



Operational Intelligence



Optimizing computing operations



*Daniele Bonacorsi, Alessandro De Salvo, Alessandro Di
Girolamo , **Federica Legger**, Lorenzo Rinaldi*



WLCG




- 200 sites
- >40 countries
- 750000 cores
- 2 million jobs/day
- 600 PB storage
- 10-100 GB links



Operations in a nutshell



Dashboards,
logs, elogs,
tickets



Email, JIRA,
GGUS, chat,
meeting

- Process information
- Triage
- Take decisions
- Escalate/fix
- Record actions

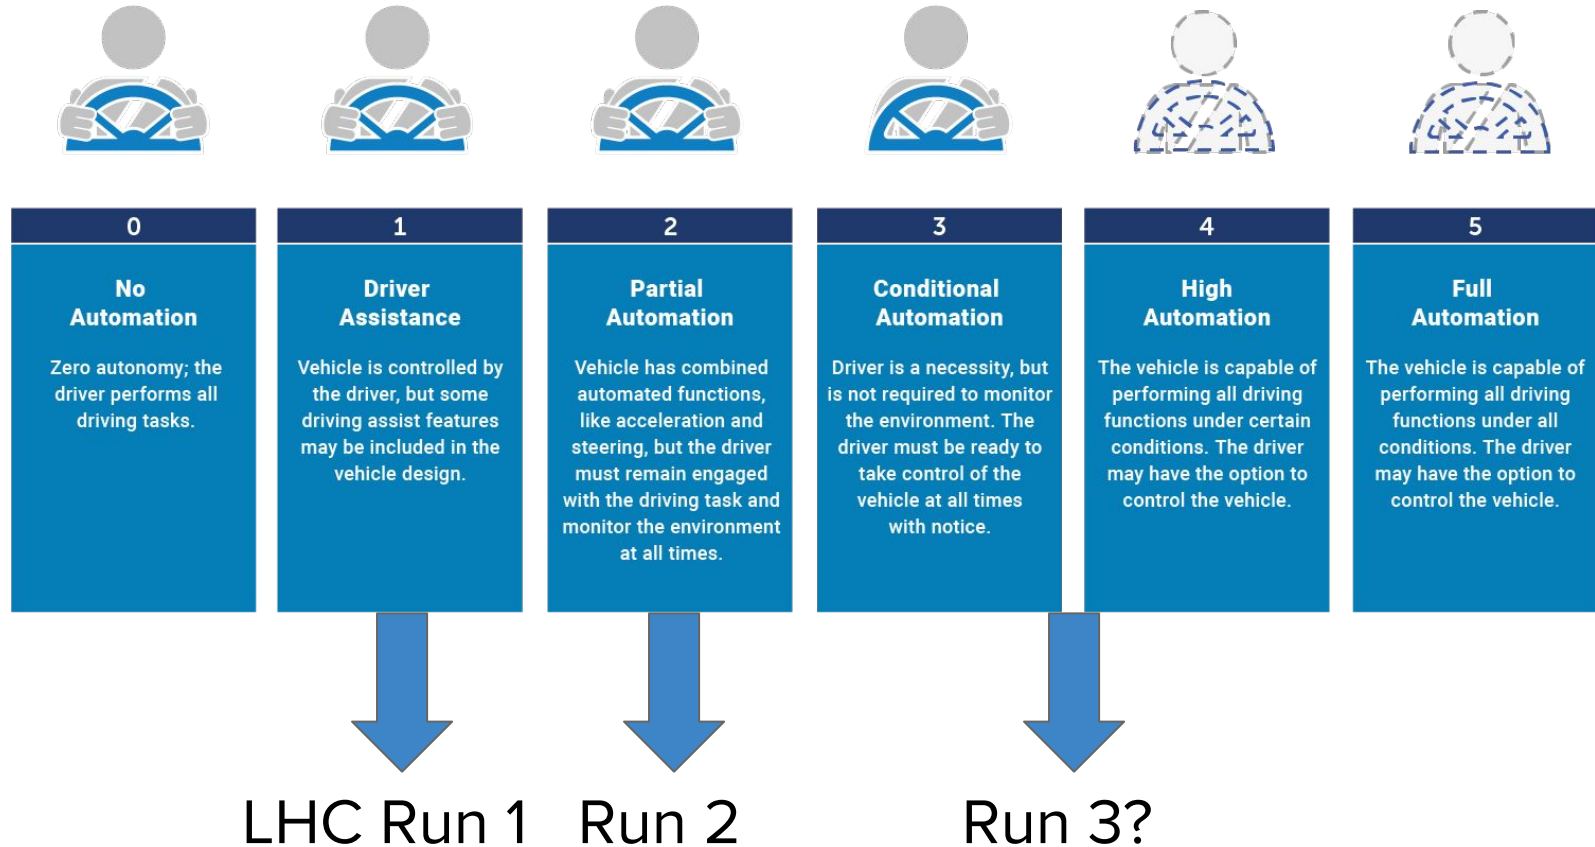
ATLAS/CMS report 100 people involved in computing operations
(50 FTEs/experiment)!

In 1 year, > 1k GGUS tickets for ATLAS, > 2k for CMS

Can we do better?



- The LHC experiments built a computing system that **worked** in LHC Run-1/2.
 - At which depth do we fully “**understand**” it?
 - Can we perform precise modelling of specific workflows / site behaviours / systems performances?
 - Can we use this modelling to make predictions (e.g. population vs pollution of Tier disks; TierX - Tier-Y data transfer patterns; ..)
 - For long, we monitored to debug in near-time, not to analyse and learn from the past to design and build what’s next.
- Computing operations (meta-)data is all archived.
 - **Only recently started to be accessed.**
 - e.g. transfers, job submissions, site performances, infrastructure and services behaviours, storage accesses, ..



- A cross-experiment effort aiming to:
 - Streamline computing operations
 - **save manpower & improve resources utilization**
 - Increase level of automation in operation tasks
 - Cost reduction metrics: needed number of operators
- By:
 - Identifying common projects
 - leveraging **common** tools/infrastructure
 - Collaborate, share expertise, tools & approaches
 - Across experiments
 - Across teams (operations, monitoring, analytics)
 - **Bottom-up** approach


- Computing systems mature and well-understood
- Clear request from funding agencies: push on commonalities
- Easier to interest students/engineers (with background different from HEP) to work on topics using industry standard tools
- ATLAS/CMS use common analytics infrastructure@CERN
- More experiments starting (or considering) using LHC-developed tools
 - for example Rucio, and FTS
 - Share efforts with wider (than LHC) community

- We successfully started OpInt activities
 - **Kickoff meeting** at HOW19
 - Regular biweekly meetings
 - For now **CMS, ATLAS, HammerCloud, Rucio, MONIT, DUNE/FNAL, LHCb**
- **Start "simple"**
 - targeting well-identified projects with precise goals/metrics
 - Must show operational cost reduction
 - Must have **operation people** on board
 - Guinea pig: **Rucio**
 - But not limited to (i.e. Data Management all round)
 - Common analytics projects: understand issues with transfers, predict latencies, popularity prediction, ...


The analytics infrastructure @CERN




Data providers



AMQ,
http,
logstash



Data sources

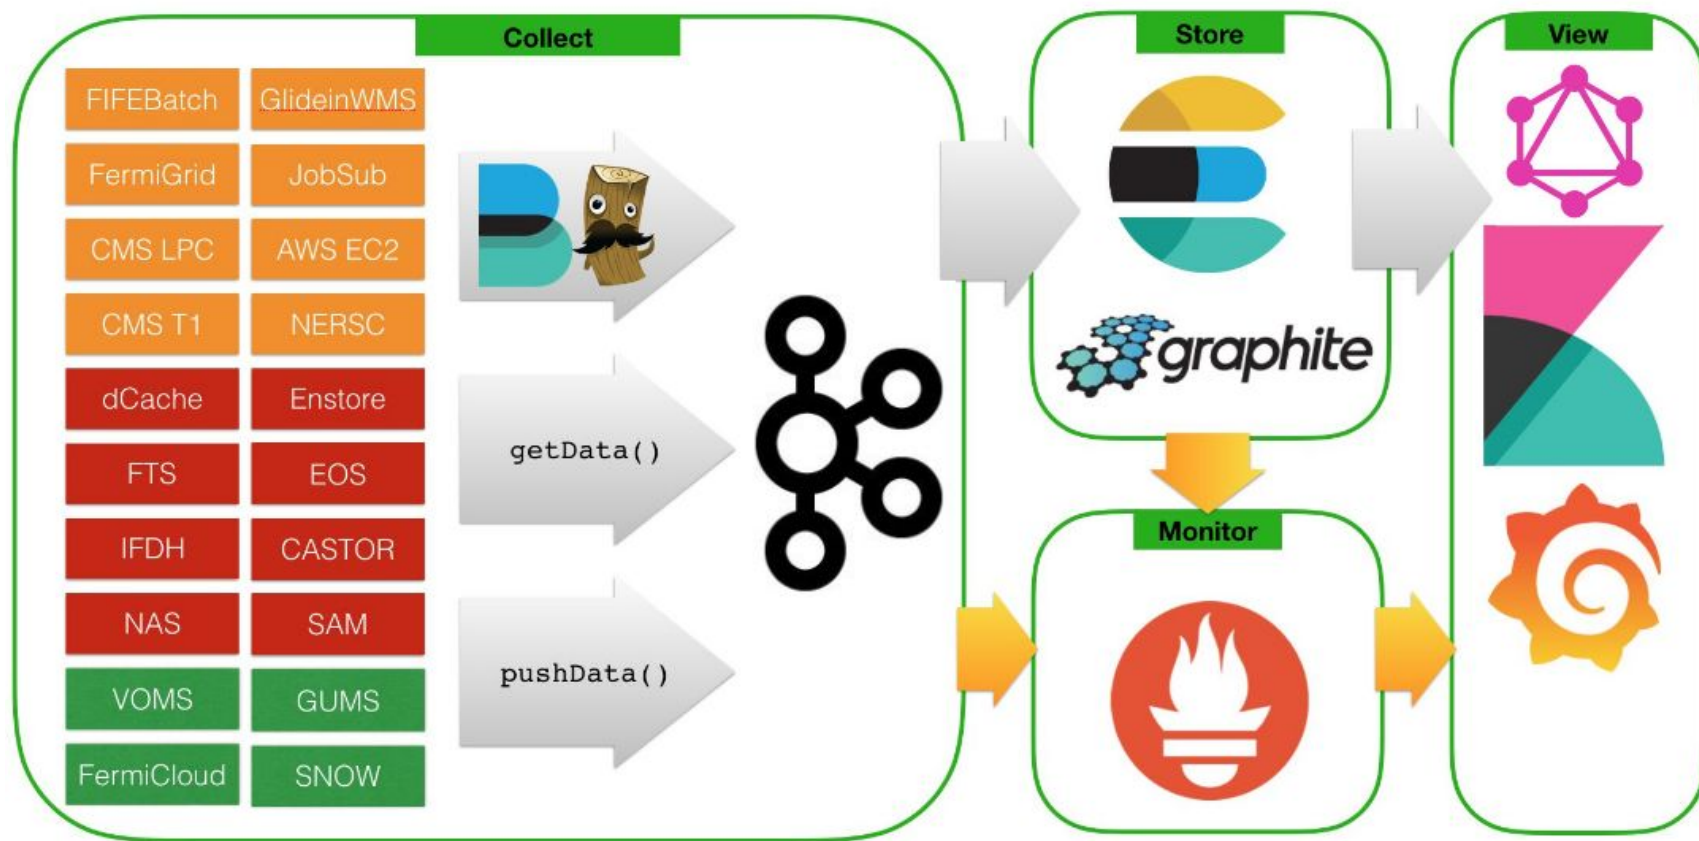


Visualization/analysis



CMSSpark, ...

Landscape @FNAL




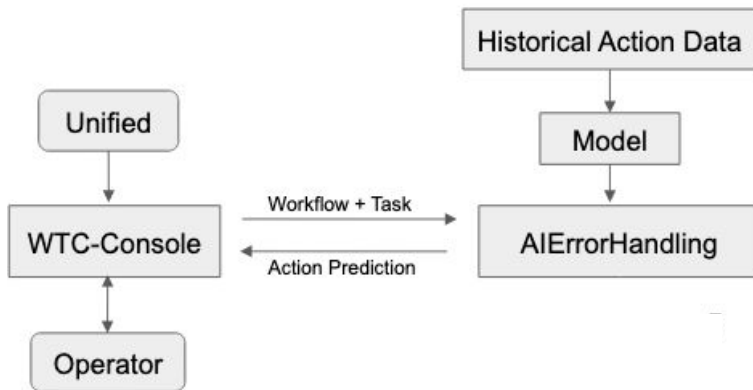
DNNs to predict the action of the operator

- Multi class and Binary (Retry/Non-Retry) classifications
- Imbalanced class distribution
 - Rate of Retry is dominant
 - Effect of resampling method studied
- AUROC ~ 70% is achieved

Possible Actions

- Retry (only failed)
- Kill and Clone
 - With new splitting
 - New settings for memory and cores
- Recovery

	T0 CH_CERN	T1 DE_KIT	T1 ES_PIC	T1 FR_CCIN2P3	T1 IT_CNAF	T1 RU_JINR	T1 UK_RAL	T1 US_FNAL	T2 CH_CERN	T2 CH_CERNBOX	T2 CH_CERN_HLT	T2 DE_DESY	T2 ES_IFCA	T2 FR_GRIF_IRFU	T2 FR_GRIF_IIR	T2 IT_Legnaro	T2 UK_London_Brunel	T2 UK_London_IC	T2 UK_SGrid_RALPP	T2 US_Florida	T2 US_MIT
-1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
84	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
85	0	2	0	0	0	1	6	0	0	0	0	8	1	0	0	0	0	0	0	0	0
92	0	0	1	0	0	0	1	0	0	0	0	0	2	0	0	0	0	0	0	0	0
134	0	0	0	0	0	0	6	0	0	0	0	0	0	0	0	0	0	0	5	0	0
139	0	0	0	0	0	0	2	1	0	0	0	0	0	0	0	0	0	0	0	0	0
8001	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0
8004	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0
50110	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0
50660	0	0	3	2	4	1	1	1	6	63	1	0	1	0	1	0	1	0	1	0	1
50664	0	0	2	7	0	0	2	18	0	32	0	0	7	0	0	4	0	0	0	0	0
71304	0	0	0	0	0	0	0	0	0	102	0	0	0	0	0	0	0	0	0	0	0
99305	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

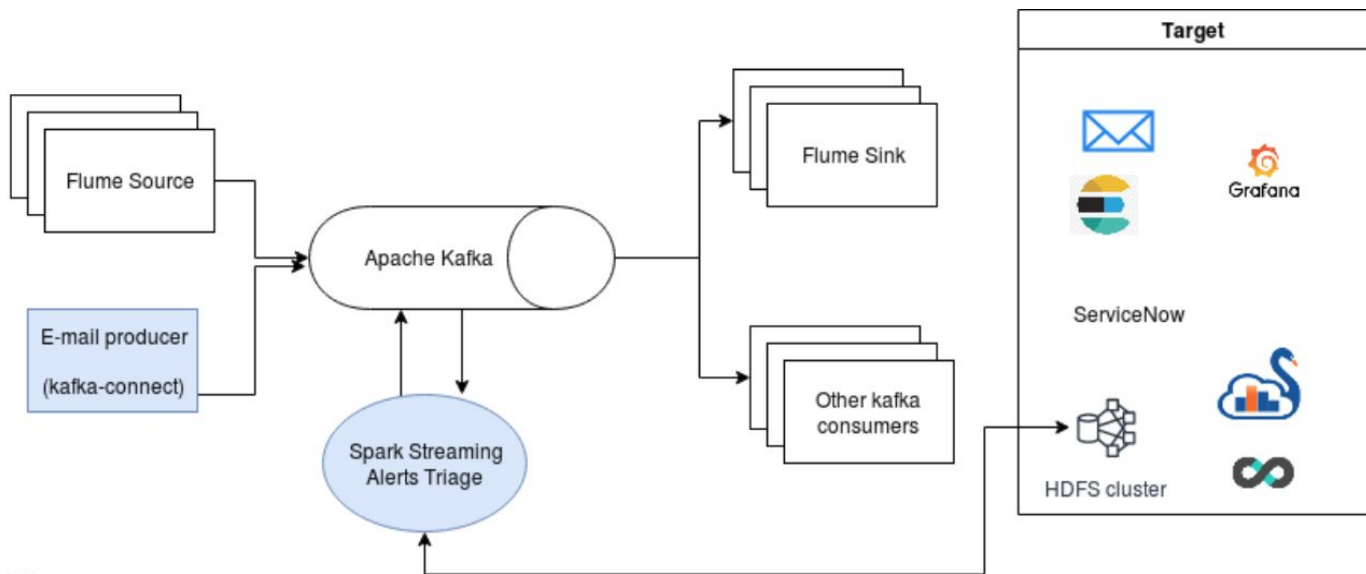


- First project completely experiment agnostic
- Several data streams from multiple subsystems going through MONIT with a common messaging service: **Kafka**

Complex patterns

E.g.:

- Timeout message has been repeated 5 times in a 1 hour window for a given node.
- Memory increase alert but there is not a request increase alert.



Goal: design a predictive intelligent algorithm to send alarms in case of steady state violations (i.e. degraded transfer performances)

- Harmonize and analyze the data transfer metrics
 - Many already collected in Analytics Platforms (ES, ...)
 - Extract new metrics from log files (Rucio, FTS)
- Look for correlations among metrics, using a reinforced learning approach
 - Start simple and then go deep
 - Use correlations to send alarms if potential anomaly are detected
 - shifter/expert will validate

01

- Very reliable dashboards and metric collectors for Data Transfers
- Few systems send notifications (based on “hard-coded” rules)

[illegible]

- 14

01

ATLAS Site	DDMF endpoint	Status																
		Delete	DELETE LAN	DELETE WAN	Fetch	Pilot Log Write	Pilot Read	Pilot Write	Read	READ LAN	READ WAN	Subscription	default SSM	Third party copy	Upload	Write	WRITE LAN	WRITE WAN
CERN-PROD	CERN-PROD_DERIVED														QUOTASPACE	QUOTASPACE		
	CERN-PROD_PFI-ICARMA														QUOTASPACE	QUOTASPACE		
	CERN-PROD_PFI-IBETS																	
	CERN-PROD_PP1PARE	manual									manual				manual			
	CERN-PROD_IAM																	
	CERN-PROD_TIG-DAD														QUOTASPACE	QUOTASPACE		
	CERN-PROD_T2DSX																	
CMS-LEG2	CMS-LEG2_LOCALGROUPSIX														DISKSPACE	DISKSPACE		
CYROMET-LEG2	CYROMET-LEG2_PPS1-IP														QUOTASPACE	QUOTASPACE		
FER-LEG2	FER-LEG2_PFI-IBTRACING														QUOTASPACE	QUOTASPACE		
HFJF-LEG2	HFJF-LEG2_PFI-IBTRON														QUOTASPACE	QUOTASPACE		
INPN-COENQA	INPN-COENQA_DATAIDSR	ACIS			ACIS			ACIS						ACIS	ACIS	ACIS		
	INPN-COENQA_LOCALGROUPSIX	ACIS			ACIS			ACIS						ACIS	ACIS	ACIS		
	INPN-COENQA_SCRATCHDISK	ACIS			ACIS			ACIS						ACIS	ACIS	ACIS		
	INPN-COENQA_T2GROUPDATA_PPS1-TOP														QUOTASPACE	QUOTASPACE		
NAOJ-T1	NAOJ-T1_LOCALGROUPSIX														DISKSPACE	DISKSPACE		
NAOJ-T2	NAOJ-T2_LOCALGROUPSIX														DISKSPACE	DISKSPACE		
RIKEN-LEG2	RIKEN-LEG2_LOCALGROUPSIX	ACIS			ACIS			ACIS							DISKSPACE	DISKSPACE		
SE-ONEC-T2	SE-ONEC-T2_LOCALGROUPSIX														DISKSPACE	DISKSPACE		
TORONTO-LEG2	TORONTO-LEG2_PFI-IBTRON														QUOTASPACE	QUOTASPACE		
T2Instream-INT	T2Instream-INT_LOCALGROUPSIX	manual			manual			manual			manual				manual +	manual +		
WTF	SACROD_DATASOURCE																	
	SACROD_SCRATCHDISK																	
pragmatic2	PRAGMATIC2_LOCALGROUPSIX	manual			manual			manual			manual				manual	manual		

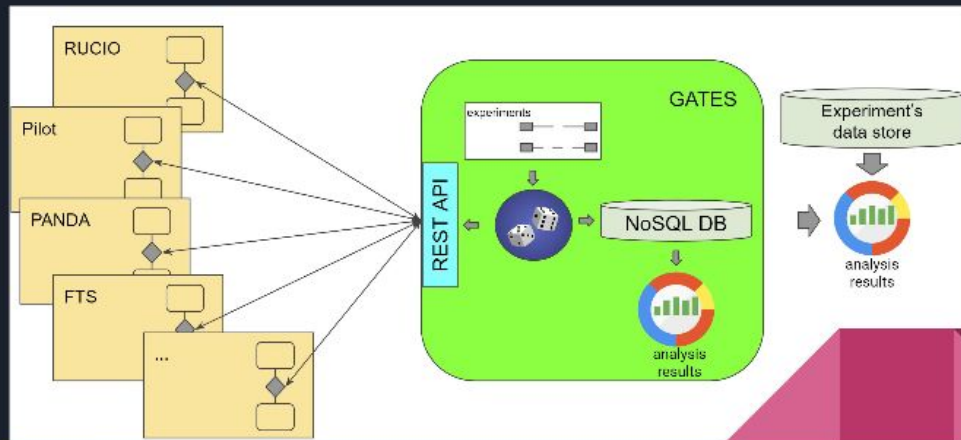
15

Experiments have large data stores collecting data from different computing systems: job scheduling, data distribution, FTS, PerfSONAR, etc.

While that is great for monitoring, accounting, and finding issues, it is not sufficient for the system optimization.

One can try to guess what kind of effect a change will made, but without validation it does not mean much.

We need a way to quickly test different options and get actionable answers.

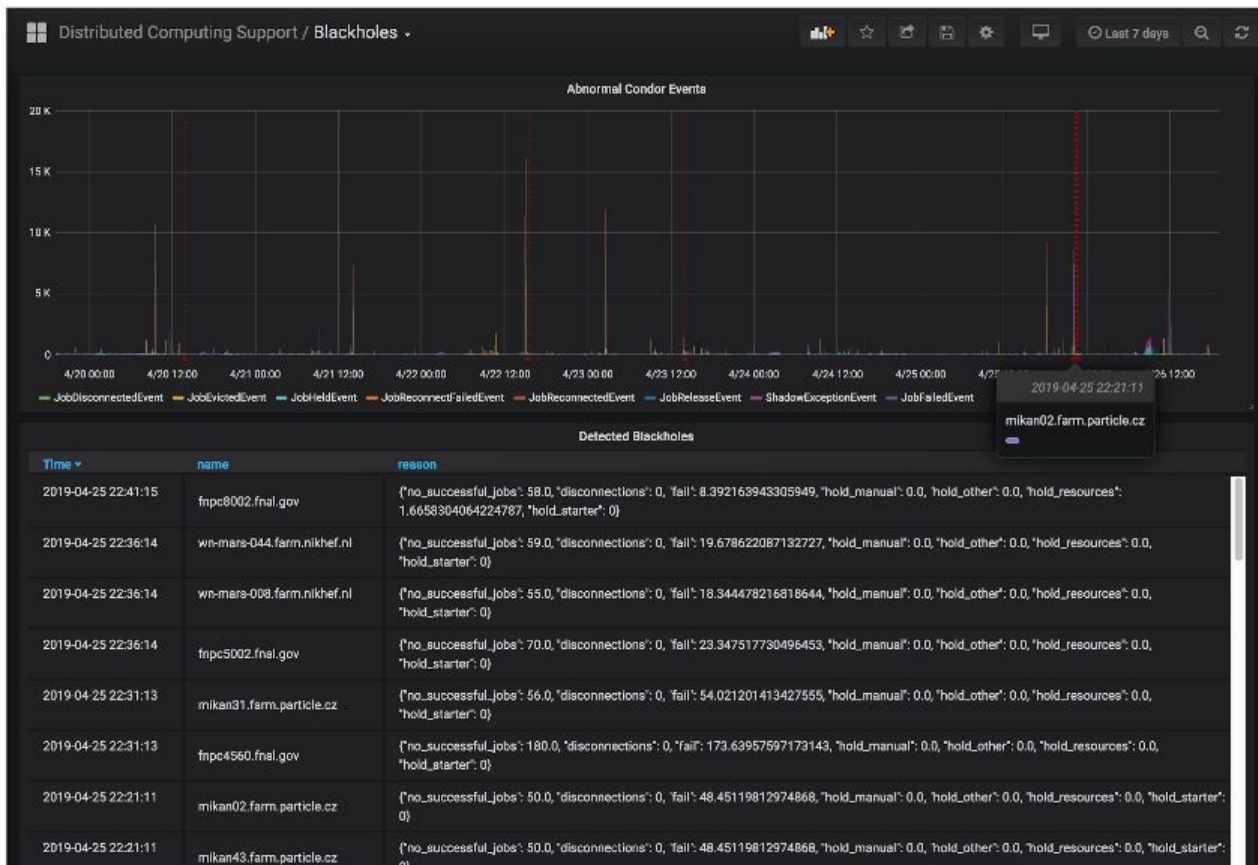


In development.
Completely experiment agnostic.
Contributions welcome.

A/B testing service

Blackhole Node Detection

Kevin Retzke et al.



- Occasionally we encounter “bad” nodes on the grid: hardware issues, missing/broken CVMFS mounts, etc.
- Analyze recent job events to look for telltale signs of bad node
- Challenges:
 - bad node vs. bad user code
 - Transient events (note and move on) vs. persistent issues (blacklist)

Smooth operations at the (multi-)experiment level rely on reliable underlying infrastructures and services

So far:

- **reactive-only** approach to problems after they show up
- rely on (whenever possible) **prompt reaction** to attack and solve issues

Work in progress, e.g. at INFN-CNAF

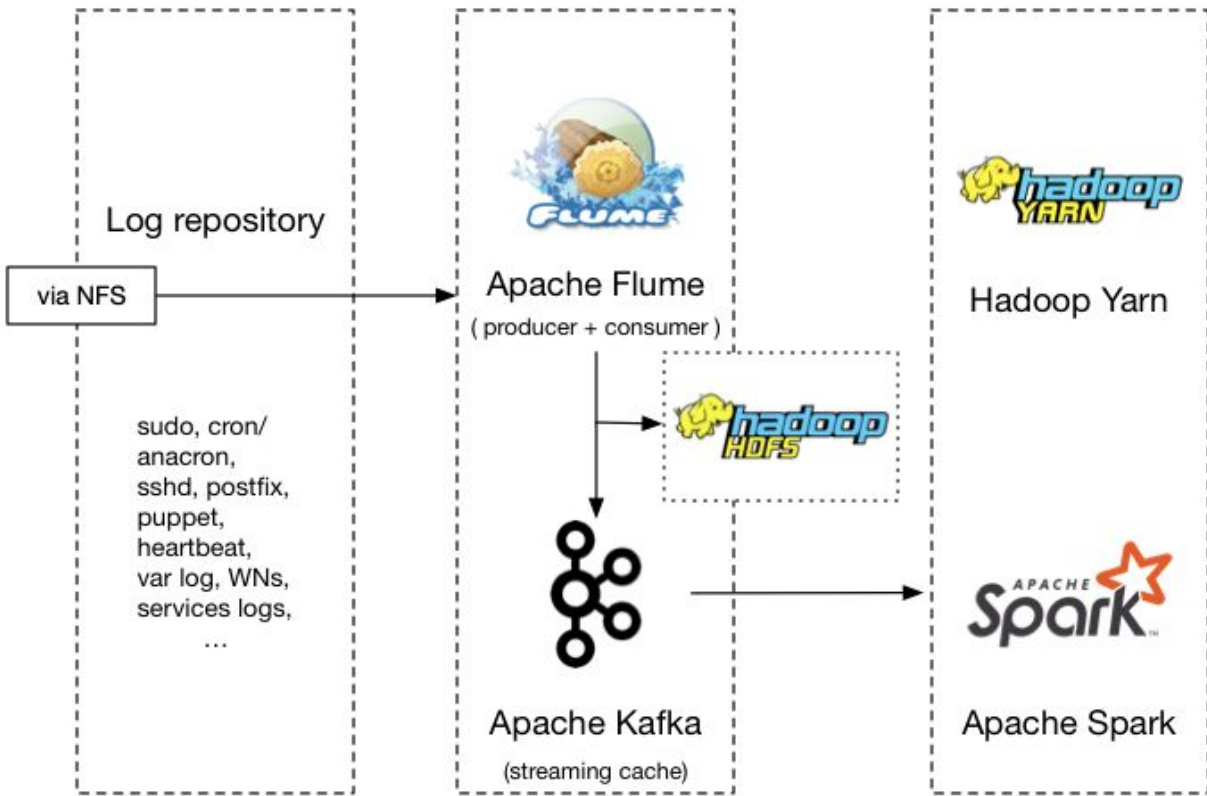
- Ongoing effort to rationalise the *collection of logs* from machines / services asynchronous log analyses (by summer students, service experts, external collaborators)
- first infrastructure work for a *long-term predictive maintenance* approach at INFN-CNAF, potentially exportable to a generic WLCG computing center

Predictive site maintenance

D. Bonacorsi et al.



Infrastructure based on Openstack, and created via DODAS



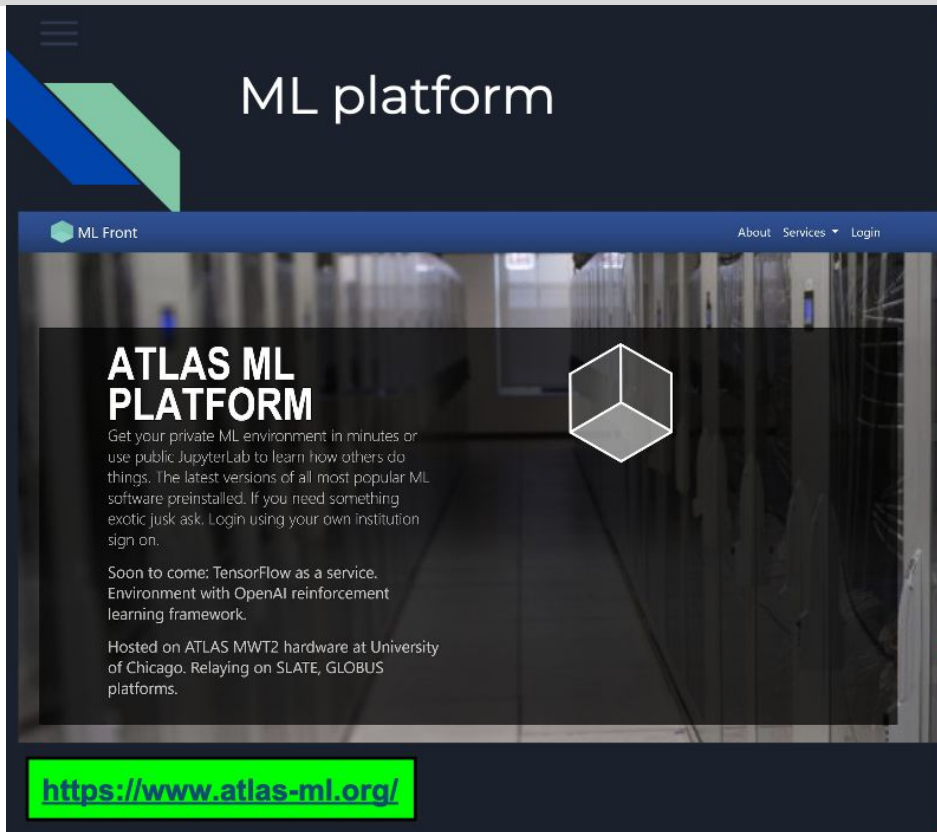
- L.Giommi et al, "Towards Predictive Maintenance with Machine Learning at the INFN-CNAF computing centre", ISGC 2019, Taipei
- T. Diotallevi et al., "Collection and harmonization of system logs and prototypal Analytics services with the Elastic (ELK) suite at the INFN-CNAF computing centre", ISGC 2019, Taipei

The Goal: create a global, distributed anomaly detection system, based on ES/Beats and DL data analysis, to monitor site activities

- Prototyping in the ATLAS Italian Tier-2 infrastructure
 - All Tier-2 sites currently sending **auditbeat** data to a global collector in Roma since april 2019 (average ~ 500 Hz of collected data, 10 GB/day)
 - start/stop processes
 - Open sockets
 - User events such as nis calls
 - Central (expandable) infrastructure currently running on an ES cluster with 7 data nodes + dual queue buffer and a totale space ~30TB, sufficient to keep several months of data online
 - Shared infrastructure, currently also collecting the ATLAS global node description data



- Data analysis and goals
 - Anomaly detection based on autoencoders being prepared, to detect problems or intrusions in the sites
 - Generally useful for many different purposes
 - Transparent distributed firewall, intrusion detections, hardware failures, process misbehaviour or malicious attacks (useful also during security challenges), etc
 - DL training and analysis can also be performed via a dedicated GPGPU nVidia facility being deployed in Roma



The screenshot shows the ATLAS ML platform website. The header has a blue and green geometric logo on the left, the text 'ML platform' in the center, and a navigation bar on the right with links for 'About', 'Services', and 'Login'. The main content area features a large image of server racks. On the left side of this image, the text 'ATLAS ML PLATFORM' is displayed in large white letters. Below this, a paragraph describes the platform: 'Get your private ML environment in minutes or use public JupyterLab to learn how others do things. The latest versions of all most popular ML software preinstalled. If you need something exotic just ask. Login using your own institution sign on.' To the right of this text is a white wireframe cube icon. Below the paragraph, it says 'Soon to come: TensorFlow as a service. Environment with OpenAI reinforcement learning framework.' At the bottom of the main content area, it states 'Hosted on ATLAS MWT2 hardware at University of Chicago. Relaying on SLATE, GLOBUS platforms.' At the very bottom of the screenshot, a green box contains the URL 'https://www.atlas-ml.org/'.

ML platform

ML Front About Services Login

ATLAS ML PLATFORM

Get your private ML environment in minutes or use public JupyterLab to learn how others do things. The latest versions of all most popular ML software preinstalled. If you need something exotic just ask. Login using your own institution sign on.

Soon to come: TensorFlow as a service. Environment with OpenAI reinforcement learning framework.

Hosted on ATLAS MWT2 hardware at University of Chicago. Relaying on SLATE, GLOBUS platforms.

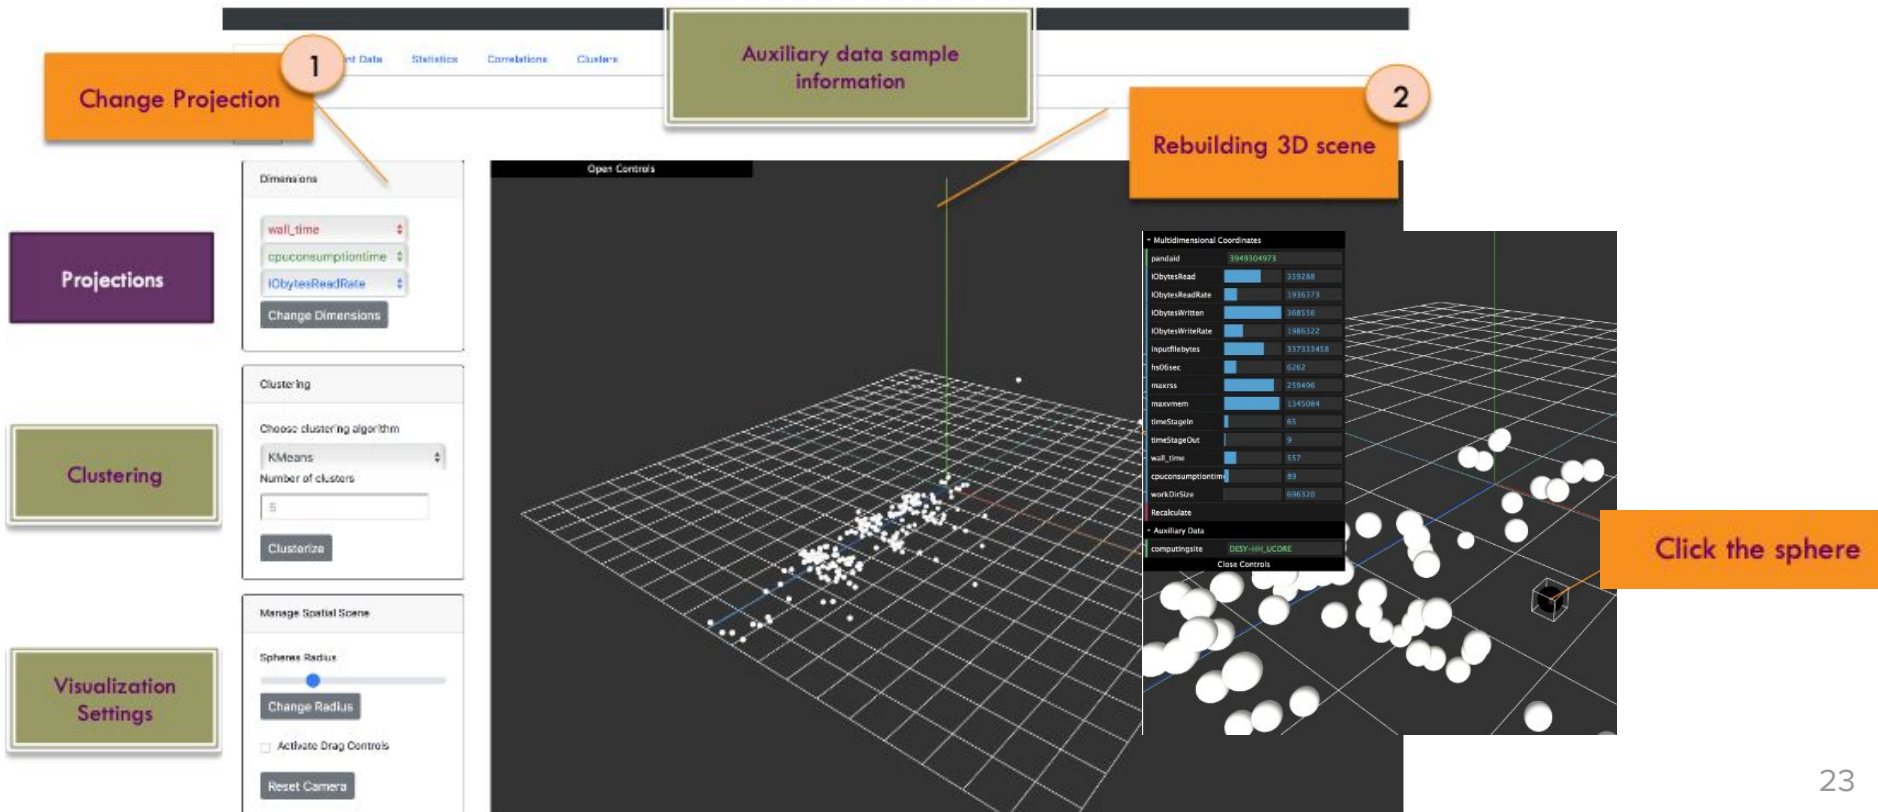
<https://www.atlas-ml.org/>

- Nodejs, expressjs site running on k8s. Has full control of k8s deployments.
- Uses Globus authentication.
- Can limit resources per instance.
- Extensible
 - Currently offers: private and shared JupyterLab instances and Spark Job submissions.
 - To come: TFAAS
- Easy to deploy elsewhere.
 - ML workshops
 - Teaching computational physics courses

Could improve:

- shared storage
- federate backend resources
- more options - CVMFS

Prototype of InVEx for ATLAS Computing <http://vap-dev.tpu.ru/>



- Anomaly detection in time series
 - Data quality, Network issues, Site performance
 - Despite importance, not many off-the-shelf tools
 - used or tried: simple hwm/lwm limits, ARIMA, plato detection, Bayesian simultaneous change point detection, ANN/BDT in time bins
 - Lack of well annotated data
- Classification
 - Data popularity prediction
 - Error classification:
 - Jobs - almost free style text - NLP?
 - FTS, Rucio, Frontier
- Need experiment-agnostic event annotation tool
 - Currently we only have tickets as a history of things that happened.
 - Not classified in any way that can be used to train any model.

Quick recap



- Operational Intelligence [website](#)
- Github repository: <https://github.com/operationalintelligence>
- E-group for communication: operational-intelligence@cern.ch
- JLAB session: <https://indico.cern.ch/event/759388/sessions/295063/#20190321>
- Google doc [draft](#)
- Regular meetings (2x/month): **3-4 pm on Mondays**
 - Indico category: <https://indico.cern.ch/category/11205/>



Today

Tomorrow

