CLOUDERA

Cloudera vs. Acceldata

	CLOUDERA	Acceldata
1 ## Different by Design **Cloudera** is a complete, enterprise data and AI platform. **Acceldata** launched its 'Autonomous Data & AI Platform' on May 19, 2026, anchored in observability and ODP.	**Large. Established. Proven.** The unified hybrid data and AI platform for the enterprise. ✓ 15+ years of data platform innovation. ✓ Thousands of enterprise customers. ✓ 30 exabytes managed in production. ✓ FedRAMP Moderate. GovRAMP. TXRAMP Level 2. SOC 2. ISO 27001. FIPS 140-2 validated cryptography. ✓ Private, public, and hybrid flexibility.	**Observability pivoting to platform pitch.** Pivoting in 2026 toward an 'Autonomous Data & AI Platform' pitch, claiming the lakehouse era is over. • Founded in 2018. • Fewer than 300 employees. • Limited public customer disclosure compared to Cloudera's thousands of enterprise customers. • Roots in observability. Layers ADM, xLake, and Agentic Runtime on top. • Relies on third-party vendors for catalog, warehouse, lakehouse, and AI.
2 ## Platform Scope **Cloudera** delivers an end-to-end platform for the entire data and AI lifecycle. **Acceldata** customers typically assemble additional vendors to achieve the same outcomes.	**One platform. End to end.** ✓ **Governance & Catalog:** Secure, trusted, compliant data with SDX. ✓ **Streaming & Ingestion:** Real-time data at any scale (NiFi, Kafka, Flink). ✓ **Data Warehouse:** High-performance SQL analytics, on-prem and cloud. ✓ **Lakehouse:** Open formats. Full Iceberg across Spark, Hive, Impala, lakehouse runtime. ✓ **ML / AI Lifecycle:** Build, train, deploy, monitor with governance. AI Inference on-prem (GA). ✓ **Unified Platform:** Shared governance, lower complexity, faster time to value.	**Built on third-party components.** • **Governance & Catalog:** Collibra, ServiceNow, Informatica typically required. • **Streaming & Ingestion:** Kafka, Fivetran, Confluent typically required. • **Data Warehouse:** Snowflake, Databricks, Redshift typically required. • **Lakehouse:** Iceberg in ODP supported only via the Spark3 path. Not across Hive or Impala. • **ML / AI Lifecycle:** MLflow, Databricks, SageMaker typically required. • **Multiple Products:** More contracts, more integration work, typically slower outcomes.
3 ## Cost & Complexity **Cloudera** reduces TCO and operational complexity with one integrated platform. **A multi-vendor stack** around Acceldata can run substantially higher than a single integrated platform.	**Platform breadth and predictability.** ✓ Broad platform scope. One contract. ✓ Unified support. One roadmap. ✓ Predictable pricing. ✓ Native security and governance through SDX. ✓ Simpler to deploy. Easier to manage.	**Multi-vendor. Multiple moving parts.** • Integration complexity and risk. • Multiple vendors, multiple contracts, higher total cost of ownership at scale. • More to manage and maintain. • More chances for gaps across tools.
4 ## Agentic Reality Check **Cloudera** AI Inference, AI Studios, and AI Agents are GA today. **Acceldata's** May 19, 2026 Autonomous Data & AI Platform centers on the xLake architecture, hosting third-party agent frameworks rather than a Cloudera-equivalent governed AI stack.	**GA today. Native and governed.** ✓ **AI Inference:** GA, on-prem and cloud. ✓ **AI Studios + Agents:** Governed agentic workflows, hybrid. ✓ **ML:** Model catalog, governance, lifecycle. ✓ **Cloudera AI Workbench:** Build, train, deploy with full governance. ✓ **Trino-Backed Warehouse:** On-prem analytics, GA.	**Marketing vs. product reality.** • **xLake architecture:** xObserve, xReasoning, xGovern, xRoute, OS Foundry. • **Agentic Runtime** hosts third-party frameworks (LangGraph, CrewAI, AutoGen, OpenAI SDK) rather than a native, governed AI service. • **Pulse Sense** remains beta in Acceldata's release notes (4.1.0, March 2026). • **G2** product description markets GDPR/HIPAA/BCBS 239 enforcement 'without manual configuration.'
5 ## Performance Claims **Acceldata's** new 'Velox-accelerated' Data Warehousing product is marketed as 10x faster than traditional warehouses. Acceldata's own published Velox benchmark hits 1.31x at 100GB, BELOW the TPC-DS 1TB minimum. xLake's 'exabyte' positioning is backed by billion-row benchmarks (terabyte scale).	**One platform. End to end.** ✓ **Real Production Scale:** 30 exabytes managed across enterprise customers today. ✓ **Streaming Throughput:** NiFi, Kafka, Flink integrated at sub-second latency. GigaOm Leader. ✓ **Hybrid AI Inference:** GPU-backed AI workloads on-prem or cloud, GA February 2026. ✓ **Open Formats:** Full Iceberg integration across Spark, Hive, Impala, lakehouse runtime. ✓ **Lakehouse Optimizer:** Iceberg performance tuning built in.	**Benchmarks don't match the marketing.** • **Data Warehousing** product is marketed as 10x faster than traditional warehouses on Velox. Acceldata's own published Velox benchmark hits 1.31x at 100GB. • **Spark/Gluten/Velox benchmark** runs at 100GB, below the TPC-DS 1TB minimum. Upstream Apache Gluten reports 3.02x at 3TB, roughly double the lift at ten times the scale. • **ORC** support is partial. JSON and CSV can fall back. Structured Streaming isn't supported. Unsupported queries fall back to the JVM. • **xLake's 'exabyte'** claim is backed by benchmarks topping out at 6+ billion rows (terabyte territory), with no customer reference at exabyte scale. • **Apache Gluten** graduated to Apache Top-Level Project status in March 2026. Acceldata bundles it rather than authors it.

All comparisons based on publicly available Cloudera and Acceldata materials as of publication date.

Cloudera, Inc. | 6220 America Center Dr, 5th floor, San Jose, CA 95002 USA | cloudera.com

© 2026 Cloudera, Inc. All rights reserved. Cloudera and the Cloudera logo are trademarks or registered trademarks of Cloudera Inc. in the USA and other countries. All other trademarks are the property of their respective companies. Information is subject to change without notice. 9287-02 | May 20, 2026