

Theseus a Composable distributed execution runtime: Performance across GPUs, Networks, and Storage (Keynote Talk)

Felipe Aramburu

Distinguished Architect, Co-Founder at Voltron Data

felipe@voltrondata.com

ABSTRACT

During this talk we will discuss how Theseus, a composable distributed execution framework, leverages composability to be scalable and performant across multiple compute accelerators and highly varied hardware configurations. We will start out by looking at two different systems. The first one is an on prem deployment with fast networking, high performant centralized storage, and larger GPUs. The second is by leveraging commodity hardware available in the cloud where networking is slower, data is read from slower object stores, and smaller GPUs are leveraged.

We will then discuss how composability allows us to map our software to the hardware we are working with. First we will discuss two of our communication protocols one leveraging UCX and the other leveraging TCP with boost asio and how these perform on our two respective systems. After this we will discuss two different abstractions for reading bytes from input files. The first leverages GPU Direct Storage with Weka in an on prem deployment and the second is used to read files that are stored in S3. Lastly we will discuss using composability to target multiple hardware accelerators showing a back end agnostic task execution framework

that allows you to create tasks that target multiple backends and memory spaces in a single process.

Speaker Biography: Felipe has been working on accelerated SQL engines and CUDA computational primitives for over a decade. At Voltron Data, he is the Architect of Theseus, a distributed SQL engine that runs on AMD and Nvidia GPUs. His main focus is optimizing DataFlow architectures in the context of distributed ETL, AI, and ML.

VLDB Workshop Reference Format:

Felipe Aramburu. Theseus a Composable distributed execution runtime: Performance across GPUs, Networks, and Storage (Keynote Talk). VLDB 2025 Workshop: Third International Workshop on Composable Data Management Systems.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment. ISSN 2150-8097.