Research track 2: Data transfer and access

Sergey Panitkin (BNL)
Danila Oleynik (UTA)

Outline

- Introduction
- CSC108 operations on Titan
- Running on the new DTN farm
- Outlook and Plans

Topics in the proposal for RT2

- Workload run-time IO optimization
  - IO related effects that can affect workload execution time and have negative impact on OLCF infrastructure
    - Workload startup
    - Workload run time IO
    - Luster MDS overload
- Move pilot operations to DTNs at OLCF
- FTS3 at ORNL
  - 3rd party transfers, etc. Pre and Post (workload execution) data management. Asynchronous IO data management for workload .

Overview of the current status

- Many issues outlined in the proposal are addressed
- Placement of ATLAS software on NFS allowed to avoid Luster MDS overload during AthenaMP/Geant4 startup and operation.
  - Current payload IO operations are within Luster operational limits (Study by Sarp Oral – RT4)
  - Can run ATLAS production simulations in parallel on up to 6000 nodes (96k cores)
- Change in pilot launcher logic (non-blocking launcher) allowed us to remove pilot's stage-out phase from the critical path and improved backfill utilization efficiency. Data stage-out does not affect anymore our ability to capture free resources on Titan
    - New launcher operational since July 2016
    - Large effect - allowed us to exceed 7M core-hours per month with backfill utilization efficiency of ~19%
- Pilots operations were moved to DTNs (twice!)
    - New DTN farm became available in February 2017
- The move to the new DTN farm allowed us to exceed ~10M core hours per month
    - For March and April of 2017 ATLAS resource utilization on Titan exceeded 14M core hours and backfill utilization exceeded 30% .
    - That puts us on track to reach and exceed 100M core hours in 2017

CSC108 performance overview

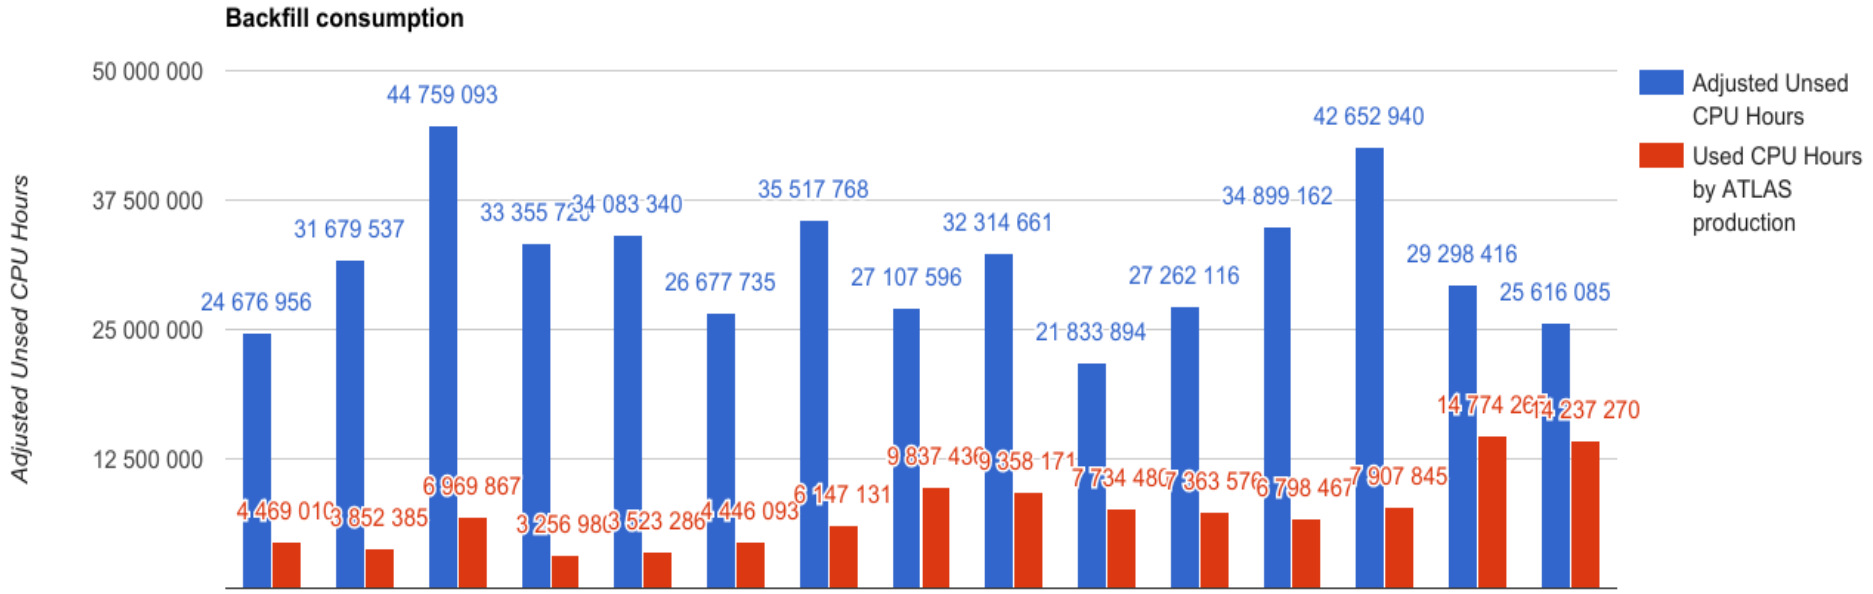

Running ATLAS production simulations 24/7 since 9/2015
Pure backfill mode. No allocation. Lowest priority on Titan
Steadily increasing CPU consumption and backfill utilization efficiency

Evolution of CSC108 performance on Titan

~74M core-hours in 2016
~37M c-h in three month of 2017

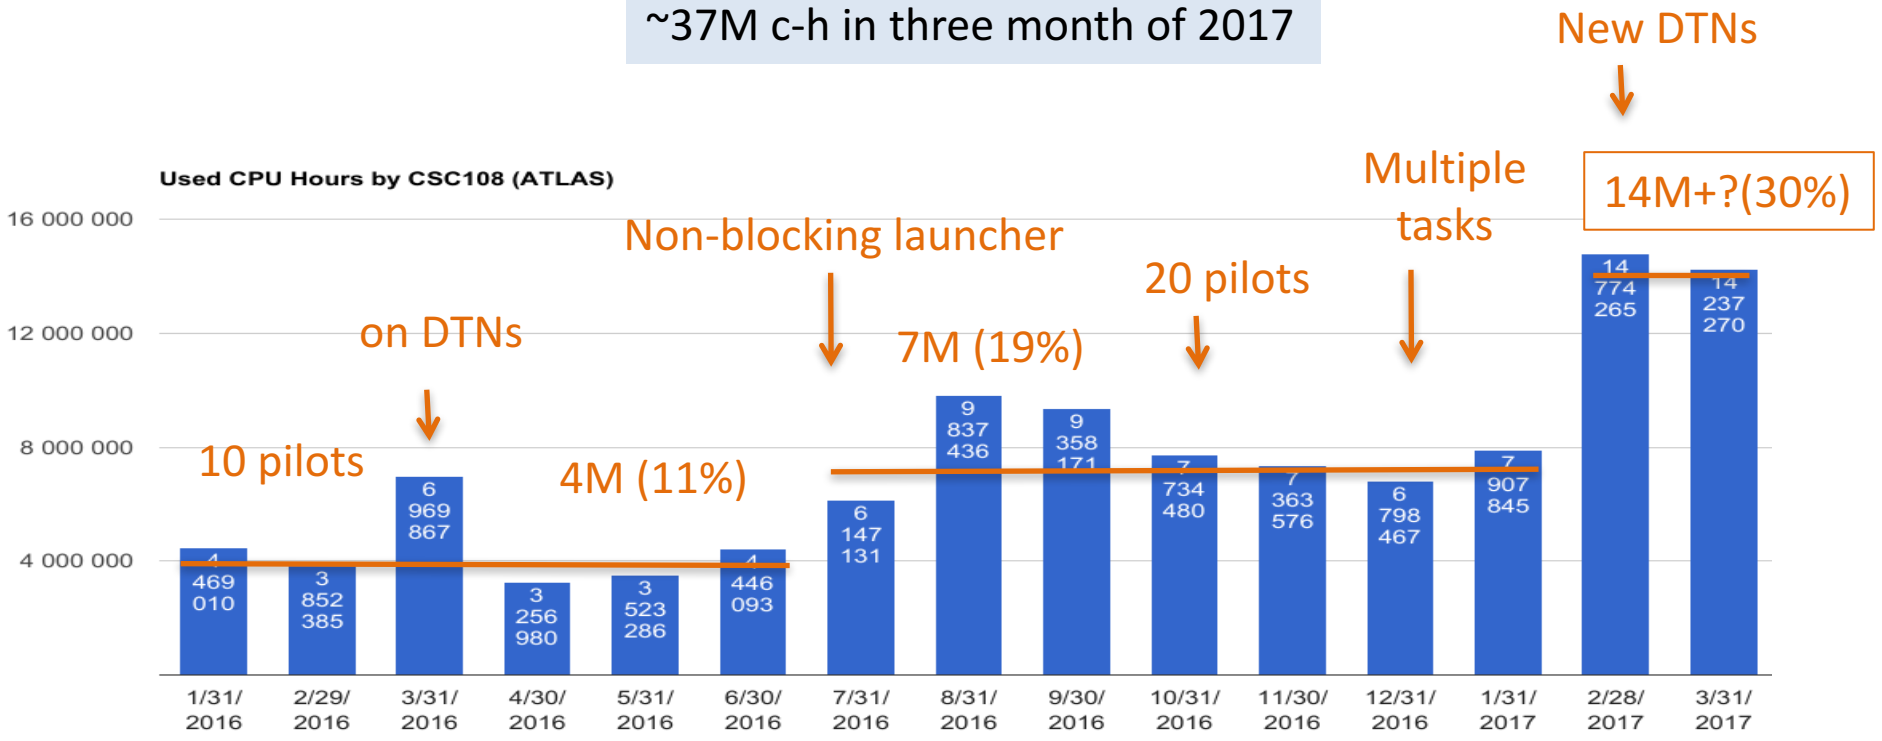

Running ATLAS production simulations 24/7 since 9/2015
Pure backfill mode. No allocation. Lowest priority on Titan
Steadily increasing CPU consumption and backfill utilization efficiency
Increases are directly related to improvements in data transfer operations

Running DTNs

- Pilot operations were moved to DTN cluster in March of 2016.
  - No big improvement in resource utilization or backfill utilization efficiency
  - Data transfer performance was no better than on login nodes
  - Reports of DTN overload at peak ATLAS (CPU bound?)
    - Limited number of total number of pilots to 20 (5 per DTN)
- In February 2017 operations were moved to the new DTN cluster
- The move to new DTNs brought improvements in data transfer, decreased pilot “reaction time” and as a result significantly improved backfill utilization on Titan

Data transfer on DTN (I). Stage-in.

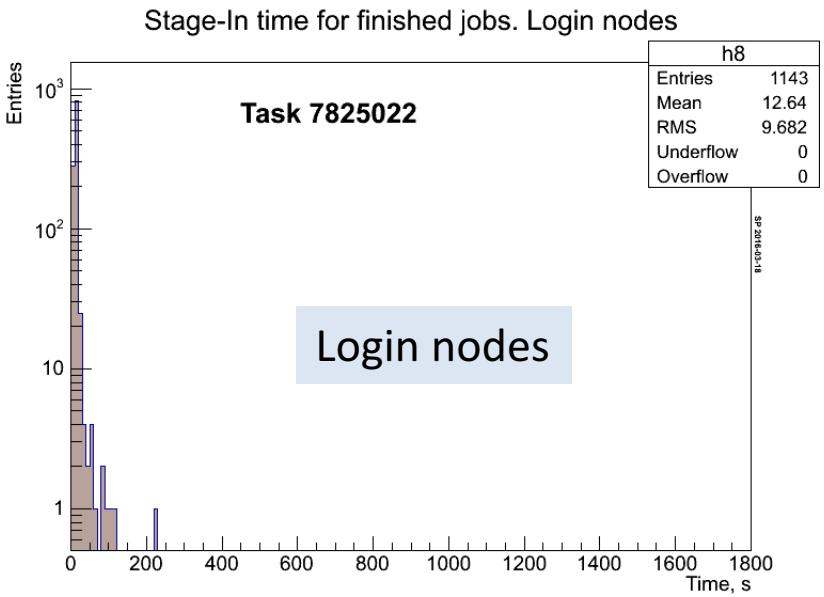

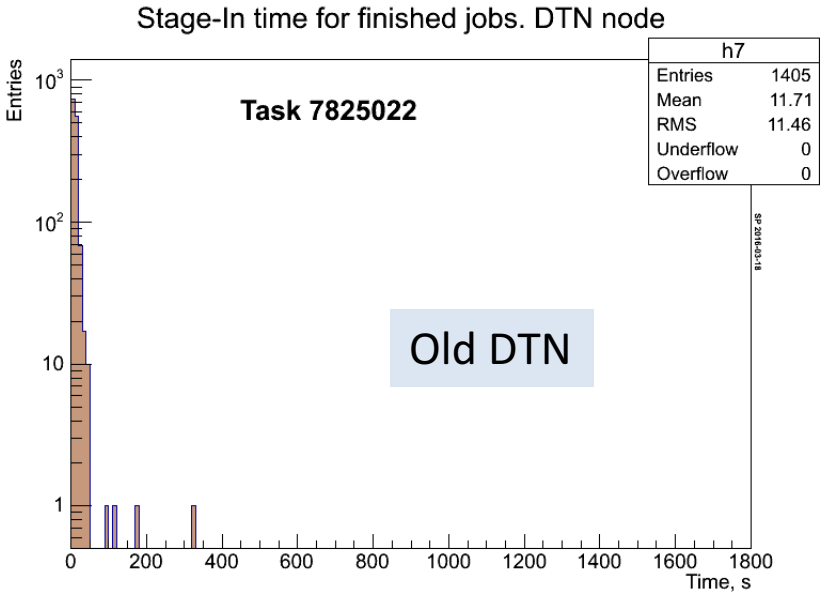

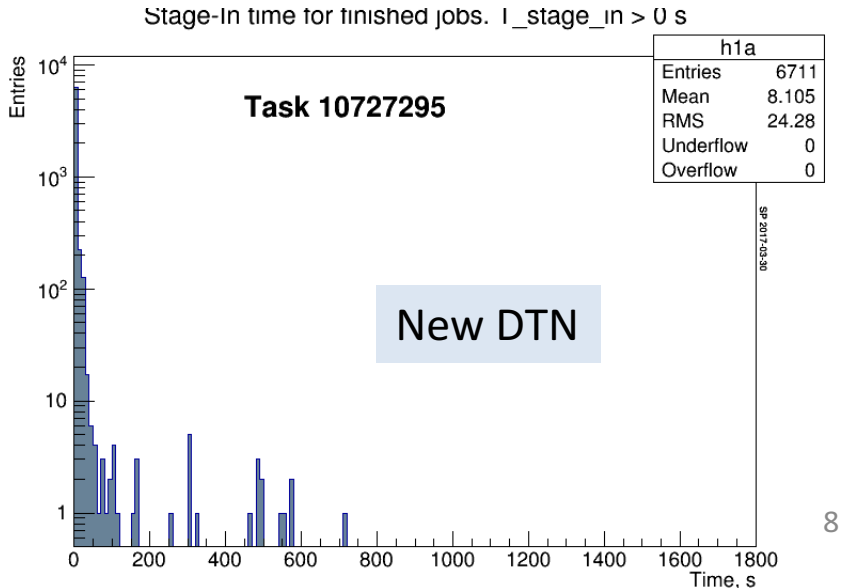

- Stage-in on the new DTN nodes seems to be faster
- This affects backfill capture speed

Data transfer on DTN (II). Stage-out.

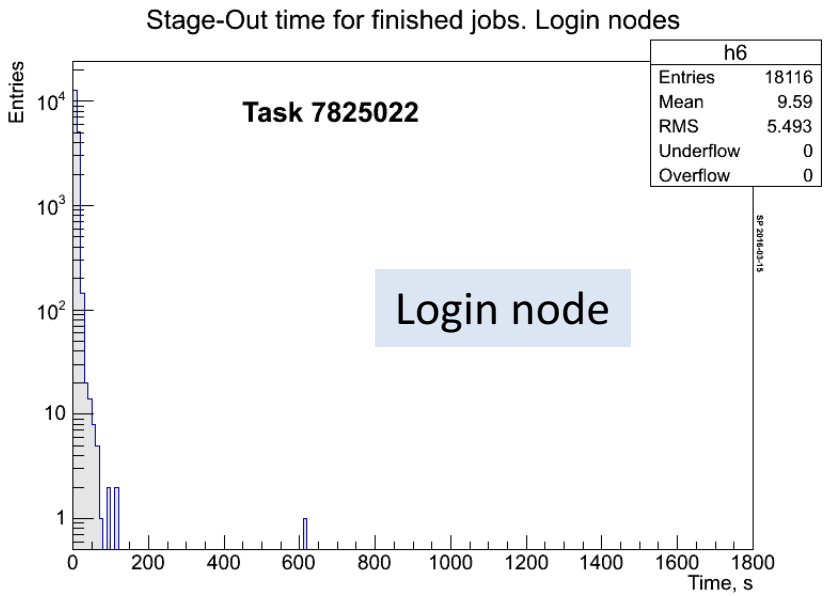

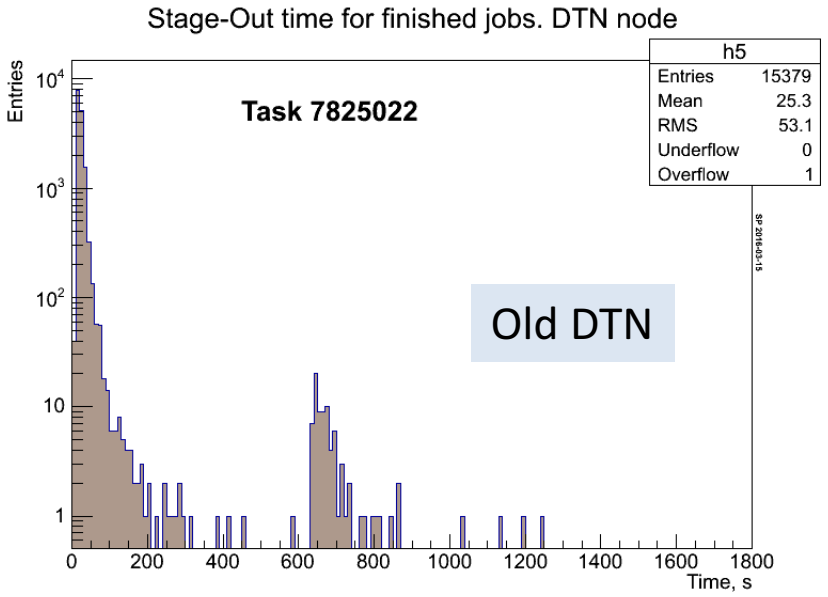

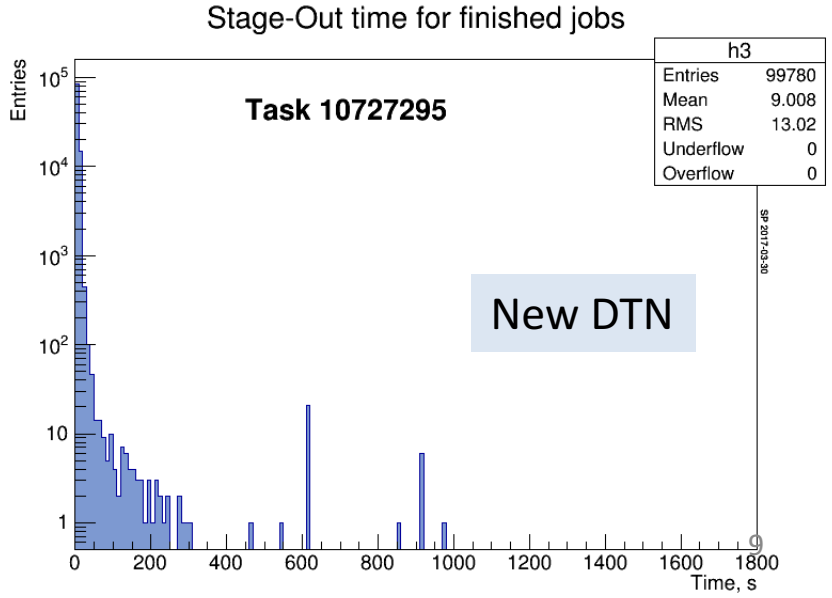

- Stage-out on the new DTN nodes seems to be faster as well
- This improves

Data transfer times for resent task

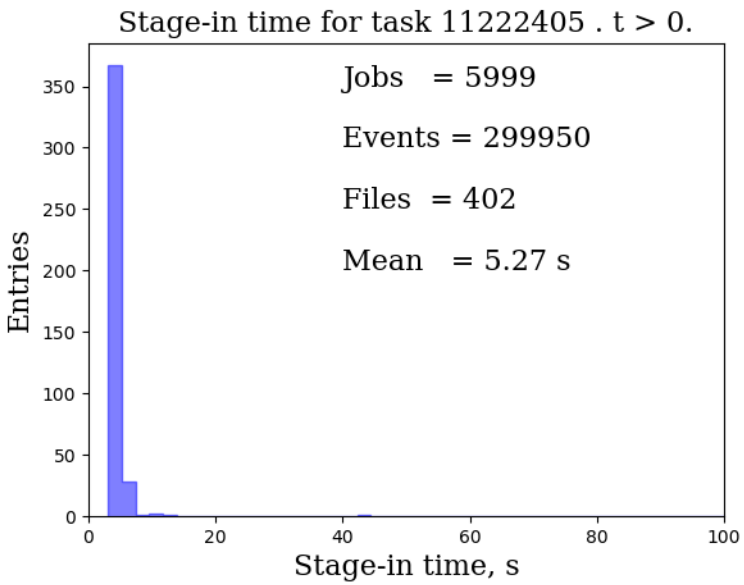

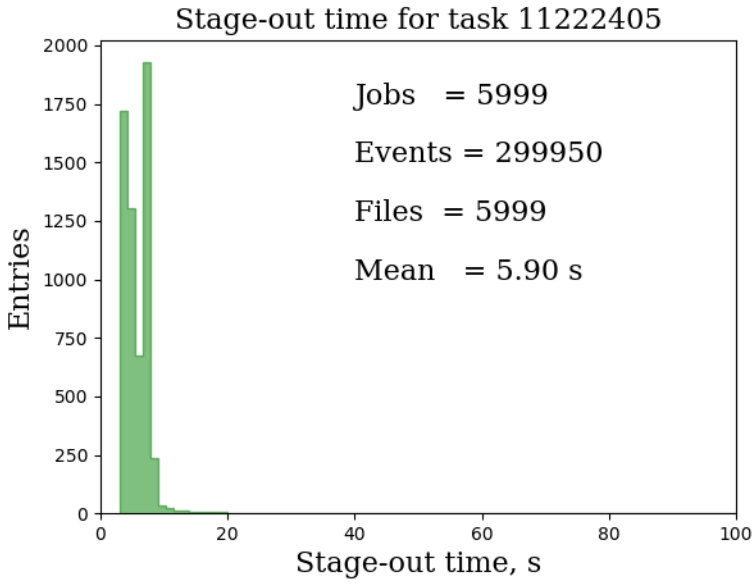

Recent task 11222405 with 50 events per job

Running on the new DTN cluster

- In February 2017 operations were moved to the new DTN cluster
  - New CPUs, more bandwidth
  - Updated version of the GFAL client libraries installed
  - Big jump in Titan resource utilization (14M core hours) and backfill utilization efficiency (33%) in the first month of running
  - Better data transfer performance than on old DTNs
  - Performed study of the CPU loads on the new DTNs
  - Based on SAR data no CPU overload on the new DTNs was observed so far
    - Average CPU load on the utilized DTNs is light and varies from 1% to 6% over the period of observation. Observed peak load ~30%
    - Steady state ATLAS production. On average ~9k Titan cores used, with max peak of ~70k cores and several peaks at ~30k cores

Outlook and plans

- CSC108 backfill utilization currently directly depends on data transfer properties
- What can be done with current setup
  - Increase number of pilots ? Increase maximum number of ranks per submission? Increase limit of stage out pilots in pilot launcher? IMHO “Yes” to all of these questions
  - Number of ranks was increased from 300 to 350
  - What are the optimal numbers? What are optimal load levels on DTNs under these conditions?
- Continue monitoring of loads on DTNs. Correlate load on DTNs with pilot activities
  - Currently only manual operations are possible
- Clear need in analytics platform that can collect relevant information from various sources (DTNs, pilot launcher, pilots, Moab, PanDA, etc) in one place
- New edge services Harvester and Pilot 2 are under development in ATLAS
  - New architecture should help to improve overall performance on Titan and data management in particular
  - Parallel data transfer instead of current sequential transfer model
  - Asynchronous IO
    - Stage-in ahead of time based on assigned tasks information?
    - Multiple stage-out sites to improve robustness of the stage-out?
  - DTNs should have enough “muscle” and bandwidth for Harvester
  - No need in FTS3 with Harvester in production ?
- Containers on Titan?
  - Configurations , IO performance, start-up times, etc. NERSC experience was positive.