MAX IV Archiving

Dmitry Egorov

Tango HDB++

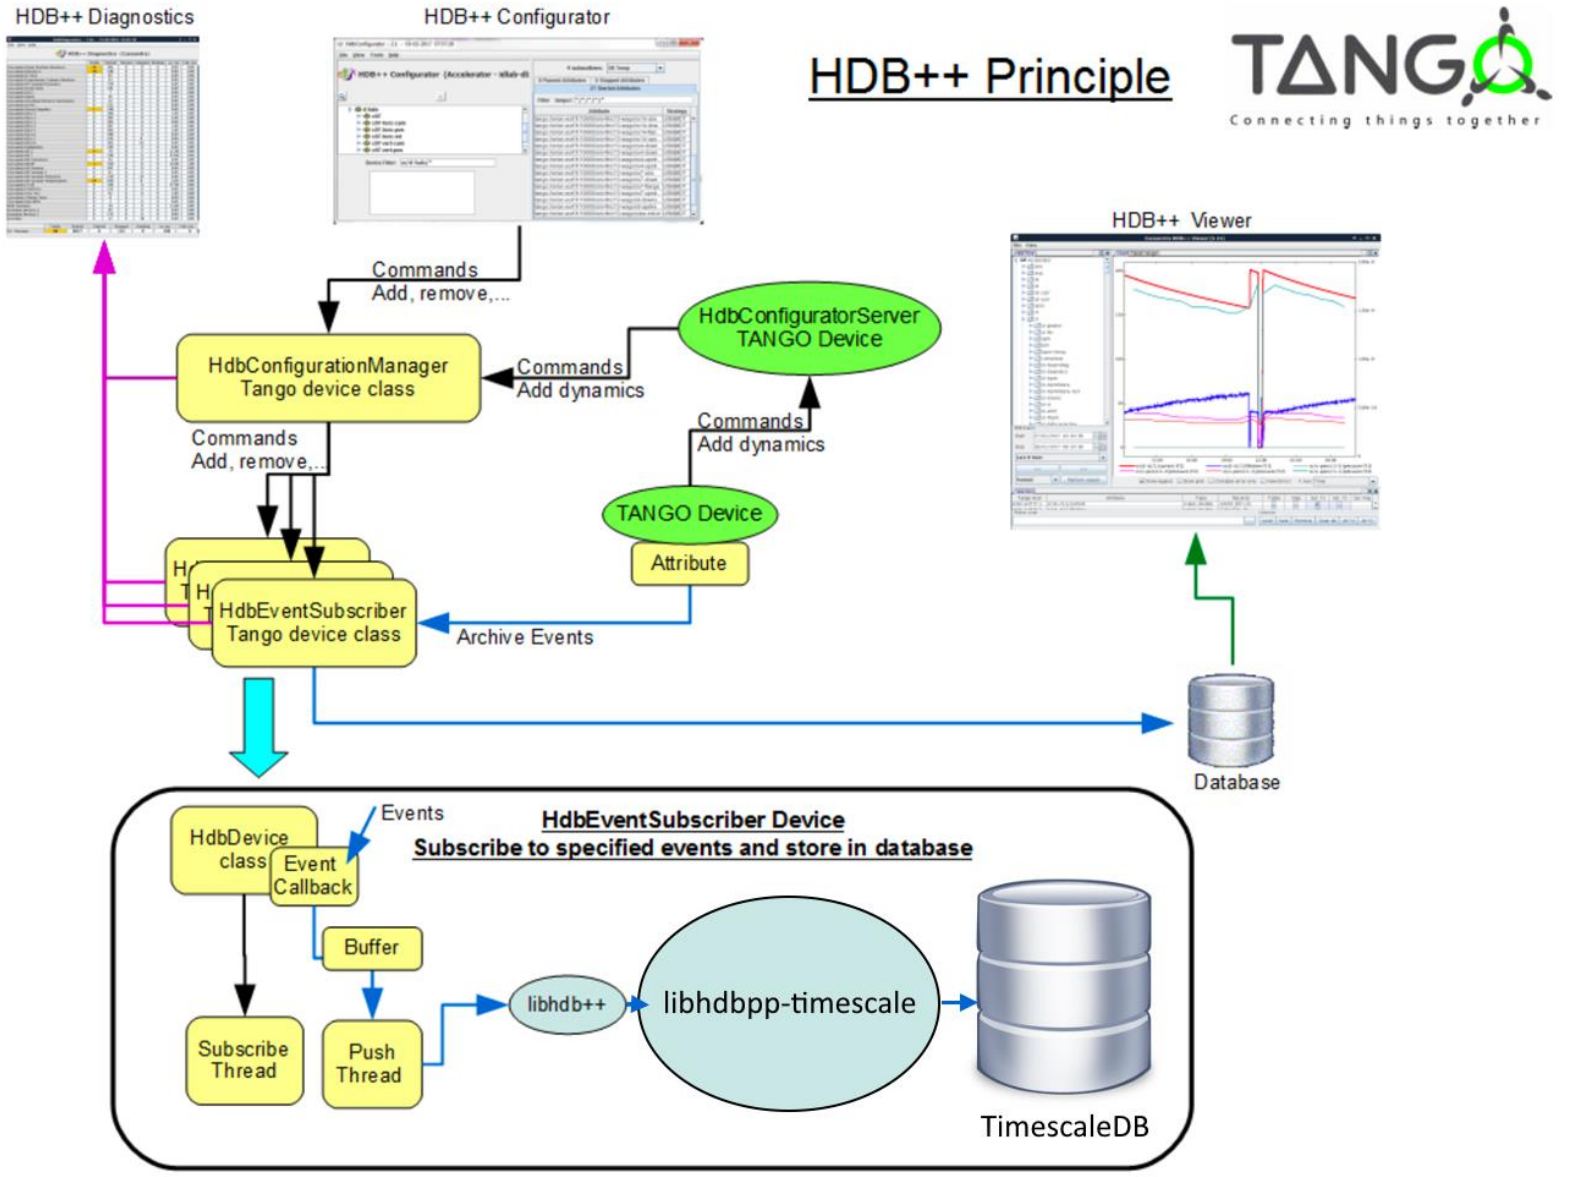

https://tango-controls.readthedocs.io/en/latest/tools-and-extensions/archiving/HDB++.html

HDB++ at MAX IV

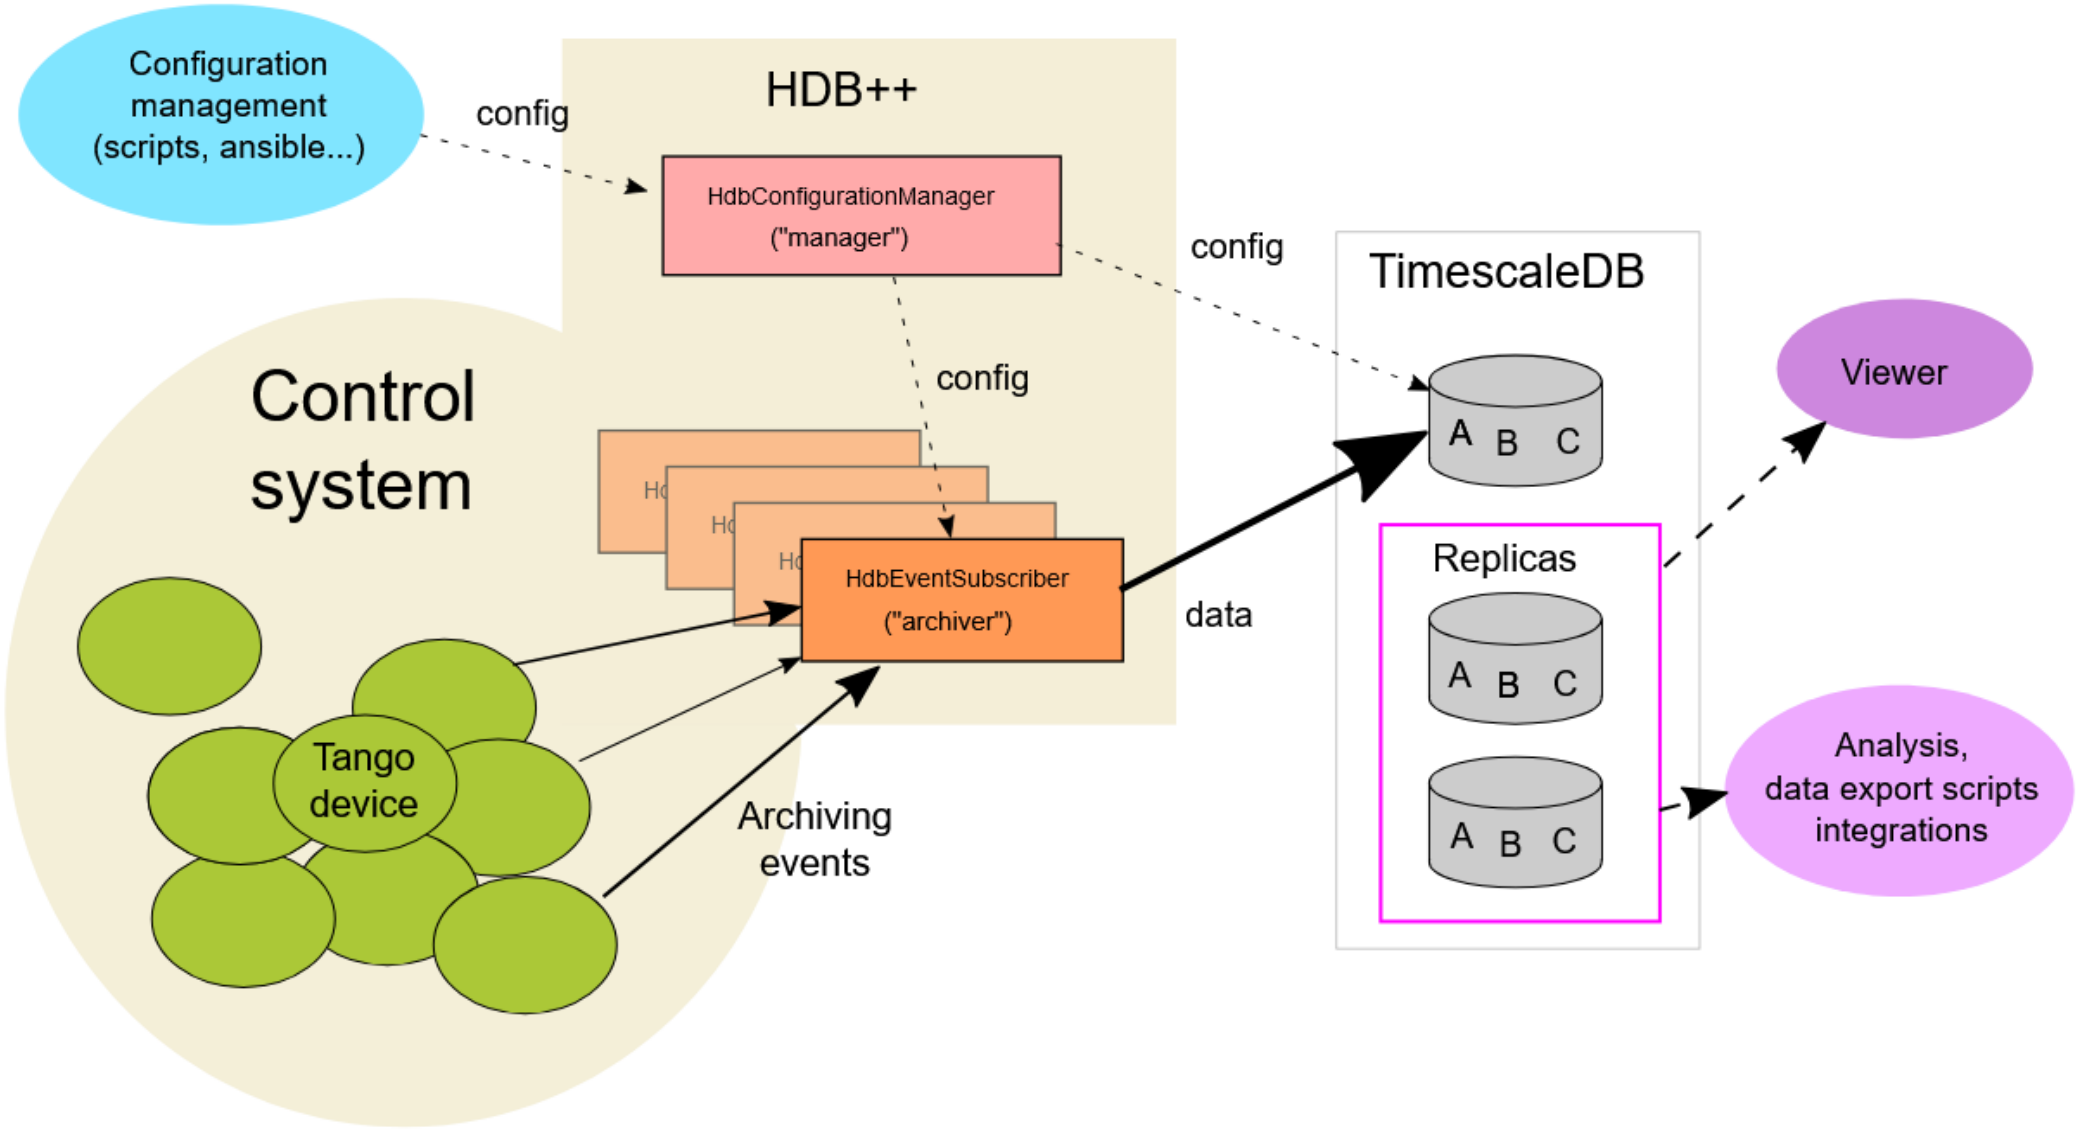

Arrow direction shows data flow
Arrow size: amount of traffic
Dotted arrow: sporadic traffic

HDB++ at MAX IV

- Running HDB++ with Cassandra back-end, since late 2016
- TimescaleDB archiving in parallel since late 2022.
- Old data migrated (apart from some “complicated” attributes)
- In late 2023 stopped archiving to Cassandra. Decommissioning…

- Configuration using “yaml2archiving”

- Interest in archiving at beamlines is growing, work on standardization

Setup

- One HDB++ setup per BL, one for accelerator (~90% of data)
- Single Postgres database cluster, 3 nodes (1 write, 2 read-only)
- Separate schema for each control system

Some statistics

~ 20000 attributes
~ 2000 events per second
~ 50 archivers across 20 control systems
~ 250 billion rows in TimescaleDB
~ 10TB disk space used

Few words about Cassandra

Advantages:

- High scalability – just add more nodes and write more data
- High write performance – efficient for low-latency writes

Disadvantages:

- Lack of in-house expertise (cluster management, CQL)
- Limited query capabilities (averaging, aggregations, etc.)
- ESRF stopped supporting Cassandra as an HDB++ backend

It has taken about 4 years to migrate from Cassandra, maintaining the old system while testing and setting up the new one.

MAX IV

…and about TimescaleDB

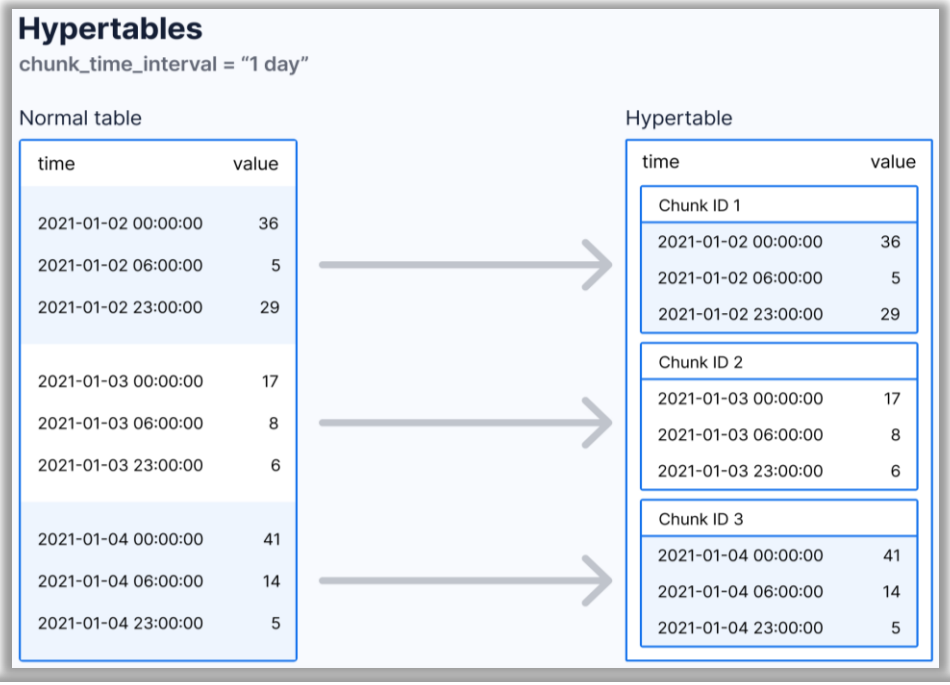

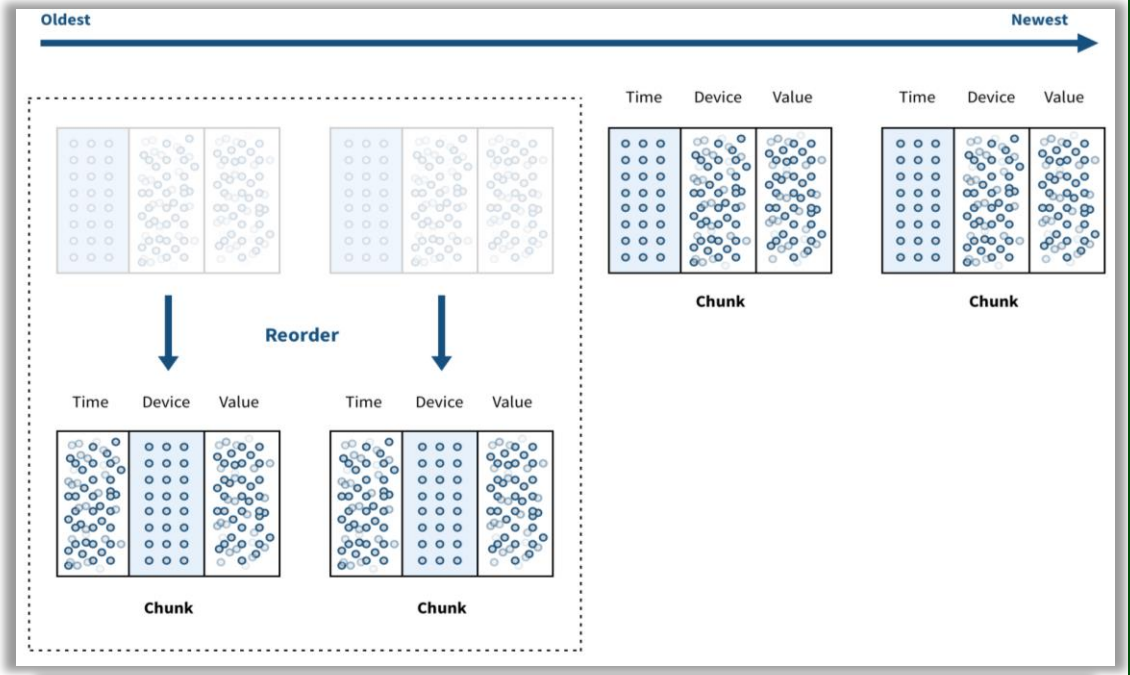

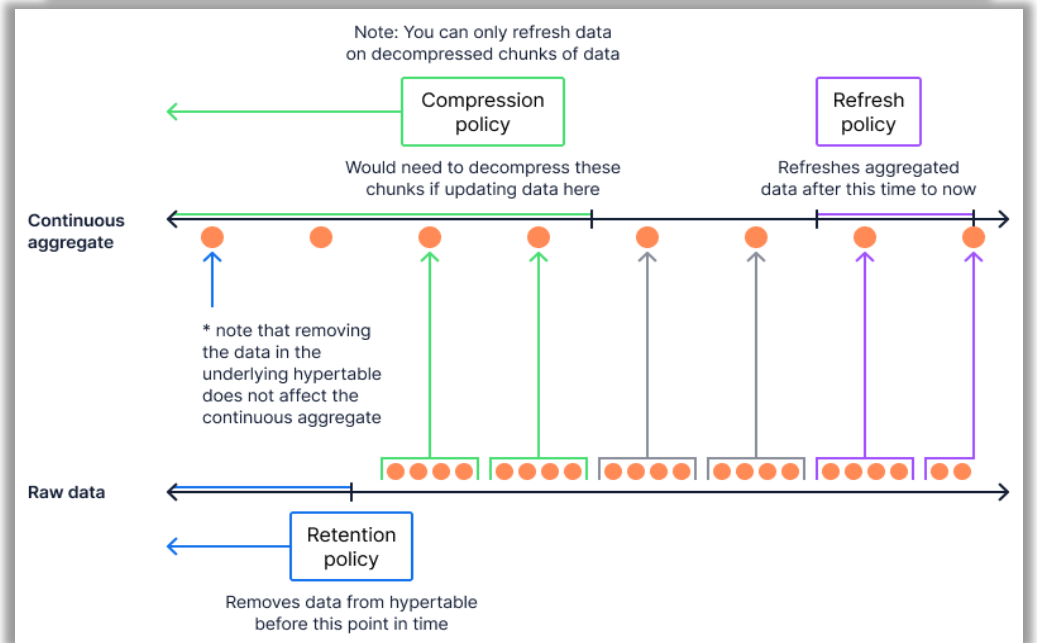

Before compression

time	device_id	cpu	energy_comsumption
12:00:02	1	88.2	0.8
12:00:02	2	300.5	0.9
12:00:01	1	88.6	0.85
12:00:01	2	299.1	0.95

After compression

time	device_id	cpu	energy_comsumption
[12:00:02, 12:00:02, 12:00:01, 12:00:01]	[1,2,1,2]	[88.2, 300.5, 88.6,299.1]	[0.8, 0.9, 0.85, 0.95]

Infrastructure deployment at MAX IV

kubernetes

See presentation by Andrii Salnikov at 15.30

Another deployment option

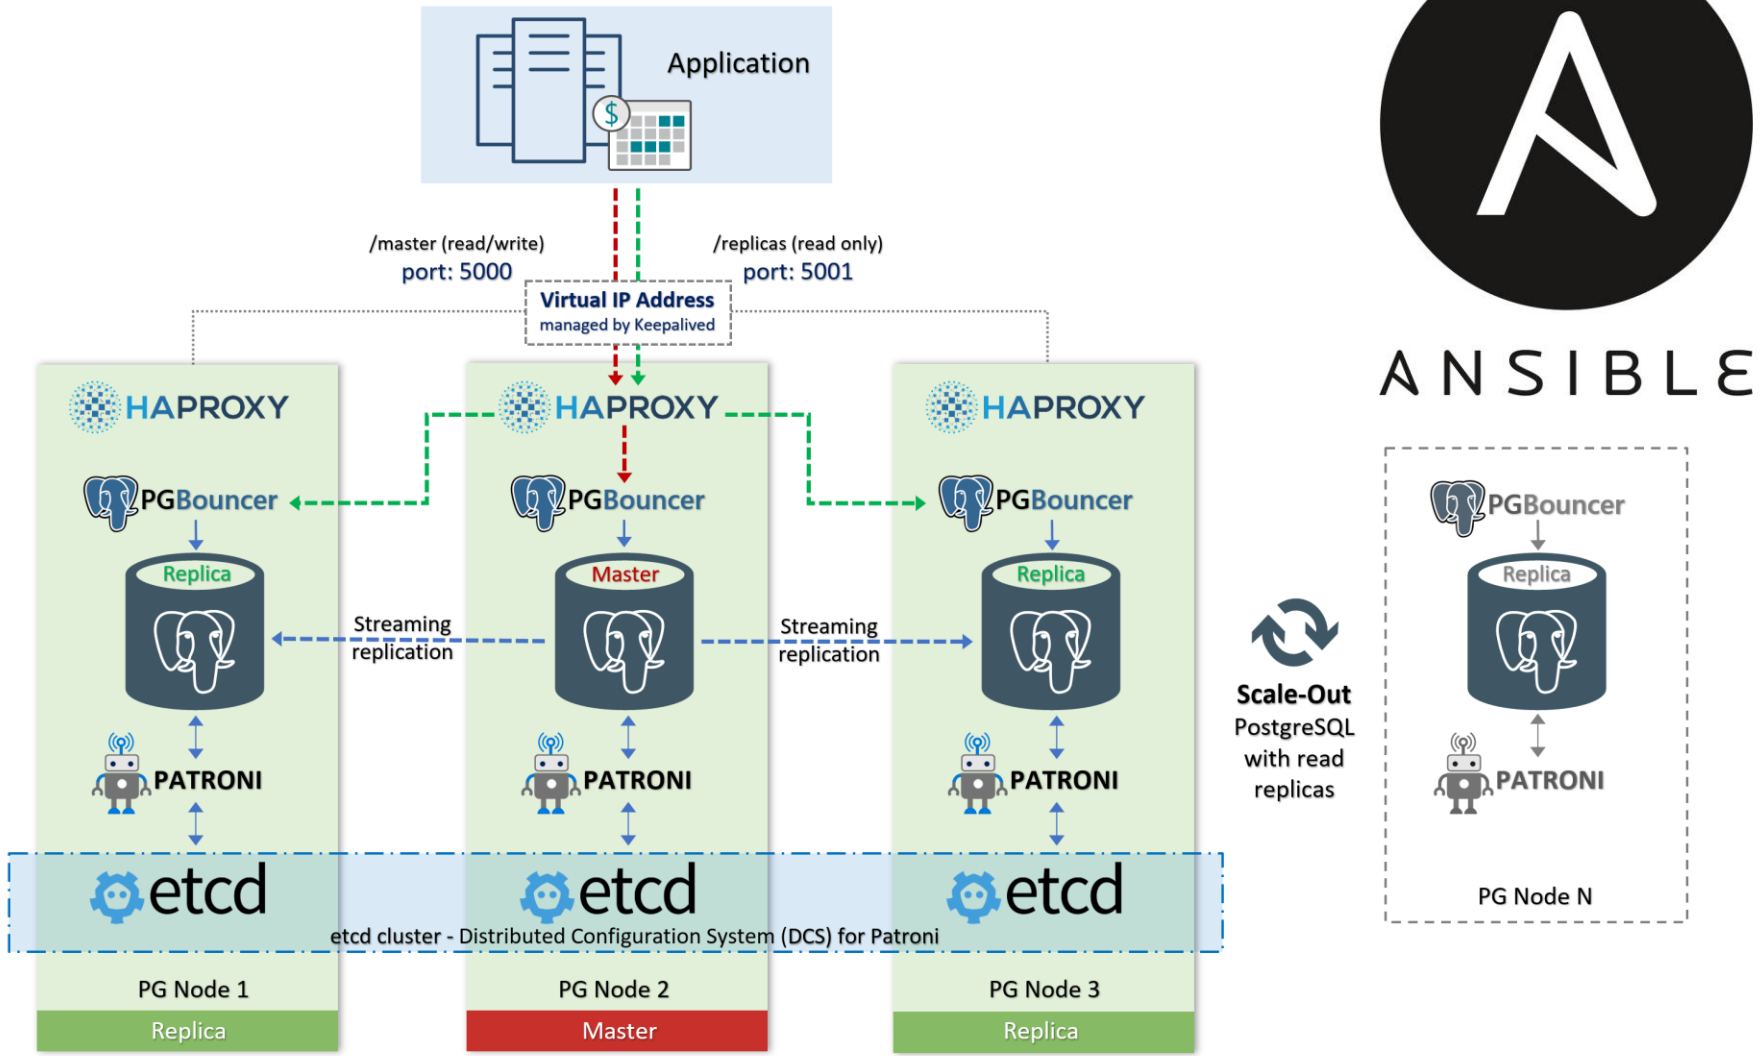

ANSIBLE

https://github.com/vitabaks/postgresql_cluster

HDB deployment at MAX IV

CONDA

ANSIBLE

AWX

Artifact (o

conda-f
HDB++ library

conda-f
HDB++ library

conda-f
Tango device

conda-f
Interface libra

conda-f
Tango device

conda-f
Tango device

conda-f
hdb++ pythor

```
tango_ds_hdbconf
  name: hdb++cm-
  conda_env_name
  conda_packages
    cpptango: "{
    pytango: "{{
    hdbpp-cm: "{
    hdbpp-es: "{
    libhdbpp-tim
  instances:
    - name: 1
      devices:
        - name:
          class:
          proper
            Arch
              -
            LibC
              -
              -
              -
              -
              -
              -
              -
```

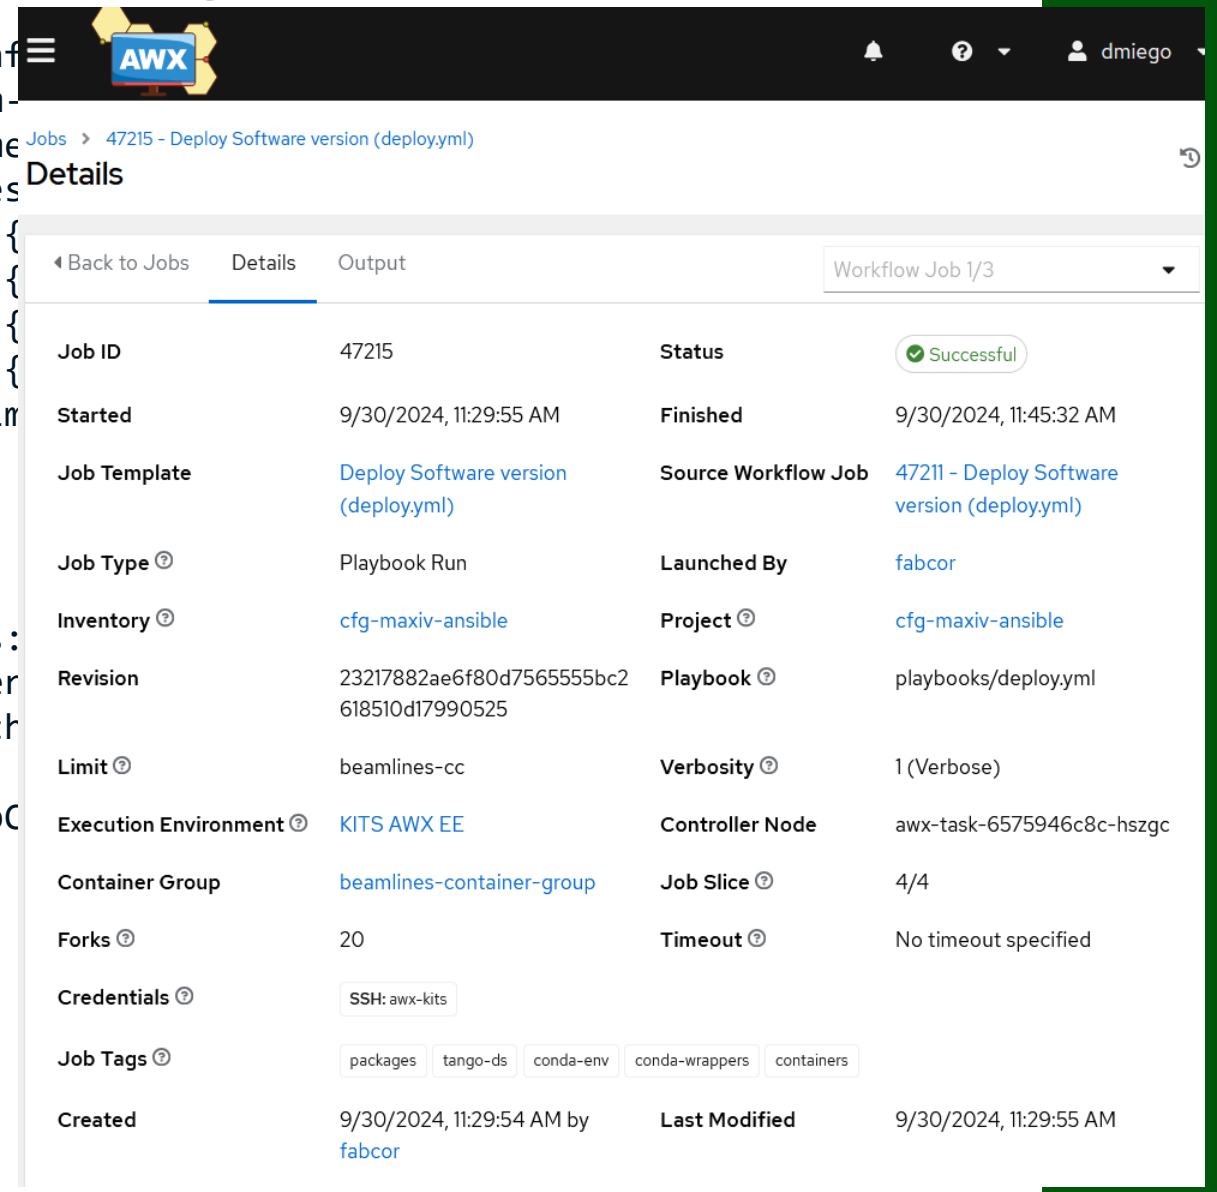

https://gitlab.com/MaxIV/lib-maxiv-dsconfig

See presentation by Benjamin Bertrand at 16.00

MAX IV

Configuration (yaml2archiving)

Yaml2archiving aims to ease the configuration of the HDB++ Archiving by using YAML format files in order to add/remove/update archived attributes. Each YAML file corresponds to one HdbEventSubscriber device and can include the state of the archived attributes with their settings.

```
db: b-v-dummymax-csdb-0.maxiv.lu.se:10000
manager: SYS/HDBPP/MANAGER-01
archiver: SYS/HDBPP/ARCHIVER-01

configuration:
  - class: TangoTest
    filtering:
      device: sys/tg_test/1
    attributes:
      double_scalar:
        polling_period: 3000
        archive_period: 10000
      short_scalar:
        polling_period: 3000
        archive_period: 10000
```

```
db: g-v-csdb-0.maxiv.lu.se:10000
manager: sys/hdbpp/manager-01-timescale
archiver: i/mag/archiver-01-timescale
defaults:
  - "../mag_defaults.yaml"

configuration:
  - class: MagnetCircuit
    filtering:
      device:
        - "i-.*"
    attributes:
      mainfieldcomponent:
        archive_rel_change: 1.0
        archive_period: 60000
        polling_period: 1000
```

https://gitlab.com/tango-controls/hdbpp/yaml2archiving

Thanks to Johan Forsberg

Monitoring (ArchWizard)

"ArchWizard 🧙 " is a simple web interface to the HDB++ archiving system.
It is intended for monitoring and troubleshooting, not maintenance such as adding and removing attributes.
It's built using Python, FastAPI, Jinja2 and PyTango.

Attribute: tango://g-v-csdb-0.maxiv.lu.se:10000/r3-b080603/pss/plc-01/ais_r3_a102011cab20_pss_bmm04_i4

Archiver: tango://g-v-csdb-0.maxiv.lu.se:10000/g/pss/archiver-01-timescale

Attribute

Label:	AIS_R3_A102011CAB20_PSS_BMM04_I4
Unit:	
Data type:	DevShort
Data format:	SCALAR

Archiving status

Event status:	The polling (necessary to send events) for the attribute ais_r3_a102011cab20_pss_bmm04_i4 is not started
Archiving:	Started
Event OK counter:	0 - YYYY-MM-DD HH:MM:SS.UUUUUU
Event NOK counter:	0 - YYYY-MM-DD HH:MM:SS.UUUUUU
DB ERRORS counter:	0 - YYYY-MM-DD HH:MM:SS.UUUUUU
Storing time AVG:	0.000000s
Processing time AVG:	0.000000s

Device

Name:	R3-B080603/PSS/PLC-01
Device class:	AllenBradleyEIP
Server:	AllenBradleyEIP/R3-PSS
Host:	g-v-ec-17.maxiv.lu.se
Last started:	2nd September 2024 at 11:34:47

Archive events

Absolute change:	1
Relative change:	Not specified
Periodic:	60000

Rendered at 2024-10-03 11:42:16.015459

https://gitlab.com/tango-controls/hdbpp/archwizard

Monitoring (hdbpp-exporter)

Prometheus HDB++ exporter
An exporter is a small webserver that publishes metrics data in a text format readable by the Prometheus monitoring service.
This exporter connects to all HdbEventSubscriber devices and exports several diagnostics attributes, such as AttributeNumber, AttributeNokNumber etc.

Example output:

```
# HELP hdbpp_eventsubscriber_AttributeNokNumber Attribute 'AttributeNokNumber' on
HdbEventSubscriber
# TYPE hdbpp_eventsubscriber_AttributeNokNumber gauge
hdbpp_eventsubscriber_AttributeNokNumber{device="SYS/HDBPP/ARCHIVER-01"} 10.0
# HELP hdbpp_eventsubscriber_AttributeNumber Attribute 'AttributeNumber' on
HdbEventSubscriber
# TYPE hdbpp_eventsubscriber_AttributeNumber gauge
hdbpp_eventsubscriber_AttributeNumber{device="SYS/HDBPP/ARCHIVER-01"} 451.0
# HELP hdbpp_eventsubscriber_error_periodic_event_timeout Attribute
'AttributePendingNumber' on HdbEventSubscriber
# TYPE hdbpp_eventsubscriber_error_periodic_event_timeout gauge
hdbpp_eventsubscriber_error_periodic_event_timeout{device="SYS/HDBPP/ARCHIVER-01"}
10.0
```

Monitoring (hdbpp-exporter)

Prometheus

Grafana

MAX IV

Data viewer (Archviewer)

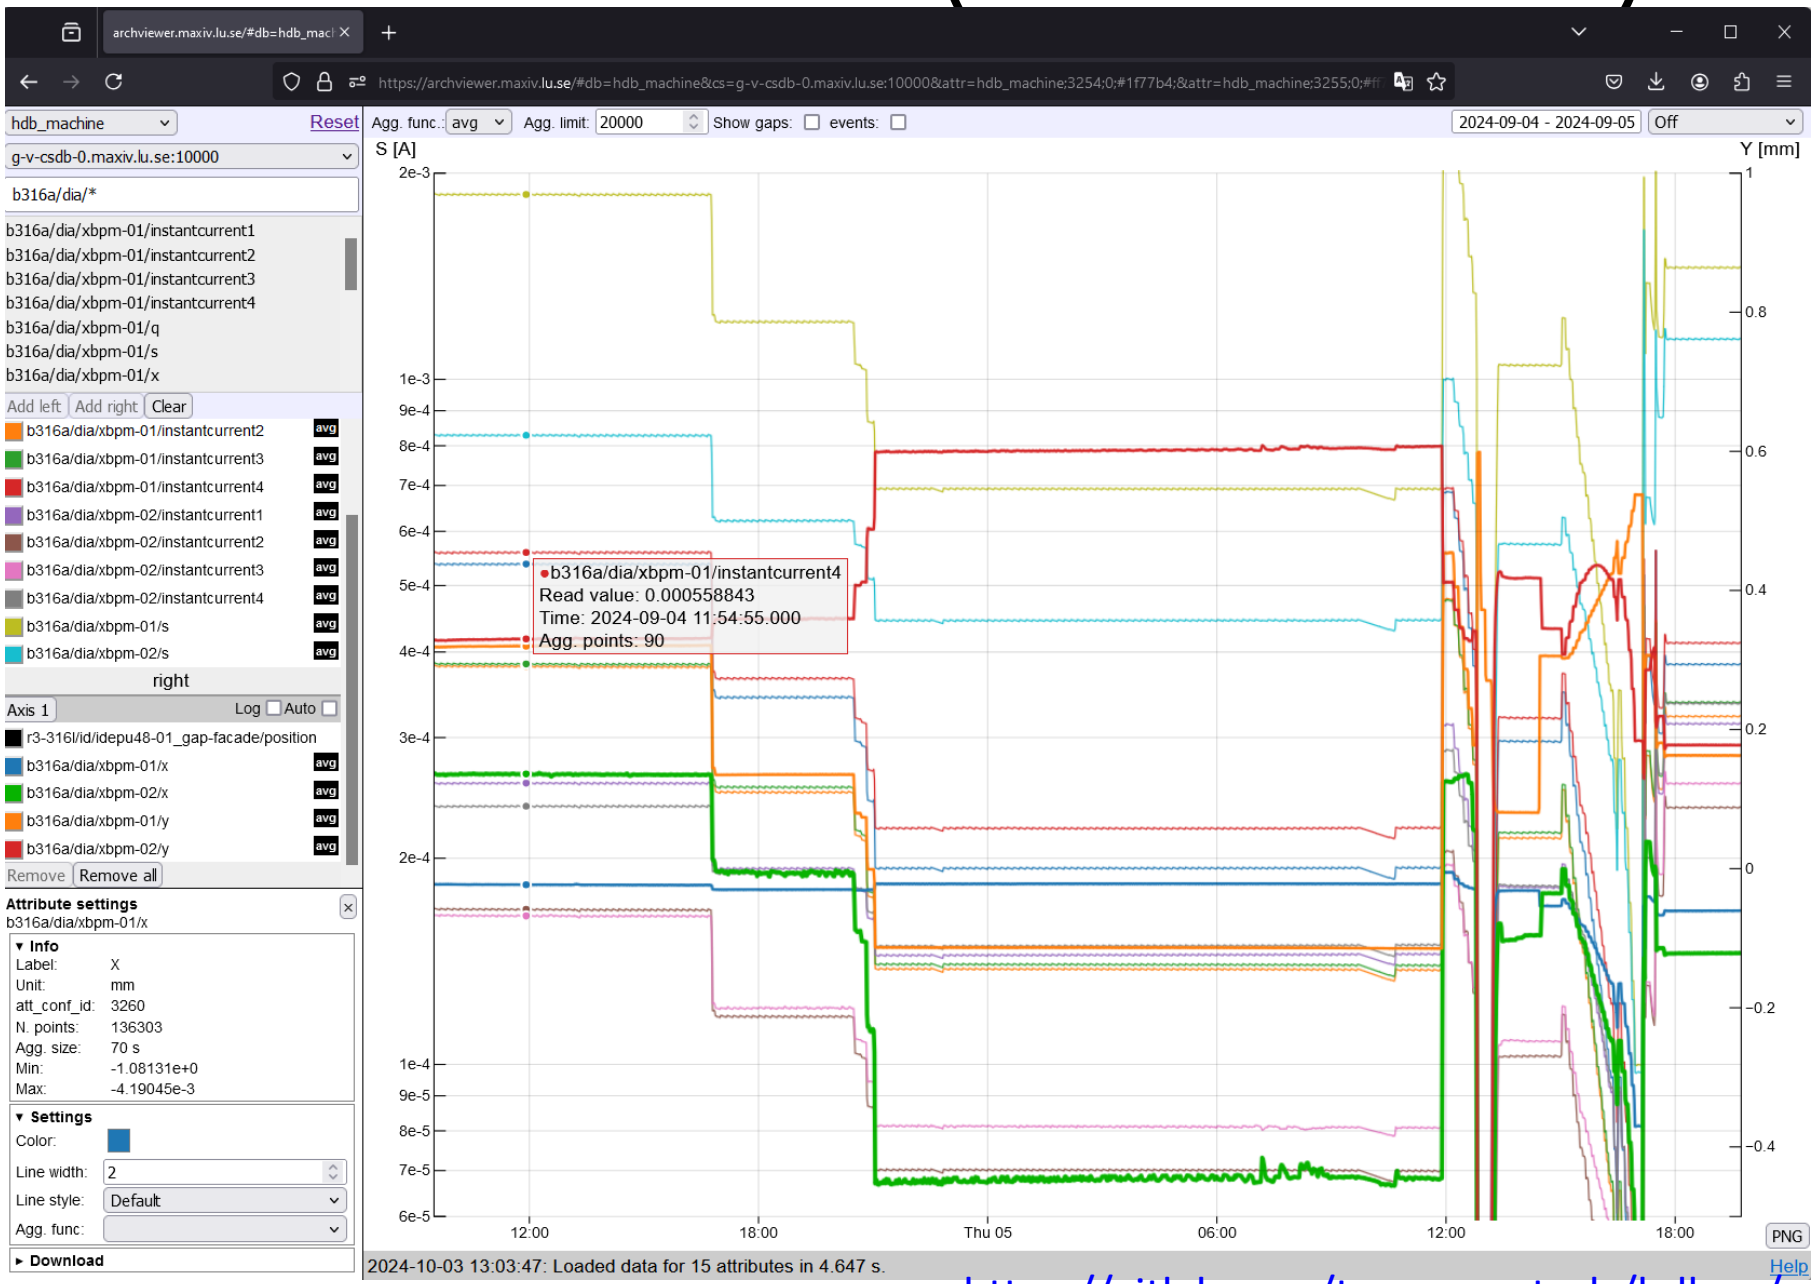

https://gitlab.com/tango-controls/hdbpp/archviewer

Data view with Grafana

```
SELECT
  data_time AS "time",
  value_r as "Beam current"
FROM att_scalar_devdouble
WHERE
  $__timeFilter(data_time)
  and
  att_conf_id = (
  att_name=tango:/
319s2/dia/dcct-01
ORDER BY 1
```

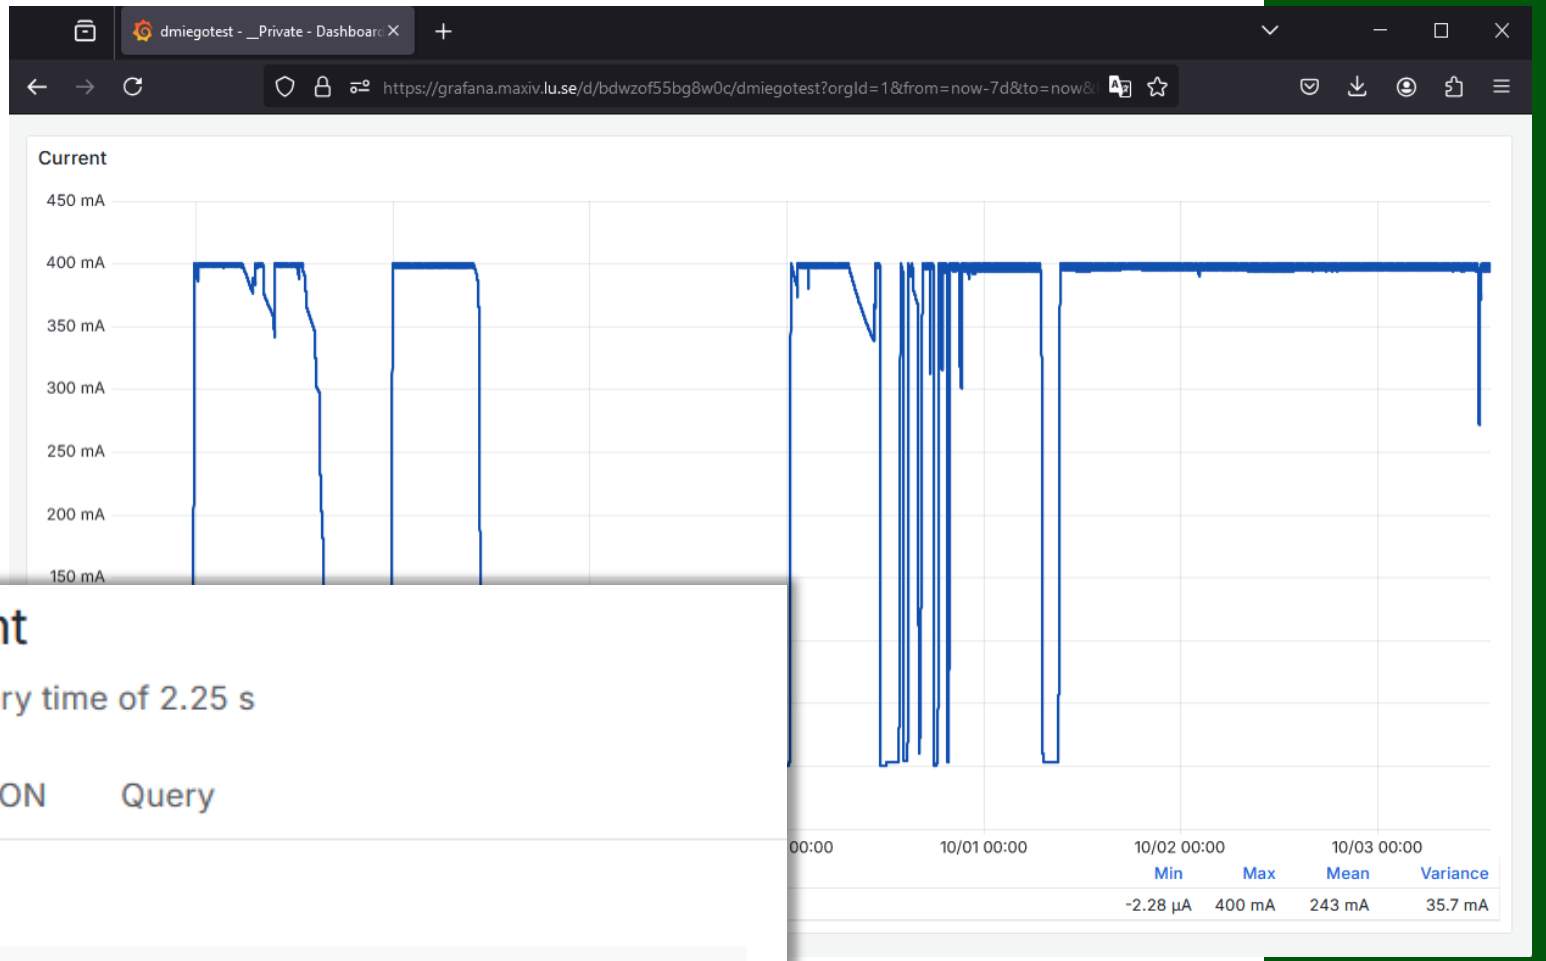

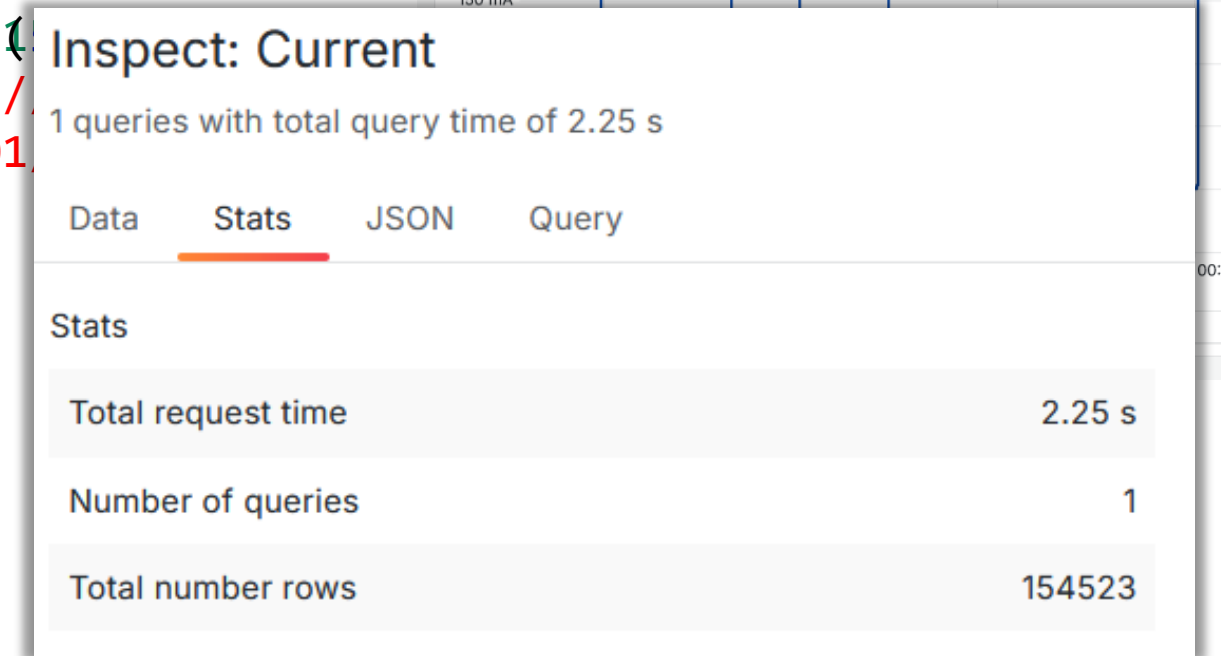

Downsampling (Timebucket)

```
SELECT
  time_bucket('1 min', data_time) AS time,
  AVG(value_r) as "Beam current"
FROM
  att_scalar_devdouble
WHERE
  $__timeFilter(data_time)
  and
  att_conf_id = (select att_conf_id from att_conf where
att_name='tango://g-v-csdb-0.maxiv.lu.se:10000/r3-319s2/dia/dcct-
01/current')
GROUP BY
  time
ORDER BY
  time;
```

https://www.timescale.com/blog/simplified-time-series-analytics-using-the-time_bucket-function/

Downsampling (Timebucket)

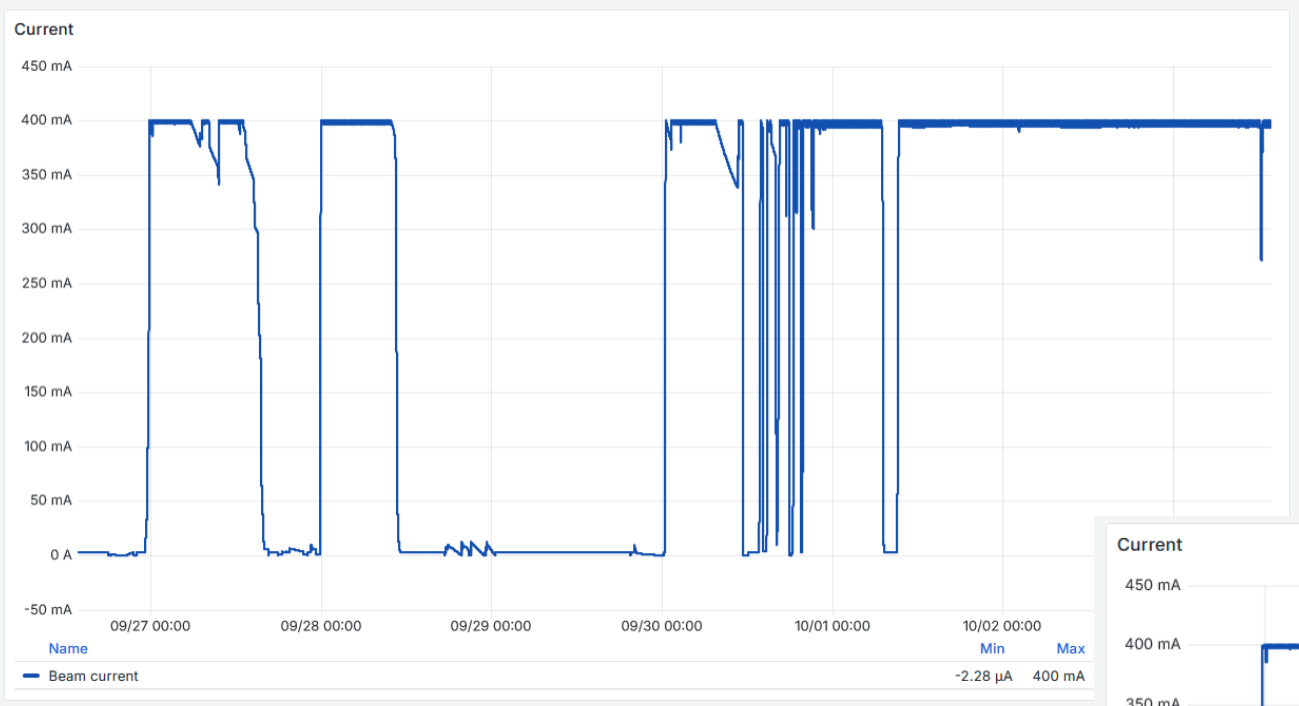

Inspect: Current

1 queries with total query time of 822 ms

Data | Stats | JSON | Query

Stats

Total request time	822 ms
Number of queries	1
Total number rows	10080

Inspect: Current

1 queries with total query time of 2.25 s

Data | Stats | JSON | Query

Stats

Total request time	2.25 s
Number of queries	1
Total number rows	154523

Downsampling (LTTB)

```
SELECT
  time AS "time",
  value AS "Beam current"
FROM unnest((
select lttb(data_time, value_r, (($__to - $__from) / $__interval_ms)::int)
FROM att_scalar_devdouble
WHERE
  $__timeFilter(data_time)
  and
  att_conf_id = (select att_conf_id from att_conf where
att_name='tango://g-v-csdb-0.maxiv.lu.se:10000/r3-319s2/dia/dcct-
01/current')
))
  ORDER BY 1;
```

https://www.timescale.com/blog/slow-grafana-performance-learn-how-to-fix-it-using-downsampling/

https://skemman.is/bitstream/1946/15343/3/SS_MSthesis.pdf

Downsampling (LTTB)

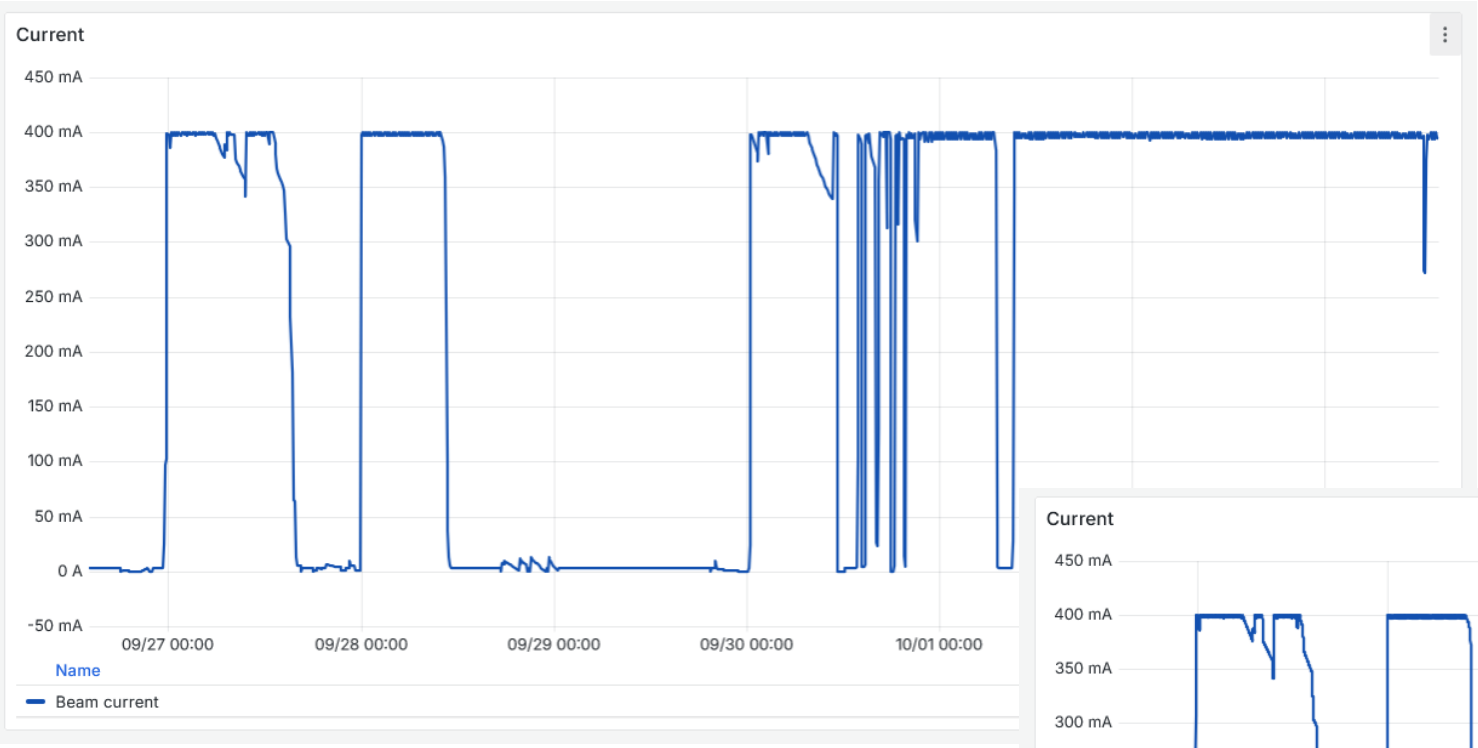

Inspect: Current

1 queries with total query time of 822 ms

Data | Stats | JSON | Query

Stats

Total request time	822 ms
Number of queries	1
Total number rows	10080

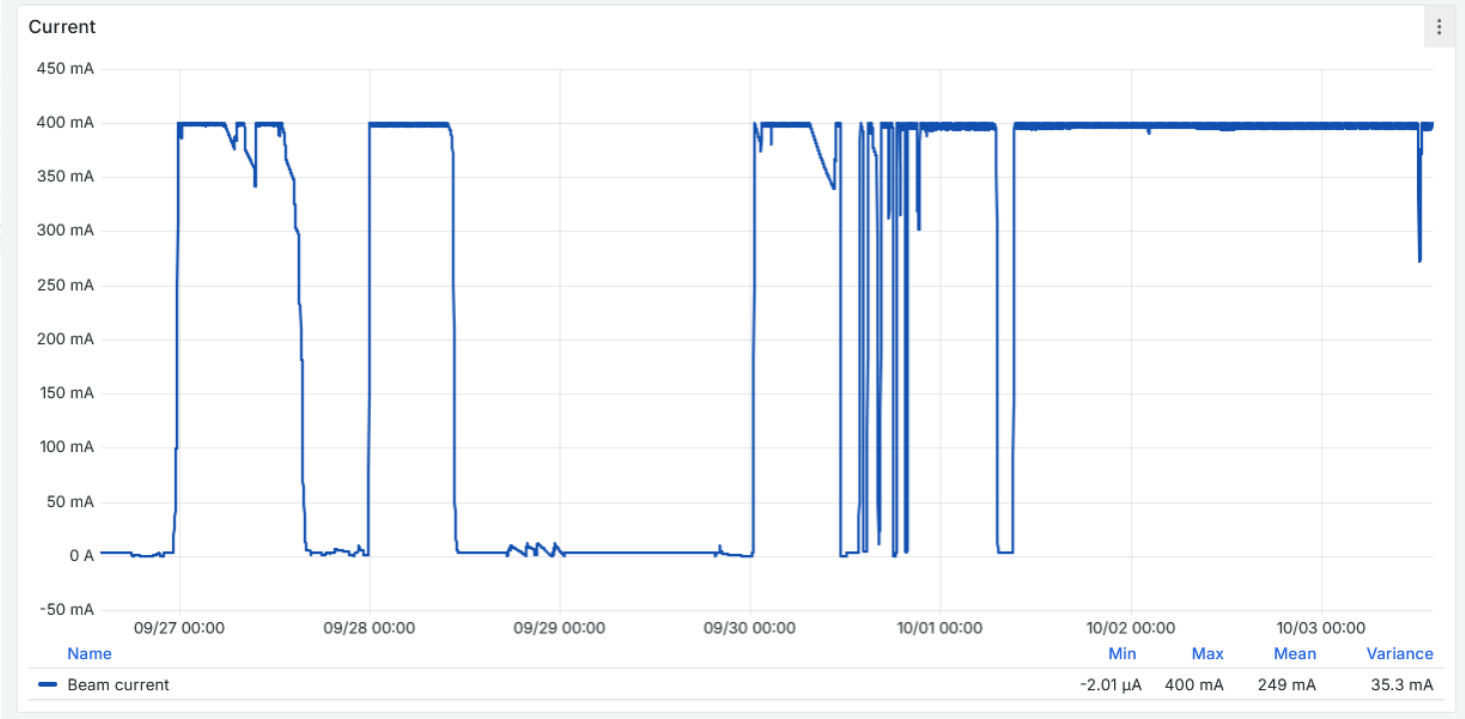

Inspect: Current

1 queries with total query time of 93 ms

Data | Stats | JSON | Query

Stats

Total request time	93 ms
Number of queries	1
Total number rows	2016

Downsampling trickiness

Downsampling (Continuous aggregations)

Not fully implemented at MAX IV (yet), example from another facility

```
WITH conf_id AS (
  SELECT att_conf_id
  FROM att_conf
  WHERE att_name='tango://bmn-sc-tangodb.he.jinr.ru:10000/fsd/plane0_top/adcsc_8_temp/temp_4'
),
varname AS (
    SELECT 'P0-4-0/22-B' AS varname
)
SELECT
  data_time AS "time",
  avg_r AS varname
FROM att_scalar_devdouble_1hour
WHERE $__timeTo()::timestamp - $__timeFrom()::timestamp >= '12 hours'::interval
  AND $__timeFilter(data_time)
  AND att_conf_id = (SELECT att_conf_id FROM conf_id)
UNION ALL

.....

SELECT
  data_time AS "time",
  value_r AS varname
FROM att_scalar_devdouble
WHERE $__timeTo()::timestamp - $__timeFrom()::timestamp < '30 minutes'::interval
  AND $__timeFilter(data_time)
  AND att_conf_id = (SELECT att_conf_id FROM conf_id)
ORDER BY 1;
```

Downsampling (Continuous aggregations)

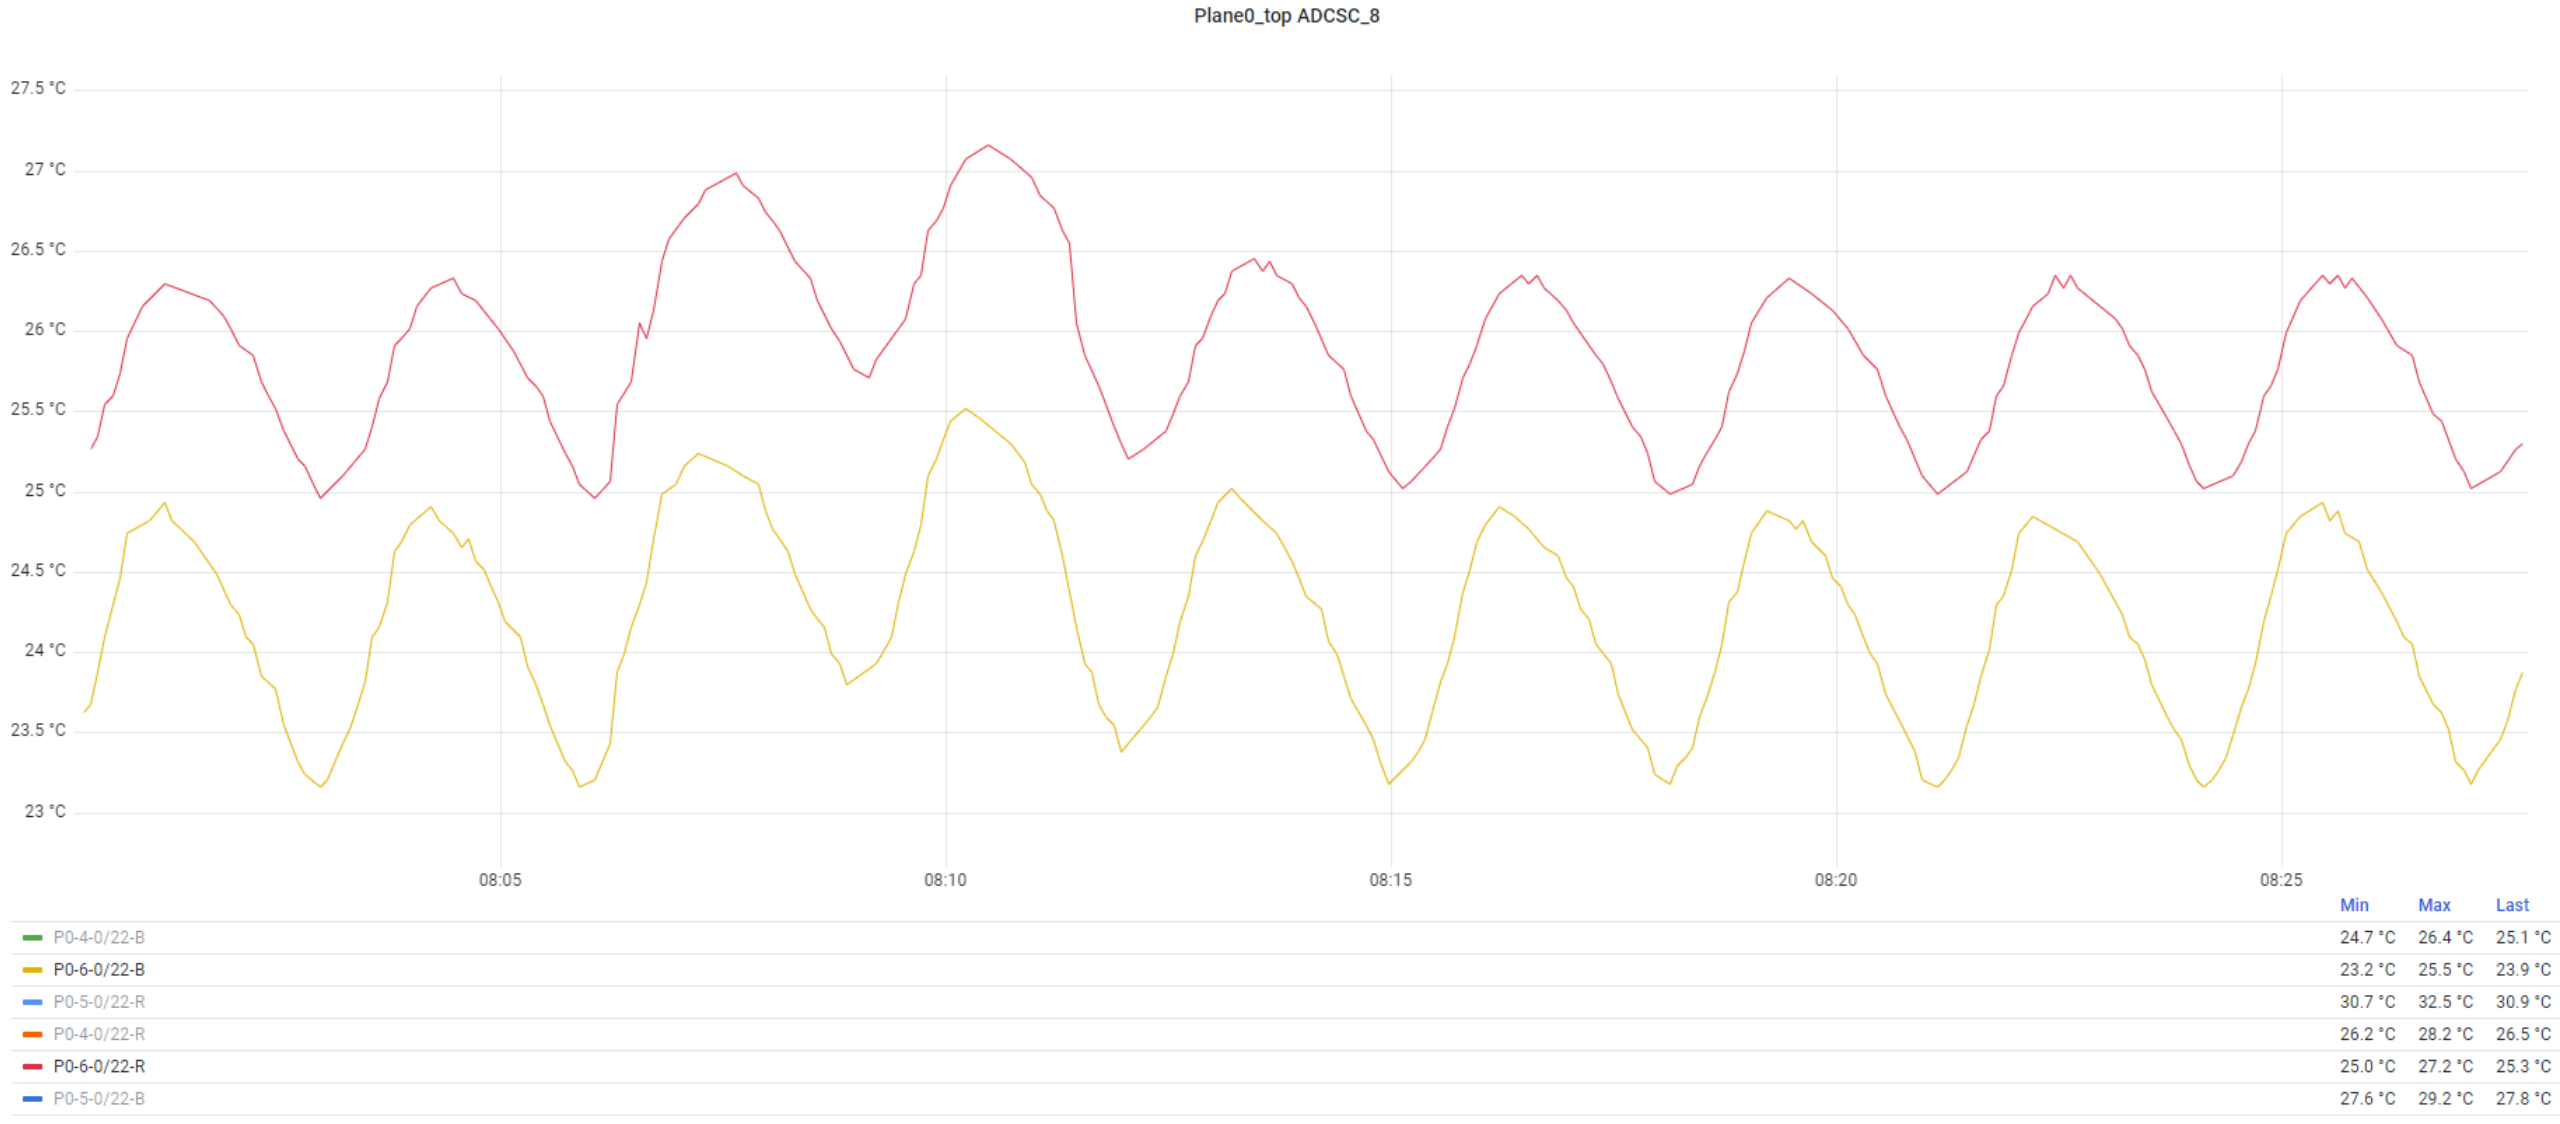

More ways to display data

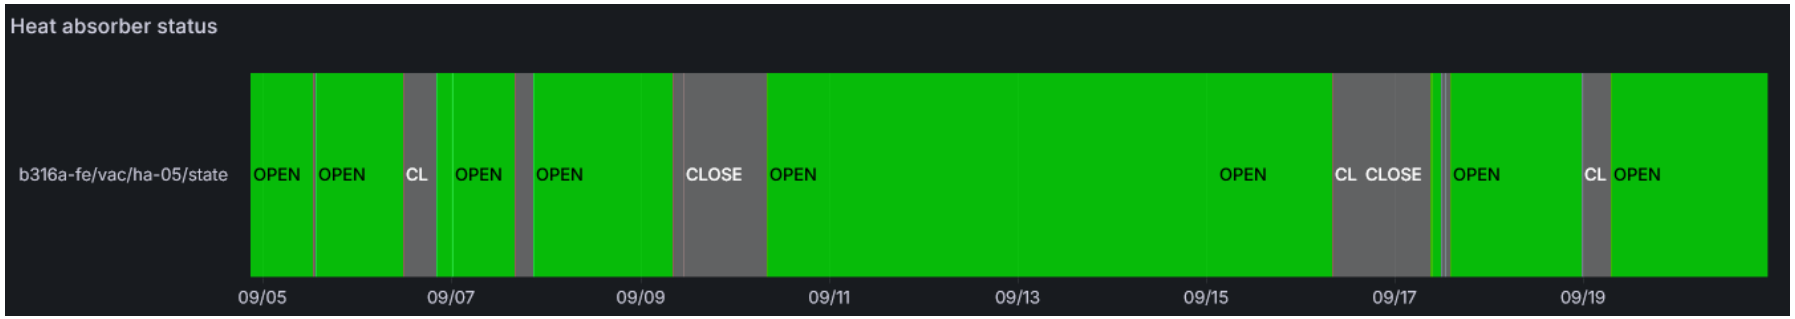

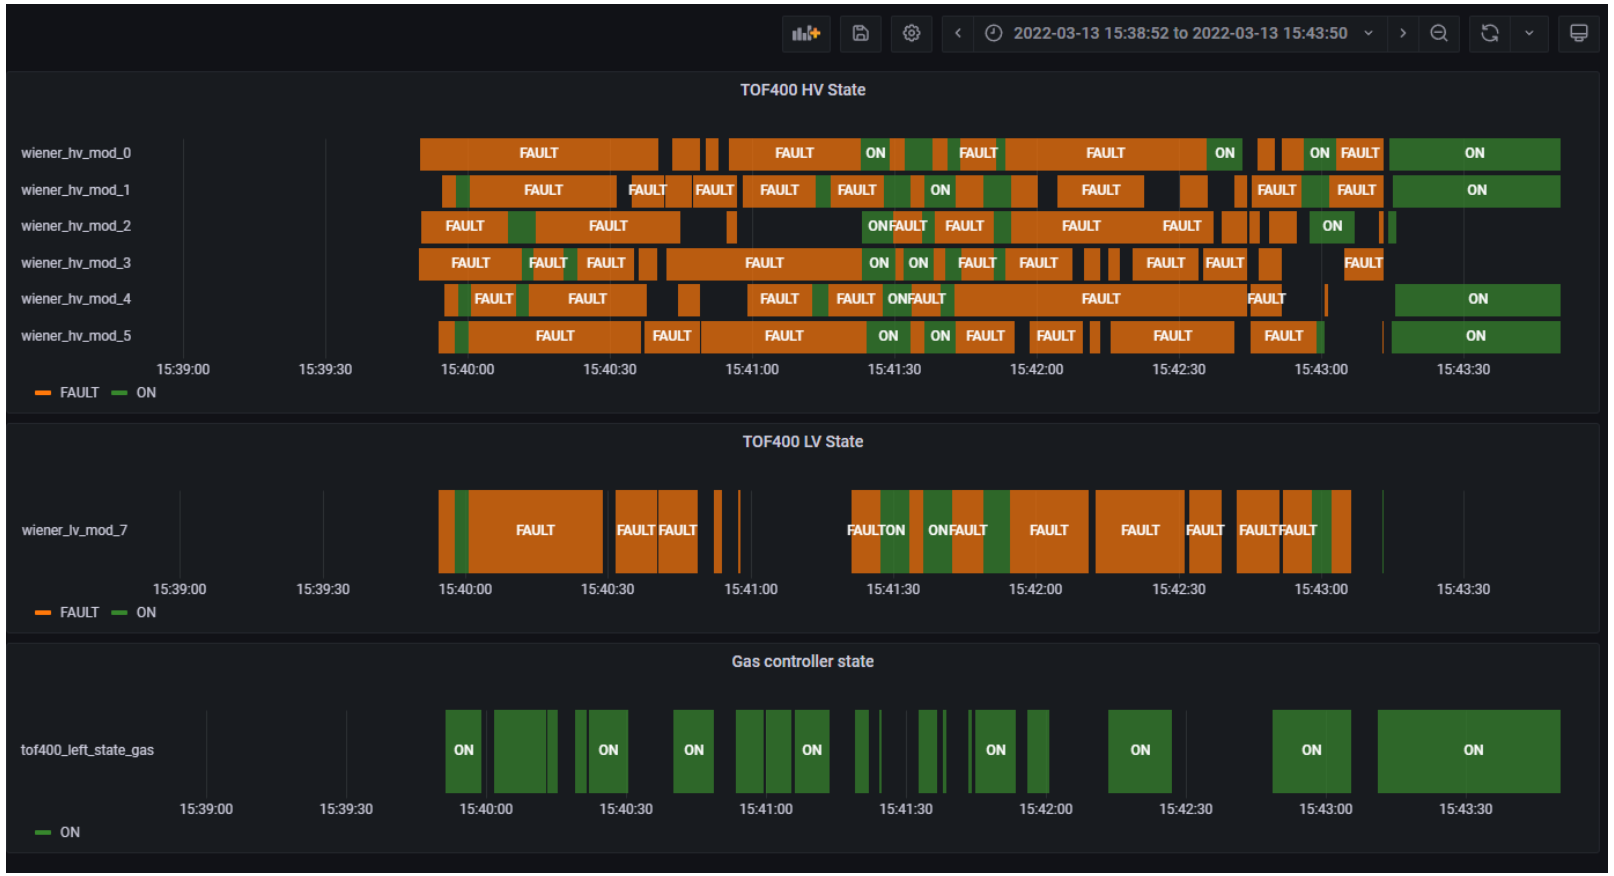

State timeline

More ways to display data

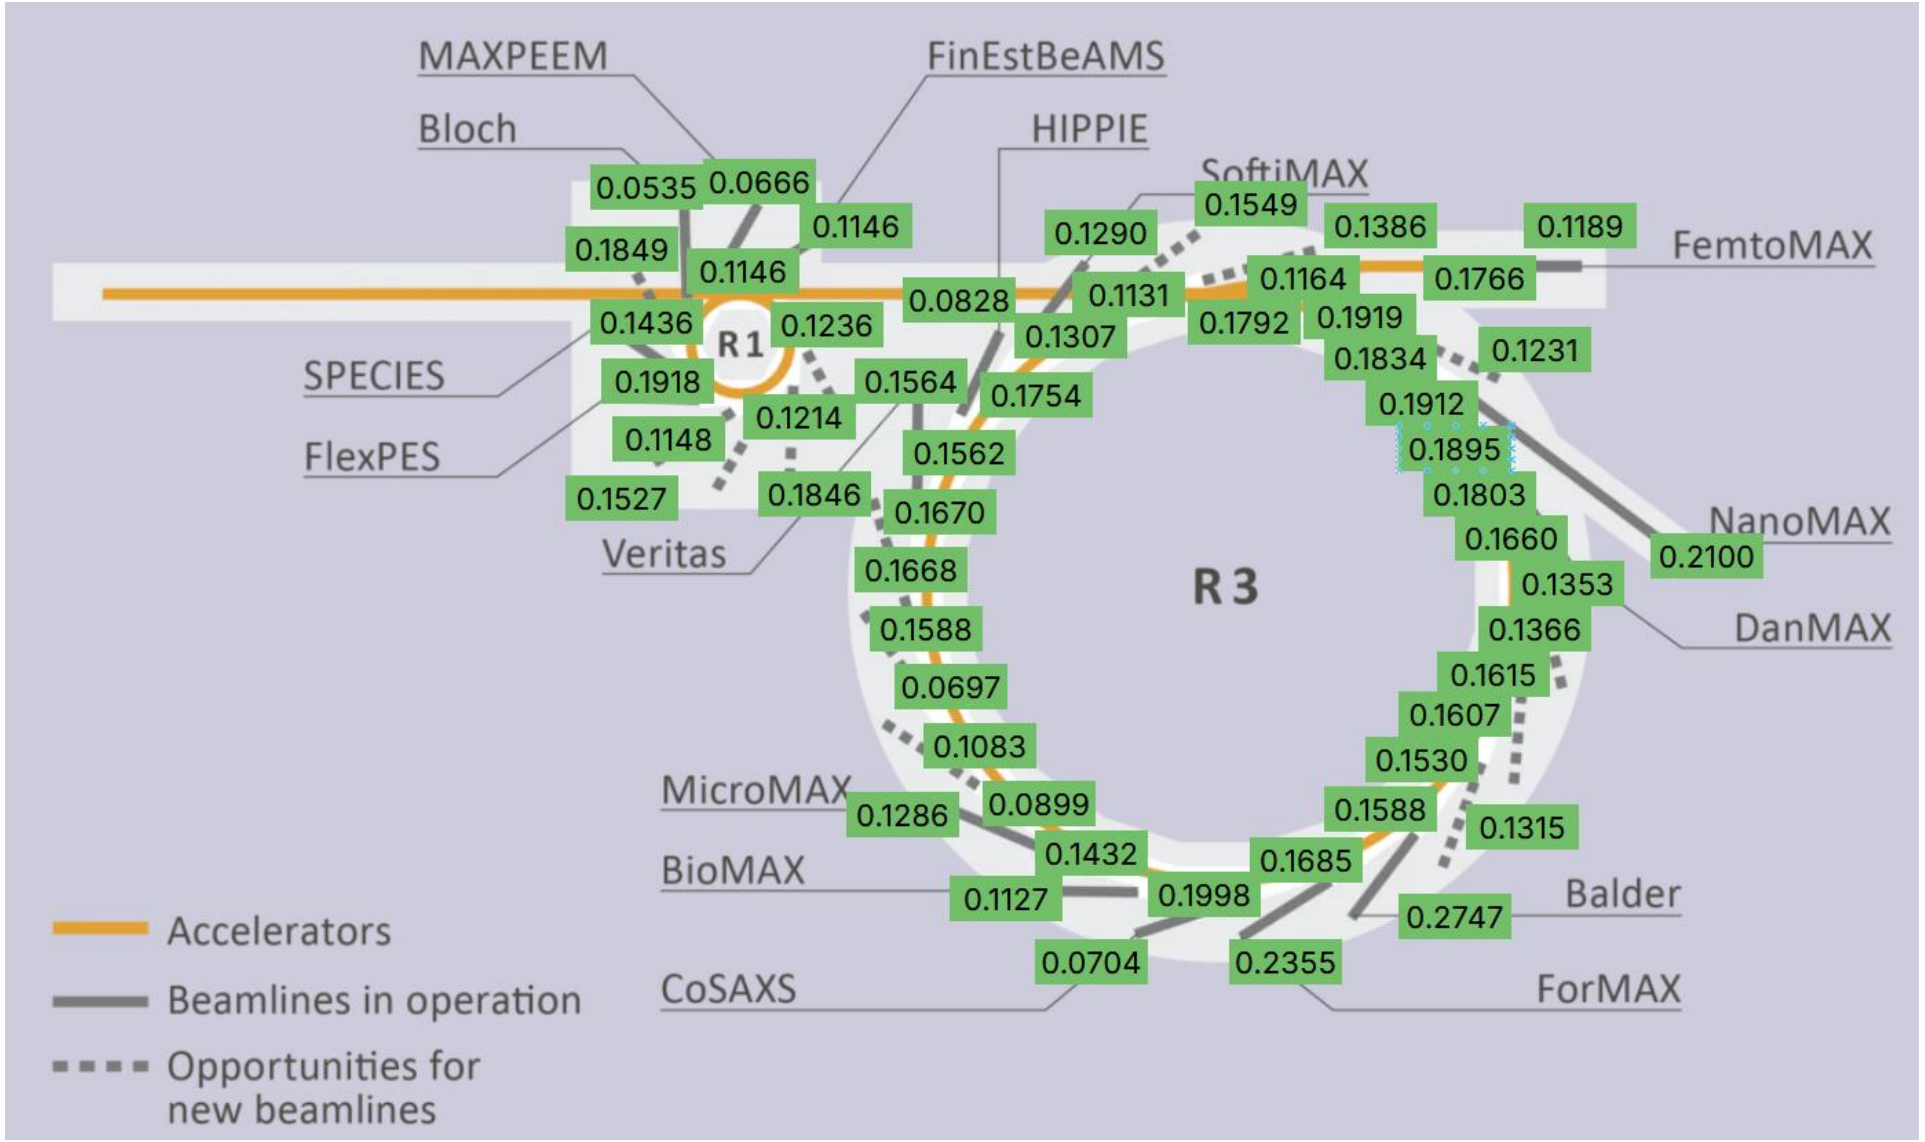

Canvas panel, radiation monitoring sensors (proof-of-concept)

More ways to display data

Display annotation when Heat Absorber is CLOSE

More ways to display data

Display annotation when Beam Current < 350 mA

Thanks to Johan Forsberg, Andrii Salnikov, Mirjam Lindberg, Vincent Hardion and everyone from KITS group who contributed to this project.