Updated CMS Computing Resource Requests 2007 – 2010

Final - 5/10/06

Summary

In response to the revised LHC start-up scenario, CMS has reviewed its overall computing requirements in the years 2007 – 10. No substantial reason has been found to revise the basic parameters of the existing computing model. An exception is that the new LHC schedule presents the opportunity to improve the high-level trigger accept rate whilst simultaneously reducing the required computing resources. This change has the potential to allow faster and better understanding of the performance of the CMS detector, allow faster understanding of the Standard Model processes and extend both the physics reach of the experiment, and the margin of safety in the computing model. We present in this document the revised resource estimates for Tier-0, Tier-1 and Tier-2 centres.

Current status

The primary input parameters to the CMS Computing Model are event rates and sizes, along with estimates of the required processing times for simulation, reconstruction and analysis. Whilst the new CMS software framework is still under development, progress to date gives us reasonable confidence that the targets set down in the computing model are achievable. Promising performance for Tier-0 functions has recently been demonstrated during the first part of the CSA06 computing challenge. We therefore see no reason to adjust our estimates of overall storage or CPU requirements per event. We note that a sensible margin of safety on all such estimates is still required, before the framework reaches full maturity.

The 'operational' parameters of the model (efficiencies of computer centres, size of disk caches, etc) have recently been discussed in the context of the LCG T1/T2 Working Group. Once again, no hard reason to adjust our basic assumptions was identified. Indeed, it became clear that a de-scoping of the computing model to meet existing Tier-1 resource pledges would not be possible without introducing significant risk to the throughput and performance of the computing system for the first year of LHC running. The data-processing efficiencies assumed in the computing model are broadly in line with our experience of data challenges over several years. The T1/T2 process also highlighted differences in assumptions between experiments regarding CPU efficiencies at Tier-1 and Tier-2 centres. In particular other experiments are using lower efficiencies than CMS assumption of 75%, leading to increasing the stated analysis CPU requirement relative to CMS. However, CMS at this point did not change its assumptions.

The calculation of the revised computing requirements therefore assumes no change to the parameter values given in the Computing TDR, with the exception of the raw input rate to the computing system – i.e. the HLT accept rate.

Changes in the Computing Resource Requests

The instantaneous HLT accept rate is limited by the ability of the Tier-0 centre to carry out prompt reconstruction of the event stream, and to some extent by the overall dataflow within the CMS readout systems. A consideration of these and other factors led to the adoption of a nominal 150Hz HLT rate in the Computing TDR. This decision implied a loss of useful physics data, which was tensioned against the required computing resources.

The revised start-up plan for 2008 entails a reduction in LHC duty cycle to around 40% of the assumptions used in the experiment Computing TDRs. This opens the possibility of an increase in HLT rate even with the foreseen reduction in the CMS Tier-0 farm, since the Tier-0 may operate in 'catch-up' mode during short non data-taking periods (an input buffer of 20 days is foreseen).

An improvement in HLT rate from the nominal 150Hz is strongly motivated by several factors:

- The reduction in LHC running time may be partially offset by roughly preserving the size of the data sample. We note that our entire sample will remain 'physics-rich'. For comparison, at the Tevatron, at lower luminosity and energy, CDF is running at a rate-to-tape of 100Hz (after a long period of fine-tuning of detector and trigger), and foresees to increase that rate to 350Hz with the planned DAQ upgrade.

- The lowering of HLT physics object thresholds while increasing the data taking rate also increases the efficiency for many important channels in both discovery and standard model physics – notably, those with b, t and τ in the final state, and/or those that depend on pre-scaled jet triggers. This will allow faster understanding of the Standard Model processes which is absolutely necessary before any discoveries can be claimed.

- The ability to take a substantial rate of technical, minimum-bias, and calibration triggers without impact upon the physics sample will greatly enhance CMS' understanding of detector and trigger performance. This will be particularly important in a start-up scenario where the LHC running conditions (bunch luminosity, bunch spacing, beam lifetime) are changing rapidly with ongoing machine development during 2008.

- The performance required of the HLT system during the early commissioning phase is somewhat reduced. Moreover, the required CPU burden on the HLT is reduced in the long term, potentially increasing the safety margin on the Level-1 trigger accept rate and/or the acceptance of the entire trigger system.

LCG is now planning to provide CMS Tier-0 resources at CERN at the level of about 80% of the CMS request for 2008. Sizing the Tier-0 centre in line with this plan implies a maximum HLT rate of around 300Hz (450MB/s data rate) for 2008. The readout system may be configured with the capacity to output event data at 1GB/s, leaving sufficient margin for safety. We therefore scale the other parts of the

computing system accordingly, allowing us to take data at rates of up to 300Hz, and providing an increased safety factor in the event of unforeseen changes to other computing model parameters.

In 2009, we foresee an HLT rate of 200Hz, giving a roughly similar reduction in Tier-0 resources. The 2010 running converges with the Computing TDR parameters, with a trigger rate of 150Hz and 10^7s of LHC running.

Revised Resource Requests

The revised request for Tier-0/CAF, Tier-1 and Tier-2 computing resources is summarised in the table overleaf, along with the reduction from the TDR requirements, and the currently pledged resources.

		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