



Resources in 2007/8

- Inputs to resource planning process:
 - WLCG pledges of CPU, disk, tape
 - (Often informal) understanding of network capability
 - The CMS computing model - with updates
- Outputs (*partially* expressed for 2007 in the 'megatable')
 - Nominal network traffic between sites
 - 'Balanced' resources at sites (for ~constant cost)
 - Some idea of the internal dataflow within sites
- Today:
 - T1/T2 resources required / planned in September 2008
 - First look at the picture in September 2007
- Ongoing:
 - Comparison of network requirements with available resources



Assumptions 2007/8

- LHC running period
 - 2008: ~100d @ 50% efficiency; 2007: ~30d @ 30% efficiency
- Data rates, event sizes, etc
 - 300Hz average rate
 - 1.5/2MB RAW/simRAW; 0.25/0.4MB RECO/simRECO; 50kB AOD
- Simulated data
 - 2008: 1-to-1 ratio with real data -> 1200Mevts
 - 2007: Assume 50% of 2008? But possibly not all retained
- Data movement
 - T0 -> T1 flow at ~constant rate for LHC running period
 - T1 <-> T1 replication over 14 days
 - T2 refreshes data samples each 30 days on average
- These basic numbers need to be discussed and checked
 - e.g. event sizes, processing times



Data sizes + Associations

Centre	Streams	Associated T2
FZK	5	German T2, Poland, Switzerland
IN2P3	6	French T2, China, Belgium
PIC	2	Spain T2, Portugal
CNAF	7	INFN T2, Hungary
ASGC	5	Taipei, India, Pakistan
RAL	5	UK T2, Estonia, Finland
FNAL	20	US T2, Brazil
CERN		Russia, Ukraine

- At least one site is missing from this list
- Reminder of CMS T1-T2 concept
 - Unique data at each T1 implies 'many-to-many' topology for RECO access
 - Presumably extend this to AOD access sooner or later (the 'mesh')



T1 Resources 2008

	Plg. CPU	Plg. Disk	Plg. Tape	#streams	CPU	Disk	Tape	Tr Buf
	kSI 2k	TB	TB		kSI 2k	TB	TB	TB
FZK	1200	650	900	5	1209	683	1280	25
IN2P3	1490	780	1180	6	1683	785	1686	29
PIC	760	350	835	2	650	442	641	19
CNAF	1925	875	735	7	1599	849	1660	29
ASGC	1530	675	585	5	1205	684	1258	24
RAL	1330	620	1280	5	1436	700	1434	26
FNAL	4256	1986	4700	20	4625	1942	4788	65
CERN					686	301	649	20
Total	12491	5936	10215	50	12407	6087	12747	218

→ Resources have been slightly rebalanced

- Typically, more tape, less CPU
- Centres are already well-balanced typically

→ Overall, 2008 T1 requirements are met

- Assuming we meet our targets for efficiency factors



T1 Networking 2008

	FZK	IN2P3	PIC	CNAF	ASGC	RAL	FNAL	CERN
OPN in	90	144	157	145	150	147	175	158
OPN out	54	129	49	123	91	110	355	438
T2 in avg	9	30	12	22	18	26	60	18
T2 out avg	95	232	122	230	182	194	600	138
T2 out peak	838	1902	680	708	1276	876	7582	456
OPN prov	2	2	3	2	3	2	6	7
T2 Prov.	9	21	8	9	15	11	82	6

- Remember to apply a factor of 2 - 3 for overhead / catch-up
- T1-centric model so far
 - This is judged to be the hardest point in the system
 - The original purpose of the megatable
- Need to move to 'real' network planning
 - Understand the site-to-site dataflows, particularly T1 -> T2
 - How does this map onto the real network infrastructure?
 - What is the appropriate target for T1 data serving capability?
- Hot spots: T1-T1 (but on the OPN); transcontinental T1-T2



T2 Resources 2008

	Plg. CPU	Plg. Disk		Plg. CPU	Plg. Disk
	kSI2k	TB		kSI2k	TB
China	500	200	UK	1500	340
Finland	900	200	Ukraine	240	30
CC-IN2P3	738	192	US/Caltech	1000	200
FR/GRIF	480	140	US/Florida	1000	200
D/CMSF	1350	300	US/MIT	1000	200
I/TIFR	450	300	US/Nebra	1000	200
INFN	1750	589	US/Purdue	1000	200
Pakistan	795	200	US/S Diego	1000	200
Poland	250	8	US/Wisco	1000	200
Portugal	225	42	Belgium	1050	270
RU/RDIG	1438	555	Brazil	500	27
SP/CMS	1036	276	Hungary	400	20
CH/CSCS	230	80	Estonia	200	100
Taipei	200	35			
			Total	10342	3117

- The 'odd' breakdown comes from the LCG tables...
 - Need to break this down per site for convergence with CMS planning
- Overall, around 20% more resources w.r.t CTDR
- Not clear that a 'uniform model' of T2 functionality is optimal



T1 Resources 2007

	Plg. CPU	Plg. Disk	Plg. Tape	CPU	Disk	Tape	Tr Buf
	kSI2k	TB	TB	kSI2k	TB	TB	TB
FZK	337	181	273	419	215	469	10
IN2P3	587	327	495	469	241	539	11
PIC	253	117	278	193	129	206	8
CNAF	840	420	350	429	262	545	10
ASGC	797	405	243	374	211	437	10
RAL	300	170	300	401	213	456	10
FNAL	1792	700	300	1240	621	1568	24
CERN				388	97	275	11
Total	4906	2321	2240	3524	1893	4220	83

→ First conclusions:

- *If we wished to keep most of our MC data, we are short on tape*
 - This is not a trivial question, some sites plan robot upgrades in 2008
- Of course, we are heavy on CPU

→ Actions:

- These pledged numbers should be checked and discussed



T2 Resources 2007

	Plg. CPU	Plg. Disk			Plg. CPU	Plg. Disk
	kSI2k	TB			kSI2k	TB
China	250	25		UK	676	95
Finland	400	100		Ukraine	180	20
CC-IN2P3	282	87		US/Caltech	470	80
FR/GRIF	168	42		US/Florida	470	80
D/CMSF	900	250		US/MIT	470	80
I/TIFR	300	200		US/Nebra	470	80
INFN	875	236		US/Purdue	470	80
Pakistan	410	80		US/S Diego	470	80
Poland	216	7		US/Wisco	470	80
Portugal	150	28		Belgium	600	150
RU/RDIG	1119	393		Brazil	250	13
SP/CMS	414	110		Hungary	200	10
CH/CSCS	72	24		Estonia	100	50
Taipei	150	14				
				Total	5707	1596

→ Around 50% of 2008 resources available in Sep 2007

→ Are these numbers correct?

- In many cases, 'back calculated' from the 2008 CMS fraction
- We know that plans are changing even now



Summary

→ 2008 resources

- We seem to be in good shape
- Some rebalancing of site resources may be desirable
- The network requires much further study and testing

→ 2007 resources

- Again, appear to be in good shape
- Around 50% of T2 capacity - fits CMS goals for 2007
- T1 centres may need significant balancing to meet 2007 needs

→ Next steps

- We should review our assumptions, data placement model
- How close to reality are the 2007 pledges?
 - Many sites are purchasing this hardware *now*
- **Feedback from sites** on resource balance, network is required
- Network requires significant work to give us confidence