Experience with Shifter Assistant - an intelligent tool to help operations of ATLAS TDAQ system in LHC Run 2

Andrei Kazarov[1,*], *Alina* Corso-Radu[2], and *Giuseppe* Avolio[3]

[1]NRC "Kurchatov Institute" - PNPI, St. Petersburg, Russian Federation
[2]Univercity of California, Irvine, USA
[3]CERN, Geneva, Switzerland

Abstract. The Trigger and DAQ (TDAQ) system of the ATLAS experiment is a complex distributed computing system, composed of $O(10,000)$ of applications running on more than 2,500 computers. The system is operated by a crew of operators on shift. An important aspect of operations is to minimize the downtime of the system caused by runtime failures, such as human errors, unawareness or miscommunication. The paper describes recent developments in one of the intelligent TDAQ frameworks, the Shifter Assistant (SA) and summarizes the experience of it's use in operations of ATLAS during LHC Run 2. SA is a framework whose main aim is to automatize routine system checks, error detection and diagnosis, events correlations etc. in order to help the operators react to runtime problems promptly and effectively. The tool is based on CEP (Complex Event Processing) technology. It constantly processes this stream of operational events (O(100 kHz)) over a set of "directives" (or rules) in the knowledge base, producing human-oriented alerts and making shifters aware of the operational issues. More than 200 directives were developed by TDAQ and ATLAS detector experts for different domains. In this paper we also describe different types of directives which were developed in course of Run 2, and present few examples of most interesting and challenging ones, demonstrating the power of CEP for this type of applications.

1 Introduction

The ATLAS experiment [1] is one of the major experiments of the Large Hadron Collider (LHC). This general purpose experiment consists of various tracking detectors (the Inner Detector), electromagnetic and hadronic calorimeters and a muon spectrometer. The detector provides many millions of read-out channels, able to capture data every 25 nanoseconds. This volume of data can not be recorded and kept for further data analysis. The ATLAS Trigger & Data Acquisition (TDAQ) system [2] is responsible for the readout, selection (filtering) and transfer of the selected physics events to the permanent storage, reducing the initial LHC collision frequency of 40 MHz to an average rate of stored physics events (1.5 MB size) of 2-3 kHz. The TDAQ hardware infrastructure includes 2,500 computer nodes where 60,000 or

[*] e-mail: Andrei.Kazarov@cern.ch
From ATL-DAQ-PROC-2018-037. Published with permission by CERN.

more supervised applications (including 47,000 high-level trigger processing tasks) run, coherently performing data acquisition, event selection, control, configuration and monitoring tasks. TDAQ software includes central infrastructure services like Controls, Configuration and Monitoring. The proceedings describe one of the TDAQ Controls tools, the Shifter Assistant, whose task is to intelligently help ATLAS operations; the proceedings also summarize the experience with the Shifter Assistant in ATLAS data taking in course of LHC Run 2.

2 ATLAS operations challenges

During the LHC Run 2 ATLAS is operated by a shift crew of 7 shifters (aka operators) in the ATLAS Control Room (ACR), supported by a larger team of subsystem experts on-call. Maintaining high data-taking efficiency (the fraction of time when ATLAS is recording data in physics configuration over the total time when LHC delivers collisions in Stable Beams mode) is the primary task of the operations crew. Operating the TDAQ system includes the following challenges:

- system complexity and heterogeneity;
- likely error conditions from custom and commodity hardware and software;
- necessity to preserve expert knowledge during long (decades) development and maintenance periods;
- big volumes of operational data to analyze: over 1M messages per 10-20 hours of a typical physics run; $O(10^5)$ rapidly changing operational parameters with a total rate of up to 200 kHz).

Operator working stations in the ACR are equipped with multiple tools, providing different views on operational monitoring data (Figure 1). However it is not trivial for a non-expert shifter to follow and react to all events happening in the system.

Figure 1. View of a shifter working environment in ATLAS Control Room

3 Shifter Assistant

Shifter Assistant (SA) is an intelligent framework, developed in order to simplify work of ATLAS and TDAQ operators by providing automatic analysis of various operational monitoring data sources [3]. SA utilizes Complex Event Processing (CEP) [4] technology allowing to formalize and store expert knowledge in form of text "directives" in EPL (Event Processing Language) and to apply this knowledge in a real-time manner to streams of TDAQ operational monitoring data flowing through the system, as it is shown on Figure 2. The processing results, called *alerts*, are usually routed to the ATLAS operators via a web application, and also can be sent to remote experts as e-mails and SMS text messages. Some alerts can be sent to another TDAQ infrastructure application - the "CHIP" Expert System [5], whose task is to take some automatic actions when specific conditions are detected. Each alert contains an assessment or a diagnosis of the detected operational condition and also an action which is to be taken by the shifter in this case. Few samples of alerts can be seen on a screenshot of the SA web application in Figure 3. Each alert must be acknowledged by the shifter, and this information is available to the experts (operations managers). Each desk in ACR has a dedicated domain or tab in SA web application (Run Control, Trigger, Shift Leader, Muons, and ID - Inner Detector). New alerts are dynamically loaded on to the page.

The SA CEP engine was implemented in ESPER - a Java-based CEP framework from EsperTech [6]. It makes use of declarative EPL language for processing "streams of events" and discovering complex patterns among the data. This technique can be seen as a mixture of Database Management System and Rule Engines: it allows applying independent SQL-like queries (*rules* or *directives*) to the streams of events as they come (in real-time manner). Generally EPL statements have the form of

```
select <event properties> from <streams>
group by <aggregating property> having <aggregation function>
where <additional selector>
```

ESPER provides a rich set of available processing functions (including time-based) for selection (filtering), aggregation, correlation, grouping, time and size data windows, defining complex patterns and contexts. It is available as a library (GPL), thus allowing straight forward integration with the existing TDAQ software framework. ESPER is a scalable, multi-threaded framework, proven to be high-performant: a single SA application can process incoming events at the rate of $O(10^5)$ Hz. The same library is also used in another knowledge-base TDAQ application - CHIP (expert-system like engine for automation of control actions in TDAQ).

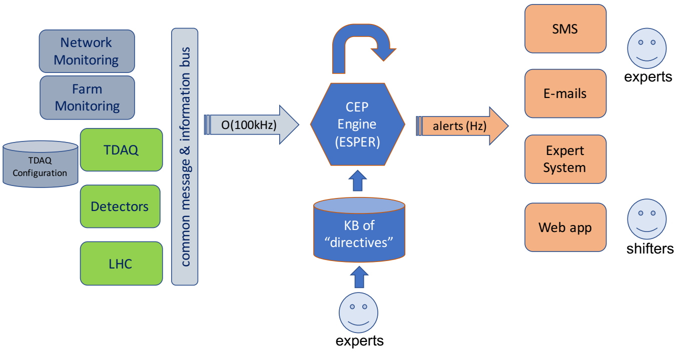

Figure 2. Architecture of Shifter Assistant

4 SA knowledge base and examples of directives

The SA knowledge base (KB) - a set of directives - was mostly developed during Run 2. Presently it contains more then 290 directives, which can be classified in the following categories:

- assessing system health, monitoring and analyzing system parameters trends, thus providing users with "focused" and pertinent monitoring data;

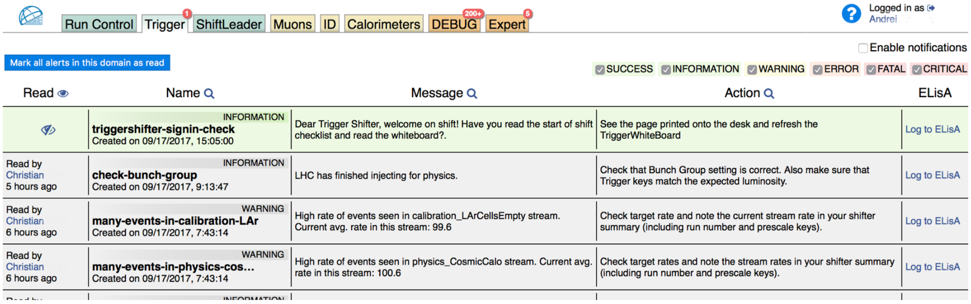

Read	Name	Message	Action	ELisA
	INFORMATION **triggershifter-signin-check** Created on 09/17/2017, 15:05:00	Dear Trigger Shifter, welcome on shift! Have you read the start of shift checklist and read the whiteboard?.	See the page printed onto the desk and refresh the TriggerWhiteBoard	Log to ELisA
Read by Christian 5 hours ago	INFORMATION **check-bunch-group** Created on 09/17/2017, 9:13:47	LHC has finished injecting for physics.	Check that Bunch Group setting is correct. Also make sure that Trigger keys match the expected luminosity.	Log to ELisA
Read by Christian 6 hours ago	WARNING **many-events-in-calibration-LAr** Created on 09/17/2017, 7:43:14	High rate of events seen in calibration_LArCellsEmpty stream. Current avg. rate in this stream: 99.6	Check target rate and note the current stream rate in your shifter summary (including run number and prescale keys).	Log to ELisA
Read by Christian 6 hours ago	WARNING **many-events-in-physics-cos...** Created on 09/17/2017, 7:43:14	High rate of events seen in physics_CosmicCalo stream. Current avg. rate in this stream: 100.6	Check target rates and note the stream rates in your shifter summary (including run number and prescale keys).	Log to ELisA

Figure 3. A screenshot of SA web application with few alerts (one not yet acknowledged)

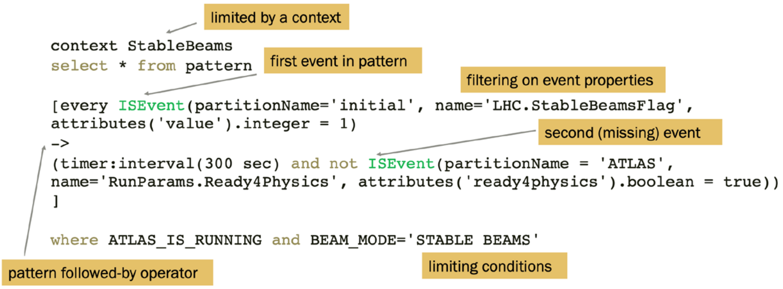

```
context StableBeams
select * from pattern

[every ISEvent(partitionName='initial', name='LHC.StableBeamsFlag',
attributes('value').integer = 1)
->
(timer:interval(300 sec) and not ISEvent(partitionName = 'ATLAS',
name='RunParams.Ready4Physics', attributes('ready4physics').boolean = true))
]

where ATLAS_IS_RUNNING and BEAM_MODE='STABLE BEAMS'
```

Figure 4. Example EPL of a SA directive using EPL pattern

- detecting and diagnosing misbehaviors: the diagnosis is based on detection of error patterns and correlation of events from different sources; no additional intrusive tests are launched. SA does not make any recovery actions (which is responsibility of another component), but can provide necessary data for that.

- focusing on important messages and events, requiring immediate action from the shifters, otherwise data or data quality loss is possible;

- implementing check-list like reminders, instructing shifters about imminent actions, sending SMS messages to relevant experts.

New directives are being added to the KB on daily basis, addressing urgent operational issues. Most of the directives can be developed in a few minutes, however in certain cases it may take hours, especially when the validation of the directives is more difficult. About 10 experts from different groups and domains were involved in the development of directives: some were providing knowledge for an EPL engineer who developed EPL code, some were writing or changing EPL code based on existing templates and directives.

In Figure 4 you can see EPL code of a directive from SA KB, called *atlas-not-ready-in-physics*. It detects a condition when ATLAS is not running in Physics mode within 5 minutes after Stable Beams are declared by the LHC, and demonstrates use of a powerful EPL *pattern* when an event is *not* followed by another event within 5 minutes.

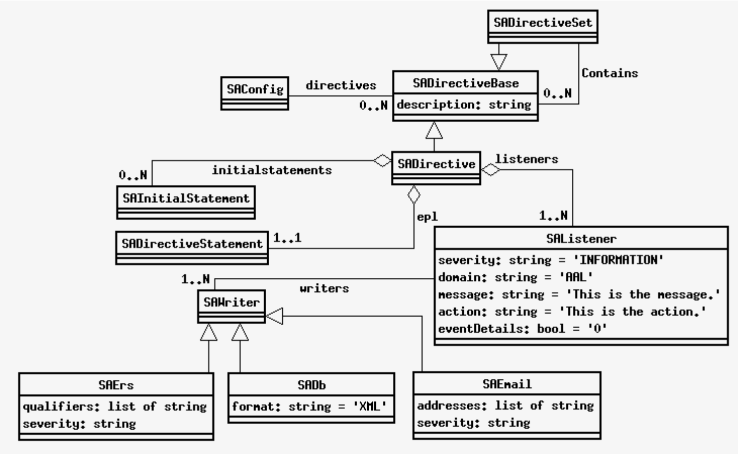

Figure 5. Object schema describing SA directives as part of a TDAQ configuration DB

5 Directives management (configuration)

Presently SA KB includes more than 290 directives, which corresponds to 2,500 lines of EPL code. In order to simplify KB management, and also to facilitate integration of SA framework with TDAQ tools, the SA KB is stored in the TDAQ Configuration service [7]. A database object schema, used to describe the directives is presented in Figure 5. The central *SADirective* class links *SADirectiveStatement* and *SAInitialStatement* classes (containing the EPL code for this directive) with *SAListener* class which defines the content and links with the destination of the alert, *SAWriter*. The latter can be one of *SAErs* (forwards alerts to common TDAQ Error Reporting Service), *SADb* (stores alerts in a DB for further use in the SA web application) or *SAEmail* (sends out alerts as e-mails to recepients). Integration with the TDAQ tools also allowed to implement a dynamic way of updating the directives: when a KB is changed, the engine is not restarted, thus providing zero downtime of the service.

6 SA Replay

Debugging and validating the directives is not trivial (especially for very rare conditions). To improve this procedure, a dedicated application called *SA Replay* [8] was developed. It uses archives of operational monitoring data for replaying the conditions and events streams of a particular run. A user can select a Run Number and time interval, and a set of directives for replay (Figure 6). During a replay execution, SA engine is controlled by the internal clock and fed by the data from the archives. Alerts are recreated in historical time-line, as if they would have happened during that period of time.

7 Conclusions

Taking data with the ATLAS detector is a complex and constantly evolving challenge. TDAQ is responsible for maintaining high data taking efficiency and requires intelligent operational monitoring tools to help human shifters. Shifter Assistant plays an important role in daily operations, applying the expert knowledge to the streams of operational data and detecting system misbehaviors. It is a performant and scalable application, matching TDAQ needs: it is

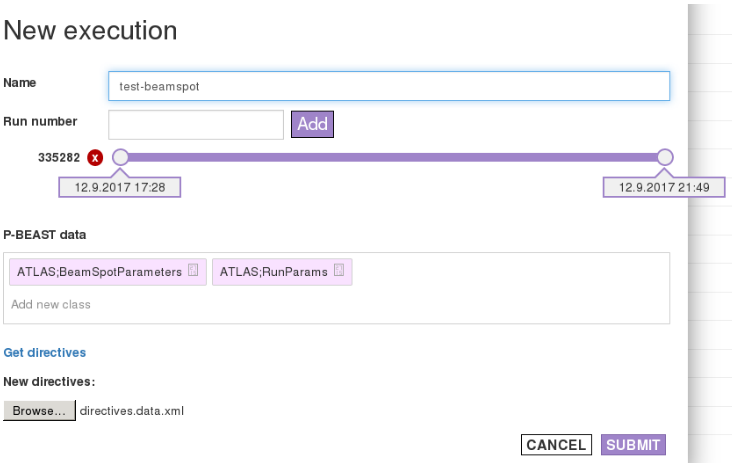

New execution

Name: test-beamspot

Run number:

335282

12.9.2017 17:28

12.9.2017 21:49

P-BEAST data: ATLAS;BeamSpotParameters, ATLAS;RunParams

Get directives

New directives: directives.data.xml

Figure 6. Screenshot of SA Replay web application

able to apply a set of 290 EPL directives to the streams of TDAQ operational data of $O(10^5)$ updates per second in a real-time manner. SA knowledge base is developed and maintained by TDAQ and subsystem experts on daily basis, addressing new operational issues as they arrive. SA Replay is an important application for validation of SA directives. Introduction of SA in LHC Run 2 greatly reduced the operational load on the collaboration, with decreased dependence on experts and removal of one control room shift task.

References

[1] The ATLAS Collaboration (ATLAS), JINST **3**, S08003 (2008)

[2] M. Abolins et al. (ATLAS TDAQ), JINST **11**, P06008 (2016)

[3] A. Kazarov, G. Lehmann Miotto, L. Magnoni, J. Phys. Conf. Ser. **368**, 012004 (2012)

[4] *Complex Event Processing*, https://en.wikipedia.org/wiki/Complex_event_processing, accessed: 2018-09-30

[5] G. Anders, G. Avolio, G.L. Miotto, L. Magnoni, J. Phys. Conf. Ser. **608**, 012007 (2015)

[6] *ESPER*, http://www.espertech.com, accessed: 2018-09-30

[7] G. Lehmann Miotto, I. Soloviev, M. Dobson, J.E. Sloper, J. Almeida Simoes, A. Kazarov, R. Coura Torres, J. Phys. Conf. Ser. **119**, 022004 (2008)

[8] A. Santos, G. Anders, G. Avolio, A. Kazarov, G.L. Miotto, I. Soloviev, J. Phys. Conf. Ser. **664**, 062055 (2015)