Ian Bird

Meeting with LHCC Referees

CERN, 3rd June 2014

WLCG status

WLCG MoU Topics

- KISTI (Rep. Korea)
  - Status as full Tier 1 for ALICE approved at WLCG Overview Board in November 2013
    - Agreed that all milestones met; performance accepted by ALICE;
    - Upgrade of LHCOPN link to 10 Gbps planned and funded
- Latin America: Federation (CLAF), Tier 2 for all 4 experiments – initially CBPF (Brazil) for LHCb and CMS
  - Since last RRB, new sites added:
    - UNIANDES Colombia (CMS), UNLP Argentina (ATLAS), UTFSM Chile (ATLAS), SAMPA Brazil (ALICE), ICM UNAM Mexico (ALICE), UERJ Brazil (CMS)
- Pakistan: COMSATS Inst. Information Technology (ALICE): MoU in preparation

- Update on progress with Russia Tier 1 implementation
  - → this meeting

Status of Wigner centre

- In production;
  - >1000 worker nodes installed
  - Disk cache (EOS) scaled with CPU capacity
- Some network problems:
  - One link was unstable every few days
  - Some firmware incompatibilities between NICs, switches and cables (!)
    - Hopefully now resolved
- Some anecdotal job inefficiencies; a systematic performance analysis was done
  - Building some additional monitoring for longer term management
    - Dedicated perfsonar testing (as deployed in WLCG)
  - Uncovered some issues with pilot jobs
  - Clear difference in efficiency between AMD and Intel – depending on workload (not new)
  - Must use correct caching algorithm in ROOT, or performance suffers for I/O bound jobs; continuing study of data transfer performance
- No significant difference in efficiency between Geneva and Wigner
  - Monitoring much improved
  - Detailed performance studies will benefit all remote I/O situations

EMI Middleware Support
Status 2014

- One year after end of EMI
  - Middleware support pledges limited to one year
- Cristina Aiftimiei (INFN /EGI) report at May GDB
  - contacted all Product Teams
  - preliminary (incomplete) status:
    - ARGUS – SWITCH, bug fixes on best effort basis only
      - needs support for evolution (identify feds, clouds, etc.)
    - gLite security products need clarification
    - INFN supported products will be supported
    - CERN products will be supported
    - NDGF products will be supported

Summary of RRB

held on 29 April 2014

Summary of RSG report

Data popularity

- CRSG asked all experiments for data-popularity information for disk use
- Minimise storage of data which is never or seldom read
- We show ATLAS plot because we found it most useful
- Pursue this with all experiments in next and future scrutinies
- We hope that revealing and monitoring this information will lead to more efficient use of disk space

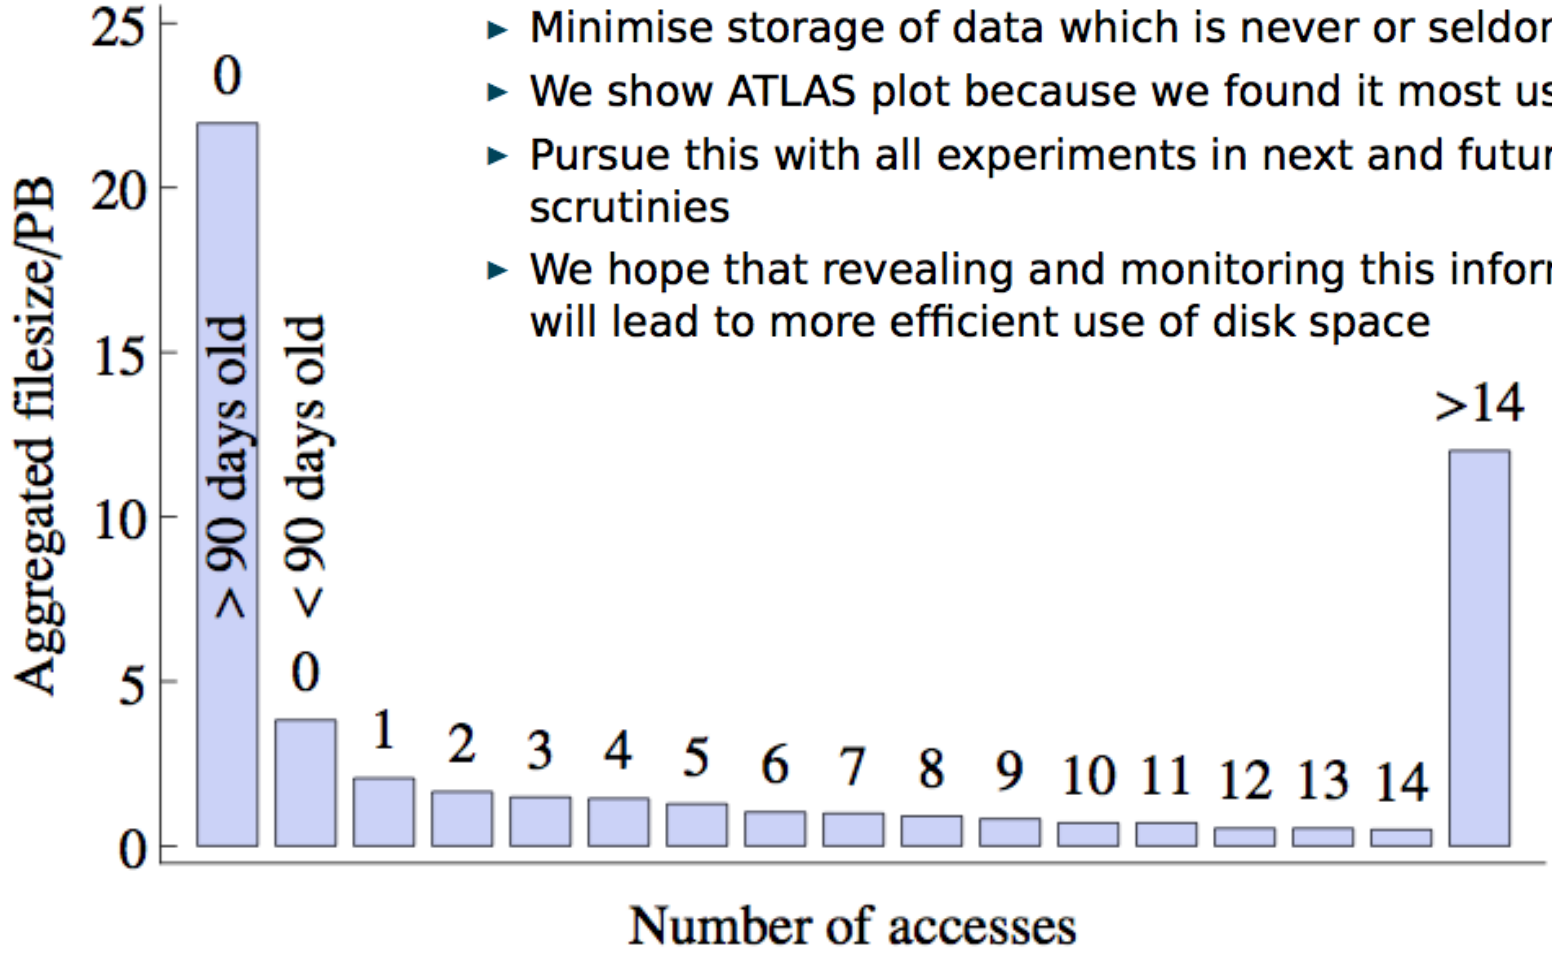

Volume of data versus number of accesses in ATLAS DATADISK at T1s and T2s for 90 days to 14 March 2014

Scrutiny of 2015 requests

ALICE

		2014 ALICE	2014 CRSG	2015 ALICE	2015 CRSG
CPU (kHS06)	T0	135	135	175	175
	T1	110	110	120	120
	T2	190	190	200	200
Disk (PB)	T0	8.3	8.3	14.5	14.5
	T1	10.1	10.1	17.8	17.8
	T2	12.8	12.8	22.7	22.7
Tape (PB)	T0	12.0	12.0	16.2	16.2
	T1	6.0	6.0	10.2	10.2

- CPU and storage for Run 2 increased by 25% (beam energy and pileup)
- PbPb and pPb events include TPC data; raises reco and sim times

ALICE

- Major demands come with heavy-ion running towards end of year.
  - T0 CPU needed for heavy-ion reconstruction before following year's pp run.
- Sum of T1 and T2 resources more important than precise division between them.
- Some significant jumps in requests for 2015 and on to 2017.
- HLT farm being upgraded; expected to be operational at end of 2014. Planned use for offline tasks in Run 2.

ATLAS

		2014 ATLAS	2014 CRSG	2015 ATLAS	2015 CRSG
CPU (kHS06)	T0	111	111	205	205
	T1	355	355	462	450
	T2	390	390	530	520
Disk (PB)	T0	11	11	14	14
	T1	33	33	39	36
	T2	49	49	55	53
Tape (PB)	T0	27	27	33	33
	T1	44	44	65	65

- Request essentially the same as last October
- Disk: reduced pre-placement and more aggressive deletion of unused data
- Multi-core capable software, new analysis format, removal of a data-copy step

ATLAS

- CRSG strongly supports software development. *Benefits needed* to constrain future resource needs
- CRSG welcomes more aggressive policy for deleting unused data, but maintains pressure to make more effective use of disk with small reduction in T1 and T2 disk
- Acknowledge successful use of HLT farm; but we think its use should be included in requests (hence CPU reduction)

CMS

		2014 CMS	2014 CRSG	2015 CMS	2015 CRSG
CPU (kHS06)	T0	121	121	271	271
	T1	175	175	300	300
	T2	390	390	500	500
Disk (PB)	T0	7	7	3+12	15
	T1	26	26	27	26
	T2	27	27	31	29
Tape (PB)	T0	26	26	31+4	35
	T1	55	55	74	74

- 2015 requests unchanged since last October
- CMS takes account of use of HLT in requests

CMS

- Efforts to constrain CPU requirements
  - Software efficiency improvements. CRSG strongly supports this.
  - T0 setup to be more like T1 to allow prompt reconstruction at T1s from 2015
  - Fewer reprocessing passes
  - Reduction in ratio of simulated to real events — may hurt physics output
- CRSG acknowledges these efforts
- As for ATLAS, still push for aggressive cleanup of unused data to make more effective use of disk → small reduction in T1 and T2 disk

LHCb

		2014 pledge	2015 LHCb	2015 CRSG
CPU (kHS06)	T0	34	36	36
	T1	110	118	118
	T2	62	66	66
	HLT + Yandex		10+10	
Disk (PB)	T0	4.0	5.5	5.5
	T1	11.7	11.7	11.7
	T2	1.1	1.9	1.9
Tape (PB)	T0	8.5	11.2	11.2
	T1	11.0	23.7	23.7

- More use of T2 for simulation and analysis; introduction of T2 disk
- Use of HLT and Yandex accounted for in request

LHCb

- Several changes for 2015
  - No further reprocessing of Run 1 data in 2015
  - Postpone *re*processing of raw data to LS2
  - Omitted reconstruction pass and reduced stripping
  - Reduced ratio of full DST to microDST to reduce storage
- Jump in tape for 2015, including significant space for data preservation
- Bigger jumps in CPU/disk/tape anticipated for 2016
- LHCb noted that common LHC running assumptions used here may be pessimistic

This led to a question about the assumptions for 2015 runnning in general:
- Main driver is livetime
- What about delayed improvements (e.g in triggers)

Conclusion was that current assumptions for 2015 are conservative – all factors conspire to make need for resources larger

C-RSG comments

- Run 2 requests made with assumption of flat budgets
- Data preservation: distinguish ability to read/analyse old data from requirements for open/public access (both tech + effort)
  - Former should be included in cost of computing
  - Latter is additional cost?
- C-RSG acknowledges use of HLT farms in LS1 and plans to use in Run 2.
  - C-RSG does not consider this to be opportunistic
  - Under control of experiment
  - Adjusted request where this had not been done by the experiment

C-RSG comments – 2

- Improving software efficiency is essential to constrain growth in requests.
  - (*Some of*) The resulting gains are already assumed.
  - C-RSG strongly supports and recommends that sufficient effort is funded
- Effectiveness of disk use only partly captured by occupancy
  - C-RSG welcomes experiments' efforts to purge obsolete/unused data, and thanks them for popularity information
- Good networking has been exploited to reduce disk use (fewer pre-placed copies) and move processing between tiers.
  - Danger that poorly networked sites will be underused and possible cost implications of providing network capacity

Flat budgets? (C-RSG view)

Meeting flat budgets?

Overview of sum of all requests from 2014 start

- CPU and disk at T0 jump above FB in 2015 but subsequent growth within FB
- Other resources, apart from T2 CPU, grow above FB earlier or later in Run 2
- For 2013 start: 2015 jump at T0, growth of tape

Full exploitation of physics potential of LHC and experiments from 2015 will require significant increase in resources.

- Meeting FB growth with FB spending depends on past funding, hardware replacement cycles, other costs (eg people, electricity)
- Might need increased budget in short term even to meet fixed-cost hardware performance increase

Combined (sum of experiments)

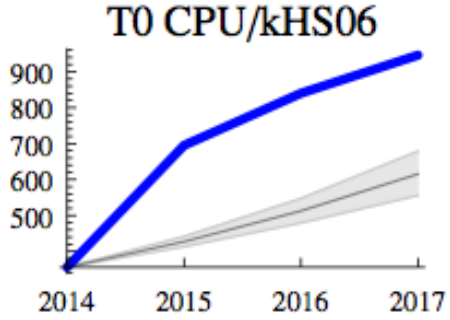

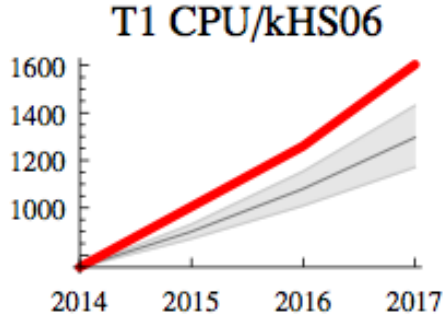

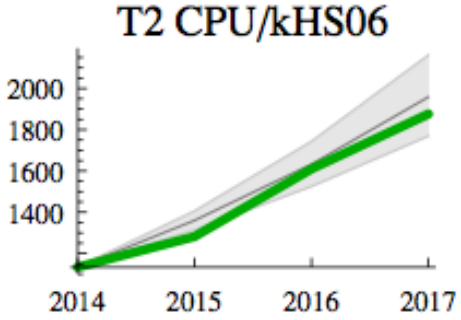

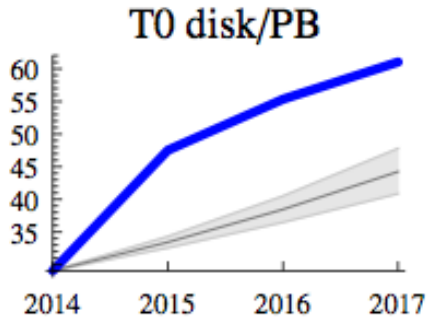

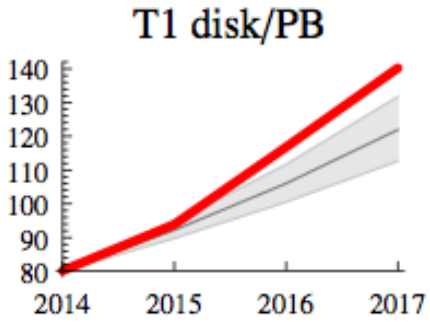

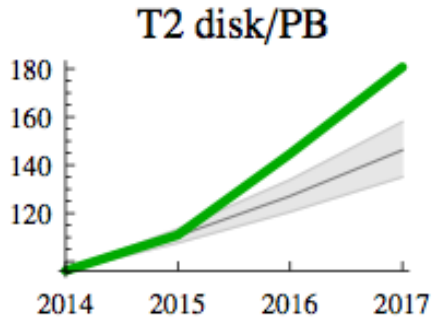

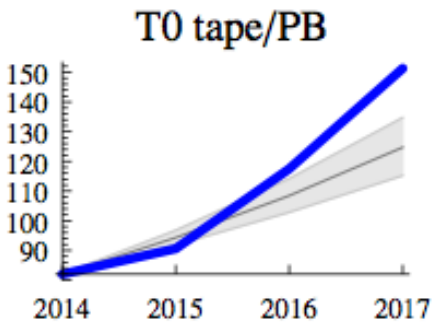

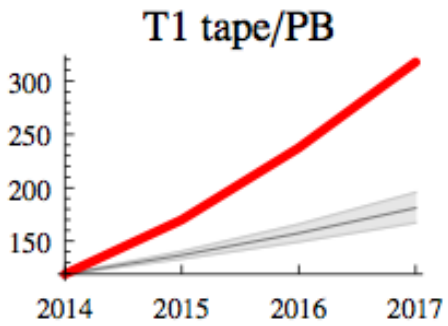

Starting from 2014 pledges

Budgets and pledges

- "Flat budget" scenario was guidance of the RRB in April 2013, reinforced in October
- Reflected in the computing requirements growth of the computing model updates (2015-17)
  - BUT: this is very gross average – we do not know the details of every site, funding agency
    - Growth figures based on understanding of market, experience at CERN and other large sites
  - Actual ability to grow strongly depends on history
    - Replacement cycle, past growth, local purchase costs and rules, etc
- Thus it is clear that the real potential growth even with flat budgets is not simple to model without feedback from sites/countries/funding agencies

Budgets and pledges

- Feedback is via the pledge process
  - Not sure what other mechanism is feasible
- Problems:
  - Pledges are given “just in time” (actually at the last minute)
  - Pledges only given in October for following year
    - For some FA’s even this is a guess and actual budgets only known later
- We really need a multi-year “estimated pledge” to match the 3-year resource outlook
  - Only with this can we model/adjust the planning
  - If some analysis has to be delayed now, can it ever be caught up?
- *How can this be achieved / approached / estimated?*

HEP Software Collaboration

Summary of the workshop held at CERN on April 3-4 2014

HEP SW: Context

- Experiment requests for more resources (people, hardware) to develop and run software for next years physics;
- Prospect that lack of computing resources (or performance) will limit the physics which can be accomplished in next years;
- Potential for a small amount of additional resources from new initiatives, from different funding sources and collaboration with other fields;
- Large effort required to maintain existing diverse suite of experiment and common software, while developing improvements.
  - Constraints of people resources drive needs for consolidation between components from different projects, and for reduction of diversity in software used for common purposes.

HEP SW: Goals

- Goals of the initiative are to:
  - better meet the rapidly growing needs for simulation, reconstruction and analysis of current and future HEP experiments,
  - further promote the maintenance and development of common software projects and components for use in current and future HEP experiments,
  - enable the emergence of new projects that aim to adapt to new technologies, improve the performance, provide innovative capabilities or reduce the maintenance effort
  - enable potential new collaborators to become involved
  - identify priorities and roadmaps
  - promote collaboration with other scientific and software domains.

HEP SW: Model

- Many comments suggested the need for a loosely coupled model of collaboration. The model of the Apache Foundation was suggested: it is an umbrella organisation for open-source projects that endorses projects and has incubators for new projects.
- Agreed aim is to create a Foundation, which endorses projects that are widely adopted and has an incubator function for new ideas which show promise for future widespread use.
  - In the HEP context, a Foundation could provide resources for life-cycle tasks such as testing etc.
- Some important characteristics of the Foundation :
  - A key task is to foster collaboration;
  - developers publish their software under the umbrella of the ‘foundation’; in return their software will become more visible, be integrated with the rest, made more portable, have better quality etc
  - it organizes reviews of its projects, to identify areas for improvement and to ensure the confidence of the user community and the funding agencies.
  - a process for the oversight of the Foundation's governance can be established by the whole community;

HEP SW: Next Steps

- First target is a (set of) short white paper / document describing the key characteristics of a HEP Software Foundation. The proposed length is up to 5 pages.
  - Goals
  - Scope and duration
  - Development model
  - Policies: IPR, planning, reviews, …
  - Governance model
  - …

- It was agreed to call for drafts to be prepared by groups of interested persons, within a deadline of ~ four weeks (i.e. May 12th.) The goal is a consensus draft, to be used as a basis for the creation of the Foundation, that can be discussed at a second workshop some time in the Fall 2014.