

## Reltio Connected Data Platform

Performance and scalability for your enterprise needs.

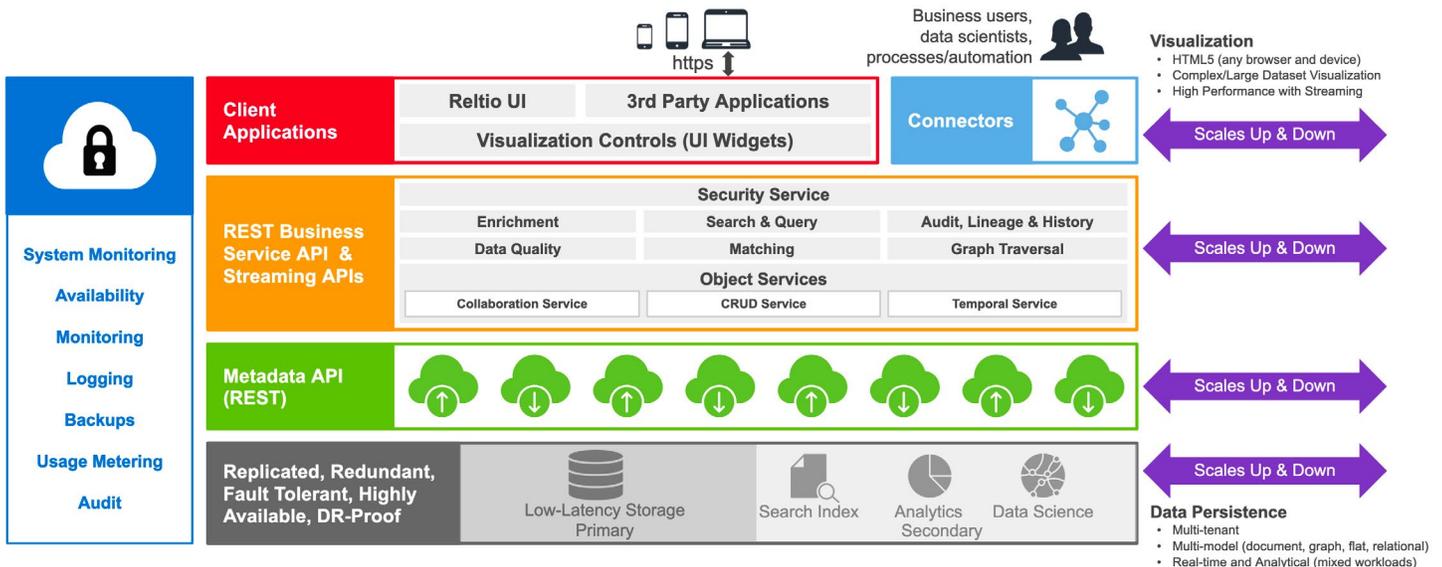


Today's most successful large enterprises are driven by data. They have harnessed their data to gain a competitive edge, make better decisions faster, and deliver a great customer experience. We help many of the largest enterprises unify, manage, and mobilize their data to accelerate growth, increase efficiency, and minimize risk.

Our AI-powered, business-responsive data unification and management solution offers the high scalability required to support the operational use cases of some of the largest enterprises for whom scalability and accelerating real-time operations are essential. Reltio Connected Data Platform is a cloud-native data management platform that supports billions of customer, supplier, location, and other profiles—enriched with thousands of attributes, relationships, transactions, and interactions from hundreds of data sources. We power enterprise-class, mission-critical applications to operate 24/7 with thousands of internal and external users. With trusted, interoperable data available to your users when they need it, you can maximize the value of your data while ensuring that it is always secure.

Our platform scales horizontally to deliver elastic performance and can support the throughput that you need for any operational or analytical use case. We support usage and consumption-based models where you can—on demand—scale up or scale down capacity based on your workloads. The scale up and down models enable you to utilize burst capacity to handle scenarios such as initial data loads and spikes in usage.

Traditional MDM systems are typically built as monoliths and are difficult to scale. Our responsive data unification and management solution is built from the ground up to be horizontally scalable at every layer. It's a microservices-oriented architecture where we can scale the independent services. As the peak capacity demands emerge with varying performance requirements, we can address those while managing infrastructure costs and giving you the best value for your dollar.



## Elastic, reliable, resilient, and secure

Our platform leverages highly scalable technologies like a NoSQL database, Elasticsearch, and Spark along with other AWS, Azure, and GCP services across multiple zones that ensure a high level of redundancy, fault tolerance, and availability. It supports high-volume transactions, high-volume API calls, as well as sophisticated analytics, and manages backend jobs for any workload. These components support an auto-scaling cloud environment.

As demand increases, the infrastructure, where possible, also automatically increases to accommodate the workloads. By taking advantage of the elastic cloud technology stack, our platform is able to scale up and meet service level agreements as well as scale down when the demand is lower, to manage your costs. And we add guardrails to make sure that there are no runaway costs by putting default thresholds on the autoscaling. Our Reltio Connected Data Platform's user-facing applications provide a modern experience optimized for different devices and seamless user collaboration. This approach ensures high performance while visualizing large datasets with complex many-to-many relationships.

## Low-latency APIs

One of the key aspects of our platform's performance is the low latency of its APIs.

We provide benchmarks and various factors that affect the latency of its seven API types. These APIs have different purposes and characteristics, such as loading data, finding entities, discovering relationships, and finding connections. The latency of each API depends on various parameters, such as the batch size, the entity size, the number of attributes, the number of RDM lookups, the overall API response size, the filter conditions, and more.

We also provide a performance monitoring application, where you can monitor the API request status, API latency, API success rate, and the number of failed API requests for each tenant. By understanding the baseline latency and the factors that affect it, you can plan your implementation and optimize your performance.

## Performance expectations

This guidance outlines factors that can contribute to variations in performance. Our platform can provide throughput and performance that may be needed by the customer by dialing it up or down. Customers can dial up the API throughput as required and have the option to request additional performance capacity. All operations are subject to varied results based on the size of entities and other factors defined in a tenant. Larger entities will result in slower throughput.

## Black Friday 2023 statistics

Looking at our platform's performance over the busy Black Friday and Cyber Monday week, see these results from our customers:

- Total of 306 million API calls
- Retail apparel enterprise: API latency was P90 at 61 milliseconds
- Vehicle retailer: API latency was P90 at 91 milliseconds
- Fast food restaurant: API latency was P90 at 98 milliseconds
- Average API availability SLA was above 99.99%

### SPEED

API Calls	Description	90th percentile baseline latency, milliseconds
POST /entities	Load data into tenant	300
GET /entities	Get entities using a filter	360
GET /entity/{id}	Find entity by ID	110
GET /entity/_matches	Find all matches for an entity	230
GET /entity/{id}_tree	Discover tree	550
GET /entity/{id}_hops	Get Hops	190
POST /entities{id}/_connections	Find Connections	155
POST /entities{id}/_searchConnections	Find objects and their relationships in a single call	930
POST /entities{id}/_findConnectedParties	Find connected objects and their references	220

Source: <https://docs.reltio.com/benchmarks/apilatenencybenchmark.html>

## Scalability by the numbers

As of December 2023, the Reltio Connected Data Platform is managing more than 8 billion entities and more than 100 billion relationships across those entities. This amounts to more than 2 PB of data under management. Since our platform supports real-time operational use cases, real-time performance is critical to your success. At Reltio, we have supported more than 56.1 billion API calls over the last 12 months.

Our Reltio team proactively monitors and manages these performance parameters, making sure you have the high performance and low latency you expect.

Data Under Management			
Consolidated profiles	Total raw profiles	Raw records processed (last 30 days)	Total API calls (last 12 months)
<b>8.1B</b>	<b>27.2B</b>	<b>17B</b>	<b>56.1B</b>
Number of API Calls (last 30 days)			
Search API	Get Object API	Find Match API	Save API
<b>422.3M</b>	<b>2.1B</b>	<b>399.3M</b>	<b>1.3B</b>

Our customers are running their mission-critical operations on Reltio Connected Data Platform, demanding consistent and reliable performance. Our customers not only bring their master data—across customer, supplier, product, and other domains. But for the customer domain, they also bring in their interaction and transactional data to learn more about their customer behavior, preferences, and intent. Blending omnichannel interactions and transactions is not typical for legacy MDM where you're usually restricted to bringing only your master data and encouraged to leave your transactional data at home. In contrast, we encourage customers to bring it all into one single system for building true connected 360 profiles.

Our Reltio Connected Data Platform delivers unified, interoperable data when and where it's needed. With the high performance and the scalability you need to power your real-time applications and downstream analytical and cognitive systems. Real-time, trusted data empowers you to improve business agility, increase efficiencies, manage risk, and drive growth.

Reltio leverages highly scalable persistence technologies like a NoSQL database, Elasticsearch, and Spark along with other cloud services that ensure a high level of redundancy, fault tolerance, and availability.

## Factors affecting performance

### Data Loads

- Entity size - larger entities slow down data load
- Match rule configuration - more matches identified by the match rules slow down performance
- Reference attributes - non-atomic attributes that reference properties of another entity—larger or more reference attributes can slow down performance

### Data Extract

- Entity size - larger entities slow down extraction

### Matching

- Match rule configuration - more matches identified by the match rules slow down performance
- Data distribution - more matches between entities will take longer to process through the match engine

### Reindexing

- Entity size - larger entities slow down reindexing

## ABOUT RELTIO

At Reltio, we believe data should fuel business success. Reltio's AI-powered data unification and management capabilities—encompassing entity resolution, multidomain master data management (MDM), data products—transform siloed data from disparate sources into unified, trusted, and interoperable data. The Reltio Connected Data Platform unifies and delivers interoperable data where and when it's needed, empowering data and analytics leaders with unparalleled business responsiveness. Leading enterprise brands—across multiple industries around the globe—rely on our award-winning data unification and cloud-native MDM capabilities to improve efficiency, manage risk, and drive growth.

To learn more, visit [www.reltio.com](http://www.reltio.com)

US +1 (855) 360-3282

UK +44 (800) 368-7643

 @Reltio

 facebook.com/ReltioHub

 linkedin.com/company/reltio-inc