Comparing the Performance and Cost of Different Cloud Service Providers for Small Applications: A Comparison between AWS and Azure

Nabiha Arif Polani
Dept. of Computer Science and Digital Technologies
University of East London
London, United Kingdom
nabihapolani@gmail.com

Umaima Haider
Dept. of Computer Science and Digital Technologies
University of East London
London, United Kingdom
umaima@uel.ac.uk

***Abstract*—This study assesses the performance and cost-effectiveness of Microsoft Azure and Amazon Web Services (AWS) for small-scale applications. As cloud computing becomes integral for developers and small businesses, this research provides detailed insights to help stakeholders select optimal cloud providers based on critical metrics, including CPU performance, memory usage, and pricing models. By conducting comprehensive performance evaluations under diverse scenarios and analyzing various pricing options—such as pay-as-you-go, reserved, and spot instances—this paper establishes a framework for maximizing cloud investments. The findings offer actionable guidance to stakeholders, enabling small businesses to make informed, scalable, and financially sound decisions regarding cloud infrastructure.**

***Index Terms*—Cloud Computing, Amazon Web Services (AWS), Microsoft Azure, Performance Evaluation, Cost Analysis, CPU Performance, Memory Usage, Pricing Models, Pay-as-you-go, Reserved Instances, Spot Instances, Small Applications**

I. Introduction

Cloud computing offers scalable, cost-effective solutions for businesses [4]. For small developers, choosing the right cloud provider is essential. This study compares Microsoft Azure and AWS, analyzing CPU performance, memory usage, and costs using Apache JMeter [12] benchmarking and provider pricing calculators. Real-world tests reveal each platform's strengths and limitations, providing insights for reliable, cost-effective cloud performance under varying workloads.

This paper is structured as follows: Section II reviews existing research; Section III outlines the methodology; Section IV presents results and analysis; and Section V concludes with recommendations for future research.

II. Existing Study

As of Q1 2024, AWS and Microsoft Azure together accounted for 56% of the global cloud market, affirming their status as leading cloud providers [13]. This section contrasts the two platforms across cost, performance, scalability, and security.

Azure tends to be more cost-effective in sectors such as education, while AWS is favored for its scalability and extensive service offerings for complex applications [1] [2] [3]. Key challenges in adoption include variable costs and rapid technological changes, with limited practical impacts from cost-optimization strategies like dynamic provisioning [4]. AWS's global reach and robust security features appeal to sectors like government and healthcare, whereas Azure's seamless integration with Microsoft services and hybrid cloud capabilities make it an ideal choice for organizations balancing cloud and on-premises solutions [5] [6]. Ultimately, platform selection hinges on specific business goals and infrastructure needs, with AWS excelling in IoT device management and Azure leading in data analysis and integration [7] [8].

III. Research Methodology

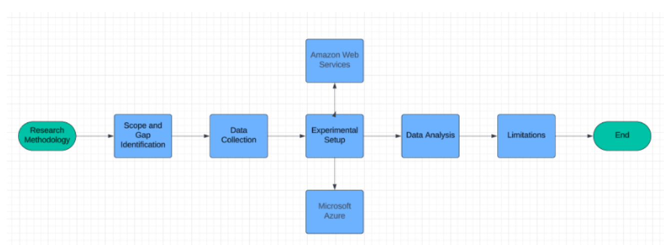

Fig. 1. Research Methodology

A. Scope and Gap Identification

This research compares Microsoft Azure and AWS for small-scale applications using Apache JMeter to evaluate CPU performance and memory utilization. Existing literature mainly addresses large-scale deployments, leaving small applications underexplored.

B. Data Collection

The analysis focuses on AWS and Azure, utilizing a Dockerized application, "Product App," built with React JS and Spring Boot as the benchmark [9]. Performance metrics are monitored via AWS CloudWatch and Azure Monitor, with cost data gathered through pricing calculators.

C. Experimental Setup

The "Product App" assesses performance through CRUD operations, using Docker containers with consistent specifications on both platforms.

1) **AWS:** Deployed in the US East region, using one EC2 instance and RDS database, monitored via CloudWatch, with costs estimated through the AWS Pricing Calculator [10].
2) **Microsoft Azure:** Runs on a single VM with SQL Server Express in the UK West region, monitored by Azure Monitor, with cost estimates from the Azure Pricing Calculator [11].

D. Data Analysis

The analysis evaluates cost-effectiveness and performance through load testing. Dashboards from CloudWatch and Azure Monitor visualize key metrics, while pricing calculators facilitate comparisons.

E. Study Limitations

The use of free-tier accounts reflects entry-level performance relevant to small-scale users. Limitations include regional differences affecting service availability and performance.

IV. Experimental Results

A. Performance Analysis

The performance analysis reveals significant differences between AWS and Azure under varying request loads. Load testing conducted with JMeter, utilizing POST (write) and GET (read) requests, offers insights into each platform's behavior as user demand escalates to 100, 300, and 600 users.

For POST requests, Azure outperformed AWS at lower loads (100 users) with 92.2 requests per second compared to AWS's 37.5. However, as user load increased, AWS exhibited greater scalability, handling 82.7 requests per second at 300 users versus Azure's 60.9. At the maximum load of 600 users, AWS maintained a higher throughput (41.6 requests per second) compared to Azure's 31.2, indicating AWS's superior consistency under high load.

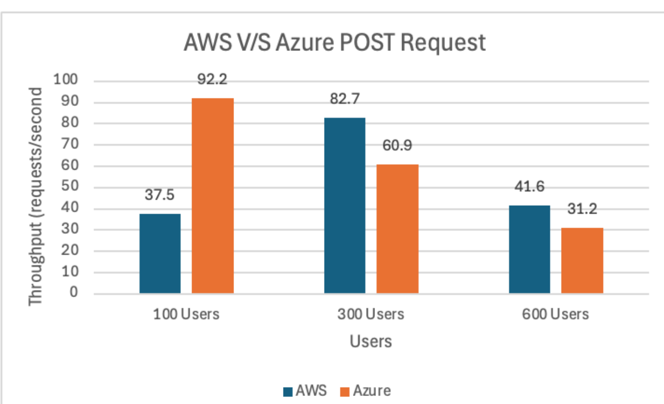

Fig. 2. Load Testing of POST Request

For GET requests, AWS led at 100 users with 54.7 requests per second, while Azure handled 37.5. At intermediate loads (300 users), Azure took the lead with 71.7 requests per second, but as the load increased to 600 users, AWS once again showed stronger performance, achieving 75.6 requests per second compared to Azure's 45.1. This indicates AWS's robustness in handling high volumes of read requests, whereas Azure performs well at intermediate levels.

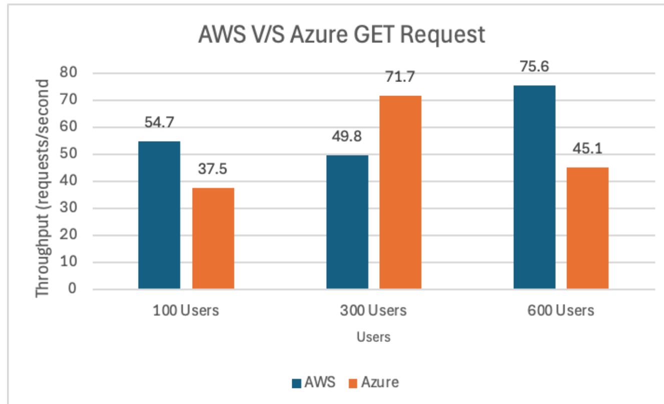

Fig. 3. Load Testing of GET Request

CPU utilization patterns showed variability based on test loads, initially rising from 0.995% to 2.17% within the first thirty minutes before dropping to 1.58%. Over an hour, it peaked at 6.74% and averaged 3.75%, indicating that increased user demand affects CPU performance. In contrast, AWS CloudWatch data during GET request testing demonstrated a gradual increase in memory usage from 23.9% to 47.2%, while CPU utilization fluctuated significantly, peaking at 6.94% before dropping back to 0.8%.

Similarly, Azure dashboards during POST requests indicated a decrease in available memory from 7.63 GiB to 0.83 GiB, alongside fluctuating CPU usage peaking at 3-4%. For GET requests, Azure maintained steady memory availability at 6.86 GiB, while CPU utilization peaked at 8% and subsequently declined to around 1%, showing effective handling of increased demand.

B. Cost Analysis

The cost analysis evaluated pricing models using calculators from both providers. Azure's Reserved Instances emerged as the most economical at $47.05 per month for stable workloads, while AWS's EC2 Instance Savings Plans start at $55.52. AWS's On-Demand Instances cost $58.73, making them more affordable than Azure's $75.39 Pay-as-you-go model. Overall, Azure is more cost-effective for predictable workloads, while AWS provides better value for on-demand flexibility.

C. Assessing Performance and Cost Efficiency

Evaluating AWS and Azure involves balancing performance and cost. AWS delivers stable performance, with POST request response times increasing from 339 milliseconds at 100 users to 2943 milliseconds at 600 users, maintaining a 0% error rate.

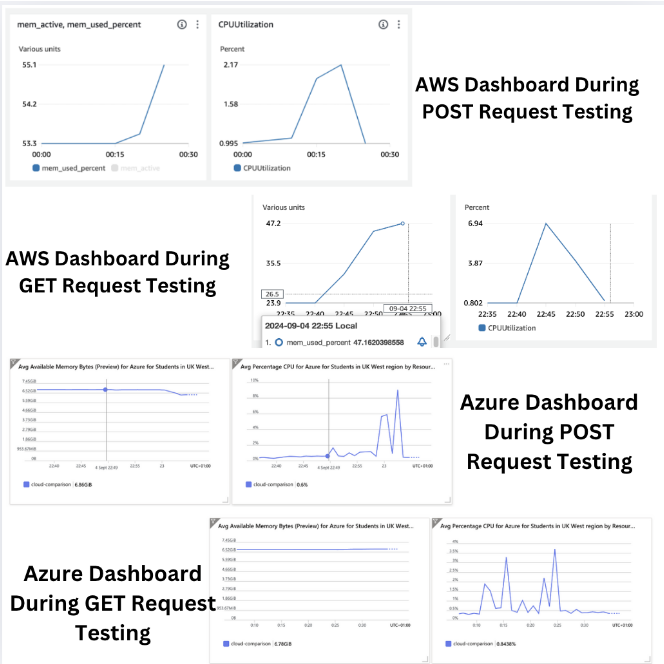

Fig. 4. Monitoring Dashboard

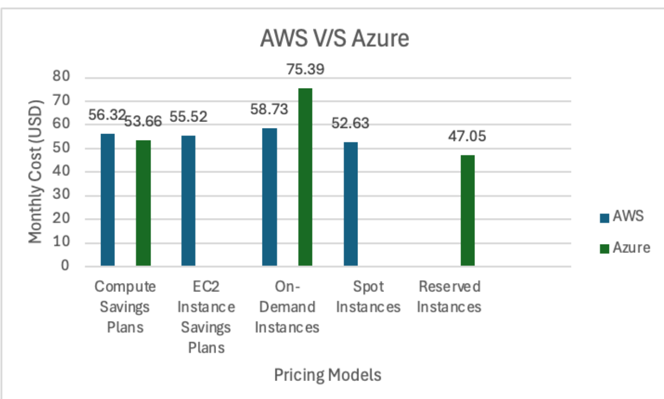

Fig. 5. Cost Comparison between AWS and Azure

Azure also achieved a 0% error rate but showed variability, with average response times of 3052 milliseconds at 600 users. Ultimately, Azure is more budget-friendly for predictable workloads, while AWS excels in stability under heavy loads.

V. Conclusion and Future Work

This analysis compares Microsoft Azure and AWS in cost, performance, and reliability for small-scale deployments. AWS achieves a 0% error rate but experiences performance degradation with increased load, while Azure displays greater variability. Cost-wise, Azure's Reserved Instances are the most economical, while AWS's Spot Instances provide savings. Azure is more budget-friendly for predictable workloads, while AWS may suit those needing reliable performance under variable loads.

Future research could incorporate the following directions:

1) Integrate Google Cloud Platform (GCP) into comparisons, particularly in the domains of machine learning and data analytics, to provide organizations with broader insights into cloud options.
2) Explore multi-cloud strategies that integrate AWS, Azure, and GCP, enhancing adaptability and resilience in cloud architectures.
3) Investigate advanced features such as load balancing and auto-scaling across all three platforms to assess their performance during traffic peaks and resource allocation.

References

[1] Ahuja, S. P., 2019. On the use of system-level benchmarks for comparing public cloud environments. In: Handbook of Research on Cloud Computing and Big Data Applications in IoT.. Hershey, PA, USA: IGI Global,, p. 24–38.

[2] J. Zhang, Z., Zheng, W., & Liu, T., 2022. Comparative Analysis of Cloud Platforms: AWS vs Azure.. Journal of Cloud Computing: Advances, Systems and Applications, 11(1), pp. 22-34.

[3] Kaushik, P. R. A. M. S. D. P. V. S. &. G. S., 2021. Cloud Computing and Comparison based on Service and Performance between Amazon AWS, Microsoft Azure, and Google Cloud.. Tashkent, Uzbekistan., pp. 268-273.

[4] Vouk, M. A., 2008. *Cloud Computing—Issues, Research and Implementations.*Piscataway, NJ, USA, p. 31–40.

[5] Babu, S. M. L. A. J. &. R. B. T., Apr. 2015. *A study on cloud based Internet of Things: Cloudio T.* p. 60–65.

[6] Salindeho, J. F. M. J. H. &. L. O., 2021. *Cost-Benefit Analysis of Cloud Computing in Education Using the Base Cost Estimation Model.*Makasar, Indonesia., pp.1-5.

[7] P. Pierleoni, R. C. A. B. a. L. P., 2020. Amazon, Google and Microsoft Solutions for IoT: Architectures and a Performance Comparison. *IEEE Access,* Volume 8, pp. 5455-5470.

[8] Zhao, W. M.-S. P. &. M. L., 2013. Low Latency Fault Tolerance System. *The Computer Journal,* Volume 56, p. 716–740.

[9] thearaseng, 2023. *React Product App.* [Online] Available at: https://github.com/thearaseng/react-product-app/tree/main

[10] Amazon Web Services, Inc., *AWS Pricing Calculator*, [Online]. Available: https://calculator.aws/#/estimatehttps://calculator.aws/#/estimate. Accessed: 2024.

[11] Microsoft Azure.[Online] Available at: https://azure.microsoft.com/en-gb/pricing/calculator/

[12] Apache JMeter, "Apache JMeter™," [Online] Available:https://jmeter.apache.org/download_jmeter.cgi.

[13] Statista, "Worldwide market share of leading cloud infrastructure service providers," Statista, 2024. [Online]. Available: https://www.statista.com/chart/18819/worldwide-market-share-of-leading-cloud-infrastructure-service-providers/