

Serverless Native Analytics Engine

Database analytics traditionally performed by analytics engines on VM clusters often result in over-provisioning due to bursty and unpredictable workloads, leading to high costs. Serverless computing, with on-demand resource allocation and fine-grained billing, is a better fit for such workloads. However, existing analytics engines designed for VMs are not suitable for serverless functions due to their unique characteristics like extreme scalability, resource constraints, limited execution time, and lack of direct communication among workers.

Previous studies have shown that utilizing serverless functions for data analytics can offer cost-performance benefits with tailored query plans. While attempts have been made to improve efficiency in various stages of query processing, the lack of an end-to-end serverless-native analytics engine has limited the scope of co-optimization across the global query-processing context, resulting in limited flexibility and applications.

The key insights are:

- Existing query planning methods from distributed analytics engines can be adapted to serverless analytics.
- A lack of a serverless-native cost model limits the scope of optimization resulting in can result in non-optimal query plans.

By leveraging these insights, the goal is to design an on-demand, end-to-end, serverless-native analytics engine. Extensive data collection and analysis have been performed to generate a specialized performance and monetary cost model for serverless functions. The model once validated will then be utilized by the analytics engine to generate optimal query plans for each input query. The generated plans will be used by the engine to execute the query. This engine aims to fill the gap in current implementations and support data analytics at scale with the benefits of serverless computing.