



THE INTERWARE

Monitoring with and for DIRAC

2nd Virtual DIRAC Users' Workshop - 9 May 2022

RUBEN POZZI

OUTLINE



- Current state of DIRAC Monitoring
 - Accounting
 - Monitoring
- Plans for new release
 - What's new, what will stay, what will change
- Elasticsearch and Opensearch
 - Benefits, why Opensearch?
- Dashboards
 - Samples of dashboards that will be used in v8

CURRENT STATE OF DIRAC MONITORING



- The monitoring of DIRAC is currently based on two backends:
 - *Accounting*: based on MySQL, and currently the default backend
 - *Monitoring*: based on Elasticsearch (ES)
- Can be interchangeable for some monitoring types, but they are designed for different purposes
- Accounting more meant for historic data, Monitoring for “real-time”

CURRENT STATE: ACCOUNTING SYSTEM



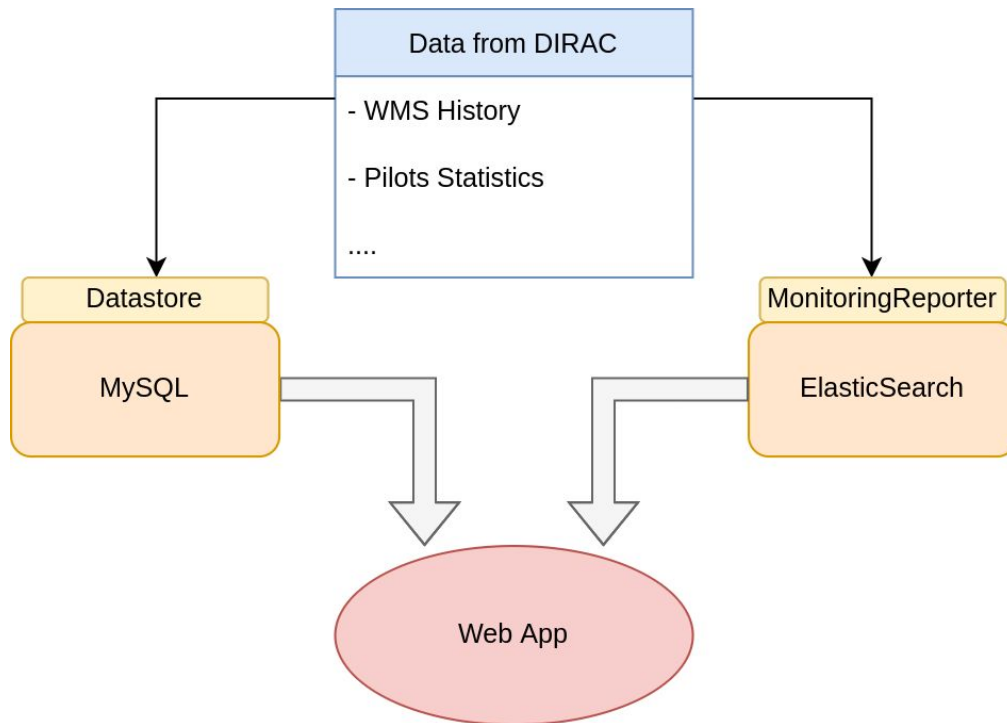
- Records are stored in the AccountingDB (based on MySQL)
- Mainly used for historical data
- Can be accessed through the DIRAC web framework in the Accounting web application
- Accounting types:
 - *Pilot*: pilot jobs historical reports
 - *Pilot Submission*: number and efficiency
 - *Job*: jobs historical reports
 - *Data Operation*: data transfer, replication, removal etc...
 - *WMS History*: monitoring of the DIRAC WorkloadManagementSystem
 - *Storage Occupancy*

CURRENT STATE: MONITORING SYSTEM



- Records are stored in the MonitoringDB (Elasticsearch-based)
- Better suited for real-time monitoring (time span ~24h)
- Can also be accessed through the DIRAC web framework in the Accounting web application
- Monitoring types:
 - *WMS History*: monitoring of the DIRAC WorkloadManagementSystem
 - *ComponentMonitoring*: deprecated (replaced by Agent- and Service Monitoring in v8)
 - *RMS*: monitoring of the requests executed by RequestManagementSystem

CURRENT STATE OF DIRAC MONITORING



PLANS FOR V8 RELEASE AND LATER




The plans for the new release, with regards to monitoring are the following

- Increase use of non-MySQL solutions, specifically ES-based.
- Dropping support for ES versions earlier than 6, and migrating to Opensearch
- Changes in monitoring (ES) types
 - Adding *PilotSubmissions* and *DataOperation* in ES
 - New *PilotsHistory* type (snapshots of *PilotAgentsDB* every 15m) in ES
 - Replacing *ComponentMonitoring* with *Agent-* and *ServiceMonitoring*
- New monitoring flag system
- MySQL types will remain untouched
- What happens to the DIRAC plotters? → no plans to expand further (new types such as *PilotsHistory* or *Agent* and *ServiceMonitoring* won't be supported on it, only on Kibana)

BENEFITS WITH ELASTICSEARCH



- Need for more scalability and less maintenance. Why not use Elasticsearch itself?
- ES offers an in-built monitoring service called Kibana, which can be used directly with a corresponding version of ES installed.
- It offers more scalability, interactivity, more insight on the data and no maintenance, compared to the WebApp plotter.
- Grafana also very useful (currently used by LHCbDIRAC in production)



ELASTICSEARCH VS OPENSEARCH



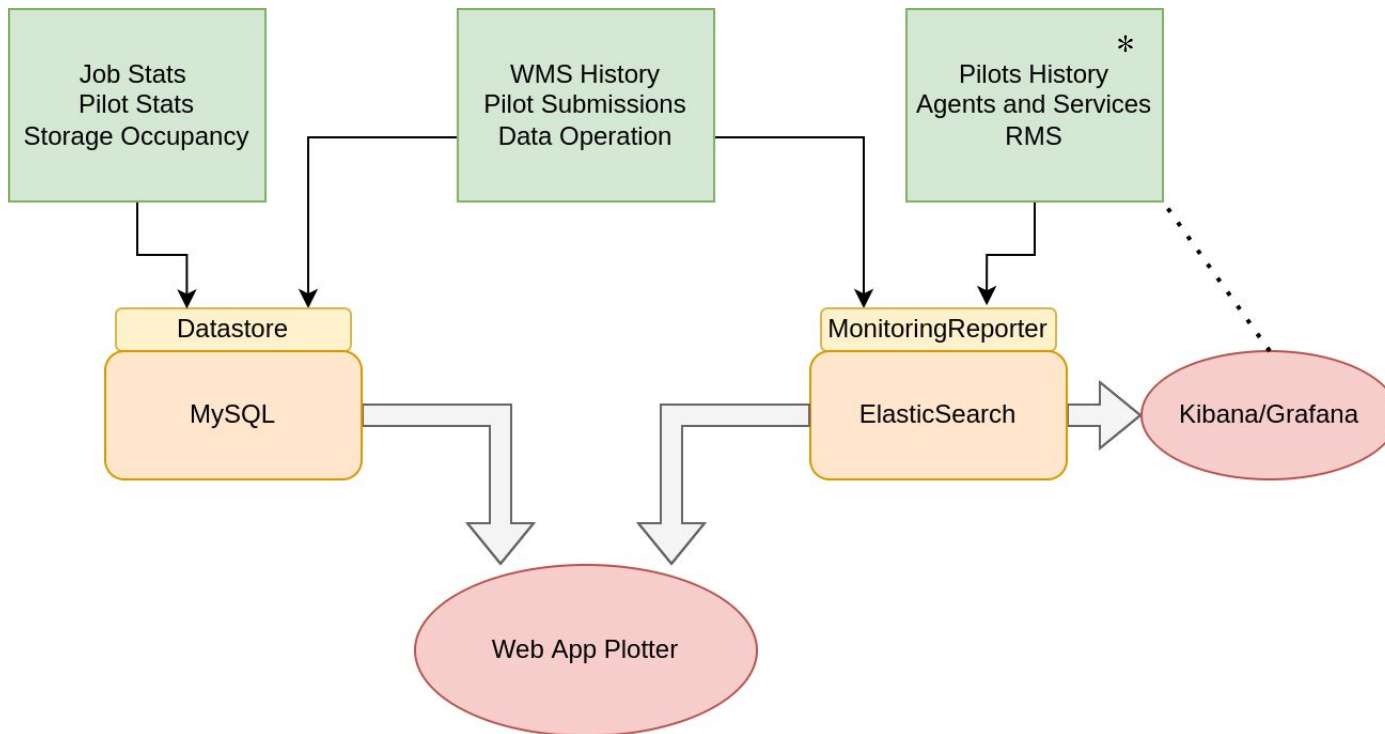
- NEW: Elasticsearch has [recently changed their a licensing terms](#), so CERN has consequently decided to migrate to Opensearch as a legitimate (and free-licensed!) substitute.
- Opensearch is an Elasticsearch fork developed and maintained by Amazon Web Services, and it is supposed to be kept free and open-source (at least in the near future).
- DIRAC already is using instances of Opendistro (an earlier version fork of Opensearch), for example one for certification ([es-dirac-certification](#)), or one that is being used by LHCbDIRAC for production.
- The plan for DIRAC v8+ is to drop all support for Elasticsearch and just switch to Opensearch an all instances

SO WHAT WILL STAY ON MYSQL?



- On MySQL, we will keep all of the types that are present now
- The following types will be present **only on MYSQL** (for the foreseeable future):
 - Pilot Stats
 - Job Stats
 - Storage Occupancy
- No plans on moving these types to ES for now
- How will the overview of types look in v8?

SUMMARY OF MONITORING TYPES IN V8



* These types are not supported in the WebApp plotter, only visible in Kibana/Grafana

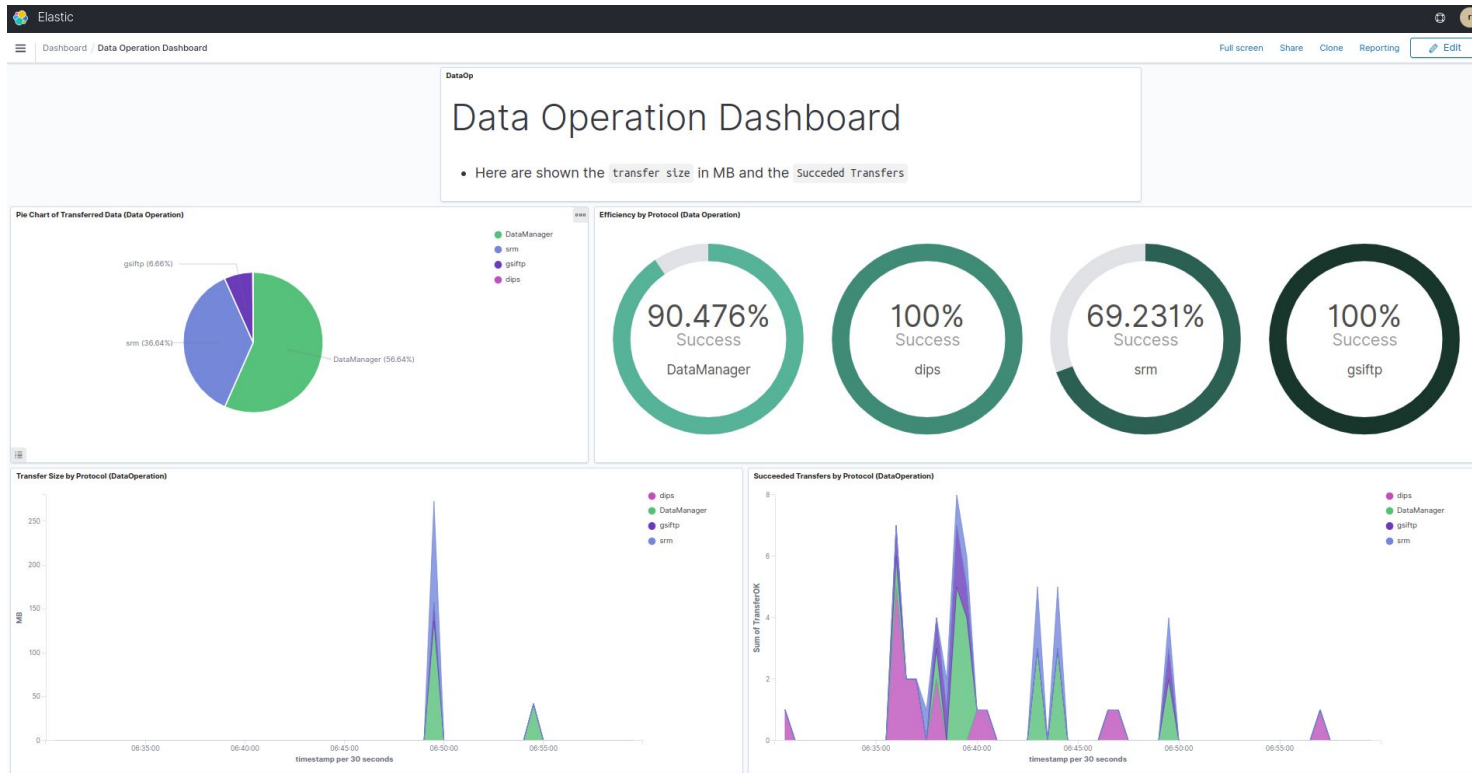
DASHBOARDS FOR DIRAC



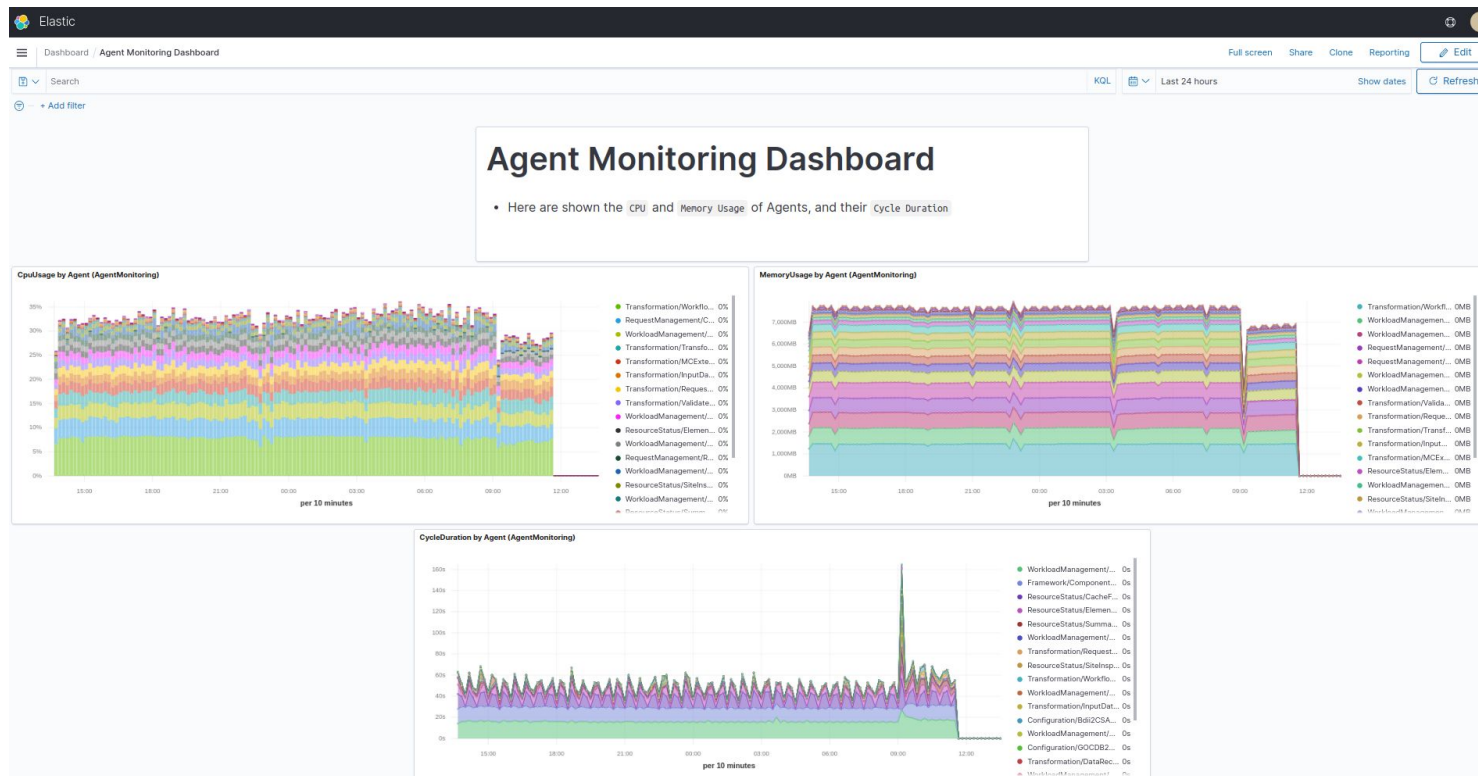
- Dashboards ready for import will be present in DIRAC under DIRAC/Dashboards
- Can be imported to a Kibana instance and used directly by adapting the index pattern to the wished one (automatic process)
- Can already be viewed and modified on es-dirac-certification.cern.ch and in [Grafana](#) (example with LHCbDIRAC)

A screenshot of a GitHub repository page for 'DIRACGrid / DIRAC'. The page shows a commit titled 'feat: Added ElasticSearch dashboards for import' by user 'rupozzi', made 5 days ago. Below the commit message, there is a list of folders representing the added dashboards: AgentMonitoring, DataOperation, PilotSubmissions, PilotsHistory, RMS, ServiceMonitoring, WMS, and diracLogs. Each folder entry includes a folder icon, the folder name, the commit message snippet, and the time '5 days ago'. The repository interface includes navigation tabs for Code, Issues (93), Pull requests (16), Discussions, Actions, Projects, Wiki, Security, and Insights. At the top right, there are buttons for 'Edit Pins', 'Watch' (16), 'Fork' (156), and 'Star' (97).

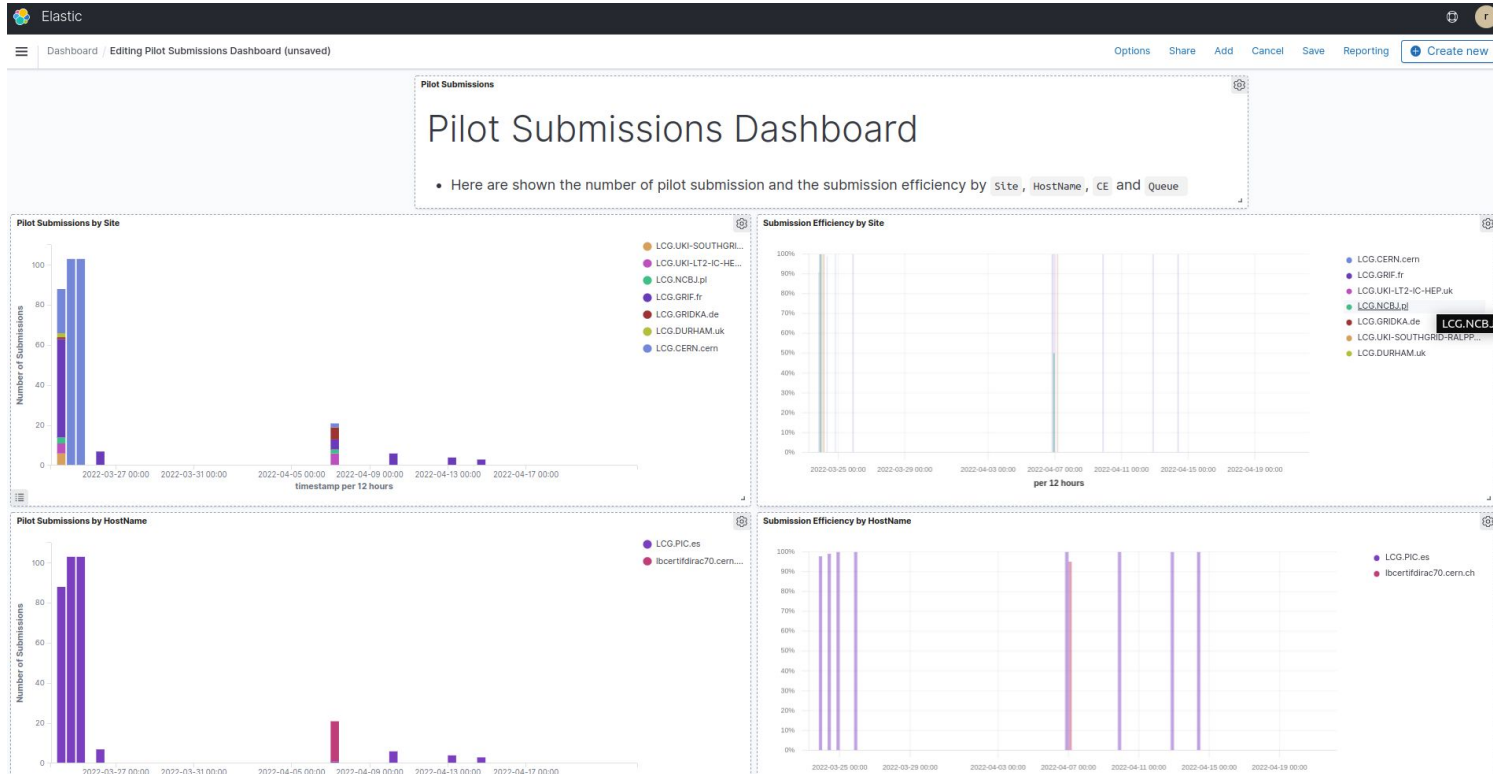
DASHBOARDS FOR DIRAC



DASHBOARDS FOR DIRAC



DASHBOARDS FOR DIRAC



Thank you for your attention!

Questions?

