


# **CMS Status**


# Resource Utilization

- ▶ CMS resource utilization is still low compared to 2012. We are still waiting on final configurations for 2011 legacy simulation
- ▶ CMS is planning a physics validation activity (CSA14)
  - ▶ Involves 2B events (~700M gen-sim events, with 3 pile-up scenarios)
  - ▶ Challenging reconstruction scenarios and an opportunity to look at physics analysis at this pile-up
  - ▶ Other upgrade samples are also needed at this time

**CMS Tier-1 Pledge Usage**



**CMS Tier-2 Pledge Usage**



# HLT Progress

- ▶ HLT declared production ready and processed first data reconstruction workflows to completion
- ▶ Only cluster phase (2) and (3) are usable for offline
- ▶ Minimum plan is a core count replacement of cluster (1) for 2015. Unfunded desired upgrade plan is a box count replacement

| cluster     | Nodes | cores (HT on)/<br>node | cores | Memory<br>(Gbyte/node) | Disk<br>(Gbytes/node) |
|-------------|-------|------------------------|-------|------------------------|-----------------------|
| (1)         | 720   | 8                      | 5760  | 16                     | 72                    |
| (2)         | 288   | 12 (24)                | 3456  | 24                     | 225                   |
| (3)         | 256   | 16 (32)                | 4096  | 32                     | 451                   |
| (1)+(2)+(3) | 1264  |                        | 13312 | 26 Tbytes              | 227 Tbytes            |

- ▶ HLT Network bottleneck observed at 20Gb/s for 6k running data reco jobs
- ▶ Negotiate for upgrade by 2015
- ▶ Use of HLT in the CMS planning for resources for 2015, so network upgrade is critical. Planned for end of year reconstruction and possibly interfill simulation.
- ▶ We believe the interfill could provide contingency that is scrubbed from the current planning
- ▶ Most of the current development activities are concentrated in reducing start-up latency and checkpointing



# Reminder of the 2015 Planning (March)

- ▶ Pile-up is expected to grow with the increased instantaneous luminosity
- ▶ Roughly a factor of 2.5 increase in the reconstruction time with the best code we currently have
- ▶ CMS currently estimates the change of the rates at similar thresholds to 2012 of 800Hz-1.2kHz of prompt reconstruction rate
- ▶ This is just the core samples, and would be another factor of 2.5
- ▶ Machine will move from 50ns running to 25ns running
- ▶ Unexpectedly this increases the reconstruction time by roughly a factor of 2
- ▶ Combination of all these effects is a factor of 12

# Progress to Mitigation

- ▶ We assume the reconstruction speed of going to 25ns can be solved
- ▶ Charge cuts look very promising, but no final validated release. Offline is confident they will get to an average 50s reconstruction per event (60% improvement and a factor 2.5 worse than the end of 2012)
- ▶ We assume that in the first year we will move about half the prompt reconstruction to Tier-I
- ▶ We commissioned this workflow for 2012, but didn't need to use it. We will push the scale during 2014 as regular replay tests
  - ▶ Tier-I capacity reduces the Tier-0 by a factor of 2 (Factor of 6  $\rightarrow$  3)
  - ▶ Tier-0 increase in 2015 of  $\sim 100\%$  (Factor of 3  $\rightarrow$  1.5)
  - ▶ Remaining Factor of 50% is taking with an expectation of less time in collision running on Tier-0 (Factor 1.5  $\rightarrow$  1)

# Mitigation Continued

- ▶ We assume that the operations model in 2015 will be more like 2012 than 2010
- ▶ More organized reprocessing, and the Tier-1s will reconstruct MC and Prompt data primarily and reprocess data in a big organized pass at the end of the year making heavy use of the HLT (HLT commissioning ongoing).
  - ▶ Some drills planned in 2014 for software and request validations
- ▶ Most of the remaining Long Shutdown activities are designed to make more flexible use of the centers and allow us to reduce pressure in the system by adding computing resources to workflows as needed
- ▶ We can gain some contingency by having maximum flexibility where we run

- ▶ Disk/Tape Separation is expected by the end of the year
  - ▶ Allows us to open Tier-1s for analysis and allows disk to be used by multiple sites (for production and analysis workflows)
- ▶ Data Federation is about 80% deployed on schedule for March completion
  - ▶ Allows Tier-2 sites to participate in simulation (or data) reprocessing, allows multiple sites to process a single workflow, and provides more transparent access to the data for analysis users
- ▶ Dynamic Data Placement makes more efficient use of primarily Tier-2 disk
- ▶ Opportunistic Computing deployment gains contingency by adding computing resources

- ▶ Is a demonstration of specific technical capabilities in Computing and Software and a general activity to work on Analysis for 2015
- ▶ Tier-0 distributed prompt reco, the CERN agile infrastructure, and the stability of the Tier-0 tools
  - ▶ Regular monthly replay tests. Working with new resources and interfaces
- ▶ Production of reprocessed datasets
  - ▶ Latency was a bigger issue in 2010 than data volume. Demonstrate essentially any workflow on any CPU resource
- ▶ Data placement and cache release and wide area access to data
  - ▶ Use information from Data popularity and cleaning tools in production since 2011
  - ▶ Includes development actions in 2014
- ▶ Transition to multi-core scheduling
- ▶ Aim for large scale analysis activity in the late summer to early fall



- ▶ Long Shutdown projects are proceeding but several are un(der)staffed
- ▶ Goal is to begin integration testing and hardening activities by March 2014 with the idea of having a year to commission
  - ▶ Some choices are intentionally conservative approaches to ensure a stable system for the beginning of Run 2
- ▶ We have had to prioritize on specific tasks for 2015, and we risk longer term upgrade activities.