Abstract

Cloud computing and related technologies such as cloud-native continue to transform ways in which applications can be developed, packaged and deployed for consumption by users. By enabling dynamic, on-demand provisioning of scalable compute, network, and storage resources over distributed infrastructure, the cloud accelerates time-to-market, improves productivity, and provides for better availability of online services. However, the adoption of cloud computing in low resource settings is still low due to infrastructural limitations, operational constraints and other technical barriers such as capacity and skills. These settings are faced with unstable power grids, heterogeneous and limited resources, poor and unreliable network connectivity, and skill gaps in management and use of cloud infrastructure. These challenges complicate the development and deployment of resilient cloud applications. While microservice architectures present a promising solution through distributed design patterns, their implementation introduces management complexities, particularly in service orchestration, placement, and network communication optimization within these settings. Additionally, distributed operations like state machine replication require a careful balance between resource availability and capability, and the unpredictable network connectivity across locations.

Through empirical experimentation, we investigate the challenges of deploying distributed cloud applications in such environments and derive key requirements for designing a resilient multi-cloud abstraction architecture. Following the Design Science methodology, and theoretical modelling and design, we develop and implement four core solutions: (1) a prototype multi-cloud Platform-as-a-Service (PaaS) for application deployment, (2) an automated container orchestration pipeline for the multi-cloud platform, (3) an adaptive state machine replication protocol for fault tolerance, and (4) a dynamic microservice placement algorithm optimized for resource-constrained setups.

We run a set of experiments for each deliverable to evaluate performance in comparison to similar approaches, protocols and algorithms. The evaluation results demonstrate that the Multi-cloud PaaS prototype delivers up to 2X better throughput performance compared to remote options with 100% transaction completion rates across cellular and wired networks; the adaptive replication protocol achieves up to 5X and 2X higher throughput than other leader-based and leaderless protocols respectively; and the placement algorithm provides for 10X better performance under certain degraded conditions when compared to other container orchestration approaches such as Random, Spread, and Binpack strategies.