Bringing the CMS Distributed computing system into scalable operations

S. Belforte (INFN), A. Fanfani (INFN), I. Fisk (FNAL), J. Flix (CIEMAT),
José M. Hernández (CIEMAT), T. Kress (RWTH), J. Letts (UCSD),
N. Magini (CERN), V. Miccio (INFN), A. Sciabà (CERN)
On behalf of the CMS Computing Group

XVII International Conference on Computing in High Energy and Nuclear Physics

Introduction

- Establishing efficient and scalable operations of the CMS distributed computing system requires proper integration, commissioning, scale testing and monitoring of
  - Underlying computing infrastructure
  - The data and workload management tools
  - The computing workflows
- Fill the gap between computing development and operations
- Computing challenges revealed the need for a sustained effort on computing integration and commissioning activities in CMS
  - Task Forces
  - Campaigns
  - End-to-End tests
  - Tools and Procedures

CMS Computing Model

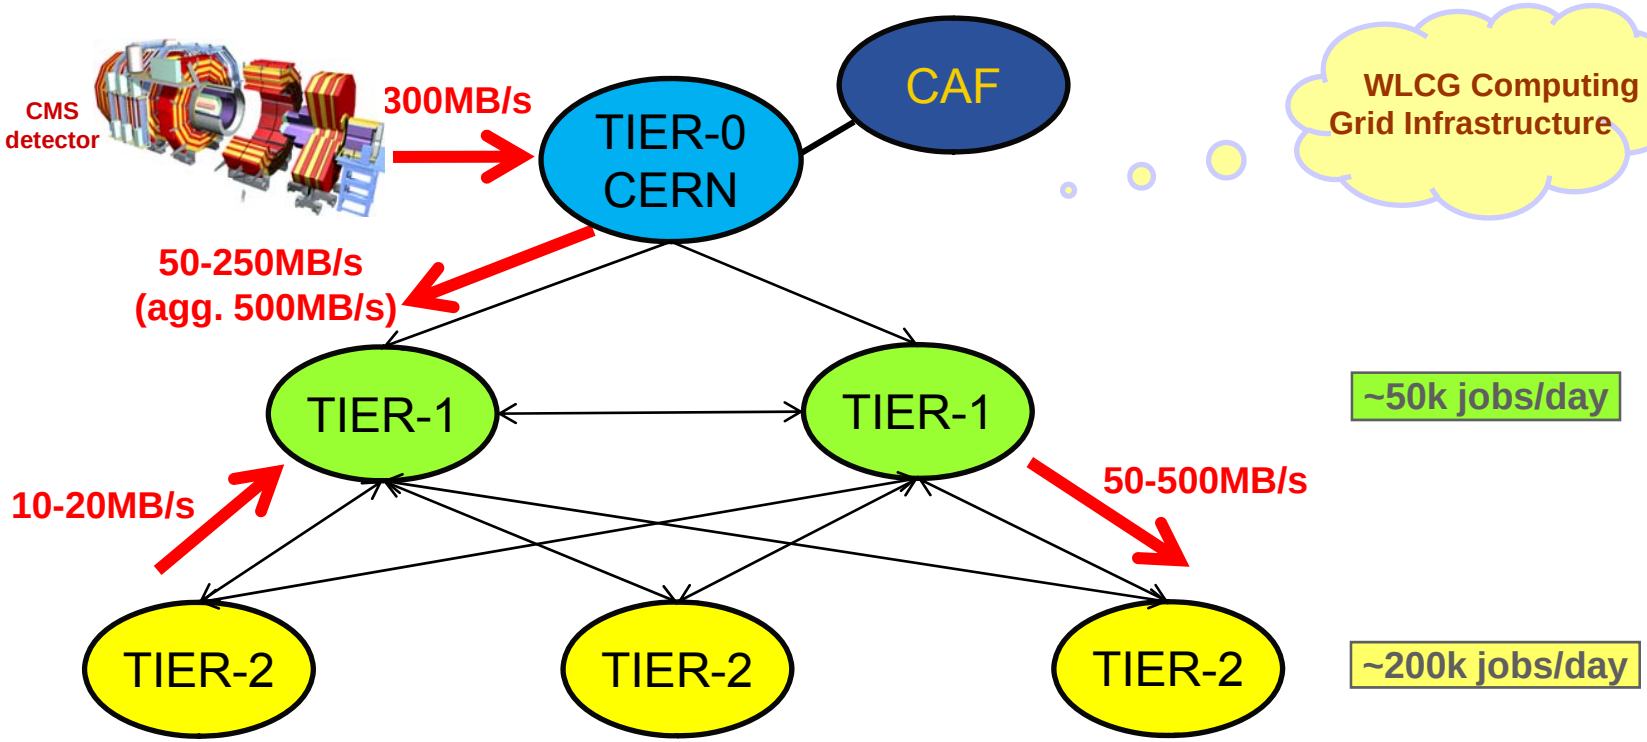

Tier-0

(the accelerator centre)

Data acquisition & initial processing
Long-term mass data storage
CMS CERN Analysis Facility
(latency critical data processing, high priority analysis)
Distribution of data → Tier-1 centres

7 Tier-1s

("online" to the DAQ)

High availability centres
Custodial mass storage of share of data
Data reconstruction and reprocessing
Data skimming & selection
Distribute analysis data → Tier-2s

~50 Tier-2s

in ~20 countries

End-user physics analyses
Detector Studies
Monte Carlo Simulation → Tier-1

Computing Challenges

- CMS has undertaken periodic challenges of increasing scale and complexity to test its computing model and Grid systems
- Performance values measured, problems identified, feedback into the design, integration and operation provided
- DC04 (2004), SC3 (2005), SC4 & CSA06 (2006), CSA07, CCRC08/CSA08
- Lots of lessons learnt
  - Sustainable operations difficult to achieve
  - Robustness to deal with intrinsic unreliability of distributed computing
  - Need continuous reliability and scaling tests
  - Continuous validation of computing infrastructure and tools
  - Tuning and finding operating points takes time
  - Better integrated monitoring
  - Need to improve reliability and performance of data access

Processing And Data Access Task Force

- Need for a sustained effort on computing integration and commissioning activities
- Processing And Data Access (PADA) Task Force
  - Validating the infrastructure for organized processing and user analysis including the sites and the workload and data management tools
  - Validating the distributed production system by performing functionality, reliability and scale tests
  - Helping sites to commission, configure and optimize the networking and storage through scale-testing data transfers and data processing
  - Improving the efficiency of accessing data across the CMS computing system from global transfers to local access

PADA Task Force Activities

- Data Transfers commissioning
  - Poster [205] J. Letts et al. “Debugging Data Transfers in CMS”
- Site commissioning
  - Poster: [292] J. Flix et al. “The commissioning of CMS sites: improving site reliability”
- Distributed Production and Analysis Systems commissioning
  - Talk: [139] G. Codispoti et al. ”Use of the gLite-WMS in CMS for production and analysis”
  - Talk: [220] S. Padhi et al. “Use of glide-ins in CMS for production and analysis”
  - Poster: [213] A. Fanfani et al. “Commissioning Distributed Analysis at the CMS Tier-2 Centers”
- Monitoring
  - Poster [275] P. Saiz et al. “Generic monitoring solution for LHC site commissioning activity and LHC computing shifts”

Commissioning data transfers

- CMS needs to have working T0 → T1 ↔ T1 ↔ T2 links
  - Tier-0 → Tier-1: to export raw and prompt-reconstructed data
  - Tier-1 ↔ Tier-1: for synchronization of AODs after data reprocessing
  - Tier-1 → Tier-2: for distribution of data to be analyzed at Tier-2 sites
  - Tier-2 → Tier-1: to upload Monte Carlo events generated at Tier-2 sites
- Debugging Data Transfers (DDT) task force
  - Active since July 2007, defined the metrics, provided a procedure and tools to test transfer links and assisted sites in solving problems
  - Use PhEDEx data transfer load generator
  - Commission links in Debug instance before activation in Production
    - 20 MB/s sustained for 24h for T0 → T1, T1 ↔T1 and T1 → T2 links
    - 5 MB/s sustained for 24h for T2 → T1 links
  - Lots of problems found and fixed: configuration (PhEDEx, FTS, SRM, gridFTP, network, SE), I/O bottlenecks
  - Extremely useful activity

DDT Task Force

- Significant increase in production data transfer quality and throughput

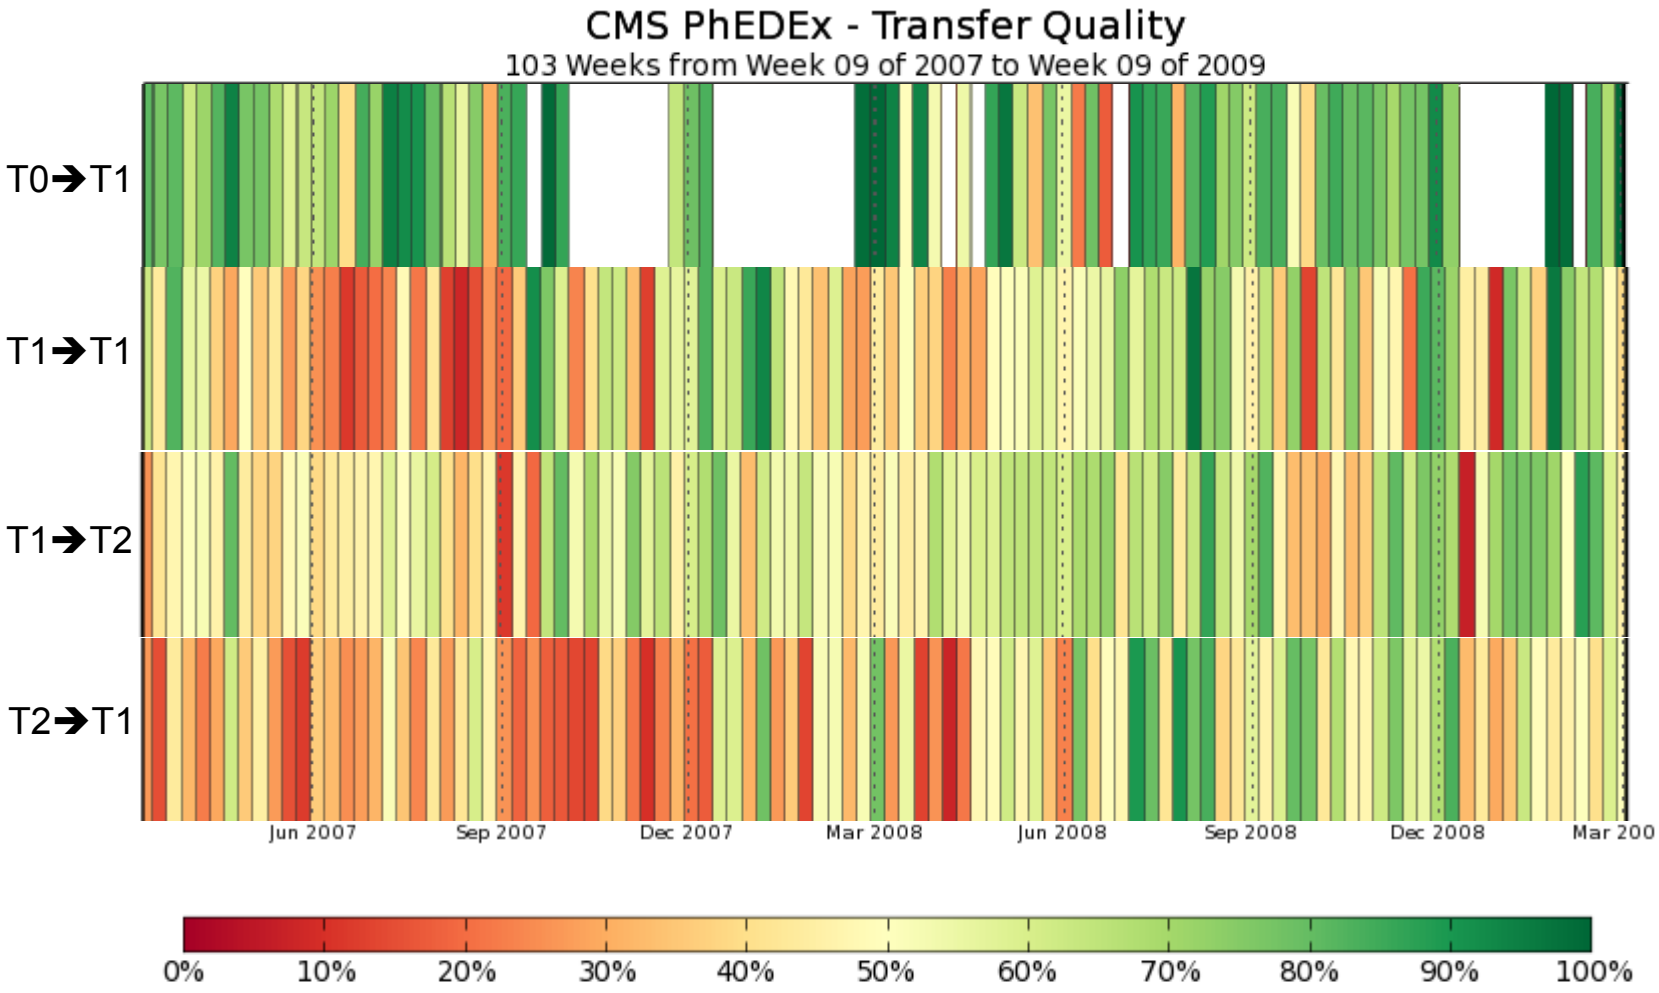

PhEDEx Data Transfer System

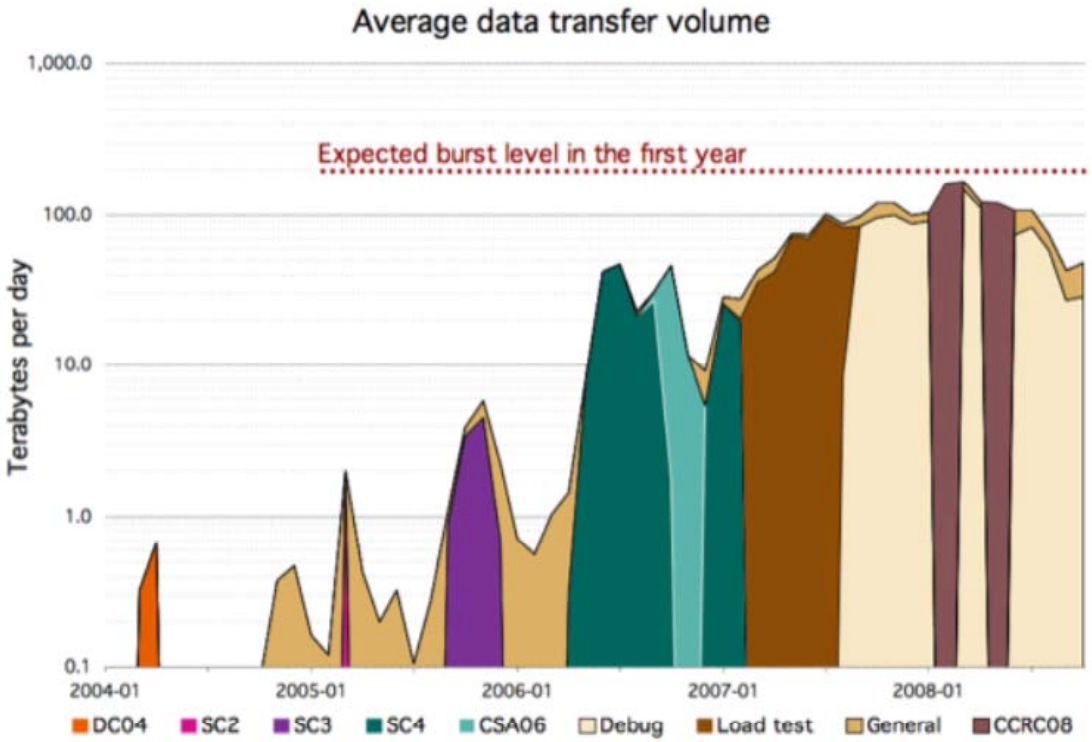

- 60 PB transferred over the WAN in the last 2 years
- Routine WAN transfers of ~50-100 TB/day (1-2 GB/s)
- ~ 30-40 TB/day backfill transfers (~500 links)
- Transfer system scales well

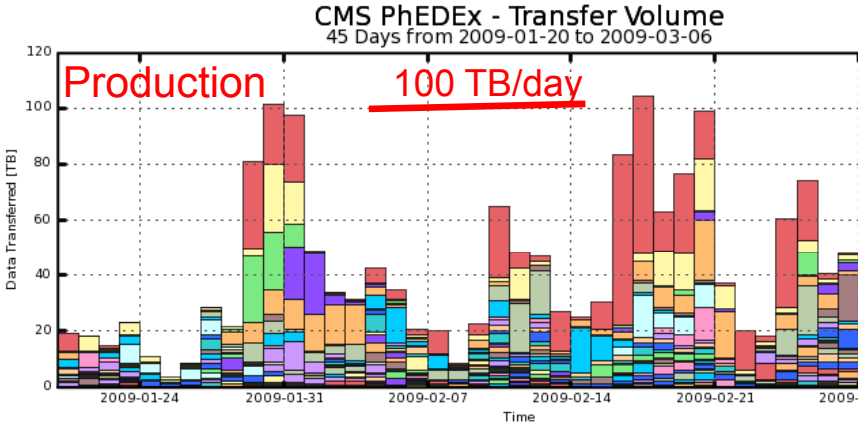

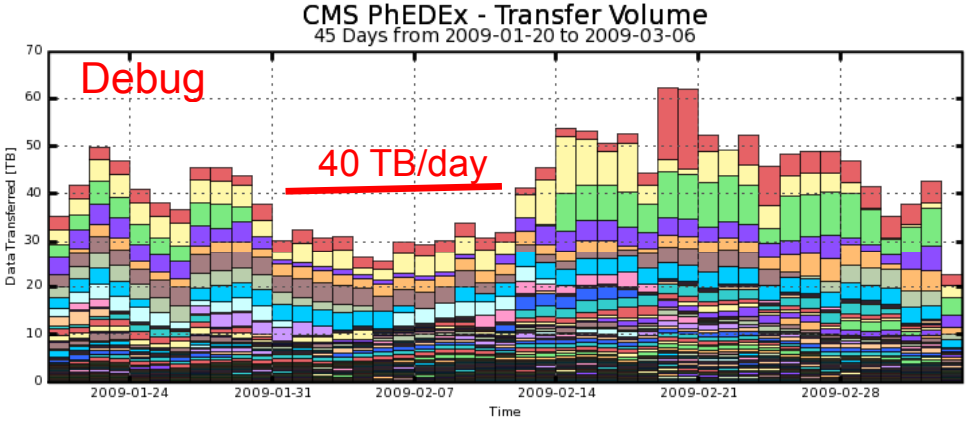

Site Commissioning

- CMS distributed computing requires stable and reliable behavior of the underlying infrastructure at all times
  - Heterogeneity and different amount of computing resources and support level
- Site Monitoring framework to track site readiness
  - CMS SAM tests
    - Jobs sent to sites to test specific services
  - JobRobot job load generator
    - Simple jobs reading data
  - Data transfers
    - Transfer quality and commissioned links
- Site readiness metrics established to guarantee data processing can be performed efficiently and reliably
  - Site Status Board
  - Production and analysis using 'good' sites

Site Readiness: SAM tests

- Site Availability Monitoring – CMS SAM tests
  - High priority jobs submitted every hour
  - Test CE, SE, experiment software, conditions cache, data read, stage out, etc
  - Require daily availability > 80% for T2s and > 90% for T1s

Sitename	Service Type	Service Name	mc	j
T1_DE_FZK	CE	ce-1-fzk.gridka.de	ok	o
		ce-2-fzk.gridka.de	ok	o
		ce-3-fzk.gridka.de	ok	o
		ce-4-fzk.gridka.de	ok	o
		ce-5-fzk.gridka.de	ok	o
	SRMv2	gridka-dCache.fzk.de		
T1_ES_PIC	CE	ce-test.pic.es		o
		ce05.pic.es	ok	o
		ce06.pic.es	ok	o
		ce07.pic.es	ok	o
	SRMv2	srmcms.pic.es		
T1_FR_CCIN2P3	CE	cclcgceli03.in2p3.fr	ok	o
		cclcgceli04.in2p3.fr	ok	o
	SRMv2	ccsrm.in2p3.fr		
T1_IT_CNAF	CE	ce04-lcg.cr.cnaf.infn.it	ok	o
		ce05-lcg.cr.cnaf.infn.it	ok	o
		ce06-lcg.cr.cnaf.infn.it	ok	o
	SRMv2	srm-v2-cms.cr.cnaf.infn.it		
		storm-fe-cms.cr.cnaf.infn.it		
T1_TW_ASGC	CE	lcg00125.grid.sinica.edu.tw	warn	o
		w-ce01.grid.sinica.edu.tw	warn	o
		w-ce02.grid.sinica.edu.tw	warn	o
		w-ce03.grid.sinica.edu.tw	warn	o
	SRMv2	srm2.grid.sinica.edu.tw		
T1_UK_RAL	CE	lcgce03.gridpp.rl.ac.uk	ok	o
		lcgce04.gridpp.rl.ac.uk	ok	o
	SRMv2	srm-cms.gridpp.rl.ac.uk		
T1_US_FNAL	CE	cmsosgce.fnal.gov	ok	o
		cmsosgce2.fnal.gov	ok	o
		cmsosgce4.fnal.gov	ok	o
	SRMv2	cmssrm.fnal.gov		

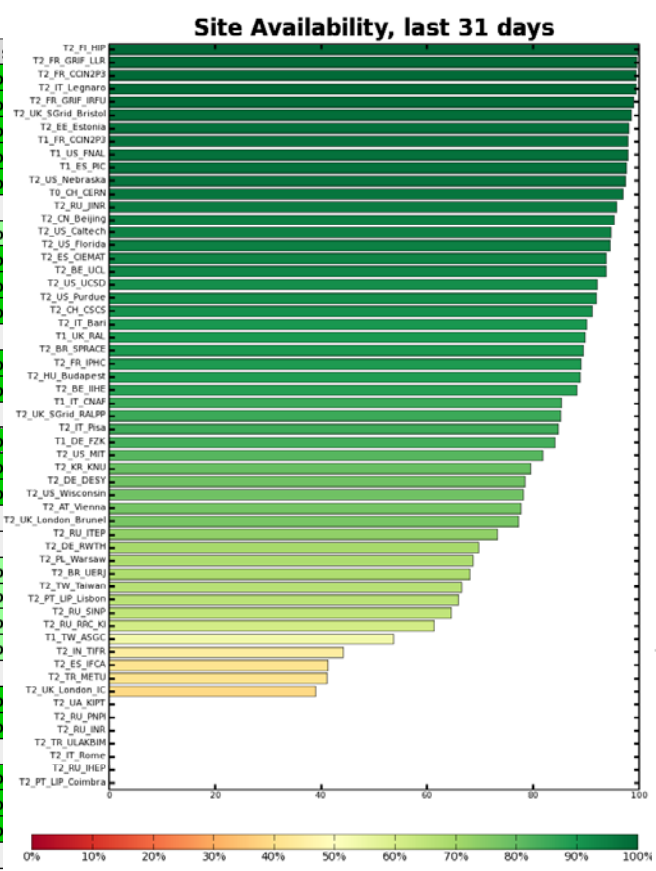

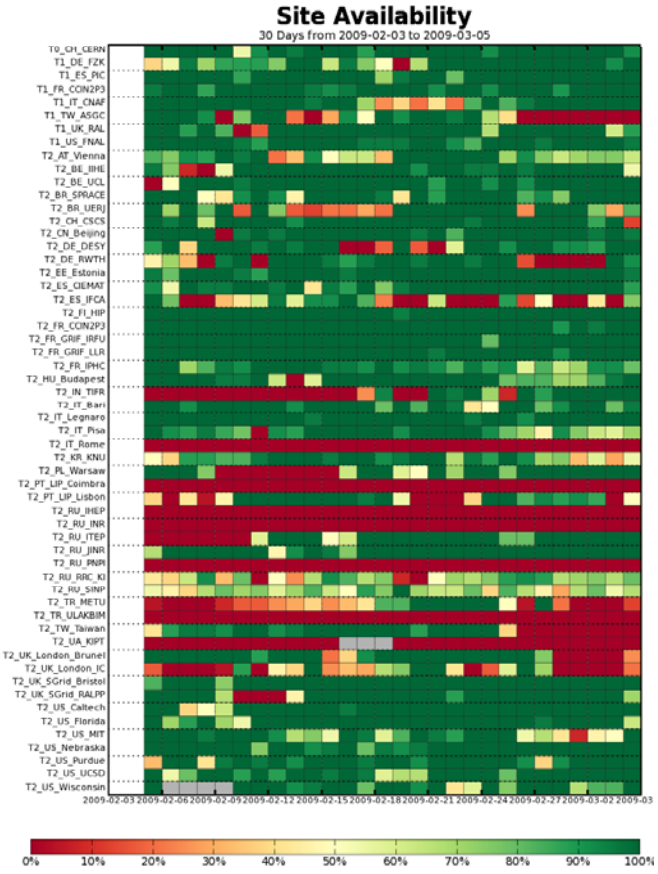

Site Readiness: JobRobot

- Job Robot load generator

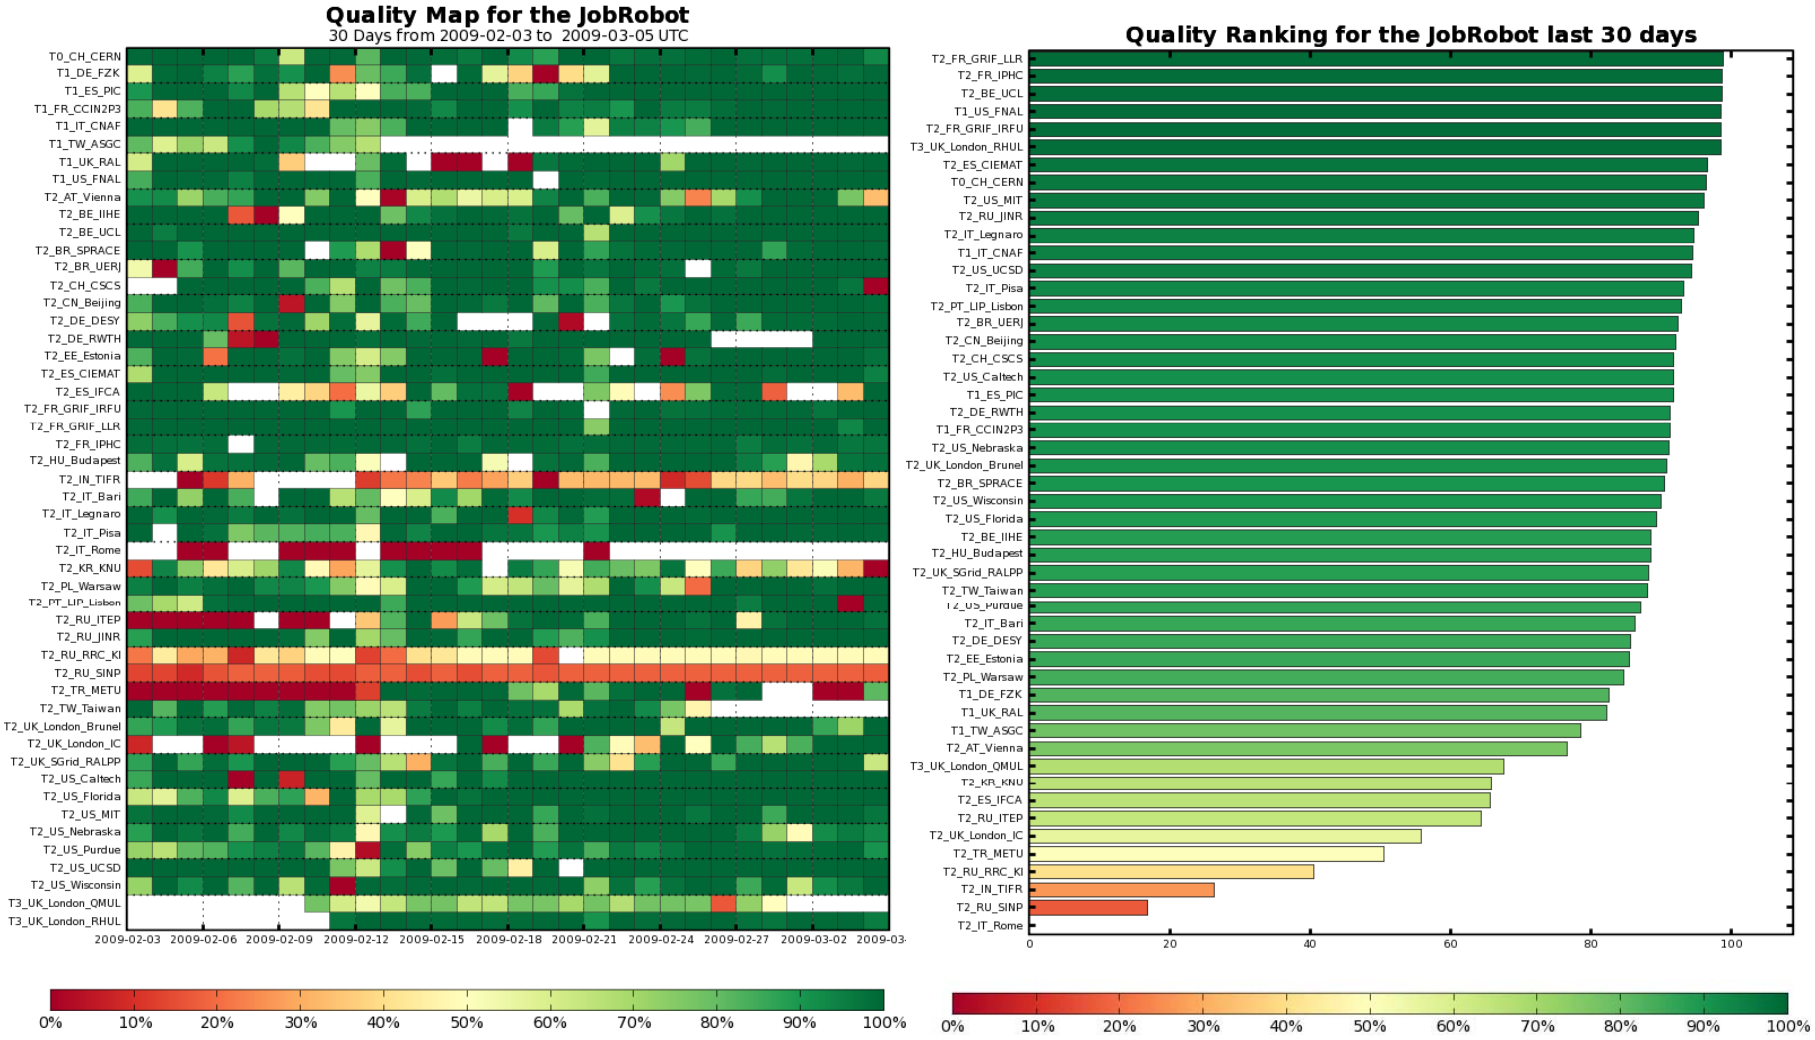

Site Readiness: data transfers

- Require minimum numbe
  - Tier-1s: ≥ 20 downlinks to
  - Tier-2s: ≥ 4 downlinks fror
- Require minimum numbe threshold
  - Count production and deb exercised at 0.5 MB/s)
  - 50% of the active links at l
    - Enough to detect system
    - Use transfer quality inforr

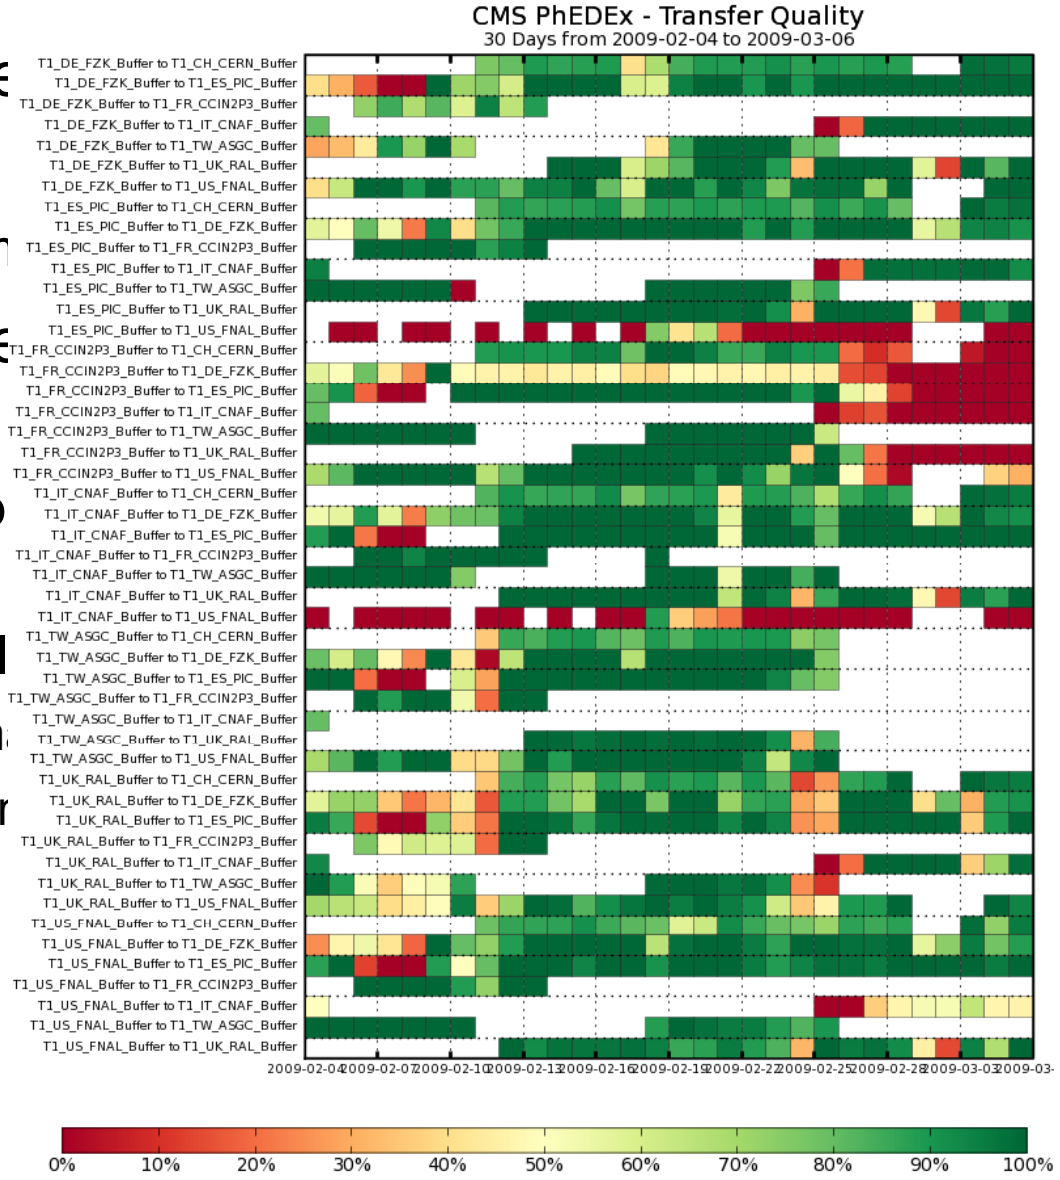

Site Readiness: putting all together

- Collect and displa…ation in Site Status Board
  - Developed by AF…eeva)
- (…
- …

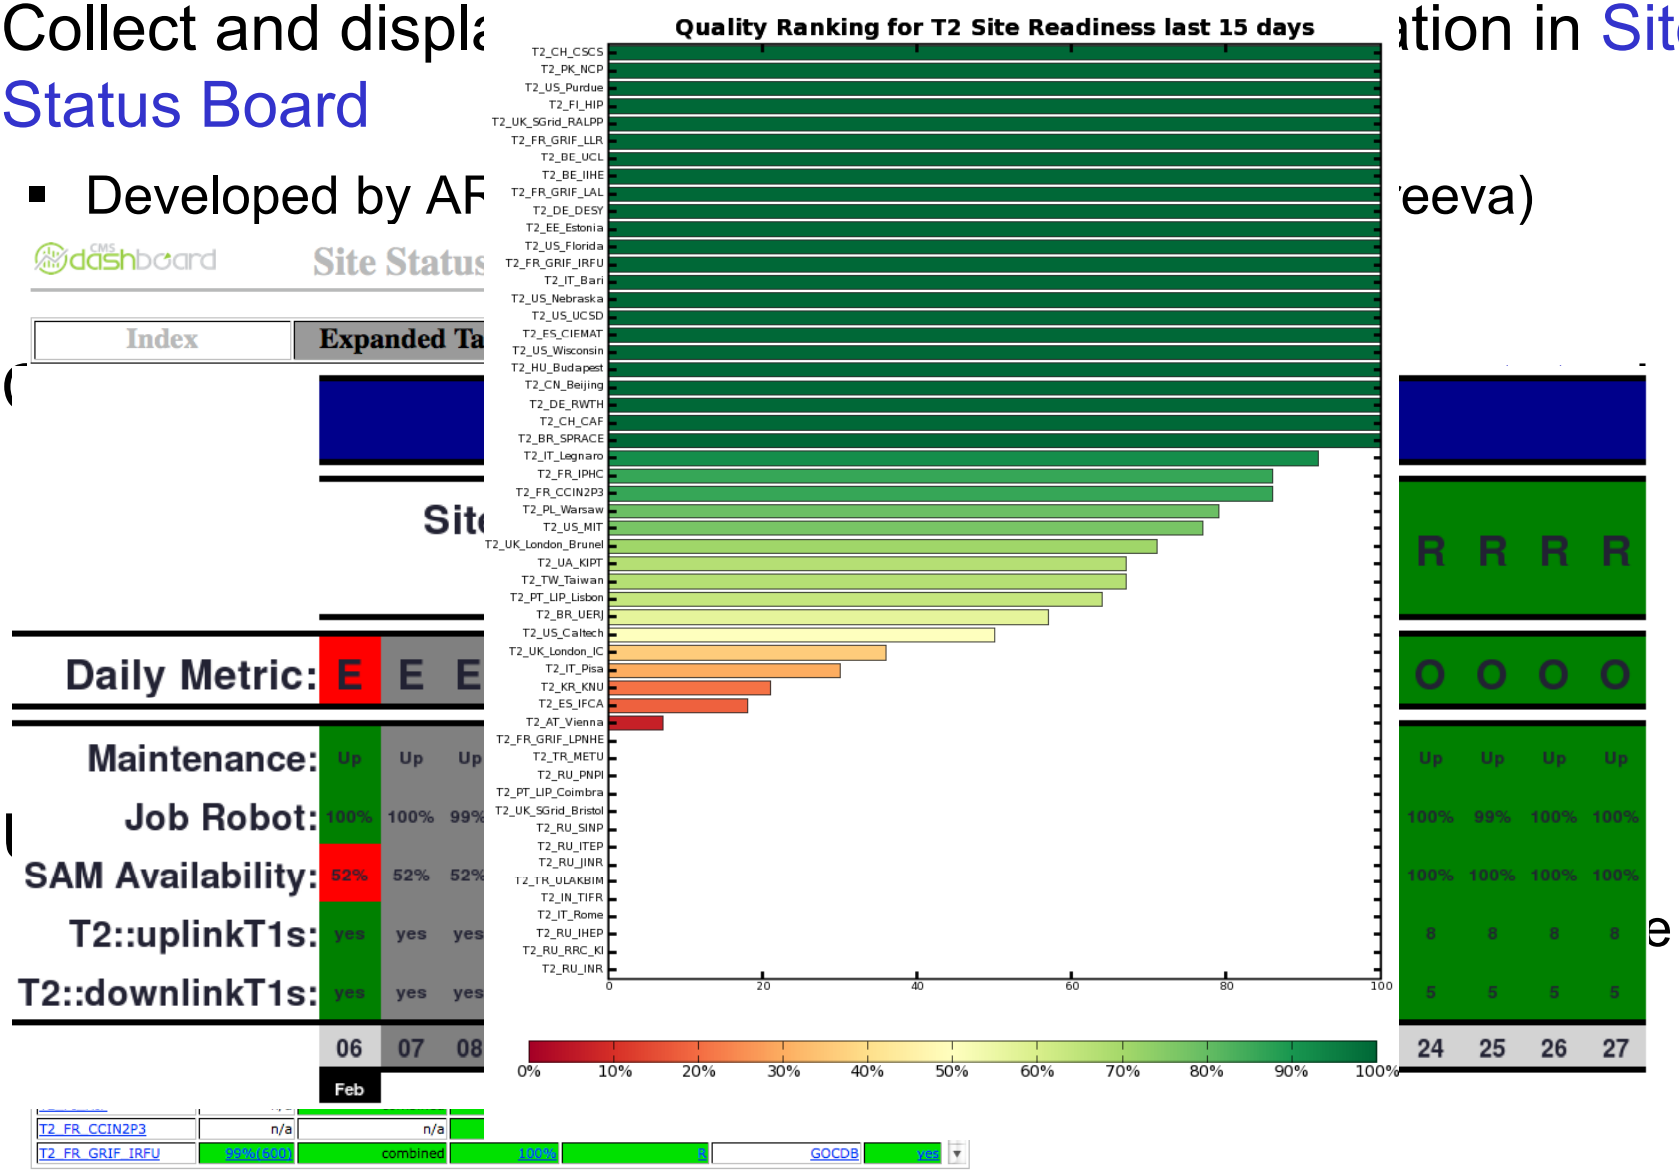

Site Readiness data

- Positive effects of site readiness program
  - Continuous monitoring of Grid & CMS services at sites
  - Helps production and users to select reliable T2 sites
- 6 months of data
- Still room for improvements
- Task Force
  - Determine usual failures
  - Help sites to improve
  - Feedback for robustness of CMS tools and services
  - Increase reliability of sites

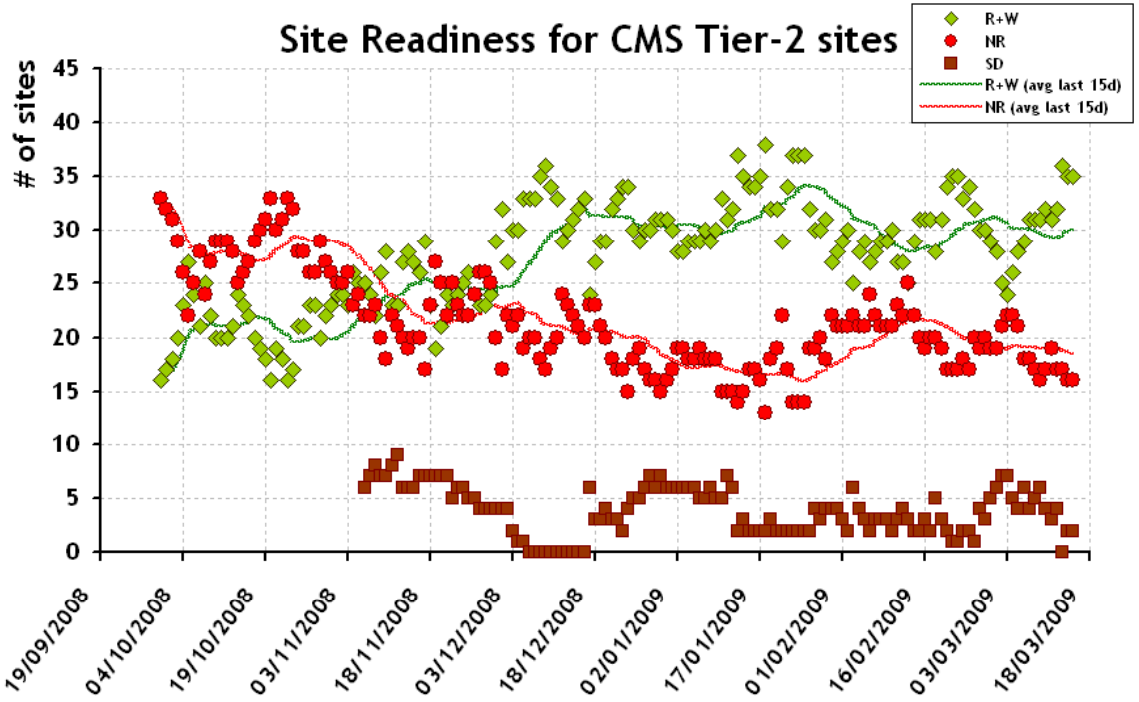

Commissioning Distributed Production and Analysis

- Systematic integration & validation of job management tools
  - ProdAgent for production activities and CRAB (Server) for user analysis
  - Various WMS used by CMS
    - gLiteWMS, condor-G and glideinWMS (pilot-based)
- Functionality, reliability and scaling tests
  - Integration of new components of the production and analysis systems
  - Scaling tests
    - 30-50 kjobs/day/instance.
    - No global scaling issues. Scaling by adding new instances
- Backfill of jobs at T1s (re)running real workflows
  - Continuous (scale) test of production tools and sites
  - Very important activity for sites and operations team
  - To be incorporated into the site readiness monitoring

Workload Management System

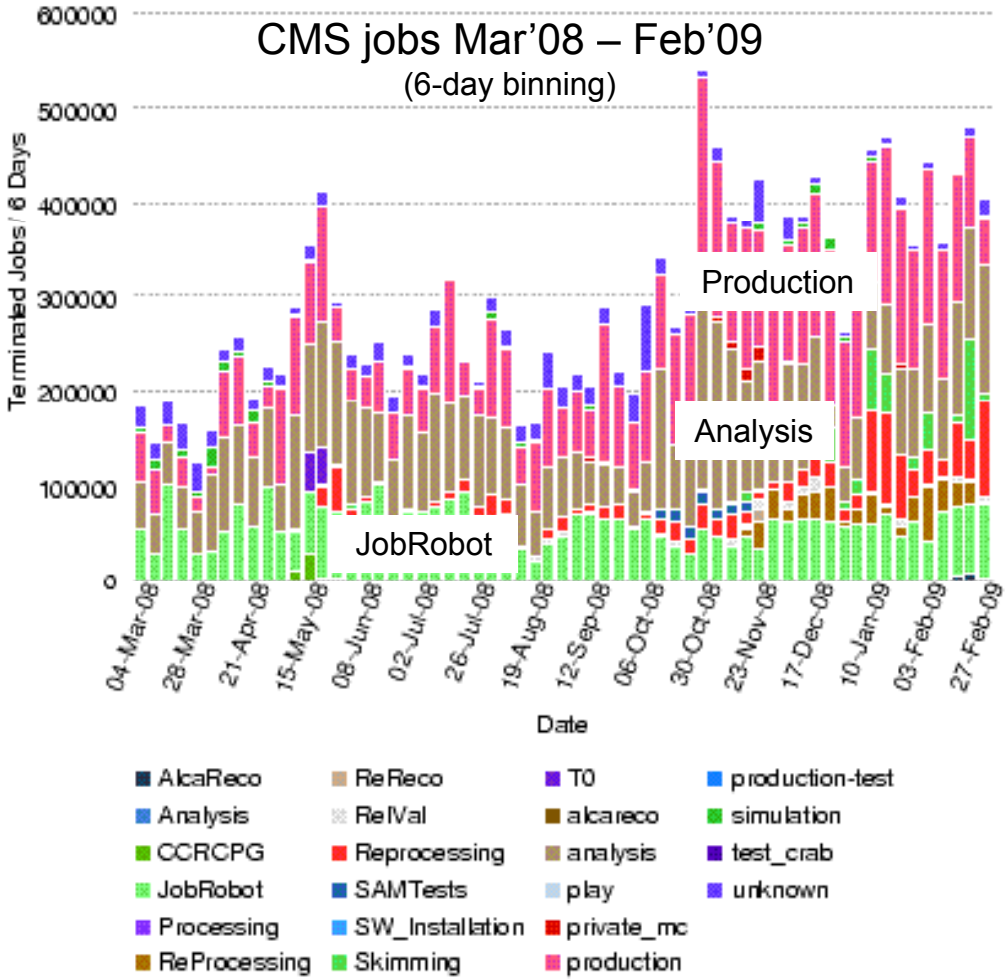

- ~ 70k/day jobs routinely handled
- ~ 150k/day reached during computing challenges
- Backfill of fake analysis jobs (~25k/day)
- No problems expected scaling up the WMS system

Job success rates last year

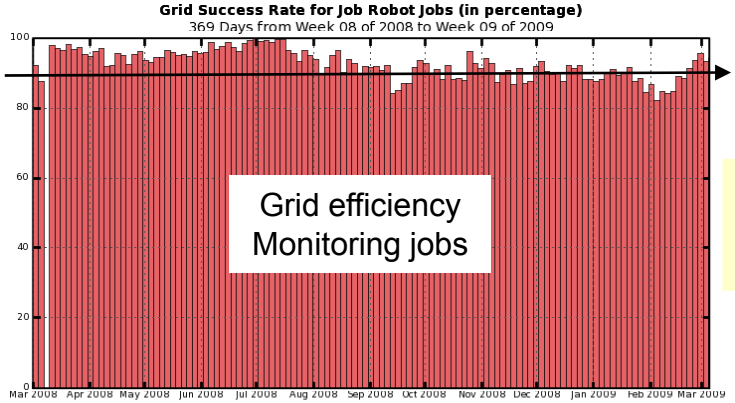

91%

94%

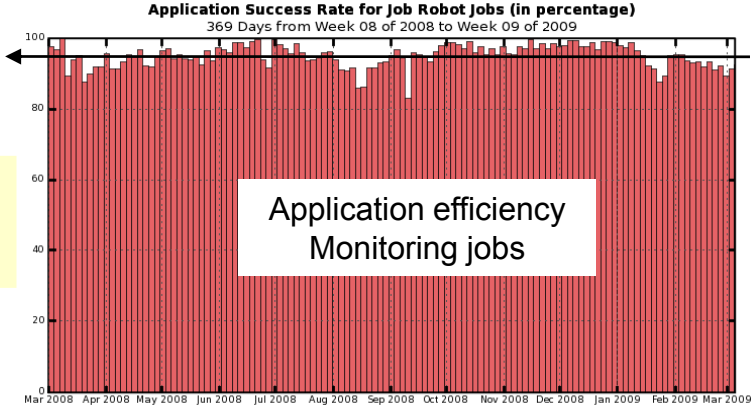

- Monitoring jobs success rate ~ 85%

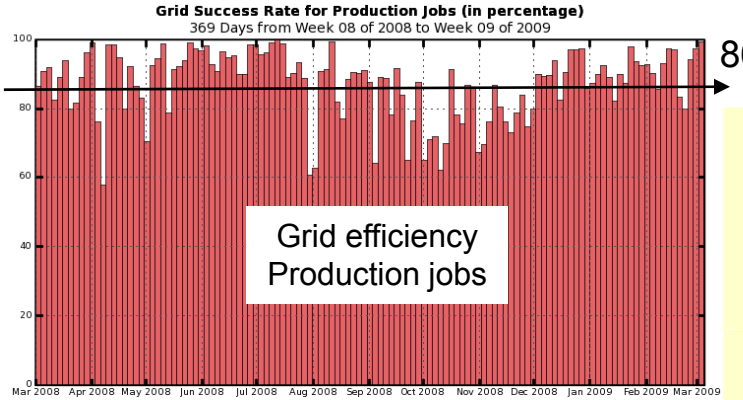

86%

94%

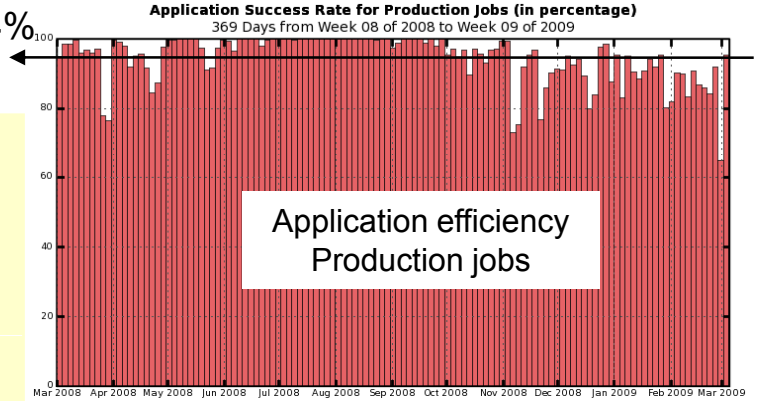

- Global production job efficiency ~ 80%
- Failed jobs automatically resubmitted
- Inefficiency impacts latency

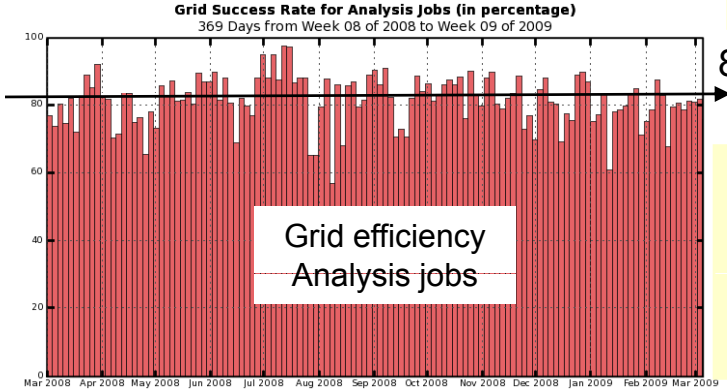

81%

71%

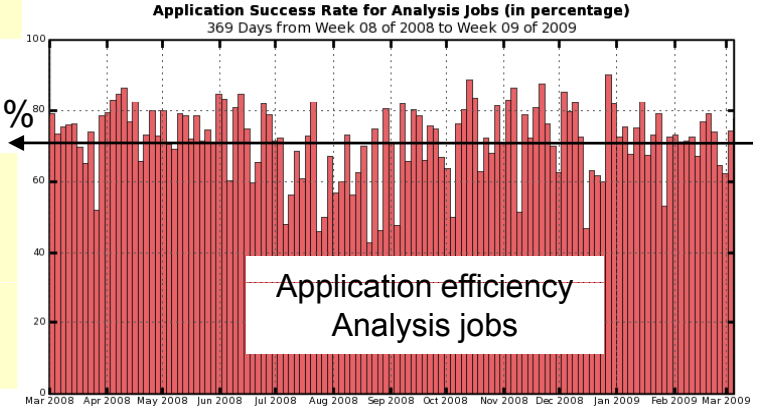

- Global User Analysis job efficiency ~ 60%
- Dominated by remote stage out

Thanks Julia Andreeva for providing the plots

Other integration activities

- Analysis Support Task Force
  - Improve reliability of analysis jobs
  - Analysis of failures, user support/feedback, improve monitoring, etc
  - Transition to Analysis Operations computing group
  - See talk: [207] J. Letts et al. “CMS Analysis Operations”
- Integration/commissioning campaigns
  - Optimize manpower and focus effort of various computing groups
  - Rollout of data consistency tools, rollout and scale testing of components of the production/analysis systems, etc
- Planned end-to-end tests for 2009
  - Consistent handling of luminosity and data quality from data (re)processing to analysis
  - Continuous reprocessing/skimming at T1s via backfill workflows
  - Reprocessing pre-staging data from tape

Summary

- Integration & commissioning activities crucial for bringing the CMS distributed computing system into scalable operations
  - Computing Challenges, Task Forces, Campaigns, End-to-End tests and continuous computing integration activity
  - Load generator tools (data transfers and jobs) very useful
- Data and Job Management systems scaling well
  - Routine large scale data transfers and job submissions
  - Continuous improvement in reliability and robustness of sites and tools
- Production/Analysis tools, Operations and Sites have greatly benefited from the Computing Integration & Commissioning Program