

Future Archiver (librarian) for WinCC OA Control Systems

Piotr Golonka, Jakub Guzik, Rafal Kulaga
CERN BE-ICS-SDS

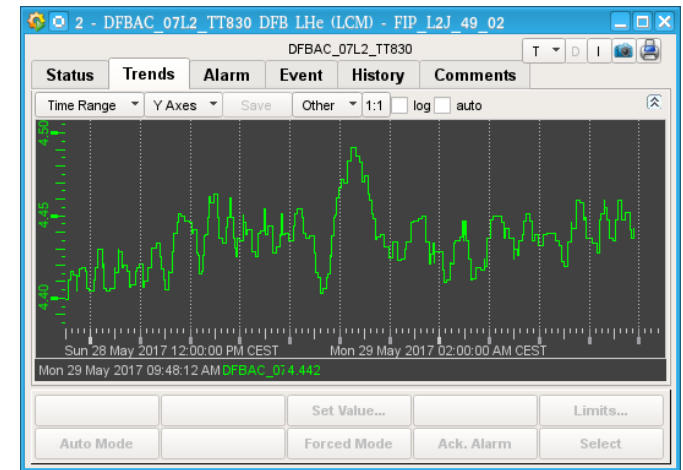
In Openlab collaboration with ETM/Siemens

Content

- **WinCC OA archiver status quo**
 - Deployments
 - Scale (example of BE/ICS applications)
 - Challenges
- **NextGen archiver**
 - Motivation
 - Architecture and technologies
 - Status

Archiver in a control system


- **Essential element of SCADA (Supervision) layer**
 - Stores the history of values and alarms (into a “database”)
 - Allows the operator to see the evolution of the process (historical trends)
 - Diagnostics and postmortem analysis “Event Screen”
- **Entry point for Data Analytics**
 - For industrial processes
 - ... but also for physics analyses (“conditions”)
- **Others**
 - “Event replay”




Fundamentally:
Time-series data

Status Quo

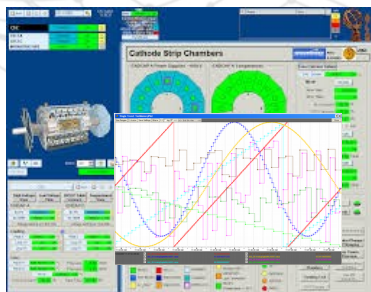

- **WinCC OA: de facto standard for slow controls at CERN**
 - Detector/Experiment Control Systems
 - the 4 LHC detectors
 - COMPASS, NA62, protoDUNE
 - Critical infrastructure of the LHC and injectors
 - QPS, Cryogenics, Vacuum, Powering
 - Technical infrastructure
 - C&V, Electrical network, ...
- **All the system use the “Oracle RDB Archiver”**
 - standard component of WinCC OA
 - Highly optimized over past 15 years
 - fruit of collaboration between CERN and ETM
 - 100 v/s → 100.000 v/s



ONLINE DB




OFFLINE DB







Low latency!

ONLINE
Operation,
Monitoring




Physics
Data
Analysis


Accelerators and tech. infrastructure (BE/ICS Applications)



Long-term
storage

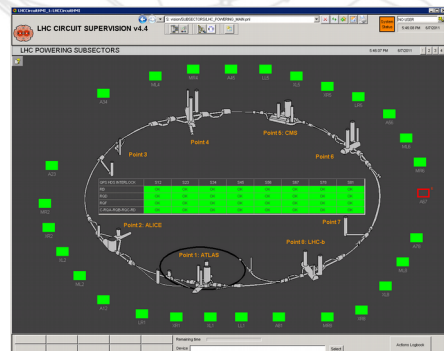


Data
Analysis



Low latency!

ONLINE
Operation,
Monitoring






Scale challenge (BE/ICS Applications)

- **LHC Quench Protection System**
 - 48 WinCC OA systems
 - 150k signals, 400M rows/hour insert rate
 - Dedicated Oracle instance: QPSR
 - Complete data set transmitted to ACCLOG
 - Dedicated optimizations in WinCC OA:
for Run II of the LHC!

Scale challenge (BE/ICS Applications)

- **“Everyone else”**
 - 170 other WinCC OA systems, different clients
 - 1 DB schema per system, some correlated
 - 2.5 million signals, 34 million rows/hour insert rate
 - Share the same common online DB: SCADAR
 - Major parts of data transmitted to ACCLOG
 - Online access with minimal latency for operators
 - Trends, event screens, event replay
 - **Statistics**
 - how many changes per signal during last 24 hours?

Data analytics at the moment



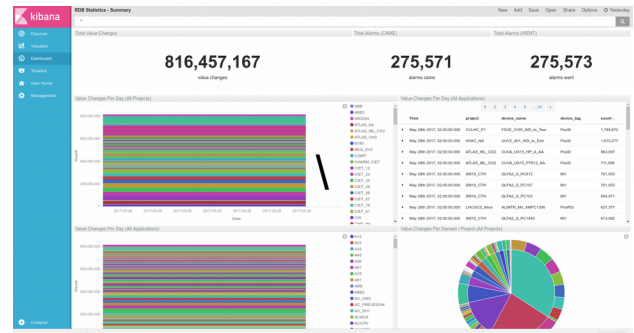
Data Analytics



Daily Statistics Calculation



elasticsearch kibana



Challenging use cases of Oracle Archiver

- **Three main use cases are in conflict:**
 - “*Trend Plot*”: history of **one** (or few) signals over a possibly long period of time
 - Access through PK on IOT table, time-based partitioning
 - “*Event Screen*”: history of **all** value changes/alarms over relatively short period of time
 - Scan over a partition
 - “*Event Replay*”: massive data retrieval for all signals over a period of time of unknown bounds
 - “give me the last value before the t_{start} , and then all following values”
 - A killer for current Oracle setup
 - Seems significantly better executed on the Hadoop data copy

Why “Future Archiver”?

- **ETMs motivation**

- Alternative(s) to very expensive Oracle
- Easy to deploy and manage
- Easy to integrate new technology
- Lookout to NOSQL and data analytics
- Opportunity for code cleanup

- **CERNs motivation**

- IT drive to alternative(s) to Oracle
- We face Big Data (3V)
 - Unknown limits really!
- Enable data analytics
- NXCALS
- New features
 - Parallel archiving, proper treatment of meta-data history, ...
 - Ready to integrate new technologies

Future Archiver for WinCC OA

- **Successor of Oracle RDB Archiver**
 - High-performance
 - Cost-effective
 - Robustness
- **Open architecture (plugins)**
- **SQL and NOSQL**
- **New architecture compatible with evolution of WinCC OA**



mongoDB



ORACLE®



Joint Openlab R&D projects with Siemens/ETM and Intel/Cloudera

> At CERN:

- Piotr Golonka (technical supervision)
- Jakub Guzik (developer of the core archiver part)
- Rafal Kulaga (“BigData” prototype backend with Apache Kudu)
- IT/DB colleagues (Luca, Prasanth, Zbyszek)

> At Siemens/ETM:

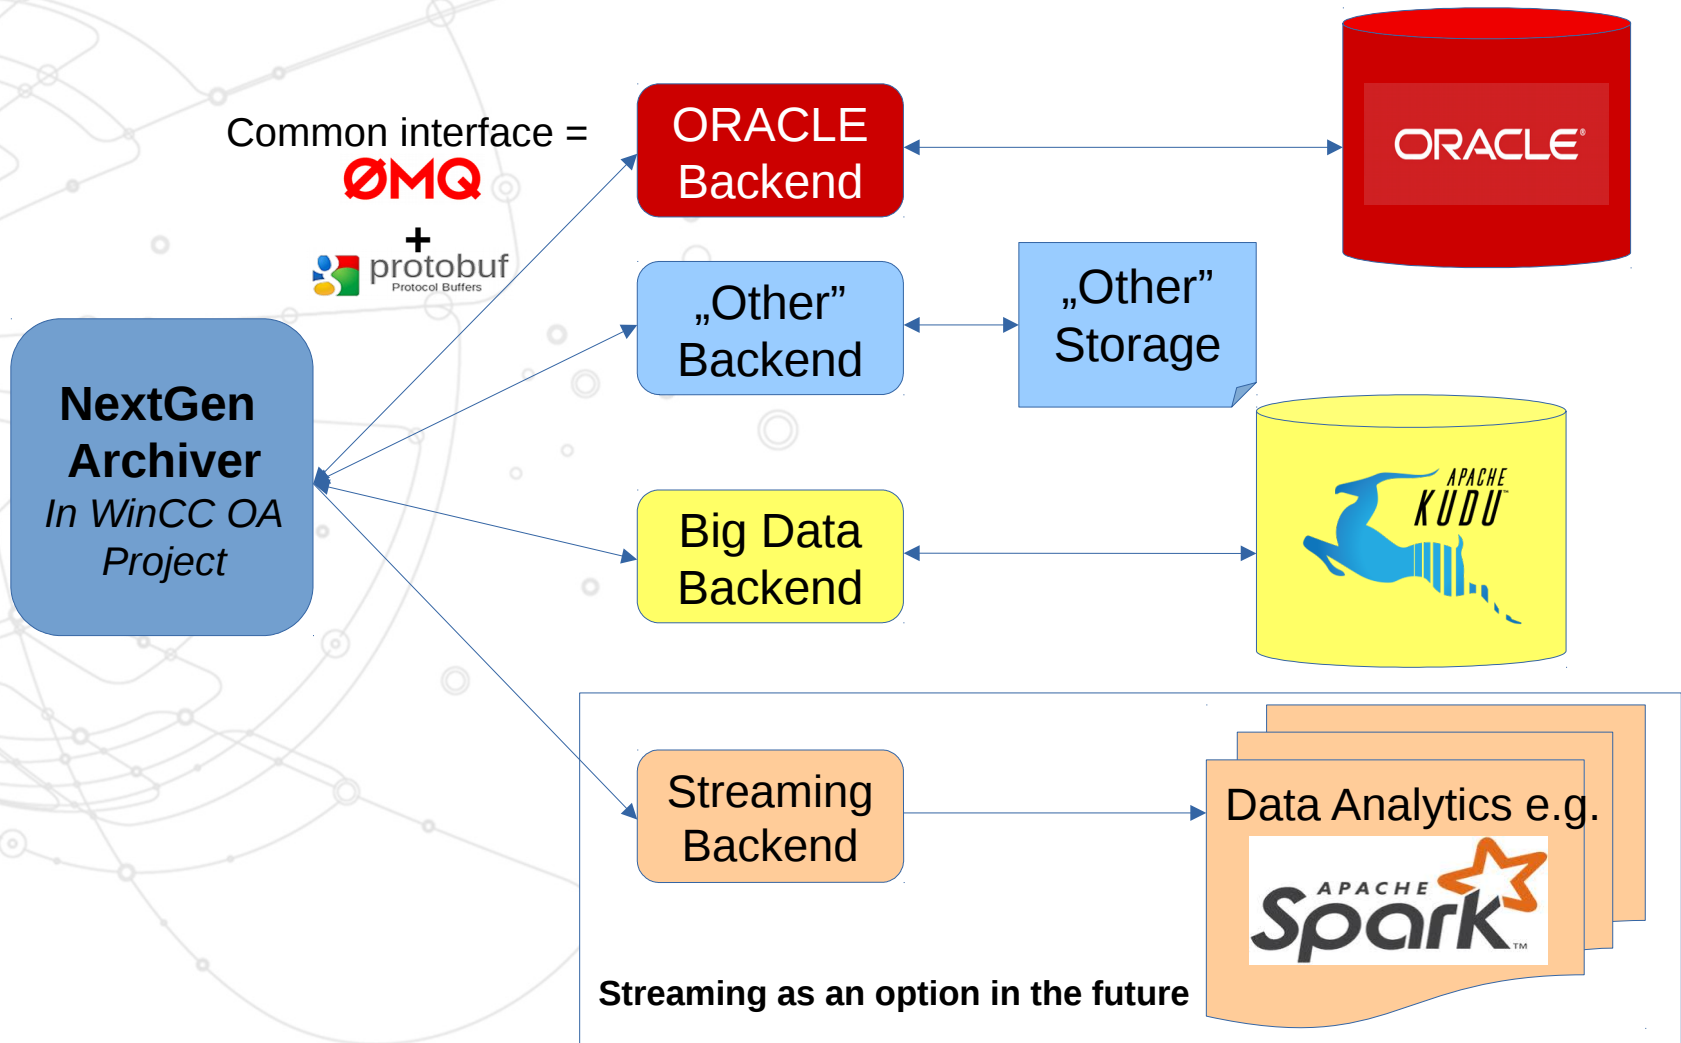
- Development of the core archiver part
- Evaluation of DB technologies
- Oracle-compatible backend
- MongoDB backend
- Project ongoing since 1 year now with first prototypes becoming available

> At Intel/Cloudera: (kickoff stage)

- Expertise for Kudu technology
 - Low latency data access
 - Quality of service and resource allocation



High level architecture „Frontend” and „Backend”



More use cases

- **“Generic data fan-out” for WinCC OA**
 - Streaming data analytics
 - Data bridge for the ALICE O2
 - 100.000 conditions parameters acquired asynchronously (on change)
 - Most recent values need to be interleaved into the online data frames every 50 ms
 - This yields 2 mln v/s to be streamed out (!)
 - “NextGen Archiver” could provide this data to a custom ALICE-developed backend
 - Open architecture (!)
 - In parallel to “standard” archiving

- **Research & Development project**
 - First functional prototype of complete architecture
 - Oracle, MongoDB and Kudu backend prototypes being implemented and tested
 - Baseline performance and functional tests ongoing
 - Likely to become available in a future version of WinCC OA
 - Use cases and requirements of the Experiments to be formulated, cross-checked and covered
 - Further synergies with NXCALS to be explored
 - No plans for deployment before the end of LS2