



Finanziato
dall'Unione europea
NextGenerationEU



Ministero
dell'Università
e della Ricerca



Italiadomani
PIANO NAZIONALE
DI RIPRESA E RESILIENZA



Enabling performance insights: generating telemetry from software services developed at INFN-CNAF

Jacopo Gasparetto (INFN-CNAF)

Workshop sul calcolo nell'INFN | La Biodola 26-30 Maggio 2025



Outline

- Introduction to OpenTelemetry
- Telemetry data: Traces, Metrics and Logs
- Instrumentation
- Collecting telemetry data
- Case study: Otello
- Examples at INFN-CNAF
- Conclusions



Introduction to telemetry: OpenTelemetry

From the [official documentation](#)

OpenTelemetry is:

- *An **observability framework** and toolkit designed **to create** and manage telemetry data such as **traces, metrics, and logs**.*
- ***Not an observability backend** like Jaeger, Prometheus, or other commercial vendors.*
- *...*

OpenTelemetry is a collection of components that provide an SDK (Software Development Kit) to enrich an existing code base to *emit* telemetry data.

OpenTelemetry **does not** store any data *per se* and some kind of database(s) is(are) needed

Telemetry data: Traces, Metrics and Logs

Telemetry data is generally grouped in three main entities, or *signals*: [traces](#), [metrics](#) and [logs](#).

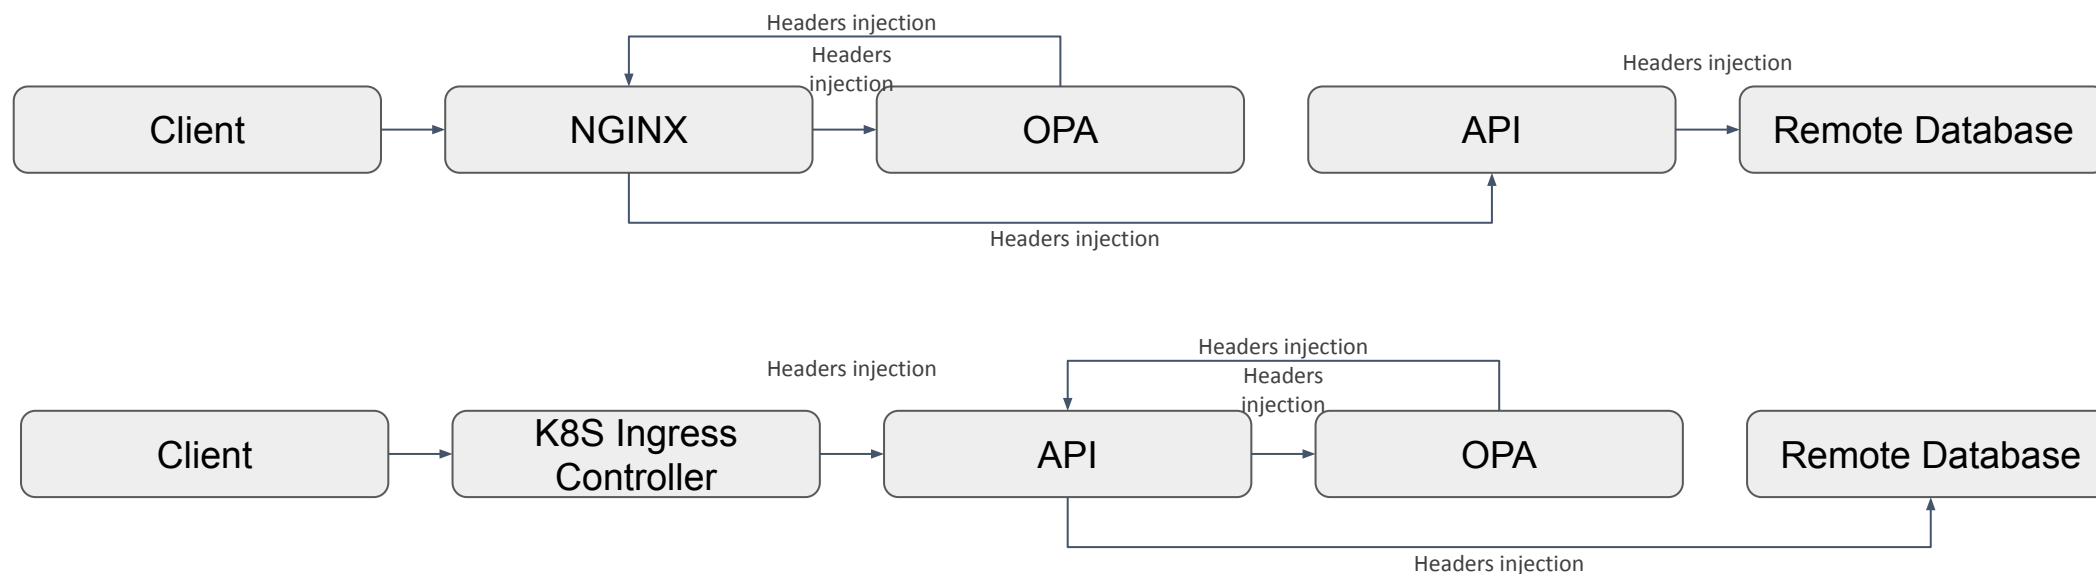
- **traces**: represent **the path of a request through an application**. They are made up by [spans](#) that represent the single units of work or operation. Each span has **start and end timestamps**, a **name**, a **parent id** (if child of another span) and **attributes**. A typical trace/span information answers “*how much time this function took to execute?*” and “*what is the runtime call stack at this endpoint?*”
- **metrics**: represent the **measurement emitted by a meter**. Meters can typically be monotonic/non-monotonic counters, histograms and gauges. In the physical world “3 kWh at 2025-02-10 12:00” is the metric produced by the electricity meter (monotonic) of our house. An example could be: “how many requests we received so far at each endpoint?”
- **logs**: time-stamped text record, either structured or unstructured, with optional metadata. We do not deal with telemetry logs, as they could easily be written with a simple access/error logger

OpenTelemetry: Instrumentation

- The process of *augmenting* a codebase to emit telemetry data is called **instrumentation**. This means that we are adding code to our existing code.
- OpenTelemetry has SDKs for basically every language on the market.
- For some high level languages, such as JavaScript/TypeScript, OpenTelemetry offers the so called [Zero-code Instrumentation](#) feature that enables telemetry with just a couple of lines of configuration.
- For other languages such as C++, there is no automatic instrumentation and the codebase must be manually instrumented at each function of interest and it is up to the developer to choose what they want to monitor.
- A combined approach of manual and automatic instrumentation is always possible

Propagation

Context propagation is performed by injecting headers by a service in the incoming HTTP request forwarded to the next service



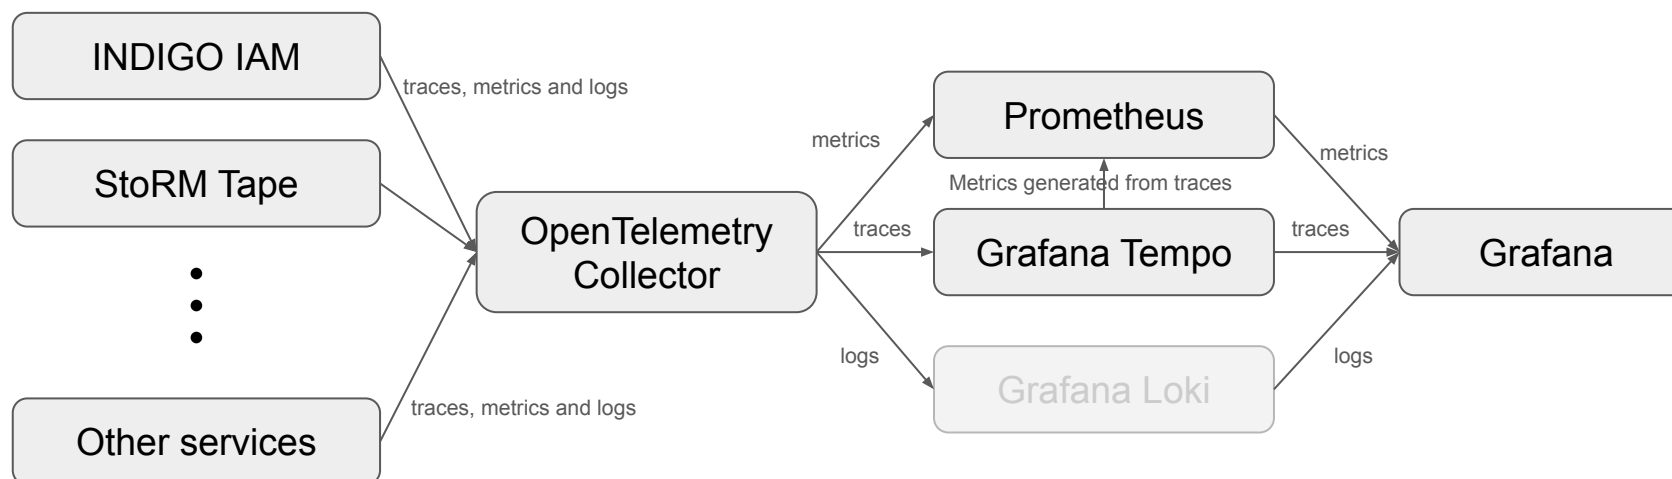
Response paths not shown in figure

Collect telemetry data

- An application can be composed of several components (micro-services), each sending telemetry data from a different endpoint but still logically representing the same application. For example, a trace with its root span can be opened by a NGINX reverse proxy that proxies the request to an API. It is important to reconstruct those traces to understand the complete path of the request.
- An application can have manyfold replicas, each sending telemetry data.
- Due to the nature of the data, different signals (traces, metrics and logs) must be stored in different kind of databases. They can also be sent through different protocols (HTTP, gRPC, kafka, ect.).
- OpenTelemetry offers a service called [OpenTelemetry Collector](#) which receives, processes and then exports telemetry data to the appropriate backends.

Otello: collecting telemetry data case study

- For our case study, we choose [Prometheus](#) as metrics backend and [Grafana Tempo](#) as traces backend
- [Grafana](#) is then used as web application to visualize KPIs, stats and plots
- A particular useful feature of Grafana Tempo is its built-in metrics generator which produces new metrics from the traces, such as the total count of requests and the latency histogram
- Otello is the prototype stack of services developed at INFN-CNAF used to collect telemetry data produced by software developed by the Software Development team





Otello: collecting telemetry data case study

Currently it is implemented as a docker compose composed by the following services:

- NGINX (reverse proxy)
- OpenTelemetry Collector (HTTP + gRPC)
- Grafana Tempo (traces database)
- Prometheus (metrics database)
- Grafana (data visualization) with custom dashboards

Website: <https://otello.cloud.cnaf.infn.it> (available only inside CNAF network)

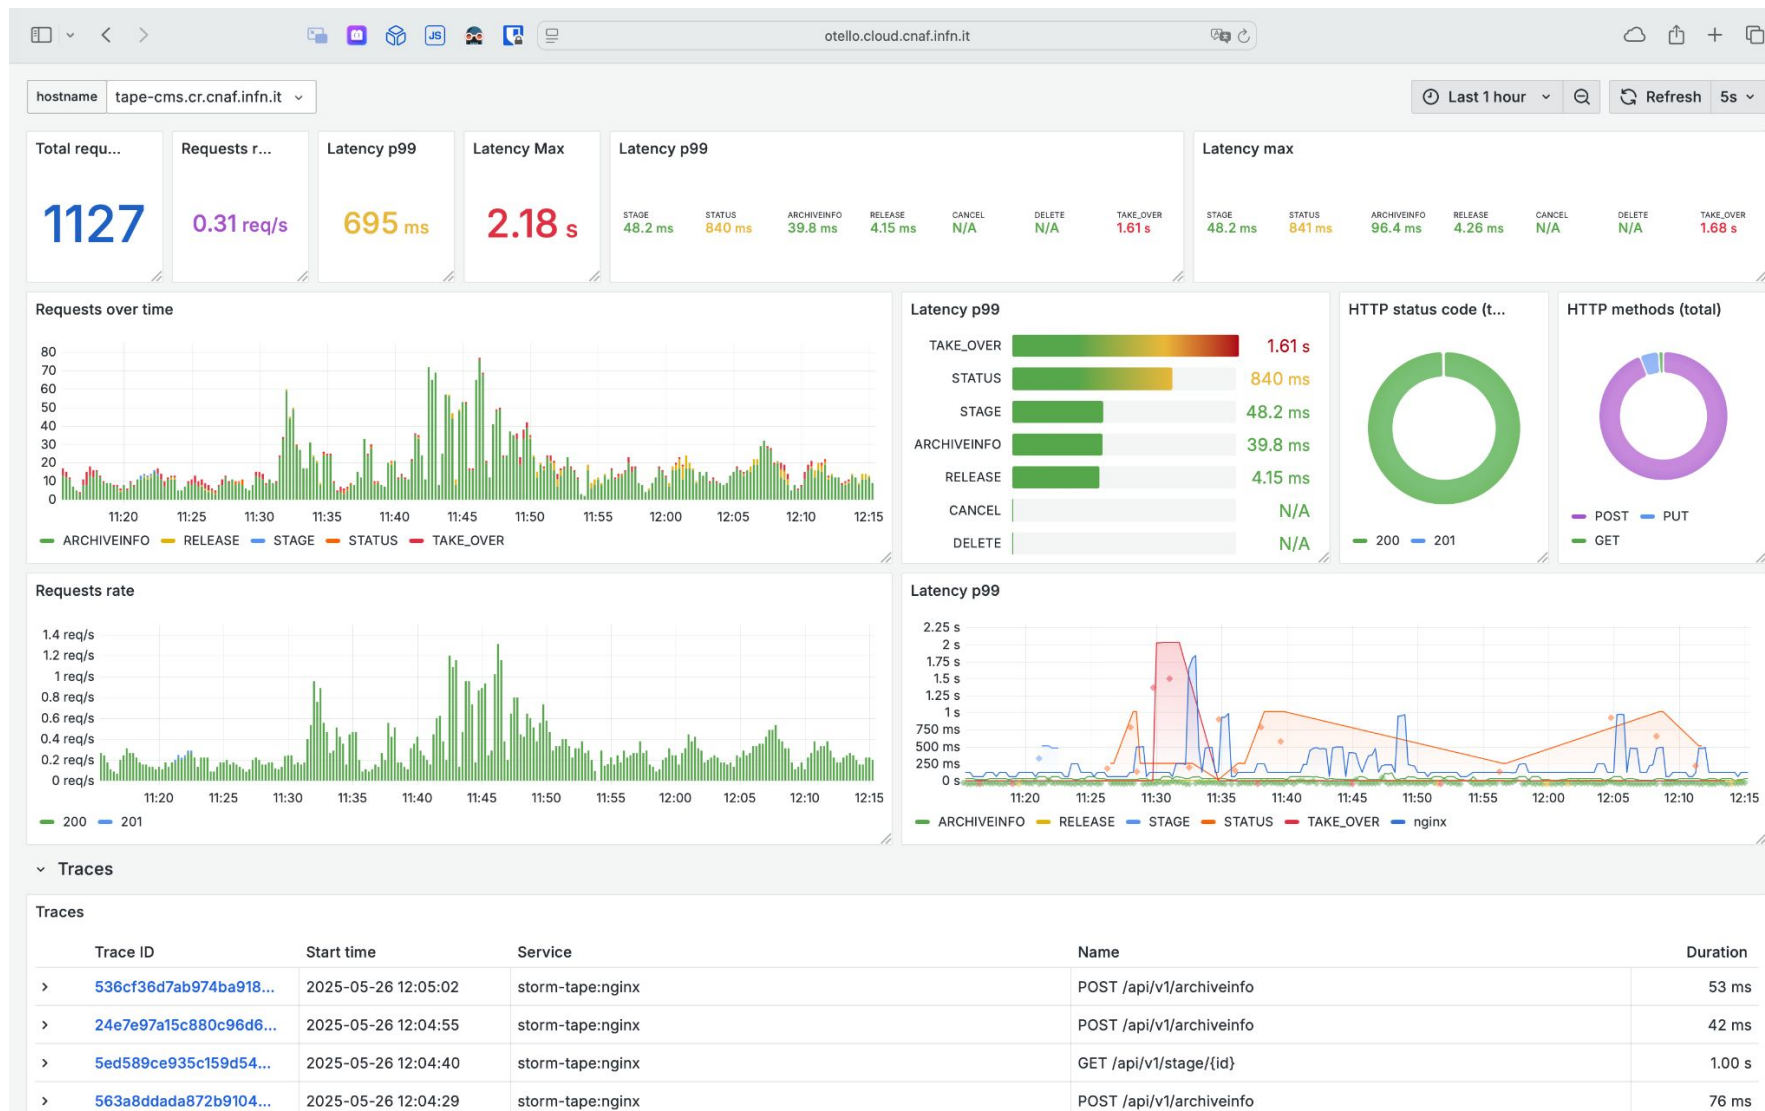
Repository: <https://baltig.infn.it/cnafsd/opentelemetry>

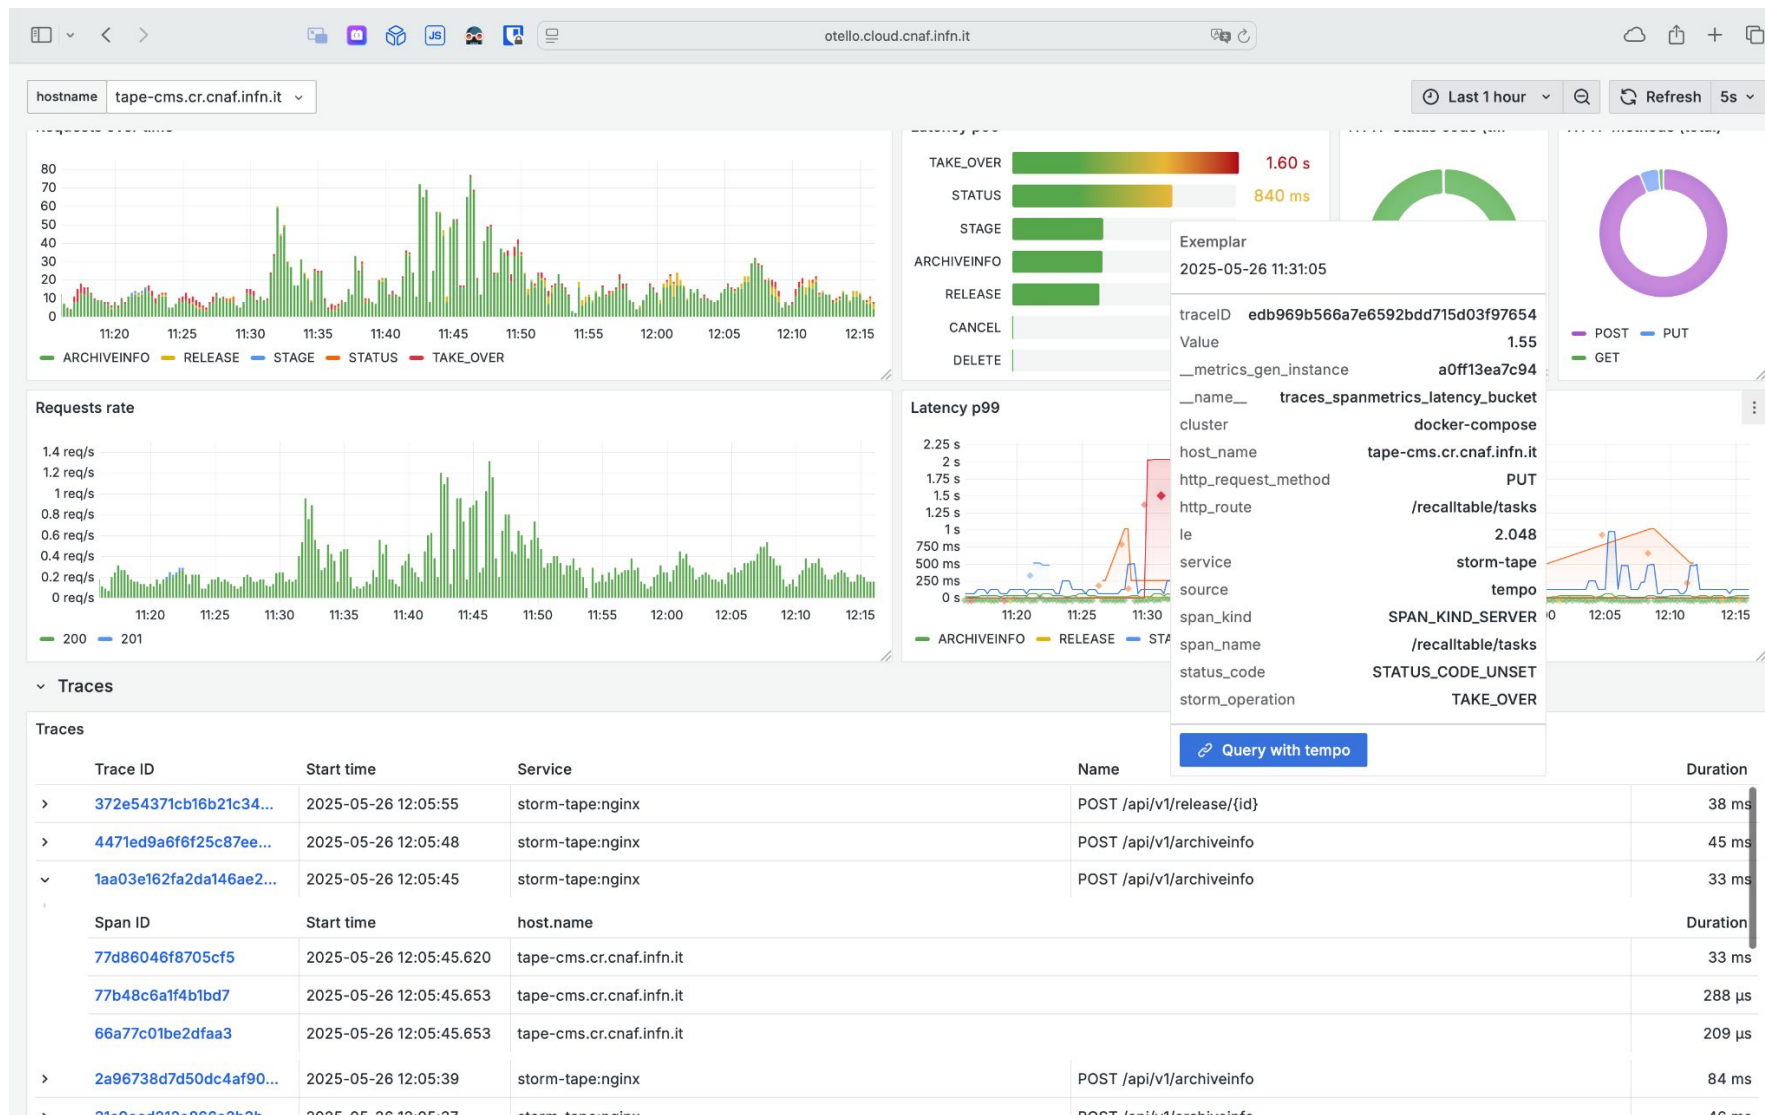


Instrumented software at INFN-CNAF

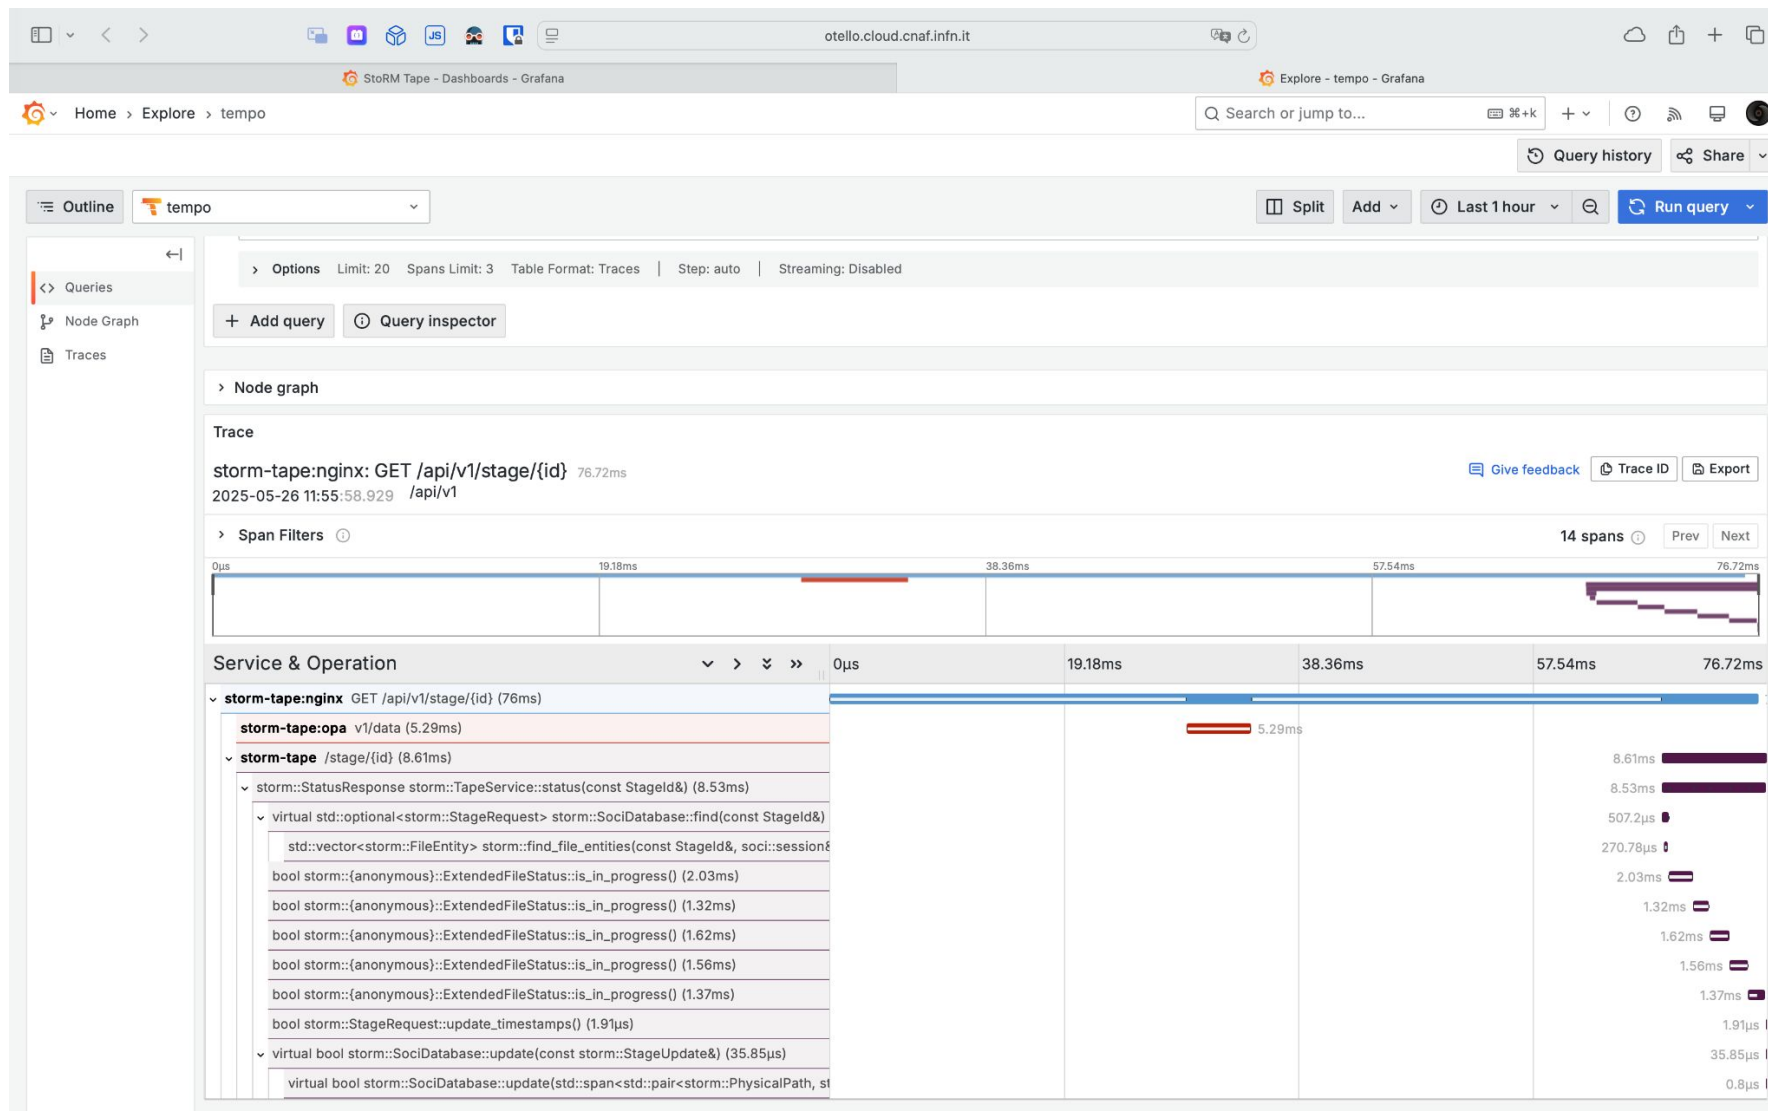
- [StoRM Tape](#) (C++)
- [StoRM WebDAV](#) (Java, testing)
- [INDIGO IAM](#) (Java, testing)
- [New IAM Dashboard](#) (Javascript)
- [Data Cloud S3 webapp](#) (Javascript)

StoRM Tape





StoRM Tape

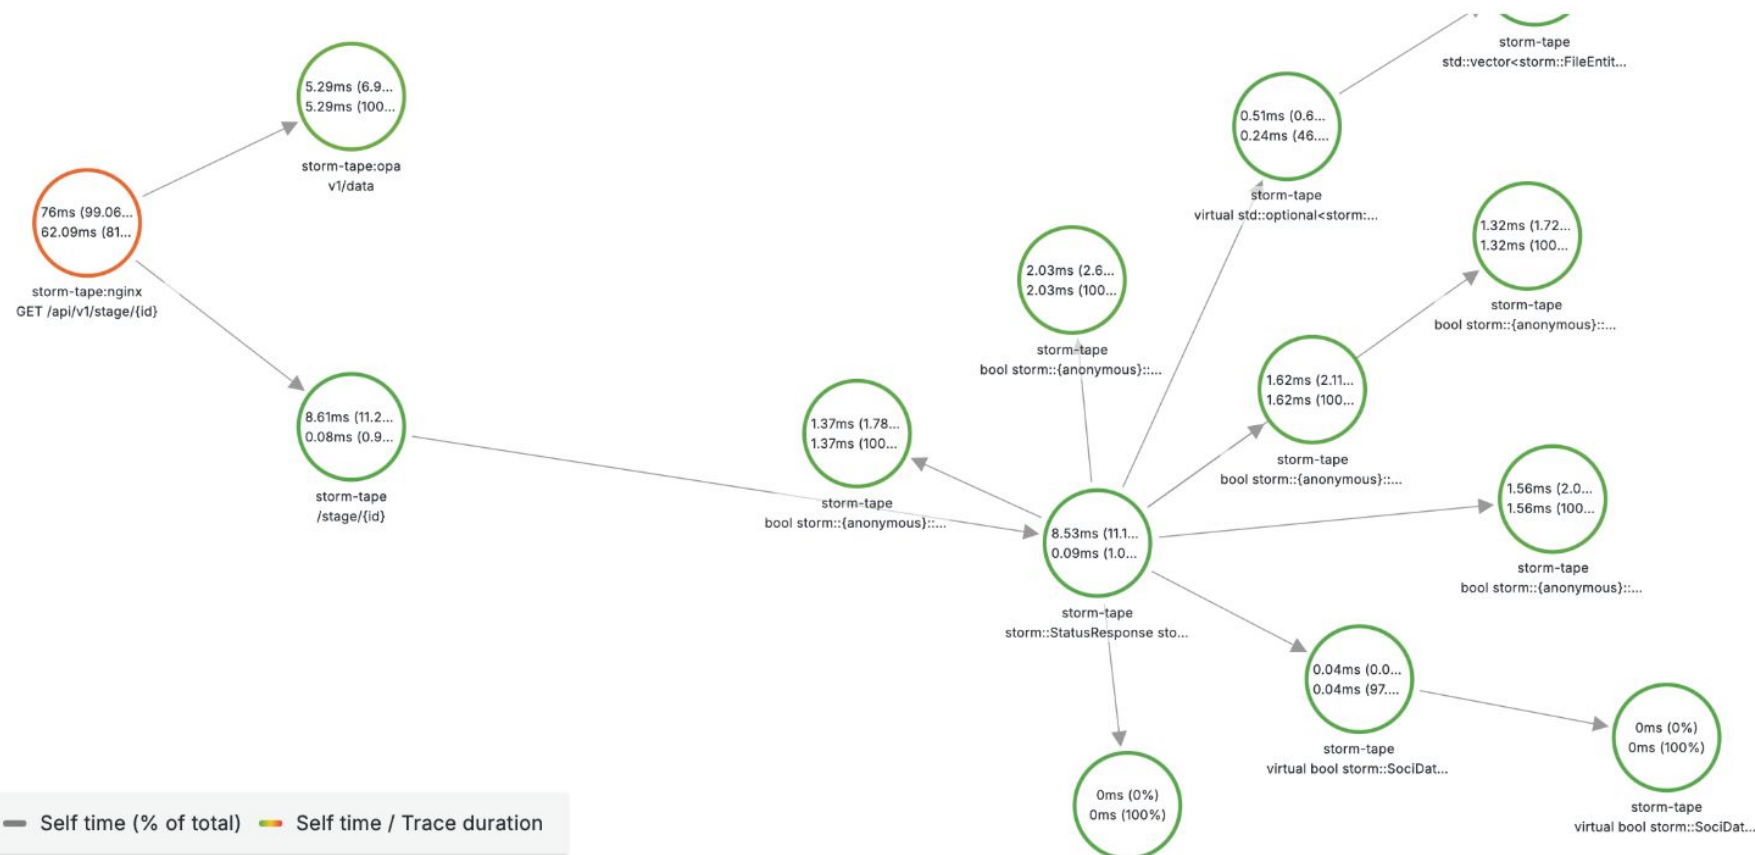


StoRM Tape

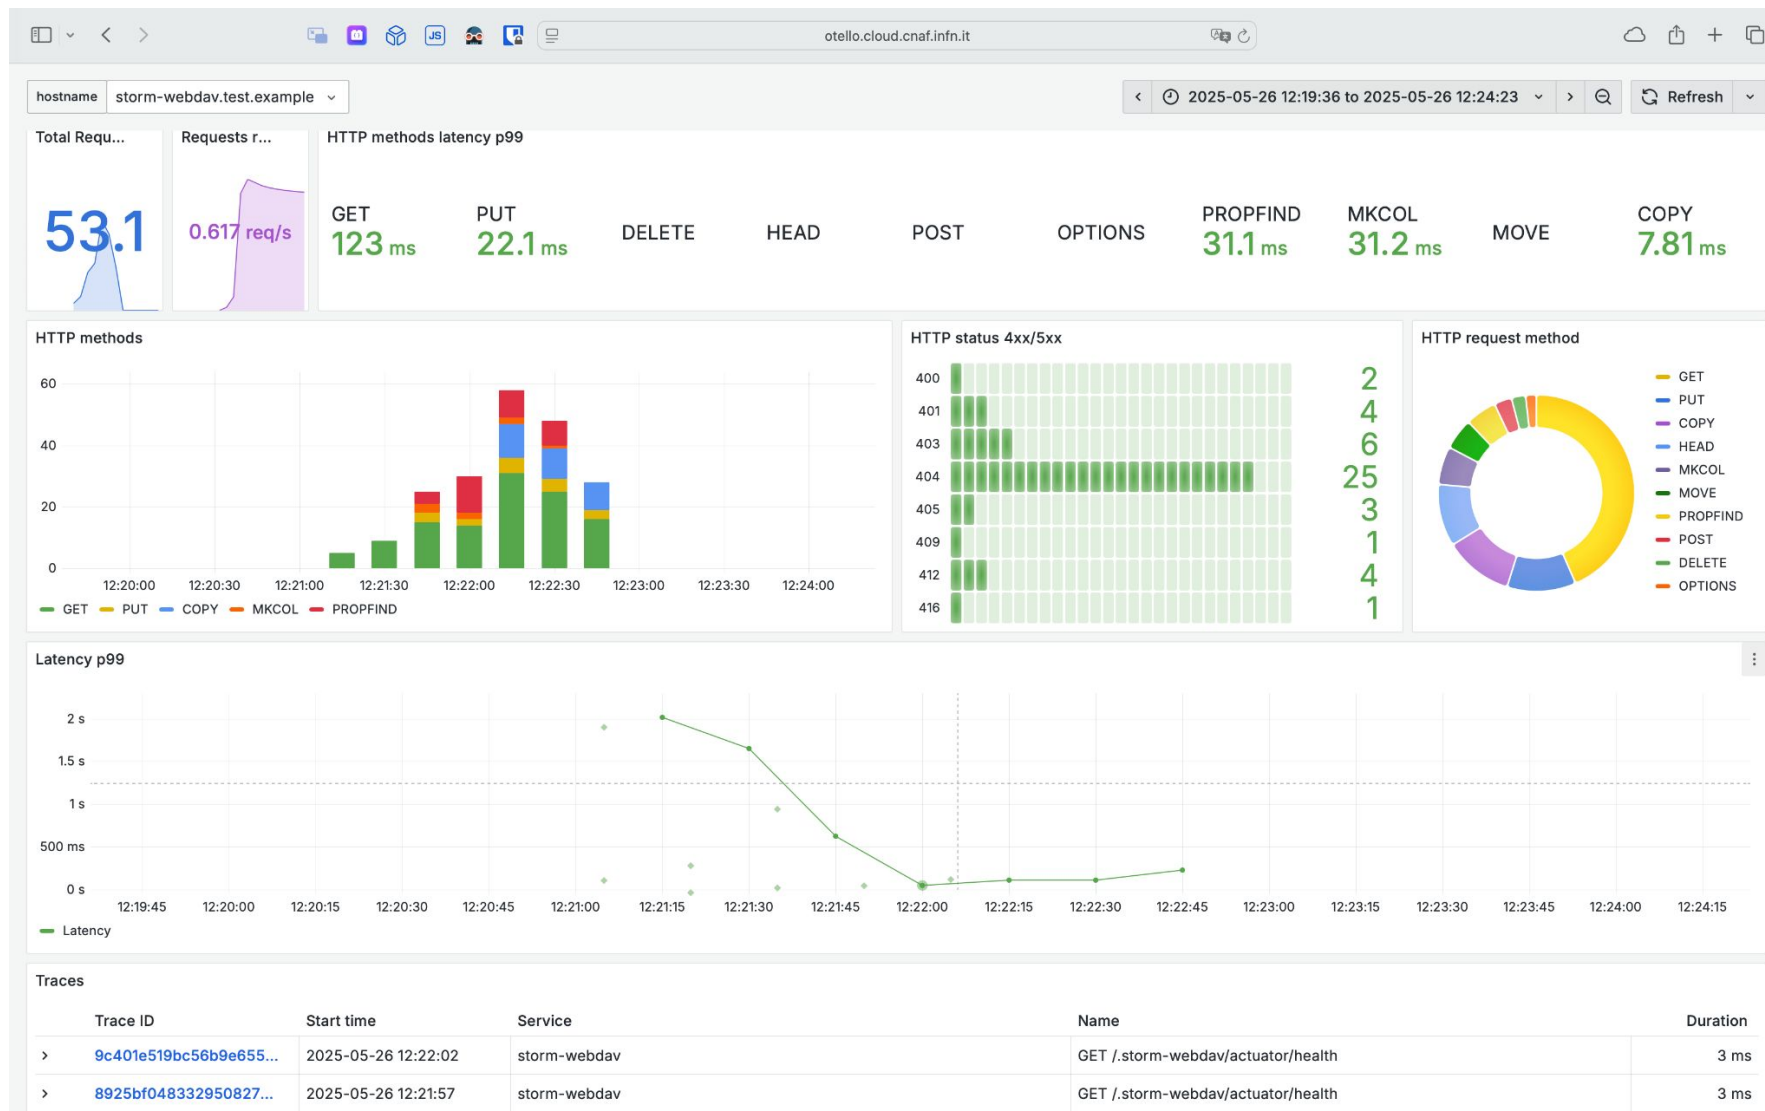


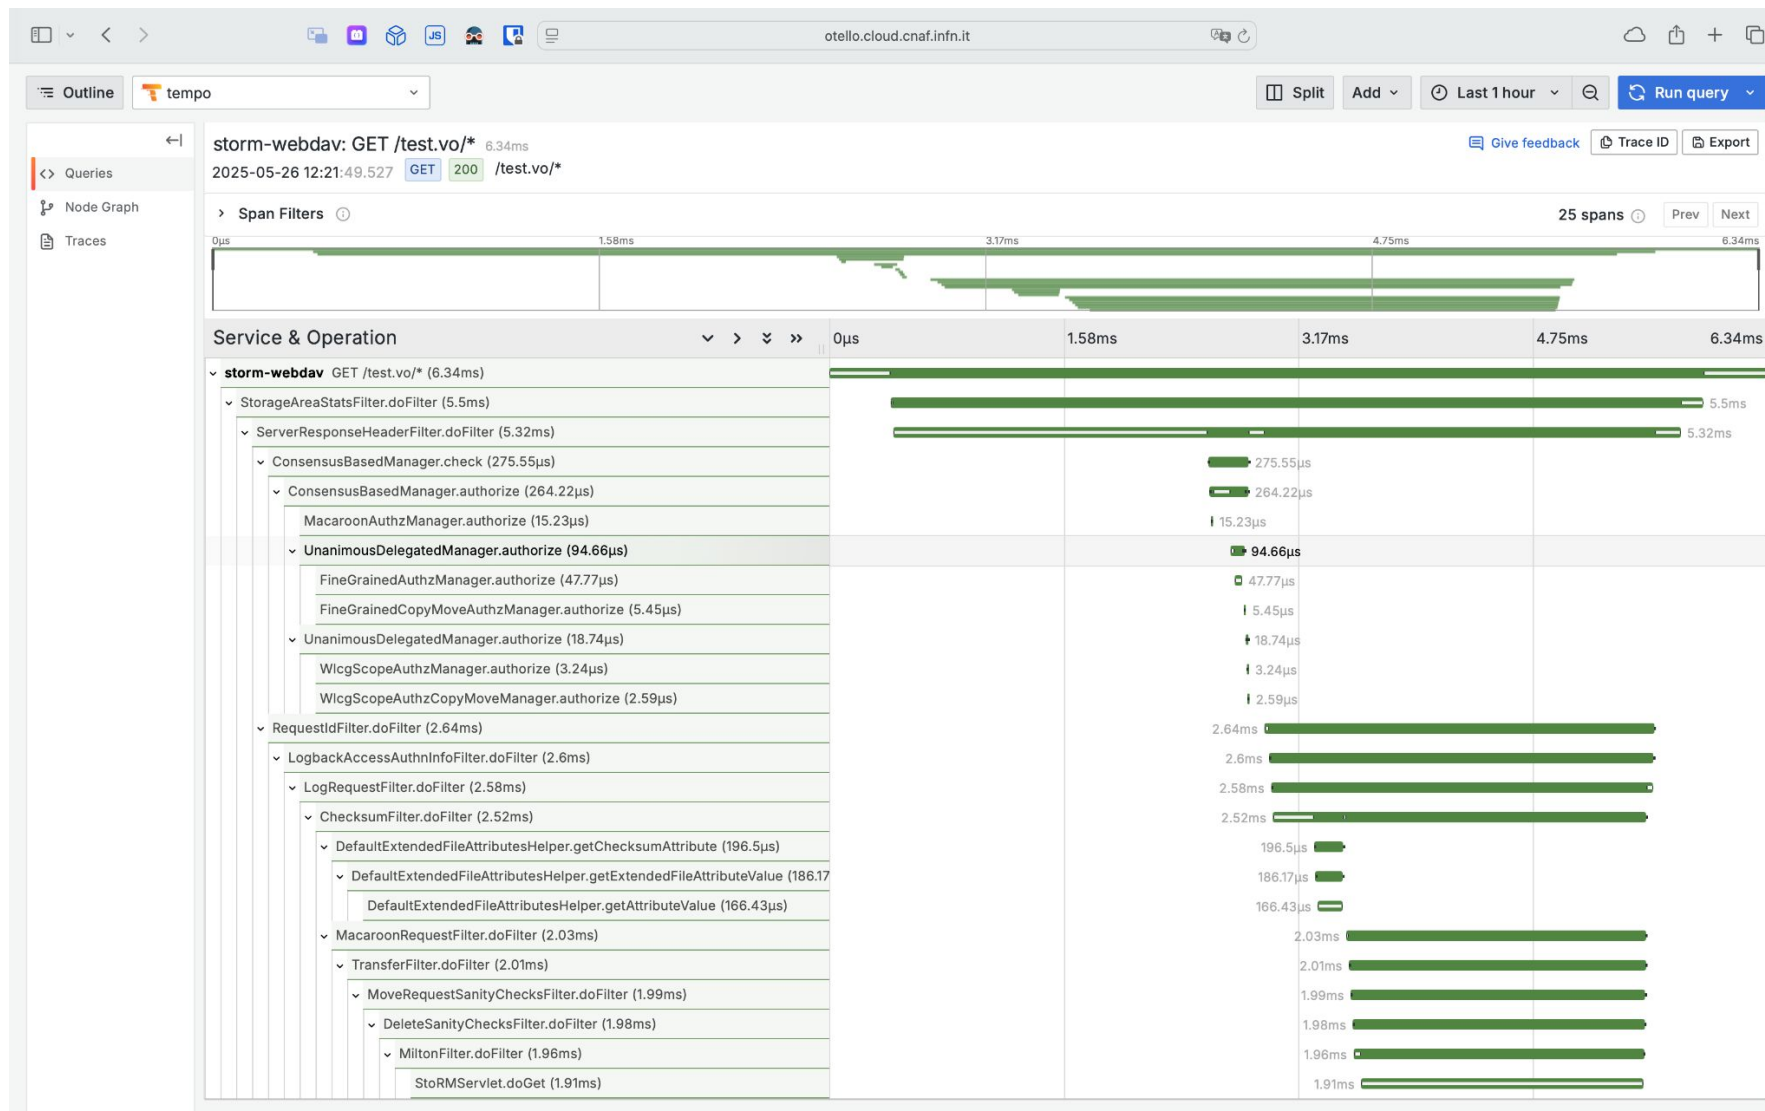
StoRM Tap

Node graph

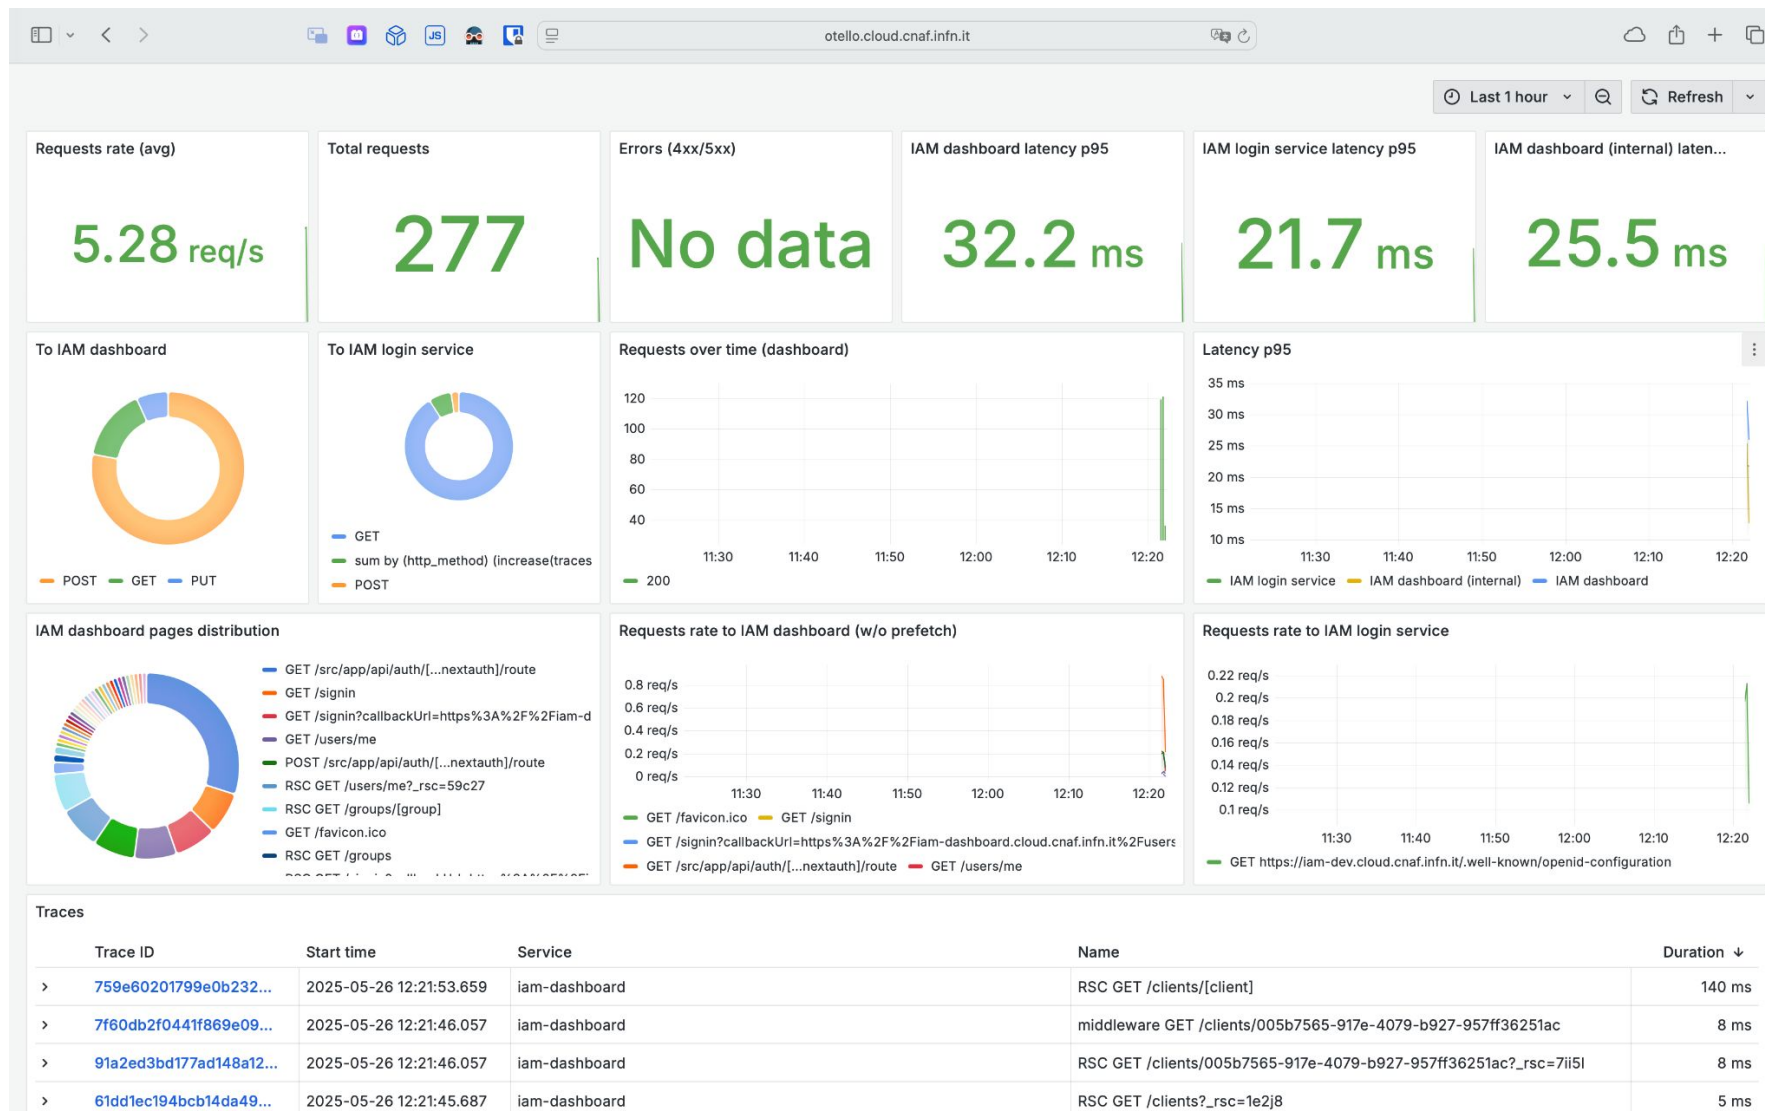


StoRM WebDAV

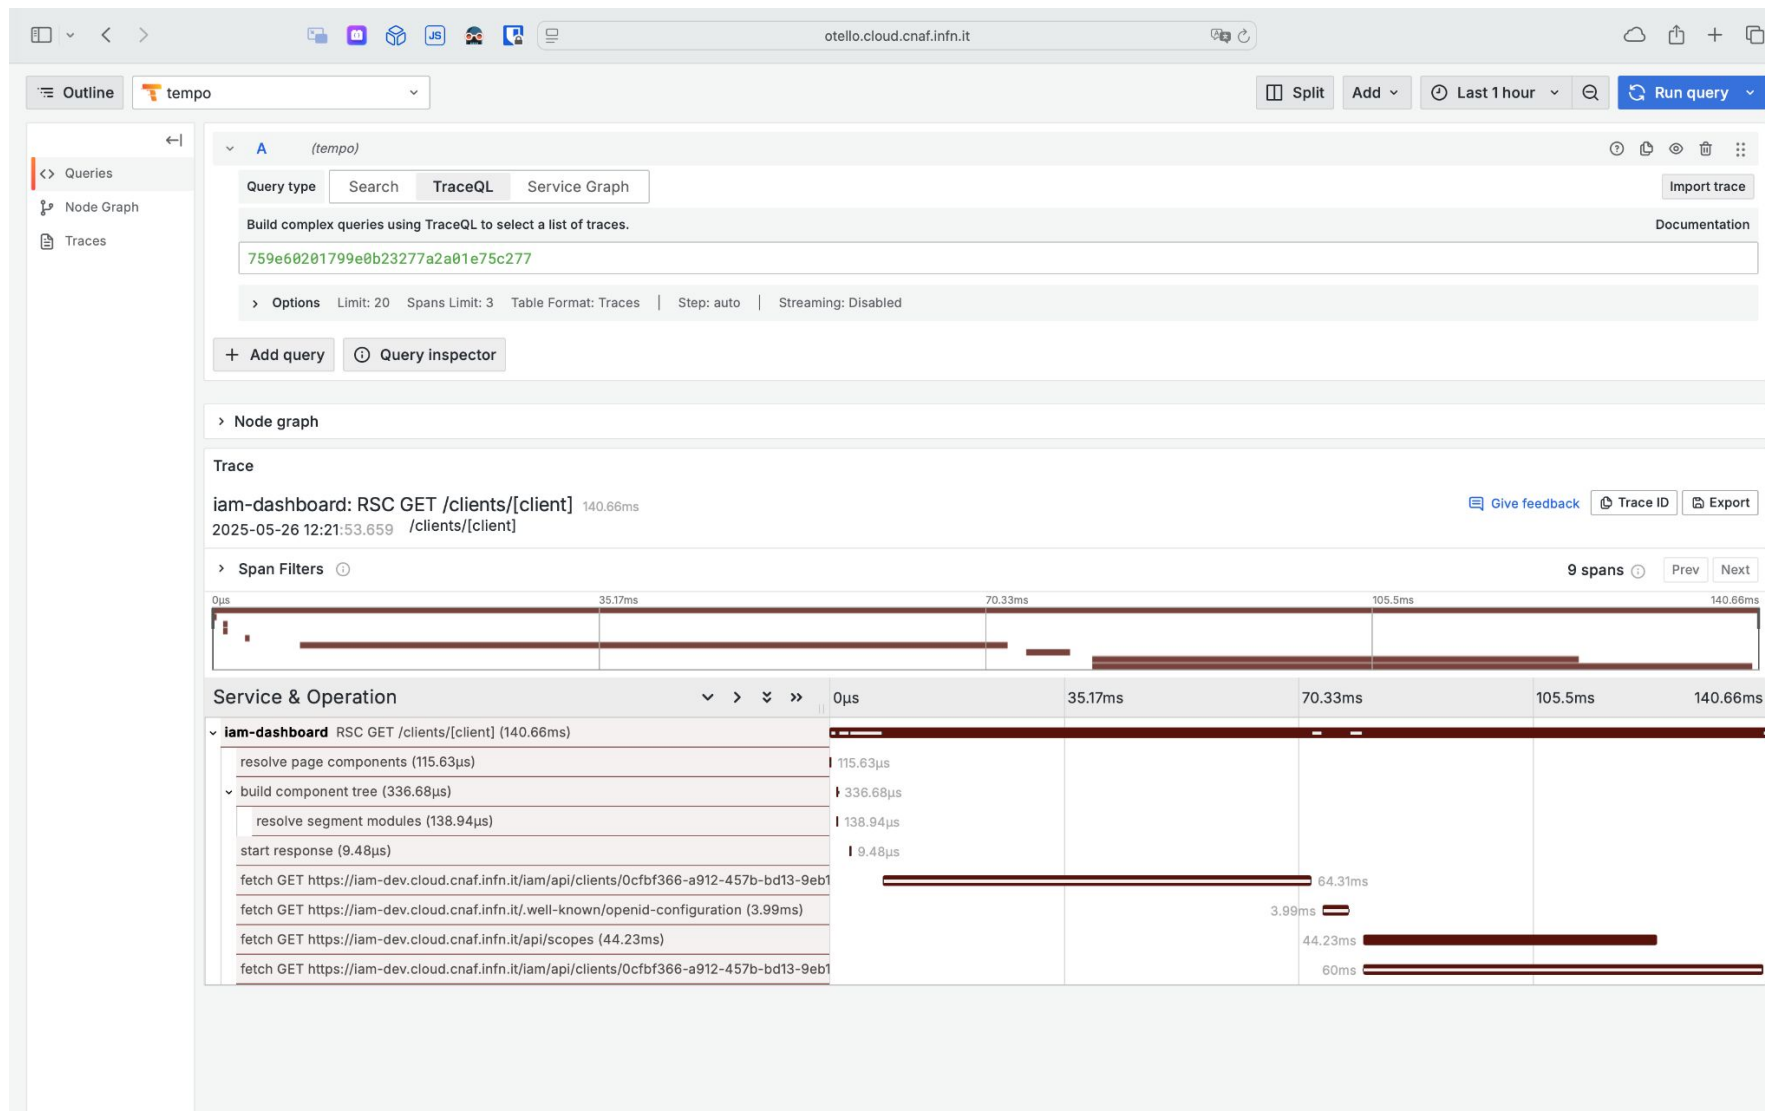




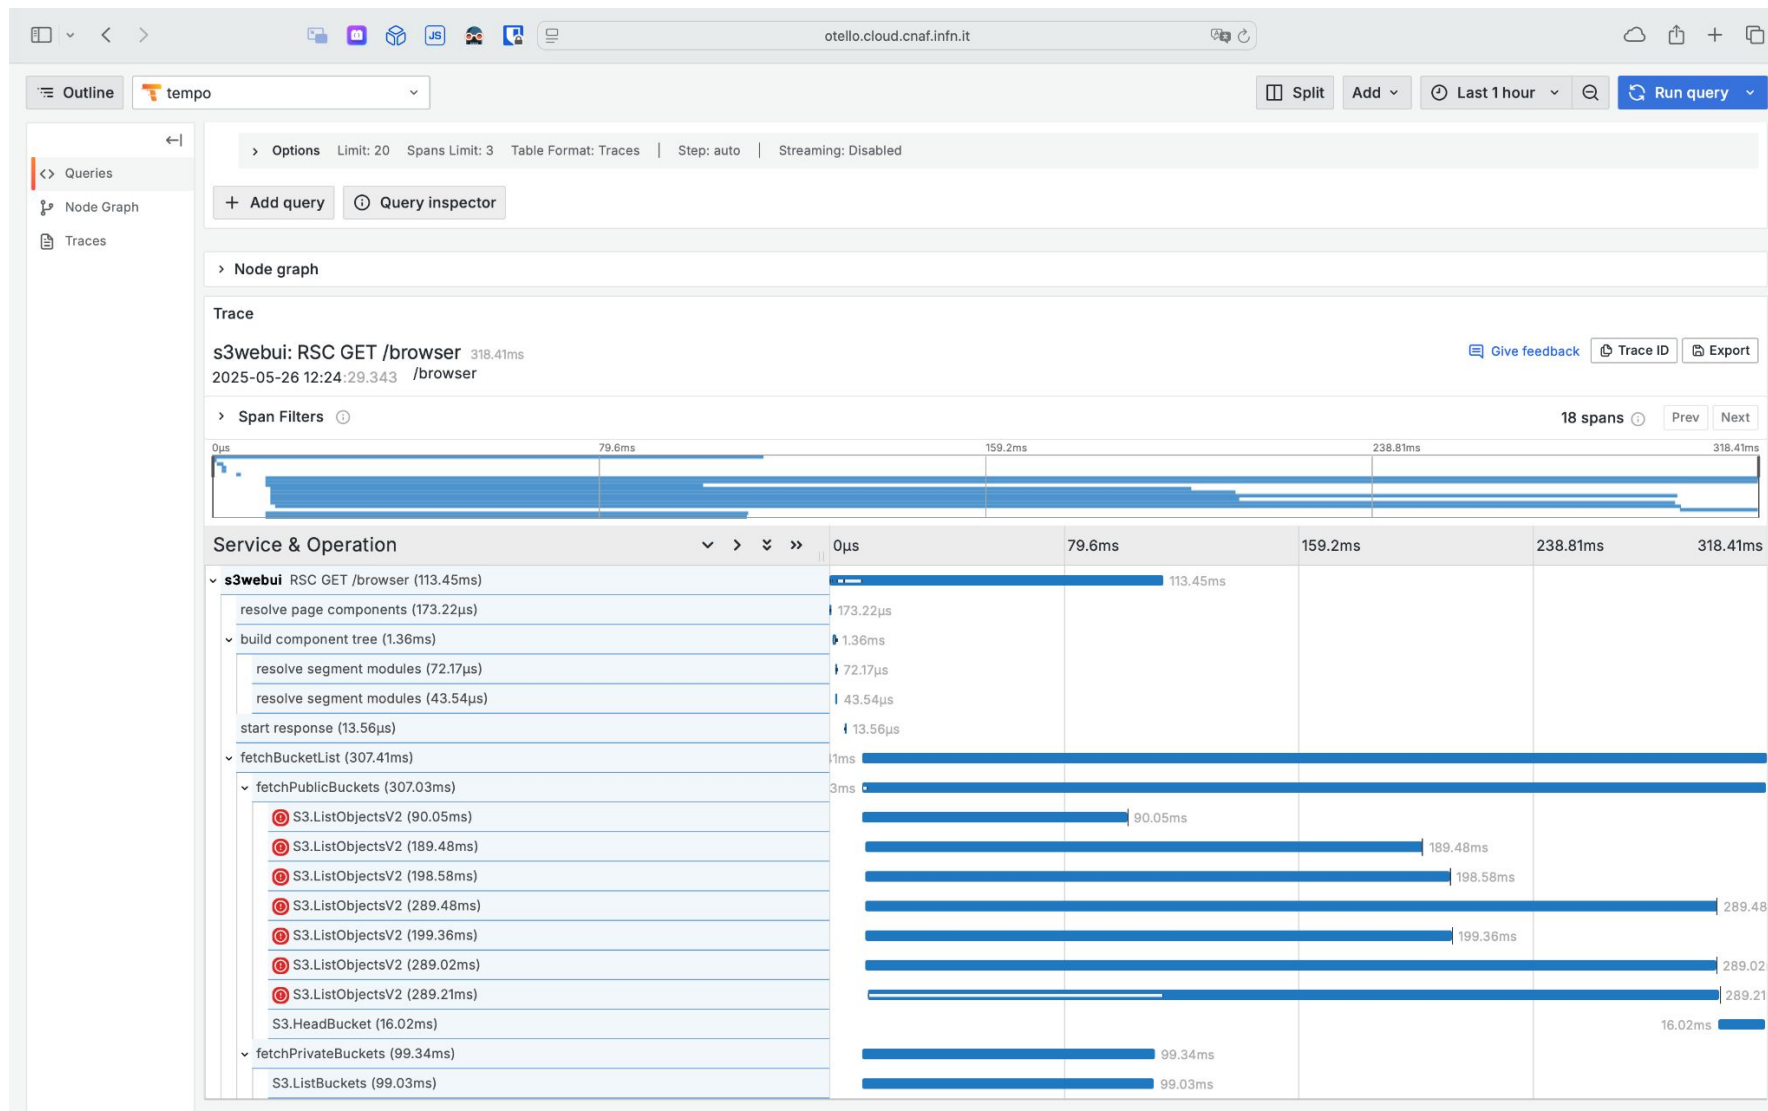
StoRM WebDAV



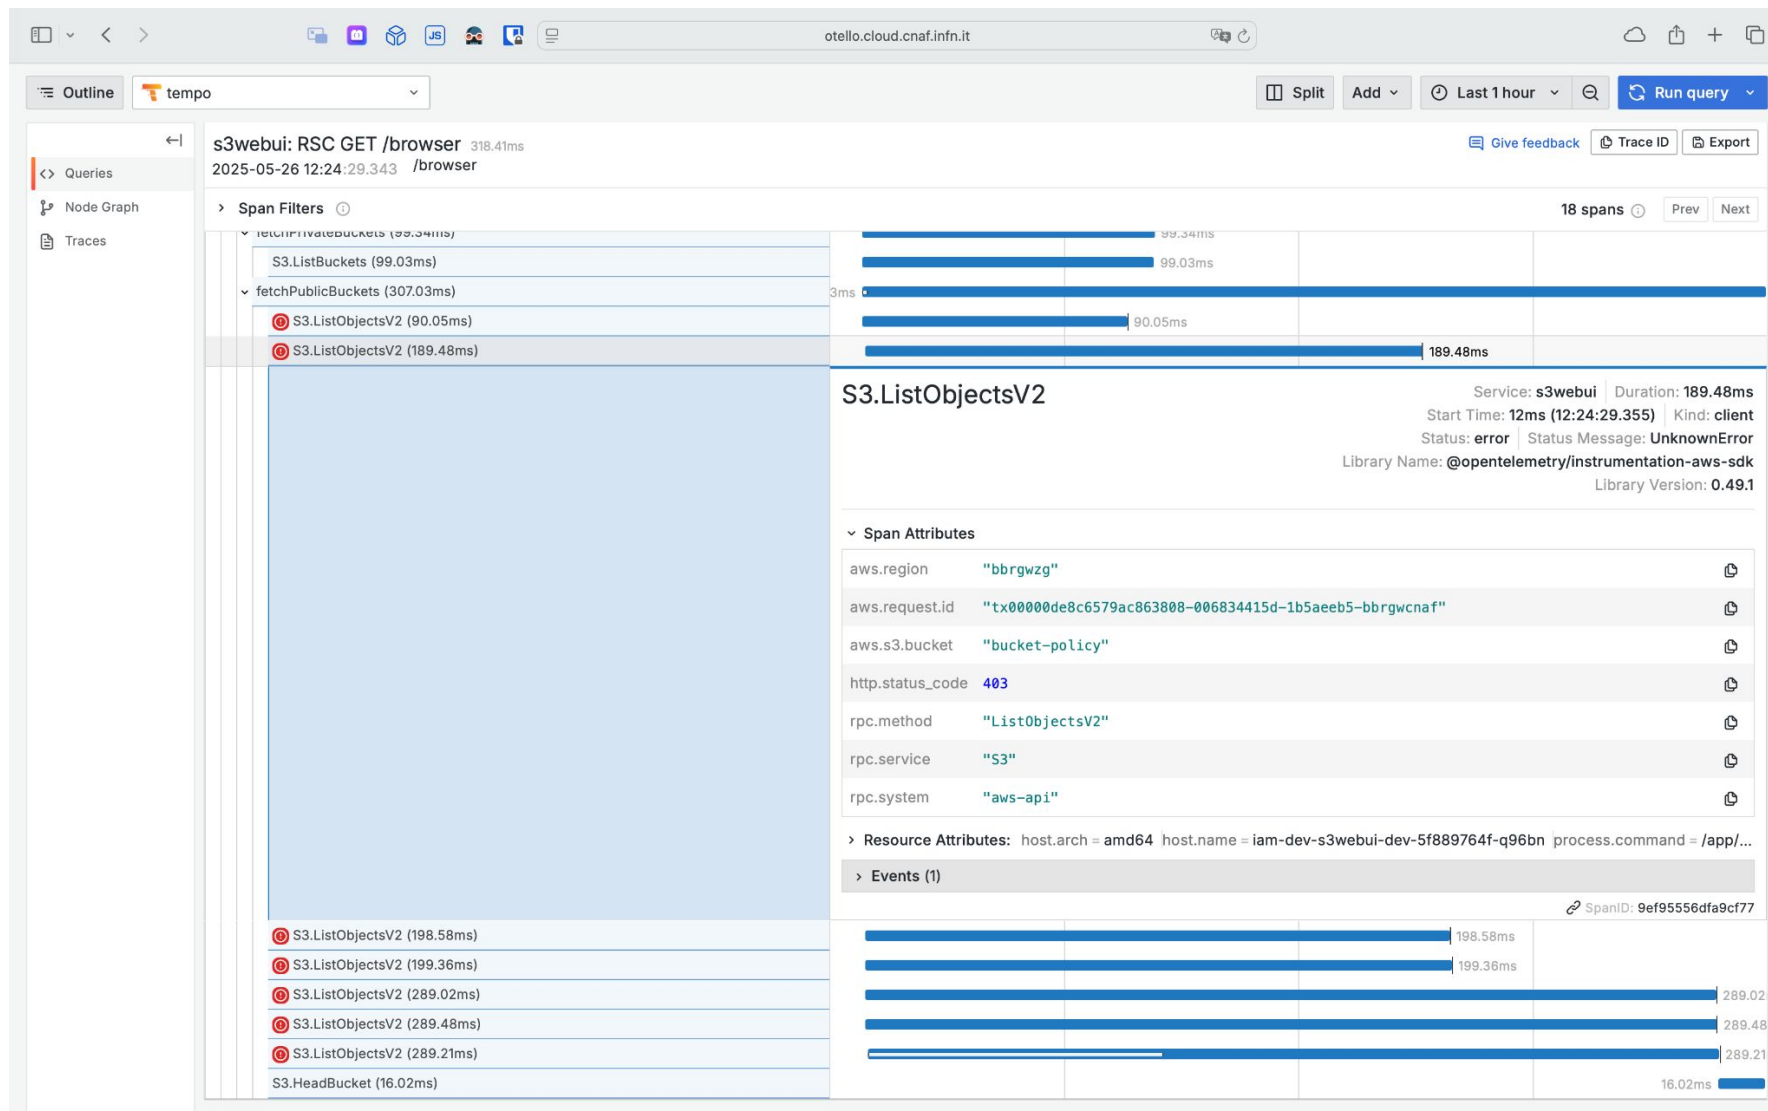
IAM Dashboard



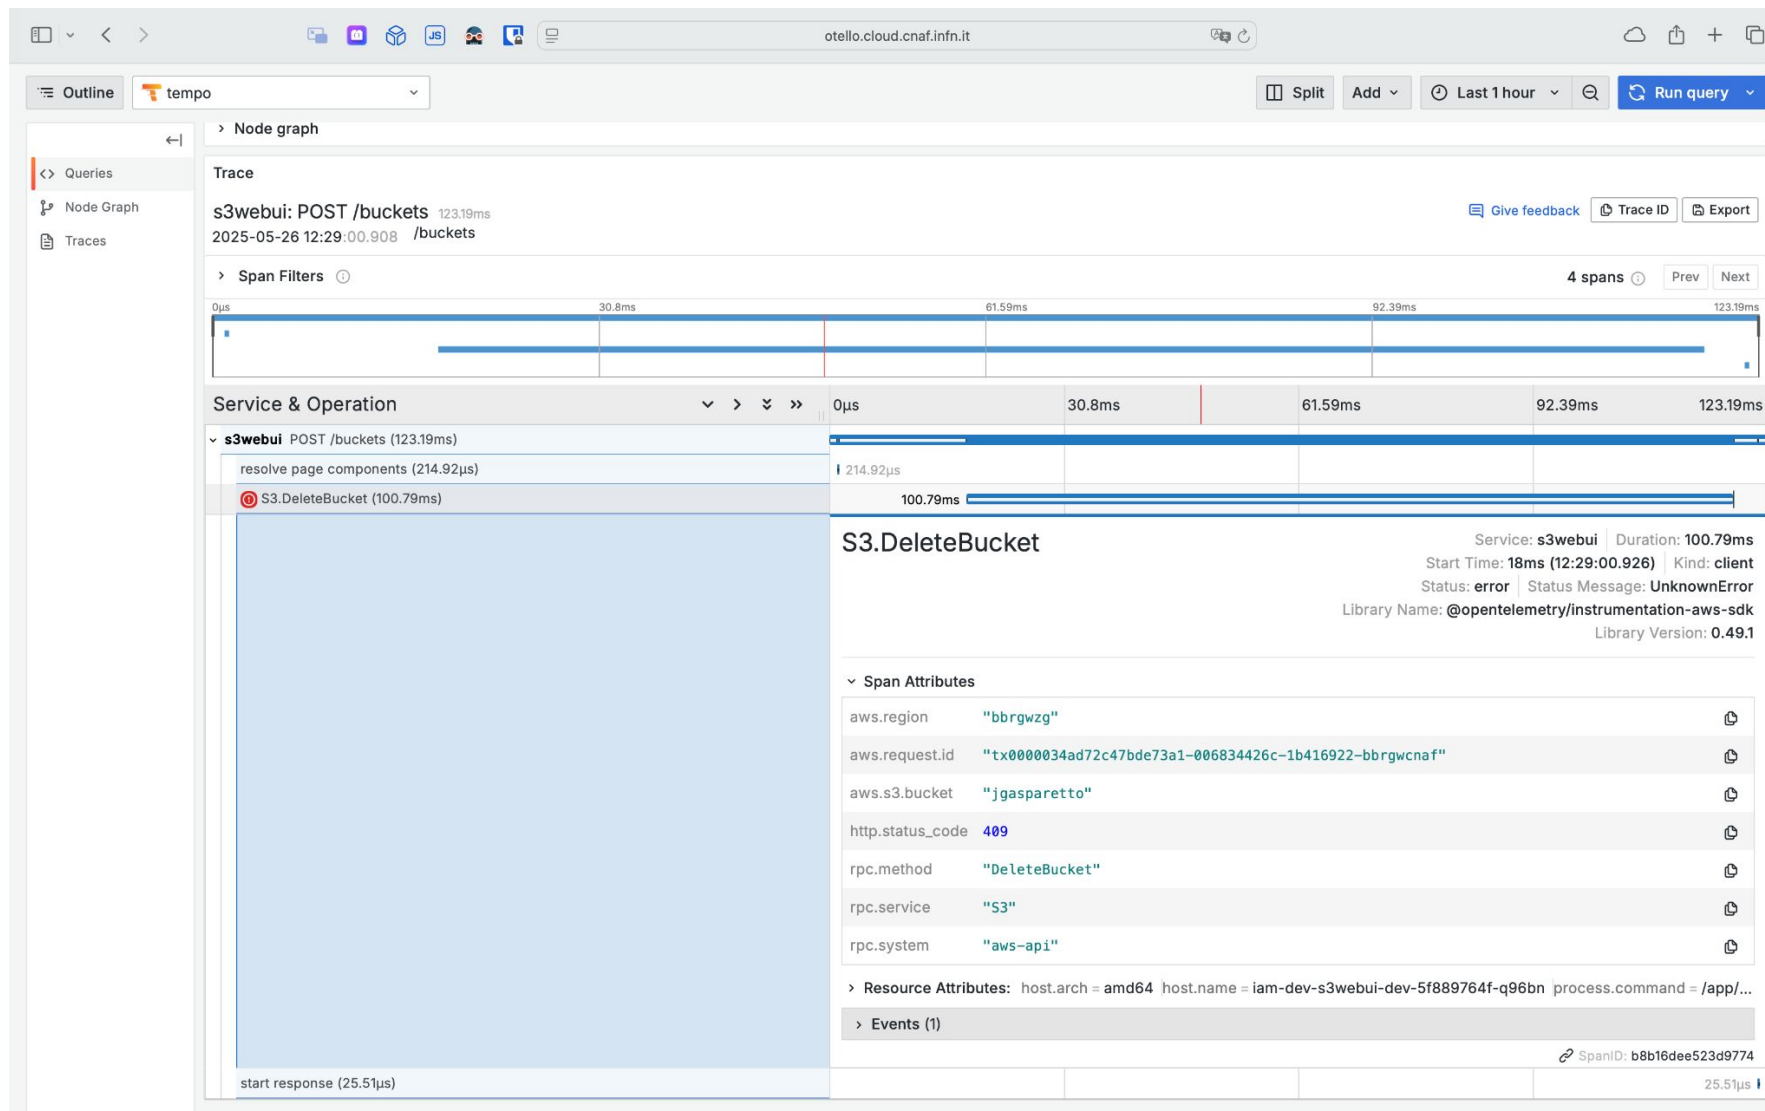
IAM Dashboard



S3 WebUI



S3 WebUI



S3 WebUI

Conclusions

- OpenTelemetry is a de-facto standard to produce telemetry data across a distributed system of any size.
- Tests (Otello) demonstrated that this technology is really promising for monitoring and debugging, especially to identify bottlenecks and unwanted path in our code.
- Even though Otello is experimental, StoRM Tape is already in production with OpenTelemetry support.
- The goal is to extend OpenTelemetry support to all software developed at INFN-CNAF, merging the Otello stack with the t1metria monitoring service.
- Proposal: integration of the OpenTelemetry support to software produced for DataCloud, perhaps integrating a stack similar to Otello and managed by WP1



Finanziato
dall'Unione europea
NextGenerationEU



Ministero
dell'Università
e della Ricerca



Italiadomani
PIANO NAZIONALE
DI RIPRESA E RESILIENZA



Backup

Propagation

- A request can travel across several distinct services, such as, for example, a reverse proxy (NGINX), an authorization engine (OPA), a REST API, a database (SQL), etc.
- OpenTelemetry fully supports the concept of [Context Propagation](#) as indicated by the [W3C TraceContext](#) specification, which *“allows traces to build causal information about a system across services that are arbitrarily distributed across process and network boundaries”*.
- This mechanism is generally established in web services using the **HTTP headers** `traceparent` and `tracestate`
- Spans are always emitted individually to the collector by the different services.
- Many third-party software, such as OPA, NGINX and K8S Ingress Controller, natively supports OpenTelemetry's Context Propagation



Sampling

- The amount of telemetry data, especially traces, can be huge, and overhead at runtime can be non-negligible. *“Sampling is one of the most effective ways to reduce the costs of observability without losing visibility”* <https://opentelemetry.io/docs/concepts/sampling>.
- Sampling can be performed probabilistically, for example: “sample only 10% of the incoming requests”
- Sampling can be controlled via `ParentBasedSamplers` that produce a span only if a parent span is detected.
- The two approaches can be combined. For example, the first service receiving the requests (e.g, NGINX) can decide to sample only a fraction of them, then the other services of the chain will sample or based on the presence or not of the parent span