

Web frontend to the Production and Distributed Analysis objects states

The ATLAS experiment uses PanDA (Production and Data Analysis) Workload Management System for managing the workflow for all data processing on over 150 data centers and opportunistic resources such as commercial clouds, supercomputers, volunteer machines.

- Since 2005 an effective interface to the heterogeneous distributed computing infrastructure
- 300k simultaneous jobs
- 2M jobs a day
- ~1500 users
- Processes over an exabyte of data per year
- Applications beyond ATLAS: COMPASS, AMS, (+ number of evaluations)

Failures investigations

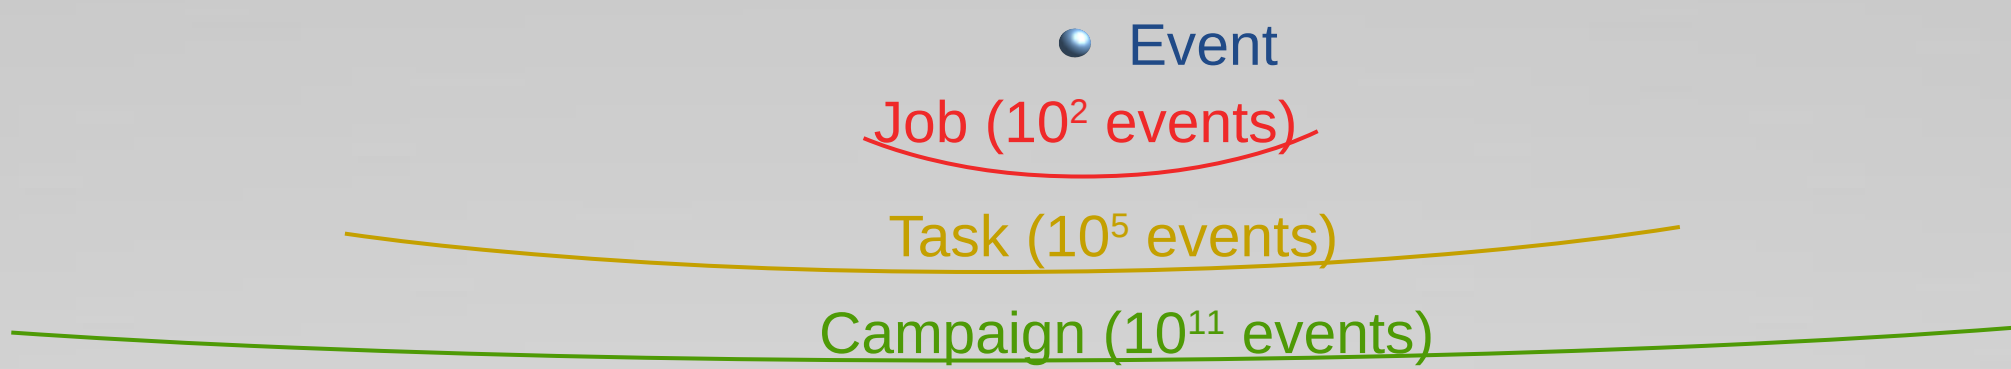
Logs acquisition and analysis together with errors message accessible from the database provides tracking of issues at different scales and views with BigPanDA monitoring:

- Job (stderr, Memory consumption, IO, IO rate)
- Jobs group (by software release, task, user, time, site, type of error, etc)
- PanDA server, Jedi errors and messages
- Payload and brokerage issues

Recommendation system analyses prior visits of bigpanda.cern.ch by a user and suggests views and objects which in the focus of his/her interests

Reports of different scale and levels of abstraction in real time mode

Production/Analysis entities: Jobs, tasks, request (as tasks filter) and campaign.



Scales covered by BigPanDA monitoring

Payload and large scope issues detection: Tasks/Jobs breakdown by Physics groups, Clouds, Sites, analytic or production, memory/CPU requirements, etc. The ~15 dashboards totally with historical data and real time activity.

Dozens of different filters: by time periods, failure reasons, users, running sites, hardware requirements, processing physics, etc.



BigPanDA.cern.ch monitoring

S.Padolski<sup>1</sup>, T.Wenaus<sup>1</sup>, A.Klimentov<sup>1</sup>, T.Korchuganova<sup>2</sup>, A.Alekseev<sup>2</sup>  
on behalf of the ATLAS Collaboration

<sup>1</sup>Brookhaven National Laboratory, Upton, NY 11973-5000, USA  
<sup>2</sup>Tomsk Polytechnic University, Tomsk, 634050, Russia

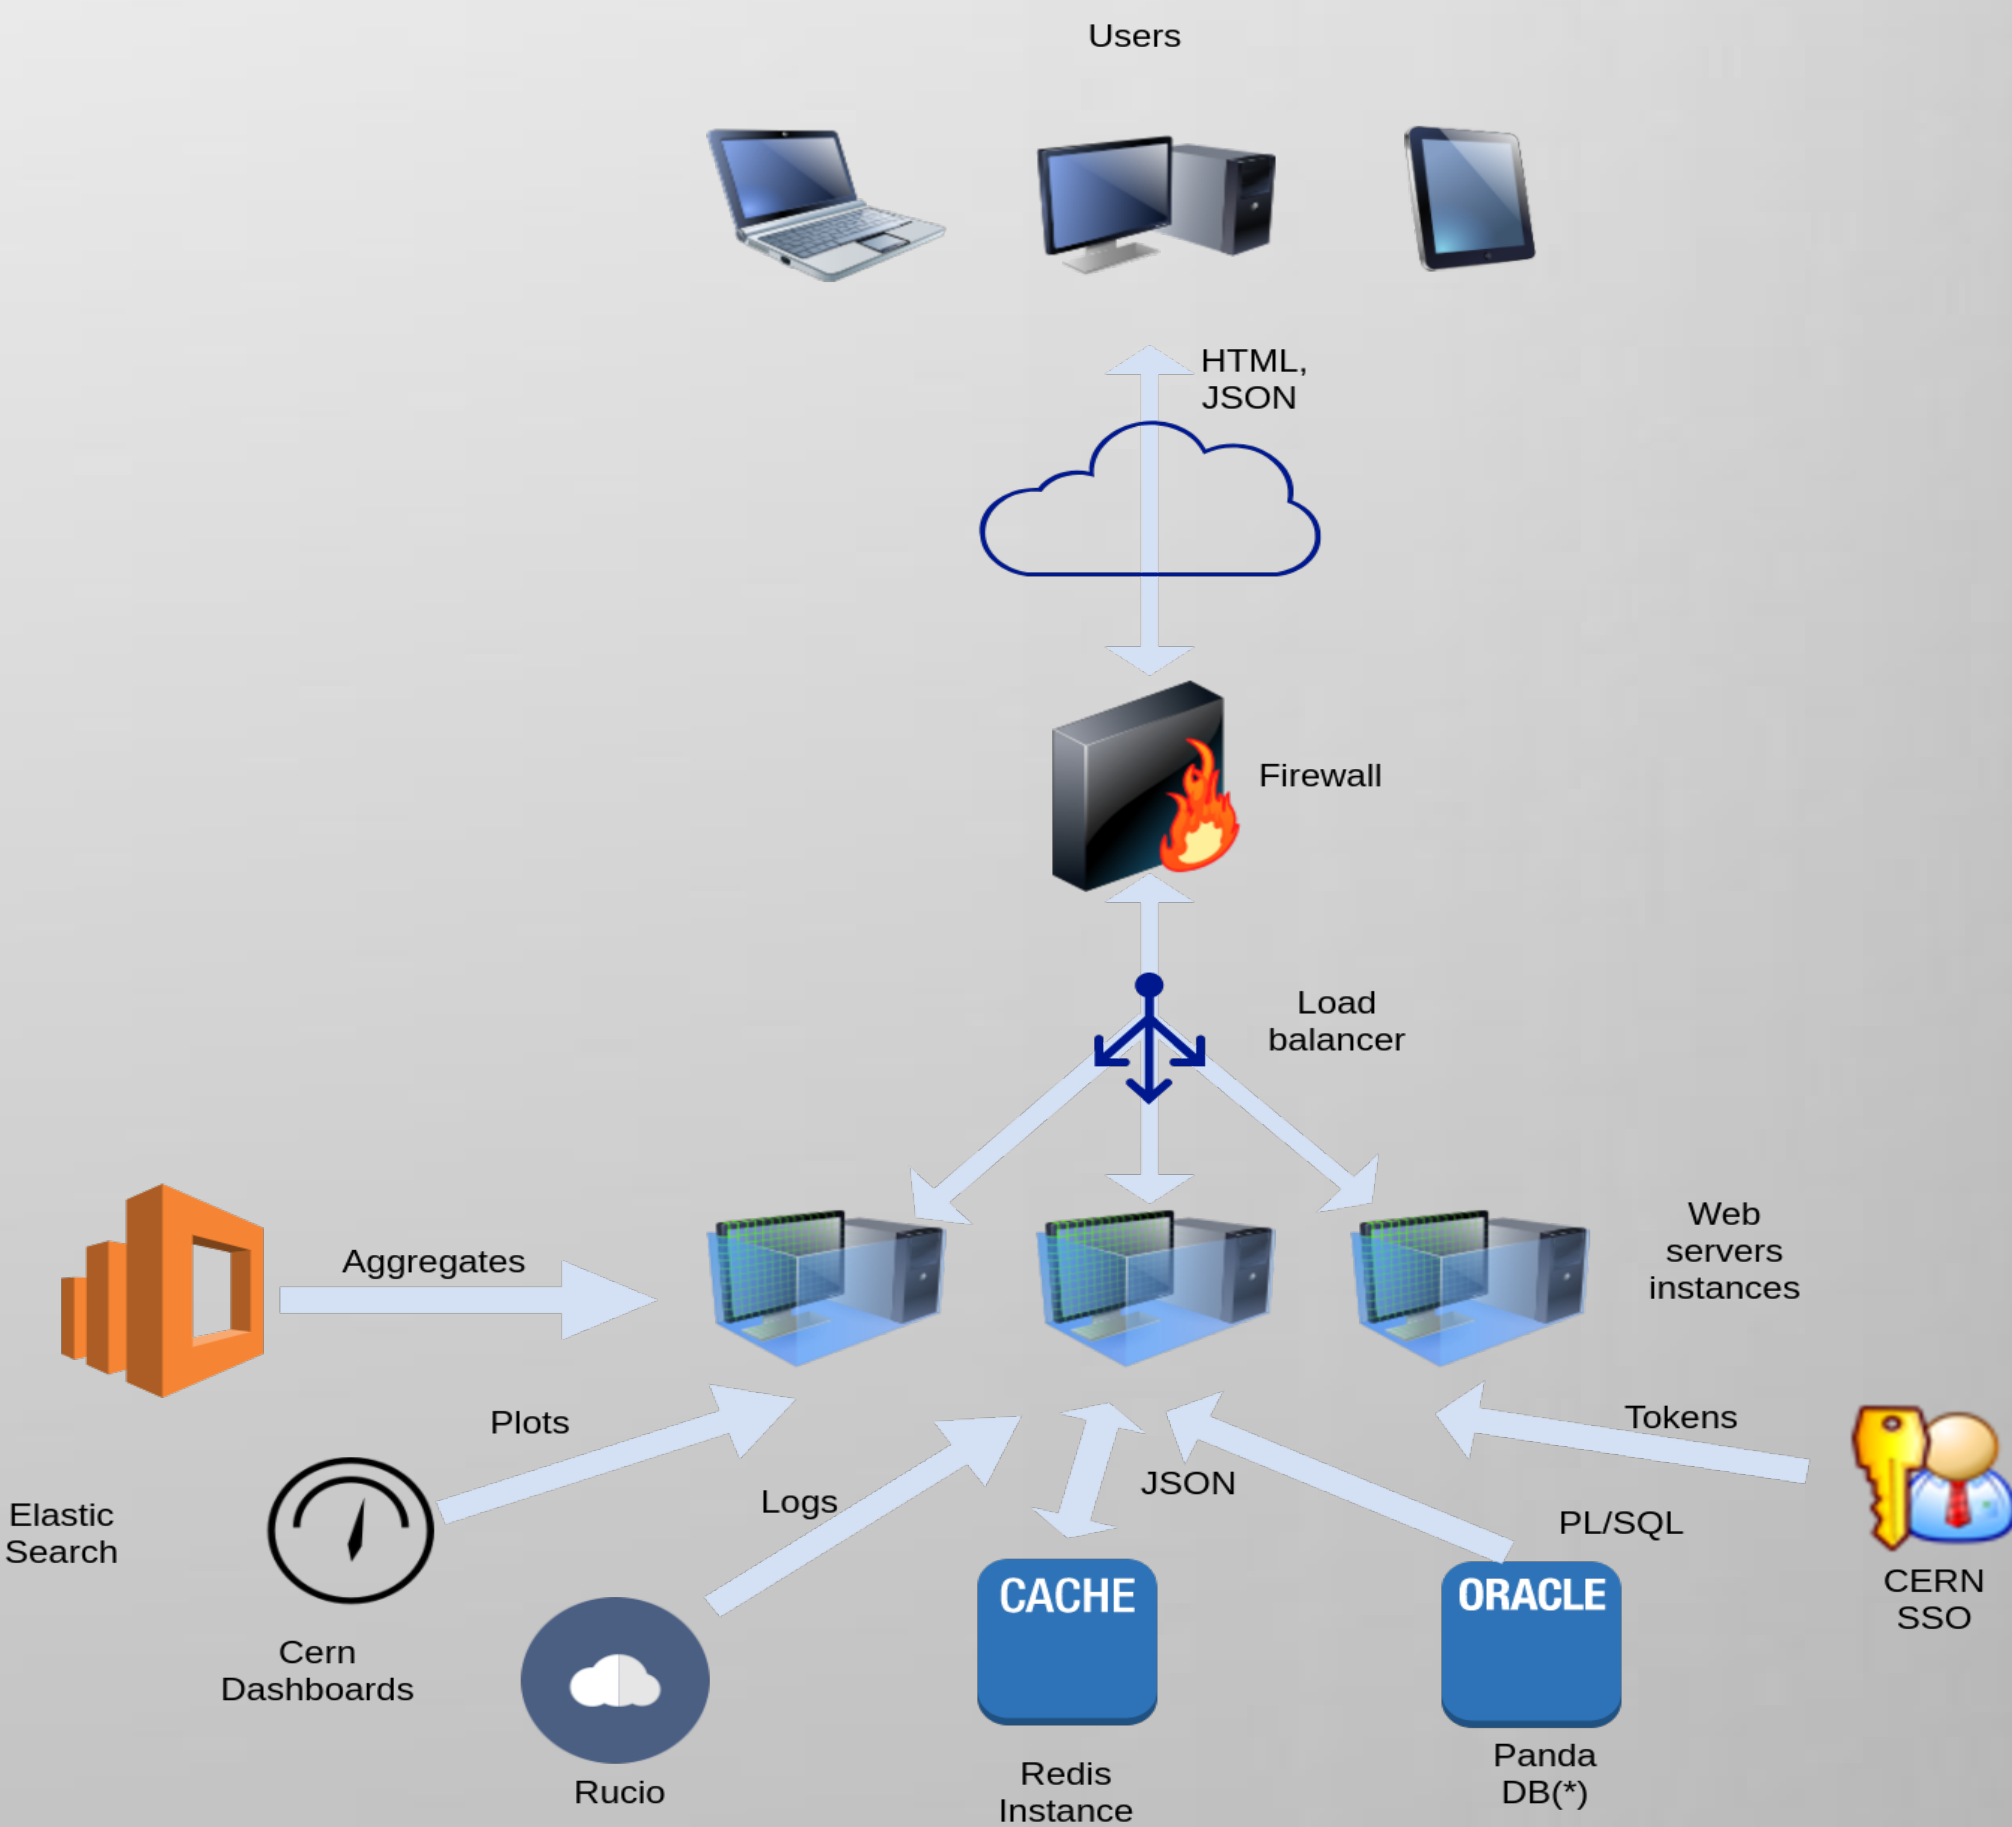


Primary source of information for ATLAS users about state of their computations

Architecture

BigPanDA daily serves more than 17k queries. This number includes requests for Web pages from ordinary users and API invocations from users scripts and other monitorings.

Each request may involve millions of DB rows to process. To keep UI of BigPanDA responsiveness a proactive caching is involved.



BigPanDA system architecture

(\*) Depending on deployment options BigPanDA monitor supports Oracle or MySQL backend.

