

Ian Bird

WLCG GDB, Amsterdam

7th March 2016

Summary of Lisbon Workshop

Introduction

- ❑ Two days devoted to medium term (Run 2-3) and longer term (Run 4) concerns
- ❑ ~140 people registered
- ❑ Aimed for more of a discussion format rather than presentations
 - (Informal) feedback from many said this was useful
 - Some aspects probably needed a bit more preparation to be more successful

Shorter term & ongoing work (key points only)

Security, AAA, etc

- ❑ Fully address traceability:
 - Freeze deployment of glexec – keep it supported for the existing use, but no point to expend further effort in deployment
 - Need full traceability solutions and tools:
 - Experiment frameworks
 - Work needed to get VO traceability info into CERN SoC (or to ...)
 - Use of VMs/containers helps
 - Invest in deploying ‘big data’ tools for managing traceability data
 - SoC capability, appliances?
 - Traceability working group needed?
 - This is a reflection of the trust of VO's now, rather than trying to trust many individuals
- ❑ Incident Response & dealing with threats (against people)
 - Invest in (coordinate in WLCG) better intelligence/trust with other communities, and with vendors
- ❑ Federated IDs: long term goal
 - eduGain etc; develop use cases within AARC2 project
- ❑ Policy work
 - Data, privacy, etc

Compute

- ❑ Lot of discussion of cloud and cloud use
 - Models of how to provision and access commercial clouds are evolving
 - HNSciCloud will explore more aspects
 - Many models of using clouds, containers, VM's
 - (vac, batch queues, & etc., etc.)
 - Probably exposure of experience in GDB is a correct way to proceed for the moment
- ❑ Lots of discussion on the use of HPC
 - Useful in certain specific circumstances or as opportunistic resources
 - Significant effort expended in this area, for few % gain in resources
 - Not to be ignored, but can we gain more in other areas for a similar effort???
 - What should our strategy be here – generalise to opportunistic resources more broadly?
- ❑ Issues of IP connectivity, lack of storage access, etc. (see these issues in HPC, cloud, etc.)
 - Addressing these fully will actually benefit our entire operation
 - Long standing concern over connectivity and implications at sites

Data

- ❑ Object Storage
 - multiple motivations
 - scalability (exploiting less meta data) as embedded storage
 - also - nicer/more modern tools
- ❑ Roles of smaller sites (or with little support effort)
 - demo - describe scenarios, ask for supporters, drop rest
 - cataloged cache (eg dpm)
 - proxy cache (eg xroot)
 - Rob (largely non hep specific components , trust?)
 - boinc (no operator, no shared storage)
- ❑ Common questions
 - prove simulation use case
 - analysis at small sites will be compressed
 - estimated impact on T1 (eg wan stageout)
- ❑ Federation of storage, desired by some experiments
 - Prefer to have a single storage endpoint to aggregate resources across several sites
 - Maintain redundancy and replica locality

Data

A possible medium term plan

- SRM: progress with decommissioning, apart for tapes
- Data access, upload, download:
 - Consolidate around the xrootd protocol (mainstream)
 - Progress with HTTP support, valuable both in the short and medium/long term
- Data Transfer
 - Investigate possible alternatives to gridFTP (e.g. xrootd like Alice, HTTP)
 - Do not forget that data deletion is as challenging as data transfer

Info sys, accounting, etc

❑ Information system:

- Too much “by hand” information in too many places – error prone
- Does WLCG need an IS? (my impression is “yes”)
 - But should be focused on as simple as possible for service discovery
 - Benchmark data should be separated
- Suggestion (and work done) to use AGIS for this
 - Needs agreement before we proceed further
 - Alternative is do nothing and let experiments gather info directly from GocDB, OIM etc

❑ Benchmarking: we need

- A real benchmark (HS06 or update) for:
 - Procurement, reporting, expressing requirements, etc
- A fast “calibration” benchmark to run e.g. at start of every pilot
 - Needed for understanding environment
 - Essential for opportunistic resources, or cloud uses
 - Ideal if we could agree a single such fast benchmark for everyone

Accounting, cont

□ Accounting

- Work has been done for cloud accounting in EGI
- Not clear how to publish accounting from commercial clouds or HPC (or non-pledged in general)
- Wallclock vs CPUtime reporting – not discussed

➤ We should review formally what each stakeholder needs from accounting

- Experiments, FA's, sites, etc
- What data should be gathered and reported?
- Today's accounting system has grown over 10 years – time to review what we want from it and how to manage it
 - Also to manage expectations of the data itself

Longer term – upgrades & HL-LHC timescale

Observations

- Probably a lack of clarity over what the situation for Phase 2 upgrades will be:
 - In terms of requirements – what is the real scale of the problem – need better estimates
 - What we can really expect from technology
 - An understanding of the real limitations of the system we have today
- We should also bear in mind that while we potentially need to instigate *revolutionary* changes in computing models, nevertheless we will have to face an *evolutionary* deployment
- Concerns over software and efficiency (in all aspects) will be a significant area of work
- Commonalities may be possible in new tools/services or next generation of existing
- Propose a number of activities to address some of these aspects

1) Definition of the upgrade problem

Set up a study group to:

□ Firstly:

- Establish and update estimates of actual computing requirements for HL-LHC, more realistic than previous estimates:
 - what are the baseline numbers for data volumes/rates, CPU needs, etc.?
- Build a realistic cost model of LHC computing, help to evaluate various models and proposals – this will be a key to guiding direction of solutions

□ Secondly:

- Look at the long term evolution of computing models and large scale infrastructure
 - Need both visionary “revolutionary” model(s) that challenge assumptions, and “evolutionary” alternatives
- Explore possible models that address (propose strawman models)
 - Today’s shortcomings
 - Try to use best of evolving technologies
 - Address expectations of how the environment may evolve
 - Large scale joint procurements, clouds, interaction with other HEP/Astro-P/other sciences
 - Possible convergence of (the next generation of) main toolsets

2) Software-related activities

□ Strengthen the HSF:

- “Improve software performance” –
 - Need to define what the goals and to define metrics for performance:
 - E.g. time to completion vs throughput vs cost
 - Continue concurrency forum/HSF activities – but try and promote more
 - And other initiatives like reconstruction algorithms etc
- Techlab
 - expand as a larger scale facility under HSF umbrella
 - Include support tools (profilers, compilers, memory etc)
 - Including support, training, etc
 - openlab can also help here
 - Should be collaborative – CERN + other labs
- Technology review
 - “PASTA” – reform the activity – make into an ongoing activity, updating report every ~2 years
 - Broad group of interested experts
 - Also under HSF umbrella – strongly related to the above activities
- What can be done about long term careers and recognition of software development

3) Performance evaluation/“modelling”

- ❑ Investigate real-world performance of today’s systems:
 - Why is performance so far from simple estimates of what it should be?
 - Different granularities/scales:
 - Application on a machine
 - Site level: bottlenecks, large-scale performance
 - Different scale sites, different workflows
 - Overall distributed system
 - At which level?
 - Are data models and workflows appropriate?
- ❑ Once we have a better handle of actual performance – can we derive some useful models/parameterisations etc?
 - Useful enough to guide choices of computing models – don’t have to be perfect or complete
 - This feeds into any cost models
- ❑ Small team in IT starting to work on this and consolidate existing efforts
 - Define a programme of work to look at current performance and concerns; define initial goals

4) Prototyping (demonstrators)

- ❑ Some specific prototyping of some of the ideas that arise from the above activities
- ❑ For example:
 - Data or storage management
 - Storage federations, caches rather than “SE”
 - Etc.
 - Optimisation of sites with little effort or expertise
 - “Site in a box” appliance,
 - What about cache, stage-out, etc
 - Others as ideas arise
- ❑ Common activity here would help to evolve into common solutions in production eventually

Summary

❑ Medium term

- A lot of work ongoing
 - Including other aspects not discussed in Lisbon (e.g. cost of operations)
- Useful to have (as discussed previously) a technical forum to coordinate all the activities?
 - Coordination: A chairperson, GDB chair, 1 per experiment (senior enough)
 - GDB and operations teams useful mechanisms for discussion/work

❑ Longer term

- 3 areas of work proposed
- Should be managed by the MB, also need to work towards a more concrete plan

❑ Prototypes/demonstrators

- A useful way to explore ideas and eventually converge on common solutions?