



Summary of revised computing requirements

Stefano Belforte
For the
CMS Collaboration

- CERN finalized LHC schedule with pilot run at end of 2007 and full physics run in 2008
- WLCG met with LHC experts to reach an estimate of machine beam time to be used as base for common planning by the experiments.

Year	Energy	Proton Luminosity	Physics beam time
2007	450 + 450 GeV	5×10^{30}	Protons - 26 days at 30% overall efficiency → 0.7×10^6 seconds
2008	7 + 7 TeV	0.5×10^{33}	Protons - starting beginning of July 4×10^6 seconds Ions - end of run - 5 days at 50% overall efficiency → 0.2×10^6 seconds
2009	7 + 7 TeV	1×10^{33}	Protons - 50% better then 2008 → 6×10^6 seconds Ions - 20 days at 50% efficiency → 1×10^6 seconds
2010	7 + 7 TeV	1×10^{34}	TDR targets: Protons - 1×10^7 seconds Ions - 2×10^6 seconds (1×10^6 in CMS TDR)



The process - 2



- In July, at WLCG request, CMS Computing Coordinators started to look at CMS computing model in view of revised schedule. Found that
 - Computing Model, carried on into Computing TDR, addresses a nominal LHC years in the context of running a well established physics program: no debug/backup triggers, no cosmics etc.
 - Understanding of startup issues in CMS is still evolving
 - Computing Model figures still have large uncertainties. Although initial indication from new CMS software are promising, it is too early to claim that old estimates will hold.
 - Could not make a bottom up revision more solid than C-TDR
- Guidance from the Collaboration at large was sought via the CMS Spokesman, high level guidelines were obtained
- In September a proposal was circulated to CMS Collaboration
- Approved by CMS Management Board and Collaboration Board at CMS week on September 22
 - Document sent to WLCG, attached to meeting agenda
- Presented to CMS and WLCG RRB in October
- Present to you in next slides



The Computing Model review



- No reason has been found to revise the basic parameters of the existing computing model. It still provides as good an estimate as we can make before being able to base requests on experience. The Computing Model has been reviewed extensively in past years by several external referee committee, no major issue was indicated.
- The amount of beam time expected in 2008 went down. This parameter had a large uncertainty anyhow, which was expected to compensate others. While this moved in the direction to reduce needs, we still have no indication about how and how much other relevant figures will move.
- The instantaneous HLT accept rate is limited in the model by the ability of the Tier-0 centre to carry out Prompt Reconstruction. Consideration of this and other factors led to the adoption of a nominal 150Hz HLT rate. This decision implied a loss of useful physics data, which was tensioned against the required computing resources.

- Revised start-up plan for 2008 has 40% beam time. This opens the possibility of an increase in HLT rate even with the foreseen reduction in the CMS Tier-0 farm, since the 20-days input buffer allows Tier0 to operate in catch-up mode over short periods.
- Improvement in HLT rate from nominal 150Hz is strongly motivated:
 - Compensate reduced running time preserving size of physics-reach data sample. Compare to CDF at lower energy and luminosity running at 100Hz to tape after years of fine tuning
 - Lower HLT thresholds to increase efficiency in channels with b, top, tau in final state and/or those depending on pre-scaled jets. This will allow faster understanding of Standard Model processes, pre-requisite to any discovery claim.
 - Can take a substantial rate of technical minimum-bias and calibration triggers to faster and better understand the detector and trigger, especially as running conditions change rapidly in 2008
 - Reduce pressure on HLT performance during early commissioning and increase the safety margin on the overall trigger system acceptance.



The new requirements



- LCG is now planning to provide CMS Tier-0 resources at the level of about 80% of the CMS request for 2008. Sizing the Tier0 in line with this implies a maximum HLT rate of $\sim 300\text{Hz}$ (450MB/sec) in 2008. HLT readout may be configured with capacity of 1GB/s, leaving sufficient safety margin.
- We therefore scale the other parts of the computing system accordingly to obtain 2008 requirements.
- For 2007 we foresee resources to be available as ramp up of computing centers to 2008 sizes and keep estimate as simple fractions of 2008 figures as in Computing TDR.
- In 2009 we foresee an HLT rate of 200Hz giving a roughly similar reduction in Tier-0 resources.
- 2010 running converges with C-TDR parameters (Tier0 keeps up with HLT in real time), with a trigger rate of 150 Hz and 10^7 sec of LHC running
- Figures may be revised as we know more from new CMSSW and develop a fuller understanding of 2008 CMS running plans

		2007		2008			2009		2010	
		Request	%TDR	Request	%TDR	Pledge	Request	%TDR	Request	%TDR

Tier-0 + CAF	CPU	1.9 + 1.9	82%	3.9 + 3.8	82%	3.9 + 3.8	6.1 + 5.8	84%	10.6 + 11.5	91%
	Disk	0.1 + 0.4	82%	0.3 + 1.3	82%	0.3 + 1.3	0.3 + 2.0	82%	0.5 + 3.3	88%
	Tape	0.8 + 0.4	76%	3.6 + 1.5	76%	3.6 + 1.5	6.7 + 2.7	78%	10.9 + 4.1	87%

Tier-1	CPU	6.2	82%	12.4	82%	12.1	16.9	82%	36.9	91%
	Disk	1.7	81%	5.6	81%	5.8	8.5	81%	13.7	87%
	Tape	3.0	79%	13.1	79%	10.3	23.5	80%	36.6	87%

Tier-2	CPU	7.6	79%	15.2	79%	18.3	25.6	79%	45.2	88%
	Disk	1.3	85%	4.2	85%	4.8	8.4	85%	13.3	90%

Units are MSI2k, PB



One liners summary



- No reason to revise basic assumptions
- Too early to reduce uncertainties on numerical estimates
- Current model was constrained by Tier0 capacity
- 2008 running will be reach in physics potential
- 2008 will have large (unforeseen) commissioning/debugging needs
- Nee to make the most of a very expensive detector and machine

- Therefore:
 - Log CMS data to tape as fast as it can be done
 - Be ready to process those along the Computing Model lines
 - For 2008 a rough 80% reduction across the board w.r.t. C-TDR
 - New requests in line with current pledges
 - May review later on in 2007 (summer ?)