



# PUSHING THE LIMITS OF TANGO ARCHIVING SYSTEM USING POSTGRESQL AND TIME SERIES DATABASES

R. Bourtembourg, S. James, J-L. Pons, P. Verdier, ESRF, Grenoble, France  
 G. Cuni, S. Rubio-Manrique, ALBA-CELLS Synchrotron, Cerdanyola del Vallès, Spain  
 G. A. Fatkin, A. Senchenko, V. Sitnov, BINP SB RAS, Novosibirsk; NSU, Novosibirsk, Russia  
 L. Pivetta, C. Scafuri, G. Scalamera, G. Strangolino, L. Zambon, Elettra-Sincrotrone Trieste S.C.p.A., Basovizza, Italy  
 M. Di Carlo, INAF - Osservatorio Astronomico d'Abruzzo, Teramo, Italy

The HDB++ Tango archiving system relies on the Tango archive events feature to collect Tango attributes values coming from one or several Tango Control Systems and then store these values in the Database back-end of your choice.

The following back-ends are currently supported: MySQL/MariaDB, Cassandra, PostgreSQL, TimescaleDB (a PostgreSQL extension) and Elasticsearch.

This list could be easily extended, thanks to the layered design of the HDB++ architecture.

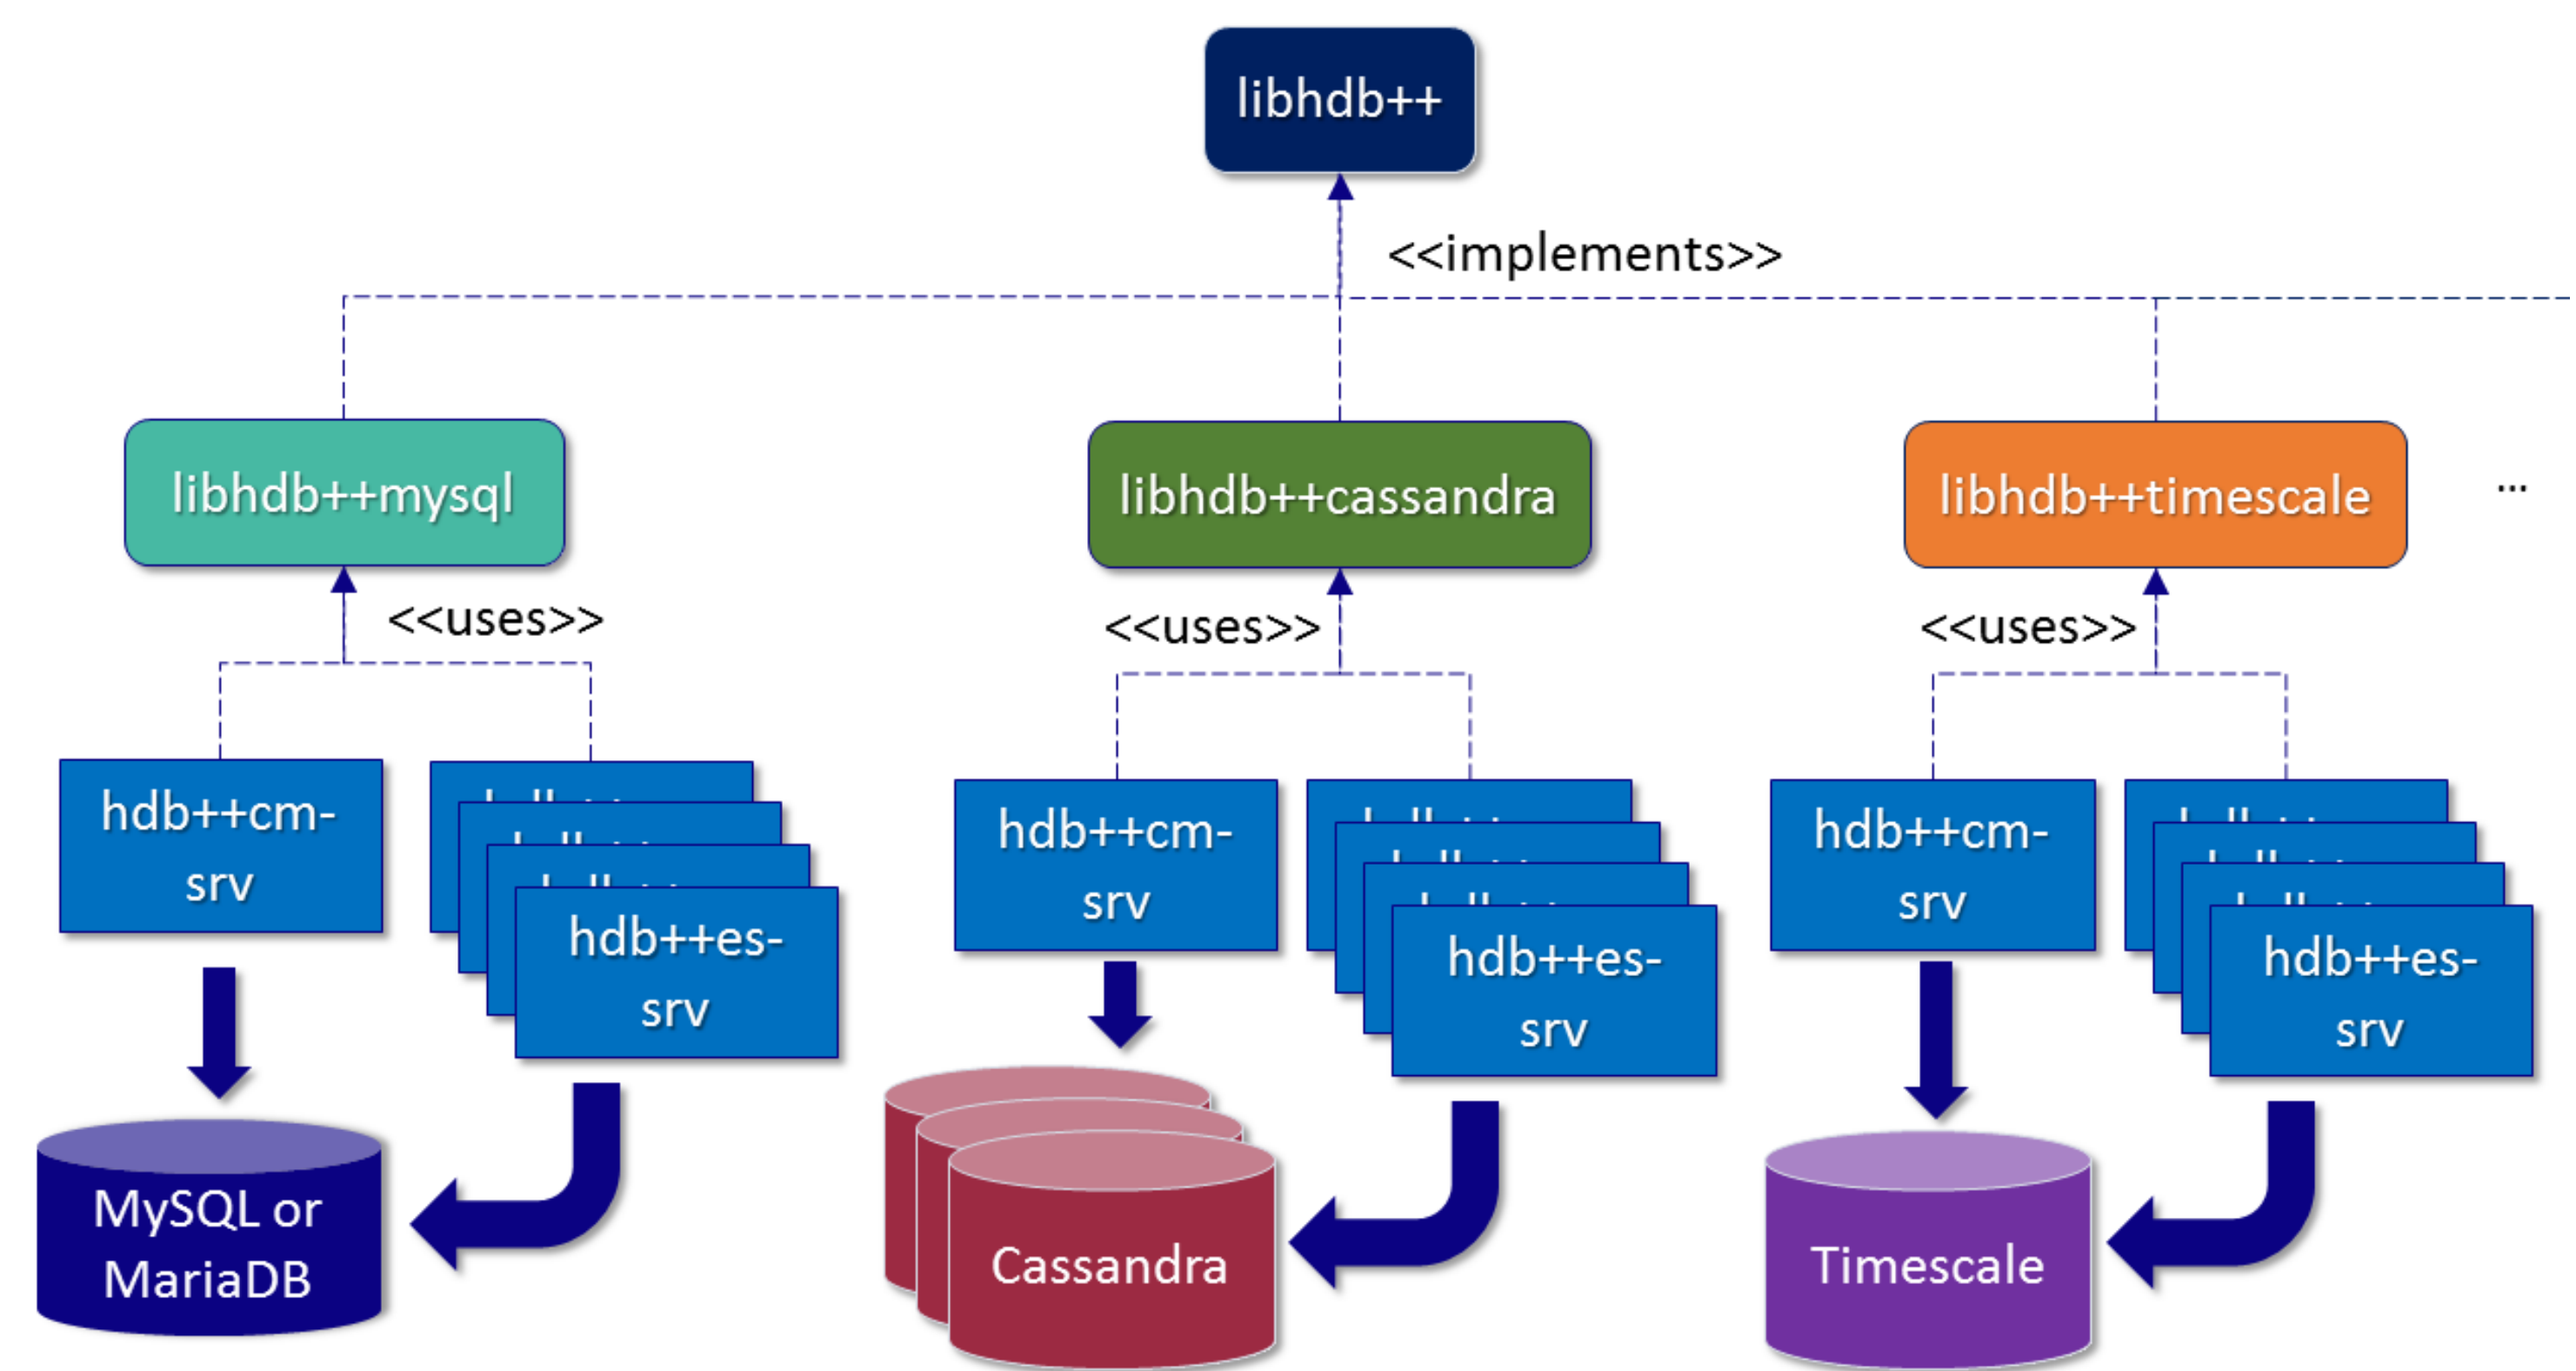


Figure 1: HDB++ Archiving Design

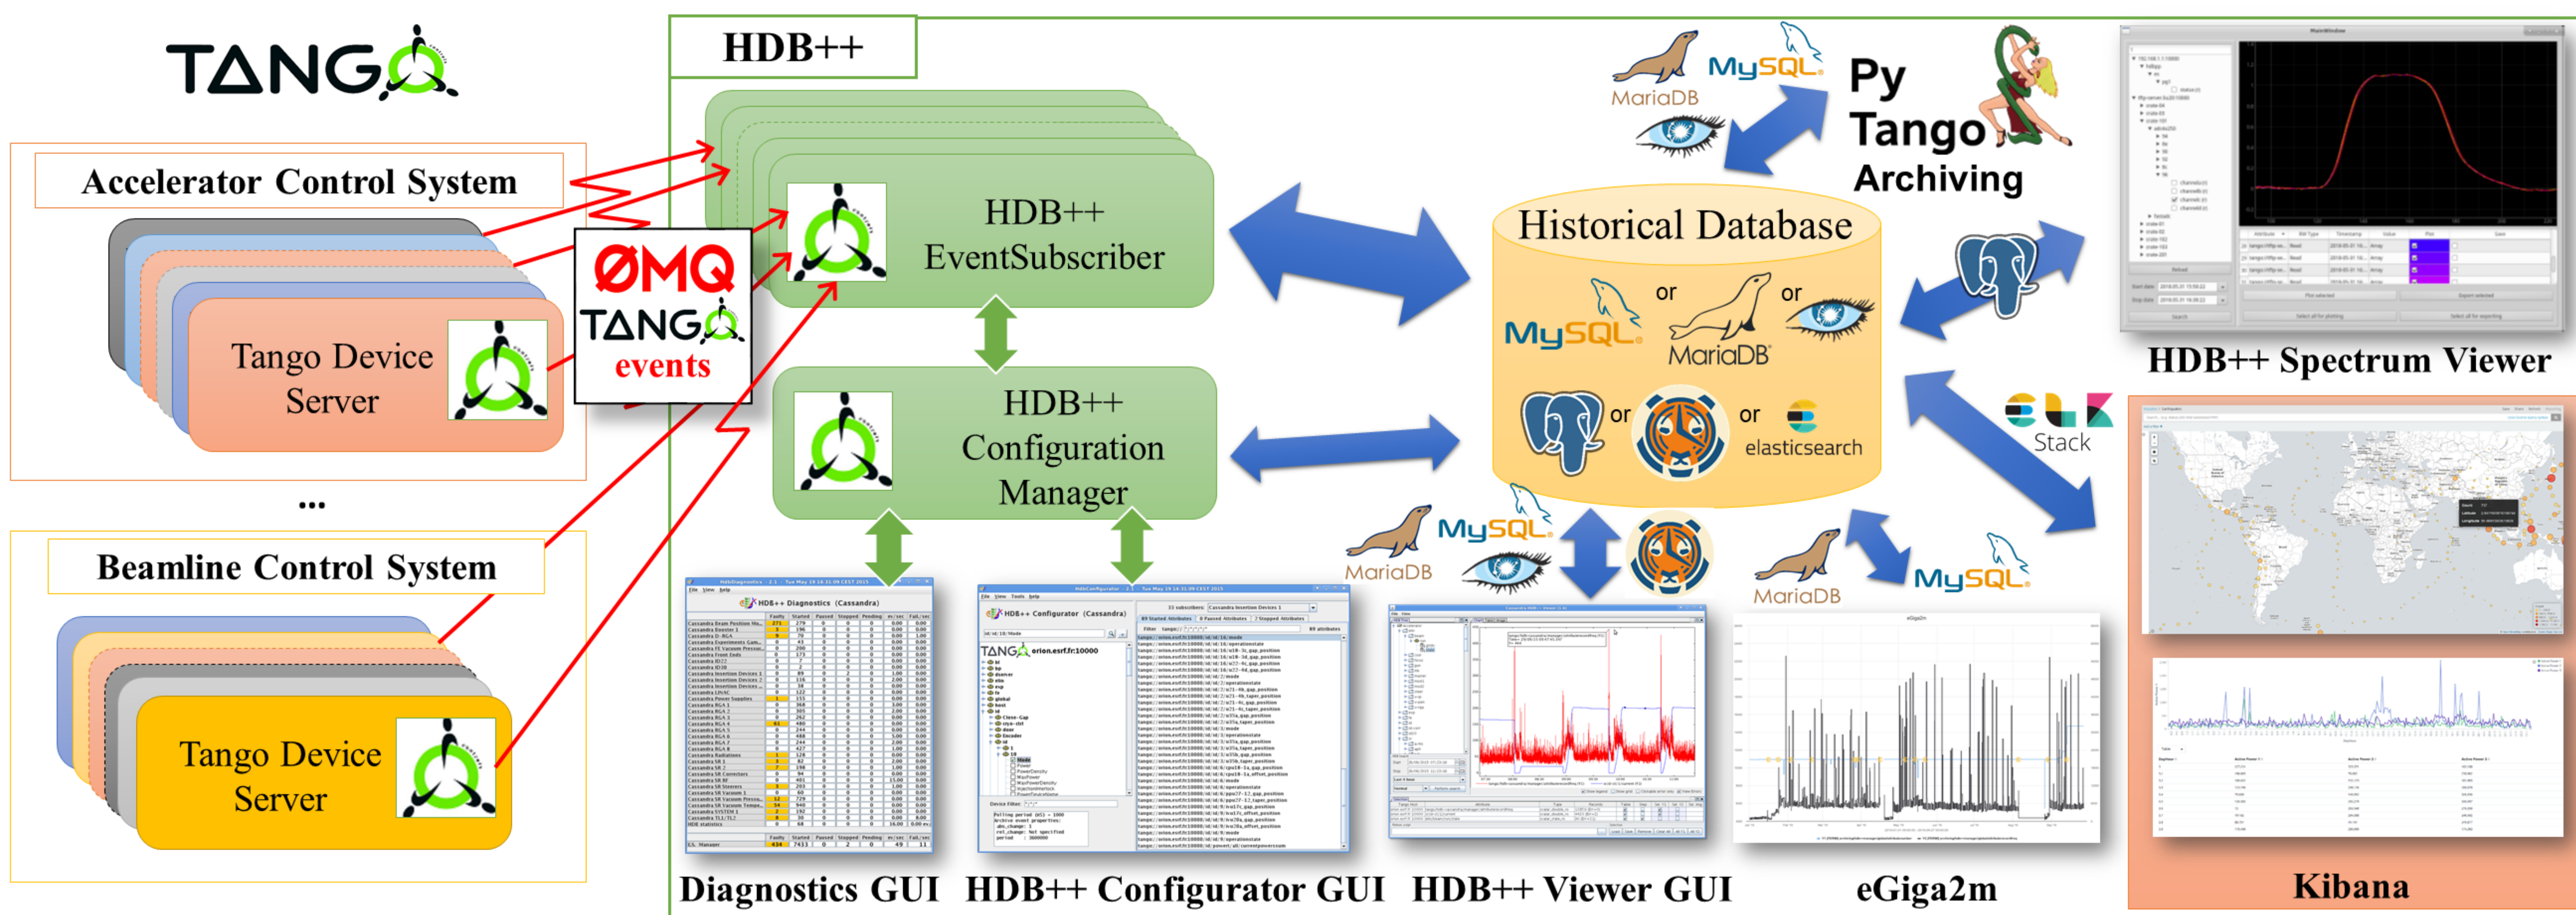


Figure 2: HDB++ Overview

| DB Backend | Pros  | Cons   |
|------------|---|--|
|            | Good for small DB<br>ProxySQL can be used to redirect queries to master slave or other MySQL/MariaDB DBs<br>Legacy DB schema compatible with old TANGO HDB tools      | Not optimized for arrays<br>Bad performance if DB is too big   |
|            | High Availability for write datacenter<br>Cassandra HDB++ backend could be used in theory with ScyllaDB   | Not good with arrays<br>Big queries could bring down several nodes because of Java Garbage Collector |
|            | Native support for arrays, so good performance of queries involving arrays<br>Ability to query individual elements of an array  | Bad performance if DB is too big   |
|            | Same as PostgreSQL + optimized for time series, scalable, automatic hypertables creation, extended API for time series, automatic continuous aggregation, gap filling | Chunk level data reordering operation<br>MUST be run regularly to guarantee good query performance   |
|            | Advantages of ELK stack (Kibana viewers,...)<br>Flexible schema   | Requires a lot of memory<br>No security  |

Table1: Comparison of the different supported HDB++ backends

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

