HDB++: What's new?

Reynald Bourtembourg
Johan Forsberg
Thomas Juerges
Damien Lacoste
Jan David Mol
Lorenzo Pivetta
Sergi Rubio-Manrique
Graziano Scalamera

HDB++ High-performance event-based Archiving System for TANGO

HDB++ is a DB-agnostic archiving system that inserts Tango Archive (or Change) events into your database of choice (TimeScaleDB, MySQL, SQLite, ...) using several dedicated Tango devices for event subscribing and configuration.

https://tango-controls.readthedocs.io/en/latest/tools-and-extensions/archiving/HDB%2B%2B.html

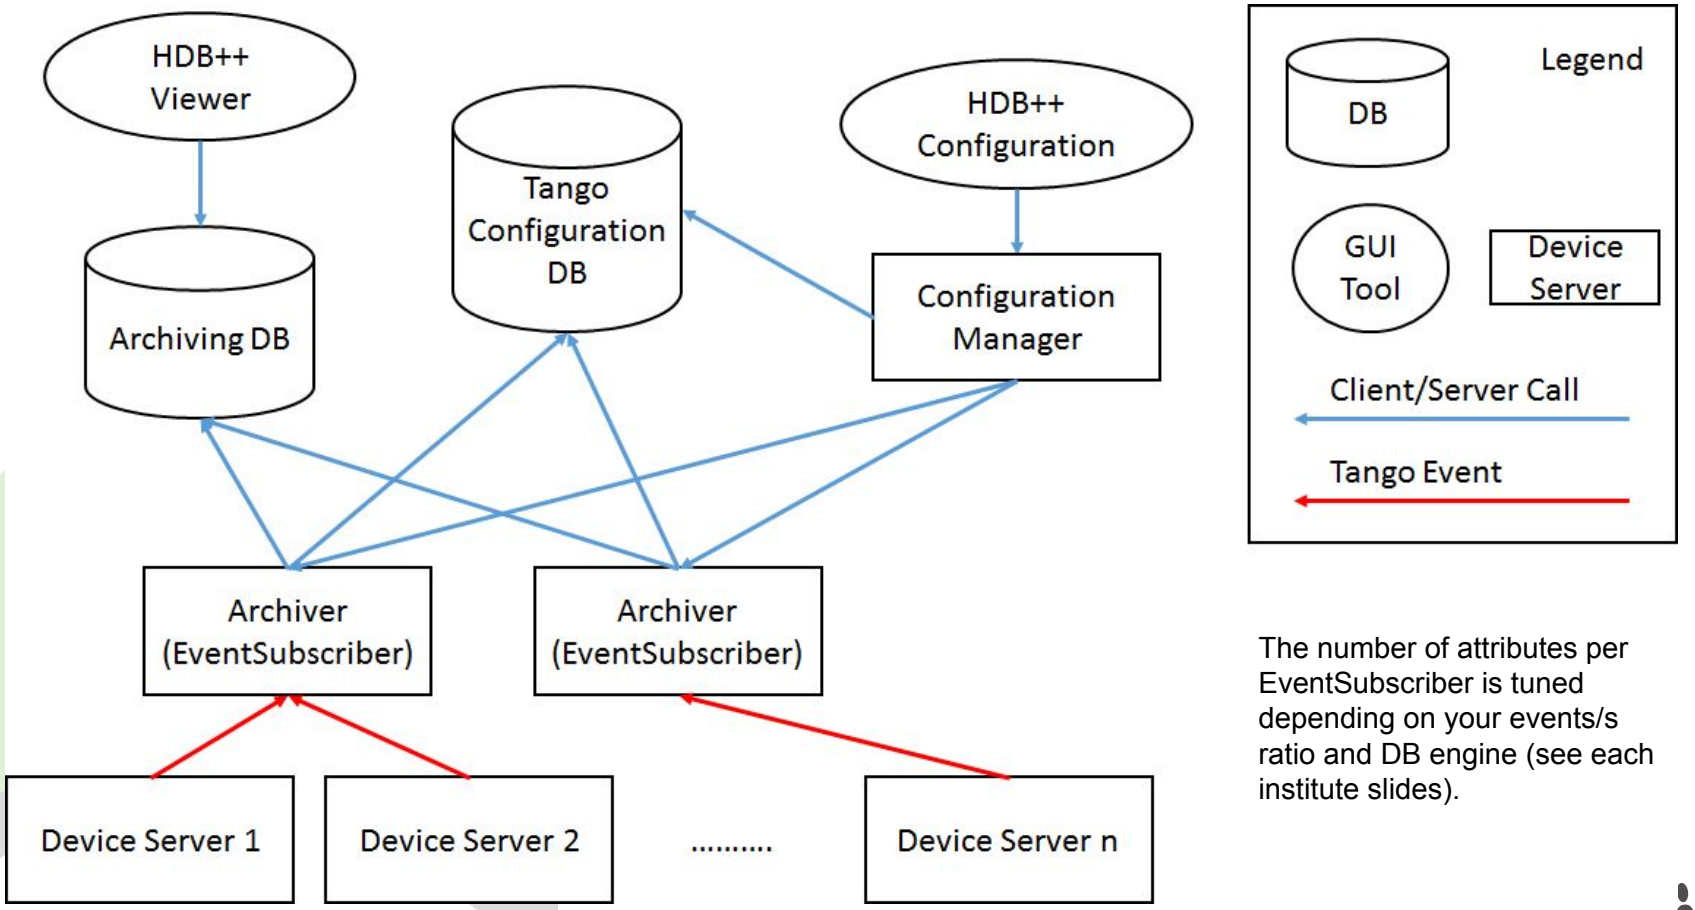

The number of attributes per EventSubscriber is tuned depending on your events/s ratio and DB engine (see each institute slides).

Supported backends and clients

Backends:

- Mysql/MariaDB
- TimescaleDB
- Cassandra Deprecated!
- ElasticSearch Status unknown!
- Mysql/MariaDB Legacy schema Deprecated
- Postgresql Status unknown, timescale library should be compatible.
- SQLite 🆕

Clients:

Extraction libraries:

- Python extraction library
- Java extraction library, a matlab binding is available.
- Cpp extraction library, not up to date.

Full visualization clients:

- eGiga (web)
- HDB viewer (java)
- Grafana (web)
- archviewer (web)
- tango_browser (PyQt, taurus)

The HDB++ Ecosystem

All libraries and tools developed by the community are accessible at https://gitlab.com/tango-controls/hdbpp

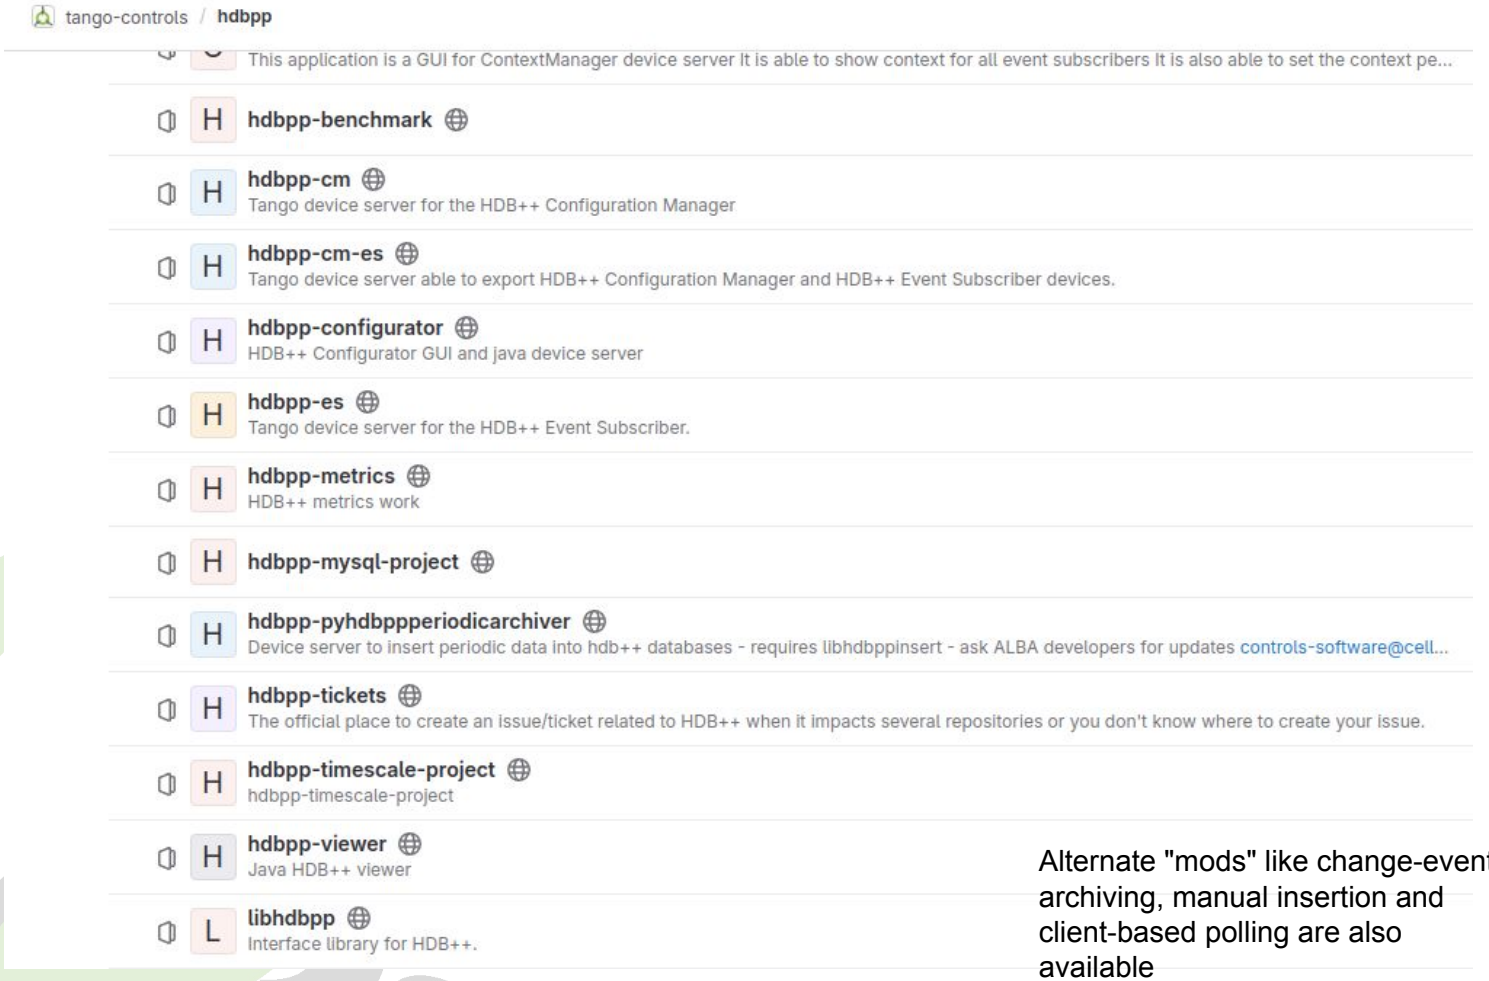

Alternate "mods" like change-event archiving, manual insertion and client-based polling are also available

Status: ALBA

ALBA - Accelerators

Running HDB++ since 2018 (HDB/TDB since 2008); configured using PyTangoArchiving, accessed using PyTangoArchiving and the new **pyhdbpp (libhdbpp-python)**

Main MariaDB host stores 6 months of undecimated data
(19779 attributes, 6 databases, 6TB in total, partitions every 15 days)

hdbacc: 943 attributes, 437 GB
hdbct: 3984 attributes, 576 GB
hdbdi: 3905 attributes, 2198 GB
hdbpc: 3127 attributes, 1112 GB
hdbrf: 3950 attributes, 525 GB
hdbvc: 4976 attributes, 1115 GB

Secondary MariaDB host stores all historical data decimated to max 1 value every 10 seconds (5TB in total since 2010).

Using 44 **CHANGE** Event Subscribers for latest devices and 52 Periodic Archivers for legacy systems.

ALBA - Beamlines

6 phase 2 beamlines (Tango9) using HDB++ (2034 attributes in 6 databases, using 189 GB)
7 phase 1 beamlines just migrated from legacy archiving to HDB++ last winter (but **notifd memleaks!**).

We use **PyHdbppPeriodicArchiver** and l**ibhdbpp-insert** for periodic/custom archiving insertion.
Old and new archiving working simultaneously while control system is migrated to Tango 9.

Status: Elettra

Fermi
HDB++ running since 2015
1 MySQL back-end, hdb++ schema
~17000 attributes from 8 Tango facilities
~9000 ev/minute; peaks up to 53.5K ev/minute
Context based archiving -> ~30 archiving strategies defined
59 EventSubscriber + 5 ConfigurationManager
~350 GB on disk - master (current + 2 previous years)
~350 + 640 GB on disk - replica

Elettra
HDB++ running since 2016
1 MySQL back-end, hdb++ schema (legacy HDB schema dropped 2021)
~6000 attributes
~5000 ev/minute
Context based archiving -> 7 archiving strategies defined
21 EventSubscriber + 1 ConfigurationManager
~330 GB on disk - master
~330 GB on disk - replica

Infrastructure (buildings facility)
1 MySQL back-end, hdb++ schema
~275 attributes (new, growing to ~1000)

ProxySQL
Used to make different DBs visible as a single one
For instance Fermi+Infrastructure, Elettra+Infrastructure, Fermi current (last 2 years) + Fermi archived (9 years) data

ESRF's Database setup (TimescaleDB)

- 1997 - 2018.
  - Moved to timescaledb backend.
  - 1 database engine with 2 databases.
  - 1.4To of compressed data + aggregates.
- From 2019.
  - Hdb++ with timescaledb backend.

- 17772 attributes, of which 17209 scalars, and 15081 doubles!
- 84 archivers.
- 3 configuration managers instances (45 devices).
- Database size ≃ 4.2To for 4,5 years.
  - 1407Go compressed data (6.5To before compression).
  - ≃ 500Go aggregates (not compressed yet).
- Stores about 700 events/s.

Status: 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 this summer.

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 (compression not fully used yet)

Status: SKAO

Engineering Data Archive (EDA) prototype deployed

- Current deployments in Australia, Canada, South Africa, UK plus developer laptops
- Based on Kubernetes & Helm/Helmfile
- Individual deployments (ITFs and PSIs for SKA-MID and SKA-LOW, no production yet)
  - Timescale in dedicated namespace + persistent volume for data storage
    - Configurable: One DB per deployment or use a shared DB (e.g. for production)
  - Pods: HDB++-CM, HDB++-ES, ArchiveViewer, ArchWizard, configurator
  - Deployment as easy as

    ```
    make k8s-install-chart ARCHIVER_DBNAME=<dbname>
    ARCHIVER_TIMESCALE_HOST_NAME=<hostname>
    ARCHIVER_TIMESCALE_PORT=<port>
    ARCHIVER_TIMESCALE_DB_USER=<dbuser>
    ARCHIVER_TIMESCALE_DB_PWD=<dbpassword>
    ```
  - Configuration (upload, download, modification) with yaml file via configurator web page: http://configurator.{KUBE_NAMESPACE}.svc.internal.skao.int:8003
  - pyhdbpp on users' computers that are on the SKAO VPN

Latest development news

- eGiga
- libhdb++
- Libhdbpp-timescale
  - Support for images merged!
- libhdbpp-sqlite
- pyhdbpp (mysql/mariadb/timescale)
- Archviewer

eGiga

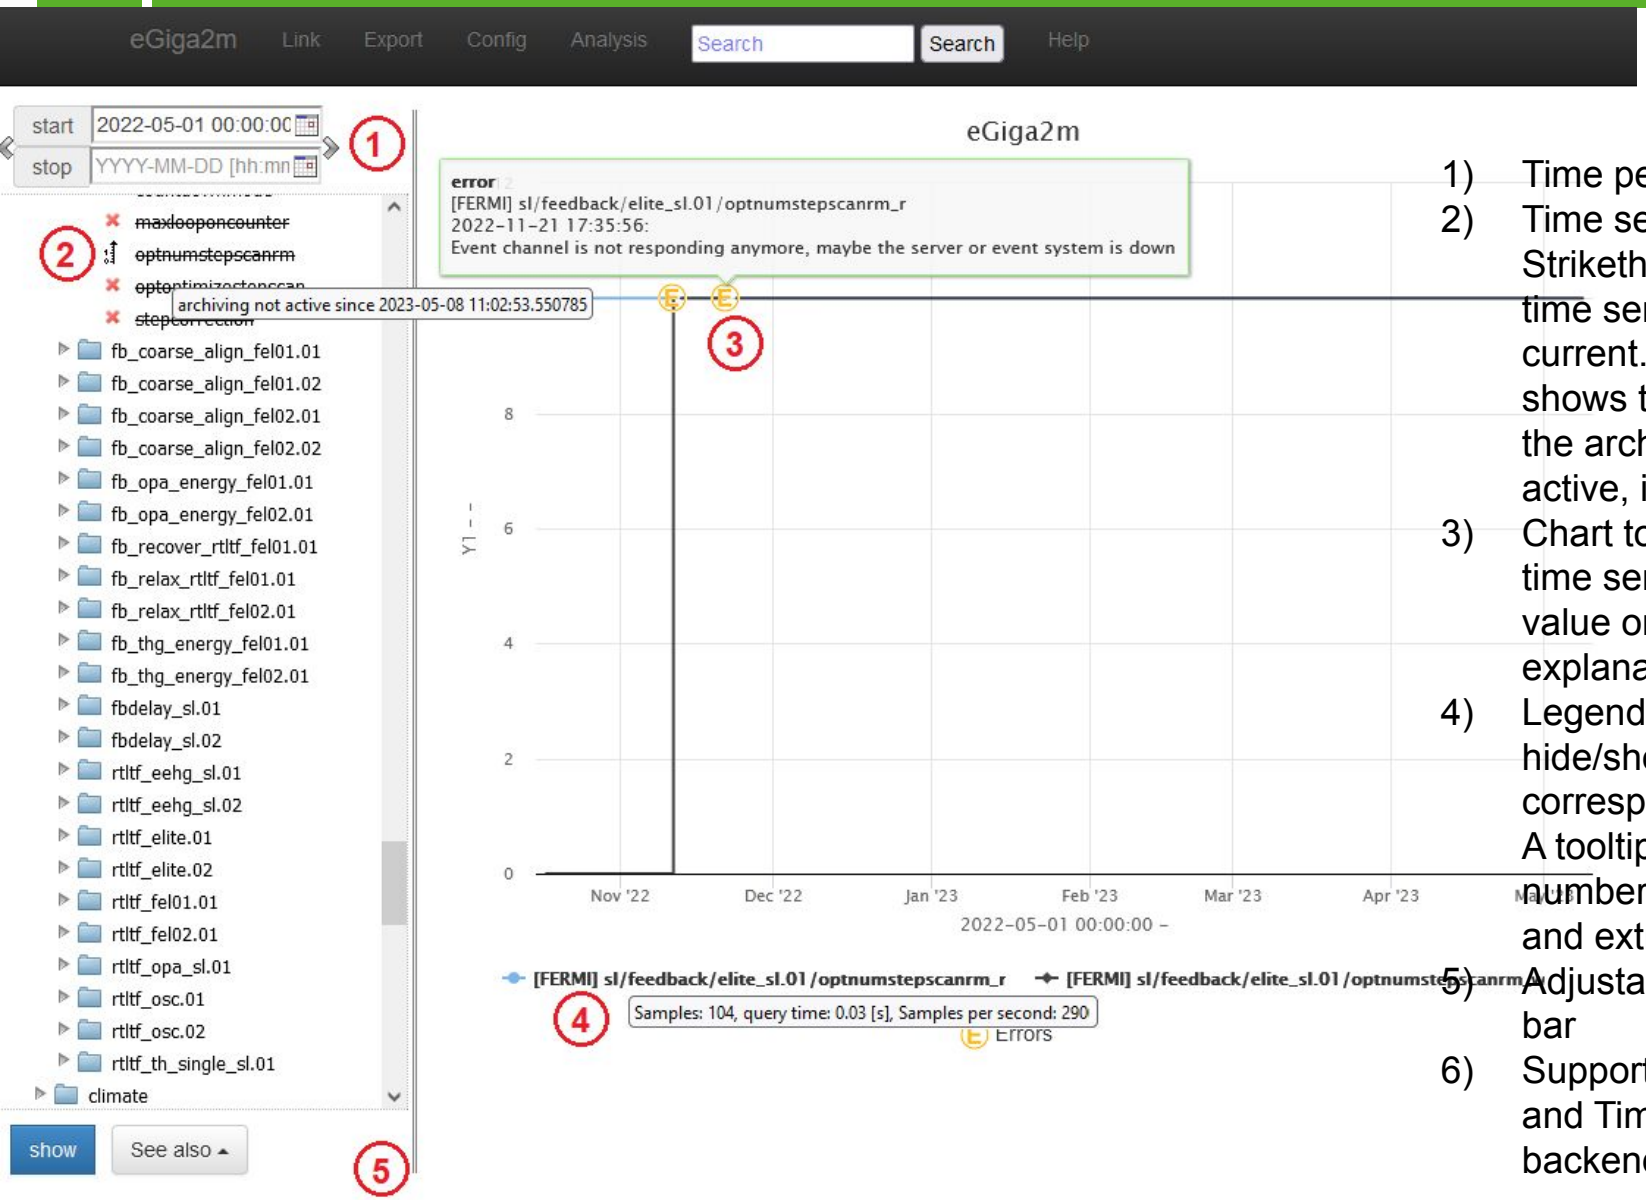

1) Time period selection
2) Time series selection. Strikethrough means time series are not current. A tooltip shows the last date the archiviation was active, if available.
3) Chart tooltip shows time series numerical value or error explanation
4) Legend. Click to hide/show the corresponding curve. A tooltip shows the number of samples and extraction time
5) Adjustable separation bar
6) Support for MySQL and TimescaleDB backends.

libhdbpp-sqlite

- libhdbpp-sqlite now available
- Supported on Linux, macOS
  - Windows unsupported

pyhdbpp

Python3 package for data extraction

> **pip3 install pyhdbpp**

- Common API for MariaDB, MySQL and TimeScaleDB
- AbstractReader object provides generic extraction interface
- MultiDB reader allows to merge data from multiple data sources
- Dedicated DB back-end implementation is loaded at runtime
- Connection setup is stored in .yaml or Tango properties
- Taurus Widget available! (pyqtgraph)

```
import pyhdbpp

rd = pyhdbpp.get_default_reader()

rd.get_attribute_values('sr/vc/01/pressure',

                        '2023-06-03','2023-06-04')


Out:  ((1685782812.169, 1.4e-08, 0),

      (1685782842.011, 1.5e-08, 0), …
```

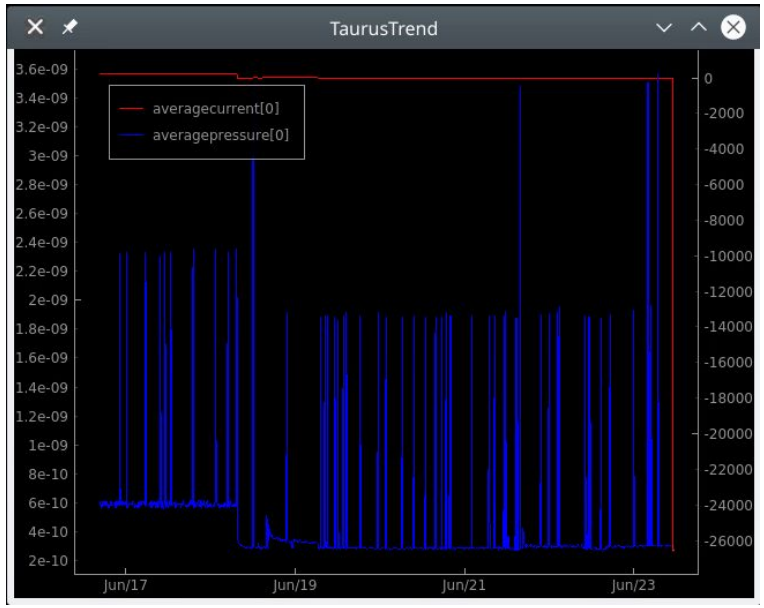

- Elastic? SQLite? … just inherit AbstractReader and contribute with your own implementation!
- New browser based on it : https://gitlab.com/tango-controls/hdbpp/libhdbpp-tangobrowser
- Configuration API still pending (PyTangoArchiving)
- Grafana proof-of-concept

pyhdbpp-tangobrowser

browser for Tango attributes with live and hdb++ plotting, using pyhdbpp, taurus, pyqtgraph

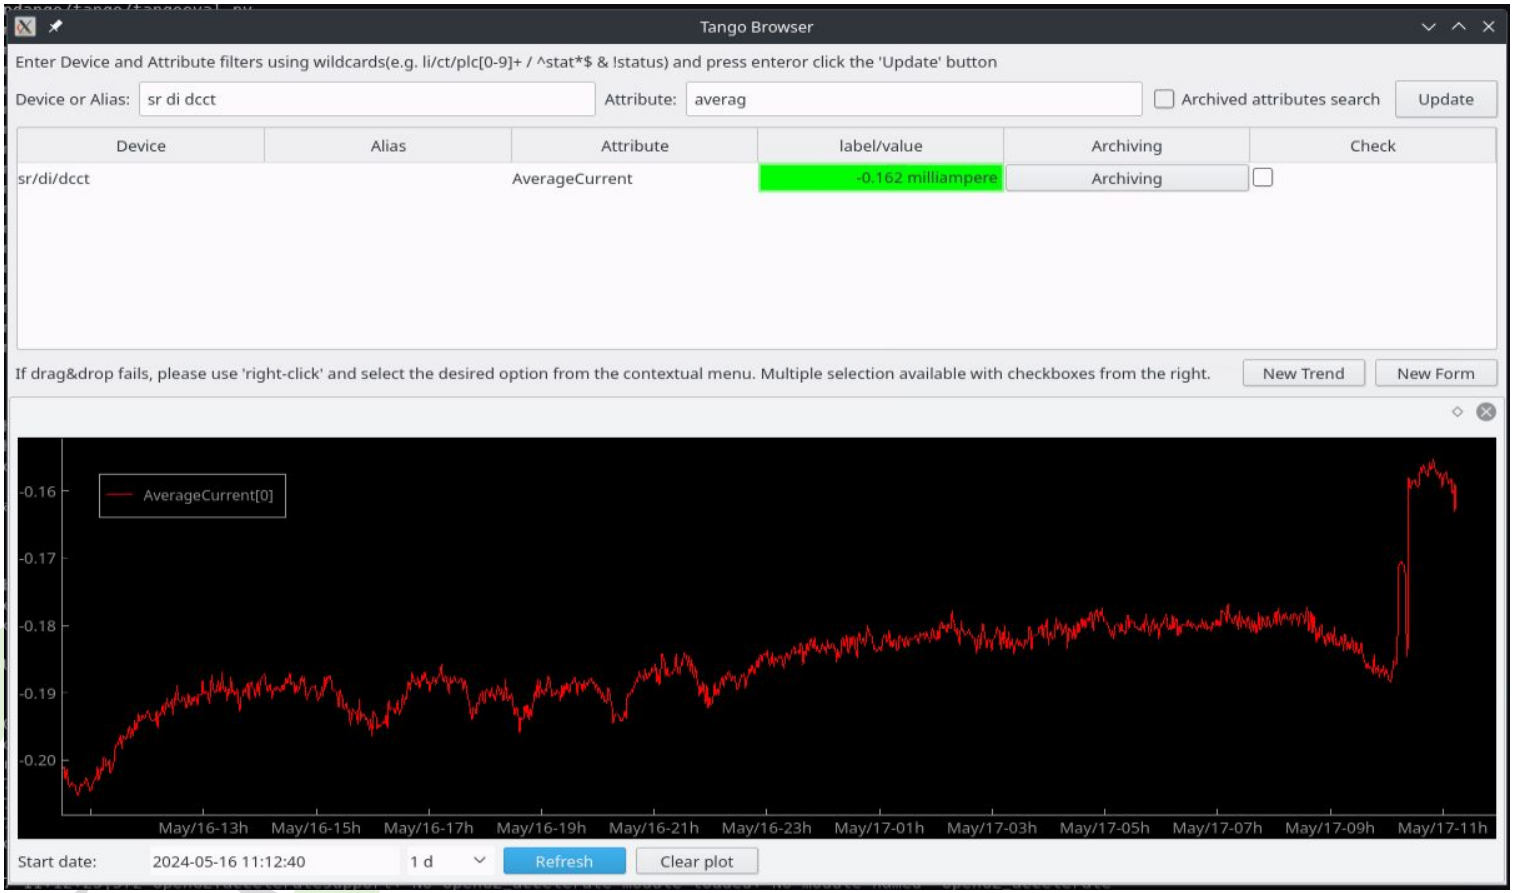


```
pip install tangobrowser
```

https://gitlab.com/tango-controls/hdbpp/libhdbpp-tangobrowser

Archviewer 1.0

Web based archive viewer, **supporting only TimescaleDB**

Frontend based on (P)React and backend in Python

https://gitlab.com/tango-controls/hdbpp/archviewer (currently a mirror only)

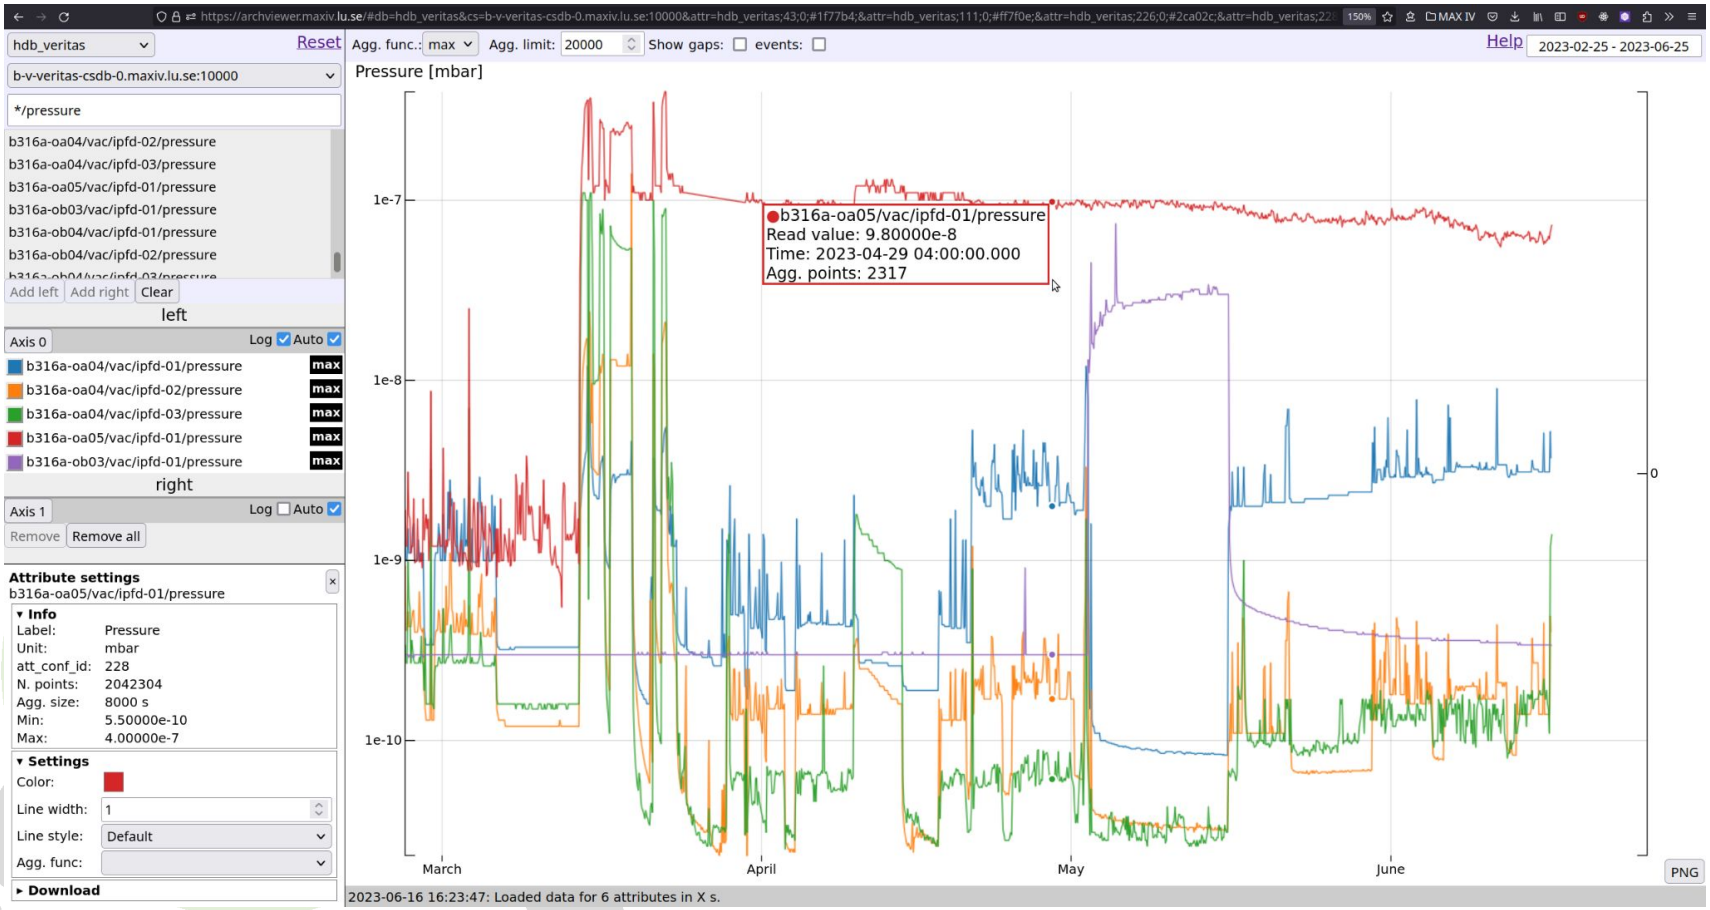

Conda packages

HDB++ archiver and configuration manager device servers (only for Timescale) available from Conda on conda-forge channel:

```
conda install -c conda-forge libhdbpp-timescale hdbpp-cm hdbpp-es
```

Debug versions also available:

```
conda install -c conda-forge libhdbpp-timescale-dbg hdbpp-es-dbg
```

Upcoming in HDB++

- **CI/CD**
  - In an effort to simplify and help installation CI will be set up on most of the hdb++ repos.

- **Support for newer cppTango versions**
  - Needed because of ABI & API changes in cppTango

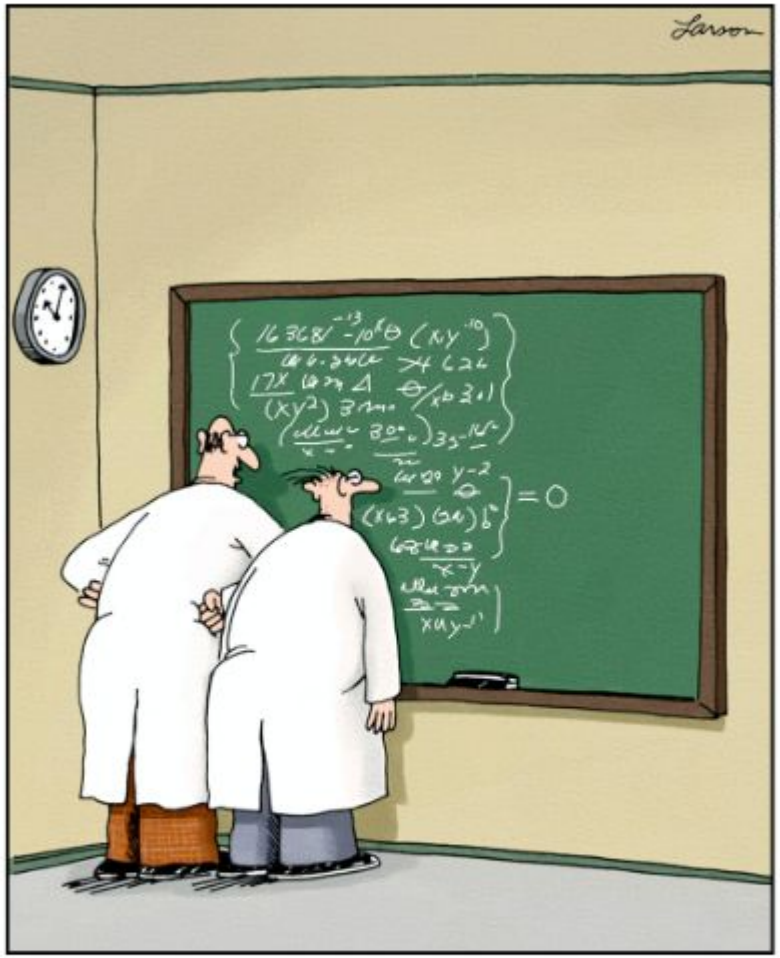

"No doubt about it, Ellington—we've mathematically expressed the purpose of HDB++. God, how I love the thrill of scientific discovery!"

Current Databases Benchmarck

Unifying the format in which we report our databases to facilitate comparison and architectural design for new facilities

(in progress)

Institute	ALBA	ESRF	SKAO MID-ITF	SKAO MID-PSI	SKAO LOW-ITF	SKAO AAVS3	ELETTRA	ELETTRA
Db Engine	MariaDB	TimescaleDB	TimescaleDB	TimescaleDB	TimescaleDB	TimescaleDB	MySQL/InnoDB	MySQL/InnoDB
Used for	vacuum	Generic	Generic	Generic	Generic	Generic	Fermi FEL	Elettra SR
mode	:hange_events, inserts	events	ents+periodic (periodic is discouraged)	+periodic (periodic is discouraged)	periodic is discouraged)	eriodic is discouraged)	archive events	archive events
Hdb++ version		1.x	current	current	current	current	current	current
Packaging/distribution	debian package	compile-in-place	docker images	docker images	docker images	docker images	compile-in-place	compile-in-place
OS	Debian 9	ubuntu 20.04	Kubernetes	Kubernetes	Kubernetes	Kubernetes	Ubuntu 18.04	Ubuntu 18.04
RAM	128 Gb	128 Gb	Varies(32G)	Varies(2x128G)	Varies(32G)	Varies(2x128G)	32 Gb master/64 Gb replica	16 Gb master/64 Gb replica
N CPUs	24	64	Varies(16)	Varies(2x64)	Varies	Varies(2x128)	8 master/4 replica	8 master/4 replica
Db Size (Gb)	514	4200	1.8	0.4	77	5	350	332
Timespan	3 months	4 years	4 months	6 months	9 months	1 month	2 years	2 years
N attributes	5159	17772					17326	6067
N attributes pushing events	4505	all?					all	all
N scalar double attributes	3461	15006						
from N devices	806	6579	~10	~10	~10	~10		
N subscribers	16	84	1	1	1	1	59	19
N attributes/subscriber	**322.4375**	**211.5714286**	**0**	**0**	**0**	**0**	**293.6610169**	**319.3157895**
partitioning	30 days	from couple of days for double to a month					2 months for double, 1 year c	2 months for double, 1 year othe
max partition size	64G	~64Gb					29 Gb	35 Gb
events/second (hdb++es)	137.133	700	Not measured yet	Not measured yet	Not measured yet	Not measured yet	150	85
events/second (stored)								
events/second/attribute (average)	0.03						0.008	0.014
events/second/subscriber (average)	**8.5708125**	**8.333333333**	**#VALUE!**	**#VALUE!**	**#VALUE!**	**#VALUE!**	**2.542372881**	**4.473684211**
max events/second/subscriber	14.36						**16.5**	**17**
max events/second/attribute	1						2.8	1.5
max events/second/attribute (stored)								
max events/second/device								
bytes/event								
bytes/second	452							
	62018							

Get in touch with the HDB++ Community

- #hdbpp channel on tango.controls.slack.com
- Report issues on https://gitlab.com/tango-controls/hdbpp/hdbpp-tickets
- Source code available under https://gitlab.com/tango-controls/hdbpp