HDB++ Meeting 09/2019

ALBA Infrastructure

Archiving infrastructure: Accelerators and Beamlines

- Beamlines use TDB/HDB in virtual machines
- Accelerators store 11000 attributes in HDB++:
  - Attributes distributed in 6 MyIsam databases based on subsystem
    - rf, vc, di, acc, ct, pc, sizes between 40G-60G per month
  - Writing database contains only last 6 monthly partitions, older data is available in another server (that can be decimated).
  - 11000 attributes stored also in legacy TDB extended to 6 months.
    - 4000 attributes still to be migrated from TDB to HDB++
  - We implement both periodic and event-based archiver for HDB++
  - Backup by mysqldump restricted by table and date.

Archiving Tasks for 2019

Archiving (Accelerators)

Archiving infrastructure 80%

All databases partitioned by subsystems to improve maintainability
Speed of searches increased by a factor of 10 (in last 3 months data) without reducing resolution (down to 10 Hz)
Maintenance scripts for backup/restore ready, **automation pending**
Two new servers for long storage archiving (archiving05/archiving06)

Archiving client api / usability 80%

API refactored to provide access to multiple databases from multiple hosts
API integrated in new Taurus version using tangoarchiving schema

Archiving plotting / usability 25%

- PyQtgraph-based plotting infrastructure
  - Support tango archiving by integrating taurus_tangoarchiving

Contributions to phase II/III beamlines

New Archiving/Plotting tools to be deployed in all new beamlines

Archiving Databases

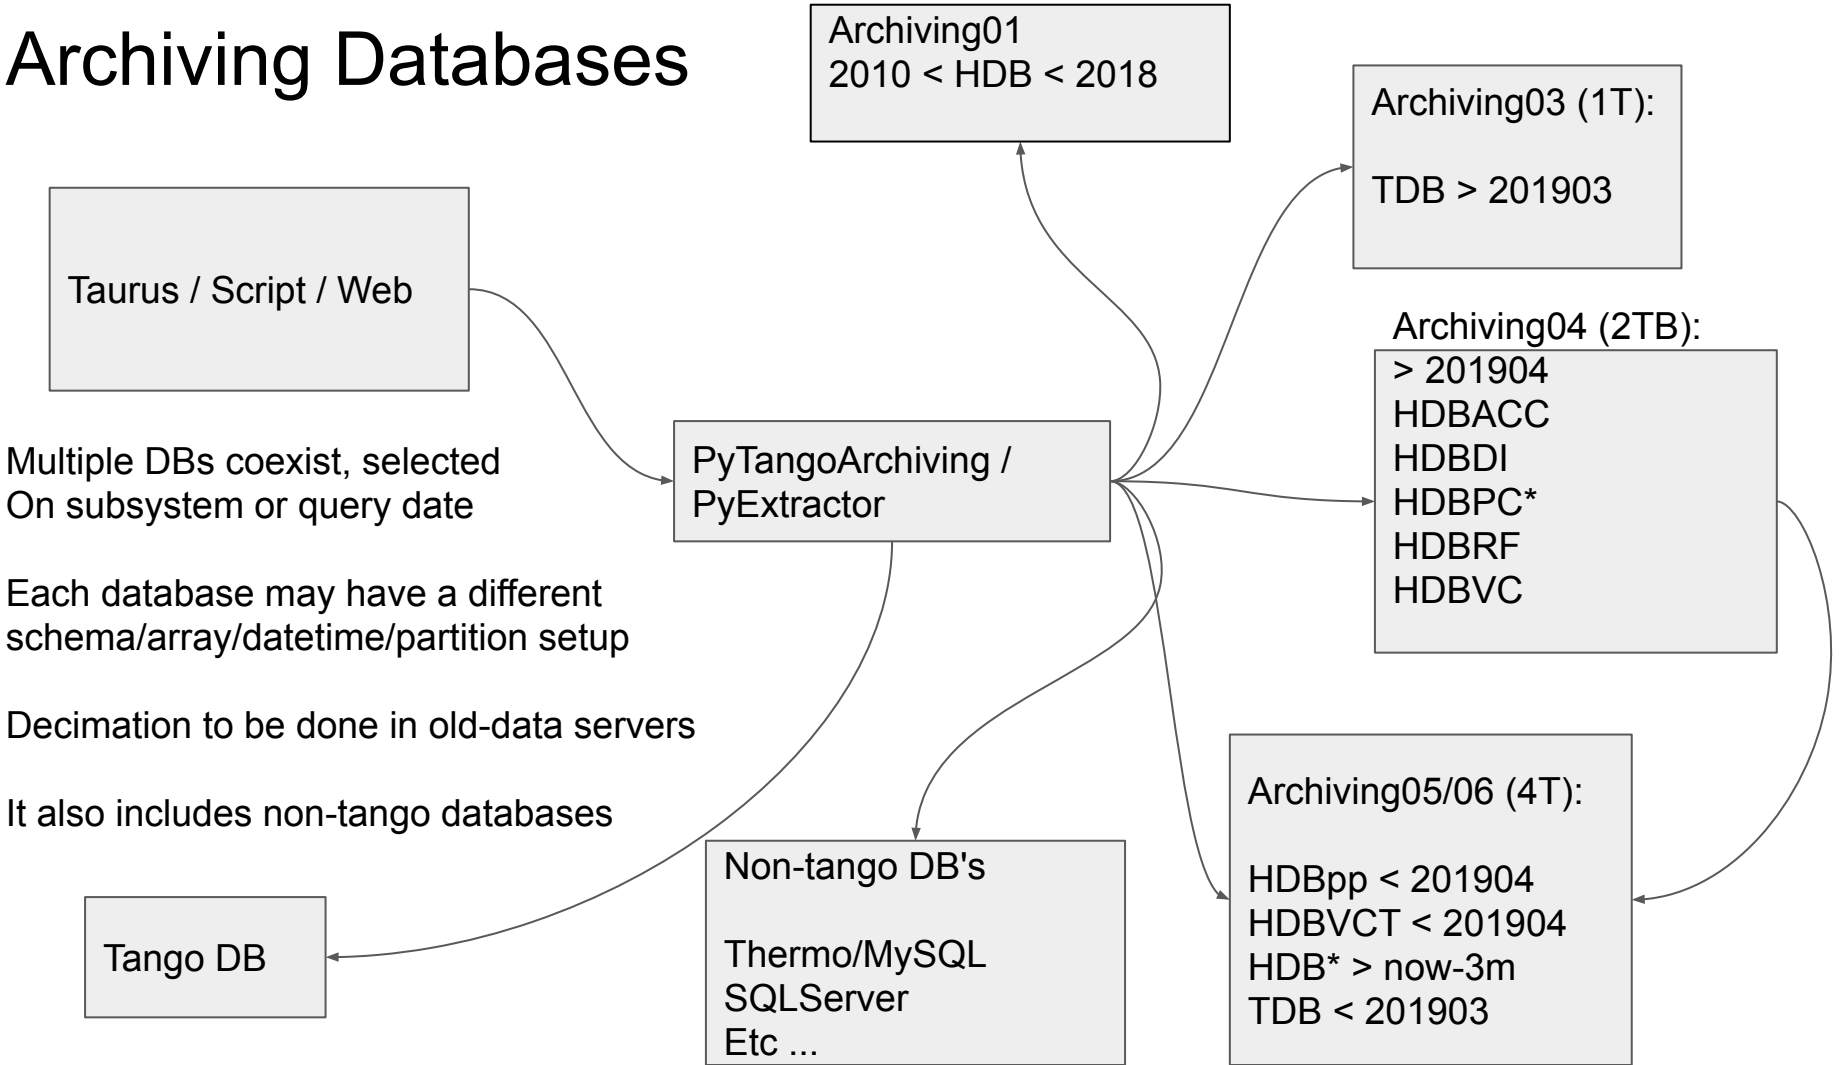

New Plotting Tool: PyQtGraph

Under test on

BL01/Miras

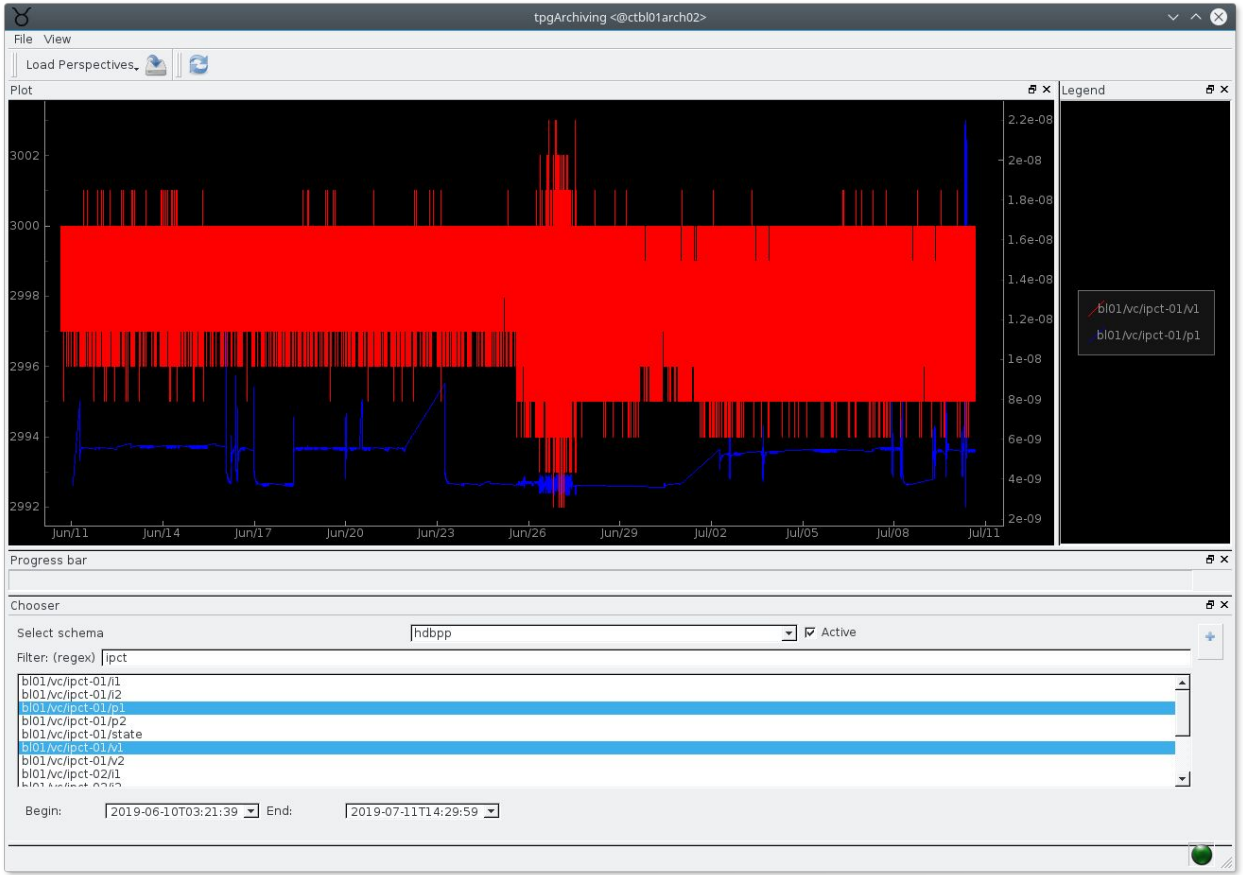

Archiving Status	Problems	Solutions adopted
• Jan 2017: HDB/TDB was the old Tango archiving system, running at Alba since 2007, storing 15000 attributes. • Resolution limited to 1Hz for last week; 0.1Hz for last 3 years.	• Values at 3 Hz were requested, to be kept for 3 months. • HDB/TDB was obsolete and hard to maintain (**~180 hours to recover after crash)**.	• The new HDB++ archiving developed by ESRF allowed acquisition at 10Hz. • Deployed at Miras/Accelerators using MariaDB servers since january 2017.
• Summer 2017: The new HDB++ required an event-based control system. • It needed to modify configuration of service area servers. • Storing at 3Hz increased the amount of data available.	• The increase of data traffic overloaded some CPU's, that we needed to optimize. • In paral.lel, the amount of data to manage for both servers and client became unmanageable.	• Beginning 2018, we started collaboration with the ESRF to solve these problems. • Modified HDB++ to reduce the CPU load. • Modified the database schema to reduce the data usage and add monthly partitioning.
• Spring 2018: Due to CPU load issues, our plan was to not shutdown HDB/TDB, but to introduce HDB++ progressively.	• HDB disks failed before summer 2018, and we had to modify TDB to keep 1 year of data during the migration. • Clients (Mambo) unable to show old data.	• Summer 2018 we accelerated the migration to HDB++. • Modified archiving tools to correlate data from old and new databases.
• Winter 2018: HDB++ timestamp had milliseconds resolution, which provides more useful information but extra load for indexing on client/server sides. • Big database indexes were too slow to backup and restore.	• The amount of memory and disk used by servers and clients grow exponentially. • Tools developed for HDB++ maintenance were not fitted to the amount of data stored (same problems for ESRF lead to full redesign of system HW).	• Winter 2018, added decimation on data and plotting to reduce clients load. • Easter 2019, reduced server load and maintenance splitting databases by subsystem (size reduced by a factor of 10).
• Summer 2019: HDB++ system currently running and stable in service area. • We store **10400 attributes, inserting 1500 values/second** on average. • This is as much data as ESRF+MaxIV+Solaris+Elettra together	• We are using 2TB disk/year; which is not a daily issue but a long-term problem for backup, restore and data availabilty. • The amount of data to be retrieved by query (2e6 points/week) is still too big for our current plotting tools (loading is slow, but still 3X times faster than using ESRF tools).	• Databases modified to reduce indexes, speeding up extraction backup and restore **(~8 hours)**. • Re-introduced periodic mode for slow values. • Two new servers for long-term storage. • New plotting tool (PyQtGraph) to be deployed in control room during summer 2019.