

Augmenting Databases-as-a-Service with Policy Support

Abstract

Cloud computing offers the potential to store, manage, and process data in highly available, scalable, and elastic environments. Yet, these environments still provide very limited and inflexible means for customers to control their data. This lack of control can hinder cloud adoption for data that falls under regulations.

As a part of our efforts in the Scalable and Secure Infrastructures for Cloud Operations (SSICLOPS) project, our work during the Spring 2017 period was focused on implementing a prototypical approach that facilitates policy adherence support based on the Compact Privacy Policy Language (CPPL) in the Hyrise-R in-memory research database scenario, depicted in Figure 1.

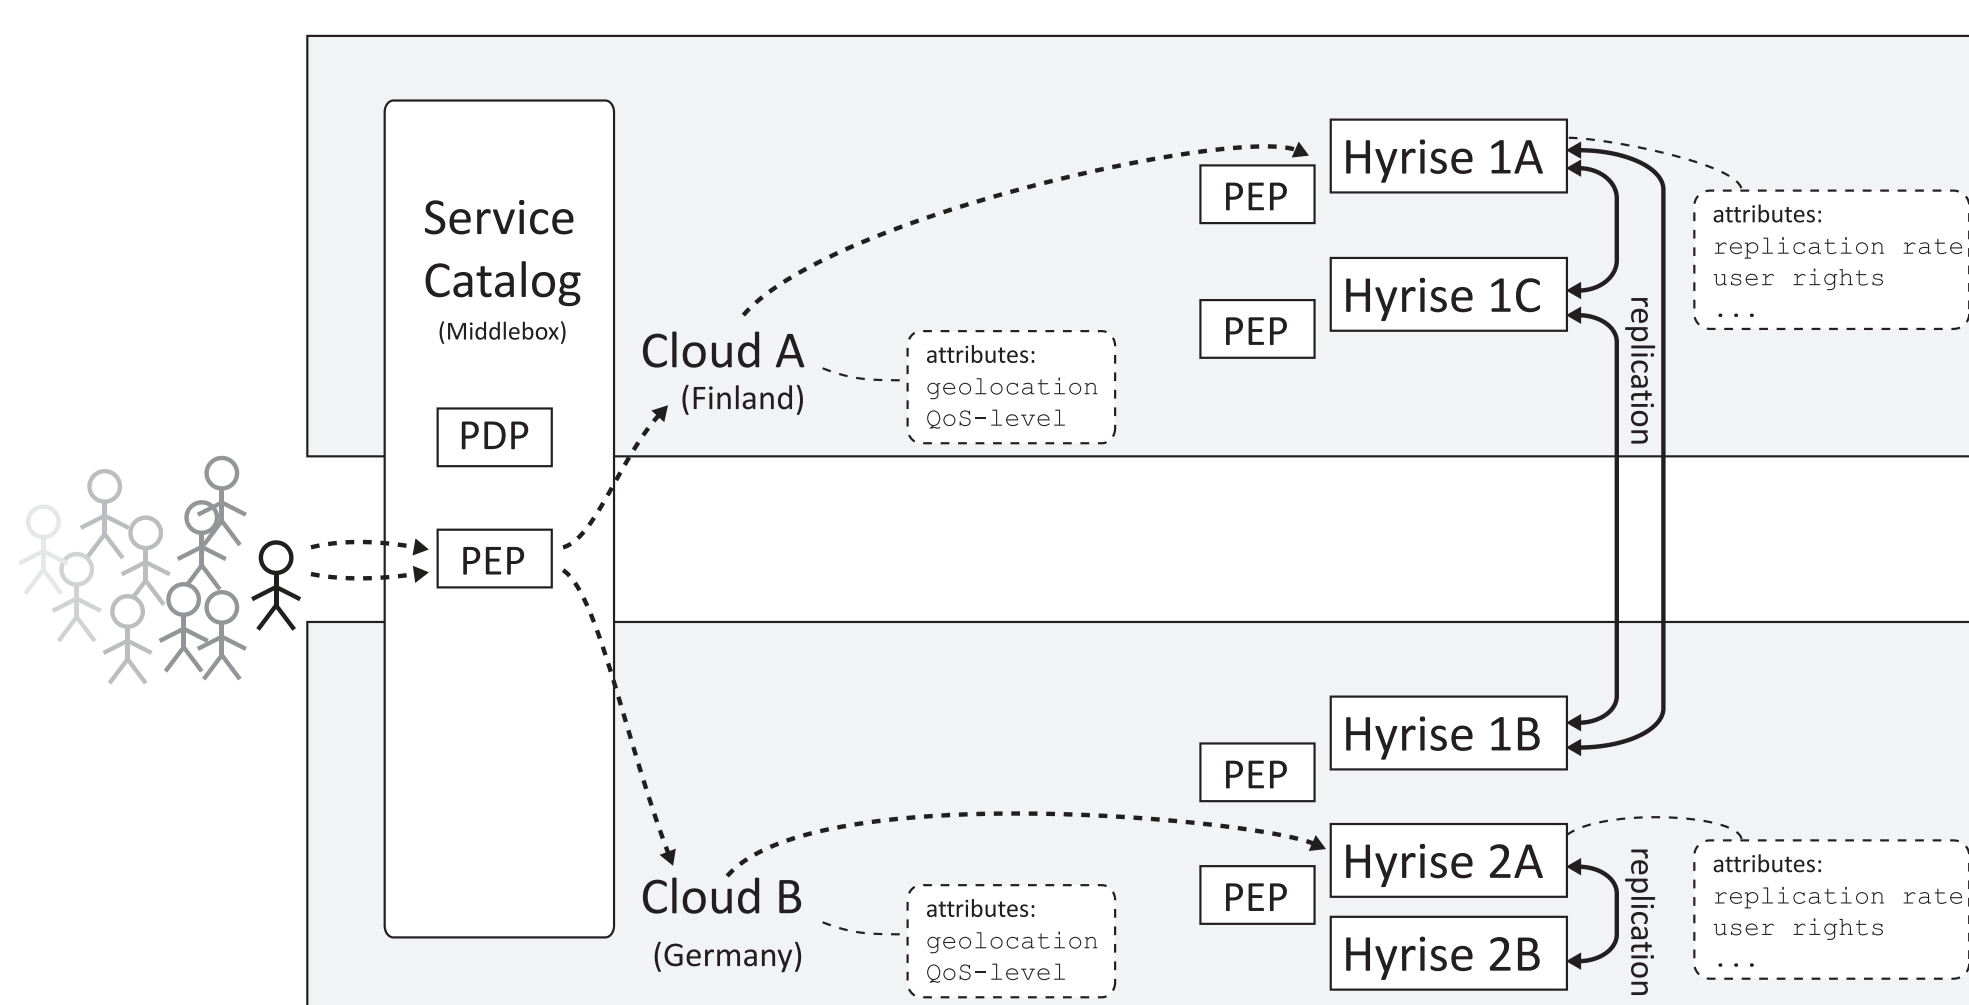


Figure 1: Use case scenario: Users request instances of the Hyrise-R in-memory database and annotate their requests with certain policy demands. The policy decision point (PDP) acts as the initial entry point and routes requests through a series of policy enforcement points (PEP) to process the requests accordingly.

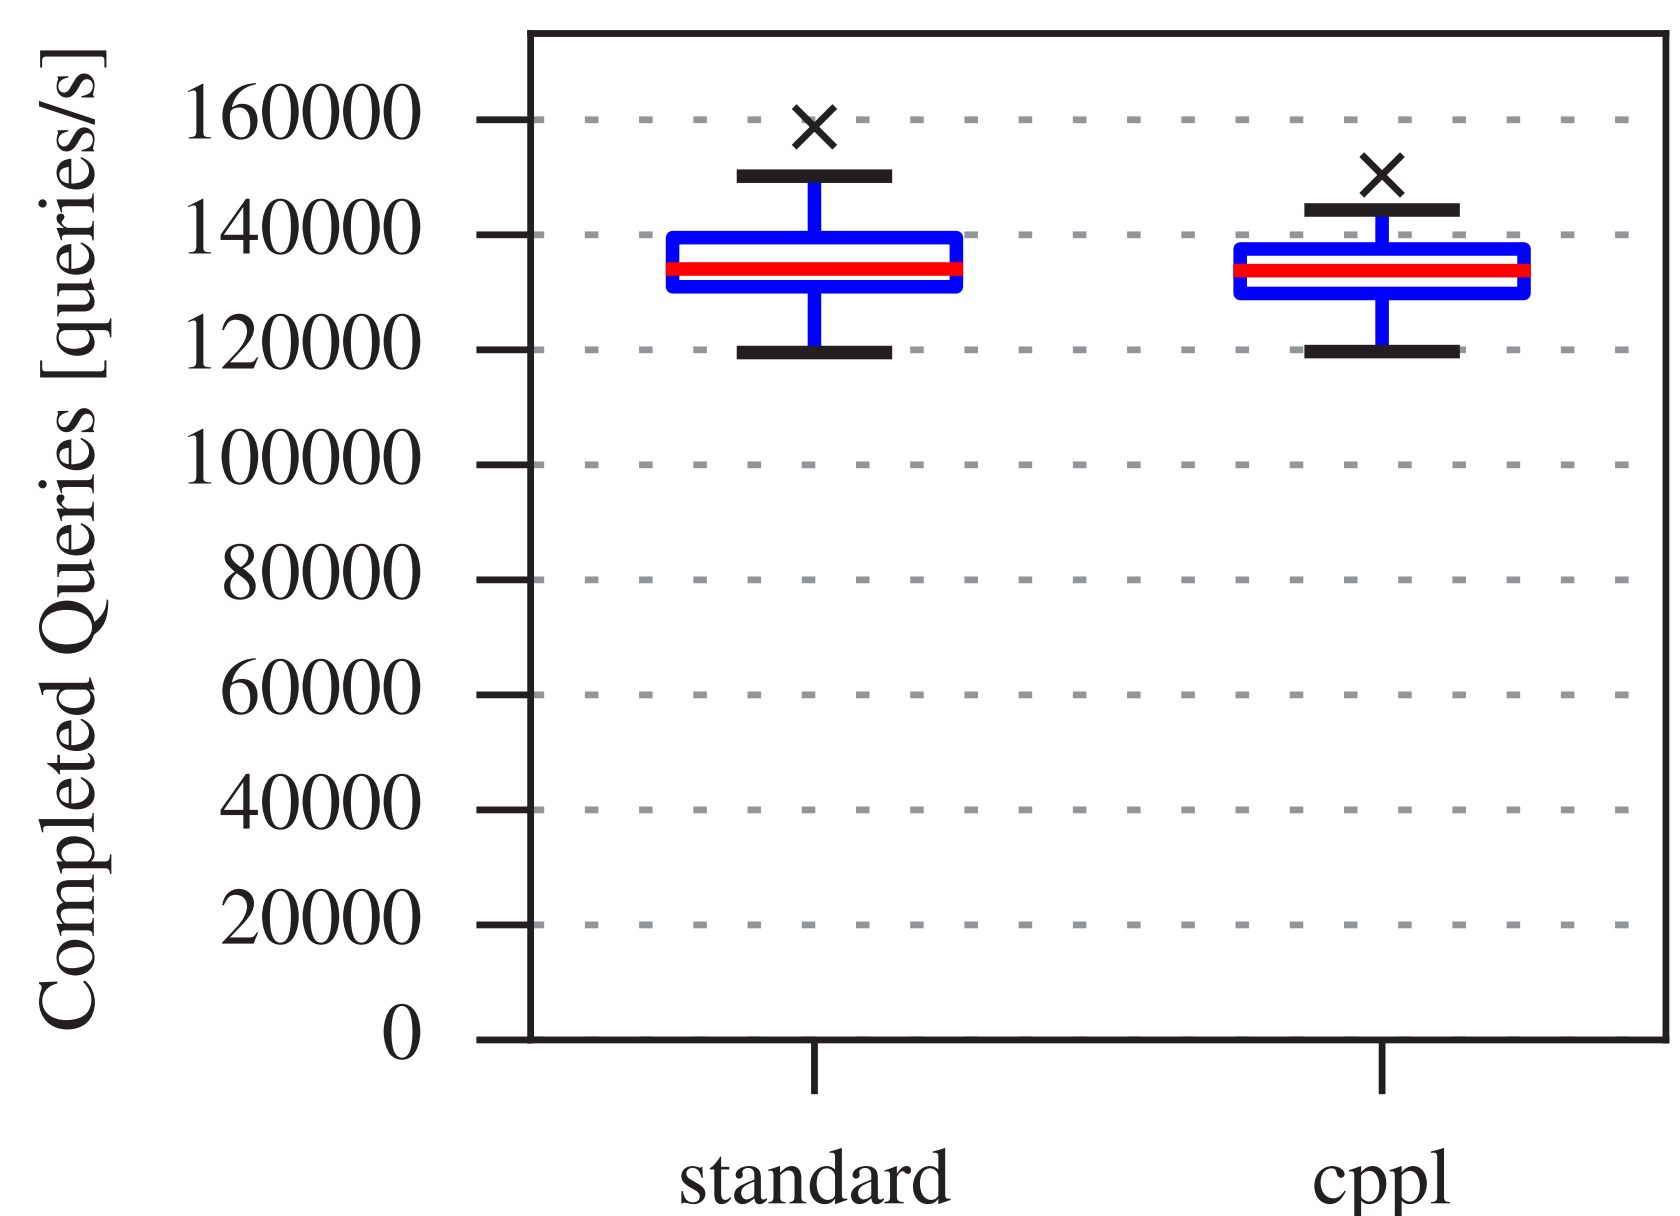


Figure 2: Performance of policy support in Hyrise-R: For high-throughput, transactional workloads, the policy evaluation incurs no notable overhead.

Contributions

Performance measurements were conducted by simulating a high-throughput transactional workload. No-op queries were used instead of actual operations to exclude confounding effects of actual operations as potential decelerating components that do not contribute to the integration of CPPL.

The results depicted in Figure 2 indicate negligible overhead for policy evaluation. Thus, based on the compact representation and the efficient matching of policies in CPPL, cloud-optimized scale-out deployments of in-memory databases can offer policy support without notable performance degradation.