dremio | **GRANICUS**

CASE STUDY

Granicus Unifies Government Data Analytics with Dremio's Zero-Copy Architecture

At a Glance

The Customer

GRANICUS

Challenge

Granicus struggled with severe data fragmentation across 40+ sources over their platform due to multiple acquisitions. Lacking unified customer identifiers, Granicus struggled to deliver real-time analytics to government agencies.

Solution

Granicus selected Dremio over Trino for its superior query performance, zero-copy architecture, and cost efficiency. This resulted in a three-tier lakehouse, unifying data for the "Citizen 360" profile. Granicus drove much success by strong data governance, strategic reflection optimizations for sub-second queries, and Dremio's Model Context Protocol enabling AI with Claude AI.

Results

- 10-1000x faster query response times at the business layer

- 40+ disparate data sources unified without data duplication

- 22 billion annual messages powered by real-time analytics

- Sub-second query performance on millions of records with optimized reflections

The Customer

Granicus is a leading cloud-based technology provider. The company's solutions connect 500,000 government employees with citizens across the United States. The company's Government Experience Cloud includes four flagship products: GovDelivery (digital communications), GovMeetings (public meeting streaming), GovAccess (website content management), and GovRecords (digital records processing). Operating at massive scale, Granicus processes 68 million forms annually, streams 200,000 hours of government meetings, and manages 300 million subscribers across 300,000 sign-up locations.

The Challenge

Granicus faced a critical challenge due to pervasive data fragmentation across its product portfolio. This prevented the company from providing their customers comprehensive citizen engagement analytics. For example, their flagship GovDelivery product, running on Oracle, held a large amount of inconsistent data, while survey data from their engagement application, eHQ, often lacked respondent identification, making entity resolution impossible. Their newer GXif product, designed to gather real time feedback from users, further complicated the situation by using a completely different set of customer identifiers.

Most critically, the absence of a universal customer ID made it impossible to unify citizen interactions across platforms. A single customer was inconsistently labeled as a "client" in GovDelivery, a "site" in eHQ, and a "customer ID" in GXif, and even the standard Salesforce ID could not be consistently located. Without the ability to connect these disparate touchpoints, Granicus was unable to provide the cross-platform analytics that government clients needed to optimize communications and measure engagement effectively.

Granicus evaluated platforms such as Starburst, Snowflake, and Databricks, but each came with limitations that didn't fit their acquisition-driven, rapidly scaling data environment. As an organization that constantly absorbs new systems and unpredictable data volumes, they needed a low cost, high performing platform that could scale elastically without forcing them to copy, centralize, or remodel data. Dremio provided the necessary performance, flexibility, and economic efficiency: fast analytics directly on lake storage, zero data movement, and rapid onboarding for scaling the business.

The Solution

After extensive benchmarking, Granicus selected Dremio based on three decisive advantages: superior query performance compared to Trino-based engines, a zero-copy architecture that enables direct source connections, and the ability to scale cost-efficiently across large workloads. These capabilities form the foundation of Granicus's new Unified Analytics Platform, developed in partnership

> "We evaluated Starburst, Snowflake, and Databricks, but Dremio's query engine was faster for our datasets, and the zero-copy architecture meant we could connect directly to sources without duplicating data. For a company built through acquisitions with unpredictable data volumes, that was critical."
>
> — Gaurav Gupta, CTO, Nuaav (Implementation Partner)

with Nuaav, which represents a significant leap forward in how the organization manages and operationalizes data.

At the core of the platform is a zero-copy semantic architecture that allows Dremio to query disparate relational databases without physically replicating data. This approach preserves data integrity and governance under data mesh principles while reducing storage overhead, minimizing ETL complexity, and ensuring that all analytical outputs reflect the freshest available information.

Data flows through a structured, multilayer pipeline designed for progressive refinement. Raw inputs first enter a staging layer for standardization before advancing to the Business Layer, where cross-source joins create a unified Citizen 360 profile. The Application Layer then transforms this curated data into dashboards, reports, and KPI-driven insights that support operational and strategic decision-making across Granicus.

The platform also incorporates advanced engineering enhancements. A custom entity-resolution layer built in Dremio reconciles complex identifier conflicts, while extensive optimization of Dremio's semantic layer—including the manual creation of raw and aggregated reflections—dramatically accelerates query performance. Governance is reinforced through a curated data dictionary, an active stewardship model focused on PII anonymization, and automated workflows orchestrated through a multi-agent LangGraph system. Granicus is further evaluating Dremio's Open MCP to extend automated metadata management across the ecosystem.

Results

The implementation of the new data platform delivered substantial improvements across operational and analytical domains. A primary achievement was the dramatic increase in query performance. Queries now run against millions of records in mere seconds, achieving a 10 to 1,000 times faster response time at the business layer. This remarkable acceleration directly enhances the speed of business intelligence and decision-making.

Furthermore, the solution has transformed data utilization by enabling true real-time analytics. The platform supports real-time reporting, turning raw data into a continuous competitive advantage, allowing the organization to react to dynamic conditions instantly. Concurrently, a significant data integration effort successfully unified disparate data sources, culminating in a comprehensive "Citizen 360" profile. This unified view provides a deeper, holistic understanding of the audience, which is crucial for targeted outreach and service delivery.

In addition to performance and integration, the new architecture significantly bolstered the organization's data governance posture. The platform provides clear and transparent data lineage capabilities. This feature ensures that the path of any data point can be tracked reliably, tracing its source from the application layer all the way back to its original system of record, such as a Microsoft SQL database. This robust tracking is foundational for ensuring data accuracy, compliance, and trust throughout the enterprise.

Finally, the initiative laid the groundwork for future advanced capabilities with the development of an Agentic AI Proof of Concept (POC). This working prototype established an agent-based system utilizing LangGraph to connect specialized agents—including those for querying, discovery, and analytics—to the Dremio Metadata Consumption Platform (MCP). This strategic development prepares the organization to seamlessly integrate and respond to sophisticated, AI-driven data requests in the near future.