

Databases: BlazingSQL, Apache Cassandra, CockroachDB

By *Roy Schestowitz*

Created *11/08/2019 - 4:16pm*

Submitted by Roy Schestowitz on Sunday 11th of August 2019 04:16:20 PM Filed under [Server](#) [1]

- [BlazingSQL, a GPU-accelerated SQL engine built on top of RAPIDS, is now open source](#) [2]

Yesterday, the BlazingSQL team open-sourced BlazingSQL under the Apache 2.0 license. It is a lightweight, GPU-accelerated SQL engine built on top of the RAPIDS. ai ecosystem. RAPIDS. ai is a suite of software libraries and APIs for end-to-end execution of data science and analytics pipelines entirely on GPUs.

Explaining his vision behind this step, Rodrigo Aramburu, CEO of BlazingSQL wrote in a Medium blog post, "As RAPIDS adoption continues to explode, open-sourcing BlazingSQL accelerates our development cycle, gets our product in the hands of more users, and aligns our licensing and messaging with the greater RAPIDS.ai ecosystem."

Aramburu calls RAPIDS "the next-generation analytics ecosystem" where BlazingSQL serves as the SQL standard. It also serves as an SQL interface for cuDF, a GPU DataFrame (GDF) library for loading, joining, aggregating, and filtering data.

- [GPU SQL engine BlazingSQL now open source](#) [3]

A new open-source project wants to take analytics to the next level. BlazingSQL is a GPU-accelerated SQL engine built on the RAPIDS ecosystem. RAPIDS is an open-source suite of software libraries for executing end-to-end data science and analytics pipelines entirely on GPUs.

According to the team, BlazingSQL was built to address the expense, complexity and sluggish pace users deal with when working on large data sets.

BlazingSQL addresses these customer concerns not only with an incredibly fast, distributed GPU SQL engine, but also a zealous focus on simplicity,? Rodrigo Aramburu, CEO of BlazingSQL, wrote in a blog post. ?With a few lines of code, BlazingSQL can query your raw data, wherever it resides and interoperate with your existing analytics stack and RAPIDS.?

BlazingSQL enables users to query datasets from enterprise data lakes directly into GPU memory as a GPU DataFrame (GDF). GDF is a project that offers support for interoperability between GPU applications. It also defines a common GPU in-memory data layer.

-

[DataStax: what is a ?progressive? cloud strategy? \[4\]](#)

With its roots and foundations in the open source Apache Cassandra database, Santa Clara headquartered DataStax insists that it likes to keep things open.

As such, the company is opening a wider aperture on its collaboration with VMware by now offering DataStax production support on VMware vSAN, now in hybrid and multi-cloud configurations.

-

[Cockroach Labs raises \\$55 million for ultra-resilient databases \[5\]](#)

Cockroach Labs, the New York-based developer of the open source distributed database project CockroachDB, today announced that it's closed a \$55 million, oversubscribed series C round co-led by Altimeter Capital, Tiger Global, and GV (formerly Google Ventures). The raise, which saw participation from existing investors Benchmark, Index Ventures, Redpoint Ventures, FirstMark Capital, and Work-Bench, brings the company's total capital raised to \$108.5 million and comes after a year in which revenue doubled quarter-over-quarter.

[Server](#)

Source URL: <http://www.tuxmachines.org/node/126834>

Links:

[1] <http://www.tuxmachines.org/taxonomy/term/147>

[2] <https://hub.packtpub.com/blazingsql-a-gpu-accelerated-sql-engine-built-on-top-of-rapids-is-now-open-source/>

[3] <https://sdtimes.com/os/gpu-sql-engine-blazingsql-now-open-source/>

[4] <https://www.computerweekly.com/blog/Open-Source-Insider/DataStax-what-is-a-progressive-cloud-strategy>

[5] <https://venturebeat.com/2019/08/06/cockroach-labs-raises-55-million-for-ultra-resilient-databases/>