


CMS Operations Update

Ian Fisk, Claudio Grandi

Machine performance has averaged to about what was planned for

- Months with technical stops are lower, others are higher



CMS Trigger performance is about what was expected as well

- This is counting all events with an expected 25% overlap
 - Fluctuating around the nominal value of 375Hz

Total events expected by this point was 1.3B events and 1.1B events are collected

Month	Average Trigger Rate (with overlap)
March	356Hz
April	334Hz
May	393Hz
June	431Hz
July	361Hz

Event size is very close to estimates

- MC Reco is larger due to out of time pile-up, which wasn't in the original planning



In general the RECO time is about 20% higher than anticipated before the technical stop

- After technical stop the reco time is essentially doubled

Tier	Observed Size	Expectation
Data RAW	230kB	390KB
Data RECO	590kB	530kB
Data AOD	165kB	200KB
MC Reco	970kB	600kB
MC AOD	250kB	265kB

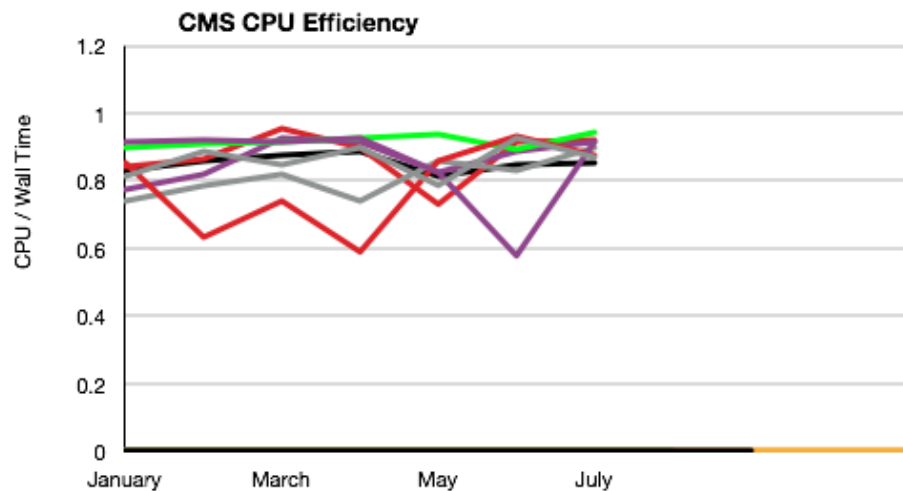

We have observed periods of 100% utilization at the Tier-0



- Memory of the application is reducing the number of cores we can run
- Newer version of the code is available



Pledge Utilization of Tier-1s has increased as we have ramped up more simulation production at Tier-1s

- High utilization and good CPU efficiency







All the Disk is used by the LHC Experiments and CMS is a reasonable trajectory to use the tape pledge for 2011

Tier-2 CPU efficiency noticeably improves with the change to CMSSW_4_2



- Improvements in the IO layer
- Good Usage over all



We've produced about 2B new events in 2011, and re-reconstructed an additional 1B



– This is a little ahead of where we expected to be

- A reprocessing of the data sample is expected in the Fall
- Increase in the Production of FastSim

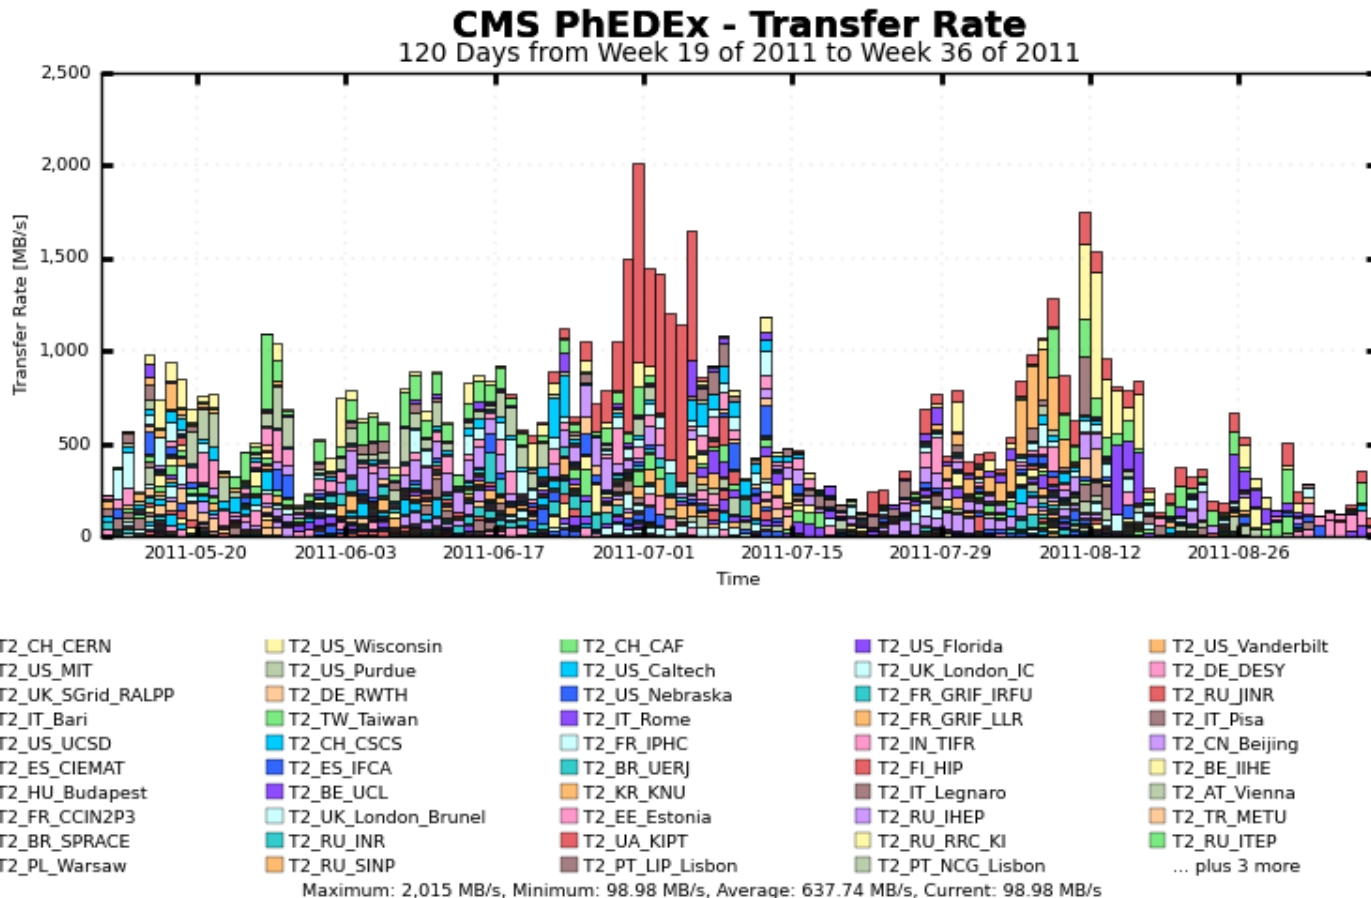



Over the summer we were at 2M terminated jobs a week (1.5M analysis jobs)

– 200k Analysis jobs a day



Looking at the destination Tier-2 transfers it's surprisingly stable over the year





Trying to continue making Tier-3s work well for analysis

– Basically this is extra computing for analysis


We have much better accounting of the data that is being accessed

- CRAB reports to the Data Popularity service for every file access, CPU hour, and individual
 - Project delivered to us by CERN IT-Experiment Support

We are using the analysis data tiers as intended and the speed at which the collaboration switches between reprocessing passes seems to be much faster

The access integrated over week of the various data tiers in CMS

- Good adoption of AOD and AOD SIM
- Low user level access may indicate a lot of that access is outside CRAB



In May we transitioned to CMSSW_4_2 and reprocessed all the data collected up to that point in 2011 with 4_2

- May10 Re-reco replaced PromptReco V1, V2, and V3
- New Data was then Promptly Reconstructed with 4_2
- If a growing consistent dataset was needed, the analyzer had to switch

Show 25 entries

Search:

DataSet	Accesses		CPU Time		Users*day	
	[N]	[%]	[h]	[%]	[N]	[%]
Spring11-PU_S1_START311_V1G1-v1	1724026	34.7	723347	37.4	1052	18.5
Run2011A-May10ReReco-v1	896330	18.1	310952	16.1	771	13.5
Run2011A-PromptReco-v2	262103	5.3	185209	9.6	747	13.1
Fall10-START38_V12-v1	254322	5.1	133816	6.9	254	4.5
Summer11-PU_S3_START42_V11-v2	175168	3.5	98953	5.1	276	4.8
Run2010B-Dec22ReReco_v1	108817	2.2	55370	2.9	197	3.5
Run2011A-PromptReco-v4	52357	1.1	51850	2.7	362	6.4
Run2010B-Apr21ReReco-v1	168452	3.4	51110	2.6	207	3.6
Run2010A-Nov4ReReco_v1	58847	1.2	46146	2.4	98	1.7
Run2010A-Dec22ReReco_v1	147199	3.0	42015	2.2	149	2.6
Run2010A-Apr21ReReco-v1	183560	3.7	34877	1.8	193	3.4
Run2011A-PromptReco-v1	49149	1.0	30875	1.6	456	8.0
Run2010B-Nov4ReReco_v1	38503	0.8	23732	1.2	75	1.3
Summer10-START36_V10-v1	13614	0.3	20871	1.1	18	0.3
Commissioning10--GoodColSlim-Sep17Skim-v1	28517	0.6	19600	1.0	19	0.3
Spring11-PU_S1_START311_V1G1-v2	14102	0.3	16393	0.8	54	0.9
Winter10-E7TeV_ProbDist_2010Data_BX156_START39_V8-v1	35152	0.7	14066	0.7	42	0.7
Run2011A-v1	11950	0.2	10718	0.6	62	1.1
Fall10-START38_V12-v2	58073	1.2	8979	0.5	85	1.5
HIRun2010-ZS-v2	11558	0.2	8261	0.4	9	0.2
Fall10-E7TeV_ProbDist_2010Data_BX156_START38_V12-v1	57900	1.2	6351	0.3	70	1.2
HIRun2010-PromptReco-v3	120608	2.4	5148	0.3	8	0.1
HIRun2010-SDmaker_3SD_1CS_PDHIAIIPhysicsZSv2_SD_JetHI-v1	3079	0.1	4339	0.2	3	0.1
Summer10-START36_V10_SP10-v1	30429	0.6	3690	0.2	13	0.2
Fall10-START38_V10-v1	3973	0.1	3394	0.2	5	0.1
Shown Sum	4507788	91	1910062	98.8	5225	91.59
Total Sum	4961795	99.99	1931718	99.69	5692	99.29
DataSet	Accesses		CPU Time		Users*day	

Show 25 entries

Search:

DataSet
Spring11-PU_S1_START311_V1G1-v1
Run2011A-May10ReReco-v1
Run2011A-PromptReco-v2
Fall10-START38_V12-v1
Summer11-PU_S3_START42_V11-v2
Run2010B-Dec22ReReco_v1
Run2011A-PromptReco-v4
Run2010B-Apr21ReReco-v1
Run2010A-Nov4ReReco_v1
Run2010A-Dec22ReReco_v1
Run2010A-Apr21ReReco-v1
Run2011A-PromptReco-v1
Run2010B-Nov4ReReco_v1
Summer10-START36_V10-v1
Commissioning10--GoodColSlim-Sep1
Spring11-PU_S1_START311_V1G1-v1
Winter10-E7TeV_ProbDist_2010Data_BX156_S
Run2011A-v1
Fall10-START38_V12-v2
HIRun2010-ZS-v2
Fall10-E7TeV_ProbDist_2010Data_BXv1
HIRun2010-PromptReco-v3

Within about 2 weeks of the processing being finished more than half of the analysis CPU is on the new sample


- 16% on May Re-reco
- 12% on the original Prompt-Reco

Number of individuals accessing the data is about the same

HIRun2010-SDmaker_3SD_1CS_PDHIAIIPhysicsZSv2_SD_JetHI-v1	3079	0.1	4339	0.2	3	0.1
Summer10-START36_V10_SP10-v1	30429	0.6	3690	0.2	13	0.2
Fall10-START38_V10-v1	3973	0.1	3394	0.2	5	0.1
Shown Sum	4507788	91	1910062	98.8	5225	91.59
Total Sum	4961795	99.99	1931718	99.69	5692	99.29
DataSet	Accesses		CPU Time		Users*day	

- *myproxy* upgrade on Aug 30th (GGUS:73926)
 - initially problems with firewall
 - Fixed in < 1 h
 - *myproxy* bug related to use of proxies created with and without the `–r` option blocked some analysis
 - crab server uses `-r`, crab standalone, via gLite WMS, doesn't)
 - 8 hours of debugging; patch requested to developers, installed the next day
 - *myproxy* bug related to the use of wildcards blocked phedex
 - 1.5 more days to debug the problem and find a workaround
 - patched version of *myproxy* put in production the week after
- Alias for `cmsdoc` broken on Sep 6th (GGUS:74099)
 - CRAB analysis blocked outside CERN
 - Fixed in 2 hours

- LSF replaced by *transfermanager* in Castor during the TS
- First stress as soon as data taking restarted. Many rfcp hanging impacting transfers from P5 and to T1s and T0 processing (GGUS:74085)
- Patch installed on 8 Sep solved the problem. Final patch expected this week
- Internal CMS problem caused by an inconsistency in the DB after the castor problem, blocked reprocessing of new runs for 24 h
- Now dealing with the backlog and with the increased processing times due to high PU
 - x2 processing time in express and 300 MB memory increase
 - x2.5 processing time in reco and 800 MB memory increase



Castor problem

Internal CMS problem
(no new runs acquired)

Reabsorbing backlog and
coping with increased PU