

ABSTRACT

Cloud-based NoSQL databases like Google Cloud Firestore are a popular choice for modern applications, yet using operational databases to handle large-scale, read-intensive regulatory reporting workloads presents significant technical challenges. The issue was identified in a case study at MedMinutes, where periodical medical reports needed to be generated. Conventional architectures that mix these workloads create workload interference, posing risks of hotspotting that threaten the reliability of critical operational services and resulting in prolonged report generation times. This research aims to design and evaluate a hybrid snapshot architecture as a solution to decouple these workloads, adopting the Design Science Research Methodology (DSRM) integrated with the ICONIX Process for software prototype design. The built prototype archives data from Firestore to Google Cloud Storage, processing it through an asynchronous ETL pipeline. Quantitative evaluation demonstrates that the architecture successfully eliminates 100% of read operations on the operational database during repeated requests, mitigating hotspotting risks, and reducing execution time from 19.26 seconds to 0.38 seconds (50 times faster). Furthermore, findings confirm that implementing asynchronous processing patterns is essential for overcoming bottlenecks in serverless environments, achieving a 9.7-fold performance increase compared to synchronous implementation.

Keywords : Workload Interference, Design Science Research, ICONIX, Asynchronous I/O, Google Cloud Platform