Americas Technology: Software: Communacopia + Technology 2026 — Key Takeaways

Gabriela Borges, CFA
+1(212)902-7839 | gabriela.borges@gs.com
Goldman Sachs & Co. LLC

Maura Hager
+1(212)9028724 | maura.hager@gs.com
Goldman Sachs & Co. LLC

Selina Zhang
+1(212)357-9979 | selina.zhang@gs.com
Goldman Sachs & Co. LLC



Presenter: Ali Ghodsi (CEO and Cofounder), Databricks

Bottom Line: The AI ecosystem is shifting from model capability to enterprise adoption, and Databricks noted it is positioning itself at the center of that transition to solve the gap between agent capabilities in theory vs. practice. Genie is helping customers operationalize AI, Unity Catalog provides the governance and cost-control layer, and expansion into CDP, SIEM and transactional databases broadens the platform beyond analytics. **We have previously written about "good sticky" vs. "bad sticky" companies, and we believe technical IP, pace of innovation, and open architecture are all key in assessing whether any one company can be successful with share gain over the next 3 years.**

Key Takeaways

- **Databricks' View: Genie may be the product to close the gap between AI capability and enterprise adoption.** While AI is already capable of performing many tasks at or above human levels, very few enterprises have successfully enabled agents into day to day workflows at scale. Databricks increasingly sees its role as helping customers close that gap. Genie is a key use case: rather than competitive alternatives that simply query and retrieve product, Genie acts as an AI analyst that computes answers directly from enterprise data, and is able to do so because of OntoRank (its proprietary data indexing on enterprise data, comparable to Google search for web). Databricks itself is already using Genie extensively across sales, marketing and finance, with internal operations changing meaningfully over the last 6-12 months as a result. Genie runs on an ontology, but rather than competing with Palantir, Databricks sees a partnership where it can work with Palantir on services-heavy infrastructure-plumbing projects that translate CEO mandates into data engineering problems. More broadly, Databricks' view is that enterprise AI is ultimately a data problem: the challenge is ensuring data is organized, governed and accessible for AI to

Goldman Sachs does and seeks to do business with companies covered in its research reports. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision. For Reg AC certification and other important disclosures, see the Disclosure Appendix, or go to www.gs.com/research/hedge.html. Analysts employed by non-US affiliates are not registered/qualified as research analysts with FINRA in the U.S.

generate business value.

- **Governance and Unity Catalog.** Governance is a growing customer priority. There are three trends converging simultaneously: 1) AI models are evolving rapidly, creating complexity around model selection and deployment; 2) token consumption is accelerating, increasing scrutiny around AI spend and ROI; and 3) enterprises have limited visibility into what agents are doing, how data is being accessed and where security risks may emerge. Against that backdrop, Unity Catalog is one of the company's most important products and increasingly a key reason customers choose the platform. Governance, access control, security and semantics need to sit alongside data and AI rather than being separate layers. Databricks also emphasized the importance of ontology and semantic layers, arguing that providing AI with the right business context may be one of the largest remaining unlocks for enterprise adoption.
- **Expansion into new software verticals.** Data engineering isn't for the sake of data engineering; the infrastructure is designed to solve business problems. There are three application software categories that F500 customers are already pulling Databricks in, by using Databricks to build their own proprietary stacks:
  - Customer data platforms (CDP): Marketing has already evolved from monolithic apps to composable apps; LLMs introduce new capabilities to level up hyper-personalization at the individual level rather than the cohort level. Databricks enables this with Infinite Campaigns, and open source models will drive the next step function change in cost viability.
  - Security information and event management (SIEM): historically, infrastructure companies have had limited success in security use cases because of the domain experience and targeted CISO selling required. However, today, the CIO and CISO roles are converging, and new AI use cases create a new set of agentic data to analyze, which may not leverage existing security domain experience and instead require new data domain experience.

  Observability is the third, but Databricks did not share any plans on timing of a product. The company noted that there are other categories that may make sense over time.
- **Lakebase** is potentially the company's largest new product opportunity. Databricks notes that more Software may be created in the next eight months than has already been created in the history of Software. The company stated that this will materially increase demand for transactional databases, creating a significant opportunity for Databricks to extend its platform beyond analytics and into operational workloads.

