



CMS 2009Q2 Report

WLCG-MB
7.7.09

Outline

- Site readiness
- STEP09
- Further tests, production

Matthias Kasemann



T1 and T2 Site Readiness

The Site readiness is closely monitored for all Tier-1 and most Tier-2 sites:

- Tools finalized early 2009
- Reports and follow-up during weekly Facility Operations meetings
- Additional meetings to focus on Asian and Russian&Turkish sites

Substantial improvement is observed for large number sites.

Sites below 60% in March with big improvements until June:

> 80% ready: BR-UERJ, KR-KNU, US-Caltech, ES-IFCA, AT-Vienna, UK-London-IC, RU-ITEP, UK-Bristol

> 60% ready: IN-TIFR, IT-Rome, RU-JINR, TR-Metu, RU-SINP




Site Readiness web page:

<https://twiki.cern.ch/twiki/bin/view/CMS/PADASiteCommissioning#ScMon>



Site readiness: Tier-1


- T1 site readiness during STEP09 was mixed:



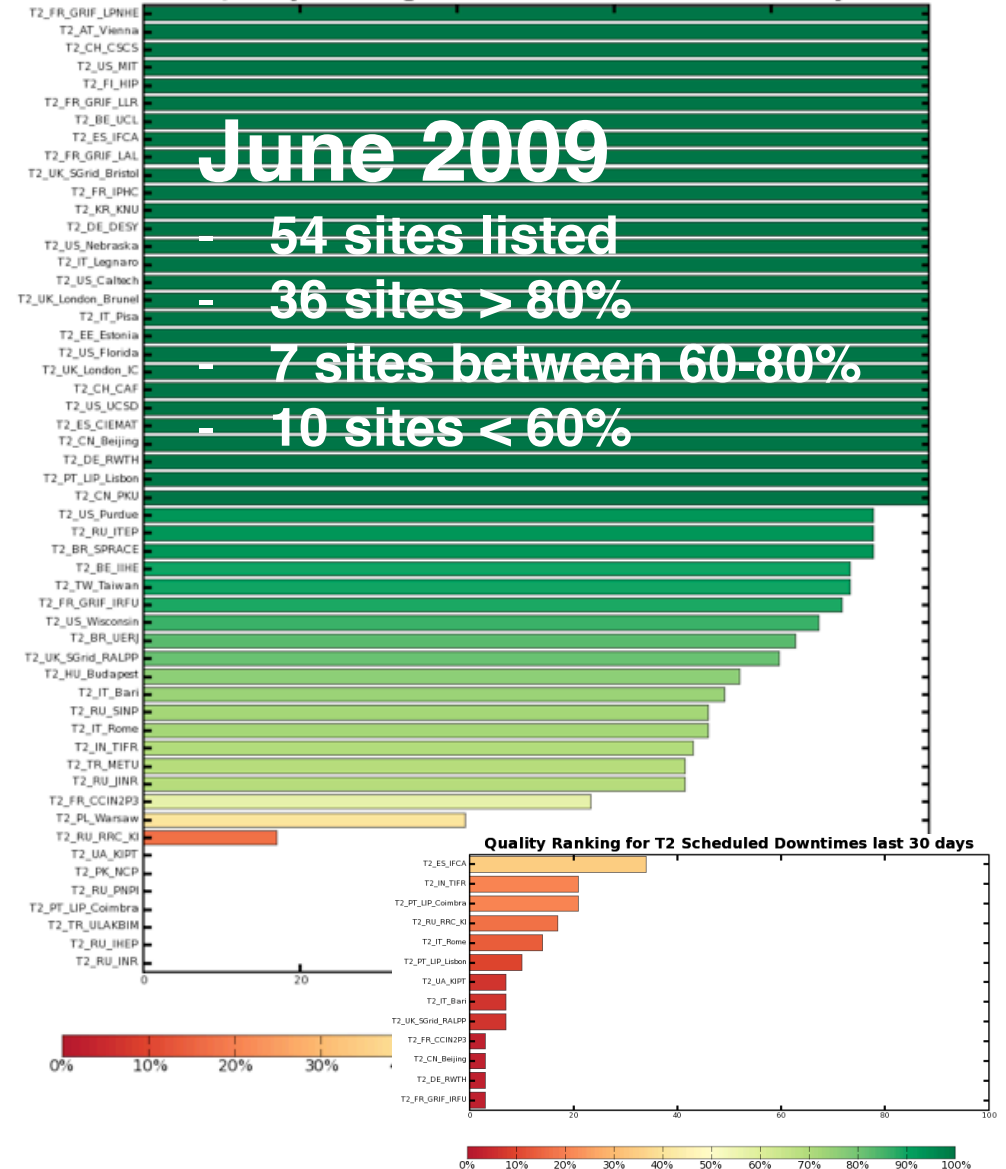


Site readiness T2: substantial improvement

Quality Ranking for T2 Site Readiness last 30 days



Quality Ranking for T2 Site Readiness last 30 days





Site Readiness Monitoring

Status last week (3.7.09):

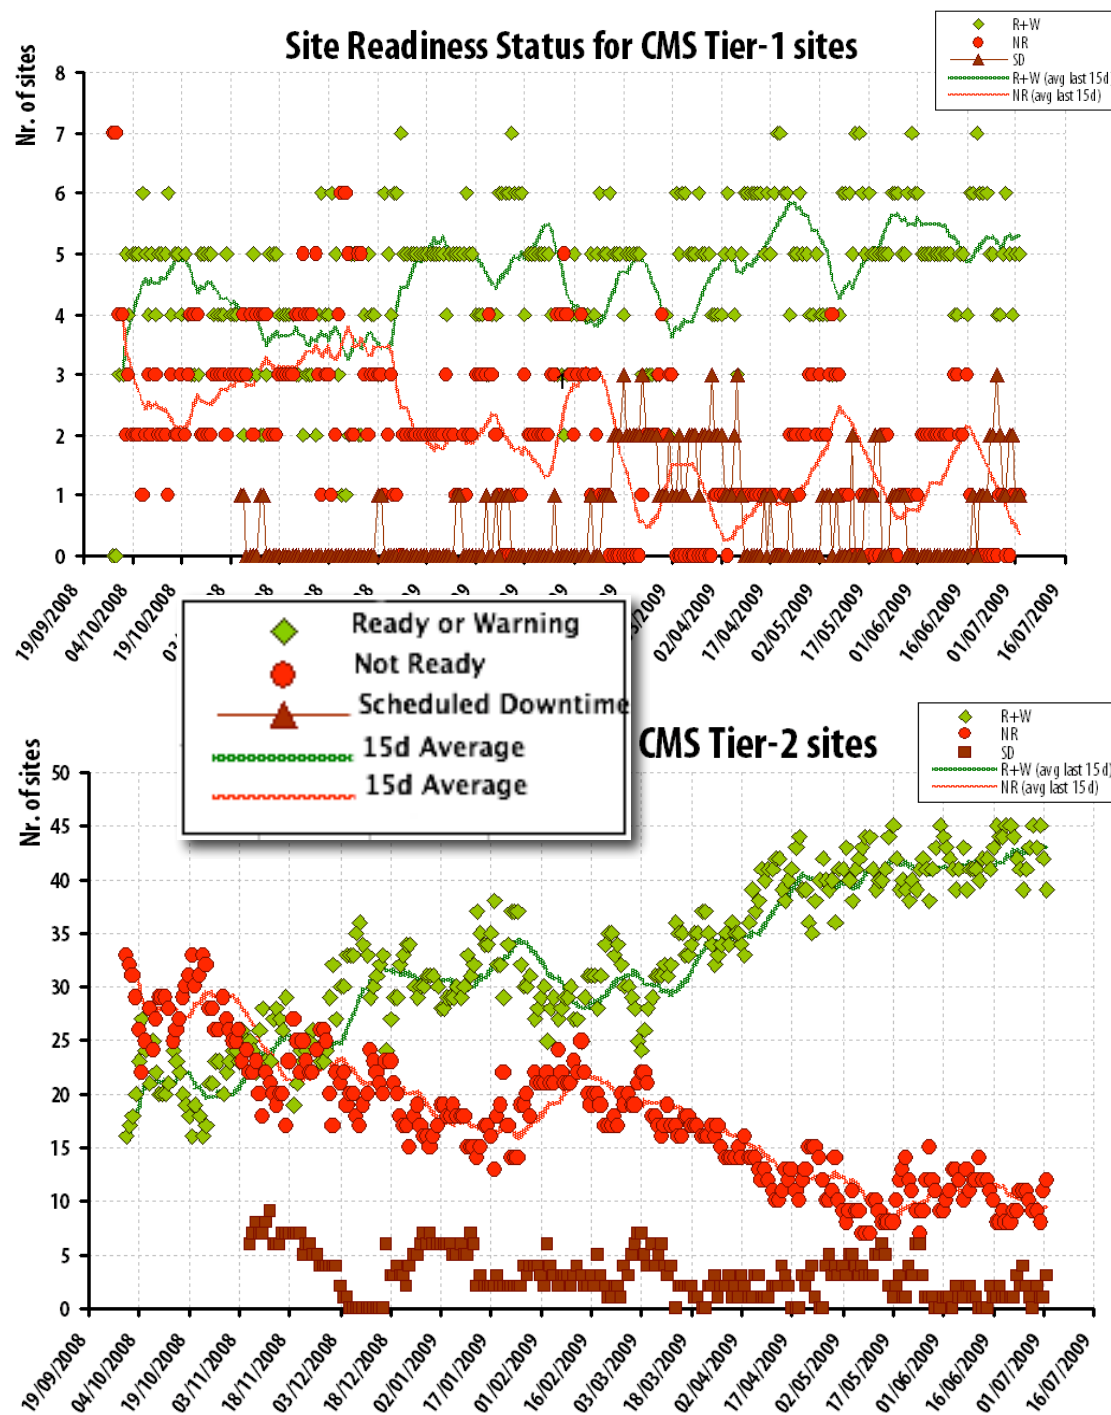
- 5 T1 sites are ready, out of 7:
70% of CMS T1s.
- 43 T2 sites are ready, out of 54
monitored: 80% of CMS T2s.

Site readiness:

Basic Services (SAM) ⊕ Successful
Testjobs ⊕
Data transfer links

The averages and the “jitter” in these
plots are:

- Tier1's: average is 5^{+2}_{-1} ,
spread is 3 = 60%
- Tier2's: average is 43, spread is 6
+ more sites will get ready





STEP'09: Scale Test of the Experimental Program




CMS Emphasis was on

- **T0: data recording in parallel with other experiments**
- **T1: tape access, testing simultaneously prestaging, processing**
- **Data transfers:**
 - T1 → T1: replicate 50 TB (AOD synchronization) between all T1s**
 - T1 → T2: stress T1 tapes, measure latency in transfer to T2**
- **Analysis at T2 :**
 - Demonstrate ability to use 50% of pledged resources with analysis jobs**

Final report during WLCG STEP '09 post mortem workshop 10./11. July



STEP'09: T0 tape writing performance



Sustained > 1 GB/s for 3 days,
No overlap with Atlas

The target of 500MB/s was exceeded in both testing periods

- Structure in first period due to problems in disk pool management
- Monitoring of tape writing and reading rates per VO can be improved



STEP'09: T1 Tape systems tests

- **For reprocessing of MC the required tape read rate between 50-250 MB/sec was tested, calculated according to the amount of data to be stored at Tier-1 center.**
 - **Overlapping tests at T1's with ATLAS was performed at some centers**
- **Preliminary results at T1 centers:**
 - Some sites met the metrics every day of the test
 - Other T1s met the metrics approximately 75% of the time.
 - At some centers the configuration and the overall stability has to improve, tests have to be repeated.
- **The bottlenecks were generally in the underlying tape systems and not in the ability of CMS to request data staging.**
 - CMS is carefully checking with the sites the implementation of tape families at each T1 center, which tends to concentrate data needed together on the same physical tapes.
 - Additionally CMS will be more actively managing the files expected to be on disk with Data Operations




STEP'09: T1 pre-staging and processing



Exercise rolling re-reconstruction:

- Pre-stage *1 day worth of data* and process it the next day
- Minimize disk consumption
- Maximize CPU efficiency because input is on disk

Pre-staging used for the first time in this planned and organized manner



- **Very good performance of all sites under multi-VO load**
- **CPU efficiency comparing w/o pre-staging measured, to be followed up**




STEP'09: Transfer tests



Emphasis on T1 → T1 transfers

- Use AOD synchronization between T1 sites after re-reconstruction

- Synchronize 50 TB of data between T1 sites, two tries




- all sites participated
- Transfer test was completed satisfactory
- links provided good rates between all sites



STEP0'09" Analysis Tests at 49 T2's + 8 T3's

Measure percentage of analysis pledge used:

- Standard analysis job: reading data, no stage-out to other T2



- Capable of filling majority of sites at their pledges, or above.
 - We used in aggregate more than the analysis pledge.
 - Roughly 80% success rate, 90% of failures are read errors



Preparation for data Taking

Computing shifts will start for CRAFT running (July 22)

- **Presence at CERN, FNAL or other CMS centre required**
- **Call for volunteers**
- **Shifts are MoA service credit task**



Computing Run Coordination, Shifts, Experts-On-Call

A Computing Run Coordinator (CRC) is responsible to oversee the operations of the CMS Computing infrastructure including the Tier-0, Tier-1, and Tier-2 centers and the services that interconnect the tiers for data transfer.

- attend run meetings and to report on data processing and data availability status and issues.
- decide when to contact expert on-call in case of emergencies.

A computing shift (CSP) will be operated from the CMS centre/FNAL ROC for at least 16h/day, potentially 24h/day, including Asia (China, India)

- follow a shift checklist, creates shift reports and Savannah tickets.
- monitor the computing infrastructure and services
- identify problems and trigger actions and calls.

Organized by Facility Operations/Peter Kreuzer

Experts-On-Call defined for computing to react to operational problems and deal with errors quickly.

March 30, 2009

Computing Preparation for Data Taking

M.Kasemann

5

For description see: <https://twiki.cern.ch/twiki/bin/view/CMS/ComputingShifts>



SLC5 Migration

The GDB recommends to start migration after STEP09, ie. **now**

- CMS is recommending to all CMS sites to migrate as soon as possible,
 - proposed migration deadline of September 1st, 2009

Status:

- 6 T1's will be done by September
- IN2p3 50% by Sept., 100% by end of 2009
- 25 T2's will be done by September, 6 not, rest: answer pending

Coordinated by Facility Operations, site Polling at:
(<https://twiki.cern.ch/twiki/bin/view/CMS/Poll-T1T2-SLC5>)
together with Offline for software validation.

CMS SLC5 migration plan at CERN

Migrate 10% by today/tomorrow, then: checked by CMS T0 teams
If ok, then migrate the rest of CMS T0 + CAF by July 19.

CMS SL5Migration documentation:

<https://twiki.cern.ch/twiki/bin/view/CMS/SLC5Migration>



Further tests and Productions

- As a STEP09 follow-up there will be targeted tests at some T1's to verify that problems are solved
 - If needed together with ATLAS
- An analysis end-to-end test is planned, goals are:
 - For computing: verify that all processing steps are 'luminosity-calculation save', ie. no un-accounted loss of events
 - These are a series of functional tests at T0, T1 and T2.
- MC production for 2009 LHC data is about to start
 - CMSSW release expected in a few days
 - Plan for 4 weeks for validation
 - An initial sample of ~200M events are required for 2009 analysis
 - Plan to finish production in September 2009



Summary

- **2-day CMS Global runs performed since March, about every week**
- **Long Cosmics run starting July 22.**
- **Monte Carlo Production at slower rate all the time.**
- **STEP09 was a valuable exercise with many tests overlapping with ATLAS and others.**
 - More information at the WLCG workshop July 9-10
- **Big improvement observed for stability and readiness of Tier2 sites.**
 - Tier-1 sites need to finish upgrades, need to show stability
 - More specific tests will be performed where needed.
- **An analysis-end-to-end test is planned in late summer.**
- **Big MC production is prepared, will start in a few days with new version of CMSSW.**