




ATLAS production @ OLCF. October 22, 2018

- 38,3M of events were delivered in October.
 - «Backfill» ~24,3M Events (15,3M Titan core*hours consumed)
 - ALCC:~14M events delivered (7M Titan core*hours)
 - Looks like that INCITE domination is over and our tasks launches more often from last week
 - Since we have stable execution with 1200 nodes per worker with up 2 concurrent submissions, we decide to try something bigger - today worker with 3800 jobs was submitted (to get longer walltime and priority boost)



Big workers. Scalability.

- Last week we got a new MC task for 25M events. Required wall time for processing of 1000 events for this task ~14-15 hours.
- We decided to try bigger workers (3800 nodes) to move to the next bin with walltime limit for 24H (and got 5-day priority boost)
 - First few runs were ok from Harvester point of view
 - At some point, few workers were failed due to a system failure,
 - That created some backlog of jobs to be processed in PanDA (updating states etc.) and cause increasing of loading to Harvester in the current configuration, and looks like that memory consumption goes to critical values
 - Ongoing work to define a proper configuration of Harvester to threat big workers safely

Harvester with backfill capability

- Testing with small (3) number of slots from last week
 - Dynamic size of workers from 10 to 450 nodes
 - Works stable, but a couple of improvements required to go to production