

A Validation System for the Complex Event Processing Directives of the ATLAS Shifter Assistant Tool

G. Anders (CERN), G. Avolio (CERN), A. Kazarov (PNPI), G. Lehmann Miotto (CERN), A. Santos (UNLP), I. Soloviev (UCI)



Outline

- **The ATLAS Trigger and Data Acquisition (TDAQ) system**
 - General schema
- **Complex Event Processing (CEP) in TDAQ**
 - Challenges
 - The Shifter Assistant (SA) tool
- **A validation tool for the SA**
 - General architecture
 - Web front-end
- **Conclusions**

Trigger and Data Acquisition System


See “The ATLAS Data Flow system for the second LHC run”
(R. Hauser - Track 1: Online Computing - 14/4)

- The **Trigger and Data Acquisition (TDAQ)** system is responsible for filtering and transferring data from the detector to the mass storage
 - **40 millions** proton bunch interactions per second
 - Most of the generated events are **totally uninteresting**
 - A filter mechanism is needed in order to select and collect the more interesting ones
- More than **20k** applications running on **2k** hosts
- **1800** point-to-point connections from detectors to TDAQ



From detector to data storage

Challenges in TDAQ Operations



- **The main goal when operating the system is to maximize the data taking efficiency**
 - Dealing fast and effectively with errors and failures
- **The system is operated by a non-expert shift crew with support by experts (from various sub-systems) providing knowledge for specific components**
 - Inefficiency may come from human interventions
- **Automating error detection and diagnosis is a key feature**
 - Introduction of the Shifter Assistant

The Shifter Assistant

What

- A tool assisting the shifter in his/her daily work
 - **Diagnosing** problematic situations and **assisting** in problem solving
 - **Remind** the shifter to (not) take action

Why

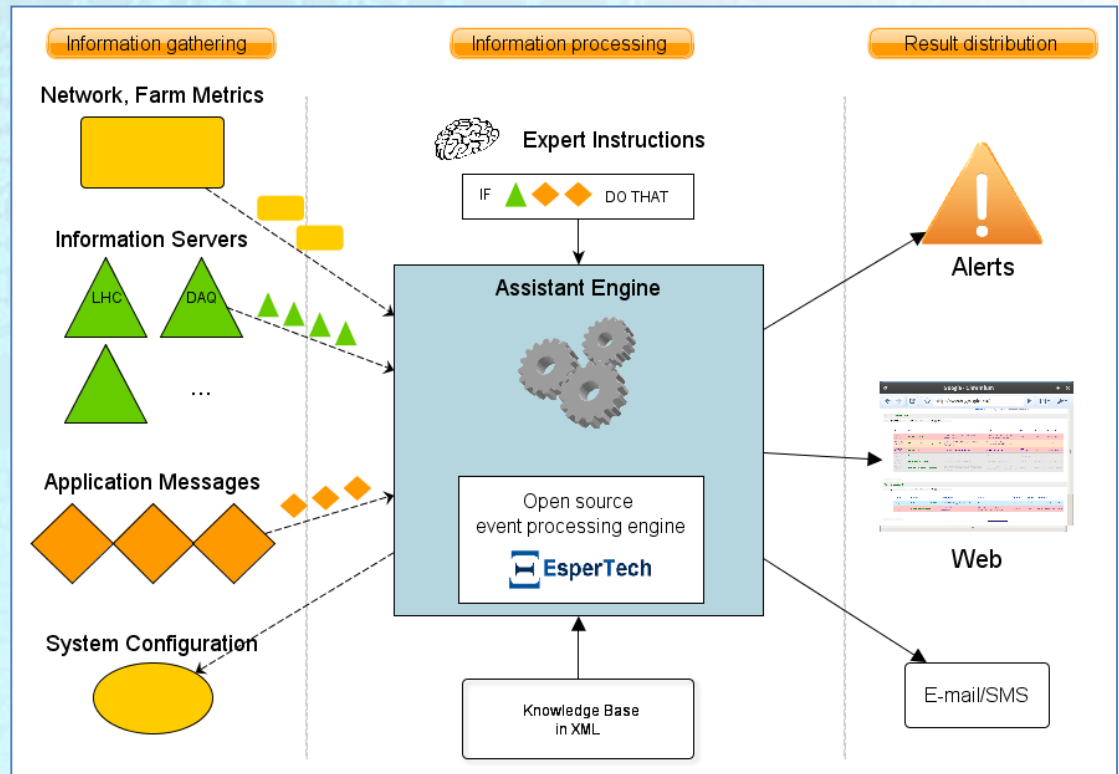
- **Reduce** and **simplify** shifter tasks
- Help shifters with more detailed and **pertinent** information
- Formalize **knowledge** from experts

How

- **Automate** checks and controls in real-time
- Process and analyze **heterogeneous** streams of information
- Provide **instructions** to the shifters on what to do and how to react
- **Promptly notify** shifters of problems and failures

The SA Architecture

- Meant to deal effectively with problems requiring the human intervention
- Based on the ESPER Complex Event Processing (CEP) engine
 - Continuous stream processing
 - Support for time/size windows, aggregation and grouping events
 - SQL-like pattern languages (EPL)
- Directives
 - Encode the knowledge from experts
 - XML structured KB



- Alerts
 - Effective and timeliness notification
 - Carry all the information needed for debug and fault diagnosis

A Validation Tool for the SA

Why

- **Maintenance** and **extendibility** of the knowledge base
- Spread **expertise** on CEP
- Writing **correct** directives is not trivial

Requirements

- Directives shall be evaluated in an environment as close as possible to the **production system**
- Verify **syntactic correctness** of directives
- Verify correctness of **“condition matching”**
- Produce meaningful **diagnostic**
- **Easy** to use

How

- Exploit the **modular** “input” of the SA
- Inject into the SA **historical archived** “real” data
- Feed the SA with the **directives** to be tested

The Working Model

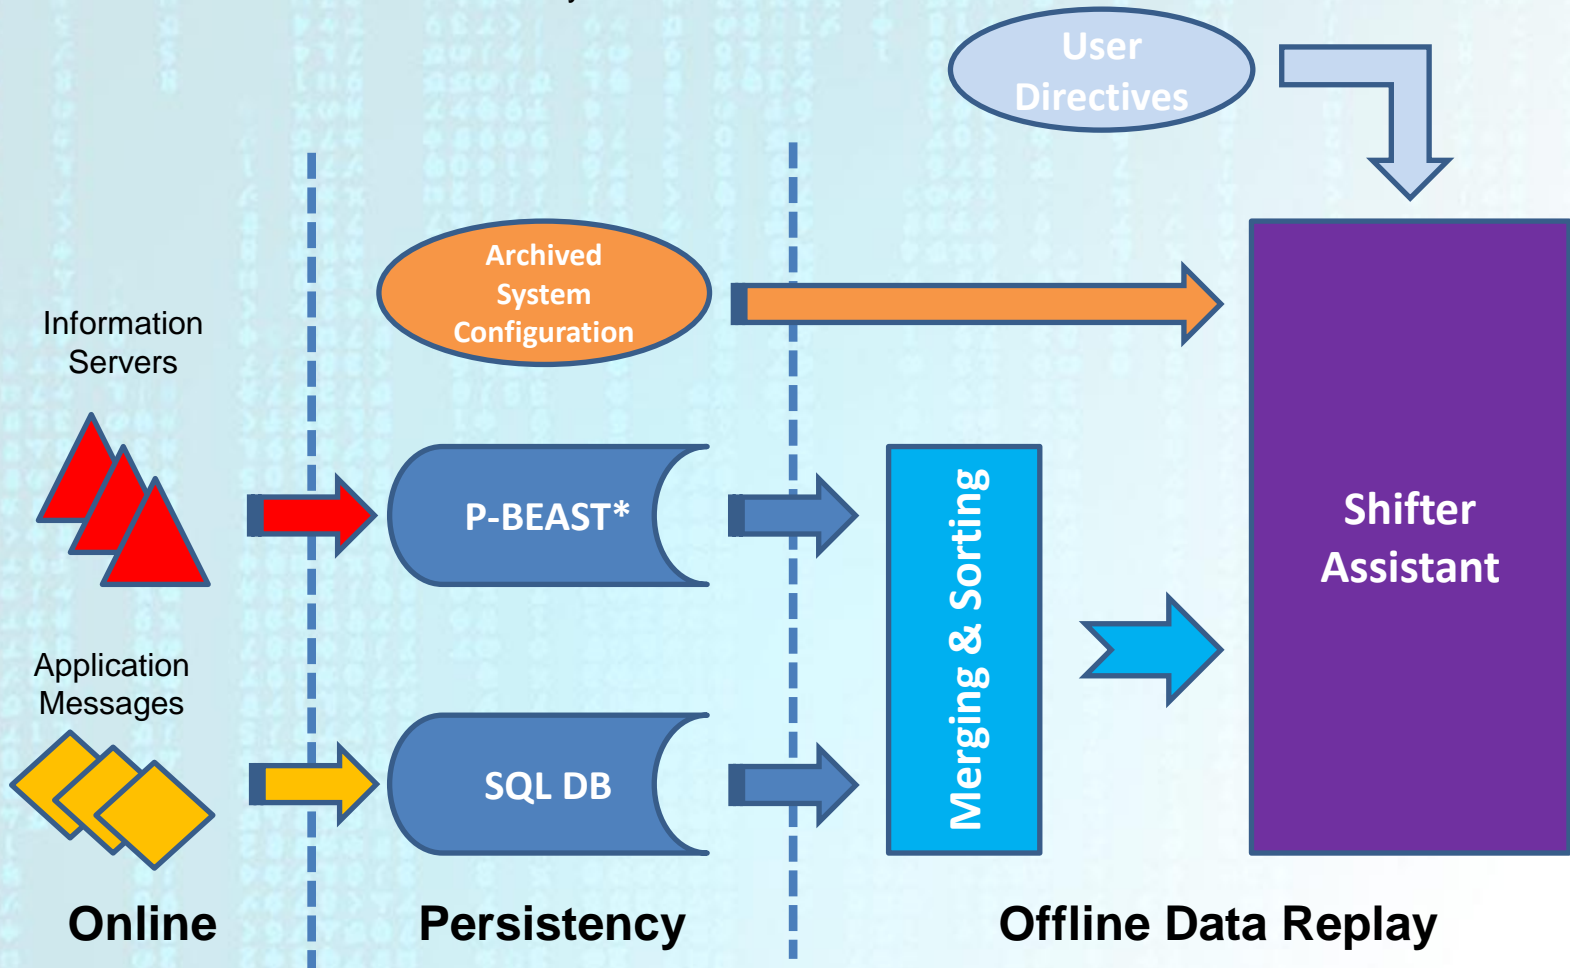


Challenges

- **Provide a proper controlled context for testing**
 - Sand-boxing the SA
- **Archive all the needed information**
 - Not archived means lost
- **Efficient data retrieval**
 - Often from custom database-like services
- **Reconstruct information that looks like the original one**
- **Merge and sort data in time**
 - Crucial for proper event correlation
- **Inject events into the SA with proper timing**

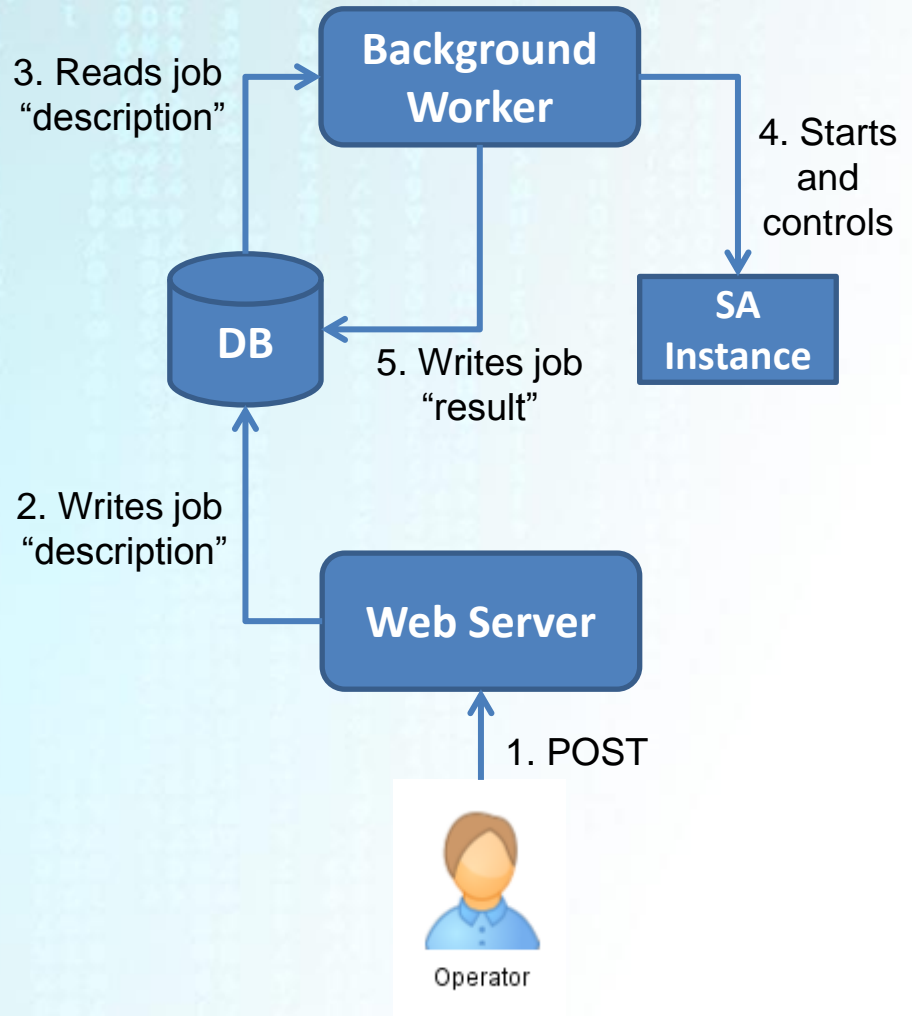
Architectural Overview

*Persistent Back-End for the Atlas Information System



Web-Based Interface

- **The user is requested to provide**
 - The run number identifying historical data
 - The directives to be tested
 - Information type to play-back
- **A server-side application takes care of running a SA instance**
 - Proper environment
 - Process management
- **Results of the SA run available on-demand**
 - Detailed activity log



Web Interface

Last executions

- GIU1, status: Finished (5)
- DQM5, status: Finished - With errors (10)
- ALE ECR3 (DEBUG), status: Finished (5)
- ALE ECR3 (DEBUG), status: Finished (5)
- ALE ECR3 (DEBUG), status: Finished (5)

[See all executions](#)

Access to previous executions

New Execution

Title

Title

Provide a run number

Run numbers:

Run numbers

Basic

Advanced

Historical data types

IS data types to read from P-BEAST

Select classes one at a time and press enter

Add new file

User directive XML file #1:

Browse...

No file selected.

User directives

Remove file #1

☐ Skip syntax errors

☐ Generate performance plots

Timeout

30 minutes

Start Replay Test

Results of a Validation Run

		Directive	Message	Action	Severity	
application-segfault-while-running	AAL.TDAQ.Expert	March 12, 2015, 5:10 p.m. (1426176602007)	application-segfault-while-running	One or more applications died while running!	Expert attention needed. Post e-log entry.	WARNING
<ul style="list-style-type: none">• Date: 12 Feb 2015 21:30:20 CET Application: ttreeDumpRC_app Host: pc-tdq-mon-69.cern.ch ErrorLog: /logs/tdaq-05-05-00/ATLAS/ttreeDumpRC_app_pc-tdq-mon-69.cern.ch_1423772536.err Signal: 6• Date: 12 Feb 2015 21:30:28 CET Application: RPC-DQAgent Host: pc-tdq-mon-79.cern.ch ErrorLog: /logs/tdaq-05-05-00/ATLAS/RPC-DQAgent_pc-tdq-mon-79.cern.ch_1423772531.err Signal: 11• Date: 12 Feb 2015 21:32:10 CET Application: HLTDQM Host: pc-tdq-mon-64.cern.ch ErrorLog: /logs/tdaq-05-05-00/ATLAS/HLTDQM_pc-tdq-mon-64.cern.ch_1423772546.err Signal: 11						
CHIP-recovery-done	AAL.TDAQ.RunControl	March 12, 2015, 5:10 p.m. (1426176622010)	CHIP-recovery-done	The expert system completed automated recoveries of type RpcTower.	Relax...	INFORMATION
<ul style="list-style-type: none">• Details: {sector=27, category=RECOVERY, tower=3, immediateRecoveryRetry=false, errorDesc=, status=DONE, name=RpcTower, action=NONE, isError=false, type=TOWER_RECOVERY} myname: RpcTower						
slow-IS-subscriber	AAL.TDAQ.Expert	March 12, 2015, 5:16 p.m. (1426176963693)	slow-IS-subscriber	An application has performance problems and can not cope with IS update rates	Experts should have a look to the application.	WARNING
<ul style="list-style-type: none">• Subscriber_Ref: Histogramming Partition: ATLAS IS_Server: Histogramming• Subscriber_Ref: Histogramming Partition: ATLAS IS_Server: Histogramming						

Conclusions

- **The Shifter Assistant proved to be an essential tool for optimal DAQ operations during LHC Run I**
- **Having a solid and robust model for the maintenance and extension of the knowledge base is crucial**
- **The developed model allows to**
 - Validate directives in an easy and “controlled” way
 - Fully expose experts/developers to the alert system
 - Better understanding of the implications of bad directives
- **Greatly improving the quality of directives**
 - The “replay” sandbox allows to experiment directives writing, and agile iterations over different approaches
 - Independently of the working system infrastructure