The ATLAS Production System Predictive Analytics service: an approach for intelligent task analysis

Mikhail TITOV, Mikhail BORODIN, Dmitry GOLUBKOV, Alexei KLIMENTOV
on behalf of the ATLAS Collaboration

CERN

ATLAS EXPERIMENT
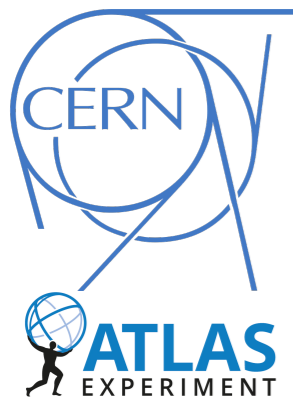
National Research Centre
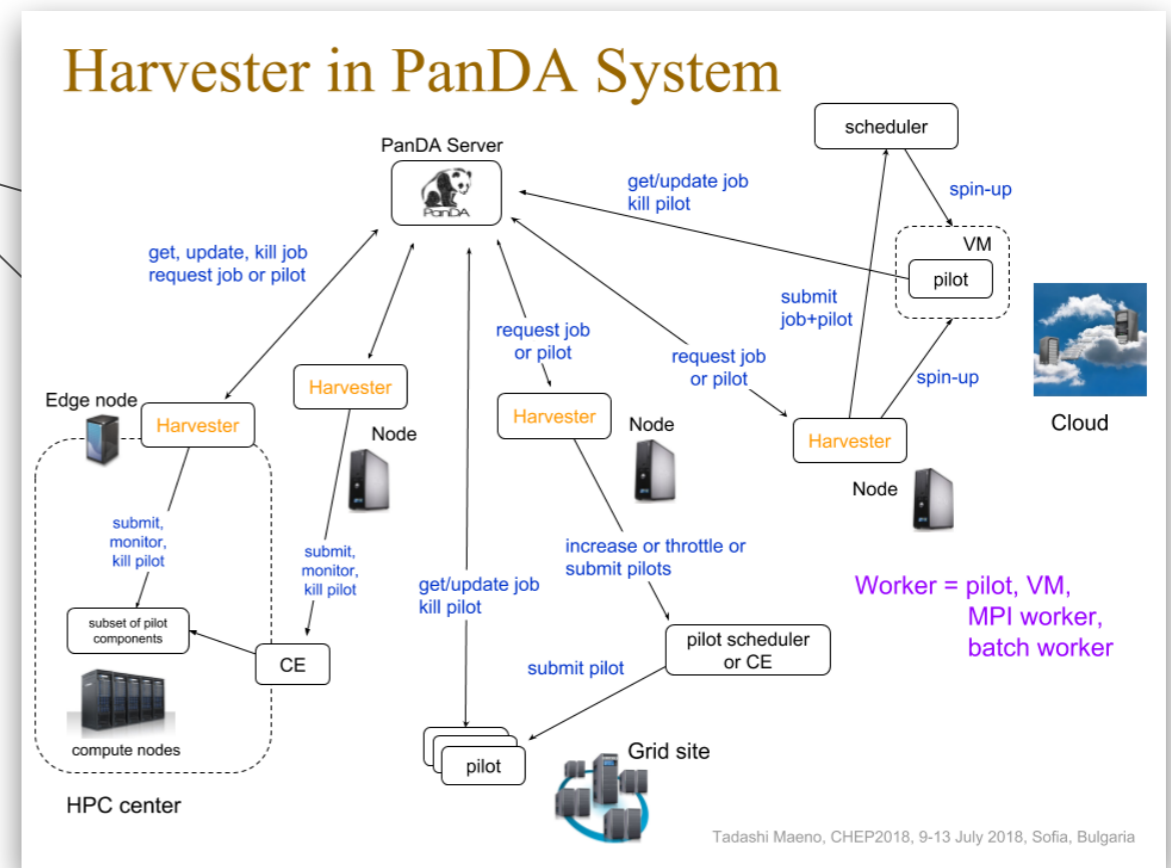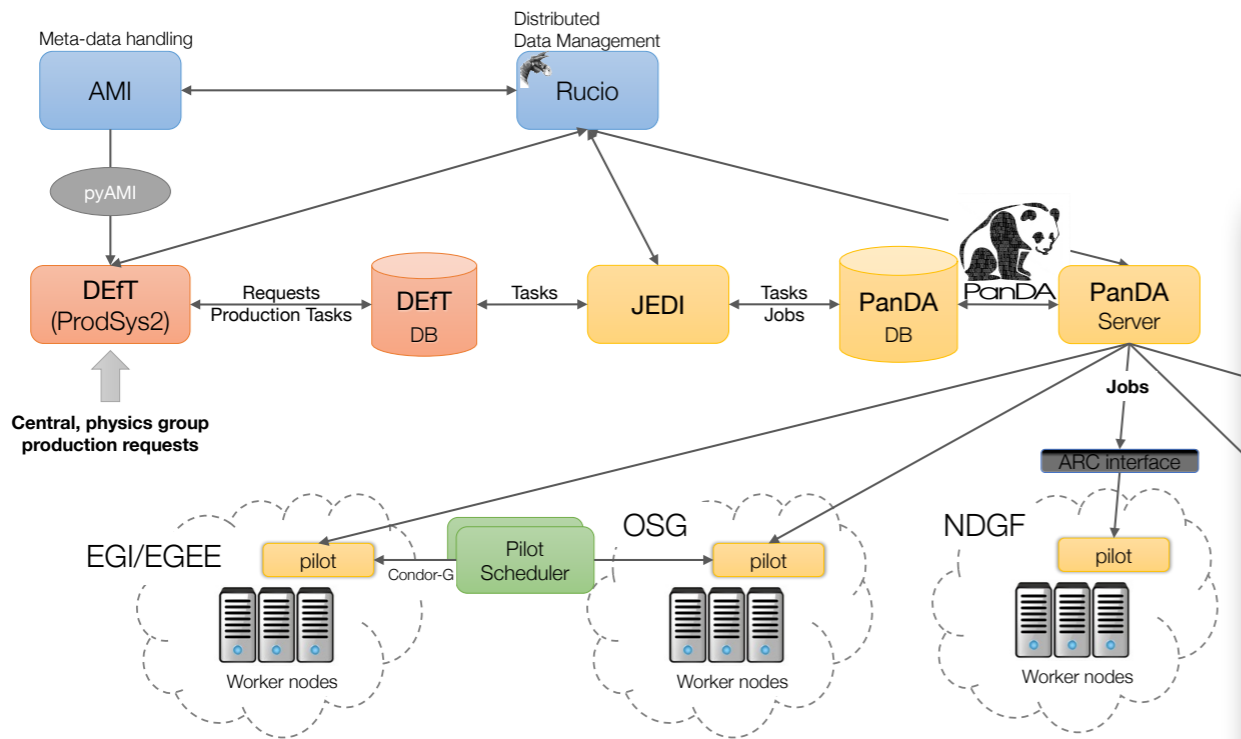**Kurchatov Institute**

THE UNIVERSITY OF IOWA

ИФВЭ
ІHEP

BROOKHAVEN
NATIONAL LABORATORY

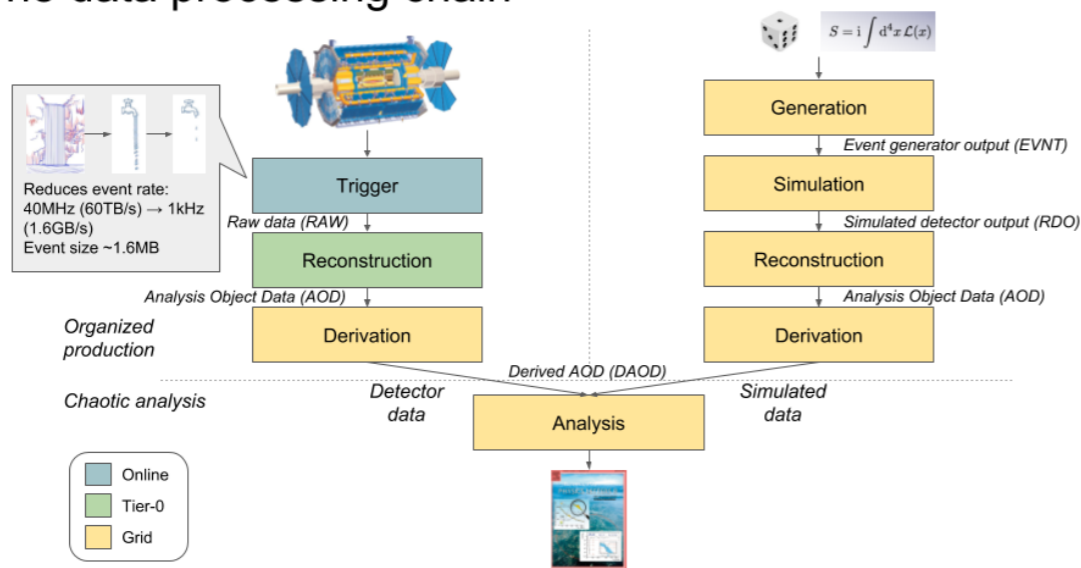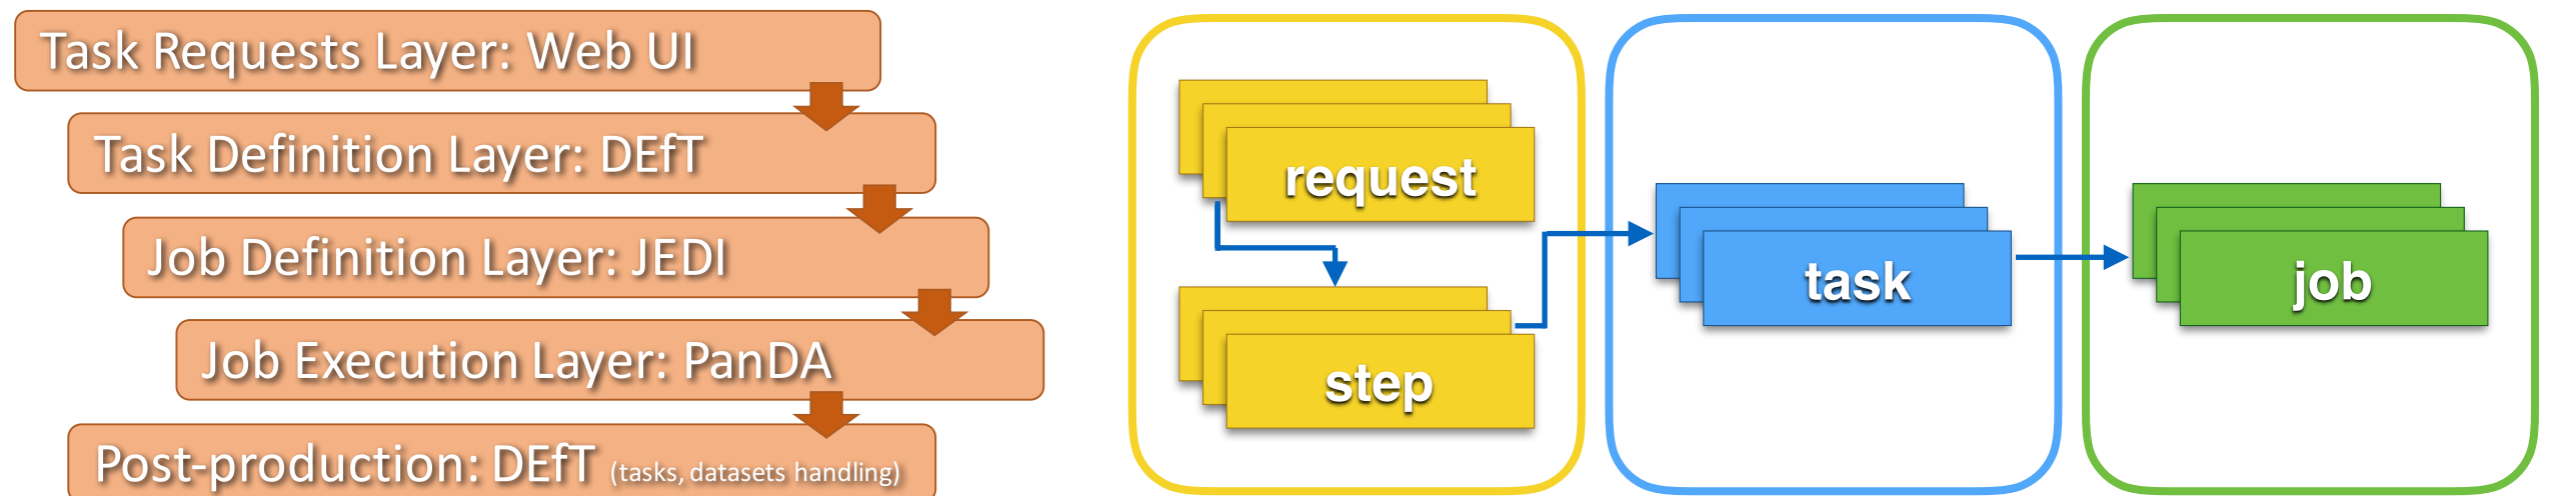# Introduction

ATLAS Workflow Management

ATLAS Production System [ProdSys2]

Database Engine for Tasks (DEfT)

- Formulate the tasks, chains of tasks and task groups (production request)
  - *Task represents a logical grouping of computing jobs, that are responsible of the execution of program/transformation on input files and generate output files*
- Complete with all necessary parameters

Job Execution and Definition Interface (JEDI)

- Task-level workload management (i.e., brokerage and execution)
- Dynamic job definition and execution (optimization of the resources usage)

Task Requests Layer: Web UI

Task Definition Layer: DEfT

Job Definition Layer: JEDI

Job Execution Layer: PanDA

Post-production: DEfT (tasks, datasets handling)

request

step

task

job

Problem statement

Analytical service focused on tasks

- Analysis of task processing
  - Example: selection and regulation of key task features that affect its processing the most

- Modeling of processed data lifecycles for deep task analysis
  - Example: generate guidelines for particular stage of data processing

- Forecasting processes with focus on data and tasks states as well as on the management system itself
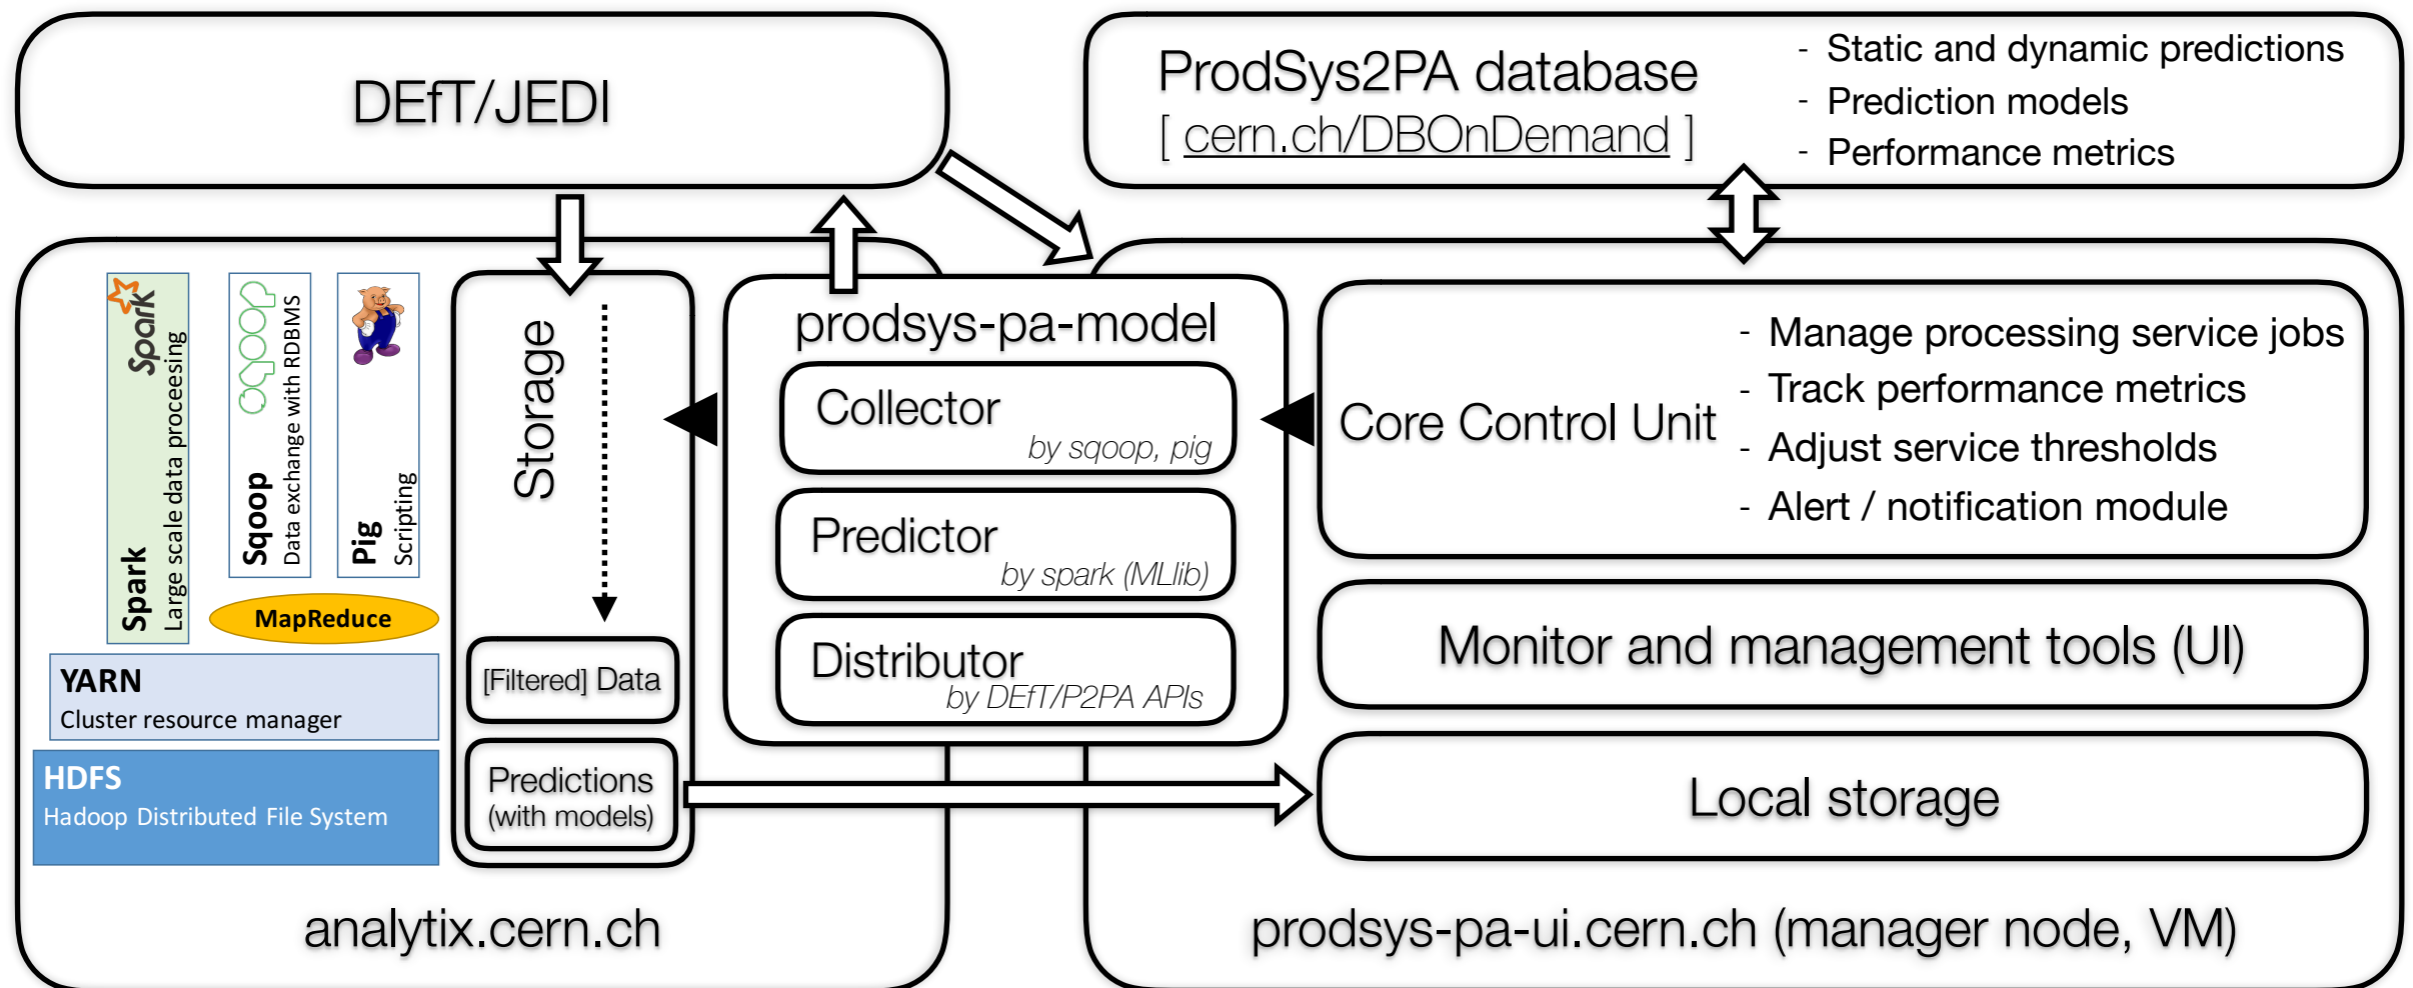 - Example: detect the source of any potential malfunction

ProdSys2 Predictive Analytics (PA) service
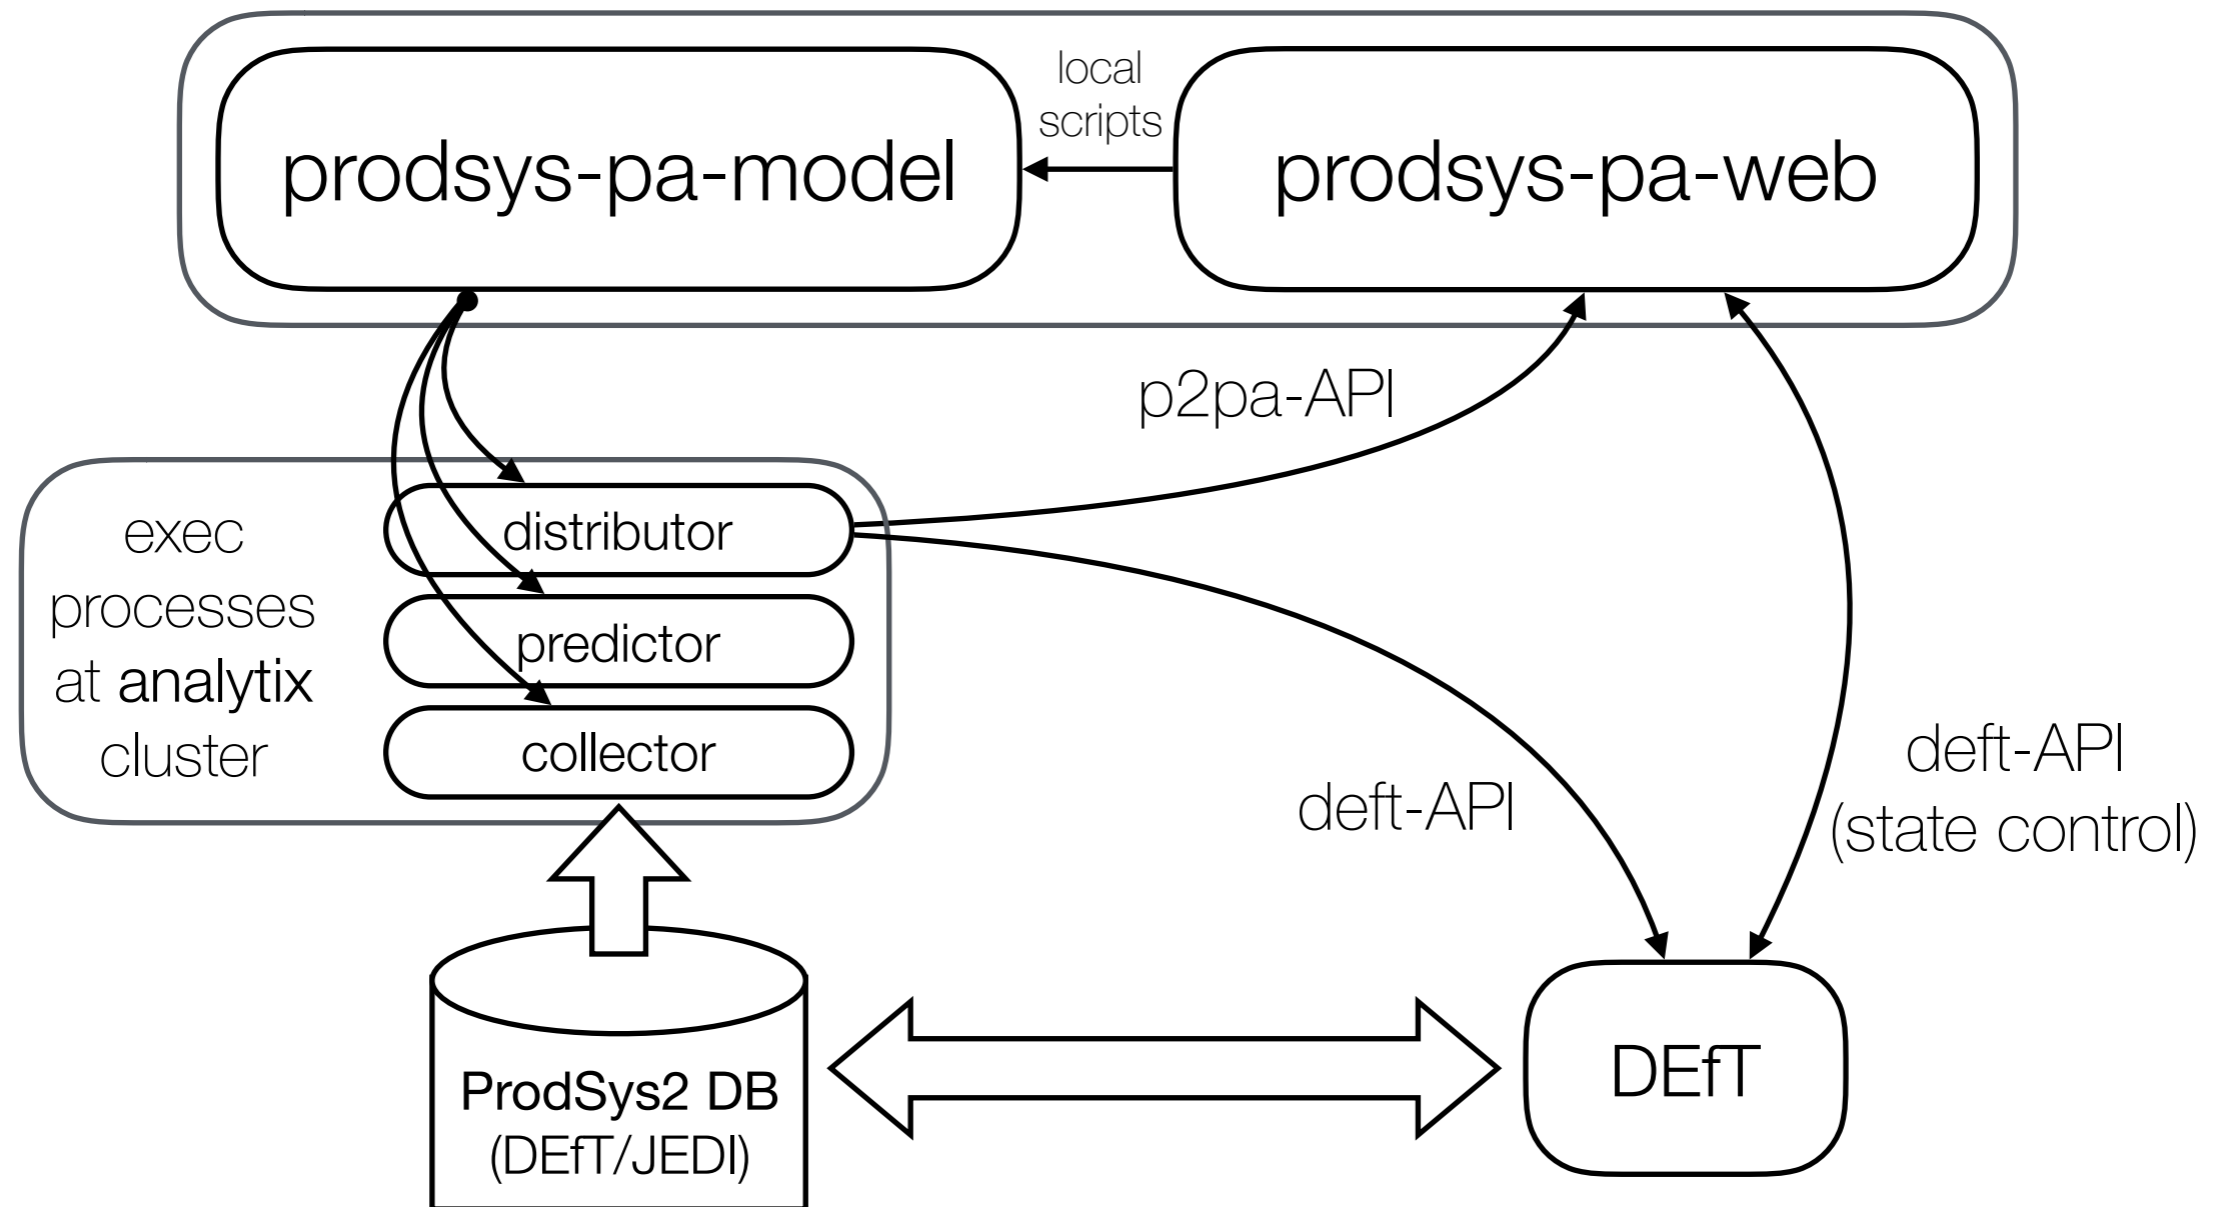
ProdSys2 PA service

Key components:

- **Predictive model handling** (*prodsys-pa-model*) - an independent package that is adjusted and integrated into P2PA

    - Data Collection
        - Source - data/mc-tasks from DEfT/JEDI

    - Data Analysis
        - Predictive model creation and usage for the process of TTC (Time-To-Complete) predictions generation

- **Web application** (*prodsys-pa-web*) - a central operation hub

    - Control (e.g., selection parameters for training and input data collections, method / technology and set of features for prediction process)

    - Monitor (e.g., exec processes that are historic data of runs, evaluation of estimated durations of task executions)

ProdSys2 PA architecture

DEfT/JEDI

ProdSys2PA database
[cern.ch/DBOnDemand]

- Static and dynamic predictions
- Prediction models
- Performance metrics

Spark Large scale data proceesing
Sqoop Data exchange with RDBMS
Pig Scripting

MapReduce

YARN Cluster resource manager

HDFS Hadoop Distributed File System

Storage

[Filtered] Data

Predictions (with models)

prodsys-pa-model

Collector *by sqoop, pig*

Predictor *by spark (MLlib)*

Distributor *by DEfT/P2PA APIs*

Core Control Unit

- Manage processing service jobs
- Track performance metrics
- Adjust service thresholds
- Alert / notification module

Monitor and management tools (UI)

Local storage

analytix.cern.ch

prodsys-pa-ui.cern.ch (manager node, VM)

ProdSys2 PA packages

prodsys-pa-model

local scripts

prodsys-pa-web

p2pa-API

exec processes at **analytix** cluster

distributor

predictor

collector

ProdSys2 DB (DEfT/JEDI)

deft-API

deft-API (state control)

DEfT

P2PA | Model handling | Analysis processes

Threshold definition

- Calculates the upper limit of the duration of tasks execution process in such a way that 95% of all tasks of the corresponding type and for the defined time period (the last 180 days) are executed not longer than the calculated value.

  - Tasks are grouped by <project.productionStep.workingGroup>

Cold-prediction generation

- Estimates task duration during the task formation process (uses descriptive data and created earlier a predictive model).

  - Apache Spark.MLlib - Random Forests regression method

P2PA | Model handling | Task TTC estimation

Legend:
- Real task durations
- Durations based on thresholds
- Durations based on cold-prediction

RMSE = 4.4
Absolute error metrics (days)
        mean = 0.47, std = 4.39
        min, max = -20.4, 38.5

Tasks are created from 12 to 26 August, 2018

Y-axis: Duration, days
X-axis: Tasks (ordered)
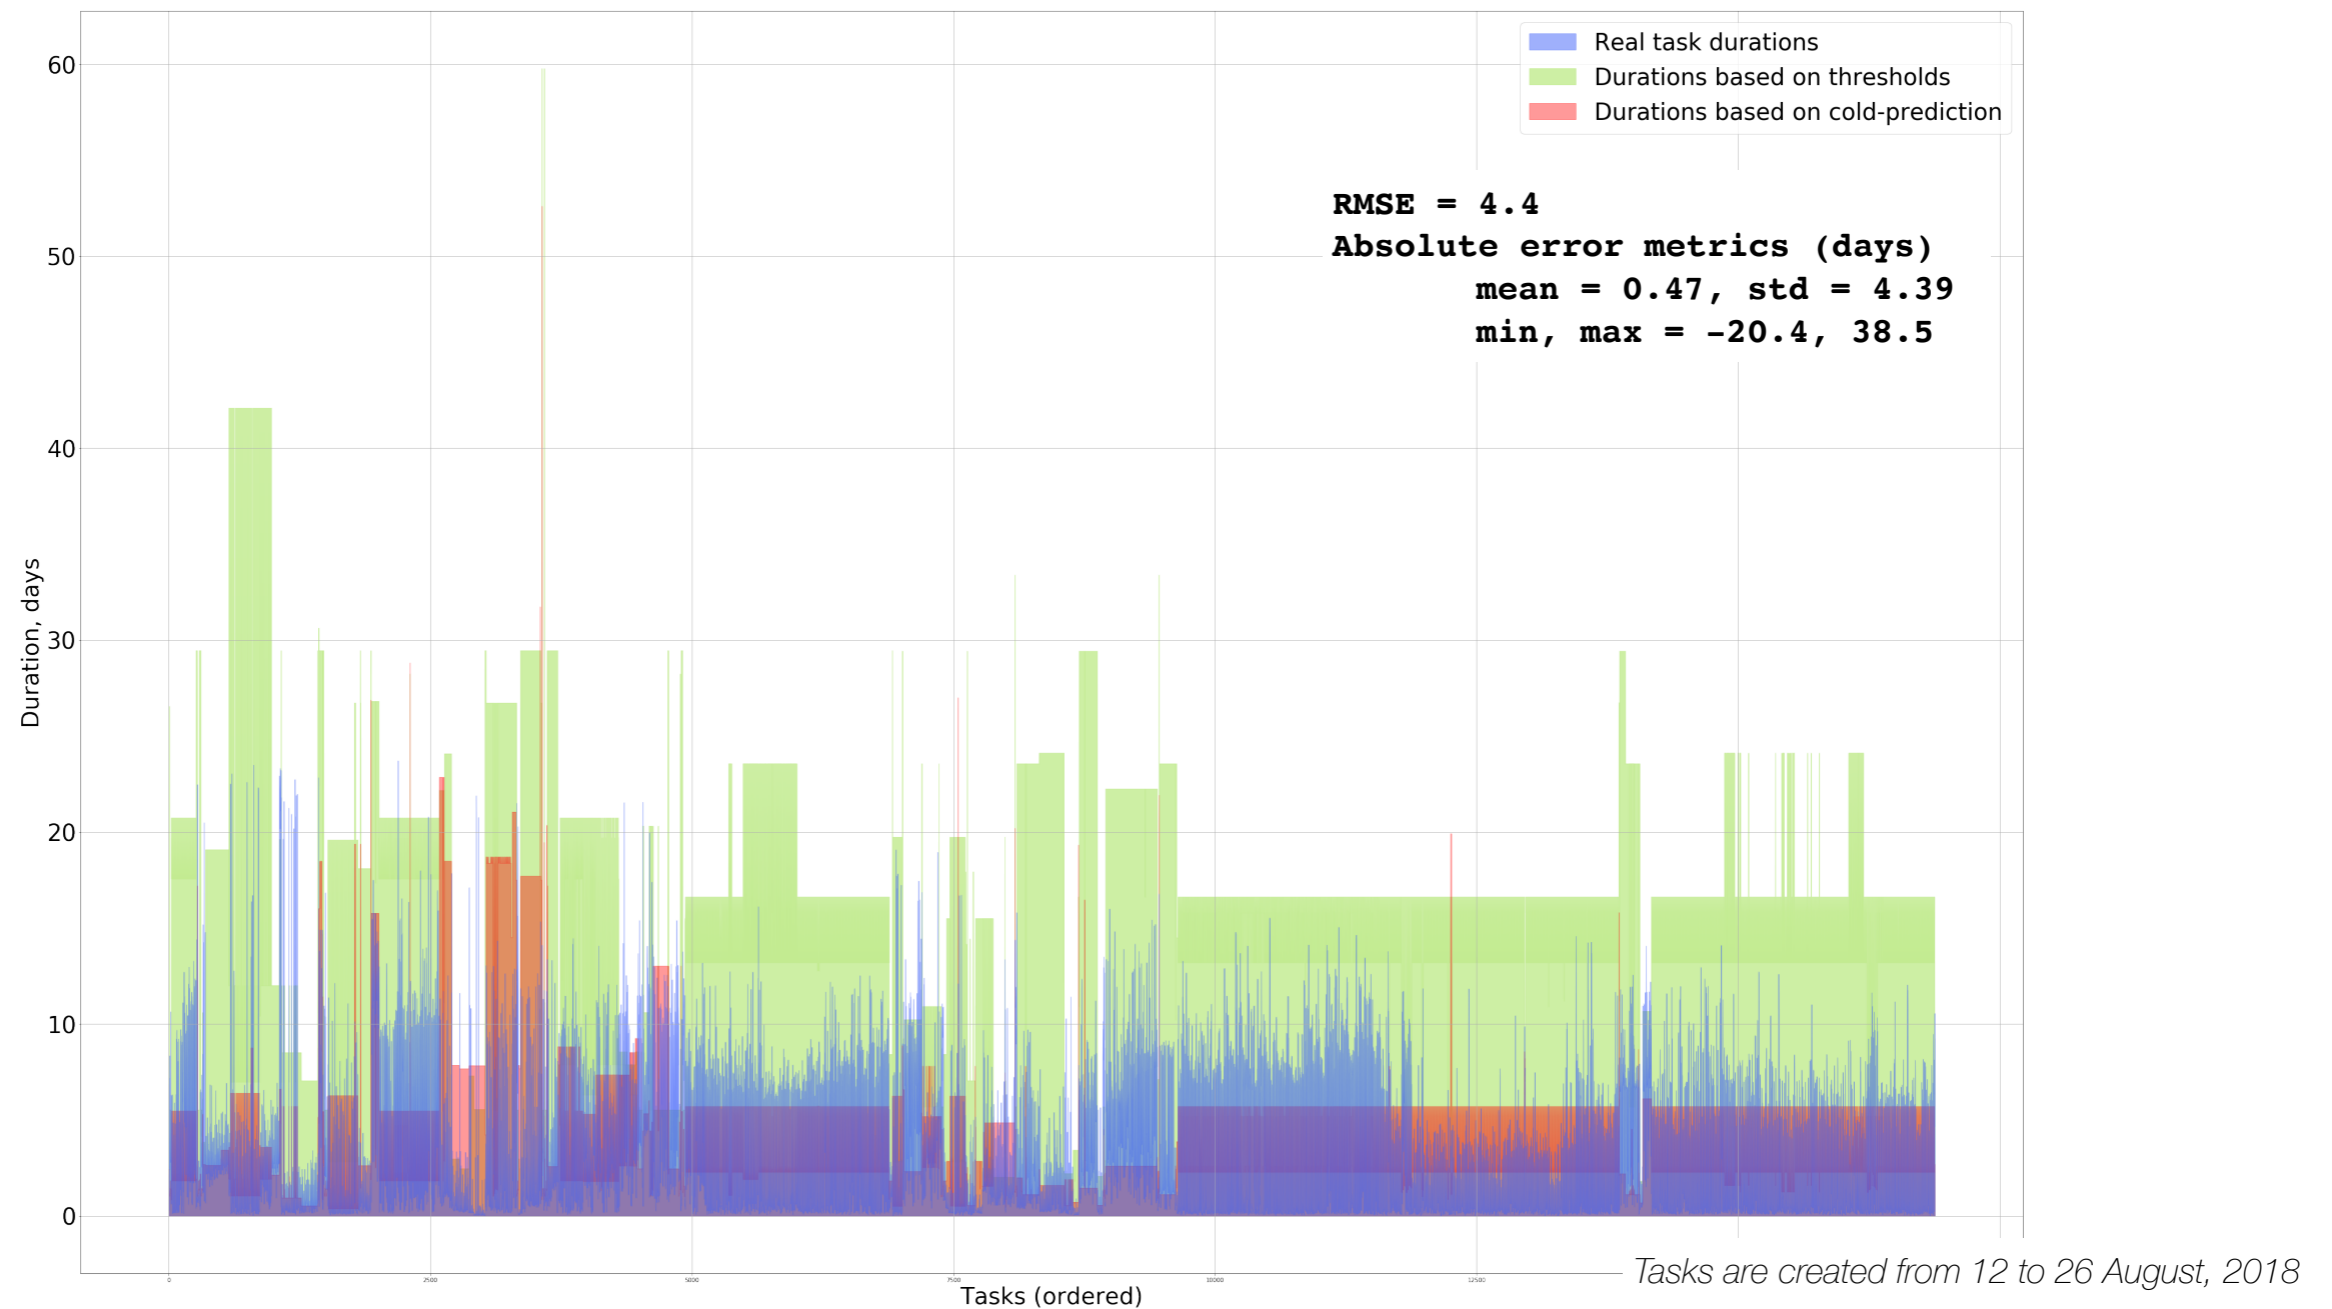
P2PA | Web UI | Infrastructure

Web application (web UI) setup

‣ https://prodsys-pa-ui.cern.ch [accessible inside the CERN network]

‣ VM by CERN OpenStack IaaS (CC7 – x86_64)

‣ Database On Demand (MySQL)

‣ Web server
  ‣ nginx
  ‣ gunicorn

‣ Web framework:
  ‣ django
  ‣ django REST framework
  ‣ celery (using RabbitMQ)
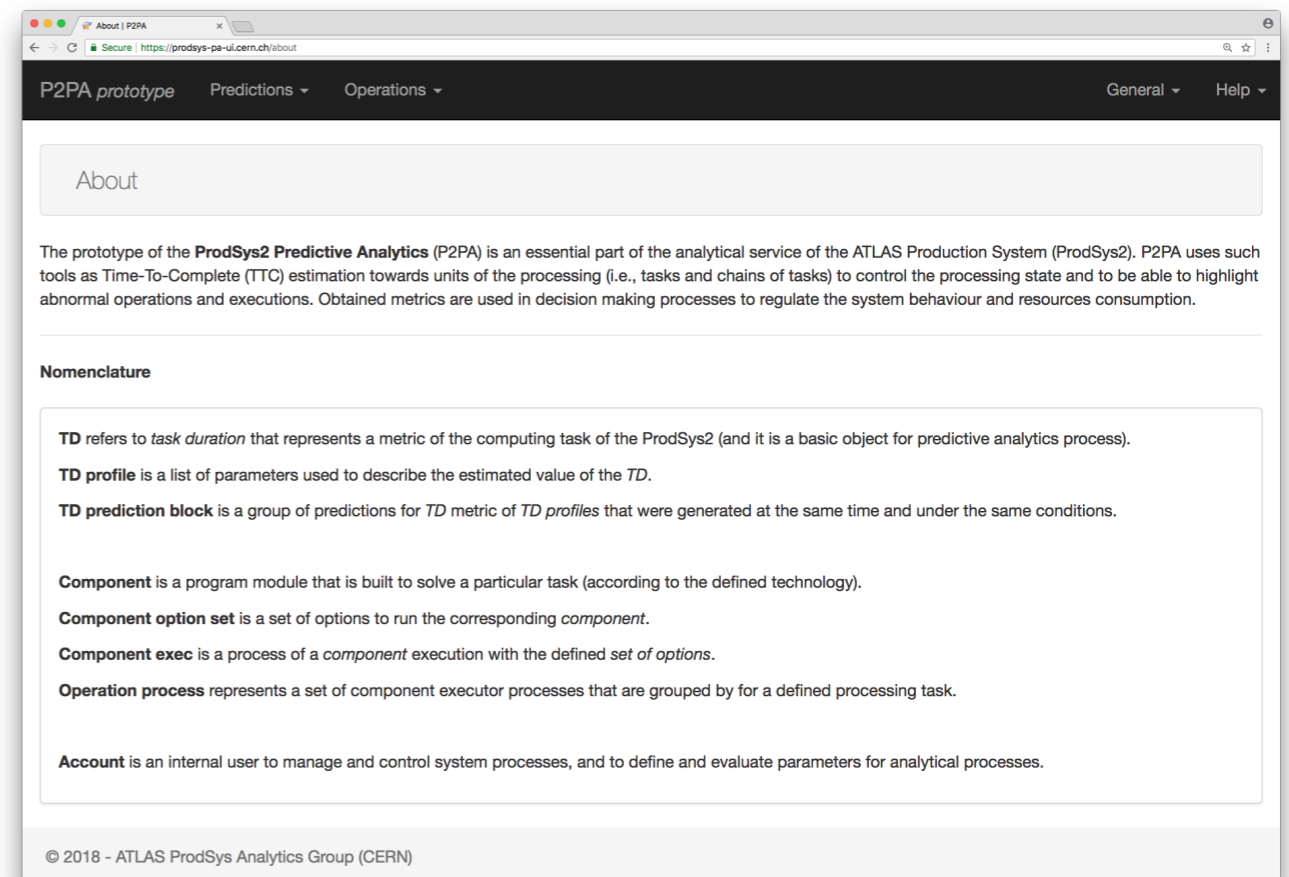
P2PA | Web UI | Operation processes

P2PA | Web UI | Predictions generation

Conclusion

Summary

- ProdSys2 Predictive Analytics service is designed to enhance workflow control at the ATLAS Production System and to be able to detect and highlight abnormal operations and executions.

- It is planned to use obtained metrics in decision making processes to regulate the system behaviour and resources consumption.

- The quality of obtained metrics (estimated values of controlled parameters) is constantly improving and new evaluation parameters and metrics will be introduced.

* If you have any questions/comments/suggestions regarding P2PA, please email: atlas-adc-prodsys2-analytics@cern.ch

Acknowledgements