CMS operational issues

[**GDB** – CERN, 9 March 2011]

Ian Fisk, Daniele Bonacorsi

on behalf of CMS Computing and Data, Facilities and Analysis Operations teams
(special thanks to Claudio Grandi, Josep Flix, James Letts for providing some specific plots)

Operational items

Not many major issues in CMS Computing Operations since last time we met.

Main focus since last report:

- The completion of the Heavy Ion zero-suppression pass
- Analysis, Analysis, Analysis.
- Preparing for 2011 data taking

CMS Site Readiness: T1s

[Credits: J.Flix from CMS Facilities Ops]

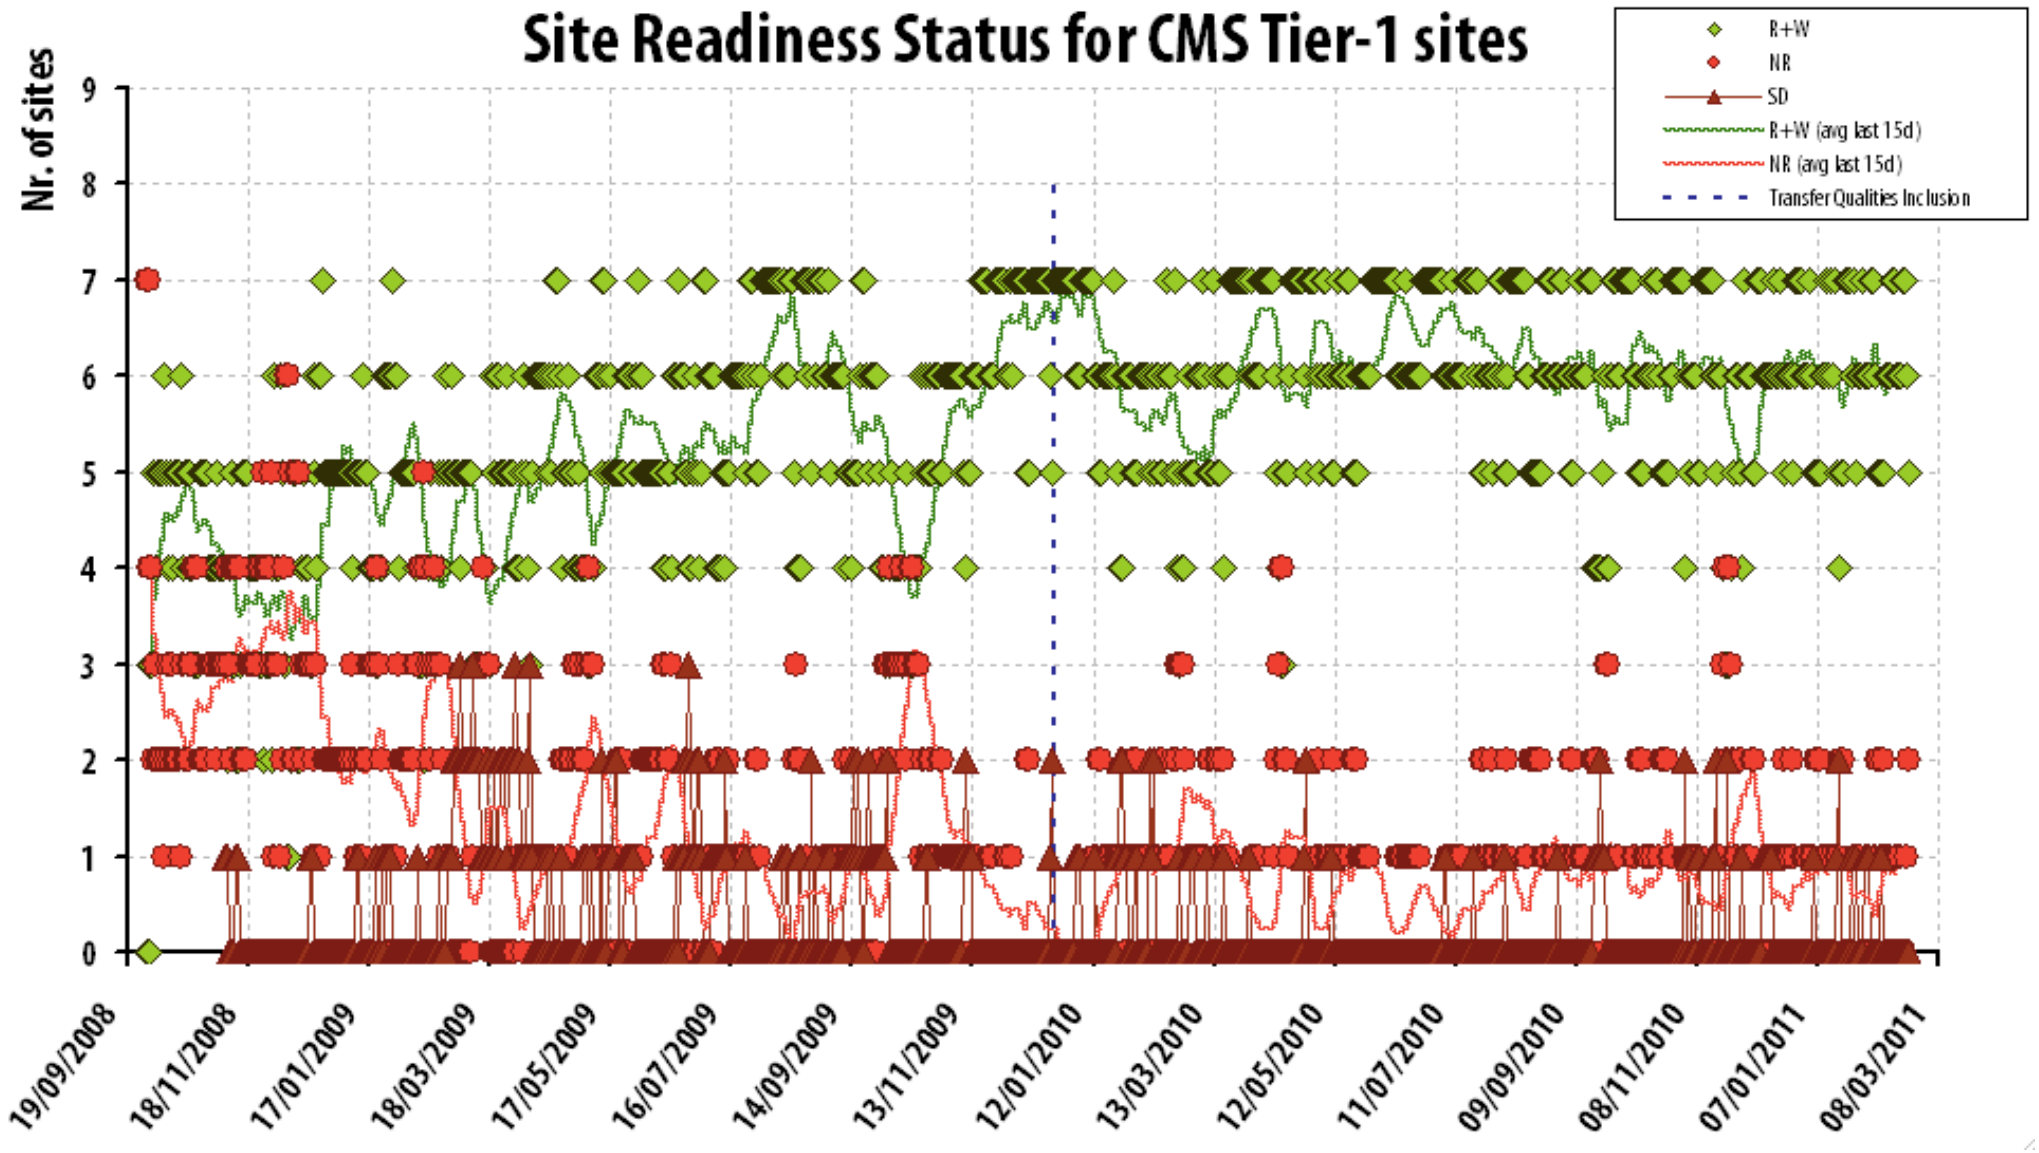

CMS Site Readiness: T2s

[Credits: J.Flix from CMS Facilities Ops]

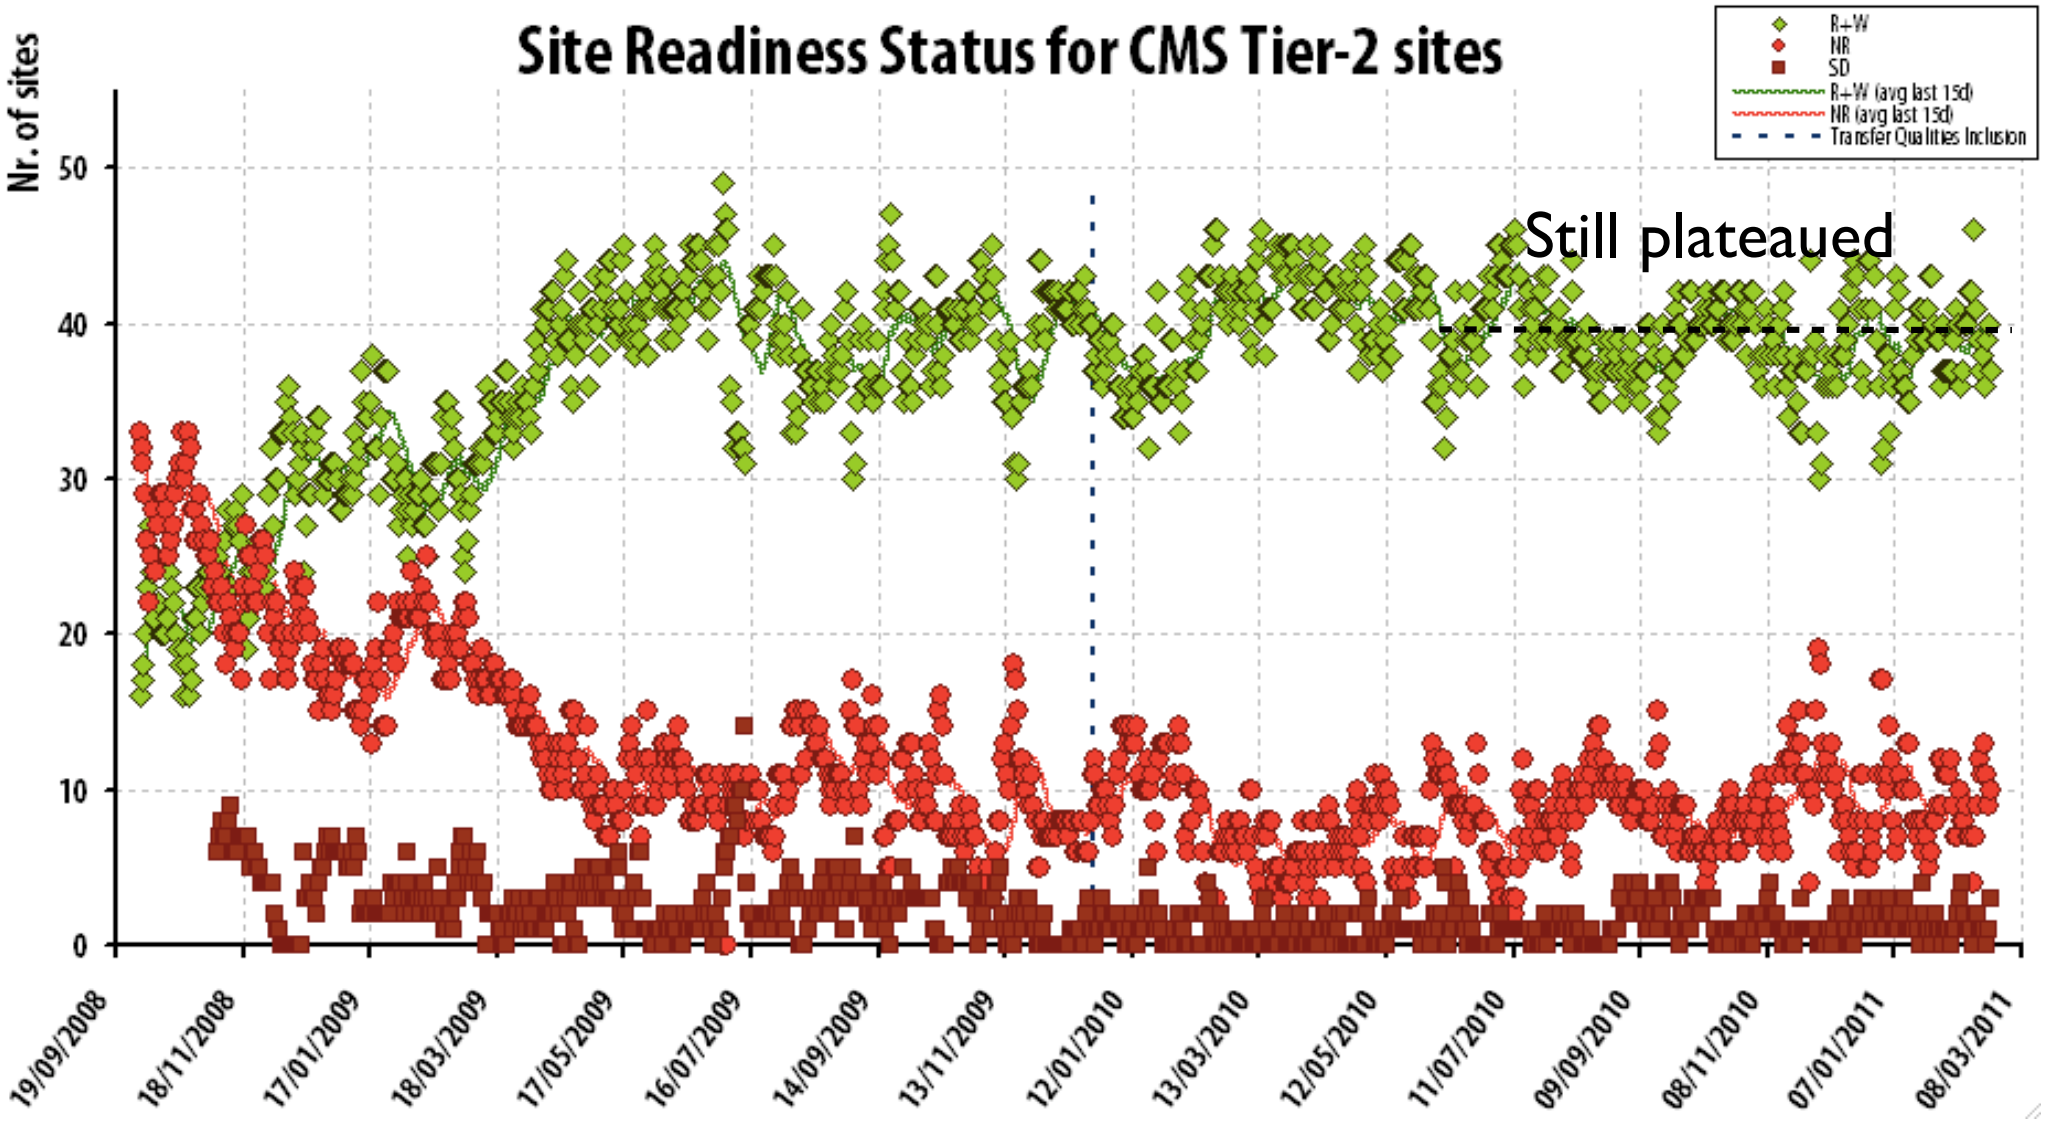

Data transfers

More traffic on non-T0 routes

- The only CERN outbound traffic is the export of HI data to FNAL
  - As high as 600-700 MB/s (daily)
- T1-T2 traffic important
- Consistent T2-T1 traffic
  - MC production upload
- T2-T2 traffic continues
  - And supports analysis needs

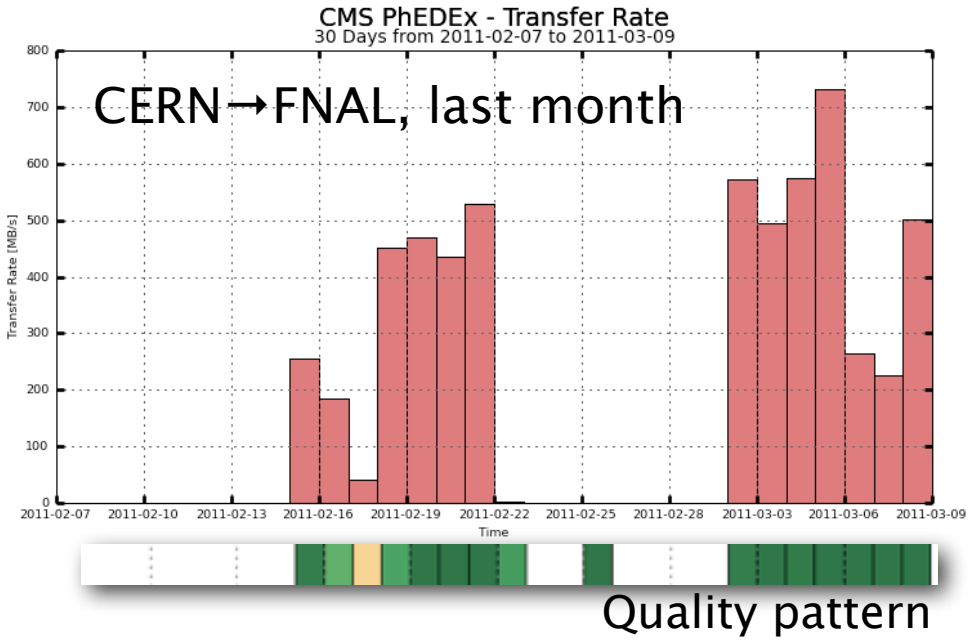

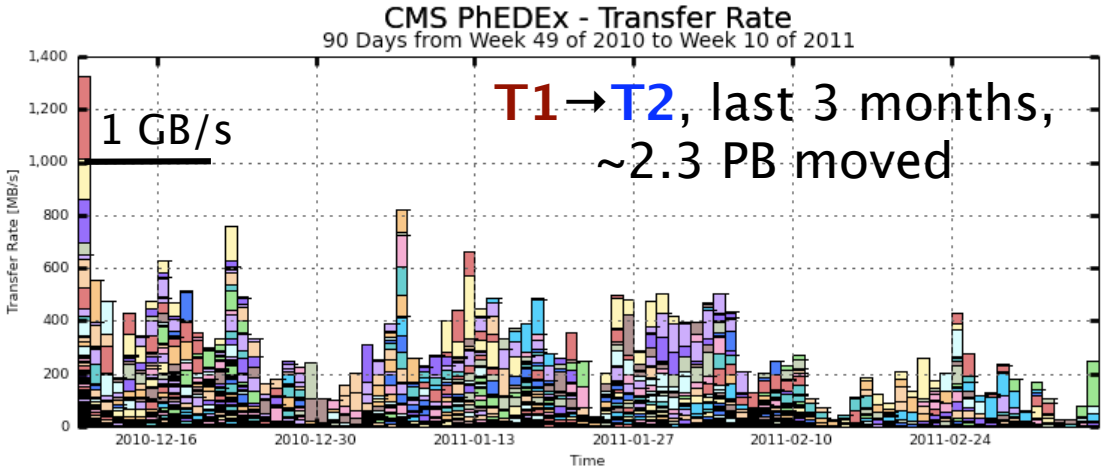

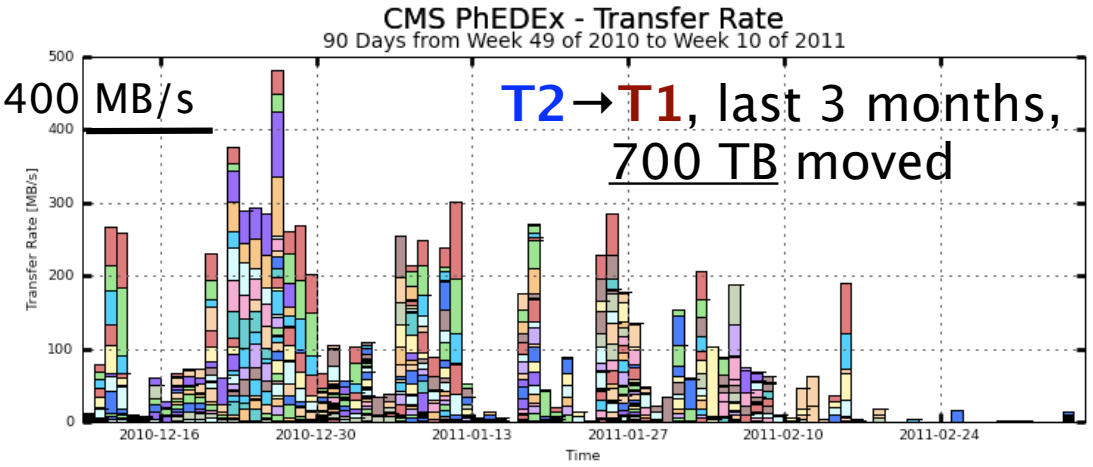

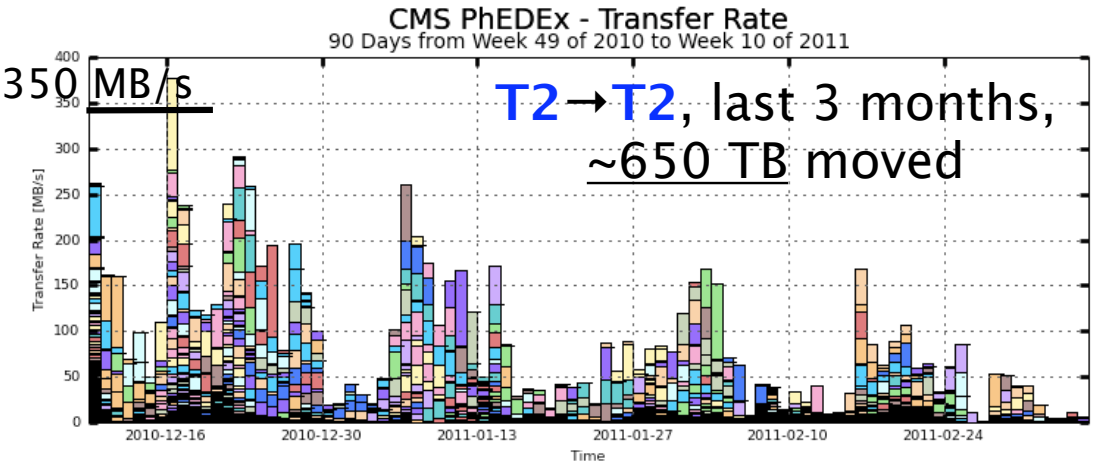

T1 processing

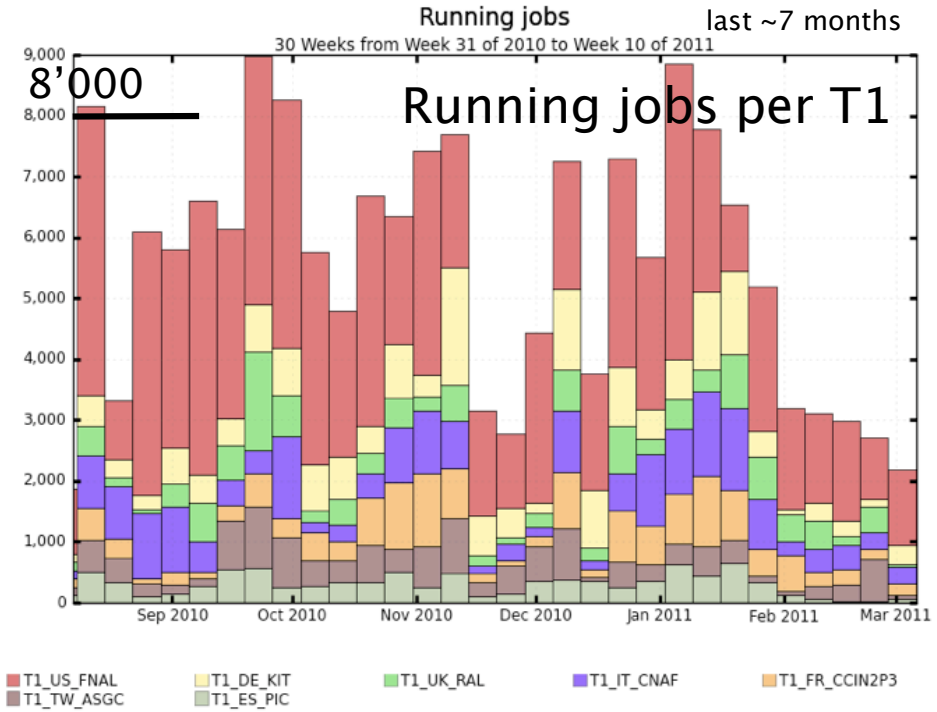

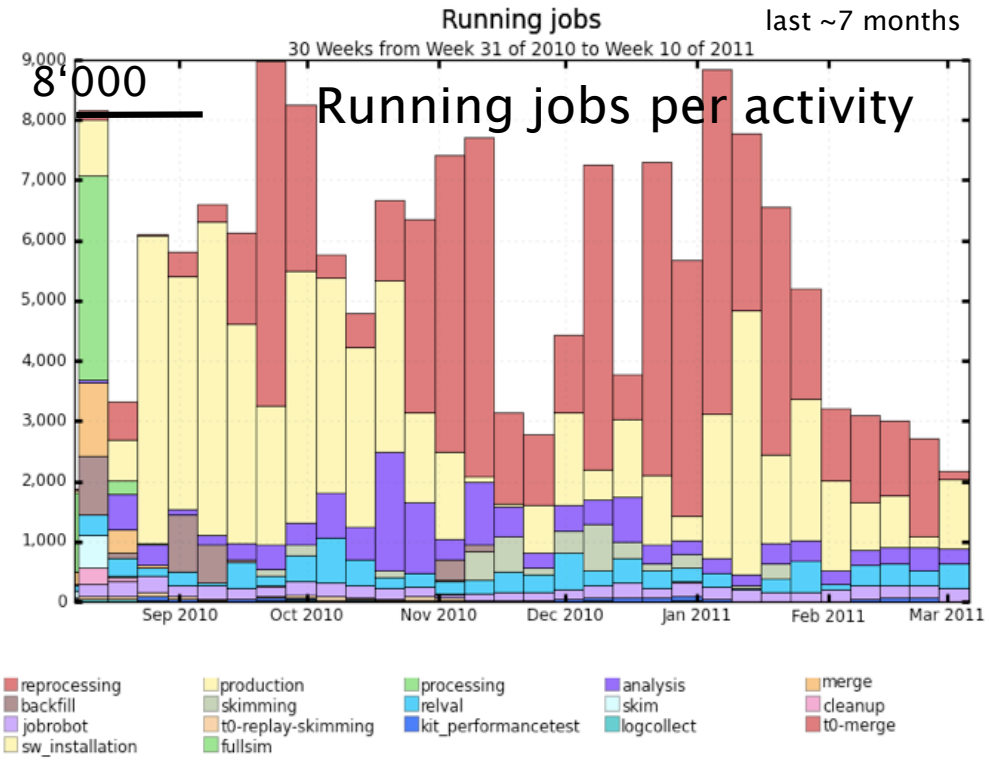

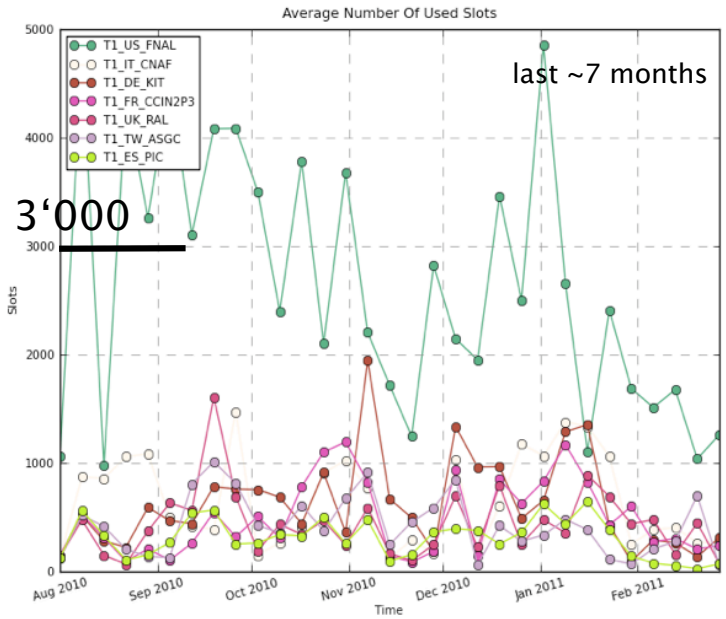

Average nb of used slot per T1

T1 processing share

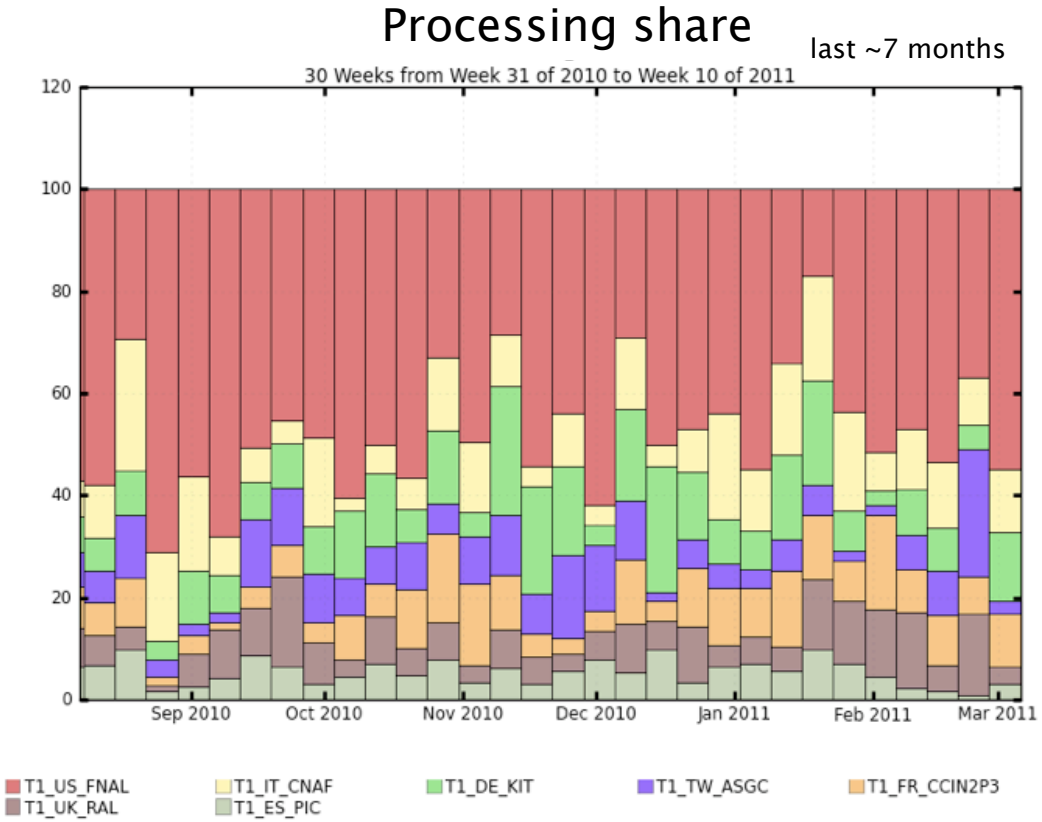

Weekly structures visible

- on average it roughly maps to the T1 pledges

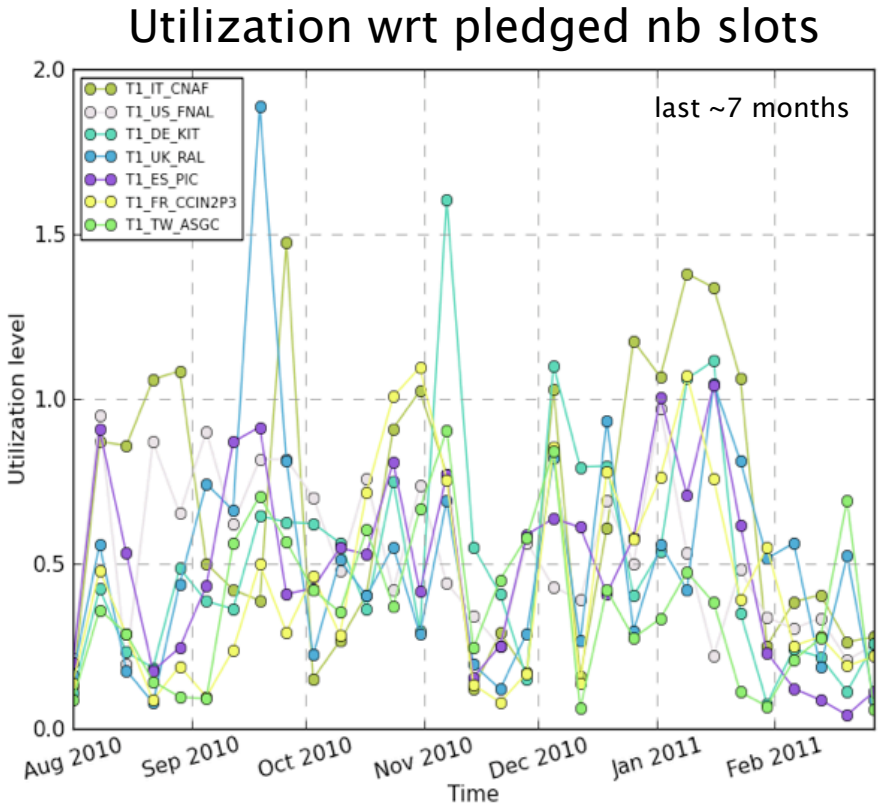

Entering soon in a resource-constraint environment

- Sites filled more, utilization would meet pledges

T1 processing “efficiencies”

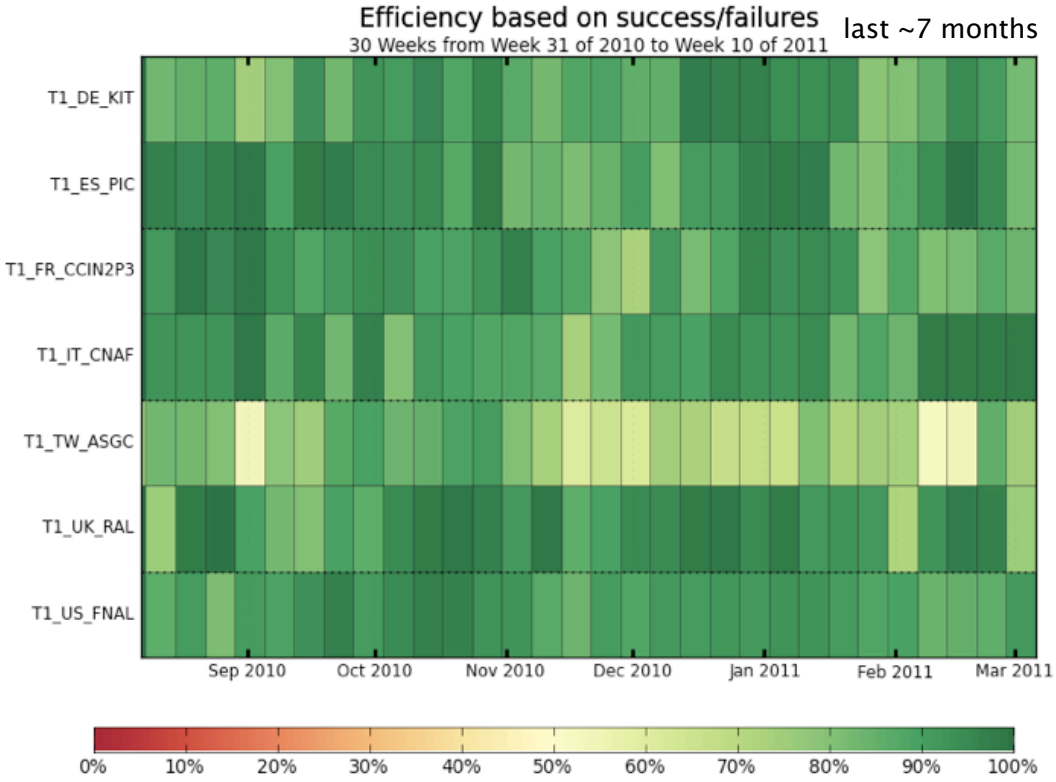

eff = successful jobs / total jobs
(per T1, all activities)

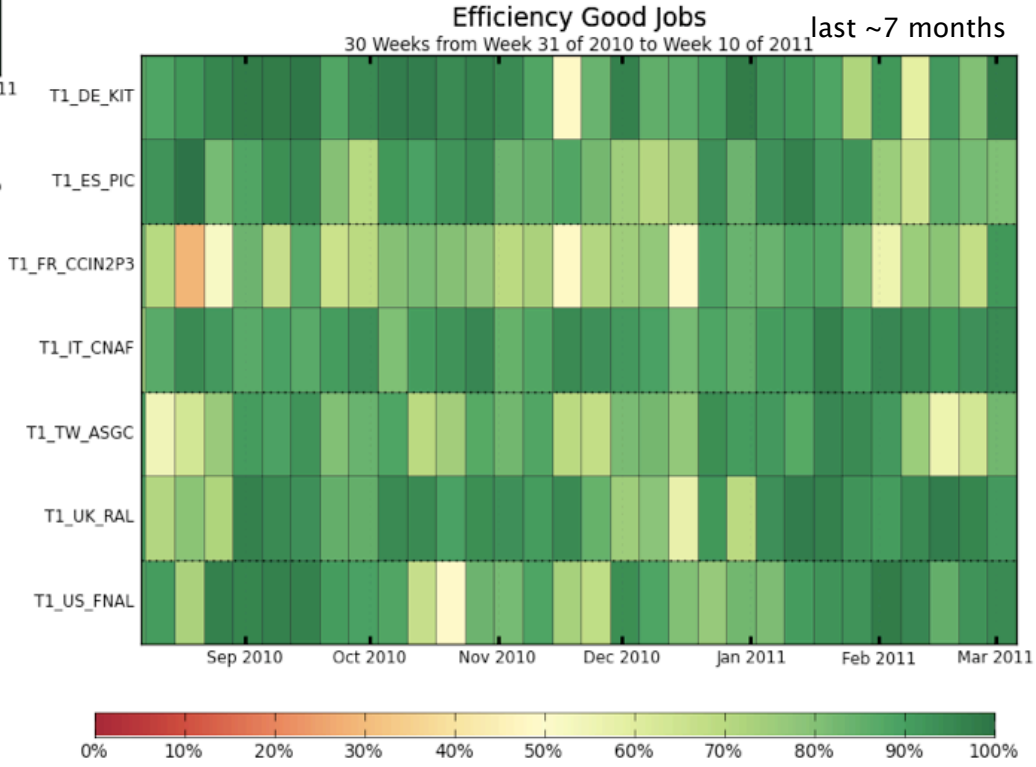

eff = CPU eff (for successful jobs only)
(per T1, all activities)

MC production (T2+T1)

Interesting times ahead...

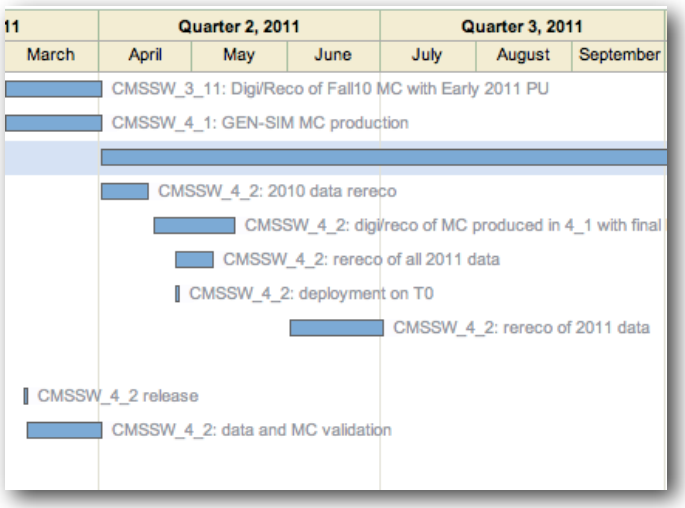

MC prod terminated jobs per site, last 3 months

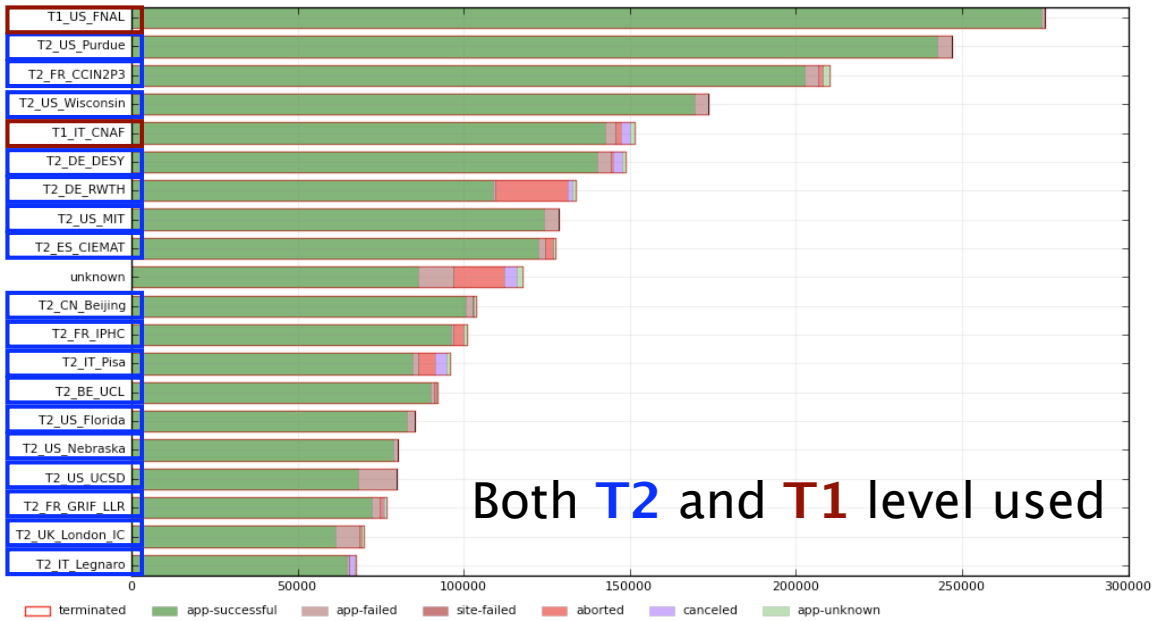

Both T2 and T1 level used

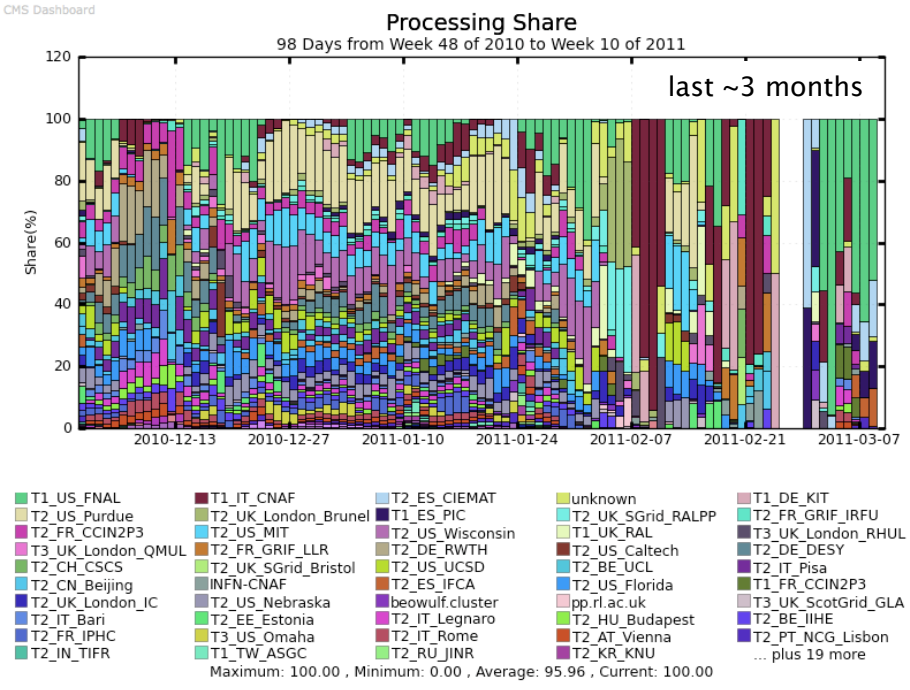

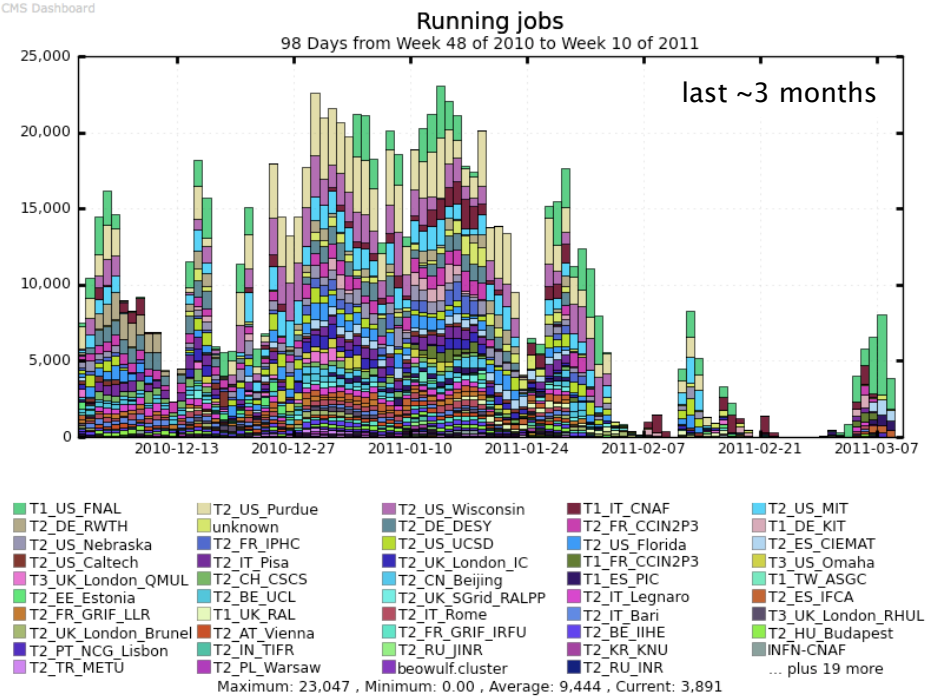

Analysis load goes on… (T2)

Jobs at T2s
(per T2, only analysis)

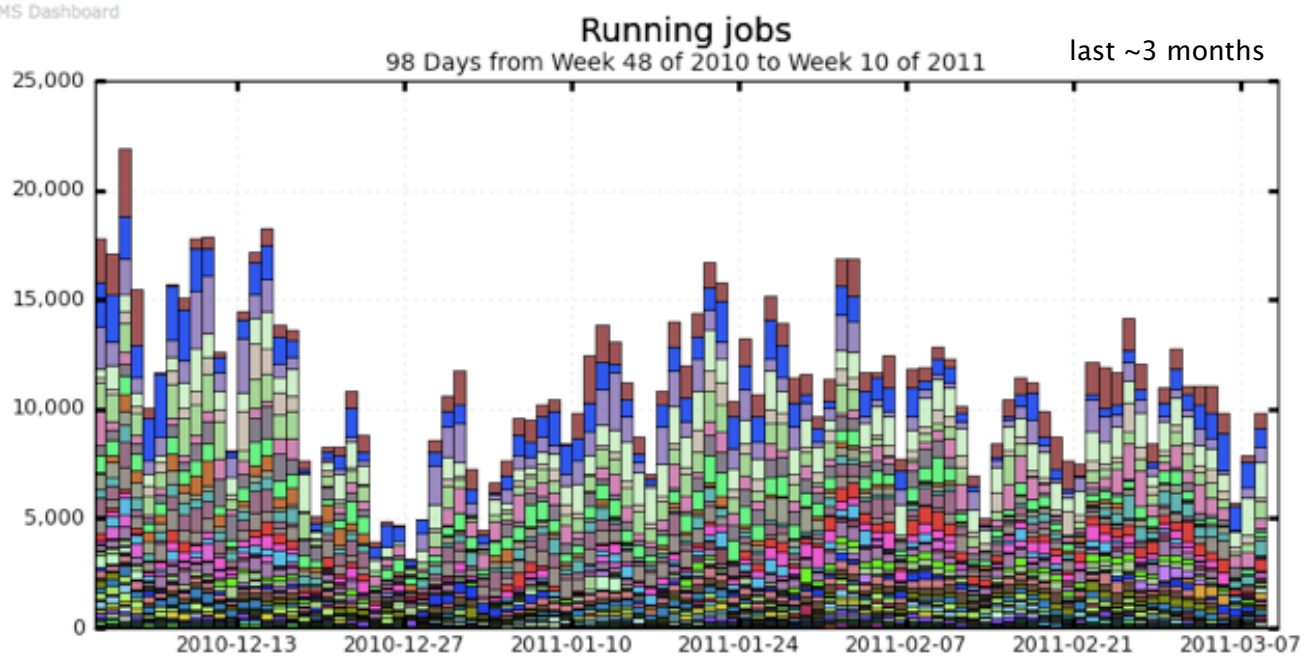

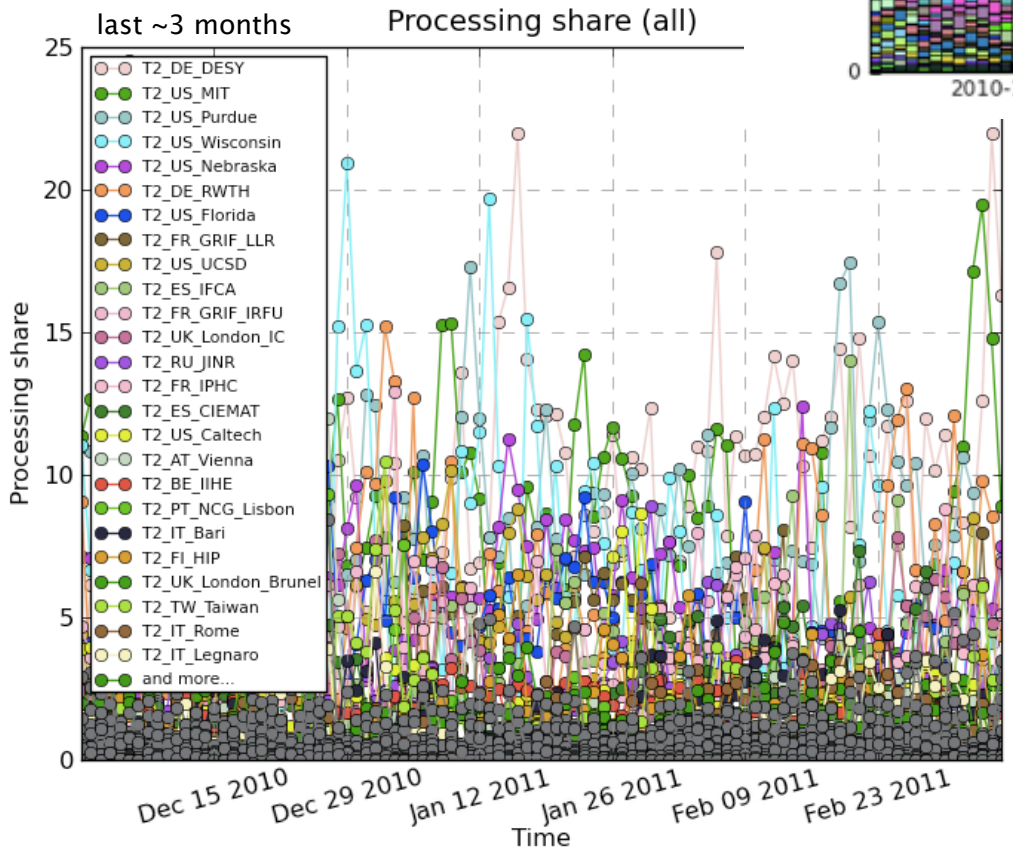

Processing share
(per T2, only analysis)

Analysis: slots, users, CPU eff, application failures

[Credits: J.Letts from CMS Analysis Ops]

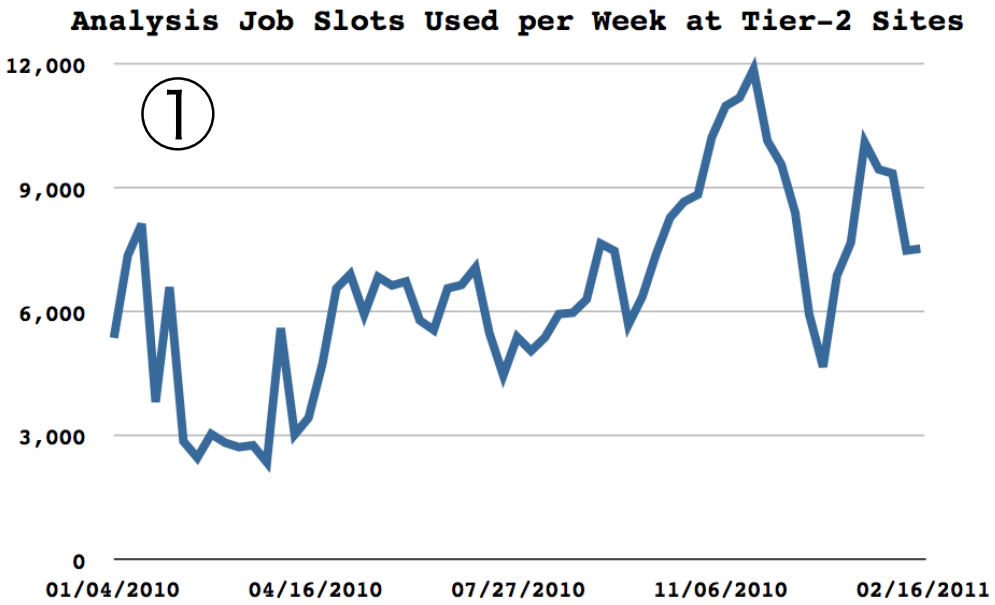

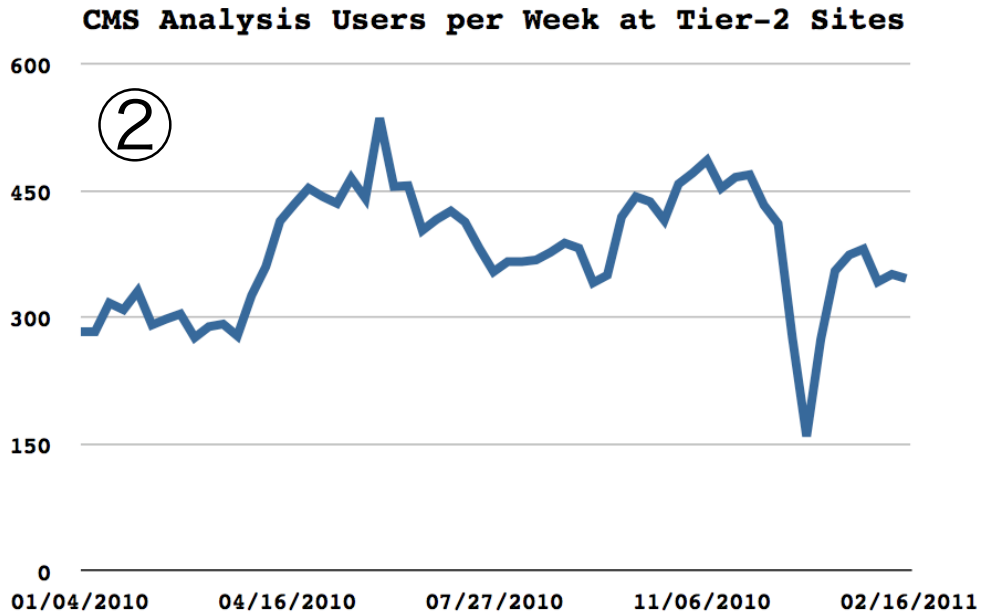

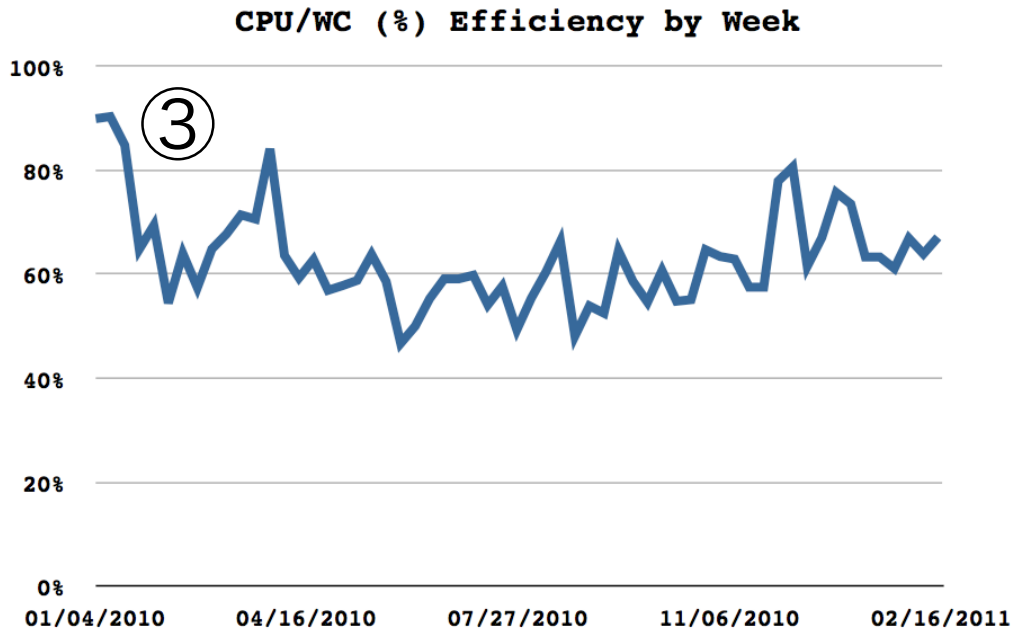

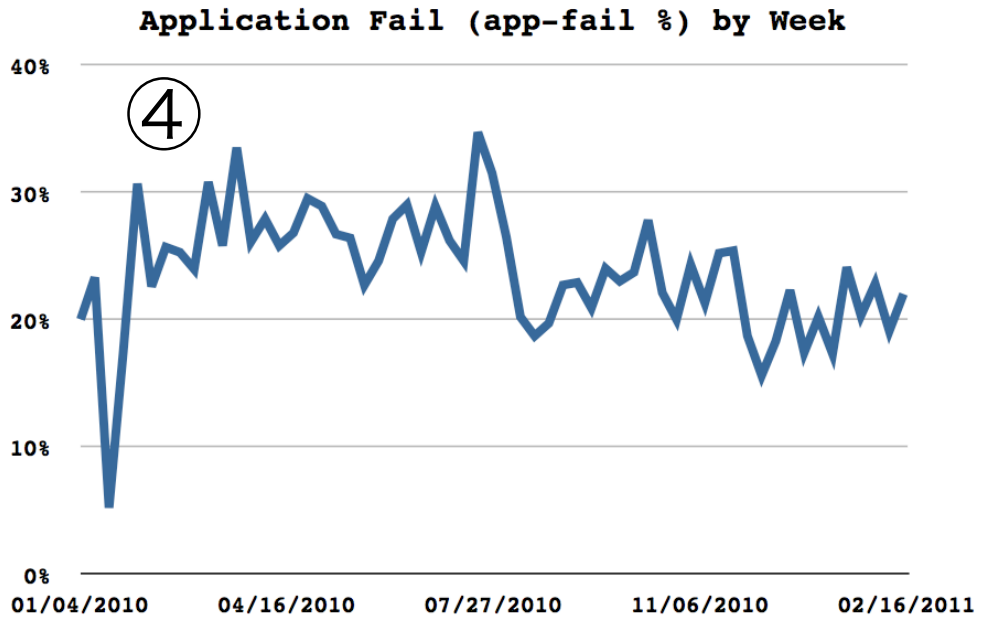

CMS is monitoring the transition to CREAM

[Credits: C.Grandi from CMS Integration]

LCG-CE support ends with the start of LHC data taking

- This is basically now!
- SLC4 support ended on 1/1/11

Production SAM Tests of CREAM Availability

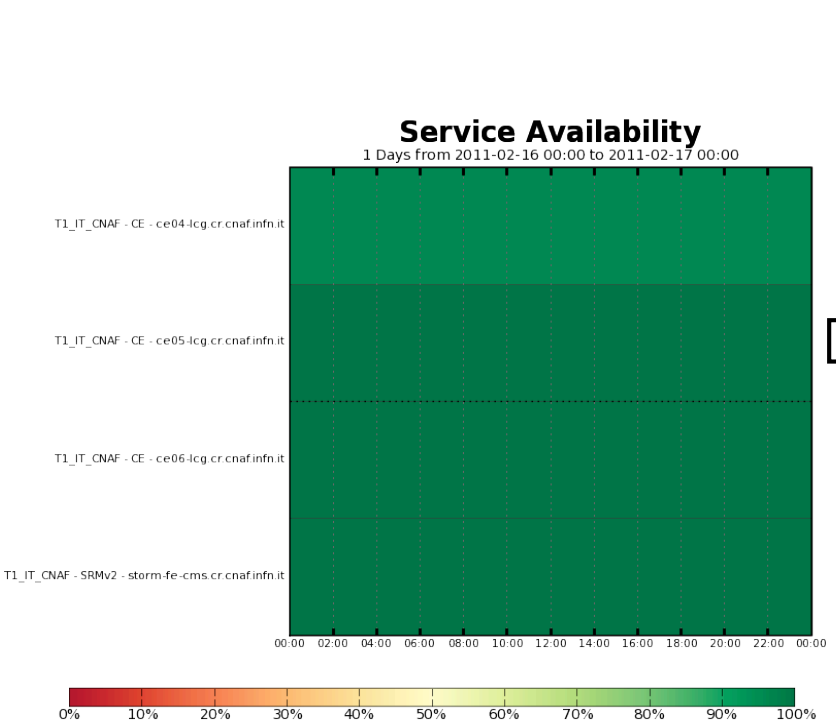

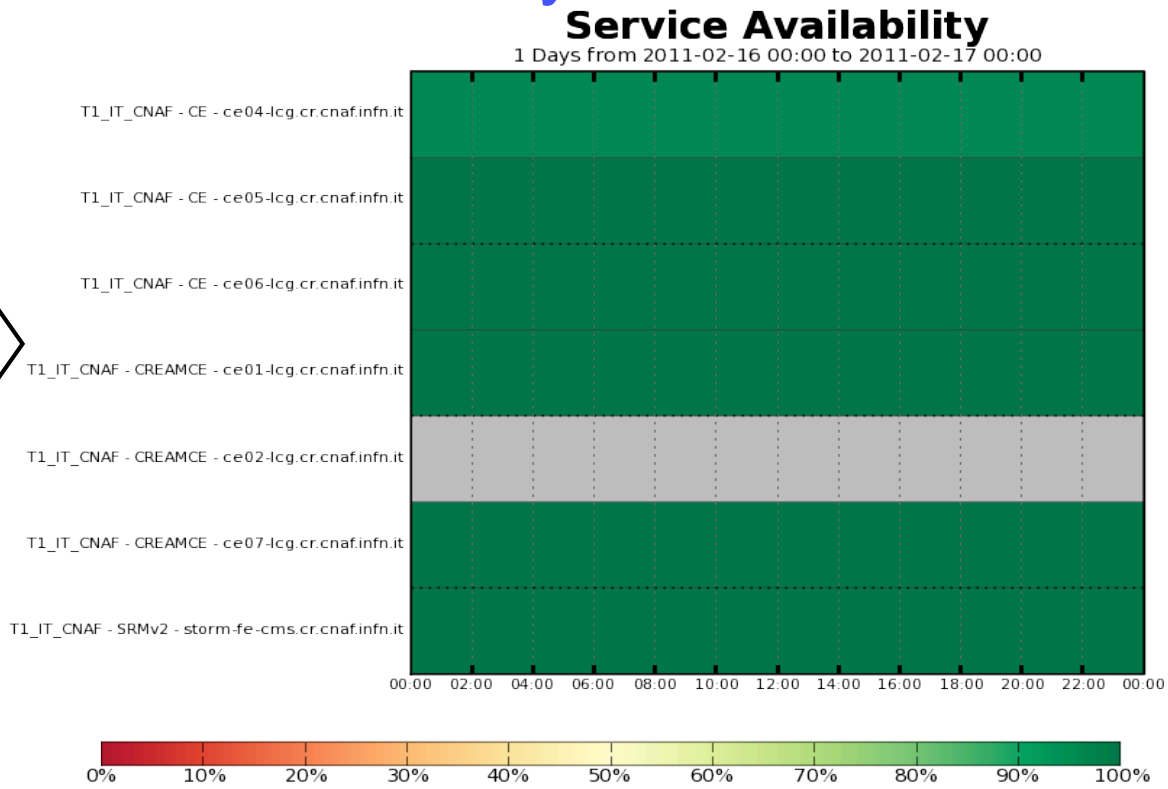

Test of /cms/Role=Pilot

[Credits: C.Grandi from CMS Integration]

17/2/2011

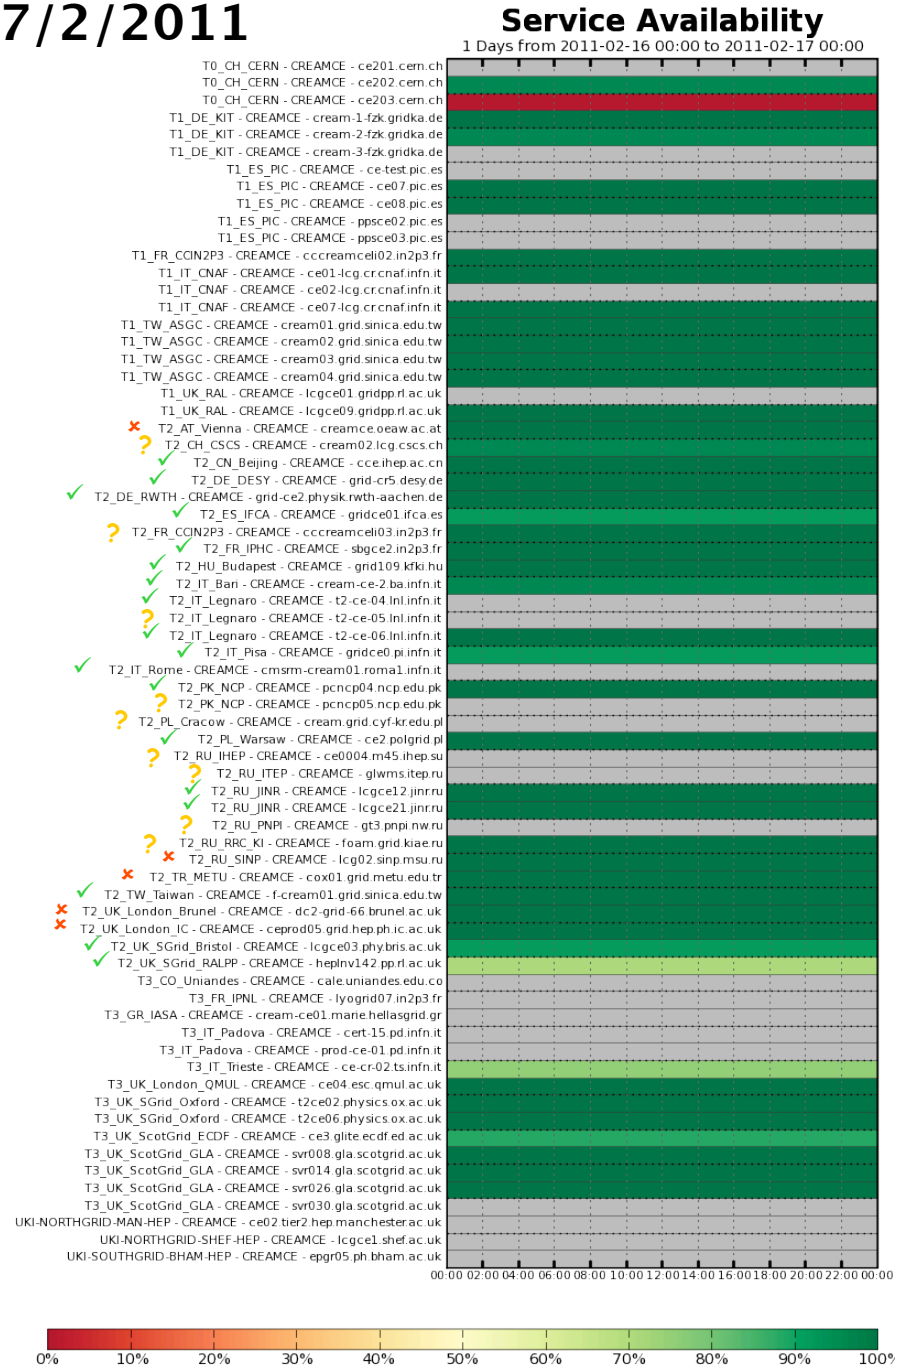

Tier-0 – Not tested

Tier-1 – Tested a few months ago

(tests done with /cms/Role=t1access)

Tier-2 – See report (next slide)

(tests done with /cms/Role=pilot)

Tier-3 – Not tested

Get latest tests from devel. version of the dashboard at http://tinyurl.com/67momta

(tests done with /cms/Role=lcgadmin)

Tier-2 Status on CREAM-CE as of February 17/2/2011

[Credits: C.Grandi from CMS Integration]

Tested successfully:

T2_BE_IIHE cream01.iihe.ac.be *(*)*
T2_CN_Beijing cce.ihep.ac.cn
T2_DE_DESY grid-cr5.desy.de
T2_DE_RWTH grid-ce2.physik.rwth-aachen.de
T2_ES_IFCA gridce01.ifca.es
T2_FR_GRIF_IRFU node74.datagrid.cea.fr *(*)*
T2_FR_GRIF_LLR llrcream.in2p3.fr *(*)*
T2_FR_IPHC sbgce2.in2p3.fr
T2_HU_Budapest grid109.kfki.hu
T2_IT_Bari cream-ce-2.ba.infn.it
T2_IT_Legnaro t2-ce-06.lnl.infn.it
T2_IT_Pisa gridce0.pi.infn.it
T2_IT_Rome cmsrm-cream01.roma1.infn.it
T2_KR_KNU cluster50.knu.ac.kr *(*)*
T2_PK_NCP pcncp04.ncp.edu.pk
T2_PL_Warsaw ce2.polgrid.pl
T2_RU_JINR lcgce12.jinr.ru
T2_RU_JINR lcgce21.jinr.ru
T2_TW_Taiwan f-cream01.grid.sinica.edu.tw
T2_UK_SGrid_Bristol lcgce03.phy.bris.ac.uk
T2_UK_SGrid_RALPP heplnv142.pp.rl.ac.uk

Tested with failure:

T2_AT_Vienna creamce.oeaw.ac.at
T2_AT_Vienna hephygr.oeaw.ac.at *(*)*
T2_RU_SINP lcg02.sinp.msu.ru
T2_TR_METU cox01.grid.metu.edu.tr
T2_UK_London_Brunel dc2-grid-66.brunel.ac.uk
T2_UK_London_IC ceprod05.grid.hep.ph.ic.ac.uk
T2_UK_London_IC ceprod06.grid.hep.ph.ic.ac.uk *(*)*

(*) In SideDB but not in dashboard

Not in SiteDB (not tested):

T2_CH_CSCS cream02.lcg.cscs.ch
T2_FR_CCIN2P3 cccreamceli03.in2p3.fr
(T2_IT_Legnaro t2-ce-05.lnl.infn.it)
(T2_PK_NCP pcncp05.ncp.edu.pk)
T2_PL_Cracow cream.grid.cyf.kr.edu.pl
T2_RU_IHEP ce0004.m45.ihep.su
T2_RU_ITEP glwms.itep.ru
T2_RU_PNPI gt3.pnpi.nw.ru
T2_RU_RRC_KI foam.grid.kiea.ru

Sites with no CREAM CE (neither in SiteDB nor on dashboard):

T2_BE_UCL
T2_BR_SPRACE
T2_BR_UERJ
T2_EE_Estonia
T2_ES_CIEMAT
T2_FI_HIP
T2_IN_TIFR
T2_PT_LIP_Lisbon
T2_PT_NCG_Lisbon
T2_RU_INR
T2_UA_KIPT
T2_US_Caltech
T2_US_Florida
T2_US_MIT
T2_US_Nebraska
T2_US_Purdue
T2_US_UCSD
T2_US_Vanderbilt
T2_US_Wisconsin

NOTE: the situation is **dynamic.**

Don't extrapolate these Feb tests as 100% valid today
(work continued on glexec, but a full new round of these specific tests has not been done again since this one)

glexec on the WN and ARGUS

For security reasons pilot jobs (glidein's) need to switch identity to the user who submitted the payload

- This is done by invoking **`glexec`** on the WN

Requires an authorization server to be running at the site

- On OSG: **GUMS** is in production at all sites
- On EGI: **ARGUS** is recommended
  - Also allows global banning of users
  - CREAM 1.7 (EMI 1 release on 1/4/2011) will also be able to use ARGUS for authorization
  - The ARGUS policies need to explicitly support all roles which users can be mapped *to* (not from, i.e. not only the pilot roles!)

Initial tests indicate we have a long way to go for full deployment.

One operational highlight: HI-ZS

Heavy Ion software Zero Suppression is an interesting stress case for Castor

- events (size: 12 MB/evt) are read and reconstructed, then new RAW and RECO events are written
- RAW is reduced by a factor of ~4
- Castor performed very well

CPU resources:

- CPU used were a combination of Tier-0 and analysis capacity
- Flexible use of LSF

CERN-outbound transfers:

- See slide 5, bottom left

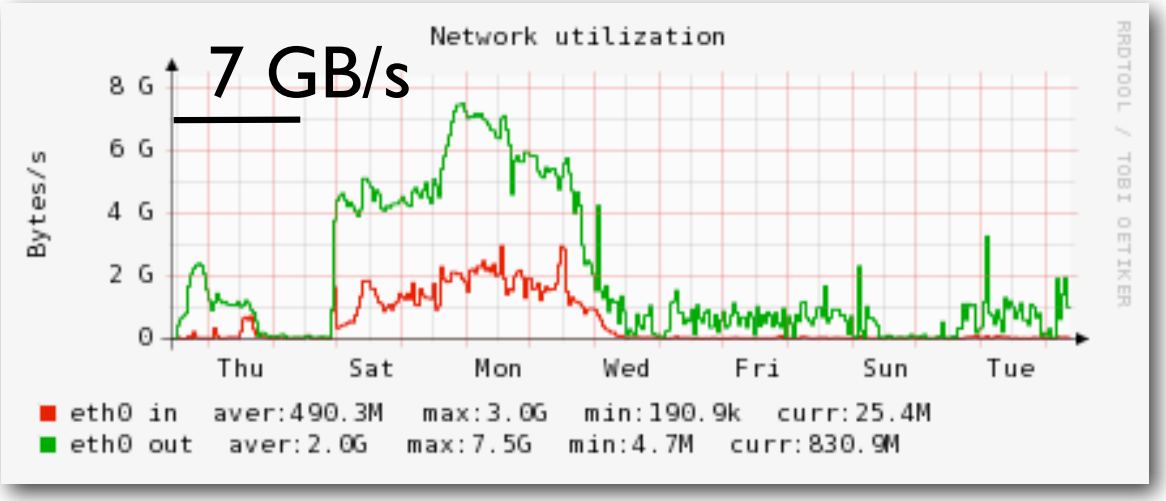

Conclusions

CMS Computing Operations OK since last time we met.

- No major operational issues
- Constant work to sustain the needs, and to support physics analyses
- We will need to be careful in an efficient use of resources in 2011
- Plenty of work in progress to prepare for 2011 data taking (and beyond)
  - This includes e,g, work on popularity, computing model evolutions, WAN access, new production frameworks, etc which we do not cover in this operations report

Ready and eagerly looking forward to 2011 pp(+HI) data.