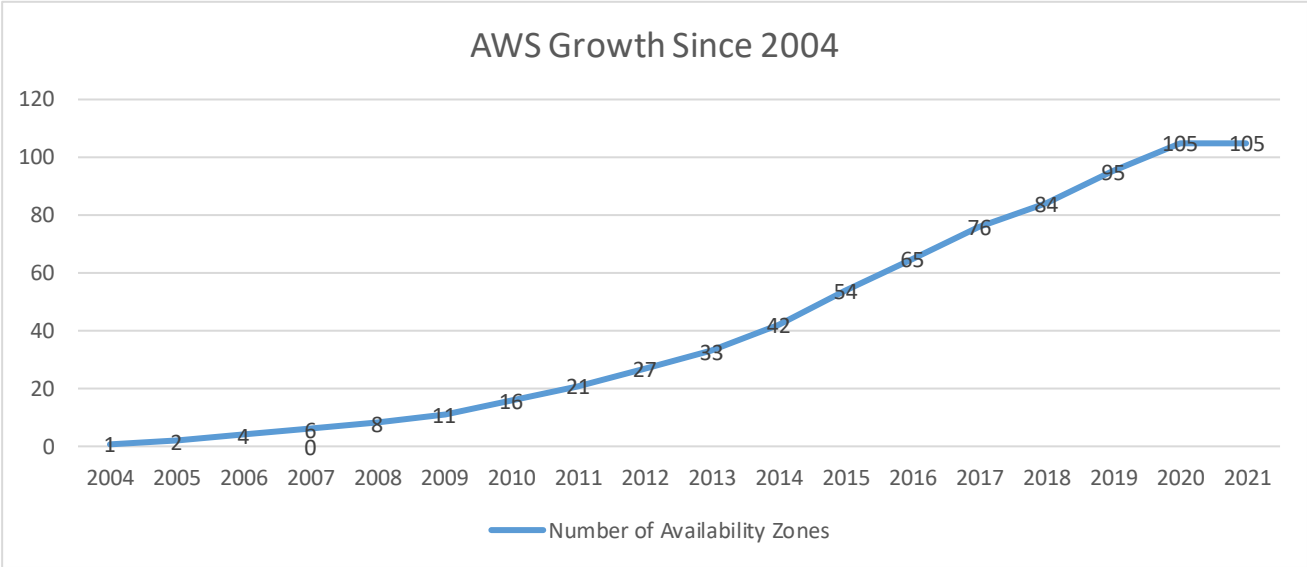
## **2. Cloud Service Providers**

Amazon Web Services (AWS) has been a key player in the cloud computing market since its inception in 2006. Offering a comprehensive suite of services, AWS caters to diverse computing needs including cloud storage, database management, analytics, networking, Internet of Things (IoT), mobile computing, and enterprise solutions [3].

These services empower organizations to accelerate their growth, optimize costs, and scale their operations effectively. With its extensive experience and robust infrastructure, AWS has established itself as a leading cloud platform globally. As a testament to its expansive reach, AWS boasts 105 availability zones spread across the world.[4] This widespread availability ensures high reliability and redundancy, enabling businesses to deploy their applications and services with confidence. With its longstanding presence and continuous innovation, AWS remains a preferred choice for organizations seeking scalable and reliable cloud solutions to drive their digital transformation initiatives.

In addition to its expansive service offerings and global reach, AWS is renowned for its commitment to security and compliance, which are critical factors for organizations operating in regulated industries. AWS provides a robust security framework that includes data encryption,

identity and access management (IAM), and compliance certifications across various international standards, such as ISO 27001, SOC 1/2/3, and GDPR. This focus on security ensures that enterprises can confidently manage sensitive data and meet stringent regulatory requirements. Furthermore, AWS's continuous investment in cutting-edge technologies, such as artificial intelligence (AI), machine learning (ML), and serverless computing, demonstrates its dedication to staying at the forefront of cloud innovation. These advancements not only enhance the capabilities of AWS but also provide businesses with the tools necessary to stay competitive in an increasingly digital world.




**Figure 1: AWS Growth Since 2004**

Google Cloud Platform (GCP) emerged in 2011 as Google's venture into the realm of cloud computing, offering an array of services to meet the diverse needs of its customers. With a robust suite of offerings encompassing storage, big data solutions, databases, analytics, cloud AI, networking, mobile computing, development tools, management tools, Internet of Things (IoT), cloud security, and data transfer, GCP provides a comprehensive ecosystem for organizations seeking scalable and innovative cloud solutions [5]

GCP's inception marked a significant milestone in Google's commitment to delivering cutting-edge technology services to its user base. In terms of infrastructure, the Google Cloud Platform boasts 121 availability zones strategically distributed across the globe. This extensive network of

availability zones ensures high availability and reliability, enabling businesses to deploy their applications and services with confidence while leveraging Google's robust infrastructure and global footprint. GCP's commitment to providing a secure, scalable, and agile cloud platform positions it as a formidable player in the competitive cloud computing landscape, catering to the evolving needs of businesses across industries[6].



**Figure 2:** Google cloud platform (GCP) since 2011

Microsoft Azure, launched in 2010, represents Microsoft's comprehensive cloud computing platform designed to provide a wide array of services to businesses and developers [7]. Azure offers a diverse range of cloud services, including computing, storage, databases, analytics, networking, machine learning, AI, IoT, security, and developer tools. As one of the leading cloud platforms, Azure enables organizations to build, deploy, and manage applications and services with flexibility, scalability, and reliability.

With a global presence spanning over 60 regions and 140 countries, Microsoft Azure provides extensive coverage and high availability to meet the diverse needs of its users worldwide. The platform's integration with Microsoft's extensive suite of enterprise products and services further enhances its appeal, allowing seamless interoperability and enabling organizations to leverage

their existing investments in Microsoft technologies [8].

Microsoft Azure's commitment to security, compliance, and hybrid cloud capabilities positions it as a preferred choice for businesses across industries seeking to harness the power of cloud computing while maintaining control over their data and infrastructure. As a testament to its scalability and innovation, Azure continues to evolve, offering cutting-edge solutions to address the evolving needs of modern businesses and developers in the digital age [7].

### 3. Cloud Service Platforms Architecture components:

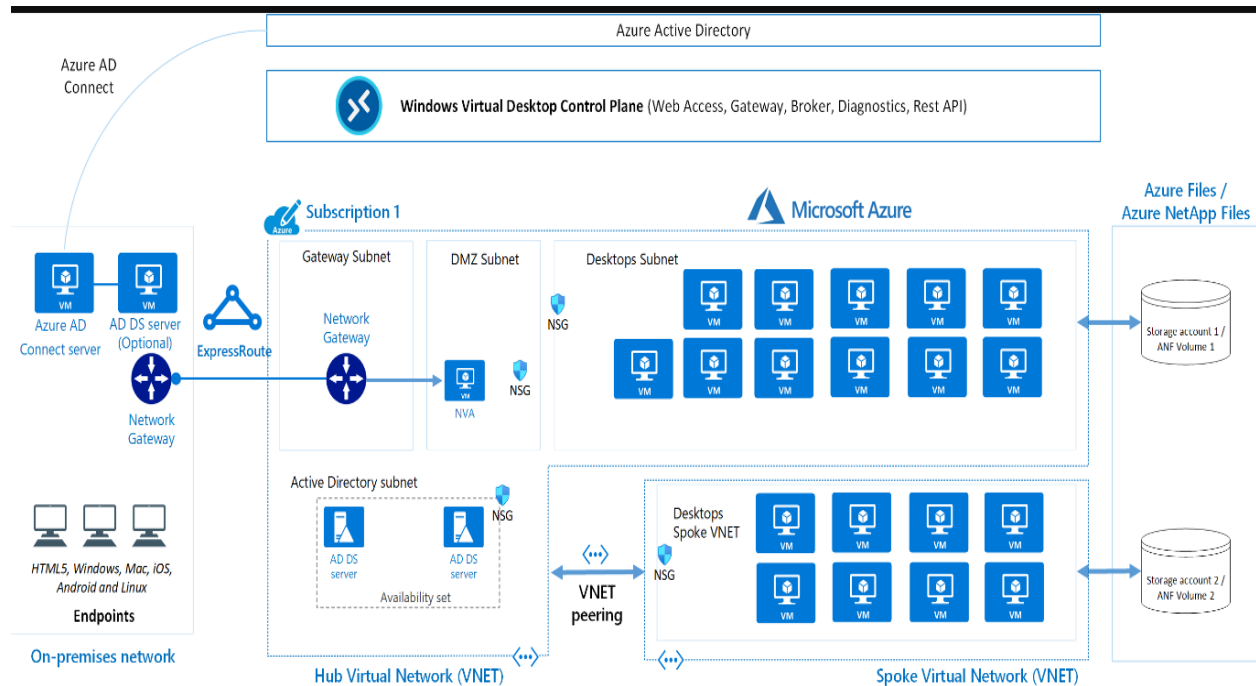


[9] Figure 3: AWS Architecture

AWS CLOUD incorporates a variety of data ingestion mechanisms to effectively manage and process user multimedia data for enrichment purposes. Specifically, Amazon Simple Queue Service (Amazon SQS) [9] is employed to oversee data-ingestion event queues, while Amazon Simple Notification Service (Amazon SNS) facilitates asynchronous communication

between Amazon EventBridge and the enrichment services. AWS Lambda orchestrates critical functions such as event queue polling, event propagation to EventBridge, and data processing across the EventBridge-enrichment services interface. Notably, [10] serves as the conduit for directing ingested data to the suite of enrichment services.

The suite of data enrichment services encompasses several specialized tools: Amazon Textract is leveraged for text extraction, Amazon Comprehend for entity identification, Amazon Rekognition for object and scene recognition in multimedia content, Amazon Transcribe for audio-to-text transcription, Amazon Translate for multilingual text translation, and AWS Elemental MediaConvert for video transcoding. In terms of data storage and streaming, the system employs Amazon DynamoDB for storing system data and metadata, while Amazon S3 serves as the repository for both the landing zone and enriched data storage. Furthermore, Amazon Kinesis Data Streams facilitates the real-time streaming of enriched data, while Amazon Kendra facilitates enriched data searching utilizing natural language processing and advanced search algorithms. Lastly, AWS Lake Formation is utilized to establish and secure the underlying data lake infrastructure[11].



**Figure 4: Azure Architecture**

The architecture of Microsoft Azure cloud operates through a sophisticated interplay of several crucial components. Firstly, Azure relies on a global network of datacenters strategically positioned across the world, comprising various regions and availability zones. These datacenters host the essential hardware resources necessary for running applications and services, including servers for computing power, storage, and memory, along with networking devices facilitating communication between components and users.

On the software and virtualization front, Azure employs a hypervisor layer to optimize resource utilization and flexibility. This software layer allows multiple virtual machines (VMs) to operate on a single physical server, ensuring efficient resource sharing and isolation. VMs, in turn, serve as software representations of physical computers, offering users dedicated computing environments pre-configured with operating systems and software. Additionally, Azure Fabric orchestrates and manages the entire infrastructure, including VMs, storage, networking, and security, performing crucial tasks like resource provisioning and load balancing[12].

Moreover, Azure provides an extensive range of cloud services built atop its infrastructure, spanning computing, storage, networking, databases, security, and AI, among others, enabling developers and businesses to deploy applications seamlessly without managing underlying infrastructure intricacies [13]. Management tools such as the Azure Portal, command-line utilities, and software development kits (SDKs) offer diverse avenues for resource management and interaction, allowing users to configure services, monitor performance, and access various functionalities efficiently.

Key attributes of Azure architecture include its scalability and elasticity, allowing for dynamic resource scaling based on workload demands. Moreover, Azure prioritizes reliability and security, offering redundant components and disaster recovery features to ensure service uptime and data protection. With its pay-as-you-go model, users only pay for the resources consumed, making Azure a cost-effective solution for a wide array of workloads[8].

#### **4. Dynamics of Major Cloud Service Providers in the Competitive Landscape**



In the contemporary technological landscape, the domain of cloud computing stands as a dynamic arena marked by intense competition among major players. Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP) have emerged as prominent contenders, offering diverse services encompassing infrastructure as a service (IaaS), platform as a service (PaaS), and software as a service (SaaS). This section provides a comprehensive analysis of their respective market standings and contributions to the evolving cloud computing paradigm [5].

As of the latest update in January 2022, Amazon Web Services (AWS) has consistently upheld its preeminence as the leading cloud service provider, boasting a robust suite of offerings ranging from computational power to storage, databases, machine learning, and analytics. AWS's global infrastructure footprint, coupled with its penchant for innovation, has solidified its dominance in the market.

Microsoft Azure, a significant player in the field, has experienced substantial growth owing to its seamless integration with Microsoft's existing suite of products and services. Azure's expansive portfolio encompasses virtual machines, application services, artificial intelligence (AI) and machine learning tools, and enterprise solutions, positioning it as a formidable competitor in the cloud computing arena. Similarly, Google Cloud Platform (GCP) [5] has emerged as a key player renowned for its expertise in data analytics, machine learning, and containerized applications. Leveraging Google's advanced infrastructure and cutting-edge technologies, GCP has garnered a substantial customer base, particularly among industries necessitating sophisticated data processing and analysis capabilities[9], [14].

Moreover, the competitive landscape of cloud computing is continually evolving, with each of these providers carving out unique niches to differentiate themselves in the market. AWS's early entry into the cloud space has allowed it to build a comprehensive ecosystem, particularly appealing to enterprises seeking a wide range of services and global scalability. In contrast, Microsoft Azure's strength lies in its deep integration with enterprise software solutions like Office 365, making it a natural choice for businesses already invested in the Microsoft ecosystem. Meanwhile, Google Cloud Platform (GCP) leverages its leadership in AI and machine learning to attract organizations focused on innovation and big data analytics. As cloud

adoption accelerates across industries, the distinct strengths of AWS, Azure, and GCP continue to shape their respective roles and competitive strategies within the cloud market.

Beyond the triad of AWS, Azure, and GCP, other notable contenders such as IBM Cloud, Oracle Cloud, and Alibaba Cloud have also made significant strides, each offering distinct features tailored to specific customer requirements. The competitive landscape among these cloud service providers remains a catalyst for innovation and cost-efficiency within the industry, fostering an environment conducive to the adoption of cloud computing across diverse applications and workloads. As businesses and consumers harness the potential of cloud technologies, the dynamics of the market continue to evolve, shaping the future trajectory of cloud computing.

**Comparative Analysis of Cloud Providers: AWS, Azure, and GCP**

**Table 1:** Theoretical Comparative Analysis

<b>Feature</b>	<b>AWS</b>	<b>Azure</b>	<b>Google Cloud</b>
Number of Regions	36	60+	27
Number of Availability Zones	105	116+	82
Pricing Model	On-demand,	On-demand,	On-demand
Compute Services	EC2 (Elastic Compute Cloud)	Virtual Machines, Azure Functions	Compute Engine, Kubernetes Engine
Storage Services	S3 (Simple Storage Service), EBS (Elastic Block Store)	Blob Storage, Azure Files	Cloud Storage, Cloud Filestore
Database Services	RDS (Relational Database Service), DynamoDB (NoSQL)	SQL Database, Azure Cosmos DB (NoSQL database)	Cloud SQL, Cloud Spanner (NoSQL database)

	database)		
Networking Services	VPC (Virtual Private Cloud), Direct Connect	Virtual Networks, ExpressRoute	VPC (Virtual Private Cloud), Cloud Interconnect
Management Services	CloudWatch, AWS Config	Azure Monitor, Azure Automation	Cloud Monitoring, Stackdriver Logging
Security Services	IAM (Identity and Access Management), Security Hub	Azure Active Directory, Azure Security Center	Cloud IAM, Security Command Center
Machine Learning & AI Services	SageMaker, Amazon Rekognition	Azure Machine Learning, Azure Cognitive Services	AI Platform, Vertex AI
Analytics Services	QuickSight, Amazon Kinesis	Power BI, Azure Synapse Analytics	Looker, Cloud Dataflow
Content Delivery Network (CDN)	CloudFront	Azure CDN	Cloud CDN
Ease of Use	Considered relatively easy for beginners	Can be complex for beginners due to extensive services	Considered developer-friendly, strong focus on open source
Best suited for	Wide range of workloads, diverse set of services	Microsoft-centric workloads, hybrid cloud deployments	Cost-sensitive workloads, AI/ML & big data projects, containerized applications
Benefits	* Extensive range of services * Mature and reliable platform * Flexible pricing options	* Strong integration with Microsoft products * Good choice for hybrid cloud deployments * Large	* Competitive pricing for sustained workloads * Strong focus on AI/ML and big data * Open-

		global footprint	source friendly
Limitations	* Can be complex for beginners * Vendor lock-in can be a concern * Not always the most cost-effective option	* Pricing structure can be complex * Can be less user-friendly than AWS * May not be ideal for non-Microsoft workloads	* Smaller global footprint compared to AWS and Azure * Limited legacy application support

**Cloud Provider VM Benchmarking: Standardizing Specifications for Comparison**

To begin the process of accurately comparing each of them using benchmarks, it is necessary to build the virtual machine on each of the three clouds, the minimum specifications that must be available in each device to provide a reasonable average service at a reasonable performance were sought, and the possibility of these specifications being available by the three providers was considered as a starting point. [15]The result is shown in Table 2

**Table 2** Virtual Machine Specifications

Component	Specification
Graphics Cards	NVIDIA GeForce RTX 30 series or AMD Radeon RX 6000 series
Processor	Intel Core i7 or AMD Ryzen 7
RAM	16GB or more
Storage	Solid State Drive (SSD) with at least 250GB
Framerate	30,60,90 fps
Latency	Ms. 100<

**Benchmark data selection**

In this step, the Benchmark data that will be used to measure and compare the performance of virtual devices were researched and selected. This involved identifying specific performance metrics and benchmarks relevant to the objectives of the study, as follows CPU utilization, RAM performance, network latency, and Framerate. The benchmark data selected, as guided by the recommendations of prior studies, data will serve as a standardized framework for assessing the quality and performance of the virtual machines deployed on each cloud provider, facilitating informed decision-making and analysis.

**Table 3** Benchmark Data Resources

<b>Component</b>	<b>Benchmark</b>	<b>Focus</b>	<b>References</b>
CPU	Geekbench6	Single Core	[16]
CPU	Geekbench6	Multi Core	[16]
GPU	UserBenchmark	Graphics processing power for gaming and demanding workloads	[17]
RAM	PerformanceTest	Memory stability and error detection	[18]
Framerate	Heaven benchmark	Real-time measurement of frames rendered per second in games or applications	[19]
Latency	Ping (Network Latency)	Delay in data transfer over a network	[20]

## Results

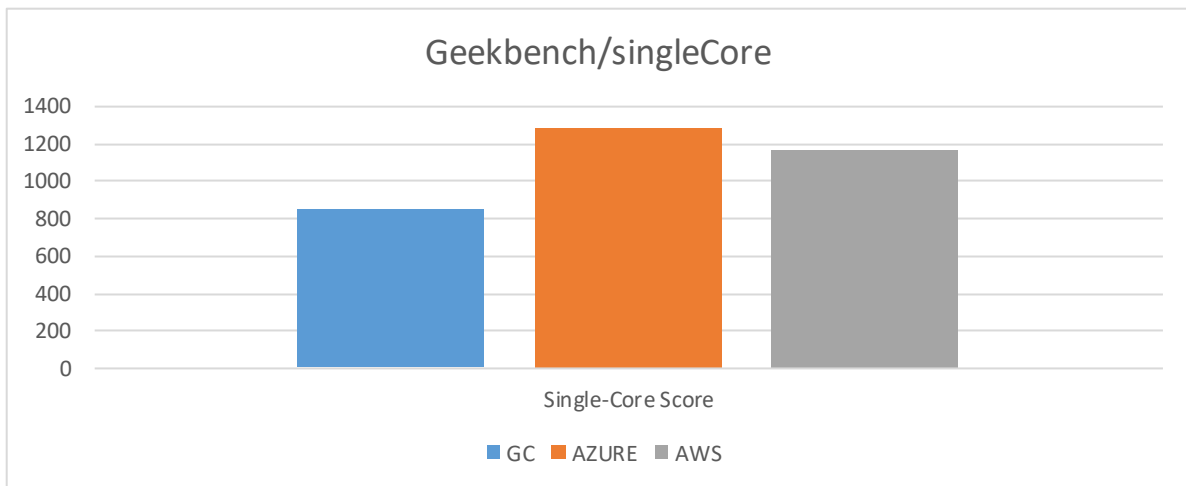
Experiments were conducted on three cloud service providers, Amazon Web Service, Google Cloud, and Microsoft Azure, to compare their performance when building a VM with the same specifications, and the same provisioning area, Qatar-Doha, as mentioned previously in **Table 2**, a set of benchmark data was chosen to examine different parts of the systems

### Comparative Analysis Result:

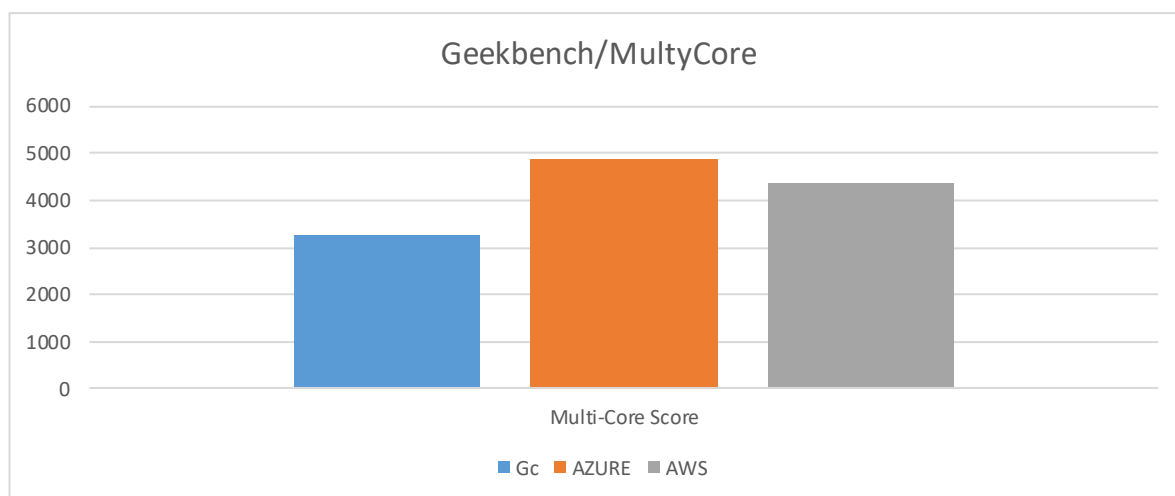
## Geekbench

Geekbench is a widely used benchmark tool that evaluates the performance of computer hardware [21], **Figure 5** shows the results of applying Geekbench on the three CPUs as mentioned earlier using a single core.

Among the three platforms, Azure exhibits the highest Single-Core Score, reflecting superior performance for a single CPU core. AWS achieves a slightly lower Single-Core Score but still surpasses the generic computer in performance. Conversely, the Generic Computer records the lowest score, implying relatively diminished performance for a single CPU core in comparison to both Azure and AWS cloud computing platforms



**Figure 5:** Geekbench/single core



The performance disparities across cloud computing providers AWS, Azure, and Google Cloud, particularly at the Single and Multi-Core Performance levels, can be linked to the underlying CPU architecture, pipeline design, and processor scheduling mechanisms. Intel Xeon and Xeon Platinum 8275CL CPUs, prevalent in AWS and Google Cloud, leverage Intel's x86 architecture, known for its balanced performance and reliability. Meanwhile, Azure employs AMD EPYC 7V12 CPUs, which harness AMD's Zen microarchitecture, emphasizing high core counts and memory bandwidth. The


**Figure 6:** Geekbench/Multi core

fferences in pipeline design, with Intel focusing on performance and efficiency while AMD prioritizes throughput and latency mitigation, along with tailored processor scheduling algorithms, contribute to Azure providing the best performance due to the advantages of AMD's architecture for cloud workloads. Conversely, Google Cloud, utilizing Intel Xeon CPUs, may experience relatively lower performance than AWS and Azure due to variations in architectural optimizations and design choices, despite all platforms utilizing CPUs with similar core/thread specifications.

**PassMark Performance Test:**

This benchmark is used to measure the performance of the CPU, 2D and 3D, Memory, and Disk, and provides numerical scores that can be compared to other systems, helping users evaluate the relative performance of their hardware. Performance Test's combination of comprehensive testing, cross-platform support, detailed results, stability, and regular updates distinguishes it as a popular choice among users for assessing computer hardware performance.

**Figure 6** shows the results of running this benchmark on the three-selected memory.




**Figure 2:** PassMark/ Memory Test

The chart highlights the differences in memory performance across Google Cloud, Microsoft Azure, and Amazon Web Services, with AWS showing the highest memory performance, followed by Azure, then Google Cloud. One reason for the difference in results is that After examining the memory performance, it was found that AWS provides the best memory operation score, which is equal to 60, followed by Azure = 54, and Google Cloud at 46.

### Heaven Benchmark



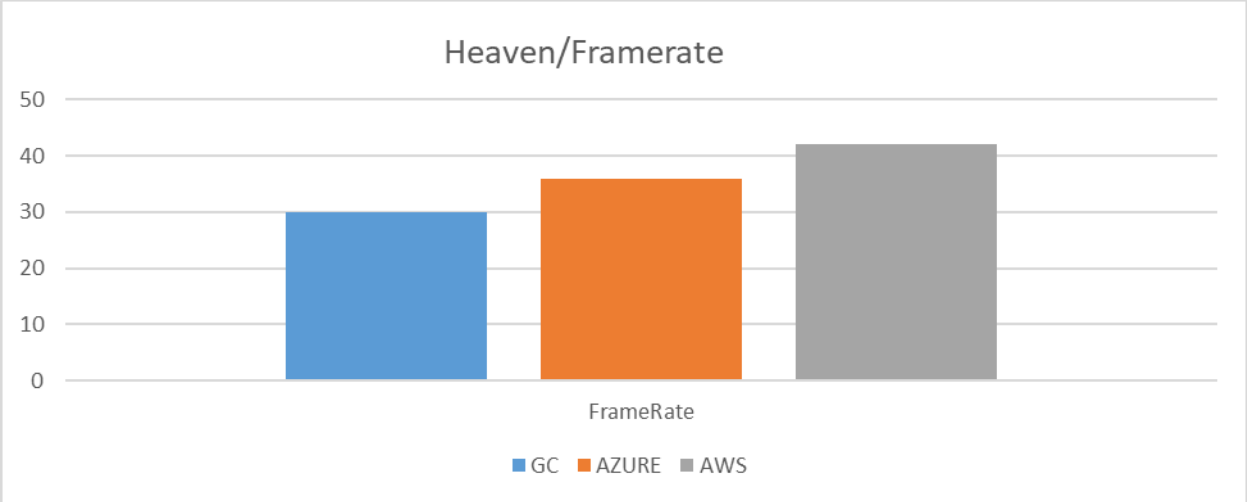
This benchmark is used to test the performance and stability of a computer's graphics processing unit (GPU). By rendering complex scenes with a high number of polygons and textures, the benchmark can identify potential bottlenecks in a system's graphics processing capabilities, **Figure 7** shows the results of running this benchmark on the three selected GPUs to



compare between GPU usage and FrameRate

The GPU usage indicates the percentage of GPU resources being utilized by applications or processes running on each cloud platform. From the data, we can see that AWS has the highest GPU usage at 21%, followed by Azure at 19%, and Google Cloud at 13%. This suggests that applications or workloads running on AWS are utilizing a larger portion of GPU resources compared to Azure and Google Cloud

**Figure 3:** Heaven/GPU Usage Test

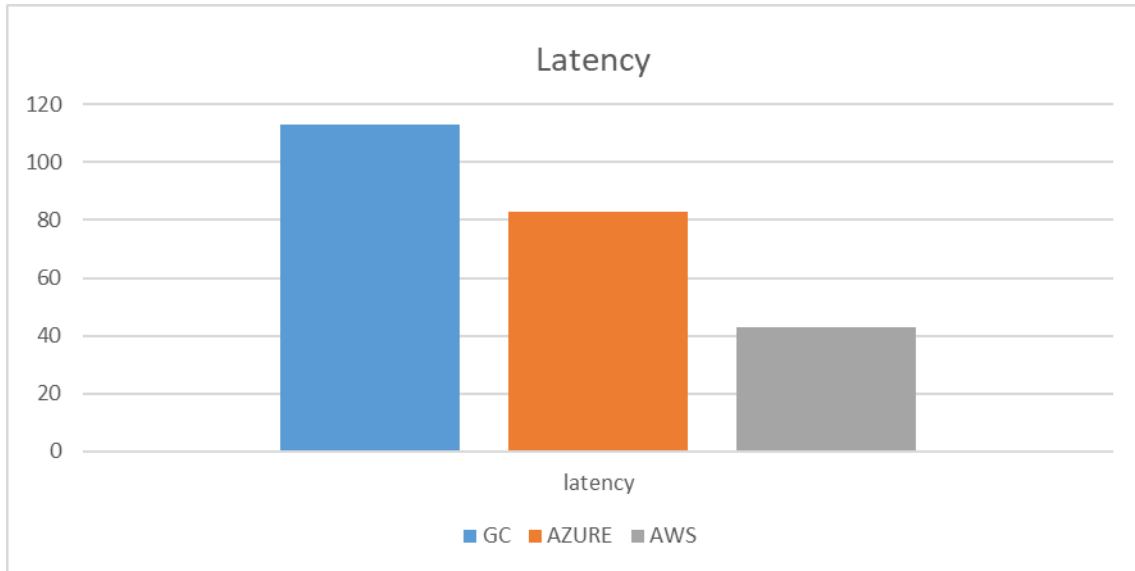


**Figure 4:** Heaven /Framerate

The frame rate represents the number of frames rendered per second in graphical applications or simulations. Higher frame rates typically indicate smoother and more responsive graphics performance. In this case, AWS achieves the highest frame rate at 42 fps, followed by Azure at 36 fps, and Google Cloud at 30 fps.

**Latency tests- 4G network**

Network responsiveness in the Middle East was assessed using ping tests on three major cloud providers: Azure, Google Cloud, and AWS. All three servers are located in the Middle East, with Google Cloud and AZURE in Doha, Qatar, and AWS in Bahrain. Since AWS is unavailable in Doha, Bahrain was chosen as the closest alternative. A ping test acts like a digital tap, sending a signal to a server and measuring how long it takes to receive a response. This response time, known as latency, indicates how quickly data travels between your device and the server. Lower latency is generally preferred.



**Figure 5:** Latency Test Result

Based on the ping test results, the AWS server in Bahrain appears to offer the best performance in the Middle East with a latency of 42.8 milliseconds (ms). This latency indicates the speed at which data travels between your device and the server. Lower latency translates to faster communication. Compared to Azure (83 ms) and Google Cloud (213 ms), AWS Bahrain exhibits the quickest response time.

**Table 4** comparative analysis result summary

Factor	Description	Test	Best Performer
Single-Core CPU Performance	Performance of a single CPU core	Geekbench6 (Single Core)	Microsoft Azure
Multi-Core CPU Performance	Performance of all CPU cores working together	Geekbench6 (Multi Core)	Microsoft Azure
Memory Performance	Speed and efficiency of data access	PassMark (PerformanceTest64)	Amazon Web Services

GPU Usage	Percentage of GPU resources utilized	Heaven Benchmark	Amazon Web Services
Frame Rate	Number of frames rendered per second by GPU	Heaven Benchmark	Amazon Web Services
Network Latency	Speed of data travel between user and server	Ping Test	Amazon Web Services

## Conclusion

This in-depth analysis compared the virtual machine performance of industry giants Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP) using a variety of benchmarks. We focused on crucial aspects of different computing needs.

For single-core workloads that emphasize processing a single task at maximum speed, Microsoft Azure emerged victorious. This indicates Azure's virtual machines excel at handling tasks like scientific calculations or single-threaded applications. However, if your applications leverage multiple cores for parallel processing, Azure maintains its lead with the top spot in multi-core CPU performance.

On the other hand, for applications requiring fast data access and manipulation, memory speed is critical. Here, AWS takes the crown. Benchmarks revealed that AWS virtual machines boast superior memory access speeds, making them ideal for tasks like in-memory databases, real-time analytics, and complex server-side applications.

Graphics Processing Units (GPUs) accelerate workloads demanding intensive graphical computations. The analysis clearly showed AWS's strength in this area. AWS virtual machines demonstrated superior utilization of GPU resources, meaning they can effectively handle tasks like video editing, 3D rendering, and machine learning with larger datasets. This translates to

faster processing times for such workloads. Additionally, the analysis highlighted AWS's lead in frame rate, signifying its virtual machines can render graphics at a smoother and faster pace, crucial for applications like video game development and simulations.

Finally, low network latency ensures swift data transfer between your application and the cloud server. Here again, AWS shines. The results confirmed that AWS offers the fastest data transfer speeds among the three platforms. This translates to a more responsive user experience for applications sensitive to latency, such as real-time collaboration tools and high-frequency trading platforms.

Choosing the right cloud platform hinges on your specific needs. If raw CPU power is paramount, Microsoft Azure stands out. However, for applications demanding superior memory speed or intensive GPU workloads, AWS might be the preferable choice. Finally, for scenarios where low network latency is crucial, AWS offers the fastest data transfer speeds. Businesses should carefully evaluate their workload requirements and prioritize the benchmarks most relevant to their use case to select the optimal cloud platform.

#### Author Contributions

All authors contributed equally to all aspects of this paper. All authors have read and agreed to the published version of the manuscript.

#### Funding

This research received no external funding.

#### Data Availability Statement

Not applicable.

#### Acknowledgments

The authors appreciate the contributions of the previous researchers whose works are referenced in this paper.

## Conflicts of Interest

The authors declare no conflict of interest.

## References:

- [1] S. Maurya, T. Mufti, D. Kumar, P. Mittal, and R. Gupta, "A study on cloud computing: a review," in *Proceedings of the 2nd International Conference on ICT for Digital, Smart, and Sustainable Development, ICIDSSD 2020, 27-28 February 2020, Jamia Hamdard, New Delhi, India, 2021*.
- [2] P. Shrivastava and M. Damle, "Investment decision in cloud gaming-based businesses opportunities: An analysis of the cloud gaming industry," in *2022 International Conference on Decision Aid Sciences and Applications (DASA)*, IEEE, 2022, pp. 1224–1228.
- [3] S. Mathew and J. Varia, "Overview of amazon web services," *Amazon Whitepapers*, vol. 105, no. 1, p. 22, 2014.
- [4] P. Dubey and R. Raja, *A Beginners Guide to Amazon Web Services*. CRC Press, 2023.
- [5] M. Ramuka, *Data analytics with Google Cloud platform*. BPB Publications, 2019.
- [6] "Mendeley Reference Manager." Accessed: May 10, 2024. [Online]. Available: <https://www.mendeley.com/reference-manager/library/all-references>
- [7] V. A. Gandhi and C. K. Kumbharana, "Comparative study of Amazon EC2 and Microsoft Azure cloud architecture," *International Journal of Advanced Networking & Applications*, pp. 117–123, 2014.
- [8] S. R. Gundu, C. A. Panem, and A. Thimmapuram, "The Dynamic Computational Model and the New Era of Cloud Computation Using Microsoft Azure," *SN Comput Sci*, vol. 1, no. 5, p. 264, 2020.
- [9] "Message Queuing Service - Amazon Simple Queue Service - AWS." Accessed: May 10, 2024. [Online]. Available: <https://aws.amazon.com/sqs/>
- [10] V. S. V. M. Krishna, T. S. R. Vamsi, and S. Kavitha, "Automation of Forensic Analysis for AWS Aurora using EventBridge and Athena," in *2022 6th International Conference on Electronics, Communication and Aerospace Technology*, IEEE, 2022, pp. 262–269.
- [11] V. S. V. M. Krishna, T. S. R. Vamsi, and S. Kavitha, "Automation of Forensic Analysis for AWS Aurora using EventBridge and Athena," in *2022 6th International Conference on Electronics, Communication and Aerospace Technology*, IEEE, 2022, pp. 262–269.
- [12] K. Perumal, S. Mohan, J. Frnda, and P. B. Divakarachari, "Dynamic resource provisioning and secured file sharing using virtualization in cloud azure," *Journal of Cloud Computing*, vol. 11, no. 1, p. 46, 2022.
- [13] S. Galiveeti, L. Tawalbeh, M. Tawalbeh, and A. A. A. El-Latif, "Cybersecurity analysis: Investigating the data integrity and privacy in AWS and Azure cloud platforms," in *Artificial intelligence and blockchain for future cybersecurity applications*, Springer, 2021, pp. 329–360.

- [14] M. Calder, X. Fan, Z. Hu, E. Katz-Bassett, J. Heidemann, and R. Govindan, "Mapping the expansion of Google's serving infrastructure," in *Proceedings of the 2013 conference on Internet measurement conference*, 2013, pp. 313–326.
- [15] "Need a gaming laptop with at least the minimum requirements to play Elden Ring : r/GamingLaptops." Accessed: May 05, 2024. [Online]. Available: [https://www.reddit.com/r/GamingLaptops/comments/157u3sy/need\\_a\\_gaming\\_laptop\\_with\\_at\\_least\\_the\\_minimum/](https://www.reddit.com/r/GamingLaptops/comments/157u3sy/need_a_gaming_laptop_with_at_least_the_minimum/)
- [16] D. Silva, J. Rafael, and A. Fonte, "Virtualization Maturity in Creating System VM: An Updated Performance Evaluation," *International Journal of Electrical and Computer Engineering Research*, vol. 3, no. 2, pp. 7–17, 2023.
- [17] S. S. Nair and T. Santha, "High availability of kernel-based virtual machine using nested virtualization," *Measurement: Sensors*, vol. 26, p. 100712, 2023.
- [18] A. M. Potdar, D. G. Narayan, S. Kengond, and M. M. Mulla, "Performance evaluation of docker container and virtual machine," *Procedia Comput Sci*, vol. 171, pp. 1419–1428, 2020.
- [19] E. Kannegieser, D. Atorf, and R. Matrai, "Improving Emotion Detection for Flow Measurement with a High Frame Rate Video based Approach.," in *CSEDU (2)*, 2020, pp. 490–498.
- [20] D. A. Popescu and A. W. Moore, "Measuring network conditions in data centers using the precision time protocol," *IEEE Transactions on Network and Service Management*, vol. 18, no. 3, pp. 3753–3770, 2021.
- [21] S. Mccarthy, "Redmi Note 9T with MediaTek Dimensity 800U spotted on Geekbench.," *UWIRE Text*, p. 1, 2021.