



# WLCG Data

preGDB 11 April  
Oliver Keeble


Today

# WLCG sites




# Pledged Resources


CPU by computer centre  
Total ~400k cores




Tape by computer centre  
Total ~400 PB



Disk by computer centre  
Total ~300 PB



# Workflows



More than half the CPU goes on simulation.

Most of the rest is reconstruction.

The remainder is analysis.

## Different probes/methods/specifications

WLCG is here!




Slide from G. Lamanna, 1<sup>st</sup> OBLICS workshop


Projects	Data Processing	Main requirements/challenges
EVENT-BASED ( $\gamma$ -rays, CR, $\nu$ ) <u>CTA, KM3Net...</u>	Evt-builder, calib. and reconstruction; reduction, real-time science.	Raw big-data. Data formats. Algorithms. On-site operation and reduction. Cooperative science tools. Observatory (A&A). Multi- $\lambda$ . [...]
IMAGE-BASED (far-IR, VIS) <u>EUCLID, LSST...</u>	Surveys/deep observation; combining photometer and spectrograph info.; Catalogue of objects.	Big-data products: data base challenges. Graphical processing, Algorithms. Images format. Catalogue preservation and query. A&A. [...]
SIGNAL-BASED (Radio, GW) <u>SKA, LIGO-Virgo...</u>	Noise cleaning; time-series, mathematical processing (FT) converting signal in images.	Algorithms. New computing architectures and data centres. Combination of HPC and HTC. Fast soft reduction. Data mining and preservation. A&A [...]

# LHC: “outstanding performance”

Disk





Tape



Estimated: Estimates made in 2014 for Run 2 up to 2017

20%: Growth of 20%/yr starting in 2016 (“flat budget”)


# 2016 data volumes



LHC data – Continue to break records:  
10.7 PB recorded in July  
CERN archive ~160 PB

June-Aug 2016  
>500 TB / day  
(Run 1 peak for HI was 220 TB)

# Castor (tape archive) at CERN




# The Future




# The road ahead

## LHC roadmap: according to MTP 2016-2020 V1



LS2 starting in 2019 => 24 months + 3 months BC  
 LS3 LHC: starting in 2024 => 30 months + 3 months BC  
 Injectors: in 2025 => 13 months + 3 months BC



# HEP Facility timescale



# High-Lumi LHC resource estimates




## Data:

- X10 from 2016
  - Raw 2016: 50 PB → 2027: 600 PB
  - Derived (1 copy): 2016: 80 PB → 2027: 900 PB


## CPU:

- x60 from 2016

# Event at HL-LHC




# Storage projections



- Areal density projections dropped from 40% to 15% pa
- Revenues are decreasing, market consolidation
- Technology evolution will not provide sufficient extra resources

# HL-LHC Solutions



11 Apr 2017


preGDB

17


# Data Management Directions

- Reduce cost/volume
  - cost of storage management
    - integrating standard (non HEP) solutions e.g. ceph
    - protocol zoo, SRM-less operation
    - T2 storage as cache
    - multi-site storage
    - regional federations
    - cloud storage
    - system manageability
  - storage overheads
    - redundancy
      - replication, erasure, RAID levels etc
    - reduce system reliability requirements?
      - reduce cost/impact of data loss
    - component technology
      - shingled disks
      - consumer/enterprise disks
- Reduce volume used
  - reduced number of global replicas
    - remote access
    - latency hiding
      - applications, overcommitting
    - global federations
    - CPU-only resources (inc cloud)
  - data formats and lifecycle, intermediate products
  - resource reporting
    - monitoring usage
    - eliminating dark data
  - data "enrichment"
    - popularity
    - caching, avoiding unused data
      - promoting locality in workflows
  - trading disk for...
    - tape
      - data parking
    - CPU
      - maintain metadata enabling regeneration of data on demand

# Networking growth (US ESnet)



# Relative Contribution to storage resources




# Evolution of computing model

- Consolidation
- WAN access
- Resource Diversity
- Standard Solutions
- Cohabitation (other big science projects coming online)

# Extras



# Disk technology



- PMR limit at 1TB/in<sup>2</sup>
- HAMR delayed
- Shingled disks are here
- Helium-filled disks are here
- New form factors?
- ...