



Ridge AI Launches Agent-Built Embedded Analytics That Run in the Browser

Abstract

Ridge AI announced general availability of its embedded analytics platform on September 15, 2026. Ridge lets teams build dashboards (called “Ridges”) for customers, executives, and other general audiences by conversing with a Build Agent, then embed them in their own applications with a paired Explore Agent that answers end-user questions in natural language.

Ridge connects and prepares data in a hosted web application, then runs analytics in the viewer’s browser. Ridge loads each dataset into DuckDB running as WebAssembly and uses the open-source Mosaic framework to compute filters and aggregations locally. Ridge’s partitioning is designed to deliver each viewer only that viewer’s slice of the data. The GA release builds on an April closed beta, extended since with a more capable Build Agent, the Embed SDK that sets initial dashboard state from user identity or context, direct editing of charts, layout, text, and sorting, and support for larger datasets in the browser.

Ridge’s co-founders are CEO Ellie Fields, a 12-year Tableau veteran and former Salesloft product leader, and chief scientist Jeff Heer, the University of Washington professor who co-created Mosaic, Vega-Lite, and D3.

Context

Embedded analytics carries a different burden than internal BI. Every chart becomes part of the end-user experience, so it must load quickly, stay inside the defined data boundary, and hold up without an analyst sitting directly behind it. Server-based embedded tools commonly execute interactions on remote compute, which adds latency and cost to every filter or query.

Conversational AI has added a new way to query data, but organizations still share data through dashboards. A chat answer serves one person and one question, while a dashboard presents a uniform view to every viewer and persists between sessions. Requests for new dashboards typically route through a central data team, and that team’s capacity sets the pace of delivery.

Key Ramifications

The following are the key ramifications of the release, informed by EMA’s hands-on testing.

Agent-built dashboards move embedded analytics from data teams to the business. Ridge’s Build Agent works from business intent: it profiles the data, proposes the questions a dashboard should answer, and builds once the owner approves. Point-and-click editing and unlimited undo then let that owner refine the draft by hand instead of re-prompting. Builders can also attach reference documents, such as data dictionaries, metric definitions, and product context, so a Ridge reflects the organization’s own terms and meanings. In EMA testing, plain-English metric definitions carried through to the generated SQL on the first build.

The dashboard and the agent ship as one unit.

Each Ridge pairs a curated dashboard with an Explore Agent working on the same dataset, with optional linked interactivity between dashboard selections and Explore. End users ask follow-up questions that would otherwise become requests for new dashboards, and Ridge says it checks the generated SQL before executing it locally in the viewer’s browser.

The browser now carries more of the workload.

Mosaic coordinates every chart and filter against a local DuckDB instance. The practical ceiling depends

on the shape of the data rather than a fixed row count: datasets in the low millions are typical, partitioning shrinks each viewer's slice, and Mosaic offers a path to push heavier work to a server. EMA built dashboards on files ranging from Ridge's 222,000-record sample up to 3 million records in a single unpartitioned dataset, with filters, transformations, and calculated metrics.

Use-case pricing and browser compute change the cost model. Ridge prices per use case: every published tier includes up to five Ridges (each one an embedded dashboard with its Explore Agent) and steps up by embeds, in-app users, monthly views, and partitioning, which starts at the Standard tier. Its browser-side engine removes the per-interaction warehouse cost of server-based tools. Teams can forecast what shared analytics will cost as audiences grow.

Ridge composes with the data stack. Ridge connects to Snowflake, Databricks, BigQuery, Iceberg, and object stores, and leaves semantics, transformation, and governance where they already live. For multi-tenancy, Ridge scopes each embedded session to its tenant with JWTs, isolates customer data at the row level, and reports a completed SOC 2 Type 2 audit. EMA connected to Parquet files on Cloudflare R2 through Ridge's native connector.

EMA Perspective

Most AI analytics products compete on how much data an agent can reach. Ridge competes on curated speed, delivering only what the agent and consumer need, in a framework that can live in the browser, compute locally, and remove the warehouse roundtrip. The platform produces dashboards built on explicit definitions, SQL a data team can audit, and interactions that resolve locally.

Mosaic is the engine behind Ridge's speed, and in EMA's testing it delivered sub-second dashboard interaction on datasets ranging from Ridge's 222,000-record sample to an EMA-supplied 3-million-record dataset, where a multi-second wait on a filter would undermine the experience. Competitors can adopt the open-source framework, but matching Ridge would mean rebuilding their architecture around it and pairing it with an agent that writes to it.

Ridge also addresses the build-versus-buy question product teams now face as AI coding tools make custom dashboards look easy. The hard part of customer analytics is tenancy, validation, and upkeep. EMA went from raw files in object storage to a working dashboard with no infrastructure to stand up and no modeling layer to build. Ridge is built for sharing data with broad audiences; open-ended exploration stays in notebooks and BI tools. The platform also draws a useful organizational line: data teams own the curated dataset and its definitions, while product and field teams build and ship the story on top of it.

Buyers should adopt Ridge the way its architecture intends. Curate the dataset upstream, sized to the audience and experience, partition by the boundaries that matter, review the generated SQL, then finish with direct edits. That structure will get the most from the platform's speed and transparency. Those replacing an incumbent embedded BI tool should prove it against their own customer data during the 30-day trial, which includes partitioning.

With this release, Ridge positions itself as the embedded analytics layer for broad data sharing: AI-built dashboards grounded in curated data and designed for speed, control, and audience trust.



About EMA

Founded in 1996, Enterprise Management Associates (EMA) is a leading IT research and consulting firm dedicated to delivering actionable insights across the evolving technology landscape. Through independent research, market analysis, and vendor evaluations, we empower organizations to make well-informed technology decisions. Learn more about EMA research, analysis, and consulting services at www.enterprisemanagement.com or follow EMA on [X](#) or [LinkedIn](#).

4802.101326