

# YODA commissioning status at OLCF

## August 06, 2018

- Commissioning started at the end of June
  - A set of issues were identified during deployment and initial runs of the application
- Some of the issues were caused by a difference in the environment and some Titan specialty
  - This types of issues didn't allow to launch Yoda properly and got fixes somehow quickly
- Yoda has launched on small scale: 3-50 nodes, to collect initial information about application behavior
  - A set of successful runs were performed, which shows that the whole ES workflow works correctly at OLCF:
    - Events were simulated
    - Staged-out to external storage
    - Successfully merged on the GRID for future usage
- Another set of issues were found, related to Yoda efficiency in terms usage of allocated resources
  - Some of 'droids' overladed with work, while some others are empty
  - All these issues are tracked
- In last month we got at least 4 major updates of application and one patch for production release of Athena MP

# Yoda: Initial IO studies

- IO profiles were collected for the set of Yoda runs.
- These profiles showed known behaviour for Python based application on the shared file system: quite intensive metadata operations during traversing of shared libraries
  - To cope this issue code of Yoda were moved to NFS space
- Due to the features of the implementation of checkpointing in ES we expect to produce of high number of small output files
  - Yoda will be needed to be optimized to perform some packing of output files on a rank level
  - IO from payloads will be decoupled from the shared file system with moving of working directories RAM disk of WN

# Harvester modification for AES at OLCF

- As our primary aim of AES at OLCF is backfill resources, Harvester should be extended with the capability of the dynamic shaping of a size of the worker
- It also may help as to finally get rid of MultiJob pilot with all his limitations on scale