



New developments on HDB++, the high-performance data archiving for Tango Controls

Damien Lacoste, Reynald Bourtembourg, Lisa Banihachemi (ESRF, Grenoble)
Sergi Rubio-Manrique (ALBA-CELLS, Cerdanyola del Vallès)
Jan David Mol (ASTRON, Dwingeloo)
Lorenzo Pivetta, Graziano Scalamera (Elettra-Sincrotrone Trieste S.C.p.A., Basovizza)
Johan Forsberg (MAX IV Laboratory, Lund)
Thomas Juerges (SKA Observatory, Jodrell Bank)

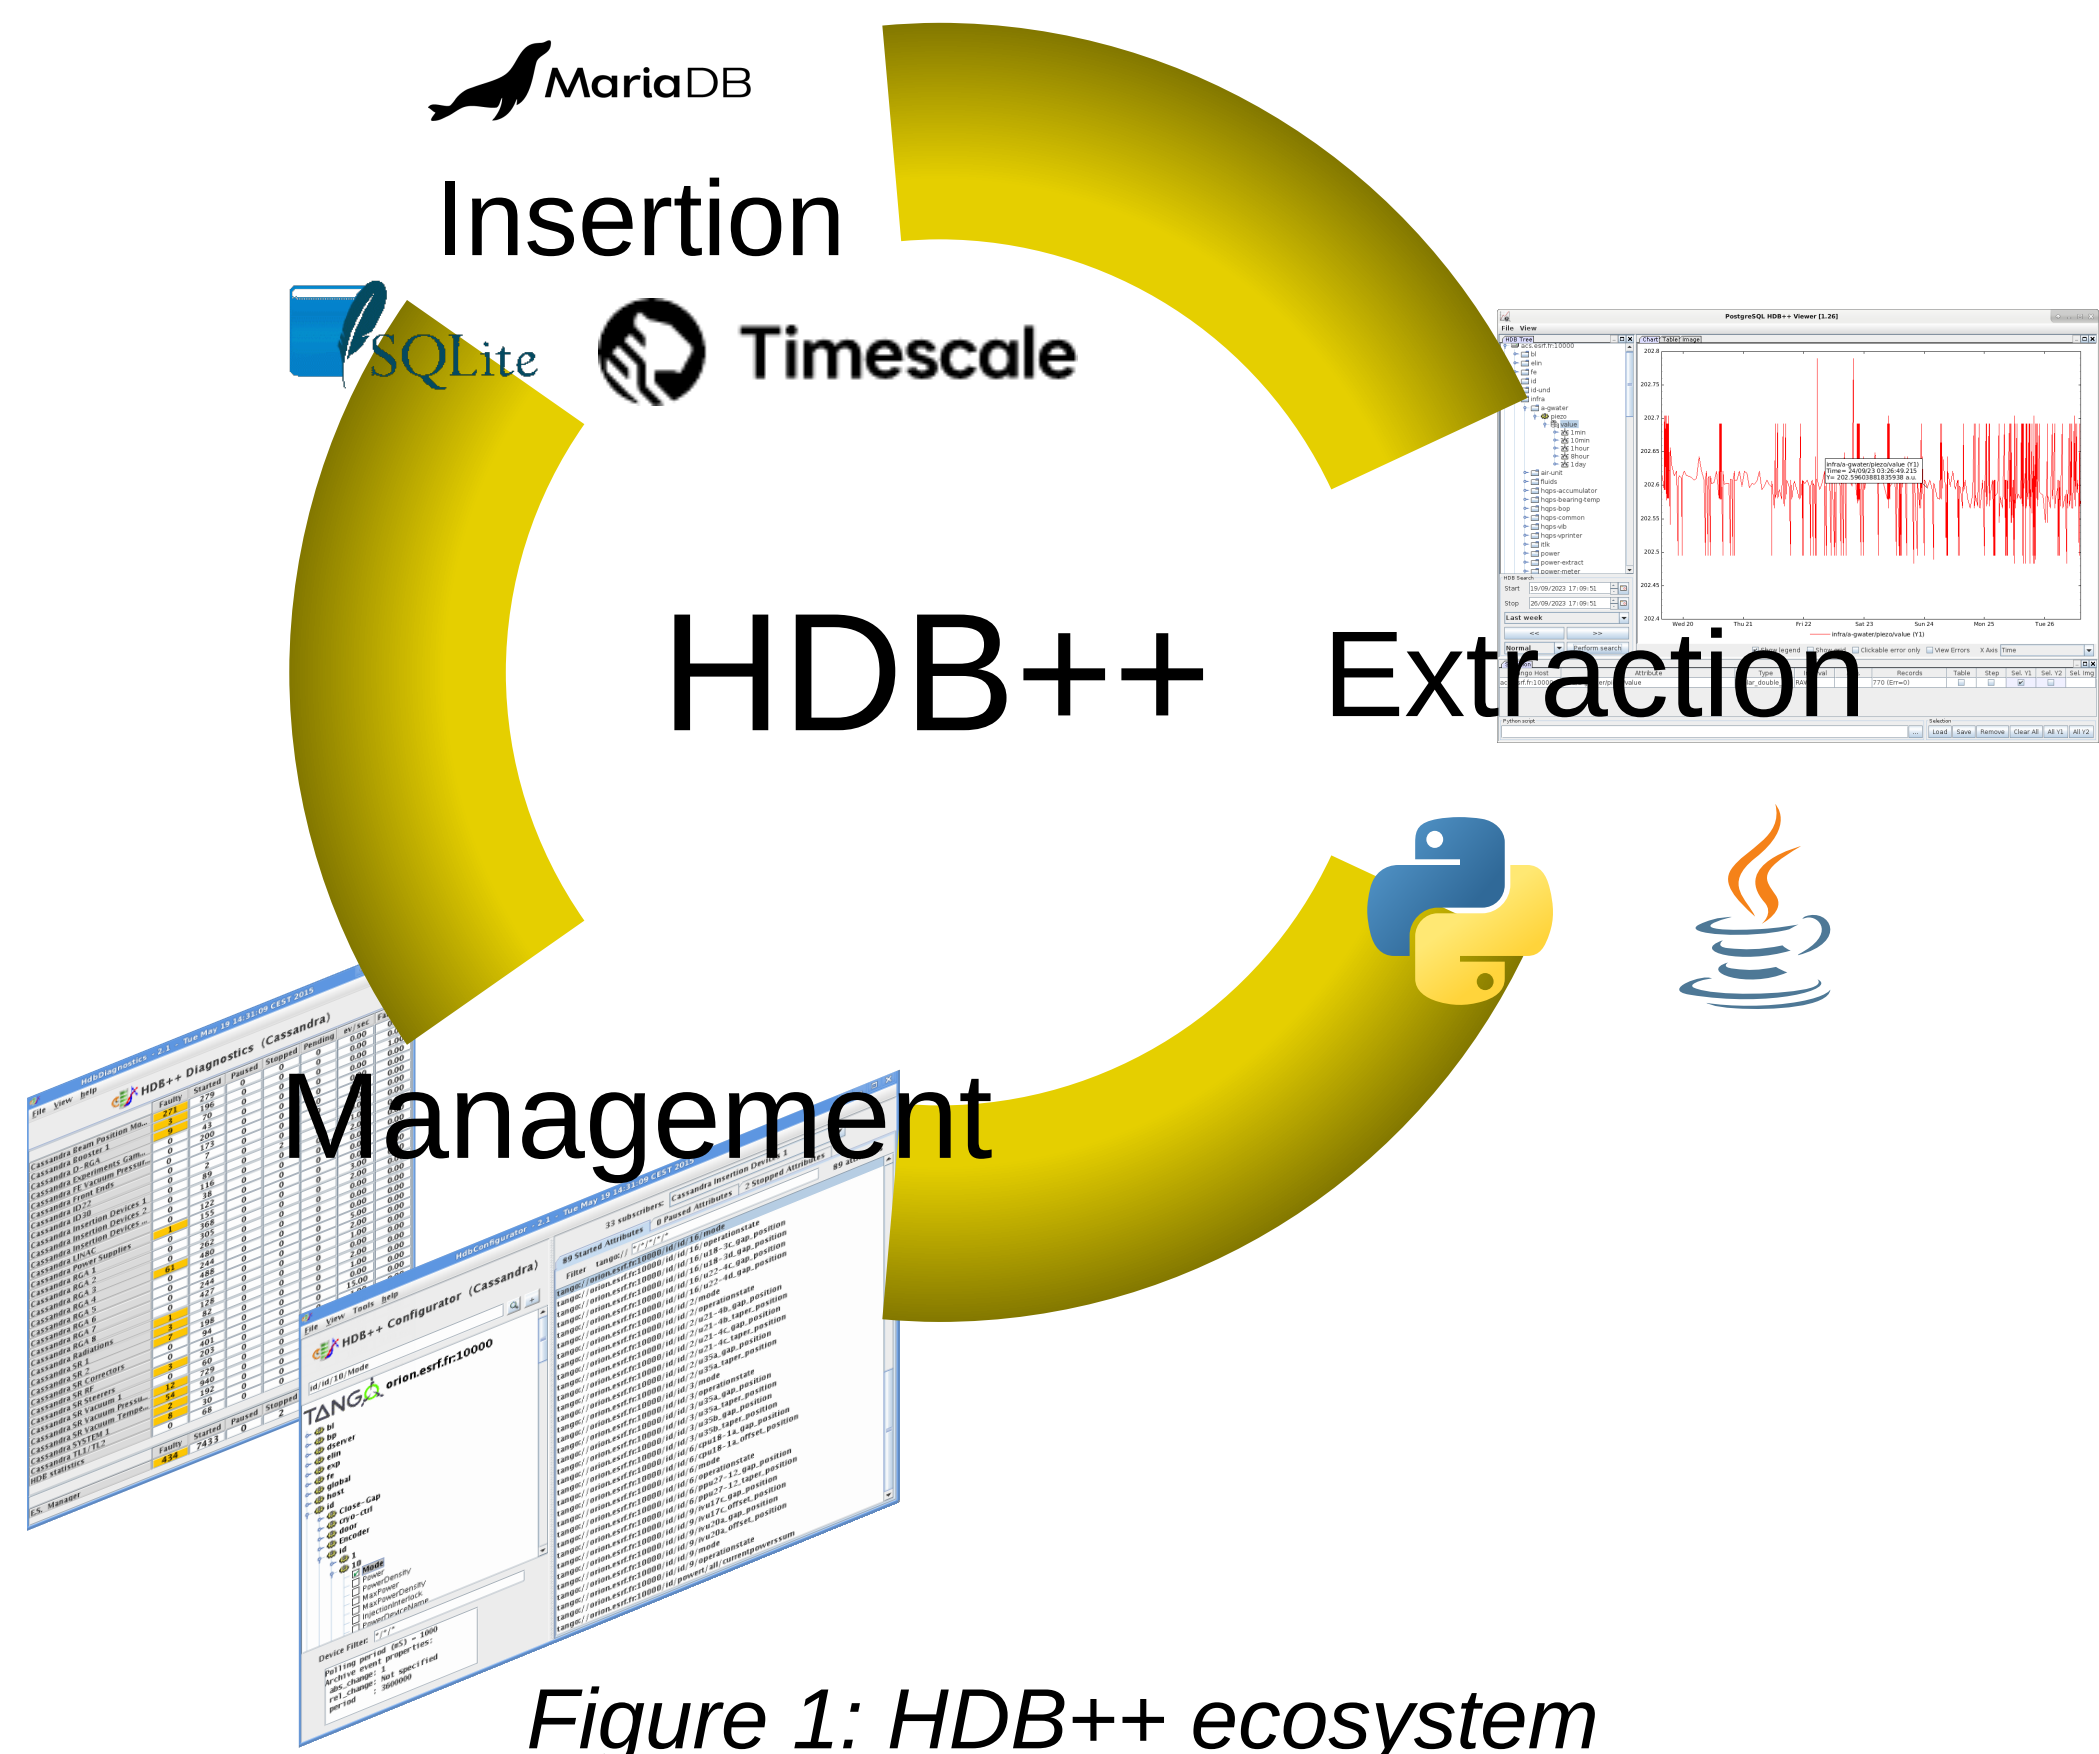
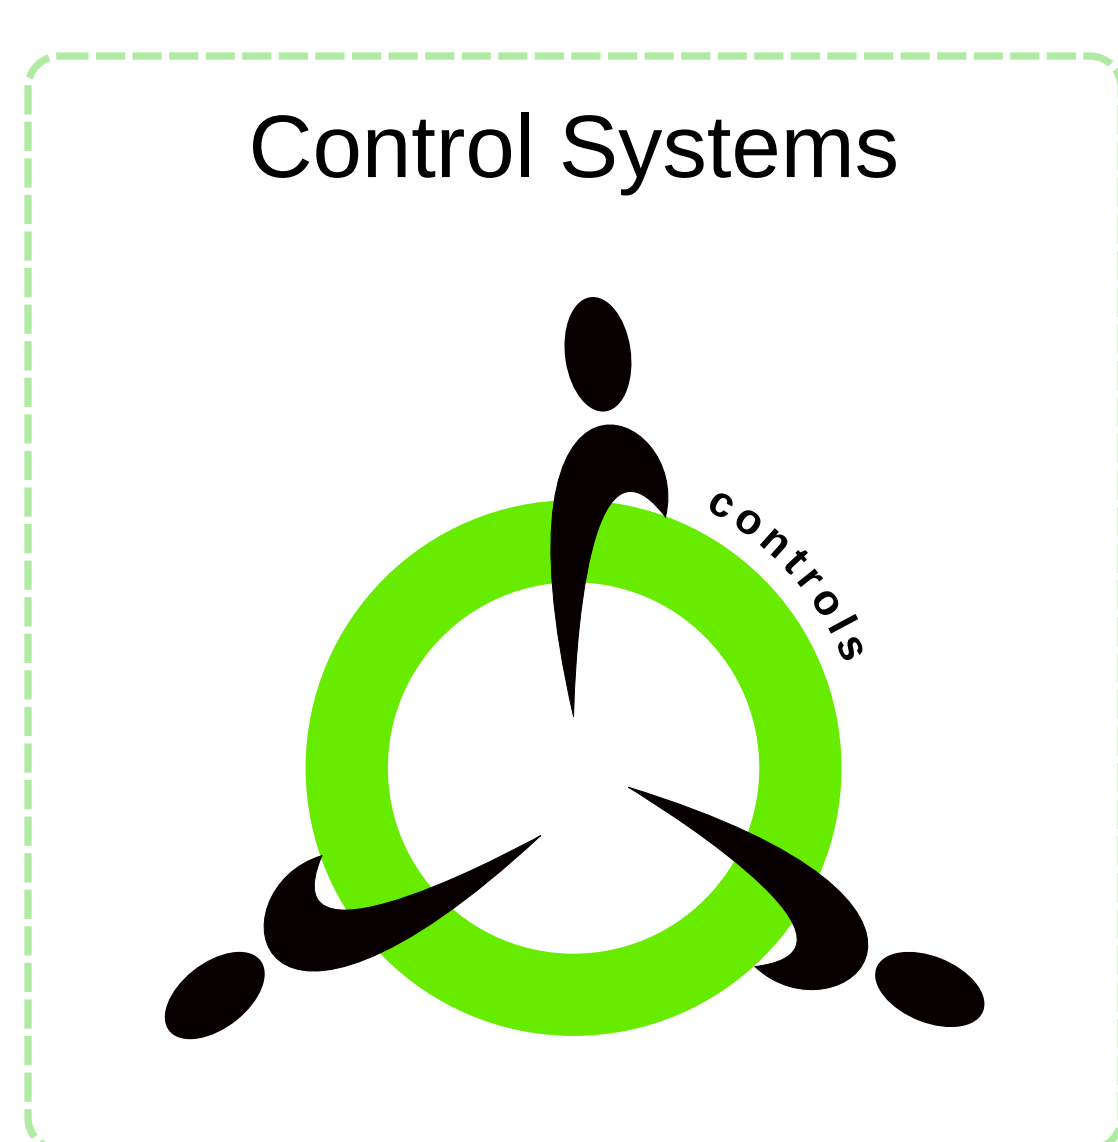


Figure 1: HDB++ ecosystem

The HDB++ project is a collaborative effort to build archiving on top of the Tango Control System. HDB++ provides a full ecosystem that leverage multiple database backends support to let you store your data however you want, but use a common set of tools for configuration, extraction and data management.

Data insertion is managed through Tango device servers, and you can get your data using your language of choice or directly use a viewer.



- 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

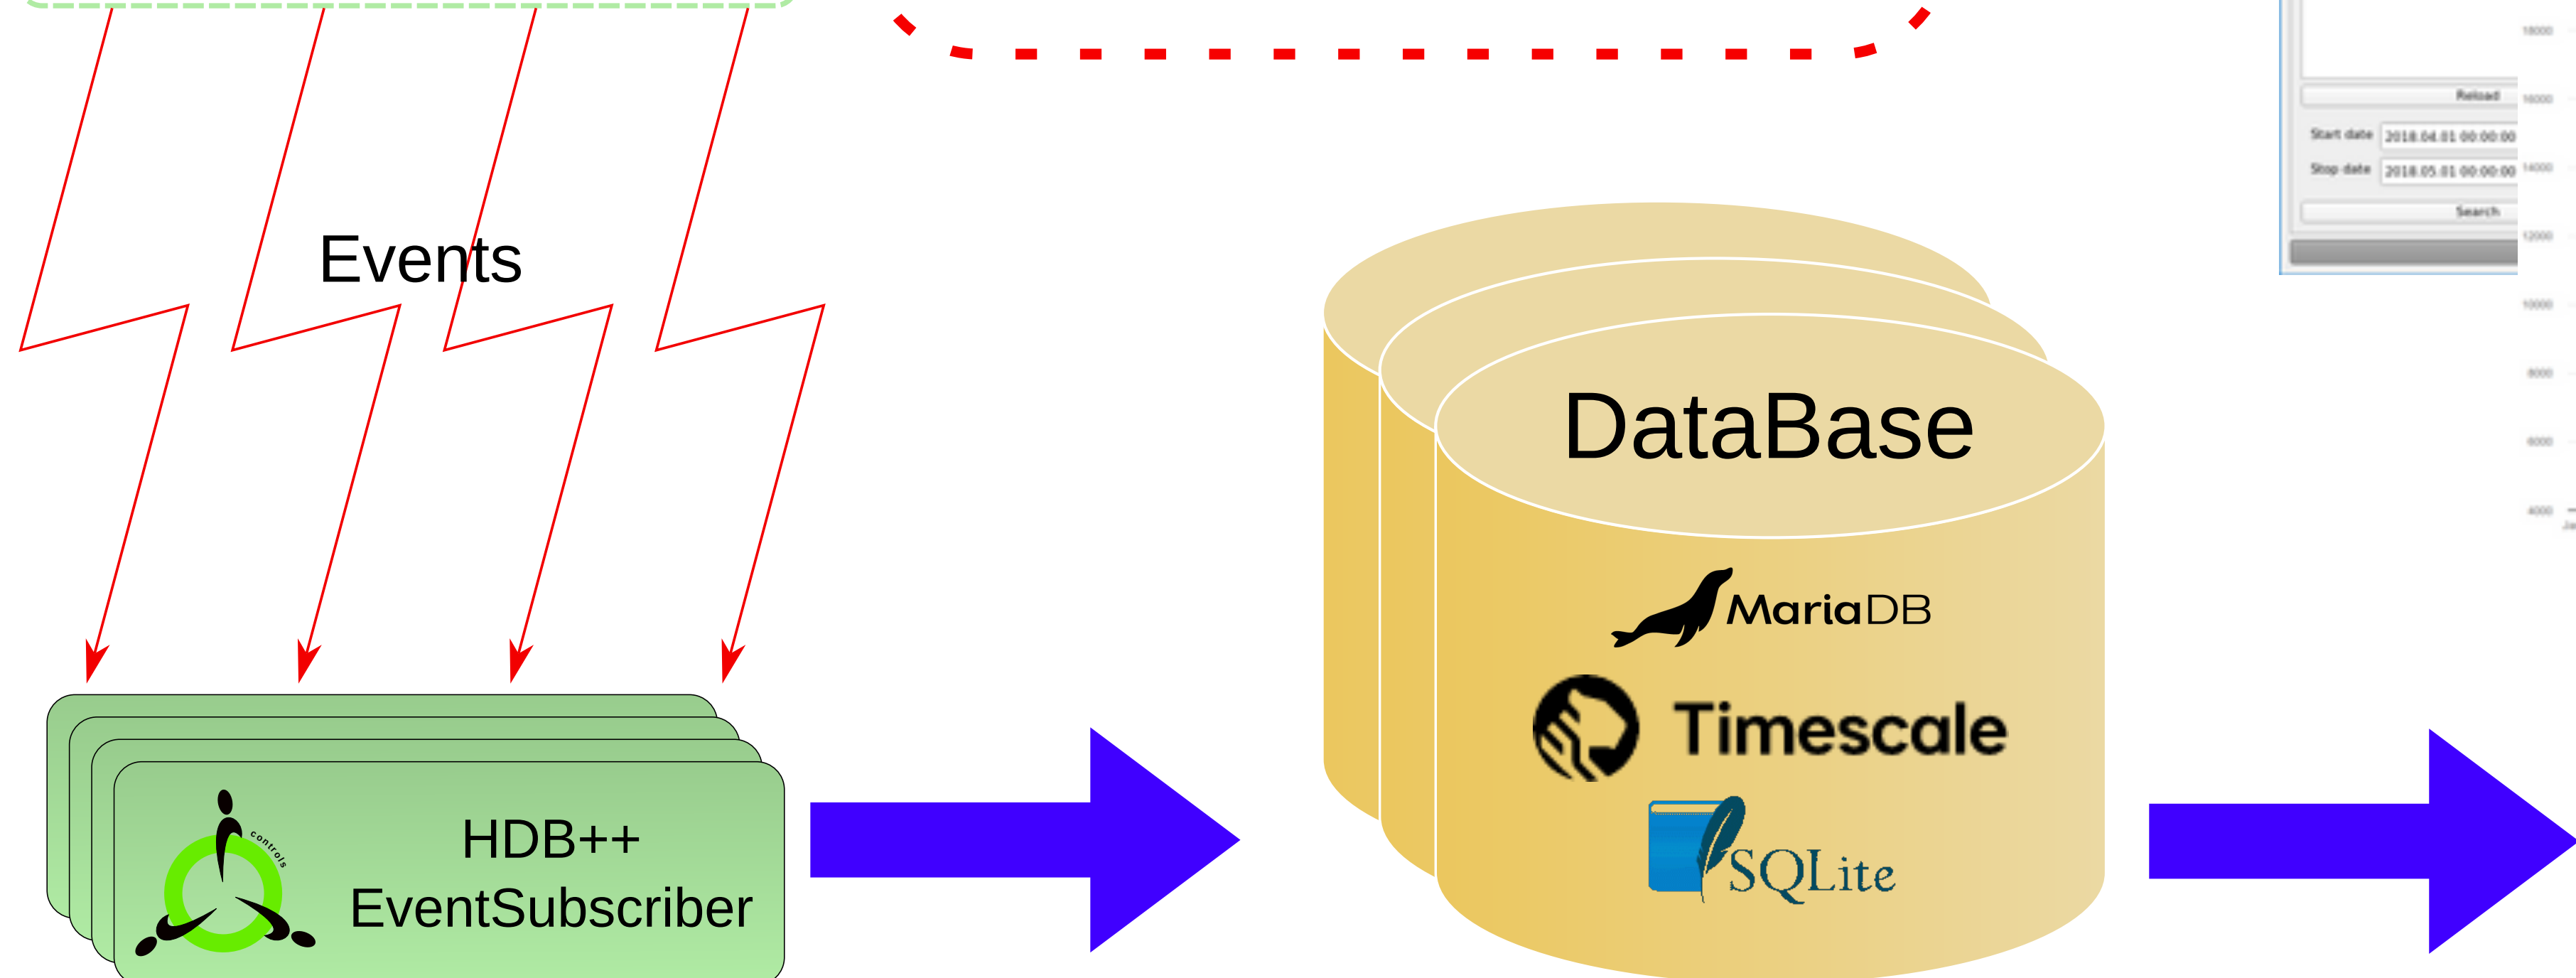
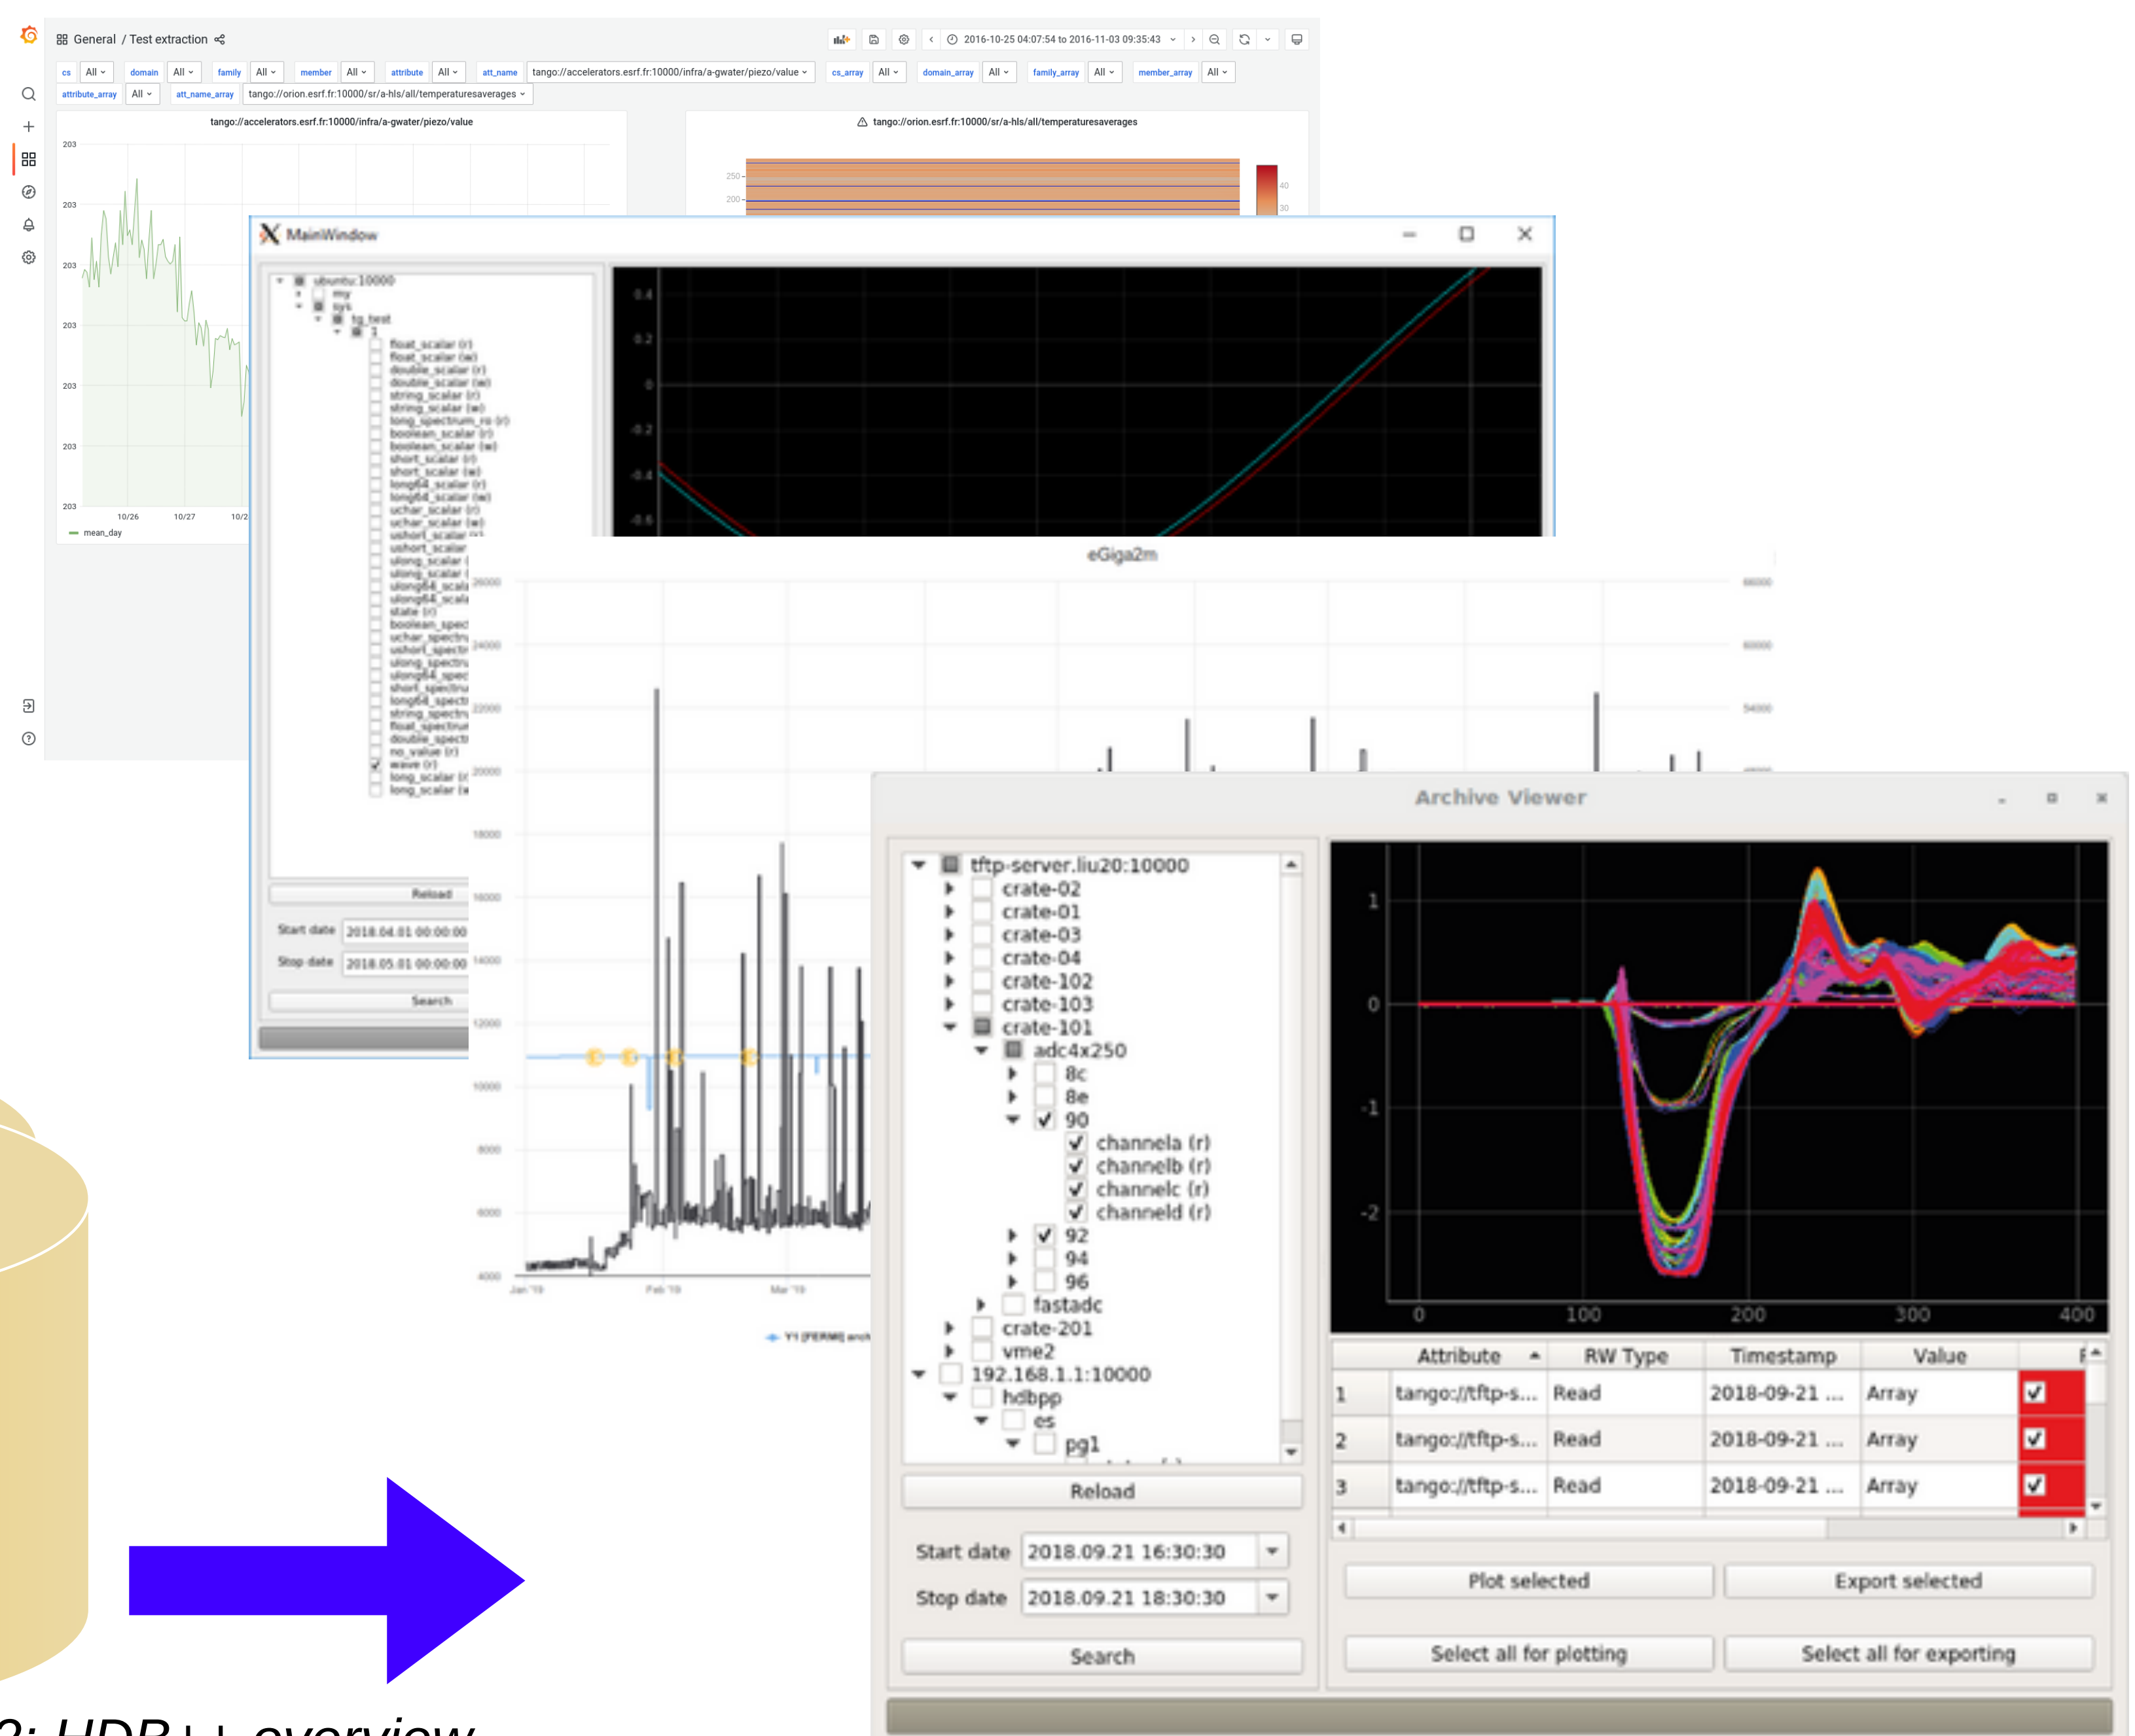


Figure 2: HDB++ overview



Long-Term data storage strategies:
Use the full power of long tested database backends to implement whichever long-term storage strategy you want.
HDB++ can use advanced data decimation, to reduce the number of samples, data compression inside the databases to further reduce disk usage, and several databases clusters to store all your data.
All these strategies can be used together or one at a time and allow to be able to keep valuable data for as long as necessary.

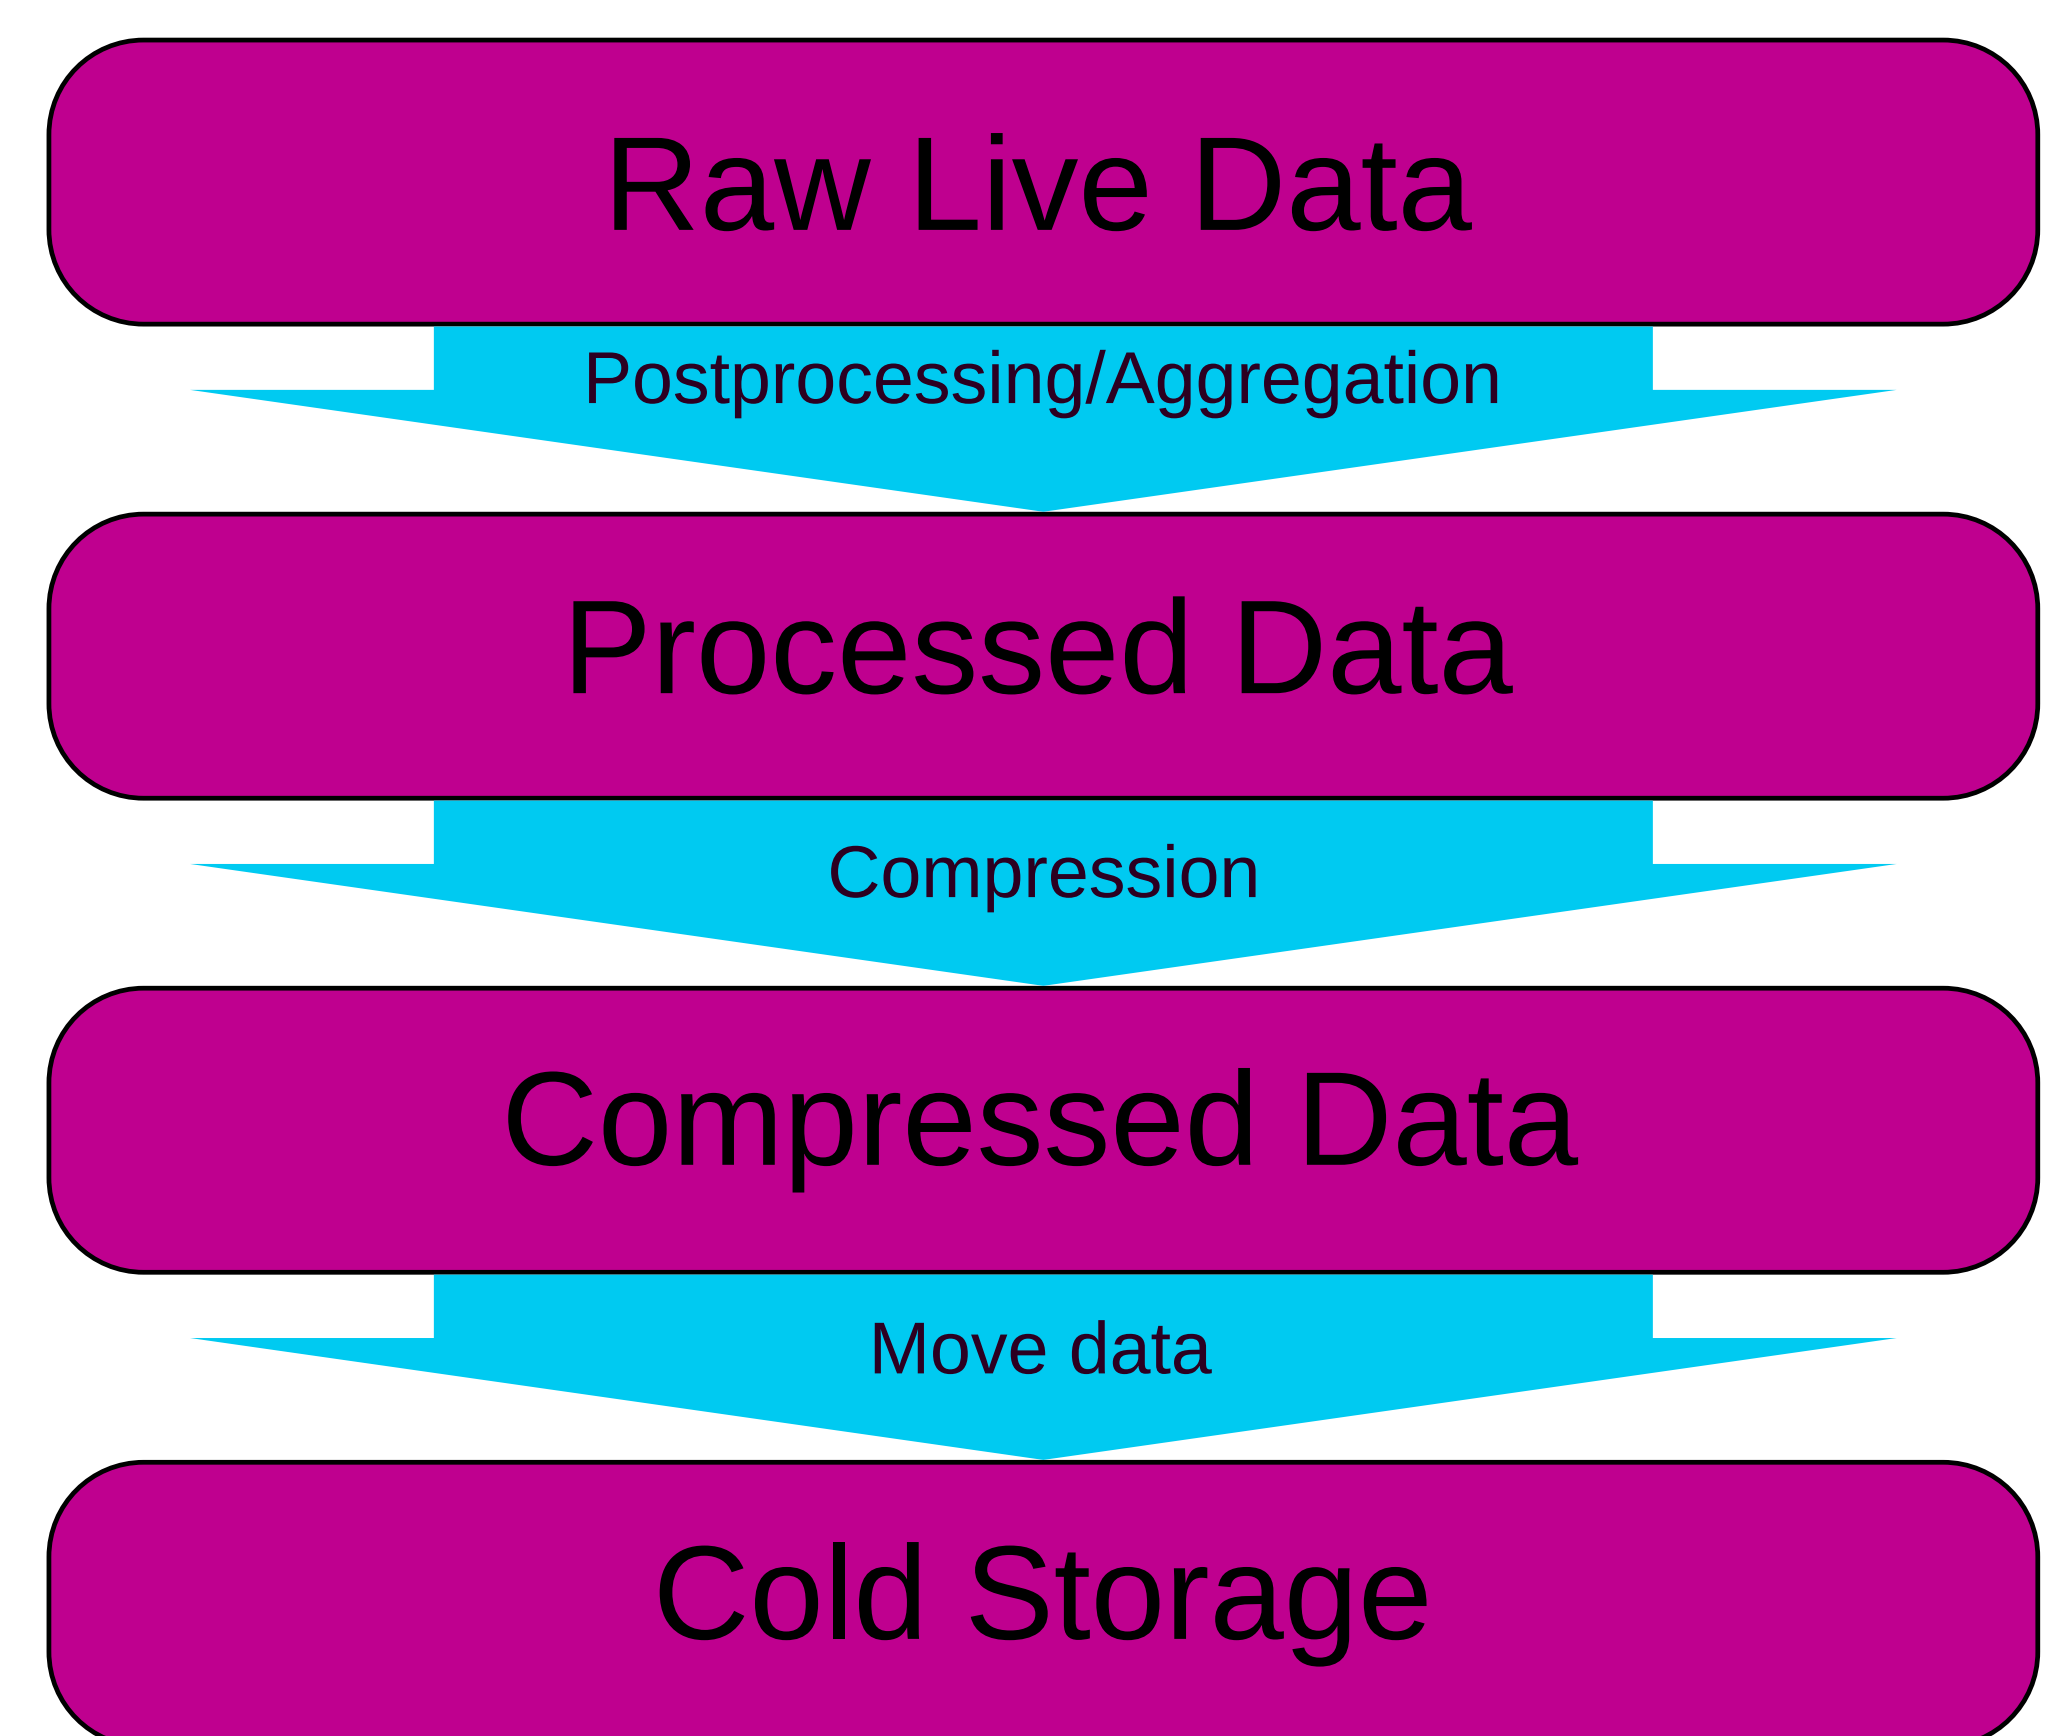


Figure 3: Long-Term archiving strategies

Acknowledgements:
Many thanks to the Tango HDB++ community for their great work and ideas

