Status of CMS Operations

WLCG Collaboration Workshop, in conjunction with CHEP 2012
19. May 2011

Oliver Gutsche
for
Computing Operations

2011 Overview

- Again a very successful year for CMS
  - Computing had its share of the success
    - Data taking was rather smooth
      - CMSSW_4_2 was the primary release for 2011 but it's memory consumption was well above 2 GB per job at high numbers of PileUp interactions (16 and more)
        - Reduces available CPU resources to 70%, can only run 5 processed on 8 core box with 2 GB per core
      - New release (CMSSW_4_4) was available that improved the situation significantly
      - Collaboration decided to stay with CMSSW_4_2 to not interrupt ongoing analyses
        - We managed to keep the Tier-0 performing sufficiently well, but it was not easy
    - We provided many, many re-reconstruction passes like last year, in addition a full end-of-year re-reconstruction pass in CMSSW_4_4
    - We produced a lot of MC, re-processed it with several PileUp scenarios including a full CMSSW_4_4 pass over Xmas
    - Analysis was performing well on the T2 level digesting all the different data and MC samples and producing over 100 papers.
- All over all, we were very busy (as expected)!

Data taking 2011: Tier-0

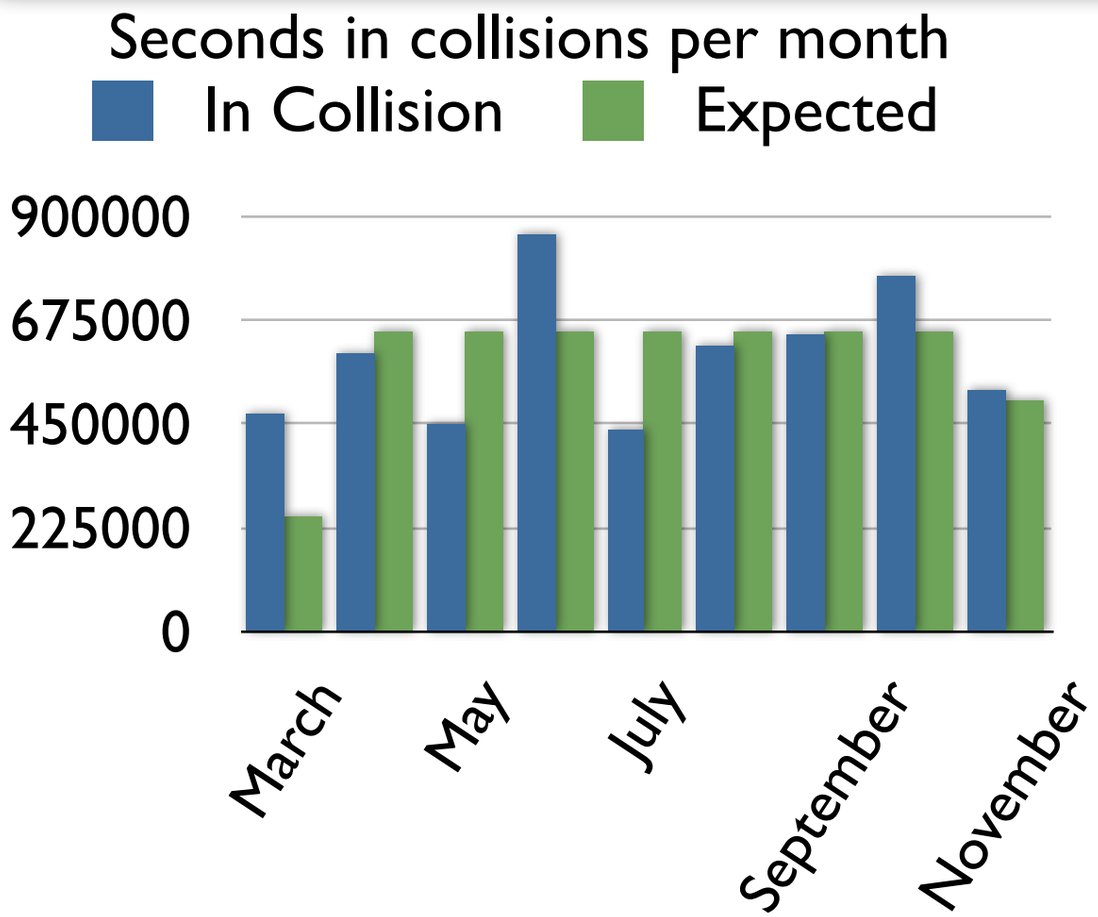

Month	Average Trigger Rate (with overlap)
March	356Hz
April	334Hz
May	393Hz
June	431Hz
July	361Hz

Tier	Observed (8 PU events)	Expected (8 PU events)	Observed (30 PU events)	Expected (30 PU events)
Data RAW [KB]	230	390	356	800
Data RECO [KB]	590	530	1316	900
Data AOD [KB]	165	200	328	250
MC RECO [KB]	970	600	-	1100
MC AOD [KB]	250	265	-	300

- 2011 data taking mostly followed the planning
  - Data taking rates of 350-400 Hz
  - PileUp of up to 16 interactions
  - Already glimpse for 2012: 30 PU interactions recorded in special runs

Data taking 2012: Tier-0

Datasets for core physics ~350 Hz

No change from 2011

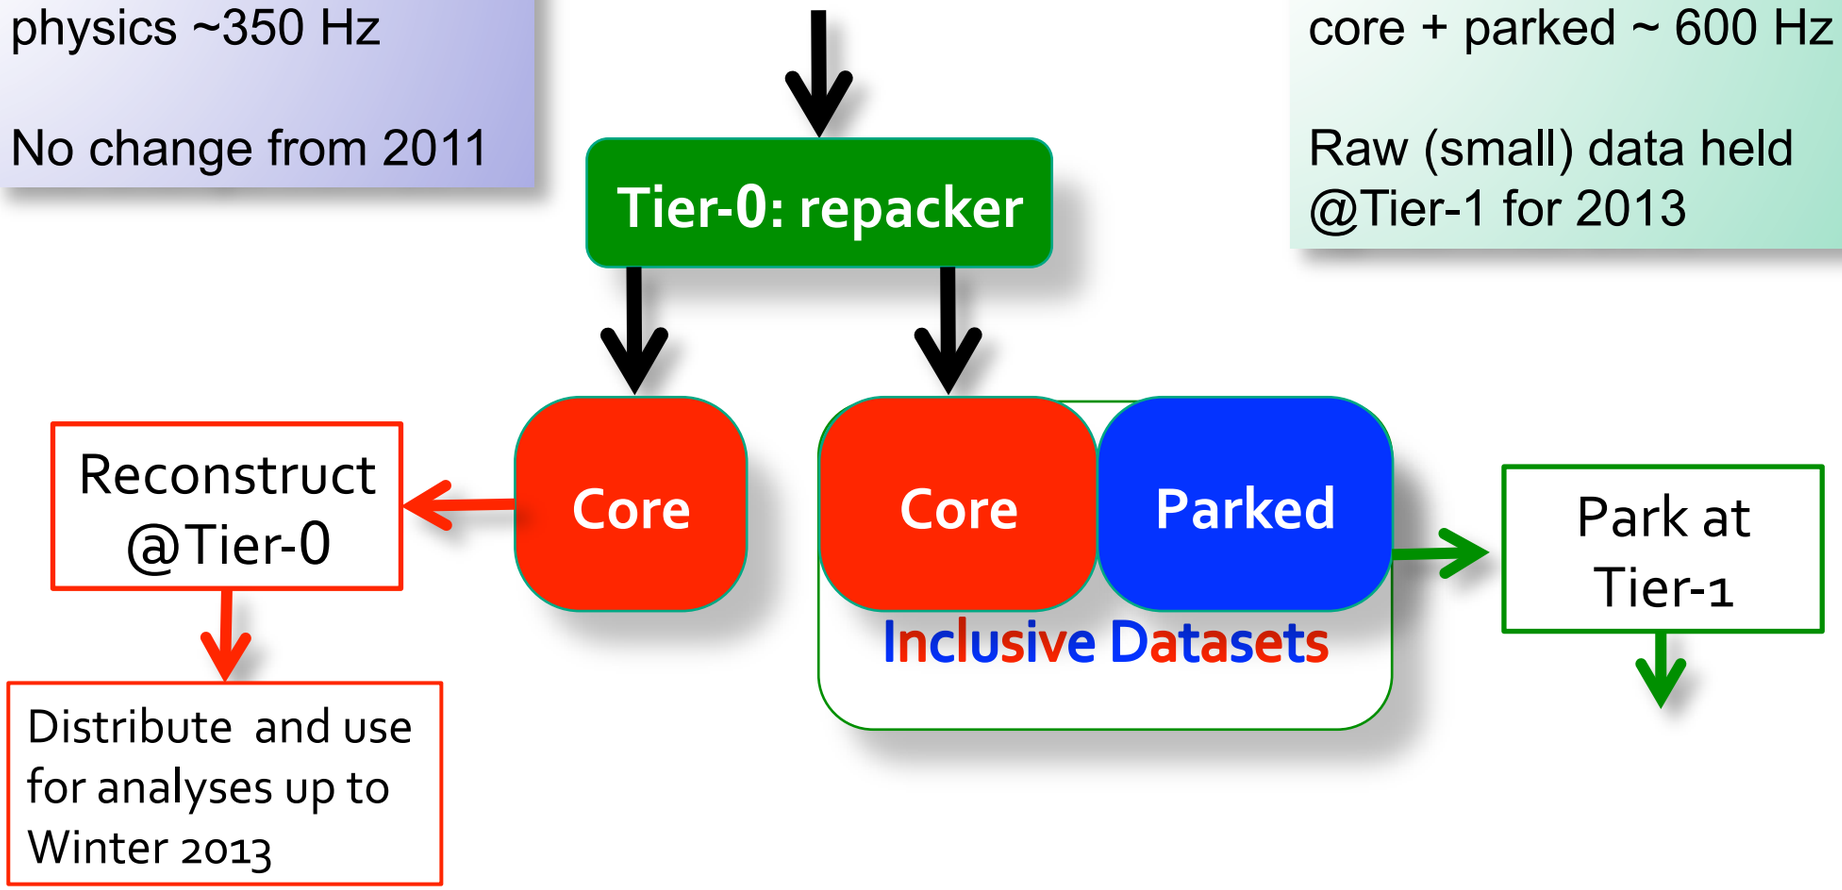

Inclusive datasets = core + parked ~ 600 Hz

Raw (small) data held @Tier-1 for 2013

- Planning:
  - Average PileUp: 30 interactions per crossing
  - Total data taking: 5.14 Million seconds pp & 700 Thousand seconds HI
- Run2012B (new era, started this week): Data taking: ~350 Hz, Data Parking: ~600 Hz

Data Re-Reconstruction 2011/2012

	AOD Events
Prompt	1,535
ReReco	3,826
Total	**5,360**

Million AOD Events in 2011

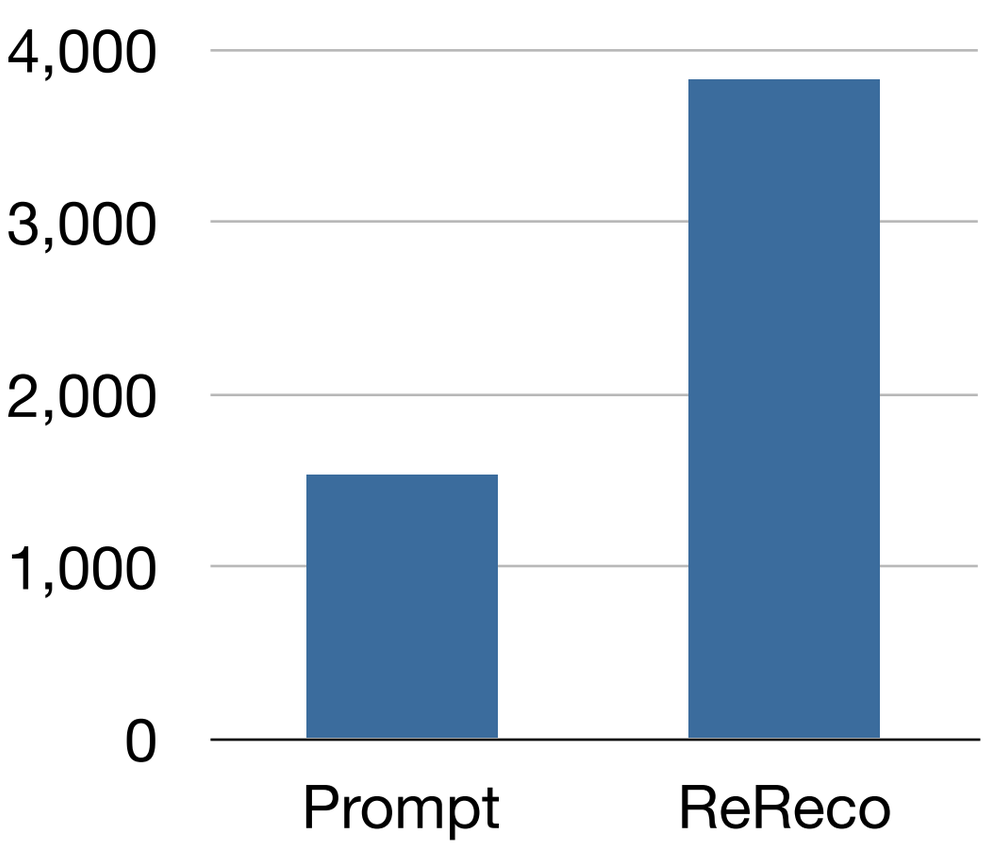

- 2011:
  - 29 individual re-reconstruction passes (complete and partial)
  - Total number of re-reconstructed events corresponds to more than 2 times of events recorded
  - Includes a complete 2011 re-reconstruction pass in CMSSW_4_4_X
- 2012:
  - Expect to follow trend and provide re-reconstruction passes during the year on the order of number of events recorded
  - No End-Of-Year re-reconstruction pass planned, rather beginning of 2013

MC: Overview Number of Events

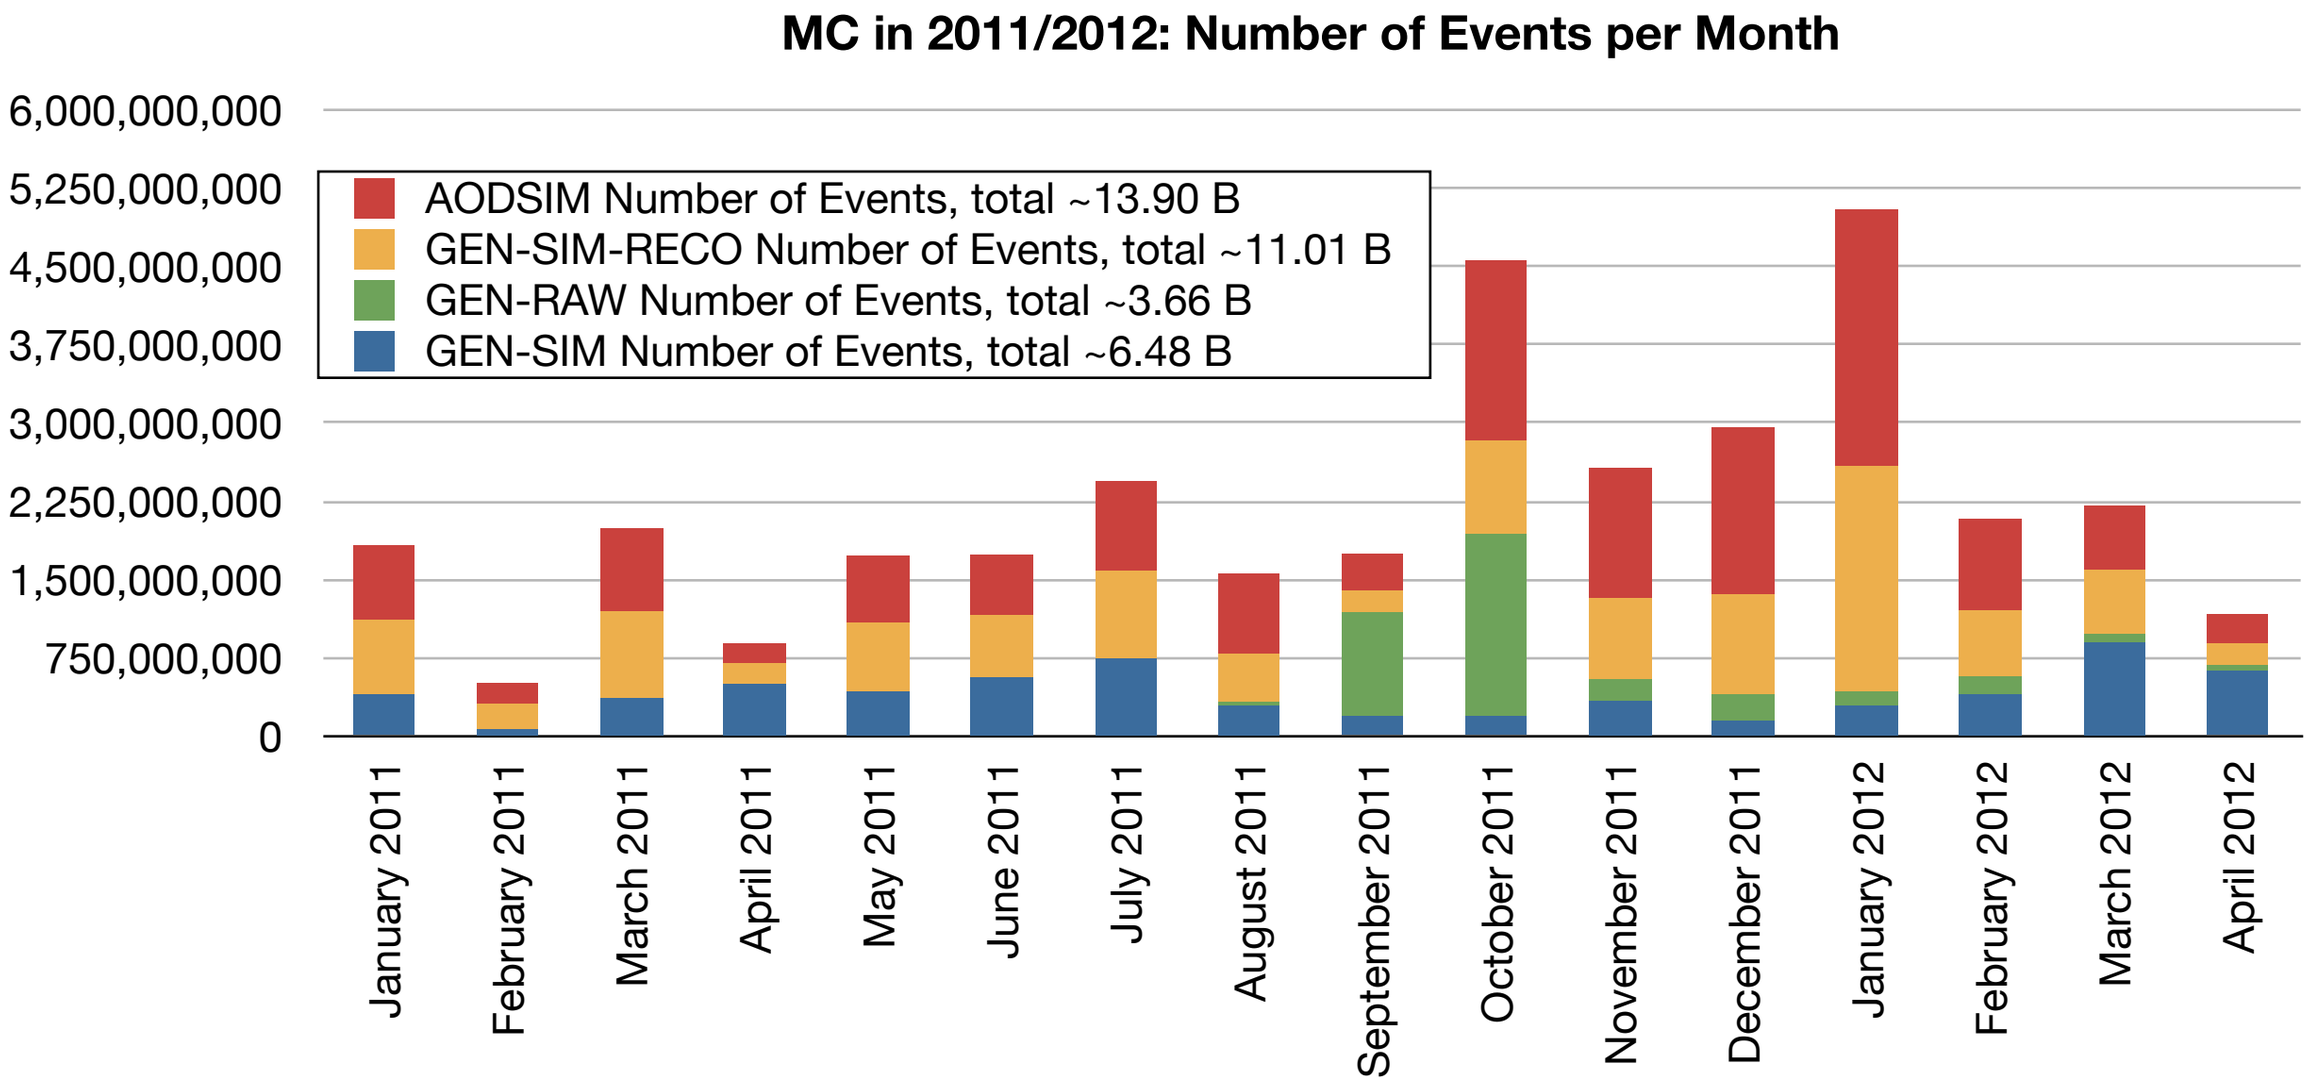

- Produced 4.25 Billion events in 2011, 2.22 Billion in 2012
- Reconstructed 2.5 Billion events twice with different PileUp scenarios plus CMSSW_4_4_X re-reconstruction pass over Xmas
- 2012: no major MC re-reconstruction pass planned (yet, used PileUp model expected to be sufficient for all of 2012)

MC: 2012 8 TeV MC

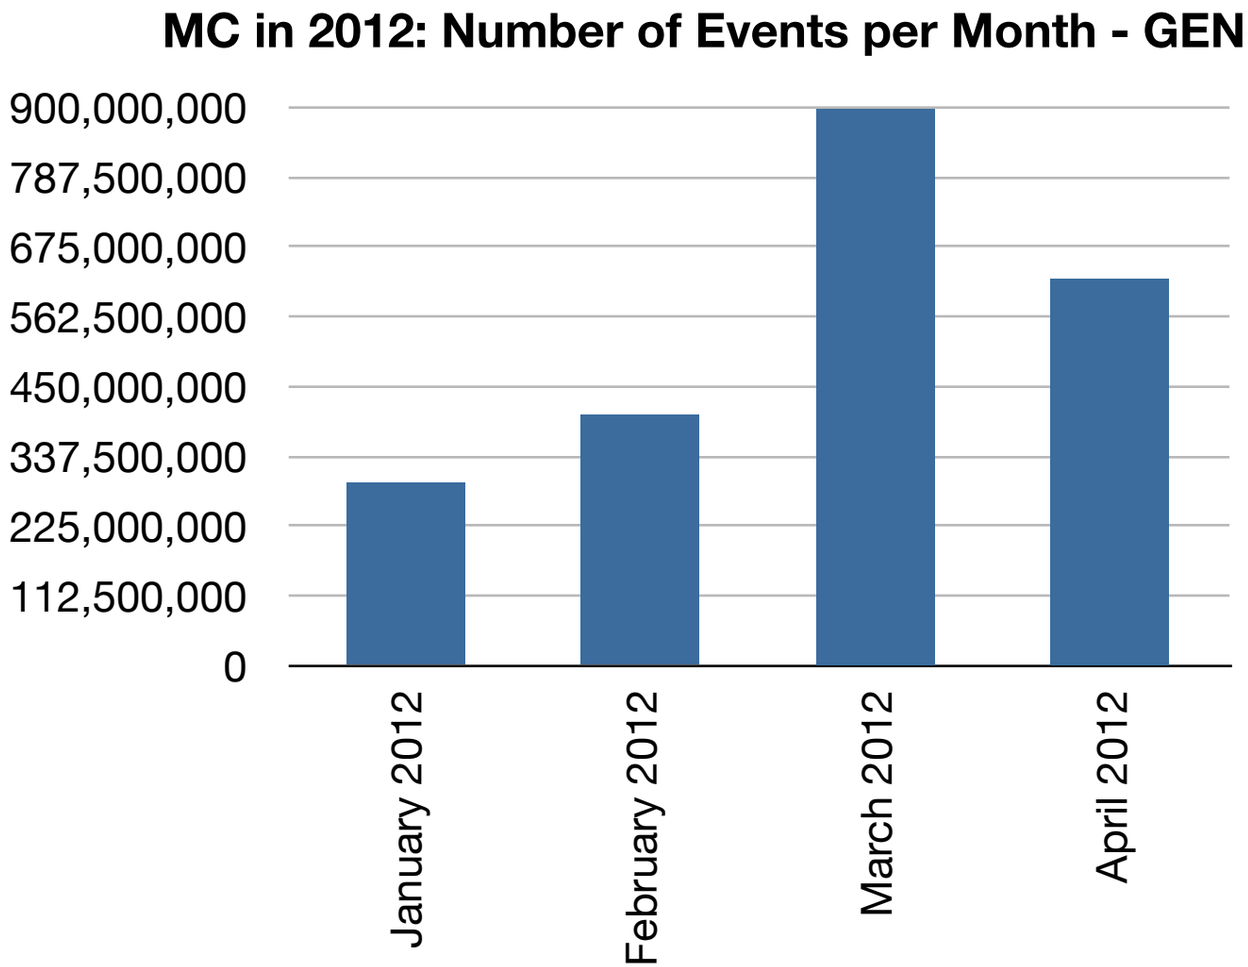


```
Summer12 GEN-SIM
---------------------------------
(DBS query string: *Summer12*GEN-SIM)
---------------------------------

PRODUCTION   :          85,686,955
VALID        :       1,614,553,607
---------------------------------
TOTAL        :       1,700,240,562
```

```
Summer12 DR52X AODSIM
---------------------------------
(DBS query string: *Summer12-
PU_S7_START52_V*AODSIM)
---------------------------------

PRODUCTION   :         188,899,455
VALID        :         413,159,385
---------------------------------
TOTAL        :         602,058,840
```

- Summer12 MC production in full swing
  - Digi+Reco started 2 weeks ago, already over 400 Million events available for analysis

From: 5/15/2012

Resource overview last week

PRODUCTION

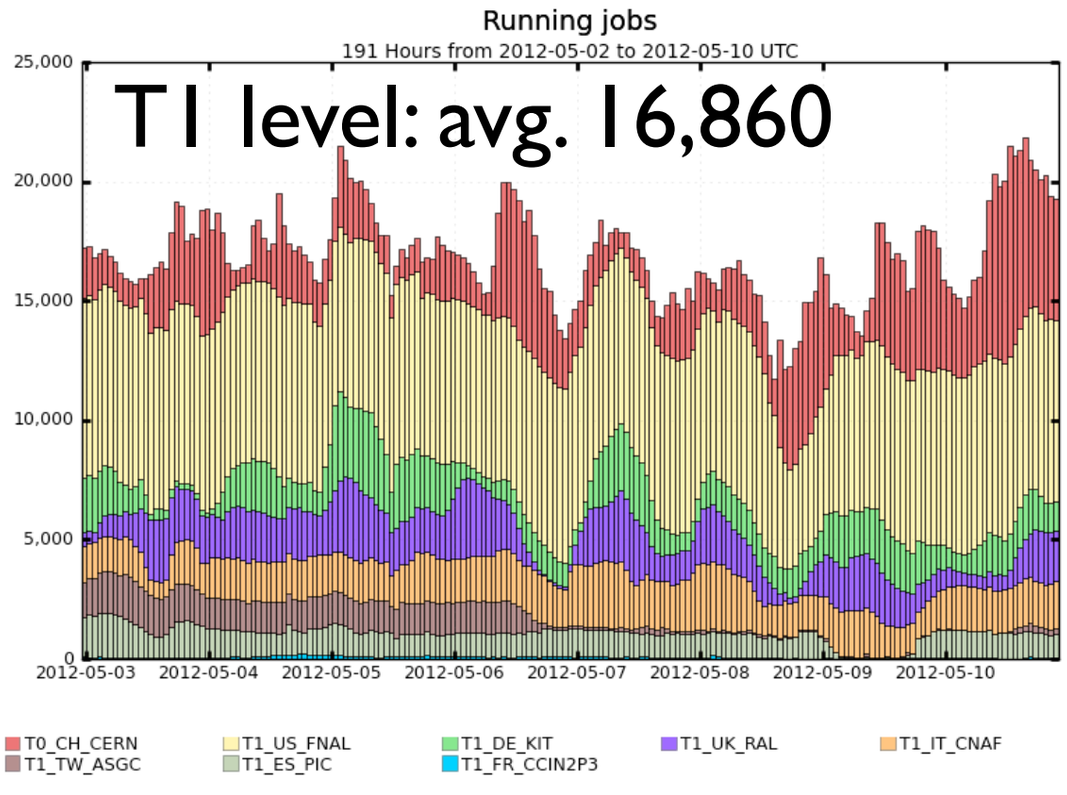

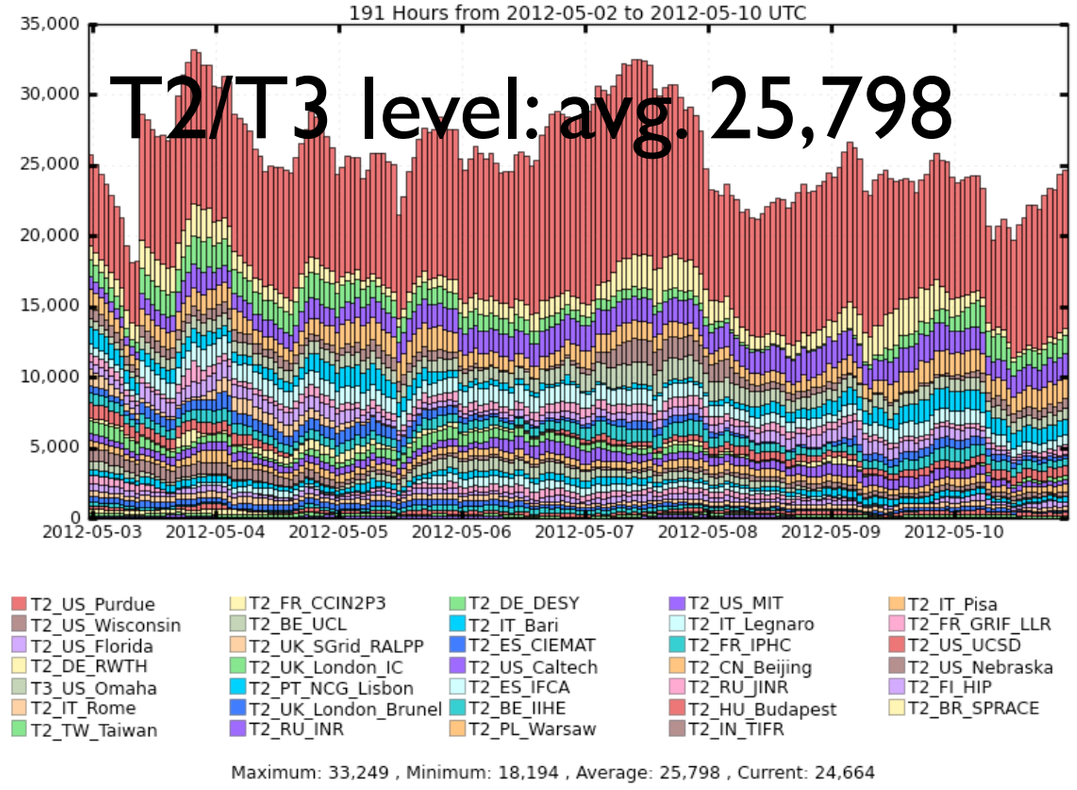

ANALYSIS

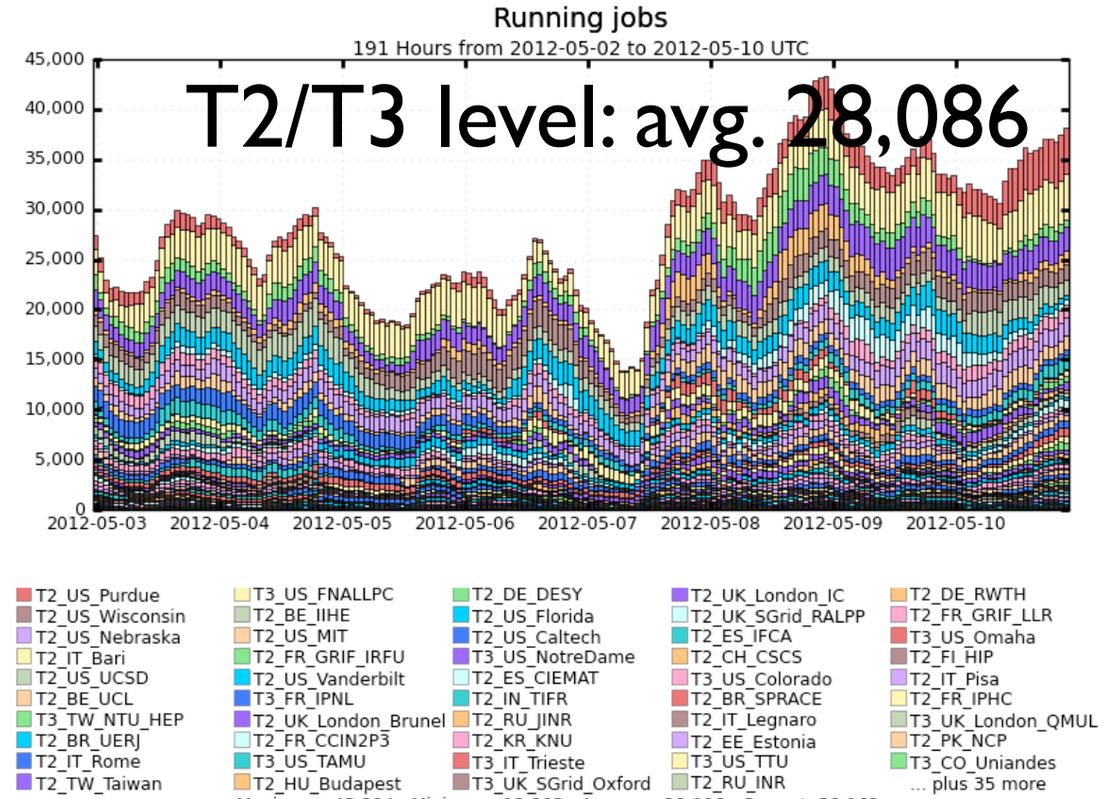

- Example week, using available resources
  - All central workflows run through glideIn WMS
  - Analysis mixture of gLite and glideIns

Storage overview

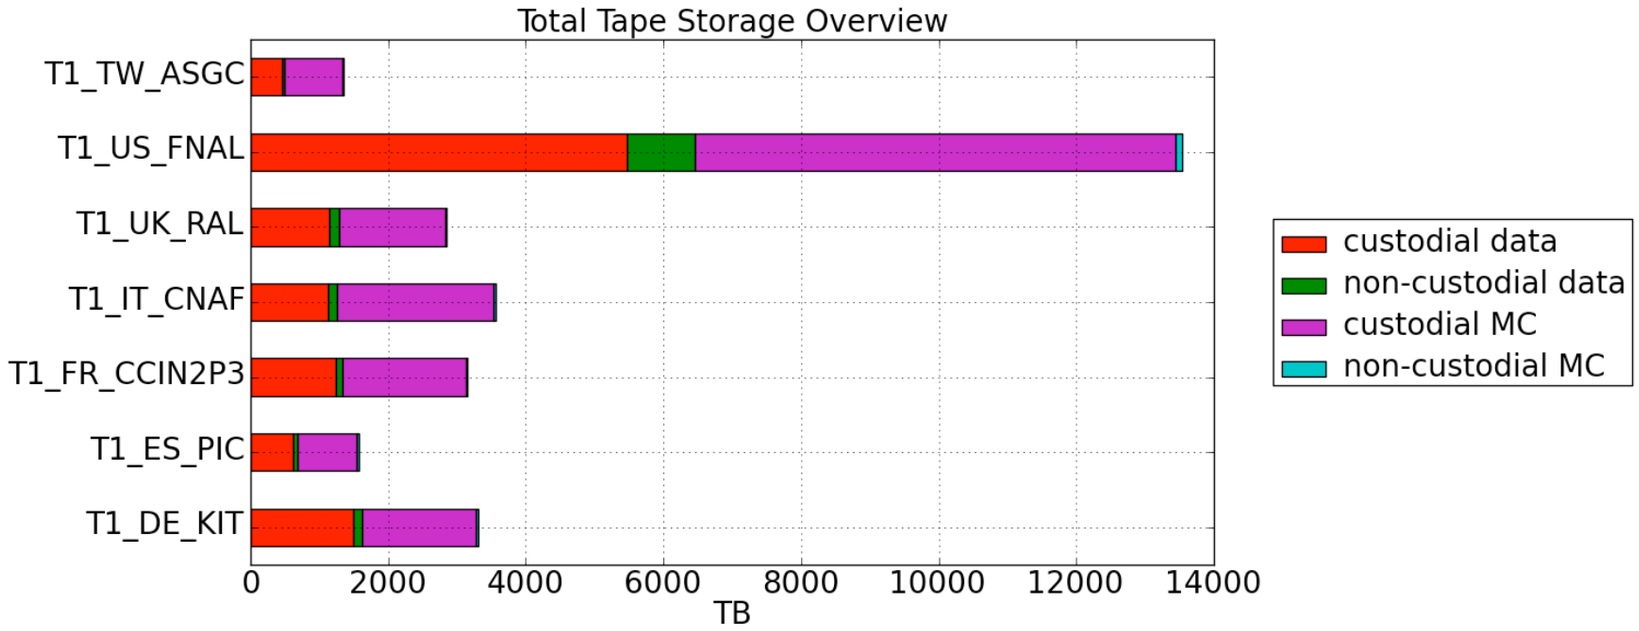

May 13, 2012	T1_DE_KIT [TB]	T1_ES_PIC [TB]	T1_FR_CCIN2P3 [TB]	T1_IT_CNAF [TB]	T1_UK_RAL [TB]	T1_US_FNAL [TB]	T1_TW_ASGC [TB]	All sites [TB]
custodial data	1,503	624	1,237	1,130	1,141	5,481	465	11,580
non-custodial data	116	59	106	128	145	986	33	1,574
custodial mc	1,664	864	1,790	2,278	1,545	6,968	836	15,946
non-custodial mc	30	21	11	22	12	96	21	212
Total	**3,313**	**1,568**	**3,143**	**3,558**	**2,843**	**13,531**	**1,355**	**29,311**
Pledges	5,100	2,601	3,600	6,600	4,080	22,000	2,550	46,531

- Storage pledges for 2012/2013: 46 PB
  - Currently using 30 PB
  - Next deletion coming: 5.2 PB old MC
  - Need to stay vigilant and delete deprecated samples as soon as possible (tape recycling takes time)

Site performance Tier-2 level

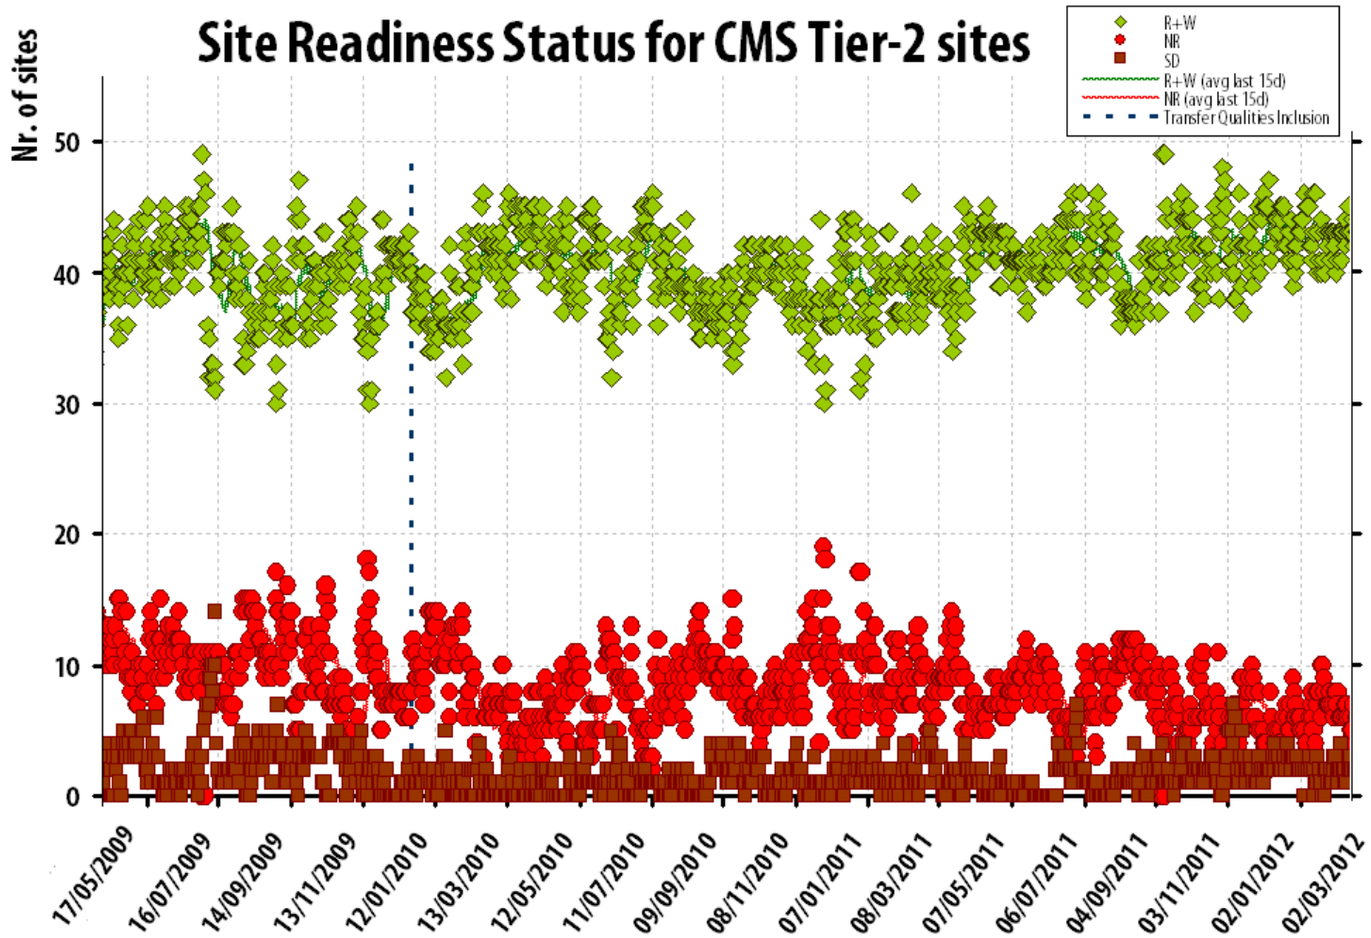

- (Tier-1 level working well, not mentioned here)
- Tier-2 level also performs a good job, supporting a lot of users and analyses
- Many thanks to all T2 admins, also for support outside their core expertise to support the community

T2 site support

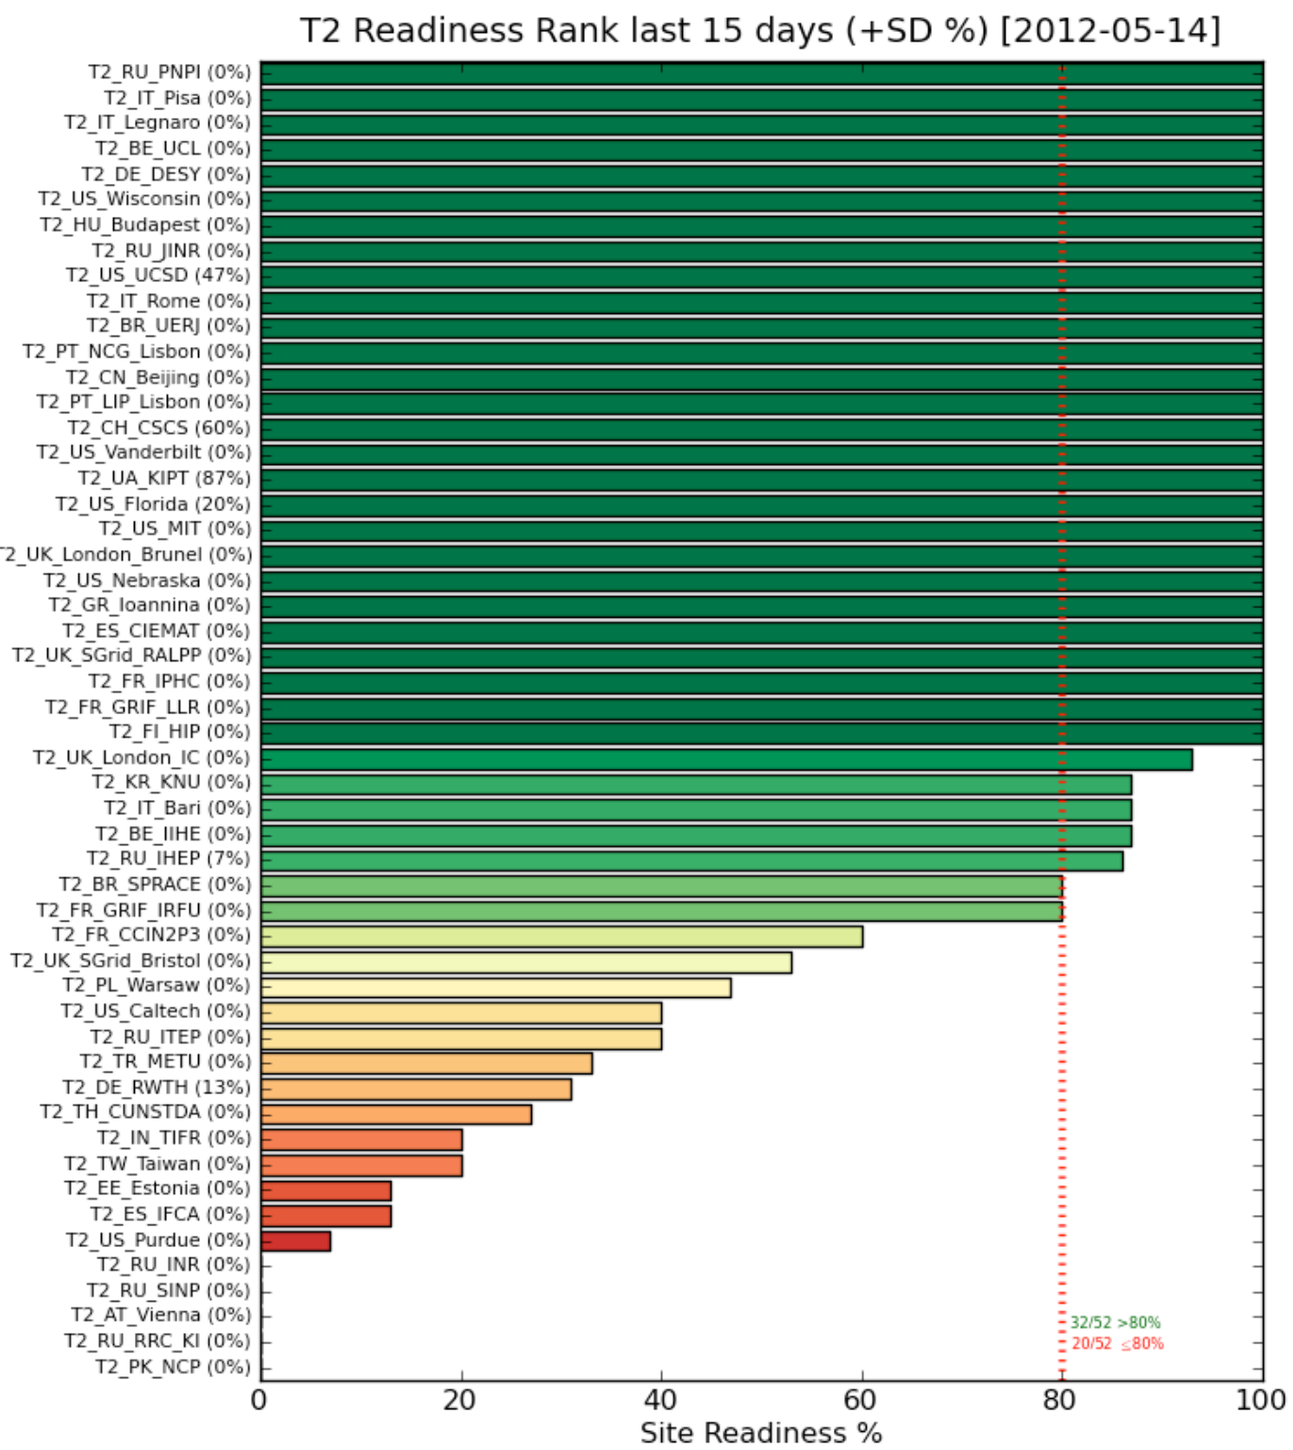

- T2 readiness ranking helps determining the overall readiness of the CMS T2 sites for MC production and analysis
  - Required is 80% readiness in the last 15 days (more details later in Andrea's talk)
- Recently saw degradation of number of ready sites
  - Our computing shifters alarm the sites via savannah tickets that specific SAM tests are failing
  - Central functions alarm sites about other problems like data inconsistencies or transfer problems

T2 support

- Goal is to help the sites to stay fully functional
- Key to this are savannah tickets: primary communication line to the T2 sites
- Most of the support is working very well, we have some cases where we can improve.

- Problems we currently see:
  - Tickets are not answered at all or after a long time
    - Expected average response time: **1 business day**, central operations cannot help you if sites don't respond.
    - It is very important that sites acknowledge tickets they receive before they start working on the problem by posting a comment.
  - Tickets are not closed when solved or re-assigned if other party needs to solve problem
    - The solution of a ticket in the responsibility of the current assigned squad. Unclear situations should be handled by the CRC.

T2 support

- We're working on procedures and measures to improve the situation
  - We will collect statistics of how long savannah tickets stay open and how long it takes to receive the first reply from sites
  - We will also more often contact the site and regional representatives and work with the CRB (CMS Resource Board)
- We're still finalizing procedures, when complete, we will start a dialog with the T2 sites to incorporate their comments and concerns.

Summary

- 2011 was a challenging year with a lot of success for CMS
  - Computing had its share of these successes
    - Kept the Tier-0 running under very difficult conditions
    - Produced and processed enormous amounts of data and MC
    - Kept the sites ready and the transfers between then going efficiently, many thanks to all of them.
    - Enabled the collaboration to analyze all the data and MC samples resulting in over 100 publications
- 2012 will be an even bigger challenge
  - Prioritization of resource usage will be the key to success
  - We need all sites work perfectly; we are there to help, but need participation from the sites