



AN INFLUXDATA CASE STUDY

Managing Satellite Telemetry at Scale: Eutelsat OneWeb + InfluxDB





Eutelsat Group, formed by the 2023 merger of Eutelsat and OneWeb, is a global leader in satellite communications. With over 40 years of experience and €1.13 billion in FY2022–2023 revenue, the company operates a hybrid constellation of 600+ Low Earth Orbit (LEO) satellites alongside a robust Geostationary Orbit (GEO) fleet. The company is a pioneer in next-generation, software-driven, all-electric spacecraft and Very High Throughput Satellite (VHTS) technology.

Eutelsat OneWeb's satellites generate an immense stream of telemetry. Today, with each satellite producing over 50,000 individual telemetry values, the operations team processes more than one million data points per second across the constellation. As Dan Kroboth, Vice President of LEO Satellite Operations, notes, this “mass quantity of data” underpins mission assurance. That data enables the team to ensure satellites remain safe and operational, and that the entire system (in space and on the ground) is performing effectively.

But reaching this scale wasn't automatic. It required a deliberate shift in how telemetry was ingested, managed, and analyzed, along with significant infrastructure investment to support proactive detection, rapid diagnosis, and continuous improvement.

Moving beyond logs: Telemetry's unique demands

Initially, the team adopted Elasticsearch and Kibana for security and application log analysis. However, when they tried to extend it to telemetry, they quickly ran into several critical problems.

- **Scalability.** As the constellation grew, the platform struggled with both the volume of telemetry and the speed at which the data arrived. Storage requirements also ballooned, driving up costs and operational overhead.
- **Data transformation.** Telemetry arrived as compact machine code. The team had to build custom logic to convert raw data into real-world values.
- **Time series visualization.** The dashboards were designed for logs, not time series. Comparing trends or correlating spacecraft and ground behavior was slow and clunky.
- **Duplicate observations.** The same data point could arrive multiple times via different channels. The team had to write and maintain custom deduplication code before any meaningful analysis could begin.

- **Tagging and cardinality.** Each point required additional metadata—satellite ID, operational phase, subsystem context—to be queryable. This exploded the number of unique series and pushed past the limits of the enterprise search platform.

Kroboth summed it up nicely: platforms like Elasticsearch are “great for JSON and logs, but aren't suitable for what spacecraft produce, which is time series.”

"With InfluxDB, we can measure everything. You never know what you'll need to ask tomorrow."

Dan Kroboth

Vice President of LEO Satellite Operations



Guiding principles for a spacecraft telemetry platform



To guide their evaluation and system design, OneWeb's operations team articulated a clear set of requirements and guiding principles to shape both their platform selection and system design.

The platform needed to scale and remain highly available under continuous load. Each new satellite added tens of thousands of telemetry signals, driving up ingest, storage, and query demands. With over a million data points per second flowing in, the system couldn't simply scale; it had to remain responsive, reliable, and fault-tolerant without architectural rewrites or tuning.

Scaling shouldn't mean runaway costs. Most traditional telemetry stacks are designed for small GEO fleets; they often don't translate economically to LEO constellations. The team needed efficient compression, affordable long-term storage, and a pricing model that made sense when multiplied across hundreds of satellites and millions of series.

The platform needed to be flexible. Incoming telemetry lacked uniformity. The series had different schemas and arrived at different resolutions, with some reporting every tenth of a second, others every hundredth, and so on. This variability is one reason why off-the-shelf relational databases and enterprise

search platforms, like Elasticsearch, have historically struggled with telemetry. That said, the bigger challenge was cardinality: the team expected to manage more than 15 million high-cardinality time series—well beyond what even some time series platforms can handle at scale.

The platform needed to be automated. The team wanted automation baked in from the start—alerting, ingestion, and monitoring had to run with minimal human input. Looking forward, they want a platform that can handle batch ingestion, deduplication, correlation, and diagnostics, all driven by time series queries and live dashboards.

The platform needed to be open. The team wanted to avoid lock-in. They needed a system built on open protocols and formats—one that could integrate easily with internal tools, commercial software, and future data infrastructure.



Embracing InfluxDB's open approach


As the team evaluated options, only InfluxDB met all of OneWeb's requirements for scalability, cost-effectiveness, flexibility, automation readiness, and openness. As Kroboth put it, Influx is "really purpose-built for our data."

InfluxDB's open source roots (and Telegraf's plugin ecosystem) fit what Kroboth describes as the team's "bias for open things." Both are available as open source distributions, the team could prototype freely on real telemetry, iterate quickly, and validate their architecture before committing. This hands-on approach to proof is one reason why, today, Telegraf instruments the ground segment end-to-end and InfluxDB is at the core of the telemetry platform.

Equally important was the InfluxData team's openness to collaboration. "The InfluxData folks were very, very helpful...helping us discover the best solutions we could apply from the toolset," Kroboth said. In particular, they helped the OneWeb team understand how the cadence and high cardinality of their set of time series affected key design choices, including data partitioning, tag strategy, metric structure, query performance, and storage efficiency.

Today, with InfluxData's guidance and InfluxDB's native support for high cardinality, Eutelsat OneWeb manages more than 15 million unique series with ease, demonstrating what's possible when an open platform meets an open partnership.

Deploying a new architecture for mission assurance



Eutelsat OneWeb's new telemetry stack is anchored by InfluxDB and Grafana, with Telegraf agents deployed across the ground segment to pull data from antennas to modems. InfluxDB serves as the centralized time series engine, ingesting more than 1 million points per second and supporting more than 15 million unique series. Grafana sits on top, enabling real-time dashboards, alerting, and cross-system correlation.



What OneWeb Can Do Now

- ✓ **Real-time dashboards in Grafana:** Monitor satellite health and track subsystem behavior in real-time.
- ✓ **Unified telemetry exploration:** Analyze spacecraft and ground-segment data side by side.
- ✓ **Proactive alerting:** InfluxDB queries trigger Grafana alerts and Microsoft Teams notifications.
- ✓ **Time series correlation:** Replay events across satellites and ground systems for root-cause analysis.

What OneWeb is Building Next

- ✓ **Machine learning pipelines:** Use time series data to train AI models that detect “previously unseen failures.”
- ✓ **Long-term retention:** Store full-fidelity telemetry using InfluxDB 3 and Apache Iceberg.
- ✓ **Fleetwide insight:** Track trends by satellite batch and surface subtle degradation patterns over years of data.

The impact goes beyond infrastructure. By centralizing telemetry and automating analysis, the team gained time to focus on what matters: debugging spacecraft, improving service quality, and evolving their next-generation platform. With years of historical data now instantly explorable, they can group satellites by batch, track long-term trends, and reveal subtle anomalies that would’ve been “hard to do any other way.” InfluxDB didn’t just replace their old system, it changed how they build and operate.

As the constellation grows, the telemetry stack continues to unlock new capabilities. The team plans to use InfluxDB to power machine learning and AI pipelines for proactive anomaly detection, training models on time series and known telemetry patterns to surface “previously unseen failures.”

That blend of flexibility, performance, scalability, and high availability is why Eutelsat OneWeb built on InfluxDB. As Kroboth describes, it’s been a “catalyst for transformation,” expanding what the team thinks is possible. With InfluxDB, the team can “measure everything”—because, as Kroboth puts it, “you never know what you’ll need to ask tomorrow.”

50K individual telemetry values
across 600+ LEO satellites

1M points ingested
per second

15M unique series
supported

Try InfluxDB

Contact us for a personalized demo at influxdata.com/contact-sales