



HDB++

*L. Pivetta
on behalf of the HDB++ collaboration*

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

- **Written in C++**
- **Event-driven:** exploit the TANGO publish/subscribe mechanism
- Architecture based on:
 - One or more archivers (EventSubscriber TANGO ds)
 - Configuration management (ConfigurationManager TANGO ds)
 - Libraries for data insertion and extraction (C++ and Java)
 - Data extraction: TANGO ds / clients
- **Fast**
 - One database for slow and fast archiving (up to 1000 samples/s, possibly more)
- **Flexible**
 - Easy to manage and maintain even without GUI front-ends
- **Self contained**
 - Single source for all configuration parameters (TANGO database)
- **Modular**
 - Abstraction+implementation libraries to support different database engines and schema
 - Support for existing HDB schema on MySQL
 - Support for **hdb++ new schema** with improved features (μ s timestamp)
 - Support for **noSQL** back-end (Apache Cassandra)
 - Easily extensible to additional database/schema
- **Scalable:** same as TANGO, deploy as many DS as needed
- **GUI:** for HDB++ configuration and data extraction as well

*Original requirements
Updates in the following slides*

Archive event

- TANGO provides specific event for archiving purposes
- The **archive** event can be sent:
 - on value change → specify absolute or relative threshold
 - periodically → specify period
- Choosing the right thresholds is mandatory:
 - if the threshold is too large no events are sent → no archiving
 - if the threshold is too small too many events are sent → “noisy” archiving
- The right threshold is **strictly related to the variable/signal** to be archived (type, bandwidth, sampling rate...)

The EventSubscriber TANGO device is the core of the HDB++ archiving system

- Event based; TANGO provides **archive events** on change and periodic basis
- Configuration stored in the TANGO database (device)
- One thread in charge of event(s) subscription and callback execution: fills a FIFO acting as producer
- One thread in charge of pushing data into the database; reads the FIFO as consumer

- Device methods allow to perform the following per-instance operations:
 - **add/remove** an Attribute to/from archiving
 - **start/stop** the archiving for all Attributes
 - start/stop the archiving for one Attribute
 - read the status of an Attribute
 - read the number/list of Attributes currently archived (started)
 - read the number/list of Attributes currently not archived (stopped)
 - read the number/list of Attributes in charge
 - read the configuration parameters of each Attribute
 - read the number/list of working Attributes
 - read the number/list of faulty Attributes with diagnostics
 - read the number/list of Attributes pending in the FIFO

- The EventSubscriber exposes some **additional figures**:
 - **for each instance**, total number of records per time
 - for each instance, total number of failures per time
 - **for each attribute**, number of records per time
 - for each attribute, number of failures per time
 - for each attribute, time stamp of last record
 - for each attribute, min and max processing and storing times

+ Manage context
+ Manage time-to-live

EventSubscriber device configuration

Jive 7.17 [acudebian7:10000]

Server: /hdb+_+_cassandra_es-srv/sr-2/HdbEventSubscriber/tango/hdb/es-sr-2/Properties

Collection Server Device Class Alias Property

Device properties [tango/hdb/es-sr-2]

Property name	Value
AttributeList	tango://orion.esrf.fr:10000/sr/d-ct/1/current;strategy=RUN;ttl=0 tango://orion.esrf.fr:10000/sr/d-halo/id7/halo;strategy=RUN toto;ttl=1 tango://orion.esrf.fr:10000/sr/d-halo/id7/beamcore;strategy=RUN;ttl=24 tango://acudebian7.esrf.fr:10000/test/universal/1/devshortrw;strategy=RUN;ttl=0
__SubDevices	sys/access-control/1 tango://orion:10000/sys/access-control/1 orion.esrf.fr:10000/sr/d-halo/id7 orion.esrf.fr:10000/sr/d-ct/1 acudebian7.esrf.fr:10000/test/universal/1

Jive 7.10 [ken:20001]

Server: /hdb+_+_es-srv/gonzalo2_b/HdbEventSubscriber/gonzalo/hdb_new/subscriber2/Properties

Server Device Class Alias Att. Alias Property

Device properties [gonzalo/hdb_new/subscriber2]

Property name	Value
LibConfiguration	libname=libhdb+_+_mysql.so host=srv-log-srf user=hdbj password=xxxx dbname=hdbpp port=3306

EventSubscriber Class configuration

Jive 7.17 [acudebian7:10000]

File Edit Tools Filter

Class:/HdbEventSubscriber/Properties

Collection Server Device **Class** Alias Property

Group3Dtm141
HakaPoC
HarmAnalysis
HdbAccess
HdbArchiver
HdbArchivingWatcher
HdbConfigServer
HdbConfigurationManager
 Properties
 Attribute properties
 Devices
HdbConfiguratorServer
HdbEventHandler
HdbEventSubscriber
 Properties
 Attribute properties
 Devices
HdbExtractor
HdbLogger
HdbManagerSimu
HdbRecordCounter
HdbSigArchiveSurvey
HdbUtilities
HlsCaptor

Class properties [HdbEventSubscriber]

Property name	Value
ContextsList	ALWAYS: Store in HDB++ under any circumstances SHUTDOWN: Store in HDB++ in SHUTDOWN period RUN: Store in HDB++ in RUN period toto: tsoin tsoin
cvs_location	/home/cvsadm/cvsroot/fermi/servers/hdb++/hdb++es/src/
DbHost	cassandra2
DbName	hdbtest
DefaultStrategy	RUN
Description	This class is able to subscribe on archive events and store value in Historical DB
doc_url	http://www.esrf.eu/computing/cs/tango/tango_doc/ds_doc/
InheritedFrom	TANGO_BASE_CLASS
LibConfiguration	keyspace=hdbtest contact_points=cassandra2 user=hdbwriter password=MyPass libname=libhdb++cassandra.so logging_enabled=true
PollingThreadPeriod	1
ProjectTitle	Tango Device Server
StartArchivingAtStartup	false
StatisticsTimeWindow	1
SubscribeRetryPeriod	10

Refresh Apply New property Copy Delete

Support several operating modes, or strategies, in the archiving engine on a per-attribute basis.

→ Modify AttributeList device property to support a name=value approach for the strategy

```
tango://srv-tango-srf.fcs.elettra.trieste.it:20000/eos/climate/18b20_eos.01/state; \
```

```
strategy=RUN|SHUTDOWN
```

```
tango://srv-tango-srf.fcs.elettra.trieste.it:20000/eos/climate/18b20_eos.01/temperature; \
```

```
strategy=RUN|SHUTDOWN
```

Labels for the strategy defined in a free property, and/or in the class property and/or in the device property, with increasing priority (ContextList)

The defaults values, as well as the DefaultStrategy, are pre-defined but can be modified by the user. The default values for ContextList are:

ALWAYS	0
RUN	1
SHUTDOWN	2
SERVICE	3

Get|SetAttributeStrategy methods to modify per-attribute strategy.

“context” r/w memorized attribute, set by higher layer logic, tells the archiver the required Context transition

To support temporary storage of historical data one new column has been added to the att_conf table:

```
CREATE TABLE IF NOT EXISTS att_conf (
  att_conf_id INT UNSIGNED NOT NULL AUTO_INCREMENT PRIMARY KEY,
  att_name VARCHAR(255) UNIQUE NOT NULL,
  att_conf_data_type_id INT UNSIGNED NOT NULL,
  ttl INT UNSIGNED NULL DEFAULT 0,
  facility VARCHAR(255) NOT NULL DEFAULT "",
  domain VARCHAR(255) NOT NULL DEFAULT "",
  family VARCHAR(255) NOT NULL DEFAULT "",
  member VARCHAR(255) NOT NULL DEFAULT "",
  name VARCHAR(255) NOT NULL DEFAULT "",
  INDEX(att_conf_data_type_id)
) ENGINE=MyISAM COMMENT='Attribute Configuration Table';
```

The ttl column defines the time-to-live in hours on a per-attribute basis. Configuring the ttl can be done from the Configuration Manager with the appropriate Method, and modified with Get|SetAttributeTTL methods; ttl is stored in the AttributeList configuration as name=value:

```
tango://srv-tango-srf.fcs.elettra.trieste.it:20000/eos/climate/18b20_eos.01/state; \
strategy=RUN|SHUTDOWN; ttl=8760
tango://srv-tango-srf.fcs.elettra.trieste.it:20000/eos/climate/18b20_eos.01/temperature; \
strategy=RUN|SHUTDOWN; ttl=8760
```

Deleting expired data delegated to the SQL backend.

The ConfigurationManager TANGO device simplifies HDB++ archiving system management:

- **handle** the request of archiving a new Attribute
 - ~~create an entry in the database if not already present~~
 - setup the Attribute archive event configuration
 - assign the Attribute to one of the archivers
- **move** an Attribute from one archiver to another
- keep trace of which Attribute is assigned to which archiver
- **start/stop** the archiving
- **remove** an Attribute from archiving
- + manage EventSubscriber context
- + manage Attribute time-to-live

moved to EventSubscriber

The Configuration manager exposes some **global statistics**:

- total number of Archivers
- total number of working/faulty attributes
- total number of events per second
- overall minimum and maximum processing and storing time

ConfigurationManager Class configuration

Jive 7.17 [acudebian7:10000]

File Edit Tools Filter

Class:/HdbConfigurationManager/Properties

Class Alias Property
Collection Server Device

- Group3Dtm141
- HakaPoC
- HarmAnalysis
- HdbAccess
- HdbArchiver
- HdbArchivingWatcher
- HdbConfigServer
- HdbConfigurationManager
 - Properties
 - Attribute properties
 - Devices
- HdbConfiguratorServer
- HdbEventHandler
- HdbEventSubscriber
- HdbExtractor
- HdbLogger
- HdbManagerSimu
- HdbRecordCounter
- HdbSigArchiveSurvey
- HdbUtilities
- HlsCaptor

Class properties [HdbConfigurationManager]

Property name	Value
cvs_location	/home/cvsadm/cvsroot/fermi/servers/hdb++/hdb++cm/src/
DbHost	cassandra2
DbName	hdbtest
Description	
doc_url	http://www.esrf.eu/computing/cs/tango/tango_doc/ds_doc/
InheritedFrom	TANGO_BASE_CLASS
LibConfiguration	keyspace=hdbtest contact_points=cassandra2 user=hdbwriter password=MyPass libname=libhdb++cassandra.so
ProjectTitle	Hdb++ configuration manager

Refresh Apply New property Copy Delete

A C++ API decouples the archiving engine (EventSubscriber) from the database back-end

- **libhdbpp** : database abstraction layer
- **libhdbpp-mysql** : implementation, HDB++ schema support, MySQL back-end
- ~~**libhdbpp-cassandra** : implementation, HDB++ schema support, Cassandra back-end~~
- ~~**libhdb-mysql-legacy** : implementation, legacy HDB schema support, MySQL back-end~~
- **libhdbpp-postgresql** : implementation, HDB++ schema support, PostgreSQL back-end
- ~~**libhdbpp-elk** : implementation, HDB++ schema support, ELK back-end~~
- **libhdbpp-timescale** : implementation, HDB++ schema support, TimescaleDB back-end
- **libhdbpp-sqlite** : implementation, HDB++ schema support, SQLite back-end

Deprecated

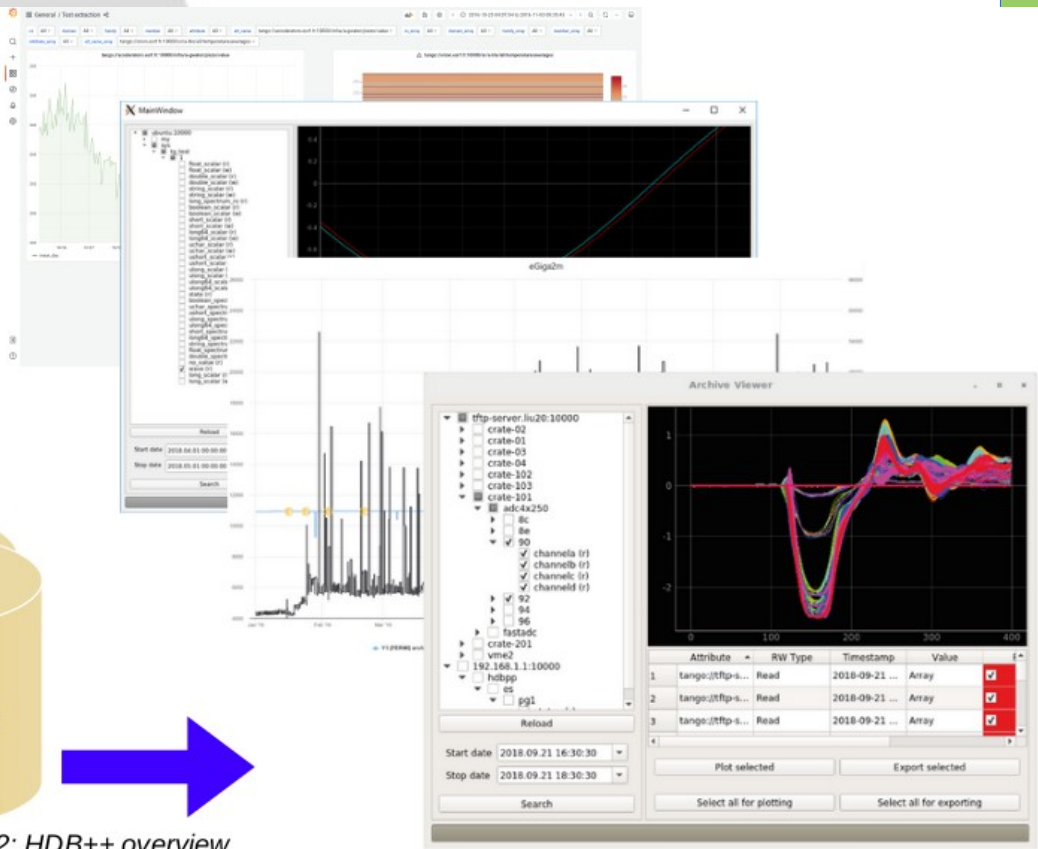

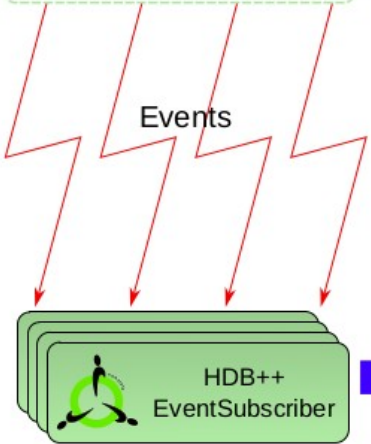
The libraries allow reusing the EventSubscriber, the ConfigurationManager and the GUIs without any change

HDB++ is easily extendable to support additional back-ends(*) just writing the specific implementation library

(*) not limited to database engines... HDF5 format on file?



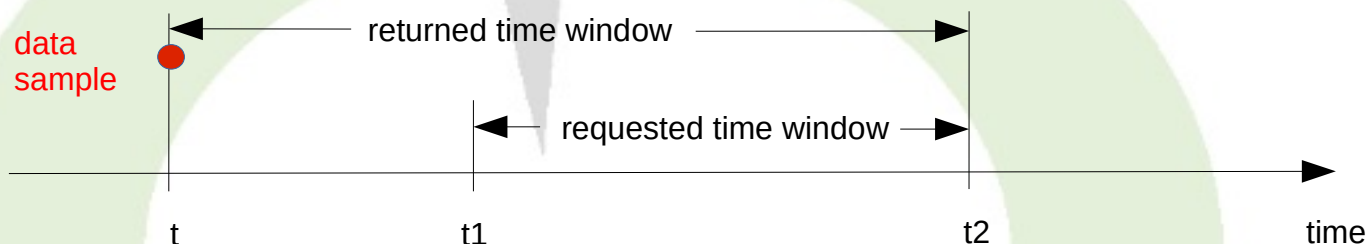
- Hdb++ overview:
- Events based
 - Scalable through the use of multiple EventSubscribers
 - Support several databases, and the use of multiple clusters to separate concerns
 - Get your data through various viewers or via libraries



C++ and Java native extraction libraries

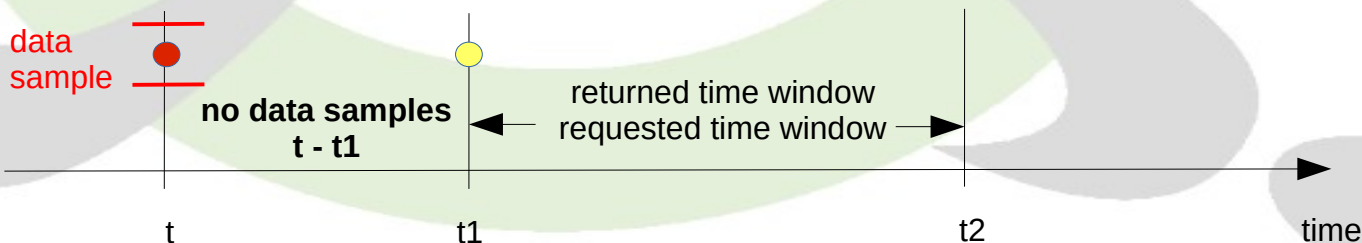
The data extraction library shall be able to **deal with event based archiving, i.e. data value change with respect to specified thresholds**; the possible lack of data in the requested time window shall be properly managed:

- returning some no-data-available error: in this case the reply contains no data
- enlarging the time window to include some archived data; no fake samples introduced




- returning the value of the last archived data anyhow; the requested time interval is kept and the last available data sample returned; the data value is guaranteed when **archiving on change**, care must be taken in case of **periodic archiving**

archive change event thresholds

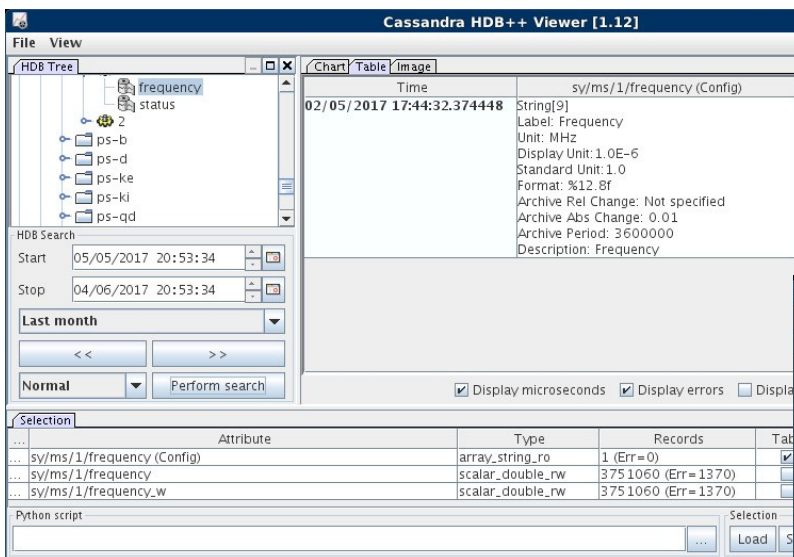


- Qt based GUI using the MathGL framework for plotting
- Exploits the C++ extraction library
- Supports multiline and surface plots




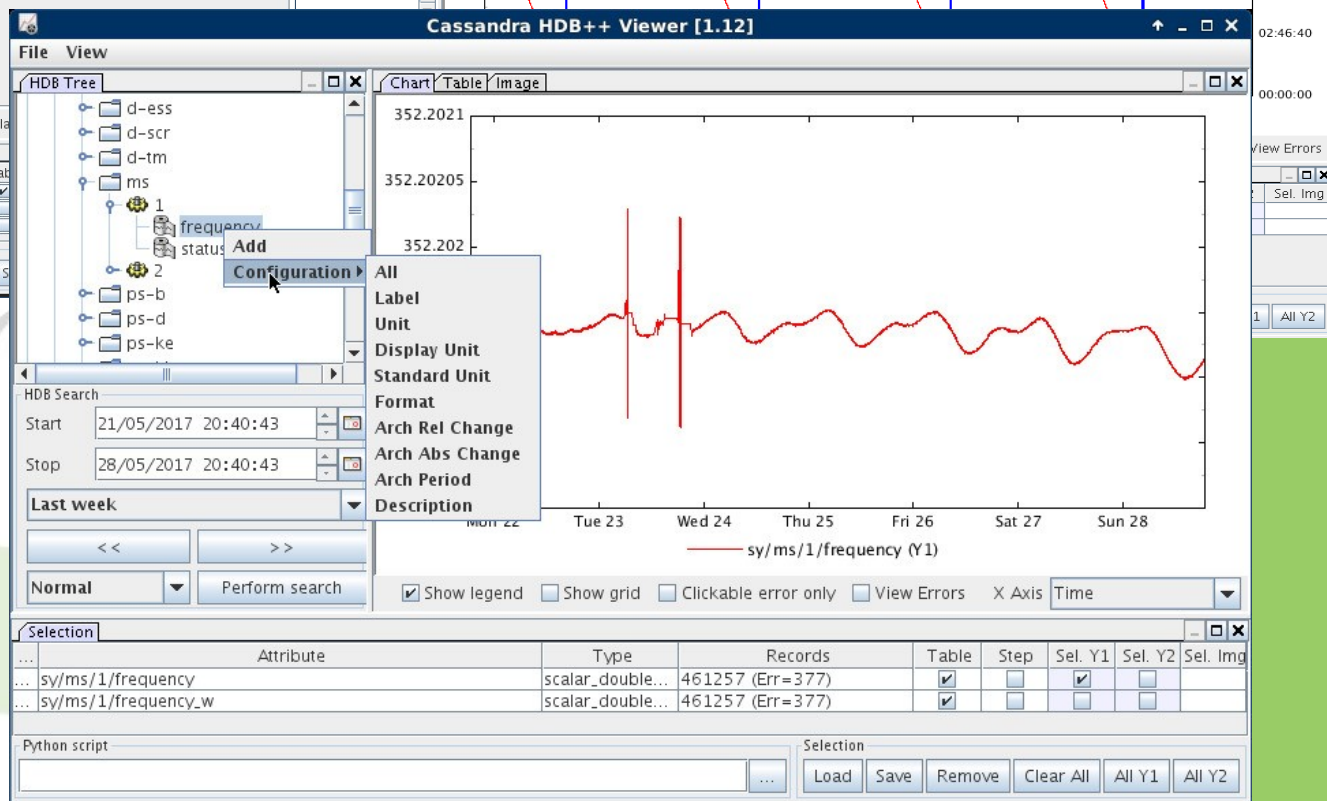
Configure Show Errors

- Java based GUI for plotting
- Exploits the Java extraction library
- Table and multiline plots



Cassandra HDB++ Viewer [1.12]

Time	sy/ms/1/frequency (Config)
02/05/2017 17:44:32.374448	String[9] Label: Frequency Unit: MHz Display Unit: 1.0E-6 Standard Unit: 1.0 Format: %12.8f Archive Rel Change: Not specified Archive Abs Change: 0.01 Archive Period: 3600000 Description: Frequency

Cassandra HDB++ Viewer [1.12]

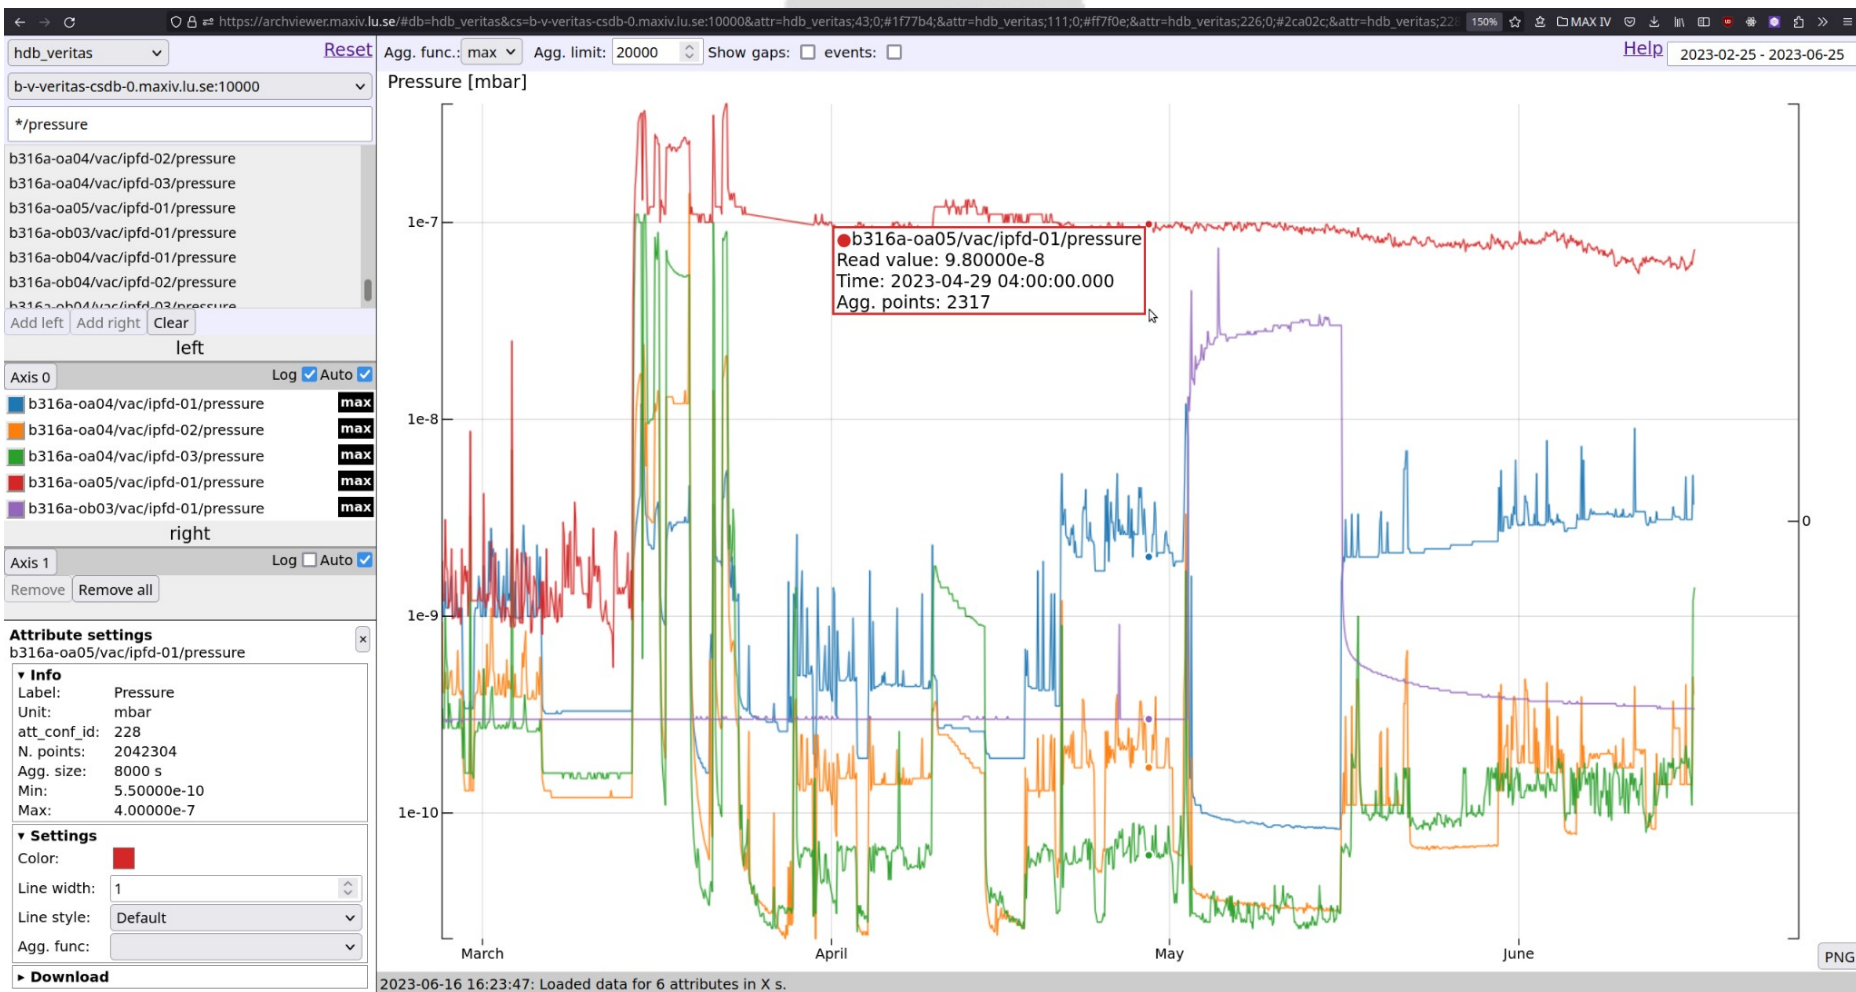
Configuration menu for 'sy/ms/1/frequency':

- Add
- Configuration
- All
- Label
- Unit
- Display Unit
- Standard Unit
- Format
- Arch Rel Change
- Arch Abs Change
- Arch Period
- Description

Chart showing current (mA) vs time (days). The plot displays a sawtooth pattern with a red signal. The y-axis ranges from 352.202 to 352.2021 mA, and the x-axis shows time from Tue 23 to Sun 28. The legend indicates 'sy/ms/1/frequency (Y1)'.

Attribute	Type	Records	Table	Step	Sel. Y1	Sel. Y2	Sel. Img
sy/ms/1/frequency	scalar_double...	461257 (Err=377)	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	
sy/ms/1/frequency_w	scalar_double...	461257 (Err=377)	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	

- Webapp for plotting HDB++ from TimescaleDB
- back-end in Python ≥ 3.8
- front-end javascript+preact+d3




- Web interface
- MySQL and TimescaleDB backend



- 1) Time period selection
- 2) Time series selection. Strikethrough means time series are not current. A tooltip shows the last date the archivation 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

HdbDiagnostics - 2.2 - 23-05-2017 13:18:05

File View help

 HDB++ Diagnostics (Cassandra)

	Faulty	Started	Paused	Stopped	Pending	ev/sec	Fail./sec	Context
"HDB++" statistics	0	78	0	0	0	12.00	0.00	Run
Beam Position Monitors	0	293	0	1	0	0.00	0.00	Run
Booster 1	0	234	0	0	0	2.00	0.00	Run
D-RGA	0	62	0	0	0	8.00	0.00	Run
Experiments Gamma Monitors	0	43	0	0	0	6.00	0.00	Run
FE Vacuum Pressures	0	375	0	0	0	0.00	0.00	Run
Fast BPM	0	4	0	0	0	0.00	0.00	Run
Front Ends	0	175	0	2	0	0.00	0.00	Run
ID22	0	7	0	0	0	0.00	0.00	Run
Injext	0	38	0	0	0	1.00	0.00	Run
Insertion Devices 1	0	94	0	0	0	0.00	0.00	Run
Insertion Devices 2	0	115	0	0	0	0.00	0.00	Run
Insertion Devices Corrections	0	32	0	0	0	0.00	0.00	Run
LINAC	0	121	0	0	0	0.00	0.00	Run
Power Supplies	0	188	0	1	0	2.00	0.00	Run
RGA 1	0	368	0	0	0	2.00	0.00	Run
RGA 2	0	305	0	0	0	0.00	0.00	Run
RGA 3	0	364	0	0	0	2.00	0.00	Run
RGA 4	0	497	0	0	0	0.00	0.00	Run
RGA 5	0	303	0	0	0	3.00	0.00	Run
RGA 6	0	427	0	61	0	0.00	0.00	Run
RGA 7	0	244	0	0	0	1.00	0.00	Run
RGA 8	0	363	0	61	0	0.00	0.00	Run
Radiations	0	199	0	0	0	0.00	0.00	Run
SR 1	3	65	0	17	0	6.00	0.00	Run
SR 2	0	200	0	0	0	1.00	0.00	Run
SR Correctors	0	94	0	0	0	0.00	0.00	Run
SR RF	7	443	0	0	0	16.00	0.00	Run
SR Steerers	0	201	0	0	0	0.00	0.00	Run
SR Vacuum 1	0	61	0	0	0	0.00	0.00	Run
SR Vacuum Pressures	0	741	0	0	0	0.00	0.00 ev/	Run
SR Vacuum Temperatures	0	934	0	1	0	1.00	0.00 ev/	Run
SY RF	1	107	0	0	0	2.00	0.00	Run
SYSTEM 1	1	215	0	0	0	0.00	0.00	Run
TL1/TL2	0	43	0	0	0	2.00	0.00 ev/	Run
test labs	0	12	0	58	0	0.00	0.00 ev/	Run
	Faulty	Started	Paused	Stopped	Pending	ev/sec	Fail./sec	Context
E.S. Manager	12	8045	0	202	0	85	0	run

Python3 package for data extraction

> **pip3 install pyhdbpp**

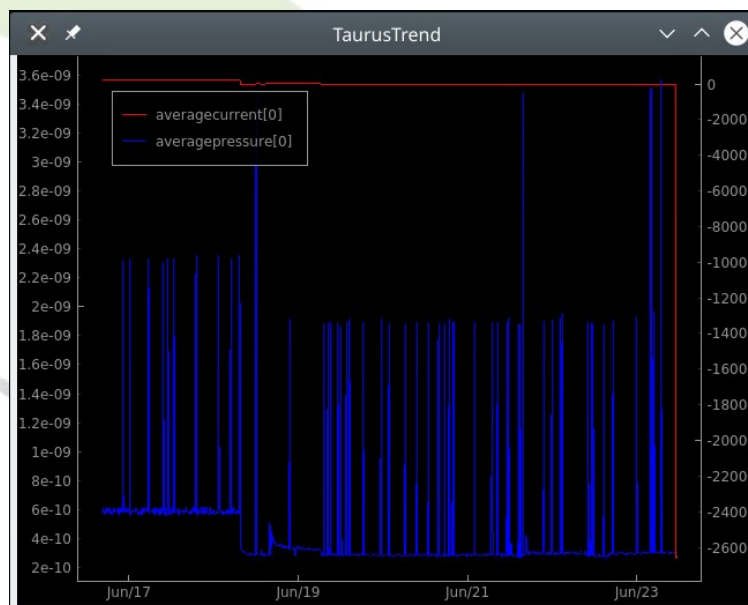
- Common API for MariaDB, MySQL and TimeScaleDB
- AbstractReader object provides generic extraction interface
- Dedicated DB back-end implementation is loaded at runtime
- Connection setup is stored in .yaml or Tango properties
- Taurus Widget available! (pyqtgraph)

```
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!
- ArchivingBrowser and SnapGUI also being migrated to python3 (ongoing)
- Configuration API still pending (PyTangoArchiving)

<https://gitlab.com/tango-controls/hdbpp>

Thank-you

...and now...