

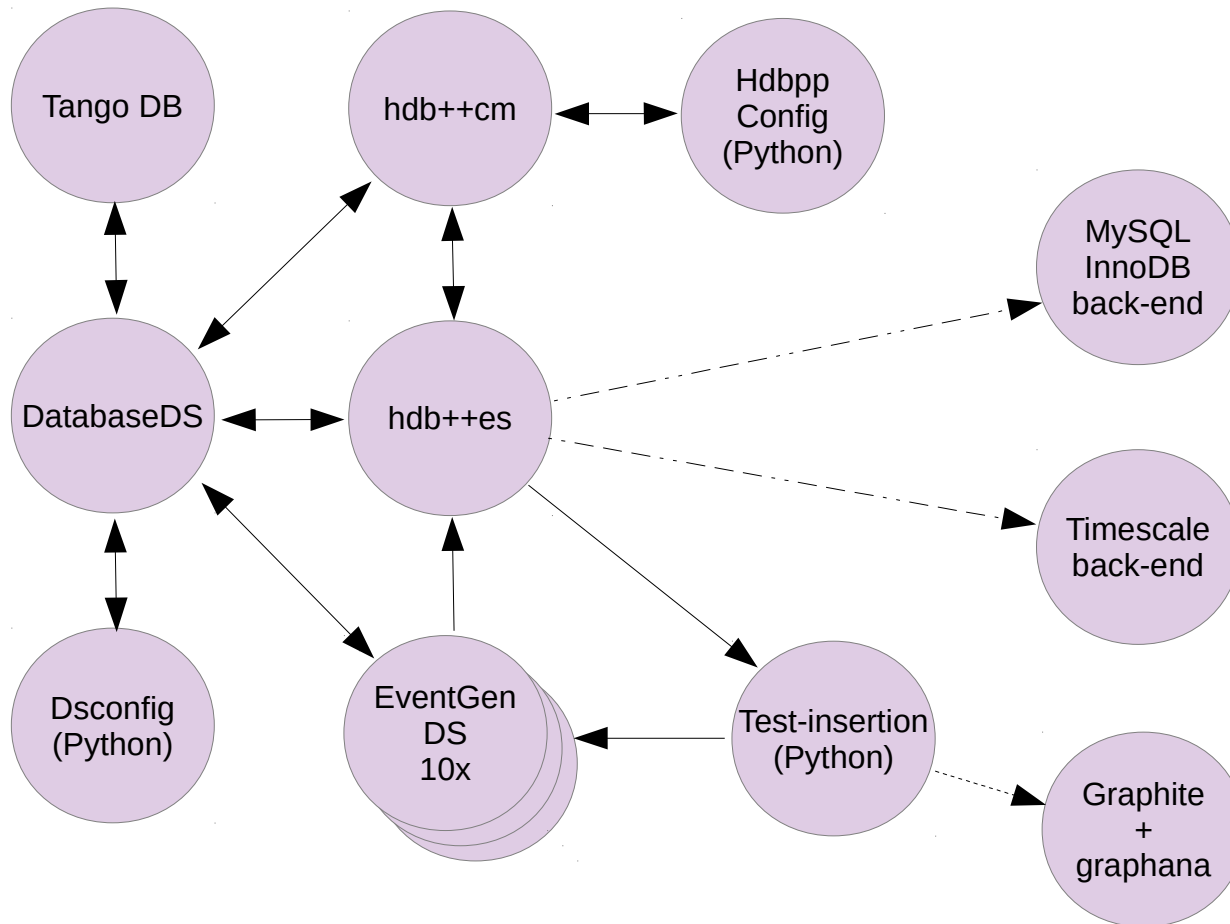
HDB++ status

Lorenzo Pivetta

on behalf of the HDB++ collaboration

- <https://github.com/tango-controls-hdbpp>
- Four meetings during 2020 (telco)
 - Last meeting: 2020.11.06
 - Improve documentation on readthedocs
 - Support TimescaleDB in PyTangoArchiving
 - Revisit work on the ELK back-end
- HDB++ still complex to build
 - Get stable master on github
 - Provide a guide to build with cmake
- Require C++14
- Hdbpp-benchmark
 - Preliminary results with batch insertion – hit ZeroMQ HWM

Benchmarking setup (Docker)



TimescaleDB

Max event rate ~1700 ev/s

Number of events archived **1341288**

Disk increase of DB partition **425 MB**

Test Duration **1700** seconds

MySQL/InnoDB

Max event rate ~4000 ev/s

Number of events archived **3463359**

Disk increase of DB partition **164 MB**

Test Duration **2500** seconds

Preliminary tests with batch insertion:

TimescaleDB / MySQL InnoDB

Max event rate ~12000 ev/s

(hit ZeroMQ HWM)

Additional investigation in progress

Fermi

- Running since 2015
- ~12300 attributes from 8 Tango facilities
- ~5000 ev/minute mean; peaks up to 48.5K ev/minute
- Context based archiving -> ~30 archiving strategies defined
- 54 EventSubscriber + 5 ConfigurationManager
- 1 MySQL back-end
- ~450 GB on disk

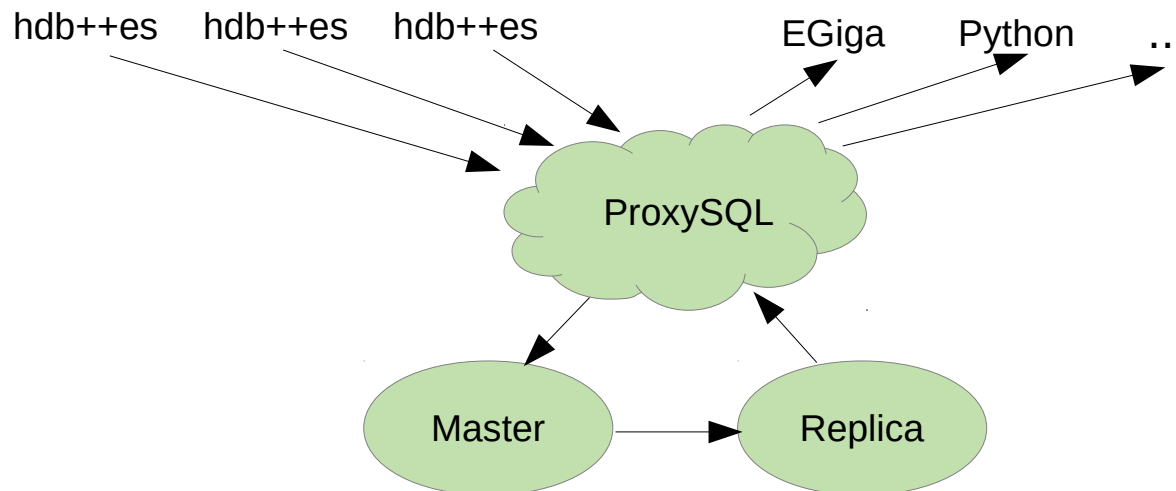
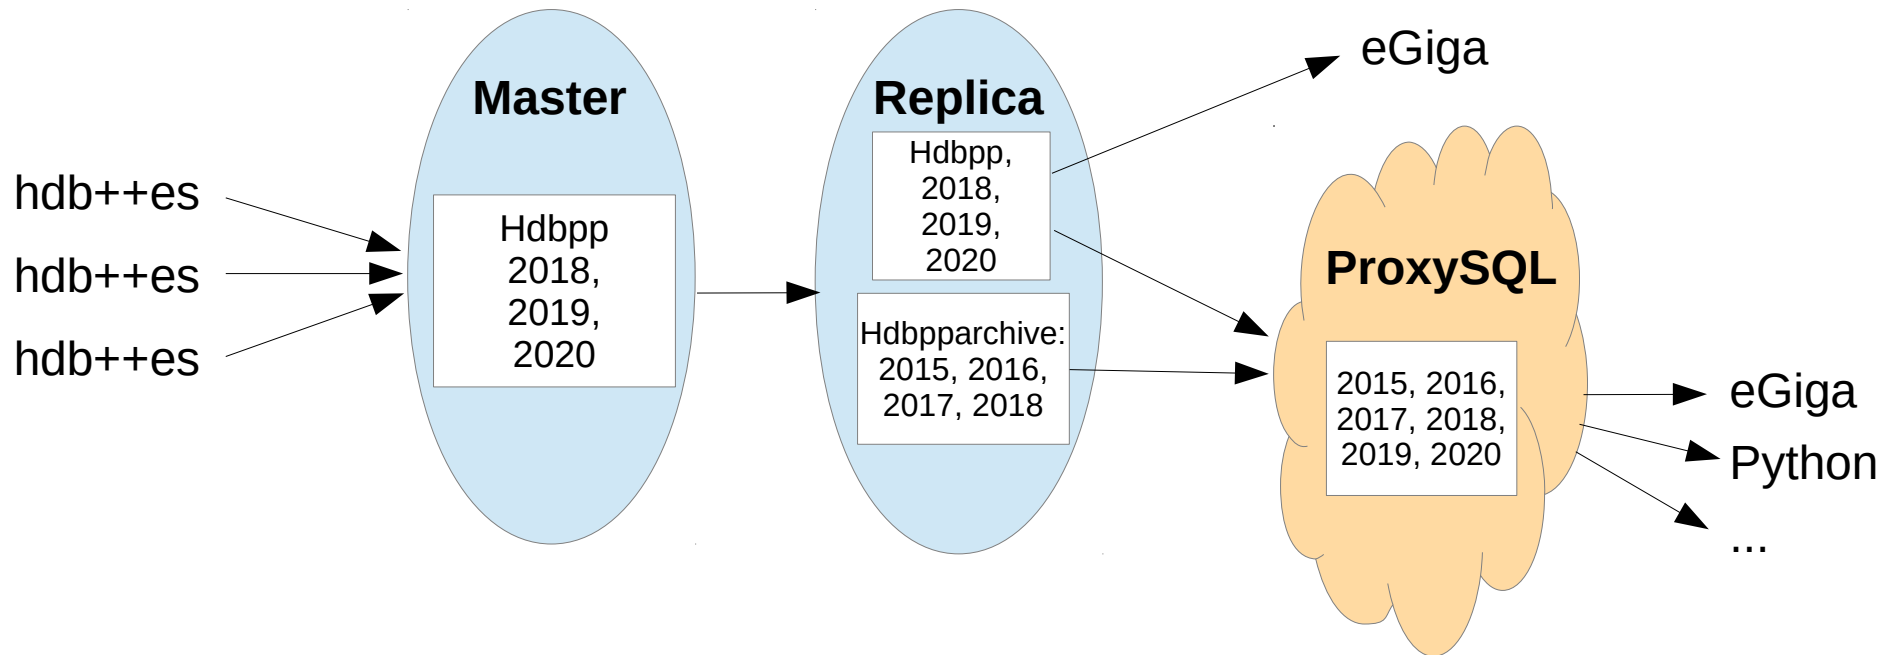
Elettra

- Running since 2016
- ~2650 attributes
- ~1850 ev/minute mean
- Legacy HDB schema, some Java HDB still archiving on the same DB
- 15 EventSubscriber + 1 ConfigurationManager
- 1 MySQL back-end
- ~700 GB on disk (includes old data)

Infrastructure (buildings facility)

- ~200 attributes (new, growing to ~1000)
- 1 MySQL back-end

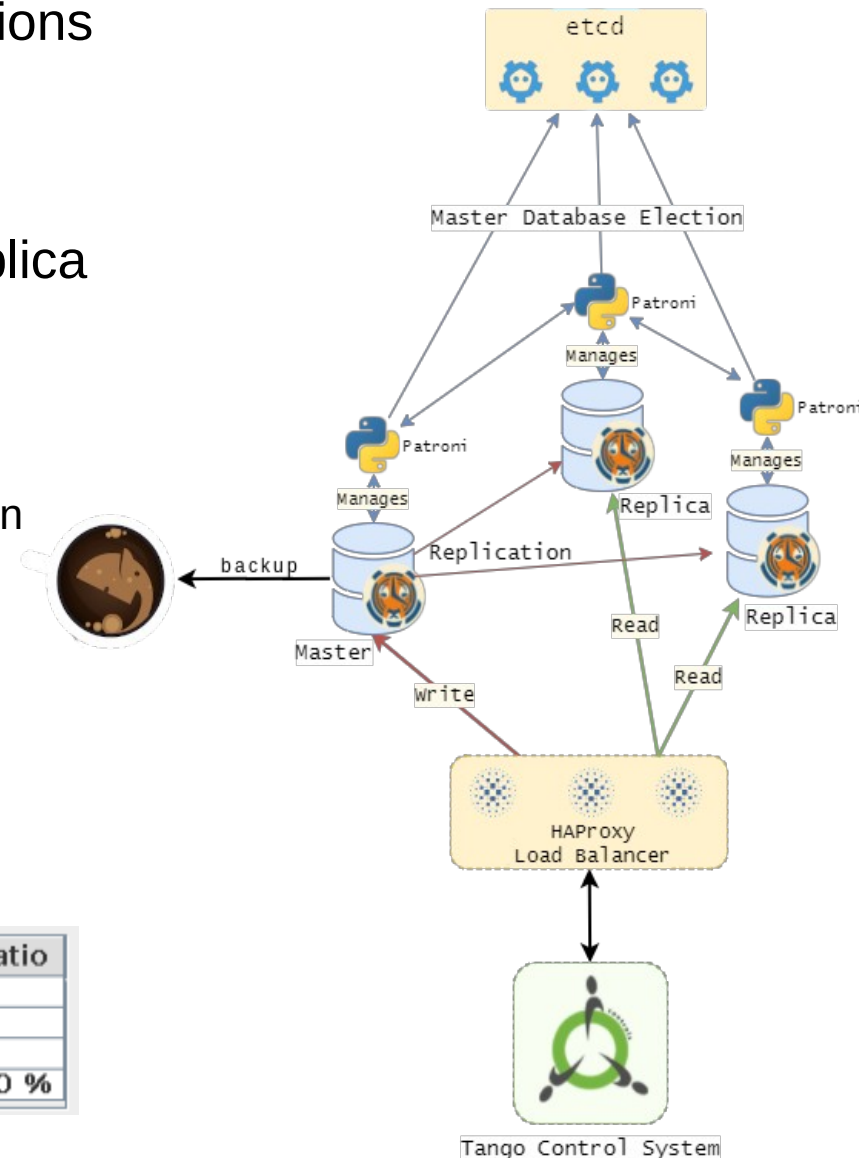
HDB++ @ Elettra

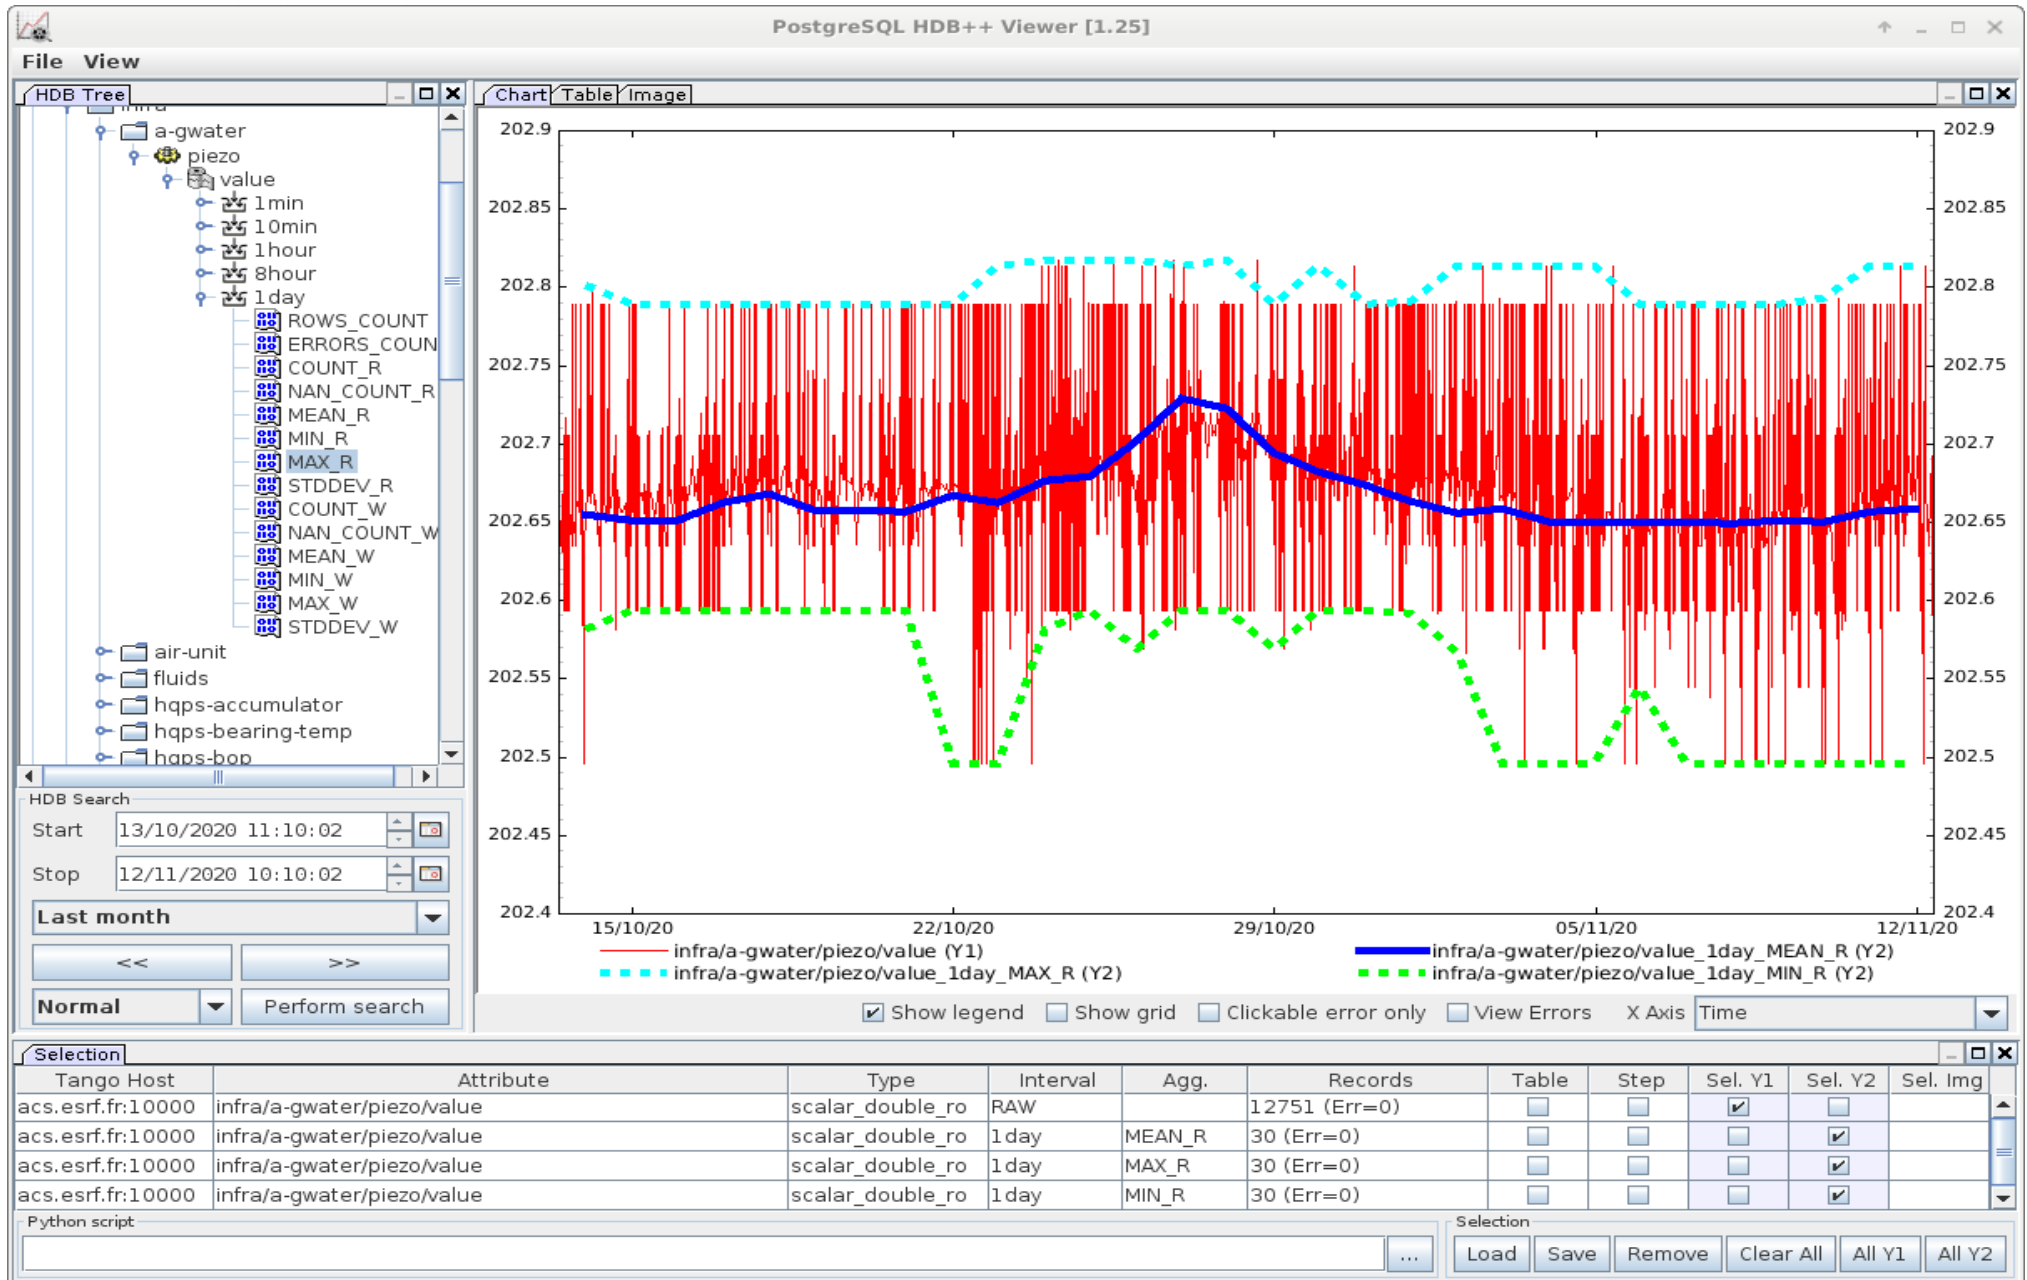


- Query caching
- Application layer proxy
- Fail-over support
- Query routing
- Query rewrite
- Advanced topology

- High availability design based on recommendations by TimescaleDB
- Fully redundant, no single point of failure
- One entry point via HA proxy
- Our design uses a Master and one or more Replica nodes
 - All writes to Master via proxy
 - All reads to Replicas via proxy
 - Reduces load on Master when users query for data
 - An extra replica is solely used for backup using barman
- Features
 - TTL
 - Data aggregation on scalar and array
 - Enum support
- Key figures
 - Database size 1.6 TB
 - Insert/s ~500

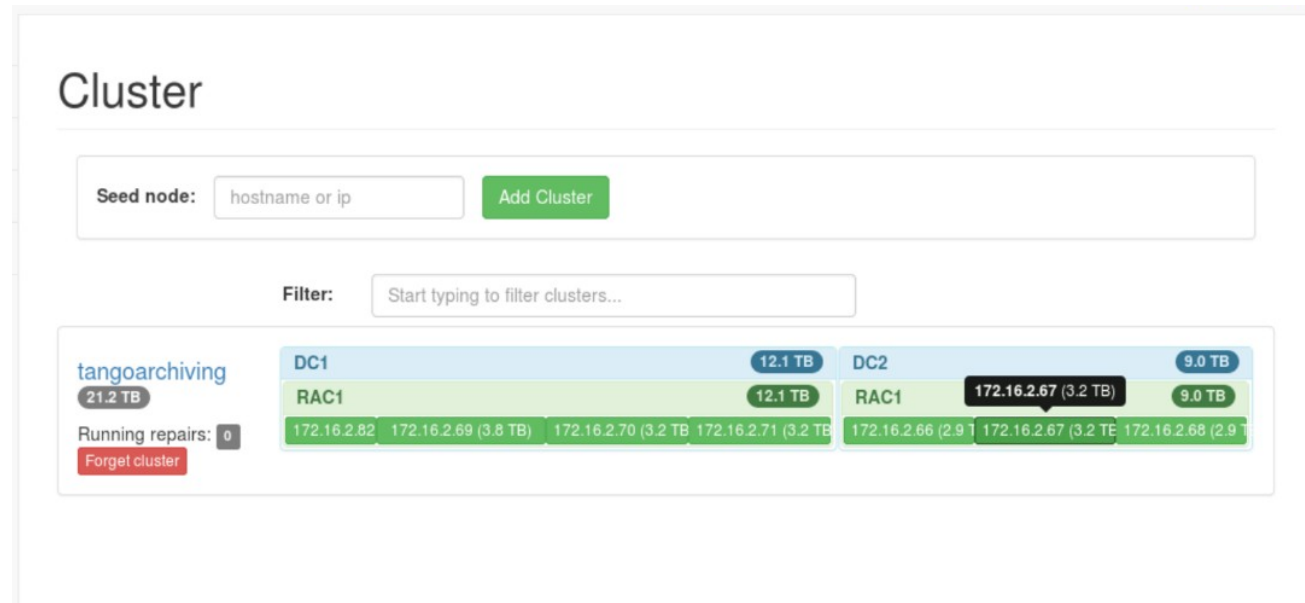
	Managers	Subscribers	Attributes	Attr Ratio
Accelerator	1	31	8898	92.50 %
ID	27	27	582	6.05 %
D	15	15	139	1.45 %
TOTAL	43	73	9619	100.00 %





- SKA-docker project contains the Docker files for running the archiver on a Mariadb backend
 - <https://gitlab.com/ska-telescope/ska-docker>
- One helm chart developed for running the HDB++ application
 - The result k8s application is composed by the following services
 - One configuration manager
 - One event subscriber
 - Mariadb
- Automatic testing implemented in BDD
- Future work:
 - Implement scalability mechanisms
 - Updates to the latest changes for the libraries
 - Have a noSql backend database (Elasticsearch)

- Central Cassandra database
- HDB++ running in accelerators and 14 beamlines
- Close to 5 TB of data
- Evaluating migration to TimescaleDB
 - + Better scalability
 - + Better maintainability
 - + Community support
 - Data migration
 - New client development



The screenshot shows the HDB++ Cluster management interface. At the top, there is a 'Cluster' header. Below it, there is a 'Seed node' input field with the placeholder text 'hostname or ip' and an 'Add Cluster' button. A 'Filter' input field with the placeholder text 'Start typing to filter clusters...' is also present. The main content area displays a table of clusters. The first cluster is 'tangoarchiving' with a total size of 21.2 TB and 0 running repairs. It has a 'Forget cluster' button. The table is organized into two data centers: DC1 and DC2. DC1 has a RAC1 with 12.1 TB and three nodes: 172.16.2.82, 172.16.2.69 (3.8 TB), and 172.16.2.70 (3.2 TB). DC2 has a RAC1 with 9.0 TB and three nodes: 172.16.2.66 (2.9 TB), 172.16.2.67 (3.2 TB), and 172.16.2.68 (2.9 TB).

Cluster	DC	RAC	Size	Nodes
tangoarchiving			21.2 TB	0
	DC1		12.1 TB	
		RAC1	12.1 TB	172.16.2.82, 172.16.2.69 (3.8 TB), 172.16.2.70 (3.2 TB)
	DC2		9.0 TB	
		RAC1	9.0 TB	172.16.2.66 (2.9 TB), 172.16.2.67 (3.2 TB), 172.16.2.68 (2.9 TB)

- Archiving split into 6 MariaDB databases (1 server) for easier and faster maintenance
- Not using `archive_event`, archiving on change instead to improve devices performance
- 13366 attributes, 340 G/month, 6 DB, 1-2 partitions per month, 10 subscribers per DB + 10 periodic archivers

Partitioning and decimation

- Tables partitioned monthly
- Data older than 6 months decimated to a separate DB
- indexing/partitions/decimation methods in `PyTangoArchiving.hdbpp.maintenance`

Data split in 6 Databases

- **hdbacc** 969 attrs / 35G month
- **hdbct** 1977 attrs / 70G month
- **hdbdi** 796 attrs / 75G month
- **hdbrf** 3493 attrs / 50G month
- **hdbpc** 1977 attrs / 10G month
- **hdbvc** 4154 attrs / 100G month

Archiving clients

- <https://github.com/tango-controls/PyTangoArchiving> library used for extraction
- eGiga used on beamlines
- shell/python tools plus qwt-based taurus widgets, pyqtgraph on development

Periodic archiving on HDB++

- ALBA developed an archiver for inserting periodical data without events
 - <https://github.com/ALBA-Synchrotron/PyHdbppPeriodicArchiver>
 - <https://github.com/ALBA-Synchrotron/libhdbppinsert>
- Configuration via PyTangoArchiving.HDBpp api (*periodic* methods)
- Use cases
 - Archiving through firewall
 - Fix period archiving required by some users
 - Legacy control system with bad polling/notifd performance