

EXPERIENCE WITH SHIFTER ASSISTANT

AN INTELLIGENT TOOL TO HELP OPERATIONS IN ATLAS
TRIGGER-DAQ SYSTEM

Andrei Kazarov (NRC “Kurchatov institute” - PNPI Gatchina / CERN)



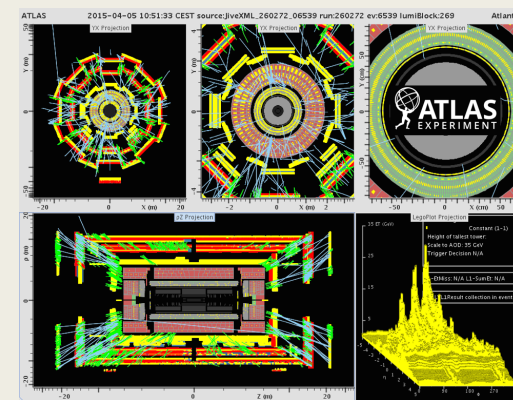
CHEP 2018, Sofia, Bulgaria

Content

- ATLAS detector and Trigger-DAQ system
- ATLAS Operations and challenges
- Technology: Complex Event Processing, ESPER engine
- Shifter Assistant: architecture
- Examples of SA directives
- Developing, debugging and validating the directives
- Conclusions

ATLAS Trigger and Data Acquisition system

- ATLAS: General purpose detector on LHC with wide physics search goals
- Trigger-DAQ system
 - buffers, transfers, selects, builds ATLAS events
 - decreasing the event rate from 40MHz of LHC collision rate down to 100kHz (level-1 trigger, h/w based)
 - high-level trigger selects interesting events decreasing the rate further down to 3kHz of recorded (1.5MB) events
 - s/w based, runs on a farm of commercial PCs
- Complex distributed computing system
 - 2500 computer nodes (including 2000 nodes of HLT farm), Ethernet network
 - 60000 supervised applications (including 47000 HLT processing tasks), coherently performing data acquisition, event selection, control, configuration and monitoring tasks



ATLAS Operations

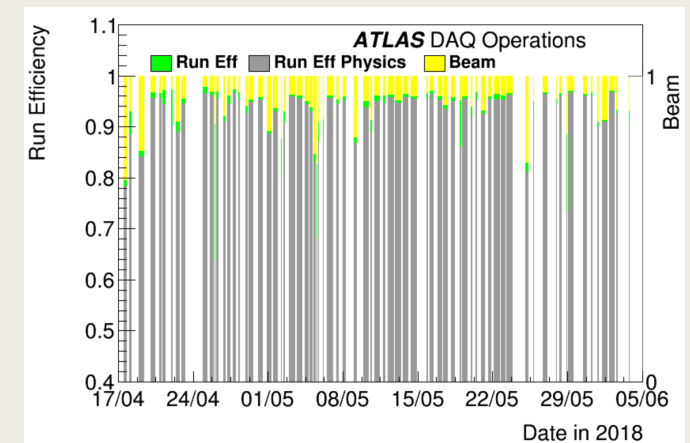
- system is operated by crew of shifters (7), assisted by experts-on-call (O(10))
- dedicated tools provided to the operators, helping them to control and monitor the system
- maintain high **run efficiency** (low dead-time) is the primary goal for the operations
 - focus on minimalizing the impact of **human errors**
- run efficiency **93.7%** in 2018 (as of June 5th) out of 97.8% achievable efficiency

Challenges

- system complexity and heterogeneity
- likely error conditions from custom and commodity hardware and software
- long development and maintenance time
 - decades of operations
 - need to preserve expert knowledge
- lots of operational data to analyze
 - ~1M messages per 10-20 hours of a typical physics run
 - O(10⁵) operational parameters, rapidly changing (rates up to 200kHz)

Solutions

- **automate** routine actions, problem diagnosis, operational data analysis
- add **intelligence** to the controls: **collect**, **store** and **reuse** expert knowledge about the system behavior



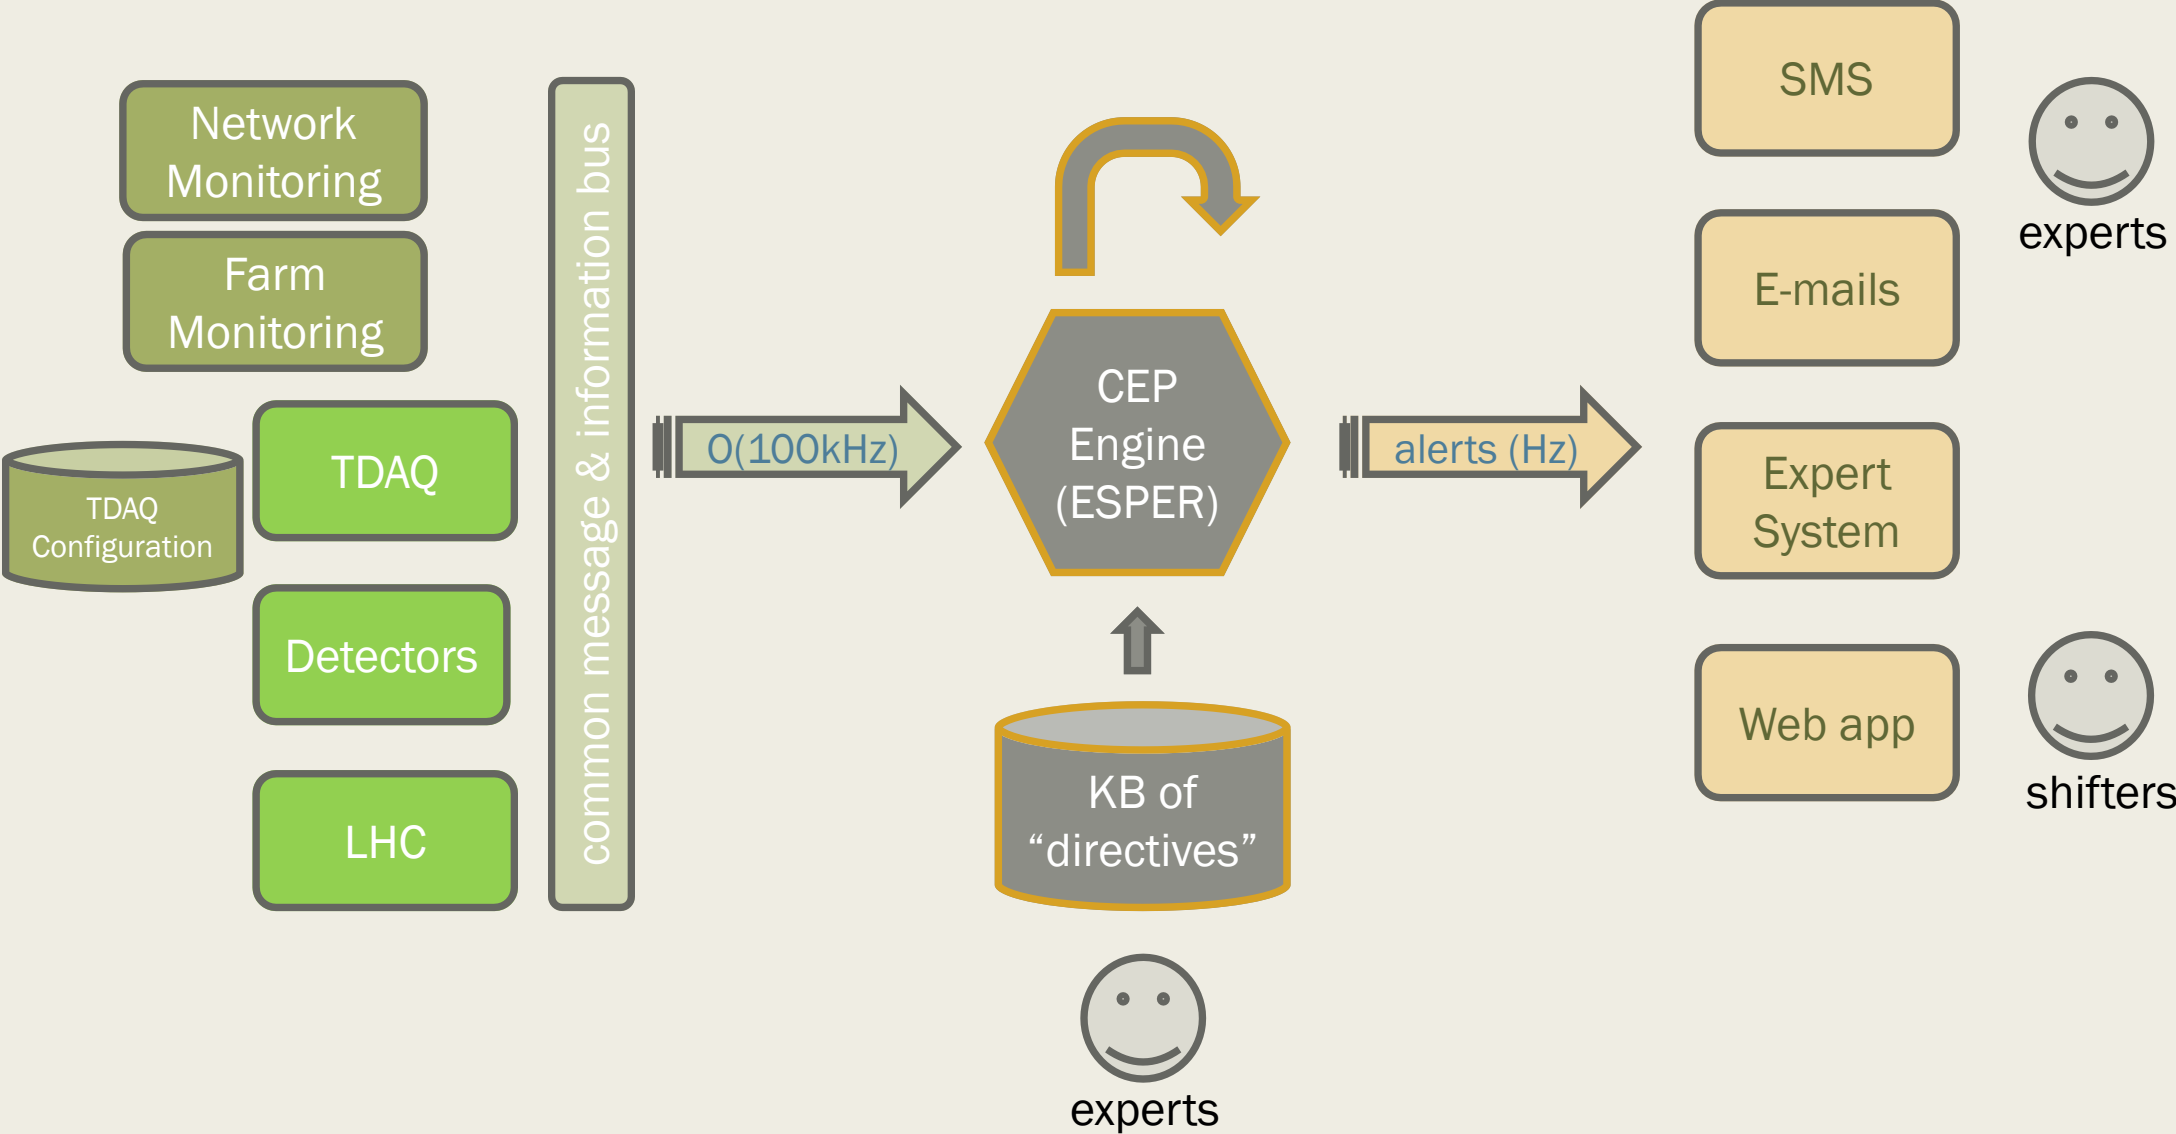
Shifter Assistant: intelligent operational monitoring

- aim is to **automate** complex operational monitoring and problem diagnosis
- apply **expert knowledge** to streams of operational data (operational monitoring events and messages)
 - leverage on Complex Event Processing (CEP) technology, implemented in ESPER (next slide)
- assess system health, monitor and analyze system parameters trends
- detect and diagnose misbehaviors
 - diagnosis is based on detection of error patterns, correlation of events
 - does not launch additional tests
 - does not make any recovery actions (responsibility of another component)
- focuses on important (requiring immediate action from the shifters, otherwise data or data quality loss is possible) messages and events, which otherwise can be missed by the shift crew
- implements check-list like reminders, instruct shifters about imminent actions
- produces operator-friendly alerts, including a message and an action to be taken by the operator
- **configurable**, fully integrated in DAQ controls & monitoring framework
- removed the need of one full time 24h/7d shifter (DAQ monitoring) in the Control Room

Technology: Complex Event Processing engine: ESPER by EsperTech

- A framework to process events and discover complex patterns among streams of events
 - aim is to allow “detecting situations”
 - formalizing it with declarative EPL language
- Used in financial analysis, business process management, network monitoring
- can be seen as a mixture of Data Base Management System and Rule Engines:
 - applies independent SQL-like queries (“rules” or “directives”) to the streams of events as they come (in real-time manner)
select ... from ... where ... group by ... having ...
- rich set of available processing functions (including time-based)
 - selection (filtering), aggregation, correlation, grouping, time/size data windows
 - patterns, contexts
- ESPER: implementation of CEP technology (java) from ESPERTECH
 - a lot of applications and customers: <http://www.esperTech.com/customers/>
- available as a library (GPL), allowing integration with DAQ framework
 - high-performant, scalable, multi-threaded: a single application can process $O(10^5)$ updates/sec
- also used in another knowledge-base TDAQ application - CHIP (expert-system like engine for automation of control actions in TDAQ)

SA architecture



SA web application

The screenshot shows the SA web application interface. At the top, there is a navigation bar with tabs for 'Run Control', 'Trigger' (with a red notification badge '1'), 'ShiftLeader', 'Muons', 'ID', 'Calorimeters', 'DEBUG' (with a red notification badge '200+'), and 'Expert' (with a red notification badge '5'). On the right, it says 'Logged in as Andrei Kazarov' with a question mark icon. Below the navigation bar, there is a button 'Mark all alerts in this domain as read' and a notification settings section with checkboxes for 'Enable notifications', 'SUCCESS', 'INFORMATION', 'WARNING', 'ERROR', 'FATAL', and 'CRITICAL'. The main content is a table of alerts with columns: 'Read' (with an eye icon), 'Name' (with a search icon), 'Message' (with a search icon), 'Action' (with a search icon), and 'ELisA'.

Read	Name	Message	Action	ELisA
	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	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	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	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

- every desk in ATLAS Control room has a dedicated tab (“domain”)
 - focus on alerts only relevant for this desk
- alerts must be **acknowledged** after the action is taken
 - helps to track shifters activity by remote experts
- available anywhere, in the control room and outside (for experts)
- pop-up notification on new alerts available in browsers (smart phones)

SA knowledge base (“directives”)

- Presently SA KB includes 290 directives, about of 2300 lines of EPL code
 - quite compact KB, not difficult to maintain
- Directives are grouped by domains corresponding to ACR desks: Run Control, Shift Leader, Trigger, ID, Calorimeters, Muons, Luminosity
- New directives are being added regularly
 - to address new issues (until a proper fix is developed) and to help operators to react quickly
 - most of the directives can be developed in few minutes, some may take hours
 - directives validation may be more difficult
- About 10 experts from different groups/domains in total were involved in development of directives
 - some were providing knowledge for an EPL engineer who developed EPL code
 - some where writing or changing EPL code
- Short overview of developed directives – see backup slide 16

EPL code example

detect when ATLAS is not running in Physics mode within 5 minutes after Stable Beams declared by LHC, i.e. a pattern when an event is NOT followed by another event within 5 minutes:

```
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'
```

limited by a context

first event in pattern

filtering on event properties

second (missing) event

pattern followed-by operator

see 2 more
examples in
backup slides


SA Replay

- Debugging and validating the directives is not trivial (especially for very rare conditions)
- SA Replay – a web application - uses archives of operational monitoring data for **re-playing** the conditions of a particular run
- User can select a RN & time interval
- User provides set of directives to replay
- Alerts are re-created in historical timeline, as if they would have happened

New execution

Name:

Run number:

335282 x 

P-BEAST data

[Get directives](#)

New directives:

directives.data.xml

child-monitors	AAL.TDAQ.RunControl	Aug. 23, 2017, 2:41 p.m. (1503492066399)	child-monitors	Number of child monitors is too high in the Global monitoring chain. This may break event sampling.	Check with DQ desk and if there is a problem with histograms, restart Global Monitoring segment.	WARNING
ChipRecovery	AAL.COM	Aug. 23, 2017, 2:48 p.m. (1503492519000)	ChipRecovery	Recovery [REMOVAL, REMOVAL,] for component(s) [AppI_PixelRCD-IBL/ROD_I1_S20, AppI_PixelRCD-IBL/ROD_I1_S20,] status: [NEW, DONE,]		INFORMATION

Summary

- Taking data with the ATLAS Detector is 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
- A performant & scalable application, matching TDAQ needs: applies set of 290 rules to the operational stream of data of ($O(10^5)$) updates/secs in real-time manner
- SA knowledge base is developed and maintained by TDAQ and subsystem experts on daily basis, addressing new issues
- SA Replay is an important application for validation of SA directives
- Greatly reduced the operational load on the collaboration, with decreased dependence on experts and removal of one control room shift task

Backup

- list of directives in SA KB
- examples of directives with EPL code

KB of “directives”: overview

■ Run Control and Data Flow:

- *monitor errors in vital infrastructure applications (not directly supervised by operators)*
- *constantly increasing memory consumption by an application*
- *detect misconfigurations early*
 - *missing parts of configuration*
 - *wrong Run parameters in physics mode*
- *high permanent average dead time (data loss) > 3%*
- *detect HLT farm saturation and lack of processing power*
- *detect busy from DAQ and determine backpressure sources (HLT processing, data collection, data recording)*
- *unbalanced SFO farm (a small farm of storage nodes for recording selected events on disk)*
- *SFO disk space occupation*
- *overloaded Read-Out system, high occupancy of detector readout links*
- *host unreachable, network switch port failures*

■ Trigger

- *unreasonable physics event rates in different trigger streams*
- *wrong trigger configuration (LHC bunch grouping scheme) at the end of LHC injection*

■ Shift Leader

- *ATLAS not in RUNNING state at LHC injection*
- *no data is taken in 5 minutes of stable beams (SMS also sent to Run Manager)*

■ ATLAS Detectors

- *misconfiguration (wrong tag in Pixel Infrastructure)*
- *Suggest restart of a LAr detector readout after 2+ processing elements were removed in the current run*
- *Luminosity ratio is outside of allowed range*

EPL (Event Processing Language) examples

detect misconfiguration of a detector after a calibration run

```
select * from  
PartitionState(partitionName in ('ATLAS'), state = 'CONFIGURED', inerror  
= false) as ATLAS_state unidirectional,  
method:ISReader.getEventByName("PixelInfr",  
"PixelRunParams.DataTakingConfig-ModCfgTag") as pixel_tag,  
method:ConfigReader.isSegmentEnabled("ATLAS", "Pixel") as Pixel_enabled  
where pixel_tag.attributes("value").string = 'PIT_MOD_STANDBY' and  
Pixel_enabled.boolean = true
```

when ATLAS is in CONFIGURED state

and Pixel detector is
enabled in ATLAS

configuration of Pixel
does not match physics

EPL examples II:

detect misbehaving element (among many), producing wrong results and requiring restart

original stream + filtering

create new stream of std.dev. of a parameter in a sliding window of 20 secs

```
insert into EbEventsStdDev
select
stddev(rack_info.attributes("EbEvents").long)/avg(rack_info.attributes("EbEvents").long) as eb_rate_dev_relative
from ISEvent(partitionName="ATLAS", type="DCM", name regexp "DF_IS:HLT.DefMIG-IS:.*:tpu-rack-.*.info").win:time(20 sec) as rack_info
group by name
having avg(rack_info.attributes("EbEvents").long) > 0
output first every 30 sec
```

aggregating: group by name

in a sliding window of 20 secs

```
select *,
min(eb_rate_dev_relative) as min_dev, count(*) as evcount
from EbEventsStdDev.win:time(180 sec) where eb_rate_dev_relative > 0.6
group by Rack having count(*) > 3
output first every 5 min
```

from the new stream, select Racks which produce wrong results for longer then 1.5 min

EXAMPLE OF A SA DIRECTIVE IN DB EDITOR

The screenshot displays four panels in a DB Editor, illustrating the configuration of an SA Directive. The panels are:

- Panel 1: SADirectiveStatement@eprcode-rc-tree-busy-3min**
 - Object ID: `eprcode-rc-tree-busy-3min`
 - Class: `SADirectiveStatement`
 - Data File: `/atlas-home/1/akazarov/runtime/directives/daq/runctrl1.data.x`
 - Attributes: `code`
 - Code:

```
context ATLASRunning
select *
from pattern [
every ISEvent(partitionName = 'ATLAS', name='RootController', type='RCStateInfo'
->
(timer:interval(3 min)
and not ISEvent(partitionName = 'ATLAS', name='RootController', type='RCStateInfo'
output every 15 minutes
```

- Panel 2: SADirective@rc-tree-busy-3min**
- Object ID: `rc-tree-busy-3min`
- Class: `SADirective`
- Data File: `/atlas-home/1/akazaro`
- Attributes: `description`
- Relationships:
 - `ep1`: `[eprcode-rc-tree-busy-3min@SADirectiveStatement]`
 - `listeners`: `[list0-rc-tree-busy-3min@SAListener]`
- Panel 3: SAListener@list0-rc-tree-busy-3min**
- Object ID: `list0-rc-tree-busy-3min`
- Class: `SAListener`
- Data File: `/atlas-home/1/akazaro`
- Attributes:
 - `severity`: `ERROR`
 - `domain`: `AAL.TDAQ.RunControl`
 - `message`: `RC tree in ATLAS stays is BUSY :`
 - `action`: `Figure out which node in the tree is stuck`
 - `eventDetails`: `true`
- Relationships:
 - `writers`: `[file0-rc-tree-busy-3min@SAFile]`, `[db0-rc-tree-busy-3min@SADB]`
- Panel 4: SADB@db0-rc-tree-busy-3min**
- Object ID: `db0-rc-tree-busy-3min`
- Class: `SADB`
- Data File: `/atlas-home/1/akazaro`
- Attributes: `format`: `XML`

Arrows indicate the following relationships:

- From Panel 1 to Panel 2: The `code` attribute of the `SADirectiveStatement` is linked to the `ep1` relationship of the `SADirective`.
- From Panel 2 to Panel 3: The `listeners` relationship of the `SADirective` is linked to the `list0-rc-tree-busy-3min@SAListener` object.
- From Panel 3 to Panel 4: The `writers` relationship of the `SAListener` is linked to the `db0-rc-tree-busy-3min@SADB` object.