

# News on monitoring for CMS distributed computing operations

Andrea Sciabà

WLCG Collaboration Workshop  
May 19-20, 2012  
New York

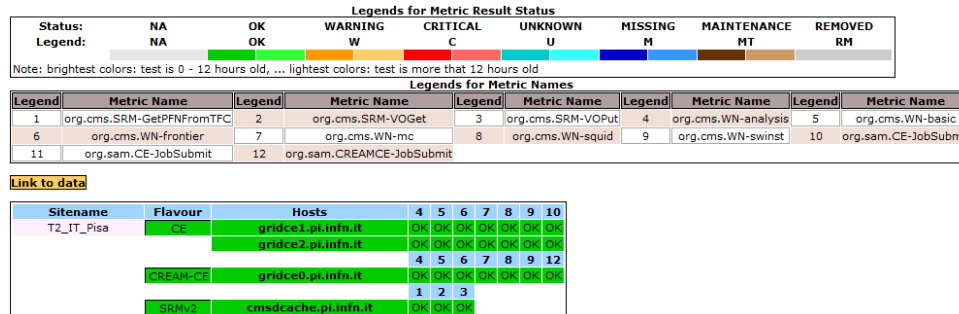
- Monitoring tools for sites
  - Tool overview: main uses and limitations
  - Recent changes
- Monitoring development highlights
- Conclusions




- Several tools available to all sites
  - Routinely used by sites to understand **what** CMS is doing and **how** it is going
  - Often complemented by local monitoring, e.g. for the batch system and the storage

Name	Description	Developed by
SUM <sup>(*)</sup>	SAM availability and test results	CERN IT
HammerCloud <sup>(*)</sup>	Functional tests	CERN IT
Site Readiness	Site quality estimator	CMS, CERN IT
Dashboard <sup>(*)</sup>	Job monitoring	CERN IT
Site Status Board <sup>(*)</sup>	Metrics aggregator	CERN IT
PhEDEx	Transfer monitoring	CMS, CERN IT
(WLCG Global Monitor) <sup>(*)</sup>	Transfer monitoring	CERN IT
Site Status Board <sup>(*)</sup>	Metrics overview	CERN IT
PerfSONAR Dashboard <sup>(*)</sup>	Network monitoring	BNL

- **SAM** is now a completely new framework based on **Nagios** (test submission), **MSG** (test publication) and **ACE** (Availability Computation Engine<sup>(\*)</sup>)
- **SUM** is a web **interface** to SAM tests and availability
- Is the site passing the SAM tests? If not, why?
  - Sensitive to a wide spectrum of site/CMS problems
  - Sometimes site configuration “optimized” for SAM ☹️



- Was the site available in the past few days? If not, why?
  - Sometimes sites become unavailable because some test does not run for > 24 h (due to priority issues) ☹️



(\*) <https://twiki.cern.ch/twiki/bin/view/LCG/ACE>

- Can the site run analysis-like jobs? If not, why?
  - Any issues in CMMSW installation, local read access to storage, publication to BDII, batch system configuration?
  - Sensitive to EMI WMS problems ☹

### Job errors summary

Select Site:  Select Regions:  Test ID:  Template ID:  Start time:  April 18, 2012, 11 a.m. End time:

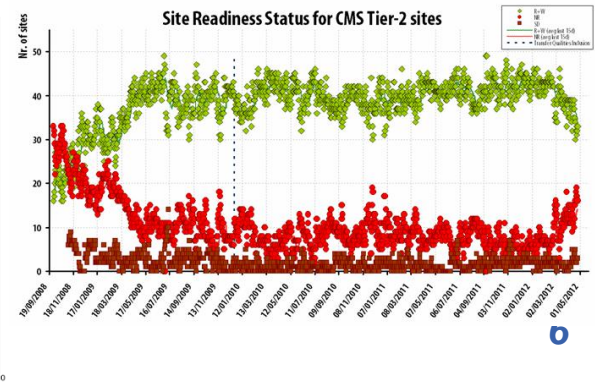

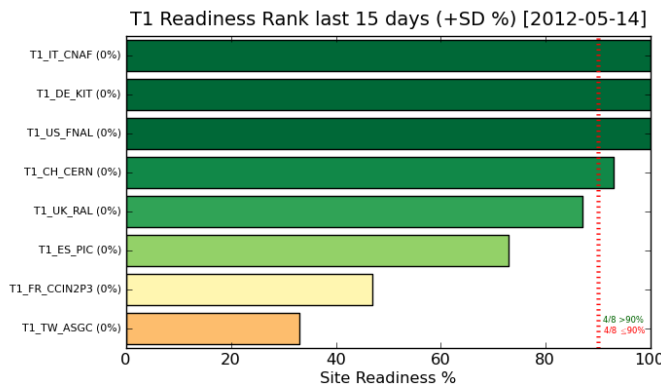
Site	Total jobs	Grid failed jobs (aborted)	Application failed jobs	Efficiency
T1_CH_CERN	144	8 »	0 »	0.944
T1_DE_KIT	211	32 »	0 »	0.848
T1_ES_PIC	573	1 »	0 »	0.998
T1_FR_CCIN2P3	84	0 »	0 »	1.000
T1_IT_CNAF	130	31 »	0 »	0.762
T1_TW_ASGC	129	1 »	0 »	0.992
T1_UK_RAL	102	0 »	0 »	1.000
T1_US_FNAL	353	0 »	0 »	1.000
T2_AT_Vienna	410	410 »	0 »	0.000
T2_BE_IHE	91	11 »	0 »	0.879
T2_BE_UCL	241	6 »	0 »	0.975
T2_BR_SPRACE	71	3 »	0 »	0.958
T2_BR_UERJ	0	0 »	0 »	0.000
T2_CH_CAF	0	0 »	0 »	0.000
T2_CH_CSCS	0	0 »	0 »	0.000
T2_CN_Beijing	0	0 »	0 »	0.000
T2_DE_DESY	2	1 »	0 »	0.500
T2_DE_RWTH	549	235 »	0 »	0.572
T2_EE_Estonia	134	55 »	50 »	0.216
T2_ES_CIEMAT	38	0 »	0 »	1.000
T2_ES_IFCA	443	4 »	0 »	0.991

- Overall, does the site have any problem? If yes, where?
  - SAM, HC, data transfers, downtimes

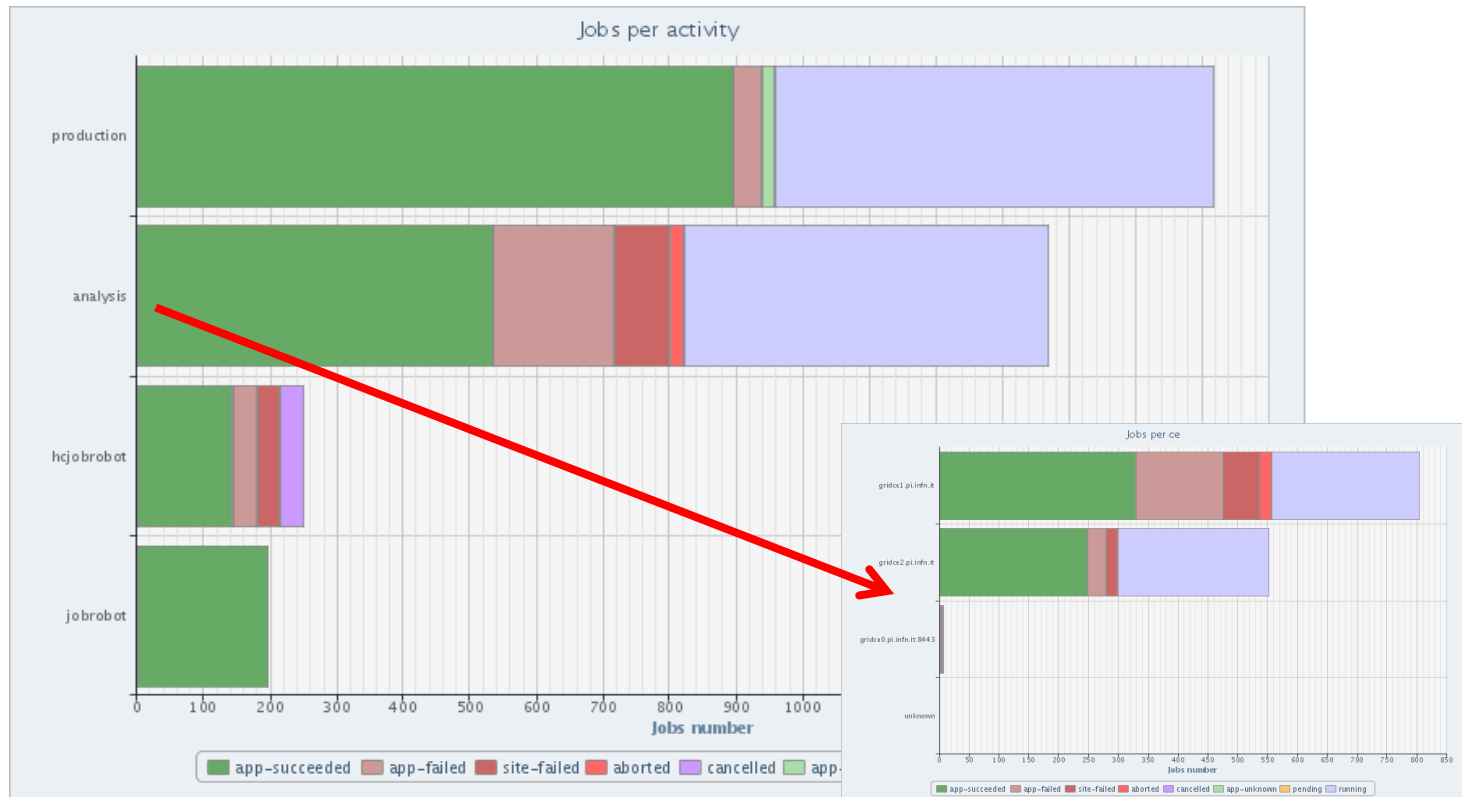
Site Readiness Status: R R R R R R R R R R R R W W R W NR																						
Daily Metric: O O O O O O O O O O O O O O O E E O E E																						
Maintenance:	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	UD	UD	Up				
Job Robot 'DB' (info):	97%	99%	99%	93%	100%	99%	100%	99%	99%	99%	100%	99%	99%	99%	99%	91%	96%	99%	97%			
HammerCloud:	94%	99%	99%	100%	98%	94%	100%	99%	99%	99%	96%	95%	97%	97%	94%	99%	62%	73%	80%	79%	79%	
SAM Availability:	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	
Good T2 links from T1s:	8/8	7/8	8/8	8/8	8/8	8/8	8/8	8/8	8/8	8/8	7/8	7/8	8/8	8/8	8/8	8/8	7/8	8/8	8/8	7/8	8/8	
Good T2 links to T1s:	7/8	7/8	7/8	7/8	7/7	7/7	8/8	7/7	7/7	7/7	7/7	7/7	7/7	7/7	7/7	7/7	7/7	7/7	7/7	7/7	7/7	
Active T2 links from T1s:	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	
Active T2 links to T1s:	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	
	25	26	27	28	29	30	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16
	Apr						May															

Report made on 2012-05-16 02:30:02 (UTC)

- How is it rated? How does it compare to other sites?
  - Sites are VERY sensitive to their ranking  $\Rightarrow$  strong incentive to be **"READY"**



- What is running on the site? Do jobs succeed? If not, on which CE do they fail? With which error?
  - Good for overview, not so good for complex troubleshooting: only Grid failure reasons and CMSSW exit codes available ☹️



- Is used as a convenient **entry point** to different metrics and other monitoring

Index Expanded Table

Show 200 entries Copy Print Save view: default Search...

Site Name	Visible	HammerCloud	SUM		Production	Analysis	Site usage		Commissioned Links (expand this column)	Under investigation	SiteIssues	In_rate_phedex	Out_rate_phedex
			SUM CE	SUM SRM			Running	Pending					
T2_IT_Pisa	OK	75.6	OK	OK	100%(2278)	98%(1130)	2230	268	2/5 combined	mark	info	3	18

Showing 1 to 1 of 1 entries DB query took 0.0601 s First Previous 1 Next Last

Shifter view

Index Expanded Table

Show 200 entries Copy Print Save view: commission Search...

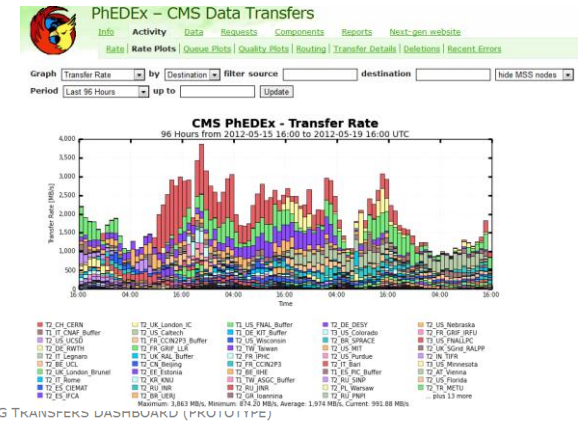
Site Name	SiteComm JR	JR (Dashboard)	Site SUM availability	HC	Commissioned Links (expand this column)	Good links combined (expand this column)	SiteReadiness Status	IsSiteInSiteDB?	TopologyMaintenances
T2_IT_Pisa	98	97.5	96	79.3	2/5 combined	2/7 combined	NR	true	1 SRMv2 1 CRAM CE 2 CE

Showing 1 to 1 of 1 entries DB query took 0.0187 s First Previous 1 Next Last


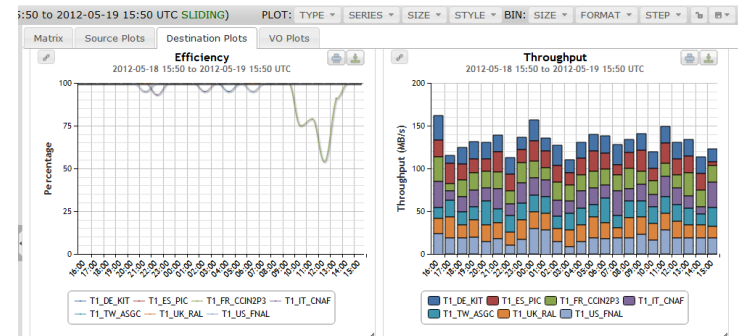
Site Readiness view



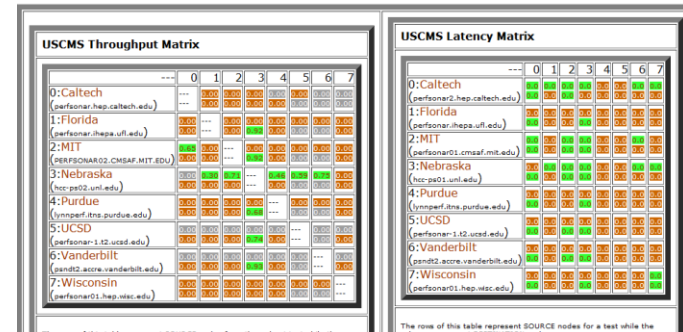
- **PhEDEx monitoring**
  - Shows status of transfers and links from a CMS point of view
  - Are transfers failing? Why?
  - What are the transfer rates?



- **WLCG Transfers Dashboard**
  - Based on information directly coming from FTS servers and sent to MSG system
  - VO-independent
  - Not yet in production but a working prototype exist



- **PerfSONAR Dashboard**
  - CMS US sites getting experience on it following ATLAS' example



- Old SAM framework decommissioned ⇒ **new SAM framework** in operation
  - Migration took a long time and a lot of work, but eventually very **smooth**
  - Nagios very **stable**, several **issues** in ACE and SUM found and **fixed** before going to production
- CMS Job Robot ⇒ **HammerCloud**
  - Many more **features**
  - Again, several bugs found and fixed (a few remain)
  - Sites not yet very familiar with it; FAQ being collected, tutorials organized

- WLCG: revive the SiteView web page as entry point from a site perspective of the VO activities
- CMS **Storage Accounting**
  - Ongoing development to produce interactive **summaries** of the **storage space utilization**
    - Sources are the standard WLCG **SE dumps**
    - Treemap-based, interactive visualisation
- New CMS **computing shift** page and new CMS **accounting** portal
  - Will be based on the CMS **Overview** framework
  - Including, but not limited to, site information
  - Development just started

- In CMS sites don't rely so much on central operations: they tend to look after themselves using common and custom tools
  - In theory, information is reasonably complete
  - in practice, see the survey results for site feedback!
- Recent convergence with ATLAS for what concerns SAM visualization and HammerCloud
- New monitoring systems being added and developed