

COMPUTES

Peer-to-peer mesh computing,
using FaunaDB.

ABOUT COMPUTES

Computes.io is building solutions for distributed applications that bypass the biggest blockchain bottlenecks, they see the market opportunity for distributed applications and are building a pragmatic platform for distributed apps. Computes.io's high performance computing infrastructure utilizes distributed web components like IPFS and IPLD to offer a simple interface for scalable function execution.

The broader Computes.io ecosystem includes multiple levels of systems. FaunaDB is used as part of a job control application, it was chosen for the easy to administer worldwide consistent cloud.

PROJECT OVERVIEW

Computes.io is building a decentralized peer-to-peer mesh computing platform. Computations can be sent into the mesh computer from any node in the private/public mesh network. Likewise, any node on the mesh network can assist with computations. This massively parallel execution engine works without any centralized administrator.

The team has a background building cloud compute platforms, so they know that existing blockchain based distributed computing is too slow. Computes.io is addressing the performance and cost of blockchain based distributed computing with a highly responsive blockless distributed ledger.

REQUIREMENTS & CHALLENGES

The market Computes.io is addressing for high-performance parallel compute includes both users of the public net as well as institutions wanting to deploy private compute environments. Because of this they need their entire software stack, from database to execution platform, to have worldwide availability on-premise or in the cloud.

One tool for accessing the Computes.io network allows users to submit job requests and retrieve pointers to job results. Because mesh members can come from all over the world and must agree on job status, they need a globally consistent database for job tracking.

They needed a cloud database, because there was no other affordable way to be globally available. Because jobs size is variable, a database with pay-per-operation pricing that can handle bursts of traffic on a budget is desired. The project is run by distributed computing enthusiasts, so they don't want to require new contributors to spend a lot on things like databases.

WHY FAUNA

Computes.io contributor Hareem Haque chose FaunaDB Cloud to store job requests as they are submitted, and results of completed compute jobs. Selection criteria included its globally distributed nature, its metered pricing model, and its operational simplicity.

FaunaDB Cloud runs in multiple cloud providers and regions, offering worldwide availability. This gives users all over the world a consistent high performance experience, and the ability to query the database from code running anywhere.

Metered pricing fits well with function execution platforms, because it is possible to track the cost of each query and bill end users. More importantly, it's affordable to run usually idle but occasionally busy workloads.

RESULTS

World wide access to Computes.io's high performance distributed computing cloud can be accessed via a managed cloud database, while developers can depend on data integrity and availability. After the initial success with the job and results tracking use case, Computes.io is expanding the use of FaunaDB to other workloads like identity and ecommerce.

FaunaDB takes the guesswork out of database management. Why struggle with Cassandra clustering when I can simply point to the nearest endpoint for FaunaDB?

- Hareem Haque, Engineer, Computes.io