

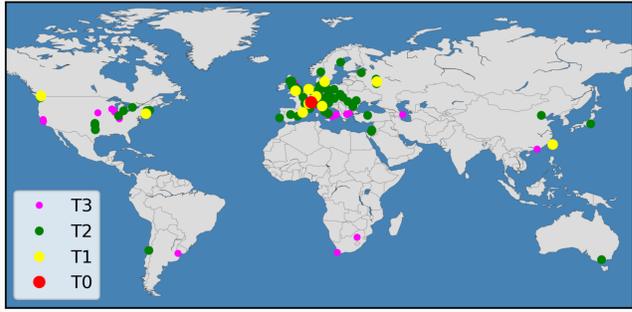
UNDERSTANDING THE EVOLUTION OF CONDITIONS DATA ACCESS THROUGH FRONTIER FOR THE ATLAS EXPERIMENT



M. Svatos, A. De Salvo, A. Dewhurst, E. Vamvakopoulos, J. Lozano Bahilo, N. Ozturk, J. Sanchez, D. Dykstra, on behalf of the ATLAS Collaboration

CHEP 2018, 9-13.7.2018

Introduction



The ATLAS Distributed Computing (ADC) runs $\mathcal{O}(100k)$ grid jobs on more than a hundred grid sites. Each job accesses Conditions, Trigger, and Geometry database data. Some jobs use only a minimal amount, some jobs access a significant amount of those data. The data are stored in the Oracle Offline Database at CERN. Access is provided by the Frontier service by means of the http protocol. Frontier servers are located at CERN and three Tier 1 sites. Tier 1 sites also store a copy of CERN Oracle database. Every site uses Squid to cache data from Frontier servers. During Run2 of the LHC, the amount of database data accesses increased. There have been multiple incidents where load generated by ATLAS grid workloads caused parts of the service to fail. There has been significant improvement in detection and monitoring of problematic tasks. Now, the incidents are quickly spotted and dealt with before they could destabilize the system. The monitoring of Frontier servers is performed by the AWStats tool. Squids are monitored by the Multi Router Traffic Grapher (MRTG) tool. Several other monitoring pages based on AWStats and MRTG data are available for shifters and experts.

Condition database request monitoring and alarms

ELK Workflow

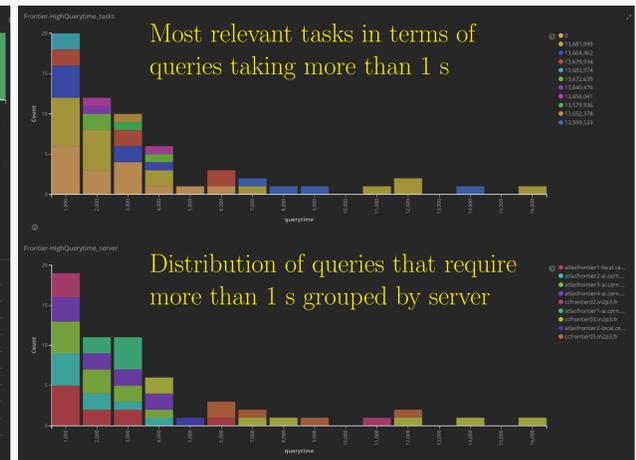
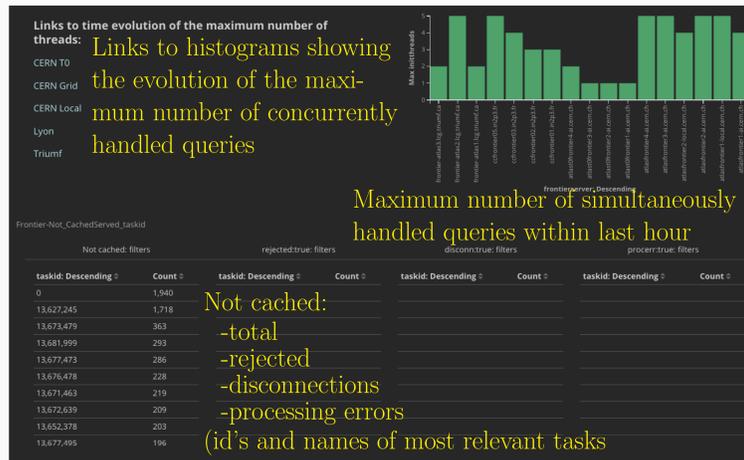
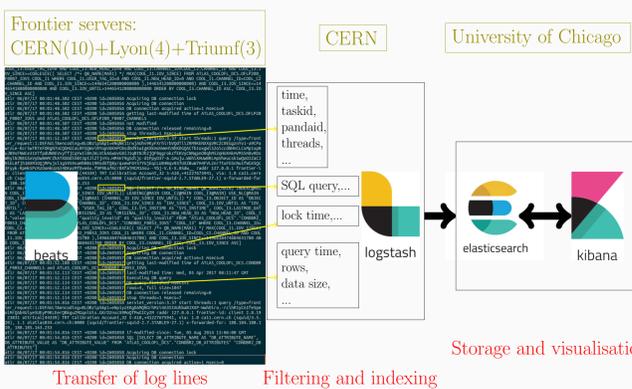
It is depicted schematically in the lower plot.

- Filebeat (Beats): collects log lines from servers and transfers them to a dedicated computer
- Logstash: filters the log lines aggregating and building tuples of variables (involved servers, user DN, task related information, SQL query, query status, times spent processing the query, etc...) associated to each query
- Elasticsearch DB: stores in a very efficient and redundant way the data coming from Logstash into indices
- Kibana: provides a web server to facilitate the visualisation of the data as histograms, tables, pie-charts, etc. Objects can then be grouped into persistent Dashboards (as seen below).

Alerts&Alarms system

When conditions indicate a foreseeable overload of the servers, an Alarms and Alerts (A&A) system informs experts. This A&A system is implemented on the basis of a Jupyter notebook executed periodically by a cron job and an electronic mail subscription service. It sends email alerts based on

- the maximum number of concurrent queries
- or the number of queries that were rejected or disconnected or that failed while being processed
- or when there is an abnormally high percentage of queries with high total execution times (above 1 s)



Frontier server monitoring

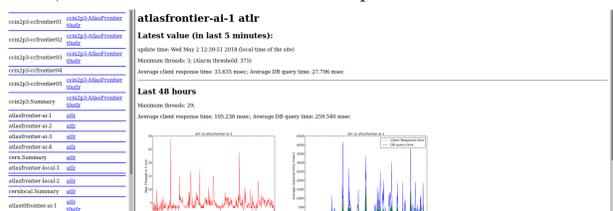
Awstats

AWStats is a tool that generates advanced web, streaming, ftp or mail server statistics. Data from squids running on each ATLAS frontier service (number of unique visitors, number of visits, pages, hits, and bandwidth) are collected and displayed by this tool.



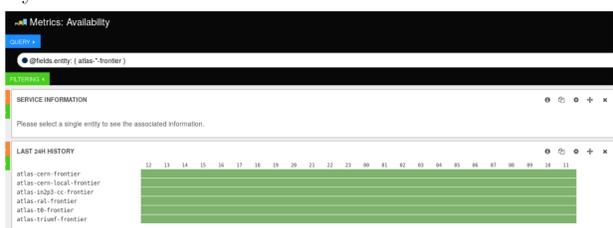
Maxthreads

Maxthreads checks the maximum number of threads used by frontier servlets in tomcat. It also monitors client response time and DB query time. In case the number of threads exceeds a predefined threshold, an alarm email is sent to experts.



Kibana

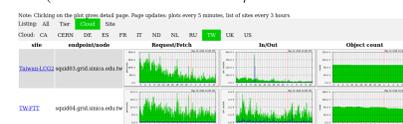
Availability of servlets is present also in Kibana monitoring of ATLAS central services. It is regularly checked by shifters.



Site squid monitoring

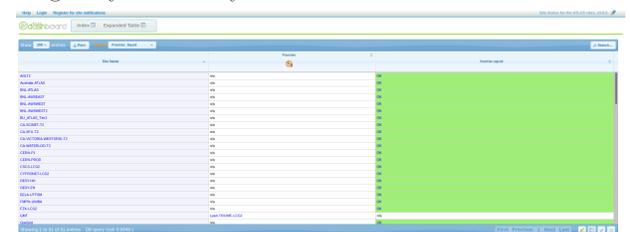
MRTG

The Multi Router Traffic Grapher (MRTG) is a tool to monitor the traffic load on network links. The newly developed ATLAS MRTG monitoring shows MRTG plots of active squids of active ATLAS sites (based on GOCDB/OIM and AGIS).



SSB

Site Status Board (SSB) is an ATLAS monitoring framework used to monitor various functionalities of sites. The squid monitoring in the SSB aggregates site status from the ATLAS MRTG page. It is regularly checked by shifters.



Failover

The failover monitor reads AWStats data of backup proxies at CERN and FNAL and of frontier servers. The framework detects and reports persistent excessive activity from worker nodes.

