

## Abstract

The CMS experiment has adopted a computing system where resources are distributed worldwide in more than 50 sites. The operation of the system requires a stable and reliable behavior of the underlying infrastructure. CMS has established procedures to extensively test all relevant aspects of a site and their capability to sustain the various CMS computing workflows at the required scale. The **Site Readiness** monitoring infrastructure has been instrumental in understanding how the system as a whole was improving towards LHC operations, measuring the reliability of sites when running CMS activities, and providing sites with the information they need to solve eventual problems.

This poster reviews the complete automation of the Site Readiness program, with the description of monitoring tools and their inclusion into the Site Status Board (SSB), the performance checks, the use of tools like HammerCloud, and the impact in improving the overall reliability of the Grid from the point of view of the CMS computing system. These results are used by CMS to select good sites to conduct workflows, in order to maximize workflows efficiencies. The performance against these tests seen at the sites during the first years of LHC running is as well reviewed.

## Site Readiness

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

The **Site Readiness** started as one of the activities of the PADA (Processing and Data Access) Task Force in 2008, whose objective was to guarantee that the data processing workflows at Tier-1 and Tier-2 sites could be performed efficiently/reliably, in an heterogeneous distributed computing resources context.

The program soon became instrumental for tracking down problems at sites, improving their performance, guiding workflows and keeping operations at stable and reliable levels.

The Site Readiness makes use of several sources of information to assess the readiness of a site to run CMS workflows:

- CMS SAM tests: jobs sent to sites to test specific services;
- job 'Robot' load generator: simple jobs reading data;
- data transfers: transfer quality and commissioned links;
- scheduled Downtimes: from GOCDB/OIM on CMS services.

Site Readiness metrics were established to guarantee data processing can be performed efficiently and reliably, and have evolved slightly since the creation of the program.

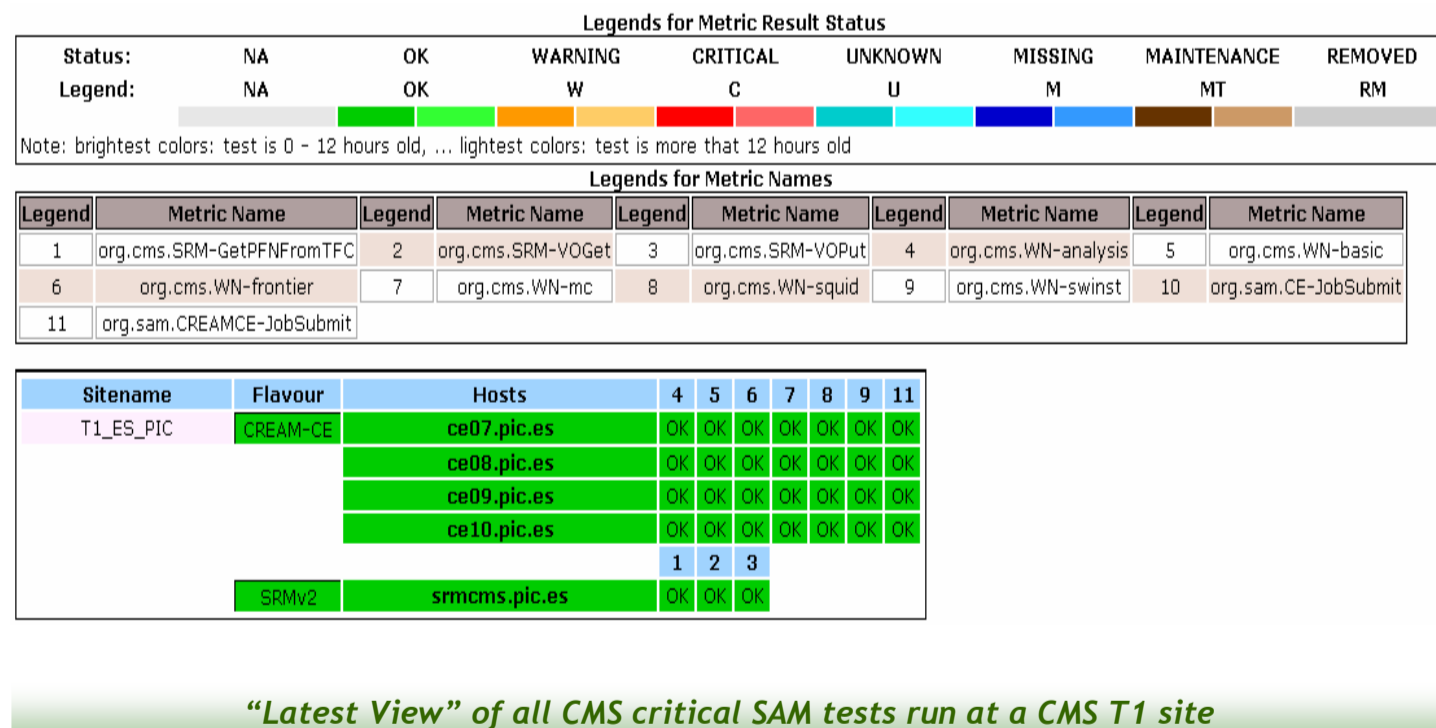
## Site Availability

<https://dashboard.cern.ch/cms/>

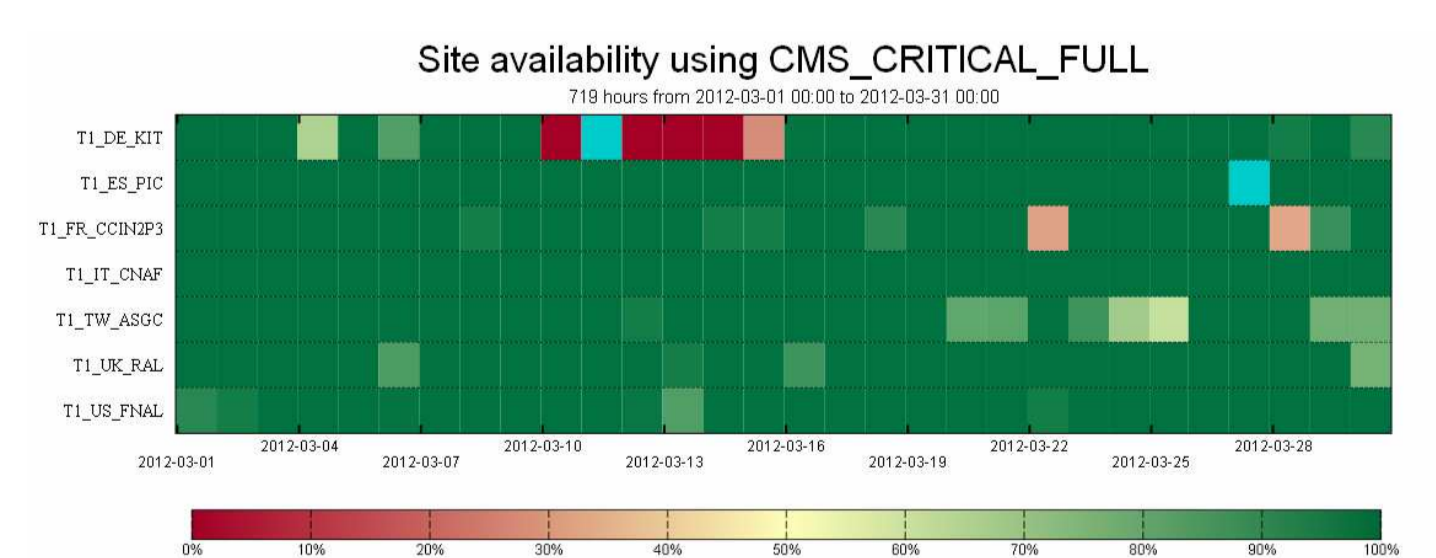
Like other LHC experiments, CMS uses the Service Availability Monitoring (SAM) framework to run specific tests on computing (CE) and storage (SRM) resources at the CMS sites. These tests allow us to determine, among other things, whether: @CHEP12 [1]

- it is possible to send and run jobs;
- the CMS software is correctly installed and configured;
- it is possible to access local CMS data in a job;
- it is possible to copy data in and out of the local storage.


A failure of a critical test means that the site is considered unavailable, until all critical tests are passed. The **site availability** has been used for a very long time to estimate the quality of the site. Currently, ACE methodology is used to compute the site availability.



"Latest View" of all CMS critical SAM tests run at a CMS T1 site



Site availability ranking using CMS\_CRITICAL\_FULL



Site Availability March 2012 history (top) and ranking (bottom) for CMS T1 sites tested

## HammerCloud

<http://hammercloud.cern.ch>

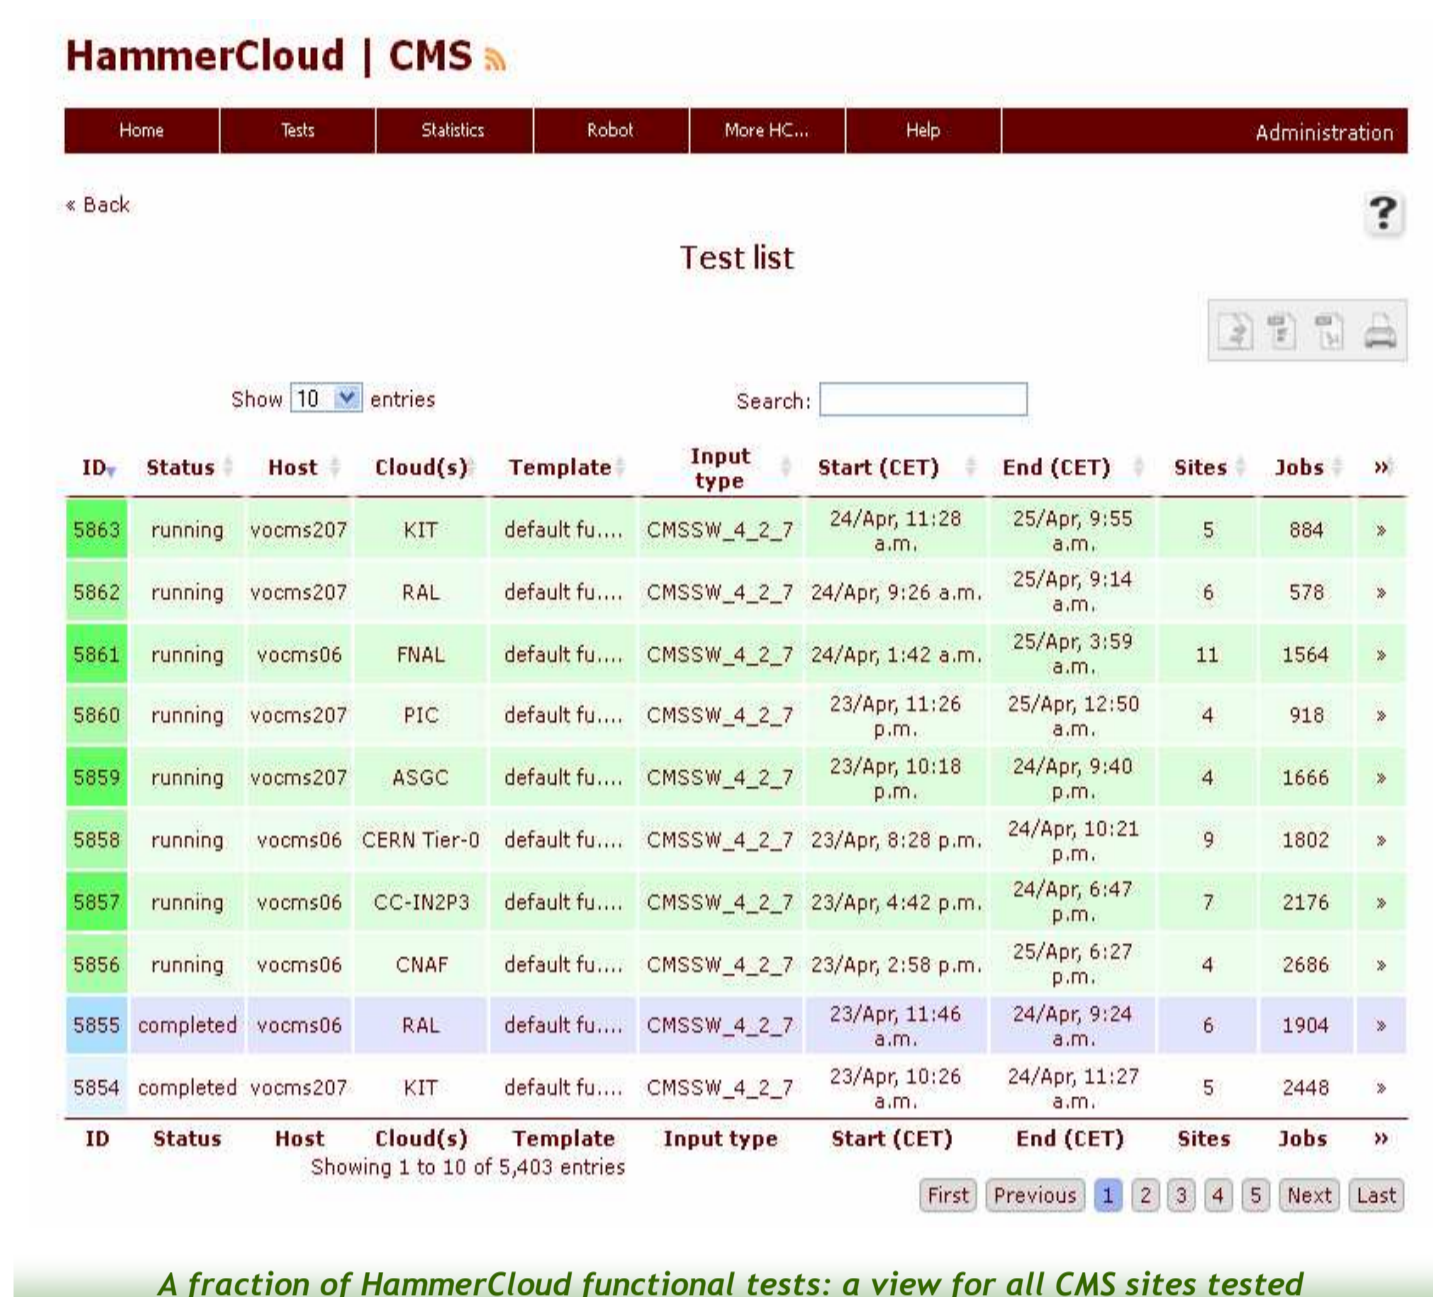
Currently, CMS uses the **HammerCloud** tool, developed by the CERN IT-ES group, to submit realistic CMS jobs with CRAB to simulate analysis on a dataset of ~50 GB available at the sites.

There are two modes of operation: @CHEP12 [2]

- **Functional tests:** continuous low rate job submission;
- **Stress tests:** filling sites with jobs on demand.

HammerCloud, which replaced the old CMS JobRobot in early 2012, shows information and statistics on its jobs via a web interface and it allows privileged users to define and submit tests.

Currently running functional tests. ~15,000 jobs/day [-60 sites]



A fraction of HammerCloud functional tests: a view for all CMS sites tested

## Data Transfers

<http://cmsweb.cern.ch/phedex>

For sites to be usable, they need to have sufficient data transfer bandwidth and connections among them. A link is a pair source→destination, and these need to be operational:

- Tier-0 → Tier-1: to export raw and reconstructed data;
- Tier-1 → Tier-1: for synchronisation and large-scale 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.

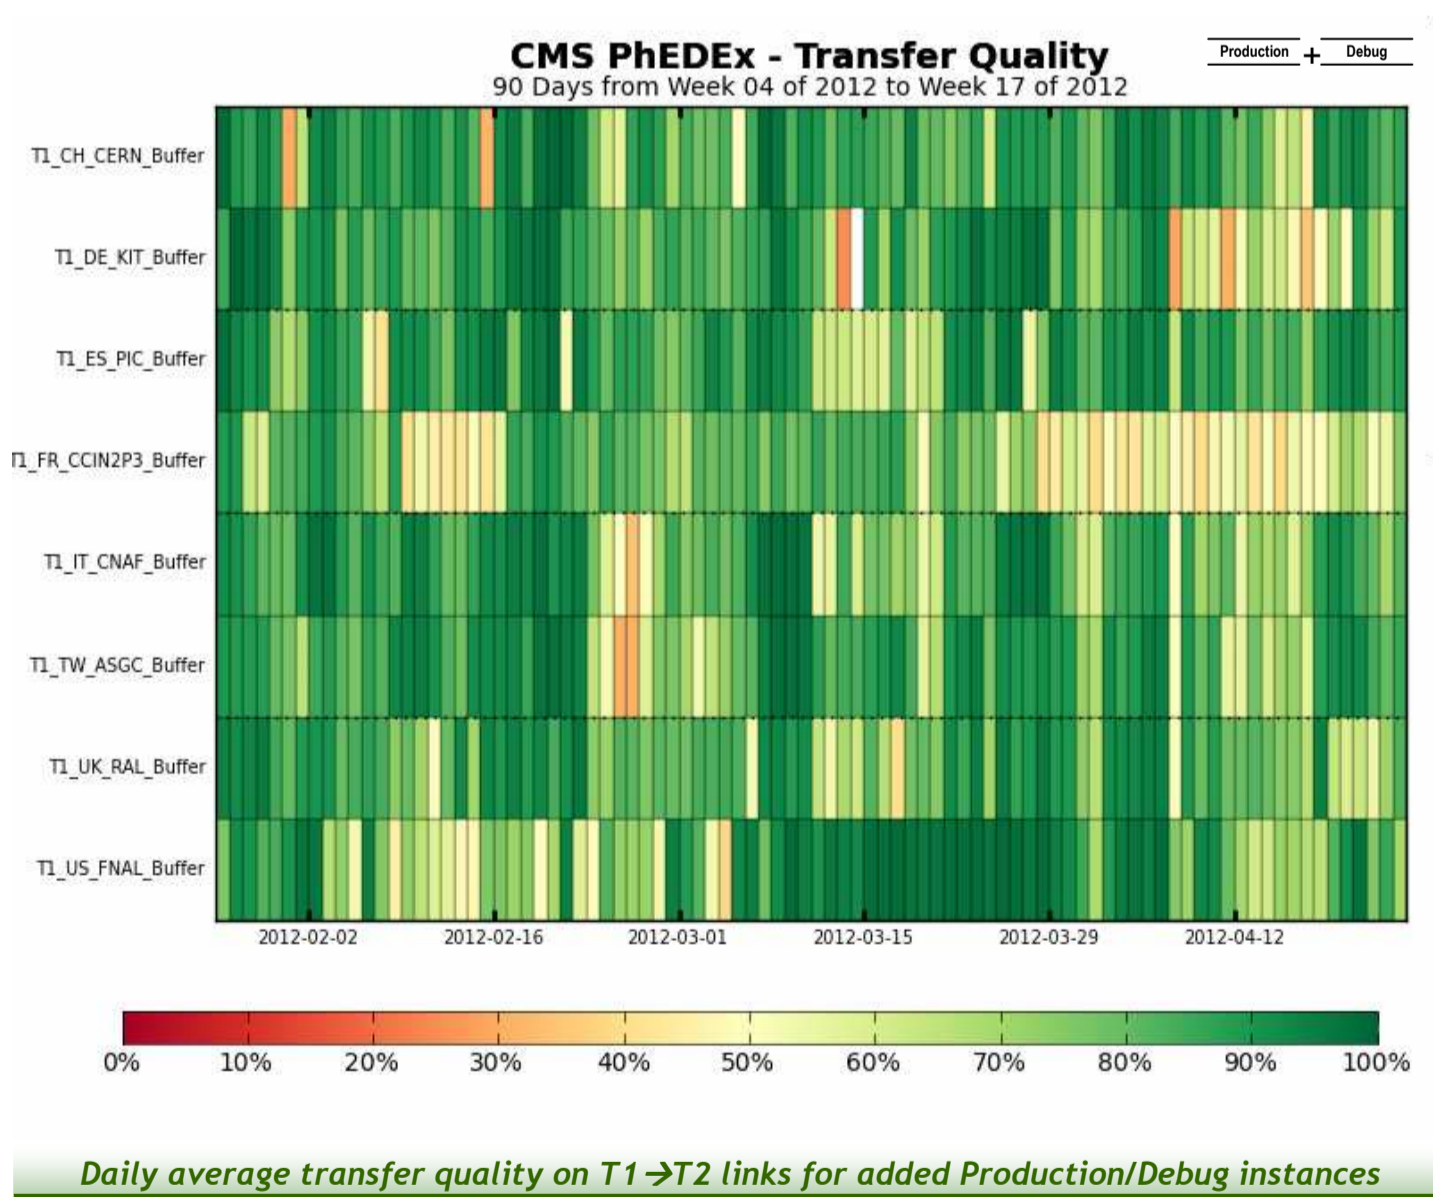
The Debugging Data Transfers (DDT) task force defined metrics, provided a procedure and the tools to test the links and assisted sites solving problems. Tests regularly ran to commission new links on Debug PhEDEx (Data Management) instance, from 2007 to 2009. Today, the tests run on demand and managed by the sites.

Each commissioned link is enabled for Production use. In March 2012:

- All 7 T0 → T1 and 49 T1-T1 cross-links in production;
- 403/416 (97%) T1 → T2 links in production;
- 330/416 (79%) T2 → T1 links in production.

All these ~800 production links are continuously exercised with test transfers at 0.25 MB/s/link (in Debug instance, except T2-T2 crosslinks, which are as well used in CMS but not monitored in Site Readiness). Both **commissioned links and transfer qualities** are used in the Site Readiness. Adding Production transfers:

- Routine WAN transfers of ~200 TB/day (~1.5-2.5 GB/s); @CHEP12 [3]
- Enough to detect systematic transfer problems.



Daily average transfer quality on T1→T2 links for added Production/Debug instances

## Declared Downtimes


CMS service downtimes are extracted from GOCDB/OIM and integrated in the Site Status Board:

- Consider services that have critical tests defined: CE/CREAM-CE and SRMv2;
- full-Maintenance when all instances of a service are in maintenance.

At Risk, Scheduled and Unscheduled downtimes are traced, for the CMS services or for the entire site.

## Put All Together

CMS collects and display all Site Readiness information in **Site Status Board**: @CHEP12 [4]



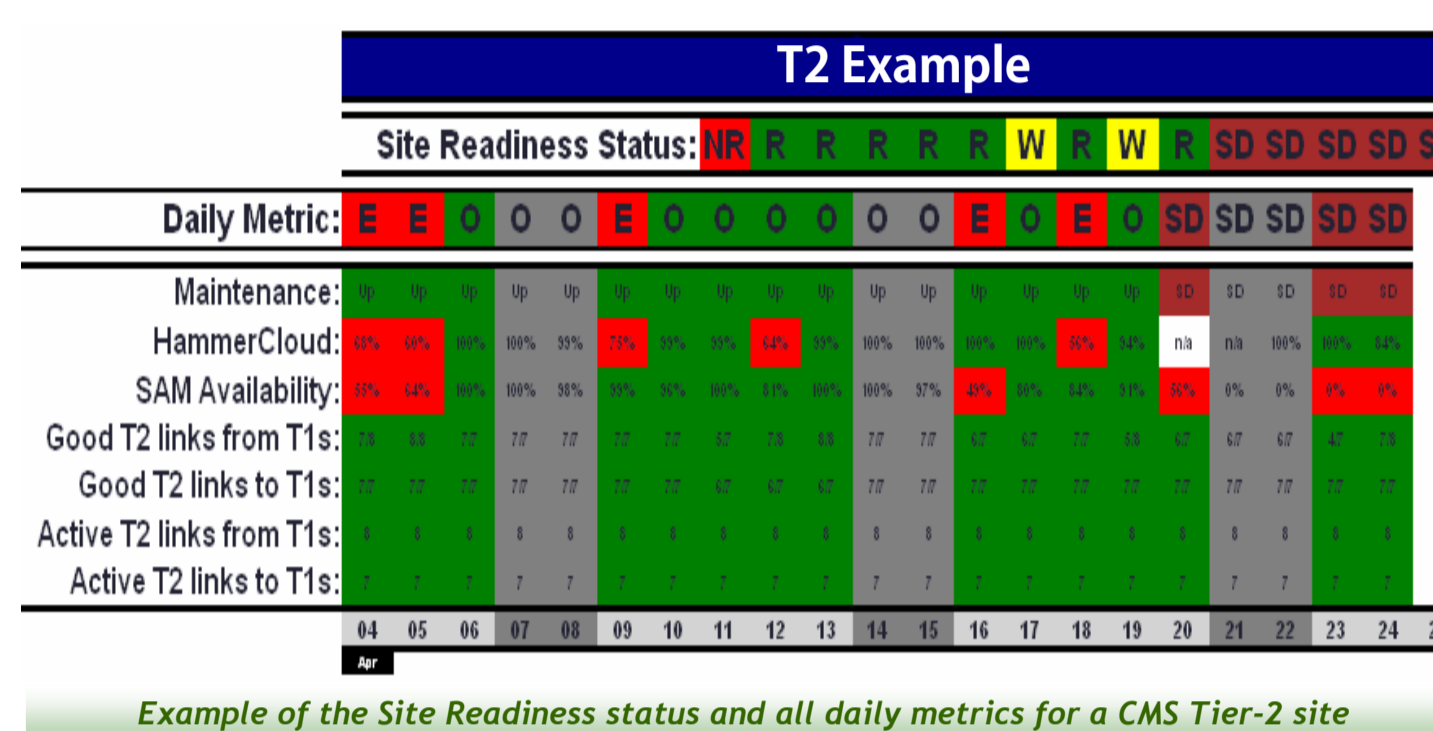
Site Status Board view used for CMS Site Readiness (Tier-1s shown)

Daily Metrics for Tier-1 sites		Daily Metrics for Tier-2 sites	
daily SAM site availability ≥ 90%	daily HammerCloud job success ≥ 90%	daily SAM availability ≥ 80%	daily HammerCloud job success ≥ 80%
link from Tier-0 commissioned	commissioned links to Tier-2 sites ≥ 20	commissioned links to Tier-1 sites ≥ 2	commissioned links from Tier-1 sites ≥ 4
commissioned links from other Tier-1 sites ≥ 4	≥ 50% of links of a group with ≥ 50% transfer quality		

Metrics are calculated for each CMS site daily. Those combine into a single daily "Site Readiness status":

READY - WARNING - NOT-READY - SCHEDULED-DOWNTIME


- Uses the history of last 7 overall Daily Metric;
- allows some degree of unreliability in thresholds and stability;
- intermediate warning state to give sites the time to recover;
- weekend failures for Tier-2s do not negatively count in the evaluation [grey].



Example of the Site Readiness status and all daily metrics for a CMS Tier-2 site

Sites have an easy way to know if CMS is finding troubles there:


- The program provides monitoring plots. Computing Shifters alert;
- results/problems are as well reported on Operations meetings weekly.



Evolution of Site Readiness status for CMS Tier-2 sites, since Site Commissioning started

Use "Site Readiness status" history to flag good/bad sites:

- Sites in **READY** or **WARNING** status are considered **stable & reliable**;
- the fraction of time a site has been stable & reliable (ignoring declared scheduled Downtimes) is shown for last 15 days;
- this is used as indication to send jobs for production/analysis on good sites.



Example of Site Readiness last 15 days ranking for Tier-2 CMS sites

## conclusions

Site Readiness activities are crucial for keeping the CMS distributed computing system stable & reliable for Computing Operations

- Continuous monitoring of Grid & CMS services at sites;
- The available information condensed in a single estimator → **positive trend**;
- Helps production and users to select reliable T2 sites;
- Regular meetings to discuss usual failures, help sites to improve, feedback for robustness of CMS tools/services, increase reliability of sites...