



Computing Operations at CMS Facilities

D. Bonacorsi

(CMS Facilities / Infrastructure Operations --- INFN-CNAF Tier-1, Bologna, Italy)

On behalf of the CMS experiment




International Conference on Computing
in High Energy and Nuclear Physics

2-7.Sep.07 - Victoria BC, Canada





Outline



- Data processing
 - ❑ aka: “how can I process 1 CMS event in a Grid-enabled world?”

- Data transfer
 - ❑ aka: “how can I transfer 1 CMS bit in a Grid-enabled world?”

- Site readiness
 - ❑ aka: “how can I help a Grid to help me by being able to do the above?”

- CMS Operations
 - ❑ aka: “No future releases exist. Only current ones. I do operations today.”

DISCLAIMERS: focus is deliberately on off-CERN Tiers

See #xyz pointers to related CHEP07 contributions



CMS Offline & Computing

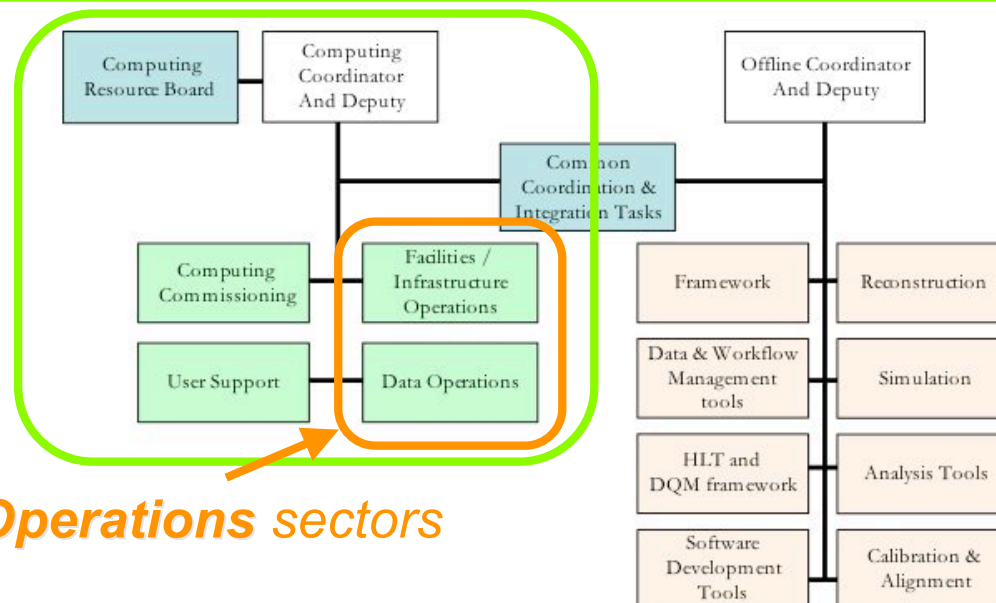


➤ Offline Project

- ❑ responsible for the overall coordination, integration and deployment of offline software developments including the application services used to manage job submission and the production and distribution of data

➤ Computing Project

- ❑ responsible for all operational aspects of CMS data processing, the operation of computing facilities, and the provisioning of the various infrastructure services



Focus on Operations sectors



CMS Computing project (1/2)



➤ Facilities and Infrastructure Operations

- ❑ provide and maintain a working distributed computing fabric with a consistent working environment for the Data operation and the users.
 - ❖ It involves coordination of facilities operation, resource management and liaison to external projects and organisations.
- ❑ Sub-tasks include:
 - ❖ *Tier-0 and CAF operations*: Coordination of systems and networking provision; Data- and Job-flow monitoring; Troubleshooting; Operational efficiency monitoring and improvement
 - ❖ *Liaisons to external projects*: coordination of operations with CERN-IT, Tier-1 and Tier-2 centres and within the WLCG and OSG projects
 - ❖ *Resource allocation*, planning and tracking by activity for CPU, storage and networking, interface to IT and Grid accounting projects.
 - ❖ *Databases operations*: provision and maintenance of DB servers, DB integrity and backup
 - ❖ *Liaison with ROCs*

➤ Strongly linked to the **Commissioning project**



CMS Computing project (2/2)



➤ **Data Operations** for physics analysis

- ❑ responsible for operational aspects of real and MC data processing and executing production workflows.

- ❖ It will ensure that detector calibration and alignment corrections are applied at prompt processing and data are properly distributed.

- ❑ **Sub-tasks include:**

- ❖ *Host lab data processing:* Prompt calibration, alignment at T0/CAF, processing of environmental data, 1st-pass reco at T0, operations for obtaining refined calibration and alignment at T0, CAF, T1s and T2s.

- ❖ *Distributed MC Production:* Distributed production and storage of MC events; coordinate MC queues, bookkeeping of MC data

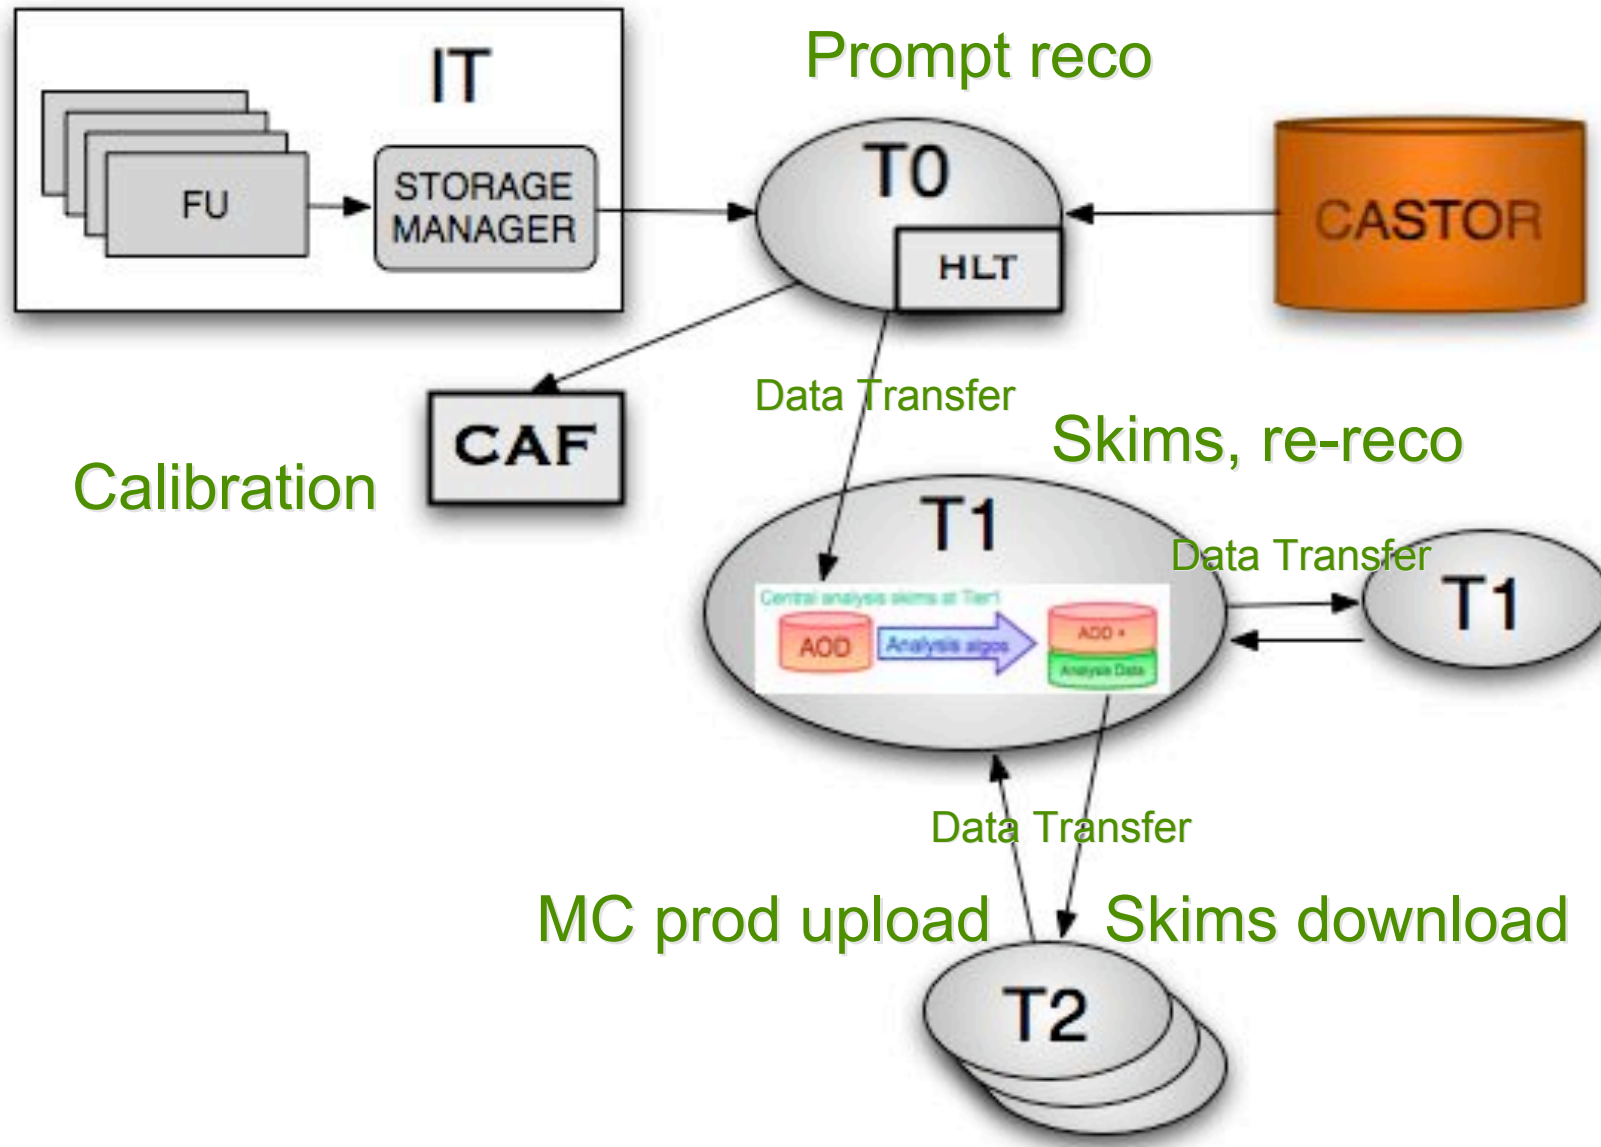
- ❖ *Data Distribution and Integrity:* Distribution of first-pass FEVT and AOD to T1s; data management and bookkeeping

- ❖ *Distributed Data Processing:* Organized large-scale analysis activities; reprocessing at T1s, RECO and AOD distribution

- ❖ *Data Certification for Physics:* Application of Data Quality Monitoring algorithms and procedures to the data; Luminosity determination; organizing offline shifts for DQM; establishing good-run lists



Main workflows





Are sites prepared to
reliably process my jobs?




JobRobot as a site commissioning tool



#240


- A simple wrapper around job submission system
 - ❑ Submits short analysis jobs using Grid (EGEE/OSG)
 - ❖ ~100/jobs/day/site to all CMS Tier-1/2's
 - ❑ Web access to statistics and log files
- Same data, same code, same users at all Tiers, every day
 - ❑ It's flat, and tuned by robot managers
 - ❑ Standard candle for user's analysis: "why then do those jobs fail?"





Site Availability Monitor (**SAM**) for CMS

#145



- SAM developed to run periodic sanity checks on Grid/experiment services
 - ❑ A CMS instance of the SAM client is installed at CERN
- SAM framework allows to plug-in new tests for existing sensors
 - ❑ CMS tests submitted from CERN every 2 hrs to EGEE/OSG sites
 - ❖ E.g.: CMS env, site local config, correct CMSSW installations, stage-out, Squid/FroNtier, ...
- Results and log files recorded in SAM DB at CERN + access via web interfaces
 - ❑ Soon: selected data will be copied from SAM DB to CMS Dashboard DB for custom aggregation and visualization
- Success in running SAM tests is an important prerequisite for sites
 - ❑ Sites are required to monitor the outcome of SAM tests
 - ❖ All failures to be understood/fixed; when sites do that, improvements over time are visible
- A good tool to focus site people's attention on quality of delivered services



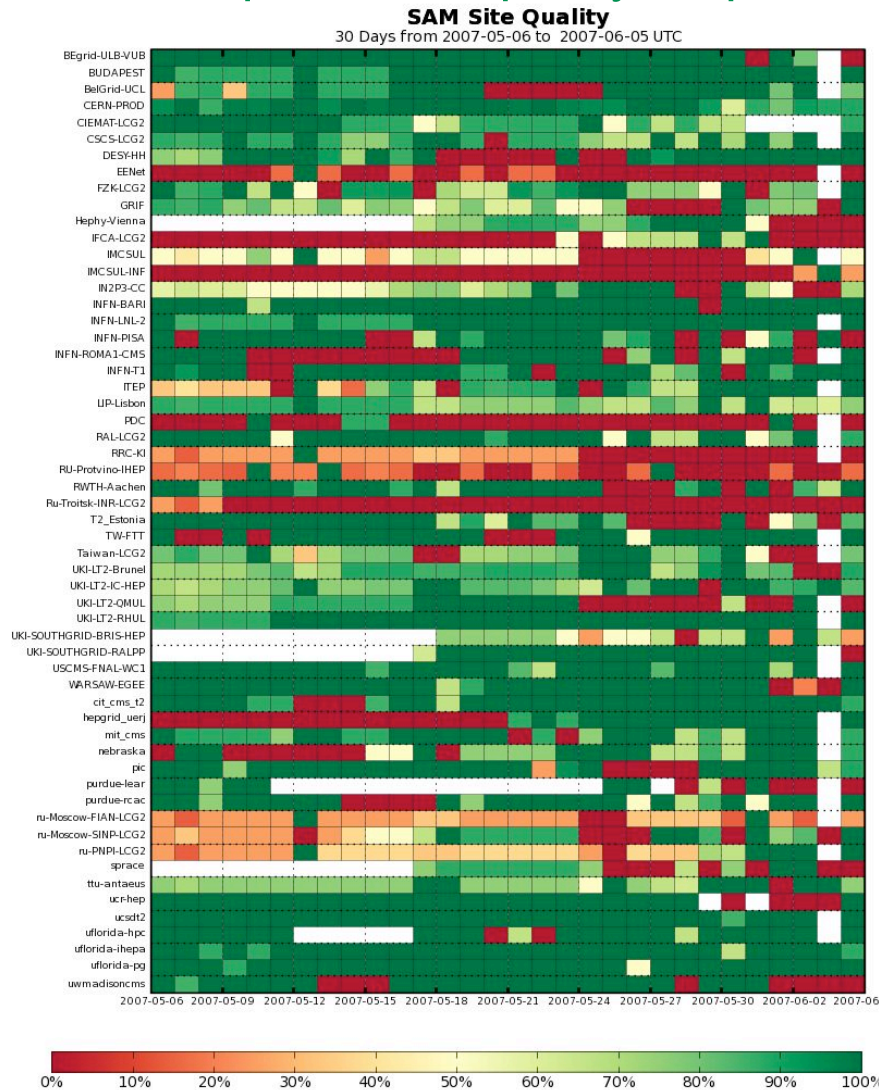
SAM for CMS: a screenshot



#145

2 compile a site quality map


1 SAM outputs on the web



No	RegionName	SiteName	NodeName	Status	cms					
					je	squid	zwnst	basic	mc	frontier
1	Unknown	cit cms t2	cit_gatekeeper.ultralight.org	INFO	ok	ok	ok	ok	error	ok
2	Unknown	hepgrid_uerj	osgce.hepgrid.uerj.br	INFO	error	ok	ok	ok	ok	ok
3	Unknown	mit cms	ce01.cmsaf.mit.edu	INFO	ok	ok	ok	ok	ok	ok
4	Unknown	nebraska	red.unl.edu	INFO	ok	ok	ok	ok	ok	ok
5	Unknown	purdue-lear	lepton.rcac.purdue.edu	INFO	ok	ok	ok	ok	ok	ok
6	Unknown	purdue-rcac	osg.rcac.purdue.edu	INFO	ok	ok	ok	ok	ok	ok
7	Unknown	sprace	spgrid.if.usp.br	INFO	ok	ok	ok	ok	ok	ok
8	Unknown	ttu-antaeus	antaeus.hpcc.ttu.edu	INFO	ok	error	ok	ok	ok	error
9	Unknown	ucf-hep	top.ucf.edu	INFO	error	ok	ok	ok	ok	ok
		t2	osg-gw-2.t2.ucsd.edu	INFO	ok	ok	ok	ok	ok	ok
		a-hpc	iogw1.hpc.ufl.edu	1						
		ihcpa	hg.ihcpa.ufl.edu	1						
		pg-shepa-ufl.edu	pg.shepa.ufl.edu	1						
13	Unknown	uflorida-pg		1						
14	Unknown	uwmadisoncms	cmsgrd02.hep.wisc.edu	INFO	ok	ok	ok	ok	ok	ok

Click for history

Click for log file



37% of sites over the 80% availability mark
62% of Tier-0/1 sites

3 define a metrics and construct an availability matrix



Data processing: concept and tools



#223

#224

#325

#352

- Computing Model: a **data-placement** based approach
 - ❑ Jobs are sent where data is, no data transfer in response to jobs
 - ❖ Data placement system calls for data replication prior to processing (if needed)
 - ❖ Data directly read from storage system via posix-like I/O access protocols
 - No stage-in from local storage to WNs

- Computing Model: **organized** vs **'unpredictable'** processing
 - ❑ Scheduled data processing at T1s
 - ❑ Organized MC production at T2s
 - ❑ 'Unpredictable' user analysis at T2s

- Tools
 - ❑ ProdAgent #235
 - ❖ unique CMS workflow management tool for MC production, data skimming and reprocessing (also reconstruction at Tier-0)
 - ❑ CRAB #314
 - ❖ CMS tool for data analysis



Data processing: roles of Tiers



- T0 prompt reco, analysis/calibration object generation, ... #290

T1

- Skimming
 - ❑ Data selection at T1s based on physics filters (followed by data transfer to T2s for analysis)
 - ❖ Exercised at large scale during CSA06 only (several filters running)
- Re-reconstruction
 - ❑ Data reprocessing at T1s with access to new calibration constants
 - ❖ Successfully demonstrated during CSA06 (0.1-1 k reco evts/T1, reading conditions data from local FroNTier)
 - ❖ Re-processing-like DIGI-RECO MC production ramping up at T1s

T2

- Simulation
 - ❑ Submission of jobs (simulation code) at T2s
 - *In the transient.* also at T1s so far (to catch all available SL4 resources)
- Analysis
 - ❑ Submission of jobs (analysis code) at T2s
 - ❖ Analysis users complemented by robotic submissions
 - *In the transient.* also at T1s so far (due to MC data placed there)

more


more



Data processing: 1 year



- Currently: >20k jobs/day
 - ① Massive MC production
 - ② Constant, significant presence of analysis jobs
 - complemented by JobRobot-driven submissions
 - ③ Middleware tests, CRAB Analysis Server tests, ...




#285




MC Production


#285



During 2006, major boost of **performance** + **scale** in CMS MC Production

- Re-factored CMS distributed production system
 - ❑ higher level of automation, scalability, robustness, efficiency
- Improvements in the organization of MC Production Operations
 - ❑ orchestrated work of developers, requestors, producers, consumers, sites
- MC production in a nutshell: 
 - ❑ Reached > 20k jobs/day, ~175 Mevts in June→now, best: 65 Mevts/month, ~23 evts/s on average.
 - ❑ resource occupation ~50%
 - ❖ Production inefficiencies, manual actions by operators, not optimized operators-to-sites associations, many sites in the production pools, many small production requests, ...
 - ❑ ~2/3 at T2s, EGEE/OSG ~50%-50%
 - ❑ <job eff> (application+Grid) ~ 75%
- Further margins of improvement being explored
 - ❖ JobQueue & ResourceMonitor, ProdManager: new components being developed and integrated to further improve automation, scale and efficient use of resources


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



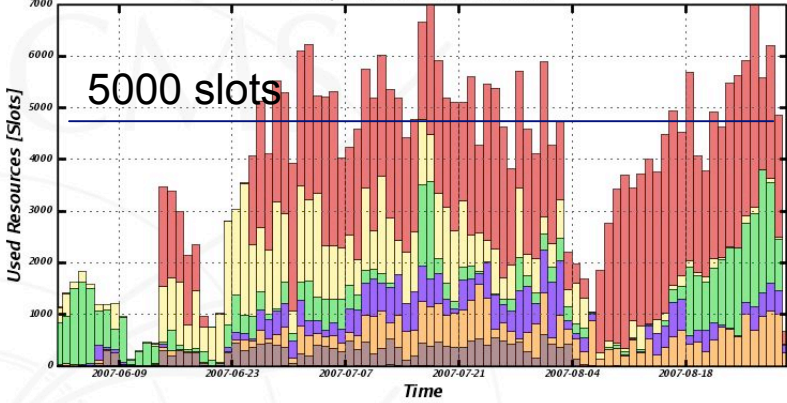
- FNAL (52415870)
- Florida (17200432)
- Wisconsin (15152024)
- Spain_CIEMAT (12905579)
- Purdue (12781786)
- Nebraska (10183793)
- CERN (7240611)
- Caltech (6611415)
- FZK (4344024)
- SPRACE (3643710)
- PIC (3104471)
- Pisa (3104471)
- IN2P3 (2695967)
- London_JC_HEP (2677084)
- RutherfordPPD (2314431)
- UCSD (1117029)
- London_Brunel (2088685)
- Taiwan (1396202)
- Spain_JFCA (1244442)
- Vienna (1206660)
- Estonia (1077859)
- Legnaro (773122)
- INFN (595739)
- WARSAW (527178)
- London_RHUL (522502)
- KNU (435035)
- TZ_DESY (434636)
- LIP-Lisbon (421406)
- BUDAPEST (375924)
- Belgium_UCL (209831)
- RWTH (99463)
- Bari (57272)
- Unknown (18033)
- ASGC (50)
- Rome1 (35939)

Cumulative Events Written (Merge)
 2160 Hours from 2007/21 to 2007/34 UTC

#285

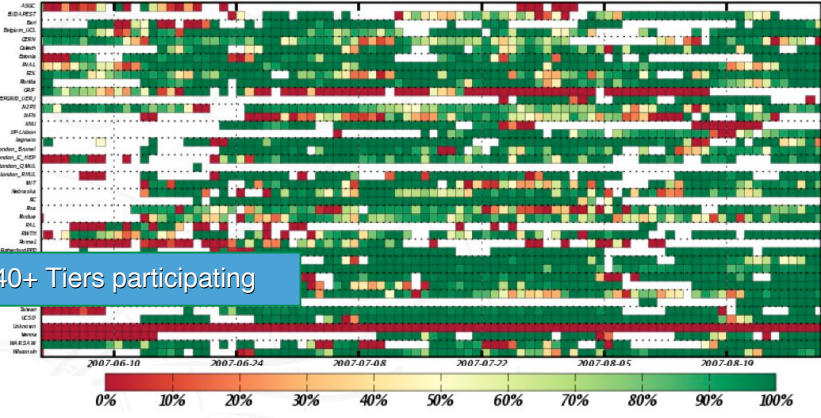


Approximate Batch Slot Count by Team
 90 Days from 2007/21 to 2007/34 UTC




~5k slots (on average) running production

40+ Tiers participating



<job efficiency> ~75%
 [failed jobs are automatically resubmitted]

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






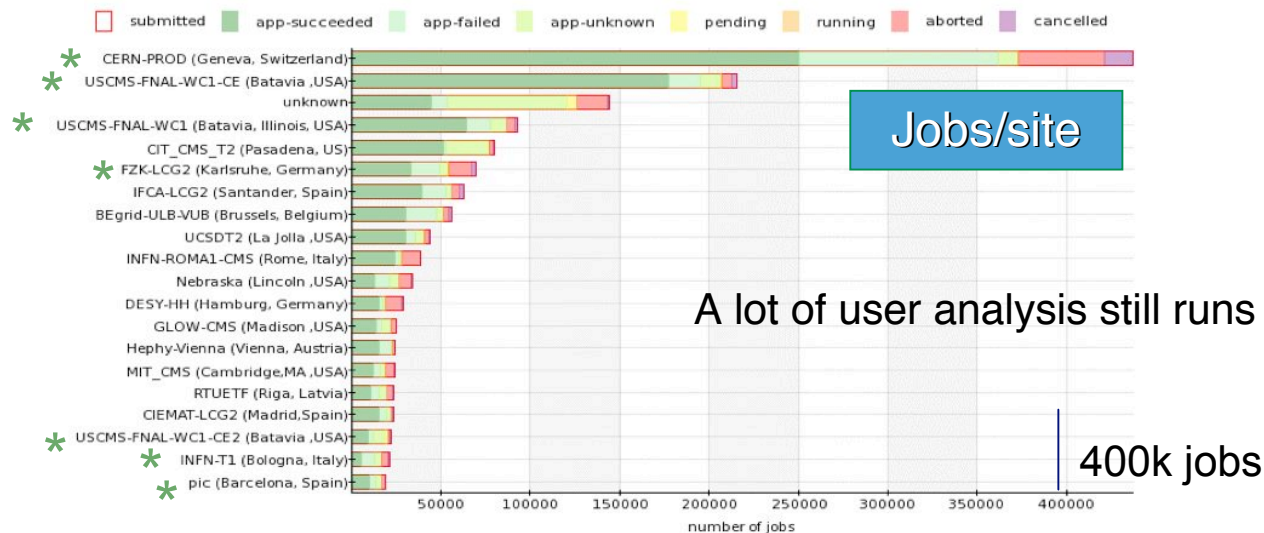
Data Analysis

#225

#314



- Significant scale of **analysis jobs** submissions reached
 - ☐ Few million analysis jobs ran during last 12 months
 - ❖ *Real* analysis users submitting *real* analysis jobs
 - ❖ Complemented by JobRobot-generated 'fake' analysis jobs




- Constant progresses of the **CRAB** project
 - ☐ being re-factored to include ProdAgent components for job tracking, automatic resub, data registration, ... (CRAB Analysis Server)
 - ☐ User data management under discussion/prototyping
 - ❖ /store/user namespace
 - ❖ CRAB may host features for user data registration in DBS/DLS, PhEDEx injections, ...



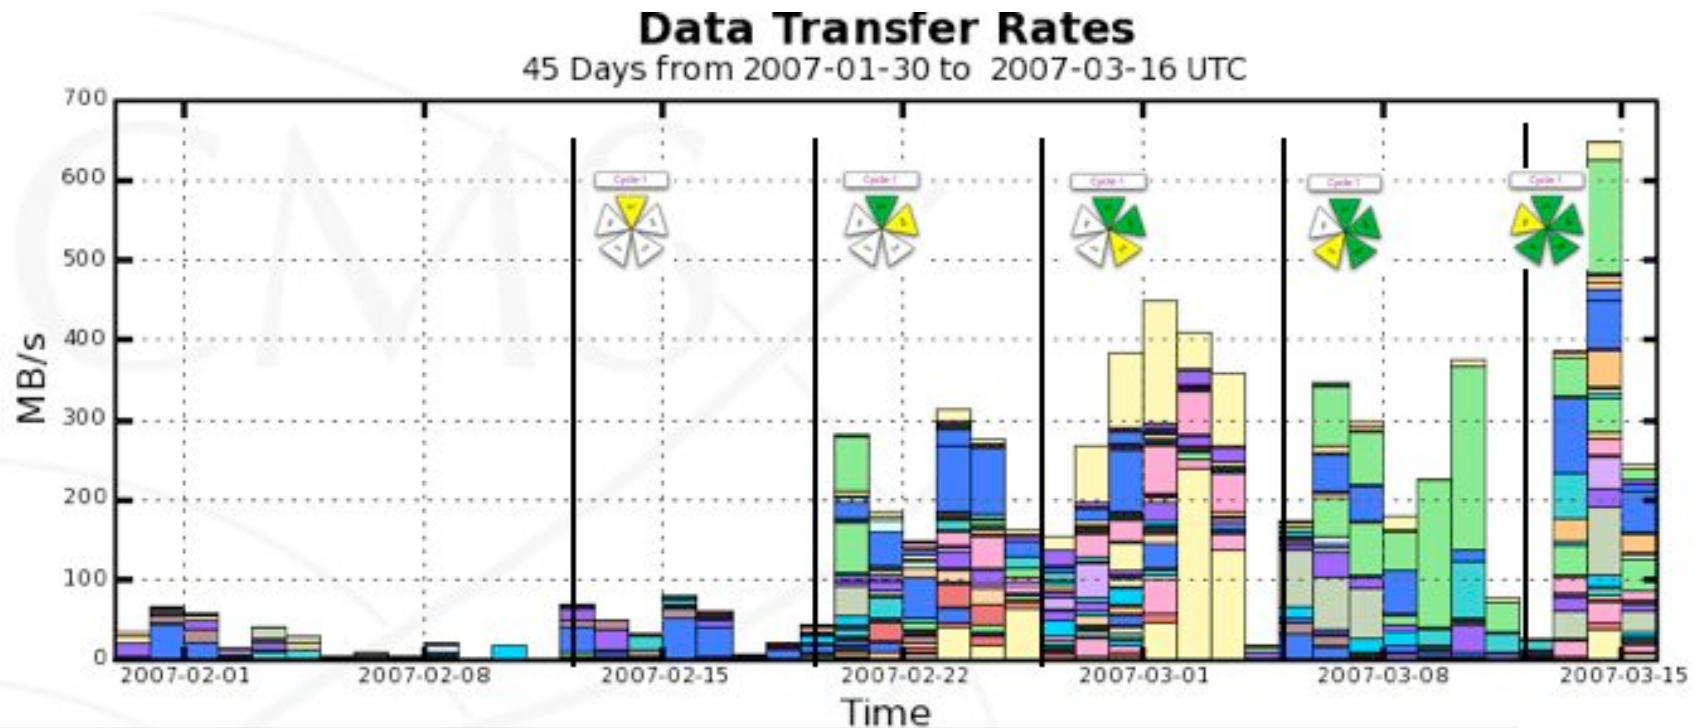
Are sites prepared to
transfer data?

Evolution of a LoadTest




Start moving

- A flexible infrastructure to generate data transfer traffic among CMS Tiers
 - ❑ “fake” but “real”
 - ❖ No real physics files; fully PhEDEx-compliant, though
 - ❑ ~24/7 activity since mid-February 2007
 - ❖ full cycles (5-weeks each) before Jun07 (then extended into DDT + CSA07 preparation activities)
 - ❖ T0→T1(tape), T1↔T1, T1↔T2 ‘regional’, T1↔T2 ‘non-regional’




Jan 2007

mid-March 2007



Jan 2007

mid-March 2007

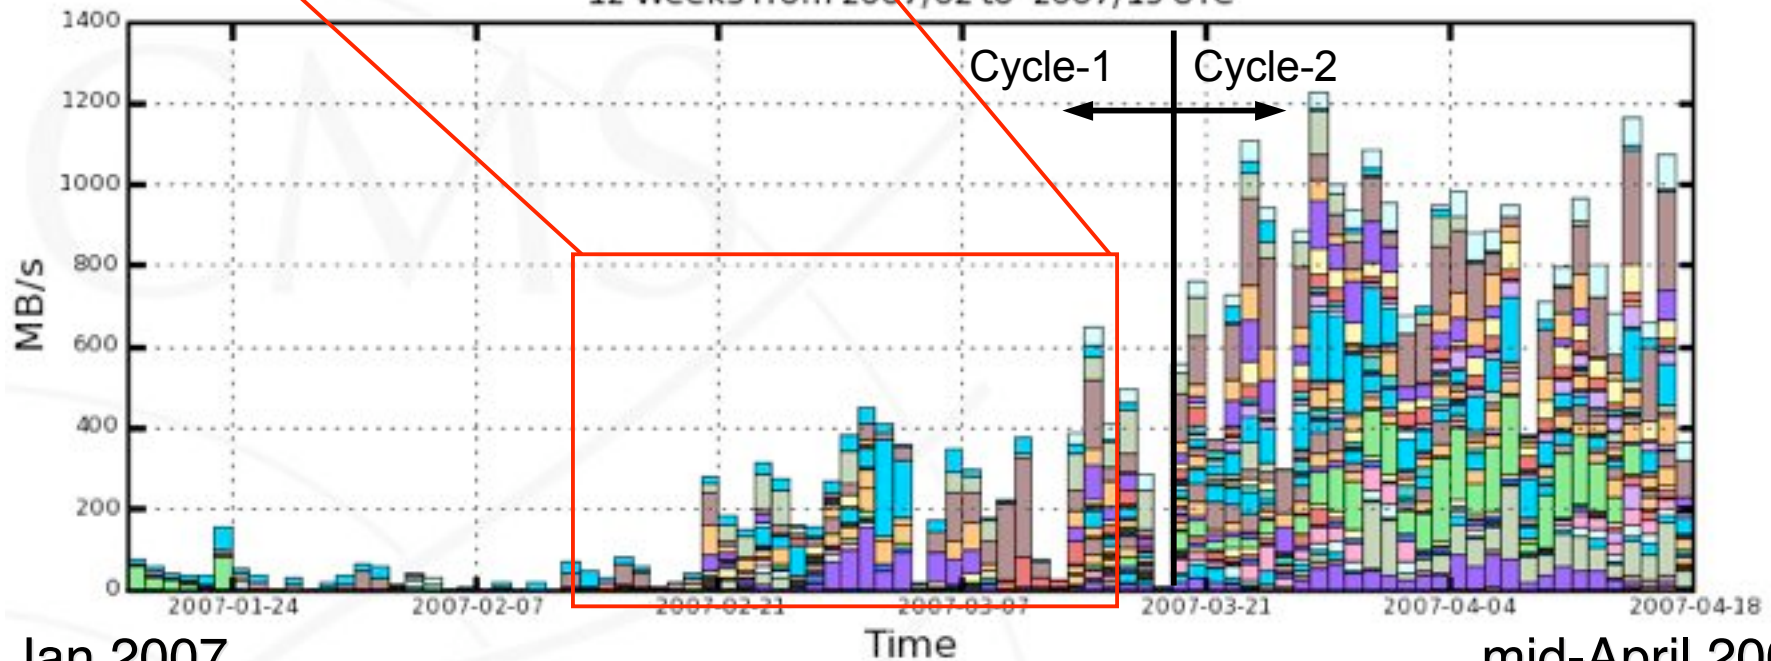


Start walking



CMS PhEDEx - Transfer Rate


12 Weeks from 2007/02 to 2007/15 UTC



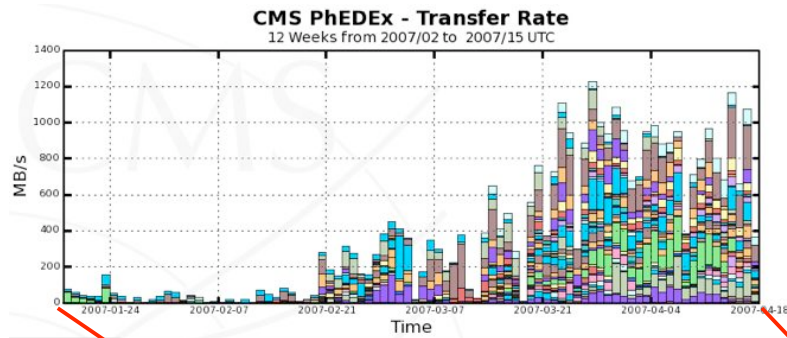
Jan 2007

Time

mid-April 2007



Walk better and faster

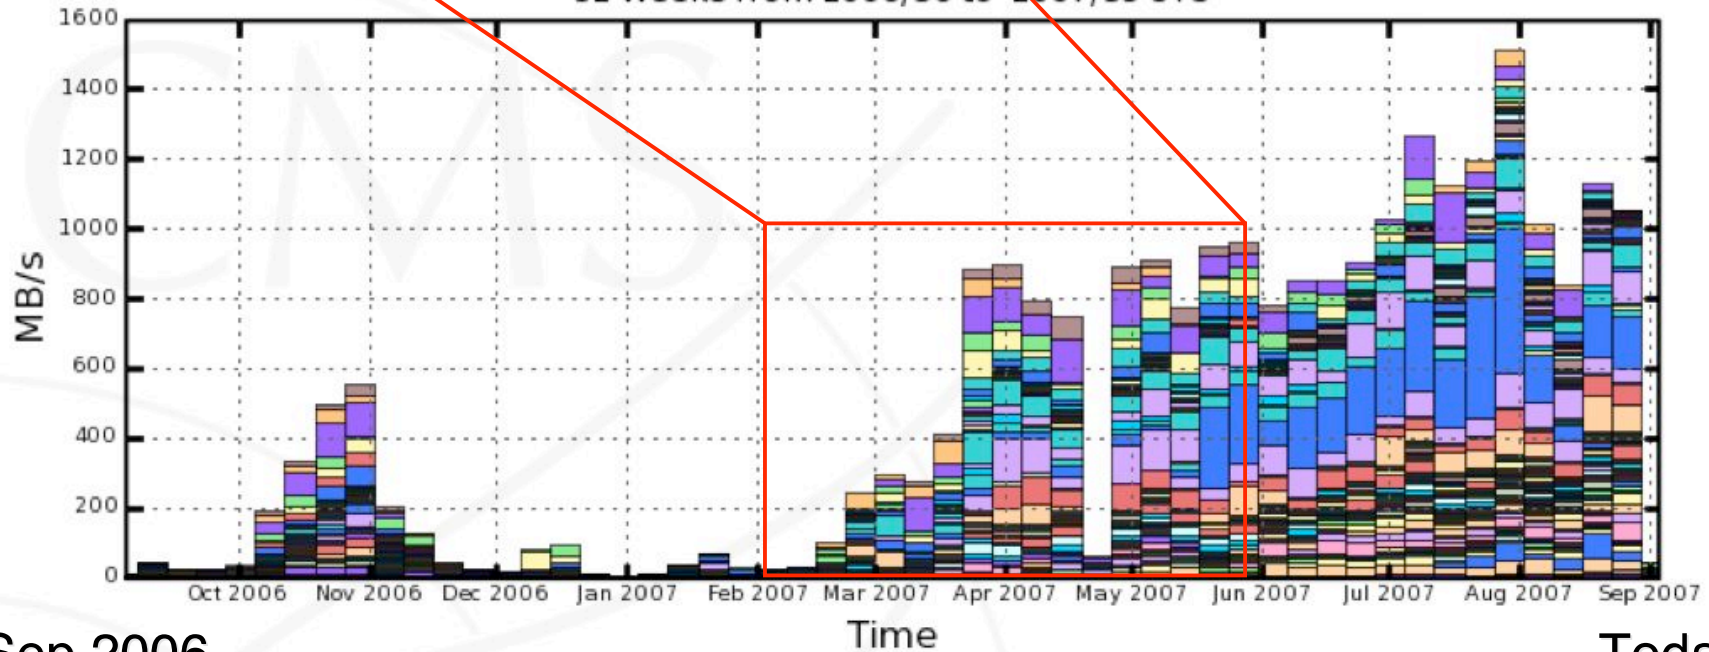


Jan 2007

mid-April 2007

CMS PhEDEx - Transfer Rate

52 Weeks from 2006/36 to 2007/35 UTC




Sep 2006

Time

Today

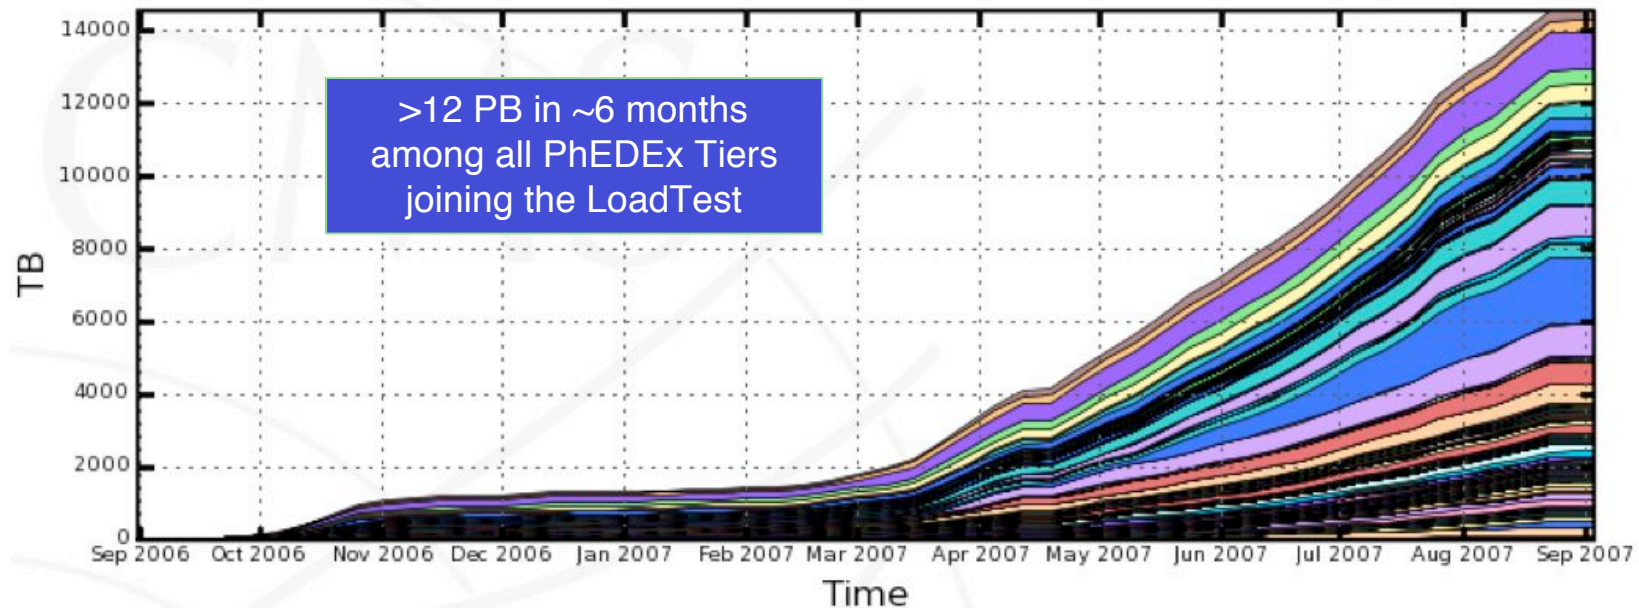
CMS PhEDEx - Transfer Rate

52 Weeks from 2006/36 to 2007/35 UTC



CMS PhEDEx - Cumulative Transfer Volume

52 Weeks from 2006/36 to 2007/35 UTC

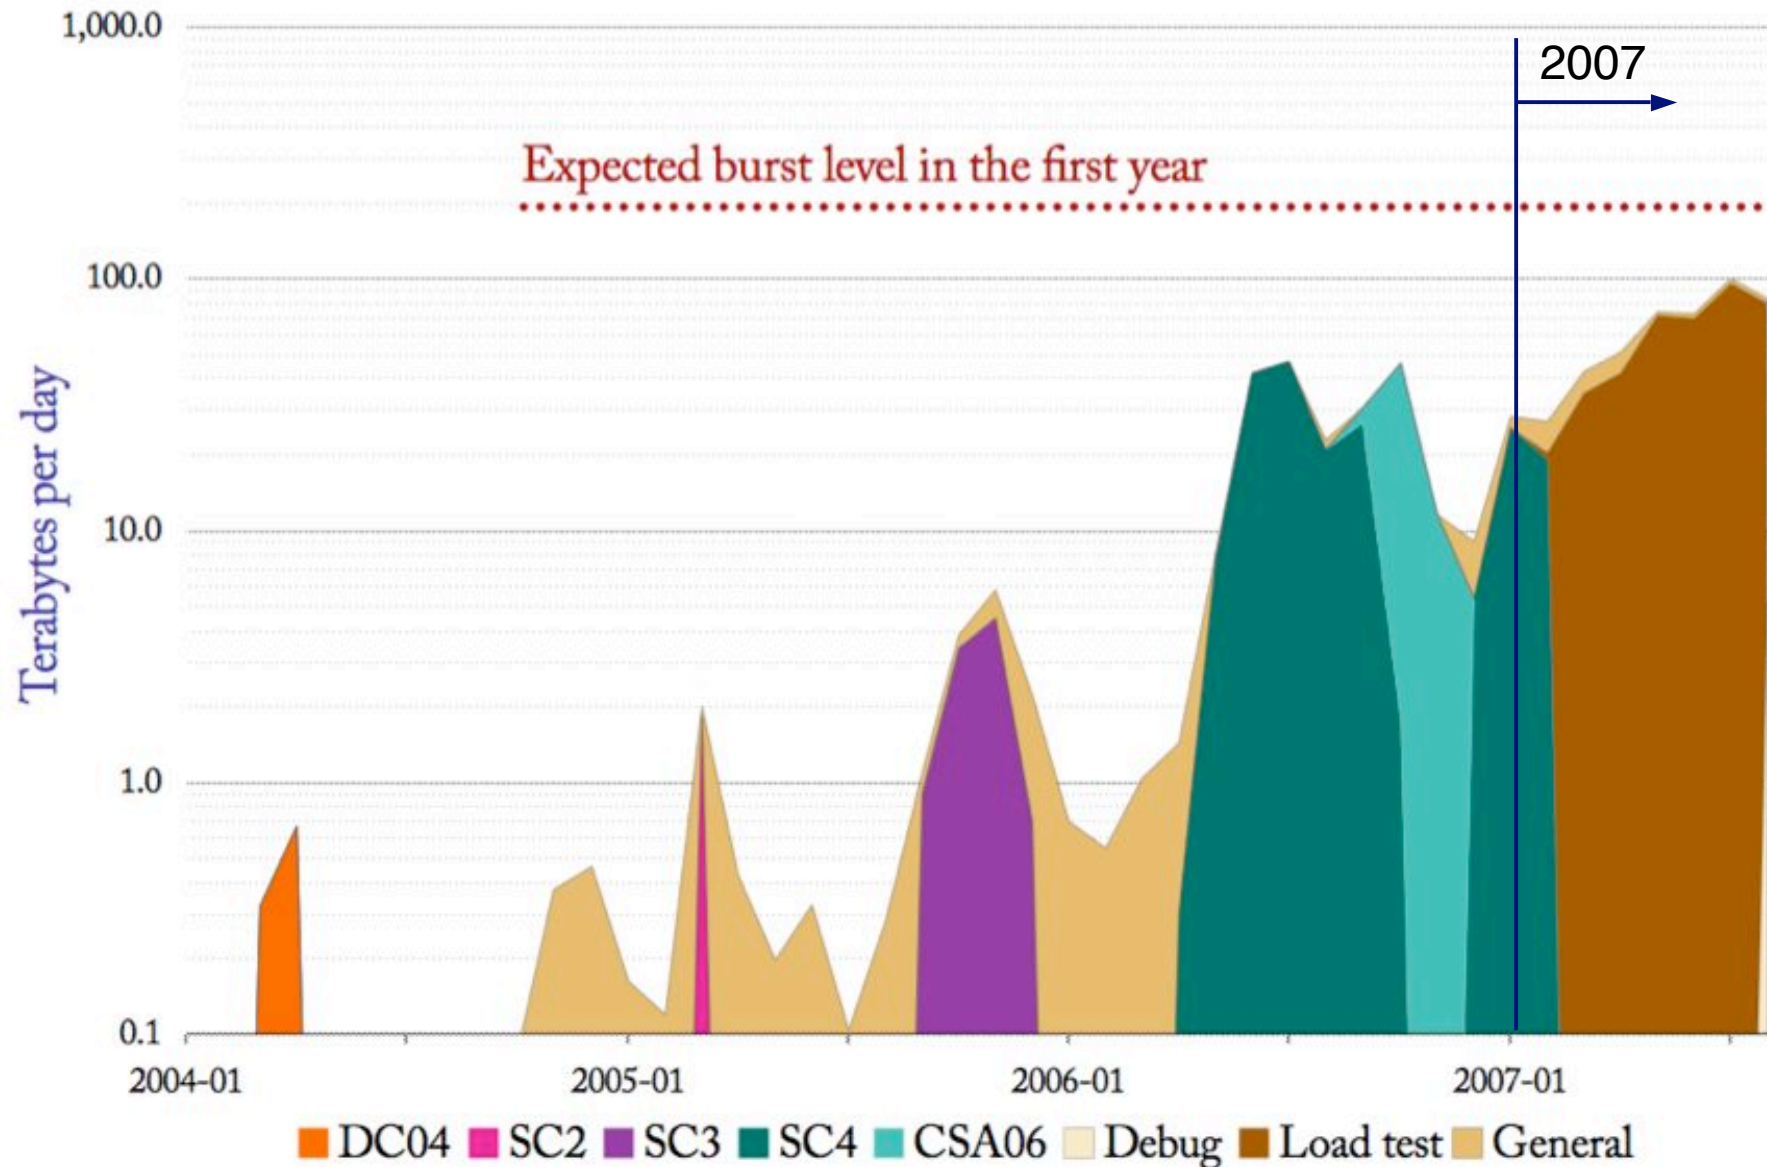


CMS PhEDEx data transfers

[courtesy of L. Tuura]

#258

Total data volume month by month





Data transfers: grandview




- CMS is exercising the data placement system since 2004
 - ❑ PhEDEx as a reliable, scalable data replication system
 - ❑ PhEDEx fully interfaced with gLite FTS since yrs (now: v2.0 ready also)

- CMS data transfers in a nutshell:
 - ❑ 7 T1s, 49 T2s, 13 T3s in current PhEDEx transfer topology (~40% of CMS Inst.)
 - ❑ Data transfers operated as if experiment was already running
 - ❖ we have had 2 service outages exceeding 24 hrs in the last 2 yrs
 - ❑ Current transfers at 1.2 GB/s \approx 100 TB/day \approx 3 PB/month global average rate
 - ❖ average CMS files size likely to be \sim 2 GB \Rightarrow \approx 50k files/day


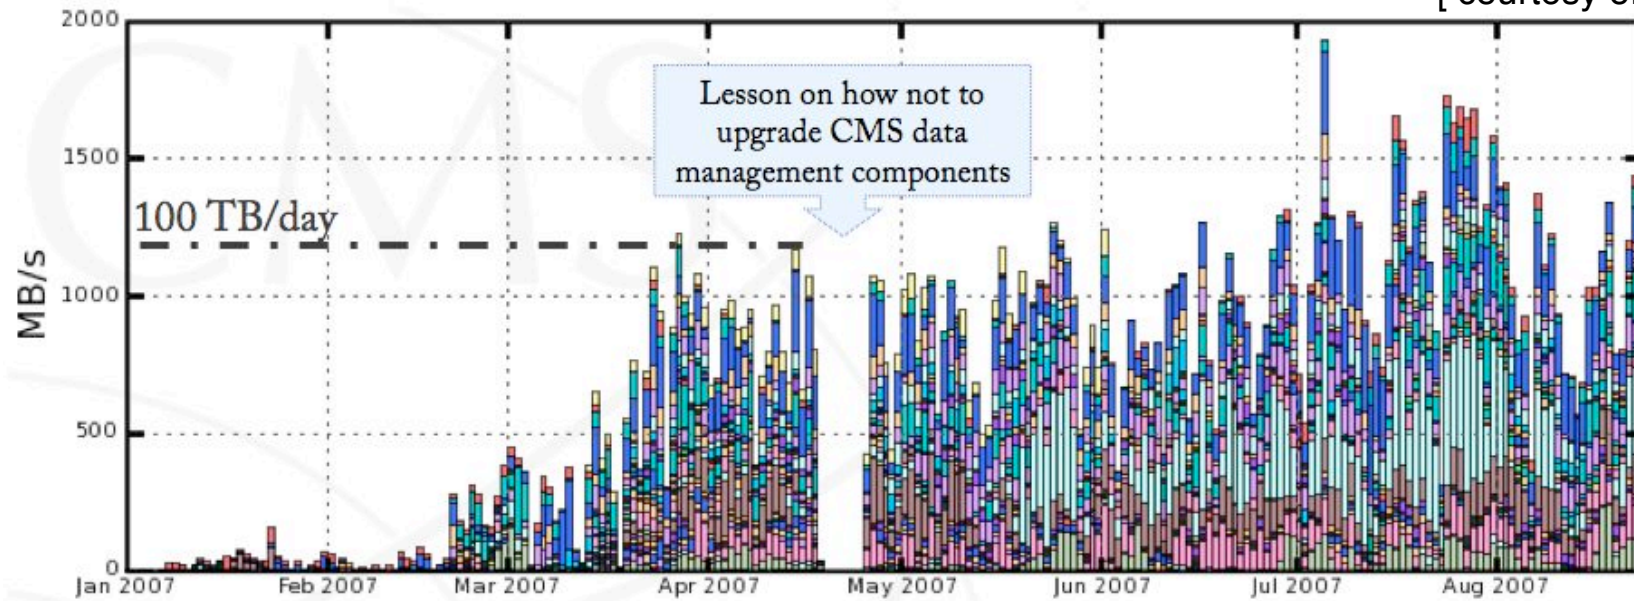
- With 1 yr to LHC start-up, CMS approaches the real transfers in scale, but not yet the full complexity →
 - ❑ From reliable transfers over the full transfer mesh, to multi-VO exercises...

- The progress is evident, though. Main sources of this are:
 - ❑ A well-designed, robust, scalable transfer system
 - ❑ A remarkable manpower investment to commission the transfer system

- Continuing efforts are needed on **debugging data transfers** more →




#258








Quality of transfers



- Evident improvement in Tiers participation to test transfers since LoadTest started
 - ❑ New PhEDEx 'probing' features also did help much
- Not evident improvement in quality, though
 - ❑ It's not one problem that lacks solution, there is a wide span of them
 - ❖ Greenish quality plots = successful transfers with fewer retrials = only when storage at both ends + network + PhEDEx set-up + site config + operators work fine... simultaneously!




Positive note: now stably at a ~challenge traffic load
Need to: Keep sites exercised + Debug and improve quality






Debugging Data Transfers




- Maintain a high-quality transfer network
 - ❑ to be handed over to CMS Data Ops
- Debug/commission *transfer links* among CMS Tiers
 - ❑ a Task Force (DDT-TF) is in charge since July 2007
 - ❖ Joined effort with CMS Facilities/Commissioning, T2 liasons, PhEDEx, Central Ops, FTS/SRM experts, network experts, site admins, ...
- Troubleshooting → by Tiers, work by milestones
 - ❑ focus on watching logs, ping site admins, fix problems
 - ❑ commission T1 ↔ T1, T1 downlinks to T2, T2 uplinks to T1, ...
- Infrastructural issues → by Facilities/Network projects
- Overall activity → DDT-TF, work on deliverables
 - ❑ E.g. a real-time status map (with reasons) of all Tier-X/Tier-Y links
 - ❑ E.g. a number of documented “success stories” in troubleshooting



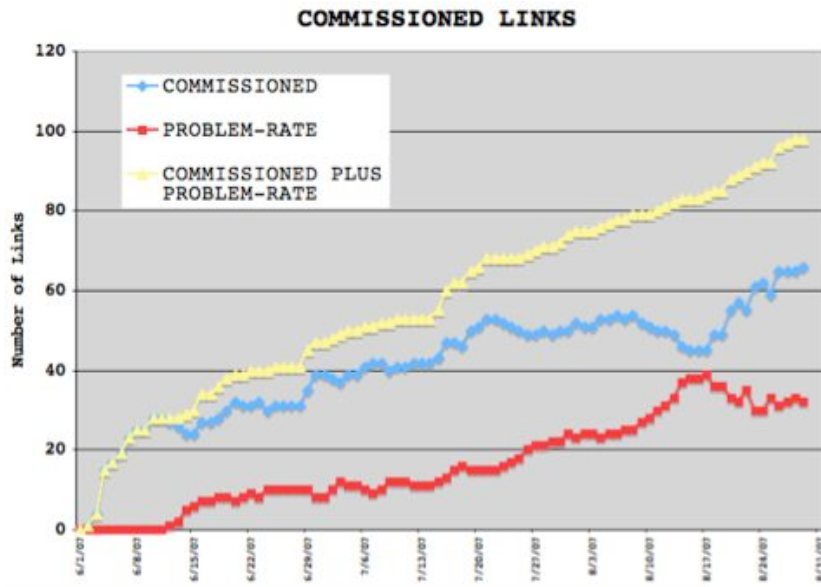
Gain confidence



Debugging Data Transfers



See CMS sessions at WLCG workshop



Steady influx of new links...

Legenda:

ROW→COLUMN: upper half of box
COLUMN→ROW: lower half of box

States:



T0↔T1

	ASGC	CNAF	FNAL	FZK	IN2P3	PIC	RAL
T0							

T1↔T1

	ASGC	CERN	CNAF	FNAL	FZK	IN2P3	PIC	RAL
ASGC								
CERN								
CNAF								
FNAL								
FZK								
IN2P3								
PIC								
RAL								


T1↔T2

	ASGC	CERN	CNAF	FNAL	FZK	IN2P3	PIC	RAL
Bari								
Beijing								
Belgium_IHHE								
Belgium_UCL								
Budapest								
CIEMAT_TMP								
CSCS								
Caltech								
DESY								
Estonia								
Florida								
GRIF_DAPNIA								

[... plus many more ...]



Summary



Thanks to all CMS people and site admins for the excellent and constant work to support CMS activities.

- Higher involvement of Tiers in operations
 - ❑ Tiers are shaken out, stability/quality concerns still drive efforts
 - ❑ “Help sites help themselves”
 - ❖ Keep checking site readiness, keep providing help
- Major boost in performance and scale of **data processing**
- Scale improvements in **data transfers**
 - ❑ Focus on achieving *stable* improvements through *clear* procedures
- CSA07 and CCRC08 stand as crucial exercises



Related contributions at CHEP07



- #145 - Testing and integrating the WLCG/EGEE middleware in the LHC computing
- #352 - The CMS Data and Workflow management system
- #369 - CMS experiences with computing, software and analysis challenges
- #280 - The CMS LoadTest: an infrastructure to exercise CMS transfer routes among WLCG Tiers
- #277 - CMS CSA06 experience at INFN
- #288 - Exercising CMS dataflows and workflows in computing challenges at Spanish Tier-1 and Tier-2 sites
- #325 - The CMS dataset bookkeeping service
- #285 - CMS Monte Carlo production in the WLCG Computing Grid
- #239 - Readiness of CMS Simulation towards LHC Startup
- #235 - CMS MC production System Development & Design
- #240 - WLCG scale testing during CMS data challenges
- #325 - The CMS Dataset Bookkeeping service
- #322 - CMS conditions data access using FroNTier
- #260 - CMS centres for control, monitoring, offline operations and analysis
- #314 - CRAB, CMS Remote Analysis Builder
- #258 - Scaling CMS data transfer system for LHC start-up
- #252 - CMS software deployment on the Open Science Grid (OSG)
- #259 - Analysing CMS software performance using IgProf, OProfile and callgrind
- #290 - CMS Tier-0: design, implementation and first experiences
- #353 - US CMS Tier-2 Computing
- #266 - CMS Offline Web Tools
- #185 - CMS Tier structure and operation of the experiment-specific tasks in Germany
- #223 - Searching for CMS data: Rapid web development using python and AJAX
- #224 - A multi-dimensional view on information retrieval of CMS data
- #225 - Analysis environments for CMS

+ more talks on facilities status and readiness (not listed here)
+ CMS sessions at WLCG workshop