GÉANT
Networks • Services • People

Detecting anomalous latency and jitter in Timemap

Fabio Farina (GARR)
On behalf of GN4-3 WP6 T1 Team

2nd Performance Management Workshop
March 8th 2022

Public

www.geant.org

Outline

- Why Timemap
- Current status
- Beyond observation: anomaly detection
- Further development

Road report: on HWY 101 there are 364 vehicles per minute

it may be nice…

Or lots of stop & go

Current weather maps are useful for bulk data transfer!

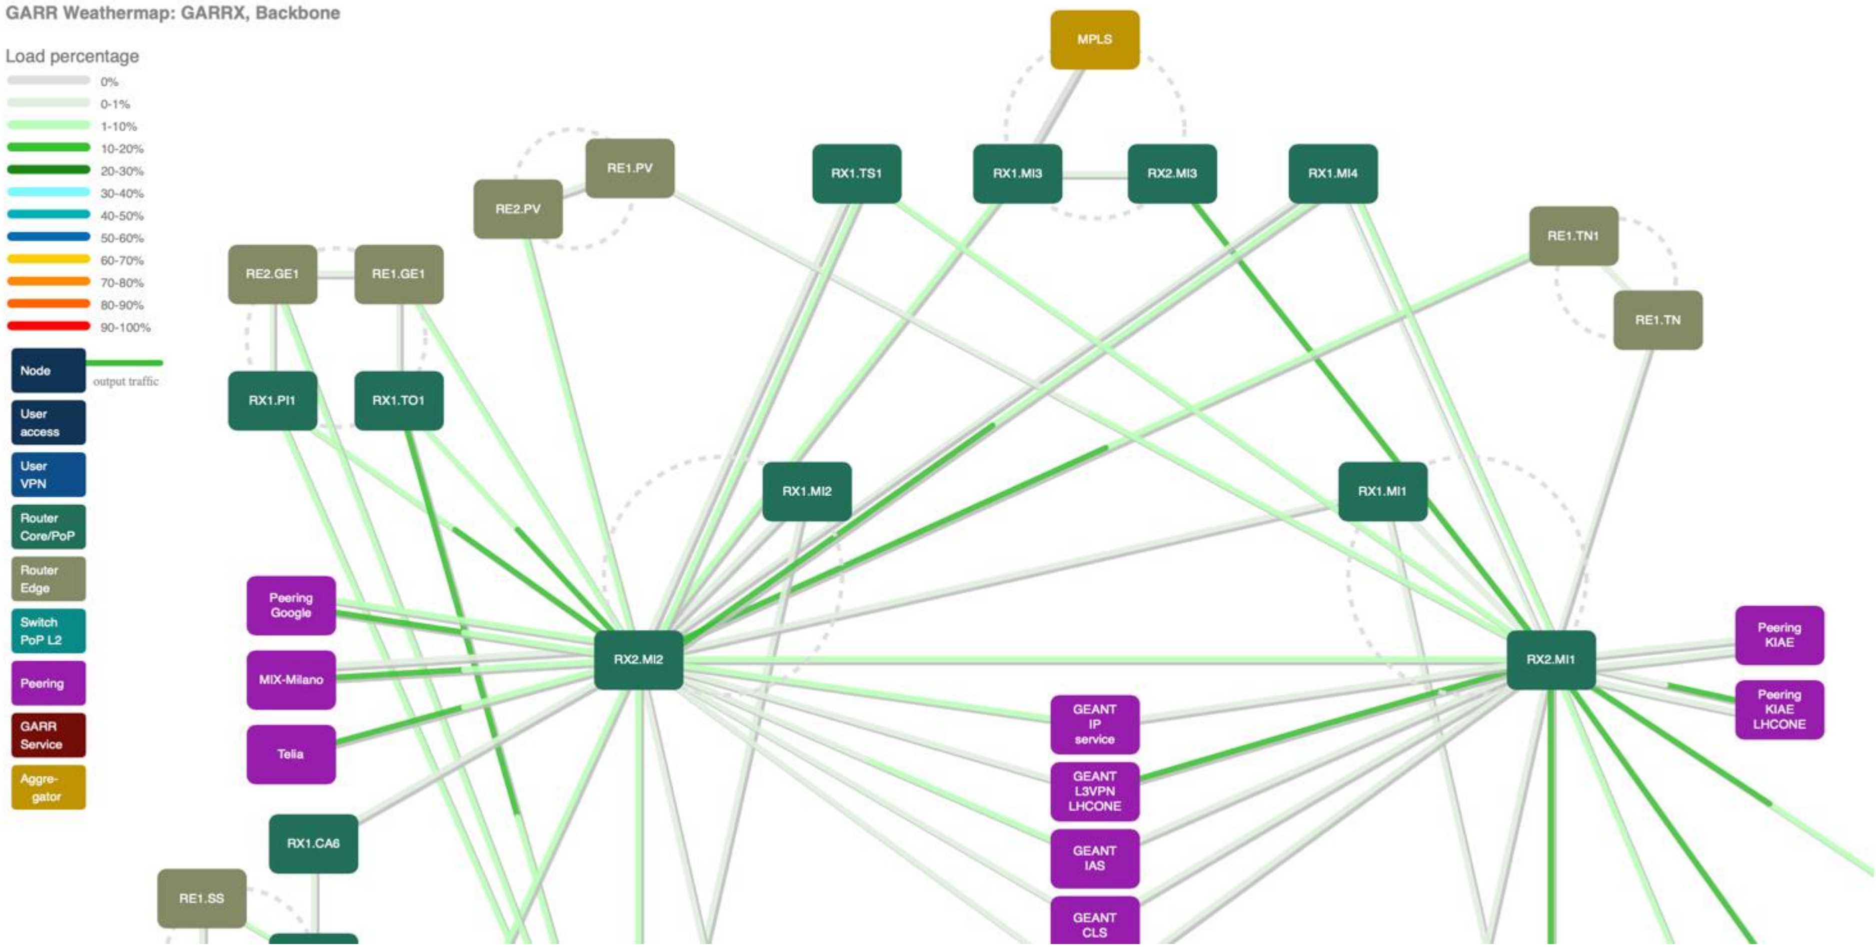

Applications which need “cruise control” on are on fast rise!

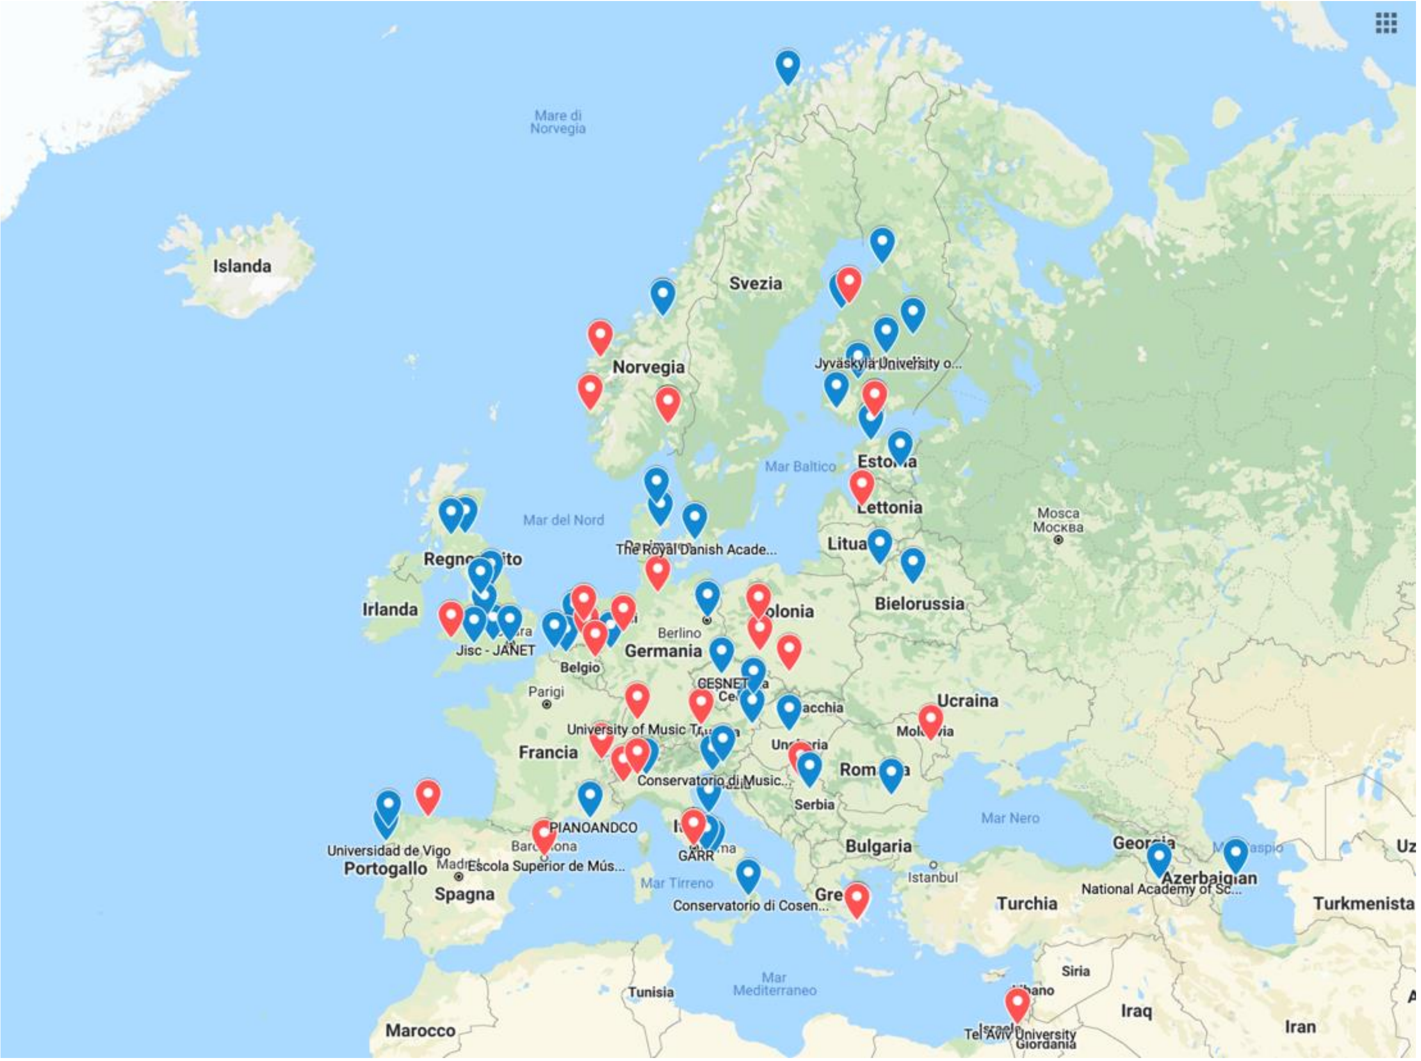

- LoLa

+ 30%

Involved NRENs: do they all know about this?

Rhnet*
Jisc
RedIris
RENATER
BELnet
SURFnet
GARR
SWITCH
DFN
Uninett*
Sunet*
DeIC*
ACOnet
CESnet

PSNC
ARNES
AMRES
NIIF
RoEduNet
GRNET
Litnet
EEnet
FUNET*
BASNET
ASNET-AM
AZscienceNet
RENAM

* Nordunet

LATNET
IUCCL

3 x 450Mbps
RTT:
- 15ms
- 25ms
- 37ms

TIMEMAP architecture and features

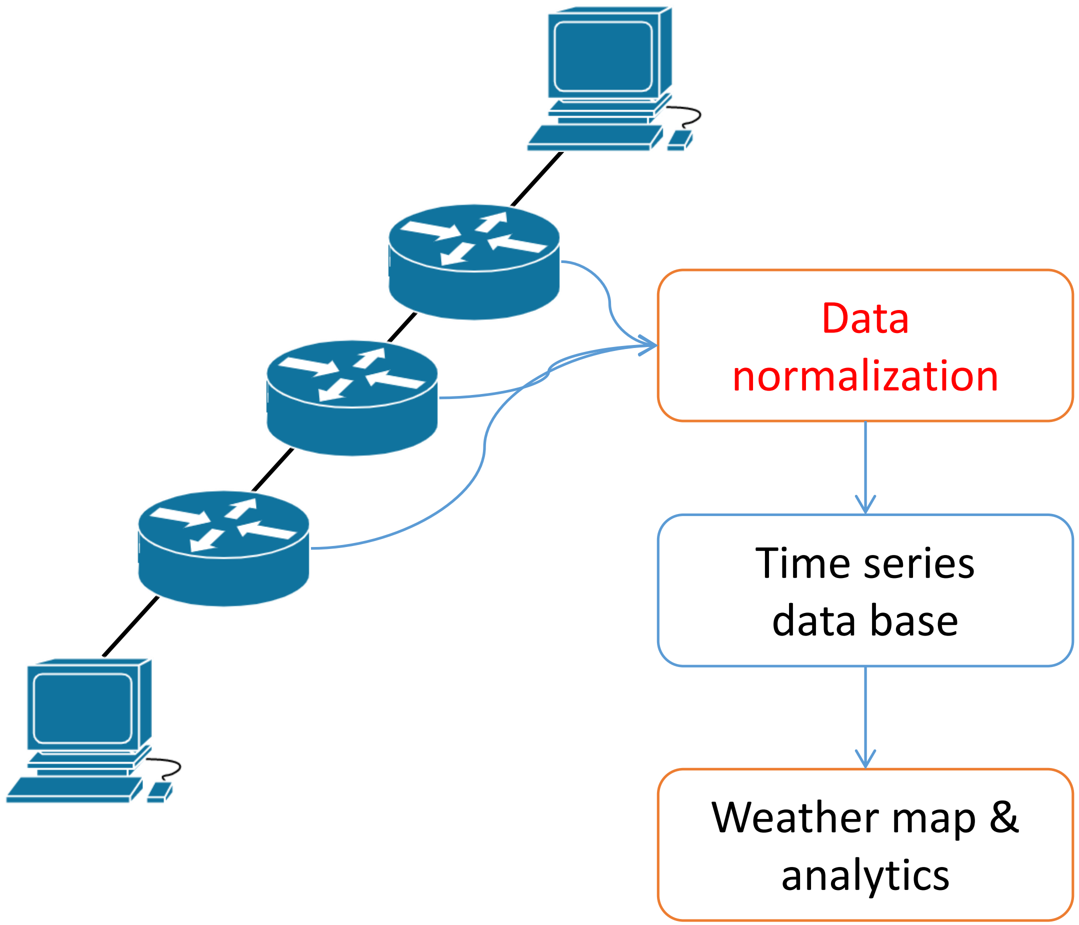

- Latency & Jitter data collection
  - RPM from all GÉANT routers
  - TWAMP from selected perfSONAR VMs and GÈANT routers
- Simplicity: almost zero footprint
  - Docker + Linux packages
  - Minimal custom code
  - Dynamic weather map GUI
- Security
  - eduGAIN authentication
  - RBAC and multi-tenancy

TIMEMAP v1 architecture – 1+ year of data taking

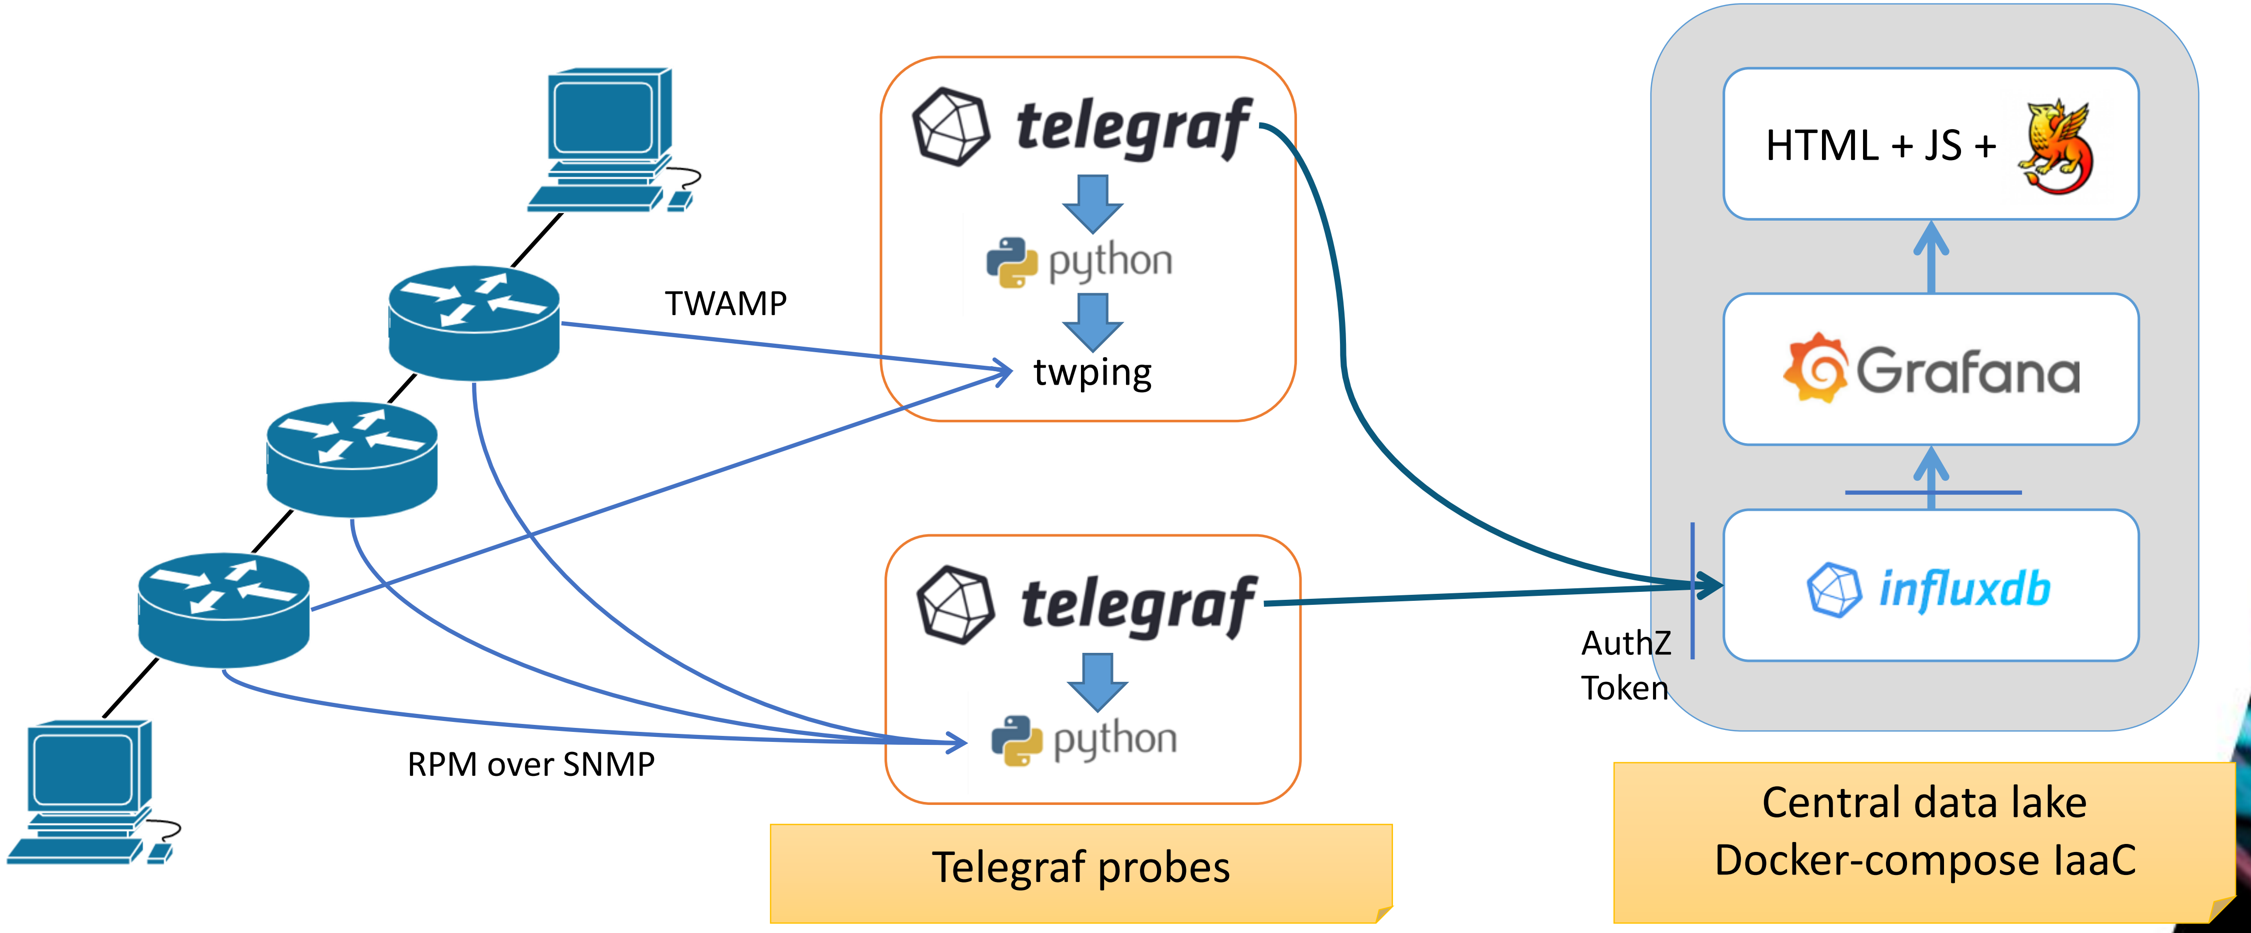

Have a look

- The service

https://timemap.geant.org/

- Documentation: source code, user and admin guides, customization

https://gitlab.geant.org/gn4-3-wp6-t1-lola/timemap_public

Timemap useful examples

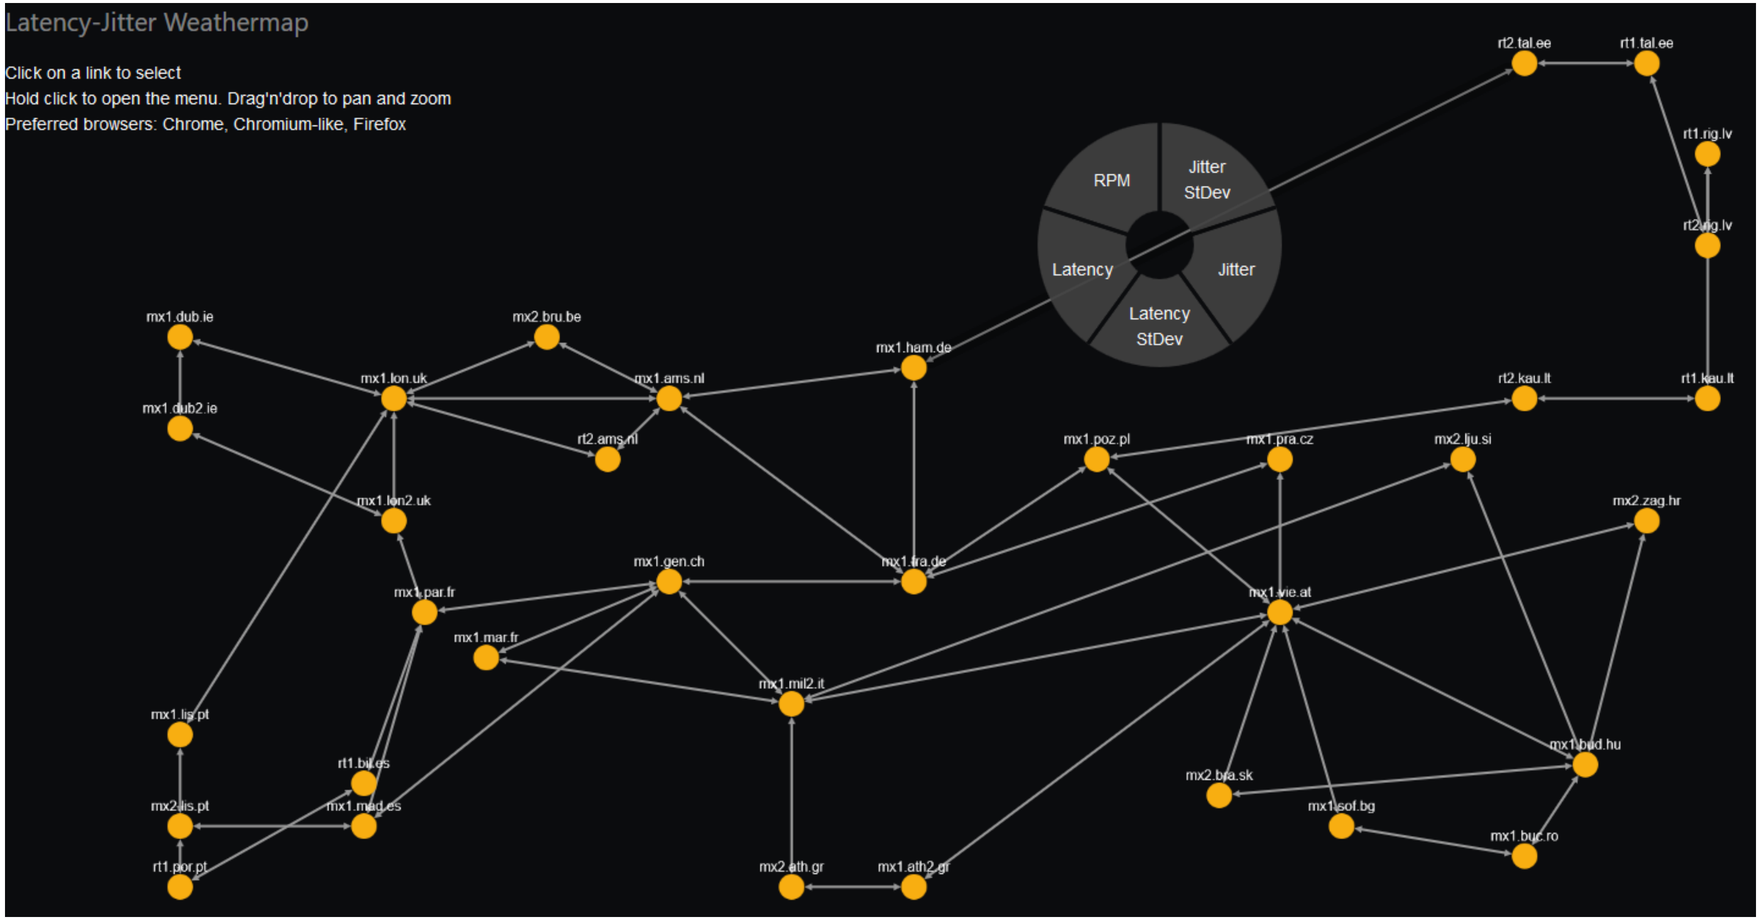

Latency-Jitter Weathermap

Click on a link to select

Hold click to open the menu. Drag'n'drop to pan and zoom

Preferred browsers: Chrome, Chromium-like, Firefox

Timemap useful examples: rerouting and periodic events

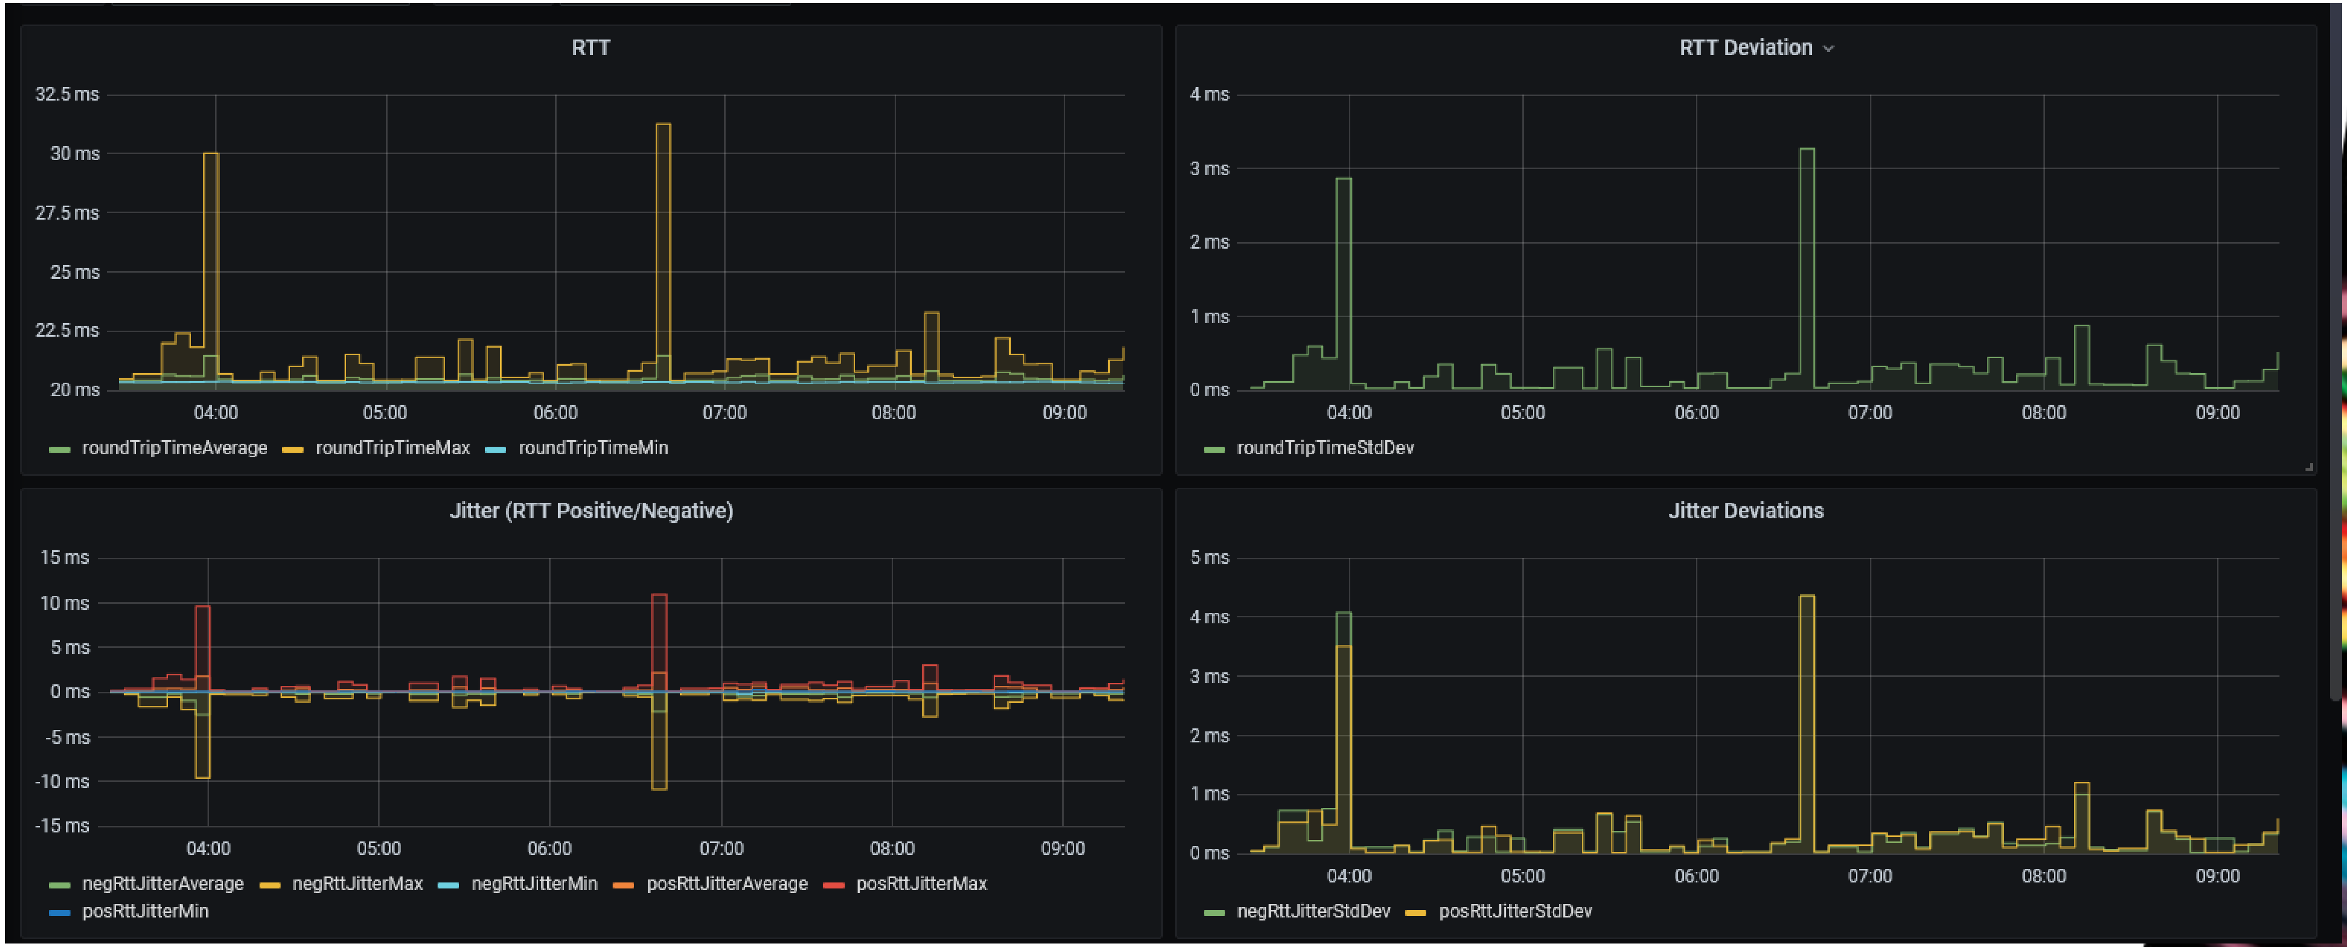

Timemap useful examples: trends

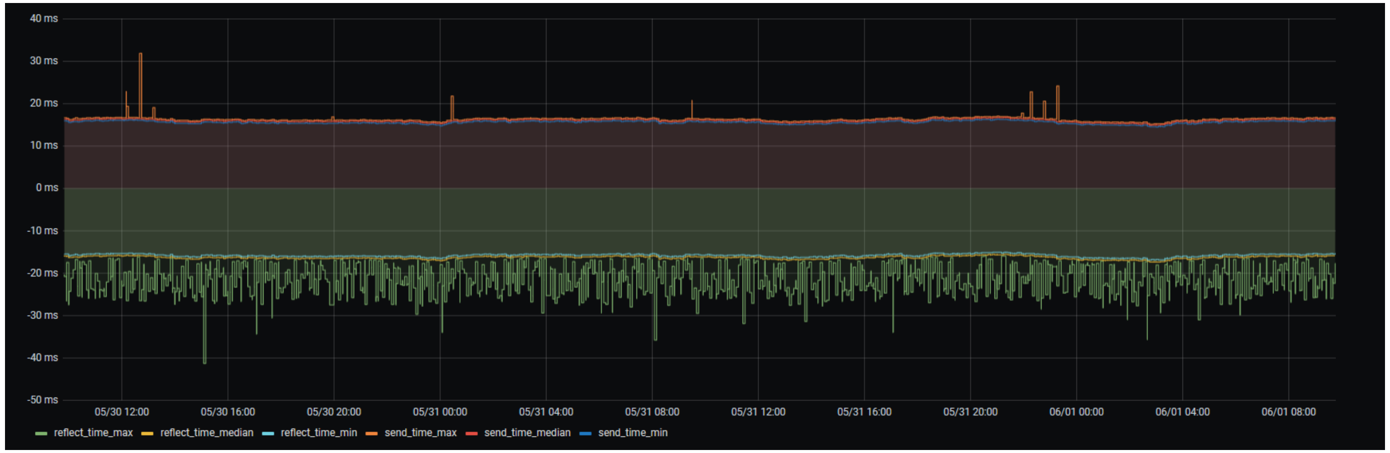

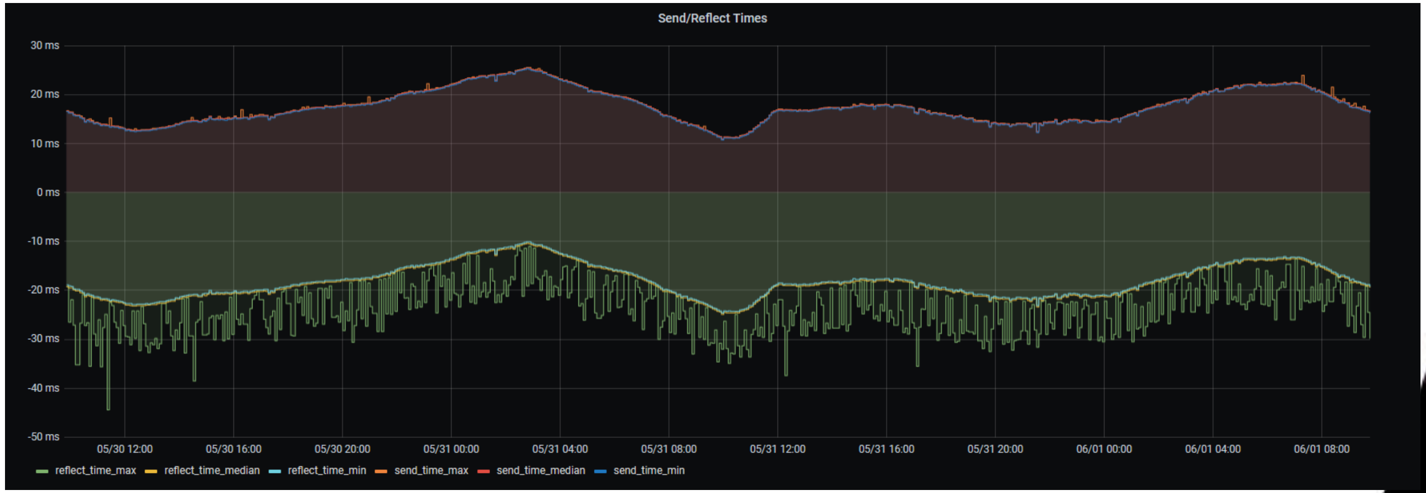

Anomaly Detection (AD) in Timemap - requirements

- Move beyond the simple observation
  - AD for Analytics and Alerting
  - Co-occurring events correlation

- Requirements on AD machine learning
  - Real-time or micro-batch learning/inference
  - Robust estimation
  - Light footprint

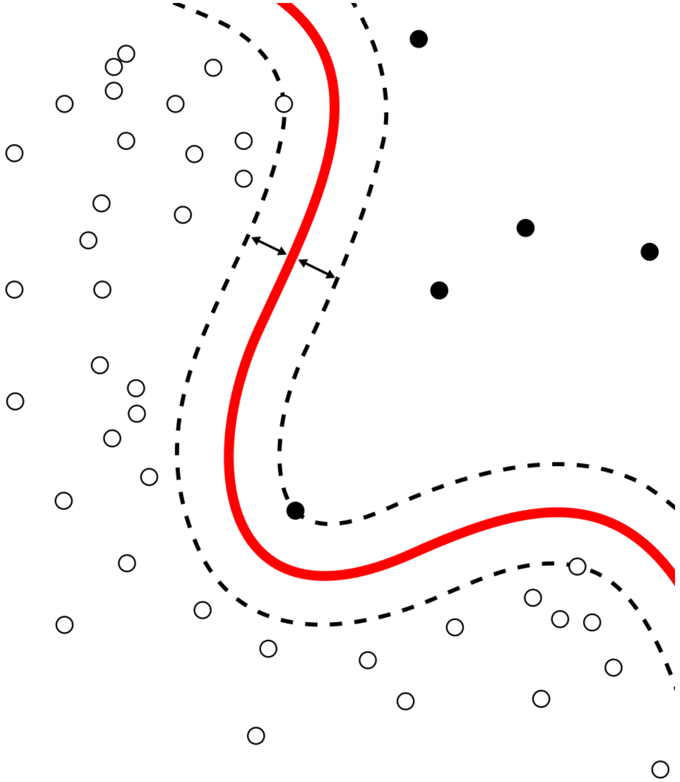

Anomaly Detection in Timemap – toolset

- Anomaly Detection, in short
  - Std.Dev classification
  - Unsupervised
  - Sensible to overfit

- Streaming ML in Python https://riverml.xyz

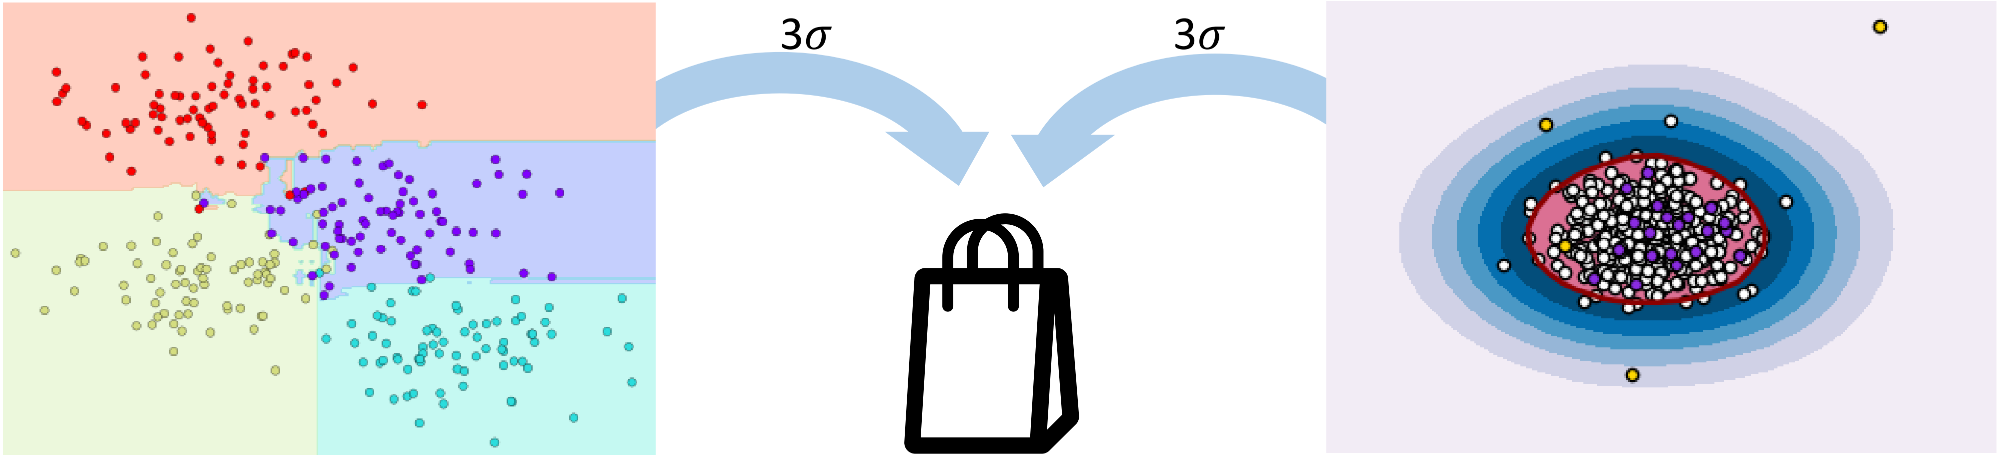

Half-space Random Trees

One-class Support Vector Machine

Model bagging

Anomaly Detection in Timemap – architecture

Software architecture extension

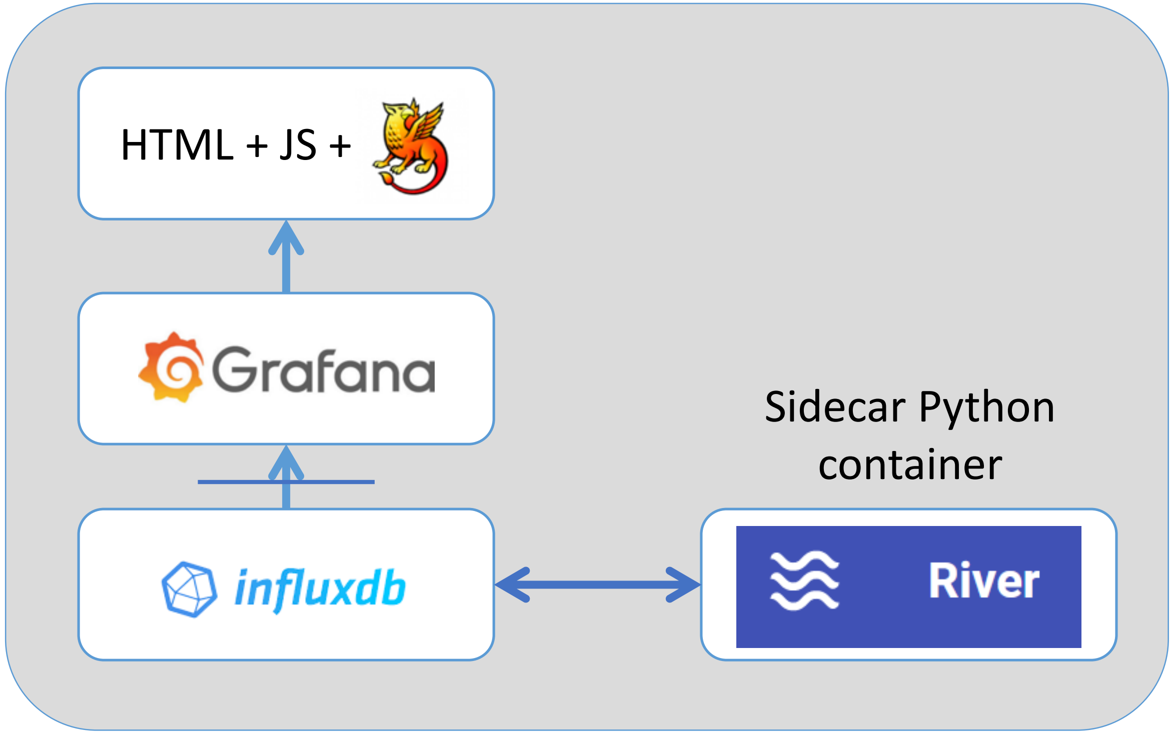

Network topology and ML models

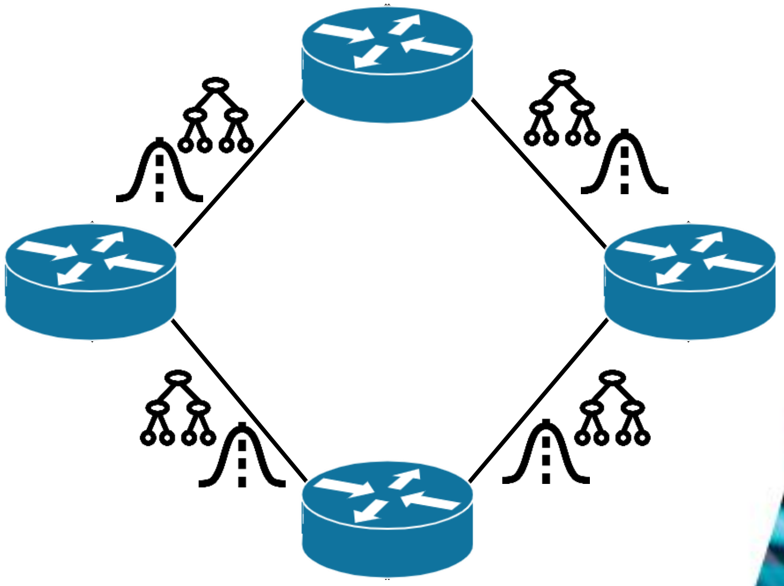

Almost the same look and feel

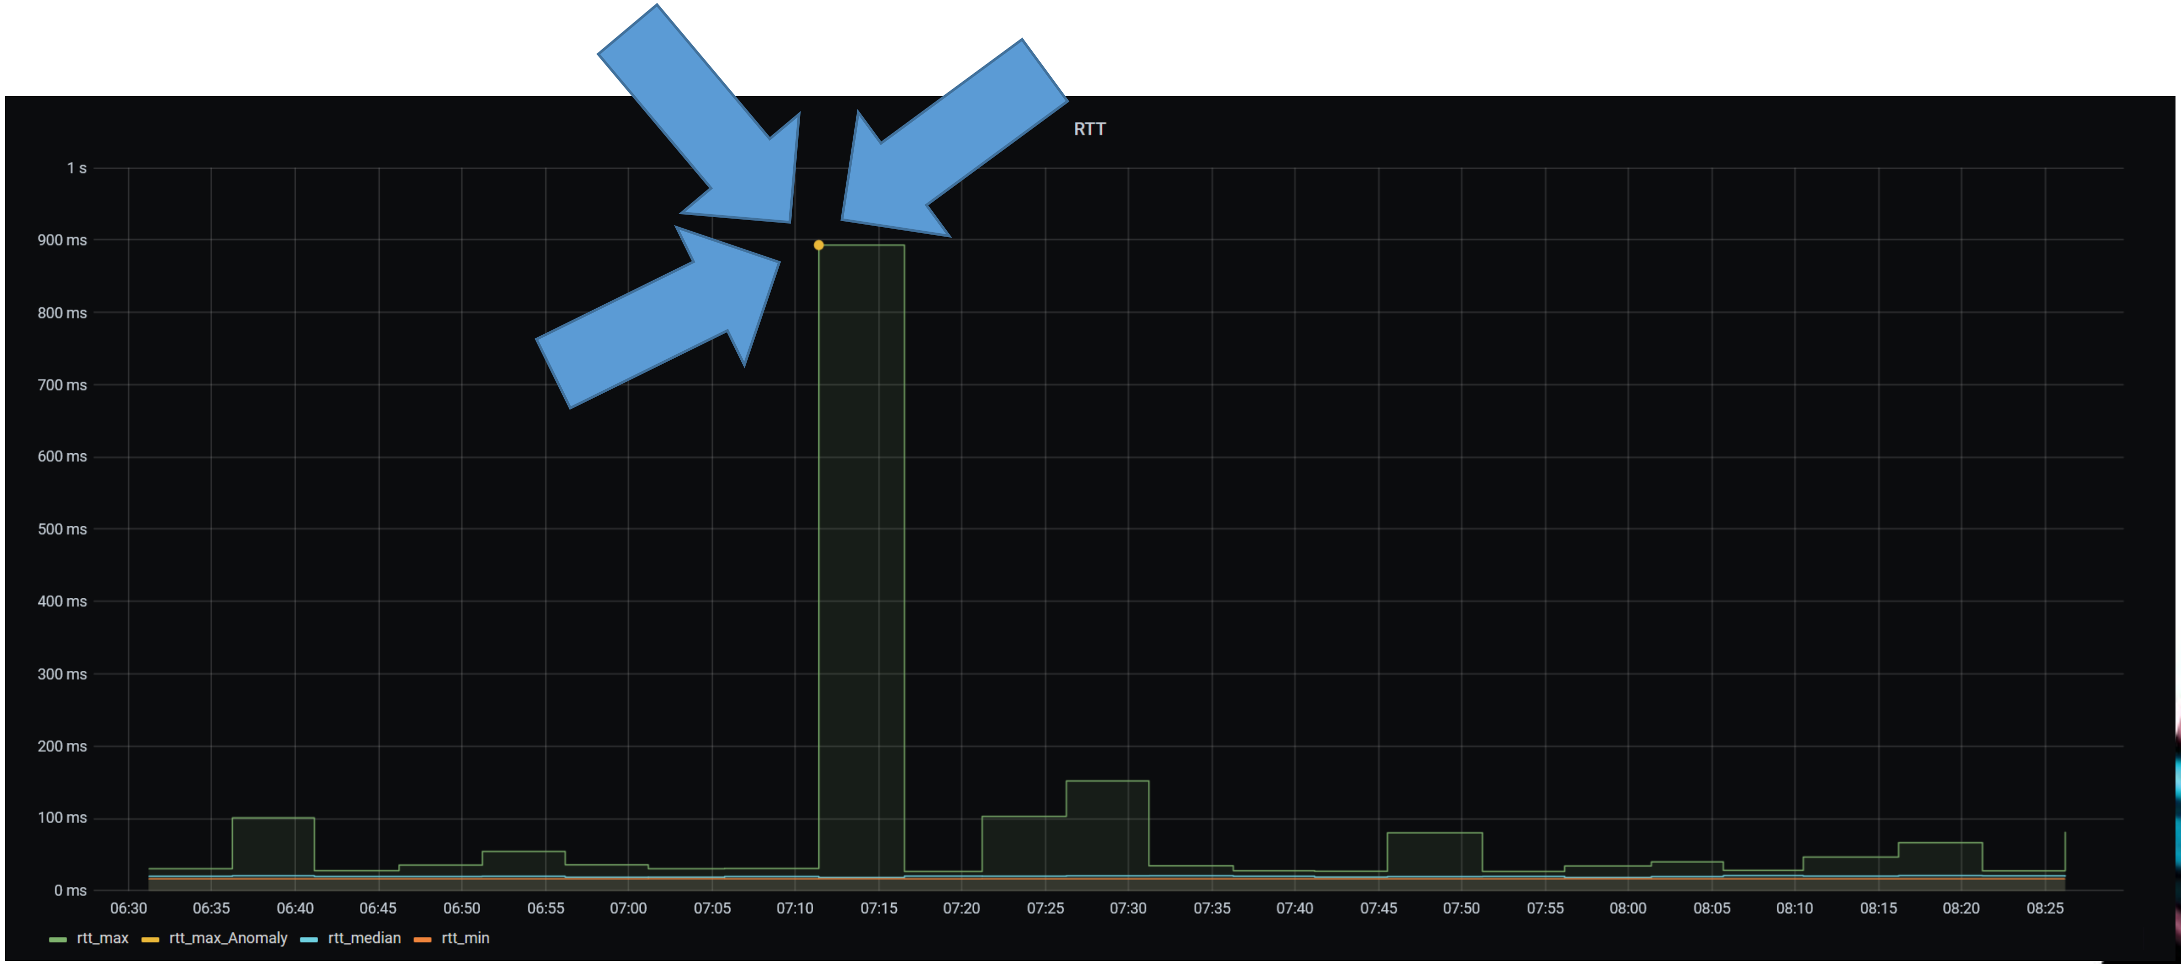

One more plot

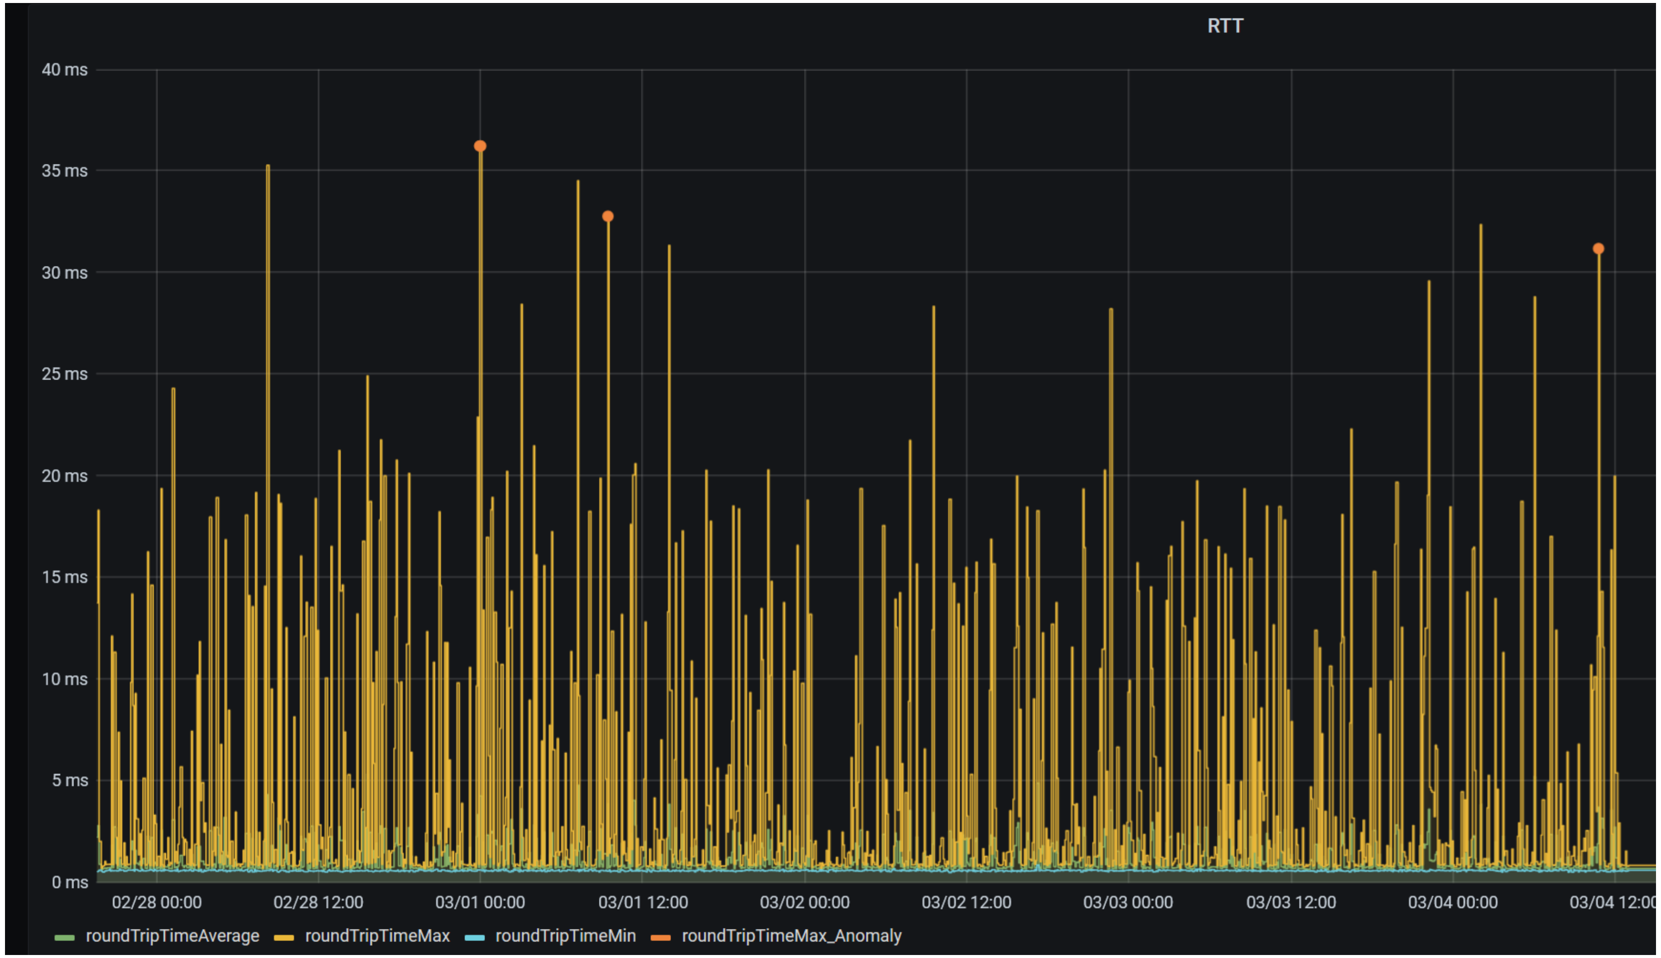

Performance and efficiency

- AD frequency == sampling frequency (5m)
- ~3500 observations every micro-batch (3s execution time)

- Learning rate

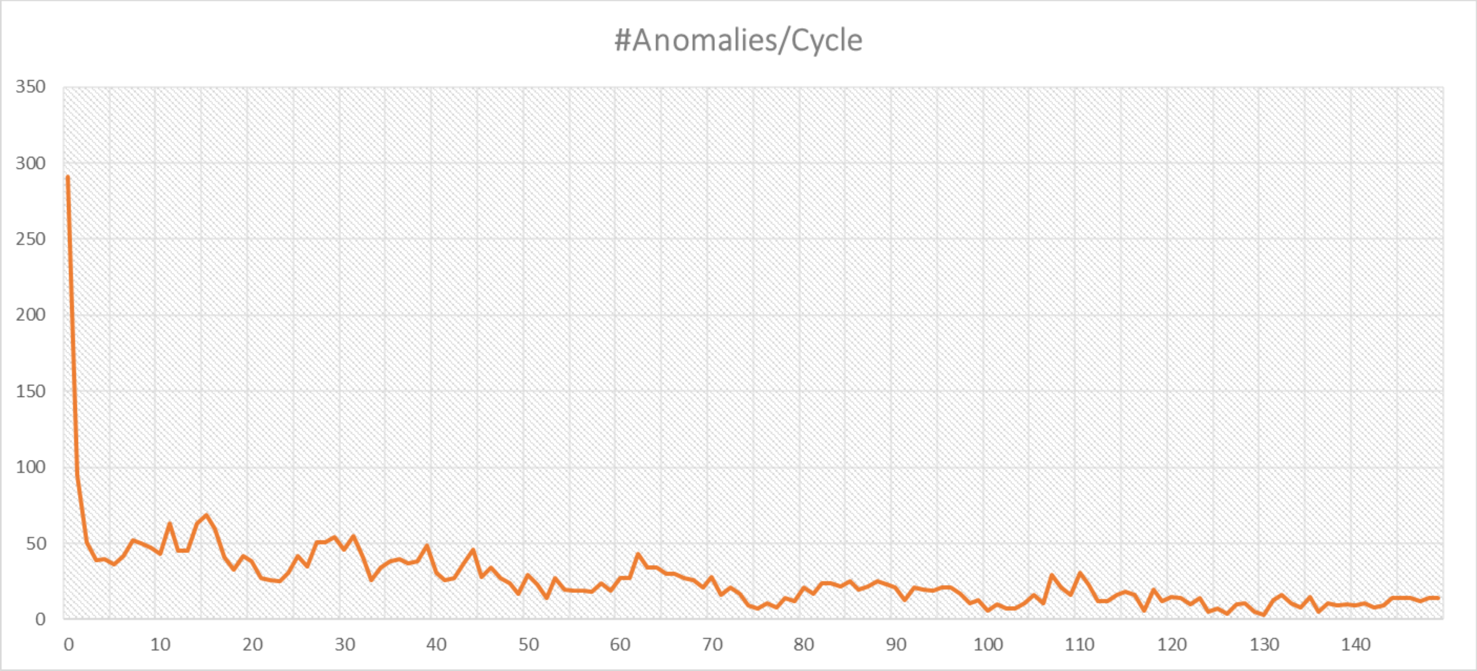

Conclusion and next steps

- Anomaly Detection
  - Up and running, requirements satisfied
  - Streaming ML, multi-model over network topology
  - About 200 lines of code in a Docker image

- Coming soon
  - Timemap hand-over to Géant Operations
  - Timemap @ GARR
  - Alerting & Events correlation

- More deployments @NRENs

Thank you

Any questions?

gn4-3-wp6-t1-lola@lists.geant.org

www.geant.org

© GÉANT Association on behalf of the GN4 Phase 3 project (GN4-3).
The research leading to these results has received funding from
the European Union's Horizon 2020 research and innovation
programme under Grant Agreement No. 856726 (GN4-3).