

At First Global StarRocks Summit, CelerData Open Sources StarOS and Multi-Warehouse and Introduces StarRocks V4.0

Milestone announcements signal a new era of openness and collaboration, inviting the global community to shape the future of real-time, AI-ready data platforms together.



Menlo Park, California Sep 12, 2025 ([IssueWire.com](https://www.issuewire.com)) - [CelerData](https://www.celerdata.com), the world's fastest, secure lakehouse engine for customer-facing and agent-driven analytics powered by StarRocks, today opened the first global [StarRocks Summit](#). More than 1,500 community members, contributors, and enterprise leaders came together to hear the latest developments in real-time intelligence on the lakehouse and to witness milestone announcements marking the next stage in the evolution of the StarRocks project.

At the Summit, CelerData announced a major commitment to openness by open sourcing StarOS and multi-warehouse features, unveiled the release of StarRocks V4.0, and introduced its vision for an architecture that is AI-ready by design. Together, these developments highlight the project's momentum and its role in preparing enterprises for an AI-driven future.

"AI is the most talked-about force of our time. It is reshaping how businesses, societies, and individuals interact with technology, yet no one knows exactly what its ultimate form will be," said Alvin Zhao, Chair of the StarRocks Technical Steering Committee and CTO of CelerData. "Organizations everywhere are grappling with the twin challenges of explosive data growth and the uncertainty of an AI-driven future. They know fresher, faster insights are essential to competitiveness, but too often they face mounting complexity and cost in trying to unify diverse data sources for both business intelligence and AI workloads. In the face of this uncertainty, the only way to prepare is to be open. Open to new ways of working, open to community-driven innovation, and open to building a platform that is ready for whatever AI brings."

StarOS and multi-warehouse go open source

Perhaps the most significant announcement of the Summit was CelerData's decision to open source StarOS and multi-warehouse features. StarOS, a fundamental project within StarRocks, enables the separation of compute and storage, the architectural shift that powers the platform's scalability and adaptability. Multi-warehouse builds on this by allowing enterprises to create independent, isolated compute clusters while still sharing the same dataset within a single StarRocks cluster.

This means organizations can run diverse workloads, from real-time analytics to customer-facing dashboards to AI inference, all on the same underlying data without duplication or compromise. By open sourcing both StarOS and multi-warehouse, CelerData is not only making these capabilities available to

the community but also inviting developers and enterprises to co-create the platform of the future.

“StarOS is the foundation that enables StarRocks to separate compute from storage, and multi-warehouse is the architecture that allows enterprises to isolate and scale their workloads without duplication,” Zhao said. “These were once commercial features. By open sourcing them, we are making a clear statement: we believe the future of data platforms cannot be built by one company alone. It must be built by a community. This decision is our commitment to that belief. We are opening the door for developers everywhere to take part, innovate, and push the boundaries of what StarRocks can achieve.”

StarRocks V4.0: a faster, more unified, and AI-ready platform

The release of StarRocks V4.0 represents the most significant leap forward in the platform’s capabilities since its inception. This version deepens StarRocks’ position as both a lakehouse and AI-ready engine, combining technical advances with community-driven optimizations that have produced a 60% year-over-year performance improvement.

Among its highlights, V4.0 introduces a faster JSON handling mechanism, critical for the dynamic, unstructured data that underpins many AI workloads. It also brings the general availability of vector search, enabling enterprises to run similarity queries that are at the core of generative AI and recommendation engines. Additionally, the new release delivers deeper integration with Apache Iceberg, ensuring users can not only query but also prepare and optimize data directly within StarRocks without the need for external tools.

“StarRocks V4.0 embodies our vision of a platform that unites BI and AI in a single, open data lake,” Zhao said. “Our focus has always been on eliminating silos, reducing complexity, and ensuring that performance does not come at the expense of flexibility. With vector search, deeper Iceberg integration, and performance gains that have been achieved together with our community, we are delivering a platform that is ready for the future.”

Preparing for an agent-driven era

The keynote spotlighted how the next wave of AI development is being defined by agents, intelligent systems capable of handling context-rich tasks at scale. Unlike human users, agents can process far more data simultaneously, require real-time access to fresh information, and demand seamless integration across diverse data types.

StarRocks has been built from the ground up with these requirements in mind. The engine already unifies data access, delivers sub-second query latency, and integrates seamlessly with open formats like Iceberg, Hudi, and Delta Lake. These principles are being extended so enterprises can thrive in an agent-driven future.

“Agents need more than just speed. They need real-time context, concurrency at scale, and access to unified data sources that can provide a rich tapestry of information,” Zhao explained. “These are the exact principles that guided StarRocks from the beginning. As AI agents evolve, StarRocks is already aligned to be the foundation they need.”

A commitment to community

Since its first release in 2021, StarRocks has evolved into a thriving open source project, trusted by

more than 500 enterprises worldwide and supported by a global community of over 500 contributors and nearly 5,000 Slack members. The decision to open source core components such as StarOS demonstrates a deepened commitment to ensuring the platform's evolution is shaped not just by CelerData, but by the wider ecosystem.

“StarRocks has always been more than just a fast SQL engine. It is a community effort, a collective project where engineers, enterprises, and innovators from around the world have come together to solve some of the hardest problems in data,” Zhao said. “Our role at CelerData is to help guide, support, and accelerate that progress, but the real power lies in the community. With today's announcements, we are giving that community more tools and more freedom to shape the future together.”

About CelerData

CelerData (powered by StarRocks) is the fastest query engine for customer-facing and AI-driven analytics at petabyte scale. Natively integrated with Apache Iceberg, Apache Hudi, and Delta Lake, it delivers low-latency, high-concurrency queries directly on open data, without ingestion delays or costly pipelines. Trusted by industry leaders like Pinterest, Tencent, and Expedia, CelerData powers the next generation of analytics on the lakehouse. Learn more at: www.celerdata.com

Contact:

Tyler Wishnoff

Marketing Lead

tyler.wishnoff@celerddata.com



Media Contact

Waters Agency

*****@watersagency.com

8183326166

1002 Sundance Ridge Rd.

Source : CelerData

[See on IssueWire](#)

