The Megatable

Christoph Eck
IT Department
CERN

GDB Meeting November 8, 2006

What is the Megatable

- The Megatable is a projection of the Computing TDRs on the current status of resource pledges for a given experiment. Taking account of:
  - The Tier-1s and Tier-2s with prioritised resources for this experiment and the amount of computing capacity pledged by them.
  - The amount of data to be transferred between given Tier-1s and Tier-2s for this experiment and the bandwidth required for this.
  - The amount and type of storage required by Tier-2s at a given Tier-1 for this experiment.
  - The amount and type of storage required at a Tier-1 to fulfil its Tier-1 functionality for this experiment.
  - The bandwidth requirements between a given Tier-1, the Tier-0 and other Tier-1s.
- In short, everything a centre would need to know in order to make an experiment happy.

How Does it Look? (Tier-2s)

		Pledge for 2008		FZK						
Tier-2	Experiment	CPU	Disk	T2=>T1		T1=>T2		Storage for T2 TByte		
		kSI2k	TByte	MByte/s aver.	MByte/s peak	MByte/s aver.	MByte/s peak	Tape1-Disk0	Tape1-Disk1	Tape0-Disk1
Tier-2s listed in the WLCG MoU										
Australia	ATLAS	300	300							
Belgium	CMS	1050	270			12.0				
China	ATLAS	500	200							
China	CMS	500	200			8.0				
Czech R.	ALICE	200	75	4.9	4.9	2.6	2.6	76.1	23.3	113.9
Czech R.	ATLAS	800	405			31.7				
Finland	CMS	900	200			9.0				
CC-IN2P3	ALICE	265	46							
CC-IN2P3	ATLAS	2595	217							
CC-IN2P3	CMS	738	192			3.0				
FR/GRIF	ALICE	192	56							
FR/GRIF	ATLAS	1056	308							
FR/GRIF	CMS	480	140			2.0				
FR/GRIF	LHCb	192	56							
FR/LPC	ALICE	90	8							
FR/LPC	ATLAS	150	12							
FR/LPC	LHCb	60	5							
FR/Suba	ALICE	200	20							
DESY	ATLAS	700	340			27.7				
D/CMSF	CMS	1350	300	6.0		32.0		242.0		29.0
D/GSI	ALICE	660	200	16.2	16.2	8.4	8.4	251.0	76.7	107.9

How Does it Look? (Totals)

TOTALS			T0=>T1	T2=>T1	T1=>T1	T1=>T2	T1<=T1	Storage for T2 TByte			Storage for T1 TByte			
	Total Tape Tbyte	Eff. Disk Tbyte	MByte/s	MByte/s aver.	MByte/s in	MByte/s aver.	MByte/s out	Tape1-Disk0	Tape1-Disk1	Tape0-Disk1	Tape1-Disk0	Tape1-Disk1	Tape0-Disk1	Cache-Disk
ALICE	1709.8	1423.1	24.0	56.9	29.3	29.5	15.3	879.0	268.7	436.0	466.5	95.5	195.9	
ATLAS	637.5	949.2	63.8	4.2	121.7	81.2		19.2	11.4	13.4	327.6	279.4	360.3	
CMS	1264.0	725.7	37.0	15.0	153.0	141.0		592.0	0.0	72.0	672.0	0.0	436.0	
LHCb	114.7	181.0	6.3	0.2	20.4	11.4	18.1	0.0	6.7	0.0	0.0	108.0	12.0	
SUM	**3726.0**	**3279.1**	**131.1**	**76.3**	**324.4**	**263.1**	**33.4**	**1490.1**	**286.8**	**521.4**	**1466.1**	**483.0**	**1004.2**	**0.0**
					With 70% Disk Efficiency				**409.7**	**744.8**		**689.9**	**1434.6**	**0.0**
					Total Storage Requ. Tape			**1777**			**1949**			
					Total Storage Requ. Disk					**1155**				**2125**
					Tape and Disk Pledges					**Tape**	**3472**		**Disk**	**3254**
					Balance						**-254**			**-25**

Status of Data Collection

- Requirement data from ALICE and LHCb are in the table.
- ATLAS and CMS data are ready, but need the final blessing by their management.
- I plan to send the Megatable to all sites early next week.

- All data are updated to the latest requirement tables of the experiments.
- Pledges correspond to the current figures in the MoU.
  - CERN Tier-0 and CAF pledges are the ones shown to the C-RRB.
- Experiments use these pledge numbers differently, as explained on the next slide.

Use of Pledge Numbers

- ATLAS and LHCb use the current pledges to evaluate the tasks they can allocate to a given site.
  - As some sites pledged capacities are not at all aligned with the CPU to disk etc. relations required by the computing models of these experiments, therefore not all pledges can be efficiently used by ATLAS and LHCb.
- ALICE requirements take account of the size variations of different sites but assume that the overall pledge level for ALICE will still increase to satisfy their overall needs.
- CMS has decided to align the pledges to their computing model and then to evaluate the tasks like ATLAS and LHCb.
  - They calculate the cost of the total capacity pledged by a given site in a given year and then create an "aligned" pledge costing the same amount of money.
  - For CMS sites I will show the original plus the aligned pledges in the Megatable.

Actions

- All sites should study the Megatable and decide:
  - Are the Tier-2 to Tier-1 relations assumed in the table acceptable to them?
  - Can they provide the required resources to the experiments asking for them?
  - Will they have the required network bandwidths?
  - Will they have the resources to manage the in some cases very high number of FTS connections required from them?
- If they can say YES to all these questions, they should tell me and Yves Schutz for ALICE, Roger Jones for ATLAS, Dave Newbold for CMS and Nick Brook for LHCb.
- If they have problems, they should contact the same people as soon as possible to start fixing these.
- This GDB will set the date for a first discussion on the outcome of the Megatable exercise.

Questions

- Shall we invite ATLAS and LHCb to "renormalise" pledges in the same way as CMS?
- CMS has calculated the tape bandwidths required by them at the Tier-1s.
  - Shall we publish these numbers in the Megatable?
  - Shall we ask the other experiments to do the same?
- What should be the deadline for all sites to have given their feedback on the Megatable?