

CMS Operations Report

8 December 2010
GDB




Operations Items

- ▶ Interesting Operations Items Since last we met
 - ▶ The completion of the Heavy Ion Run
 - ▶ Reprocessing Plans
 - ▶ Continued increases in analysis


HI Run

- ▶ Since our last report CMS completed the Heavy Ion Run
- ▶ Due to resource constraints this is a different workflow that pp running
- ▶ Ran smoothly for the month



Statistics



- ▶ HIAAllPhysics RAW: 882.2TB, 129132259 events (7.2MB/event)
 - ▶ checking single good runs it's something like 11.2 MB
- ▶ HICorePhysics RAW: 85.3TB 8621432 events



- ▶ Event size is large, but lumi section of 23s is the same
 - ▶ Average file size in CMS grows substantially

Processing

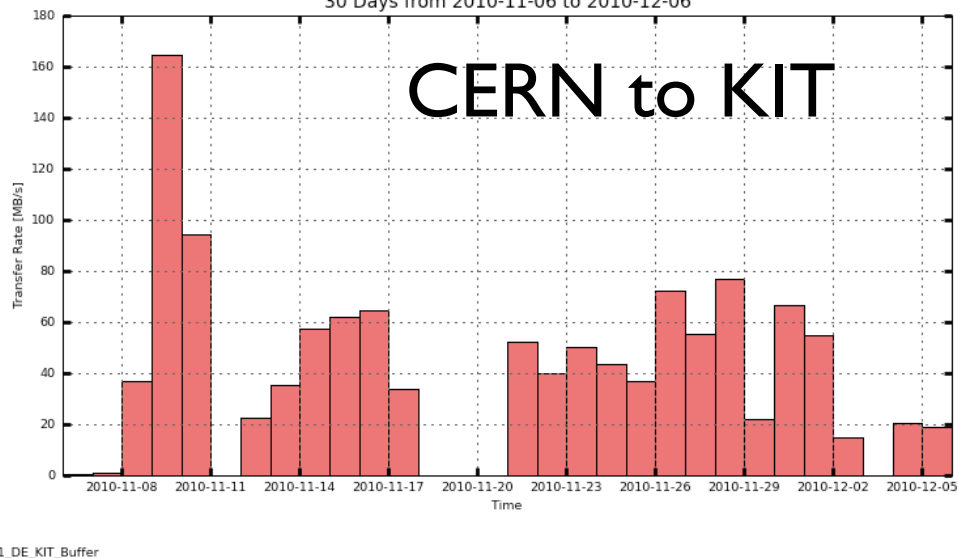
- ▶ Tier-0 was able to keep up with 100% of the HI events
- ▶ IO very high due to the large events
- ▶ Performance of CMSSW similar in HI events
- ▶ Great performance by Castor




Transfers

- ▶ 10% of the most central data was transferred to KIT and then served to Tier-2 sites

CMS PhEDEx - Transfer Rate
30 Days from 2010-11-06 to 2010-12-06





CMS PhEDEx - Transfer Rate
30 Days from 2010-11-09 to 2010-12-09



Maximum: 246.70 MB/s, Minimum: 7.22 MB/s, Average: 89.87 MB/s, Current: 36.32 MB/s

Remote Site Readiness

- ▶ Facility Ops monitors T1s and T2s.
- ▶ Site readiness is the “AND” of several experiment and WLCG monitored tests.
- ▶ Hard to be ready stably
- ▶ Tier-1s are doing a good job, but room for improvement
- ▶ Tier-2s do not have the same expectations of uptime, but about 20% of the sites are in a not ready state




Reconstruction

Pass	Identifier	Size [TiB]	Status
1	FIRSTCOLL	0.146	DELETED
2	GR09_P_V7	0.731	DELETED
3	Dec9	1.540	DELETED
4	Dec14	7.952	DELETED
5	Dec19	20.812	
<hr/>			
6	Jan23	10.613	DELETED
7	Jan29	13.347	
8	Feb9	11.561	DELETED
9	Mar1	0.694	DELETED
10	Apr1	1.626	
11	Apr20	26.180	
12	May6	70.836	
13	May27	249.313	
14	Jun9	170.321	
15	Jun14	181.600	
16	Jul6	7.588	
17	Jul15	11.341	
18	Jul16	4.490	
19	Jul23	9.000	
20	Jul23	9.000	
21	Jul26	3.307	
22	Sep17	327.174	
23	Nov4	466.858	
		1606.028	

2009



2010

Re-reconstruction passes 2010




Reconstruction

- ▶ CMS has been able to reconstruct the data many more times in the first year than was planned.
- ▶ About 17 passes of the data were made at various points
- ▶ The reconstruction of the full proton-proton data is about 10 days and another week of debugging and postmortem
- ▶ Good CPU efficiency for reprocessing





Maximum: 13.872 , Minimum: 0.00 , Average: 7.725 , Current: 13.872



Reconstruction

- ▶ Nice balance of jobs across the Tier-1s during Nov. data processing





- Total cpueff
- DE-KIT cpueff
- IT-INFN-CNAF cpueff
- TW-ASGC cpueff
- US-T1-BNL cpueff
- CA-TRIUMF cpueff
- ES-PIC cpueff
- NDGF cpueff
- UK-T1-RAL cpueff
- CH-CERN cpueff
- FR-CCIN2P3 cpueff
- NL-T1 cpueff
- US-FNAL-CMS cpueff

- ▶ CPU Efficiency is within expectations
- ▶ Improved with 3_8 IO improvements

Data Transfer

- ▶ Nearly 5PB transferred from CERN and other Tier-1s to Tier-1s
- ▶ Smooth running



Site	Share
TI_DE_KIT	0.11
TI_ES_PIC	0.06
TI_FR_IN2P3	0.19
TI_IT_CNAF	0.09
TI_TW_ASGC	0.08
TI_UK_RAL	0.08
TI_US_FNAL	0.39



Total: 4,858 TB, Average Rate: 0.00 TB/s



Tier-2 Traffic

- ▶ CMS is able to replicate data from most Tier-1s to most Tier-2s
- ▶ The majority of links between Tier-2s are also tested
- ▶ Full mesh allows flexibility and speed for replicating samples.



Tier-2 Hosted Data For Analysis



- ▶ Nearly 7PB of official samples hosted at Tier-2s sites for analysis
- ▶ Largest single group managed by Analysis Operations, but space also managed by local sites and physics groups
- ▶ Need to work in 2011 on how we allocate space
- ▶ Big variation on how much data is accessed
- ▶ Migration to AOD in 2011 for central space





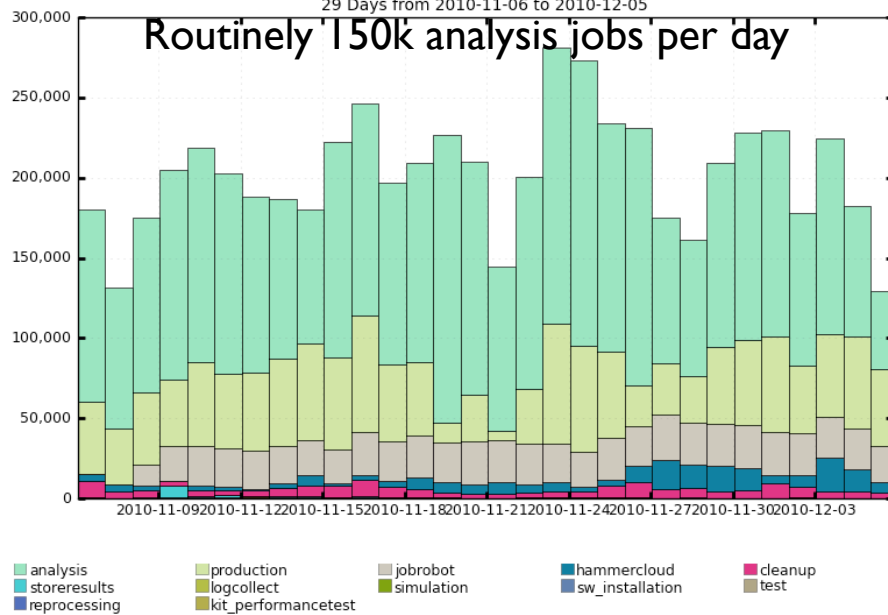
CMS Analysis Activity

Mail volume (#messages) handled by Analysis Operations on the CrabFeedback forum in 2010



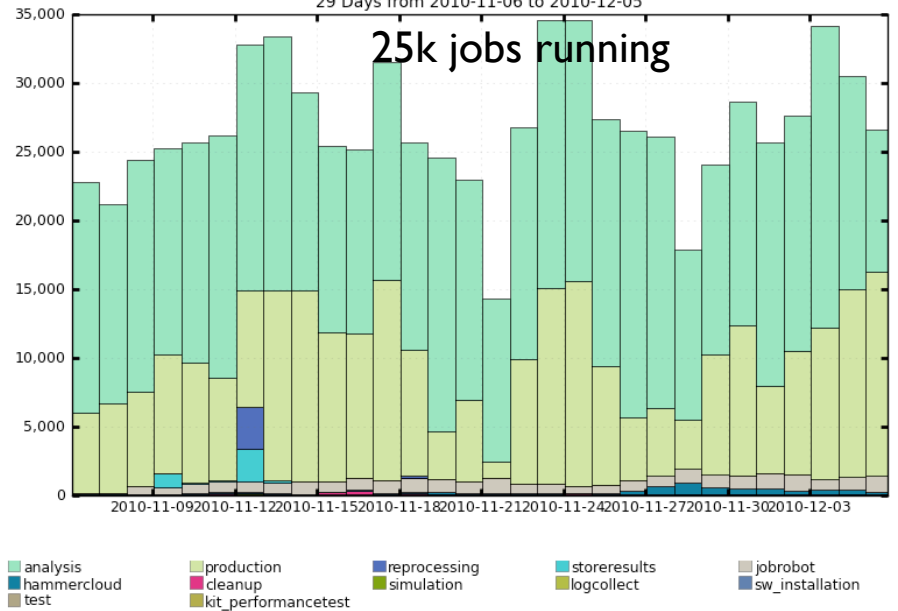
Terminated jobs
29 Days from 2010-11-06 to 2010-12-05

Routinely 150k analysis jobs per day



Running jobs
29 Days from 2010-11-06 to 2010-12-05

25k jobs running






Analysis Activity

- ▶ Analysis Jobs and users are ramping up
 - ▶ Current global priorities are coarse
 - ▶ Transition to resource constrained environment will not be painless


Analysis Job Slot Usage

- T2 analysis pledge for Q4 2010 is 13,600 slots
- Now almost 90% used, and increasing rapidly!




Analysis Jobs Per Day

- Over 140K jobs/day submitted in week 47.
- Testing benchmark was 100K jobs/day.



Stage Out Errors Reduced


- (Mostly) T2-T2 CRAB job stage out errors reduced over the same period.






Migration to AOD

- ▶ Moving to AOD has started
 - ▶ Good trend, but total number of jobs on AOD in November was ~200k, about 4% of the total
- ▶ Most AOD samples are from the Nov. 4 R-reco



Number of jobs (Top-20)



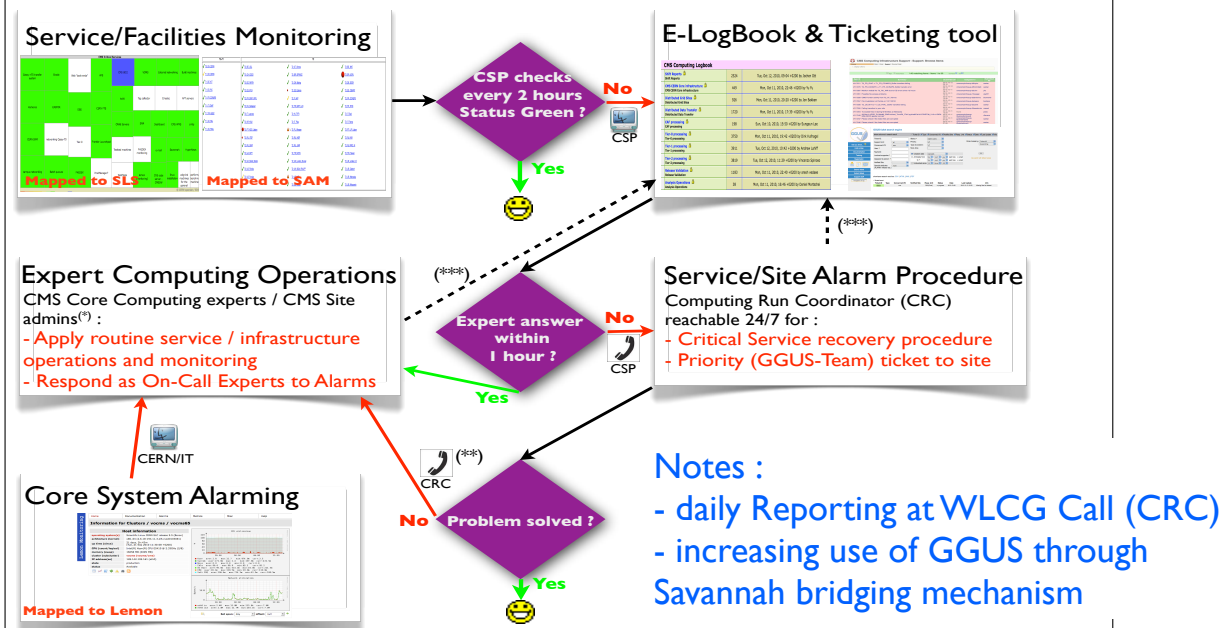
Operations

▶ Central shift procedures working in operations

- **Central Services :**

- Went into 24/7 operations mode since the LHC collisions start

- ➔ Experienced only a few outstanding service incidents with “hero-mode” recovery
 - ➔ Revisited Criticality Level Documentation and Service Recovery procedures : <https://twiki.cern.ch/twiki/bin/view/CMS/CMSCriticalServicesDocumentation> (see next slide)
 - ➔ Critical Services Alarm/Recovery included in Computing Shift Procedures. Successfully tested for DBS&Frontier, with involvement of Computing Run Coordinator on duty



Critical Service Recovery

Documentation for CMS Critical Services

Complete:

The present page shows the Criticality Level (CL) of CMS Services and those of CERN/IT services they depend on. It is a number between 1 and 10 that describes the impact of losing a given service has on CMS and defines the time a response should take. Work is still in progress to finalize all the information (Dec 2, 2010).

Standard CMS Services Operation and Administration

- The [Central CMS Service Operations Documentation](#) contains all standard Service Operations and Administration procedures

CMS Service Criticality and Alarm Recovery Procedures

Definition of Criticality Rankings

The ranking system used to differentiate the services is:

Criticality Level	Meaning	Response Time
10	CMS stops operating	0.5
9	CMS stops transferring from Cessy	0.5
8	T0 Production stops	0.5
7	T1/T2 Production/analysis stops	2
6	Services Critical when needed but not needed all the time (currently includes documentation)	4
5	A service monitoring or documenting a critical service	8
4	CMS development stops if service unavailable	12
3	Services not critical for CMS	24
2	Services not required for CMS	72
1	Used by an insignificant fraction of CMS	72
0	Not used or discouraged by CMS	Forever

CMS Services and Alarm Recovery Procedures

The following table includes all CMS Services, their CL and a link to the Alarm Recovery procedures. The latter are readable ONLY by the concerned operators and application

Service	Impact	Service Description	Complete
Data Management			
DBS	8	DBS used for simulation and reconstruction	Complete:
Frontier Launchpad	8	Used for serving Oracle information from Conditions into the distributed SQUID infrastructure	Complete:
PhEDEx Agent	8	Used to support local PhEDEx agents.	Complete:
Workflow Management			
T0 Bookkeeping	9	Central T0 Submission Service	Complete:
T1 ProdAgent	7	Used for submitting re-reconstruction and skimming jobs	Complete:
T2 ProdAgent	7	Used for submitting simulation jobs	Complete:
T0 Development	4	T0 development and tests platforms	Complete:
Release Validation	6	Used to verify workflows and CMSSW releases	Complete:
CRAB	7	CMS Analysis Submission tool	Complete:
CRAB@CAF	8	CMS Analysis Submission tool for calibration	Complete:
Web Tools Services			
Frontend	8	Web Server used for WebTools Backend Applications	Complete:
Web Tools Base	3	Masthead component used by WebTools Backends	Complete:

Frontier Launchpad Critical Service Documentation

Name: Frontier Launchpad

- Frontier Launchpad Critical Service Documentation
 - Name: Frontier Launchpad
 - Service Contacts
 - Manager: Dave Dykstra
 - Technical Description
 - List of instances
 - List of Dependencies
 - Recovery procedures
 - Monitoring info
 - Description of Lamon Monitoring
 - System
 - Exceptions and Alarms
 - Description of SLS Reporting

Service Contacts:

- If any part of the Frontier Launchpad service is down, see [Recovery Procedures](#) at the bottom of this page

Name	Role	Contact info	Hours (CERN time)	On Call
Weizhen Wang	operator	j cell: +41 767744208 office: 72167	M-F 8am-6pm	M-F 6am-8am S-S 6am-10pm
Dave Dykstra	developer	office: +1 630 840 3227 cell: +1 630 485 1223 home: +1 630 264 0296	M-F 4pm-12am	M-F 12am-6am S-S 10pm-6am
Barry Blumenfeld	developer/ operator	cell: +1 410 252 5814		If can't reach Weizhen or Dave

NOTE: In case none of the experts On Call above are reachable AND if the recovery procedures for the CRC (see below) are not sufficient, please call the Facilities Operations Level 2 coordinators Peter Kreuzer (+41 78 882 23 11) or Page PII

Manager: Dave Dykstra

Other Contacts: Weizhen Wang, Barry Blumenfeld

Technical Description:

Tomcat web services pass web queries to oracle databases.
Squid caches recent queries.

List of Instances:

Instance	Host	In SDB
Prod	vooms03, vooms04, vooms126	Yes

List of Dependencies:

Service	Consequence
Oracle	Queries cannot be answered


Recovery procedures

- First check to see if the CMDB database in SLS is having errors. If it is, then follow the database procedures instead.
- Next, send a message to Front-End-support@cern.ch.
- If it's no response in that within a reasonable amount of time:
 - If it is during working hours of someone in the above contact list: call that person.

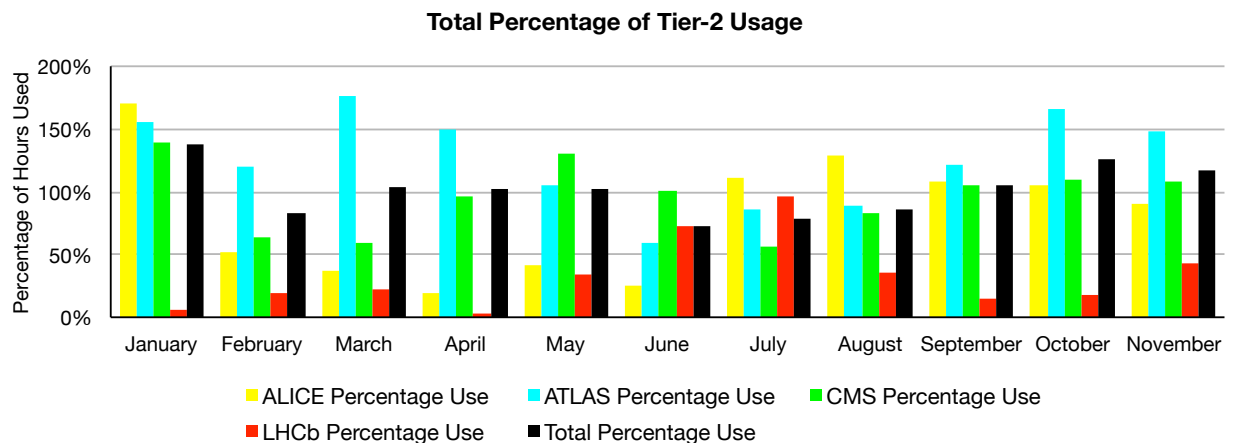
- Service Application Managers responsible for maintenance of Critical Recovery Procedures + List of On-Call Experts (read-restricted to CRC and Computing Management)
- VOC responsible for overall documentation coherence and for mapping CMS CL's to CERN/IT Service Data Base (SDB)

Resource Utilization

- ▶ Utilization of Tier-1 resources grows in the second half of 2010
- ▶ Expected to continue with reprocessing of data and 8TeV simulation



- ▶ Tier-2 resources hit 100% in September and are staying there



Challenges Ahead in 2011/12

- ▶ CMS had a lot of flexibility this year due to the smaller number of pile-up events and the lower than expected live time
 - ▶ We could take higher trigger rates than planned for in the CTDR
 - ▶ Events we did take were smaller and reconstructed faster
 - ▶ Preliminary numbers for 16 interactions per crossing put us very close to the computing model expectations
 - ▶ Early numbers for RECO size are large and may need to be reduced
 - ▶ We have to assume in 2011 that the live time will be higher
 - ▶ Introduce prompt calibration loop and reco delay. Improve usefulness of PromptReco data
- ▶ CMS is rolling out an upgraded workflow management system for reprocessing to start and eventually analysis
 - ▶ Will impact the ways jobs are run on Tier-1s and Tier-2s
 - ▶ Lots of commissioning work

Looking forward to 2012

- ▶ Original Plan was that 2012 is a shutdown year
 - ▶ Resources were requested last April to reflect this
 - ▶ Leveling of most resources, with a small increase in tape to handle a reprocessing and more simulation
- ▶ If 2012 is a very busy year at 8TeV, all bets are off
 - ▶ Expect 20-30% increases in many quantities to deal with large datasets and lots of activities.
- ▶ Computing will generate 2 scenarios for 2012 on the time scale of Feb 1.

Outlook

- ▶ It has been an exciting year
 - ▶ Most elements work well
 - ▶ Thanks to the hard work of development and operations
 - ▶ CMS has been able to do things more frequently than we planned, which has given good flexibility for analysis and results
 - ▶ 2011 and 2012 look very busy
 - ▶ The 2011 resources are based on the Computing model requests made last April
 - ▶ DMWM Offline is rolling out a new workflow system. Impacts reprocessing, simulation, and analysis.
 - ▶ Needed to reduce the operations effort and deal with the large volumes of data. Lots of commissioning work.