



CMS

CMS Status Report 2008-Q4

WLCG MB
January 27, 2009

Matthias Kasemann

Compact Muon Solenoid

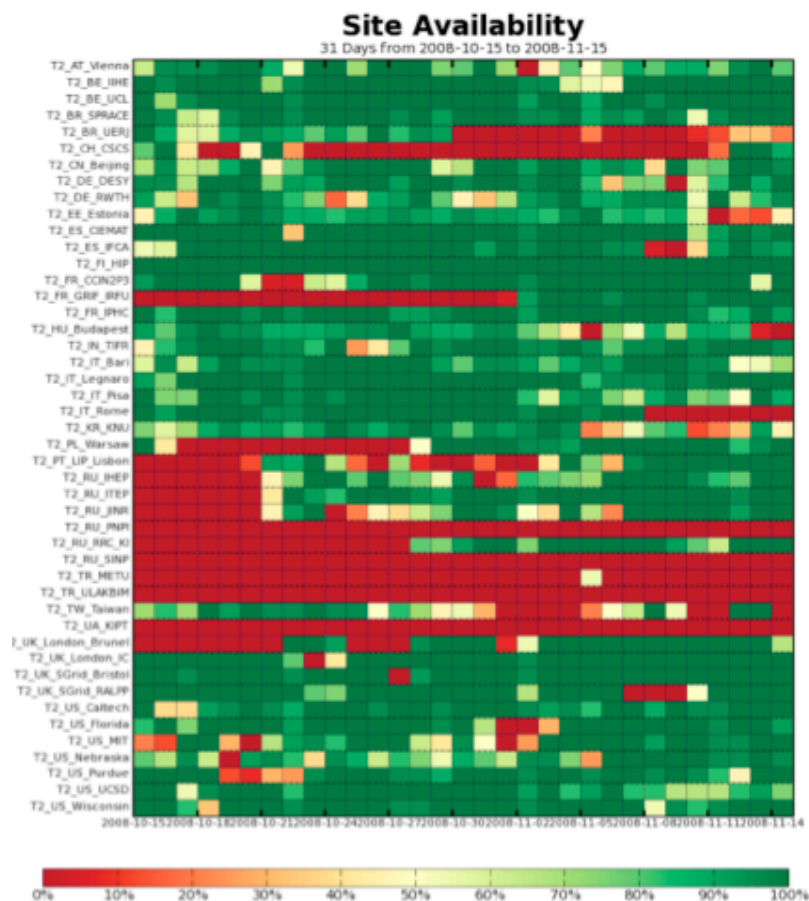
Outline

- CMS infrastructure
- MC production
- CMS data taking in 2008
- Outlook 2009
- SAM results Dec'08




Site Availability

CMS T0+T1s during CRAFT



CMS T2s during CRAFT

Site Availability, 2008-10-15 - 2008-11-15



Site availability is tracked at weekly facilities meetings and is part of shift duties.

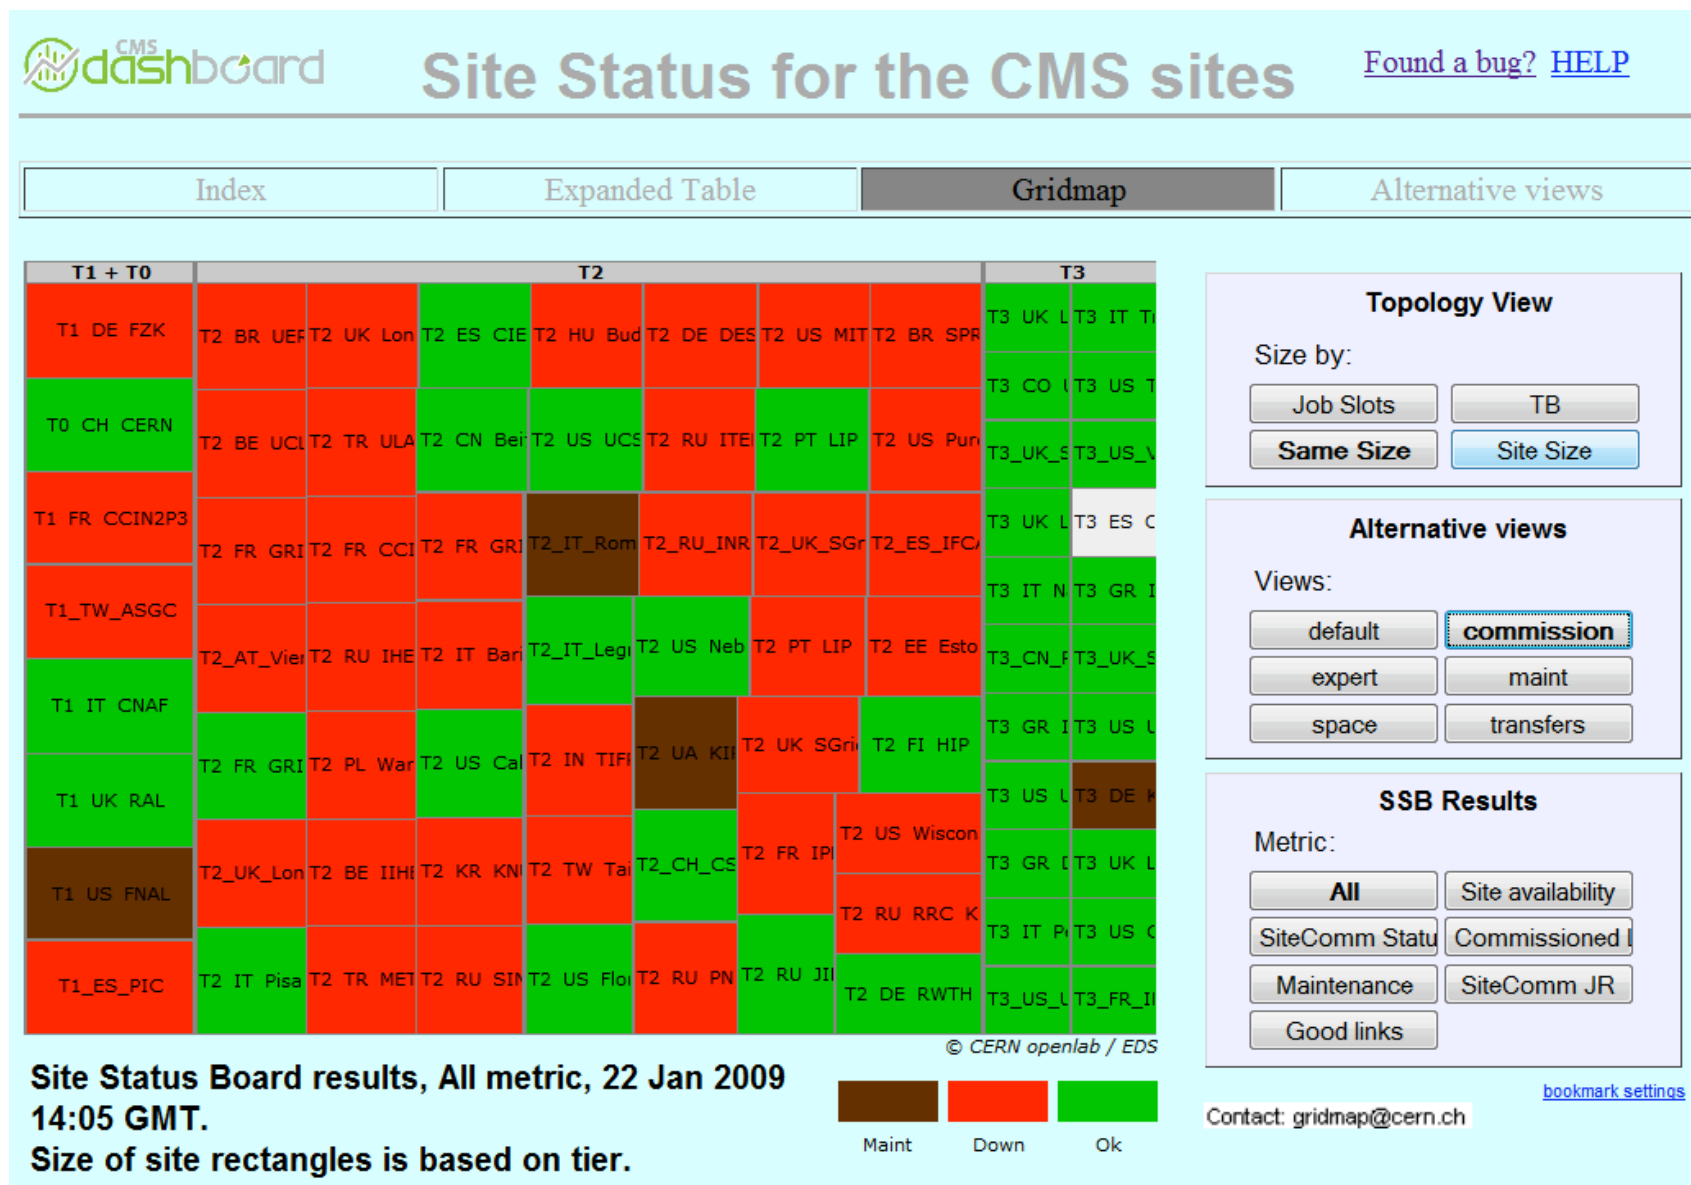


Site commissioning

- **We plan to declare working sites as “Commissioned”, i.e. ready for production and analysis activities**
 - Based on JobRobot and SAM monitoring
 - Inclusion of downtimes
 - Number of “good” transfer links per site
- **Goals reached**
 - The tools developed are now reasonably stable
 - We have several months of monitoring data to study
 - The readiness status flag is probably ready to be used
 - The new thresholds on the no. of commissioned links are
- **A taskforce started to complete:**
 - Increase the reliability of sites
 - Determine main failure reasons and help sites to improve
 - Understand how to automatically recover from most common failures
 - review importance of middleware bugs and configuration/support issues
 - Assess the evolution of failure rates
 - Measure the “best possible reliability”




Site commissioning: Snapshot





Monte Carlo Production

- Continuous request of MC samples produced to support commissioning and studies for Primary Dataset definitions
- Current round of MC production (started 15.9.):
 - DPG requests
 - For commissioning and detector studies
 - Physics requests
 - For physics performance studies
 - Production is running smoothly
 - Storage space at Tier-1s will become an issue...
- Production will continue to run constantly
 - Full simulation
 - Fast Simulation (>500M)



CMS Centres Worldwide

Major Centres (CERN, FNAL, DESY...) are working very well:
LHC first beam, cosmics, computing grid operations, outreach ...



CMS Experiment Control Room



Remote Operations Centre at Fermilab



**Permanently running video
links to operations centres**



CMS Centre at CERN: monitoring, computing operations, analysis

CMS Centres Worldwide



Major Sites are fully operational for monitoring, operations, shifts ...

- CMS Centre @ CERN
- LHC @ FNAL
- CMS Centre @ DESY

Modest Centres at CMS institutes for monitoring, analysis, students, outreach, video-conferencing ...

- CMS Centre @ Adana
- CMS Centre @ Dubna
- CMS Centre @ Mumbai
- CMS Centre @ New Zealand
- CMS Centre @ Pisa
- CMS Centre @ Rio
- CMS Centre @ Sao Paolo ...



Interested in having a CMS Centre at your institute?


Lucas.Taylor@cern.ch



CMS Data Taking '08

CMS Global runs since '07 (weekly)

- To prepare data taking with all components



CRUZET4 (Cosmics Run at Zero Tesla)

- First Global run with final CMS configuration (including newcomers EE and Pixels) accumulating data more stably. 38 M cosmic triggers logged. A total of ~ 300 million cosmic triggers recorded.


CRAFT (16 October-11 November @3.8T)

- Global run with final CMS configuration at Operating field (3.8T)
- Around 370 million cosmic triggers recorded.



CRUZET runs

Events collected by CMS during 2008 at B=0)



January 27, 2009
December 8, 2008
2008

CMS report, 2008-Q4T. Camporesi, D. Acosta,
T. Camporesi, D. Acosta, CMS gen meeting
CMS gen meeting




CRAFT (16.10.-11.11.08)

Goals:

- ◆ Run CMS for 4 weeks continuously to gain operational experience
- ◆ Study effects of B field on detector components (since MTCC)
- ◆ Collect 300M cosmic events with tracking detectors and field ($\approx 70\%$ livetime)

Achieved:

- ◆ Ran 4 weeks continuously
 - 19 days with $B=3.8T$
- ◆ 370M cosmic events collected in total
- ◆ 290M with $B=3.8T$ and with strip tracker and DT in readout
 - 194M with all components in






Data Taking Lifetime

Run Length

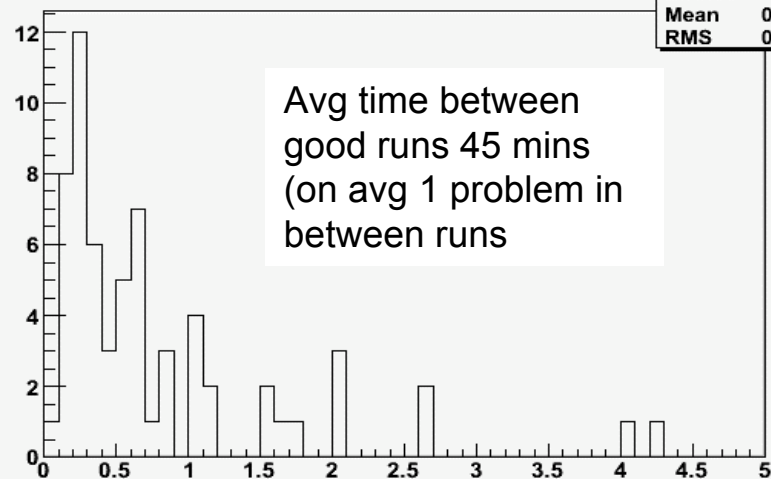
hRunLength	
Entries	90
Mean	3.262
RMS	4.218



hours


Time Between Runs

h1timeDown	
Entries	63
Mean	0.7895
RMS	0.8628



Problem occurrences per system
in last 11 days of CRAFT

Problems at Run Start/Stop
29-Oct thru 8-Nov





Alignment/calibration

- Data used with the magnetic field to calibrate and align the detector beyond what achieved with CRUZET
 - Momentum cuts, rigorous multiple scattering treatment
- Event-based alignment & calibration workflow strategy commissioned
 - Event-based calibration skims → “AlCaReco” streams
 - Specialized high rate HLT triggers → “AlCaRaw” streams
- Regular validation and sign-off procedures
 - Done weekly
- Laser pulsing of ECAL calibration workflow components commissioned
- Optical alignment of Tracker / Muons

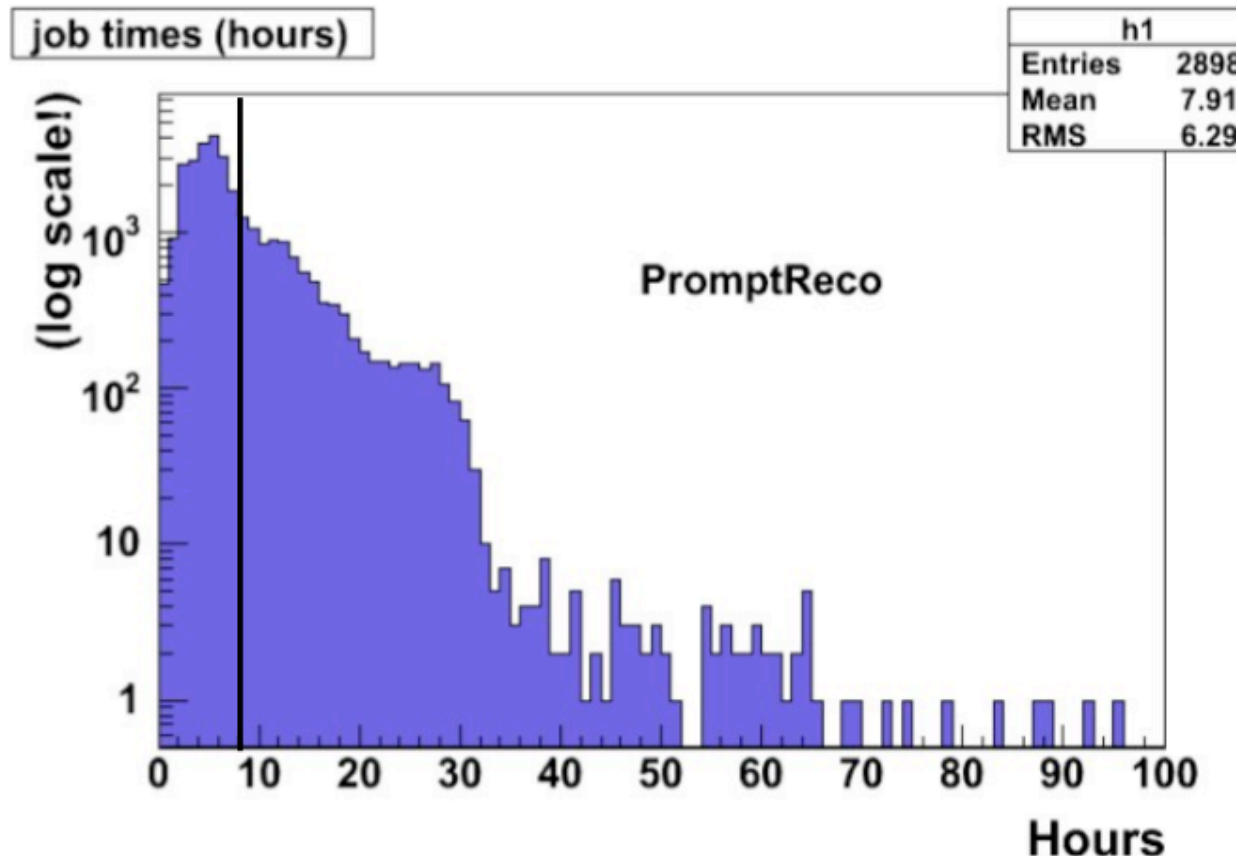


Tier-0 Operations

-Tier-0 in CRAFT

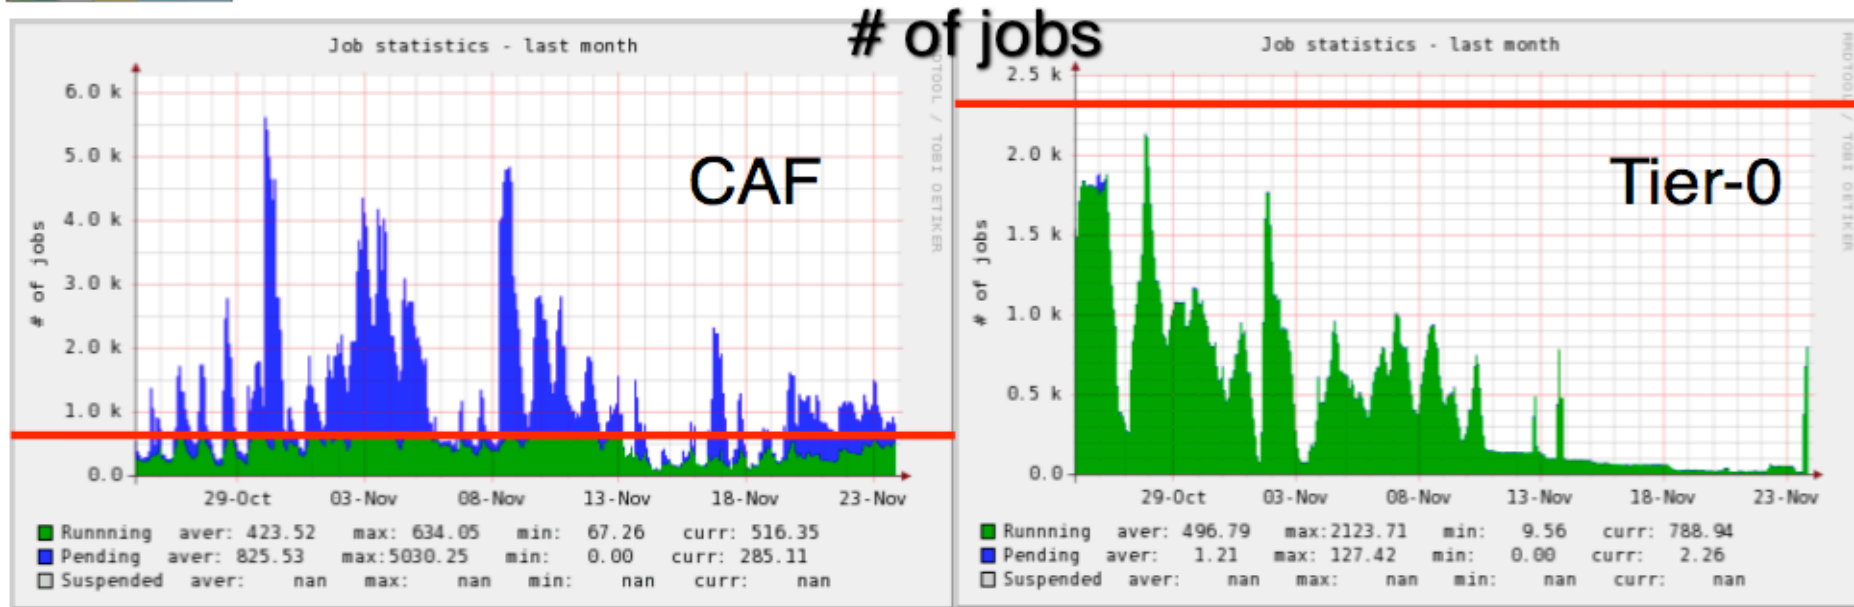
- Run re-packer and prompt reco; store data on tape
- Send raw and reco data to custodial Tier-1s
- Run AICaRECO/RAW and DQM harvesting

Most files reco'd
within 8 hours; tails
are understood





CAF at CERN (CRAFT)



CAF

- 632 slots available
- Average # slots: 67%
- Slot saturation ~ 1/2 of time during CRAFT

Tier-0:

- 2324 slots available
- Average # slots: 23%
- No saturation during CRAFT, but possible



CRAFT at the Tier-1s

Tier-1 tasks for CRAFT

- run re-reco of custodial data at each Tier-1
- produce skims of initial prompt and re-reco samples
- produce AODs during re-reco, in the future

Status and Lessons

- Skims on prompt reco were attempted but mostly failed (need better tools)
- Re-reco - one false start - was started on Dec 2 at RAL, FZK and IN2P3
 - Now partially complete, expect to complete by Dec 19
 - Cosmics 30%, MinBias 95% and Calo is somewhere in between.
- Skims of re-reco are in preparation: looks good this time :-)
 - RAW+RECO skims can only start once re-reco is complete (tools)
- DQM harvesting tools not yet ready at the Tier-1s



- **Software & Computing Infrastructure prepared:**
Simulation (full) of the 10 TeV datasets completed.
- **Preparation for physics prior to Sep 19: everything and everyone had been directed towards the imminent arrival of data. Now:**
 - Continue work on early publications
 - Continue the 900 GeV analyses – to completion
 - Start analyses with the 10 TeV samples
 - Restart Monte Carlo analysis approvals
 - Update some of our results (especially for 10 TeV)
- **Continue Physics Studies**
Fastsim: (~ 0.8 Bevs) assume 10 TeV, inst. Lumi up to 10^{32} , 50pb^{-1} .
Prepare CMSSW3 and generate samples for LHC re-start



Planning for 2009

- **Plans for 2009 operations**
 - Maintain Monte Carlo production at the sites
 - Support data taking for CMS
 - Support analysis
 - Monitor the sites/services/transfers
- **Test the systems**
 - End-to-end functionality tests
 - Production tool testing and validation
- **General planning for CMS computing**
 - **Update ECOM phase II report from summer 2008 before April/May.**
 - more disk/CPU resources will be available in 2009
 - **Resource planning for 2010 ~~needed by April~~ reviewed now**



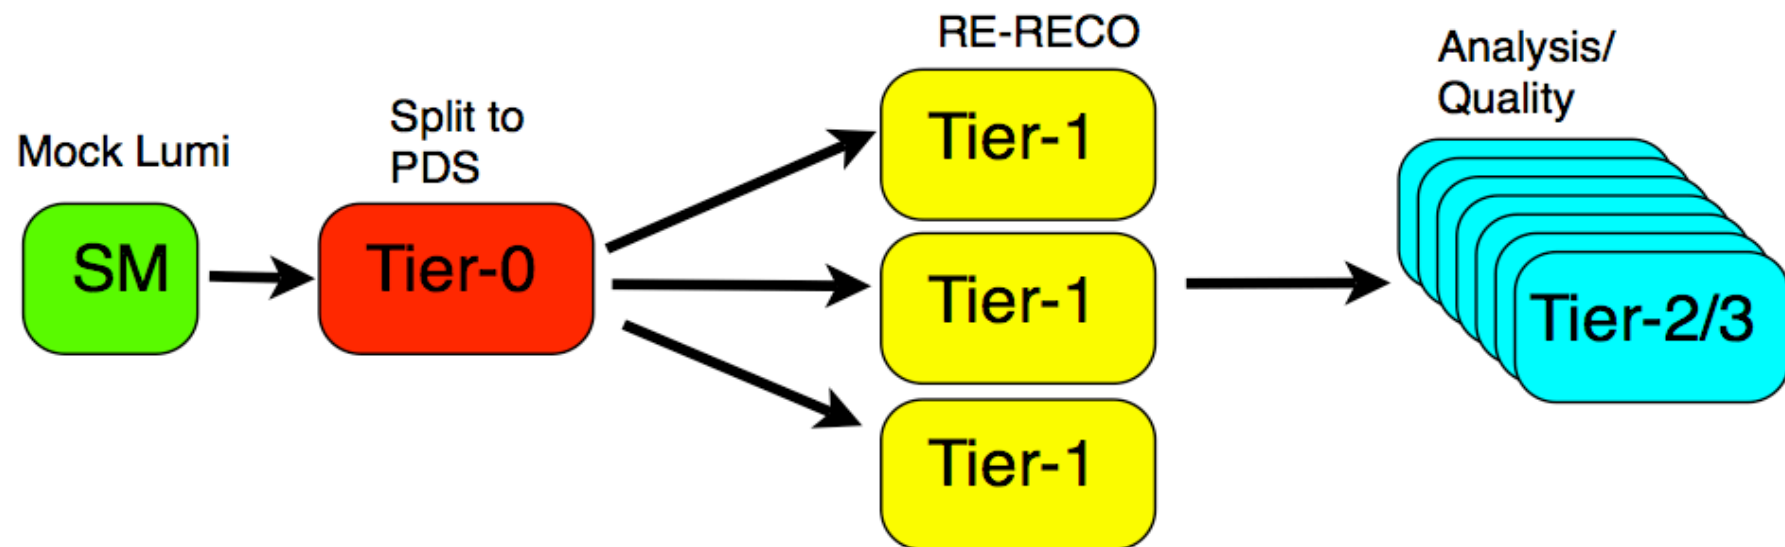
Tier-1 work

- **Dedicated to production activities**
 - custodial data
 - Reprocessing, Alignment/Calibration and Skimming
- **Will enforce authorization to reserve Tier-1s to the production users**
- **Working to improve the handling of custodial data**
 - Tape families / space tokens
 - Integration with PhEDEx
- **Improvements needed in pre-staging techniques**



End-to-End Functional Testing

- We propose to have an exercise to verify that luminosity, conditions, and quality information are consistently handled through the Tier-0 to the Tier-1s and Tier-2s and Tier-3s
 - We plan to run through the standard workflows at all tiers until we are convinced ourselves we handle the lumi and quality information consistently and correctly.



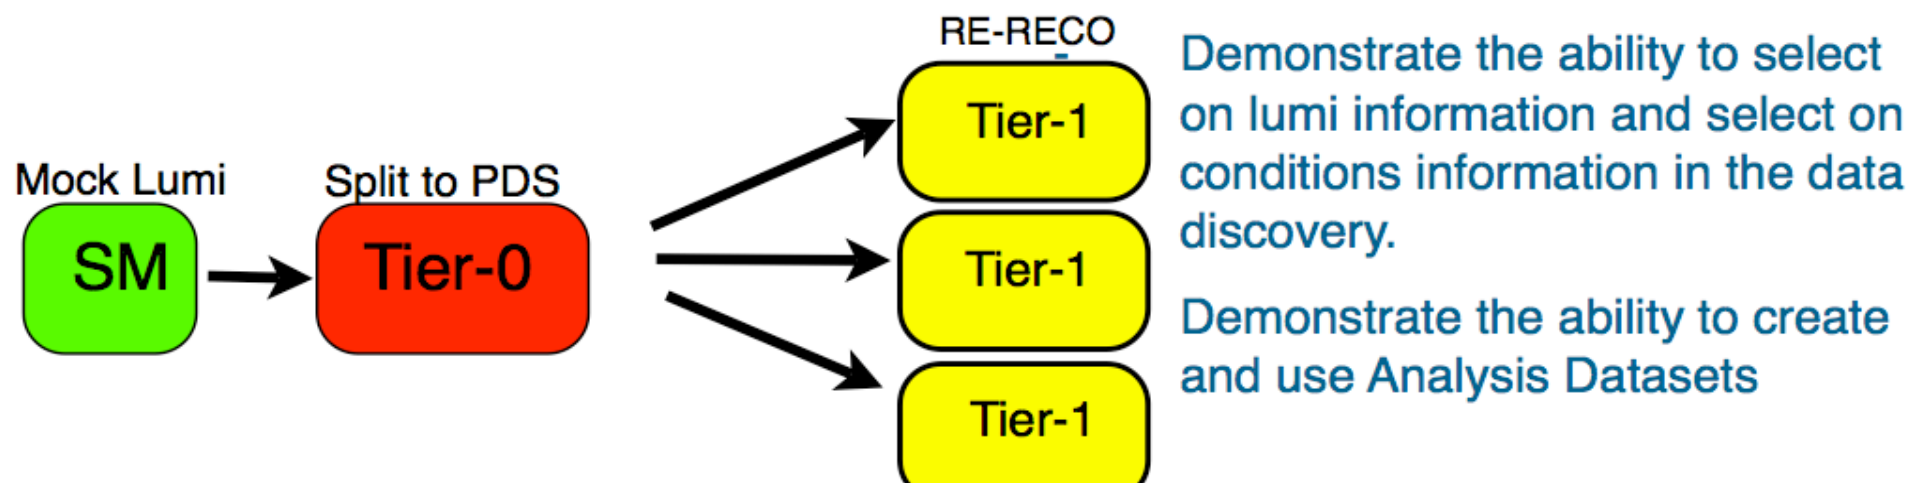


2009 functionality test - in steps

- ➔ **Produce Pythia events with mocked up luminosity information**
 - 100M Events in a single dataset
 - If possible introduce basic conditions or quality information
- ➔ **Stage the events through the repacker and divide into Primary Datasets**
- ➔ **Export the PDS to the Tier-1 sites**
- ➔ **Re-Reconstruct the data at Tier-1s**

- Demonstrate that we can run over every event and ensure that the luminosity information is valid at the end of a Re-Reconstruction step.
- Verify complete dataset processing and split by lumi section for Skim and PAT creation

➔ **Transfer Data to Tier-2s and access them with CRAB**





Moving towards Analysis Operations

- **Task Force on Analysis Support is starting;**
 - F2F meeting following CMS week
- **Goals of task force (still a draft):**
 - Improve the success rate for CRAB submitted analysis jobs
 - Determine the tools needed to assess, categorize and reduce the number of non- zero exit code applications submitted through CRAB
 - Work with the dashboard developers to get (new) views to identify issues quickly
 - Develop techniques and tools for providing direct user support
 - Formalize the collection of user feedback on functionality and experience
 - **Transition to steady analysis operations and evolve beyond community support. This task force is a first step.**
- **Deliverables:**
 - Assessment of effort and tools required to do analysis support
 - Document tools and techniques in preparations for analysis operations
- **Time scale** - March 2009 for transition to operations

1) Maintenance & Operation

ating

2) Software, Computing & Physics Analysis

DRAFT
For discussion

Services available

Install ES1

Install ES2

Tracker Cooling Plant Operational

Close CMS

Contingency on closure or CRAFT

CRAFT

CMS READY for Beam

Dec

Jan

Feb

Mar

Apr

May

Jun

Jul

Aug

Sep

Fullsim (CMSSW2.2.x), then Fastsim

Release CMSSW3_0
(limited validation, step towards 3_1)

End-Feb: Input for 3_1:
CRAFT results, Trigger Review (menu) etc.

Release CMSSW3_1 (LHC Startup)

↕ Full validation of 3_1 (incl.
production and physics)

↕ Start Fullsim production 3_1

← Deploy 3_1 widely
CMS gets familiar with 3_1

↕ Start Fastsim

CMS READY for Beam



Summary

- Data collected in the CMS Global Runs, CRAFT and MC data are being analyzed by physicists
 - A 1st reprocessing with latest software has been completed
- CMSSW 3.0 in Feb with new OS/compiler (32bit) is being prepared
- MC prod continues (standard + >500 M fast simulation)
- CAF/Tier-0/Tier-1/Tier2 infrastructure defined
 - Monitoring tools, commissioning policies, ...
- Upgrades in CMS tools (CRAB, PA, PhEDEx, ...)
- Luminosity-safe analysis exercise planned
- Resource requirements for 2009/10 are being re-discussed based on latest information on LHC schedule









SAM results: CMS availability in December









- Critical services
 - CE, SRMv2 (since December)
- Critical tests
 - CE: job submission (run by CMS), CA certs (run by OPS)
 - SRMv2: "lcg-cp" (copies a local file to the SRM in the provided tokens and default space and then it copies it back) (run by CMS)
- Short Summary:
 - Excellent results for CERN, CNAF, FNAL*, PIC and RAL
 - Sites with downtime or problems: ASGC, IN2P3 and FZK (see next slides)







- Just a few SRM failures; downtime on 9/12 due to Oracle upgrade
- Excellent reliability







- Oracle problem caused a very high load on CASTOR database. Debugged with help from CERN CASTOR and DB team
- Poor reliability; no downtime was scheduled to fix the problem









- Excellent reliability
- A problem not shown by any CMS critical test: software area was unavailable on 1-2/12 due to GPFS problems









- CE at FNAL still ignored in GridView...
 - CE and SRMv2 are in two different BDII sites, GridView does not aggregate them
- Very good reliability, but risk of overestimating it due to CE missing









- CE: electrical maintenance on 1-4/12; jobs stuck on 26/12
- SRMv2: downtimes declared 17-19/12 (why barely visible on Gridview?). SRM overloaded due to recursive srmls issued by CMS ProdAgent
- Very poor reliability, diagnosis took long



- Maintenance from 14 to 16 due to storage service upgrade
- Unreliability on 16/12 "amplified" due to CE test failure in a short time window not in maintenance?
- Modulo that, excellent reliability



- CE: extended downtime of all CEs from 25/12 to 29/12, then frequent job aborts ("unspecified gridmanager error")
- SRMv2: several dCache errors, most about full devices
- Poor reliability during the holidays



- Excellent reliability

