LCG+experiments

with an adagio on data preservation

A.Boehlein, T. Mori, J-C Brient, C.Diaconu

June 13, 2013

Resources

- "Flat budget" is fashionable
  - Flatness with respect to what is more disputable
- Pledges mechanism is not always fully transparent, there is a (non)safety margin of ~20%
  - Scrutiny process to be improved (CRSG will work with WLCG MB towards a consolidated procedure to asses the needs)
- Funding by (EU) projects may disrupt personnel continuity
  - next H2020 not earlier tha one year
- Opportunistic usage continues to occur, experiments better and better prepared to use ad-hoc environements

Resources 2013/2014/2015

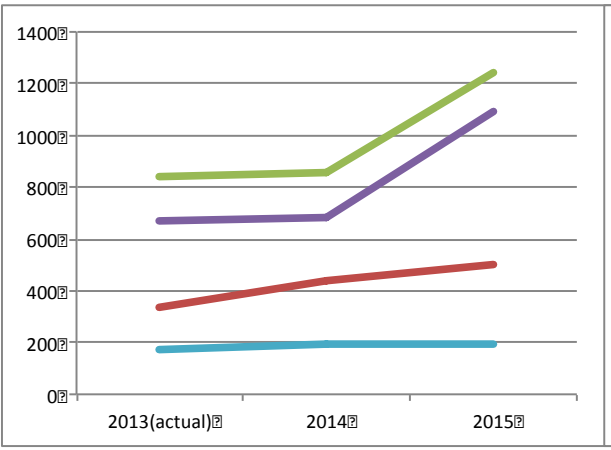

CPU

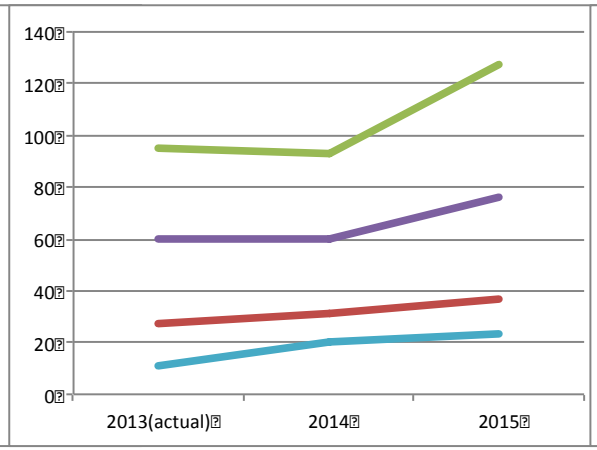

Disk

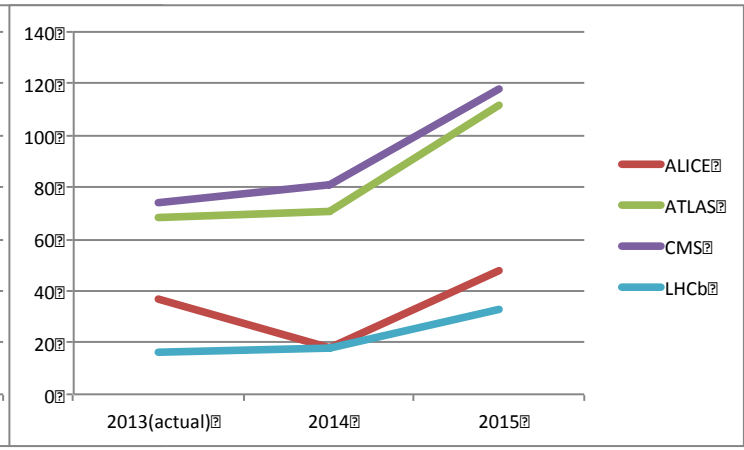

Tape

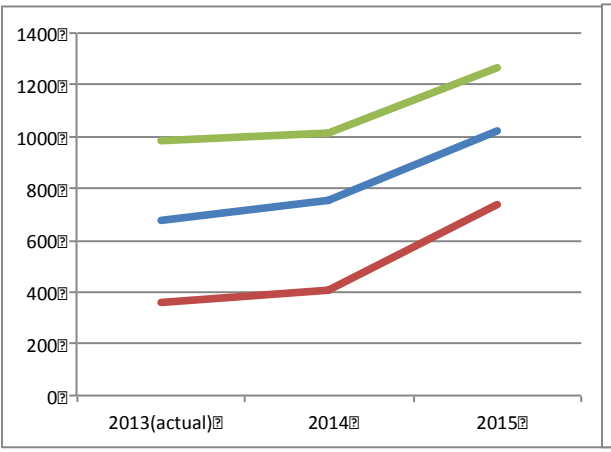

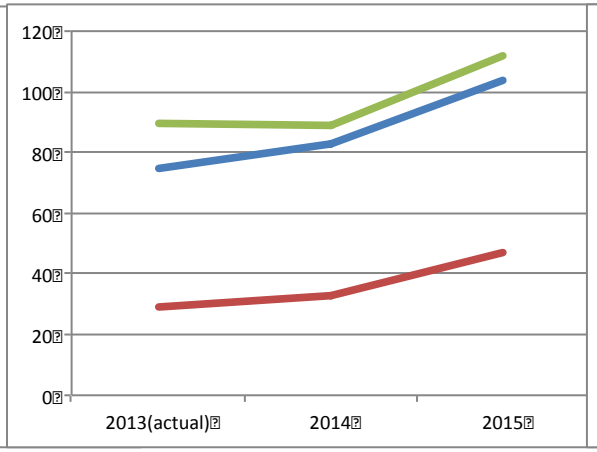

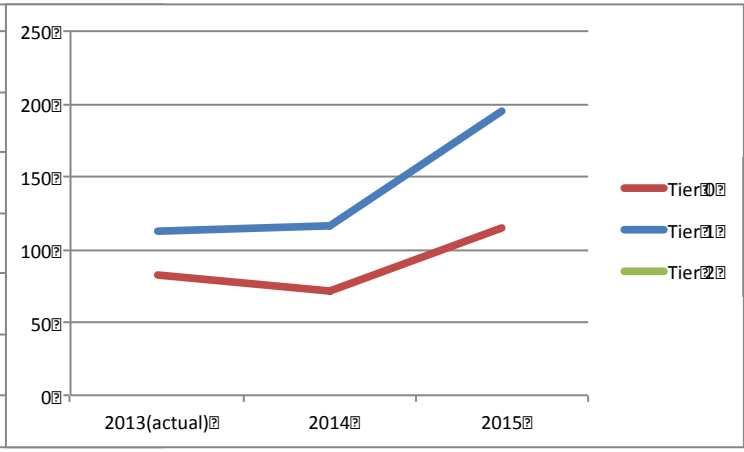

Situation following April 2013 RRB – RSG report

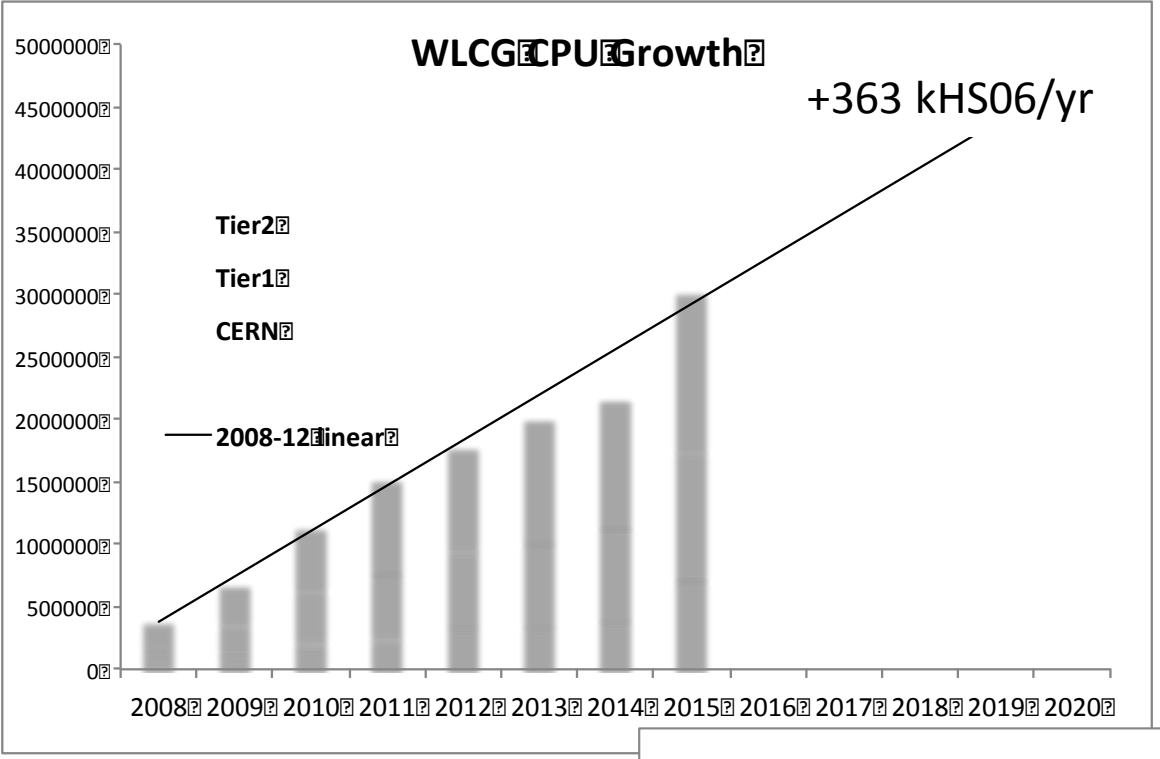
WLCG CPU Growth
+363 kHS06/yr
Tier2
Tier1
CERN
2008-12 linear
5000000
4500000
4000000
3500000
3000000
2500000
2000000
1500000
1000000
500000
0
2008 2009 2010 2011 2012 2013 2014 2015 2016 2017 2018 2019 2020

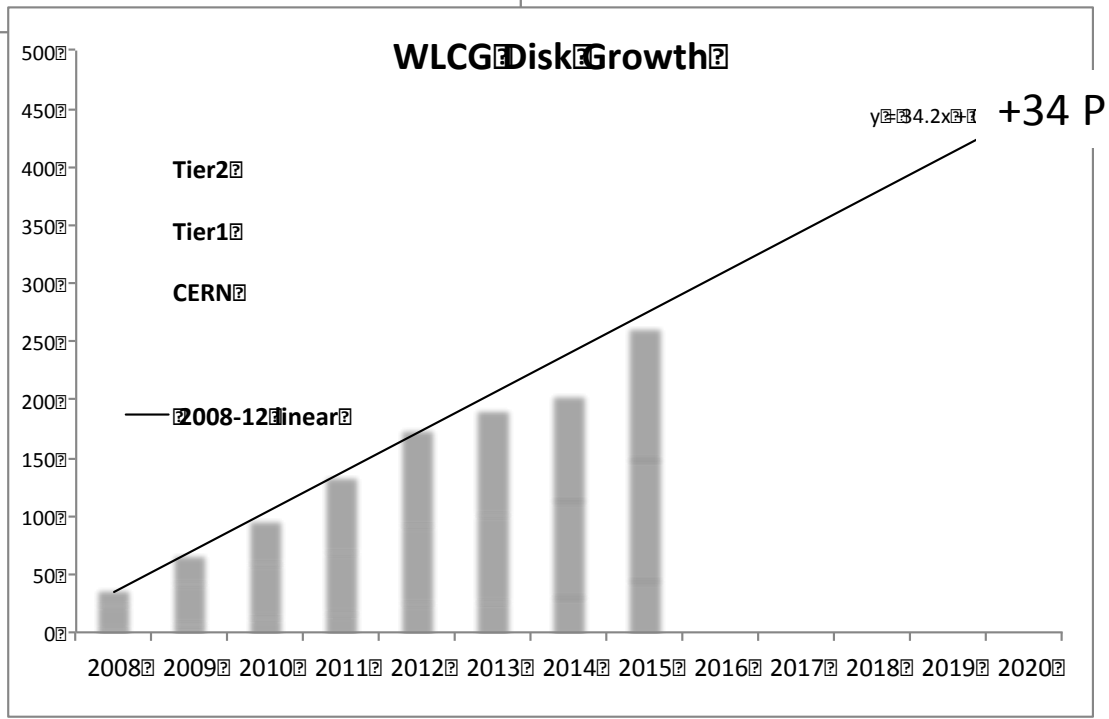
WLCG Disk Growth
+34 PB/yr
Tier2
Tier1
CERN
2008-12 linear
500
450
400
350
300
250
200
150
100
50
0
2008 2009 2010 2011 2012 2013 2014 2015 2016 2017 2018 2019 2020

Activity

- Eficient usage of ressources, no significant sign of slow down (except data processing activities, as expected)

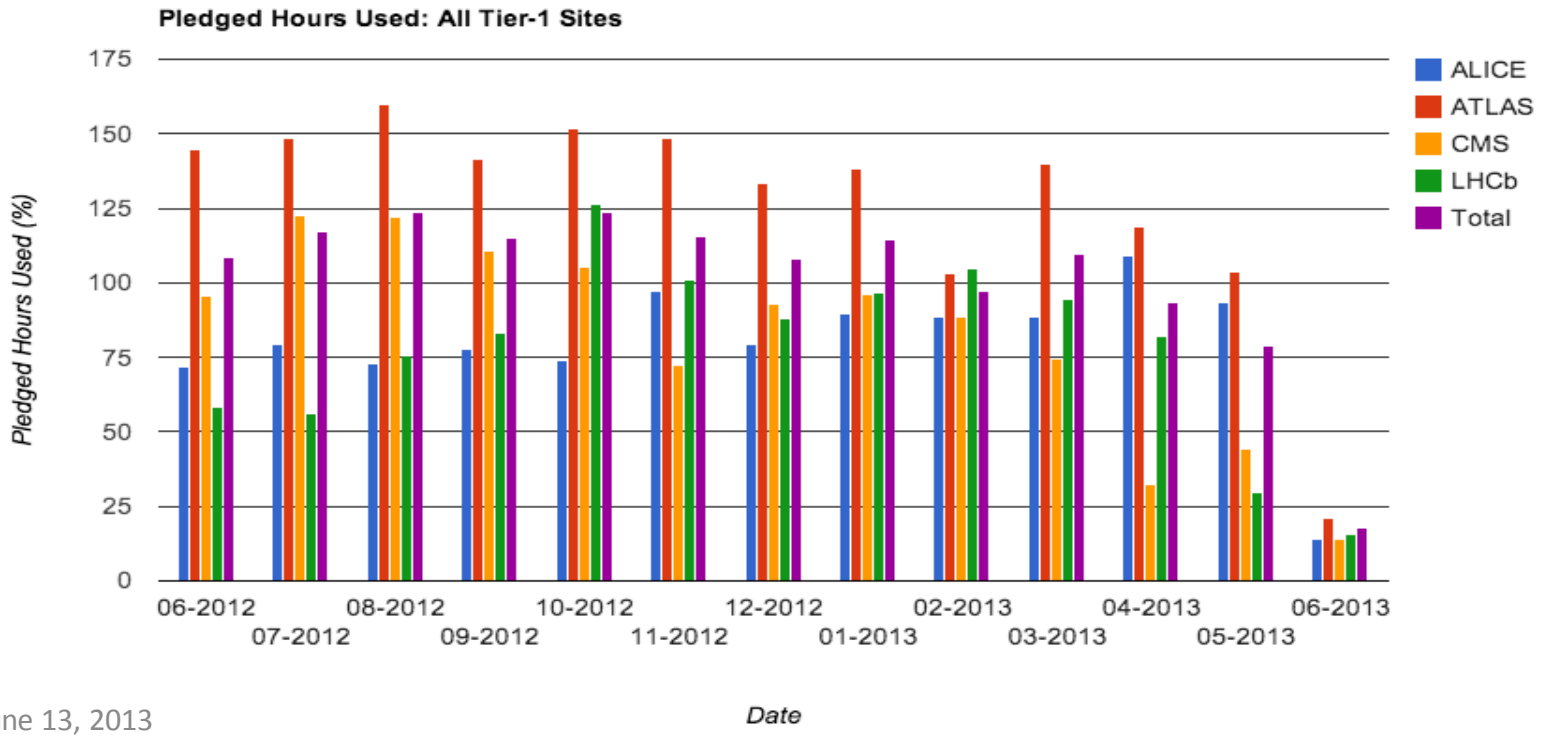

CERN IT

- CERN@Wigner installed
  - 5000 cores, few PB disk, 2x100 Gb/s
- CERN-IT: EMI and EGI-Inspire ended, new organisation
  - ES and GT merged, group size will decrease by 50% in the next year (Eudat, iMarine, HelixNebula)
  - Three sections:
    - Operations and Liaison
    - Monitoring Infrastructure
    - Information and Data

ATLAS

- Plans for improvements approved

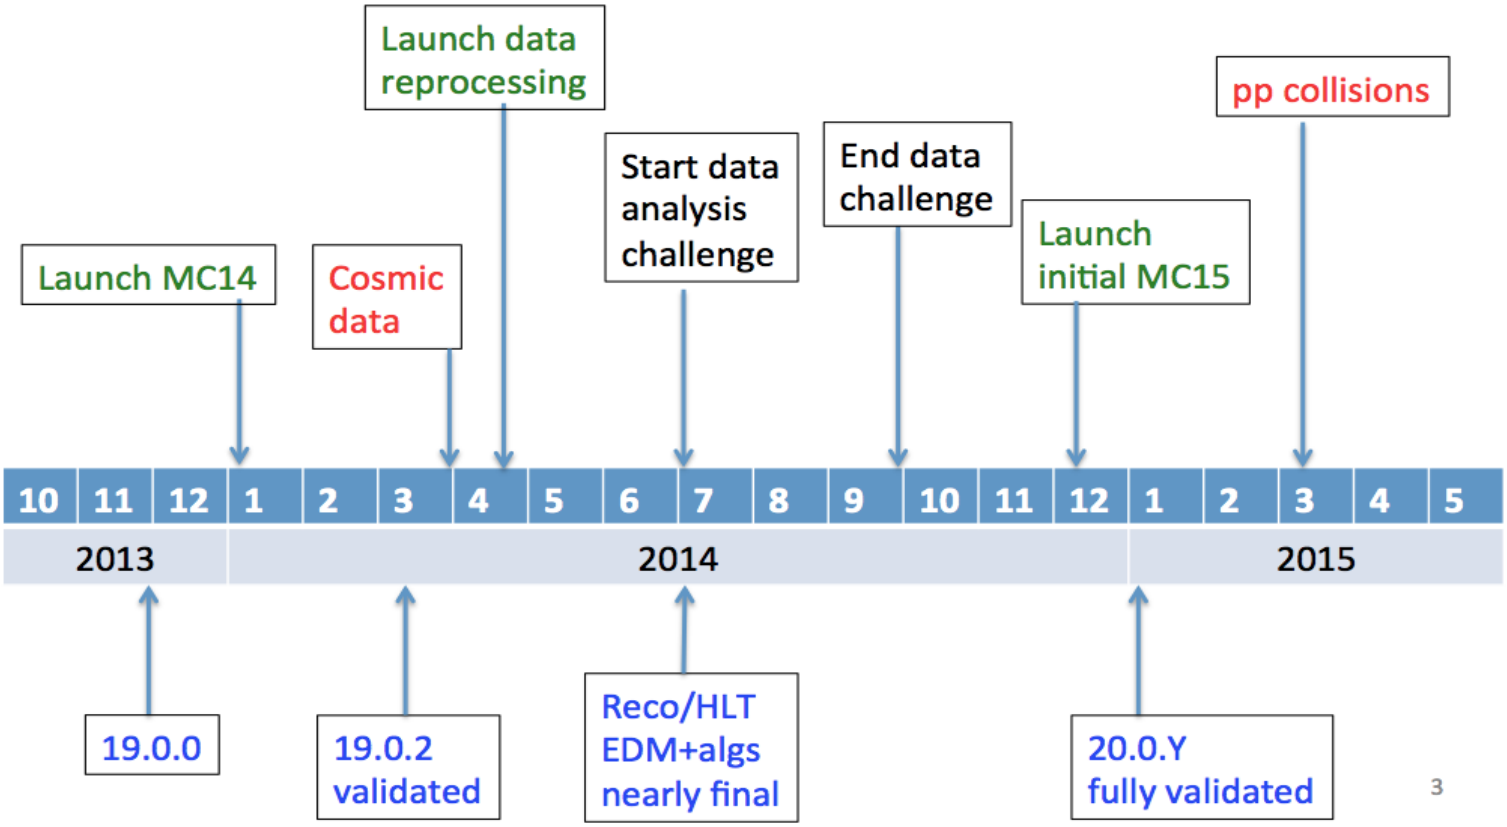

ATLAS: resources 2014/2015

- To be or not to be “flat”
  - Technological improvements are part of the “cash”

ATLAS RESOURCE PROJECTIONS COMPARING TO ‘FLAT BUDGET’ (Moore) AND 2012 USAGE

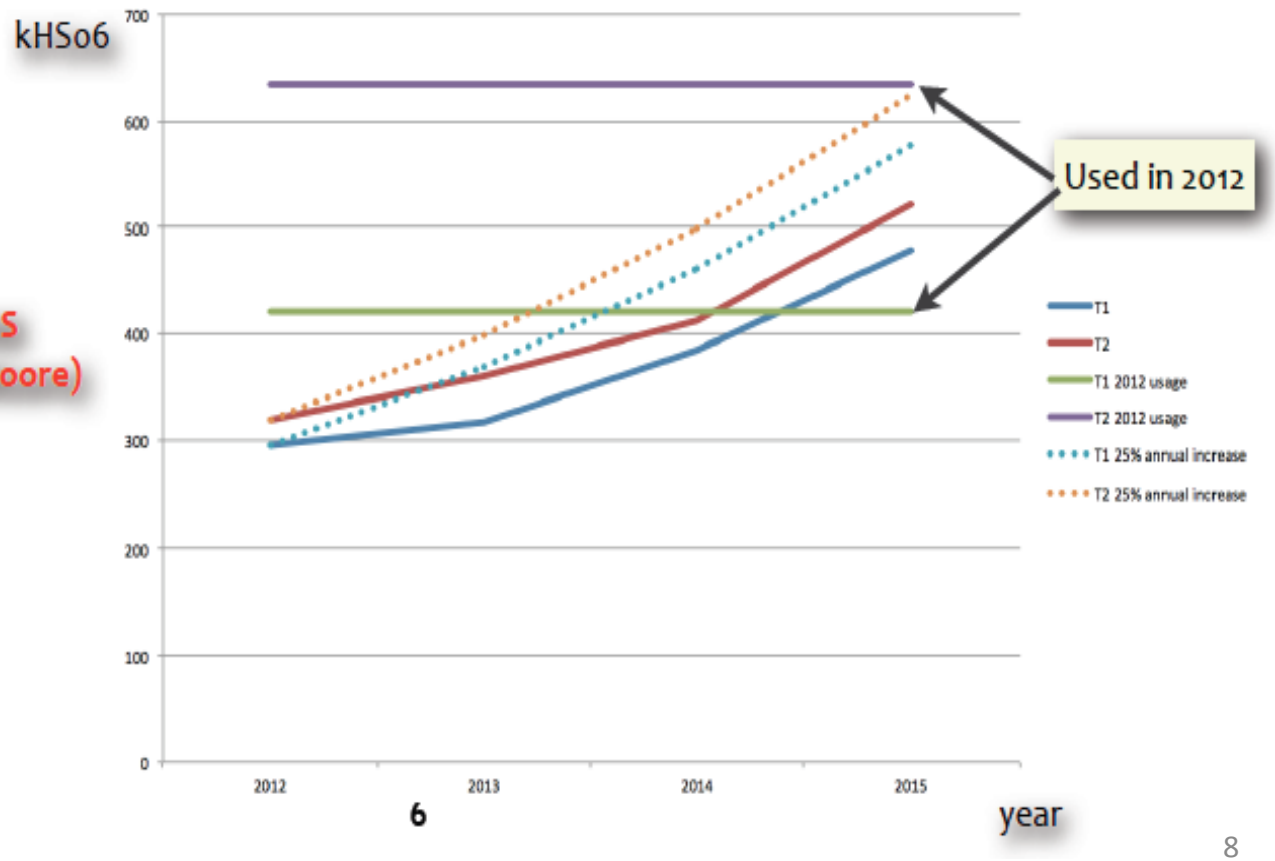

ATLAS:Resources and CRSG

- CRSG recommendations
  - aligned to pledges (different from the observed real usage figures)
  - In particular the tape situation worries ATLAS: efforts to recover space in some T1s

Resource	Site	2013 ATLAS	2013 CRSG	2014 ATLAS	2014 CRSG
CPU(kHS06)	T0+CAF	111(111)	111	111(111)	111
	T1	316(333)	319	373(327)	355
	T2	360(396)	350	408(399)	390
Disk(PB)	T0+CAF	10(10)	11	11(11)	11
	T1	35(36)	33	36(33)	33
	T2	51(49)	49	56(47)	49
Tape(PB)	T0+CAF	25(27)	23	31(31)	27
	T1	42(41)	40	53(43)	44

Table 13 ATLAS resources request. The numbers in parentheses () are pledged resources and the table contains our revised recommendations for 2014.

ATLAS: opportunities and challenges

- HLT commissioned, h/w to be installed then start
- Opportunistic usage of “atypical” resource
  - Google “pocket calculator” borrowed to ATLAS:
    - 4000 cores for 2 months (cloud-like environement)
- Explore HPC-like resources as well
  - Important source (traditional model)
  - Challenges and costs should be (at some point) properly estimated

ALICE

- Successful computing since the last meeting
  - 40kjobs routinely
  - CPU efficiency up to 90-95% (former worry now stably solved)
- Analysis 25% (12.5% in trains)
- Repro proceeds according to the defined strategy
    - pA (ongoing) then 2010/2011 data (pp and PbPb) – including MC – using 2012 calibration scheme

ALICE: plans for future

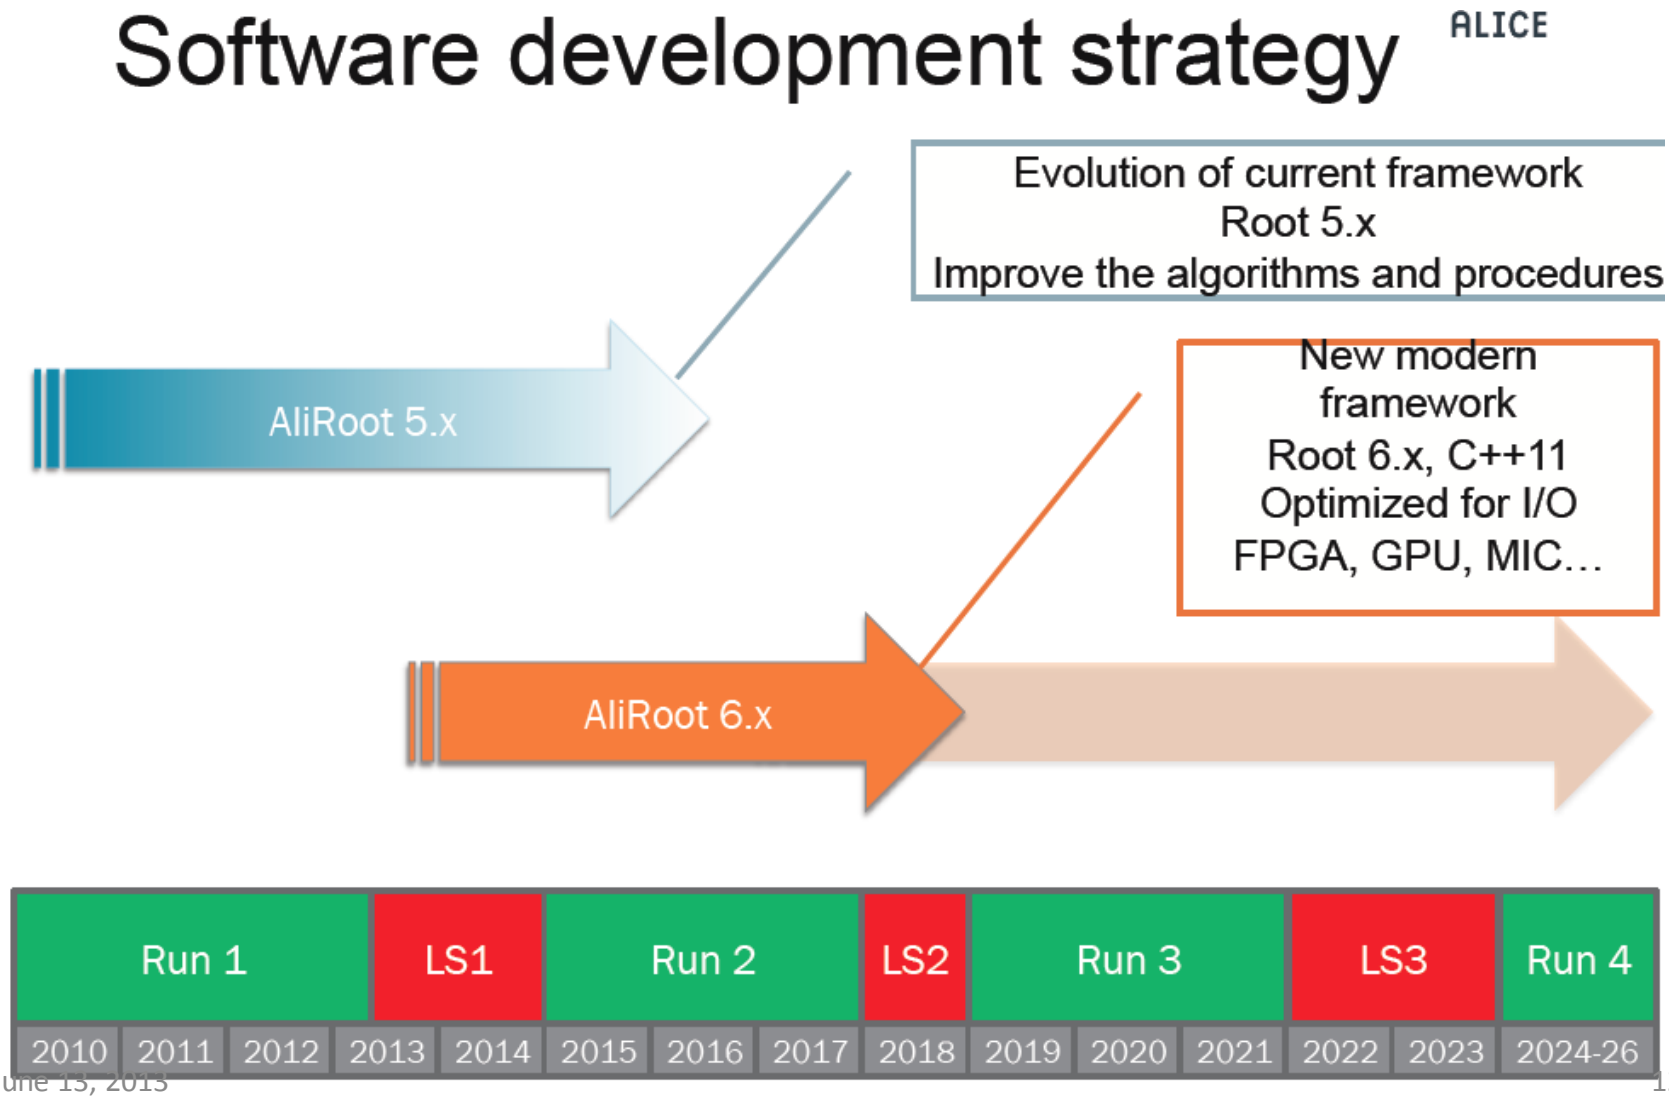

ALICE: Improvements

- Targeted for Run3 (Computing TDR in October 2014)
  - Will be incorporated in the existing framework if/when appropriate
- Simulation: G3 => **Geant 4**
    - Speed issues are being solved by close collaboration with G4 experts
    - Plan to profit from further improvements in G4(multi-thread, GPUs etc)
- Reconstruction: use TPC/HLT tracks as seeds
- Analysis: reduce the turnover to 12 h, include most of the users
- HLT farm: used for offline processing
    - currently 2500 cores (108FPGA+ 64 GPGPU)
    - Foresee a massive increase for RUN3 for online reconstruction
      - 20X reduction in event size

ALICE goes virtual

- Use CERNVM family to enable various use-case (i.e./and/or resources)

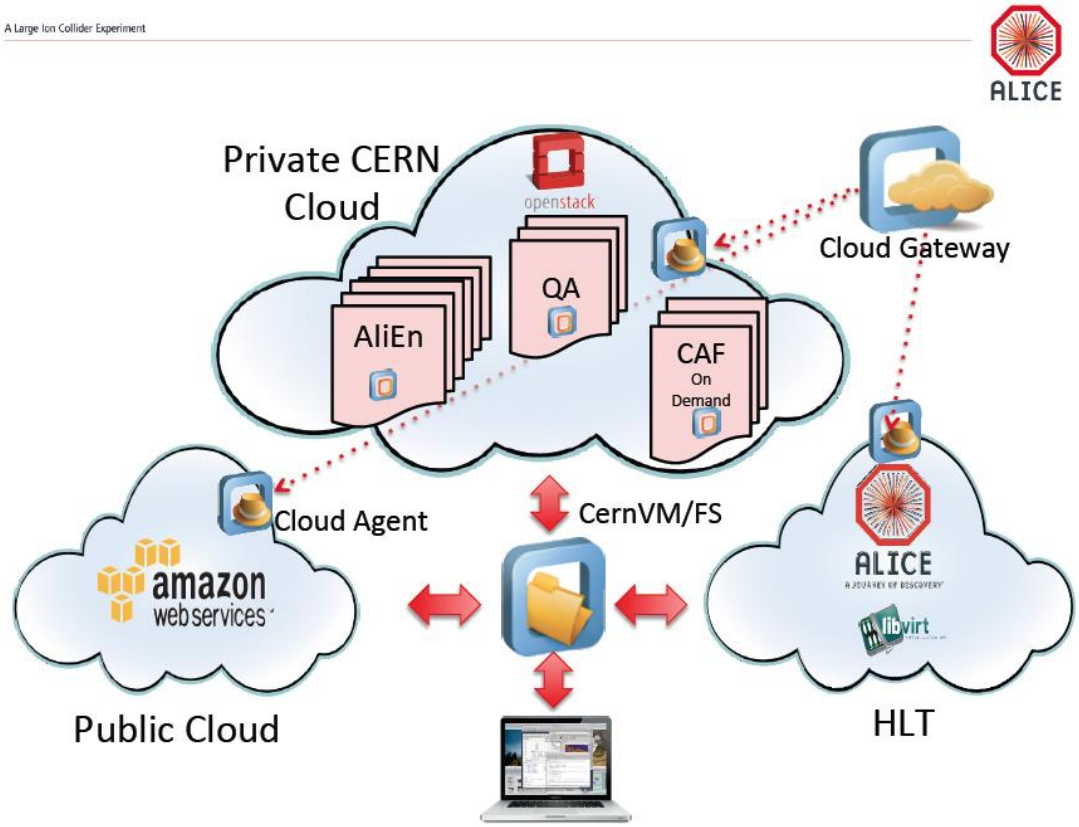

CMS

- Repro ended on T1's, MC requests not yet ready
- 2015 preparation, LHE and upgrade studies will take significant resources and organisation in the second half of the year
- Flexible usage of resources
  - Move to CVMFS for s/w distribution, simplifies multiple platform usage and enable opportunistic computing
  - Disk/tape separation (archive versus disk-only at T1s)
  - Xrootd deployment ongoing
  - Multi-core project: pilot system configured
  - HLT comissioning is ongoing
  - Data Management:
    - work to develop common algorithms and frameworks within the effort from CMS and Operations Coordination
    - Implies changements in T2 space management and in the security model (PhEDEx)
  - Analysis framework: CRAB3 (with Panda back-end) preproduction in August
- Data management schemes and (user) analysis workflows to be validate by mid 2014.

CMS: HLT comissioning

- 4000 jobs running, some stability issues to be solved

LHCb

- Computing activities according to plans
  - Incremental re-stripping of 2010+2011
  - Processing and stripping of 2013 pA data
  - Simulation for 2012 commissioned
- HLT plays a significant role

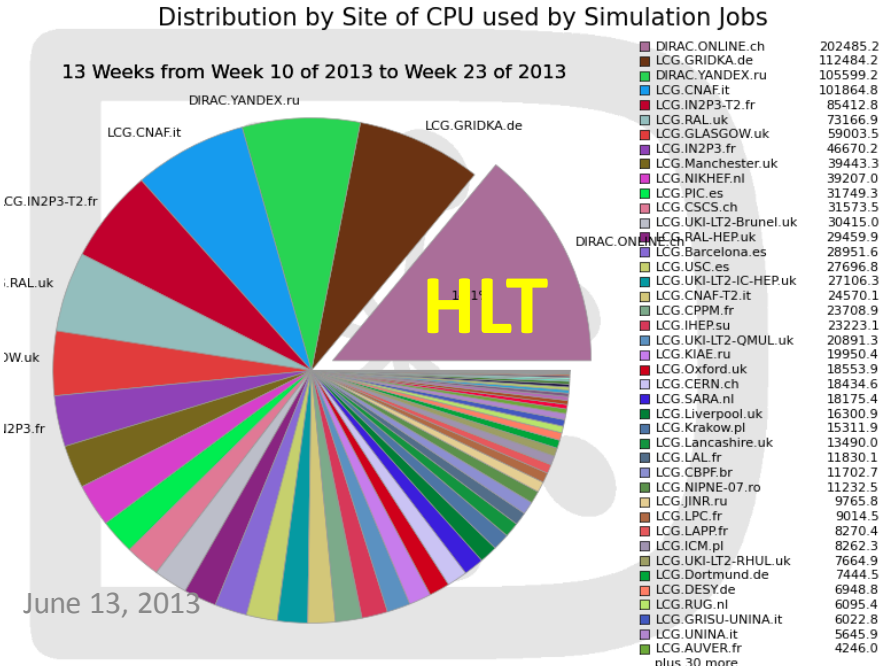

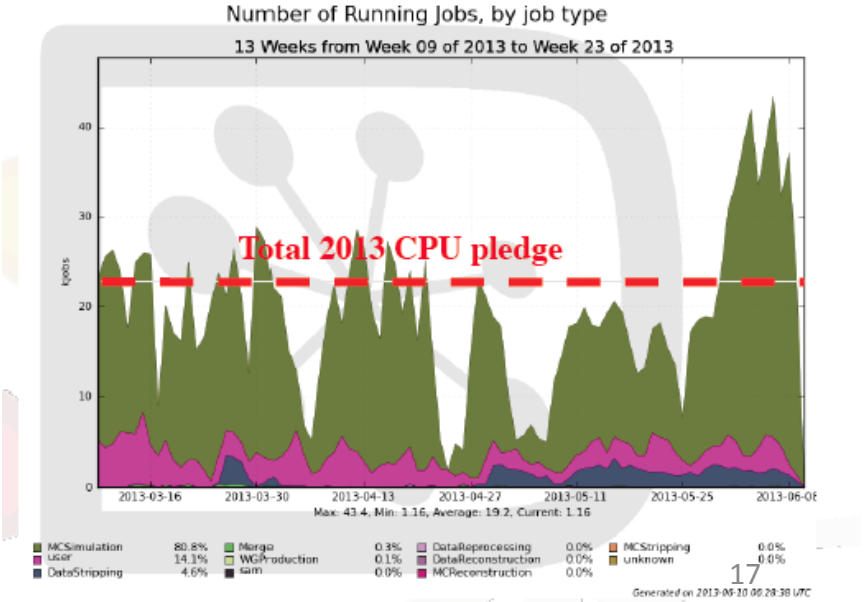

LHCb: Disk and Tapes

- Disk
  - Deficit in disk in 2013 (11/15 PB)recovered due to delays
  - Change in computing model:
    - Allow disk (and analysis) at selected Tier2 ('T2-D')
    - Goal is ~10 sites with >300TB per T2-D, but start with 100TB
- Tapes
  - Massive access (restripping 3.8 PB) problematic for some sites
  - Some tapes lost, recovery procedures not fully defined

LHCb: software improvements

- LS1 is window of opportunity for adopting latest software
  - Must be in place by the early 2014, to allow validation and commissioning (in particular for HLT)
- Several technical steps forward
  - Compiler (gcc 48), support for C++11
  - Adopt ROOT6
- Internal computing reviews
  - On distributed computing and new data model
  - Person power is an issue

Data Preservation

- DPHEP ICFA Study Group formed in 2009
  - Major labs and collider experiments
- LHC experiments joined in 2011
  - Local/CERN-LHC task force on data preservation and access
  - Contact persons/working groups in all experiments
- DPHEP Blueprint published in mai 2012
- October: CERN appoints a DPHEP Project Manager (Jamie Shiers)
- Feb. 2013 ICFA Statement => DPHEP Collaboration

ICFA Statement – March 15, 2013

- *The International Committee for Future Accelerators (ICFA) supports the efforts of the Data Preservation in High Energy Physics (DPHEP) study group on long-term data preservation and welcomes its transition to an active international collaboration with a full-time project manager.* ***It encourages laboratories, institutes and experiments to review the draft DPHEP Collaboration Agreement with a view to joining by mid- to late-2013.***
- *ICFA notes the lack of effort available to pursue these activities in the short-term and the possible consequences on data preservation in the medium to long-term. We further note the opportunities in this area for international collaboration with other disciplines and encourage the DPHEP Collaboration to vigorously pursue its activities. In particular, the effort required to prepare project proposals must be prioritized, in addition to supporting on-going data preservation activities.*
- ***ICFA notes the important benefits of long-term data preservation to exploit the full scientific potential of the, often unique, datasets****. This potential includes not only future scientific publications but also educational outreach purposes, and the Open Access policies emerging from the funding agencies.*

- DPHEP Collaboration Implementation Board installed
- Collaboration agreement: first draft circulated, signature campaign will start soon

Data Preservation: LHC perspective

- Long term perspectives should be taken into account early enough
  - Some Data sets already “old” (2010, soon 2011)
- Working/task groups in place in experiments
  - Act within the existing activities
  - Implement concepts and define policies
- Projects emerge within new funding opportunities or as a part of upgrades/model improvements
  - **DPHEP-CERN H2020 Vision Document in work**
- Clear multi-disciplinary potential, to be exploited within the future collaboration

DPHEP-CERN Vision Document

#	Summary of Recommendation
R1	CERN Technical Student (TECH) for DPHEP Portal – Fall 2013 committee
R2	Small fraction of future TECHs for DPHEP work
R3	DPHEP activities integrated into CERN School of Computing
R4	Experiments / institutes to use archive material for training
R5	Formalise and increase collaboration with Digital Libraries
R6	HEPiX and IEEE for coordination of archive issues – DSS and PES
R7	Pursue H2020 and NSF projects via RDA and other channels
R8	Calculate total cost of ownership of data – WLCG and DSS + 4C
R9	Very long term data archives ("Beyond H2020")
R10	Collaborate with KT and Outreach activities
R11	Project funds for the above
R12	Foster exchange of information on migration issues / experiences
R13	Validation system – support implications for hosting institute(s)
R14	HEP Software Collaboration – Coordinate with and through WLCG
R15	Project funds for above
R16	CERN Fellow for "last gasp" LEP Data Preservation activities – Fall 2013 committee preferred
R17	"CERNLIB consortium"

Connections and opportunities

- Mobilize resources through existing structures:
  - **Research Data Alliance:**
    - Funding / strong interest from EU, US, AU, others
    - Part of roadmap to "Riding the Wave" 2030 Vision
    - **STFC and DCC personnel strongly involved in setup**
  - **WLCG:**
    - Efforts on "software re-design" for new architectures
    - Experiment efforts on Software Validation (to be coordinated via DPHEP), building on DESY & others
  - **DPHEP:**
    - Coordination within HEP and with other projects / disciplines
- National & International Projects
  - H2020 / NSF funding lines
  - National projects also play an important role
    - France, Italy, UK
    - US/DASPOS

General comments

- Computing is an very good shape
  - pressure is less visible in the last weeks, but this can change over LS1
- Resources
  - provisioning relies on technological improvements, in particular for 2015
  - Opportunistic may become routine , costs benefits to be followed, new ideas may emerge (HPC, private/industrial clouds etc.)
  - Pledges versus real allocation and true usage need a better predictability: LHCC and CRSG should continue the good cooperation in order to ensure a robust and efficient resources provision
- Work on improving the models started, with very good prospects for 2014
  - Collaborations on methods/software across experiments
  - New paradigms may emerge (virtuallisation, clouds etc.)
  - Document expected before the next LHCC meeting: prediscussed with the referees?
- Data Preservation is now also considered by the LHC experiments
  - International Collaboration being build around DPHEP, with a strong potential for interdisciplinary