TANGO ARCHIVING SERVICE STATUS

G. Abeillé, S. Pierre-Joseph, J. Guyot, M. Ounsy (Synchrotron Soleil),
S. Rubio (ALBA), G. Strangolino, R. Passuello (ELETTRA)

Abstract

In modern scientific instruments like ALBA, ELLETRA or Synchrotron Soleil, the monitoring and tuning of thousands of parameters is essential to drive high-performing accelerators and beamlines. To keep track of these parameters and to easily manage large volumes of technical data, an archiving service is a key component of a modern control system like Tango [1].

To achieve this, a high-available archiving service is provided as a feature of the Tango control system. This archiving service stores data coming from the Tango control system into MySQL [2] or Oracle [3] databases.

Tree sub-services are provided:

- An historical service with an archiving period up to 10 seconds,
- A short term service providing a few weeks retention with a period up to 100 milliseconds,
- A snapshot service which takes "snapshots" of Tango parameters and can reapply them to the control system on user demand.

This paper presents how to obtain a high-performance and scalable service based on our feedback after several years of operation. The deployment architecture in the different Tango institutes will then be detailed. The paper concludes with a description of the next steps and incoming features which will be available in the next future.

CONCEPTS

The Tango Archiving service is a central diagnosis tool for large scientific instruments. It can be used for daily operation such as:

- Vacuum, temperature follow-up
- Insertion devices follow-up…

But also to analyse complex scientific phenomena:

- Beam orbit stability
- Post-mortem analysis…

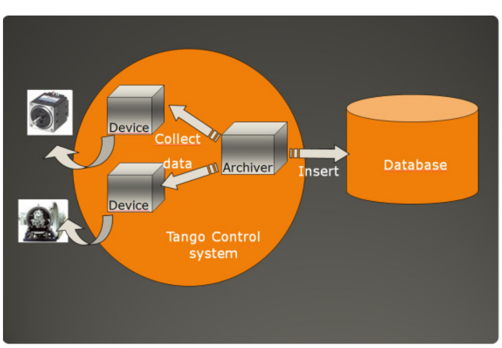

Figure 1: Tango archiving data collection.

The archiving service collects Tango control data (also called Tango attributes) to persist them in databases. The data collection is done by Tango devices called Archivers (see Fig. 1). A complete explanation of Tango concepts is detailed here [1].

All archiving components are based on a Java [4] API (Application Programming Interface), designed with JDBC [5]. Upon this API, a set of modules is rendered:

- Extractor Tango device: service to retrieve inserted data from the database.
- Watcher Tango device: monitor the archiving. Report alarms for not inserted values.
- Manager Tango device: configure, start or stop the archiving.
- GUIs in Swing: End-user interfaces to configure, start or stop the archiving; data plotting of archived values.

The archiving service proposes several data collection policies as detailed here in.

Historical/Temporary DB

The historical and temporary DB services can record data at a fixed frequency. The minimum period is 10 s for HDB and 100 ms for TDB. It is also possible to filter the data to be inserted i.e. to insert data only when it changes or when it goes over some thresholds. Due to the amount of data, TDB tables may be truncated on a regular basis (every 12 weeks at Soleil).

Furthermore, a GUI called Mambo (see Fig. 2) can configure, start or stop the archiving. It is also possible to extract data and plot it.

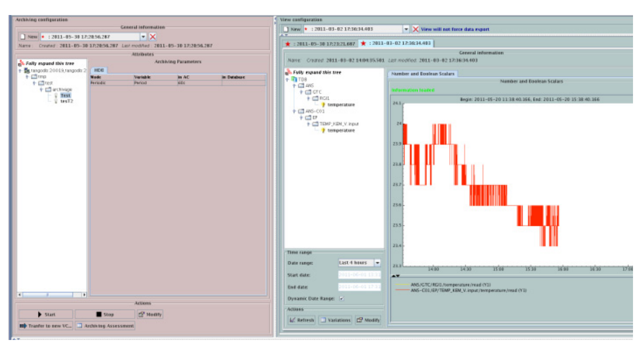

Figure 2: Mambo.

To achieve scalability, a collection of archiver devices may be deployed. Each archiver is in charge of collecting a set of Tango attributes; they may be set up in two different ways:

- Dedicated mode: each archiver collects a pre-defined list of attributes.
- Load-balancing: when a data collection is started, its archiving is allocated to the least loaded archiver.

Snapshot Database

The Snapshot service takes on-user demand "pictures" of a control system for a pre-defined set of Tango parameters (or Tango attributes). The persisted parameters may subsequently be reapplied on the Tango control system.

With the Bensikin (Fig. 3) GUI, the user can define the sets of Tango attributes. For each context, he can launch snapshots, display their contents, and reapply values to the Tango control system.

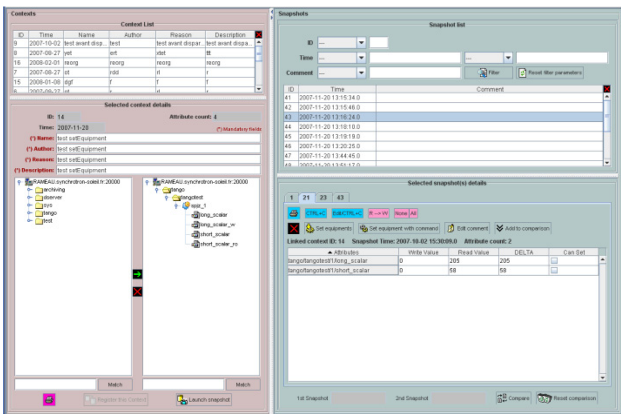

Figure 3: Bensikin.

STATUS

The archiving service is aimed to be used by any institute using Tango as a control system.

The project was initiated by Soleil in 2002. The initial design was jointly done with the ERSF [6] control team who in charge of the Tango core development. All the API, devices and GUIs were developed under the responsibility of Soleil. It was built up during the Soleil construction phase with continuous deployment on production platforms. At this time, the Soleil IT team was constantly under users' pressure to obtain a stable and high-performance system. Due to a lack of resources, most of the development was subcontracted.

Then ELETTRA joined the Tango collaboration at the end of 2003 followed by ALBA in late 2004.

Owing to different contexts and technical backgrounds, the archiving service is not used in the same configuration in all institutes. The actual deployments are described below.

ALBA

The ALBA synchrotron uses MySQL for all databases. Around 8000 HDB attributes are archived at periods between 10s and 1 minute and 2000 TDB attributes are archived at periods between 1s and 1 minute.

95% of the Archivers are deployed on the same host as the Tango archived devices in "dedicated mode".

On beamlines, it is under deployment with a centralized load-balancing mode using virtual servers.

The client applications were developed in Python using the PyTango binding.

Elettra

The TANGO archiving system is used on both Elettra and FERMI@Elettra. Each accelerator is running its own instance of the Archiving System.

Common characteristics to the two archiving infrastructures are:

- MySQL version 5.5.13 with MyISAM indexed and partitioned tables.
- Centralized infrastructure with 30 archivers running for HDB and 10 for TDB
- Management of archiving performed with an Elettra GUI based in QT, MANGO [1],
- Retrieving and data plotting done using E-Giga [1], based on PHP and MySQL.

Elettra figures are:

- 1561 attributes configured in HDB
- 7 attributes configured in TDB

Archiving system for Elettra running on a dedicated hardware based on an Intel(R) Xeon(TM) QuadCore CPU @3.20GHz with 4 GB of RAM and 1 TB of storage capacity on a RAID 6.

FERMI@Elettra figures are:

- 1798 attributes configured in HDB
- 401 attributes configured in TDB

Archiving system for FERMI@Elettra running on a Xen virtual machine using 4 cores of an Intel(R) Xeon(R) CPU E7330 @2.40GHz with 4 GB of RAM and 1 TB of storage capacity on a RAID 6 fibre channel connected external storage.

Both databases are replicated with MySQL replica system on a different machine with the same performances which is used by E-Giga to access data, so the archiving system is not loaded by such kind of requests.

Soleil

The archiving service is deployed on accelerators and beamlines. All databases are Oracle Enterprise Edition 11.1.0.7.0 with RAC (Real Application Cluster) [7] and partitioning options. There are 2 database clusters with each machine having 16 GB of RAM and 2 network interface cards (fiber channel):

- One RAC of 3 machines for the accelerators.
- One RAC of 2 machines for all beamlines.

A Disaster Recovery Plan (DRP) has been established to obtain a high-available service and no data loss:

- A degrade server can take over the RAC.
- Several backup strategies are active, e.g. RMAN (Recovery Manager) [8] does an incremental backup every day and a full back-up every two months. Data restoration is tested twice a year.
- Data are replicated in two buildings, on different media (hard disk and tape).

On top of this infrastructure, the Tango Archiver devices for data collection are configured in "load-balancing" mode and are distributed as followed:

On the accelerators, they are shared out on 4 machines (see Fig. 4 for an overview of the whole infrastructure):

- 270 HDB archivers.
- 300 TDB archivers.
- 1 SNAP archiver.

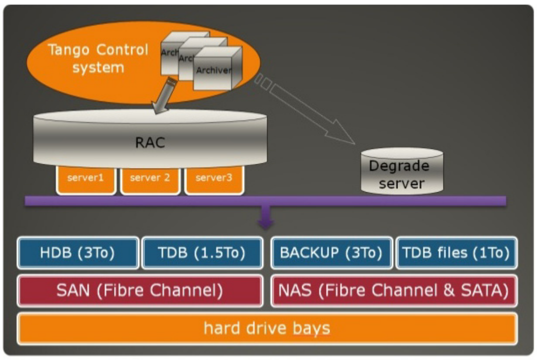

Figure 4: Soleil accelerators archiving infrastructure.

HDB and SNAP is also operational on all 25 beamlines. Some beamlines are archiving up to 350 attributes in HDB; one machine per beamline (with other services installed):

- 5 HDB archivers.
- 1 SNAP archiver.

Table 1 shows some figures for the Soleil accelerators. The numbers are growing constantly as the operators configure the archiving by themselves. Soleil machine operators are planning to collect up to 20,000 attributes in HDB and TDB.

Table 1: Soleil Accelerator Loads

	HDB	TDB	Snap
attributes	10 607	7 527	2700
Insertion/sec	480 (10TPS)	4 200	NA
size	80Go/month	20Go/day	2Go

CHALLENGES

Technical Challenges

Until recently, the main concerns at Soleil have been performance, scalability and stability issues. Various actions have resolved these issues:

- Change the whole infrastructure; install a cluster database with several storage systems.
- The Archiver devices were unstable and their performance was poor. As the code was subcontracted, it suffered from "quick and dirty" fixes, "copy-paste"… To overcome all these issues, the code had to be re-implemented.

The Archivers are now very stable at Soleil. There are no more failures after months of running. However, the code still needs to be cleaned up as has been pointed out by Soleil's quality tool, SONAR [9]. See [10] for how this tool has been used at Soleil.

The current watcher device is only able to report failures, but it is not yet able to diagnose the root cause. Indeed, there may be many reasons for a failed data insertion; for instance the archived Tango attribute may have been renamed, or its device may have been shut down. Consequently, it must be extended to identify and report the origin of a failure.

Data extraction and plotting is pretty basic today. A customisable reporting tool should be integrated in the archiving standard toolkit. The user should be able to simply configure and create reports in any format, such as pdf, csv, and so on. This tool may be also integrated into a scheduler so that machine operators can have an automatic follow-up every shift.

As all data extraction is done manually, all diagnoses are also manual. At Soleil, a project has recently been started to automate the diagnosis of orbit instability with a rule engine [11]. As soon as an orbit deviation is recorded, this tool will automatically search in archived data to identify the causes of instability, such as insertion movements, temperature drifts, power supply misbehaviour, etc.

Collaboration Challenges

A major challenge for Soleil is to fulfil all needs of the Tango community while also remaining responsive to internal user requests, under the constraint of limited human resources. The Tango collaboration has no dedicated resources, so it is at the discretion of each institute to bring some resources to the collaboration. All institutes have different technical backgrounds, it is sometimes difficult to put everything is common as some tools were designed for a specific context, but we must do our best to have the maximum in common.

Effort has recently focused on packaging in order to deliver an easy-to-install archiving system [12]. But there are still many things to improve to efficiently collaborate. The archiving service is now stable, but many evolutions may be added, so a common specification must also be written up. Moreover, this organisation must be also ready to provide support for new incoming Tango users like MAX IV [13].

To conclude, the archiving project is on the one hand a matter of technical challenges, on the other hand it is a major human challenge. It is now a mature project, but it is far from coming to an end as many improvements remain to be done. Besides, this Tango add-on reinforces the added value of the Tango framework.

REFERENCES

[1] http://www.tango-controls.org/
[2] http://www.mysql.com/
[3] http://www.oracle.com/us/products/database/index.html
[4] Java: http://www.oracle.com/technetwork/java/index.html

[5] JDBC: http://fr.wikipedia.org/wiki/Java_database_connectivity
[6] ESRF: http://www.esrf.eu/
[7] http://en.wikipedia.org/wiki/Oracle_RAC
[8] http://en.wikipedia.org/wiki/RMAN
[9] SONAR: http://www.sonarsource.org/
[10] THBHMUST02. Assessing software quality at each step of its lifecycle to enhance reliability of Control Systems. V.Hardion
[11] WEPKS008 Rules-based Analysis with JBoss Drools: Adding Intelligence to Automation. E. De Ley
[12] http://www.tango-controls.org/download/archiving-1
[13] http://www.maxlab.lu.se/maxlab/max4/index.html