

Application of rule-based data mining techniques to real time ATLAS Grid job monitoring data



Fachhochschule Köln
Cologne University of Applied Sciences

R. Ahrens^{1 2}, T. Harenberg¹, S. Kalinin¹,
P. Mättig¹, M. Sandhoff¹, T. dos Santos¹, F. Volkmer¹
on behalf of the ATLAS Collaboration



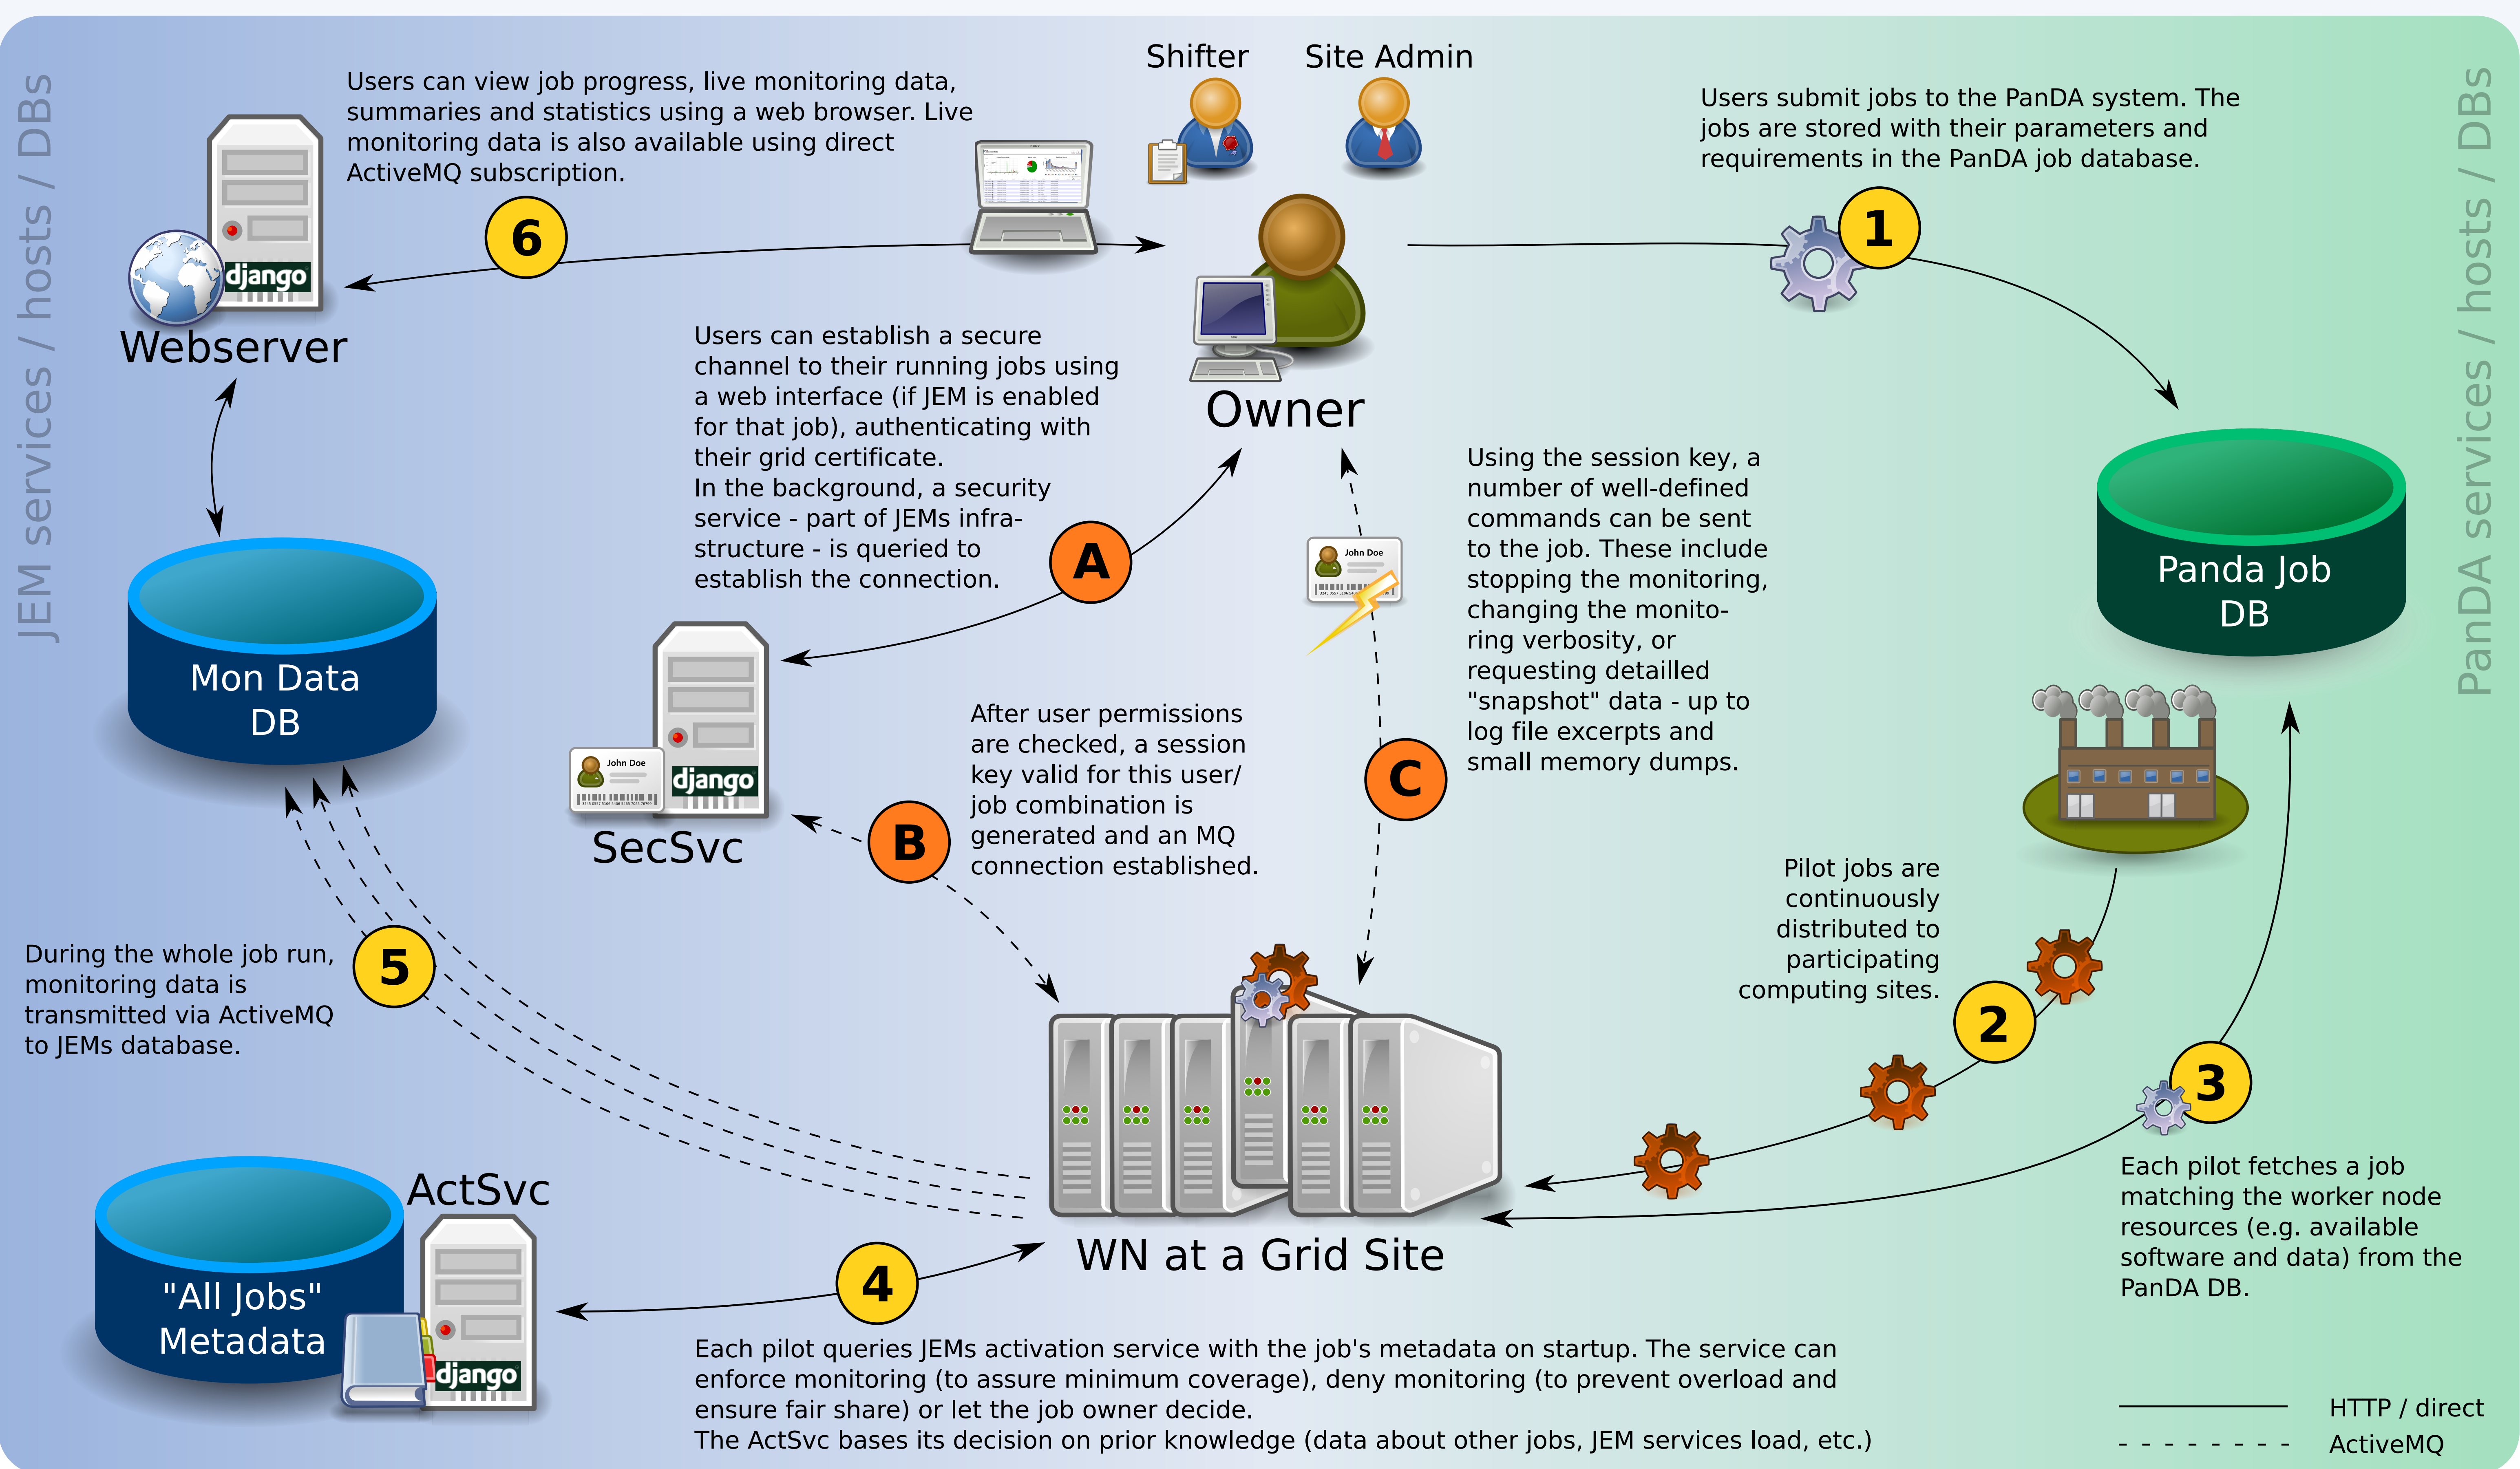
¹ Bergische Universität Wuppertal
² Fachhochschule Köln

Job centric live monitoring ①②③④⑤⑥

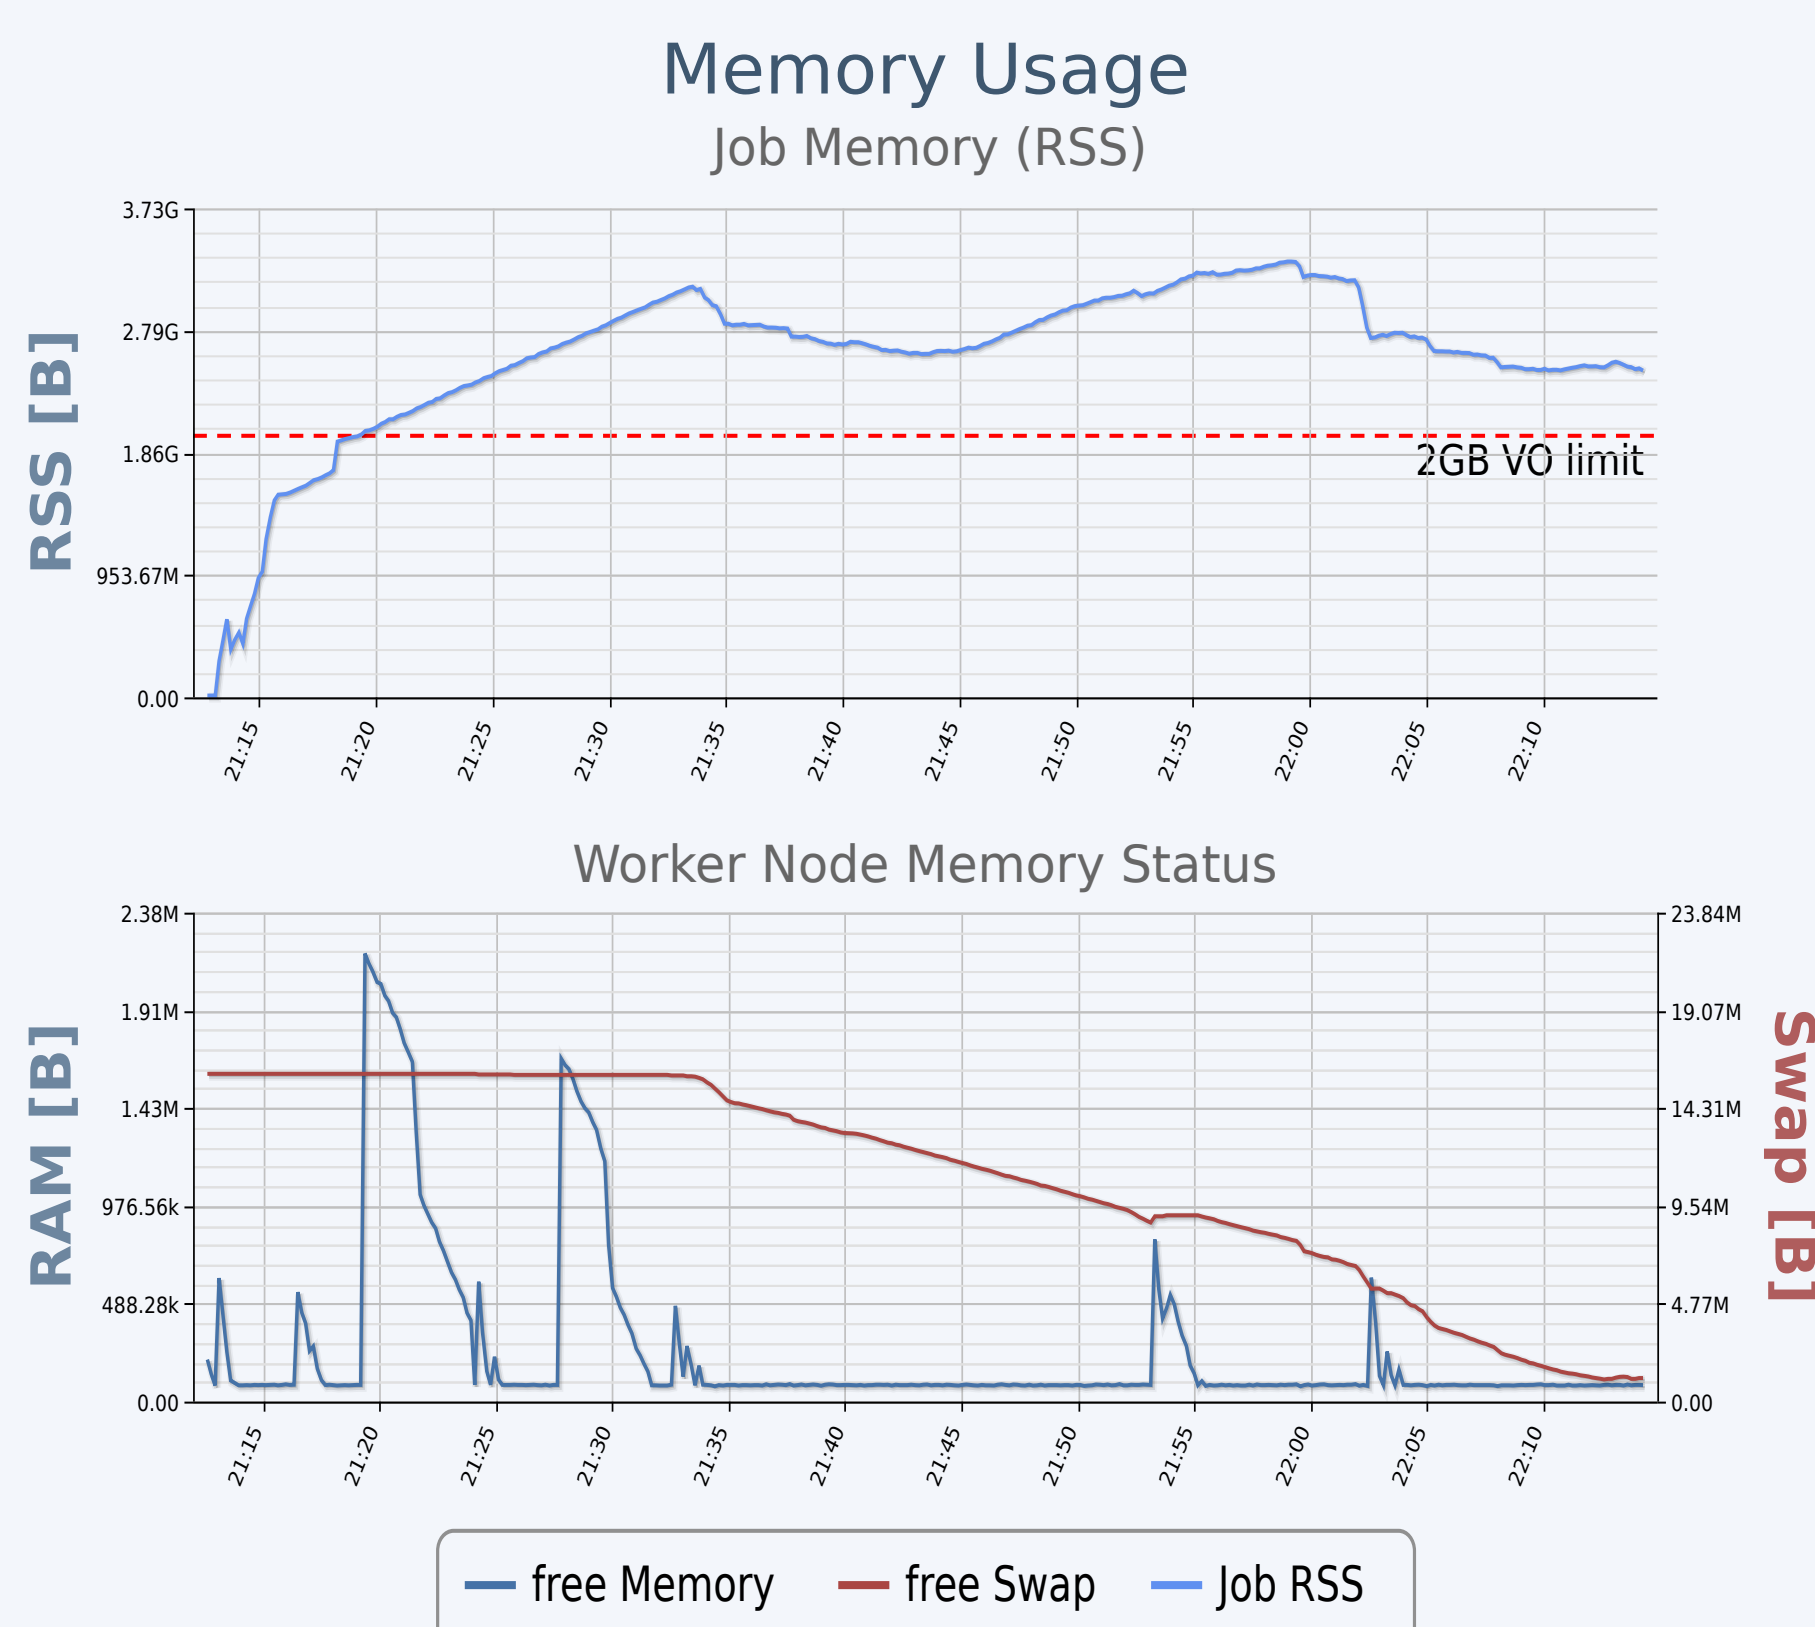
The Job Execution Monitor (JEM) is a toolset that allows to follow running Grid jobs' progress by providing system metric and job status data. JEM is integrated into ATLAS' distributed job brokerage system PanDA and can be enabled by users in their jobs without the need to modify them. The monitoring information is transferred using a secure channel to a central database that provides web-based access, plots and statistics.

Real time monitoring configuration A B C

To control the monitoring detail level and coverage of their jobs, a facility is being developed to allow users to send commands to the running JEM instance (in their job). These commands will allow to enable JEM only on request in the midst of a running job, to disable it if nothing interesting happens or to control its verbosity. The system is designed to provide a secure and authenticated channel between the job and the job owner.




Failure condition inference using data mining techniques



The resident set size of the job increased during its runtime, well over the assigned 2GB, while the free RAM and the free swap of the WN decreased significantly as other jobs of the same user

Automatic analysis of JEM live data checks all monitored jobs periodically for faulty behaviour and can generate simple rule based warnings to single jobs. For example, violations of thresholds and extrapolated trends of single or combined metrics as well as aggregations of many jobs running for the same user or on the same site / worker node can be detected. Outliers can be identified and corresponding actions taken (e.g. notifications to the owner of the job or the site admin). The example on the left shows the metrics of a typical misbehaving job versus the corresponding WN metric. Such faulty resource usage has the potential of affecting other jobs on the same WN, and can be detected early using JEM.

The distribution of all monitored jobs' average memory consumption is shown below for one site. Faulty job patterns can be detected, allowing site admins to take measures without the need to regularly check manually. Various other metrics also provided by JEM, like CPU consumption or network traffic, can also be checked.



Average job resident set size of all monitored jobs running on the Pleiades grid cluster in Wuppertal.