

Testing CMS Computing Model in DC04

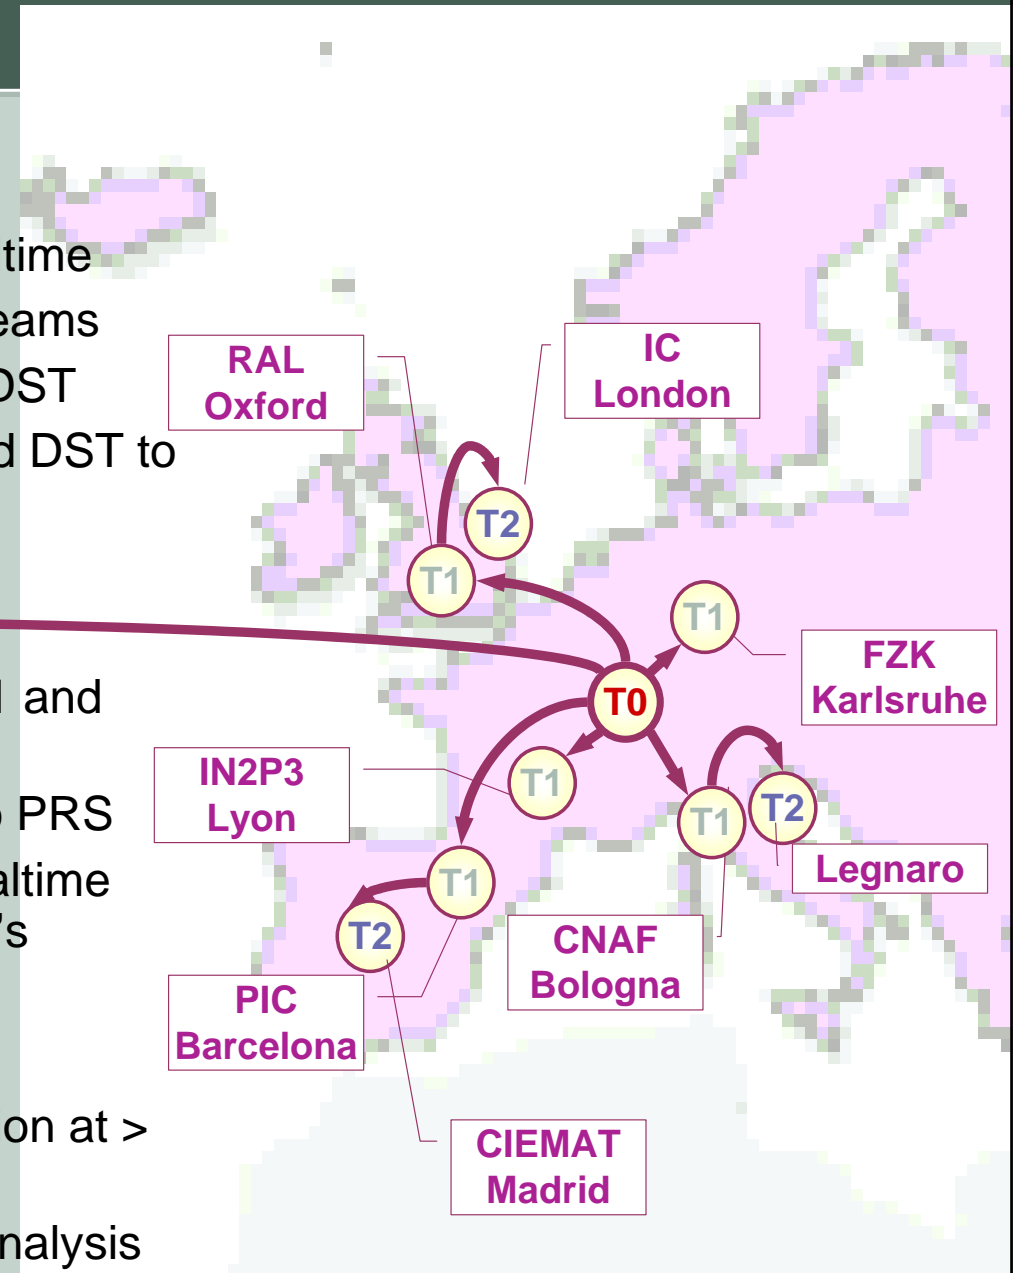
- Focus (for now) on organized (CMS-managed) data flow/access
- Functional DST with streams for Physics and Calibration
 - DST size ok, almost usable by “all” analyses; development in progress
- Tier-0 farm reconstruction
 - 500 CPU. Ran at 25Hz. Reconstruction time within estimates.
- Tier-0 Buffer Management and Distribution to Tier-1's
 - TMDB: a CMS-built Agent system communicating via a Central Database.
- Tier-1 Managed Import of Selected Data from Tier-0
 - TMDB system worked.
- Tier-2 Managed Import of Selected Data from Tier-1
 - Meta-data based selection ok. Local Tier-1 TMDB ok.
- Real-Time analysis access at Tier-1 and Tier-2
 - Achieved 20 minute latency from Tier 0 reconstruction to job launch at Tier-1 and Tier-2
- Catalog Services, Replica Management
 - Significant performance problems found and being addressed

DC04 Data Challenge

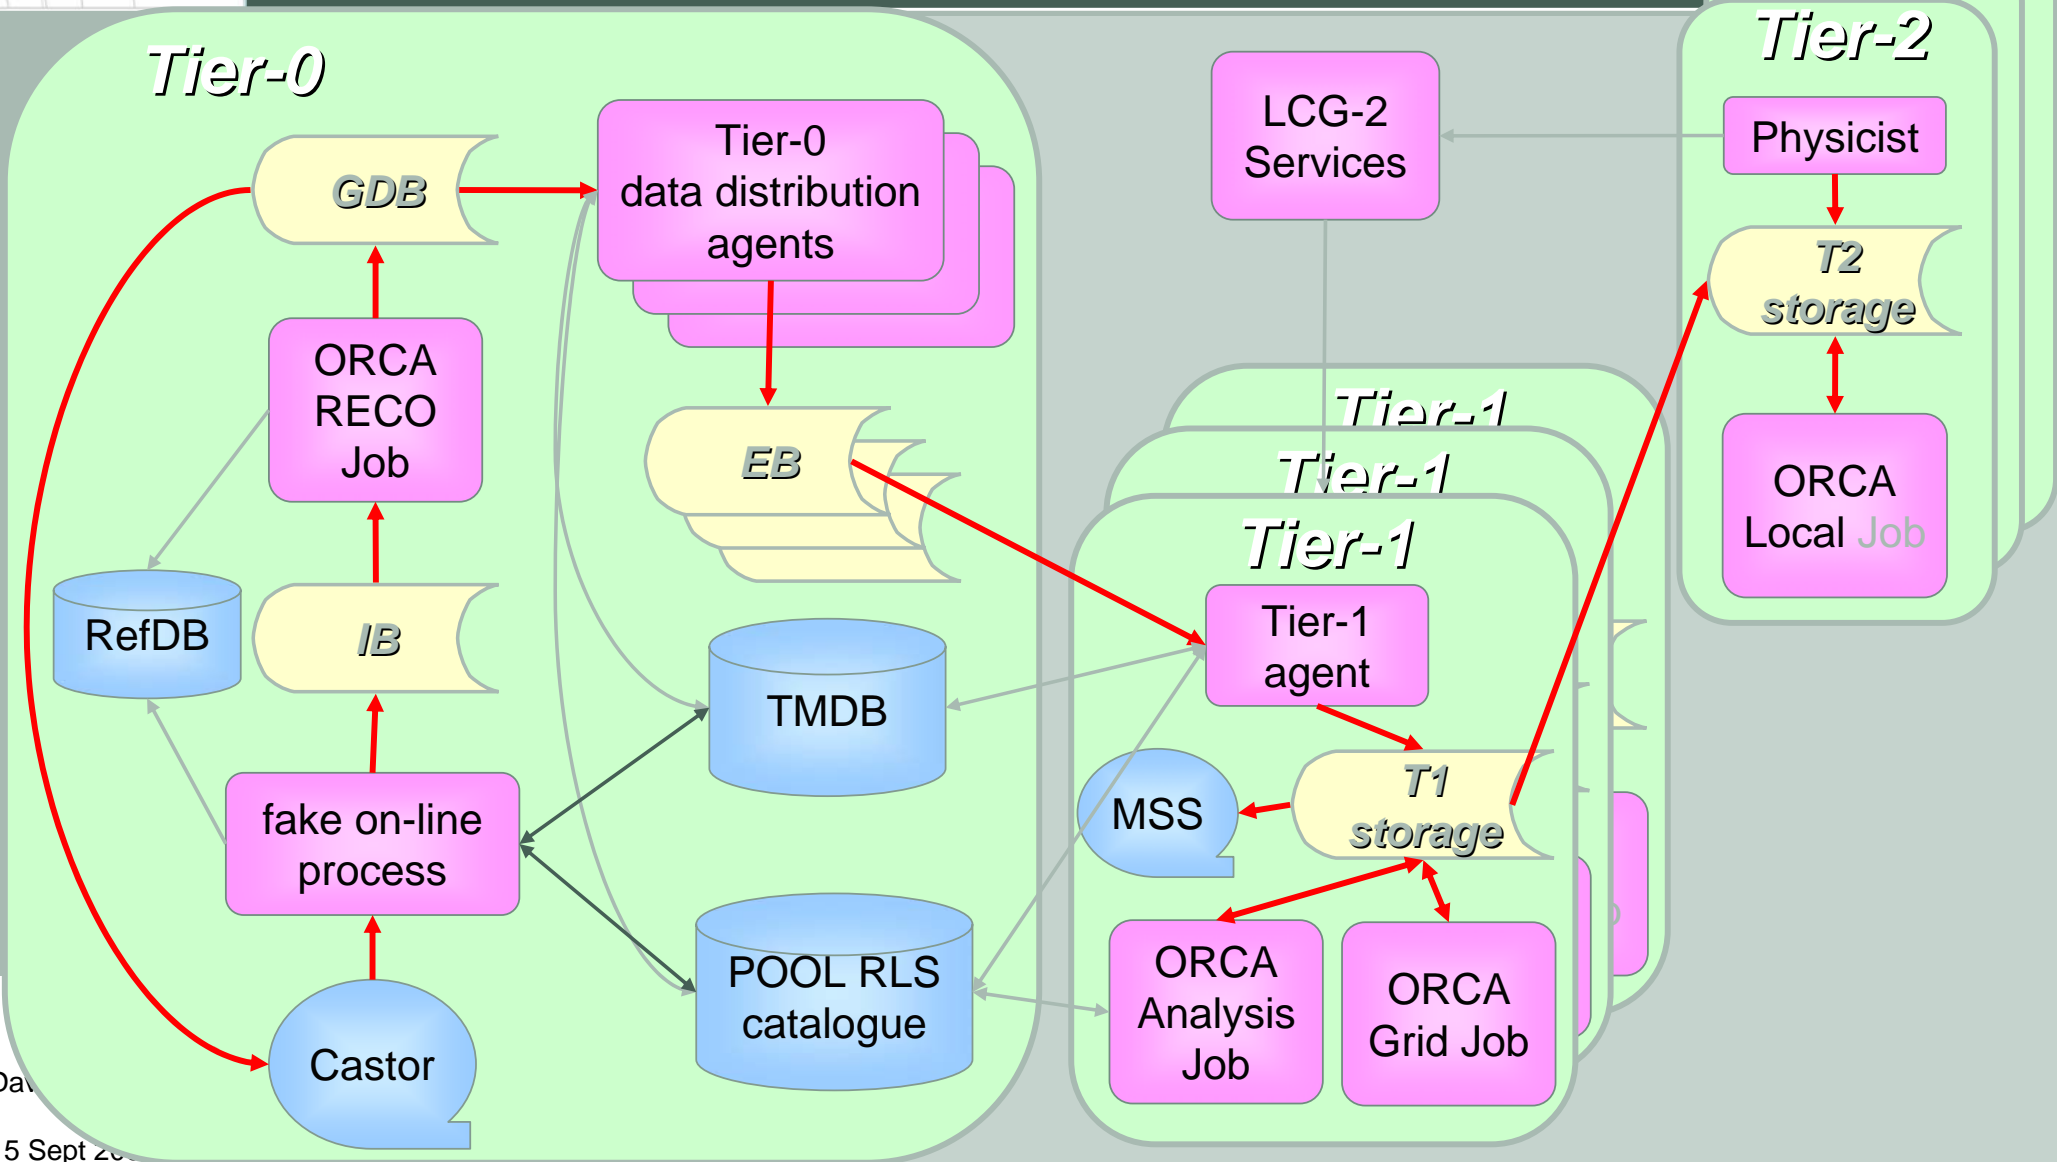
- T0 at CERN in DC04
 - 25 Hz input event rate
 - Reconstruct quasi-realtime
 - Events filtered into streams
 - Record raw data and DST
 - Distribute raw data and DST to T1's

- T1 centres in DC04
 - Pull data from T0 to T1 and store
 - Make data available to PRS
 - Demonstrate quasi-realtime “fake” analysis of DST's

- T2 centres in DC04
 - Pre-challenge production at > 30 sites
 - Modest tests of DST analysis




DC04 layout



DC04 Processing Rate

T0 Events Per Time




- Processed about 30M events
 - First version of DST code was not so useful for Physicists
 - Post-DC04 2nd version ready for production in next weeks

- Generally kept up with SRM based transfers. (FNAL, PIC, CNAF)

- Got above 25Hz on many short occasions
 - But only one full day above 25Hz with full system

- RLS, Castor, overloaded control systems, T1 Storage Elements, T1 MSS, ...

Event Processing Rate




Da
15

Aspects of DC04 involving LCG-2 components

- register all data and metadata to a world-readable catalogue
 - RLS
- transfer the reconstructed data from Tier-0 to Tier-1 centers
 - Data transfer between LCG-2 Storage Elements
- analyze the reconstructed data at the Tier-1's as data arrive
 - Real-Time Analysis with Resource Broker on LCG-2 sites
- publicize to the community the data produced at Tier-1's
 - Not done, but straightforward using the usual Replica Manager tools
- end-user analysis at the Tier-2's (not really a DC04 milestone)
 - first attempts
- monitor and archive resource and process information
 - GridICE

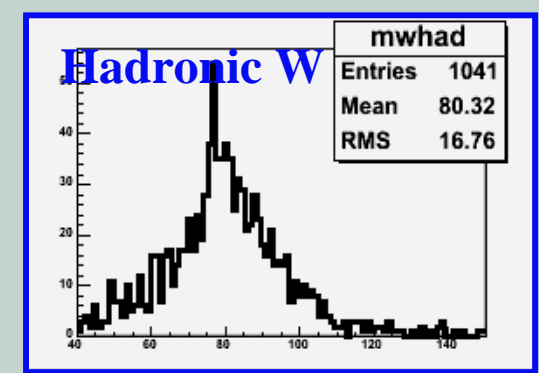
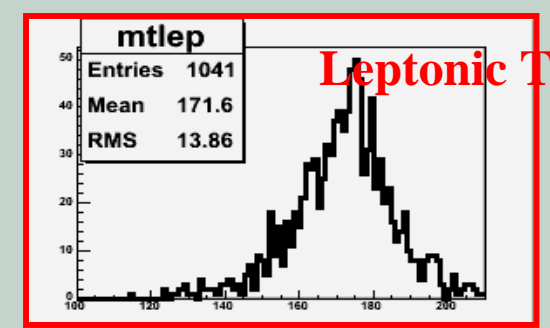

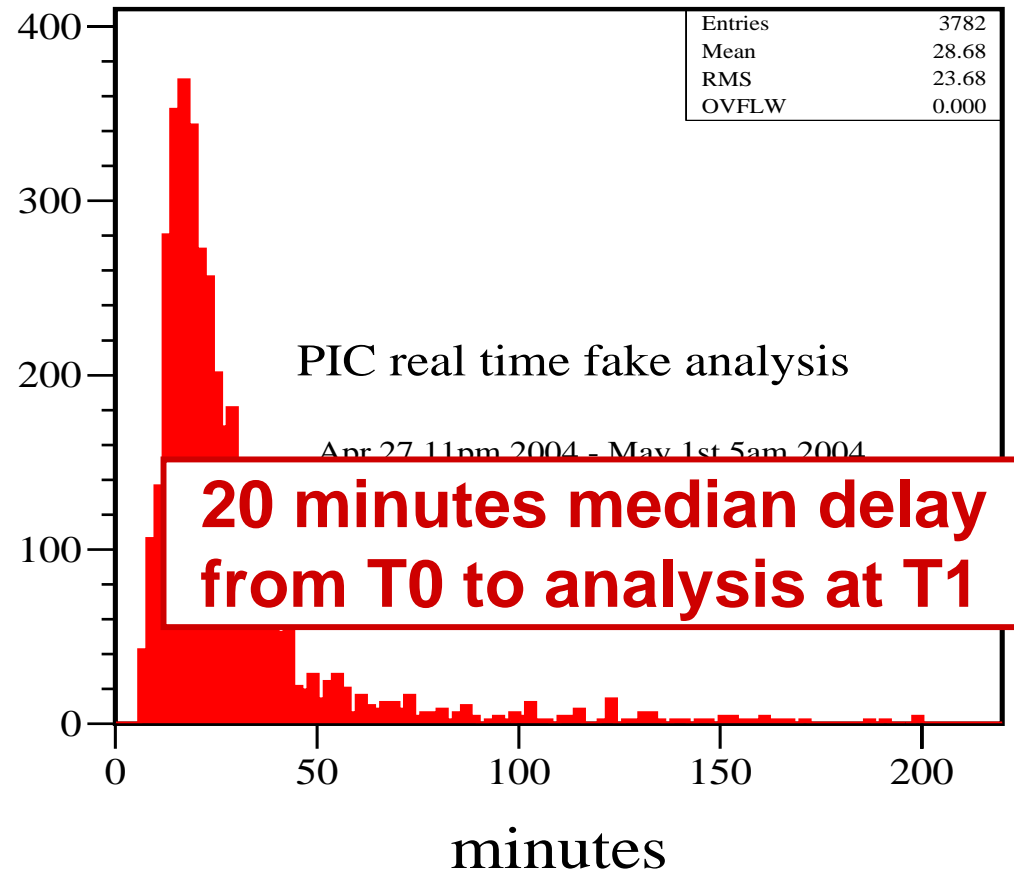
- Full chain (except Tier-0 reconstruction) could be performed in LCG-2

From Tier-0 Reconstruction to Analysis at Tier-1



Swiss Grid CMS Report

Time delay (File analyzed at T1) - (File on GDB)





Swiss Grid

CMS Report

David Stickland

15 Sept 2004

Page 7/12

Physics TDR

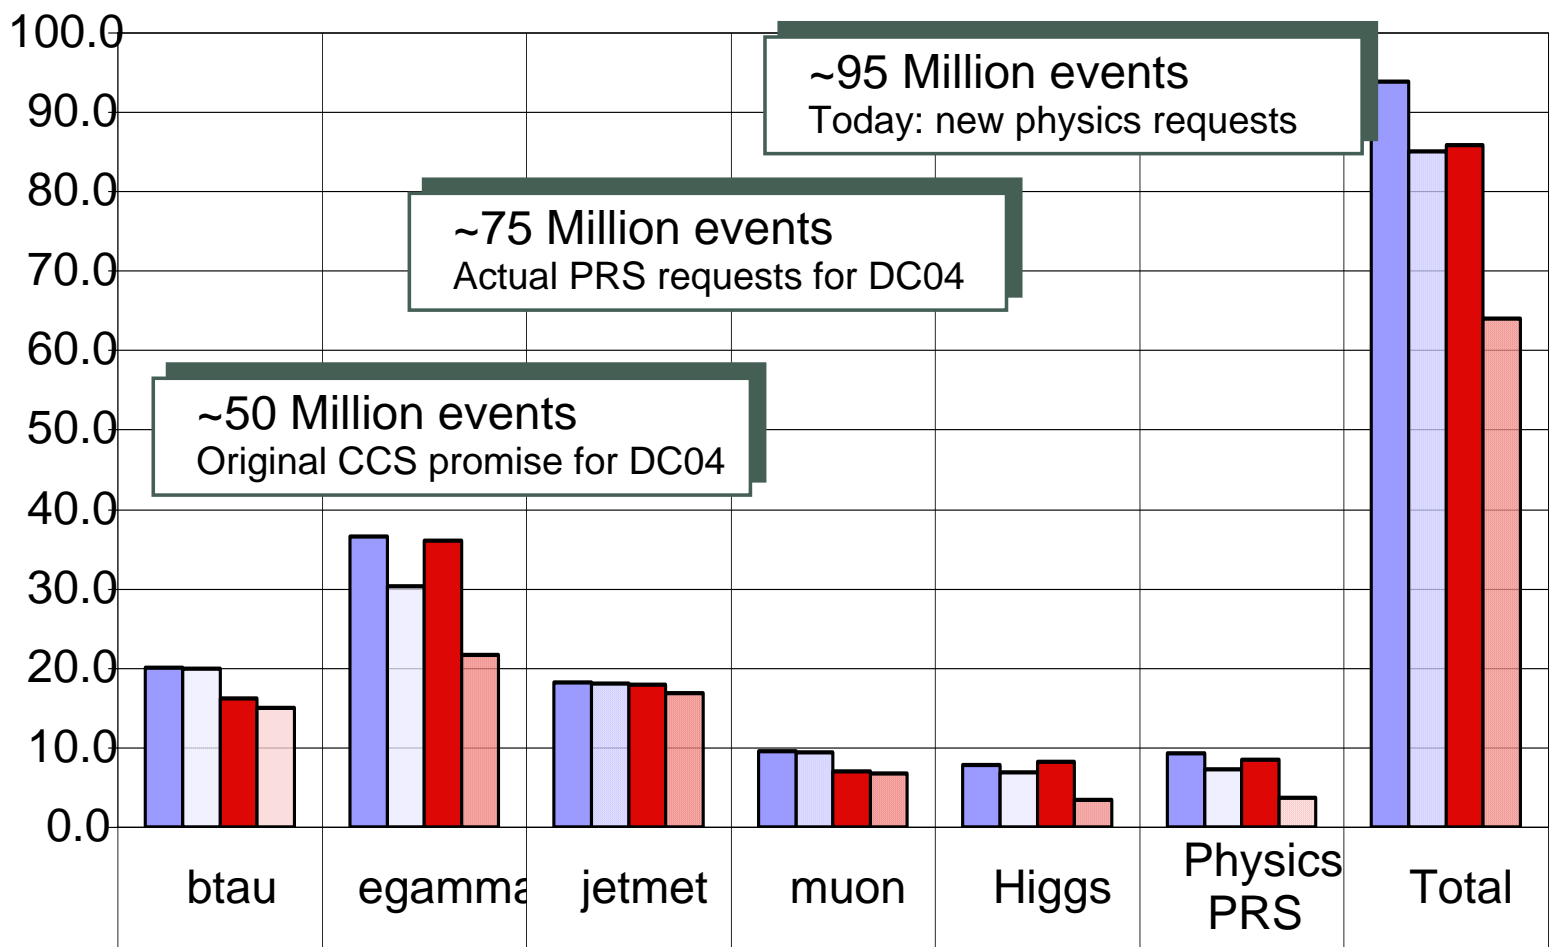
- **Physics TDR scheduled for December 2005**
 - Not a “yellow report”, but a detailed study of methods for initial data taking and detector issues such as calibration as well as physics reach studies.
- **Current Simulated samples more or less adequate for low luminosity**
 - About to enter re-reconstruction phase with new DST version
 - Estimate similar sample sizes for high luminosity
- **Target 10M events/month throughput**
 - Generation, Simulation, Digitization, Reconstruction, Available for analysis
 - New production operations group in CMS to handle this major workload
- **In light of DC04 experience, DC05 is cancelled as a formal computing exercise**
 - Not possible to serve current physics requirements with data challenge environment
 - However, specialized component challenges are foreseen

Continuous Operation

- Instead of DC05 we need to progressively bring up a full time operation
 - Not a test or a challenge.
 - Physicist access to data required.
- Generic GRID resources (for Generation/Simulation)
 - ~750 CPU Continuous
 - CPU means current generation ~2.4+GHz
- CMS T1 resources (Grid with significant Disk and MSS)
 - Needed for data intensive Digitization and Reconstruction steps
 - ~750 CPU Continuous
 - Now 60TB +20TB/month (Spread across T1 centers)
- T1/T2 resources (probably not generic)
 - 150C CPU Continuous
 - Now 40TB Analysis disk space to grow by about 2-5TB/month (T1+T2)
- We intend to run all this within LCG
 - “LCG” being all those resources available to CMS being steady migration from/between LCG2, GRID3, gLite, ...
- We need to reach the 10M event/month level soon (Autumn)
- We need to make the resources available to a growing CMS user base in the same time scale

CMS Plan / Status

Simulation / Digitisation Production



	btau	egamma	jetmet	muon	Higgs	Physics PRS	Total
■ Simulation requested	20.1	36.6	18.3	9.5	7.8	9.4	93.9
■ Simulation done	19.9	30.3	18.1	9.4	7.0	7.3	85.1
■ Digitisation requested	16.3	36.0	18.0	7.0	8.3	8.6	85.9
■ Digitisation done	15.0	21.7	16.9	6.8	3.5	3.7	64.1

Tier-2 Computing in CMS

- We are documenting our Computing Model now
 - Christmas 2004, reviewed by LHCC in Jan/Feb
 - But current understanding is...
- Tier-2 centers
 - Will be the major resource for Simulation in CMS
 - Will be the major analysis sites for Physicists once main data reduction has been carried out
 - All physicists will need access to at least one Tier-2
 - Local copies of selected data
 - Disk is getting much cheaper
 - Local MSS of user data (or contract with someone else)
 - But CMS data need not be on local MSS, but sent to Tier-1(s)

Current Tier-2 Scales

HLT Rate	100Hz	150Hz	200Hz	
CPU	310	380	450	KSI2k
Disk	190	220	250	TB
Archive Tape	115	130	150	TB /year
Tape I/O	100			MB/s
WAN I/O	50	60	75	Mb/s