



The Compact Muon Solenoid Experiment

# Conference Report

Mailing address: CMS CERN, CH-1211 GENEVA 23, Switzerland



11 May 2012

## Towards higher reliability of CMS Computing Facilities

Jose Flix Molina for the CMS Collaboration

### 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 paper 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. Based on these results, CMS automatically excludes 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 will be as well reviewed.

Presented at *CHEP 2012: International Conference on Computing in High Energy and Nuclear Physics*

# Towards higher reliability of CMS computing facilities

**G Bagliesi<sup>1</sup>, K Bloom<sup>2</sup>, C Brew<sup>3</sup>, J Flix<sup>4,5</sup>, P Kreuzer<sup>6</sup>, A Sciabà<sup>7</sup>**

<sup>1</sup> University of Pisa, Italy

<sup>2</sup> University of Nebraska-Lincoln, USA

<sup>3</sup> Particle Physics Department - Rutherford Appleton Laboratory (STFC), UK

<sup>4</sup> Port d'Informació Científica (PIC), Spain

<sup>5</sup> Centro de Investigaciones Medioambientales y Energéticas (CIEMAT), Spain

<sup>6</sup> RWTH Aachen University, Germany

<sup>7</sup> CERN, Switzerland

E-mail: [jflix@pic.es](mailto:jflix@pic.es), [Andrea.Sciaba@cern.ch](mailto:Andrea.Sciaba@cern.ch)

**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 contribution 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.

## 1. Introduction

The Large Hadron Collider (LHC), located at CERN, became operational in spring 2010 with p-p collisions at centre-of-mass energy of 7 TeV, increasing up to 8 TeV in 2012. In 2011, the collected data was about two orders of magnitude larger compared to 2010. The Compact Muon Solenoid (CMS) [1] is one of the four detectors that observes the collisions and it generates PetaBytes of data per year. To scientifically exploit the data, the data processing requires the use of computing and storage resources from several centres outside CERN. These are in fact coordinated by the Worldwide LHC Computing Grid (WLCG) [2], which at most sites exploits the computing infrastructure provided by other Grid projects, like EGI, Open Science Grid and NorduGrid.

In the case of the CMS collaboration, around 60 sites from about 20 countries are involved, ~100 sites if considering small university computing facilities. They are organized with a tiered structure, where different tier levels correspond to different functions. The Tier-0 site is CERN, and takes care of the prompt event reconstruction and detector calibration, the distribution of raw and processed data to

external sites and the backup storage of the raw data. Seven Tier-1 sites run the subsequent reprocessing, including data skimming, keep an active copy of the raw data and store the Monte Carlo (MC) generated at Tier-2 sites, and also generate MC samples. Finally, around 50 Tier-2 sites get samples of the skimmed data for analysis and are used to run the MC simulation. A complete description of the CMS computing model and its services can be found elsewhere [3].

Given the complexity of the infrastructure, it is important to measure its performance in a continuous way, in order to inform the sites of any problem CMS is encountering, or will encounter, when having activities at that site. The Grid projects operating the infrastructure have their own procedures to identify and correct problems, but these do not necessarily reflect the usage CMS does of the resources. For this reason, CMS has established a set of techniques and tools intended to provide a better picture of the site performance and reliability. The following sections describe how this is performed: the procedure to test sites in an automatic and continuous way, the results obtained so far, and the implications for site usage.

## 2. Site evaluation techniques

Well before datataking, CMS undertook periodic computing challenges of increasing scale and complexity to test its computing model and the Grid infrastructure. Performance values were measured, problems were identified and feedback into the design, integration and operation was provided (for example [4]). This allowed the collaboration to quantify the readiness of the multi-tiered sites and provided invaluable input to the operations of the computing centres.

However, the operation of the complete CMS computing system requires stable and reliable behavior of the underlying heterogeneous (in computing resources, regions and support) infrastructure at all times. The Site Readiness [5] [6] activity 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 a heterogeneous distributed computing resources context, and help sites commissioning their facilities prior data taking. This information is used by the sites to become aware of problems, and by CMS to plan the distribution of the workload such that temporarily unreliable sites are not used. The program soon became instrumental for tracking down problems at sites, improving their performance, guiding workflows and keeping operations at stable and reliable levels.

In order to accomplish that, custom tests are regularly run at each site, and they are conceived to check every possible functionality exploited by CMS, aiming at having the highest possible correlation between failures of these tests and of real CMS jobs. Sites must satisfy certain lower limits on the success rate of these tests to be considered reliable.

The following information is used to evaluate the readiness of a CMS site:

- The results of the CMS SAM tests, via site availability, running on Grid site resources to check their functionality and the local CMS software and configuration;
- the success rate of the HammerCloud jobs, a job load generator simulating user data analysis;
- the number and the quality of the data transfer links used in production;
- the downtimes scheduled by the site.

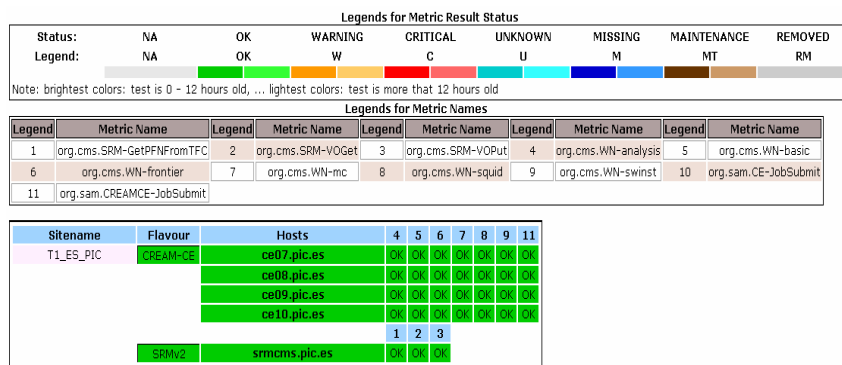
### 2.1. SAM tests: the site availability

In the WLCG, all Grid services are periodically tested via Nagios probes in a framework called SAM (Site Availability Monitor) [7], which executes periodic tests on all the Grid services within the infrastructure. SAM provides one of the main sources of information for the Grid operations and is used to measure the availability of Grid services. ACE component is used for computing status and site availability.

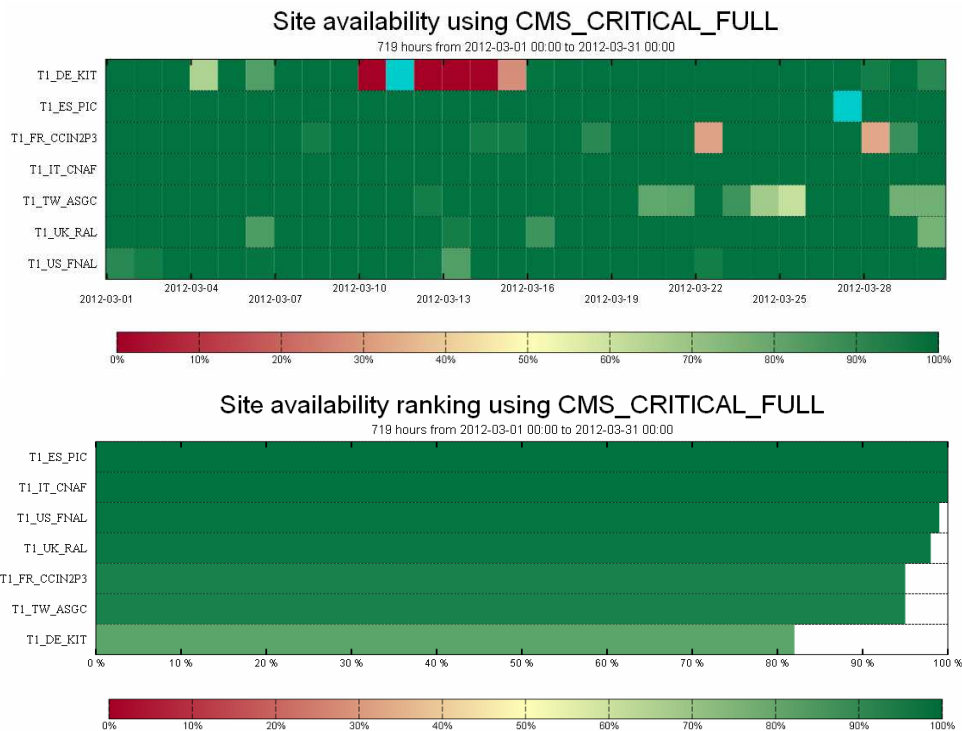
CMS has adopted the framework to run custom tests on the Computing Elements (CE) and Storage Resource Manager (SRM) instances at the sites. These tests allow CMS to determine, among other things, that:

- It is possible to send and run CMS 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 CMS data in and out of the local storage.

These tests are detailed in [8] and are run once per hour on all CE and SRM instances accessible by CMS.



**Figure 1.** “Latest View” of all CMS critical SAM tests run at a CMS T1 site.



**Figure 2.** Site Availability March 2012 history (top) and ranking (bottom) for CMS T1 sites tested.

A failure of any of these tests determines the "unavailability" of the service instance where the test ran. If all instances of a given service type in a site are unavailable, the service itself is considered unavailable. Finally, if either the CE or the SRM service is unavailable, the site itself is considered unavailable.

The site availability over a time interval is the fraction of that time interval when the site was available. An example of the daily evolution of this metric and ranking for CMS sites tested, for a month period, is shown in Figure 2.

## 2.2. Job load generator: the HammerCloud

Another complementary testing method consists of regularly submitting jobs similar to real analysis jobs. The difference with respect to the SAM tests is the fact that the statistics are much higher, ~250 jobs/(site×day), the fact that the accessed data can be spread on several disks, as it sizes 0.5 TB, and a higher load on the site storage system. Currently, CMS use the HammerCloud tool [9], developed by the CERN IT-ES group, to submit these realistic CMS jobs with CRAB [10], the CMS analysis job submission tool. 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.

At regular time intervals, a new analysis task is created for each site, to be run on a specific dataset. The task is then split into several jobs, which are submitted as a collection to the gLite WMS. Each job performs a trivial data analysis on a fraction of the dataset. All submitted jobs are classified as successful, as failed at the application level or as aborted at the Grid level.

The HammerCloud daily statistics are used to measure the success rate for each site. HammerCloud has two modes of operation: *functional tests*, a continuous low rate job submission, and *stress tests*, filling sites with jobs on demand. CMS is currently running functional tests, however the stress tests saturating all CMS slots at the sites can be used to compare the results to the resource pledges to CMS, to uncover possible bottlenecks or scaling problems in the site services, although this has not been tested yet.

## 2.3. Data transfer links

A site needs to have sufficient data transfer connections and bandwidth to other sites in order to perform CMS workflows. In the CMS computing model the reconstructed data is distributed between Tier-1 sites; for analysis based on reconstructed data, transfers from all Tier-1 to all Tier-2 sites are required. Transfers between Tier-1 sites at very high rate will be needed every time there is a global replication of reduced data, as a result of different reprocessing passes occurring at the Tier-1 sites. For uploading MC data produced at Tier-2 sites, typically the regional Tier-1 is used but transfers to other Tier-1 sites are also required. Recently, Tier-2 to Tier-2 connections to share reduced data became important as well. Hence, data is massively moved in the organized system for further processing and the data placement needs to work efficiently.

In 2007, a Debugging Data Transfers (DDT) task force was created to design and enforce a procedure to debug problematic links [11][12]. A clear certification procedure was set, using a traffic generator to test the quality of a link and considering a link to be commissioned when it demonstrates:

- for links with source at Tier-0 or Tier-1 sites, 20 MB/s averaged over 24 hours;
- for links with source at Tier-2 sites, 5 MB/s averaged over 24 hours.

Only links which are commissioned are used to move data for production usage. The DDT activities, from 2007 to 2009, were extremely useful and the data transfers increased in number and improved in quality. Today, the tests run on demand and managed by the sites.

At end of March 2012, and ignoring Tier-2↔Tier-2 connections, a total of 789 links are enabled for Production use in CMS:

- 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.

It is worth to mention that missing commissioned links are typically the ones not used for production (for example, some non-regional Tier-2→Tier-1 links). They are not failing the commissioning process; they are simply not being used at all.

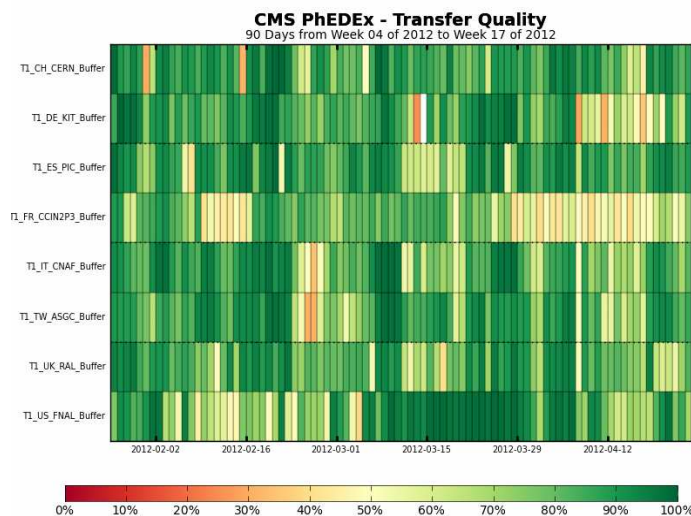
Moreover, all these ~800 production links are continuously exercised with test transfers at 0.25 MB/s/link (except T2-T2 crosslinks, which are used in CMS as well but not monitored in Site Readiness). Both commissioned links and transfer qualities are used in the Site Readiness. Adding Production transfers, we routinely have WAN transfers of ~200 TB/day (~1.5-2.5 GB/s), which is more than enough to detect systematic transfer problems, not only at the network level, but also in the data transfer services and the storage infrastructure. Figure 3 shows the transfer quality for all transfers from Tier-1 to Tier-2 sites occurring in a three month period on the first 2012 quarter. Occasionally, some periods in which Tier-1 sites had export problems to the Tier-2 sites are clearly seen.

#### 2.4. Scheduled Downtimes

When evaluating the reliability of a site, one needs to properly trace the scheduled downtimes of Grid services used by CMS. These downtimes are published in GOCDB/OIM and can concern individual service instances as well as the whole site. In a given day, a site is in a scheduled downtime from the point of view of CMS if either:

- The whole site had a scheduled downtime during the day;
- the CREAM-CE, or all the SRM instances used by CMS at the site, had a scheduled downtime during the day.

Currently, At Risk, Scheduled and Unscheduled downtimes are traced, for the CMS services or for the entire site. This downtime monitoring is embedded into the Site Status Board.



**Figure 3.** Daily average transfer quality on T1→T2 links for added Production/Test transfers.

## 2.5. The Site Status Board (SSB)

The Site Status Board (SSB) is a synoptic view of the status of all CMS computing sites [13]. It is designed to allow users to correlate the output of their workflows with known problems at sites, and to provide experts with a single entry point to the full suite of CMS monitoring tools. The provided information is often changed as the understanding of what is most relevant for making a good diagnosis of problems improves.

The SSB is a flexible presentation layer above a dynamic framework where information is stored in the columns of a database table, having the site name as key. These columns are filled by processes collecting data from the internal CMS dashboard database, the WLCG information system and ASCII files on the web. New columns can be freely added and defined via a web interface and grouped in “views”, or collections of columns. The time history of any column can be graphically displayed or retrieved in XML format.

Most of the *Site Readiness* plots and the entire results are kept and accessed from the SSB.

Site Name	Site SUM availability	HC	Commissioned Links (expand this column)	Good links combined (expand this column)	SiteReadiness Status	IsSiteInSiteDB?	Topology/Maintenances
TI_CH_CERN	98	100.0	1/5 combined	5/7 combined	R	true	1 SRMv2 . 6 CREAMCE .
TI_DE_KIT	94	99.5	2/5 combined	5/7 combined	R	true	1 SRMv2 . 3 CREAMCE .
TI_ES_PIC	100	99.9	2/5 combined	5/7 combined	R	true	1 SRMv2 . 4 CREAMCE .
TI_FR_CCINZP3	0	100.0	2/5 combined	5/7 combined	R	true	1 SRMv2 . 2 CREAMCE .
TI_IT_CNAF	100	99.3	2/5 combined	5/7 combined	R	true	1 SRMv2 . 8 CREAMCE .
TI_TW_ASGC	100	100.0	2/5 combined	5/7 combined	R	true	1 SRMv2 . 1 CREAMCE . 4 CE .
TI_UK_RAL	100	100.0	2/5 combined	5/7 combined	R	true	1 SRMv2 . 3 CREAMCE .
TI_US_FNAL	98	100.0	2/5 combined	5/7 combined	R	true	1 SRMv2 . 3 CE .

**Figure 4.** Site Status Board view used for CMS Site Readiness (Tier-1s shown).

## 3. Site Readiness evaluation

### 3.1. Site Readiness criteria

The quantities defined above must satisfy some constraints to consider the site as ready. Ideally, these constraints should be defined in such a way as to *a)* allow temporary glitches, *b)* enforce a reasonable level of reliability over a period of time and *c)* allow sites to quickly recover their ready status when problems are solved. In addition to that, downtimes due to scheduled maintenance and failures during weekends (for Tier-2 sites) should not be negatively considered in the site evaluation.

Currently, each day a site is evaluated as good (tagged as 'O', Ok) if the conditions in Table 1 are satisfied, and as bad if at least one metric is not satisfied (tagged as 'E', Error). Additionally, the scheduled downtimes are accounted as well (tagged as 'SD'). In order to take into account the stability of a site, a readiness daily status of a site is evaluated using the history of the last 7 days overall daily metrics (ignoring downtimes) and it is expressed by a flag with four possible values: *Ready (R)*, *Warning (W)*, *Not-Ready (NR)* and *Scheduled-Downtime (SD)*, which means respectively that the site is fully usable, that it is usable but suffering from temporary problems, that the site is unusable and that the site is under a maintenance period.

The transition rules between these states are shown in Table 2. Weekend daily failures do not negatively count for Tier-2 sites in the evaluation of the Site Readiness daily status. If the site is in maintenance, its daily readiness status is 'SD', regardless the last 7 days history of daily states. Note that the intermediate warning state gives sites reasonable time to recover.

Tier-1 sites	Tier-2 sites
Daily SAM availability $\geq 90\%$ Daily HammerCloud job success $\geq 90\%$ Link from Tier-0 commissioned Commissioned links to Tier-2 sites $\geq 20$ Commissioned links from/to other Tier-1 sites $\geq 4$ $\geq 50\%$ of links of a group with $\geq 50\%$ transfer quality	Daily SAM availability $\geq 80\%$ Daily HammerCloud job success $\geq 90\%$ Commissioned links to Tier-1 sites $\geq 2$ Commissioned links from Tier-1 sites $\geq 4$ $\geq 50\%$ of links of a group with $\geq 50\%$ transfer quality

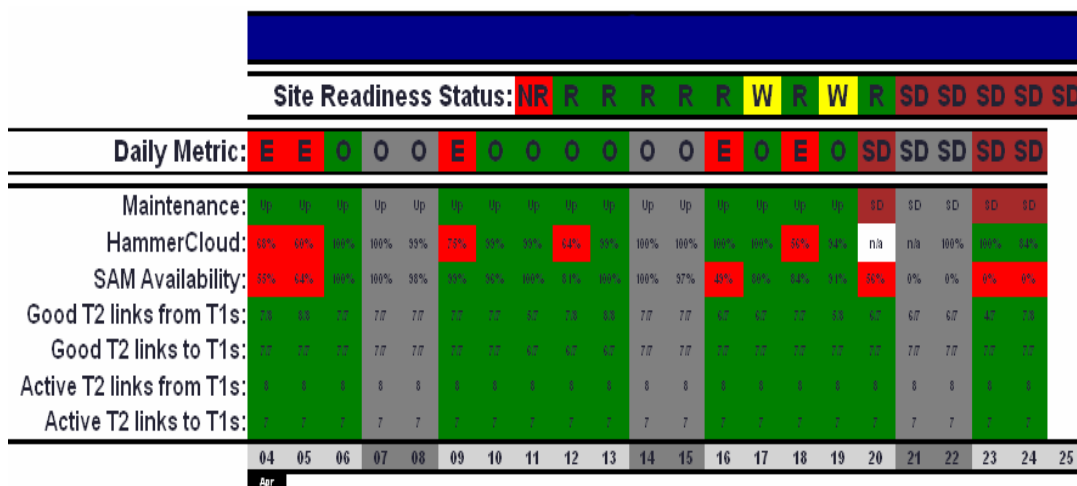
**Table 1.** Site Readiness daily metrics required for Tier-1 and Tier-2 CMS sites.

↓from/to→	NR	R	W
NR	-	O for last 2 days	-
R	E for >2 days over last 7	-	E but not E for >2 days over last 7
W	E for >2 days over last 7	O and not E for >2 days over last 7	-

**Table 2.** Site Readiness daily status: transition rules.

### 3.2. Site Readiness results

The Site Readiness program has been active since October 2008. It started as part of computing commissioning and turned into a regular tool used in Operations. Apart from consulting the results and status in the SSB, each site is provided with an overall picture of its last 15 days Site Readiness status and daily metric status, as well as the independent daily measurements. Figure 5 shows an example of such view for a Tier-2 at the end of April 2012.



**Figure 5.** Overall picture on Site Readiness results for a CMS Tier-2 site, at the end of April 2012.



Figure 6 shows the results of all the individual metrics for Tier-1s for the last year, i.e. from May 2011 to May 2012. The metrics related to transfers qualities are indicated by *GoodT1linksfromT0*, *GoodT1linksfromT1s*, *GoodT1linkstoT1s*, *GoodT1linksfromT2s*, and *GoodT1linkstoT2s*; the metrics related to enabled data transfer links (links used in production) are indicated by *T1ProdlinksfromT0*, *T1ProdlinksfromtoT1s*, and *T1ProdlinkstoT2s* (in both cases, the names point to how we group the links to evaluate the metrics); the metric related to the *Job 'load' generator* is shown, as well as the *Site Availability*. Note that there were changes last year and these are displayed on the graphs: *a*) almost the whole year CMS used the *JobRobot* tool, a specific CMS tool to send the load jobs, *b*) from 17/04/2012 CMS migrated to *HammerCloud*, and *c*) a migration from *SAM Nagios* to *SAM Nagios* using ACE component to evaluate availability occurred on 13/02/2012. Similar plots are provided in Figure 7 for Tier-2s, for the metrics which are evaluated on these sites.

For Tier-1s, all relevant data transfer links are enabled in production and there is a good transfer quality on those links, which is reflected on small quantity of failures for the *GoodT1links\** metrics. Most of the failures are related to the *Site Availability* or *Job 'load' generator* metrics, although the Tier-1s are shown to be pretty stable and usable, as the instabilities have small durations in time given the commitment and responsiveness of the Tier-1 sites to promptly solve issues and problems.

For Tier-2s we have a similar view on data transfer links enabled in production: almost all the sites have all their relevant links commissioned, except for one Tier-2 with poor network connectivity and new Tier-2s added to the system, which need to commission their data transfer links according to DDT rules (as clearly seen in December 2012 and in May 2012). The transfer quality is as well good, with best transfer qualities of links from Tier-2s to Tier-1s than the download of data from Tier-1s. Normally, the transfers from Tier-2s to Tier-1s are MC uploads, on top of the test transfers. The MC files are sitting on disk on Tier-2s and the upload goes more smoothly than downloading datasets from Tier-1s, for which some files may sit on tape systems, and the transfers could timeout until the files are on disk, ready for transferring. Additionally, the configurations of PhEDEx agents in Tier-1s are pretty more stable and sized. These effects could explain the slight difference, although the overall export/import quality is good, in average. The *Site Availability* has recently degraded due to some problems at Tier-2 sites, and this has translated to poorer readiness of Tier-2 sites, as we will see later.

Figure 8 shows an historical view of the number of Tier-2 sites in each readiness status, since the Site Readiness activities were set. A trend towards increasing numbers of good sites was evident, from 15 sites at the beginning of October 2008 to an average of 35 end March 2009. This commissioning period was extremely important to get the majority of the sites into good shape. In January 2010, a more restrictive metric was included for the data transfer quality links, and some Tier-2 sites were negatively affected. They rapidly corrected, adapted, and became “ready” again. Since then the number of Ready sites has stayed at about 40 sites (out of ~50). This has been the trend since we started data taking: of about 80% of Tier-2 sites were achieving the Site Readiness goals and considered Ready for undertaking analysis and production activities.

CMS does not black-list sites upon these measurements: the measured Site Readiness is used as a guide to flag sites as good or bad for running computing activities, like data reprocessing, MC simulation and analysis. For example, at Tier-1s, trends will be used to determine data placement and long term job routing, as the typical length of a workflow is of the order of days and the readiness of the site is a useful indicator to schedule the activities. To help the guide, the fraction of time a site has been stable and reliable is daily estimated based on the last 15-days history of Site Readiness daily status. Figure 9 shows the ranking plots based on this Site Readiness for Tier-1 and Tier-2 sites, on 10 May 2012. The goal is to chose those sites for operations above a certain threshold on Site Readiness: placed at 90% for Tier-1 and 80% for Tier-2 sites.

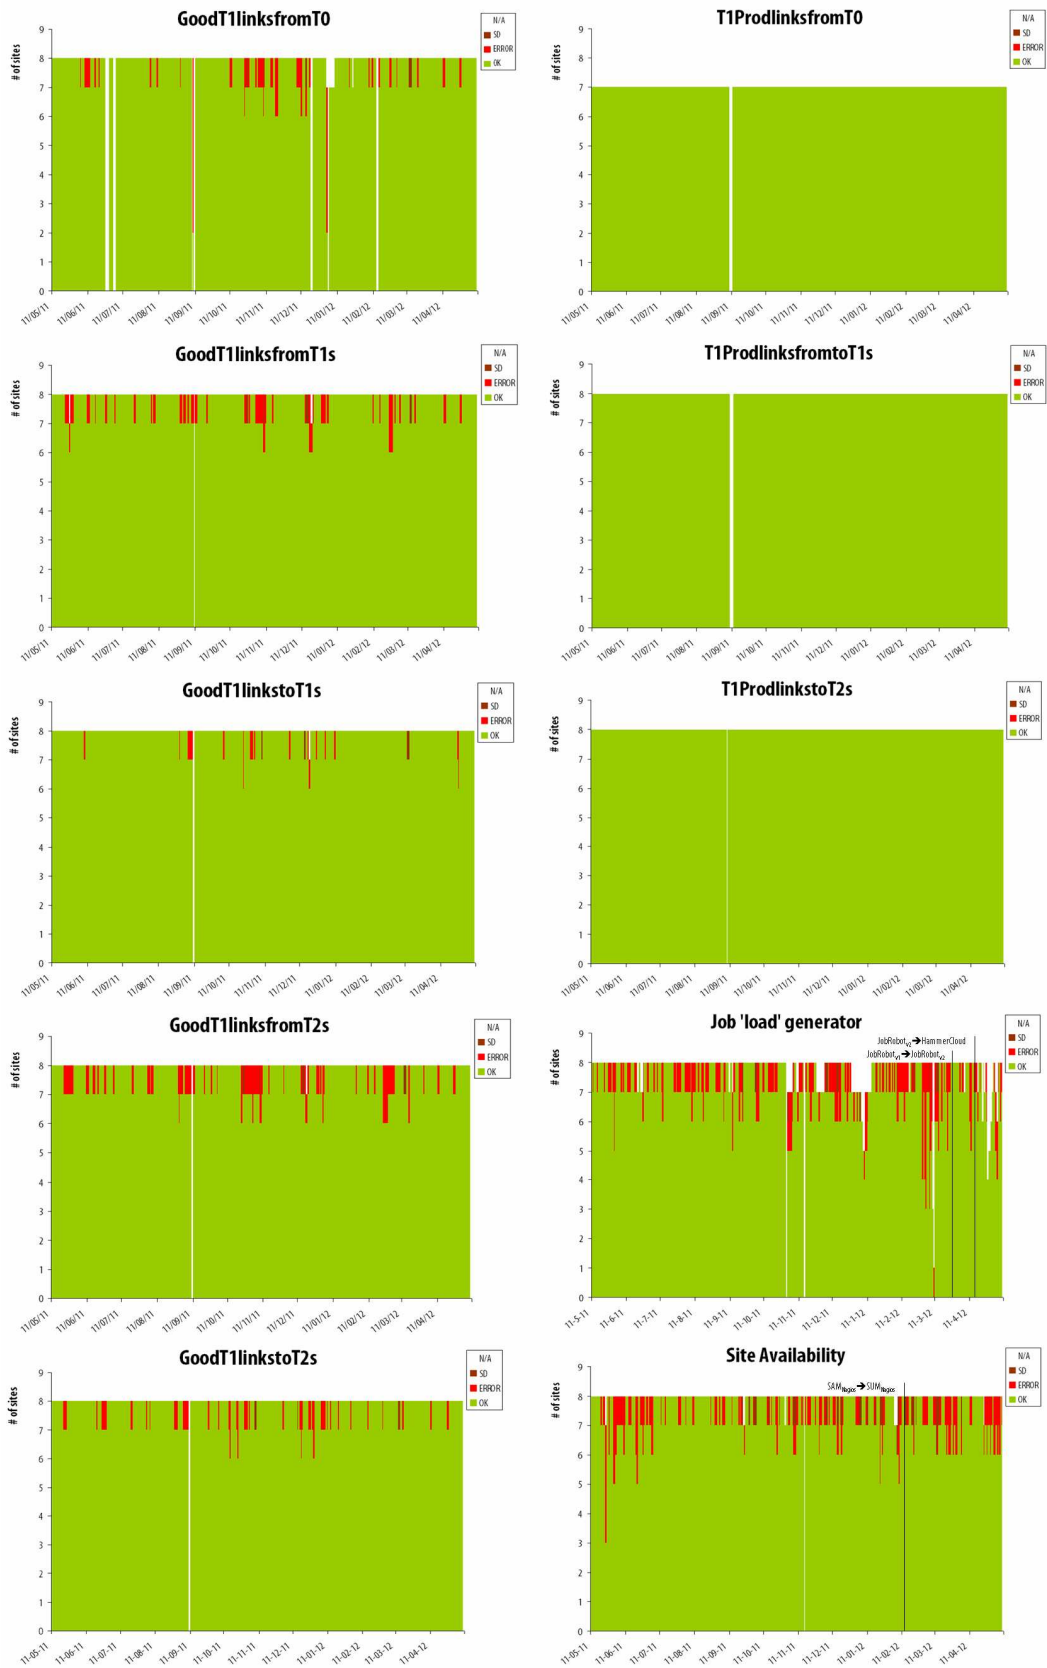
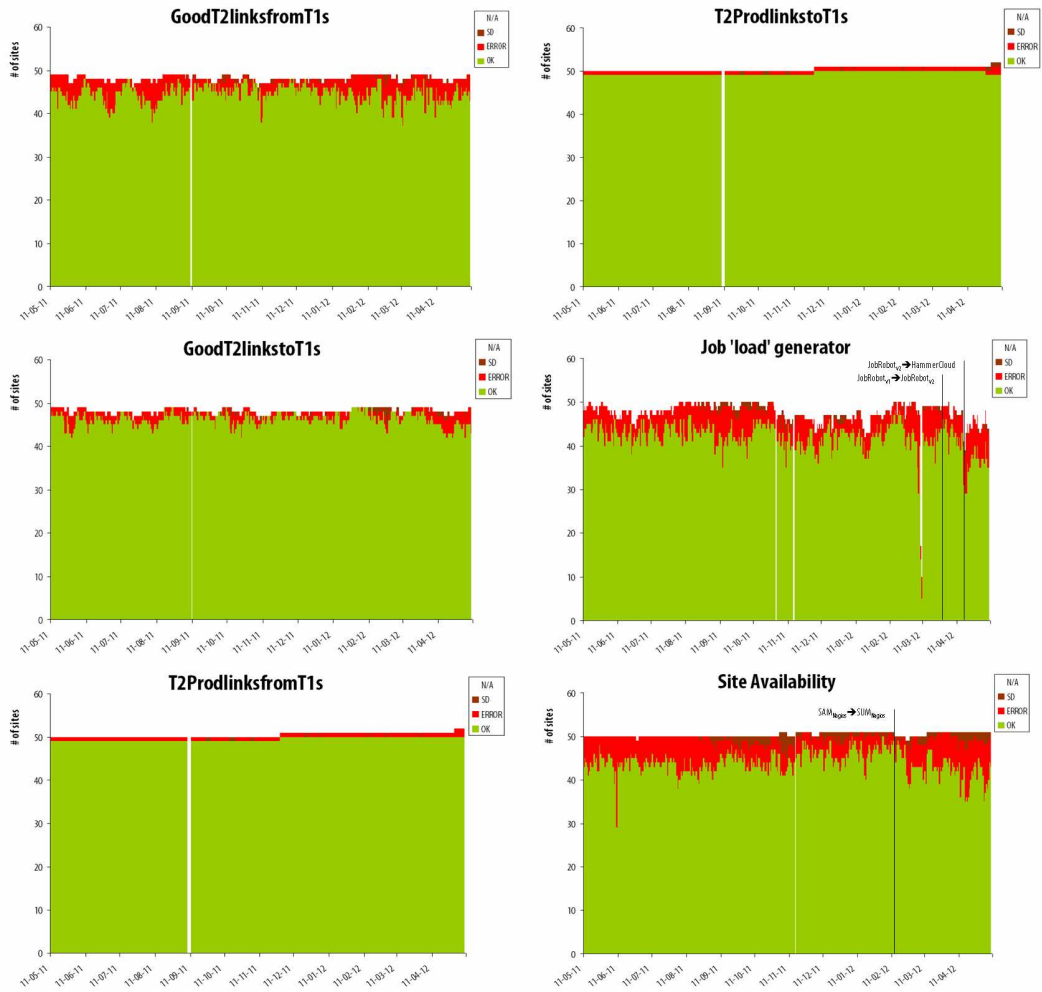
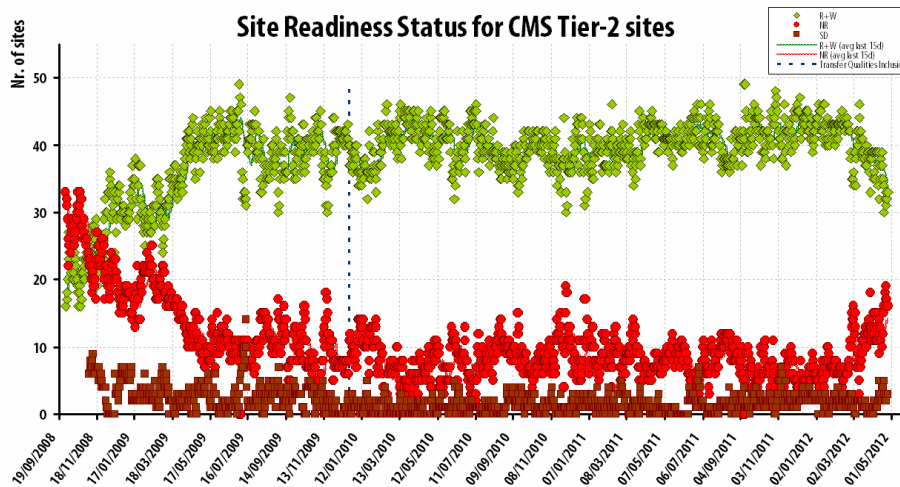


Figure 6. Daily metrics status for all metrics evaluated in Tier-1s, from May 2011 to May 2012.




**Figure 7.** Daily metrics status for all metrics evaluated in Tier-2s, from May 2011 to May 2012.



**Figure 8.** Evolution of Site Readiness status for CMS Tier-2 sites, since Site Commissioning started.

By means of this readiness evaluation, sites have an easy way to know if CMS is seeing problems at them. The program provides all kind of monitoring plots, XML feeds, widgets, and automatic alerts to stable sites prior being Not-Ready via Savannah ticketing system. Site performances according to the Site Readiness criteria are periodically reviewed in the CMS collaboration on a weekly basis.



**Figure 9.** Example of Site Readiness last 15 days ranking for Tier-1 and Tier-2 CMS sites.

#### 4. Next steps

The Site Readiness programme is quite mature by itself. The next steps include the inclusion of some monitoring plots directly into the Site Status Board, the ability to create more customized views of Site Readiness status for bigger time periods, moving *HammerCloud* to a pilot-based framework (CMS *glidein factory* [14]), and considering the inclusion of real production/analysis activities performances as a metric to be added to the existing ones.

#### 5. Conclusions

Site Readiness activities were important for commissioning the CMS sites prior data taking, and nowadays the tool is crucial to keep the CMS distributed computing system stable and reliable for daily operations. We have continuously monitored Grid and CMS services at sites for about 3.5 years with this tool. All the available information is condensed in a single estimator, whose value also takes into account the stability of the site. This helps production teams and users to select reliable CMS sites. A positive trend in reliability for Tier-1 and Tier-2 sites was observed, and the sites stayed at reasonable levels of readiness for the whole periods of LHC data taking. More metrics could be included, in the hope of providing the CMS sites with sufficient feedback and monitoring results so they can make their sites more reliable to CMS, and for CMS to use them in the most efficient way.

### Acknowledgments:

The authors wish to thank especially the Dashboard Team at CERN, for the fruitful collaboration towards a common tool development (SSB) for the benefit of the whole WLCG collaborators.

The Port d'Informació Científica (PIC) is maintained through a collaboration between the Generalitat de Catalunya, CIEMAT, IFAE and the Universitat Autònoma de Barcelona. This work was supported in part by grants FPA2007-66152-C02-01/02 and FPA2010-21816-C02-01/02 from the Ministerio de Educación y Ciencia, Spain.

### **References**

- [1] CMS Collaboration, The Compact Muon Solenoid Technical Proposal. CERN/LHCC 94-38, 1994.
- [2] Worldwide LHC Computing Grid. Technical Design Report. E-ref: <http://lcg.web.cern.ch/LCG/TDR/LCG TDR v1 04.pdf>
- [3] CMS Collaboration, "CMS Computing Project: Technical design report". CERN-LHCC 2005-023, 2005.
- [4] J.D. Shiers et al., "The (WLCG) Common Computing Readiness Challenge(s)", CCRC08, contribution N29-2, session Grid Computing. Nuclear Science Symposium, IEEE (Dresden), October 2008.
- [5] J. Flix et al., "The commissioning of CMS computing centres in the worldwide LHC computing Grid", contribution N29-5, session Grid Computing. Nuclear Science Symposium, IEEE (Dresden), October 2008.
- [6] J. Hernandez et al., Monitoring the Readiness and Utilization of the Distributed CMS Computing Facilities during the first year LHC running, J. Phys. Conf. Ser. 331 072020, 2011.
- [7] A. Duarte, P. Nyczyk, A. Retico, D. Vicinanza, "Monitoring the EGEE/WLCG Grid Services", J. Phys.: Conf. Ser.119 052014, 2008.
- [8] A. Sciabà et al., "New solutions for large scale functional tests in the WLCG infrastructure with SAM/Nagios: the experiments experience", in these Proceedings.
- [9] D. van der Ster et al., "Experience in Grid Site Testing for ATLAS, CMS and LHCb with HammerCloud", in these proceedings.
- [10] D. Spiga et al., "The CMS Remote Analysis Builder (CRAB)", LNCS vol. 4873, pp. 580-586, 2007.
- [11] G. Bagliesi et al., "The CMS Data Transfer Test Environment in Preparation for LHC Data Taking", contribution N67-2, session Applied Computing Techniques. Nuclear Science Symposium, IEEE (Dresden), October 2008.
- [12] J. Letts et al., "Debugging Data Transfers in CMS", CHEP09, Prague, Czech Republic, March 2009.
- [13] P. Saiz et al., "Collaborative development. Case study of the development of flexible monitoring applications", in these proceedings.
- [14] I. Sfiligoi et al., "The benefits and challenges of sharing glidein factory operations across nine time zones between OSG and CMS", in these proceedings