International Journal of Cloud Applications and Computing
Volume 16 • Issue 1 • January-December 2026

Evaluating Multi-Cloud Portability and Elastic Scaling Across Serverless-Container Platforms

Dac Tuan Thanh Nguyen
http://orcid.org/0009-0009-3700-3866
Department of Computer Science and Information Technology, School of Computing, La Trobe University, Melbourne, Australia

Tony de Souza-Daw
Department of Computer Science and Information Technology, School of Computing, La Trobe University, Melbourne, Australia
Received: January 30th, 2026 | **Accepted:** March 13th, 2026

ABSTRACT

Multi-cloud adoption is increasingly driven by resilience, regulatory compliance, and vendor independence. Serverless-container platforms such as AWS Fargate, Google Cloud Run, and Azure Container Instances promise portability through similar abstraction models, yet they exhibit differences in configuration semantics and runtime behavior. This study presents a metric-driven evaluation of multi-cloud compatibility using the Cloud Compatibility System, distinguishing between feature-level configuration coverage and empirical execution reliability. Unified container configurations are deployed across providers under controlled conditions to measure execution success, startup latency, and runtime variability. Results show high execution reliability despite limited feature compatibility, revealing a weak correlation between documented configuration support and runtime behavior. These findings highlight the limitations of feature-centric portability assessments and establish an empirically grounded framework for evaluating interoperability in serverless-container environments.

KEYWORDS

Multi-Cloud Compatibility, Cloud Architecture, Serverless-Container Platforms, Cloud Elasticity, Compatibility Metrics, Cross-Provider Deployment, Serverless Computing, Cloud Interoperability

INTRODUCTION

Serverless-container platforms have become a foundational execution model for modern cloud applications, providing benefits such as managed scaling, reduced operational overhead, and portability across cloud providers. While services such as AWS Fargate, Google Cloud Run, and Azure Container Instances share similar abstraction goals, they expose substantial differences in execution semantics, including cold-start behavior, scaling thresholds, concurrency policies, and resource allocation strategies. These discrepancies introduce hidden incompatibilities that remain largely unquantified in existing multi-cloud research, limiting our understanding of true portability across platforms. Prior studies, including previous work on the Cloud Compatibility System (CCS), have primarily focused on provisioning-layer interoperability, such as configuration translation and feature-level compatibility. However, provisioning success alone does not ensure consistent runtime

DOI: 10.4018/IJCAC.406105

This article published as an Open Access article distributed under the terms of the Creative Commons Attribution License (http://creative-commons.org/licenses/by/4.0/) which permits unrestricted use, distribution, and production in any medium, provided the author of the original work and original publication source are properly credited.

behavior, especially in serverless-container environments, where elasticity and execution dynamics are critical for real-world workload portability. This gap in existing methodologies necessitates a more comprehensive approach that incorporates runtime execution characteristics.

This paper extends the CCS metric-based evaluation framework by integrating runtime performance with configuration compatibility. Specifically, unified container deployments are evaluated across AWS, Azure, and GCP, focusing on execution success, latency characteristics, scaling responsiveness, and the deviation between feature-level compatibility and observed runtime behavior. The primary contribution of this work lies in its provision of the first quantitative and reproducible analysis of runtime compatibility and elasticity for serverless-container platforms. By addressing a critical gap in multi-cloud portability research, this study explores the dynamic execution behavior and elastic scaling of containerized workloads, thereby offering a methodologically rigorous foundation for performance-aware multi-cloud portability assessments.

BACKGROUND AND PROBLEM STATEMENT

Multi-cloud computing has evolved into a strategic architectural approach for organisations seeking resilience, cost optimisation, regulatory compliance, and vendor neutrality. However, the increasing diversity of cloud services has amplified the longstanding challenge of interoperability across heterogeneous providers. Foundational studies show that cloud platforms expose distinct resource models, configuration semantics, and operational behaviors, creating fundamental barriers to portability even when high-level service categories appear equivalent (Petcu, 2013; Zhang, Wu, & Cheung, 2013; Varghese & Buyya, 2018; Wang et al., 2017). The NIST definition formalised the conceptual characteristics of cloud systems, but it does not prescribe how these characteristics should be expressed or executed across different providers, leaving ambiguity in cross-cloud alignment (Mell & Grance, 2011). As cloud ecosystems expanded, containerisation technologies such as Docker were introduced to promote packaging portability (Merkel, 2014; Boettiger, 2015); yet cross-provider behavioral divergence persists because container runtimes, networking policies, and execution environments differ substantially across platforms (Burns et al., 2016). Interoperability research subsequently progressed toward semantic abstraction frameworks—including ontologies, cross-cloud libraries, and multi-cloud adapters—that attempt to homogenise the representation of cloud resources (Bassiliades et al., 2018; Achilleos et al., 2019; Kamateri et al., 2013). These tools enable provider-agnostic descriptions of compute, storage, and networking components, offering a partial solution to the complexity of heterogeneous cloud APIs. Nonetheless, they generally operate at the metadata or schema level and are not designed to evaluate whether abstracted resource definitions behave consistently when executed in practice. As a result, existing approaches cannot determine whether semantic portability translates into functional compatibility. The emergence of serverless and serverless-container execution models further compounds this challenge. Studies on Function-as-a-Service (FaaS) demonstrate that performance varies with cold starts, scheduling delays, regional load distribution, and platform-specific optimisations (McGrath & Brenner, 2017; Schirmer et al., 2023; Ebrahimi et al., 2024). Similar variability is observed in container-based cloud services, where execution paths, resource isolation, and underlying infrastructure significantly impact runtime behavior (Ruan et al., 2016; Leitner & Cito, 2016; Huang et al., 2024). Even when identical images or configurations are deployed, cloud providers frequently exhibit divergent latency profiles, scaling triggers, and reliability patterns (Wang et al., 2018). Prior research has characterised cloud variability and resource elasticity (Al-Dhuraibi et al., 2018; Calheiros et al., 2015; Zhao et al., 2021). However, these findings have not been integrated into a broader framework that links configuration-level definitions with real execution outcomes.

To address these challenges, this paper builds upon the Cloud Compatibility System (CCS), a modular provisioning framework introduced in earlier work. CCS provides a provider-agnostic command-line interface that abstracts cloud-specific deployment logic and enables users to express

20 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage: www.igi-global.com/article/evaluating-multi-cloud-portability-and-elastic-scaling-across-serverless-container-platforms/406105

Related Content

Marketing Analytics for Analyzing Social Media and Branding Efforts of Indian Telecom Operators
Sujata Joshi, Sindhu Koparratiand Vaishali Singhal (2022). *International Journal of Cloud Applications and Computing (pp. 1-18).*
www.irma-international.org/article/marketing-analytics-for-analyzing-social-media-and-branding-efforts-of-indian-telecom-operators/297092

RDF Query Path Optimization Using Hybrid Genetic Algorithms: Semantic Web vs. Data-Intensive Cloud Computing
Qazi Mudassar Ilyas, Muneer Ahmad, Sonia Raufand Danish Irfan (2022). *International Journal of Cloud Applications and Computing (pp. 1-16).*
www.irma-international.org/article/rdf-query-path-optimization-using-hybrid-genetic-algorithms/284492

Secure Data Deduplication of Encrypted Data in Cloud
Sumit Kumar Mahanaand Rajesh Kumar Aggarwal (2019). *Handbook of Research on the IoT, Cloud Computing, and Wireless Network Optimization (pp. 196-212).*
www.irma-international.org/chapter/secure-data-deduplication-of-encrypted-data-in-cloud/225719

Information Dispersal Algorithms and Their Applications in Cloud Computing
Makhan Singhand Sarbjeet Singh (2018). *Critical Research on Scalability and Security Issues in Virtual Cloud Environments (pp. 72-94).*
www.irma-international.org/chapter/information-dispersal-algorithms-and-their-applications-in-cloud-computing/195343

A Cloud-Based Environment for Collaborative Resources Management
Anna Goy, Giovanna Petroneand Marino Segnan (2014). *International Journal of Cloud Applications and Computing (pp. 7-31).*
www.irma-international.org/article/a-cloud-based-environment-for-collaborative-resources-management/122592