CMS Computing Model

General Changes

- Nearly all the changes proposed by CMS are intended to make more flexible and efficient use of resources
  - Workflows and data paths are largely the same as Run1.
    - A possible exception is the change to a transient RECO, that saves writing large quantities of derived data to tape. RECO is written to disk storage for a window of time. By default we use 3 months, but can be optimized by dataset
  - Optimizations are in code performance and in a general decoupling of CPU and storage resources, so there is more flexibility in where workflows run and how computing resources are used

Computing Model Changes

- The largest modifications are in capabilities and expectations of the sites, only the Xrootd data server and CVMFS are added as services
  - Archival and disk storage are separated so the Tier-1 sites are similar to Tier-2s in functionality, but have higher availability and support
  - The Tier-0, Tier-1s and Tier-2s become active members of the data federation
  - These changes allow
    - Tier-1 sites to participate in analysis and Prompt Reconstruction
    - Tier-2 sites to perform simulation and data reprocessing
    - The HLT to become a massive reprocessing resource
    - Opportunistic resources can be built up
    - Differences between the capabilities of the sites are diminished, but availability of the Tier-1s become at least as important

Software Changes

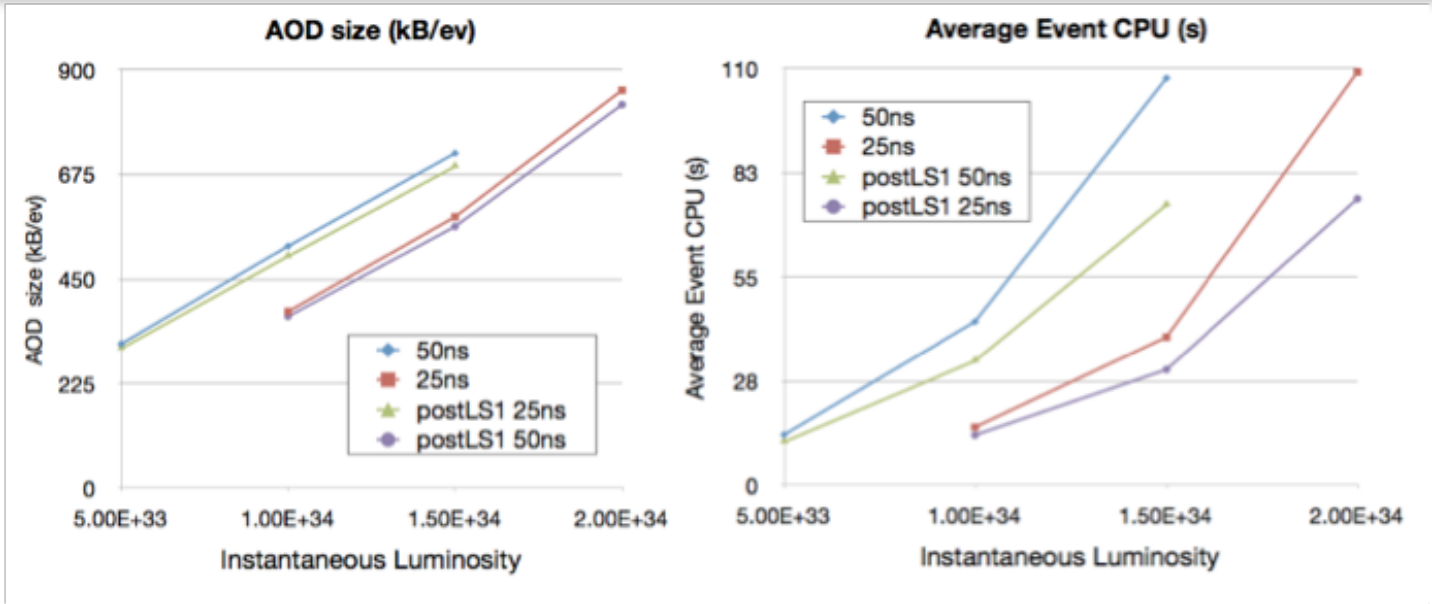

- The software continues to improve during the long shutdown, but we are seeing diminishing returns for changes that completely maintain physics performance.
- The production version of the event level parallel version of CMSSW is expected in a few weeks
- Work on porting the code to alternative architectures is ongoing (ARM and GPU investigations), and will open new types of computing resources and potentially big improvements.