



WLCG Status Report

18th November, 2008
LHCC Referee
Meeting

Ian Bird
LCG Project Leader





Agenda

General status of project – see recent reports:

- Last quarterly report (June – Sept)
- Status report to Overview Board (27th Oct)
- Status report to C-RRB (11th Nov)
- Status report to LHCC open session (tomorrow)

Topics today:

- Planning for LHC shutdown (last meeting was only days after Sep 19)
- Report from C-RSG – how to proceed



Consequences of LHC shutdown

- The present shutdown of the LHC has a number of consequences for the planning of WLCG:
 - Capacities and Procurements for 2009
 - Software and service upgrades during the shutdown
 - (Re-)Validation of services for 2009 following changes



Capacities and procurements

- The WLCG MB has agreed that with the information currently available to us and the present understanding of the accelerator schedule for 2009:
 - The amount of data gathered in 2009 is likely to be at least at the level originally planned, with pressure to run for as long a period as possible this may be close to or exceed the amount originally anticipated in 2008 + 2009 together
 - The original planning meant that the capacity to be installed in 2009 was still close to x2 with respect to 2008 as part of the initial ramp up of WLCG capacity
 - Many procurement and acceptance problems arose in 2008 which meant that the 2008 capacities were very late in being installed; there is a grave concern that such problems will continue with the 2009 procurements
 - The 2009 procurement processes should have been well advanced by the time of the LHC problem in September
- The WLCG MB thus does not regard the present situation as a reason to delay the 2009 procurements, and we urge the sites and funding agencies to proceed as planned. It is essential that adequate resources are available to support the first years of LHC data taking.



Upgrade plans

- Since several software upgrades were postponed in anticipation of LHC start-up, we propose that the following changes are addressed in the coming months:
 - SRM – agreed list of “short term” changes; available by end 2008
 - FTS on SL4 (+available for SL5?) – deployment was postponed
 - WN on SL5 to be available for deployment
 - glexec/SCAS to support pilot jobs with identity changing
 - CREAM CE – make available in parallel to existing CE which is known to have scaling issues when there are many different users;
 - needs Condor_g client
 - WMS: must be able to submit to CREAM
- + a few other smaller changes ...

- Many of the above are deployments in parallel to existing production services and so non-disruptive



Re-validation of the service

- All experiments are continually running simulations, cosmics, specific tests (and have been since CCRC'08) at high workload levels – this will continue
- A full CCRC'09 in the same mode as 2008 is not regarded as useful
- But, we will perform specific tests/validations:
 - Service validation if software is changed/updated
 - Specific tests (e.g. throughput) to ensure that no problems have been introduced
 - Tests of functions not yet tested (e.g. Reprocessing/data recall at Tier 1s)
- Details of the test programme were discussed in the workshop last week
 - To be summarised and milestones proposed ...



Resource Scrutiny Group

- Process: scrutinise the requirements (normally) presented in Spring RRB:
 - Resource accounting for previous year
 - Use made of the resources
 - Overall requirements for the next year (and for +2 years)
 - Examine the match between requirements and pledges
 - Make recommendations in case of apparent under-funding
- This is the first scrutiny, requirements based on the TDRs (2005) + changes since as elements of the computing models have been tested
 - In particular ATLAS doubled the CERN (Tier 0 + CAF) request for 2009
 - Next reassessment of requirements should be with 6 months of beam experience
- Scrutiny took a standard set of assumptions on beam time
 - 3 months 2008 + 7 months 2009 → 7 months 2009 (after Sep 19); includes 1 month AA (0.9×10^7 s pp + 10^6 s AA)
 - Implemented simplified versions of the Computing Models for comparison (therefore cannot just take the numbers from RSG at face value)



Summary of conclusions

- Given that the ...
 - Experiment models are accurate only to ~10-20% before we have more experience
 - The RSG models are quite simplified compared to the experiment models
- ... the scrutiny agrees with the experiment requirements at this level (including the increased ATLAS request at CERN)
 - ALICE:
 - Reasonable; disk requests underestimated at Tier 1s, overestimated at CERN; tape request too large as assumed 2007 start
 - Unlikely that ALICE requests will be met; recommend ALICE make statement to LHCC how physics will be affected
 - ATLAS:
 - TDR model was optimistic for event sizes, data formats, etc. Parameters of the model should be re-examined.
 - Concern over event sizes and proliferation of data formats
 - Worry that resource needs for >-2010 may be hard to find
 - Agree that CERN request is justified but concern over de-emphasis of Tier 1s
 - CMS:
 - Good agreement with CMS model; CMS made good progress in understanding model details
 - RSG compute far less need for Tier 2 resources (but assumptions not agreed)
 - LHCb:
 - LHCb model is valid and solid
 - Slight reduction in disk requirement without 2008 data; similar conclusion for tapes



General recommendations

- Prudent to scrutinise usage after some months of data taking; also need a scrutiny for 2010 requirements asap for procurement cycles
- Scale of WLCG is unprecedented; uncertainties remain despite many tests being done → (particularly some) experiments should do risk analysis in order to cope with shortfalls
- Information about AA program is largely missing for some experiments – will have an impact on future needs
- Experiments should actively reduce raw event sizes and other derived data ... as detectors are better understood
- Recommend experiments keep models updated – clear that some assumptions are no longer realistic
- CERN resources – make clear separation between Tier 0/CAF and local analysis



RSG comments for LHCC

- Most experiments propose using increased trigger rates as compared to the ones stated in the TDR reviewed by the LHCC. We feel we are not sufficiently competent to review the need or convenience to do so.
- ALICE wants to increase substantially their amount of pp data; in particular they stress the benefit of acquiring data at 10TeV. We have not assessed these needs from the physics point of view and we do not know whether such lower energies will be available in the 2009 run or anytime in the future.
- One of our conclusions is to recommend that ALICE undertakes a full assessment of how their physics reach might be affected by requested computing resources not materializing.
- The event size has a very direct impact on the computing requirements. Some experiments, such as CMS, have made an effort to reduce the event sizes by establishing a reduction profile after startup. We believe that this example should be followed by all experiments.
- We take note of potential modifications of the computing models due to the use of different data formats serving the same purposes, not always well justified.
- The realization of the computing model for ATLAS seems to differ slightly from the implementation originally envisaged in the TDR for reasons discussed in the report. This implies, in particular, heavier demands on CERN resources. We believe these demands are largely justified, however.
- Cosmic data taking is now much emphasized by experiments; while it is clear that cosmics are extremely useful in commissioning for calibration, this data is by nature transient and it seems somewhat questionable to us to support substantial requests based on cosmic runs, but we do feel we have not sufficient insight to make a definite scientific judgement on this.



Experiment requirements - next steps

- The C-RSG report validates the experiment requirements – within the limits of how well the models are currently understood
 - No useful re-evaluation is possible without ~ 6 months real data taking experience
- ATLAS request at CERN is now doubled compared to pledged resources
- How should this be managed? Require guidance on priorities.
 - The existing CERN resources (and MTP) based on the understanding at the time of the TDRs
 - Increasing the ATLAS allocation would be at the cost of other experiments ...



Tier 0/CAF resource balance

CERN Tier0	Split 2009	ALICE	ATLAS	CMS	LHCb
CPU (kSI2K)	Offered	9000	4058	9800	1050
	Required	9000	7600	9800	1050
	% of Req.	100%	53%	100%	100%
Disk (Tbytes)	Offered	4200	265	200	991
	Required	4200	650	200	991
	% of Req.	100%	41%	100%	100%
Tape (Tbytes)	Offered	7300	5562	7300	2270
	Required	7300	8560	7300	2270
	% of Req.	100%	65%	100%	100%

CERN Analysis Facility	Split 2009	ALICE	ATLAS	CMS	LHCb
CPU (kSI2K)	Offered	2600	2562	3900	0
	Required	2600	5800	3900	0
	% of Req.	100%	44%	100%	100%
Disk (Tbytes)	Offered	300	1809	2300	0
	Required	300	3300	2300	0
	% of Req.	100%	55%	100%	100%
Tape (Tbytes)	Offered	0	651	2000	0
	Required	0	1130	2000	0
	% of Req.	100%	58%	100%	100%



Tier 1+2 Pledge Balance in 2009

- The table below shows the status at 17/11/08 for 2009 from the responses received from the Tier-1 and Tier-2 sites
 - The Total 2009 pledge from Russia is included but not the split across the experiments

	ALICE	ATLAS	CMS	LHCb	Sum 2009
T1 CPU	-49%	6%	-2%	2%	-12%
T1 Disk	-43%	-5%	-13%	-2%	-13%
T1 Tape	-50%	-7%	7%	6%	-13%
T2 CPU	-41%	0%	-3%	-33%	-8%
T2 Disk	-42%	-19%	38%	-	-1%



Pledge Balance 2008-2013

- The table below shows the global picture for 2008-2013, status as of 17/11/08. % indicates the balance between offered and required
- Some Federations have recently signalled a change to procurements for 2009, not supported by WLCG Management or Overview Boards

	2008	2009	2010	2011	2012	2013
T1 CPU	-5%	-12%	-11%	-15%	-20%	-26%
T1 Disk	-12%	-13%	-15%	-18%	-24%	-29%
T1 Tape	-13%	-13%	-16%	-22%	-24%	-23%
T2 CPU	-2%	-8%	-29%	-31%	-32%	-37%
T2 Disk	-12%	-1%	3%	-6%	-6%	-17%