CONTROLLING THE GOOD GIANT PEEKING DATA...

CONTROL AND CONFIGURATION OF THE ATLAS TRIGGER AND DATA ACQUISITION SYSTEM DURING DATA TAKING ACTIVITIES

Riccardo Maria BIANCHI (CERN)[1],

on behalf of the ATLAS TDAQ Collaboration[2]

WAKING UP ATLAS THE GOOD GIANT IN A SHORTER TIME...

Improving Operation timings through parallelization and network usage optimization

To wake up the ATLAS Giant in order to take data, we have to go through a Finite-State-Machine, up to the "Running" state, configuring all the services and applications which let ATLAS take data.
Each step is now optimized to efficiently use CPU and Network resources.

- we optimized the state transitions
- we parallelized the handling of commands to applications
- faster launching
- optimized usage of both CPU and network

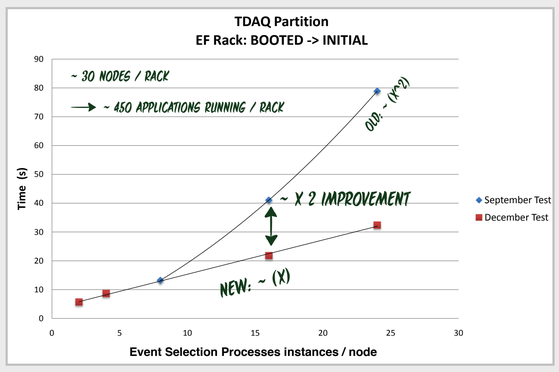

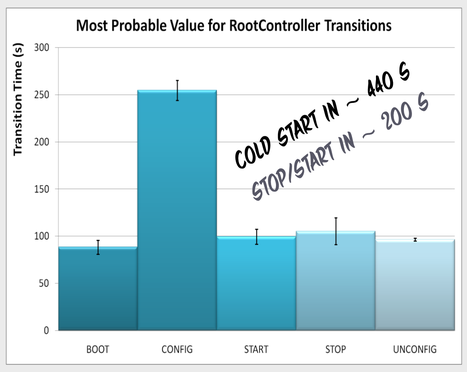

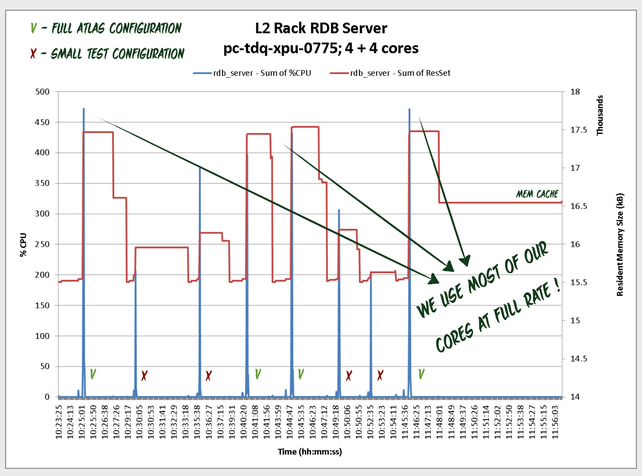

We are now capable of using most of our CPU cores with a safety margin for extra work, despite the huge load due to the large size of the full ATLAS configuration.

The memory usage during the configuration is stable.

NOW ATLAS REMEMBERS THINGS BETTER AND FASTER...

Optimization of the Configuration DataBase access

The run configuration of the ATLAS detector is stored in a custom distributed database (OKS). Configuration data have to be promptly accessible, when needed, to all applications which are responsible to run the detector.

Bottlenecks have been spotted and a more effective code has been implemented:

- multi-threaded DB loading
- rewriting of slow algorithms
- more effective memory usage
- usage of more effective object collections (e.g. maps versus ordered lists)

The configuration can be modified now by many users at the same time; all of them being able to modify the ATLAS configuration. The reload is then centrally managed to ensure a safe and stable data-taking.

Configuration database Read / Write operations are now ~6 times faster.

...AND IT MAKES USE OF ALL THE ROADS HE CAN FIND...

Optimization of the Network usage

During operations, configuration and commands have to be sent to many systems and applications, through a dedicated 1Gbit network link. We recently upgraded our code to be able to effectively use all the bandwidth.

Now we use our 1Gbit network link in an effective way, feeding up only while we need it.

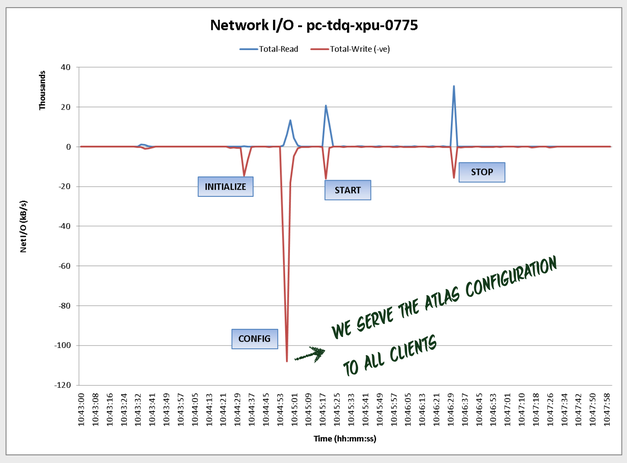

OPTIMIZATION

ATLAS is up and running, taking data successfully since more than one year.

The first goal of the TDAQ achieved --- recording data from the experiment in a safe way --- it was time to improve the Core packages to optimize the usage of the resources; and to explore new technologies, developing ancillary tools to get better system monitoring and to help the human operators during the configuration and data-taking operations.

NEW TECHNOLOGIES

LET THE GIANT THINK SMARTER...

Automation via Artificial Intelligence: the Shifter Assistant

The ATLAS detector system is operated by a non-expert shift crew, assisted by a set of experts providing knowledge for specific components. Operational tasks can be divided in:

- Operational procedures to run the system;
- A set of periodic checks and controls;
- Notify experts in case of problems

But computers are better than humans in automation. Checks and controls can be easily automated in order to: reduce and simplify shifters tasks, help shifters with more detailed and pertinent information, be more efficient avoiding repetition, formalize and store knowledge from experts, minimize system down-time, deal fast and effectively (and possibly automatically) with errors and failures.

Shifter Assistant uses an open-source Complex Event Processing (CEP) engine from "EsperTech" to process streams of information in real time, combining information from several sources, to detect patterns in time window, and reacting to them producing alerts and messages. The knowledge base is a list of directives.

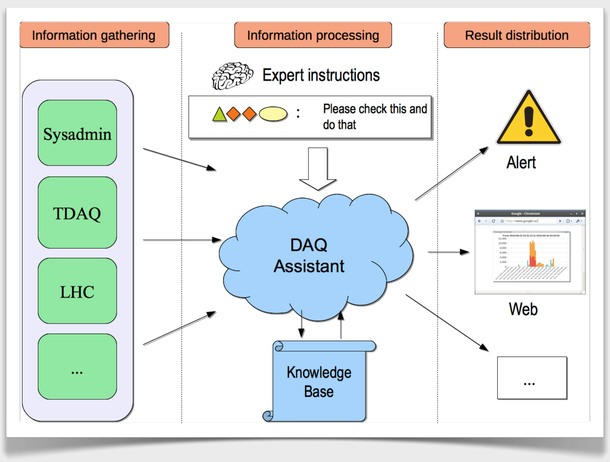

Shifter Assistant gathers information streams from various systems and returns alerts based on patterns found in data

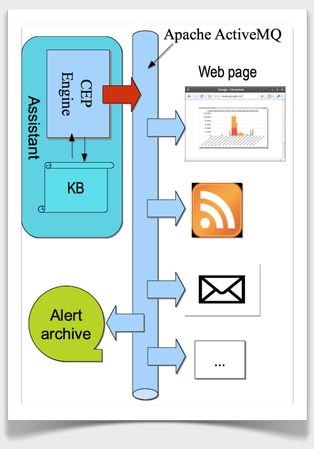

Shifter Assistant uses ActiveMQ dispatcher system to send notifications and alerts

LET'S CONFIGURE THE GIANT AT A FASTER PACE...

DBE: the new Configuration Database editor

The old editor used to edit the ATLAS Configuration database so far, had been conceived and implemented many years ago upon the Motif graphical libraries, and it lacked of many modern features.

The new Configuration editor is based on the "QT" graphical libraries[3], and it has all the features that one could expect from a modern editor

Features of the new database editor:

- Read/Write mode to remote DB servers (RDB) and to local database files;
- Multithreaded DB objects access for faster page rendering;
- Bunch editing of multiple DB objects at the same time; Drag & Drop with type check capability;
- Undo/Redo actions;
- User customized test plug-ins;
- Interoperability with db generation python scripts;
- Customizable user view, and user settings storage;
- Improved help to the user, through extensive checks, tooltips and messages;
- User's Guide accessible from the editor.

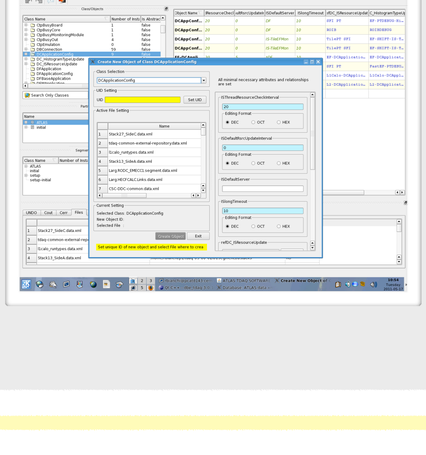

3) "QT" is a © by Nokia Corporation

GIANT'S MASTERS MESSAGES HANDLED WISER...

ELisA, the new interface to the ATLAS operation log messages DB

During operations, when the detector is running, a lot of messages and logs are exchanged among shifters, experts and automatic services. Real-time analysis let us to spot and understand problems.

The previous implementation, the ATLOG, had many major problems:

- very slow and unstable
- not scalable with the number of e-log entries
- monolithic package: same application for server and web interface
- load balancing not possible
- not multi-threaded
- not well tailored for Oracle interaction

The new implementation uses the Spring Java web application framework: a mature, stable framework with an optimized interface to Oracle.

We adopted:

- client-server architecture;
- a Model-View-Controller design pattern
- JSP to dynamically create the "View" pages
- JDBC used for data access in Oracle

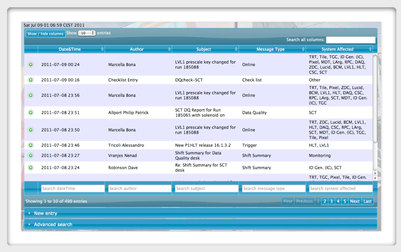

Speed and scalability test

The old ATLOG takes a lot of time to fetch and display e-log data (~x^2) putting a huge overhead on top of Oracle operation time; and it could not handle more than few thousands of entries.

The new ATLAS E-Log "ELisA" can handle tens of thousands of entries in a short time, putting only a light overhead on top of Oracle transactions.

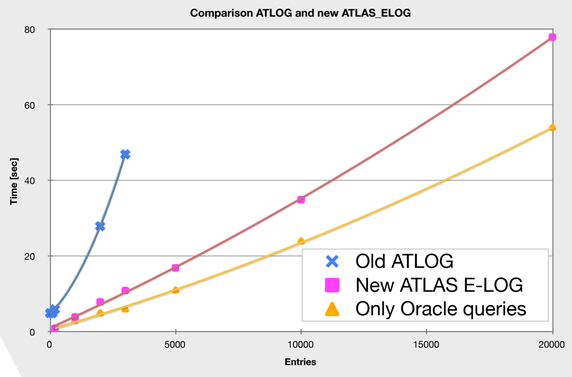

References:
Giovanna Lehmann Miotto, *et al.*, *"Present & Future of the Configuration Control of the Atlas TDAQ"*, TIPP 2009, Tsukuba, Japan. Nuclear Instrumentation and Methods in Physics Research, Section A: Volume 623, Issue1, 1 November 2010, Pages 549-551. See http://dx.doi.org/10.1016/j.nima.2010.03.066
ATLAS Collaboration, J. Instr. 3 (2008) S08003, 437pp.
ATLAS Collaboration, *"ATLAS High-Level Trigger, Data Acquisition and Controls Technical Design Report"*, CERN/LHCC/2003-022, June 2003. See http://cdsweb.cern.ch/record/616089/.

1) Corresponding author: Riccardo Maria BIANCHI (*CERN*), rbianchi@cern.ch, tel: +41-22-76-71196

2) ATLAS TDAQ Collaboration, The ATLAS Trigger/DAQ Authorlist, version 4.0, ATL-DAQ-PUB-2010-002, CERN, Geneva, 2010, http://cdsweb.cern.ch/record/1265604" *(to be updated)*