ML Payloads at OLCF

Danila Oleynik, Sergey Panitkin

“TO DO List” 1: Validation of container with ML software on Titan environment

- As expected, the default container did not work properly on Titan:
  - The major reason: “Creating containers for OLCF HPC systems requires special consideration due to proprietary system software such as MPI and CUDA”
  - To build compatible with Titan images special service “container builder” was deployed by OLCF team
    - For using this tool for building user containers we need to have a 'recipe' - in general, it’s a list of software which should be included in the container
- Sergey had to prepare a container with coreML software (Tensorflow, Keras, Horovod) compatible with Titan
  - Basic tests was passed after a set of iterations

“TO DO List” 1: Containers in production

- As we can’t use same containers for HPCs and GRID sites we need to have a solution which will allow properly match containers and jobs for different infrastructures
  - Naming convention?
  - Registry in AGIS?
- Containers at OLCF contain only core software, without code of transformation itself
  - How transformations will be distributed? Special releases?

“TO DO List” 2: Adoption of Pilot2 HPC workflow to manage (ML) containers

- On the current step i don’t see complicated changes in HPC workflow of Pilot 2 to support ML containers, and most of changes will be in the configuration
  - On Titan we will need to use dedicated PanDA queue for ML payloads, because of difference in environment setup against production jobs

Summit

- Recently BigPanDA got Director's Discretion (DD) allocation on Summit
  - This allocation will allow performing some RnDs including the study of ML payloads and building of ATLAS Releases for PowerPC
- For the moment there is no support of containers on Summit, but there is the plan to have production support of Singularity ecosystem before the during summer 2019

Containers in context

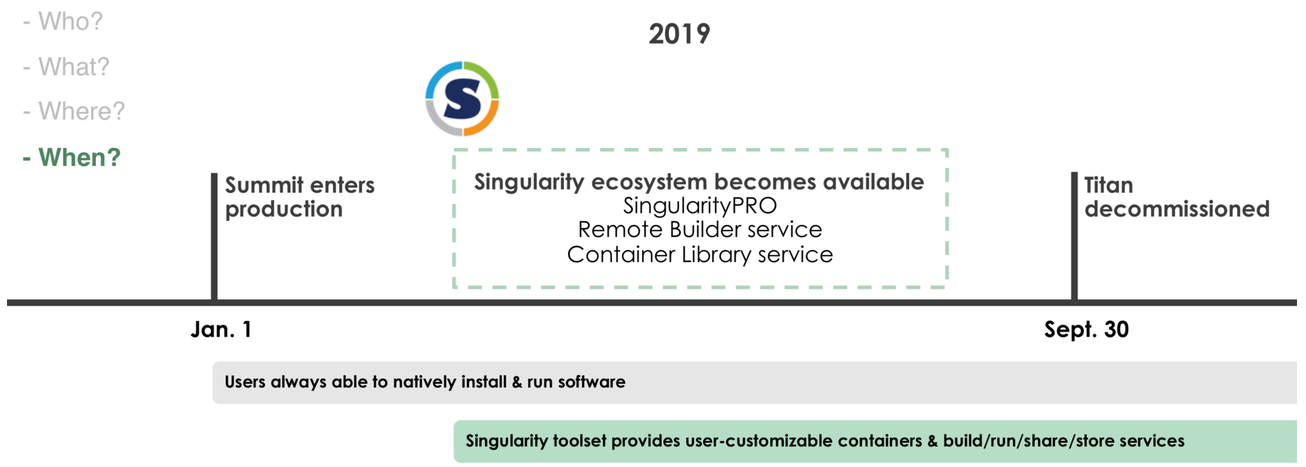

ML payloads on Summit

- Software stack for distributed ML (TF, Keras, Horovod) was deployed on Summit.
  - A set of test with ML+ Horovod was ran with up to 200 V100 GPUs utilized per job.
- Harvester instance was installed on Summit, now working on integration with local batch system (Spectrum LSF)