



CMS Operations

[GDB – CERN, 14 Oct 09]

Oliver Gutsche, Markus Klute
[CMS Data Ops]

Daniele Bonacorsi, Peter Kreuzer
[CMS Facilities Ops]

on behalf of CMS Computing and all Data Operations and Facilities Operations teams

Special thanks to the SiteReadiness team (Andrea Sciabà and Josep Flix)



Operations project in CMS Computing

Data Operations (coords: Gutsche / Klute)

- ◆ responsible for central data processing and transfers: RAW data repacking and prompt reconstruction at T0, RAW data and MC re-reconstruction and skimming at T1's, MC production at T2's
- ◆ it will ensure central data consistency and data distribution to T0/T1s including custodial storage of primary datasets

Facilities Operations (coords: Kreuzer / Bonacorsi)

- ◆ Responsible of providing and maintain a working distributed computing fabric with a consistent working environment for Data Operations and Analysis users
- ◆ It involves coordination of facilities operation, resource management and liaison to external projects and organizations

Analysis Operations (NEW) (coords: Wurthwein, Belforte, D'Hondt)

- ◆ Responsible for central data placement at T2 level, CRAB server operations, validation, and support, and for metrics, monitoring and evaluation of the distributed analysis system



Operations meetings

Data Operations

- ◆ Weekly meetings, Monday afternoon, 3pm GVA time
 - Data Certification; Processing at CERN; MC production; Re-processing at T1's; Data Transfer and Integrity

Facilities Operations

- ◆ Weekly meetings, Monday afternoon, 5pm GVA time
 - Operational achievements and issues at Tier-0/1/2
 - Report by CMS contacts at T1's; overview of SAM and SiteReadiness status for all sites
 - CMSSW deployment status on EGEE and OSG
 - Invited talks on hot-issues at CMS Facilities issues (challenges, storage, monitoring, etc)
- ◆ Weekly (rotating real/virtual) Asian meetings and Russian/Turkish meetings
 - Minimal effort from sites needed: no reports asked
 - Maximum efforts come from CMS central Ops teams (reports, overview, ...)
 - Time is good for Asian colleagues, and central CMS Ops people
 - Other CMS people (depending on topics and availability) sometime attend

CMS attends WLCG Ops daily calls, 3pm GVA time

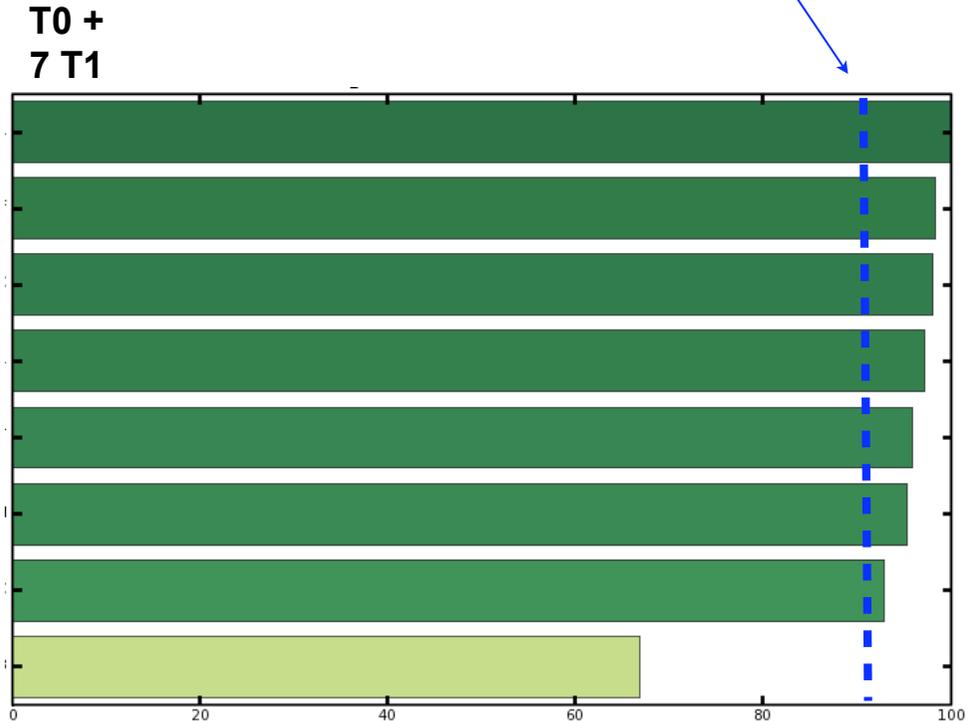
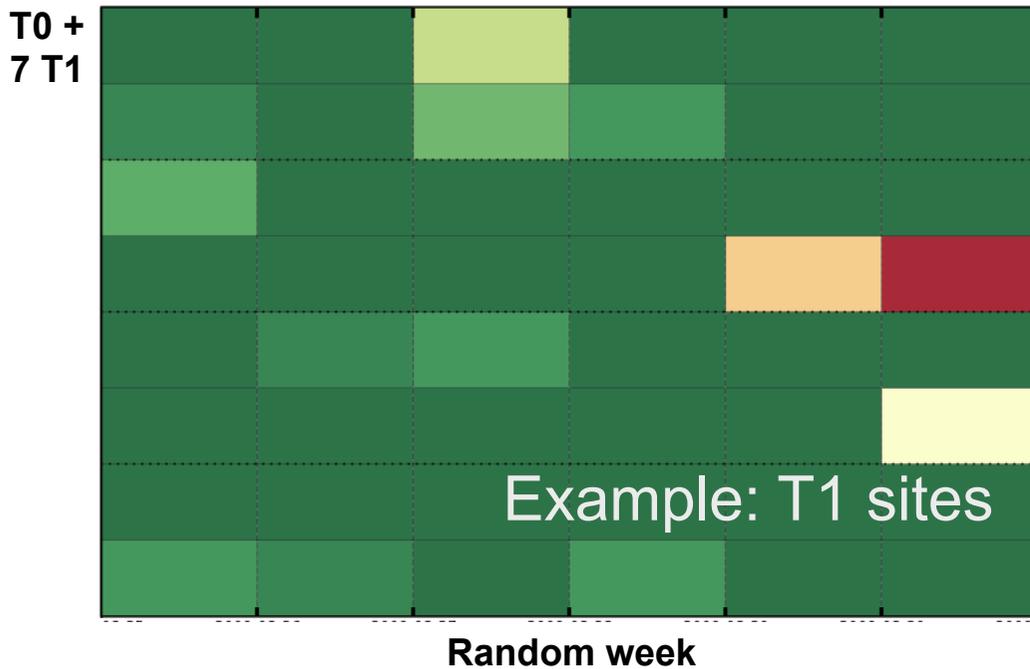
- ◆ twiki-based CMS reports since Feb09
 - https://twiki.cern.ch/twiki/bin/view/CMS/FacOps_WLCGdailyreports



Readiness of sites: CMS requirements on Tiers [1/4]

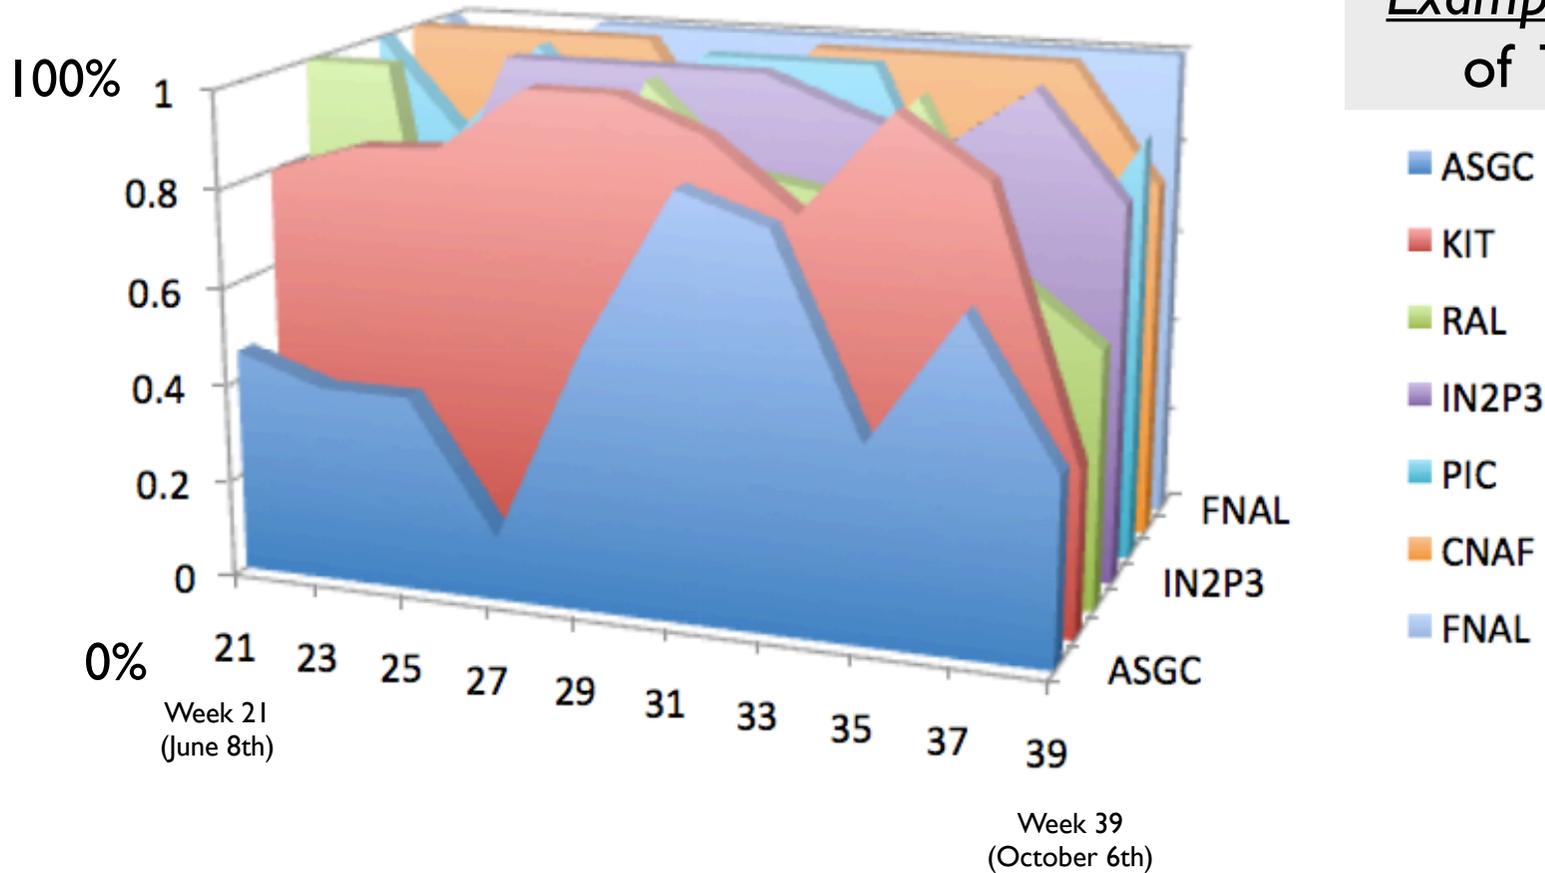
SAM availability for CMS Tiers

- ◆ CMS-specific SAM tests
 - Complementary to WLCG SAM, to mimic real CMS workflows
 - Widely documented elsewhere
- ◆ Overall SAM Availability ranking for CMS {T1,T2}'s: goal is {90%, 80%}
 - For all **orangish/redish** boxes we discuss at FacOps weekly meetings





Readiness of sites: CMS requirements on Tiers [3/4]



SiteReadiness goal for T1's: **90%**

Achieved averages in *Jun-Oct 2009*:

- ✦ {FNAL, CNAF} at {**99%**, **95%**}
- ✦ {PIC, IN2P3, KIT, RAL} at {**87%**, **86%**, **85%**, **73%**}
- ✦ ASGC at **50%**

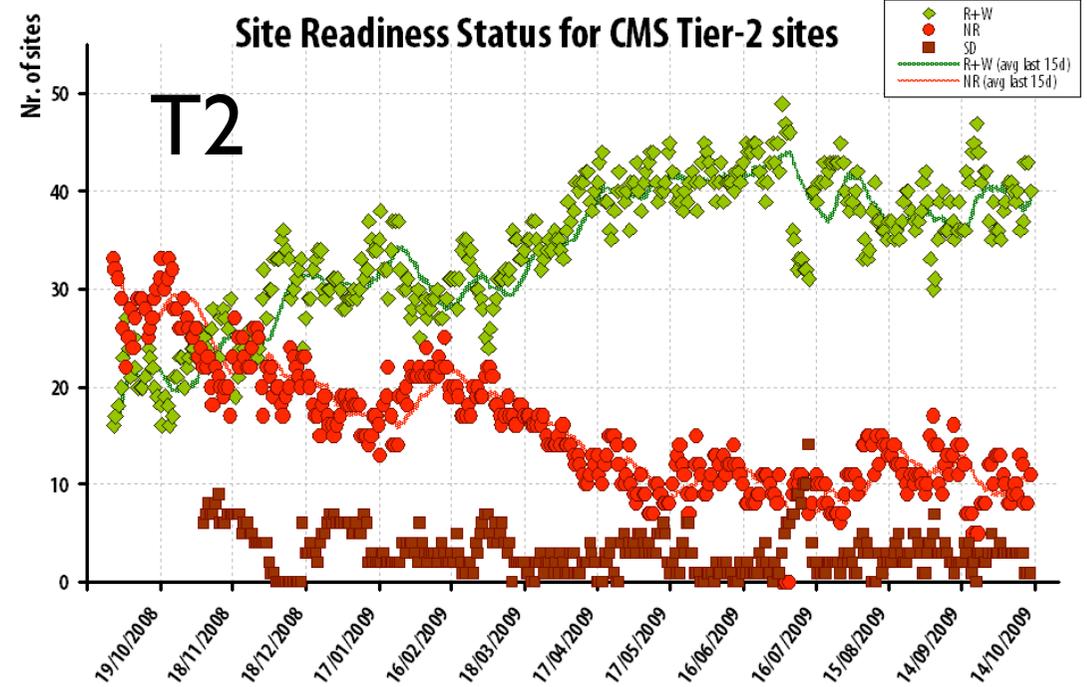
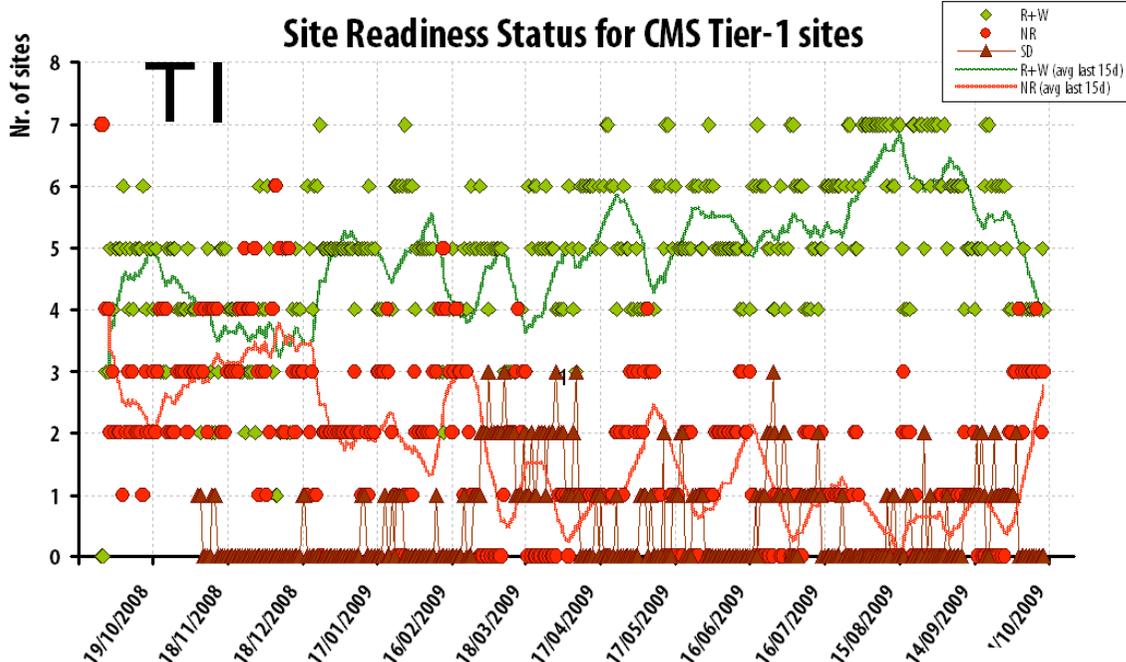
WLCG SAM (ops) not the full picture
CMS-specific SAM not the full picture
SiteReadiness (even!) not the full picture

Need high CPU eff, disk stability, MSS solidity and performance, ...



Readiness of sites: CMS requirements on Tiers [4/4]

Example: historical data on T1 and T2 sites





CMSSW deployment

CMSSW installed via Grid job on EGEE and OSG sites

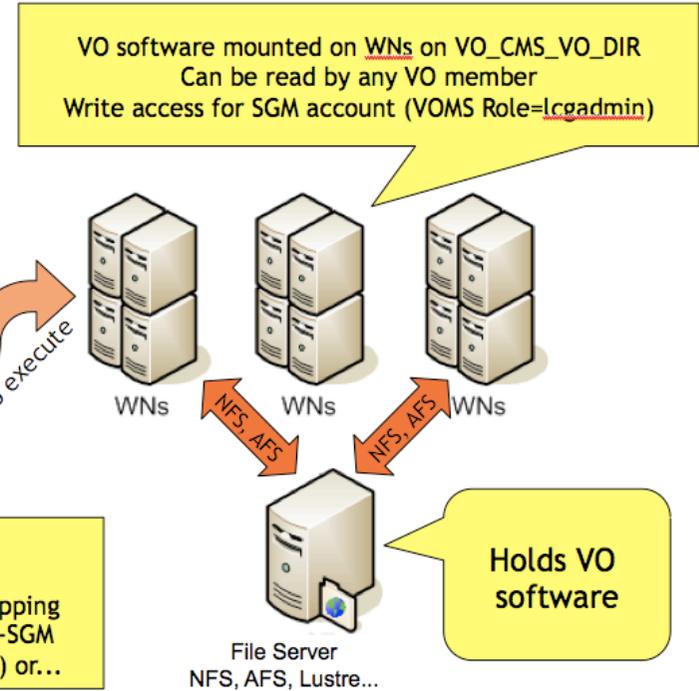
- Basic strategy: use RPM (with apt-get) in CMS SW area

CMSSW_314 deployment

- release announcement on Saturday Oct 3rd at 13h22
 - first installation jobs submitted at 13h41
- status: CMSSW_314 release deployed and ready for Oct-X start-up on da
 - EGEE: submitted to 51 Computing Elements (CE), 44 were DONE after few hrs (see plot); started to follow up on tails over the weekend already
 - OSG: release not tagged into the tag collector, so installed manually; smoothly and quickly completed in most OSG T2/T3



Credits:
 CMSSW deployment teams in Facilities Ops:
 Bockjoo Kim,
 Joris Maes,
 Lukas Vanelderden,
 Petra Van Mulders,
 Ilaria Villella,
 Christoph Wissing



On EGEE and OSG:

- CMSSW releases get routinely installed smoothly in most sites within few hrs from the release

OSG	Year 2009
Installed Releases (2_2_4 TO 3_3_0)	33
Total number of CEs	25
Total Installations	621
Total Removals	365

Credits: Bockjoo Kim



Sites responsibilities and CMS expectations [1/3]

INFRASTRUCTURE.

Example: T I sites

- ◆ CMS software server
- ◆ Frontier / Squids
- ◆ Local WN disk space (needed for caching input files - LazyDownload - and writing output), needs to be sufficiently dimensioned
 - CMS is working hard to restrict *all* file sizes to <10 GB and optimize workflows not to overflow WNs at the sites
- ◆ Correctly working TFC for CMS application and PhEDEx
- ◆ Role and cleanup of /store/unmerged
 - temporary small files, don't have to go to tape, should be cleaned up automatically by production systems, not perfect, need automatic cleanup by sites for files older than 30 days)
- ◆ FTS server for regional site support
- ◆ Data incoming from T0 should stay on disk for a period of time
 - in the first year all data coming from CERN should stay on disk.
- ◆ Deploy improvements
 - checksum verification of incoming files via PhEDEx may soon be in production



Sites responsibilities and CMS expectations [2/3]

Example: T I sites

OPERATIONAL.

- ◆ Pre-staging via different methods (tested in STEP'09)
- ◆ Tape family creation ('manual' communication with the sites).
 - In case of transfers into the sites: requests are not auto approved
 - additional safety net if tape families are not setup.
 - In case of output produced at sites, tape family creation is checked before workflows are started
- ◆ Transfer request approval within 24 hrs during business hrs
 - Does not work reliably at all times, if necessary this is overruled by central approval



Sites responsibilities and CMS expectations [3/3]

SUPPORT.

Example: T I sites

Facility and all central services: 24/7 (for T2's: working hrs, Mon-Fri)

- ◆ In case of alarm tickets, response within 1 hour
- ◆ All central services (we expect: CEs, SEs, MSS, batch system, WNs, Frontier/Squids, access to installed CMS software, ...)

CMS contact: nominally business hours (valid also for T2's)

- ◆ Contact has to be familiar with CMS workflows (valid also for T2's)
- ◆ Organize tape family setup
- ◆ Approve transfer requests (valid also for T2's)
- ◆ Regularly check PhEDEx logs, proactively take actions (valid also for T2's)
- ◆ Regularly run PhEDEx consistency tools to check for orphaned files, and inconsistencies in bookkeeping systems (valid also for T2's)
- ◆ Follow up on Savannah tickets and triage to facility if needed (valid also for T2's)
- ◆ Make sure that local hardware situation (available disk, tape) is close to pledges and update SiteDB regularly (valid also for T2's but less critical)
- ◆ Proactive information about site issues (valid also for T2's but less critical)



More on: **General vs exp-specific support issues**

The WLCG MoA's cover the general support by the sites

- ◆ Normally the exp-specific support is much more limited
- ◆ Impact on the site performance (as the site quality is measured on the overall system)

Higher level of integration is desirable

- ◆ Experiments have specificities that is impossible to avoid
- ◆ Nevertheless there are areas where the situation may be improved

Where WLCG can help?

- ◆ Common interfaces for monitoring
 - So that problems can be spotted by a non-experiment person
 - The Dashboard already helps a lot
- ◆ Enforce correctness of publication of resources in the BDII
 - So that exp people can more easily identify anomalies
 - Where are we with the implementation of:
 - https://twiki.cern.ch/twiki/pub/LCG/WLCGCommonComputingReadinessChallenges/WLCG_GlueSchemaUsage-1.8.pdf
- ◆ Guidelines and monitoring of authorization
 - Support for implementation of AuthZ in computing and data access using VOMS groups / roles
 - Monitoring of access rights (e.g. verify that normal users cannot delete experiment data)
- ◆ Tutorials by CERN-IT / WLCG :
 - How they would like CMS to correctly communicate issues and problems
 - Guidance of what would be GGUS Alarm ticket worthy would be appreciated by DataOps



Ticketing systems [1/2]

GGUS

Savannah

Item ID	Summary	Submitted On	Assigned To	Submitted By
#110213	Commissioning Muon T2-T2 links	2009-09-29 13:42	cmscompinfrasup-ddt	afanfani
#110216	WMS setup for cms Role=priorityuser and pool account limit	2009-09-29 16:25	cmscompinfrasup-glitewms	afanfani
#110274	October exercise: Fix required on Padova CE (deny /cms/Role=priorityuser) to avoid running on Padova instead of Legnaro	2009-10-02 09:43	cmscompinfrasup-t2itlegnaro	afanfani
#110291	wms202 and wms218 at CERN show problem in proxy delegation	2009-10-04 11:19	cmscompinfrasup-glitewms	afanfani
#110309	error with working dir, nomore priorityuser mapping, accessing a datasetat Purdue	2009-10-05 09:34	cmscompinfrasup-t2uspurdue	afanfani
#110386	Aborted jobs at Padova for users accessing dataset in Legnaro	2009-10-07 15:59	cmscompinfrasup-t2itlegnaro	afanfani
#110539	Grid Aborted at T2_FR_CCIN2P3	2009-10-13 09:47	cmscompinfrasup-t2frcin2p3	afanfani
#110555	failures accessing MC files (8001/8020) for instance ppMuX/Summer09-MC_31X_V3_SD_DoubleMu-v1/GEN-SIM-RECO	2009-10-13 15:13	cmscompinfrasup-t2uspurdue	afanfani
#110517	Some CMSSW version published but not present	2009-10-12 16:57	cmscompinfrasup-cmsswdeploy	aholguin
#110518	SRMV2 for T3_CO_Uniandes	2009-10-12 17:00	cmscompinfrasup-sam	aholguin
#110550	Problem with CC-IN2P3 AF	2009-10-13 13:27	cmscompinfrasup-t2frcin2p3	andriusj

GGUS

- ◆ Long tradition of the standard Global Grid User Support system
- Reaches the WLCG site-admins and the fabric-level experts

Savannah

- ◆ Problem tracking, troubleshooting reference, statistics, ...
- Reaches 'squads' easy to define: CMS contacts at Tiers, tools/services experts, ...
- More: baseline tool for Offline Computing shifts, integrated with other CMS projects, ...



Ticketing systems [2/2]

Wouldn't a single ticketing system be preferable?

- ◆ Of course. BUT: is there one with all the features CMS uses for Ops?

CMS requested a Savannah-to-GGUS bridging

- ◆ Work finalized. Now ready to be used. Start soon to gain experience in Ops
 - Thanks to Guenter Grein (GGUS), Yves Perrin (LCG/SPI) and Simon Metson (CMS) for their great efforts in the technical implementation and testing

- Main features of the tool
 - those sites who want GGUS should add to their "site-squad" the "GGUS Robot" account.
 - For these sites, when a new Savannah ticket is opened, e.g. by the CMS Computing shifter, the bridging will be activated (It can be manually turned off via Savannah option "Use GGUS == NO")
 - The GGUS ticket is opened by the GGUS Robot. The GGUS cross-reference URL and a "GGUS on hold" status are automatically set in the Savannah ticket
 - Once the GGUS ticket is closed, the Savannah ticket is automatically closed

To Do:
- Provide instructions to CMS sites and for CMS Operators

Note:
- GGUS Alarm/Team tickets treated separately



CMS SL5 migration status

- CMS sites were encouraged to migrate their Worker Nodes to SL5 **after STEP09**, based on slc4_ia32_gcc345 production architecture compatible with both SL4 and SL5
- 2 **unexpected problems** encountered during migration :
 - dCache/srmcp issue with GLite 3.2 and SLC5
 - Globus library issue with DPM / rfio**→ (Temporary) workarounds found** for both issues
- **WN migration status :**

	Sites migrated Today	Expected end-Oct 09
CMS Tier-0	T0 nodes : 238 / 400 (60%) CAF nodes : 0 / 94 (0%)	T0 : 100% CAF : 100% (mid-Nov 09)
CMS Tier-1s	7 / 7 (*)	7 / 7
CMS Tier-2s	37 / 50 (74%)	46 / 50 (92%)

(*) Multiple VO Tier-1s essentially migrated the CMS share, while keeping a variable fraction on SL4 for other VOs

- **Phase II** : CMS can now run proper native-SLC5 integration builds. Once a large majority of sites will have WNs ready, will **switch to native slc5_ia32_gccXXX builds**, which will only run on SLC5 machines, of course. During that time CMS will have for some time a number of legacy slc4_ia32_gcc345 releases to support
- In parallel (now !), CMS started to **validate all DMWM services** (PhEDEx, DBS, ...) on SLC5
- Estimated time scale for above 2 bullets : **Nov-Dec 09**
- In order to be on the safe side, CMS has requested to CERN-IT to keep default **interactive resources** to SLC4 until end-2009



CMS Computing shifts

Started in Fall 2008 with 16/7 coverage

- ✦ paused during Winter/Spring 2009
- ✦ resumed in August 2009, with 24/7 coverage (3 timezones)

Currently CMS Offline/Computing shift crew pool

- ✦ 35 people in 3 timezones, mainly non-computing experts

Shift roles

- ✦ Computing Shift Personnel (CSP) monitor systems and raise alarms
 - contributes to the standard MoA service work defined by CMS, accounted as the Online Central CMS shifts
- ✦ Other roles support the CSP: see box

Shift procedures and checklists

- ✦ <https://twiki.cern.ch/twiki/bin/view/CMS/ComputingShifts>

More CSP tools/info

- ✦ Stable EVO in ad-hoc virtual room: <http://evo.caltech.edu/>
- ✦ Shift Sign-Up tool: <http://tinyurl.com/r78zsb>
- ✦ IM under generic “FacOpsShifter” account (instructions on a twiki)
- ✦ Computing Plan of the Day: <http://cmsdoc.cern.ch/cmscc/shift/today.jsp>
- ✦ Account in the CSP E-log: <https://prod-grid-logger.cern.ch/elog/>
- ✦ Savannah: <https://savannah.cern.ch/projects/cmscompinfrasup/>

Computing Run Coordinator (CRC)

- Subscribes to all CSP E-log sub-sections
- Assists CSP in raising alarms/tickets for complex cases
- Calls EOC during off-working hours (see below)

Core Computing Operator (FacOps, DataOps Experts)

- Subscribes to relevant CSP E-Log sub-sections
- Supports CSP during working hours

Computing Expert On Call (EOC)

- Responsible of a particular service
- Alarmed by CSP via Email/IM/Tel during working hours
- Alarmed by CRC if really needed off-working hours

CMS Site Contact Person

- Responds to alarms (e.g. Savannah, GGUS tickets)

Other shifters (DQM, Online, Detector, ...)

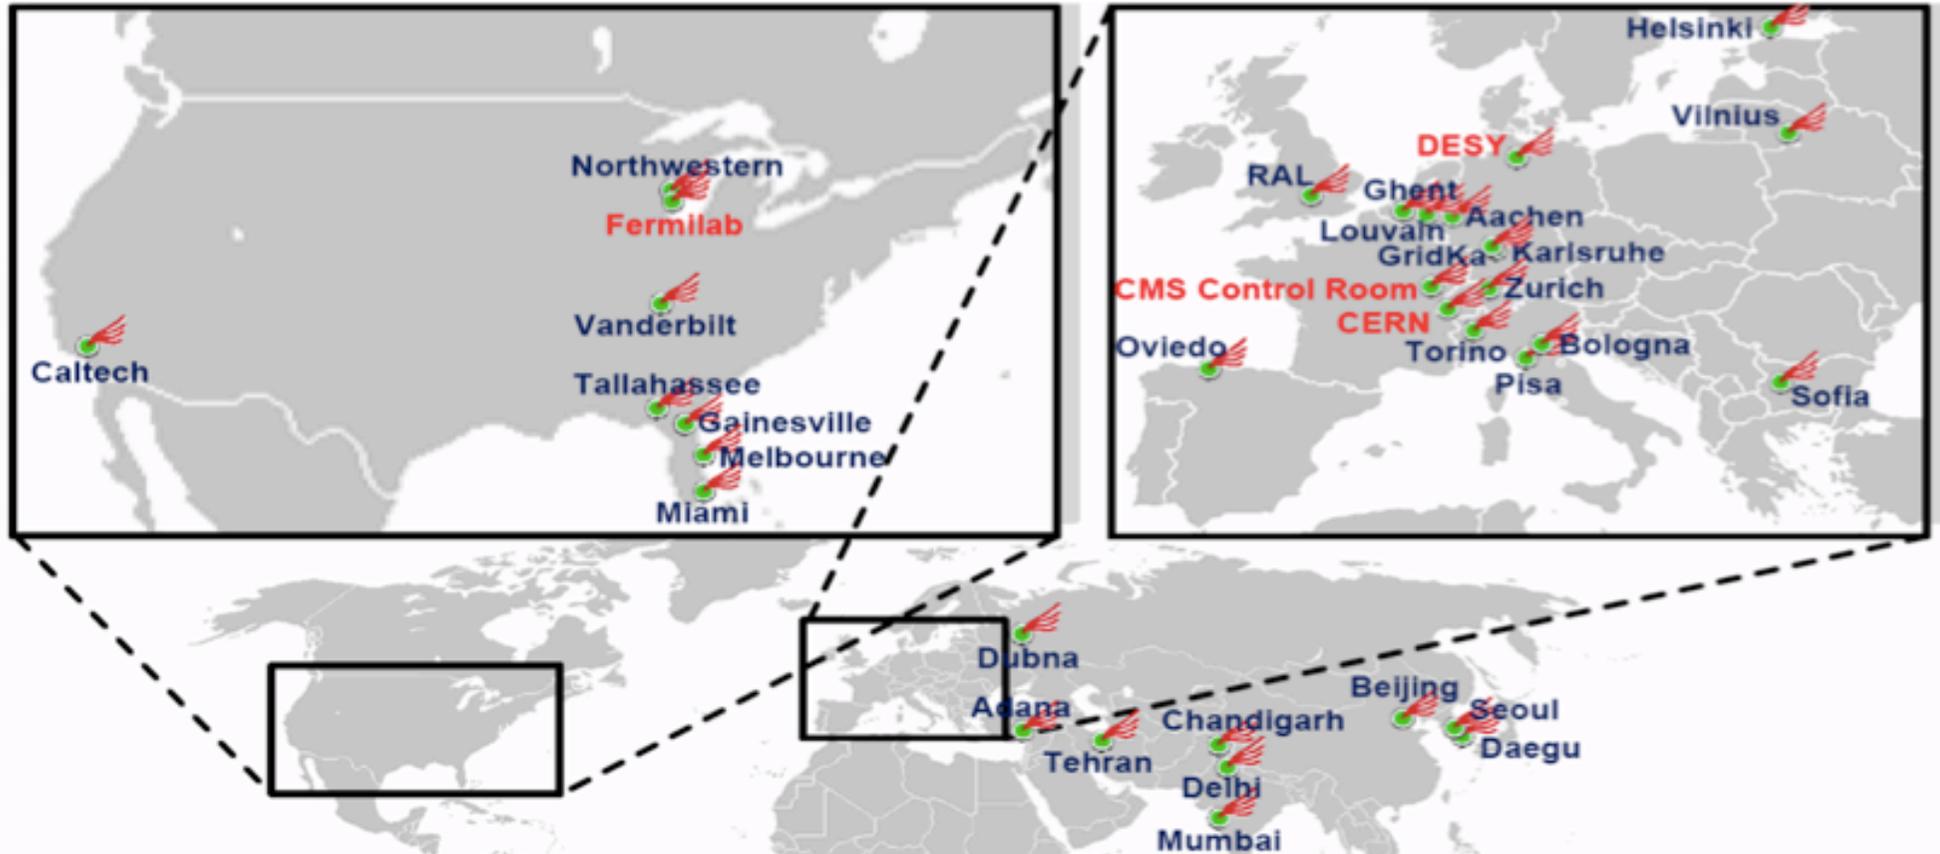
- In temporary absence of CRC, the CSP is the [Core Computing contact](#) for any shifter at P5/CMS Center/FNAL ROC

CSP procedure responsible

- Assigns CSP shifts



CMS Computing shifts and CMS Centres



More details on Computing Centers here :

<http://lucas-nice.web.cern.ch/lucas-nice/cms-centre/www/Publications/CHEP07-paper/CHEP07-LucasTaylor-CMSCentre-FinalPaper.pdf>

- 3 “**Primary CMS Centers**” : P5 Control Room, CMS Center CERN, FNAL ROC
- Many existing or upcoming “**Secondary CMS Centers**”
- Primary Centers permanently connected via **Tandberg Video** system
- Secondary centres using **permanent EVO** room
- Both primary and secondary centres used for Offline Computing Shifts



Critical services and GridMap

Computing systems are built on a large number of services

WLCG needs to

- ✦ periodically evaluate the readiness of services
- ✦ have monitoring data to assess the service availability and reliability

CMS needs to:

- ✦ provide a list of “critical services”
- ✦ Define (and keep up-to-date!) the “ranking” from 1 to 10
 - Matching of ranking values to call-outs, support actions, etc

CMS gives a special meaning to all ranking values

All services are monitored during Computing shifts

- ✦ GridMap as the interface

Rank	Definition	Max. downtime per incident (Hrs)	Comment
11	CMS Stops operating	0.5	Not covered (yet) here
10	CMS stops transferring data form Cessy		Cessy output buffer time
9	T0 Production stops		min(T0 input buffer/CESSY output buffer) or defined time to catch up
8	T1/T2 Production/analysis stops		defined time to catch up
7	Services critical when needed but not needed all the time (currently includes documentation)	0.5	
6	A service monitoring or documenting a critical service	8	
5	CMS development stops if service unavailable	24	
4	CMS development at CERN stops if service unavailable	24	
3	Services not critical for CMS	24	
2	Services required for CMS	72	
1	Used by a significant fraction of CMS	72	
0	Not used or discouraged by CMS	forever	

Rank 10: 24x7 on call
Rank 8,9: expert call-out

