U.S. CMS Operations Program

The next five years: U.S. CMS

Ken Bloom

8 October 2015

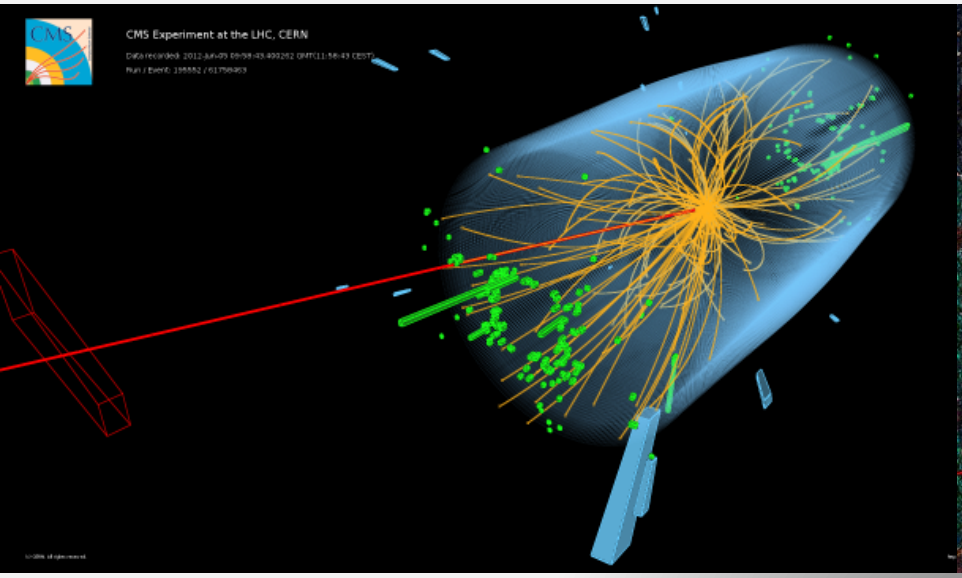

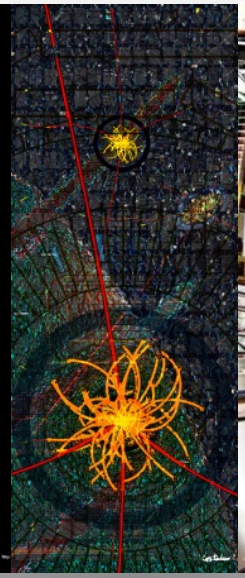

CMS science needs computing

- The scientific problems of particle physics are naturally addressed through high-throughput computing
  - Has been true since long before we knew what “high-throughput computing” actually was!
  - Individual particle interactions (“events”) are statistically independent, and can thus be analyzed in an embarrassingly parallel fashion
- The scale of the data naturally lends itself to HTC
  - CMS expects to record 9.2 PB of raw data in 2015-17
  - Data are then reconstructed, secondary datasets created
  - Simulations data have similar size and processing needs
  - Overall CMS computing needs by 2017: 1692 kHS06 ≃ 155,000 batch slots, 115 PB disk, 193 PB archival storage

On the grid

- The political nature of the LHC collaborations lends itself to *distributed* HTC
- 78% of the processing and 85% of the storage live outside of CERN, the host lab for the experiment.
  - 7 “Tier-1” sites in 7 countries
    - 24x7 support, archival tape
    - Fermilab T1 on the OSG, largest single CMS computing site
  - ~50 “Tier-2” sites in 24 countries
    - Business hour support
    - 8 U.S. and 2 Brazil sites on OSG
  - And numerous “Tier-3” sites at universities
- The U.S. sites comprise ~40% of the CMS computing resources outside of CERN
- And U.S. sites provide (28+10)% of OSG computing hours
- U.S. CMS is a very important customer, and OSG provides critical support to fulfilling CMS science goals

Usage in the last year

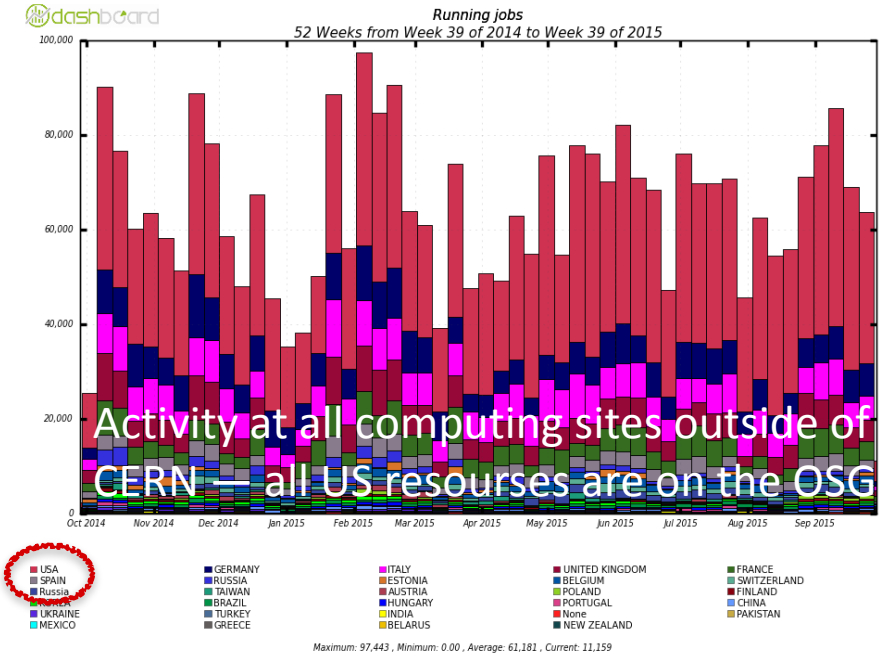

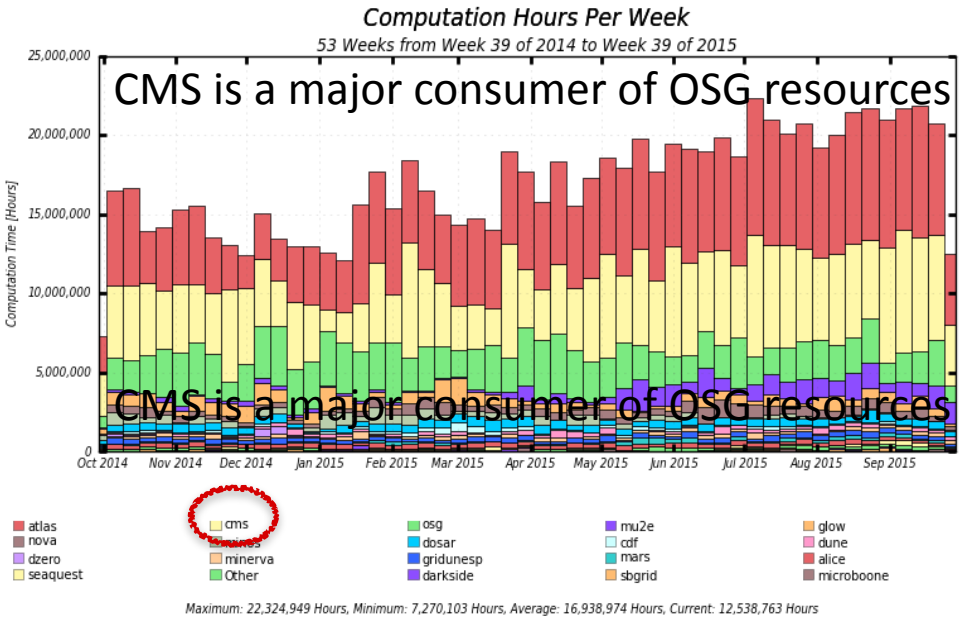

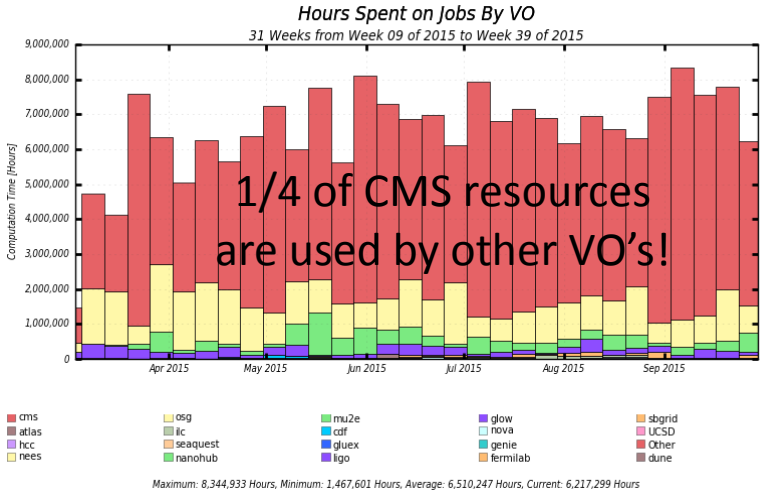

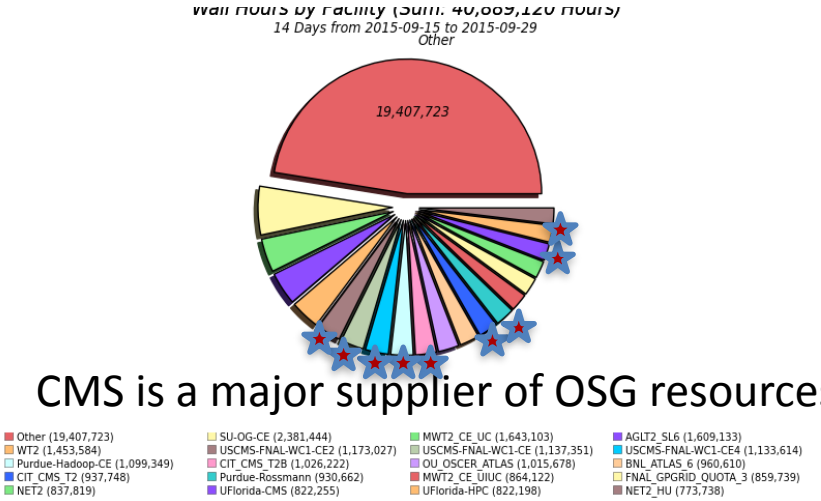

Only the beginning

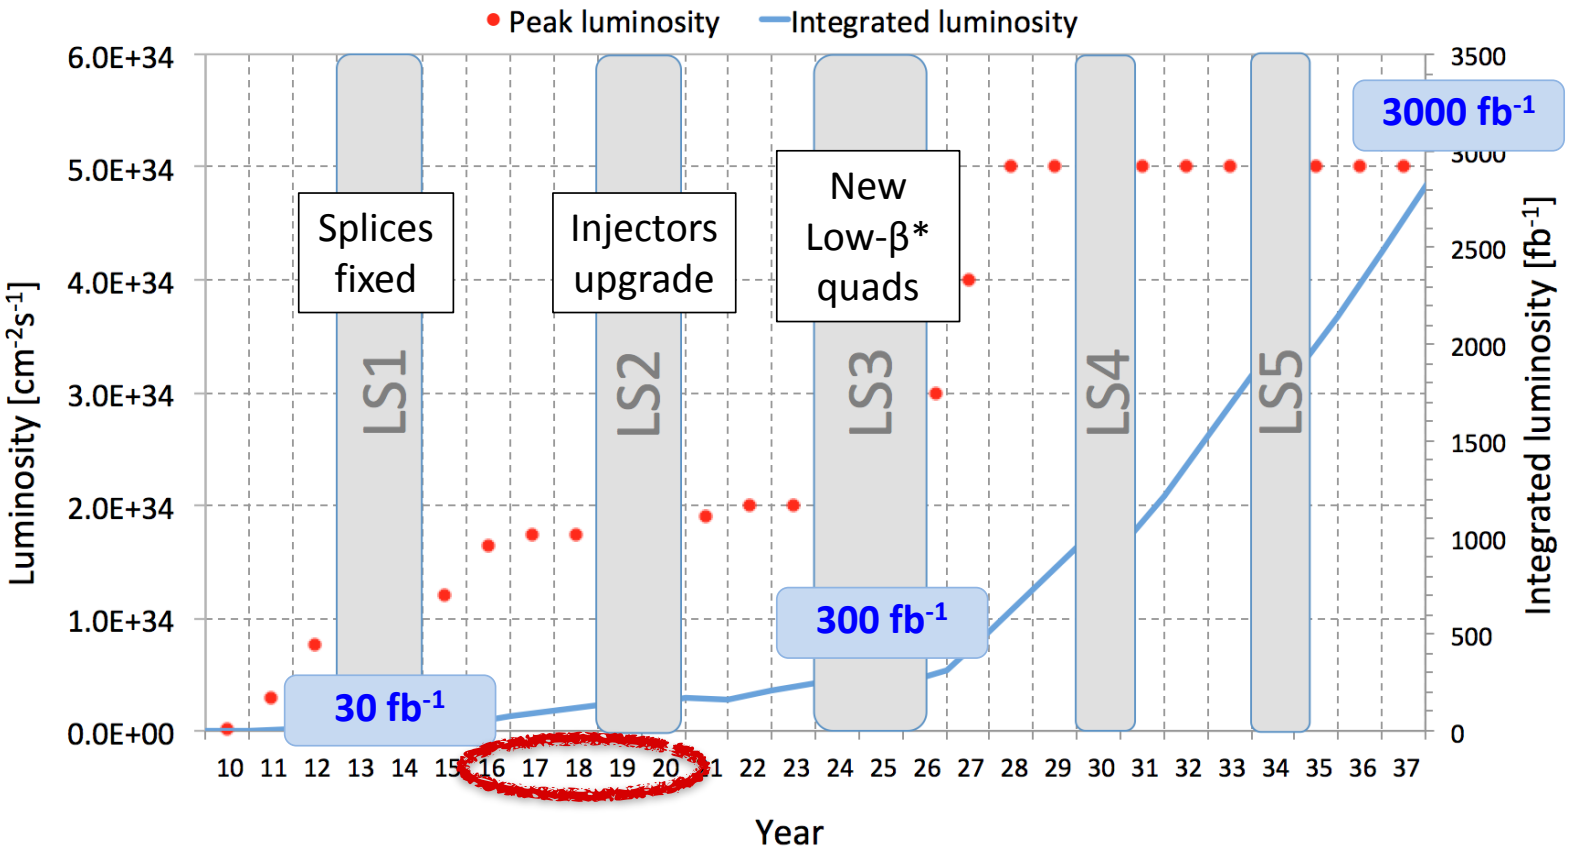

- The LHC will operate for at least 20 more years, with ever-growing datasets of ever-growing complexity, requiring growing computing resources and calling for continuous improvements in computing ingenuity!

Evolution on the horizon

- Original LHC computing models were very static
  - Known resource providers with well-understood computing tasks
  - Fixed amount of computing resources “pledged” to experiment
  - Resources themselves were very standardized
  - Activity of a given job largely contained within a given site
- Changing to a much more dynamic model
  - Much more fluidity in the tasks performed at a given site, so site capabilities must be more generic
  - Resources owned by the experiment might not be sufficient to support the physics needs, so need to explore other providers
  - Those “other providers” might run a greater diversity of computing architectures
  - Growing use of data federations will lead to greater activity across wide-area networks — the many distributed computing sites will act more and more like one single facility?

Where can U.S. CMS work with OSG?

- Access to a greater variety of resources outside the WLCG system
  - Resources where we may get a special allocation, like SDSC
  - Cloud-like resources, be they opportunistic or for rent
  - HPC architectures, such as those being made available to us by DOE ASCR (e.g. NERSC/ANL/ORNL)
- Elastic resources
  - Sites that are “ours” may want to provision resources differently
  - Fermilab HEPCloud project seeks to straightforwardly expand their resources by adding in e.g. cloud resources when it is favorable to do so to meet demand
    - Will be testing it this fall with expansion into AWS
  - Gives the T1 site more flexibility for how they provide resources
- Pilot job systems such as glideinWMS will enable CMS to take advantage of both of these cases
  - CMS relies on glideinWMS in many places and I expect that we will continue to do so; we invest some of our own effort in development

Where can U.S. CMS work with OSG?

- Individual universities are important players for U.S. CMS and have resources to use and perhaps share
  - The “Tier-3 sites”, resources controlled by and/or made available to individual small research groups for their own benefit
  - Can range from small clusters maintained by physicists, to portions of massive campus-wide centers
    - Could be guaranteed share, or opportunistic
- Want to enable campus physicists to take advantage of any resources made available to them
- CMS usage of campus resources can be a beachhead for OSG — a way to introduce university IT organizations to the power of DHTC and then add them to the OSG community

Outlook

- CMS remains an example of what the OSG was built to serve, and OSG has served it extremely well
  - Provides the underlying computing fabric that has enabled physics discoveries
- CMS has also made great contributions to OSG in the form of resources, expertise and innovation
- The next five years will probably see much greater fluidity and diversity in available computing resources and how they might be used by CMS
- Perhaps this presents a challenge to OSG, but the opportunities presented are surely even greater!