Operational Monitoring Persistence and Dashboard: Status and Outlook

Igor Soloviev
TDAQ Week
September 2018

Introduction

- **P-BEAST (Persistent Back-End for the ATLAS Information System of TDAQ)**
  - Time-series database to implement operational monitoring data persistence
    - Home-made on top of Google protobuf, CORBA and libhttpd
    - Is used since end of Run 1 replacing Cassandra prototype
  - Data sources: IS, network monitoring (REST)
  - Data retrieval: REST, CORBA, C++, Java, Python APIs
  - Visualisation GUI:
    - Grafana Dashboards using P-BEAST data plug-in
    - SWAN Beauty using P-BEAST APIs

Data Insertion Rates

- **Store all raw IS data but histograms from ATLAS, few infrastructure detector partitions and network monitoring**
- **ATLAS has average 180 KHz insertion rate during data taking**
- **Network monitoring has constant rate around 10 KHz**
- **Other partitions may contribute significant data as well**

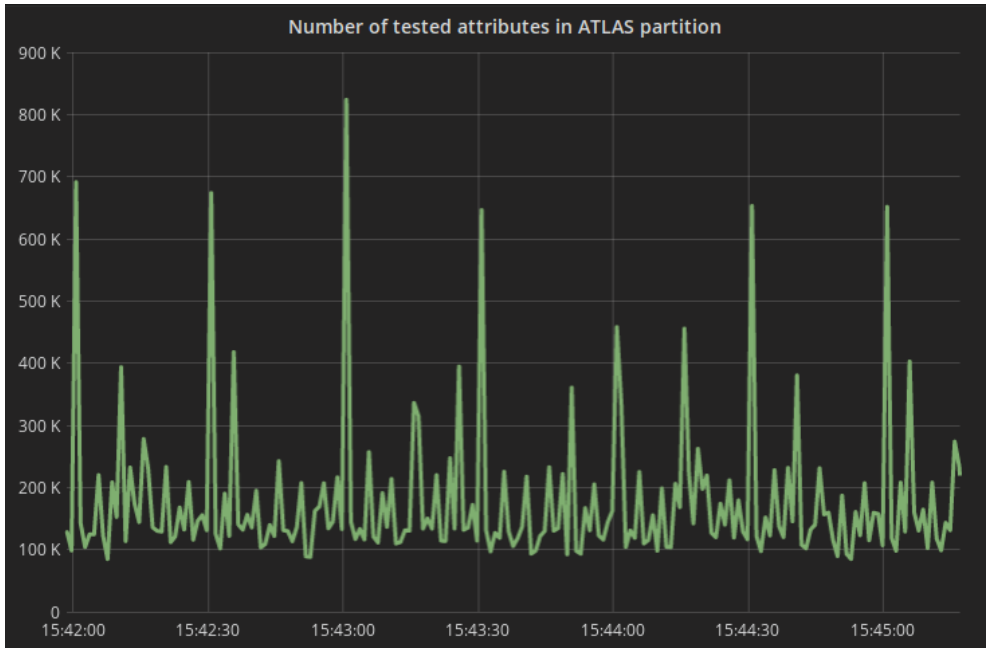

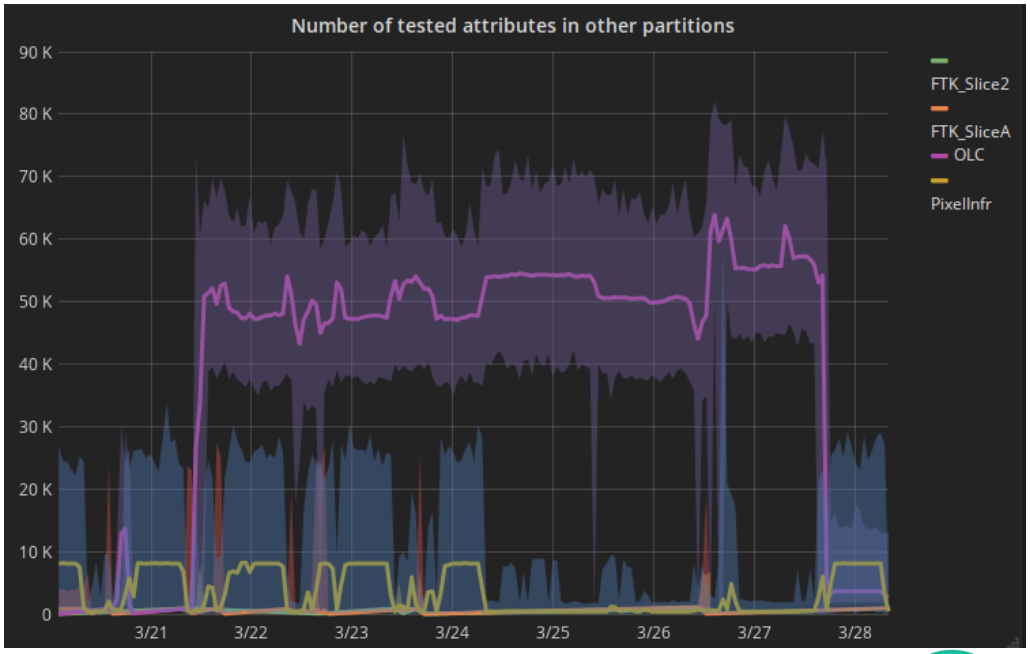

Storage Resources

- **Two 2015q4 32 TB and three old 2012q1 4 TB nodes**
- **1.5 TB / month during data taking periods**
- **Store raw data with compaction and compression**
- **Safe for Run 2**

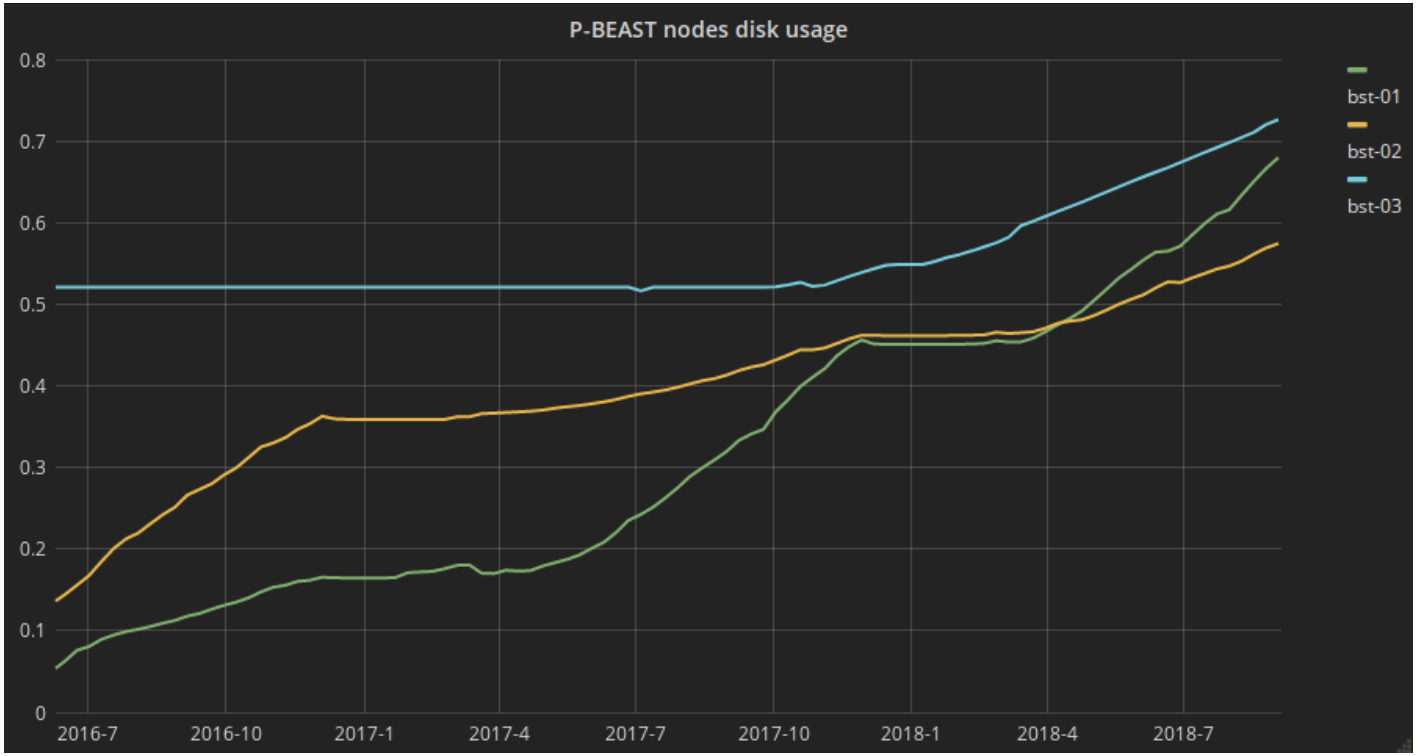

- **Weekly archive data on EOS (increase quota to 60 TB)**
  - The EOS is too unstable for production DB use, but suitable for archive
- **Plan to store raw data for lifetime of ATLAS**

Queries Statistics

- 75 Grafana dashboards
- ~30 queries per minute (dashboards auto-refresh)

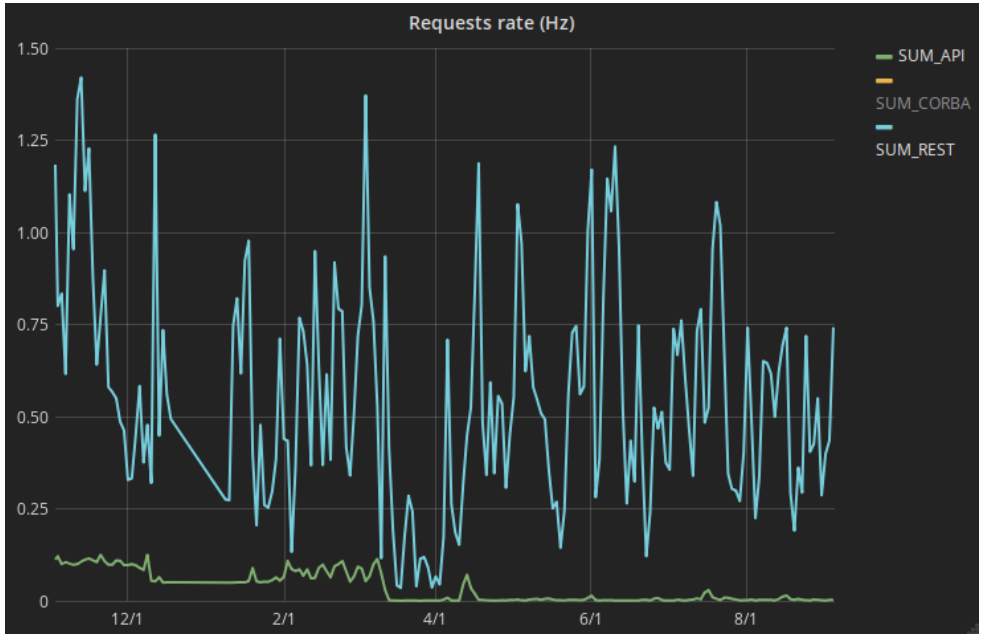

- There are periods with much higher request rates (above 50 queries per second)

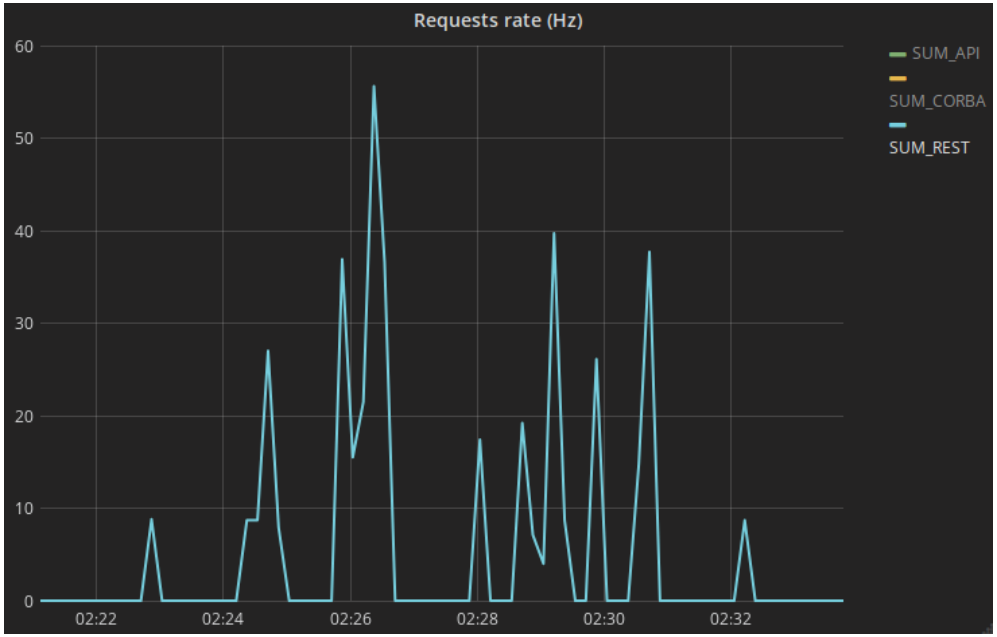

- P-BEAST is an essential tool for ATLAS operations and analysis. Provides access to live and historical data inside and outside P1.

Use Grafana 4

- **Migrated from Grafana v1.9 to v4.6**
- **Re-implemented and improved P-BEAST plugin design and revamped UI**
- **LDAP integration**
- **Use internal database for dashboards**
  - Every authorized user may create, save and update dashboards
- **Many new types of plots and panels are available:**
  - Histograms
  - Pie charts
  - Tables
  - Heatmaps
  - ...

Grafana 4: Scatter Plots

- **Show correlations between P-BEAST data (similar to functions available in Beauty)**

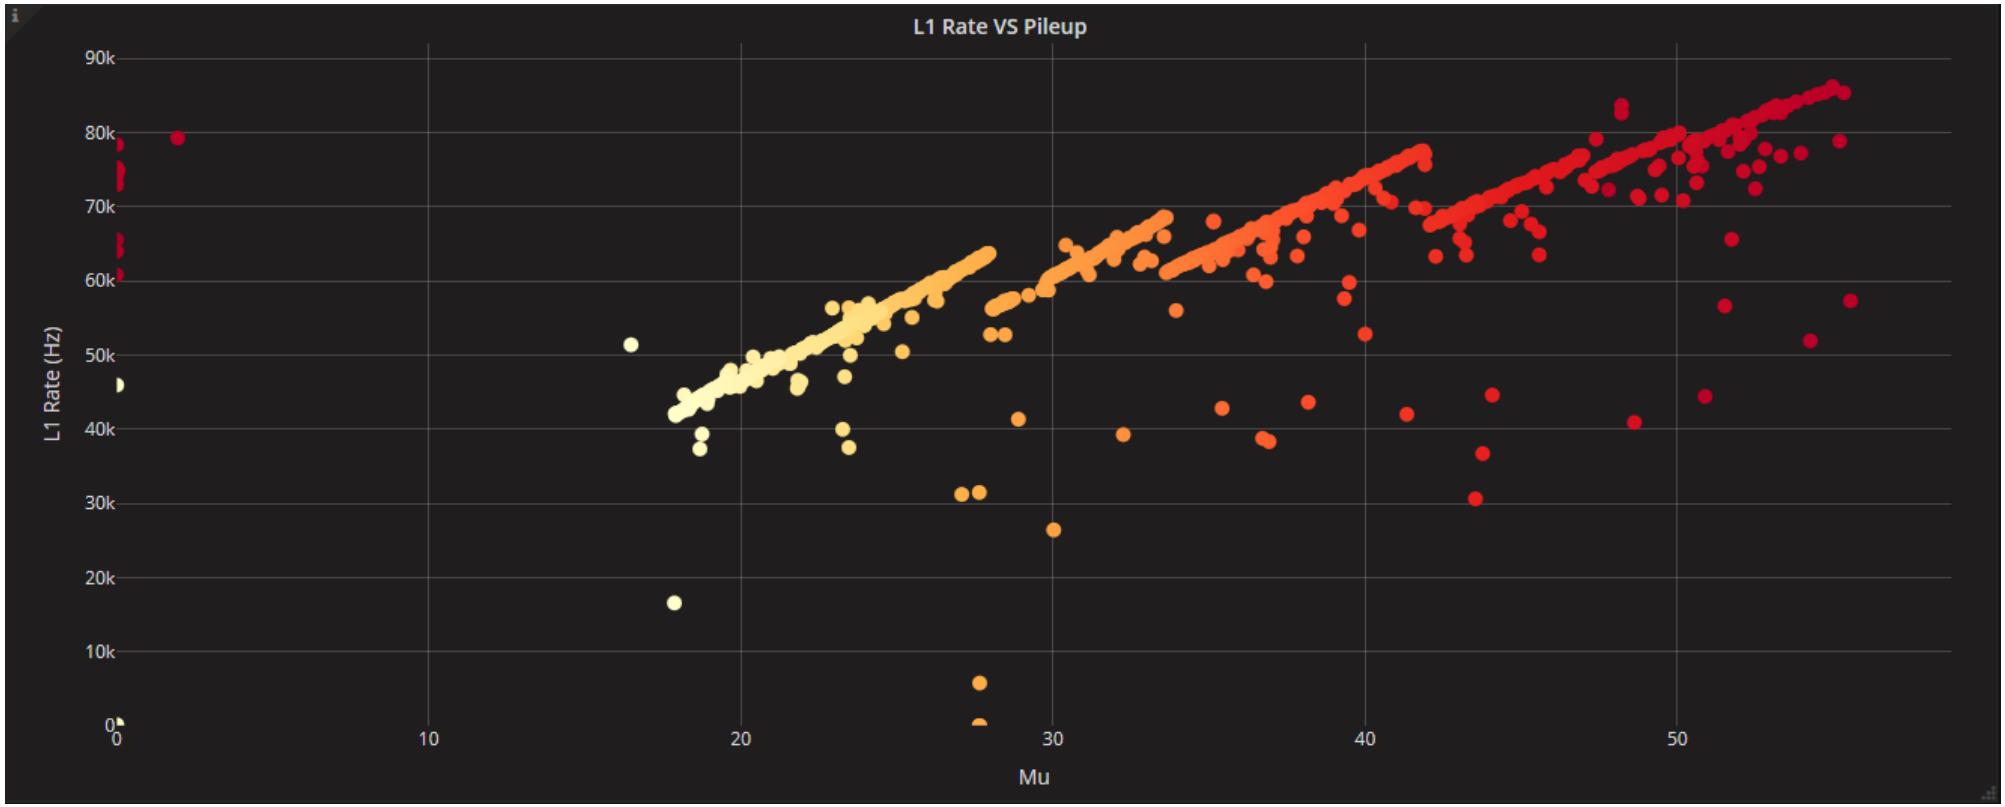

Grafana 4: Strings Data View

- **It is possible to display non-numeric data**

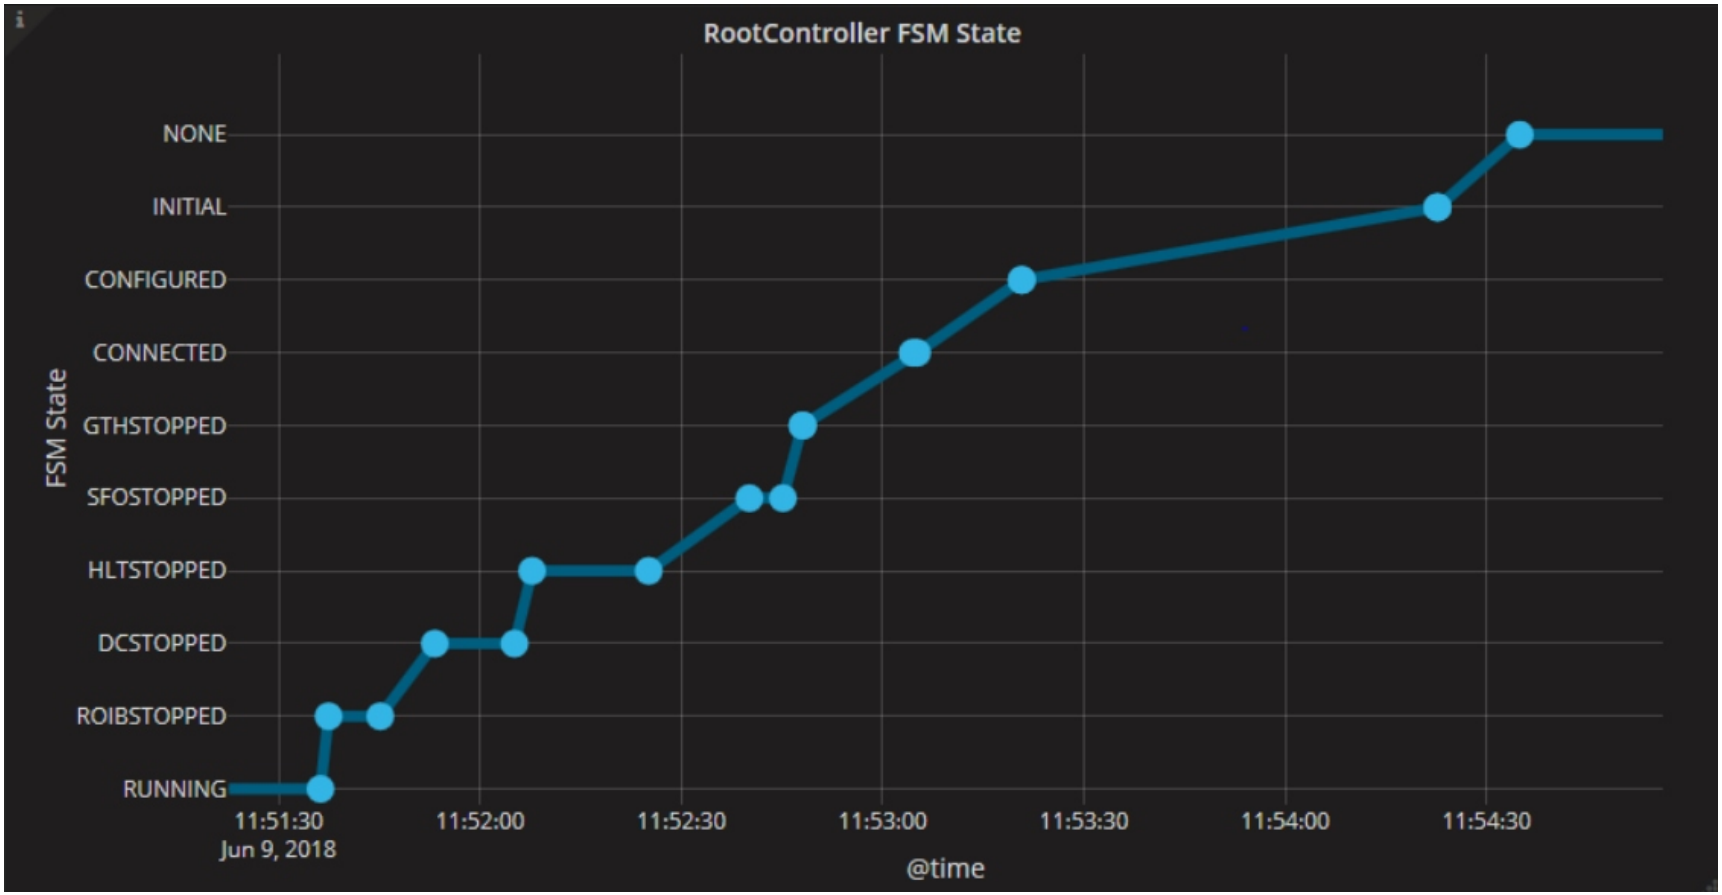

Next Grafana Version

- **Version 5.3 is going to be released**
  - v5 was not stable when we moved to version 4
  - Plan to move P-BEAST to new Grafana during LS2
- **New Dashboard Layout Engine**
  - Provides much easier drag, drop and resize features and new types of layouts
- **The UI has big improvements in both look and function**
- **Dashboard Folders**
- **Permissions**
  - On folders and dashboards helps manage larger Grafana installations
- **Group users into teams**
  - Use them in the new permission system

P-BEAST Service: Averages

- **Replace *random samples* by *averages* for down-sampling**

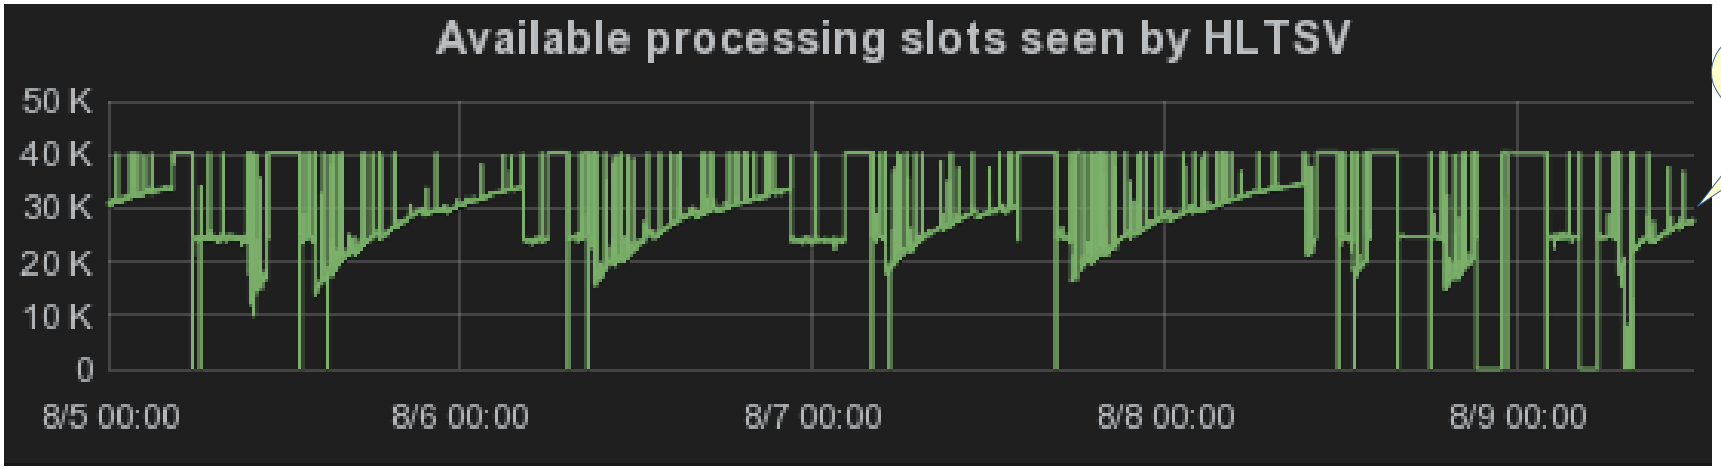

old

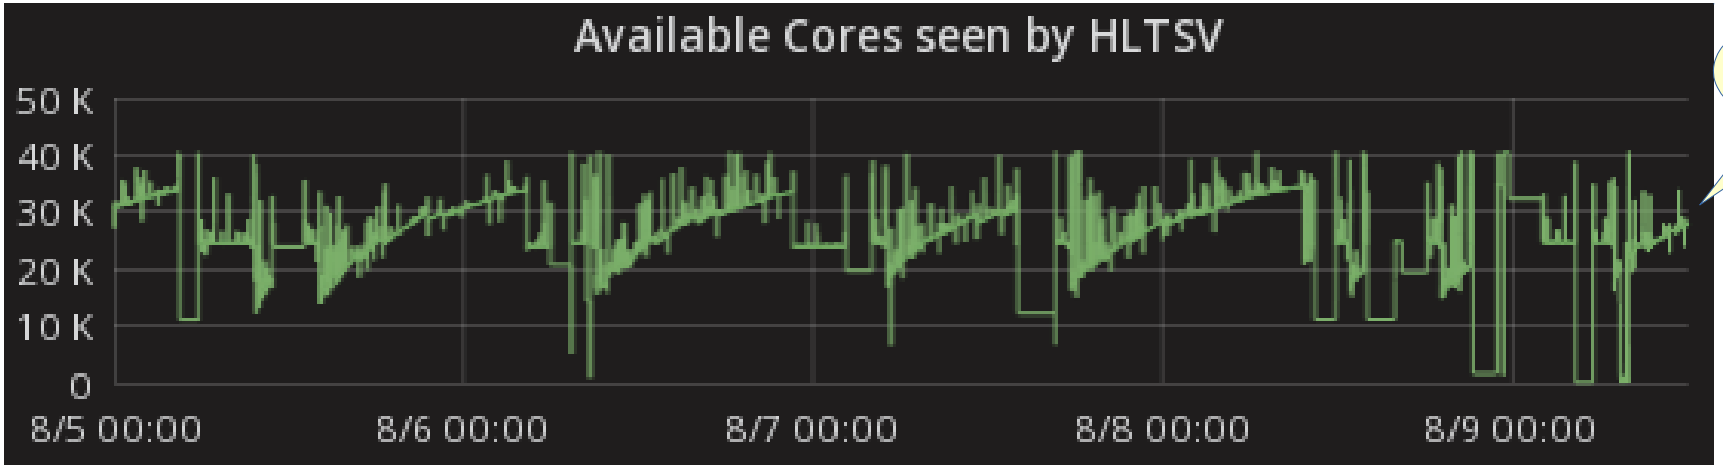

new

P-BEAST Data Functions

- **Aggregations: min, max, sum, average**
- **Data Transformation: sumArray, summarize**
- **Series select filters:**
  1) average/ current / highest / lowest
  2) above / below / max / ...
- **Data Filters:**
  - Remove data above / below some value
- **Keep Graphite names of functions**
- **Available in P-BEAST query editor**
- **See few examples on next slides ...**

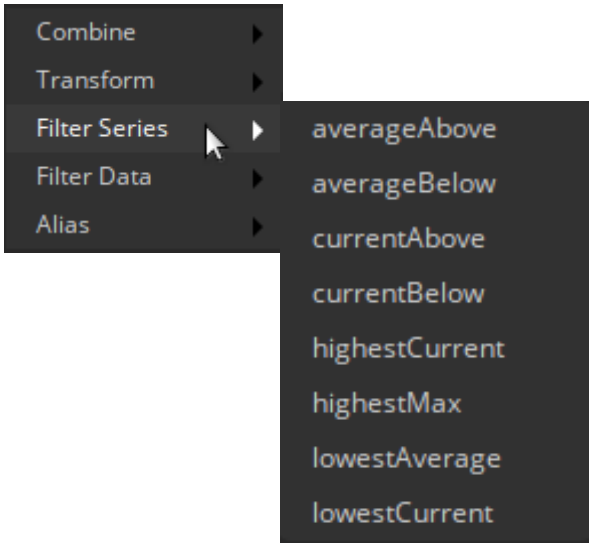

P-BEAST Service: Aggregation Functions (1/3)

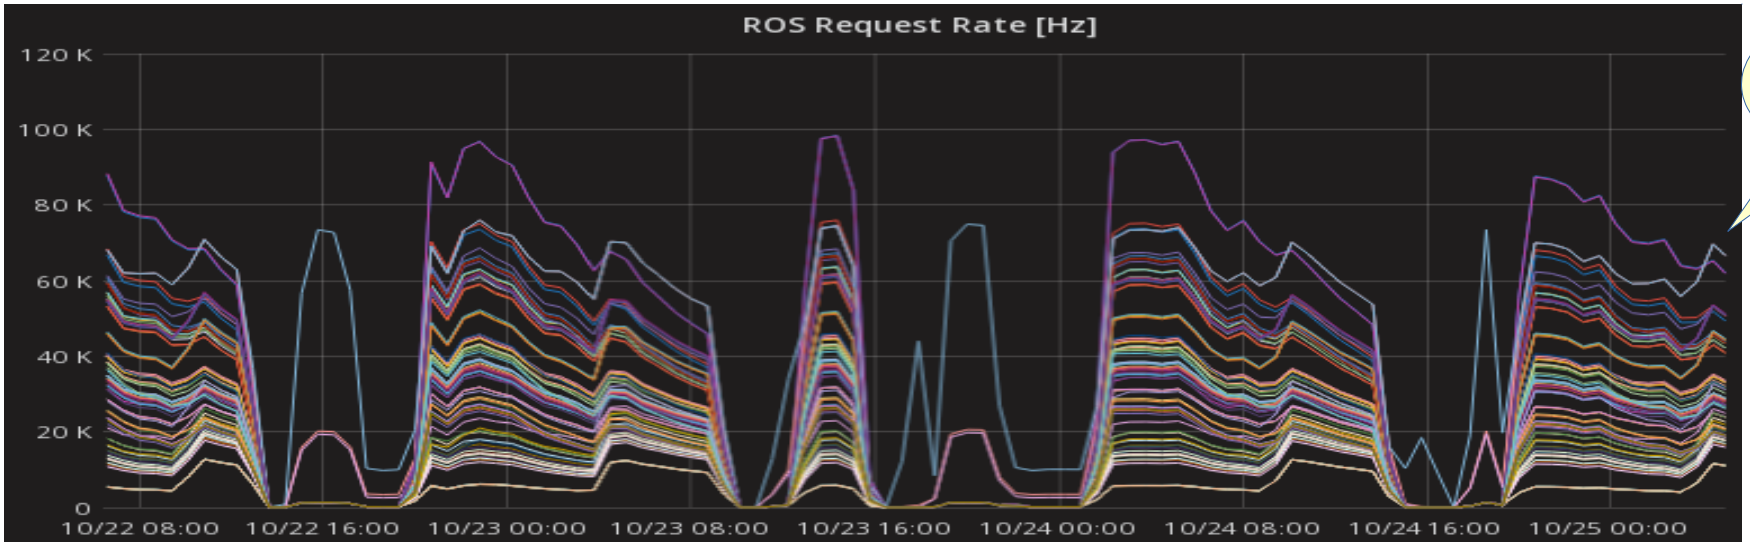

query data

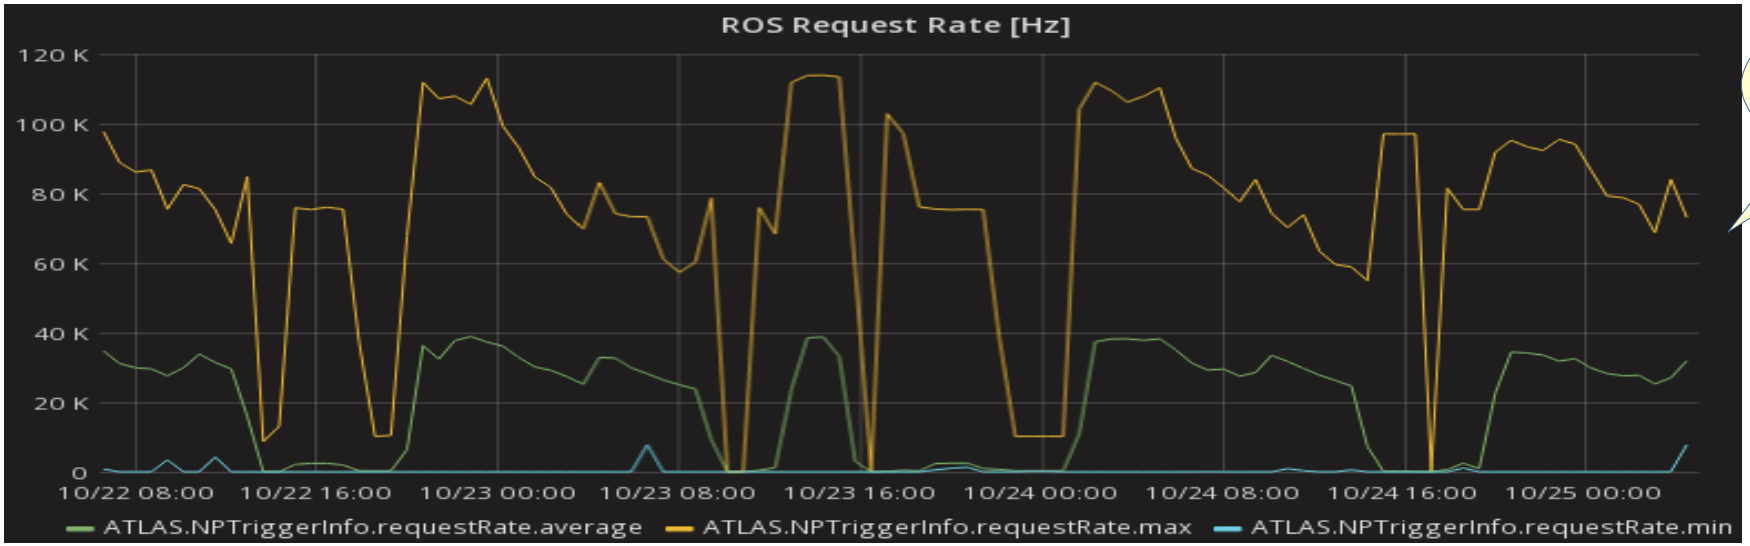

min,max, average

P-BEAST Service: Aggregation Functions (2/3)

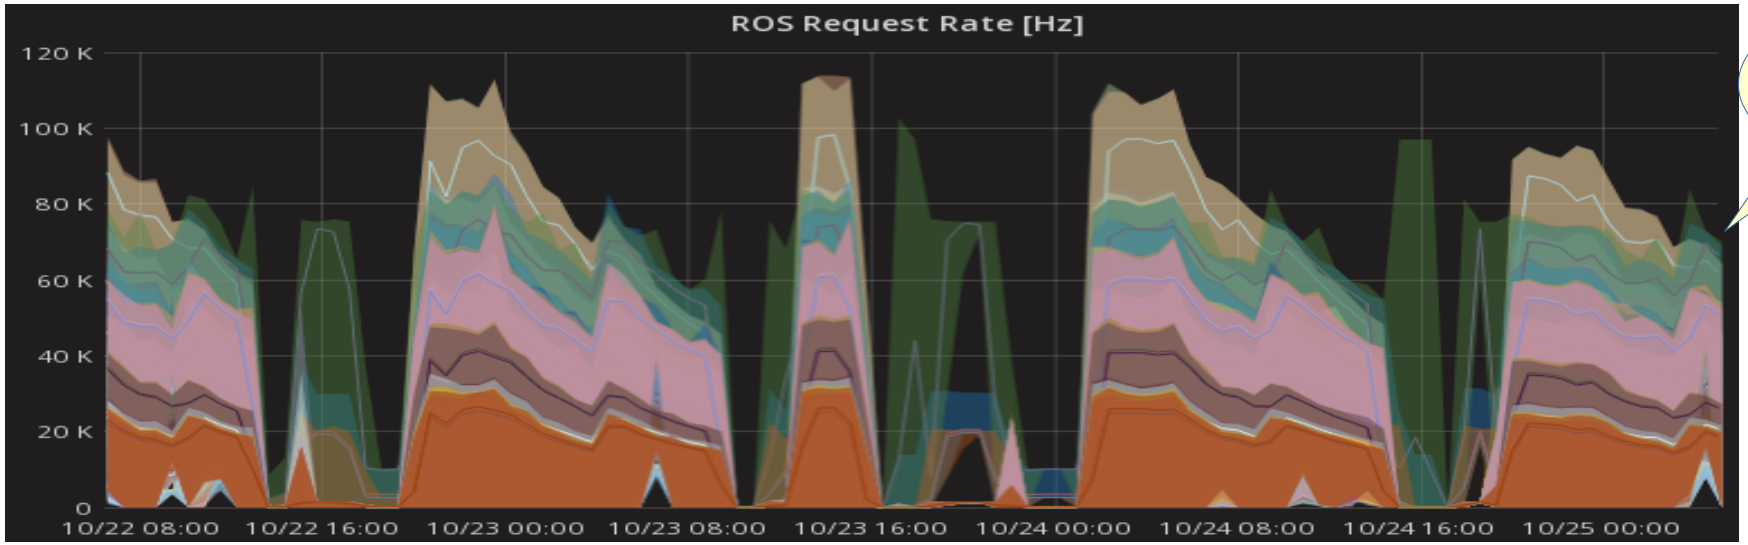

+ error bars

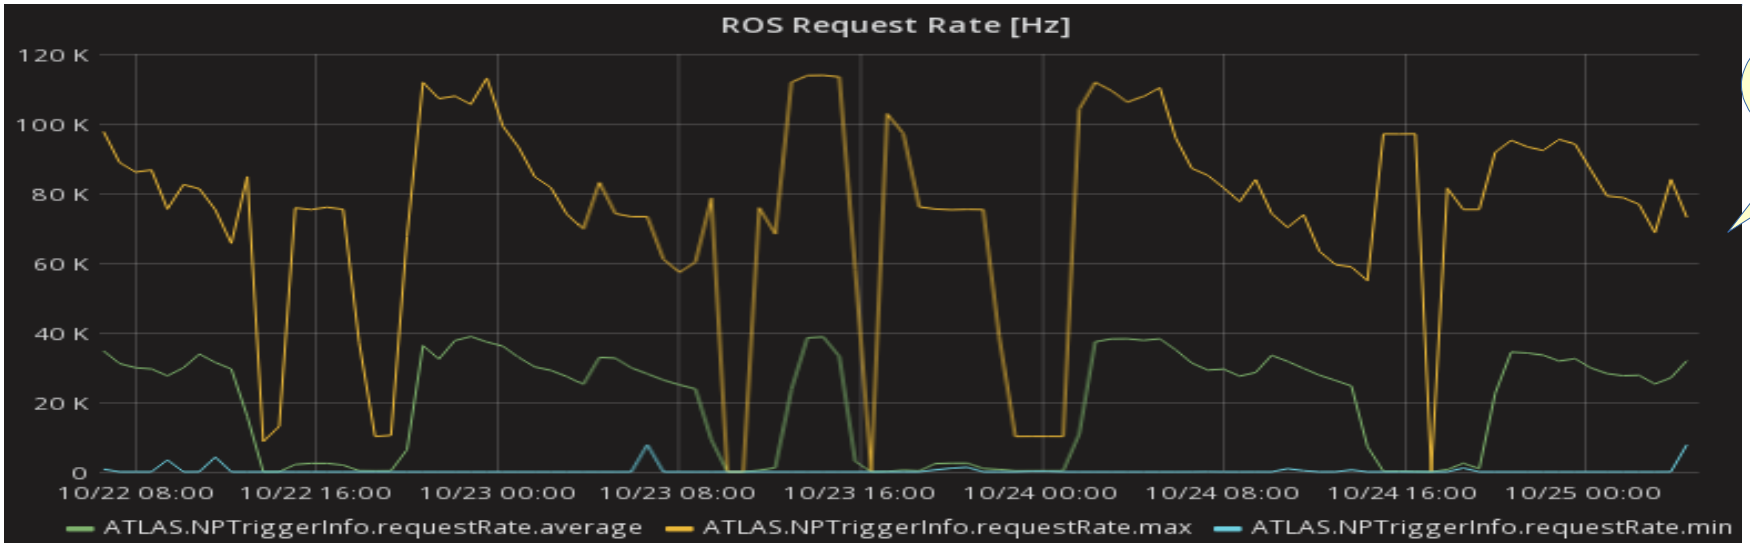

min,max, average

P-BEAST Service: Aggregation Functions (3/3)

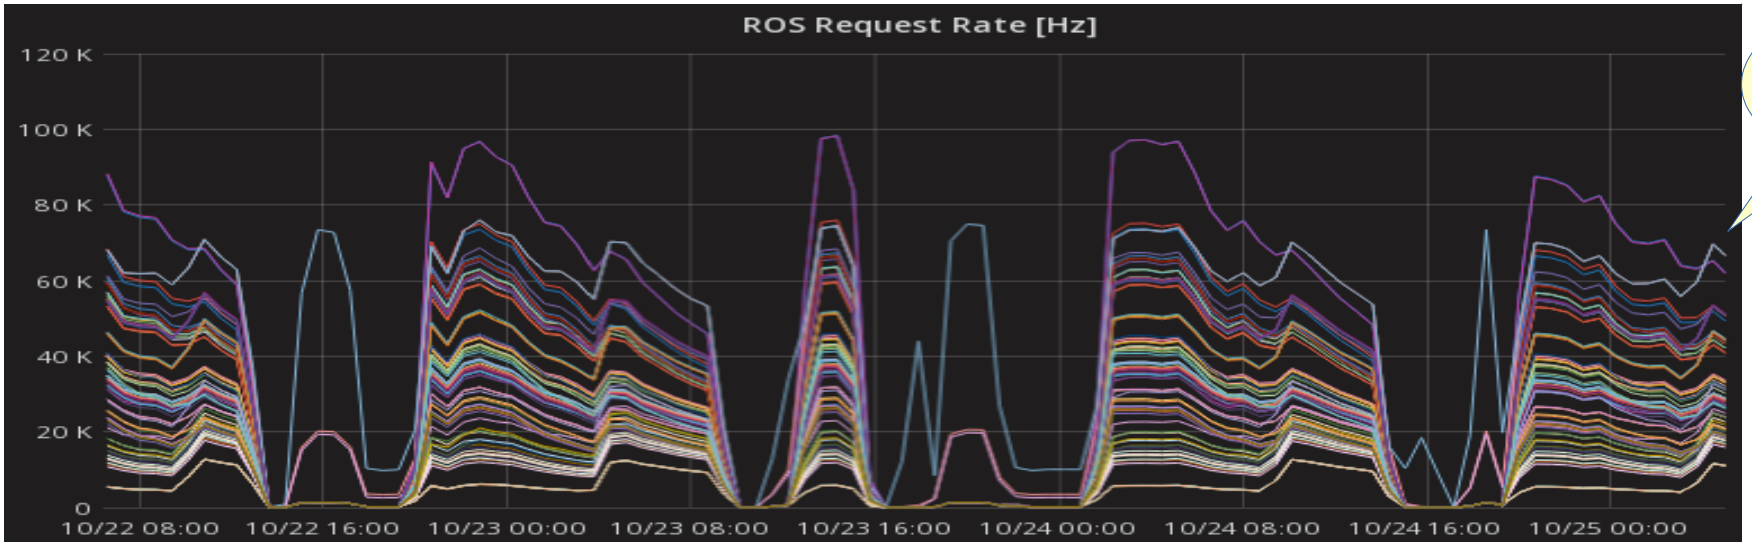

query data

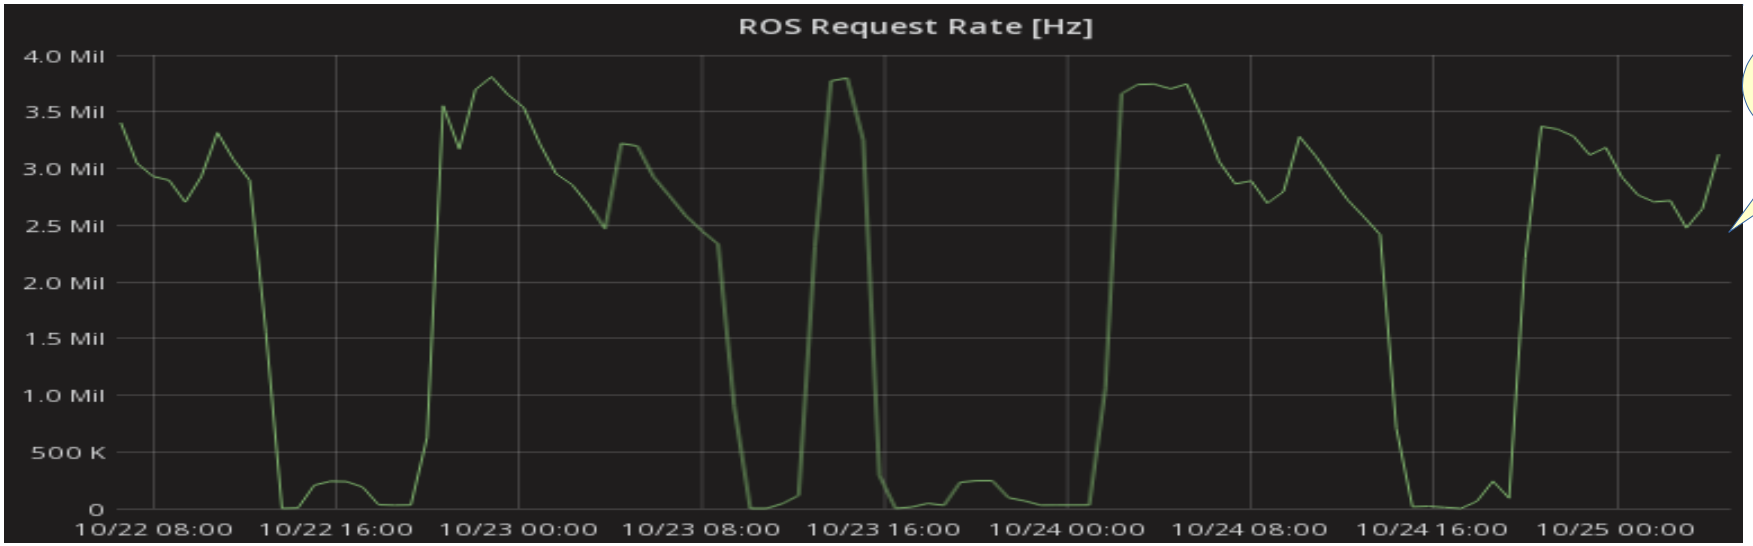

sum

P-BEAST Service: SumArray() Transformation Function

- ***SumArray()*: for every metrics of array type sum values of all indexes and store into first index**

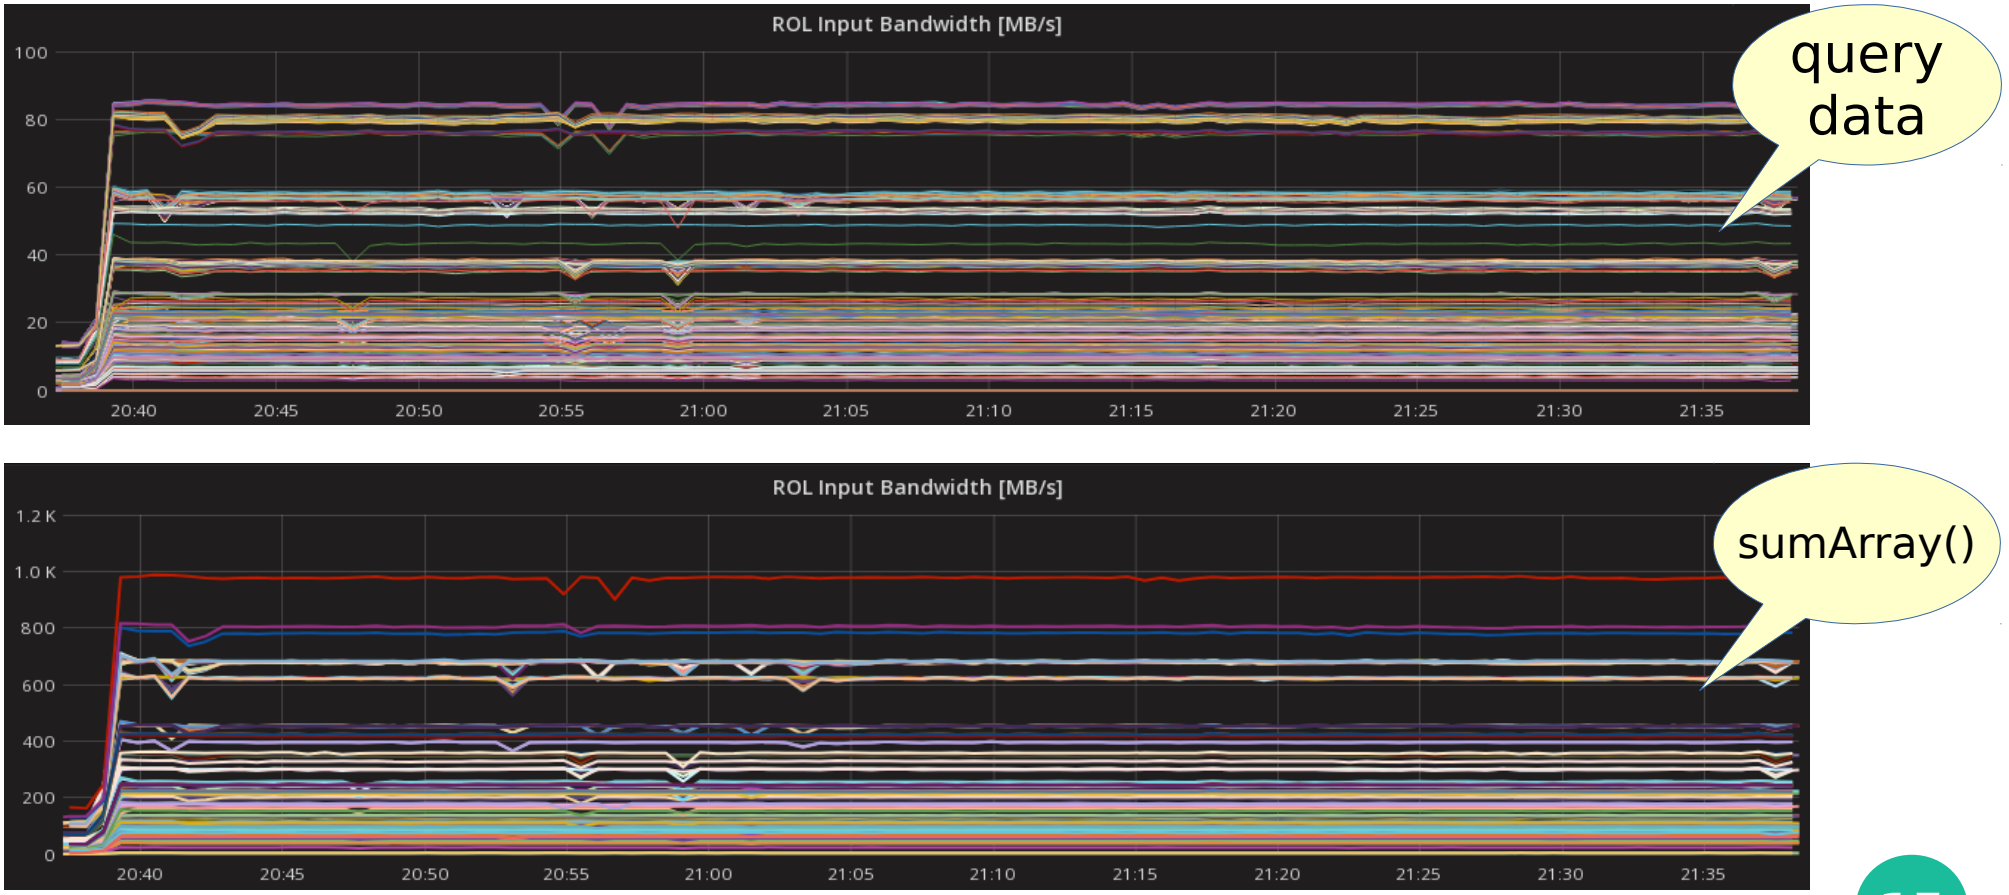

P-BEAST Service: SumArray() and aggregation sum() functions

- **Convert three dimensions into one**
  - E.g. transform ROLs throughput into DAQ system one

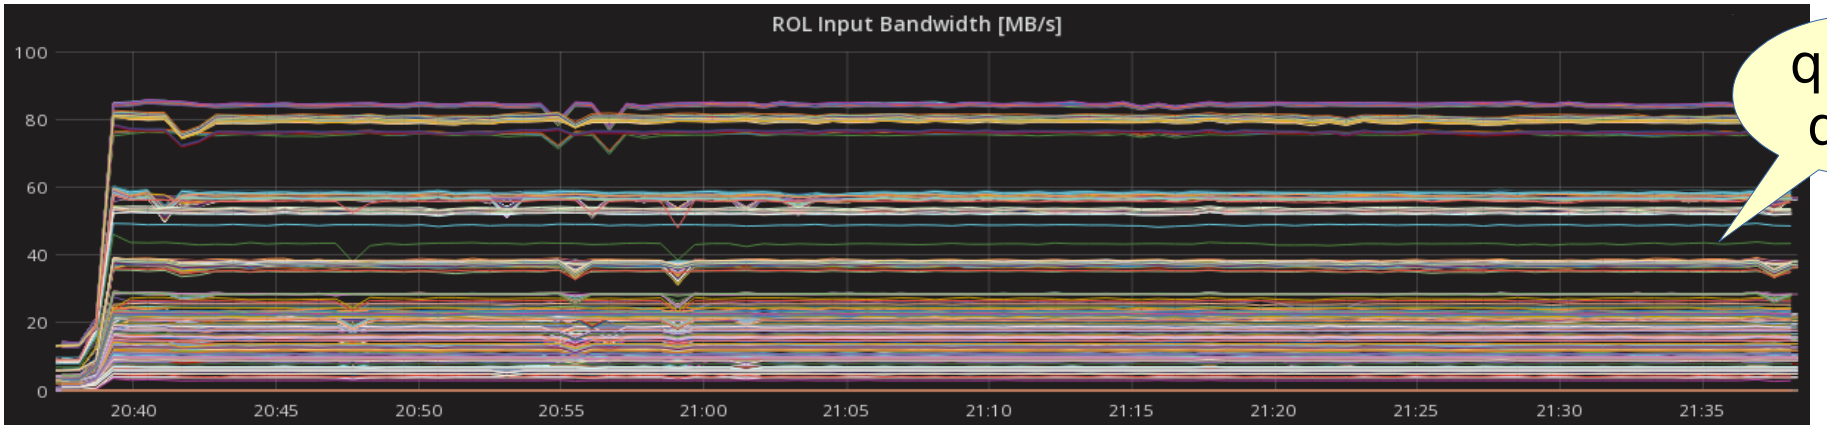

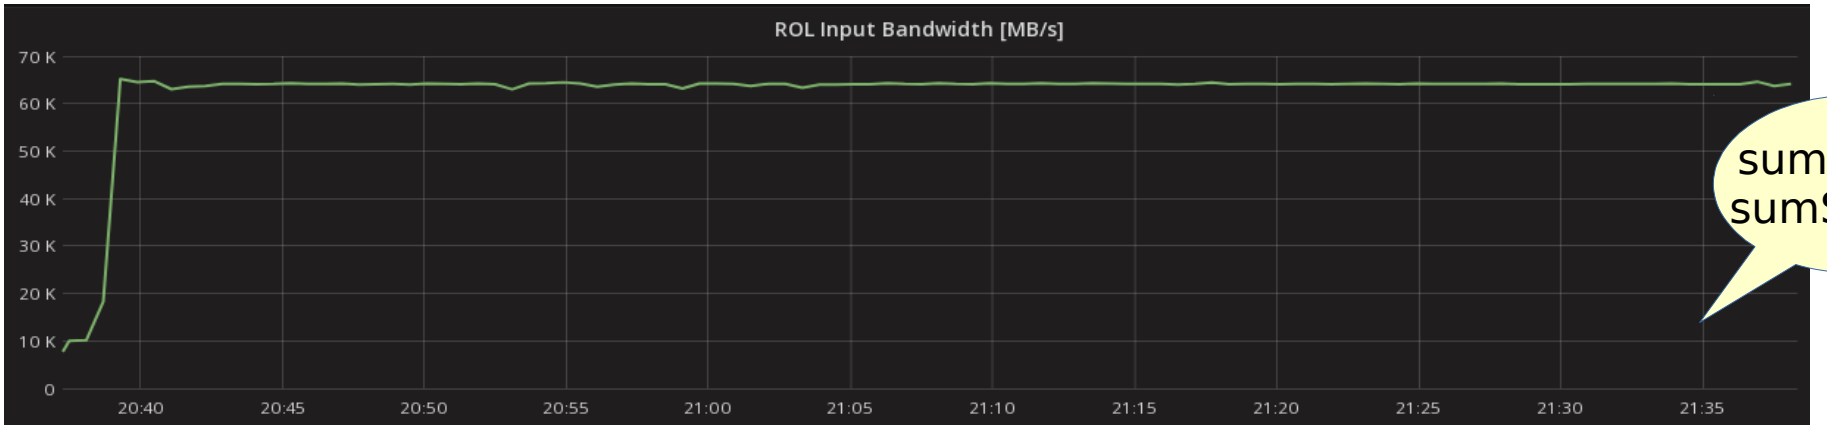

P-BEAST Alias Functions

- **Change names of data series in accordance with some rules:**
  - alias(new_name) – rename data serie
  - aliasByNode(idx[,idx2,...]) - select name segments
  - aliasByMetric() - last segment name
  - aliasByObjectID() – no partition / class / attribute names
  - aliasSub(regex,replace) – use regex substitution
- **The alias functions can be combined**
- **They allow to make names shorter and more readable, and the legends more usable**

P-BEAST Annotations

- **Grafana provides a way to mark events**
  - In our case such events could be run number and luminosity block, run control state, LHC state, ready for physics flag changes, etc.
  - Can be defined on the level of dashboard using query editor

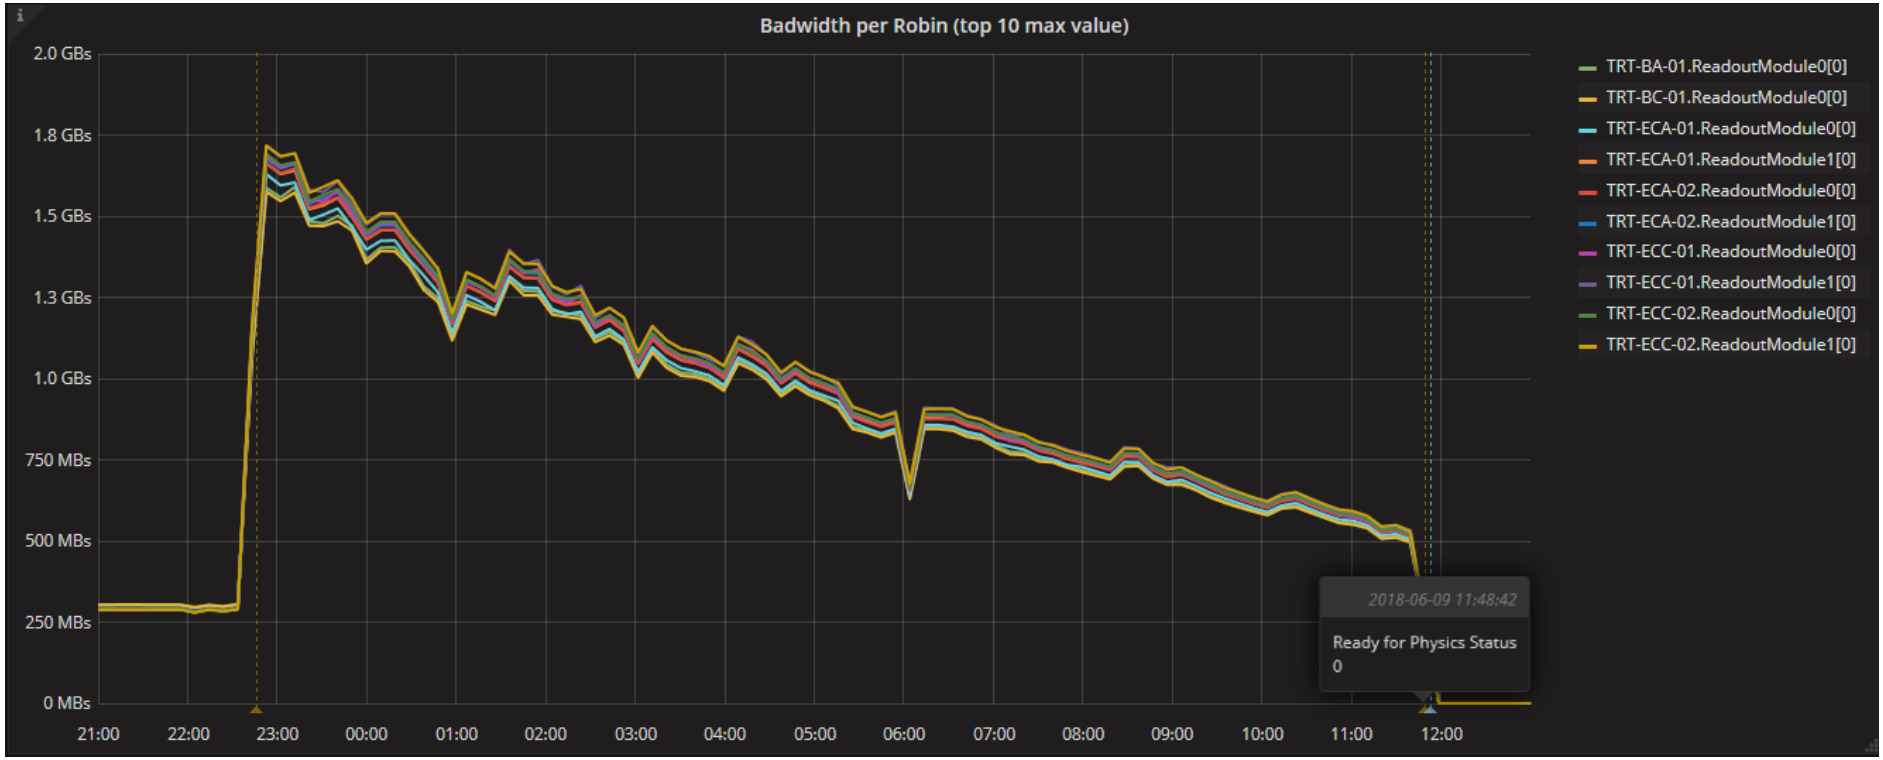

Technology Evaluation

- **We are performing technology evaluation to see possible candidates for P-BEAST service implementation or ideas how to improve it**
  - InfluxDB is under evaluation
  - ClickHouse (yandex) column-oriented DB will be next candidate

Status and Plans

- **P-BEAST is successfully used during Run 2**
  - Stable despite of development is on-going
- **During last year implemented a lot of new features and improvements. Have in mind new ones:**
  - Migration to Grafana 5 next year
  - Trends prediction and alarms would be next possible features to be implemented
- **P-BEAST could be used for Trigger Presenter GUI and FELIX monitoring**
- **User feedback and comments are welcome:**
  - See Atlas/PbeastDashboard Twiki page for dashboard user guide
  - Use ADAMATLAS Jira project for general discussions and issues