



A Comparative Study of Cloud Computing Platforms for Enterprise Applications

Pyarali El-Masry

Department of Technology and Innovation Management, Visayan Maritime Polytechnic University, Philippines

pyarali.el.masry@vmpu-ph.net

Peer Review Information

Submission: 03 Jan 2016

Revision: 18 Feb 2016

Acceptance: 05 March 2016

Keywords

*Cloud computing,
Enterprise applications,
AWS, Microsoft Azure,
Google Cloud Platform*

Abstract

Cloud computing became a practical enterprise delivery model by combining on-demand resource provisioning, virtualization, measured usage and elastic scaling. Yet enterprise adoption is not simply a choice of the least expensive provider; platform fit depends on workload behavior, service quality, governance, security, interoperability and operational effort. This paper presents a methodology for comparing three representative public-cloud platforms—Amazon Web Services (AWS), Microsoft Azure and Google Cloud Platform (GCP)—for enterprise application deployment using evidence and evaluation concepts available up to 2016. The method defines a common application profile, deploys equivalent compute, storage and network configurations, measures repeatable technical indicators, and combines them with enterprise-oriented qualitative criteria. Metrics are normalized and analyzed through a weighted multi-criteria model rather than a single benchmark. The comparative analysis indicates that provider selection is workload- and organization-dependent: differences in elasticity, service integration, management, security assurance and cost behavior can change the preferred platform even when raw compute performance is similar. The proposed procedure therefore offers a reproducible framework for enterprise cloud selection without assuming that one provider is universally superior.

Introduction

Cloud computing changed the traditional enterprise model in which servers, storage and application platforms were purchased, installed and maintained inside an organization. In the cloud model, resources can be requested on demand and released when they are no longer required. Early research described the cloud as an evolution toward utility-style computing and emphasized virtualization, elasticity, service-level management and usage-based charging as central characteristics (Buyya et al., 2009; Armbrust et al., 2010). Vaquero et al. (2009) similarly showed that a useful definition of cloud computing must include dynamically scalable

and virtualized resources delivered as services. These ideas are particularly relevant to enterprises because application capacity can be aligned more closely with changing business demand.

The enterprise decision, however, is more complex than deciding whether cloud computing is useful. Public cloud providers expose different combinations of infrastructure, platform services, management tools and pricing mechanisms. Prodan and Ostermann (2009) showed that cloud offerings could be distinguished by service type, resource deployment, hardware, security, business model,

middleware and performance. Li et al. (2010) later demonstrated through CloudCmp that public-cloud services can differ substantially in compute, storage and network behavior, and that a provider that performs well for one application may not be the most appropriate for another. Therefore, platform selection should be treated as a structured engineering decision rather than a brand-based choice.

For enterprise applications, the relevant criteria extend beyond raw processing speed. Marston et al. (2011) highlighted the business value of rapid scaling and reduced capital expenditure while also identifying governance, regulation and stakeholder concerns. Security research has shown that moving data and applications to a third-party infrastructure introduces issues involving confidentiality, privacy, availability, virtualization and compliance (Subashini and Kavitha, 2011; Hashizume et al., 2013). At the same time, energy-aware resource management and elasticity influence both provider efficiency and the cost ultimately experienced by the consumer (Beloglazov et al., 2012). A credible comparison must therefore combine performance measurement with security, manageability, availability, interoperability and cost-related criteria.

This study proposes a repeatable methodology for comparing AWS, Microsoft Azure and GCP from the perspective of a medium-scale enterprise web application. The aim is not to declare a universal winner. Instead, the objective is to define a controlled comparison process in which equivalent workloads are deployed, metrics are collected using the same procedure, and enterprise requirements are translated into explicit weights. The study uses research published no later than 2016 so that the conceptual basis matches the requested period. The resulting framework can be adapted to other cloud providers, application types or organizational priorities.

Literature Review

Buyya et al. (2009) presented cloud computing as the delivery of computing resources as a utility and discussed market-oriented resource allocation, service-level agreements (SLAs) and representative commercial platforms. Their work established a foundation for evaluating a cloud not only through hardware capacity but also through service management and economic mechanisms. For enterprise applications, this supports the inclusion of elasticity, SLA behavior and usage cost in a platform comparison.

Prodan and Ostermann (2009) surveyed infrastructure and web-hosting providers and proposed a taxonomy covering service type, deployment, hardware, runtime tuning, security, business model, middleware and performance. Their taxonomy is especially useful for comparative methodology because it avoids reducing cloud evaluation to a single benchmark. It also identifies interoperability, sustained performance and service-level guarantees as important areas for users selecting remote infrastructure.

Armbrust et al. (2010) examined the practical opportunities and obstacles of cloud computing, including elasticity, pay-as-you-go economics, availability, data transfer and performance unpredictability. Their analysis explains why enterprise workloads should be tested under both steady and changing demand. A platform can appear economical at low utilization but become less attractive when data movement, persistent storage or peak-resource requirements are included.

Li et al. (2010) developed CloudCmp to compare public cloud providers using application-relevant measurements of elastic computing, persistent storage and networking. The study emphasized fairness, representativeness and low measurement cost. Its main methodological lesson is that identical measurement logic should be applied across providers and that provider choice should reflect the application's resource profile rather than a generic performance label.

Iosup et al. (2011) empirically evaluated commercial cloud services and showed that workload type strongly affects observed performance. Although their focus included many-task scientific workloads, the result generalizes to enterprise evaluation: compute-intensive, I/O-intensive and latency-sensitive services should not be represented by one synthetic benchmark. Multiple workload classes are therefore included in the methodology proposed here.

Calheiros et al. (2011) introduced CloudSim as a repeatable environment for studying resource provisioning and application behavior in cloud systems. Their work demonstrates the value of controlled modeling before or alongside costly real deployments. This paper adopts the same principle by defining workload parameters, repetition rules and metric normalization explicitly so that the comparison can be repeated or simulated when direct provider access is constrained.

Subashini and Kavitha (2011) analyzed security issues across SaaS, PaaS and IaaS delivery models and emphasized data privacy, data protection, application security and virtualization risks. Hashizume et al. (2013) further related vulnerabilities and threats to cloud architecture and outsourced service delivery. Together these studies justify treating security and governance as first-class comparison dimensions rather than as descriptive appendices.

Garg et al. (2013) proposed a framework for measuring and ranking cloud services using quality-of-service attributes and user preferences. Their work shows that service selection is fundamentally a multi-criteria decision problem. Later studies continued this direction: Wagle et al. (2015) proposed an evaluation model for commercially available providers, while Baranwal and Vidyarthi (2016) developed a ranked-voting model using QoS expectations and user experience. Cayirci et al. (2016) added risk-oriented provider assessment. These studies support a weighted comparison model in which enterprise priorities determine how technical and non-technical metrics are combined.

Methodology

The proposed methodology is a controlled comparative evaluation designed for enterprise applications. It follows five stages: requirements definition, equivalent platform configuration, workload execution, metric collection, and normalized multi-criteria comparison. AWS, Microsoft Azure and GCP are treated as representative public-cloud environments. The method is deliberately provider-neutral; the same application image, workload mix, observation period and measurement rules are used wherever technically possible.

1. Enterprise Application Profile and Test Environment

A reference three-tier web application is used because it represents a common enterprise pattern: a web/API tier, an application-processing tier and a persistent database/storage tier. The test environment

should use comparable general-purpose virtual machines with the same virtual CPU count, similar memory allocation and the same operating-system family. Each provider configuration is placed in a single region for baseline testing. The application dataset, software versions and test scripts are identical. Before measurement, each deployment receives a warm-up period so that installation, caching and first-run initialization do not distort the results.

Three workload classes are executed: (i) steady transactional load, representing routine business traffic; (ii) burst load, representing sudden demand increases; and (iii) data-intensive load, representing report generation and file/database transfer. Each workload is repeated at least five times. Median values are used for comparison, while dispersion is retained as an indicator of performance stability. This approach follows the repeatability principle emphasized by CloudCmp and performance-analysis studies (Li et al., 2010; Iosup et al., 2011).

2. Evaluation Criteria and Metrics

The criteria combine measurable technical indicators with enterprise requirements. Performance includes response time, request throughput and processing completion time. Elasticity is represented by resource-provisioning delay and the time required to recover a target response level after a load increase. Cost is calculated from compute time, persistent storage and data-transfer components for the same test workload. Reliability is represented through successful-request ratio and observed service continuity during the test window. Security and governance are assessed through available identity controls, isolation, auditability, data-protection mechanisms and administrative visibility, consistent with the concerns identified by Subashini and Kavitha (2011), Hashizume et al. (2013) and Cayirci et al. (2016). Manageability and interoperability include deployment effort, automation, monitoring, portability and integration with enterprise systems.

Table 1. Evaluation criteria and example enterprise weights

Criterion	Representative measures	Weight	Direction
Performance	Response time, throughput, completion time	0.25	Higher performance preferred
Elasticity	Provisioning delay, recovery after burst	0.15	Shorter delay preferred
Cost efficiency	Compute + storage + transfer for same workload	0.20	Lower normalized cost preferred
Reliability	Success ratio, continuity, variation	0.15	Higher stability preferred

Security & governance	Identity, audit, isolation, data protection	0.15	Higher assurance preferred
Manageability & portability	Automation, monitoring, deployment effort	0.10	Lower effort / higher portability

3. Normalization and Comparative Decision Procedure

Because the metrics use different units, each observed value is normalized to a 0–1 scale. For a benefit metric such as throughput or success ratio, the normalized value is calculated as $(x - x_{min}) / (x_{max} - x_{min})$. For a cost metric such as response time, price or provisioning delay, the direction is reversed using $(x_{max} - x) / (x_{max} - x_{min})$. A composite platform score can then be calculated as the sum of each normalized criterion score multiplied by its enterprise weight. The weights in Table 1 are an example profile and must be changed when business priorities differ. This explicit weighting prevents

the analyst from hiding subjective preferences inside the final comparison, consistent with multi-criteria cloud-selection research (Garg et al., 2013; Wagle et al., 2015; Baranwal and Vidyarthi, 2016).

Sensitivity analysis is performed by increasing and decreasing each major weight by 20% while renormalizing the remaining weights. A platform recommendation is considered robust only when the result remains stable across reasonable weight changes. If the preferred platform changes easily, the correct conclusion is that the decision is requirement-sensitive and should be validated with organization-specific pilots rather than treated as a universal ranking.




Figure 1: Cloud Platform Evaluation and Benchmarking Workflow

Results

The proposed procedure produces a comparative profile rather than a single isolated benchmark number. The principal result from the literature-informed analysis is that cloud platforms exhibit meaningful trade-offs across service behavior, cost structure and enterprise controls. CloudCmp reported wide differences

among providers and showed that application-specific measurements can alter provider choice (Li et al., 2010). Garg et al. (2013), Wagle et al. (2015) and Baranwal and Vidyarthi (2016) likewise demonstrate that user preferences and QoS priorities materially affect service ranking. Consequently, the most defensible result is a criterion-by-criterion interpretation followed by a weighted decision for the target enterprise.

Table 2. Comparative enterprise-oriented interpretation using the proposed framework (2016 context)

Dimension	AWS	Microsoft Azure	Google Cloud Platform
Compute / service flexibility	Strong IaaS flexibility; suitable for varied VM-based enterprise workloads.	Strong fit where enterprise applications align with Microsoft platform tooling and hybrid integration.	Strong fit for cloud-native and managed application patterns; evaluate VM and data-service behavior for the target workload.
Elasticity	Evaluate scale-out time and auto-scaling response under burst load.	Evaluate scale-set / platform scaling behavior with identical thresholds.	Evaluate instance and managed-service scaling under the same burst profile.
Cost behavior	Assess compute, storage and data-transfer together;	Include licensing and enterprise	Include sustained resource use, storage and network behavior in total cost.

	avoid comparing headline VM price alone.	integration effects in total cost.	
Security / governance	Assess identity, audit, isolation, encryption and responsibility boundaries.	Assess the same controls plus integration with enterprise identity and governance processes.	Assess the same controls with emphasis on centralized identity, auditability and data access policy.
Management / portability	Broad automation possibilities can increase flexibility but require governance discipline.	Potentially lower operational friction for Microsoft-centered enterprises.	Potentially efficient for teams emphasizing managed, developer-oriented cloud services.

For a transaction-oriented enterprise application, performance should be interpreted through median latency, throughput and variability rather than peak throughput alone. If the three platforms deliver similar response time, the decision can shift toward manageability, governance or predictable total cost. For a burst-oriented service, elasticity becomes more influential: the platform that restores the target response level with less delay may justify a higher unit price. For data-intensive applications, storage and network transfer can dominate the cost/performance profile. These findings are consistent with the argument of Armbrust et al. (2010) that data transfer, availability and performance variability can be significant cloud obstacles and with Iosup et al. (2011), who showed that cloud suitability depends strongly on workload characteristics.

The comparison also shows why security should be scored as a structured requirement rather than a binary feature. All major platforms can provide security mechanisms, but the enterprise outcome depends on configuration, auditability, identity integration, data location and the division of responsibility between provider and customer. Security surveys before 2016 repeatedly identified privacy, virtualization, application security and outsourced control as significant concerns (Subashini and Kavitha, 2011; Hashizume et al., 2013). Cayirci et al. (2016) further demonstrated the value of risk-based provider selection. Thus, a platform with high technical performance may still be unsuitable when governance evidence does not satisfy organizational requirements.

Finally, the weighted model should be viewed as a decision aid, not a permanent ranking. When performance receives the largest weight, benchmark differences dominate the final score. When security, interoperability or management receives greater weight, the preferred platform may change. This sensitivity is not a weakness of

the method; it reflects the real nature of enterprise cloud selection. A reliable decision is one whose criteria and weights are explicit, repeatable and traceable to business requirements.

Conclusion and Discussion

This paper presented a structured methodology for comparing AWS, Microsoft Azure and Google Cloud Platform for enterprise applications using research and evaluation concepts published up to 2016. The framework begins with an explicit enterprise application profile, establishes equivalent cloud configurations, executes multiple workloads classes, collects technical and operational metrics, normalizes the observations and combines them through a weighted multi-criteria model. The approach responds to a central finding in prior research: public cloud services differ in performance, cost, risk and management characteristics, and no single measure can represent overall suitability. The comparative analysis indicates that platform selection should be driven by application behavior and organizational constraints. Compute performance is important but must be considered together with elasticity, storage and network behavior, reliability, total cost, security, governance and operational effort. A platform that is attractive for a bursty customer-facing application may not be preferable for a data-intensive analytical service or a tightly governed enterprise system. Likewise, strong integration with an organization's existing technology stack can reduce operational cost even when raw infrastructure pricing is not the lowest. The main limitation is that cloud services evolve rapidly and region-specific pricing or features can change; therefore, historical literature cannot substitute for a current pilot when an organization is making a real procurement decision. The methodology remains useful because it separates stable evaluation principles

from provider-specific values. Future work can extend the framework with larger benchmark suites, hybrid-cloud scenarios, application migration cost, failure-injection experiments and organization-specific compliance controls. The most important practical conclusion is that enterprise cloud evaluation should be repeatable, multi-dimensional and requirements-driven rather than based on a universal provider ranking.

References

- Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., Lee, G., Patterson, D., Rabkin, A., Stoica, I., & Zaharia, M. (2010). A view of cloud computing. *Communications of the ACM*, 53(4), 50–58. <https://doi.org/10.1145/1721654.1721672>
- Baranwal, G., & Vidyarthi, D. P. (2016). A cloud service selection model using improved ranked voting method. *Concurrency and Computation: Practice and Experience*, 28(13), 3540–3567. <https://doi.org/10.1002/cpe.3740>
- Beloglazov, A., Abawajy, J., & Buyya, R. (2012). Energy-aware resource allocation heuristics for efficient management of data centers for cloud computing. *Future Generation Computer Systems*, 28(5), 755–768. <https://doi.org/10.1016/j.future.2011.04.017>
- Buyya, R., Yeo, C. S., Venugopal, S., Broberg, J., & Brandic, I. (2009). Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility. *Future Generation Computer Systems*, 25(6), 599–616. <https://doi.org/10.1016/j.future.2008.12.001>
- Calheiros, R. N., Ranjan, R., Beloglazov, A., De Rose, C. A. F., & Buyya, R. (2011). CloudSim: A toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms. *Software: Practice and Experience*, 41(1), 23–50. <https://doi.org/10.1002/spe.995>
- Cayirci, E., Garaga, A., Santana de Oliveira, A., & Roudier, Y. (2016). A risk assessment model for selecting cloud service providers. *Journal of Cloud Computing*, 5, 14. <https://doi.org/10.1186/s13677-016-0064-x>
- Garg, S. K., Versteeg, S., & Buyya, R. (2013). A framework for ranking of cloud computing services. *Future Generation Computer Systems*, 29(4), 1012–1023. <https://doi.org/10.1016/j.future.2012.06.006>
- Hashizume, K., Rosado, D. G., Fernández-Medina, E., & Fernandez, E. B. (2013). An analysis of security issues for cloud computing. *Journal of Internet Services and Applications*, 4, 5. <https://doi.org/10.1186/1869-0238-4-5>
- Iosup, A., Ostermann, S., Yigitbasi, M. N., Prodan, R., Fahringer, T., & Epema, D. H. J. (2011). Performance analysis of cloud computing services for many-tasks scientific computing. *IEEE Transactions on Parallel and Distributed Systems*, 22(6), 931–945. <https://doi.org/10.1109/TPDS.2011.66>
- Li, A., Yang, X., Kandula, S., & Zhang, M. (2010). CloudCmp: Comparing public cloud providers. *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*, 1–14. <https://doi.org/10.1145/1879141.1879143>
- Marston, S., Li, Z., Bandyopadhyay, S., Zhang, J., & Ghalsasi, A. (2011). Cloud computing—The business perspective. *Decision Support Systems*, 51(1), 176–189. <https://doi.org/10.1016/j.dss.2010.12.006>
- Prodan, R., & Ostermann, S. (2009). A survey and taxonomy of infrastructure as a service and web hosting cloud providers. *10th IEEE/ACM International Conference on Grid Computing*, 17–25. <https://doi.org/10.1109/GRID.2009.5353074>
- Subashini, S., & Kavitha, V. (2011). A survey on security issues in service delivery models of cloud computing. *Journal of Network and Computer Applications*, 34(1), 1–11. <https://doi.org/10.1016/j.jnca.2010.07.006>
- Vaquero, L. M., Roderio-Merino, L., Caceres, J., & Lindner, M. (2009). A break in the clouds: Towards a cloud definition. *ACM SIGCOMM Computer Communication Review*, 39(1), 50–55. <https://doi.org/10.1145/1496091.1496100>
- Wagle, S. S., Guzek, M., Bouvry, P., & Bisdorff, R. (2015). An evaluation model for selecting cloud services from commercially available cloud providers. *2015 IEEE 7th International Conference on Cloud Computing Technology and Science (CloudCom)*, 107–114. <https://doi.org/10.1109/CloudCom.2015.94>