

Summary of CMS throughput requirements on UK T2 – 2008

DRAFT 1.0 – DMN 16/1/07

Scale of T2 centres

Within the GridPP3 planning for 2008, CMS has been allocated substantive resources at four sites, as indicated below (numbers from the GridPP3 proposal). We assume that access to resources at all other UK sites will be possible on a load-sharing basis, but that these sites will typically be used for Monte Carlo production as opposed to analysis; the throughput requirements on such sites are therefore far more relaxed.

Institute	Frac	CPU (kSI2k)	Disk (TB)	% nominal T2
Bristol	12.5%	225	50	25%
Brunel	7.5%	135	30	15%
Imperial	67.5%	1215	270	135%
RAL PPD	12.5%	225	50	25%

WAN network requirements

The CMS CTDR states that a nominal T2 will require an overall WAN throughput of 5TB/day (roughly 1Gb/s at 50% efficiency). This is almost entirely download of samples from T1 centres. However, this figure is an average over centres concentrating on Monte Carlo production and analysis. We therefore prorate the figure by 100% to allow for the use of the UK centres almost exclusively for analysis.

T2 traffic is by nature extremely bursty. In our network planning, we assume the traffic to be ‘maximally bursty’ – that is, a T2 centre will transfer at maximum rate allowed by the WAN connection with a duty cycle corresponding to the average throughput. In addition, we specify a connection speed of at least 1Gb/s for any CMS T2 wishing to take part in analysis activities, ideally sustainable to any CMS T1 centre.

We therefore conclude that CMSUK T2 sites other than Imperial require the minimal 1Gb/s connection speed in 2008, which will be substantially occupied for Bristol and RAL PPD. Imperial requires around a 1.5Gb/s connection. The T2 should be capable in principle of pulling data at full rate from all CMS T1, though clearly the connection to RAL is the short-term priority.

Local network requirements

Each T2 centre should clearly be internally arranged to sink at least the WAN bandwidth to disk on a continuous basis, whilst also sustaining traffic to and from the CPU farm.

The CTDR specifies two parameters for the disk-CPU connection at a nominal T2. Firstly, that each node (assuming a node performance of ~8kSI2k) should be connected with 100MB/s capability from disk. It is not expected that all nodes will

simultaneously make use of such rates, but this allows for a mix of high- and low-IO analysis jobs on the farm. Secondly, that a nominal T2 should be capable of sustaining 1GB/s from disk to the farm on a continuous basis. However, this implies a mixture of Monte Carlo and analysis use, whereas the UK T2 are intended substantially for analysis. We therefore suggest that this is prorated to 2GB/s, and then scaled by the T2 size.

Non-CMSUK Tier-2

UK T2 sites other than those specified above may be used for Monte Carlo production or low-IO analysis, and therefore require a nominal network connection to the RAL T1 of 10MB/s, principally for upload of produced data.

Summary

Institute	WAN throughput	LAN throughput
Bristol	100MB/s	500MB/s
Brunel	100MB/s	300MB/s
Imperial	150MB/s	2700MB/S
RAL PPD	100MB/s	500MB/s
Others	10MB/s	100MB/s

2009 and beyond

Network requirements are expected to scale as the size of the total CMS dataset, rather than as the resources at any given centre. A steep rise in network requirements is therefore to be assumed from 2009 onwards, particularly as high luminosity running commences from 2010. We would anticipate the total international traffic to/from the UK for CMS analysis to exceed 10Gb/s aggregate by 2010.