

The first results for jobs preemption on Titan

Danila Oleynik
WFMS meeting 11/04/2019


Reminder

- ATLAS was allowed consuming the “backfill” walltime on Titan. “Backfill” means that this walltime was not distributed among other users by the local resource management system.
- PanDA components were designed for fast harvesting of resources. Harvester monitor available resources on the machine, and grab acceptable resources as only appear. Algorithm does not allow increasing of priority of ATLAS jobs with ageing.
 - To minimize impact to other users of Titan, ATLAS associated jobs have minimal possible priority in the system (-367-day penalty by priority)
 - ATLAS is allowed to use a higher number of job slots than other users: up to 20 (8 used for the moment)
 - Dynamic size of job slots:
 - 15-150 nodes
 - one PanDA job per node
 - 50 Events per job
 - Without preemption, ATLAS compete for resources only during scheduling. ATLAS jobs was not interrupted during execution
- Preemption allows the local resource manager to shoot out ATLAS jobs during execution, if job from another user, with higher priority, will come

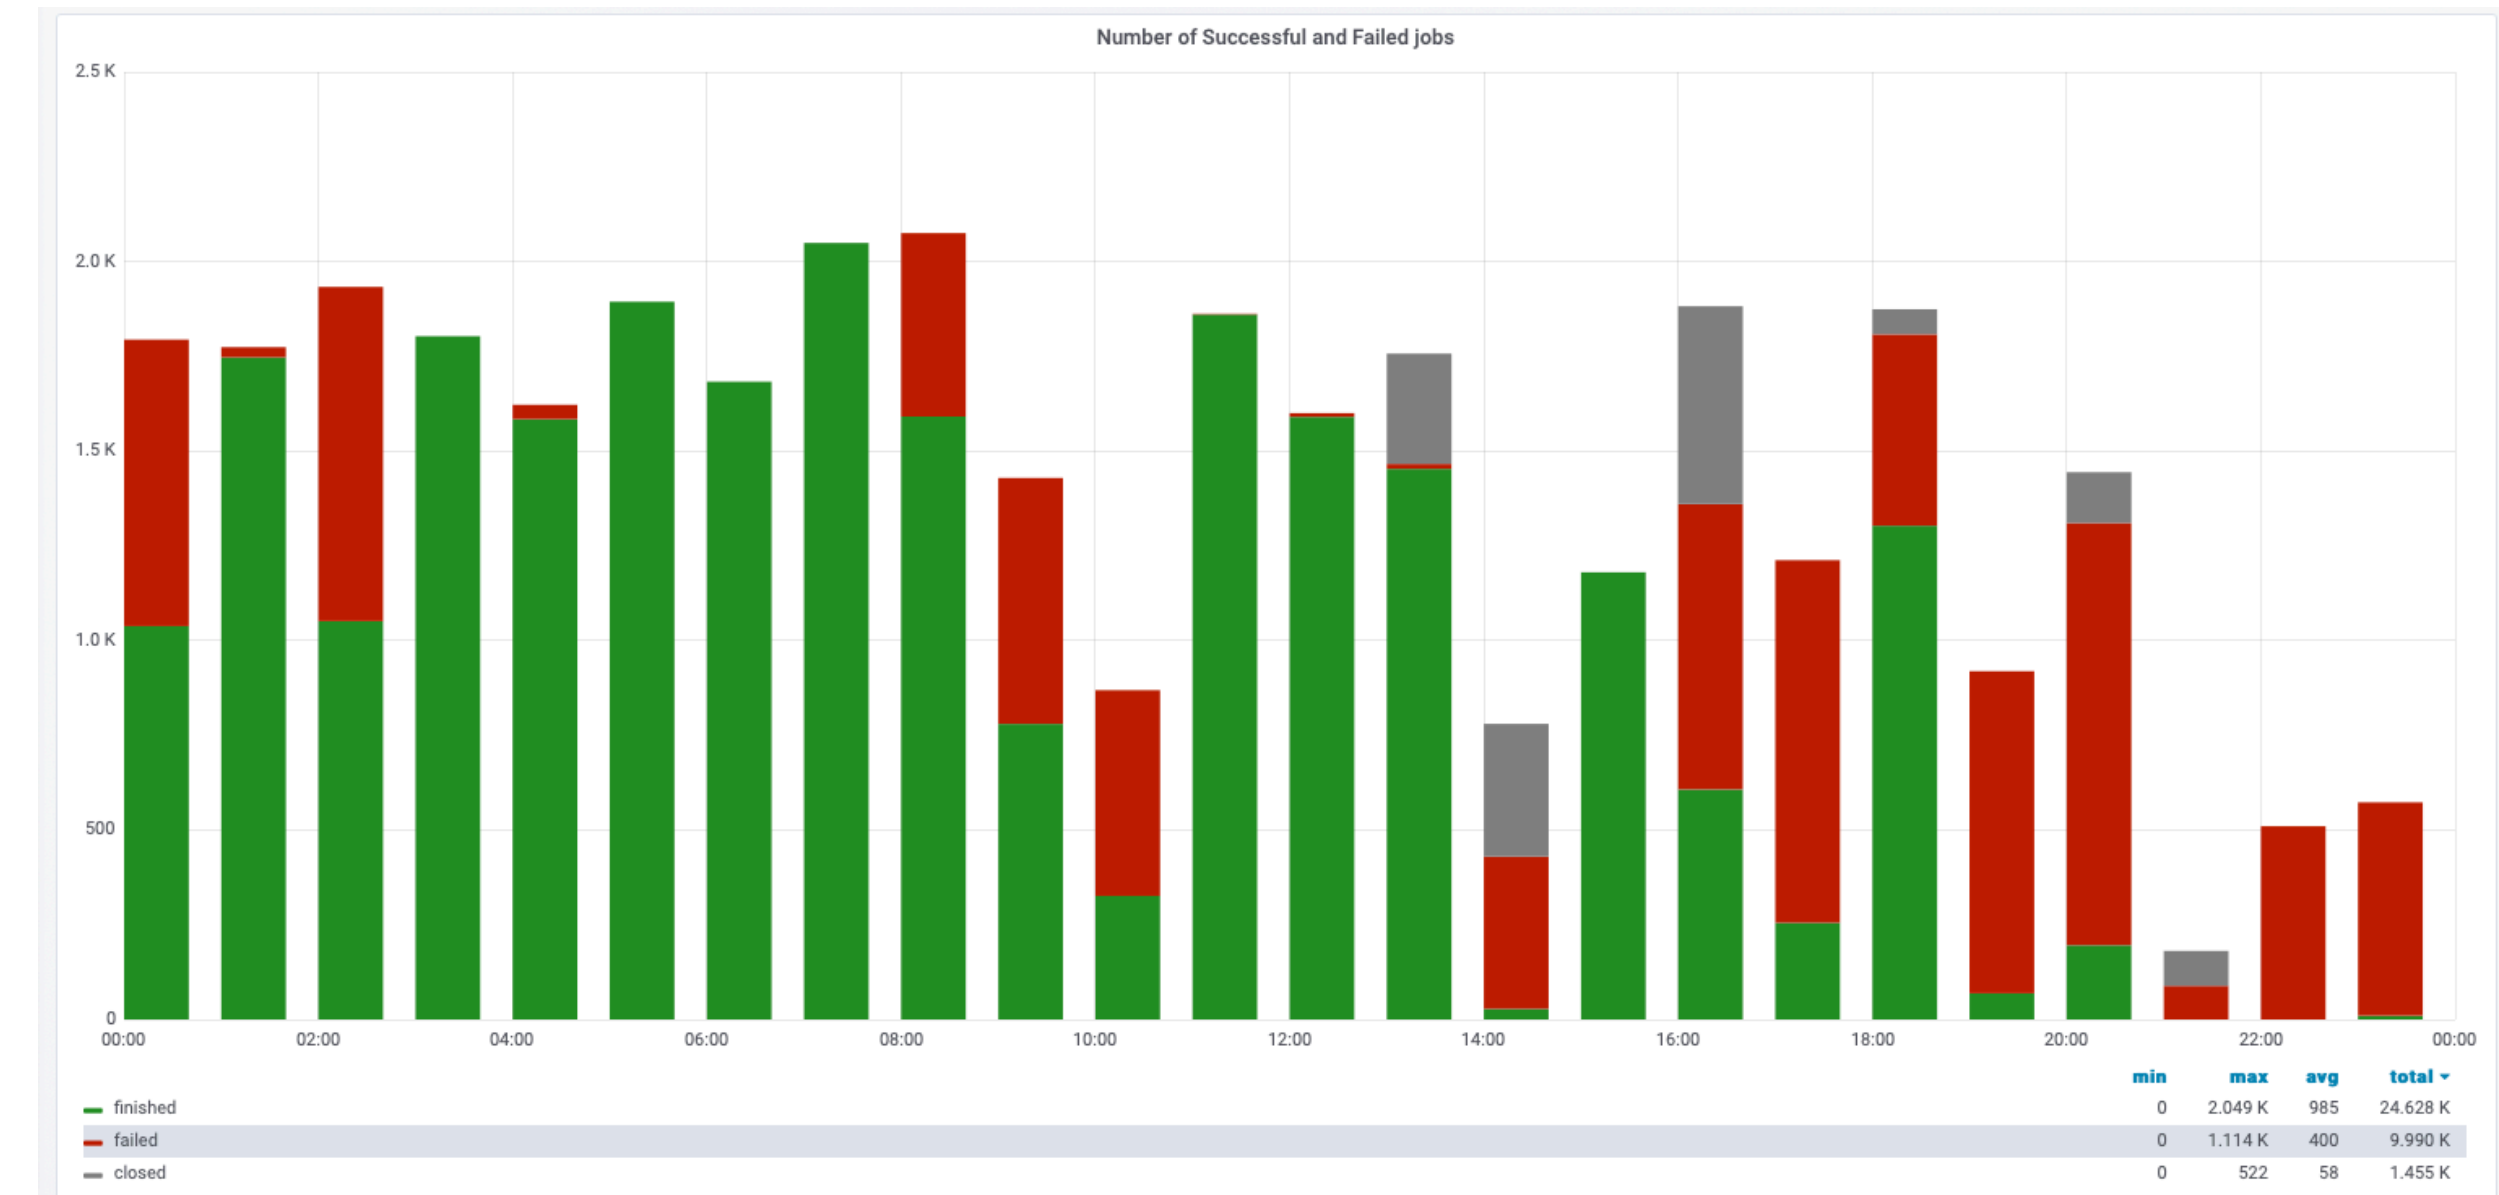
Preemption on Titan

- ... was not used/activated before
 - First experience for OLCF team
 - After set of discussions in BigPanDA project, we were chosen as "guinea pig"
 - A set of different test will be performed next days/weeks
- Two main parameters for preemption were identified:
 - Job priority
 - "Grace period" - time, while job can't be preempted
- First round of test started last Monday
 - Job priority - for ATLAS did not changed (still minimum in the system)
 - "Grace period" - 20 minutes

First observation... I thought it would be worse




- April 3. Preemption on.
- 26K competed jobs, 5 failed




- April 9. Preemption on.
- 24K competed jobs, ~10K failed

Wasted walltime



- 28% failed jobs



- 19% wasted wallatime (failed jobs)

- I have an impression, that wasted wallatime can be decreased if “grace period” will set to 0

Timing

- Current simulation task: <https://bigpanda.cern.ch/task/17584907/>
- Job required waltime: 27 ± 3 Min.
- Pilot overhead: 1,5-2 Min.
- Increasing of "grace period" to 30 Min will minimise number of interrupted jobs
- Increasing internal priority for our submissions will decrease probability of interruption of jobs
- Tests will continue with changing of different parameters

