CMS Computing Operations

Stephen J. Gowdy
Grid Deployment Board
21st March 2012

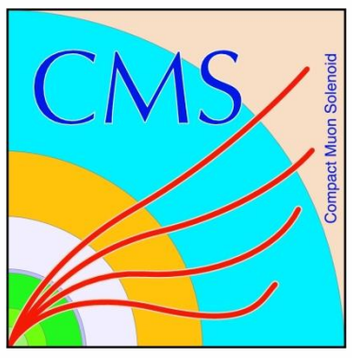

Special thanks to Oliver Gutsche for providing the material

Computing Operations in CMS

- CMS Computing Operations was re-organized Beginning of 2012
  - Reflect the change from **Commissioning** to **Operations and Optimization** of the **Computing systems and infrastructure**
- Data Operations & Facilities Operations were merged → Computing Operations
  - Lead by: Markus Klute (MIT), Oliver Gutsche (FNAL), Pepe Flix (PIC)

Overview

Leadership

L2

Operations Office

Experience

L3

CERN | **MC** | **Services** | **Tier-1 Sites**

Contact | **Processing** | **Tier0** | **Tier-2 Sites** | **Transfer**

CORE

Tier-0	Infrastructure	Workflow	Transfer	Monitoring
Team	**Team**	**Team**	**Team**	**Team**
Team Leader	Team Leader	Team Leader	Team Leader	Team Leader
Operators	Operators	Operators	Operators	Operators

Status

- Promising start of new project with a lot of gained synergy
  - Focus is on communication between the different parts of the Computing Project and to other projects of CMS
- Still places where we can improve processes and smoothen out operational procedures

- Team is not complete, following positions are not filled:
  - L3_CERN (1): service and infrastructure coordination at CERN
  - L3_Services (2): global services & infrastructure coordination
  - L3_Tier1Sites (1): liaison to Tier-1 sites
  - L3_Transfers (1): Transfer monitoring and trouble shooting

Review 2011: data

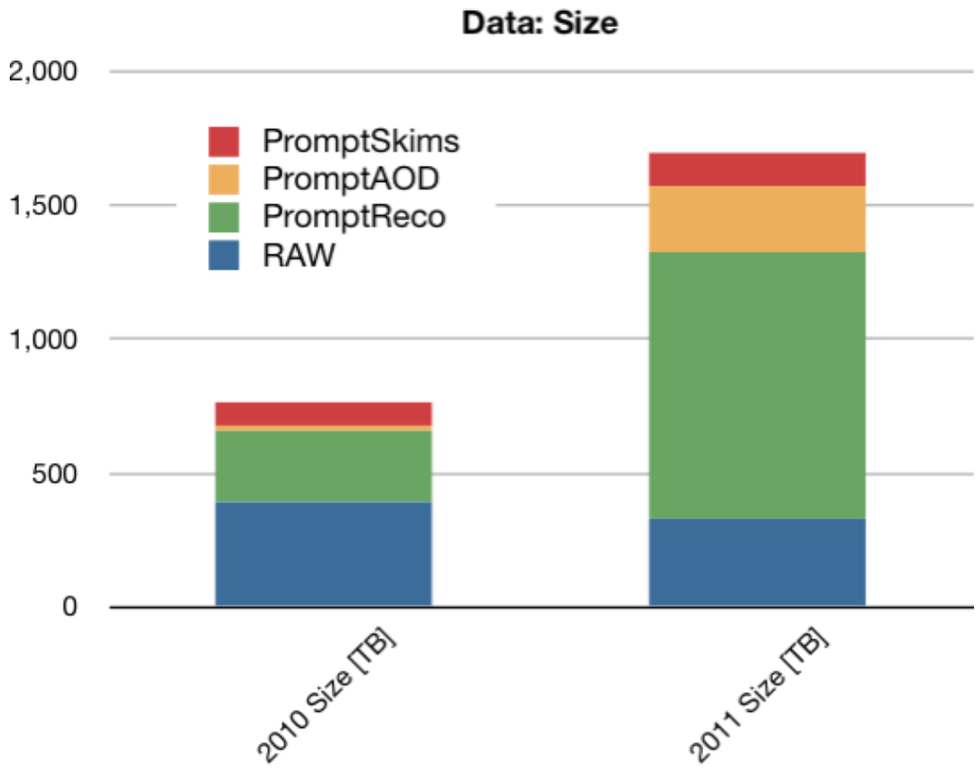

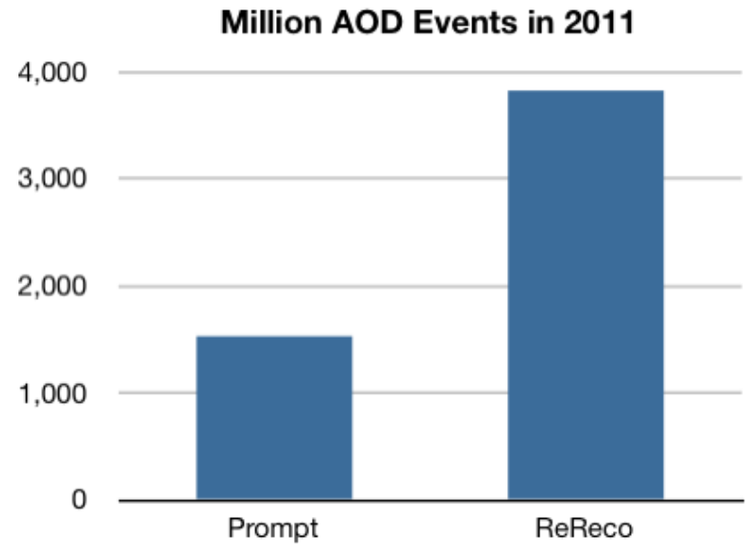

	RAW	PromptReco	PromptAOD	PromptSkims	Total
2010 Size [TB]	393	266	19	87	765
2011 Size [TB]	331	995	247	123	1,696
Total Size [TB]	**724**	**1,261**	**266**	**211**	**2,462**

	RAW	PromptReco	PromptAOD	PromptSkims
2010 [Million Events]	1,536	1,186	184	222
2011 [Million Events]	1,535	1,525	1,443	284
Total [Million Events]	**3,071**	**2,711**	**1,627**	**506**

- Prompt data: 1.7 PB
- 2011 re-reconstruction passes
  - 27 partial
  - 1 complete
  - In total, re-reconstructed more than twice the events we recorded (once during the year, one time at the end of the year)

Review 2011: MC

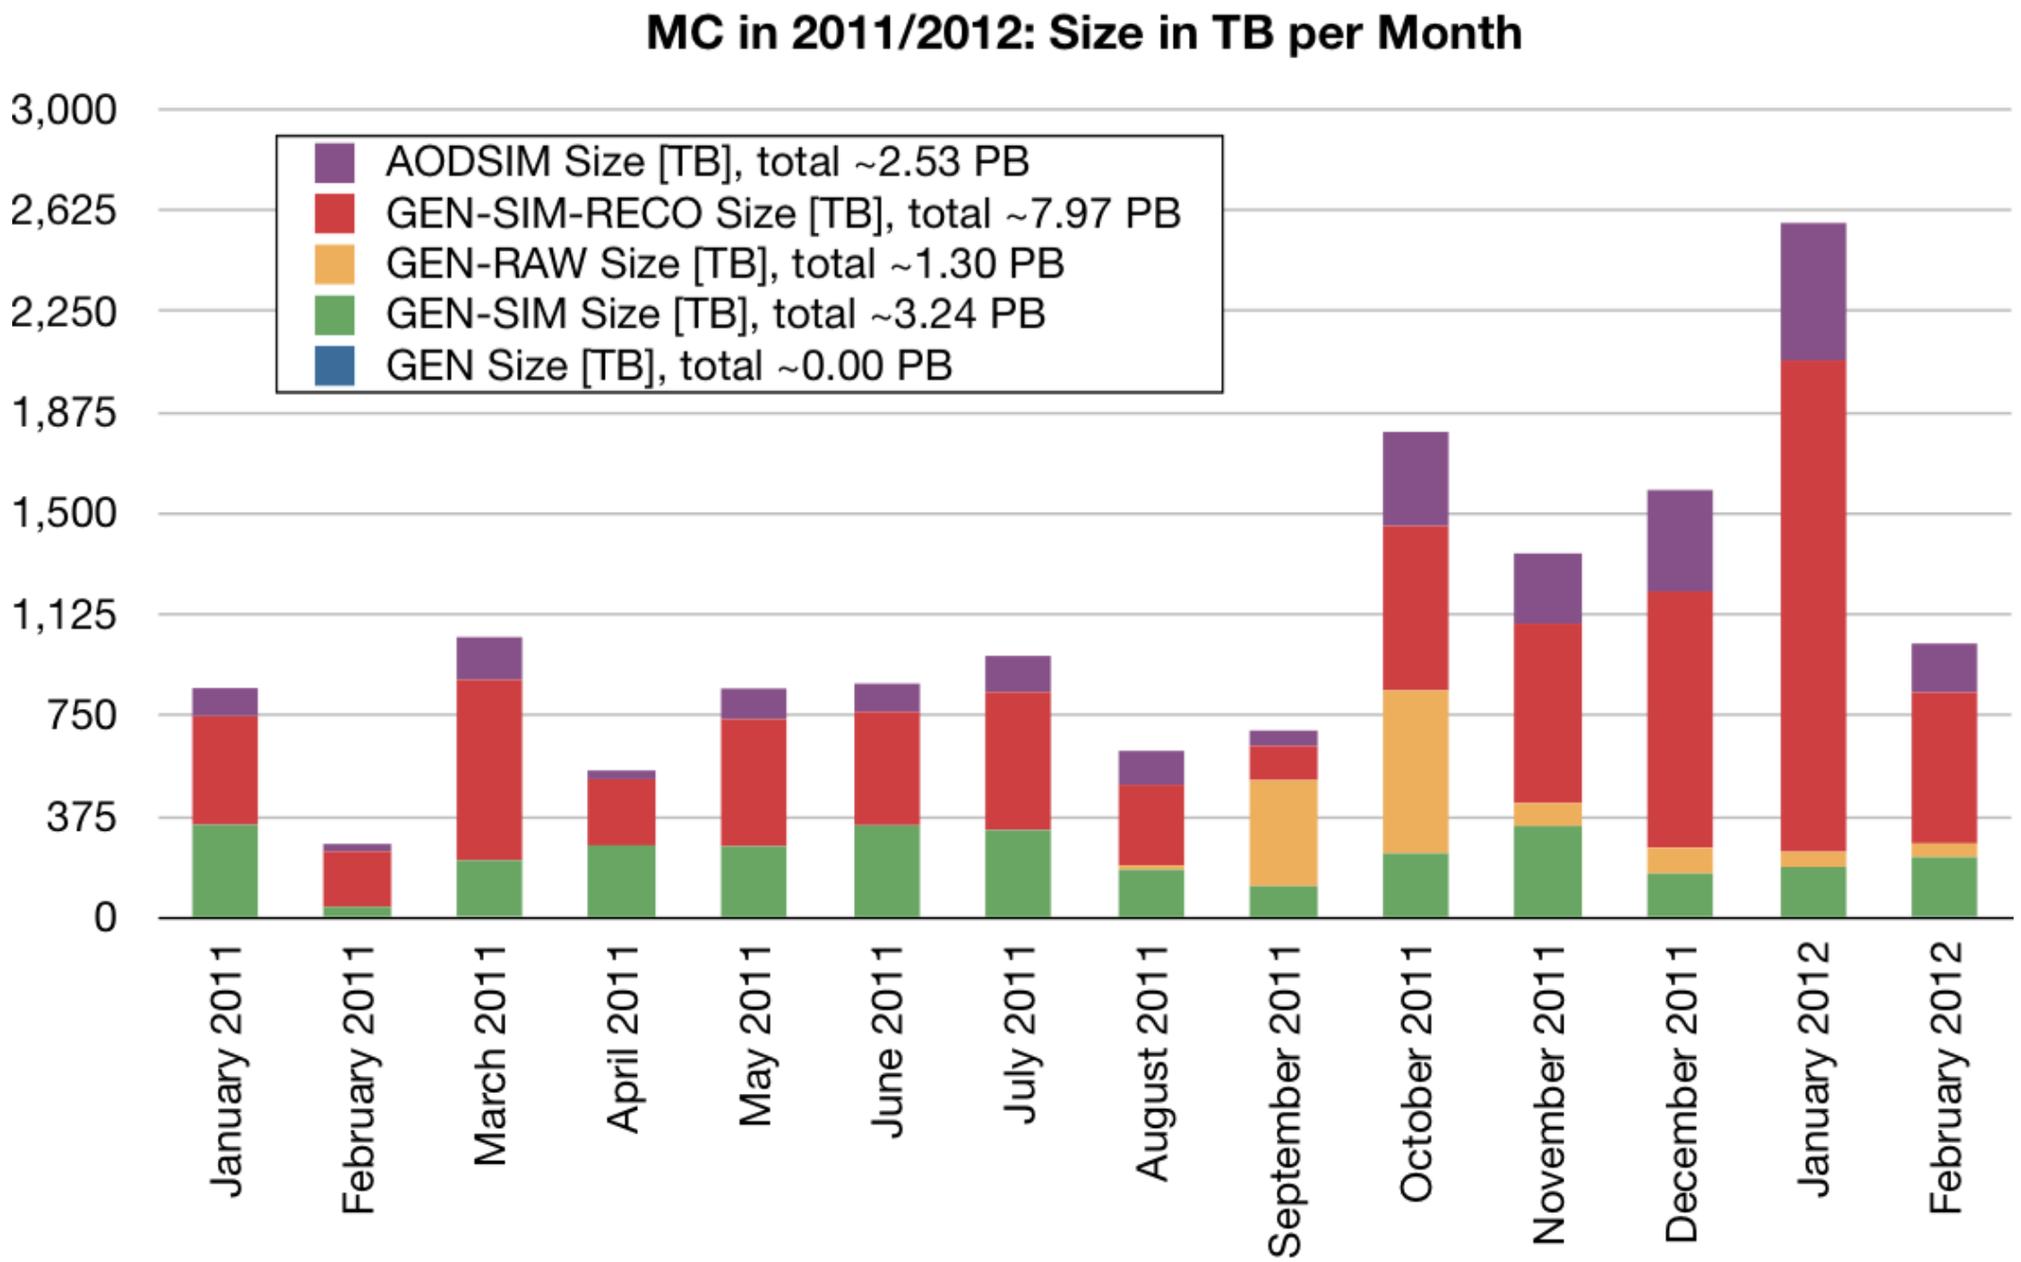

- Total in 2011: 11.5 PB
- Total in 2012: 3.6 PB

Review 2011: MC

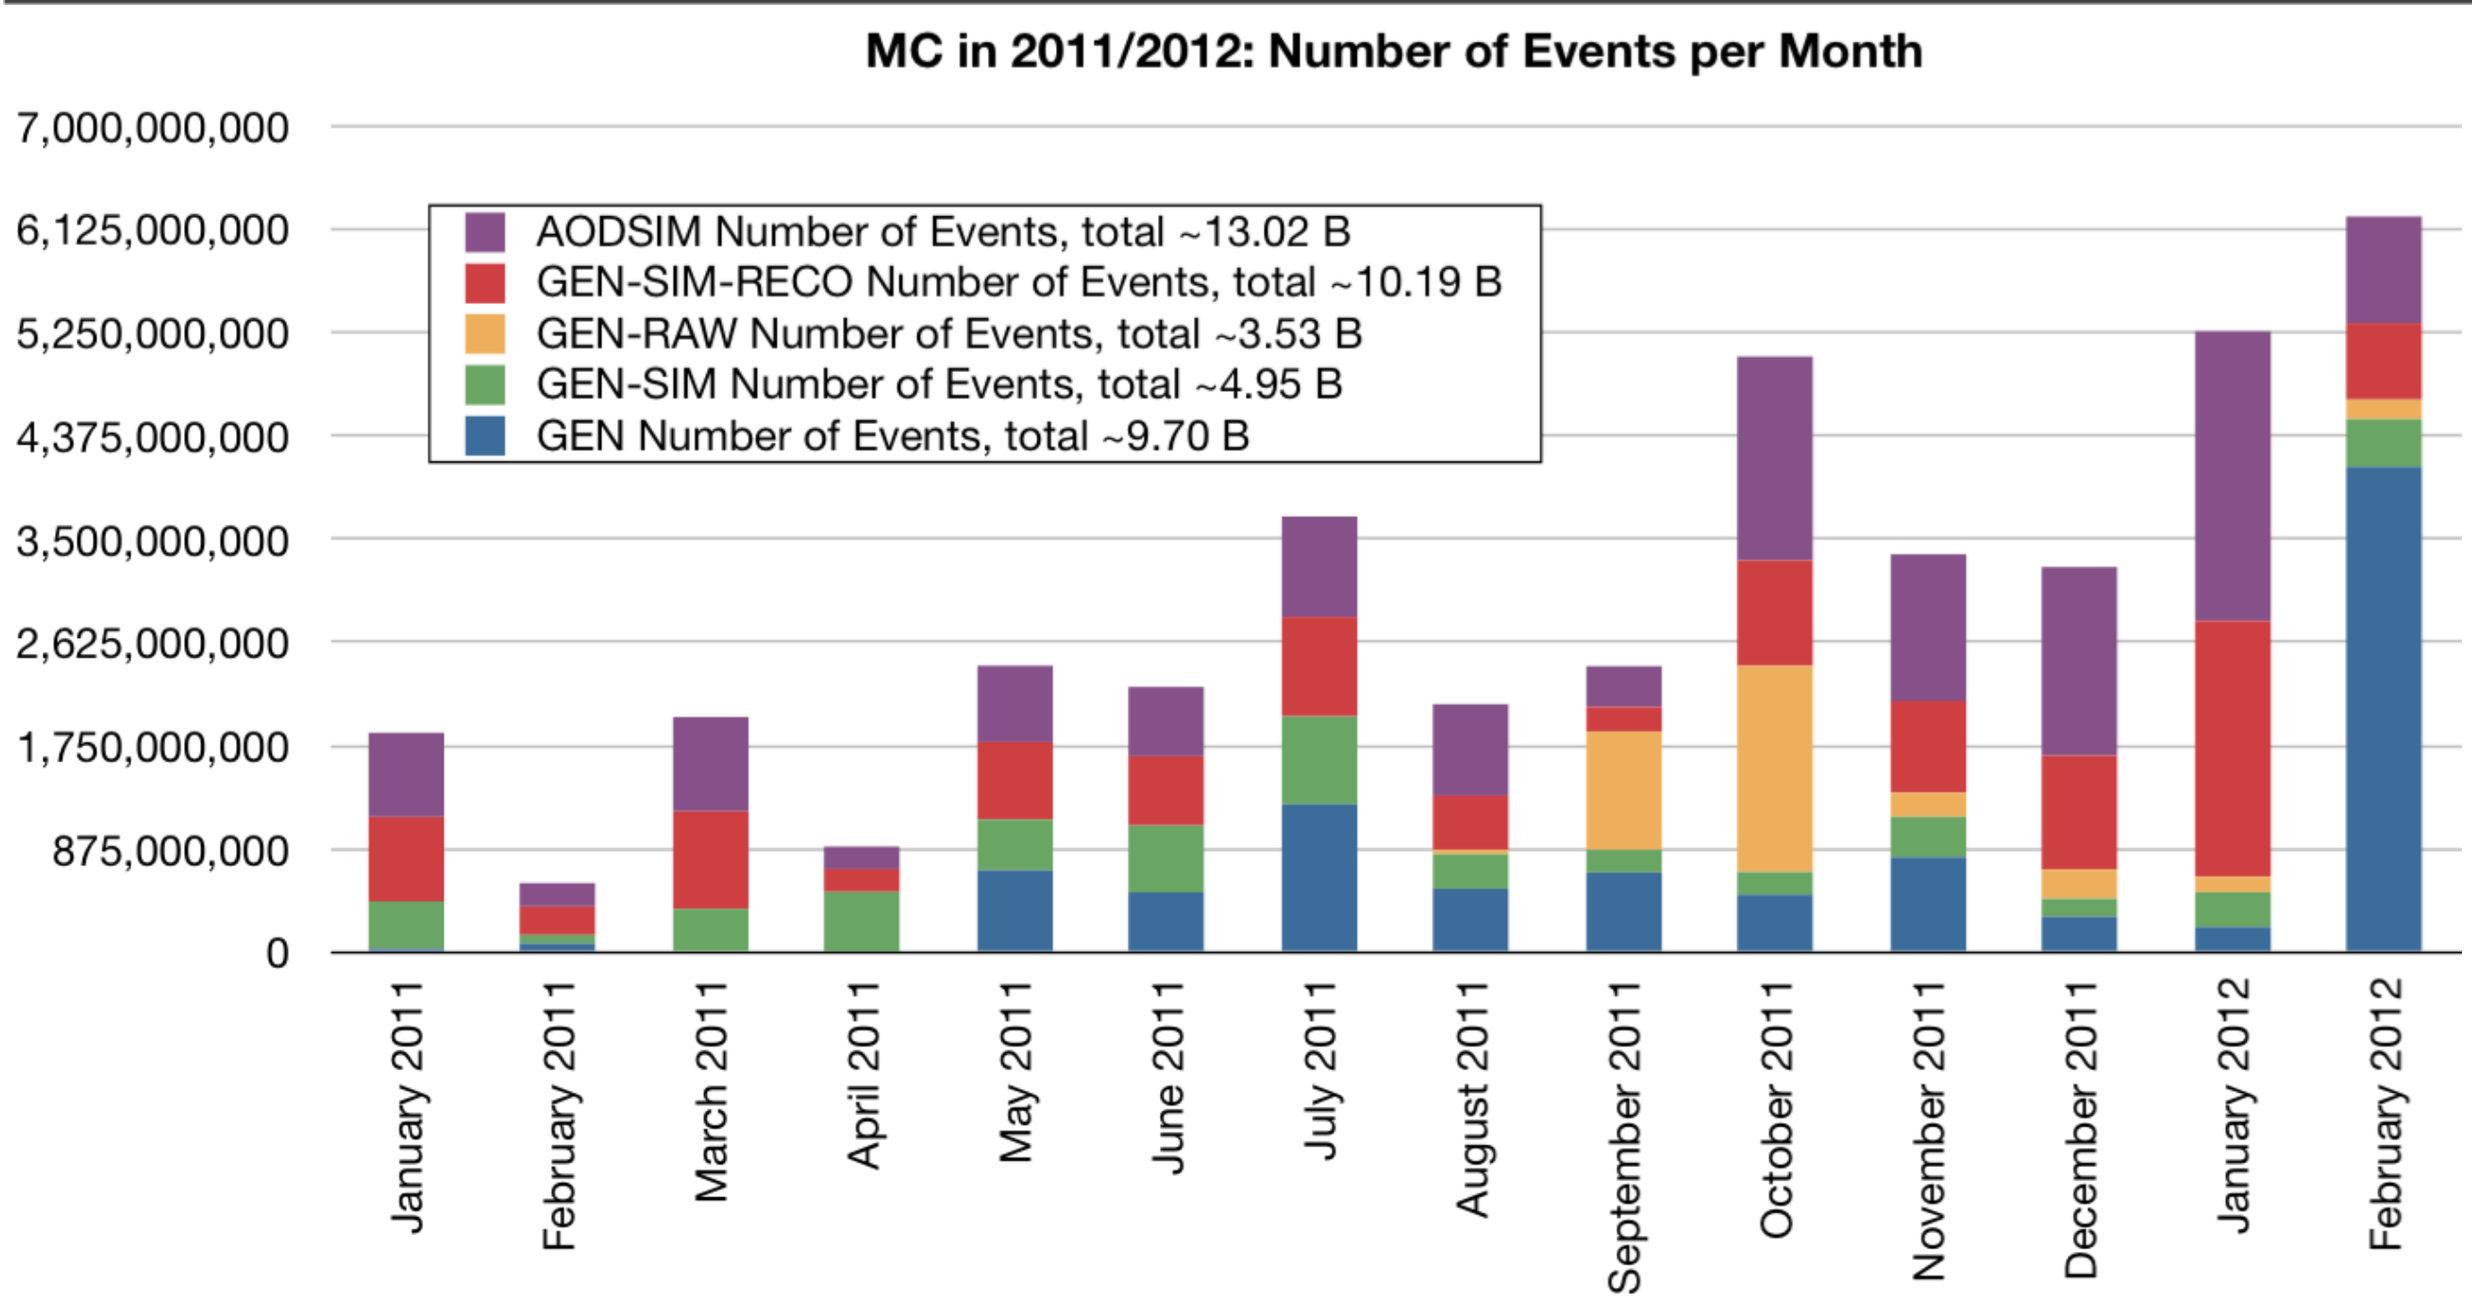

- 2011 GEN-SIM: 4.3 Billion ➔ 2011 AODSIM: 9.7 Billion
- Re-reconstructed and partly re-digitized every GEN-SIM event more than twice

Popularity Service

- https://cms-popularity.cern.ch/
  - Statistics about dataset access through CRAB since June 2011
- Analysis since June 2011:
  - **Summer11_R1**: CMSSW_4_2_X & PU_S4
  - **Fall11_R2**: CMSSW_4_2_X, PU_S6
  - **Fall11_R4**: CMSSW_4_4_X, PU_S6
- Datasets without any access:

	Summer11_R1			Fall11_R2			Fall11_R4	
	GEN-SIM-RECO	AODSIM		GEN-SIM-RECO	AODSIM		GEN-SIM-RECO	AODSIM
Valid Datasets in DBS	2160	2182		1651	2612		2636	
Datasets accessed once	240	1946		799	3		118	
Datasets not accessed at all	1920	236		852	2609		2518	

- Observations:
  - AOD switch worked very well! according to the RECO numbers
  - Sizable number of AOD datasets not accessed at all, also in Fall11_R2
- Caveat:
  - Only looked at number of datasets, not at number of events (could be that many small datasets have not been accessed)
- **Popularity Service very powerful, also for site admins!**

Plans 2012: data taking

- Max possible data taking rate:
  - 300 Hz Prompt reconstructed at Tier-0
  - ~300 Hz additional reconstructed at Tier-1 sites
    - 300 Hz parked and reconstructed in 2013
  - **Important:** these are the planning numbers, the current estimates of used bandwidth are somewhat lower.
- 2012 will see more than 1½ times the data events of 2011
  - Impact for analysis
- 2012 will see a physics MC sample of 3 Billion events compared to 2.5 Billion in 2011

	Plan 2011	Plan 2012
Data taking [s]	5,200,000	5,200,000
Data taking rate [Hz]	300	600
Overlap	1.25	1.25
Data events	2,000,000,000	3,100,000,000
MC events	2,500,000,000	3,000,000,000

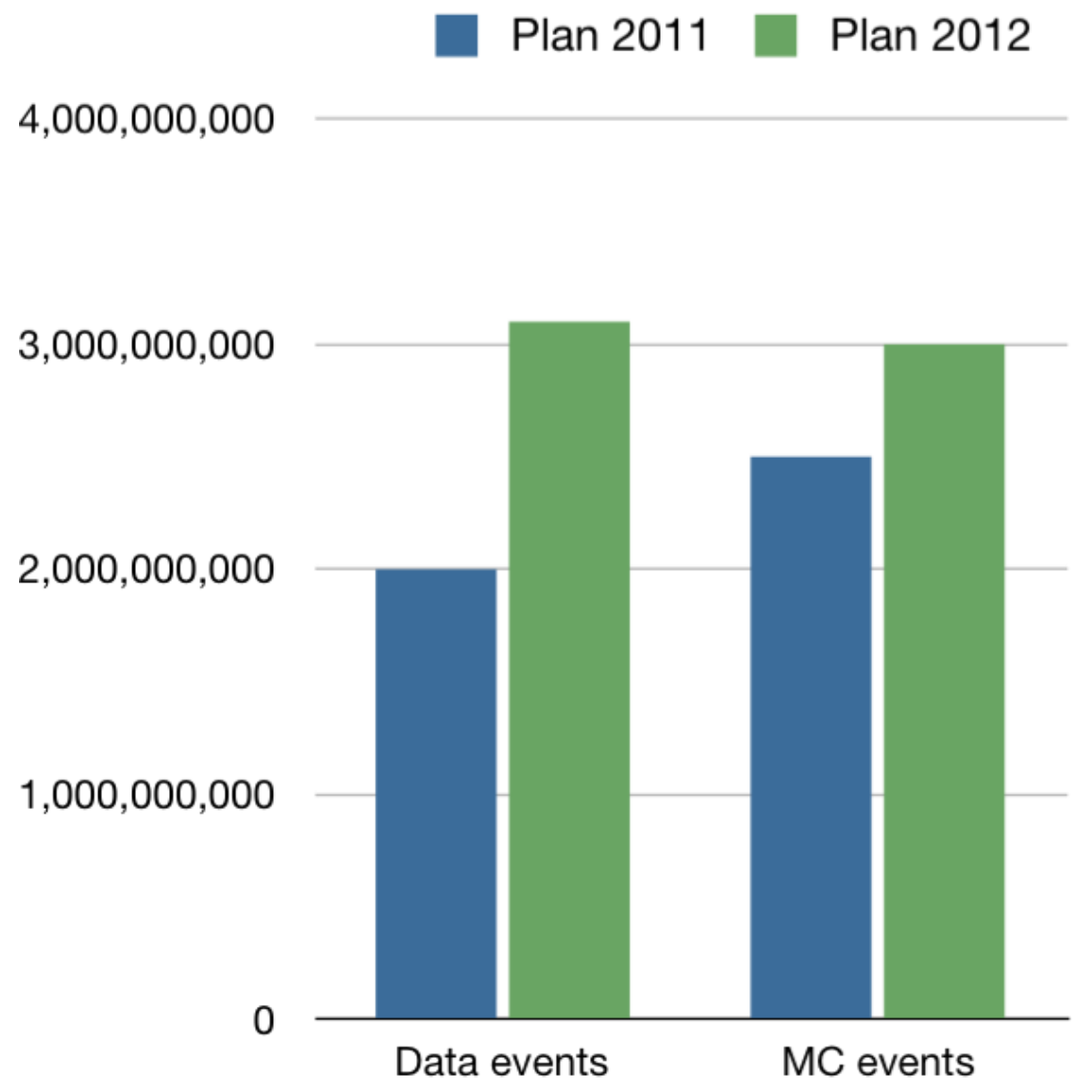

Plans 2012: Re-Processing

- Data re-reconstruction
  - Assume that small re-reconstruction passes like in 2011 will re-reconstruct the whole 2012 dataset
  - No plan for End-Of-Year re-reconstruction pass
- MC re-digitization/re-reconstruction
  - 2011 saw 2 complete MC re-digitization/re-reconstruction campaigns
  - 2012 will only see one complete campaign
- PileUp will increase from 16 to 30 PU events
  - Large impact on data reco and MC re-digitization/re-reconstruction times
  - Small impact on AOD analysis times

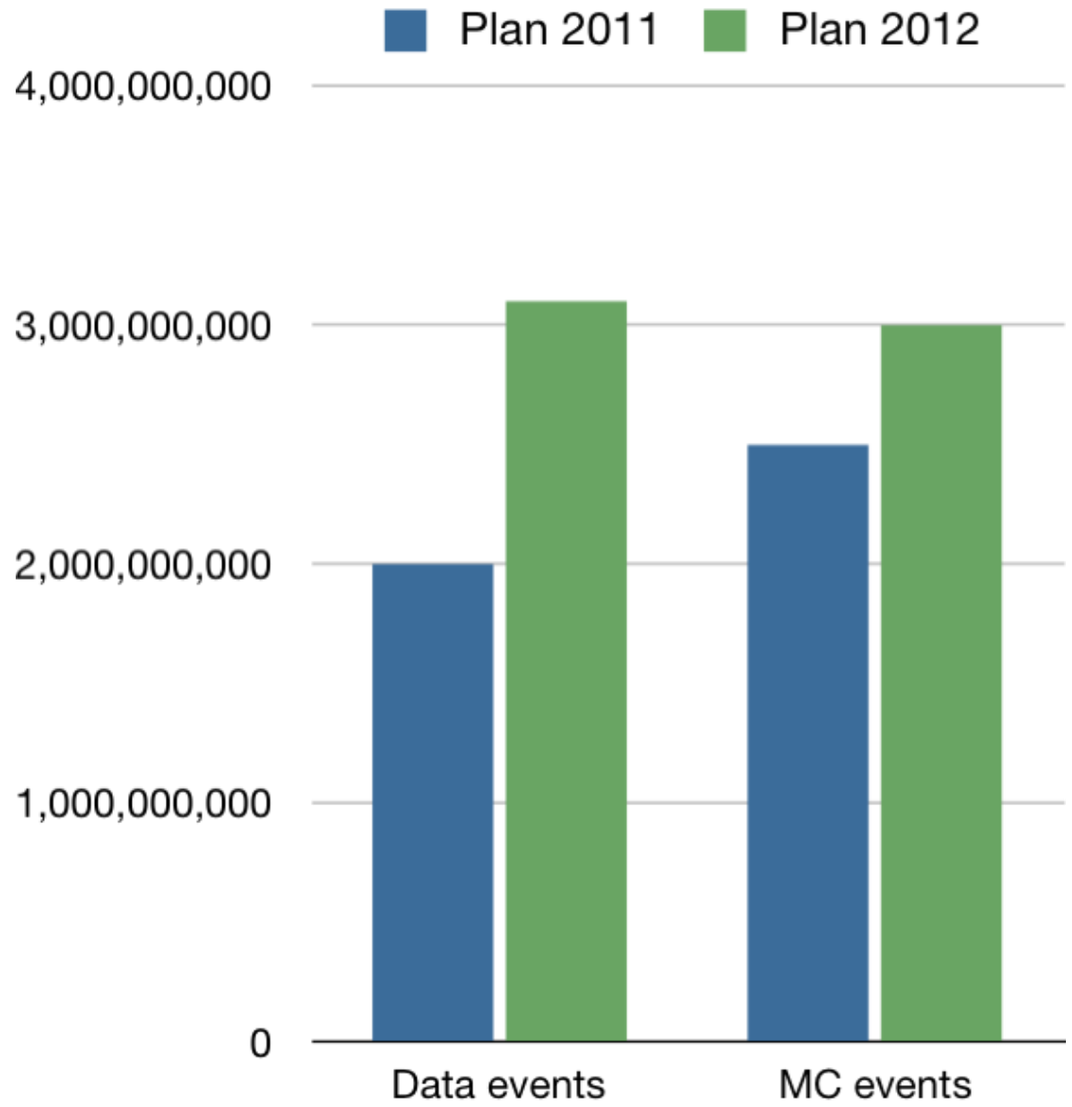

	Plan 2011	Plan 2012
PileUp	16	30
RECO Time Data (HS06s)	92	280
Re-digi/Re-RECO Time MC (HS06s)	164	400
RECO Analysis Time	35	50
AOD Analysis Time	11	12

2012: where we are and where we go

- Christmas 2011/2012
  - Produced GEN events using higher order Generators: 4.5 Billion events
- February 2012:
  - Started MC Simulation in CMSSW_5_0_X
- Mid March 2012:
  - Started MC Digitization/Reconstruction in CMSSW_5_1_X
- April 2012:
  - Switch Tier-0 to CMSSW_5_2_X
  - Re-start MC Digitization/Reconstruction in CMSSW_5_2_X
  - Goal: 1 Billion MC events for ICHEP conference analysis (taking place Beg. Of July 2012)
- May 2012:
  - Possible checkpoint to decide if sufficiently large progress in physics performance of software was made, if yes:
    - Validate CMSSW_5_3_0 and deploy at Tier-0 during technical stop around ICHEP
    - After ICHEP, start 2012 data re-reconstruction
    - Decide if MC has to be re-digitized and/or re-reconstructed
    - Decide if 2011 data and MC has to be re-digitized and/or re-reconstructed
- Tier-2 sites can expect to see significant fluctuation of samples going in the site for analysis!

Tier-2 sites

Job Submission: last week

T1 processing FNAL glideIn WMS factory
Full scale: up to 20k running jobs

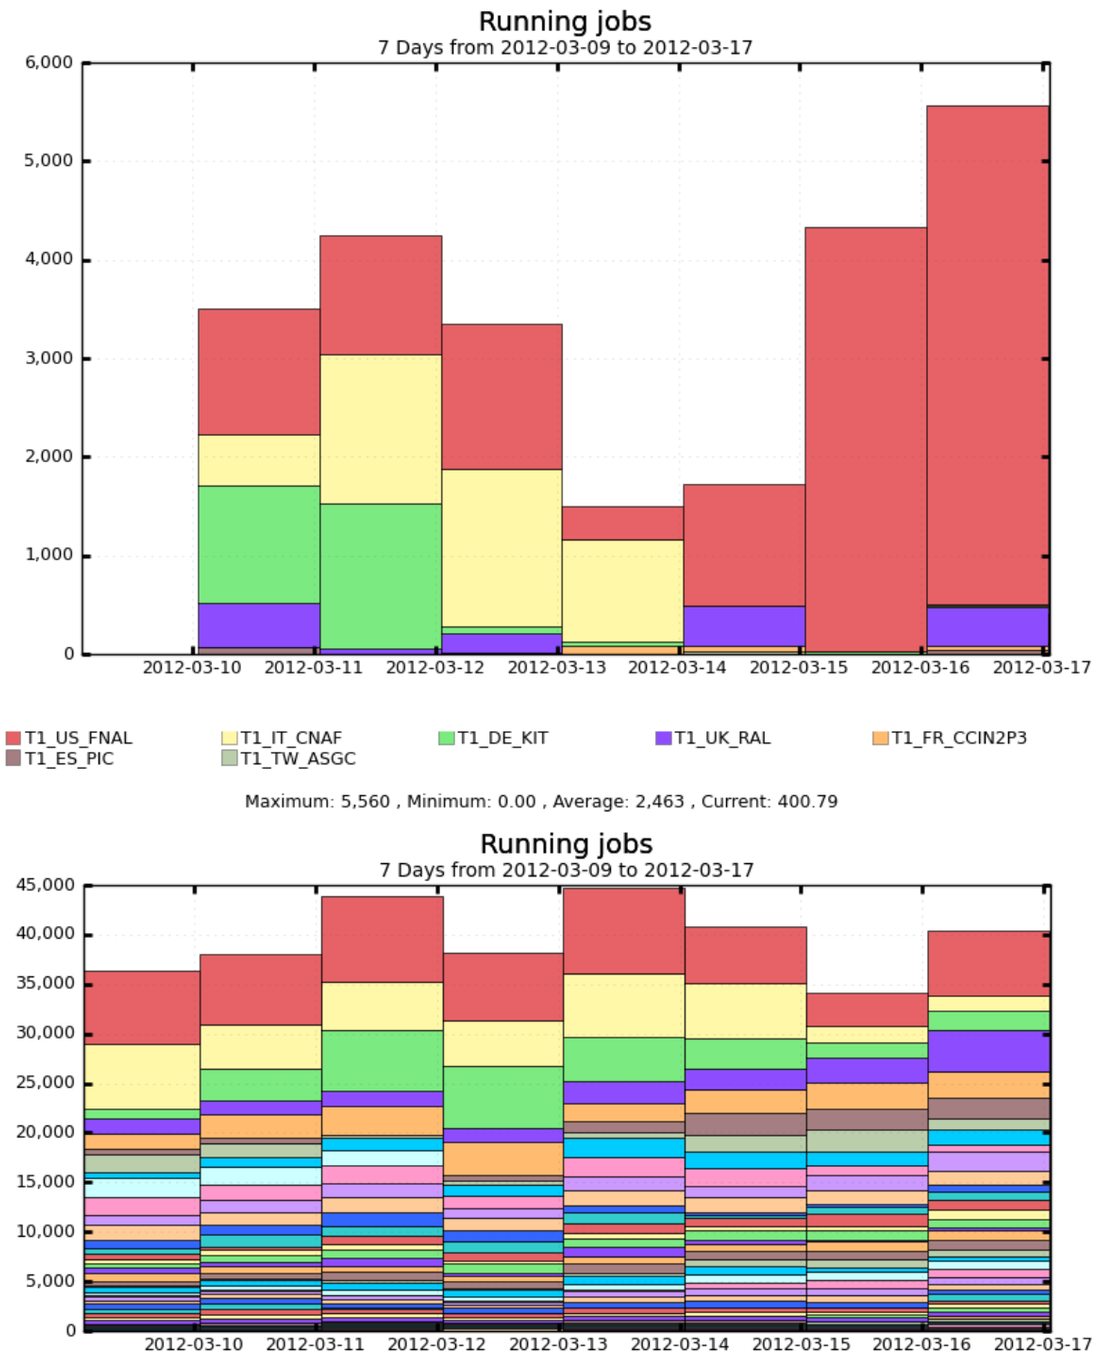

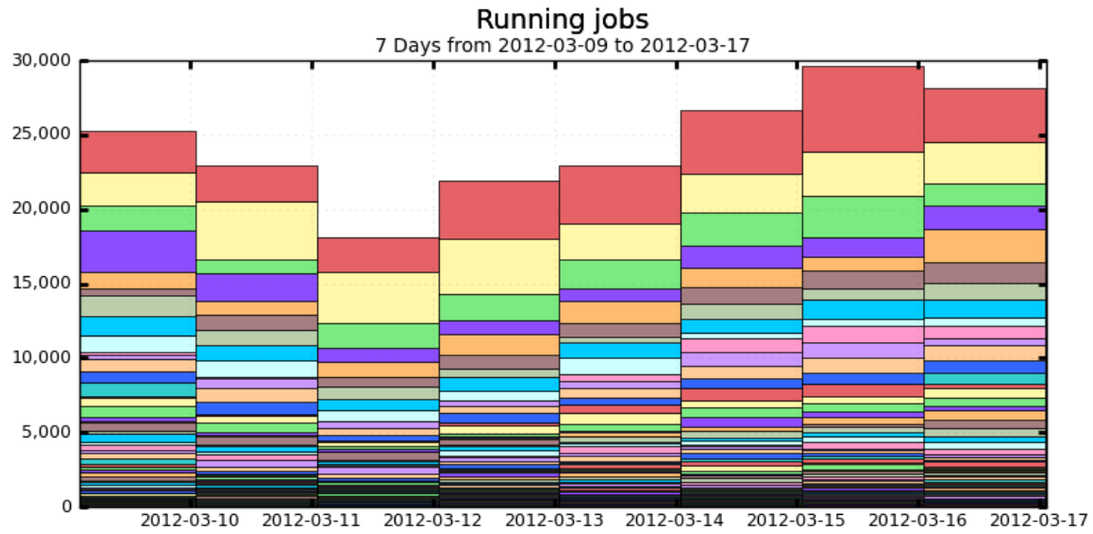

MC production CERN & UCSD glideIn WMS factories
Reached: 45k running jobs on all tier levels

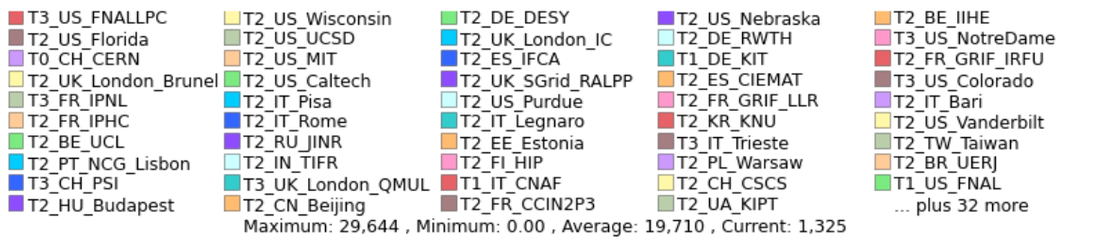

Analysis
gLite WMS & CERN & UCSD glideIn WMS factories
Routinely 30k++ running jobs

Job Submission

- glideIn WMS in CMS
  - Used in Central Operations since Spring 2011
    - Since 2012 also MC production via glideIns
  - Using for Analysis since the CRAB servers at UCSD came online in 2011
- gLite WMS
  - Used only for analysis
- 2012/2013:
  - Hope to migrate analysis to glidein-WMS by the end of 2012
  - Atlas will be moving to glideIn WMS
- Currently using 2 main VOMS roles important for Tier-2 sites
  - MC production
    - 50% of the resources of a Tier-2 site
    - Production role
  - Analysis
    - 50% of the resources of a Tier-2 site
    - Priority role is used to prioritize analysis jobs before jobs with no role
- Important for Tier-2 sites
  - Correct prioritization of these roles and following the respective resource percentages!

Storage

Tier-1 level

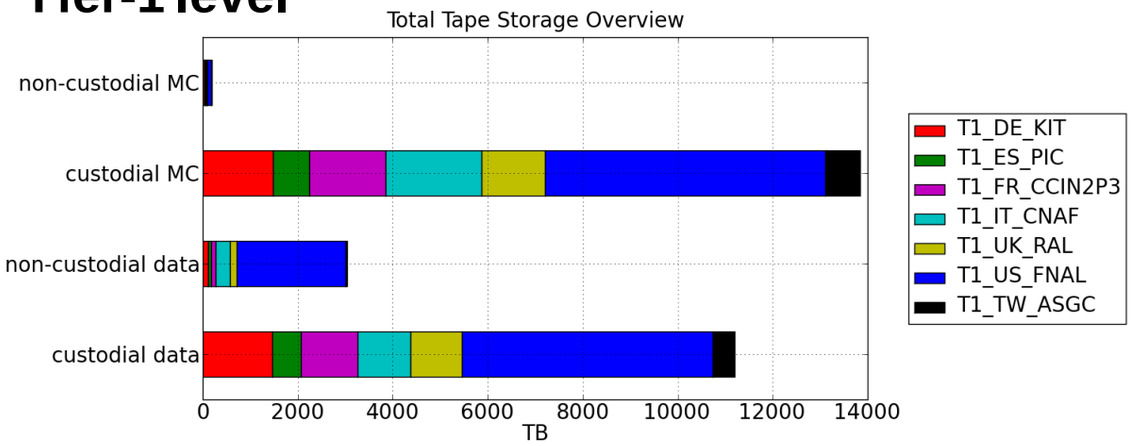

Tier-2 level

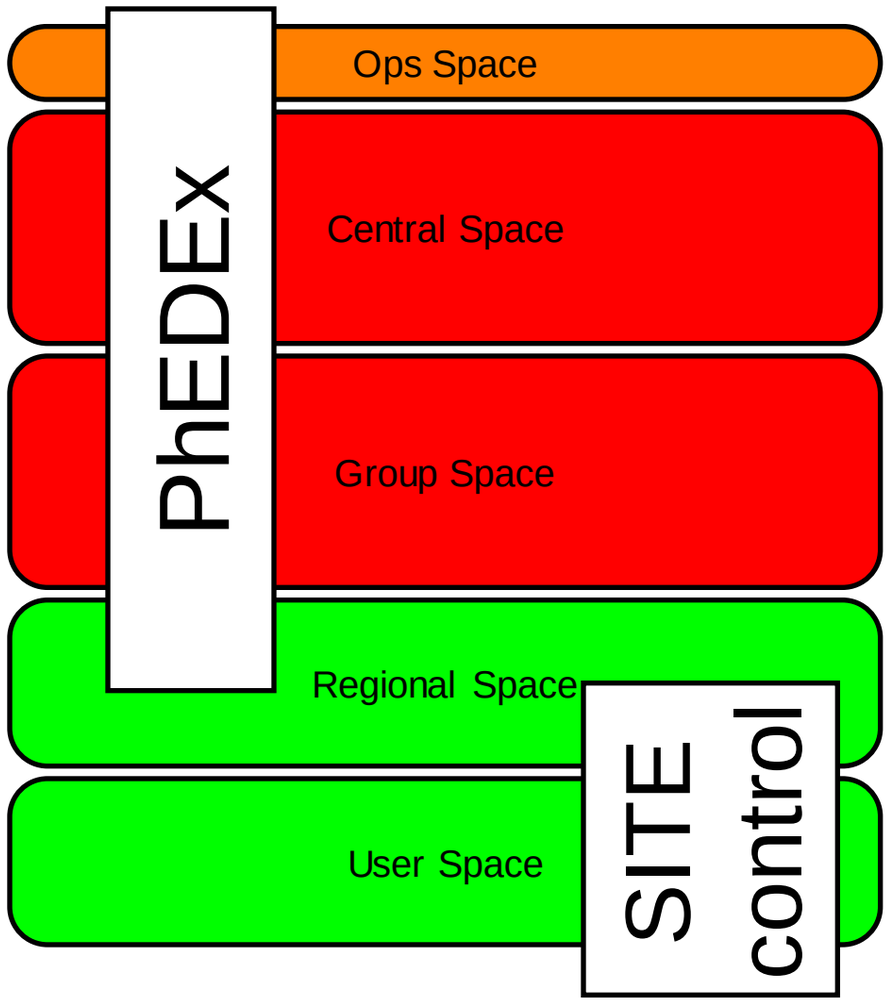

- Data Manager responsibilities very important
  - Ops, Central and Group space needs to be guaranteed
  - Regional and User Space have to be checked regularly and overfilling of the site has to be avoided
  - Data Manager is responsible that central PhEDEx spaces can be used at agreed quotas at all times
- CMS is working on tools to allow central accounting of used space outside PhEDEx control (Regional and User space)
  - Based on regular complete storage dumps at sites
- CMS is also establishing consistency checks for missing files and orphans
  - Tier-1 sites are checked every month
  - Pilot program with test sites on Tier-2 level running
  - Planned to have monthly consistency checks on Tier-2 level as well

/store/user

- Every CMS user has at least one /store/user area where he can store ntuples through GRID access
  - Handled through regional/national organization, ask your institute leader where your /store/user area is hosted
  - Every Tier-2 sites supports a number of users
  - Usually 1-3 TB of space is allocated per user
- CERN:
  - CERN still provides tape backed user directories for CMS
    - /castor/cern.ch/user
  - Last year, CMS users wrote 2.7 PB (compared to 1.7 PB prompt data!)
  - Access to these stored files has impacted data taking and transfers
    - Files need to be staged from tape, available disk pool very small, significant activity impacting the whole system
  - CMS will work in the next months to transition users to official /store/user storage and close /castor/cern.ch/user
  - Expect increased demand for /store/user areas at your sites, CERN users can also qualify for space on EOS at CERN (T2_CH_CERN)

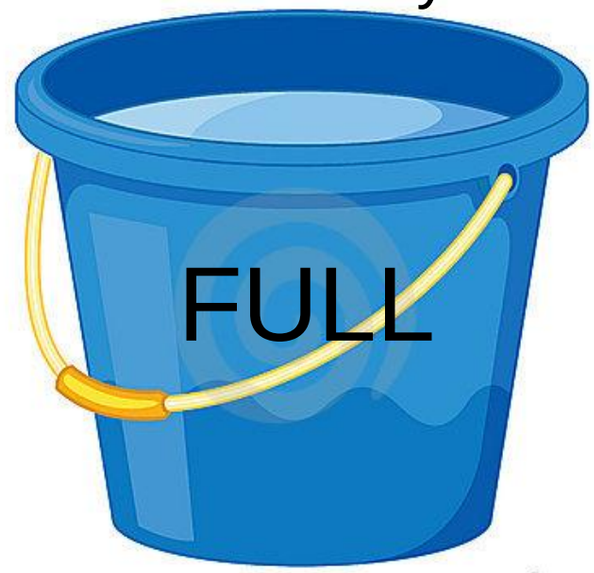

CVMFS

- Software deployment
  - Transition to model where external software installation is not necessary anymore ➔ access to new software releases is provided through self-updating systems
    - 2 solutions for sites: CVMFS or CRON installation
- CRON:
  - Site runs CRON job that installs new software releases automatically
- CVMFS (CERN virtual file system based on SQUID caches)
  - Site mounts CVMFS on all worker nodes
  - Software is installed centrally at CERN and distributed through CVMFS to all sites
  - Status: Preparation of central CERN installation about to me migrated to final production hardware.
  - First Tier-2 sites in UK tests CVMFS served from the current installation base at CERN, then we expand testing to more sites and in the end migrate all sites to CVMFS or CRON

Multi-Core CMSSW jobs

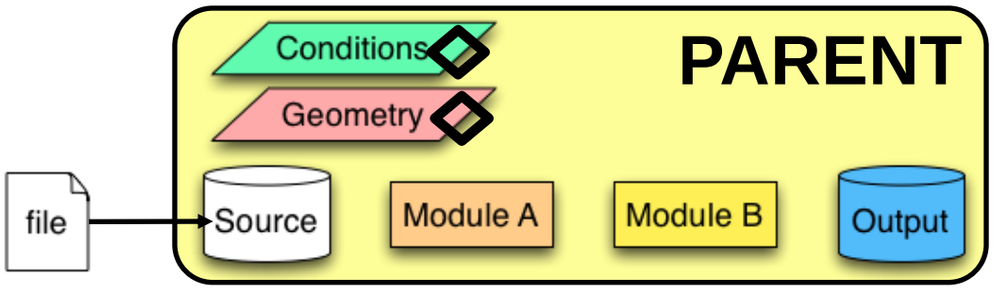

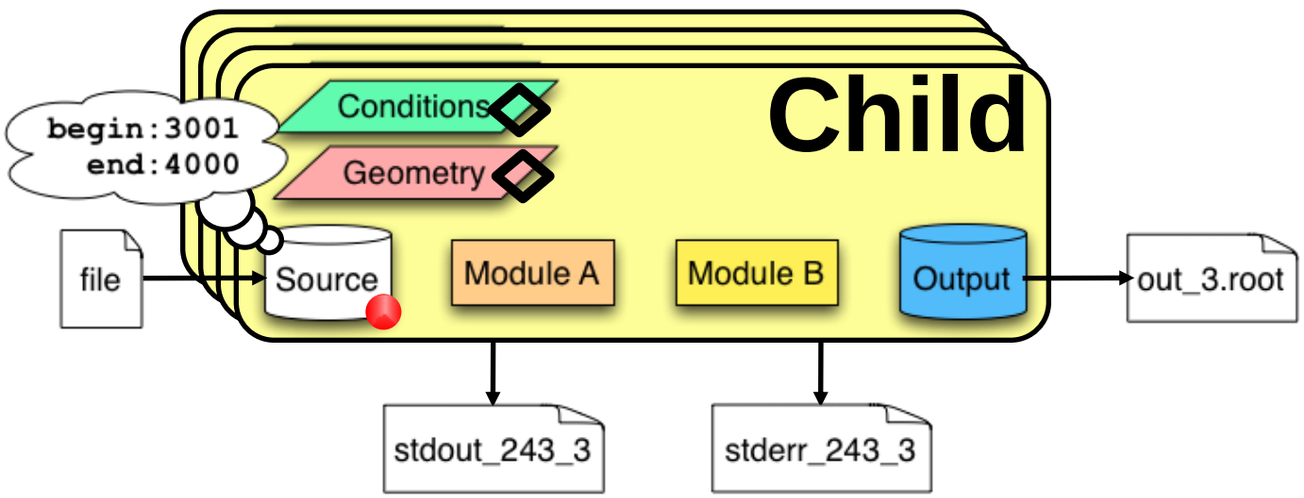

For the Expert:

-one CMSSW configuration file

-Select number of children via parameter

-Write out one file per child

-Provide one FrameworkJobReport.xml per child and one master xml

- **Parent**
  - Reads configuration and loads modules
  - Configuration says how many children and # events/child
  - Opens input file and reads first run
  - modules are not called
  - Pre-fetches conditions, calibrations and geometry
  - Sends message to all modules that forking is going to happen
  - source closes file
  - Forks
- **Children**
  - Redirects stdout and stderr to own files whose names contain parent PID and child #
  - Send messages to modules saying process is child X
  - Output modules append child # to file names
  - Sources calculate their event ranges to process (no IP communication) and re-open the file
  - Process events in child's start/end range normally

Multi-Core WMAgent job

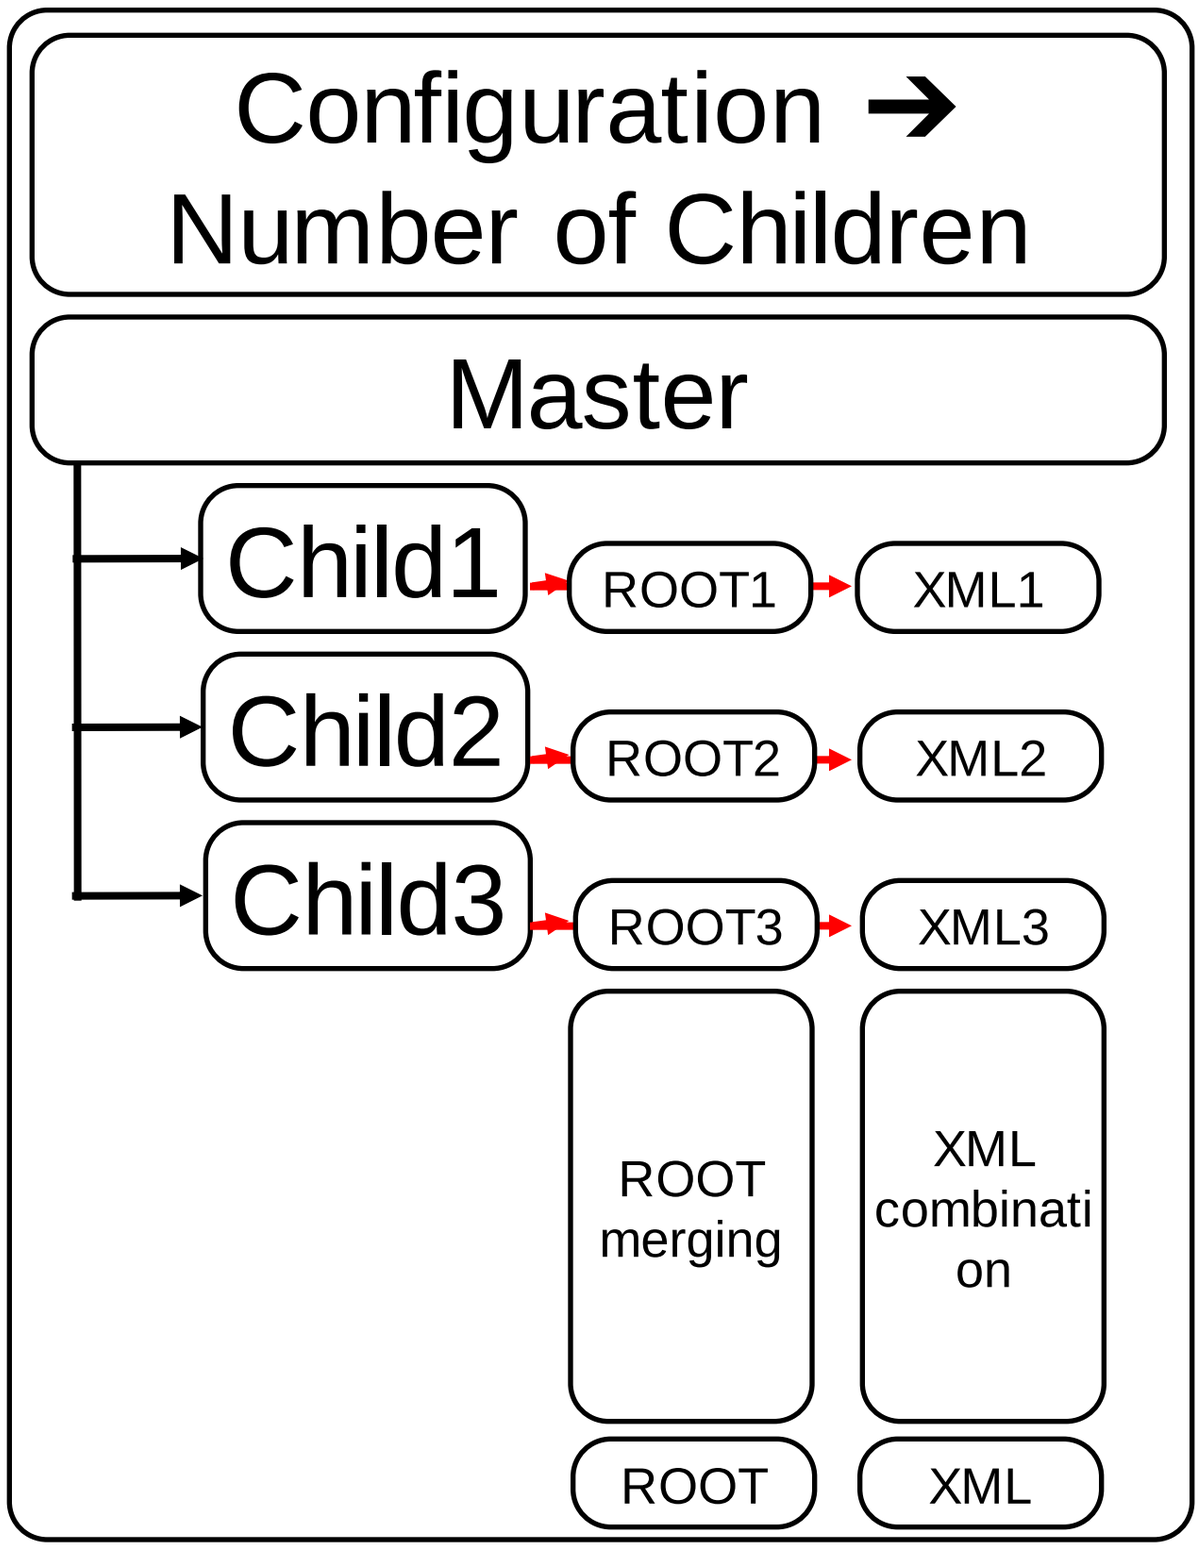

- JobWrapper configures number of children
  - Either via workflow settings
  - Or using /proc/cpuinfo to use the whole node
- JobWrapper executes single CMSSW job producing master xml file and multiple FrameworkJobReport.xml and output files
- JobWrapper merges all ROOT files and stages it out to MSS and also combines all xml into one

Multi-Core Memory Performance

8 Single-core jobs

One 8-core job

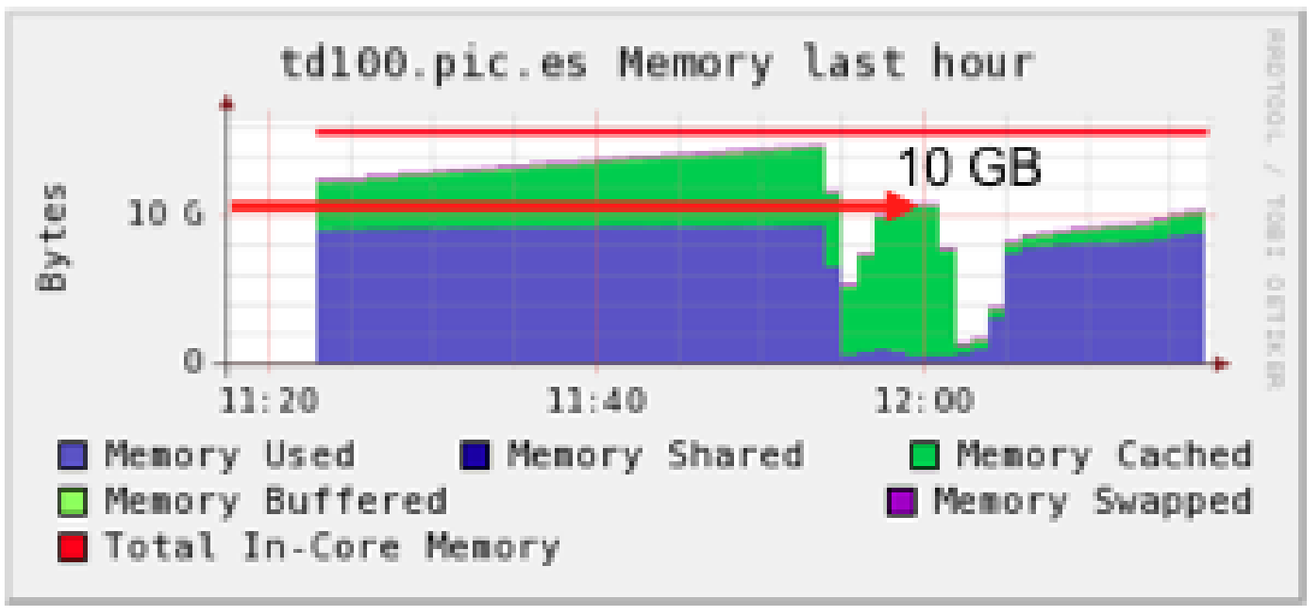

- Clear memory gain (~20%) with multi-core processing
  - 8-core job, ~2 hour long, ~9 GB total memory used by the machine
  - Reported in framework job report for each processing child: VSIZE: 2 GB, RSS: 1.5 GB, PSS: 900 MB
  - Parent process also consumes some memory
- To compare with 8 simultaneous single-core jobs, same workflow, ~11 GB

```
top - 09:20:35 up 1 day, 15:37,  3 users,  load average: 7.96, 5.07, 2.20
Tasks: 158 total,   9 running, 148 sleeping,   0 stopped,   1 zombie
Cpu(s):  0.7%us,  0.1%sy,  0.0%ni, 99.1%id,  0.1%wa,  0.0%hi,  0.0%si,
Mem:  16437844k total,  9414440k used,  7023404k free,   239148k buffers
Swap:  4192924k total,        0k used,  4192924k free,  1234428k cached

  PID USER      PR  NI  VIRT  RES  SHR S %CPU %MEM    TIME+  COMMAND
13411 cmsdc04  418   0 1917m 1.5g  91m R 101.0  9.6   4:49.59 cmsRun
13414 cmsdc04   25   0 1946m 1.5g  91m R 101.0  9.7   4:49.08 cmsRun
13407 cmsdc04   18   0 1935m 1.5g  91m R  99.0  9.6   4:47.88 cmsRun
13408 cmsdc04   20   0 1934m 1.5g  91m R  99.0  9.7   4:48.14 cmsRun
13409 cmsdc04   21   0 1980m 1.6g  91m R  99.0  9.9   4:41.31 cmsRun
13410 cmsdc04   18   0 1946m 1.5g  91m R  99.0  9.7   4:45.93 cmsRun
13412 cmsdc04   18   0 1947m 1.5g  91m R  99.0  9.7   4:49.49 cmsRun
13413 cmsdc04   25   0 1917m 1.5g  91m R  99.0  9.5   4:49.41 cmsRun
13404 cmsdc04   22   0 1184m 964m 168m S   0.0  6.0   0:43.70 cmsRun
```

Lazy Download

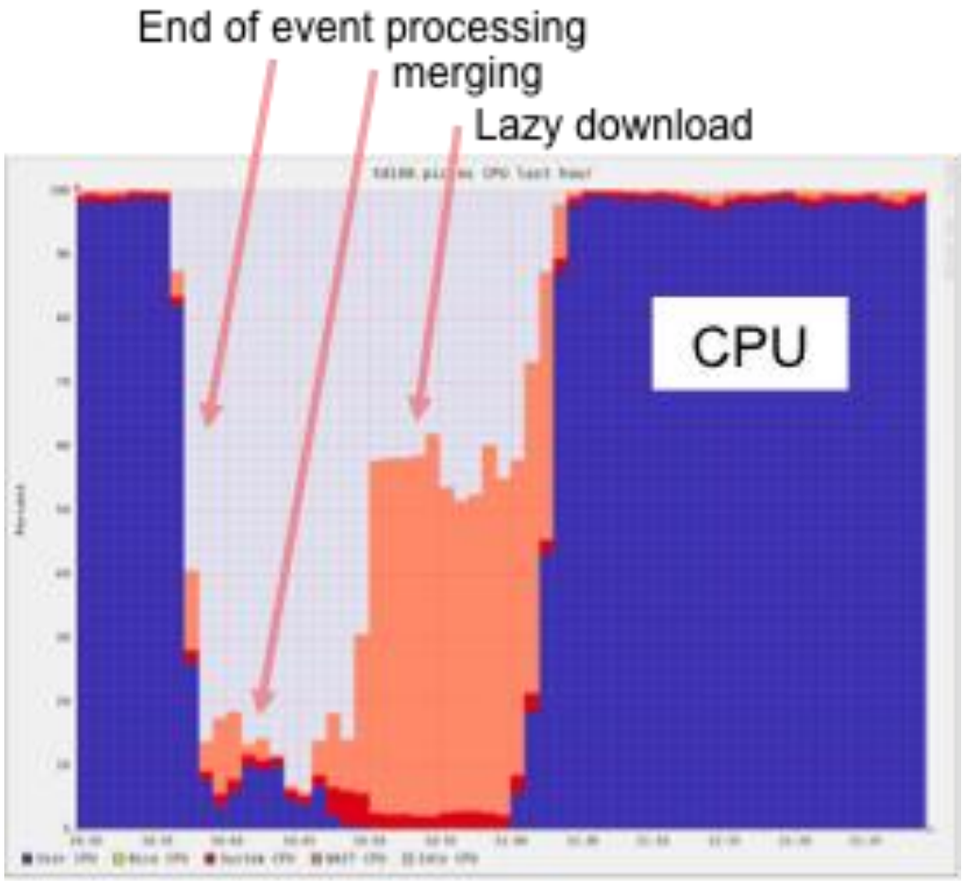

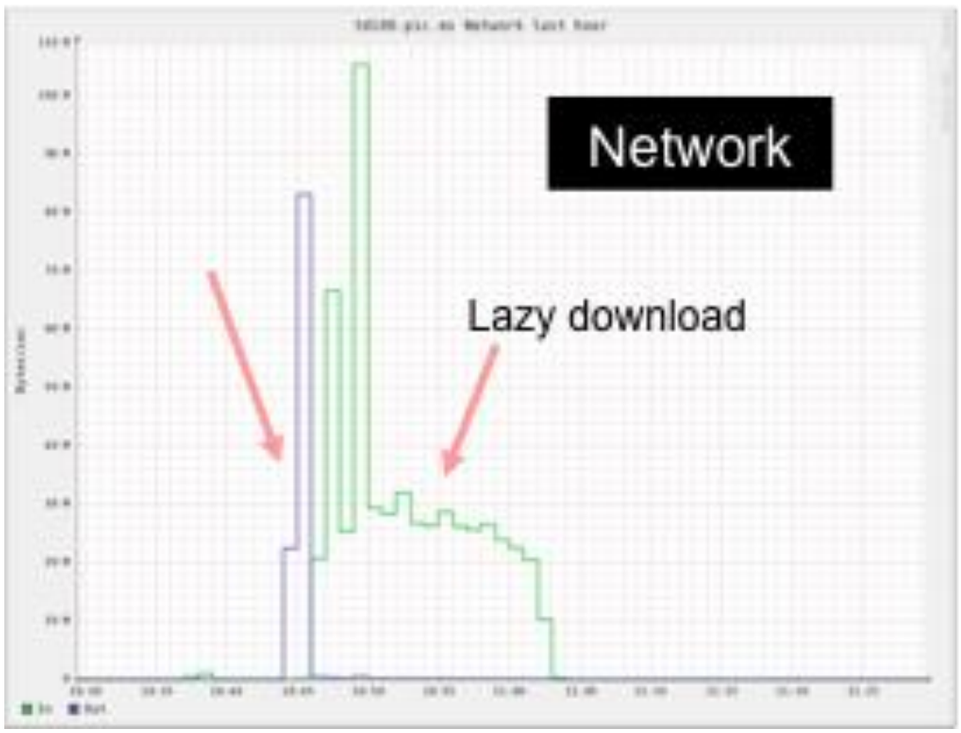

- Lazy download produces a large startup overhead for multi-core jobs
  - Large IO/wait, local disk hammered by processing children downloading input file(s)
- With lazy download off no startup overhead
  - Small overhead due to children processing dispersion, file merging and stage-out (~10 minutes all cores ~idle)

Without Lazy download

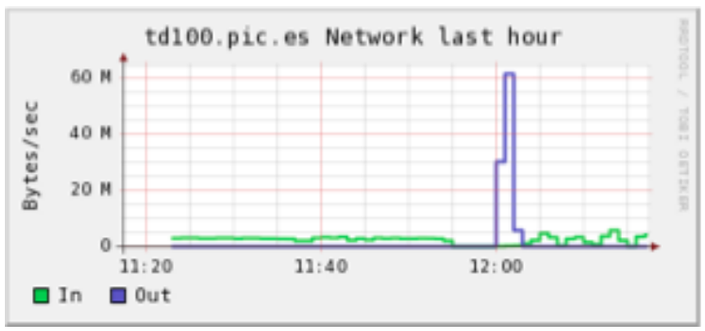

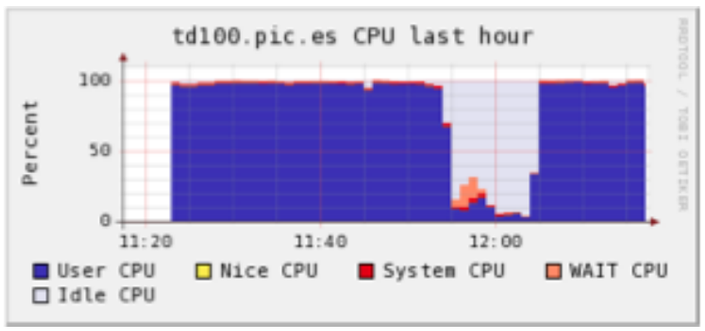

Tests

- Setup
  - Dedicated small whole-node queues at all T1 sites and one T2 site
  - Configured as separate HTPC queues in glideIn factory
  - New: Dynamic 8-core queue at Purdue

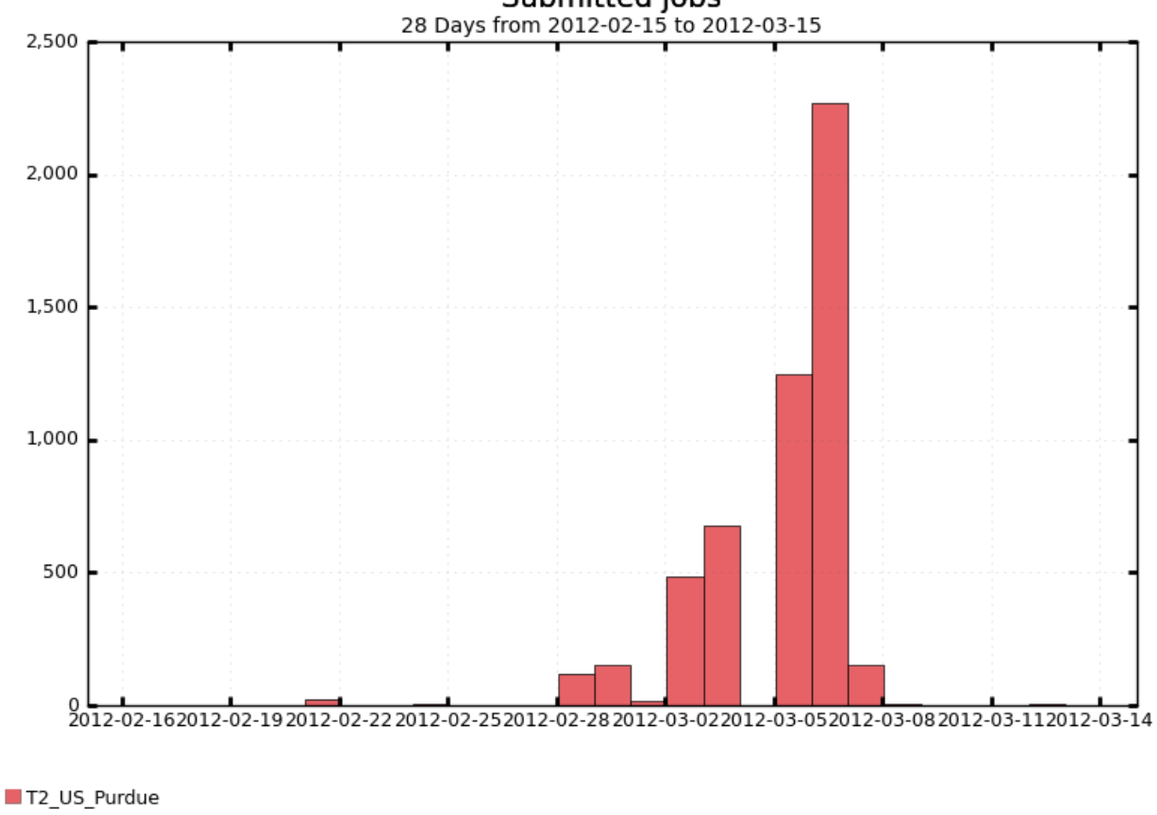

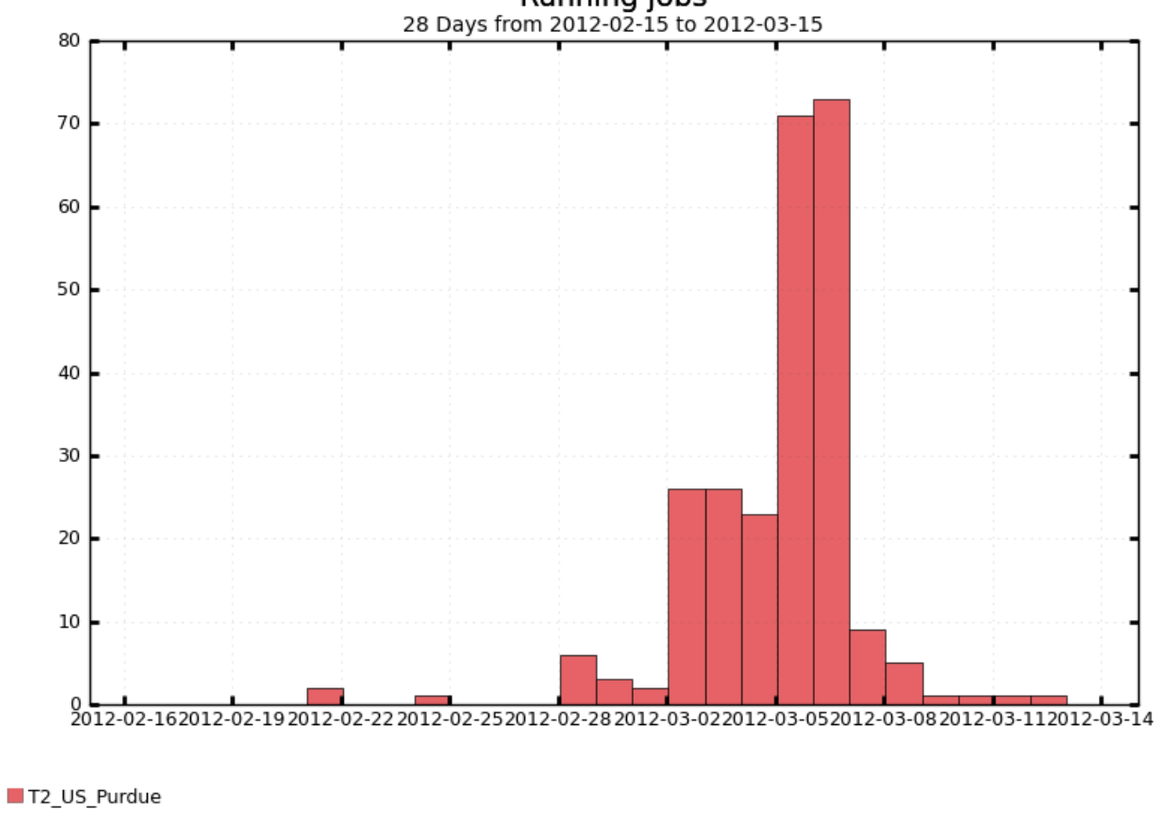

Summary

- Multicore provides 20% memory gain compared to single core jobs
  - Asynchronous merging very much reduced
  - Number of processing jobs very much reduced

- Dedicated queues at Tier-1 sites used for initial tests
  - Tier-1 sites will not like to move parts of their resources to multi-core usage

- Dynamic multi-core slots at Purdue are working and simple to use
  - ~5k jobs run with about 70 jobs in parallel (70x8 cores!)
  - Preferred solution of Tier-1 sites to use multi-core jobs, but still questions about accounting (for example when draining a node to have enough cores for 1 job)
- T1_DE_KIT will be providing similar queues with 4 and 8 cores available per slot very soon
- WLCG TEG recommendation
  - Number of cores configurable during job submission and site provides dynamically access to multi-core slots

Summary (Cont.)

- Computing Operations reorganized beginning of 2012

- 2011 saw a lot of activity, Tier-2 sites backbone of CMS analysis and a vital part of the whole analysis chain
  - Thanks for exceptional performance in 2011!

- 2012 will be a busy year with even more data as 2011 and a lot of analysis activity

- Data manager role is very important to guarantee space availability and avoid overfilling of site storage

- /store/user areas become more and more important, especially when CERN Castor user area is closed

- 2012 will be very exciting, so expect a lot of activity!