



CMS

Tier-1/Tier-2 operation experience

José M. Hernández
CIEMAT, Madrid

WLCG Collaboration workshop

Compact Muon Solenoid




Outline

- CMS Computing model, data flow and T1/T2 data processing workflows
- Data transfer operations
- Scheduled processing (data skimming and reprocessing) at Tier-1's
- Monte Carlo production at Tier-2's
- Data analysis at Tier-2's
- Monitoring
- Summary



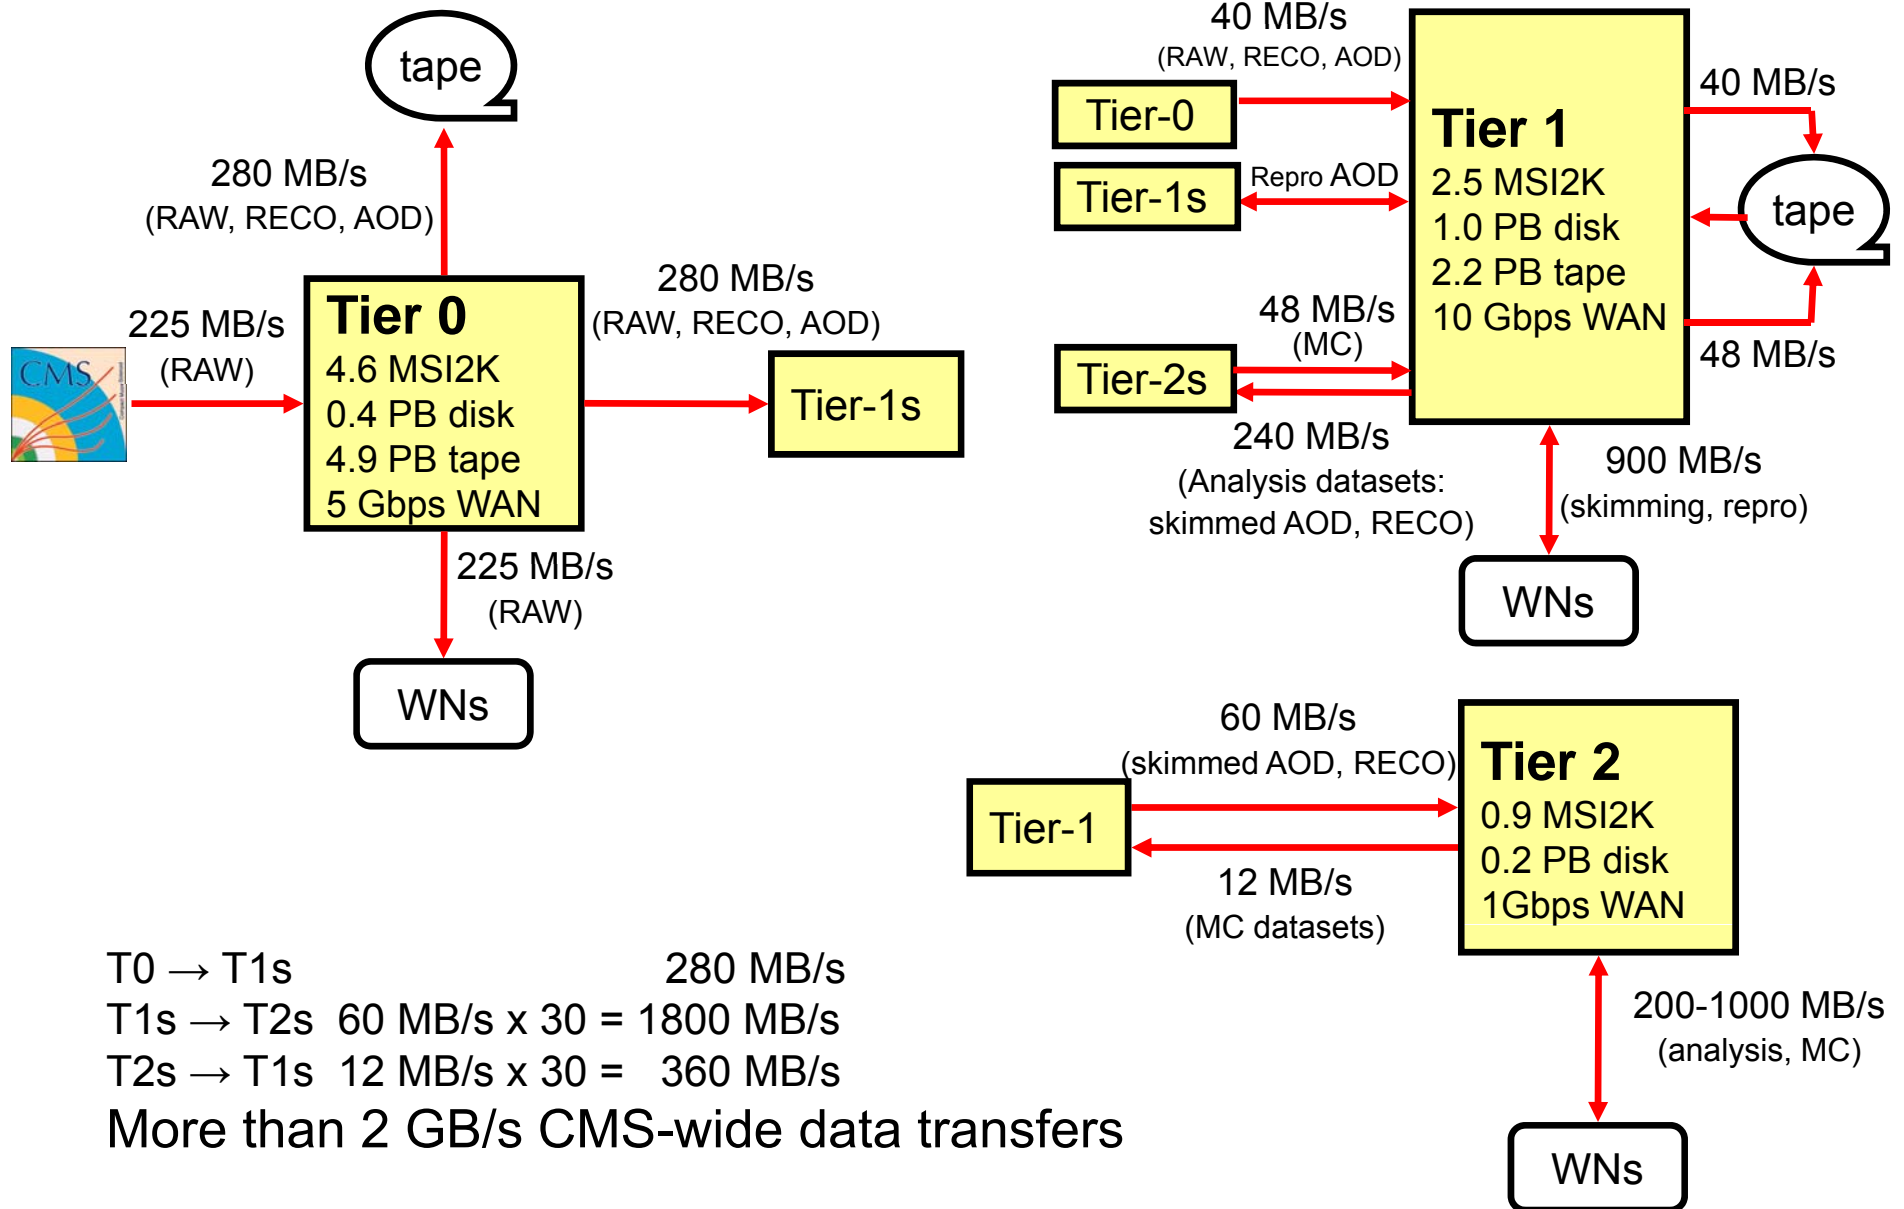
CMS Computing Model

- Distributed computing model for data storage, processing and analysis
- Grid technologies (Worldwide LHC Computing Grid, WLCG)
- Tiered architecture of computing resources
- Several Petabytes of data (real and simulated) every year





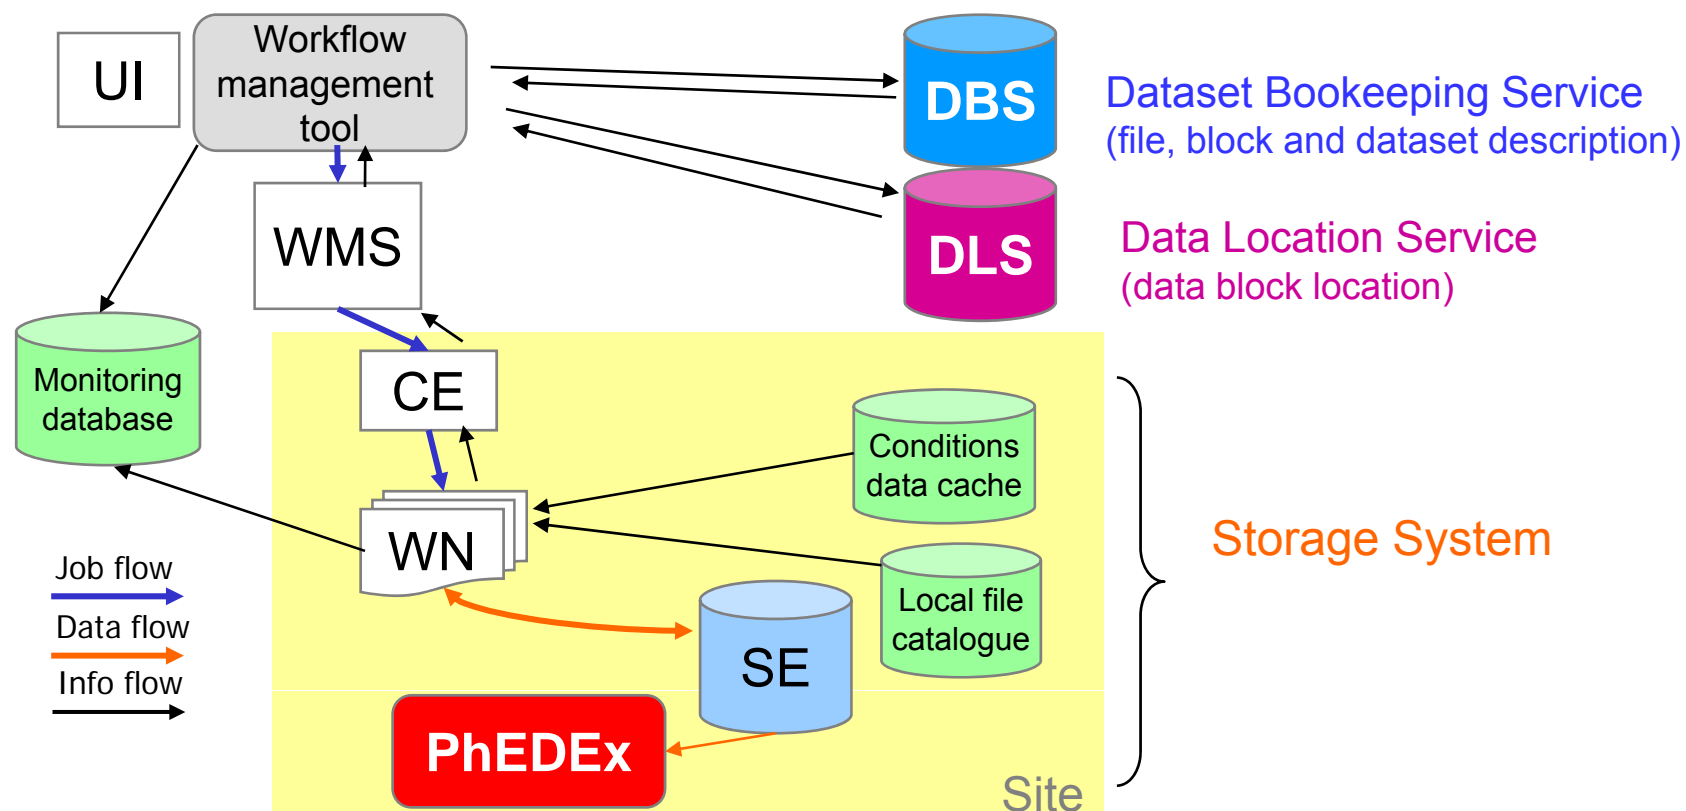
CMS data flow





Data processing workflow

- Data Management System to discover, access and transfer event data in a distributed computing environment
 - Fileblock-based data replication and tracking
 - Local (trivial) file catalogue
 - Merging processing step
 - Data location based processing





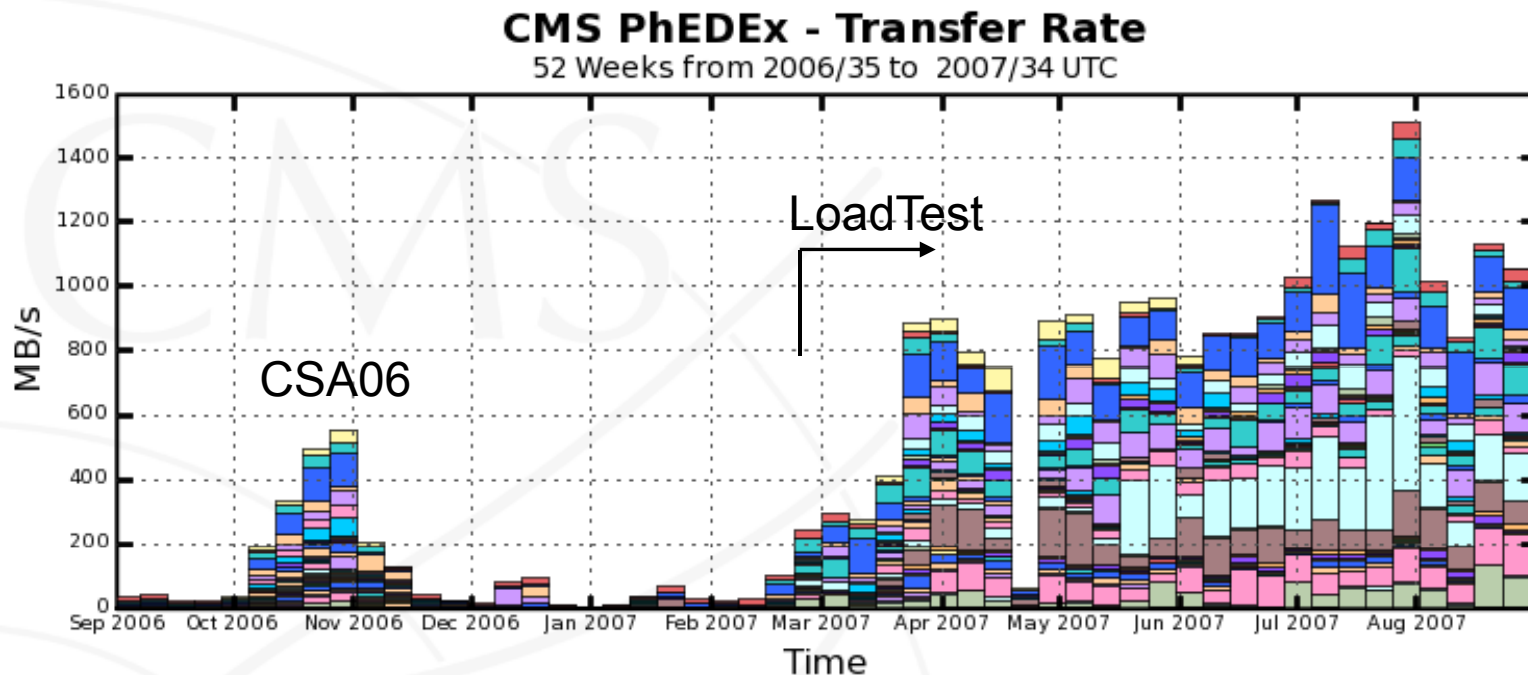
Data transfer operations

- PhEDEx is the CMS data transfer and placement system
 - Reliable end-to-end large scale data transfers
 - Global and site local agents
 - Multi-hop routed transfers. Dynamic routes
 - On top of LCG middleware (FTS, SRM, gridFTP)
 - Implements CMS policies and priorities
 - Includes CMS data organization (files, blocks, datasets)
- Experience from site operation
 - Very mature, robust, scalable, good monitoring and support
 - Needs contact person at sites running local agents (pros and cons)
 - Data management features (subscriptions, deletions, mass prestage)
 - Sometimes difficult to figure out what is going on (error logs and routing information available via global monitor helps)
 - FTS server architecture, FTS server at destination or source T1 ?



Data transfer operations

- LoadTest data transfer load generation machinery
 - Production transfers (MC data) too small in scale to stress the system
 - Routinely keeping a level of ~ 1 GB/s CMS-wide (~ 0.5 PB/week)
 - Mainly T0 \rightarrow T1s and regional T1 \leftrightarrow T2s transfers
 - Need to very much improve T1s \leftrightarrow T1s and non regional T1 \rightarrow T2s transfers
 - Big effort by DDT team to commission all links
 - Link commissioning takes a lot of time (many layers involved)





Storage systems

- CASTOR-2, dCache, DPM
- rfiio and dcap posix I/O access protocols
- Painful migration to CASTOR-2 (CNAF, RAL, ASGC)
 - Much heavier to tune, operate and maintain than CASTOR-1 (LSF, ORACLE)
 - Melt-down by heavy load
- dCache seems to work well for the other Tier-1s (FNAL, PIC, FZK, IN2P3)
- At Tier-2s DPM and dCache
 - Tier-2 sites with dCache generally happy. Heavier to manage than DPM
 - Problem in DPM with rfiio libraries with SL4 CMSSW versions...
 - Xrootd with DPM being investigated
- Would be good to share information and experience on configuration and tuning between sites (via infrastructure operations program)
 - See example of US sites



Data management and access

- Data management at Tier-1 sites
 - In 2008 all data fit on disk but in later years data retrieval from tape for organized data processing will be needed
 - Data pre-staging? Needed? How? PhEDEx? srm-bring-online? Coupling with workflow management tool? Pre-stage pilot jobs?
 - ▶ under debate
 - Storage organization
 - Storage classes in SRM v2.2 (disk0Tape1, disk1Tape1, disk1Tape0, disk0Tape0)
 - Disk space managed by SE according to space tokens or experiment
 - How many disk pools (WAN, processing, etc)
 - Automatic inter-pool copies by storage system
 - dCache namespace mapping to different disk pools
 - No policy dictated by CMS. Sites free to organize storage
 - Unique big disk0tape1 disk pool might be enough at the beginning
 - We have to see with experience



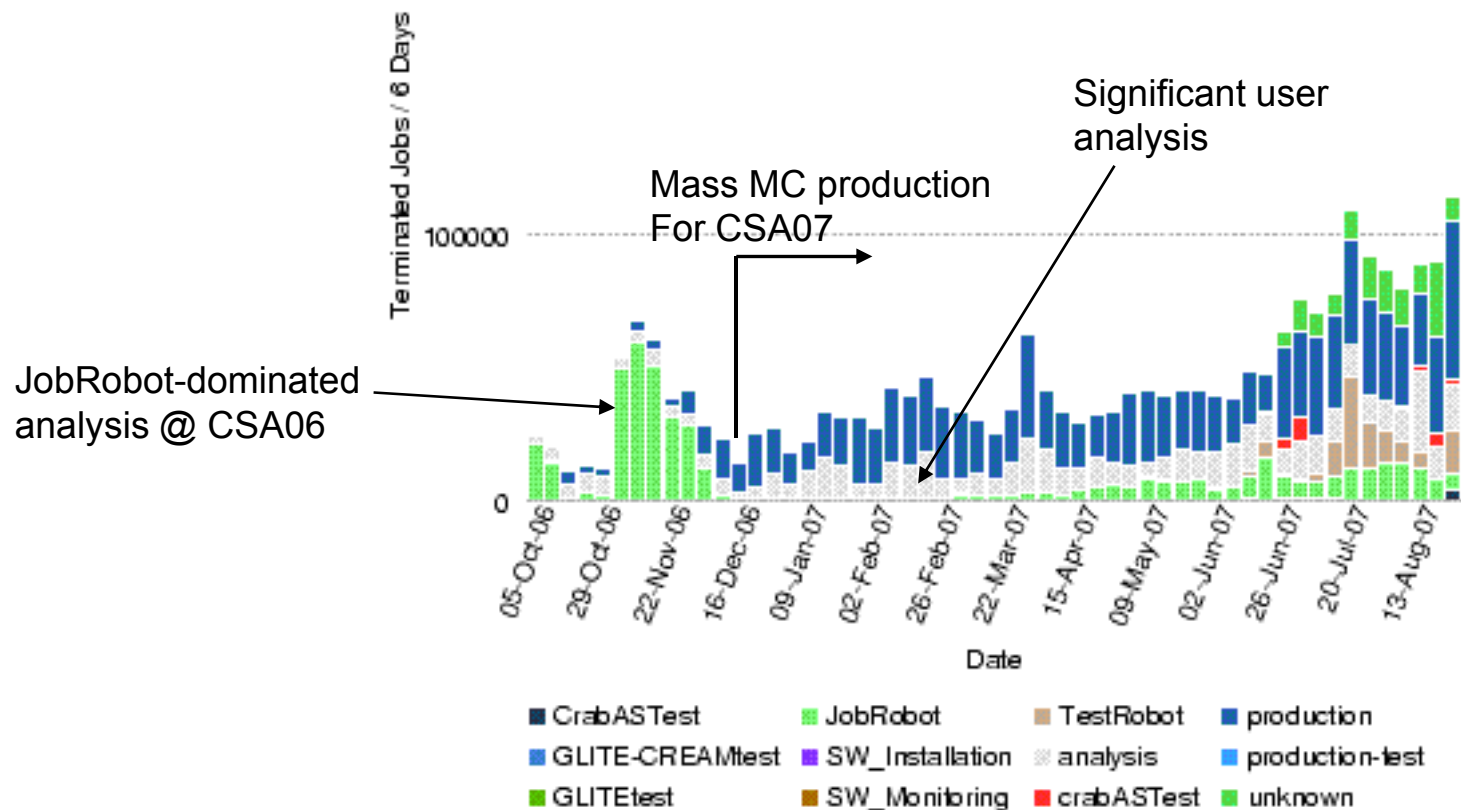
Data processing

- Data placement based
 - Jobs are sent to data, no remote stage-in
 - Data placement system replicates data if necessary prior to processing
- Data directly read from storage system via posix I/O access protocols
 - No stage-in from local storage into Worker Node
- Organized (scheduled) data processing at Tier-1s (data skimming and reprocessing)
- Organized MC production at Tier-2s (also at Tier-1s so far due to availability of resources)
- 'Chaotic' user analysis at Tier-2s (also at Tier-1s so far due to MC data placed there. Moving to proper Tier-1 processing workflows)
- Unique workflow management tool (ProdAgent) for MC production, data skimming and reprocessing (also reconstruction at Tier-0)
- CRAB for data analysis



Data processing

- Currently more than 20K jobs per day
- Mass MC production
- Significant user data analysis complemented by JobRobot jobs





Data processing at Tier-1's

- Data skimming only exercised during CSA06 at large scale
 - Several filters run concurrently producing several output files
- Successful demonstration of data reprocessing at Tier-1s during CSA06
 - 100-1000k reconstructed events at every Tier-1
 - Reading conditions data from local Frontier cache
 - Reprocessing-like DIGI-RECO MC production ramping up at Tier-1s
 - Only significant production at CERN and FNAL so far but other Tier-1s ready (SL4 WNs and CMSSW releases available)
- Organized data processing workflows at Tier-1s will be exercised during CSA07 at a larger scale
 - Recalling data from tape?




MC production

- Major boost in performance and scale in CMS MC production during last year
- Re-factored production system has brought automation, scalability, robustness and efficiency in handling the CMS distributed production system
- Much improved production organization, bringing together requestors, consumers, producers, developers and sites, has also contributed to the increase in scale
- Reached scale of more than 20K jobs/day and 65 Mevt/month with an average job efficiency of about 75% and resource occupation ~ 50%
- Production is still manpower intensive. New components being integrated to further improve automation, scale and efficient use of available resources while reducing required manpower to run the system
 - JobQueue & ResourceMonitor
 - ProdManager



Production performance: yield




~175 M events in last 3 months
~ 2/3 done at Tier-2 sites
~ 50% done at OSG sites



Total: 174953424.00 Events, Average Rate: 22.50 Events/s

Events Written by Site (Sum: 174953424 Events)
12 Weeks from 2007/21 to 2007/34 UTC

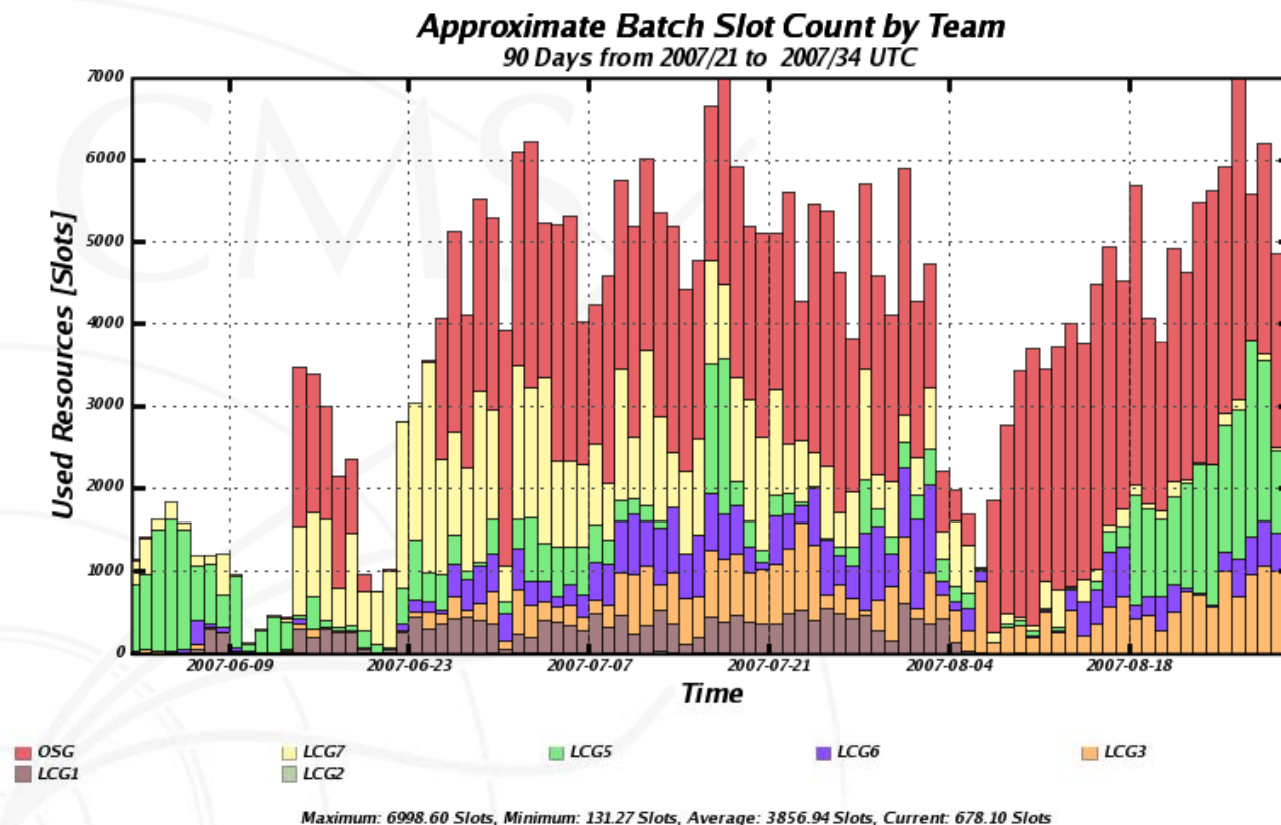


FNAL (52415870)	Florida (17200432)	Wisconsin (15152024)	Spain_CIEMAT (12905579)	Purdue (12781781)
Nebraska (10183793)	CERN (7240611)	Caltech (6611415)	MIT (5003822)	FZK (4344924)
SPRACE (3643710)	PIC (3633444)	Pisa (3104471)	IN2P3 (2695967)	London_IC_HEP (1206660)
RutherfordPPD (2314431)	London_Brunel (2088685)	Taiwan (1396202)	Spain_IFCA (1244442)	Vienna (1206660)
UCSD (1117029)	Estonia (1077859)	Legnaro (773122)	INFN (595739)	WARSAW (527178)
London_RHUL (522502)	KNU (435035)	T2_DESY (434636)	LIP-Lisbon (421406)	T2_Belgium (4071)
BUDAPEST (375924)	Belgium_UCL (209831)	RWTH (99463)	Bari (57272)	Rome1 (35939)
Unknown (18033)	ASGC (50)			



Production performance: resource usage


- Approximate slot usage (average number of slots continuously running production)
- Ramp up from June (~ 5000 slots used in average out of the 10K slots available)
- In average ~50% resource occupation (production inefficiencies, no automatic resource management, manual job release by operators, many sites, many small production requests)






Production performance: job efficiency

Job Success Rates
89 Days from 2007/21 to 2007/34 UTC





- 40+ sites used continuously for production with good job efficiency
- Average job efficiency ~ 75%, including application and Grid. (Failed jobs are automatically resubmitted and get done)



- > 20K/day production jobs

Job Exit Code
89 Days from 2007/21 to 2007/34 UTC



Maximum: 32996.00 Jobs, Minimum: 3957.00 Jobs, Average: 16226.37 Jobs, Current: 12955.00 Jobs

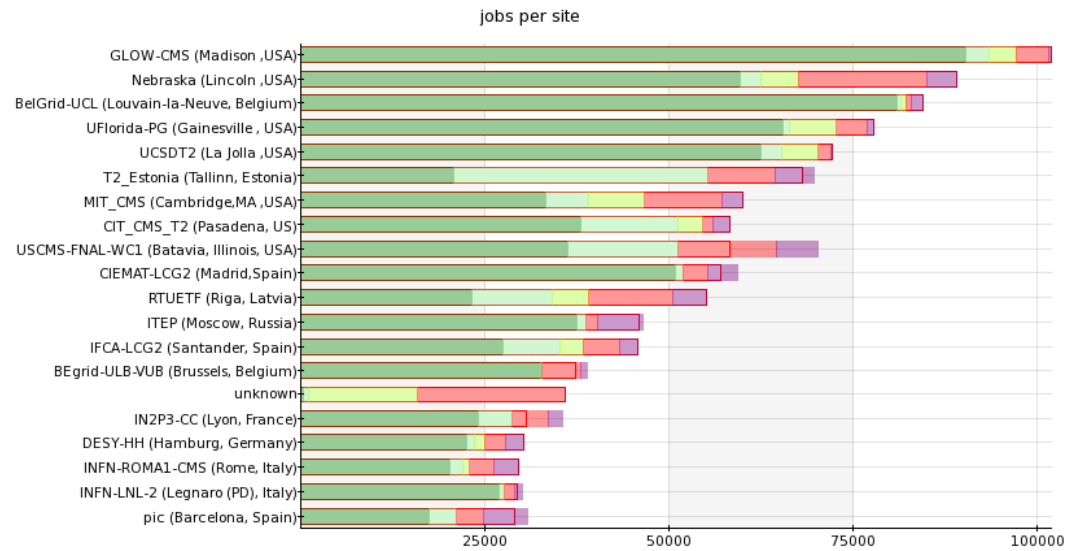


Data analysis


- Significant scale of job submission
 - Real user analysis jobs
 - Complemented with JobRobot fake analysis jobs (job load generator)
- User data management under discussion/prototyping
 - /store/user namespace
 - Possibility in CRAB of user data registration in DBS/DLS
 - Possibility in CRAB of PhEDEx injection for transfer
 - User data management at sites (disk0tape0?)
- CRAB being re-factored to include Prodagent components for job tracking, automatic resubmission, data registration, etc (CRAB server)



Data analysis



- Few million analysis jobs (real user analysis + JobRobot fake analysis) run during last year
- A lot of user analysis still at Tier-1s



submitted app-succeeded app-failed app-unknown pending running aborted cancelled



Site DMS and WMS monitoring

- Several tools to monitor infrastructure and DMS/WMS services at sites to make sure CMS workflows will run correctly
- SAM jobs to monitor site configuration and services very useful
- JobRobot jobs to monitor job submission and data access
- Sites should check SAM, JobRobot monitors to find and fix problems
- PhEDEx monitor for data transfers
- Very important prompt reaction from sites
 - Data access problems, wrong configuration, WNs black holes, BDII, etc
- Savannah tickets for data operations team, and GGUS tickets for sites for operational problems
- CMS site contacts sometimes directly contacted by operations team
- Hypernews fora for site contacts
- Need involvement from sites



Outlook

- Tier-1/Tier-2 CMS data and workflows well established
- Large and sustained scale
 - 1 GB/s data flows CMS-wide
 - ~20k jobs/day MC production and data analysis
- Big effort ongoing in data transfer and site commissioning which requires commitment from sites
- Still a lot of work to do
 - T1s ↔ T1s and non-regional T1 → T2s transfers
 - Exercise workflows involving data recall from tape
 - Storage organization at sites