HDB++, a retrospective on 5+ years using Timescale

D. Lacoste, R. Bourtembourg, ESRF, Grenoble, France
S. Rubio-Manrique, J. Ramos, ALBA-CELLS, Cerdanyola del Vallès, Spain
L. Pivetta, G. Scalamera, Elettra-Sincrotrone Trieste S.C.p.A., Basovizza, Italy
J. Forsberg, D. Egorov, MAX IV Laboratory, Lund, Sweden
T. Juerges, SKA Observatory, Jodrell Bank, United Kingdom
G. Jourjon, CSIRO, Canberra, Australia

HDB++ Overview

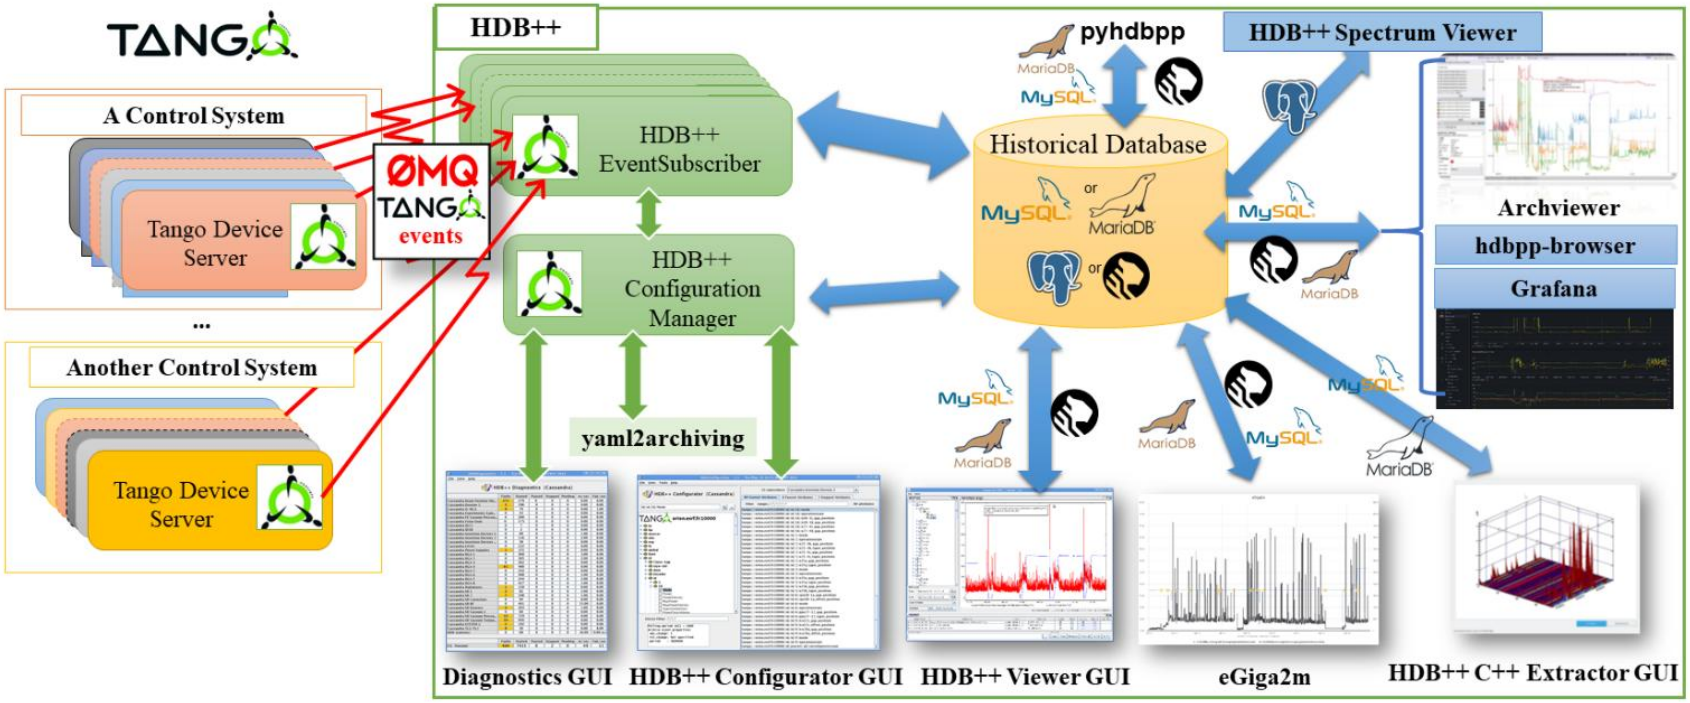

Timescale Deployment @ ESRF

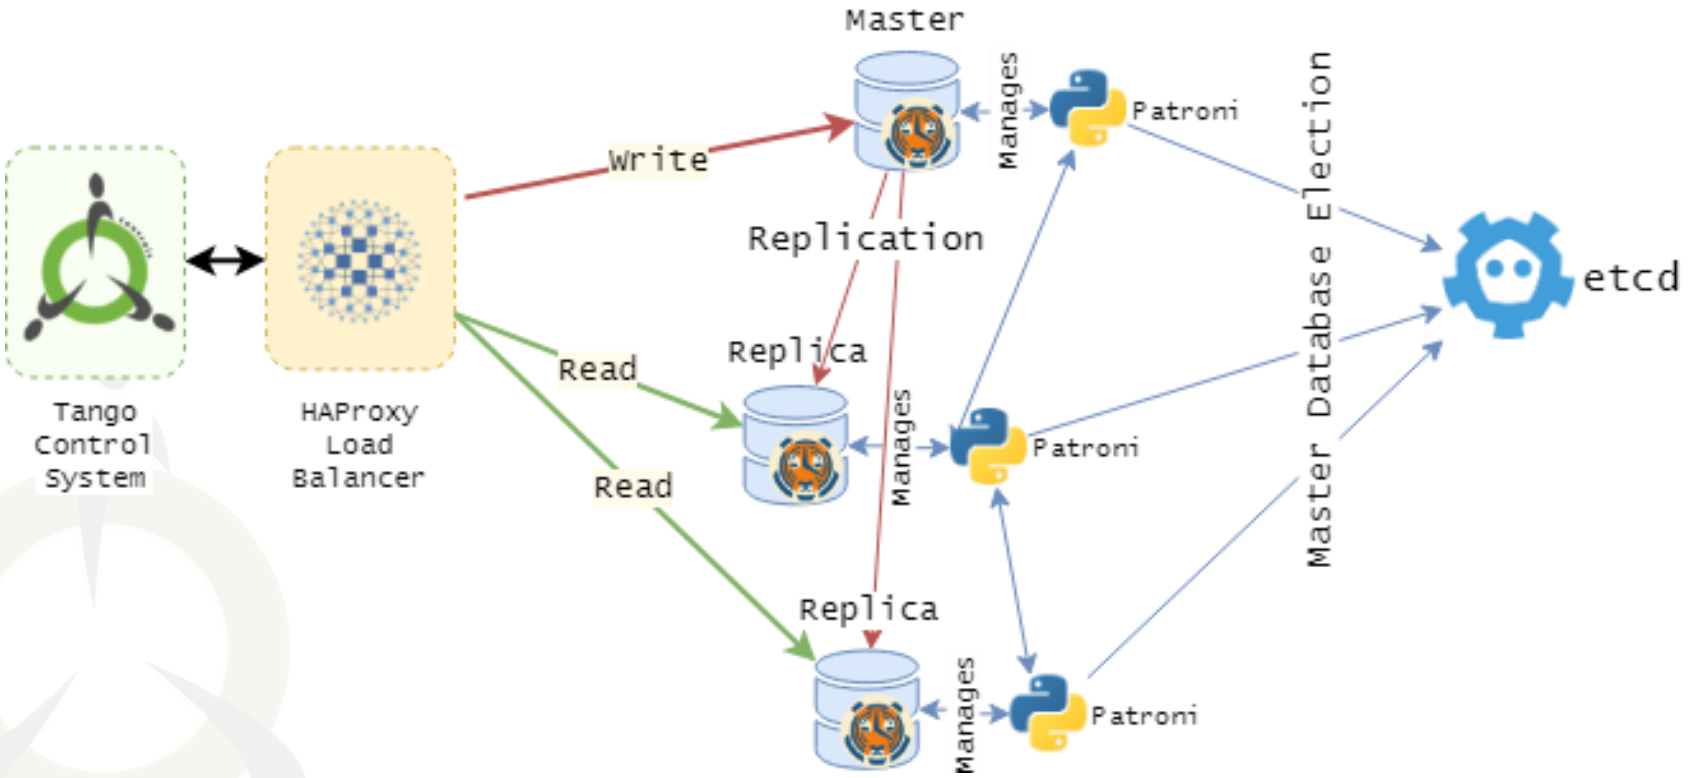

5 years using Timescale

- ➔ Timescale works **great** and is reliable, considered as default HDB++ DB Backend nowadays.
- ➔ Use the Timescale compression!
  But DELETE/UPDATE not possible in compressed chunks when using Timescale < 2.11 / PostgreSQL < 14
- ➔ **Great automation** features (aggregates, sorting,) and **great** performance
- ➔ HDB++ DB Backend abstraction is **great**!
- ➔ Migration from Cassandra to Timescale backend done at ESRF and MaxIV
- ➔ HDB++ Community is **great**!
- ➔ Many thanks to all the **great** HDB++ contributors!

For more details please come and see the **great** poster **TUPD054!**