# CMS Offline and Computing Status and Plans

LHCC Referees Meeting 03.03.2015

Maria Girone, David Lange

# Outline

- Computing Model Changes in Ls1
- Usage of resources and previsions for Run2
- Run2 preparation, milestones, updated plans
- Computing for the Upgrade
- Outlook

# Introduction

- In Run2 CMS expects to increase the HLT output to ~1kHz
  - will promptly reconstruct more than twice as many events as the final year of Run2
  - expected pile-up will require twice as much processing time per event
  - The budget and evolution of technology permits less than doubling of the computing resources between 2012 and 2015
  - The main focus of Long Shutdown 1 has been finding ways to do more with less and to look for efficiency improvements in every system

# Computing Model Changes in Ls1

- Evolution out of the MONARC model had already started in Run1. The LS1 focused on
  - Additional Flexibility on the way resources are accessed
  - Improved Performance on the way resources are used
  - Optimized Access to data for analysis and production

- Instrumental for these changes have been the adoption of
  - Logical separation between the Disk and Tape storage systems
  - Dynamic Data Placement
  - Data Federation
  - One Central Condor Pool for all types of resources and applications
  - New Distributed Analysis Tool (CRAB3)

# Reducing Resource Needs

- Operational improvements
  - Reducing the number of expected reprocessing passes
  - Commissioning the use of the HLT farm for the main data reprocessing in winter shutdowns
  - Constraining the simulation budget and developing techniques to allow simulation reconstruction to be run on more resources
  - Distributing the prompt reconstruction between CERN and the Tier-1 centers

- Technical improvements
  - Reconstruction improvements and the development of a multi-core application
  - Commissioning of the multi-core queues at Tier-0 and Tier-1s
    - Multi-core decreases the number of processes that need to be tracked and reduces the overall operational load

# Blurring the Site Boundaries

- In Run2 CMS computing resources are intended to work more like a coherent system than a collection of sites with specialized functions
  - Improved networks have been key to this
- Data Federation will make CMS datasets available transparently across the Grid
- One central queue for all resources and all workflows
- The HLT farm is now an integrated resource when we are not in data-taking mode

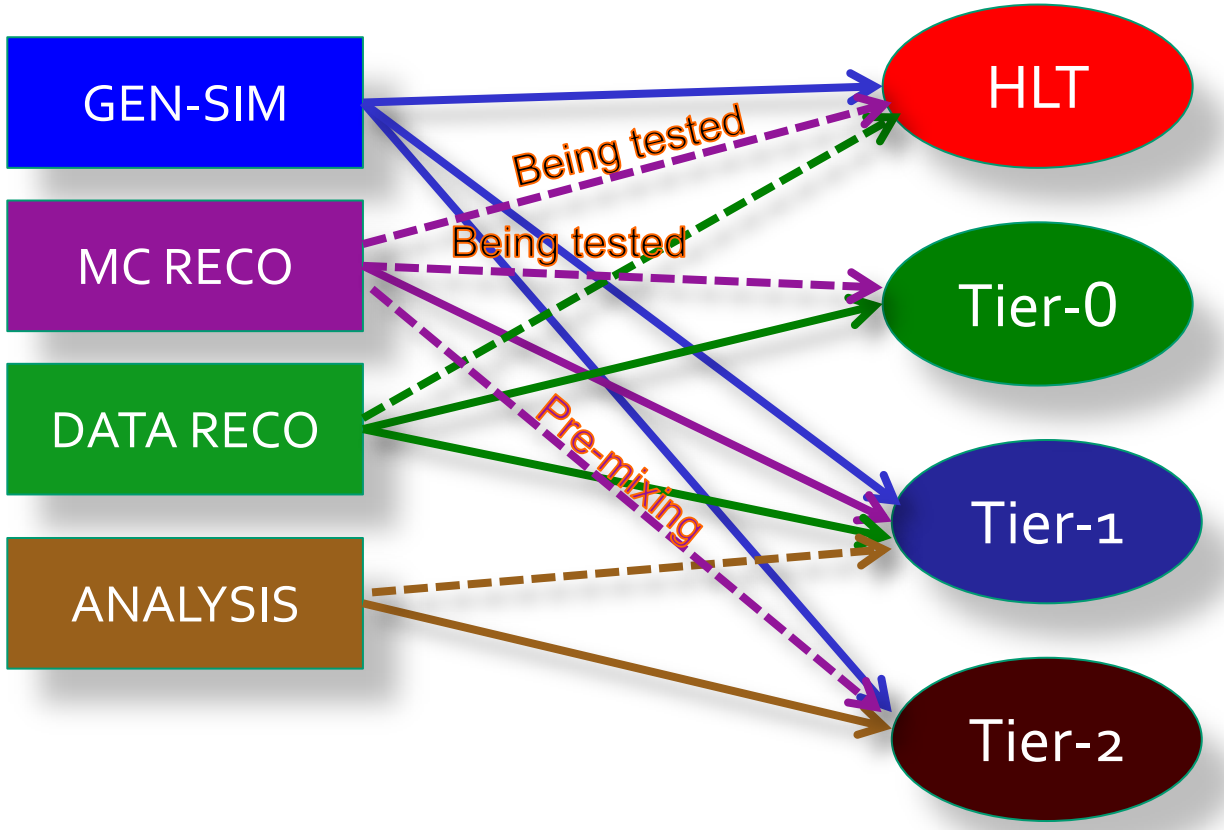

# CMS Resource Utilization

- CMS Resource utilization was high for most of LS1
  - There is a general drop after the beginning of the year. This is a software development period and Run2 simulation samples will be launched in April
  - Samples production for the Upgrade TP currently ongoing

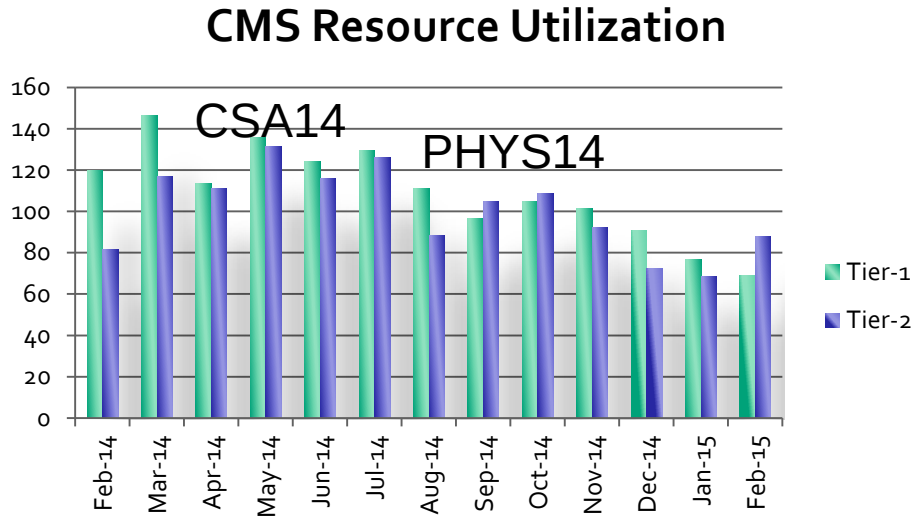


- The trend down is also in Tier-2 utilization
  - This is partially caused by fewer simulation requests, but also indicates that Run1 analysis activity is completing, which we relied on in the resource planning

# Resources in 2015

- We will increase the CPU by ~100% at Tier-0, by 70% at Tier-1, by 25% at Tier-2
  - About 50% of the Tier-1 resources will be dedicated to prompt-reco when the machine reaches the target instantaneous luminosity
- Disk space has smaller increases in the first year of Run 2
  - only a year of accumulated data and data federation will improve access
- Tape increases but less derived data is planned to be written to tape
  - RECO will be kept transiently on disk: typically ~4 months
- We are about 10% low in the initial pledges for Tier-1
  - We will review after deployments in April

| | 2014 | Increase from 2013 | 2015 | Increase from 2014 |
|---|---|---|---|---|
| Tier-0 CPU (kHS06) | 121 | 0% | 256 (256) | 111% |
| Tier-0 Disk (TB) | 7000 | 0% | 3250 (3000) | Reallocated to CAF |
| Tier-0 Tape (TB) | 26000 | 0% | 31000 (31000) | 31% |
| CAF CPU (kHS06) | 0 | 0% | 15 (15) | - |
| CAF Disk (TB) | 0 | 0% | 12100 (12000) | - |
| CAF Tape (TB) | 0 | 0% | 4000 (4000) | - |
| T1 CPU (kHS06) | 175 | 0% | 300 (300) | 71% |
| T1 Disk (TB) | 26000 | 0% | 27000 (26000) | 4% |
| T1 Tape (TB) | 55000 | 11% | 73500 (74000) | 34% |
| T2 CPU (kHS06) | 390 | 14% | 500 (500) | 25% |
| T2 Disk (TB) | 27000 | 4% | 31400 (29000) | 16% |

# Work in LS1

- Commissioning of the Higher Level Trigger Farm for Production
- Refactoring the Tier-0 for Distributed Prompt Reconstruction (will include 50% of the Tier-1s resources)
- Multi-Core Commissioning
- Software Improvements
- Improvements in Data Access
- Improvements in Data Management
- Work toward Opportunistic Computing

# The HLT Farm

- An addition for Run2 is the use of the High Level Trigger (HLT) farm for offline processing
  - It is a large computing resource (15k cores) that is similar in size to the Tier-0 in terms of number of cores, but we cannot reach this scale until March
  - Successfully interfaced using cloud computing tools. It is similar to the Tier-0 AI
- In 2014 the network link P5 to the computing center was upgraded from 20 to 60Gb/s
  - Larger than needed for data taking but necessary to access the storage in the computing center for simulation reconstruction
  - **Will be upgraded to 120Gb/s before the 2015 run starts**
- Production workflows have been commissioned including the HI reprocessing, Gen-Sim, and Simulation reconstruction
  - All access to data is through the data federation and primarily served from CERN


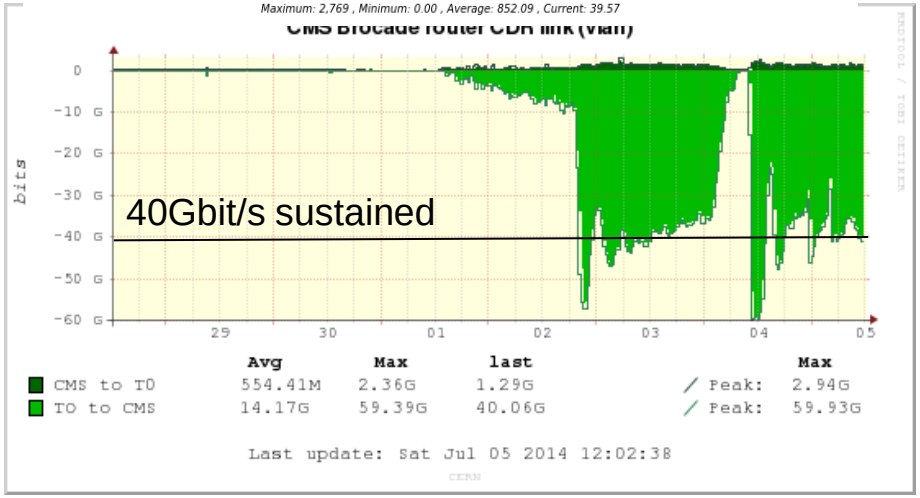

# Tier-0 Commissioning

- The Tier-0 submission infrastructure was reworked over LS1
  - The capability of distributing Prompt Reco was added to allow the Tier-1 centers to contribute (with up to 50% of resources)
  - The Tier-0 processing predominately come from the CERN Agile Infrastructure with direct resource provisioning through Condor
    - CMS was an early adopter of CERN AI
  - The Tier-0 infrastructure was implemented using same components as the rest of the computing reprocessing system (using GlideInWMS)
    - This reduces the software maintenance of the system and reduces the overall operations load

- A global scale test will be performed as soon as the 2015 resources are deployed - April 2015 - we some worries on this late schedule

# Multi-Core Commissioning

## Multi-threaded CMSSW Framework

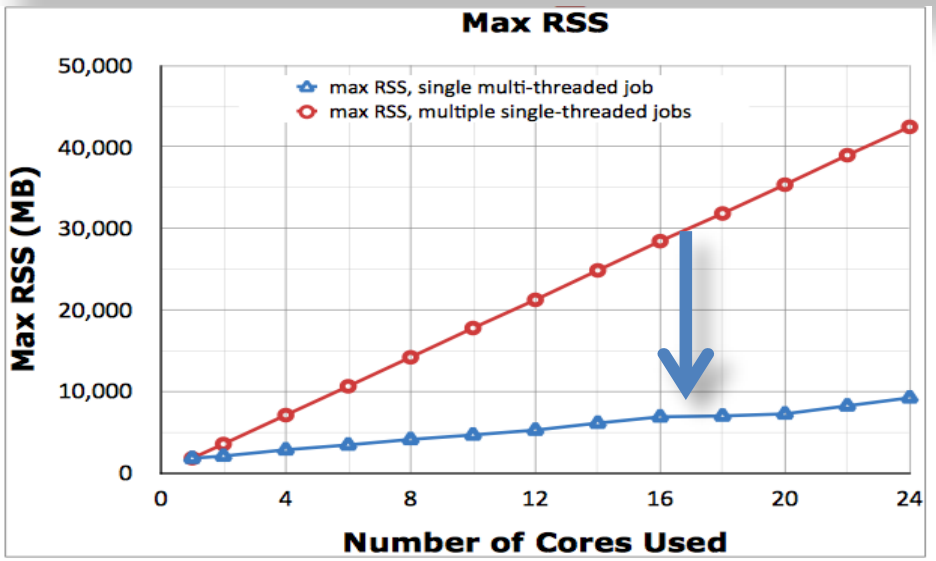


- Developed next-generation framework for CMSSW based on a multi-threading approach
- This Framework allows us to process higher Run 2 trigger rates efficiently and to adapt to computing technology trends
- Observe substantial application memory savings in CMS reconstruction 0.35 GB per additional thread instead of 1.8GB/job (targeting >=4 cores in 2015)

## Multi-threaded Resources

- We progressed on the commissioning of multi-core queues at Tier1s
  - **This was done in collaboration with WLCG and the other LHC experiments**
  - Functionality fully validated BUT a complete scale test will only be possible in April, when resources at CERN and Wigner will become available
  - A Central Condor Pool has also been commissioned for all types of resources and applications, allowing easier prioritization of analysis and production workflows

# Finalizing Software Release for Run 2 Startup

## Reconstruction improvements

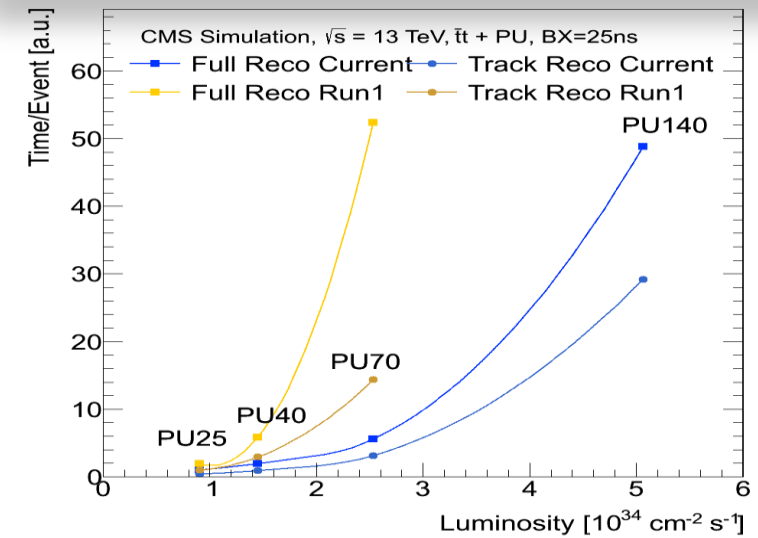


- Substantial collaboration effort to develop robust algorithms for 25 ns bunch spacing and Run 2 conditions
- Meanwhile, we achieved needed gains in technical performance to keep up with increased trigger rate and event complexity in Run 2 given pledged computing resources

## Simulation improvements

1. New pileup simulation improves physics and technical performance
   - Memory and I/O use of pileup simulation on GRID drastically reduced to allow Run 2 and Phase-II upgrade simulations
2. Technical performance improvements have gained 50% in CPU since Run 1 without loss of accuracy in physics simulation
1. Upgraded to Geant4.10, which was critical for use of multi-threaded Framework

Major MC data production will start in April 2015 (@ 50ns and 25ns)

# Improvements in Data Access

- Any Data, Anytime, Anywhere (AAA) has been a primary focus area in 2014
  - CERN, all Tier-1s,most of the Tier-2 sites serve data in the federation
    - More than 90% of all current CMS data should be accessible
    - Sufficient I/O capacity to provide 20% of the total access (~100TB/day)
    - Enable a system of hierarchical redirectors to maintain access within geographic regions when possible
  - Nearly all sites are configured to use the federation to access samples, if they aren't available locally

- CMS has developed a new user analysis data format (MiniAOD) that is 5-10x smaller than the Run 1 format
  - Design and content based on Run 1 experience across analysis groups
  - Targets most analysis needs (70-80%)
- Potential for big analysis improvements in Run 2:
  - Increased analysis agility in limited computing resources: Recreating miniAOD is much faster than rerunning the reconstruction for a vast range of performance improvements

- Optimization of the IO has been an ongoing activity for several years, which has paid off in high CPU efficiency over the wide area
- Big push in 2014 to commission sites to measure IO and file open rates to understand the sustainable load and to deploy and use advanced monitoring

# Creation of the Transitional Federation

- The addition of the production workflow puts additional constraints on the required reliability of the Federation

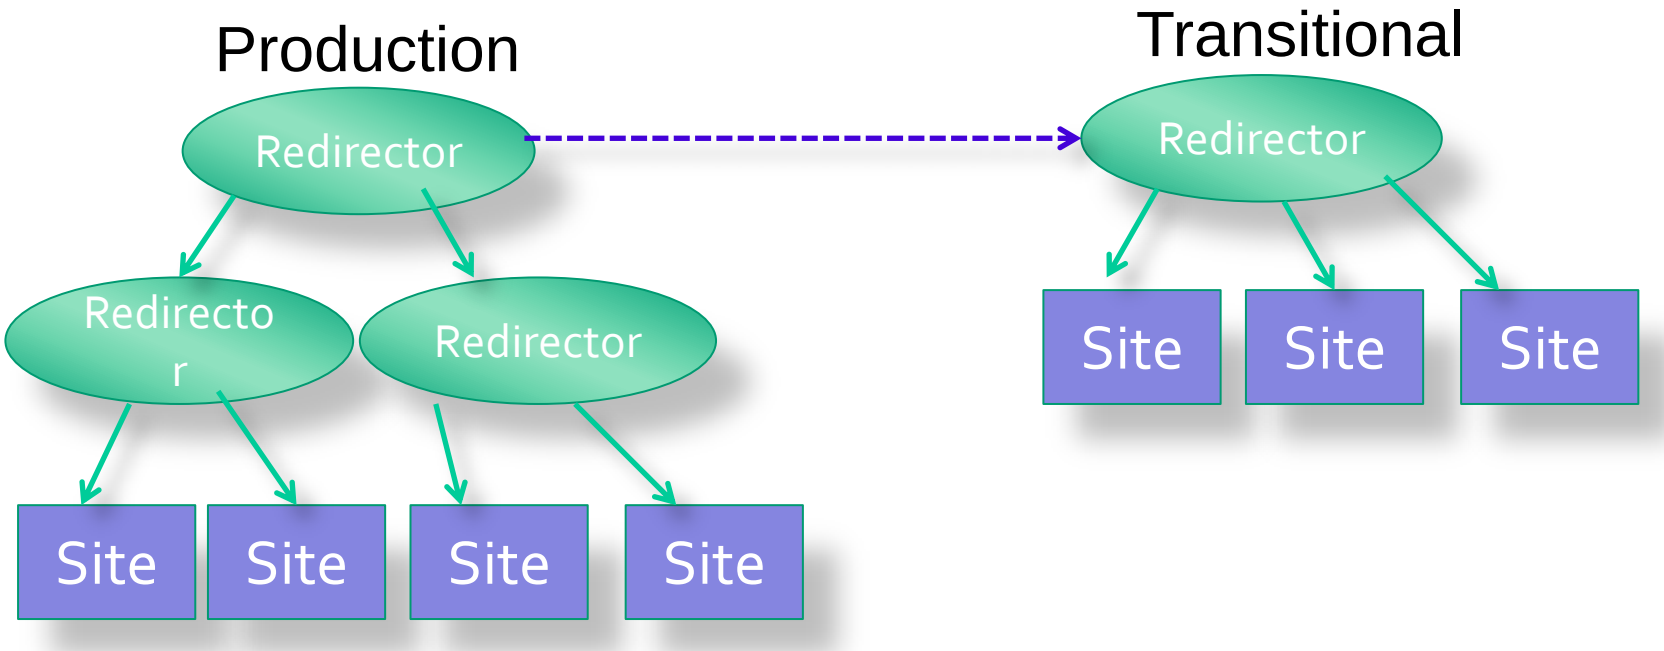


- Validated sites are in the production federation, sites being commissioned are in an independent federation and only when a sample cannot be found in production are they used
- **This new concept has been a result of a collaborative effort between ATLAS and CMS federation developers**

# Improvements in Data Management

- In addition to work on data federation we have tried to improve our traditional data placement and access
- The use of samples is continuously monitored through the data popularity
  - Number of replicas depends on the popularity of the datasets
- Samples will be replicated dynamically if the load is high and replicas removed if they are not accessed for a period of time
- Samples are only deleted when there is new data to replicate, disks are kept full
  - "*0 old*", shows the volume data that were last accessed prior to the period covered by the plot; the "*0*" bin stands for no access in the period selected
  - The zero bin includes un-accessed replicas and datasets that have only **one** copy on disk

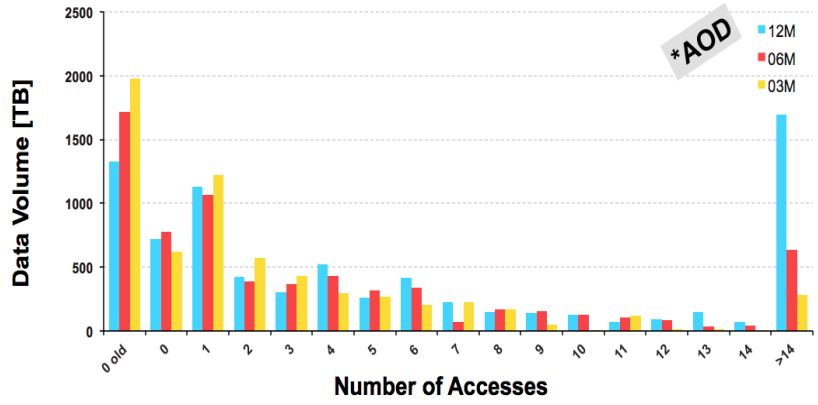


- As of today, the system has triggered the deletion of roughly 5 PB of the least popular samples residing at Tier-2 sites, and it is being enabled at Tier-1s

# Towards Opportunistic Computing

- CMS has almost no contingency to deal with something unexpected or having to redo something that goes wrong
  - We would need to make painful choices or find new resources
    - Buying new dedicated computing is probably unaffordable, but borrowing opportunistic computing or renting from commercial providers is a possibility
- We were awarded an allotment of 3000 processor cores on Gordon at SDSC for a month
  - This is a good exercise for demonstrating the deployment of the environment, but is a 5% increase in the Tier-2 resources for that month, so not a scale test for submission

# Resources for Contingency

- If we want to introduce contingency to solve a problem, we need to be able to double a class of computing for the period of a month
  - This is only globally ~10% of the yearly resources, but requires that the submission system scale to a factor of 2-3 for extended bursts
  - Demonstrating improved scaling is useful also for growth
    - Dedicated Computing resources will continue to grow in Run2

- Doubling a class of resources for a month would be adding 40-60k cores
  - ATLAS just received a grant from Amazon Web Services for 80-160k cores for 1 month on the spot market, and 1PB of space
  - Building on the ATLAS findings, CMS would like to apply for a similar grant and work in collaboration with them on a common testing program

- We are examining real workflows that could be good for commissioning
  - A difficult test like premixing sample production would stress all parts of the system and would free up Tier-1 resources

# 2016 Request

- Tier-0 + Tier 1: evolution with the expected live time
- Tier-0 tape: ~25 %
- Tier-2 (CPU & Disk) ~ 40%

| | 2014 | Increase from 2013 | 2015 | Increase from 2014 | 2016 (C-RSG Oct 14) | Increase from 2015 |
|---|---|---|---|---|---|---|
| Tier-0 CPU (kHS06) | 121 | 0% | 256 | 111% | 302 | 18% |
| Tier-0 Disk (TB) | 7000 | 0% | 3000 | Reallocated to CAF | 3250 | 0% |
| Tier-0 Tape (TB) | 26000 | 0% | 31000 | 31% | 38000 | 23% |
| CAF CPU (kHS06) | 0 | 0% | 15 | - | 15 | 17% |
| CAF Disk (TB) | 0 | 0% | 12000 | - | 13100 | 9% |
| CAF Tape (TB) | 0 | 0% | 4000 | - | 6000 | 50% |
| T1 CPU (kHS06) | 175 | 0% | 300 | 71% | 400 | 33% |
| T1 Disk (TB) | 26000 | 0% | 26000 | 4% | 35000 (33000) | 30% |
| T1 Tape (TB) | 55000 | 11% | 74000 | 34% | 100000 | 35% |
| T2 CPU (kHS06) | 390 | 14% | 500 | 25% | 700 | 40% |
| T2 Disk (TB) | 27000 | 4% | 29000 | 16% | 40000 (38000) | 37% |

# Upgrade Activities

- A lot of Computing and Offline upgrade activity has been focused on simulation for detector studies. Computing and Offline have been fully engaged with software development and samples preparations for the TP Upgrade 1 and 2

- But there is also a developing program of R&D and upgrade activities (Computing model evolution to HL-LHC and evolutionary changes in S&C) , open to outside CMS
  - Tracking (reco) on heterogeneous many-core resources, on heterogeneous worldwide resources
    - We have ported the CMS track reconstruction to the Intel Phi
  - Trigger on many-core resources, potentially homogeneous at P5, heterogeneous elsewhere
  - Simulation (G4 or beyond) on heterogeneous resources
  - Resource scheduling, I/O, event processing frameworks for heterogeneous resources, parallel processing models
    - Developing a Computing Model simulation to study ways to optimize data access
  - Tracking and other reconstruction algorithms at high pile-up
    - Continuous improvement activities
  - Tools and infrastructure (profilers, development tools and models, etc.)
    - Prototyping container technology (e.g., DOCKER) for CMSSW
  - More focus on power use, in addition to raw performance

# Outlook

- A lot of work has been done in LS1
  - We have a functional data federation, which we expect to declare production ready before the start of the run
  - We have much better flexibility in where workflows can run
  - We have reworked the Tier-0 infrastructure for distributed prompt-reconstruction
  - We deployed a central Condor pool for easier prioritization of analysis and production workflows
  - We have implemented a new distributed analysis tool with faster client, improved automatic resubmissions and centralized handling of the user outputs
  - We have improved reconstruction and simulation software and a more efficient analysis data format

- We needed to make big increases in operational efficiency to survive the conditions in Run2 with the resources that could be afforded
  - There is not a lot of contingency

- We believe CMS Computing and Offline is **ready** for start up
  - We need to execute a full scale test of the Tier-0 distributed prompt reco

- We expect to learn a lot in a short period of time about operating the new services during the beginning of the run


# BACKUP SLIDES

# What We Produced in 2014

- Considering the CSA14 production
  - Fall13GS (CMSSW6.2) and Spring14DR (CMSSW7.0)

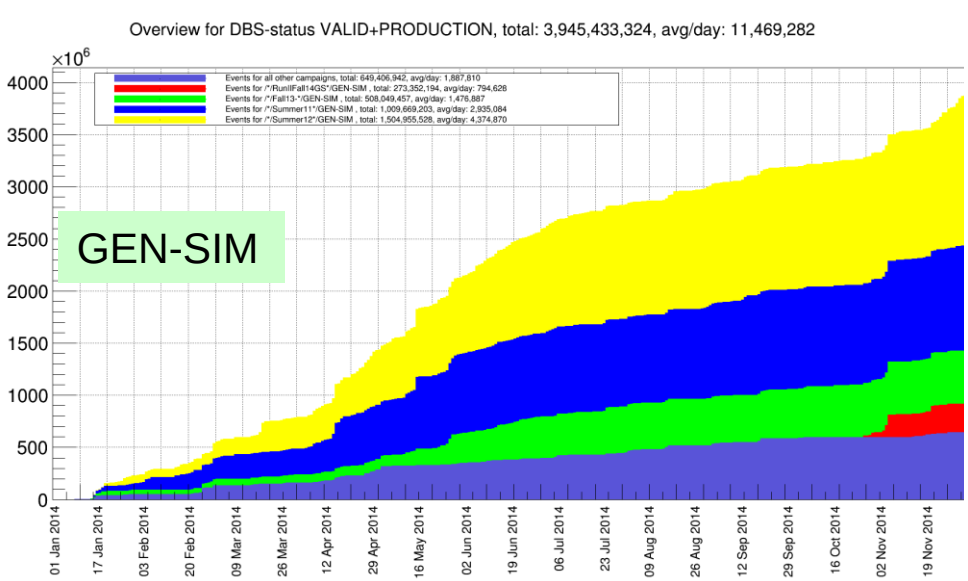


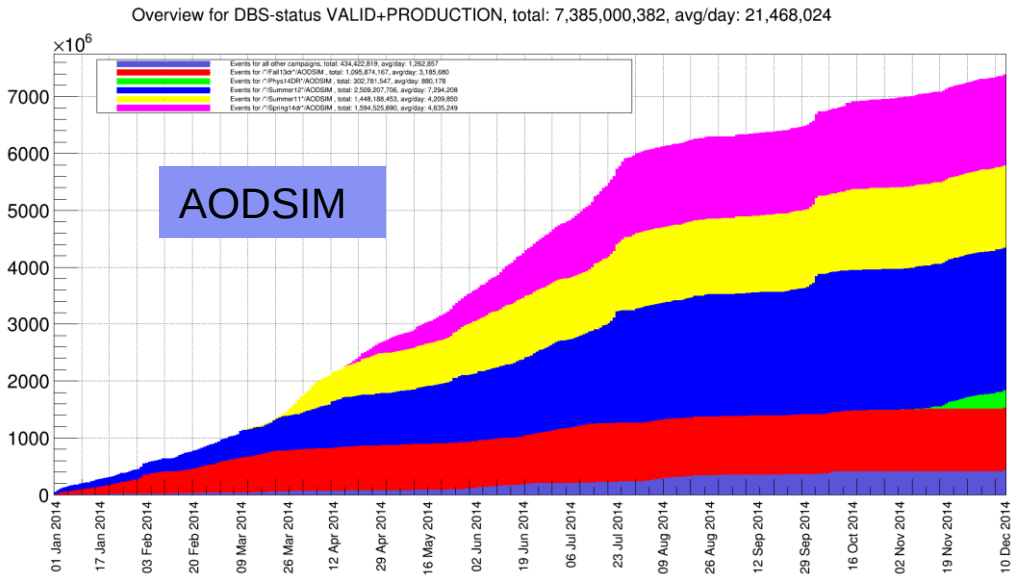


- During summer we demonstrated we could sustain: 500M/month Gen-Sim and 1B/month Digi-Reco
- We should expect event time simulation improvements with next release