RELTIO

DATASHEET

Why data warehouses and data lakes need Reltio for trusted AI

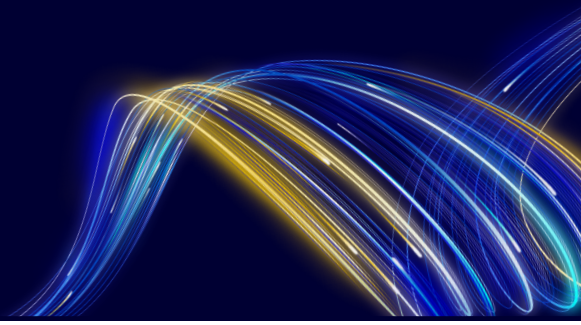

Cloud data warehouses and data lakes provide the backbone for analytics and AI/ML use cases. They collect and aggregate data from source systems to enable large-scale analysis and model training, but **collection is not the same as unification.** These platforms do not natively unify core business entities in real time. Customers waste time and resources to either build these capabilities themselves or to purchase and integrate with third-party tools for essentials like entity resolution, cleansing, governance workflows, and standardization. Without unification, they lack the knowledge graph and relationship intelligence needed for agentic AI. Highly critical to enable reasoning, personalization, and explainability.

This is where Reltio Data Cloud™ adds value. As a cloud-native data unification platform, it continuously unifies and curates data into trusted 360-degree profiles—accessible in milliseconds across operational systems, AI agents, and cloud data warehouses and lakes. So people, digital systems, and AI operate with the right context and deliver positive business outcomes. These platforms are complementary: Cloud data warehouses and lakes scale analytics and AI model development, while Reltio fuels them with unified, high-quality, and connected data for greater business impact.

	Cloud Data Warehouses and Data Lakes	Reltio Data Cloud
Primary purpose	Analytics first: Designed for large-scale storage and computation to drive BI dashboards, ML training, and advanced analytics. Best suited for use cases where aggregated, denormalized raw and historical data are needed for trend analysis, reporting, model training, and retrospective insights. **Limitation:** Does not natively deliver accurate, consistent data. Requires additional components.	Continuously ingests, unifies, standardizes, and enriches siloed data into clean, context-rich profiles mobilized across operational applications, AI agents, and analytical environments in milliseconds. Best suited for use cases where insights and actions are linked to unique entities and their relationships with the business; e.g., personalized offers for customers, supplier performance tracking, fraud detection. **Benefit:** Provides high-quality, real-time data to agentic AI, analytics, and operational systems for accurate decisions and actions.
Data model	Optimized for SQL queries, joins, and data science pipelines. Data warehouses provide data in rows and columns. **Limitation:** Focuses on raw and analytical data. Complex to model real-world entities and relationships due to handling one-to-many and many-to-many relationships using relational data models. Creates processing dependencies and often requires extra tools and manual steps.	Unifies multidomain data (e.g., customers, products, suppliers, locations, assets, and more) and multi-model data through a schema-flexible structure combined with a graph-based approach. Its simple but powerful logical data model is easy to understand by both humans and AI agents. Captures hierarchies from structured and unstructured data across entities, interactions, and transactions, and connects them across domains. **Benefit:** Provides 360° views for operational insights and AI-ready, context-rich data that power personalization, fraud detection, regulatory compliance, and AI models.

	Cloud Data Warehouses and Data Lakes	Reltio Data Cloud
Support for agentic AI	Provides AI agents with large amounts of data in batch mode, along with strong tools for model training. Offers AI agents for AI/BI. **Limitation:** Does not unify, standardize records, or show relationships between entities out of the box to provide the context needed for reliable insights. Can't adequately deliver on real-time, low-latency agentic scenarios.	Fuels AI agents with continuously curated, unified 360° profiles—interactions, unstructured data related to the entity, relationships across customers, products, suppliers, and more—with millisecond latency. LLMs and other AI systems can access trusted data securely with role-based access control (RBAC) and audit logging. Provides purpose-built AI agents for data steward and business workflows. **Benefit:** With Reltio, agentic AI runs on unified, clean, and context-rich data, enabling effective decisions and actions. Can serve data to an AI agent in 50 ms or less to fuel operations at scale.
Entity resolution	No native support for entity resolution. Requires building pipelines with custom AI/ML models, utilizing specialized technical skills, or integrating with third-party applications. **Limitation:** DIY approach demands heavy engineering effort, ongoing tuning, and significant compute resources. Difficult to adapt to multiple enterprise use cases or new requirements. Integrating third-party tools often involves: - Complex custom development - Managing multiple vendors - Higher costs - Increased project and security risks - Steeper learning curve	Built-in entity resolution augments the rule-based approach with pretrained ML-based matching (Flexible Entity Resolution Network). Easily unifies records from multiple sources with a semantic understanding to increase accuracy and reduce manual effort. Offers configurable survivorship rules to decide the "best version of truth" for different departments or systems. **Benefit:** Provides reliable and rich profiles across all domains with auditability and automated governance. Breaks down silos with minimal effort and flexibility to adapt to changing needs.
360° views	Do not natively create or maintain persistent 360° views. Relationship awareness and graph-based modeling require external tools. Building 360° views with relationships is highly complex due to the relational data model, which involves numerous 1:M and M:M tables. Data models are often designed for analytical use cases and lack survivorship logic or dynamic updates. Difficult to change these data models for other projects. No built-in audience segmentation capability. **Limitation:** Complex and costly to build and manage 360° views. Lack of agility to add use cases.	Delivers always-on 360° profiles for all data domains through continuous data quality validation, monitoring, and Reltio Intelligent Data Graph™, capturing and updating cross-domain relationships. Unique and rich profiles are enriched, connected in real time, with built-in audience segmentation capabilities. **Benefit:** Highly flexible, graph-powered 360° views that evolve as data changes—supporting personalization, agentic AI, regulatory needs, and operational precision across all systems.

	Cloud Data Warehouses and Data Lakes	Reltio Data Cloud
Data quality & trust	Relies heavily on ETL/ELT pipelines to cleanse and standardize data before loading. Bidirectional zero-copy sharing makes that data broadly and quickly available for analytics and collaboration without duplication. Limited ability for users to review data quality issues due to the lack of prebuilt interfaces. Instead, logic is embedded in the data pipelines, making it difficult for end users to diagnose and remediate issues. **Limitation:** Data issues go unnoticed and unresolved, eroding confidence in insights. Data governance becomes a strictly IT-led initiative—with insufficient business-level involvement and context to drive fast and positive business outcomes.	Continuous data validation, curation, and monitoring—including automatic checks and updates for emails, phone numbers, and addresses—so that data stays unified, clean, and explainable—not just passing a one-time quality check. Dashboards reveal patterns, anomalies, and metrics instantly. Reltio AgentFlow™ AI agents monitor data quality, suggest or execute fixes, and drive continuous improvement. Offers zero-copy data sharing of comprehensive profiles downstream cloud data lakes, and analytics platforms in milliseconds without the need for pipelines. Simple yet powerful user interfaces (and agents) for identification and remediation of data quality issues. Embeds multiple capabilities so that everything is available on day 1, without having to bolt on additional capabilities as new issues are found. **Benefit:** Reduces errors, rework, and risk. Clear visibility of data health, while operational, AI/ML, and agentic AI systems operate with greater accuracy, compliance, and confidence.
Integration	Ingests and stores data from many sources (databases, SaaS apps, APIs, streaming pipelines). Provides a robust ecosystem of partner ETL/ELT tools and zero-copy sharing for efficient distribution across teams, business units, and even external partners. **Limitation:** Pushing records back into operational systems (CRM, ERP, marketing platforms) typically requires heavy custom engineering or third-party solutions. Outbound integrations primarily serve advanced analytics/ML use cases.	Bidirectional integration: ingests data from multiple systems, unifies continuously, and provides accurate profiles back into operational and analytical platforms. API-first architecture includes no-code/low-code integration development via Reltio Integration Hub. Zero-copy inbound integration from Snowflake, outbound data sharing to Databricks, and prebuilt integrations for Salesforce, Veeva Vault, and third-party enrichment sources. **Benefit:** Enables all operational and analytical systems to consume the same trusted version of the truth in real time. Eliminates silos and reduces engineering effort via low-code/no-code integration.
Data governance	Provides data catalog primarily through native products. Native data governance capabilities include managing metadata, permissions and lineage. No user interfaces or data governance workflows available to remediate possible matches, merge duplicates, manage data conflicts, or unmerge when needed.	Built-in capabilities to govern core data entities, enforce general data policies, and support data stewardship activities. Some of the capabilities include: - Data lineage and audit trails show how customer or product records changed over time - Role-and attribute-level access controls, data masking

	Cloud Data Warehouses and Data Lakes	Reltio Data Cloud
Data governance (cont'd)	Absence of audit trails and lineage to track decisions made over time regarding merge/unmerge operations results in the inability to properly answer compliance questions. **Limitation:** No enforcement of governance policies on critical business data. No native stewardship workflows to validate or correct data. E.g., track a column called *customer_id* back to its source, but no assurance that ID is the same customer as in another system.	- Policy enforcement for GDPR, HIPAA, and other privacy regulations - Native integration with best-of-breed enterprise governance ecosystems—e.g., Microsoft Purview, Collibra, Alation—to exchange metadata and provide a unified governance fabric. **Benefit:** Enables access to high-quality data across the enterprise with compliance and auditability. Business users, data stewards, and AI agents can use the system to define and execute governance rules
Data latency	With streaming pipelines and bidirectional zero-copy integration, cloud data warehouses can make data available to other systems in near real-time for analytics and decisioning. **Limitation:** Ingestion and data sharing can be low latency, but manually coded data quality steps are batch processed. Unable to provide low-latency APIs with clean data needed for agentic AI.	Continuously ingests, unifies, enriches, and shares high-quality data with low latency across operational, AI/agentic systems, and cloud data warehouses. Offers streaming APIs and zero-copy sharing of 360° profiles downstream across cloud data warehouses and data lakes. Reltio Lightspeed™ Data Delivery Network further reduces latency by serving accurate core data with consistent, ultra-low response times (sub-50ms for many queries). **Benefit:** Enriched, high-quality data in milliseconds so operational workflows always act on reliable, context-rich information.
User audience	Primarily used by data scientists, data engineers, and ML engineers, who build data pipelines, train models, and run advanced analytics. Data teams also leverage this to ingest, store, and share large datasets. SQL is widely required—or the default—for querying, ETL pipelines, advanced data processing, BI and reporting, optimizations, and more. **Limitation:** Dependency on technical expertise—business users depend on curated datasets or dashboards prepared by IT or analytics teams.	Used by business teams, data stewards, governance teams, and IT to manage and govern core enterprise data. Increasingly used by AI/ML teams and agentic AI applications for auditable and context-rich data. **Benefit:** Democratizes access to high-quality, connected data for both business and IT teams, eases collaboration, and enables rapid innovation.
Process scalability	Data quality development introduces error risk, tech debt, and rigidity. Each new use case adds complexity, requires changes to the original solution, and increases the effort to maintain it. **Limitation:** Cannot support multiple and evolving use cases as fast as the business needs.	Flexible, graph-based data model, AI-powered automation, prebuilt integrations, and purpose-built AI agents for data unification and governance enable unlimited use cases and rapid solution modifications. **Benefit:** Supports evolving and increasing business needs for all parts of the enterprise, enabling competitive advantage in the market.

	Cloud Data Warehouses and Data Lakes	Reltio Data Cloud
Time to value	To provide accurate, context-rich, and real-time data for agentic AI, these systems require significant development efforts with technical skills or the acquisition and integration of multiple third-party products. **Limitation:** Building or integrating multiple third-party tools delays execution and distracts from higher-impact, strategic initiatives that drive growth and competitive advantage.	Purpose-built for data unification and to provide a system of context for agentic AI. With end-to-end capabilities—ingestion, entity resolution, continuous data quality management, 360° views with relationships and interactions, reference data management, segmentation, zero-copy sharing, and integration with AI and other systems. With prebuilt industry- and domain-specific components, MVPs can go live in 90 days and undergo rapid modifications to add new use cases. **Benefit:** Fast time to value and agility to add new use cases and modify as needed. Enables competitive advantage to rapidly innovate and stay ahead of the market.
Total cost of ownership	Building and maintaining homegrown extensions increases operational costs as business needs and technologies change. Requires new, specialized skills and introduces business risk to maintain a fragile solution. **Limitation:** Ever-increasing cost and risk to the business. SaaS cost unpredictability with consumption-based pricing.	Automated, SaaS-based comprehensive capabilities, including AI agents, built-in security, and high availability, minimize system engineer, data steward, and integration development resource needs. Enables the use of critical human talent in initiatives that drive the business forward. **Benefit:** Reduced operational costs, ability to use experts in critical business initiatives. Predictable pricing based on SaaS model using unified profiles.

Better together

Reltio complements cloud data warehouses and data lakes by providing high-quality, governed data—rich in relationship intelligence and with built-in AI/ML for entity resolution and stewardship, as well as continuous data quality management. This enables analytics, agentic AI, ML models, and operational systems to all run on consistent, reliable, and auditable data. With Reltio, cloud data warehouses can focus on delivering insights, rather than having to first unify, govern, and build trust in data from multiple disparate systems.

By using Reltio Data Cloud for data unification, you can avoid the lengthy and costly process of building the necessary capabilities to provide the agentic AI and other operational systems with unified, reliable data in milliseconds. Using the latest innovation via a cloud-native SaaS platform, you can focus your experts on using high-quality data to drive business impact

US +1 (855) 360-3282
UK +44 (800) 368-7643

@Reltio
facebook.com/ReltioHub
linkedin.com/company/reltio-inc