Ian Bird

LHCC Referees meeting

CERN; 22nd September 2015

WLCG Report

LHC Running conditions

The outline LHC schedule out to 2035 presented by Frederick Bordry to the SPC and FC June 2015 can be found here

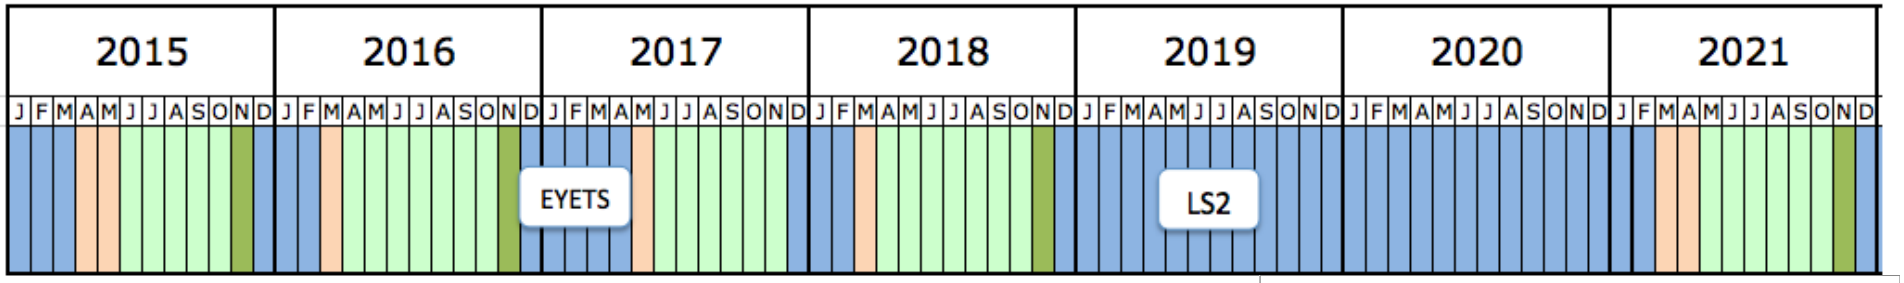

- LHC estimates of live time for 2017 reduced wrt previous
  - ➔ 5 M secs live time for pp in 2016 and 2017
  - Has changed resource estimates
- Less data in 2015, LHCb reduced 2016 requirements
- Likely to reflect in updated pledges

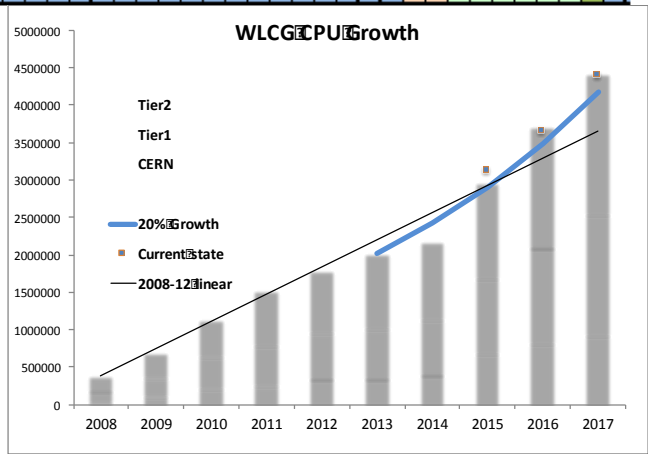

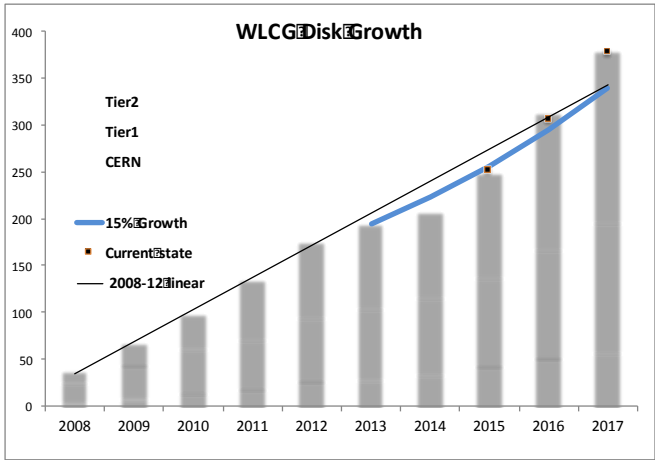

Current state of 2016 pledges

- Based on April RRB approved requests, and pledges as updated in preparation for October RRB (still some updates missing, biggest missing is Russia)
  - May change following RRB

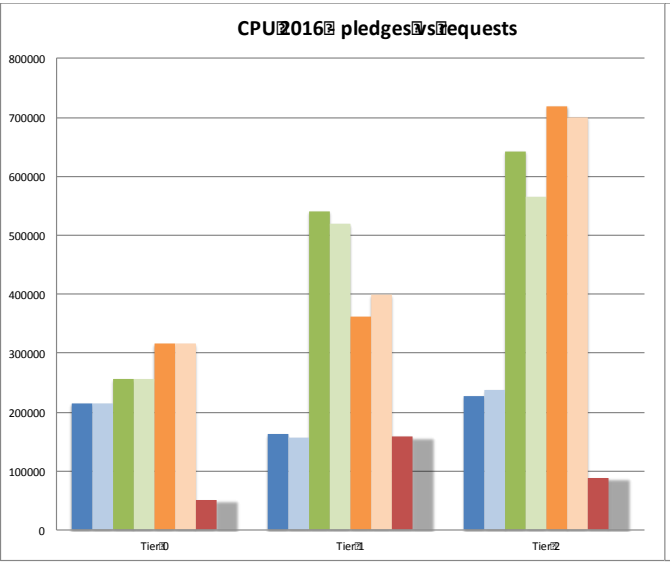

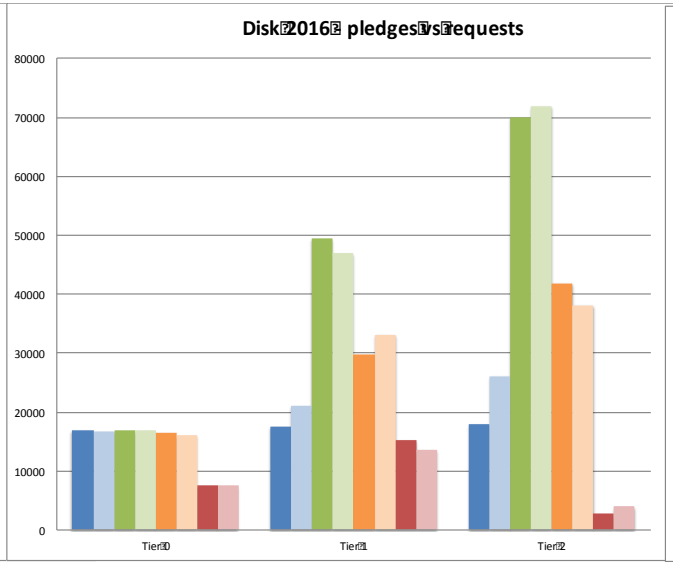

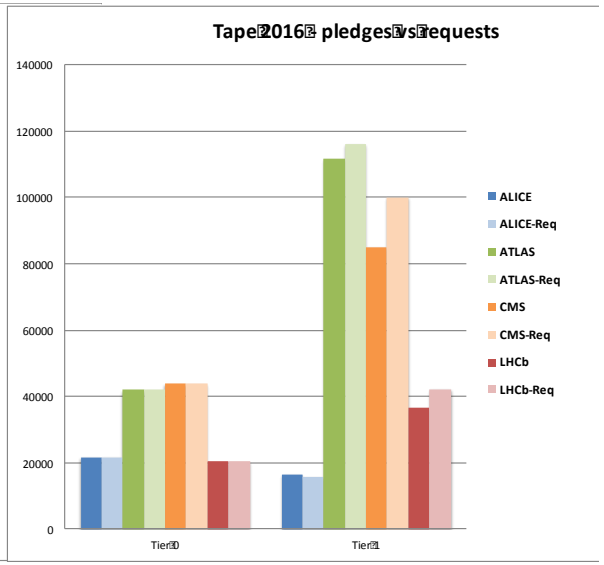

Status of Tier-0 Resources (1)

- Pledge fulfilment
  - Disk and tape provided early in 2015 as pledged
  - CPU slightly late (firmware issue), provided as pledged in July
- General concerns:
  - Capacity provided as virtual machines in CERN private cloud; currently performance hit of ~10%
    - Already compensated for regarding installed capacity
    - On-going work on further optimisations (~5% or less)
  - With settings deployed on part of capacity, worker nodes can get into state of very slow processing
    - Initially reported by LHCb
    - Reason found, will be fixed with on-going optimisations

Status of Tier-0 Resources (2)

- ALICE concerns
  - Same class of jobs 15…20% less efficient (CPU/wall) than elsewhere – under investigation
  - Sporadic high I/O wait of worker nodes – under investigation
- ATLAS concerns
  - Rocky start of dedicated LSF Tier-0 instance; okay as of mid-July
  - Changed ATLAS Tier-0 workflow means that worker nodes can only be used ~50% sometimes
- CMS concerns
  - Slow Tier-0 nodes due to memory over-commitment: in process of being resolved
- For 2016 we will increase installed capacity (~10%) above pledges to compensate
  - Also planning to test deploying containers or “bare-metal” where performance should not be penalised
- NB: the use of openstack helps to manage the much larger Tier 0
  - Size has doubled wrt 2012, and significant increase again for 2016

VM performance

- Non-uniform memory architectures in modern CPUs
- Problem arises if hypervisor allocates virtual cpu cycles on different physical cores
- Effect is worse for larger VM's
- Pinning and tuning can reduce effect to few %

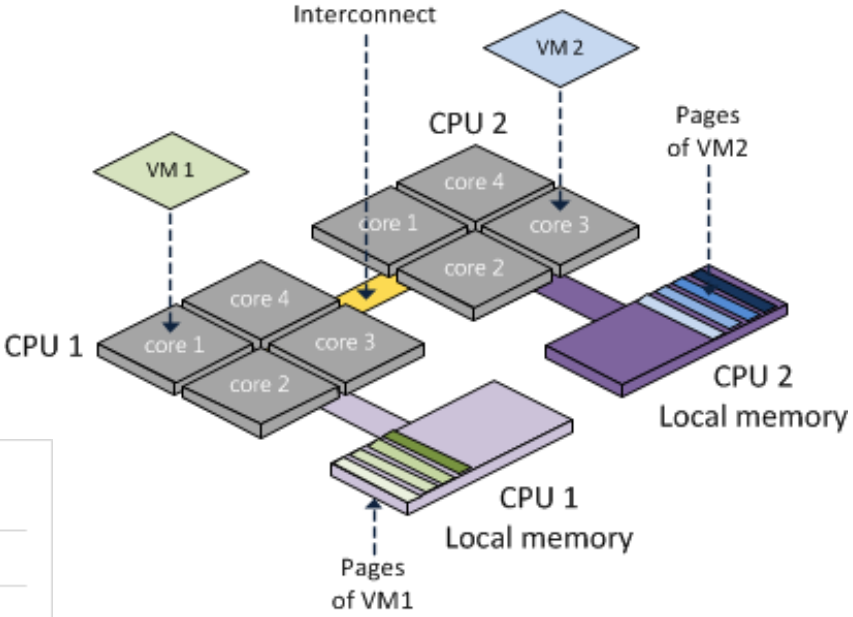

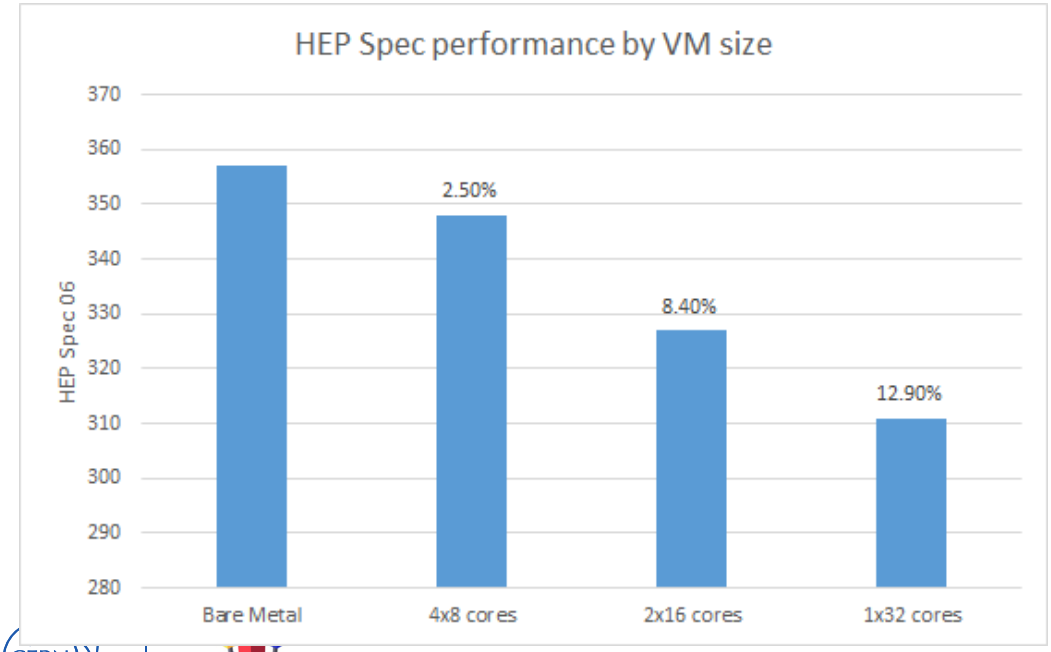

Upgraded connectivity P5 → Tier 0

- 10 Gb/s connection was limiting CMS ability for Run 2 and in being able to make full use of HLT
- Now upgraded to 100 Gb/s – data rates of 800 MB/s achieved

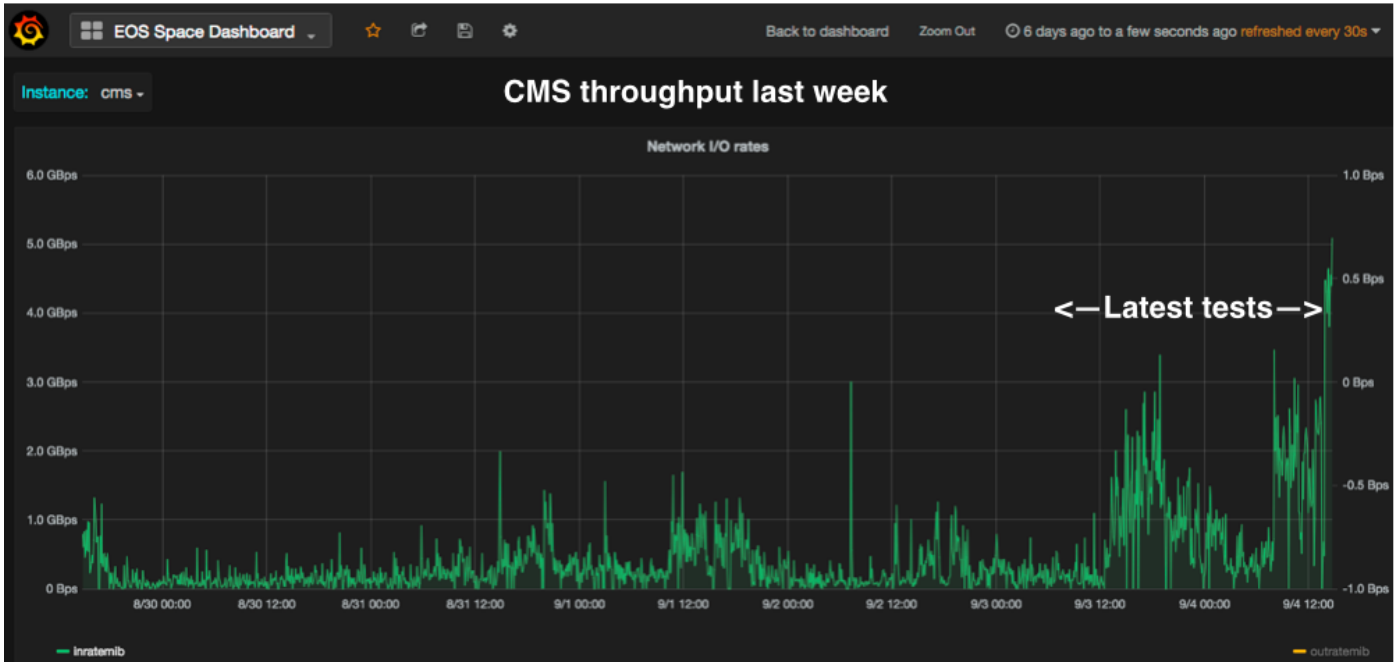

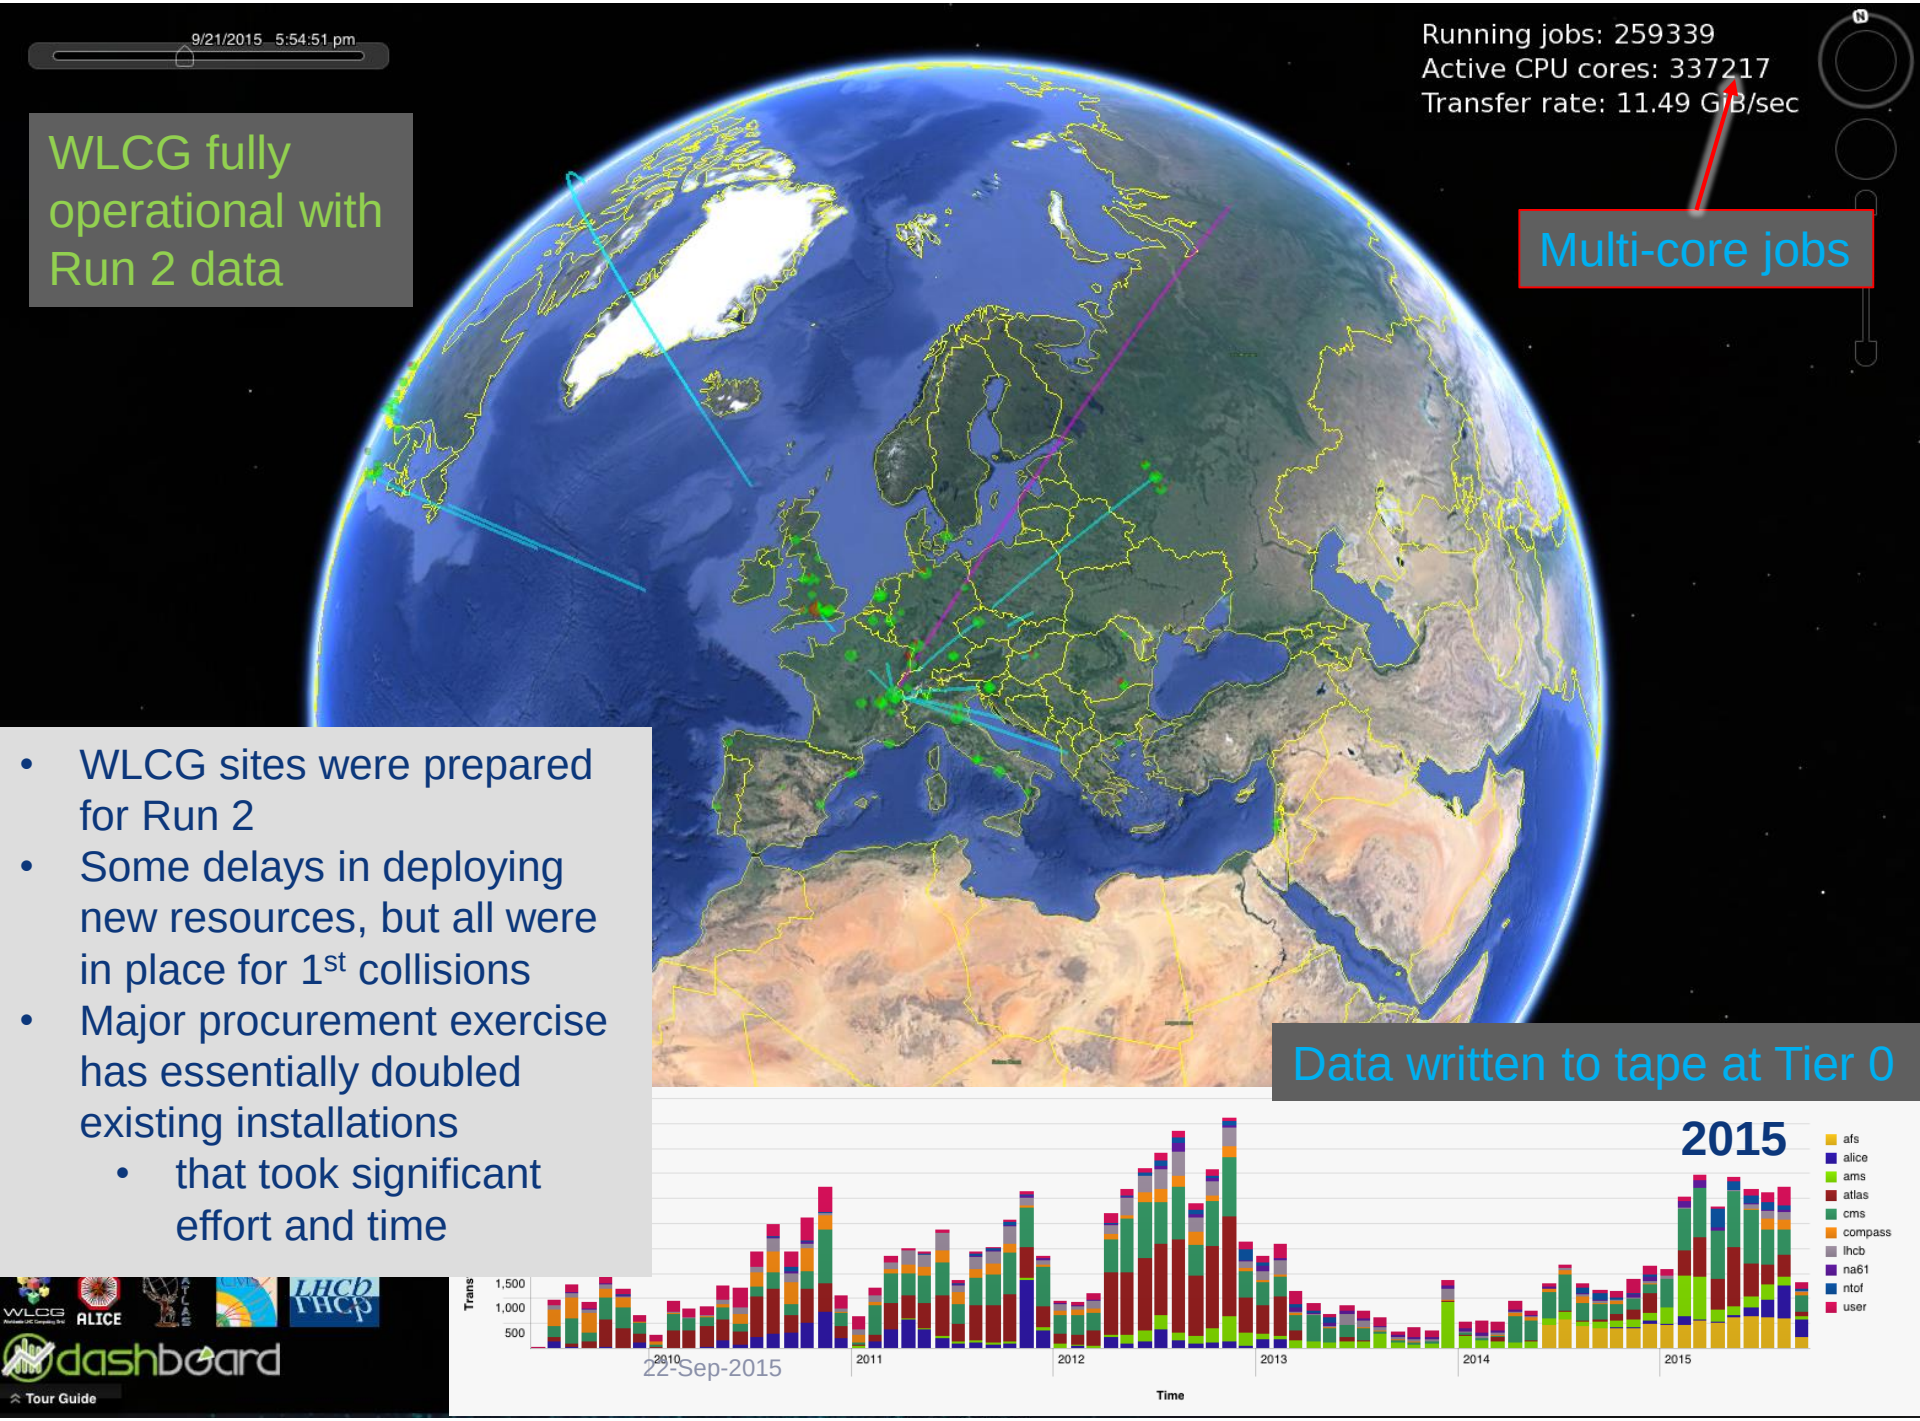

- WLCG sites were prepared for Run 2
- Some delays in deploying new resources, but all were in place for 1st collisions
- Major procurement exercise has essentially doubled existing installations
  - that took significant effort and time

Longer term planning

- Putting in place a “WLCG Technical Forum” to explore possible computing models for the HI-LHC era
  - Initial action will be to document some outline ideas for investigation
  - WLCG workshop in Feb 2016 is an opportunity for further input on this
  - In parallel look at a number of evolutionary topics
- Must be done in parallel with ramping up effort on software improvements (HSF)

Science Cloud

- PCP project has been approved and will start in January 2016
  - Derogation of CERN procurement rules to allow compliance with EC procurement (EC member states) agreed by CERN Council last week
  - ➔ see next few slides

- Independent of the PCP project, CERN is doing 2 procurements to understand market prices and usage models:
  - 200 KCHF for compute (essentially simulation)
  - 400 KCHF for data intensive use cases (needs market survey)
  - To be procured as extensions of the CERN cloud
  - Eventual use cases – investigate alternatives to “Wigner-style” extensions for next round; dynamic expansion of CERN resources
  - This capacity is in addition to pledges

Open Science Cloud and the EC

Commissioner Moedas accompanied by CERN Director-General R.Heuer at the LHC facility © CERN

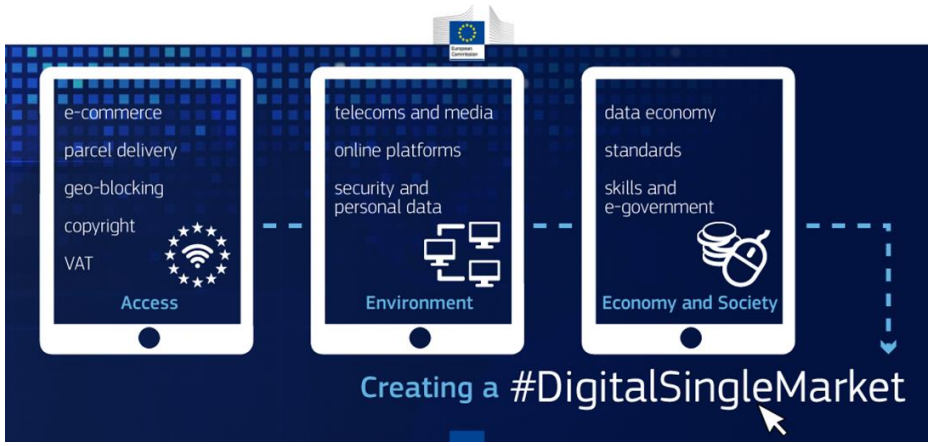

Questions posed by RJ Smits during visit to CERN with Commissioner Moedas in January 2015:

1. Why should Europe develop its own European Research and Science Cloud?
2. What would it take from a technical point of view to set up such a cloud?
3. What would be the range of services that could be offered?
4. What would be the time frame?
5. What would be the costs?

CERN produced a paper *Towards the European Open Science Cloud* that answered these questions **http://dx.doi.org/10.5281/zenodo.16140**

Commissioner Moedas announced in June 2015:

*"The Commission will launch a European Cloud initiative including cloud services certification, contracts, switching of cloud services providers and a **research open science cloud**."*

The draft H2020 Work Programme 2016 – 2017 (annex 4) includes funding call INFRADEV-04-2016 *European Open Science Cloud for Research* (10M€)

HNSciCloud H2020 PCP Project

The group of buyers have committed

- ~1.6M€ of funds (generating ~6M€ total funds)
- Manpower
- Applications & Data
- In-house IT resources

To procure innovative IaaS cloud services integrated into a hybrid cloud model

- Commercial cloud services
- European e-Infrastructures
- In-house IT resources

Procured services will be made available to end-users from many research communities