HTCondor: CMS Experience

- *Farrukh Aftab Khan*
- *(on behalf of CMS Submission Infrastructure team)*

- May 19, 2016

Compact Muon Solenoid experiment at CERN's LHC

Collaborators

Apologies for the lack of logos!

CMS is a huge collaboration of multiple institutes and universities all over the world

We wouldn't have been able to do what we do without their support and effort!

HTCondor: CMS Experience

This talk covers an overview of how CMS utilizes HTCondor in its grid computing infrastructure

So basically…

What is CMS?! What is our use case? – i.e. what are we trying to solve? What is our solution? – i.e. the CMS global pool! How well is the solution working?

HTCondor: CMS Experience

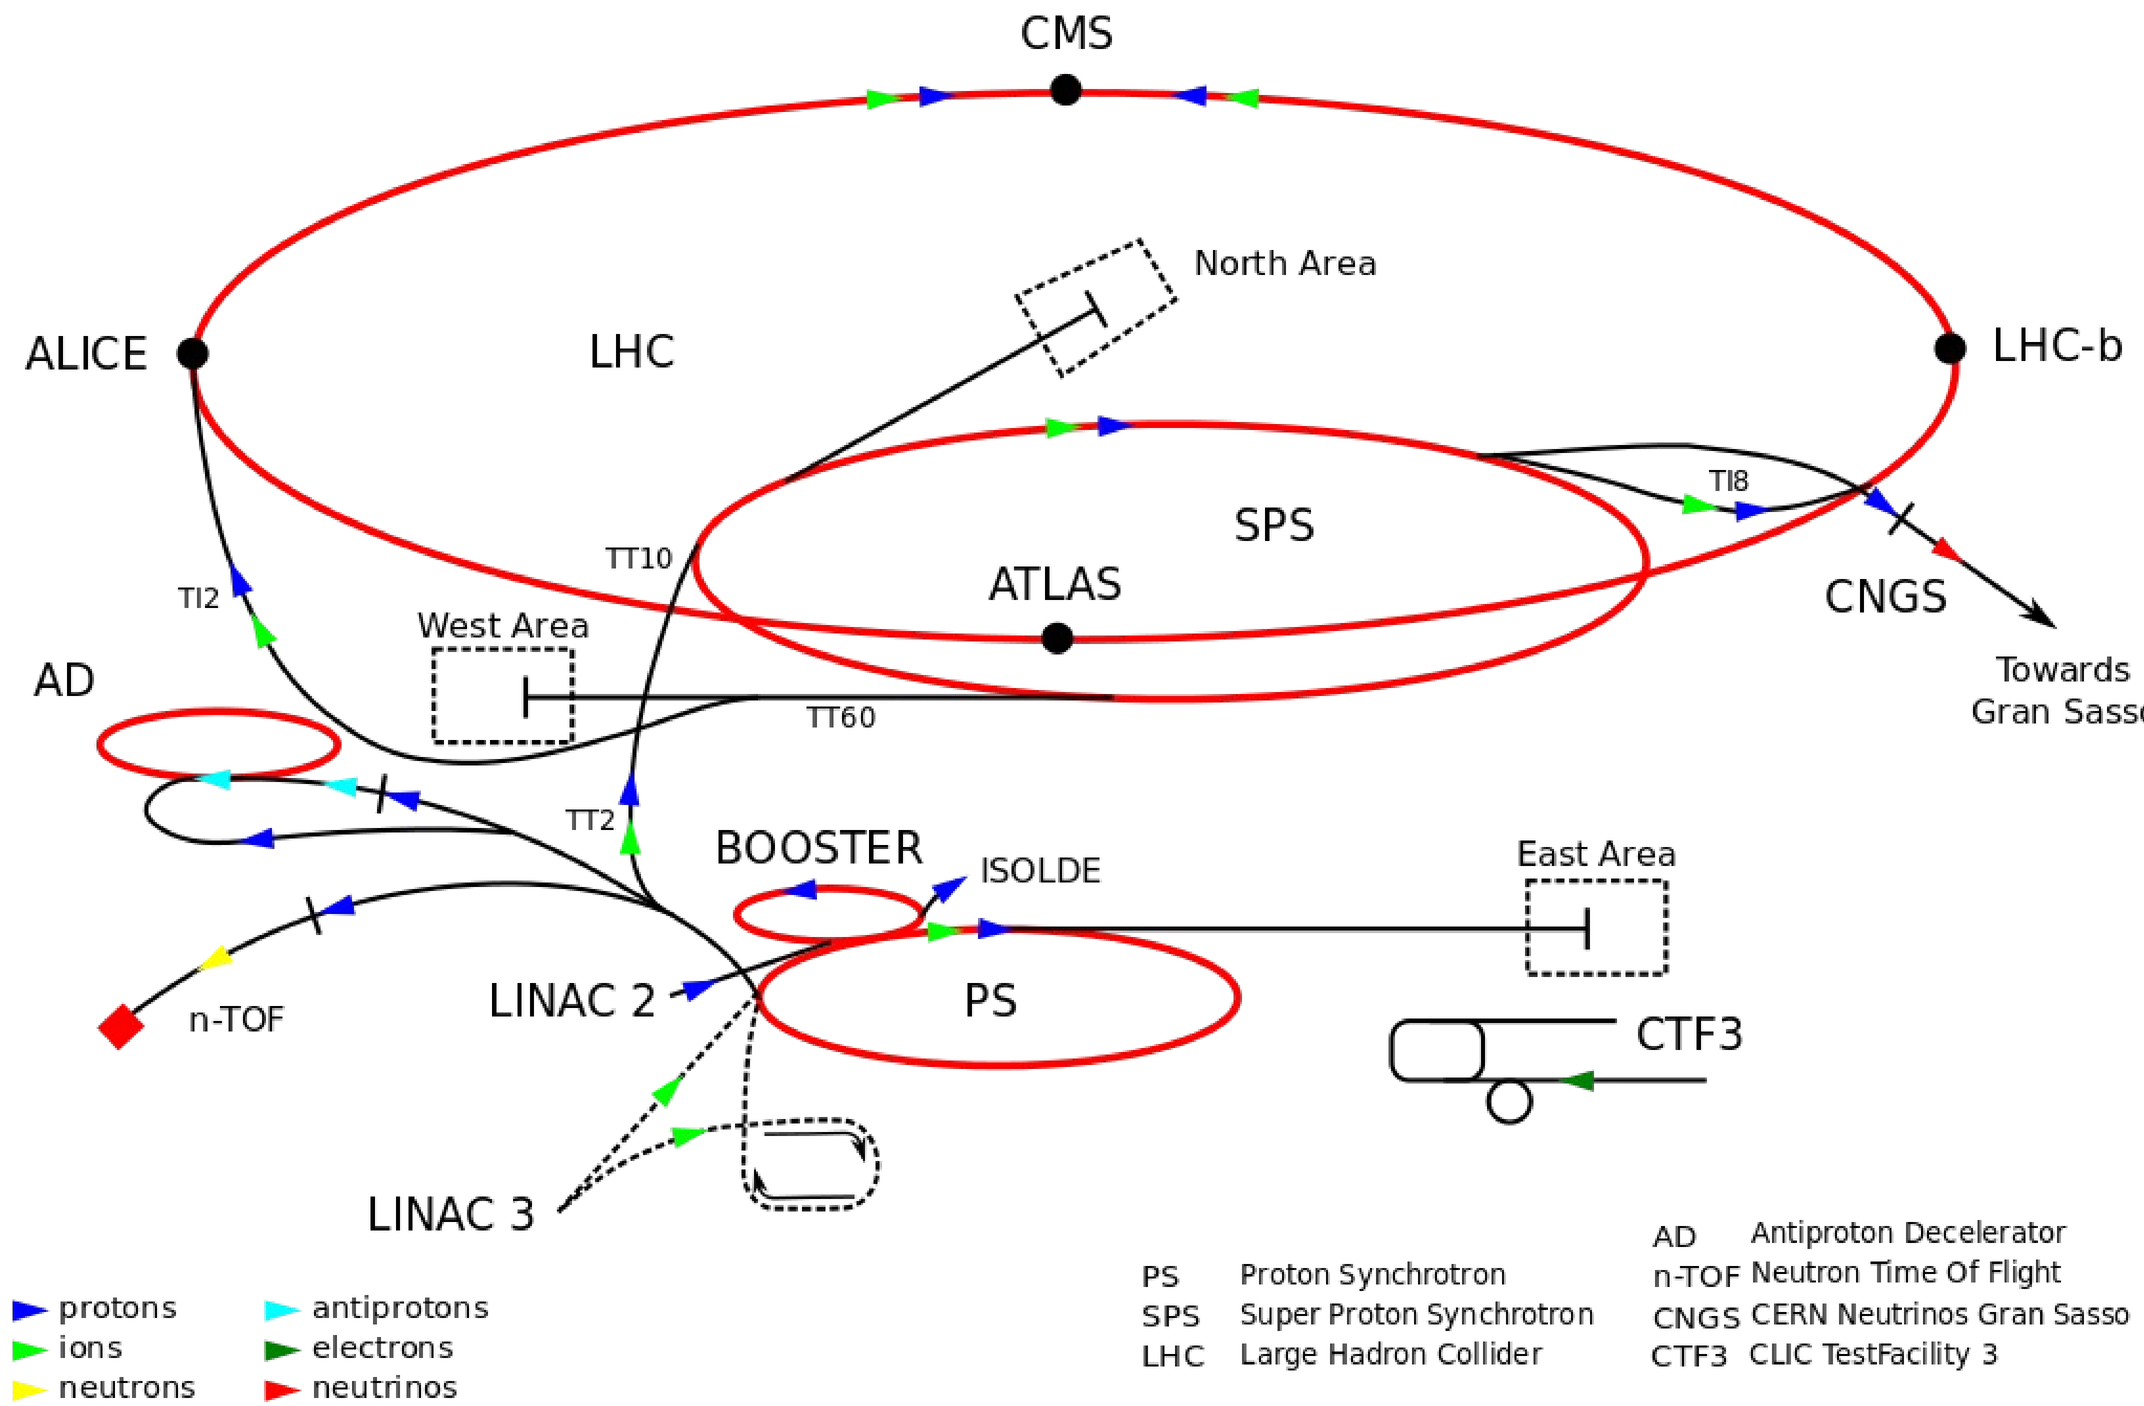

HTCondor: CMS Experience

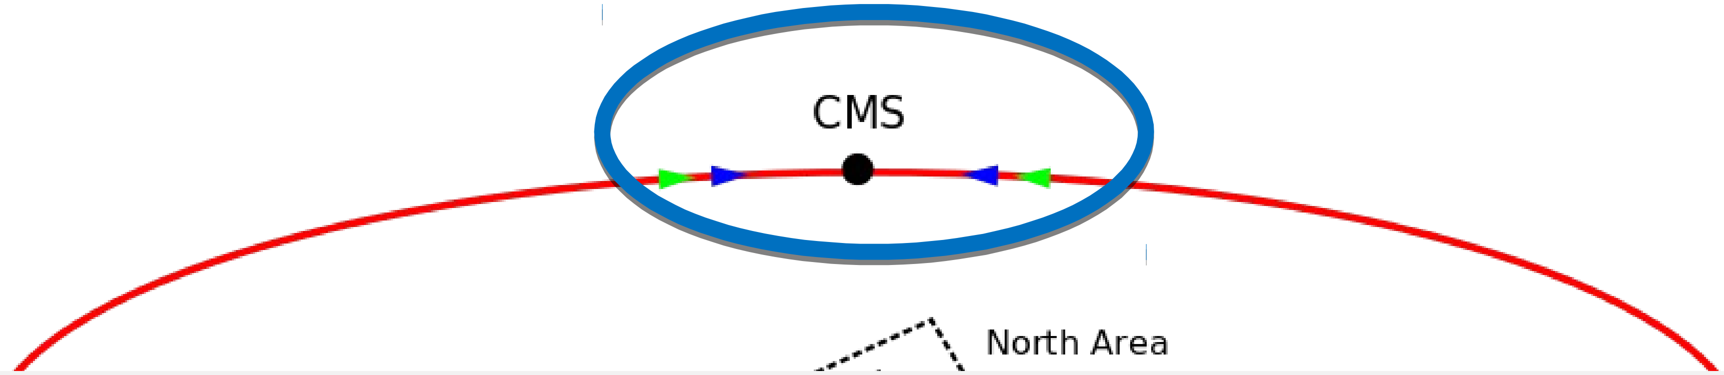

- CMS (Compact Muon Solenoid) is a particle physics experiment at CERN LHC
- **Millions** of protons collide every second at a very high energy
- CMS detector records collisions which the experiment deems interesting. These are called “events”

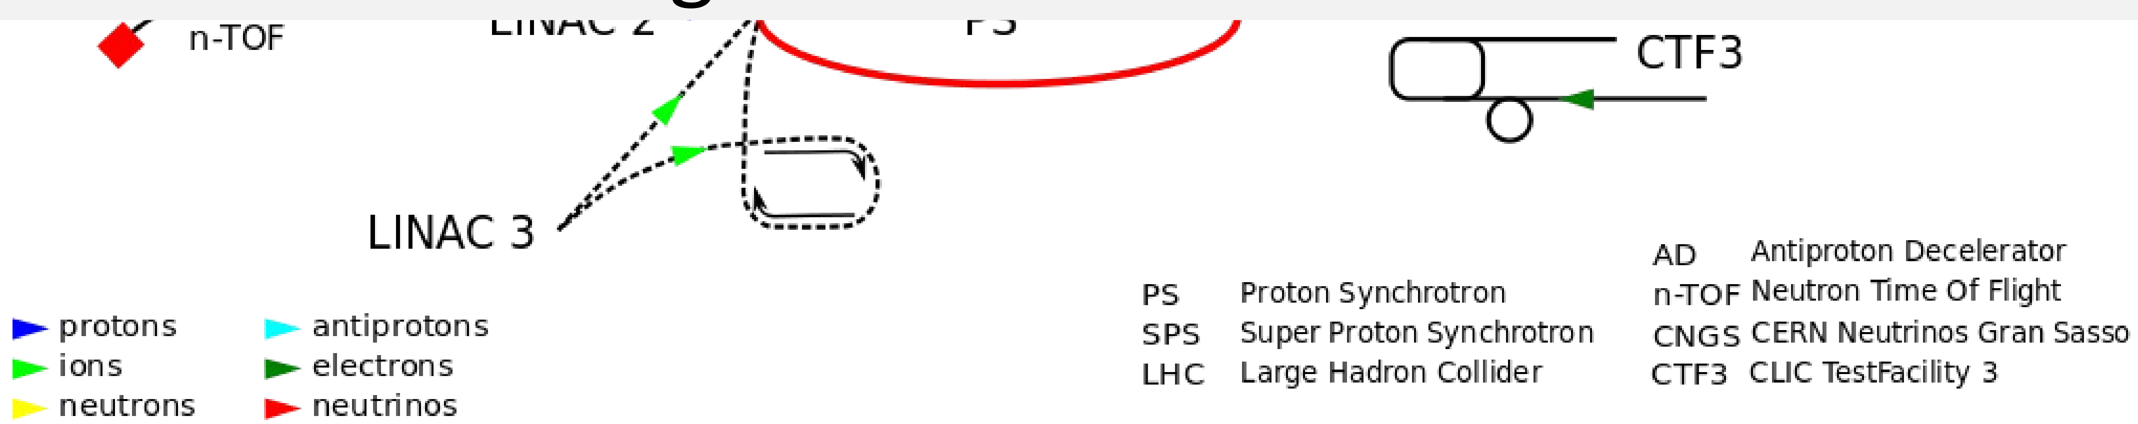

Compact Muon Solenoid experiment at CERN's LHC

HTCondor: CMS Experience

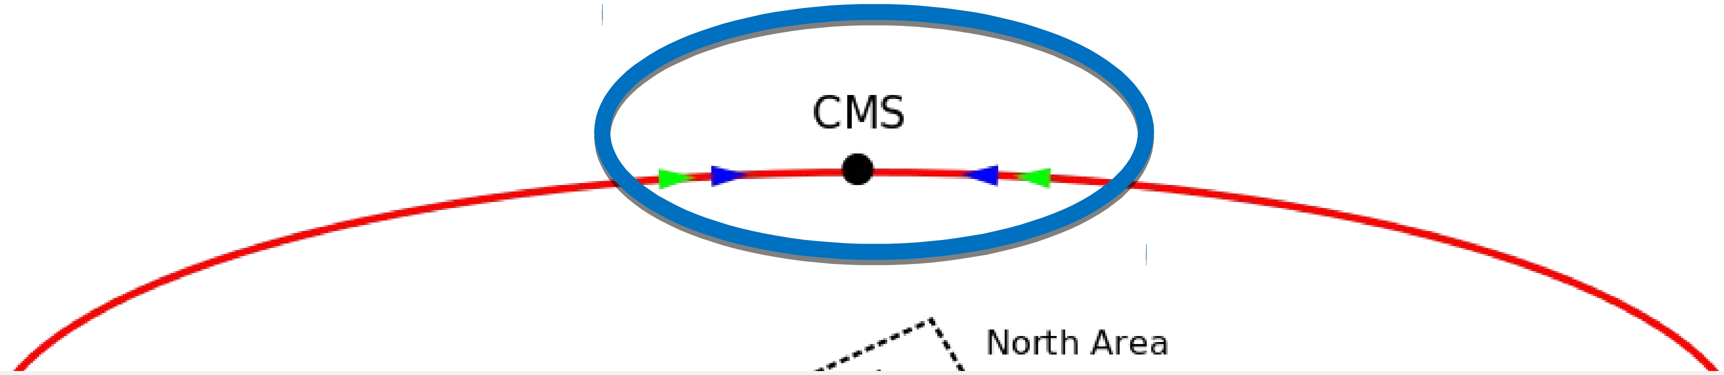

- CMS keeps **billions** of such events
- Disentangling what happened in **each individual** event takes ~minutes of CPU time

Billions of events each needing several {tens, hundreds, thousands} of minutes = need for high throughput computing!

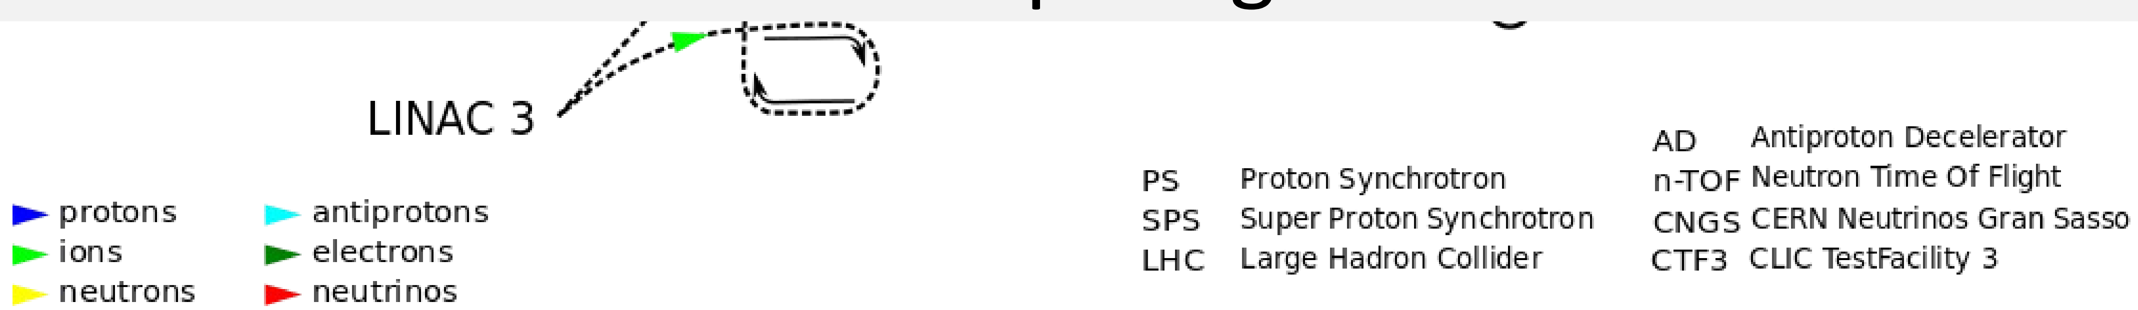

HTCondor: CMS Experience

- CMS runs various different types of workflows (GEN-SIM, MC-RECO, DATA-RECO, USER Analysis) submitted
- Compared to the last run, the beam energy is almost double i.e. more collisions, more data to analyze/process, more jobs
- More jobs also imply the need for more 'machines' to process them

HTCondor: CMS Experience

- CMS runs various different types of workflows (GEN-SIM, MC-RECO, DATA-RECO, USER Analysis) submitted
- Compared to the last run, the beam energy is almost double i.e. more collisions, more data to analyze/process, more jobs
- More jobs also imply the need for more 'machines' to process them

Enforcing global priorities

Flexible provisioning of resources

CMS Global pool

A global queue was formed in 2014 to simplify operations and enable global prioritization, especially in order to flexibly use the entire set of resources available to CMS for any kind of workflow

The submission infrastructure is pilot based and is a combination of:

- HTCondor
- GlideinWMS (Glidein based Workflow Management System)

Compact Muon Solenoid experiment at CERN's LHC

HTCondor: CMS Experience

CMS Global pool

- There are two different tenants of the global pool: central production and user analysis.
- The analysis and central production use cases rely on agents (ultimately daemons written in python) collecting, building and submitting jobs
  - CRAB3 (CMS Remote Analysis Builder) collects user jobs and handles job submission, retries, data stage out
  - WMAgent (Workflow Management agent) handles requests from physics groups for simulation or data reprocessing
- Site whitelists, requested memory, requested number of cores and required wall time determines where a job gets to run. Later, a global priority determines who runs first.

HTCondor: CMS Experience

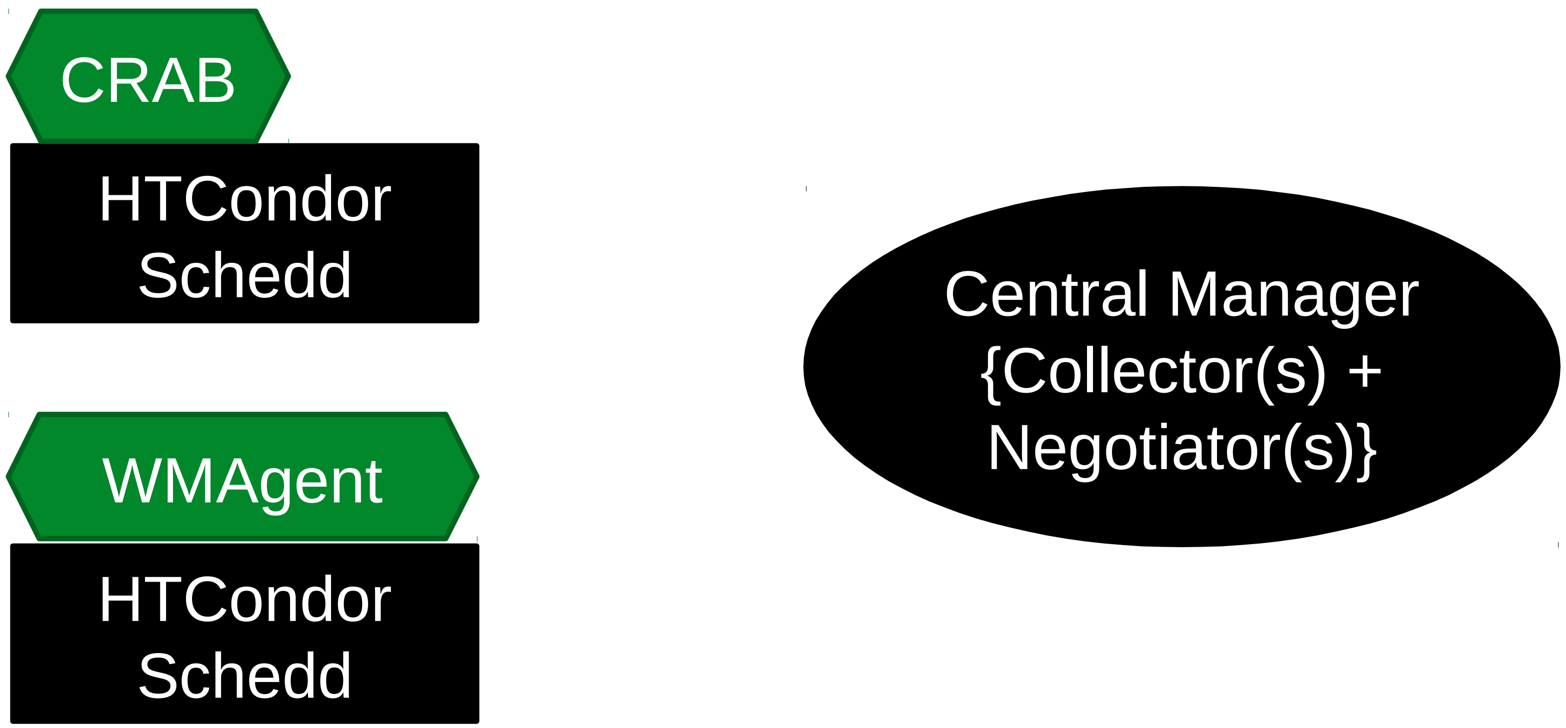

CMS Global pool - GlideinWMS

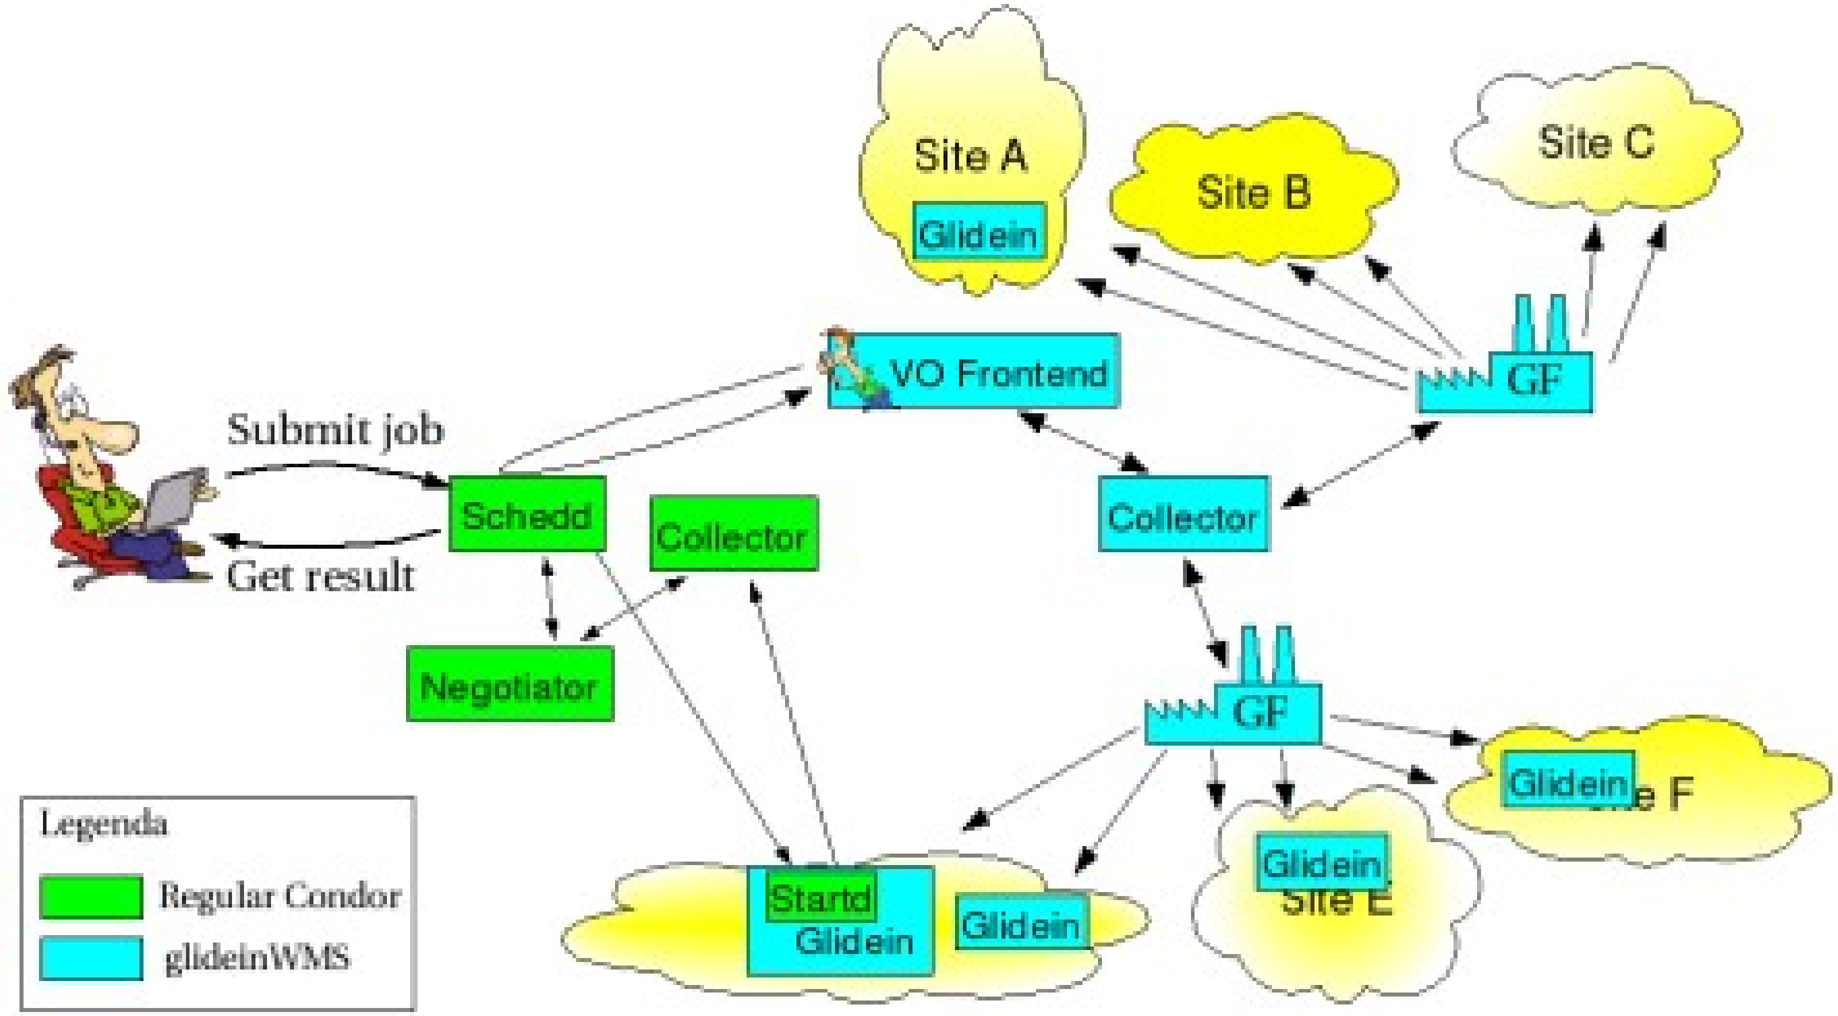

GlideinWMS is a pilot based workflow management system that works on top on HTCondor . Independent of the underlying batch system of a site, from the VO perspective glideinWMS constructs a uniform HTCondor pool — essential for making a global pool.

(*) glideinwms.fnal.gov

HTCondor: CMS Experience

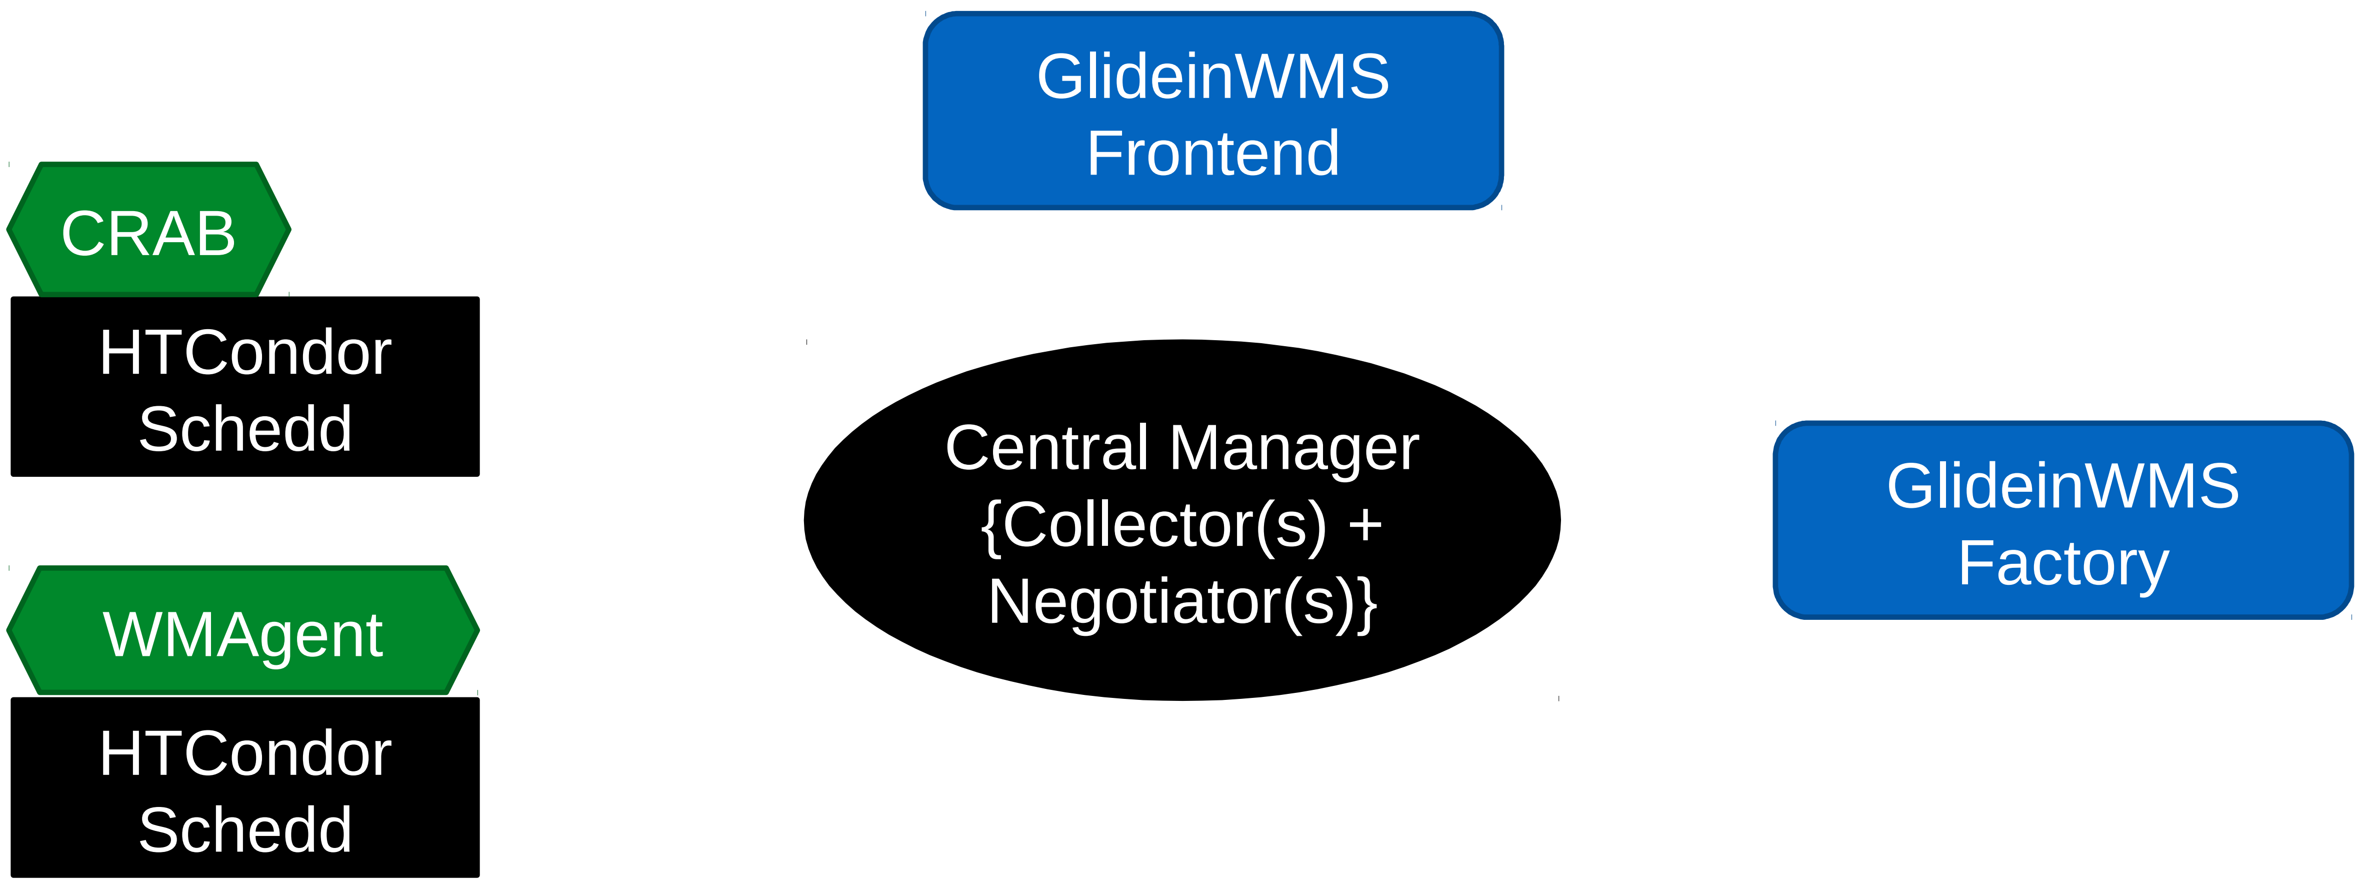

HTCondor: CMS Experience

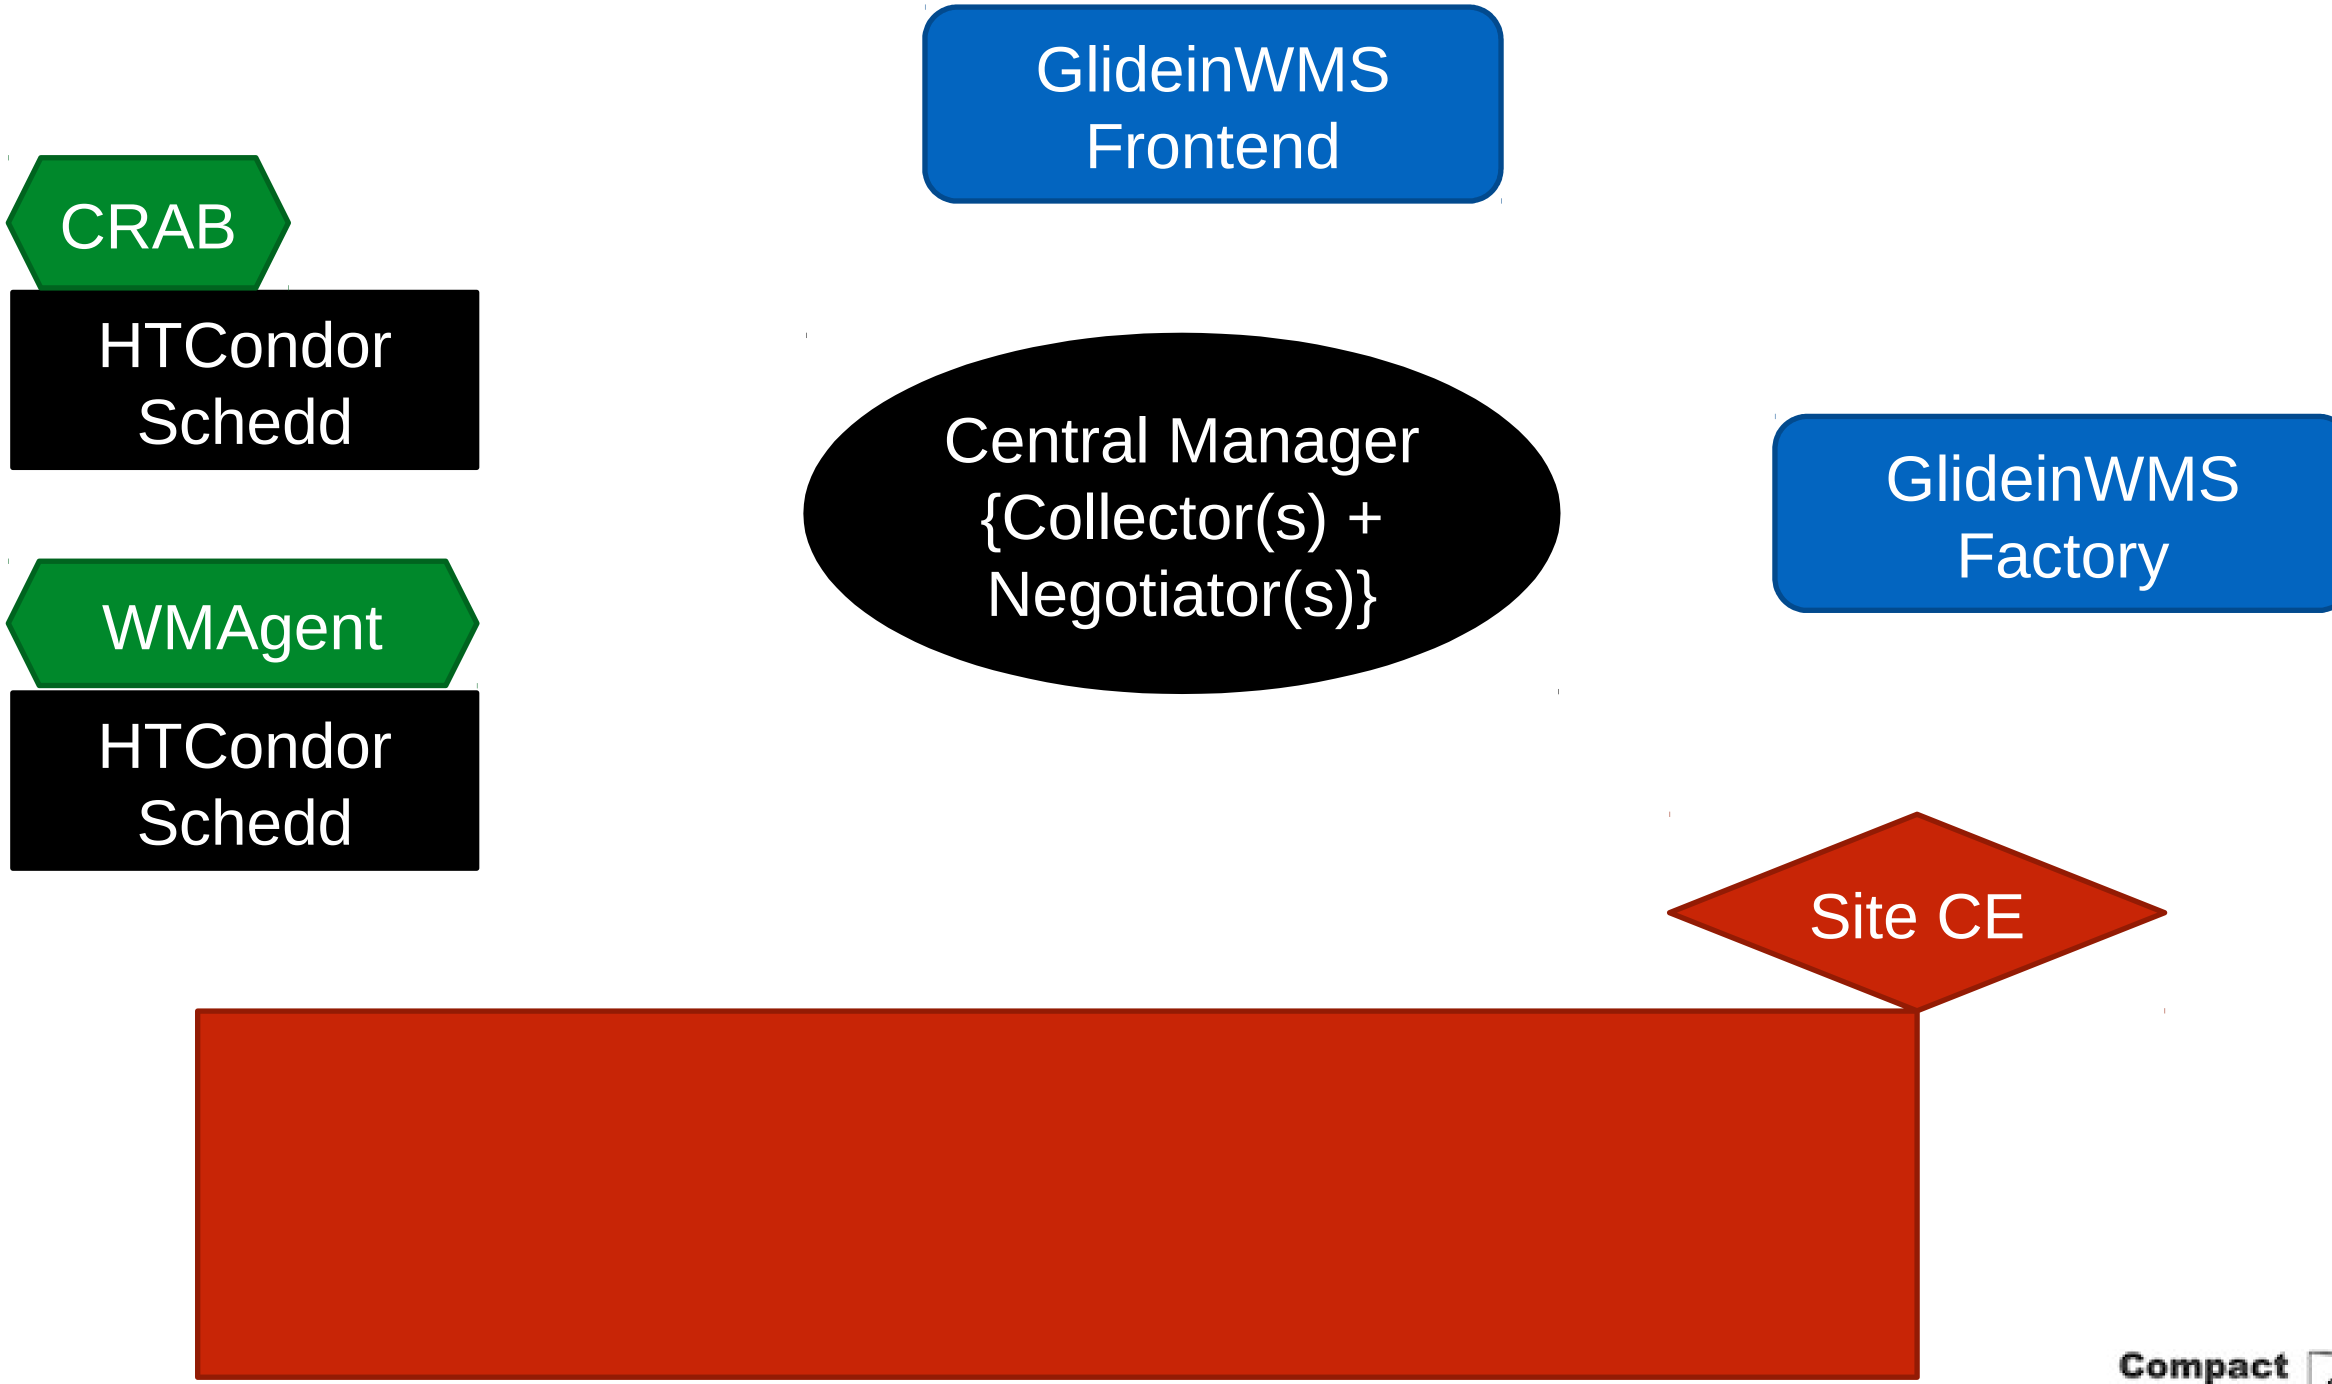

Compact Muon Solenoid experiment at CERN's LHC

HTCondor: CMS Experience

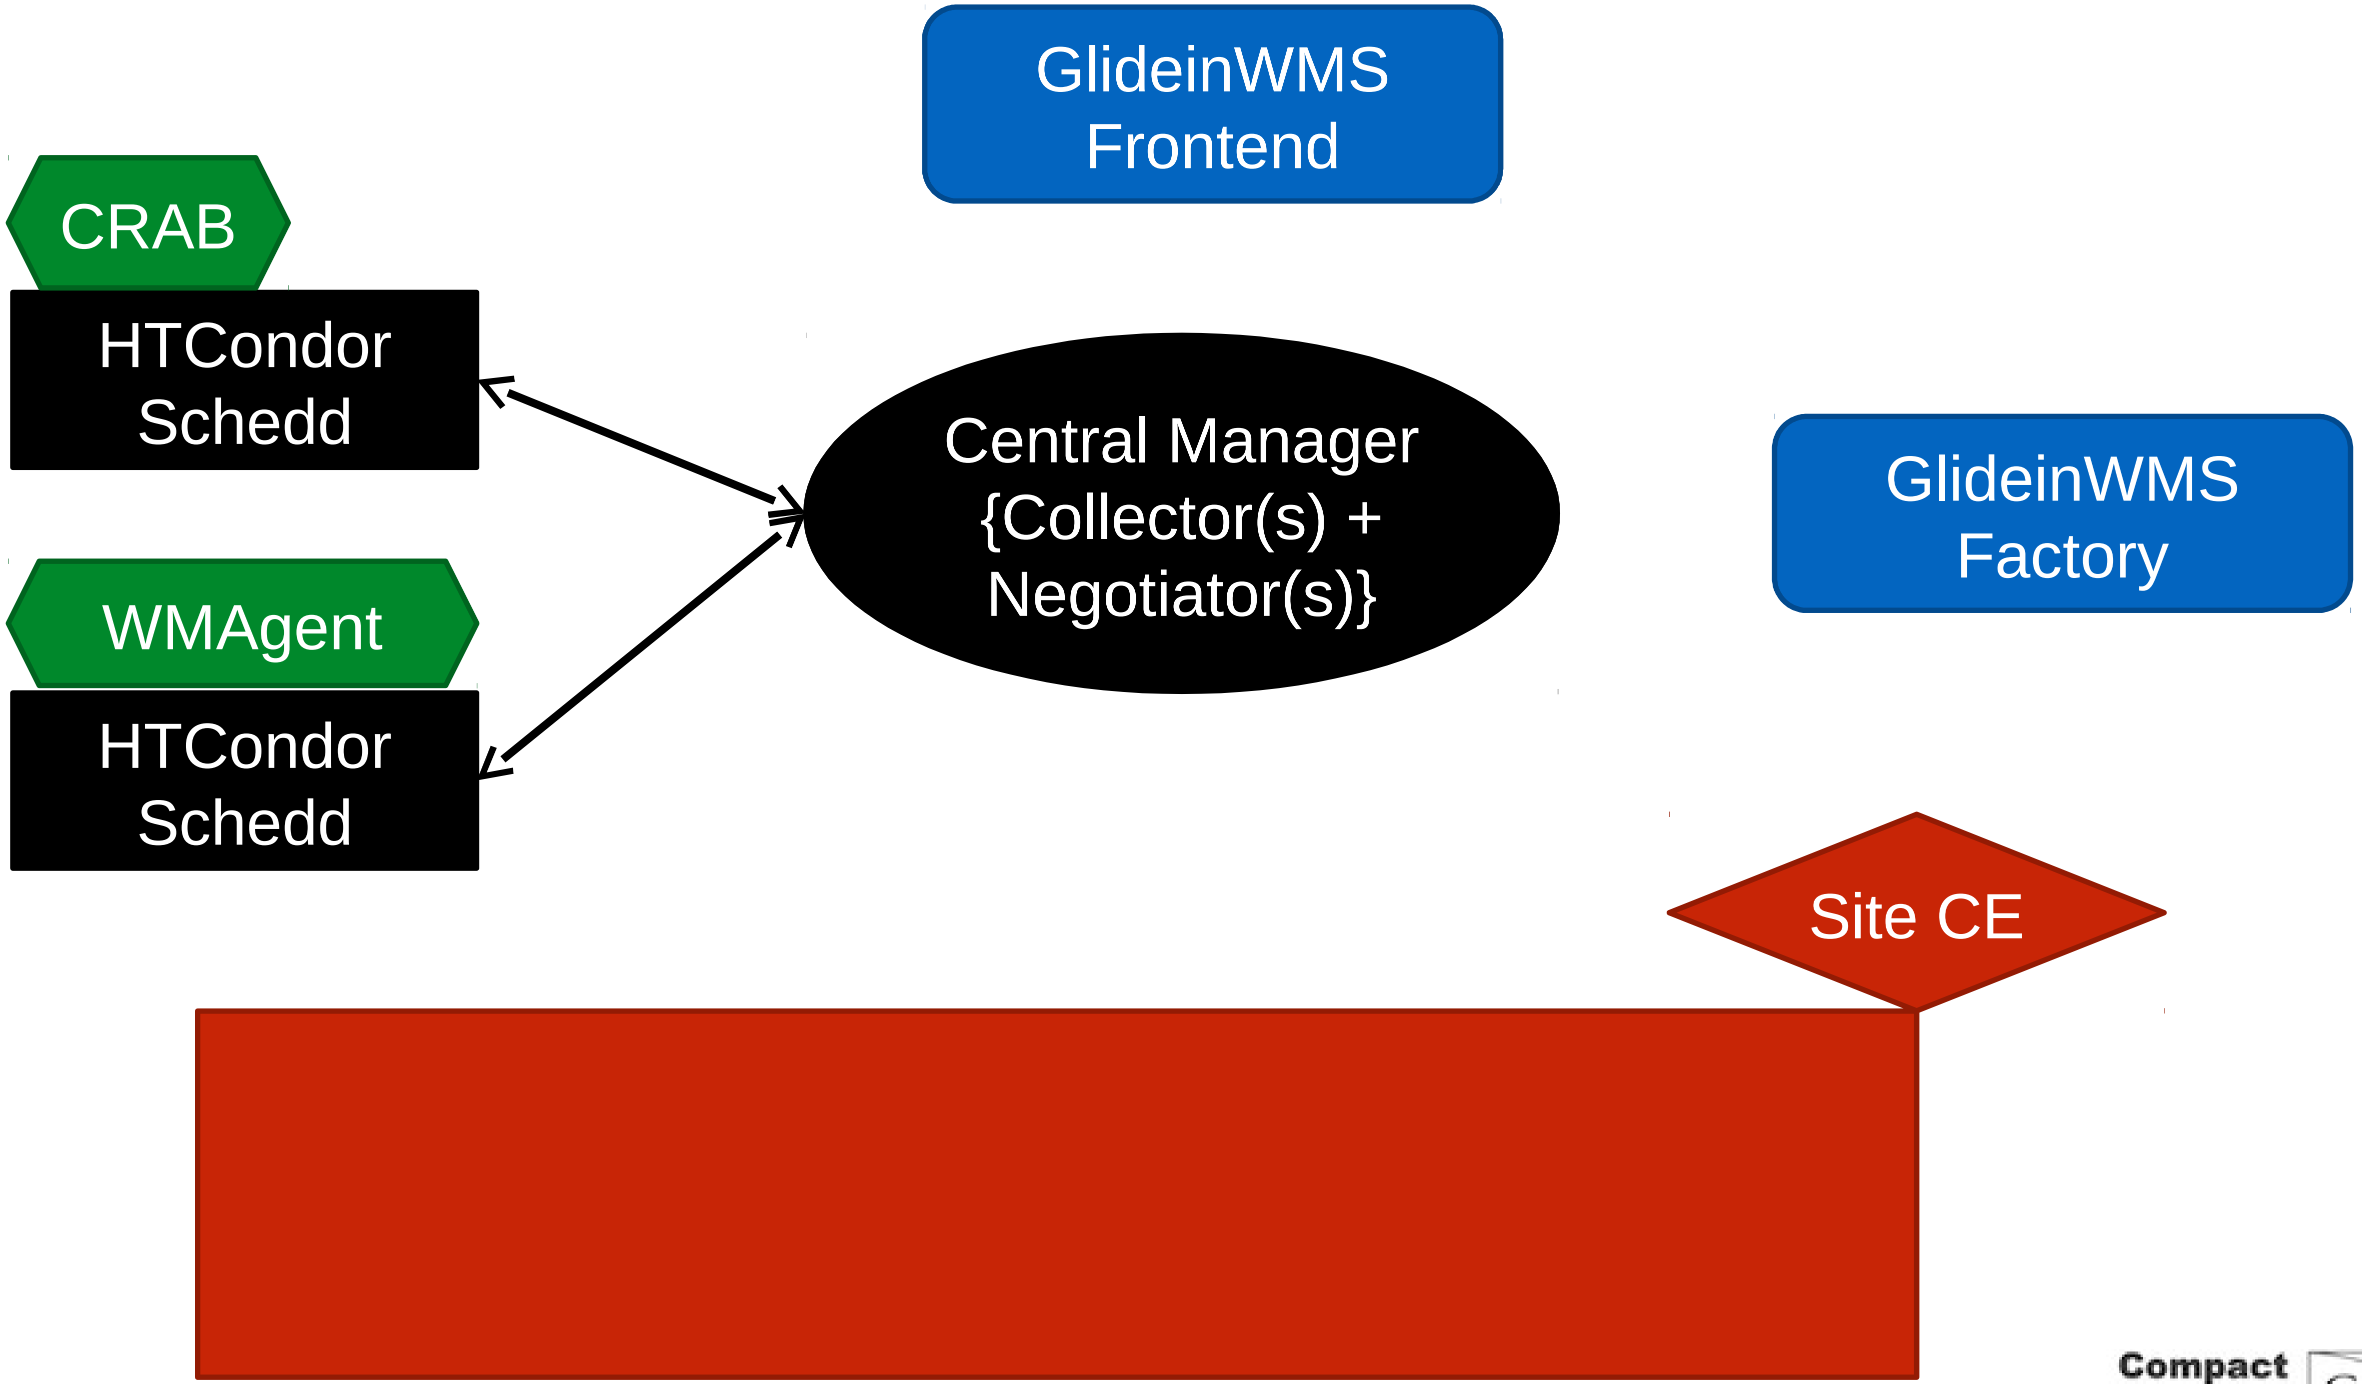

HTCondor: CMS Experience

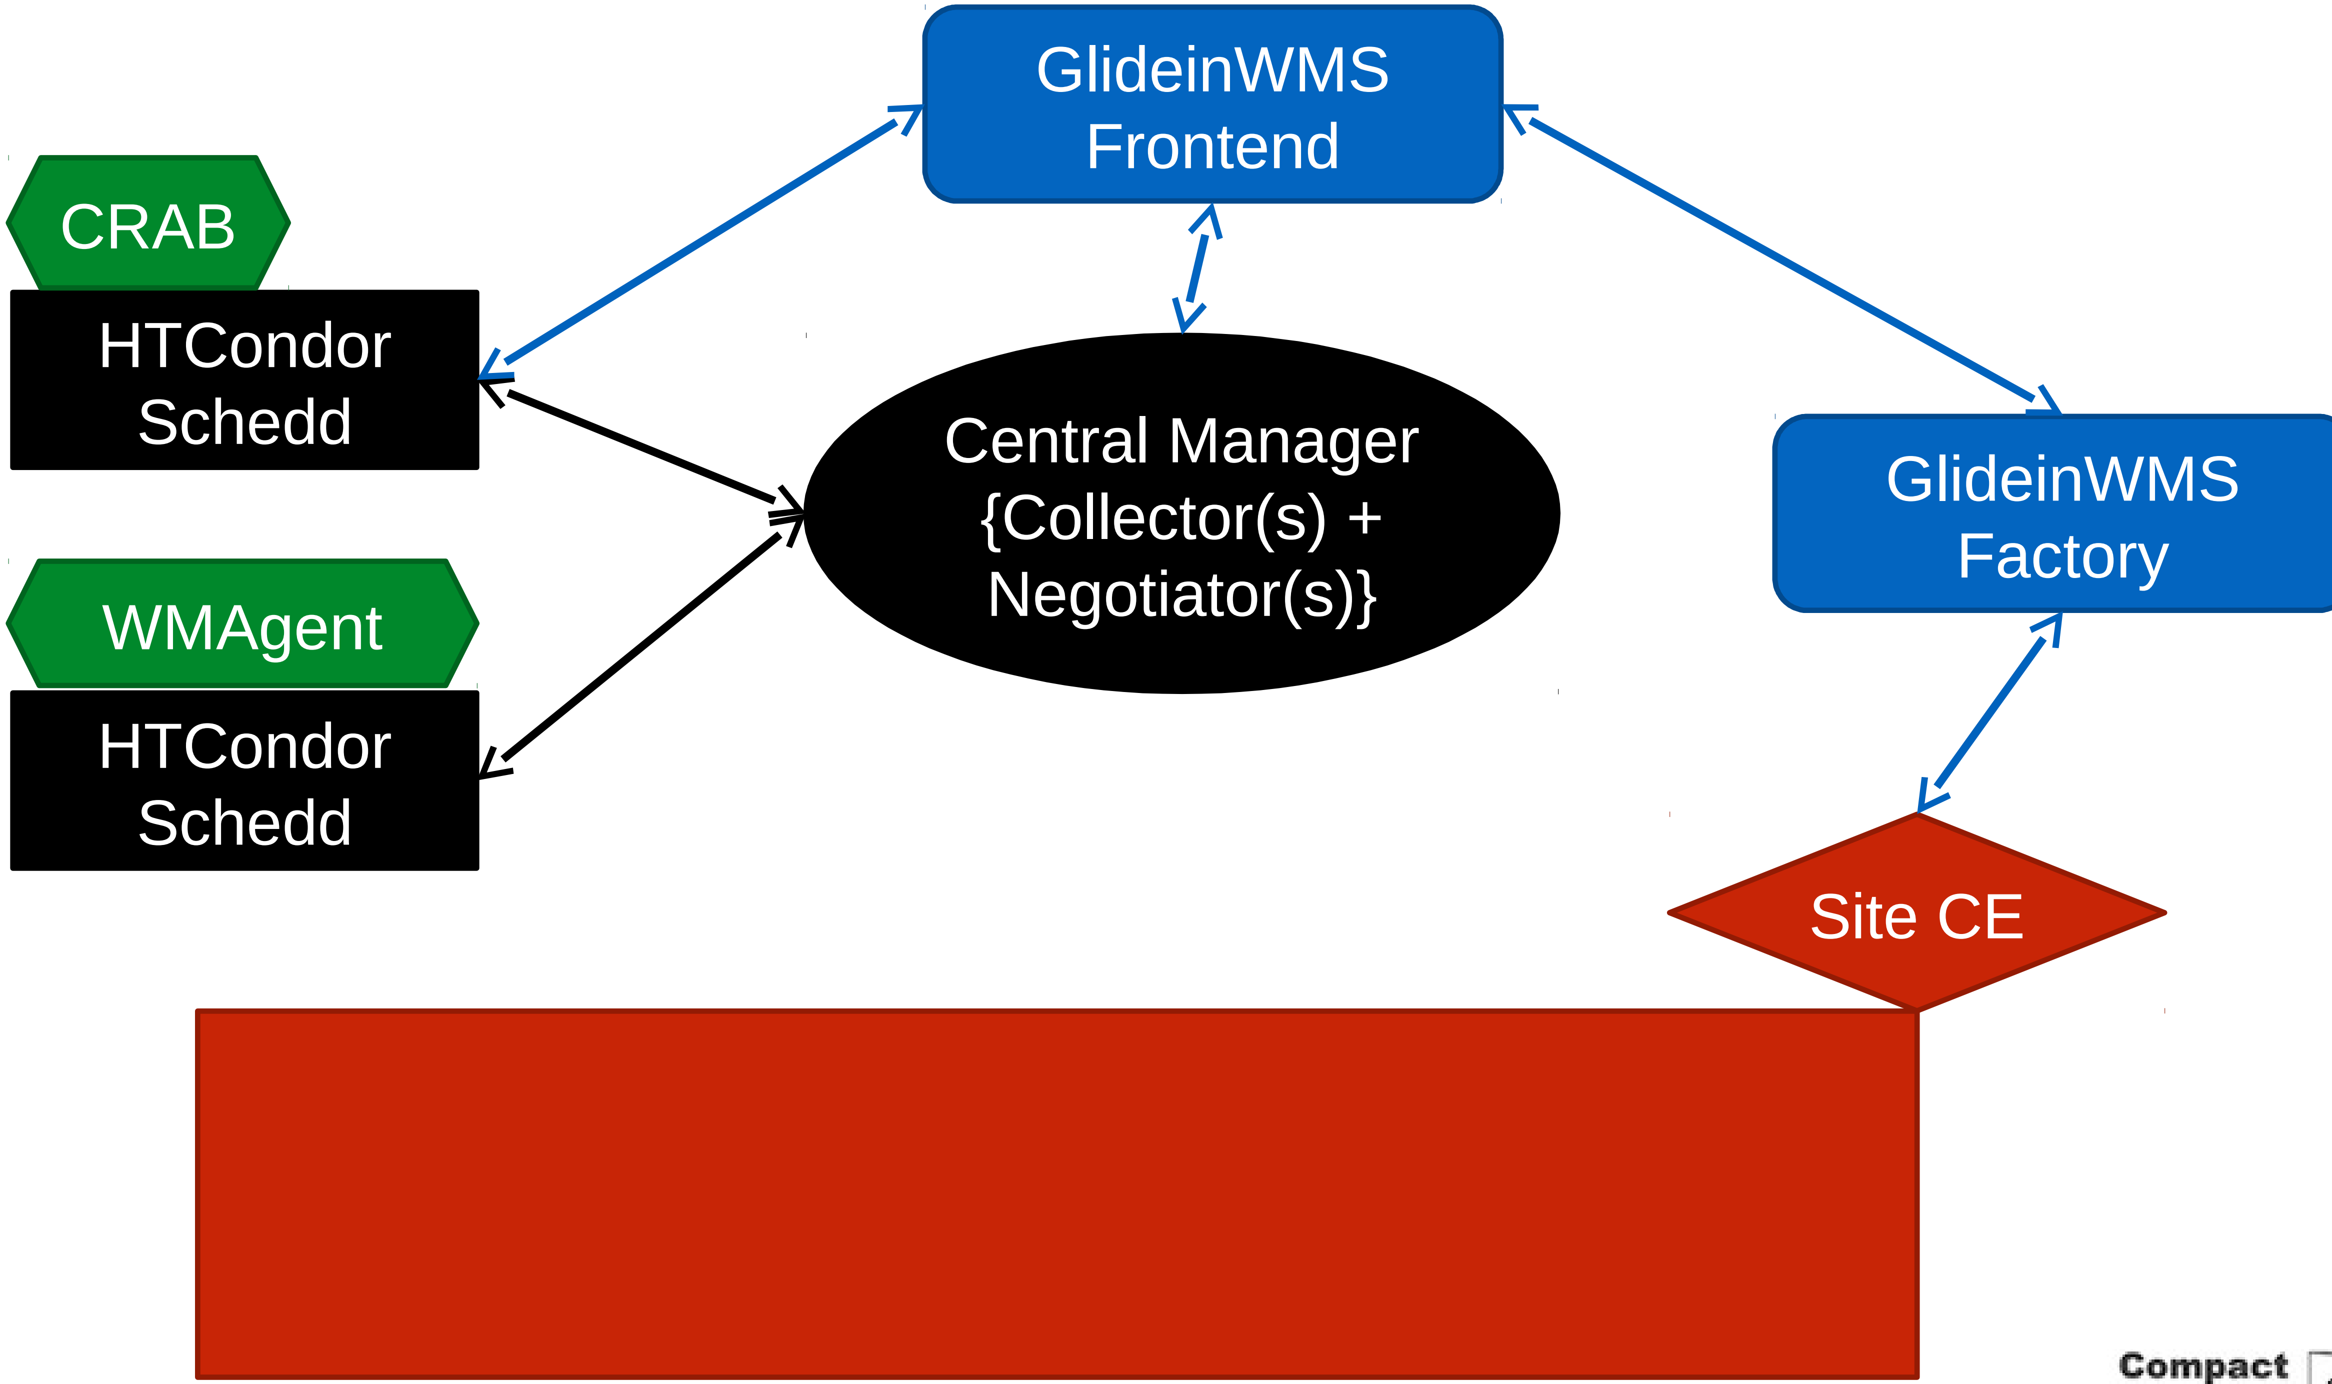

HTCondor: CMS Experience

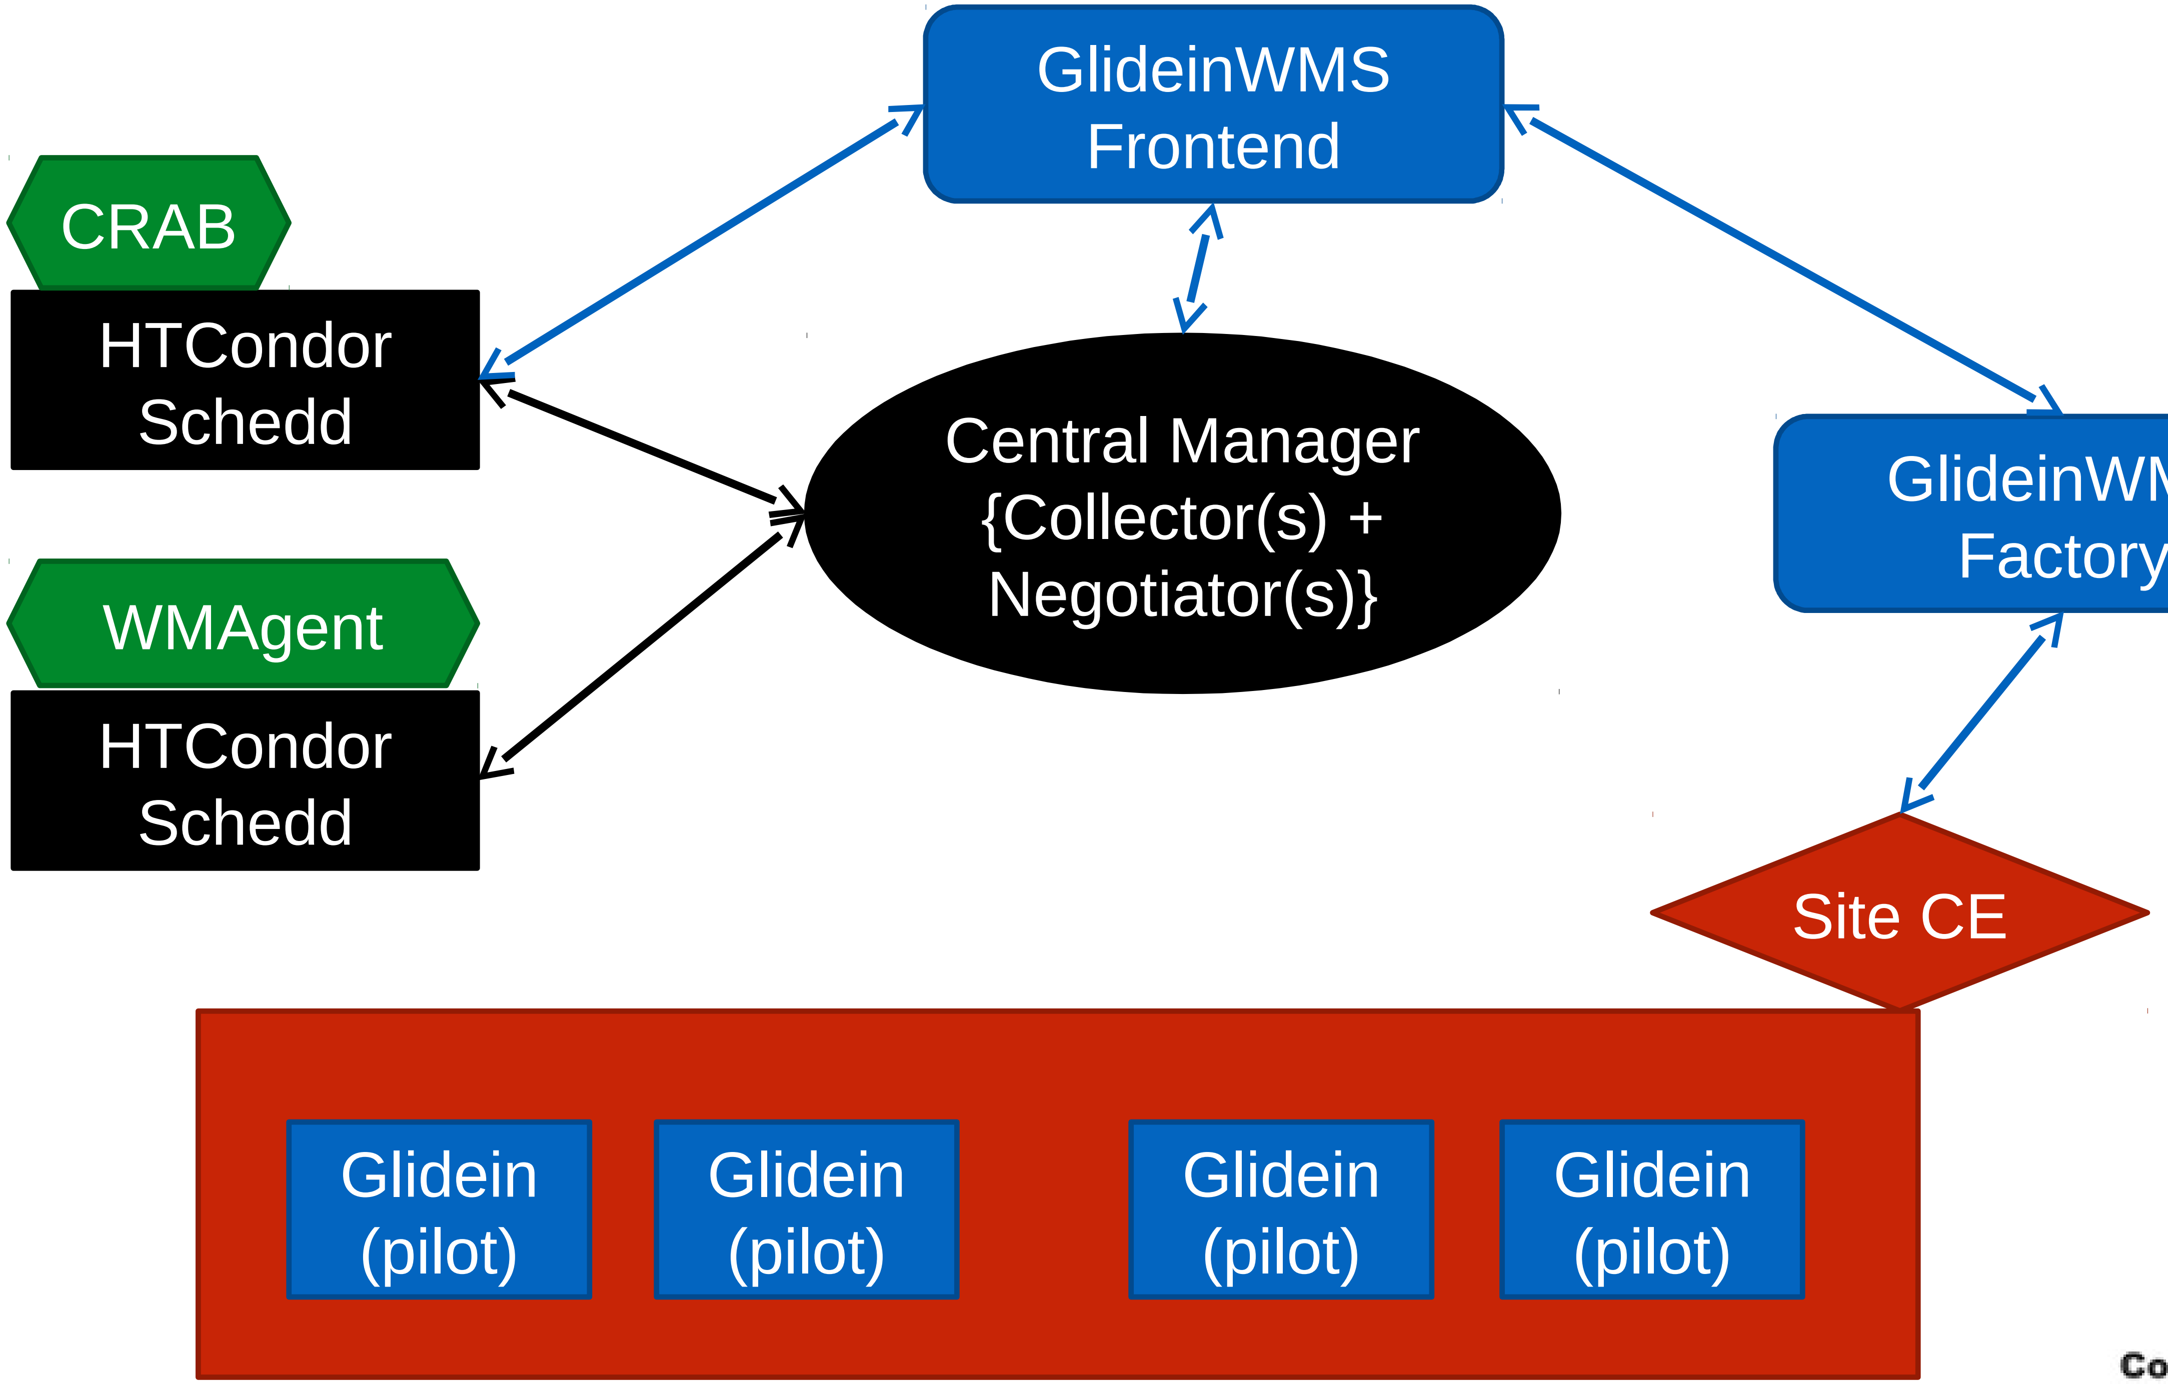

HTCondor: CMS Experience

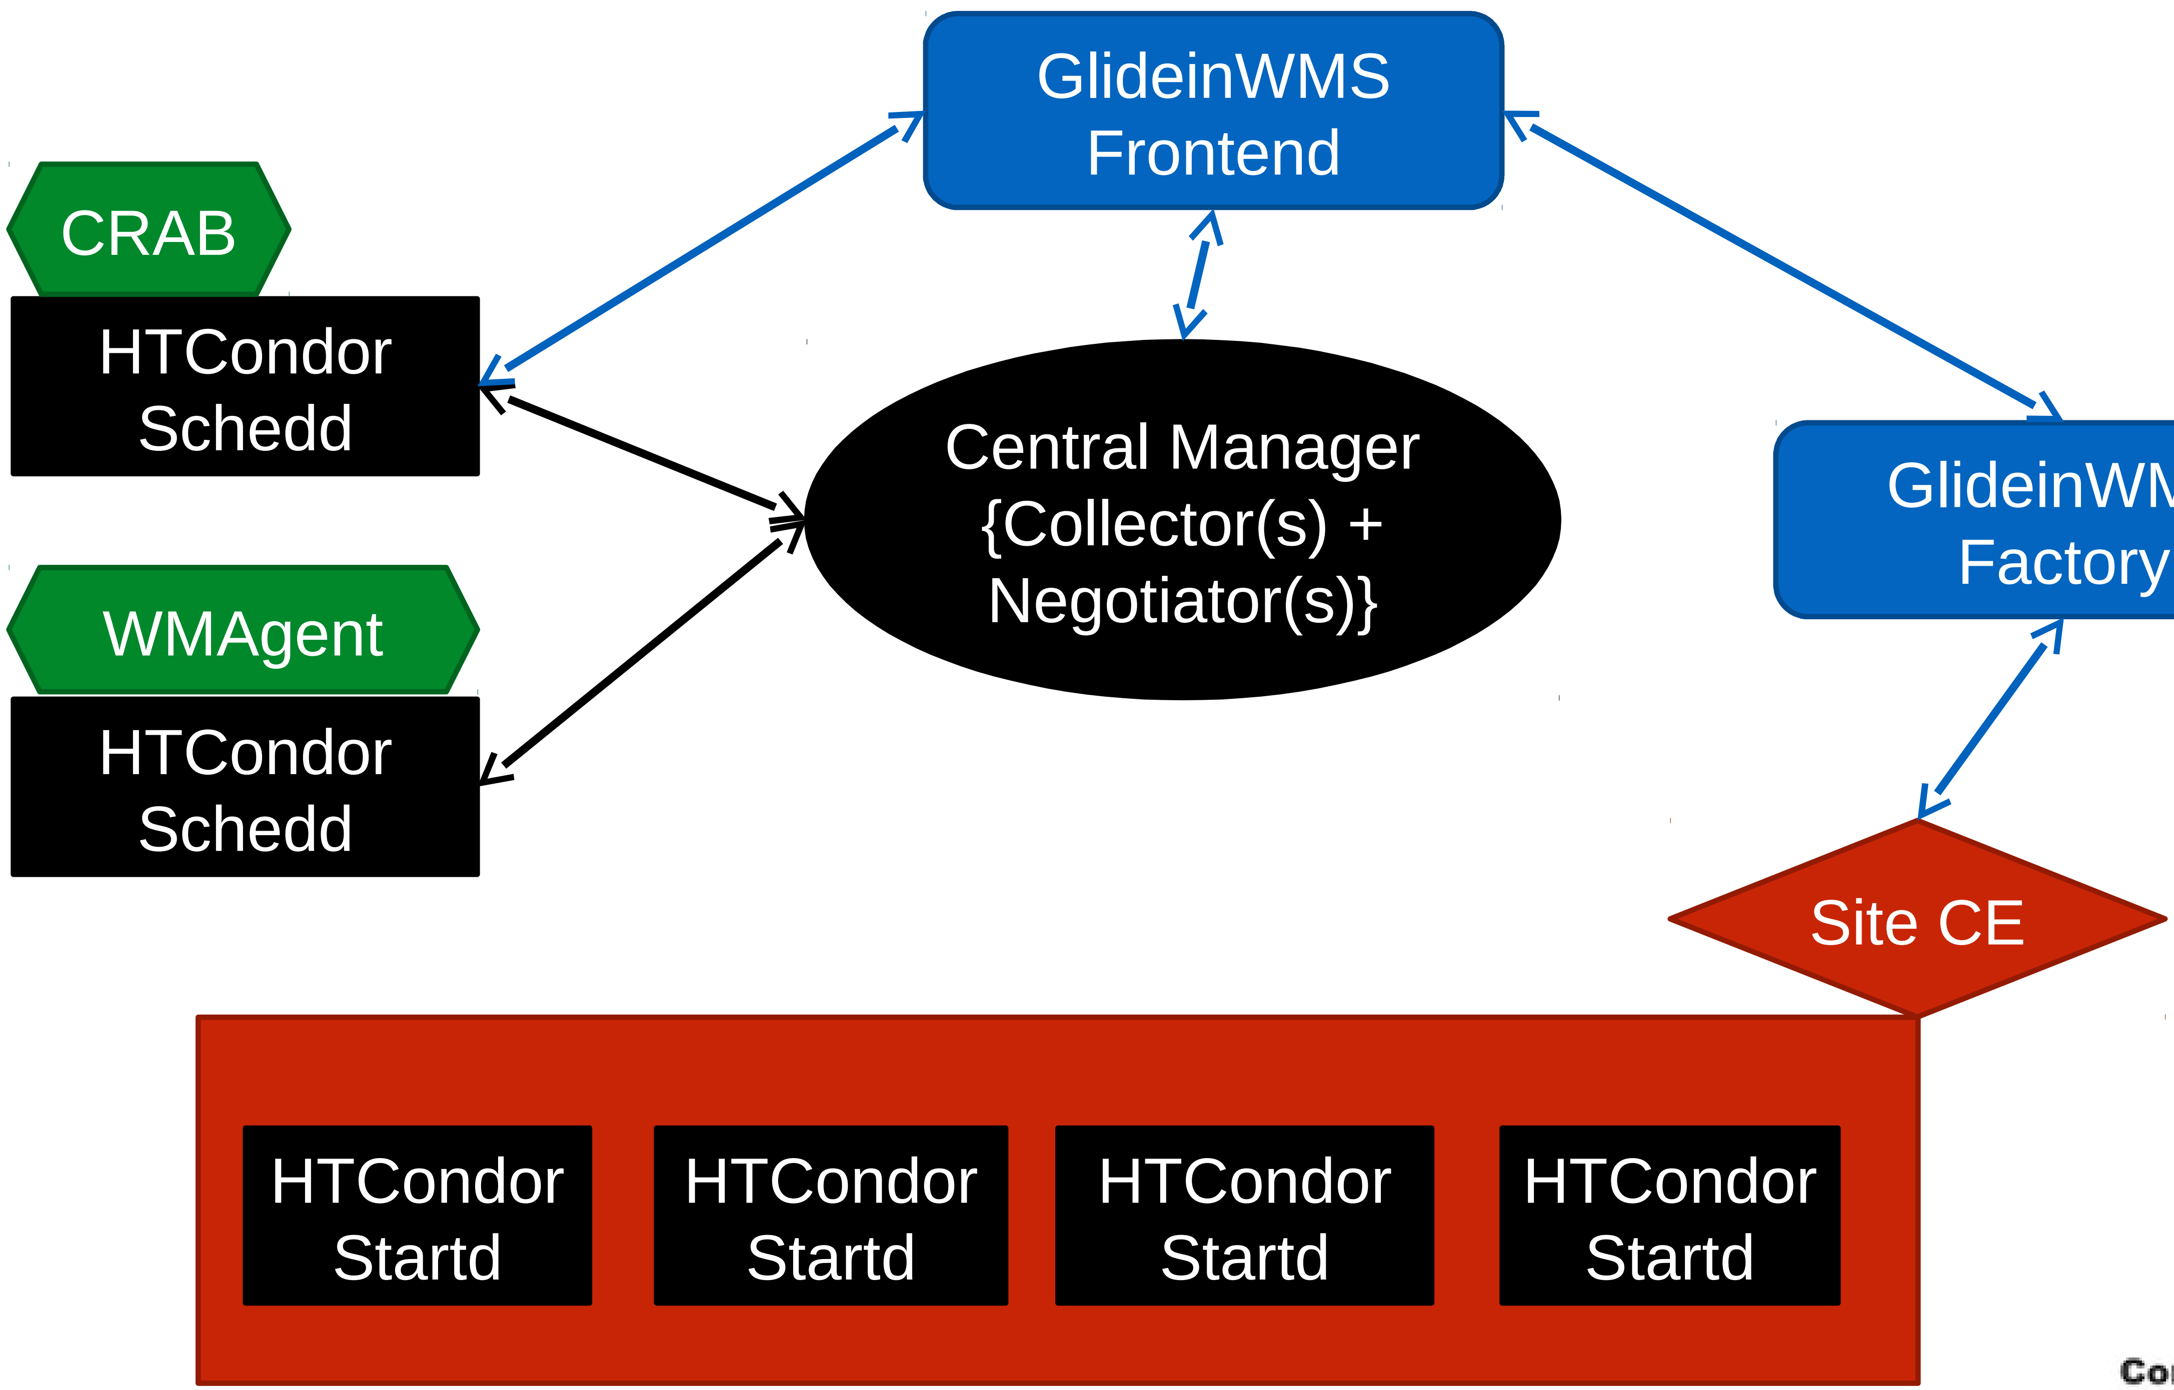

HTCondor: CMS Experience

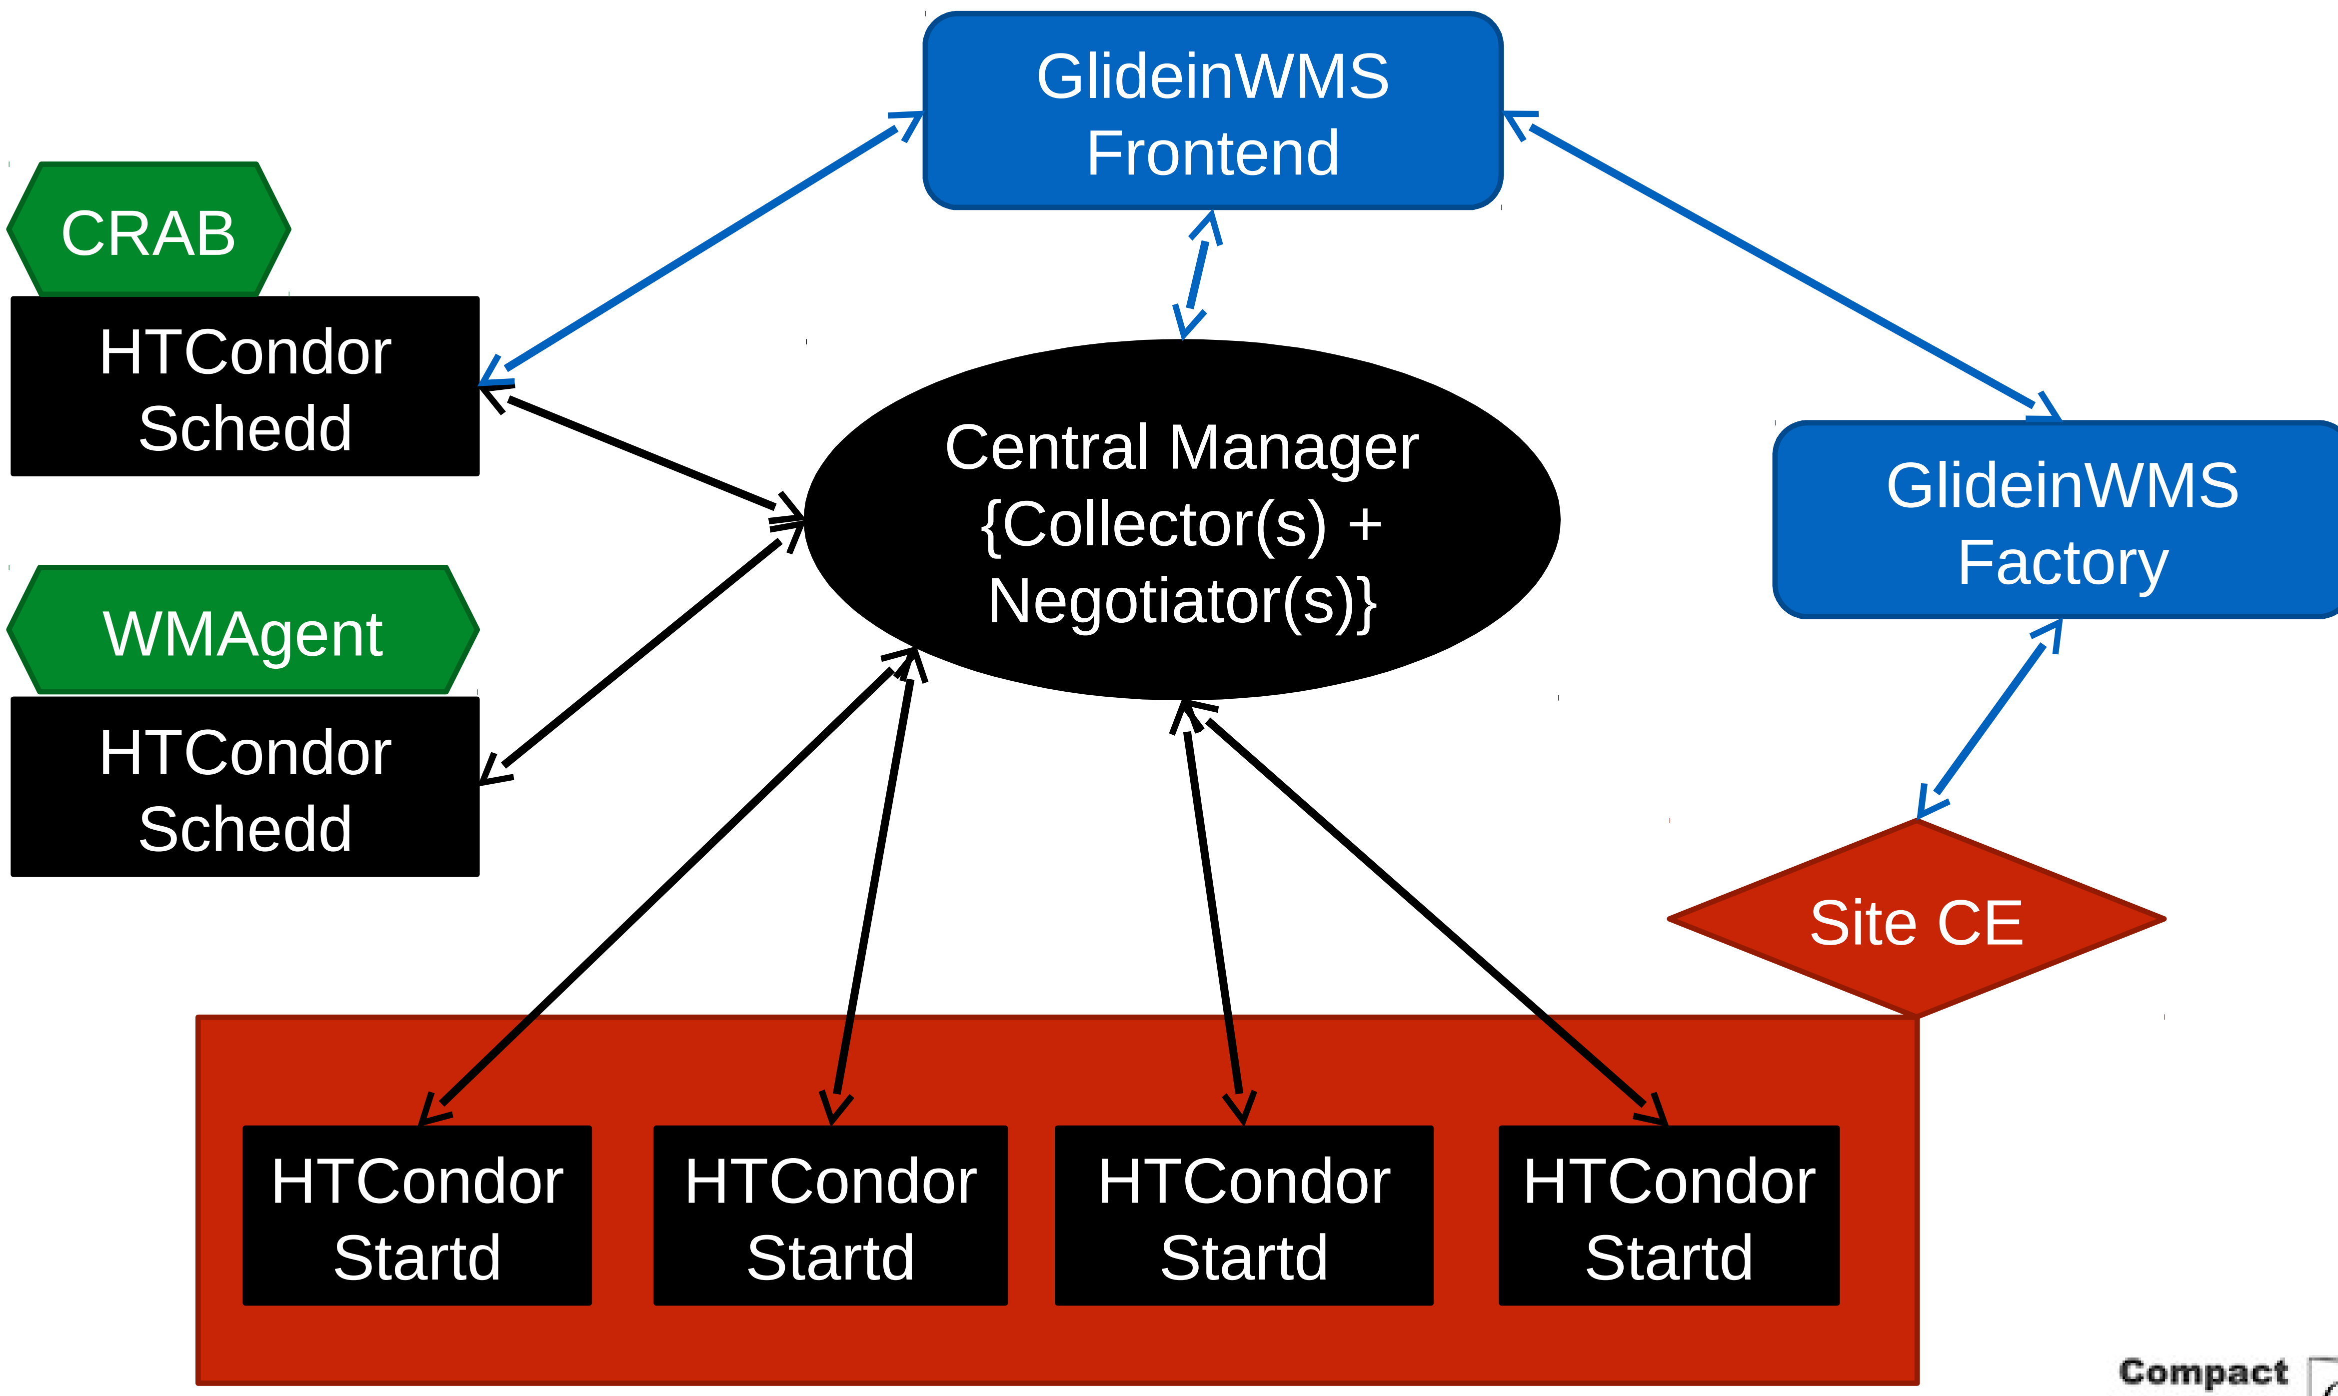

CMS Global pool

- In the first stage of matchmaking, glideinWMS frontend matches jobs to their desired sites and requests the glideinWMS factory to send glideins (properly configured condor tar ball)

- The 2nd stage of matchmaking is when a job gets matched to a slot once the condor starts on the worker node and makes itself available in the pool

- Glidein pulls in the job and then GLExec is used to switch to central production or analysis user's credentials

CMS Global pool

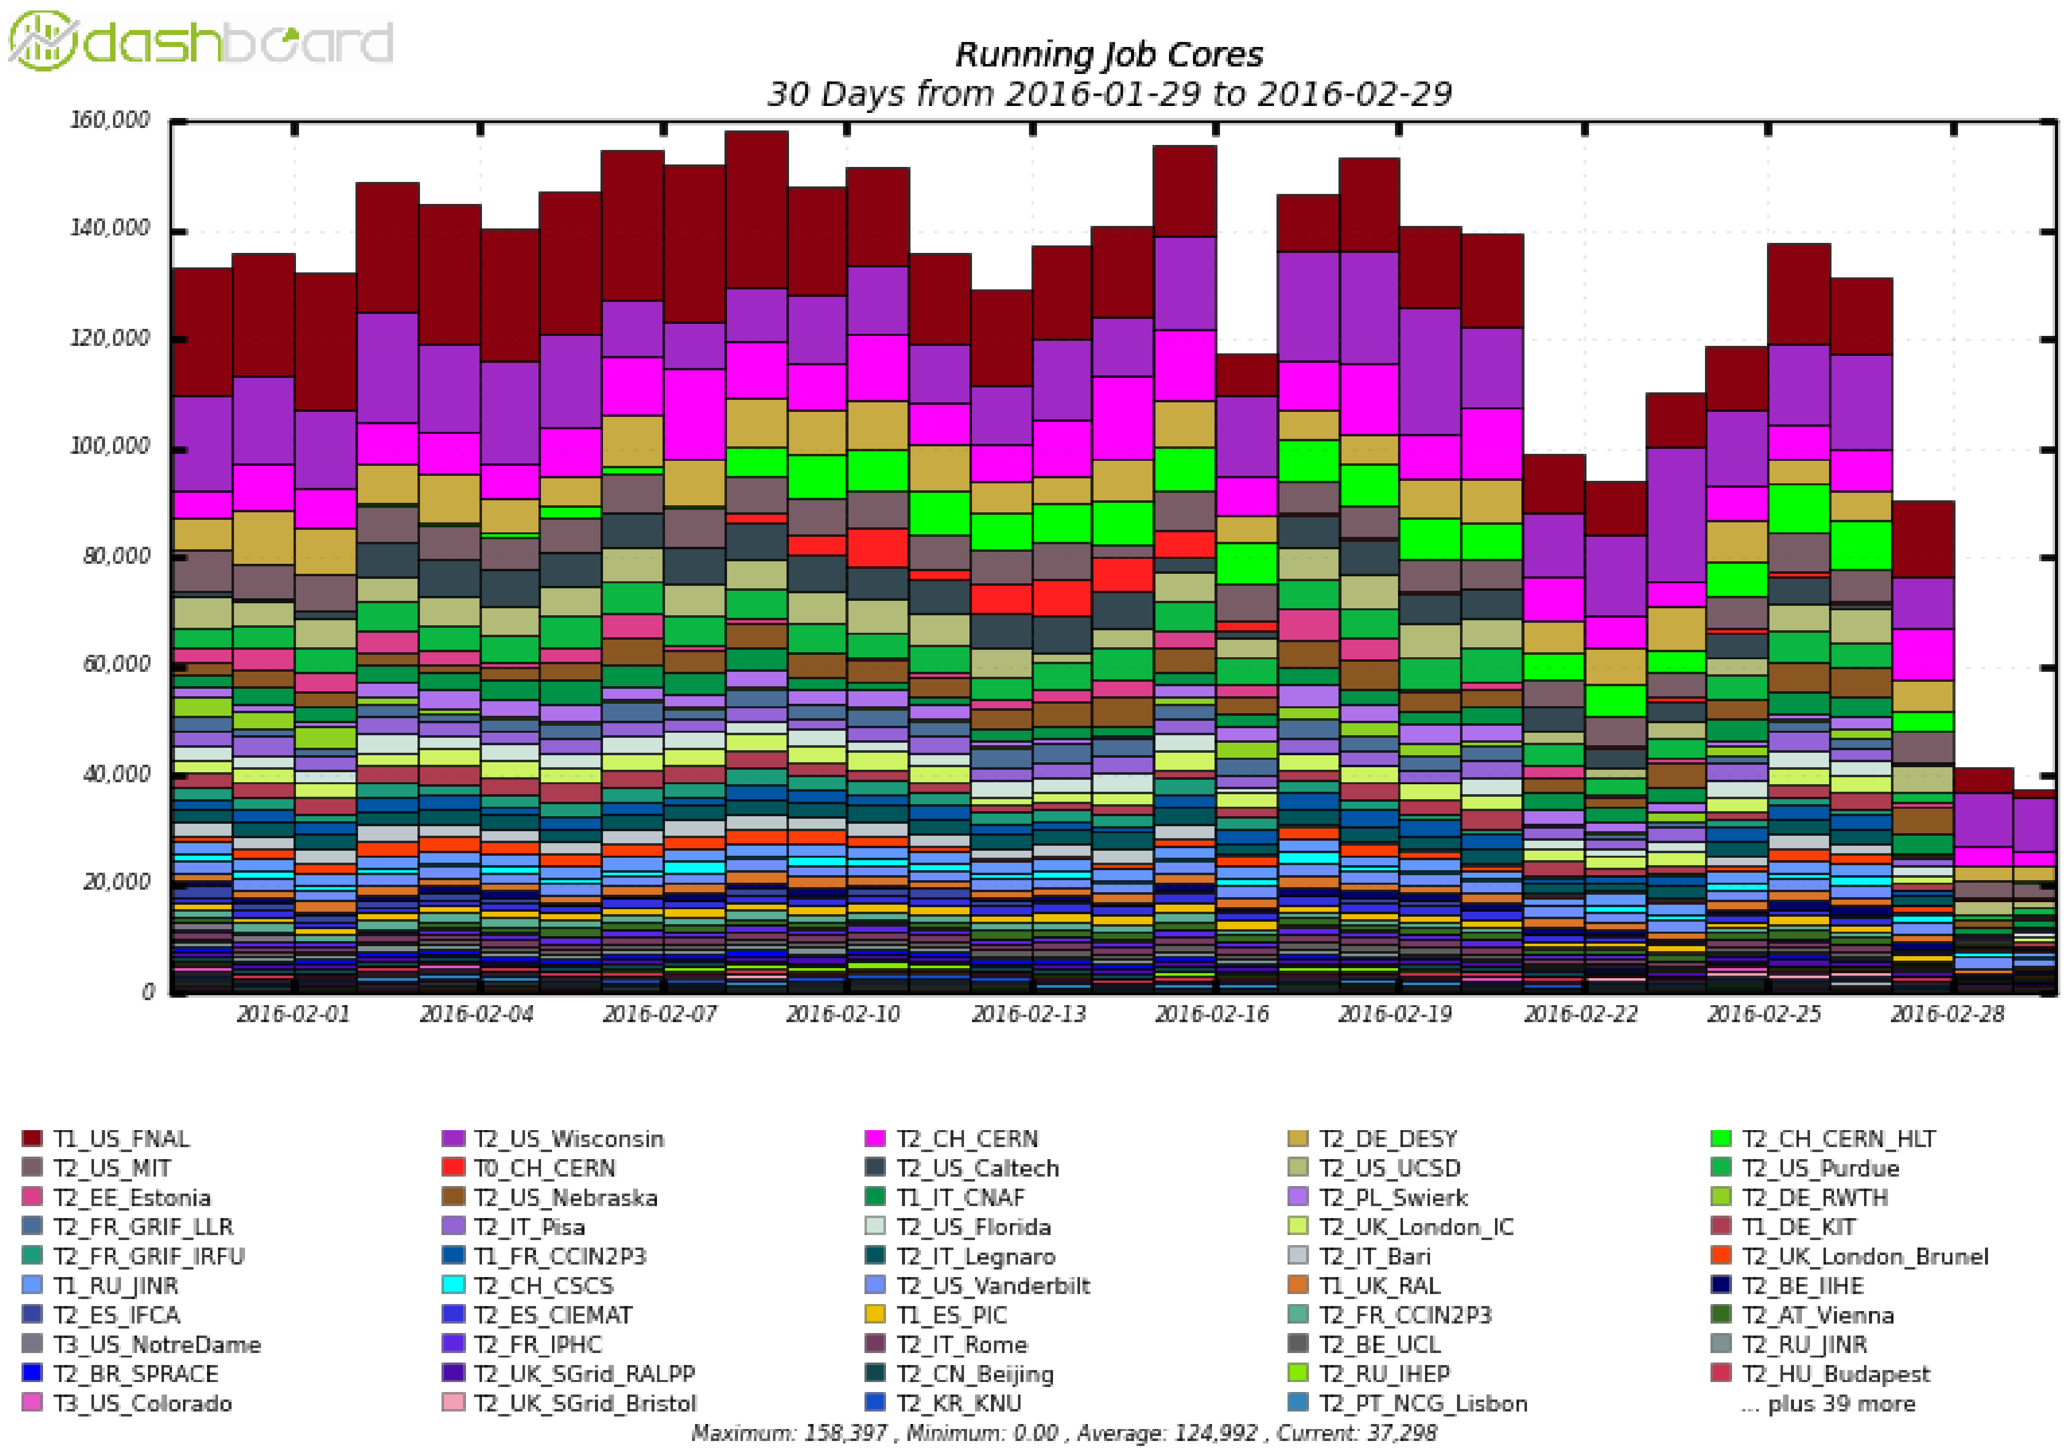

In the global pool we have successfully been able to ramp up beyond 150k occupied cores

Compact Muon Solenoid experiment at CERN's LHC

CMS Global pool

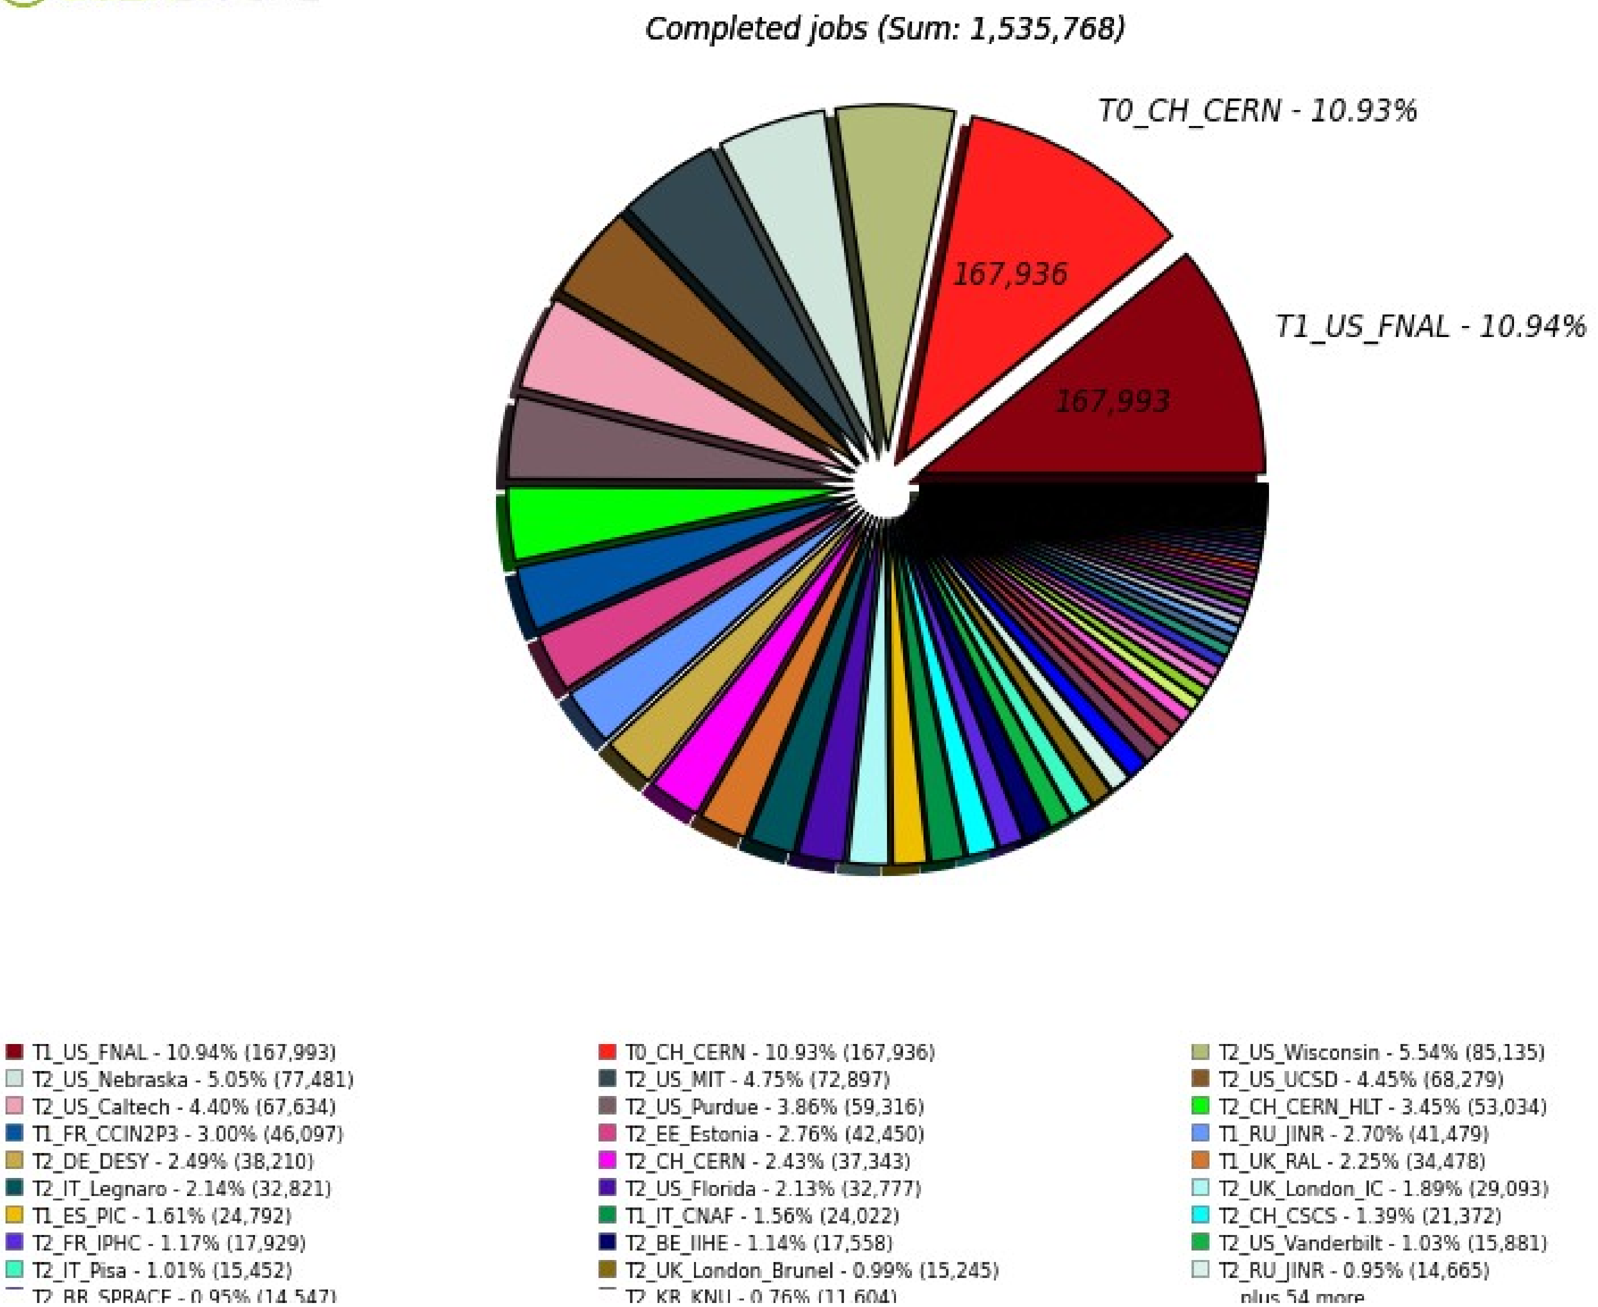

Over 1.5 million jobs completed in the last 48 hours

Compact Muon Solenoid experiment at CERN's LHC

CMS Global pool

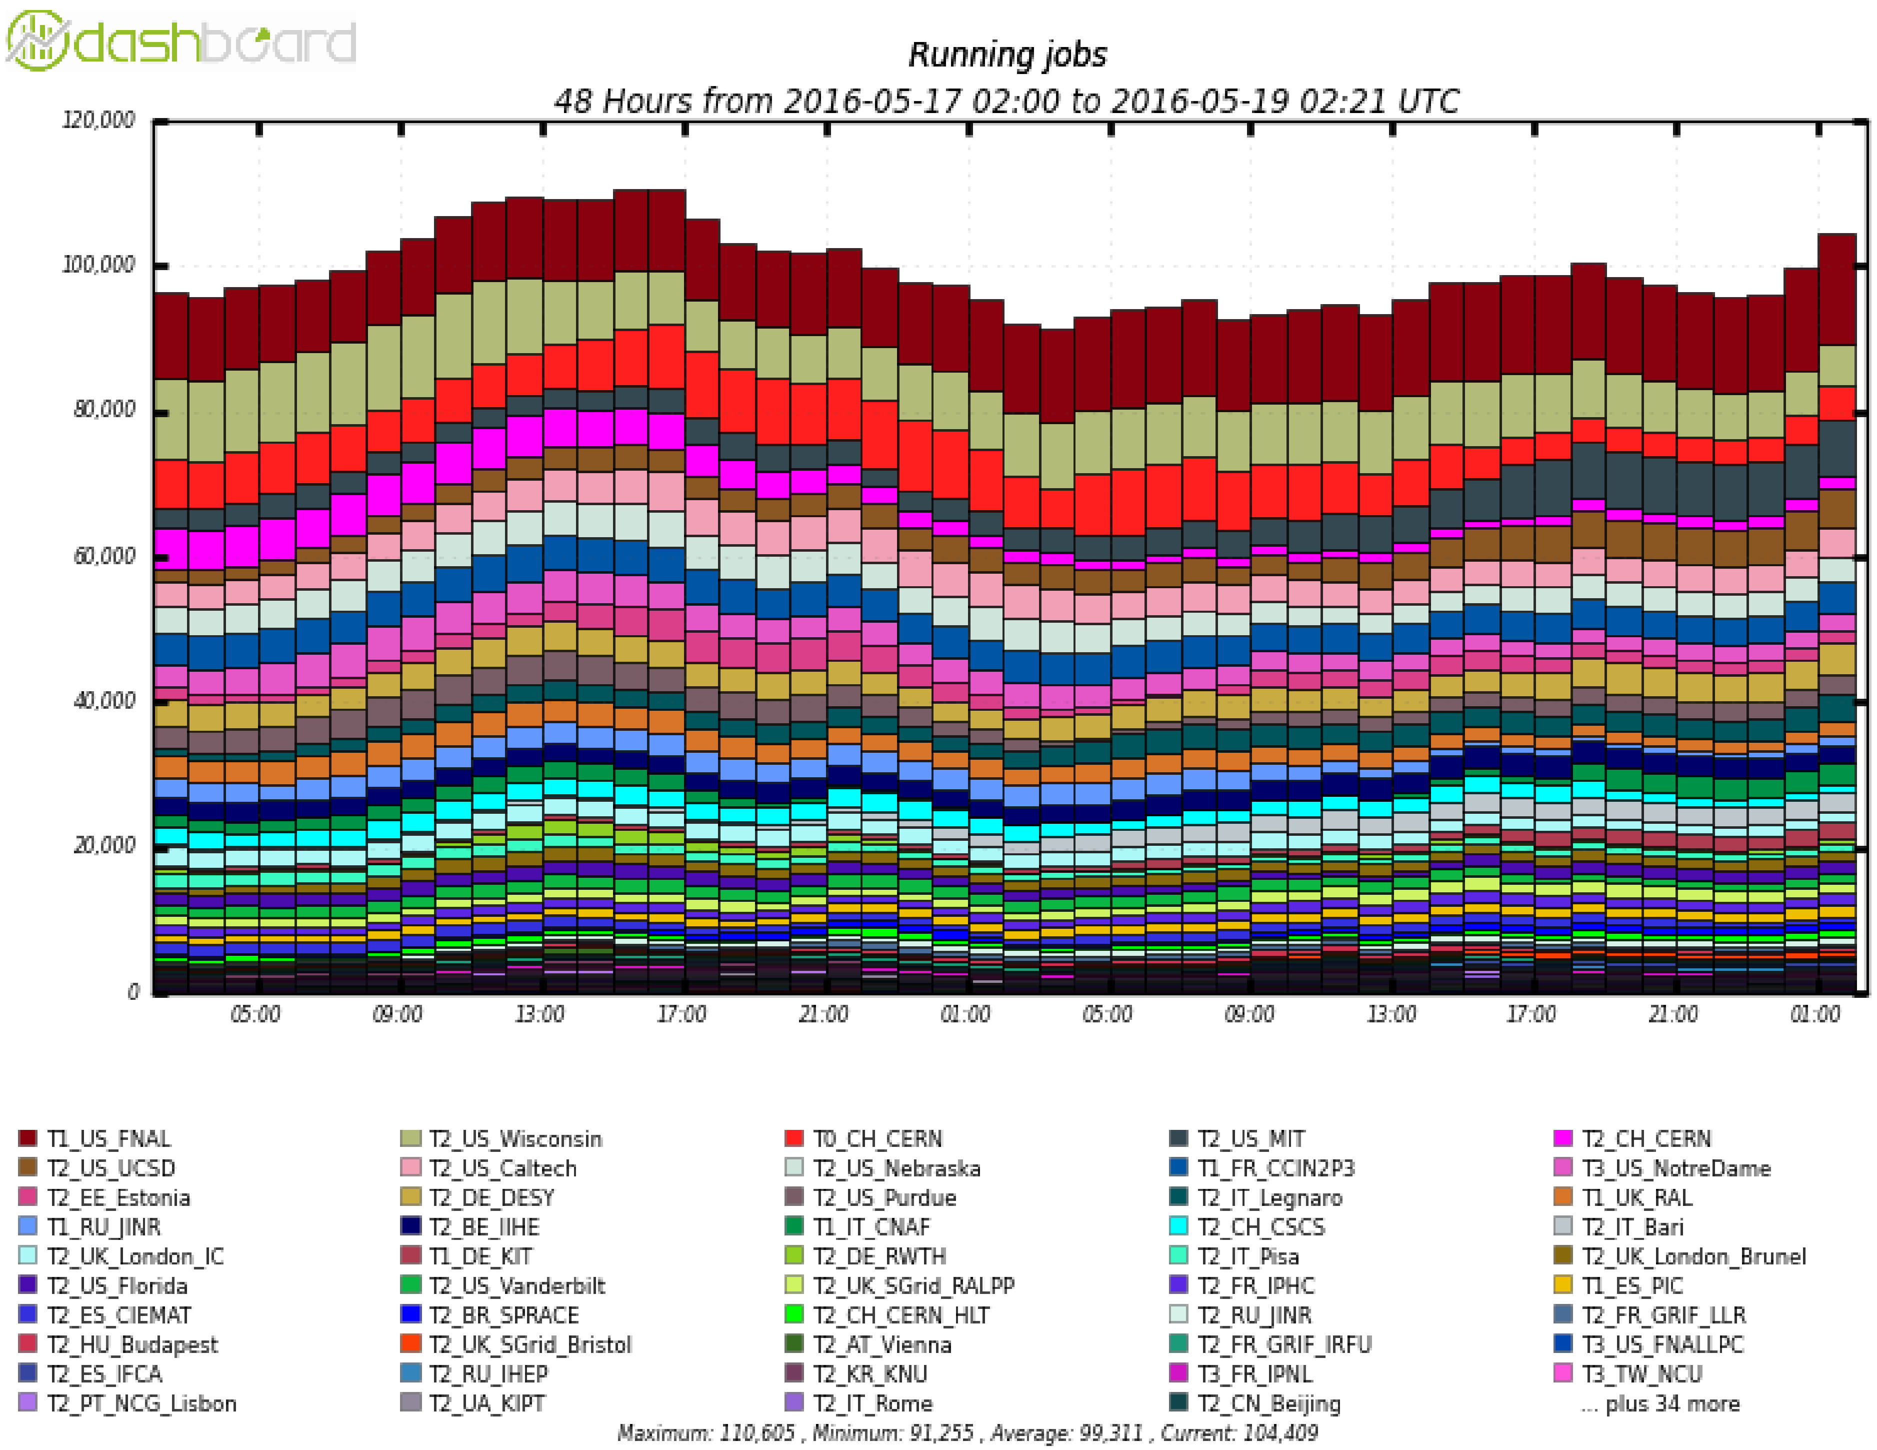

The number of parallel running jobs for the last 48 hours

Compact Muon Solenoid experiment at CERN's LHC

CMS Global pool

- CMS Global pool has over 30 schedds connected to it in total, and over a dozen of them are active at the same time. The schedds are divided into three different availability zones for redundancy

- Central managers and frontend run in HA mode with automatic failover

- Four factories for redundancy

- Infrastructure at CERN and FNAL managed via puppet

Compact Muon Solenoid experiment at CERN's LHC

CMS Global pool - HA

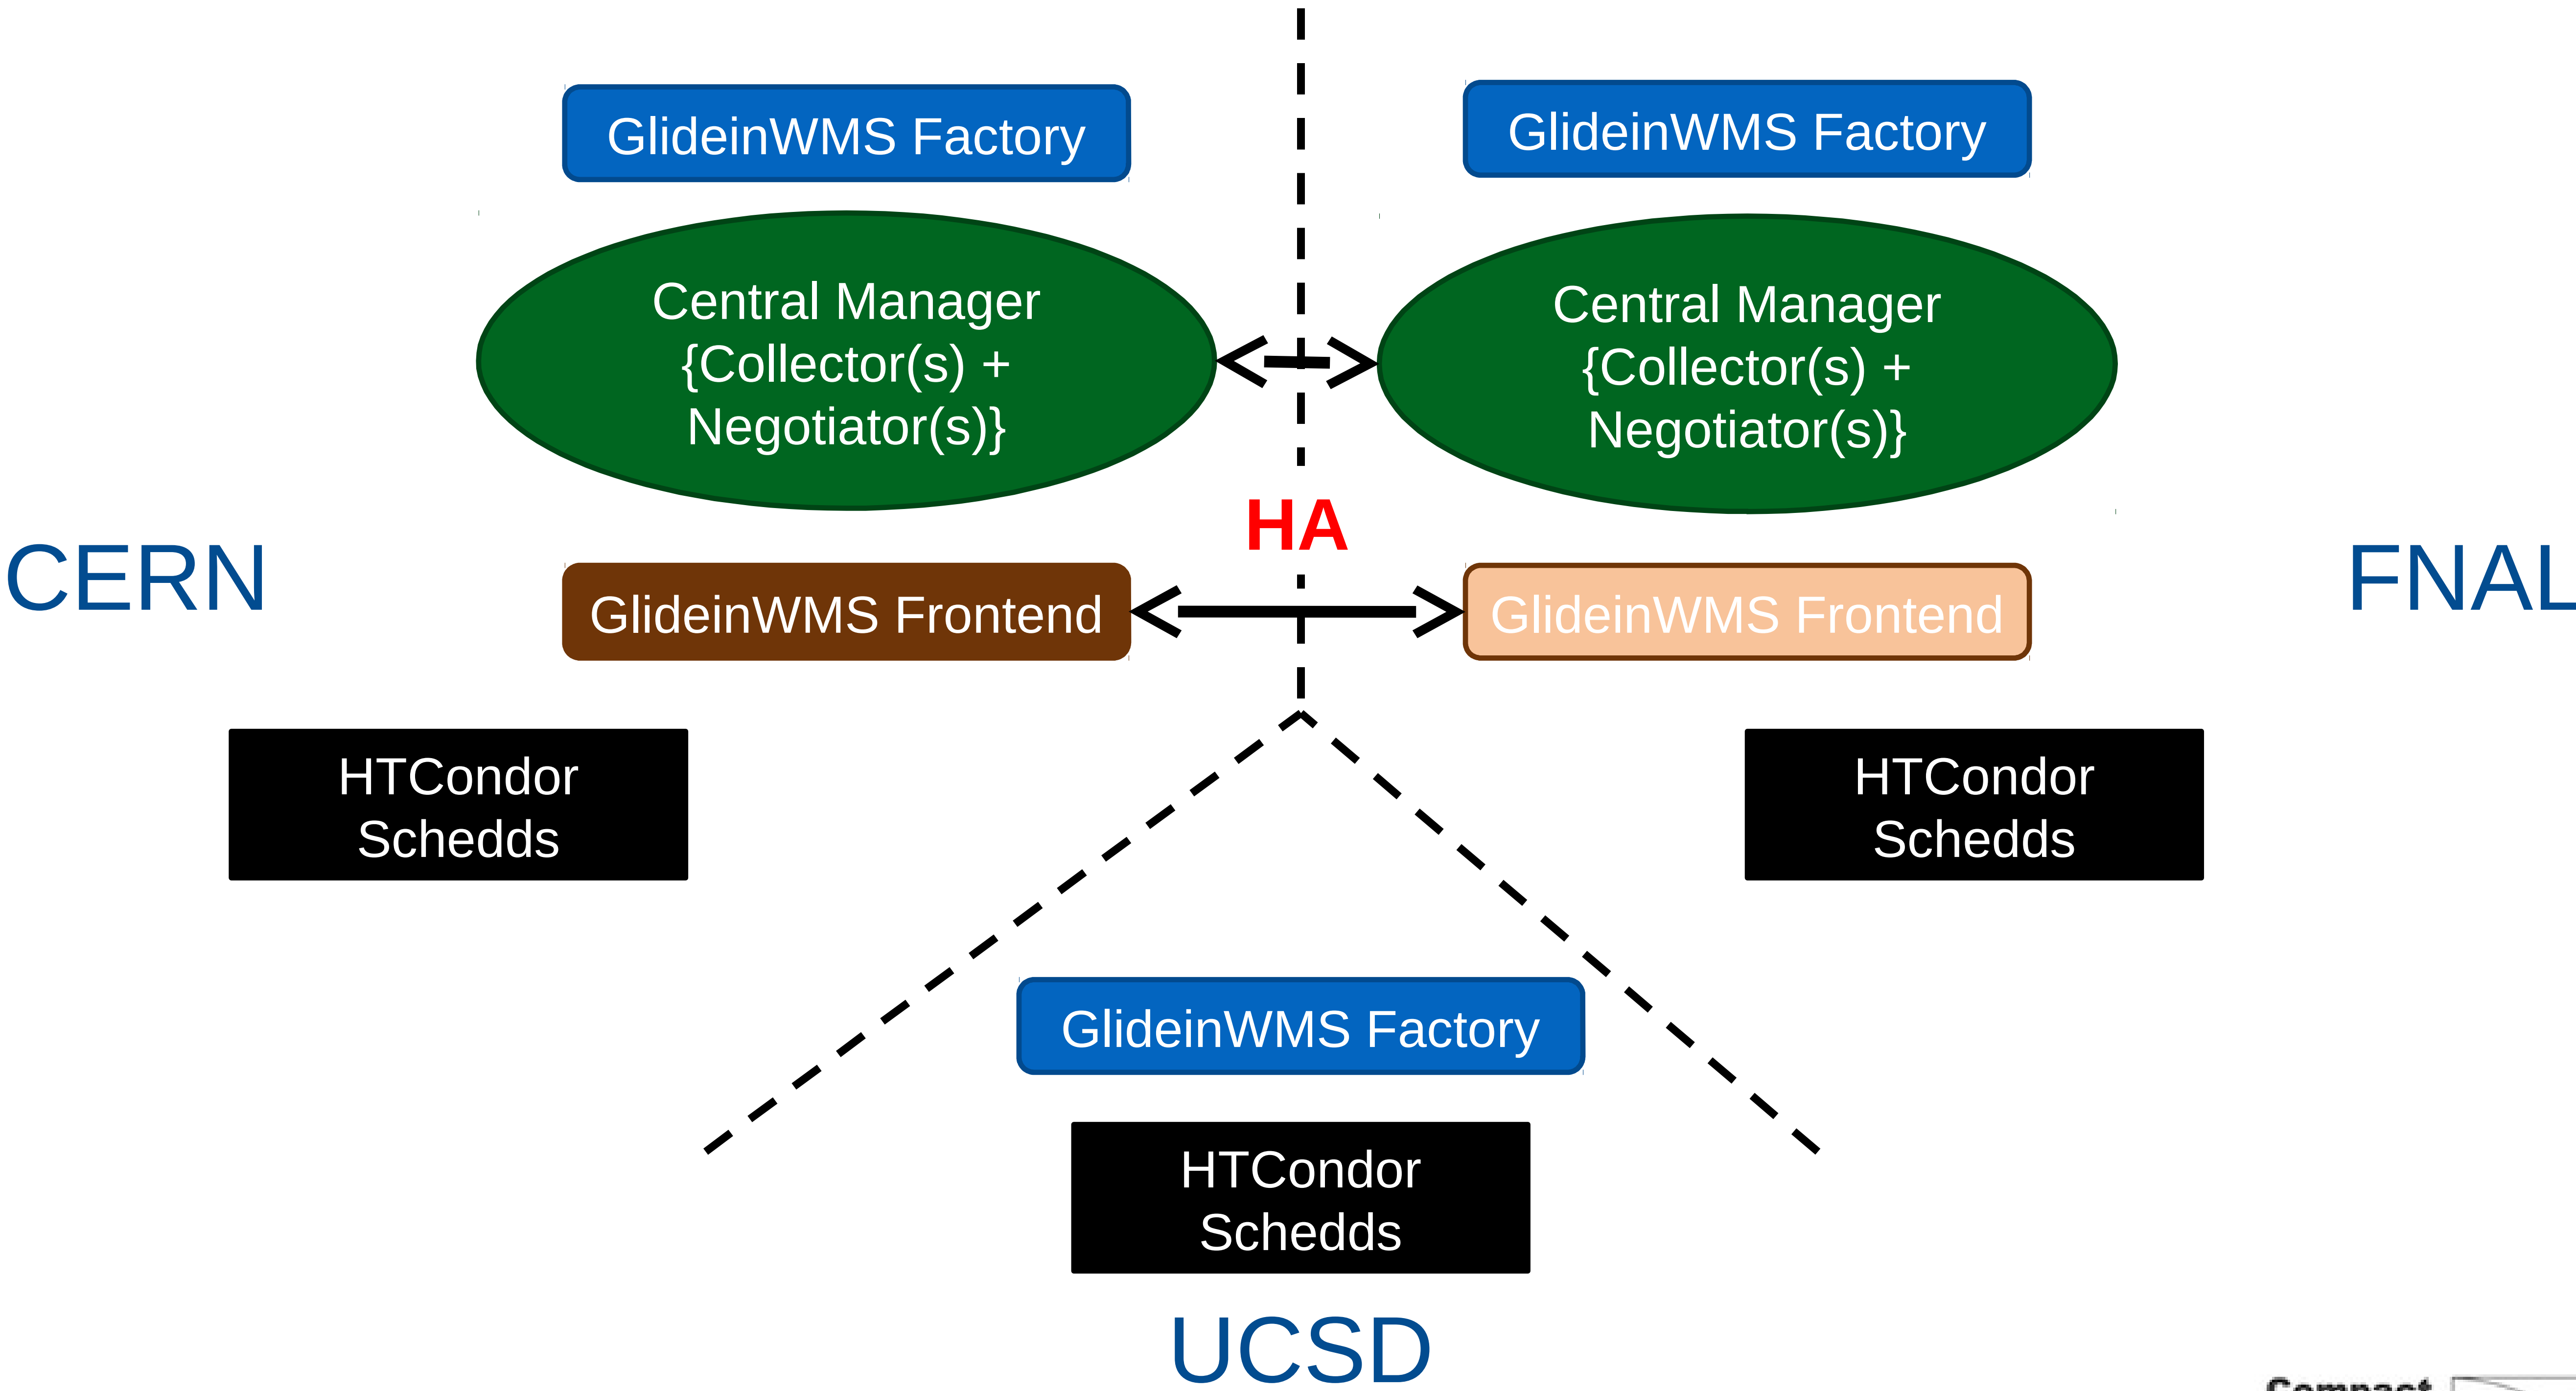

CMS Global pool - Flocking

With tier0 being so critical to CMS, we decided to move tier0 operations to its own pool to mitigate risks from any issues in the global pool.

Tier0 schedds utilize HTCondor flocking to access tier1 resources via the global pool (and vice versa)

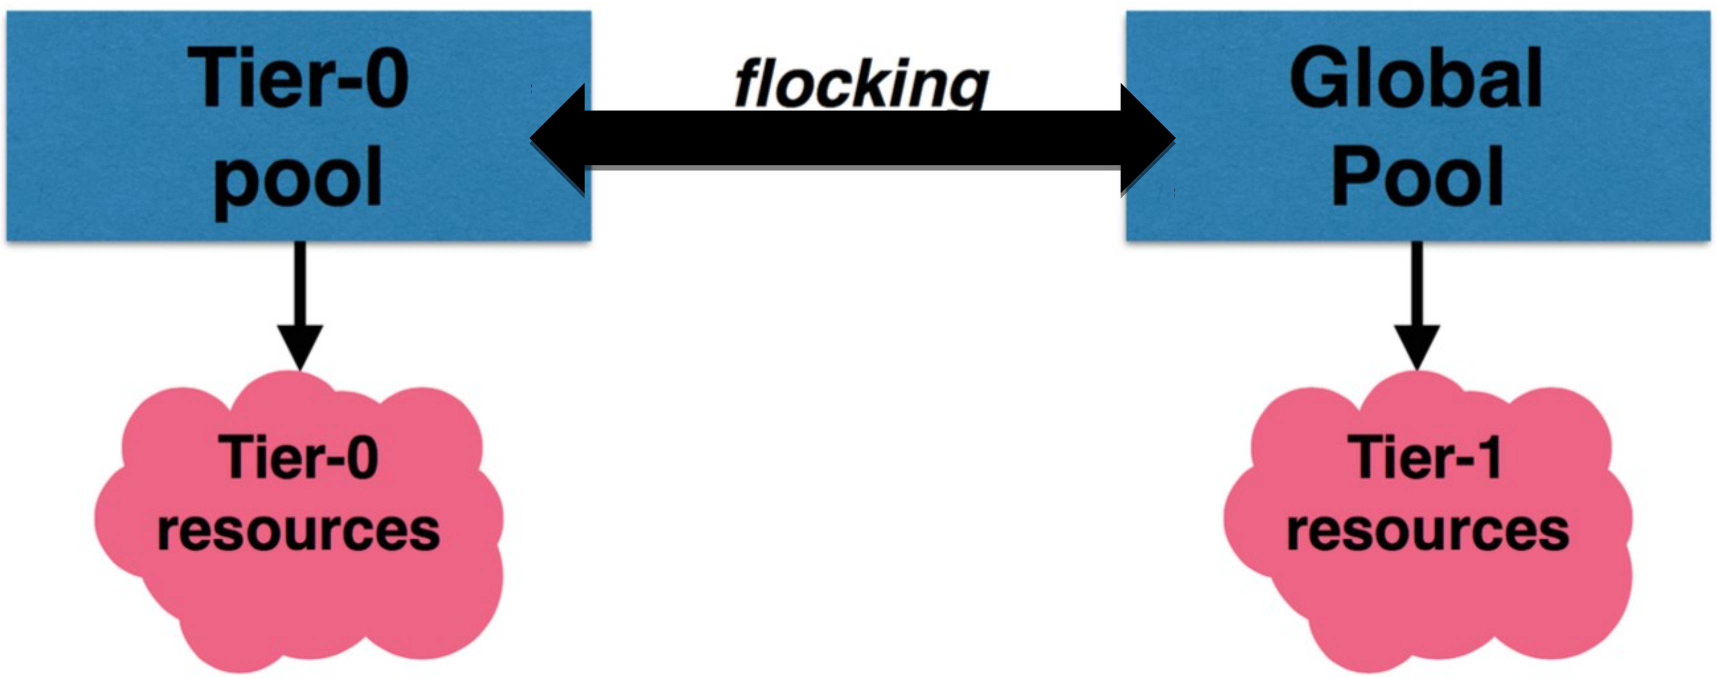

CMS Global pool – Prioritization

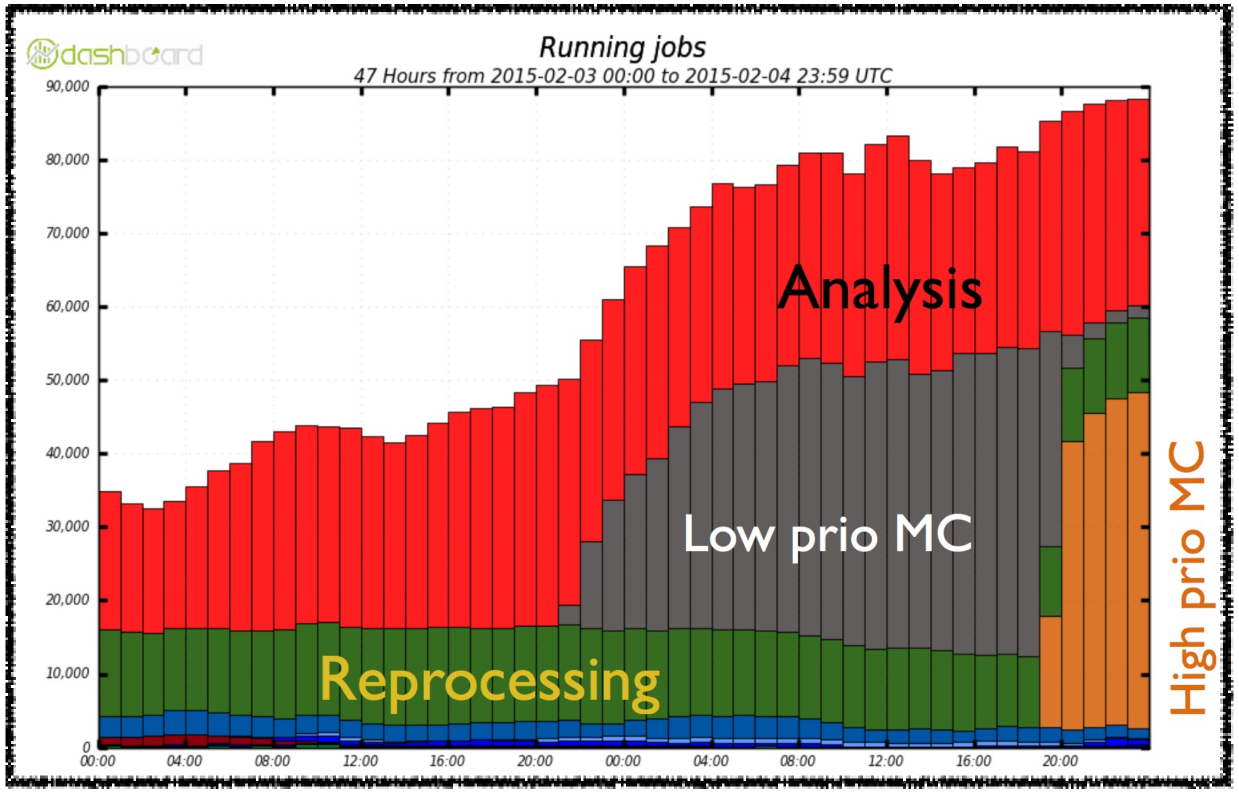

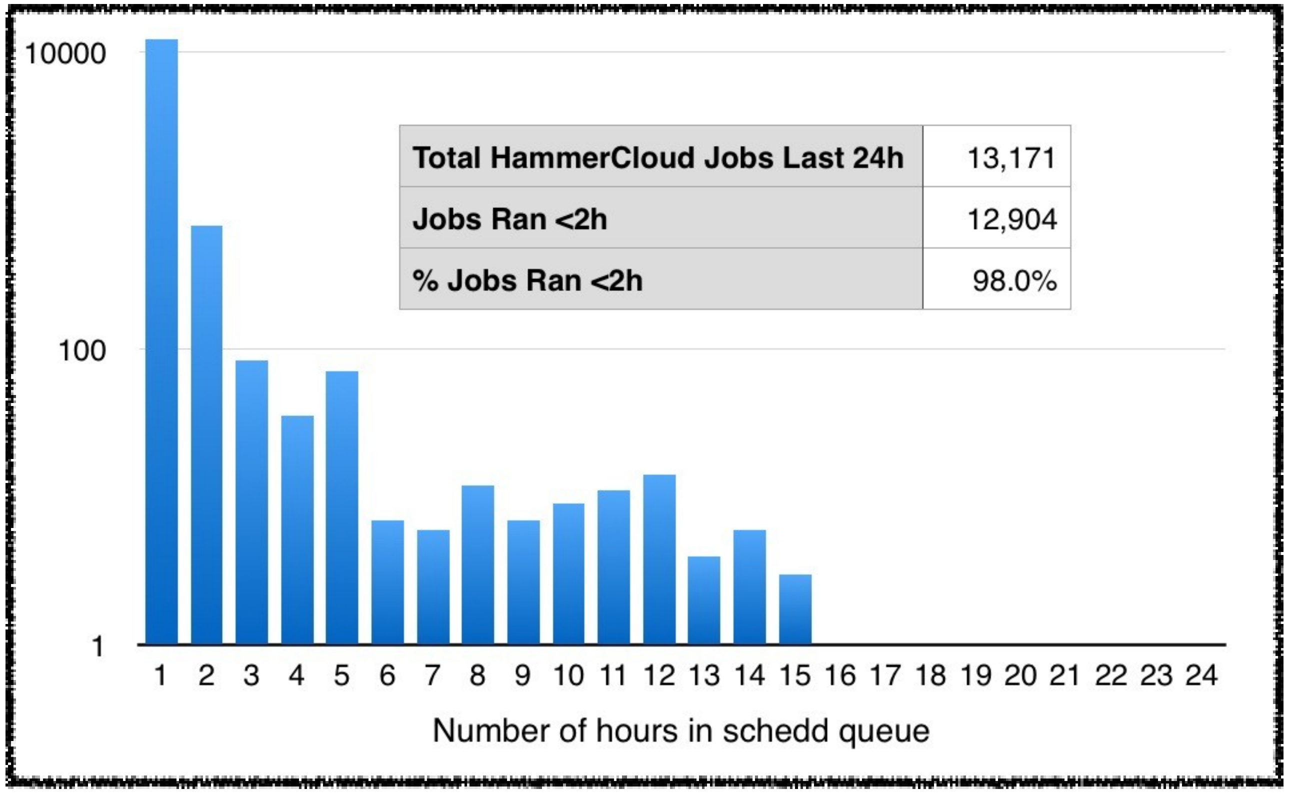

Job priority and user priority features provided by HTCondor utilized to ensure global prioritization

Accounting between central production and user analysis also managed via HTCondor

Compact Muon Solenoid experiment at CERN's LHC

CMS Global pool - Slots

- CMS uses both single core static slots and 8 core partitionable slots. Eventually the idea is to run multi core pilots everywhere

- User analysis is single core

- We have been working on dynamic resizing of jobs in CMS (See Brian's talk for more details!)

- Still need to work on multi core pilot efficiency with the job mix that we have

CMS Global pool - Monitoring

Various different monitoring pages available to monitor the pool

Name	Machine	TotalRunningJobs	TotalIdleJobs	TotalHeldJobs
cmsgwms-submit1.fnal	cmsgwms-su	6446	2850	1
cmsgwms-submit2.fnal	cmsgwms-su	7431	2784	0
cmssrv113.fnal.gov	cmssrv113.	0	0	0
cmssrv214.fnal.gov	cmssrv214.	0	1	0
cmssrv217.fnal.gov	cmssrv217.	0	0	0
cmssrv218.fnal.gov	cmssrv218.	668	122	0
cmssrv219.fnal.gov	cmssrv219.	4518	2783	0
cmssrv248.fnal.gov	cmssrv248.	877	0	0
crab3-3@submit-4.t2.	submit-4.t	7559	19212	2398
vocms005.cern.ch	vocms005.c	0	0	0
crab3-2@vocms0109.ce	vocms0109.	7157	14986	1219
crab3-7@vocms0114.ce	vocms0114.	5719	3267	1235
vocms0230.cern.ch	vocms0230.	0	0	0
vocms026.cern.ch	vocms026.c	0	0	0
vocms0303.cern.ch	vocms0303.	289	54	0
vocms0304.cern.ch	vocms0304.	0	0	0
vocms0308.cern.ch	vocms0308.	0	0	0
vocms0309.cern.ch	vocms0309.	329	40	15
vocms0310.cern.ch	vocms0310.	0	0	0
vocms0311.cern.ch	vocms0311.	49	25	7
vocms053.cern.ch	vocms053.c	0	0	0
crab3-5@vocms059.cer	vocms059.c	3166	3645	916
crab3-4@vocms066.cer	vocms066.c	3673	6560	1611
crab3test-2@vocms095	vocms095.c	12273	11856	3300
crab3test-3@vocms096	vocms096.c	12592	5409	2992

	TotalRunningJobs	TotalIdleJobs	TotalHeldJobs
Total	72746	73594	13694

```
Negotiation time (ideally under 300s) T1    = 160s
Negotiation time (ideally under 300s) T2_US = 137s
Negotiation time (ideally under 300s) other = 181s

Difference between collector total job counts: 72699 (vocms099.cern.ch) - 72727 (cmssrv221.fnal.gov) = -28

Frontend Group glidein counts:
      2 local_users
  70254 main
   2223 overflow
     61 overflow_conservative
   6267 t1prod
```


Compact Muon Solenoid experiment at CERN's LHC

CMS Global pool - Monitoring

Various different monitoring pages available to monitor the pool

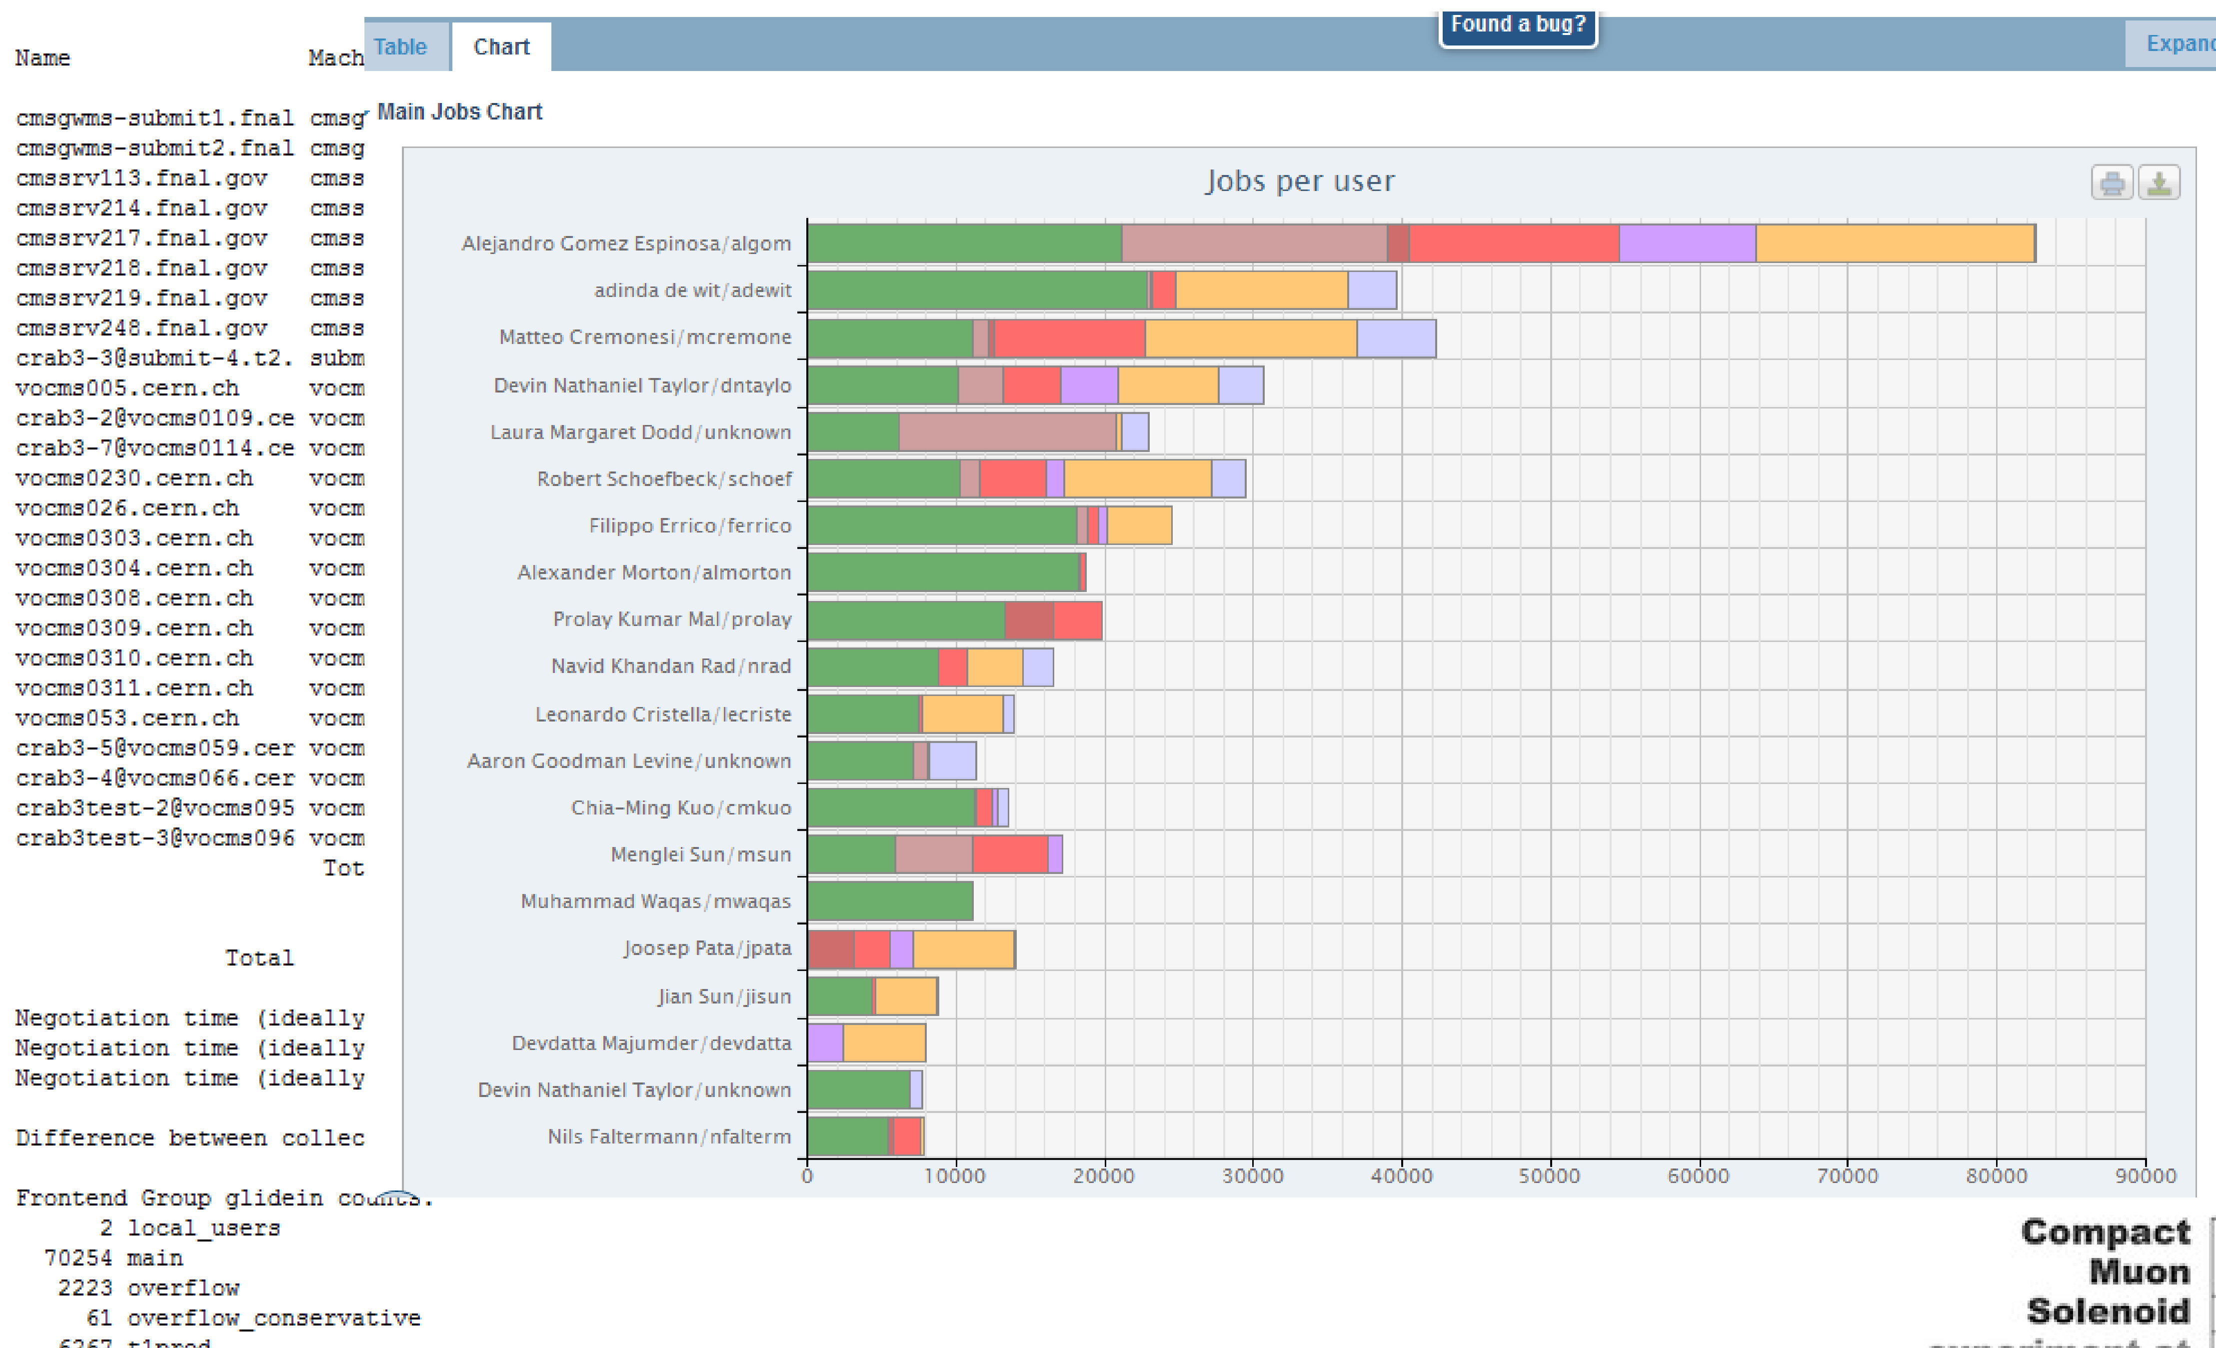

Compact Muon Solenoid experiment at CERN's LHC

CMS Global pool - Monitoring

Various different monitoring pages available to monitor the pool

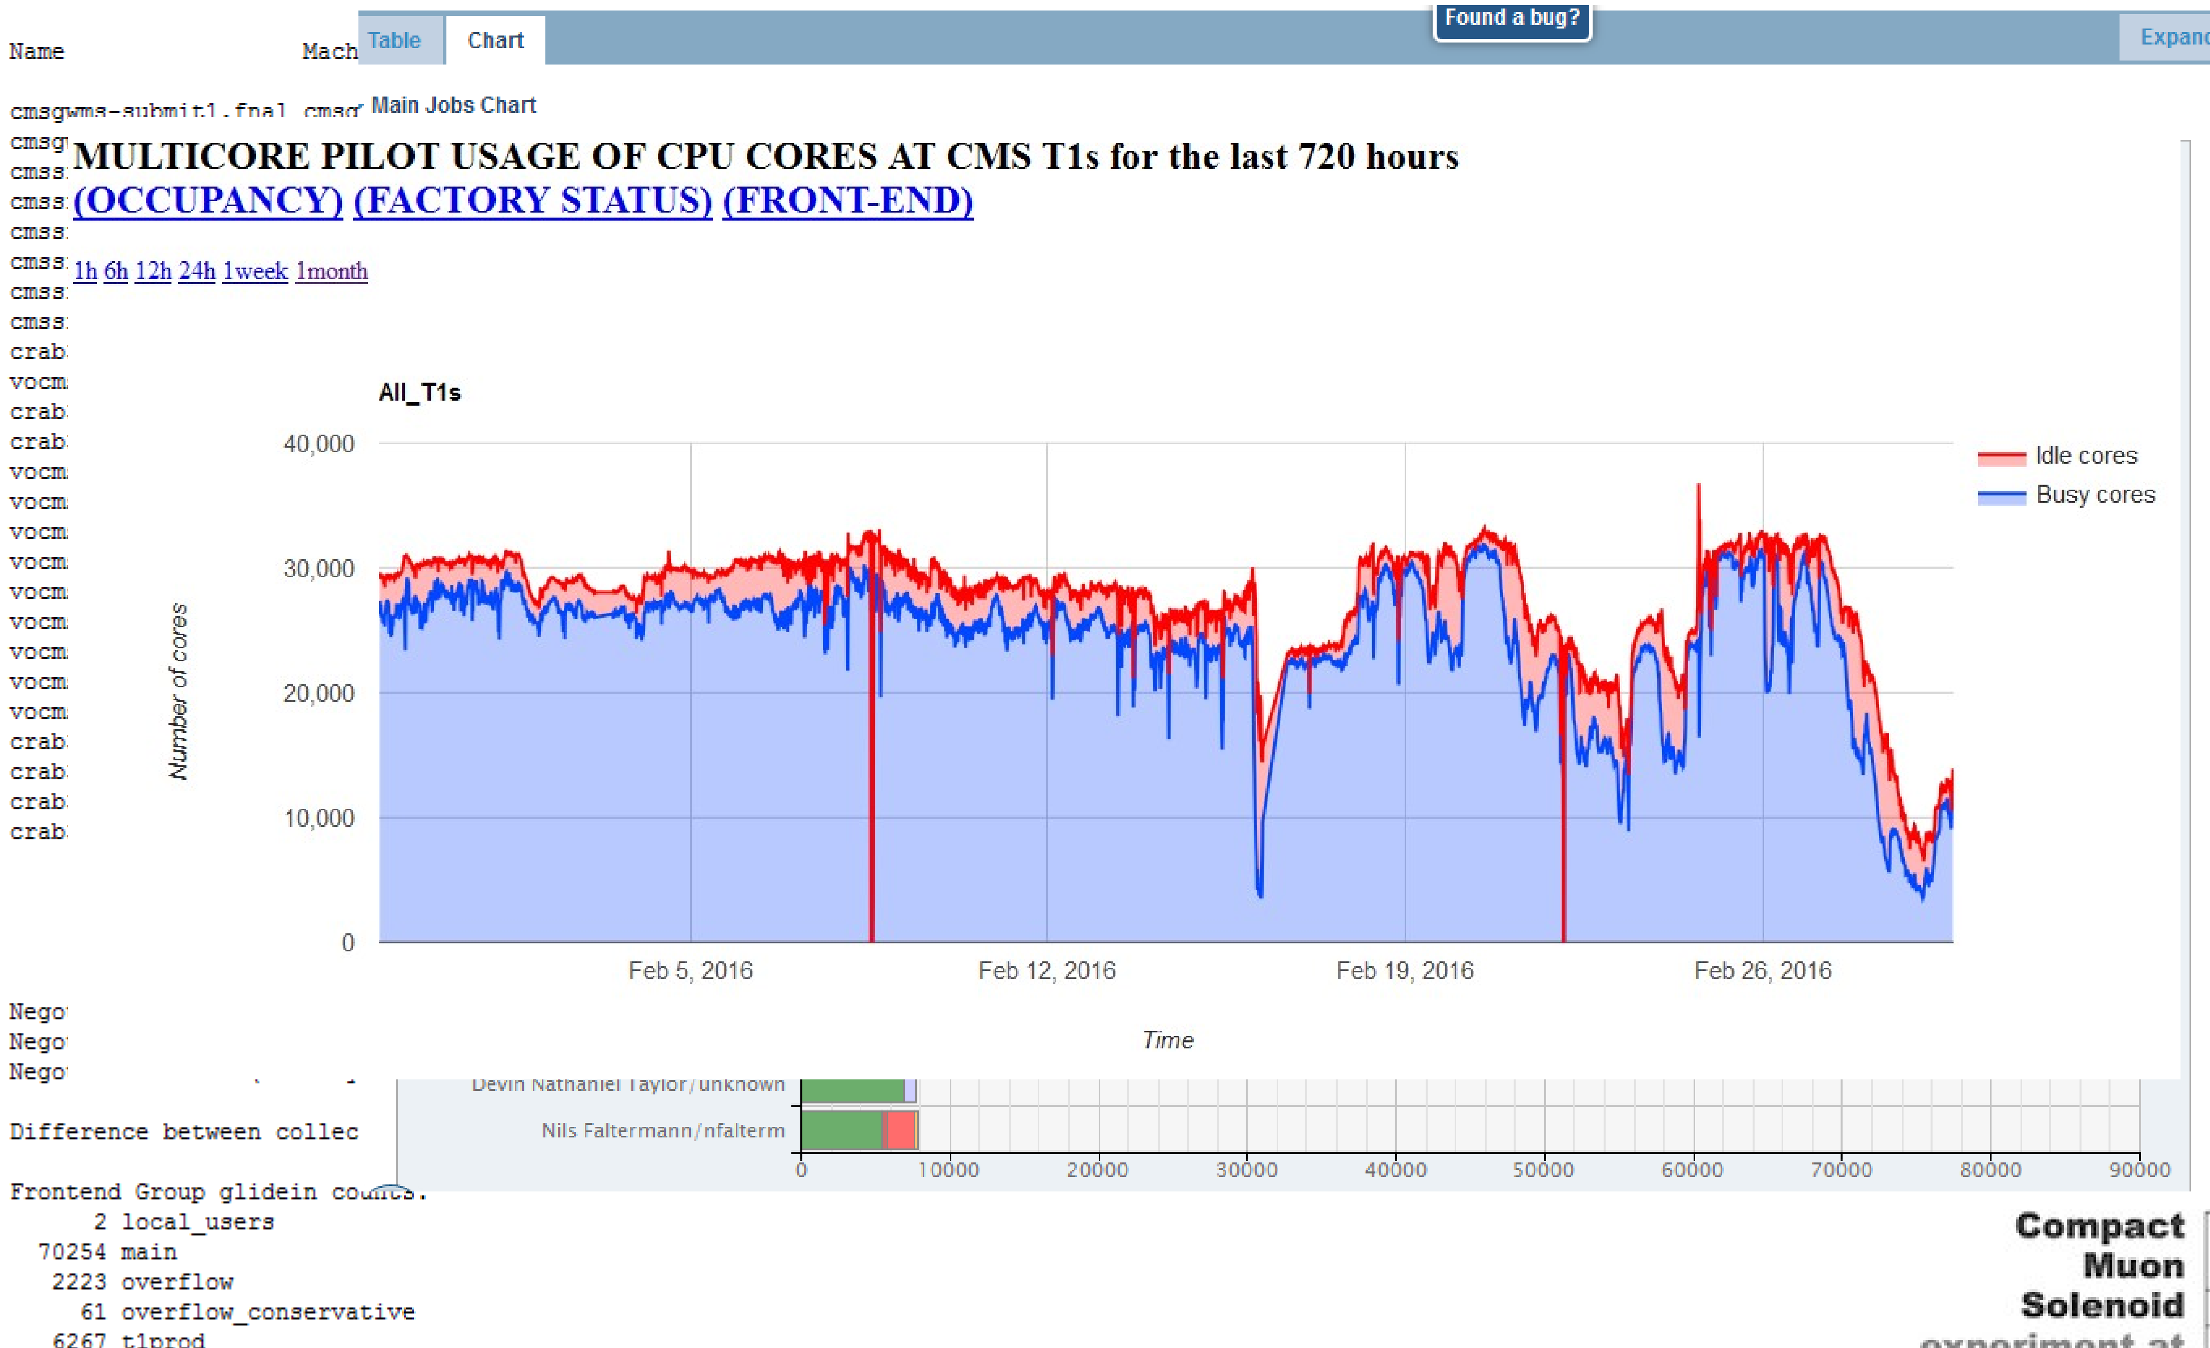

Compact Muon Solenoid experiment at CERN's LHC

CMS Global pool - Monitoring

Various different monitoring pages available to monitor the pool

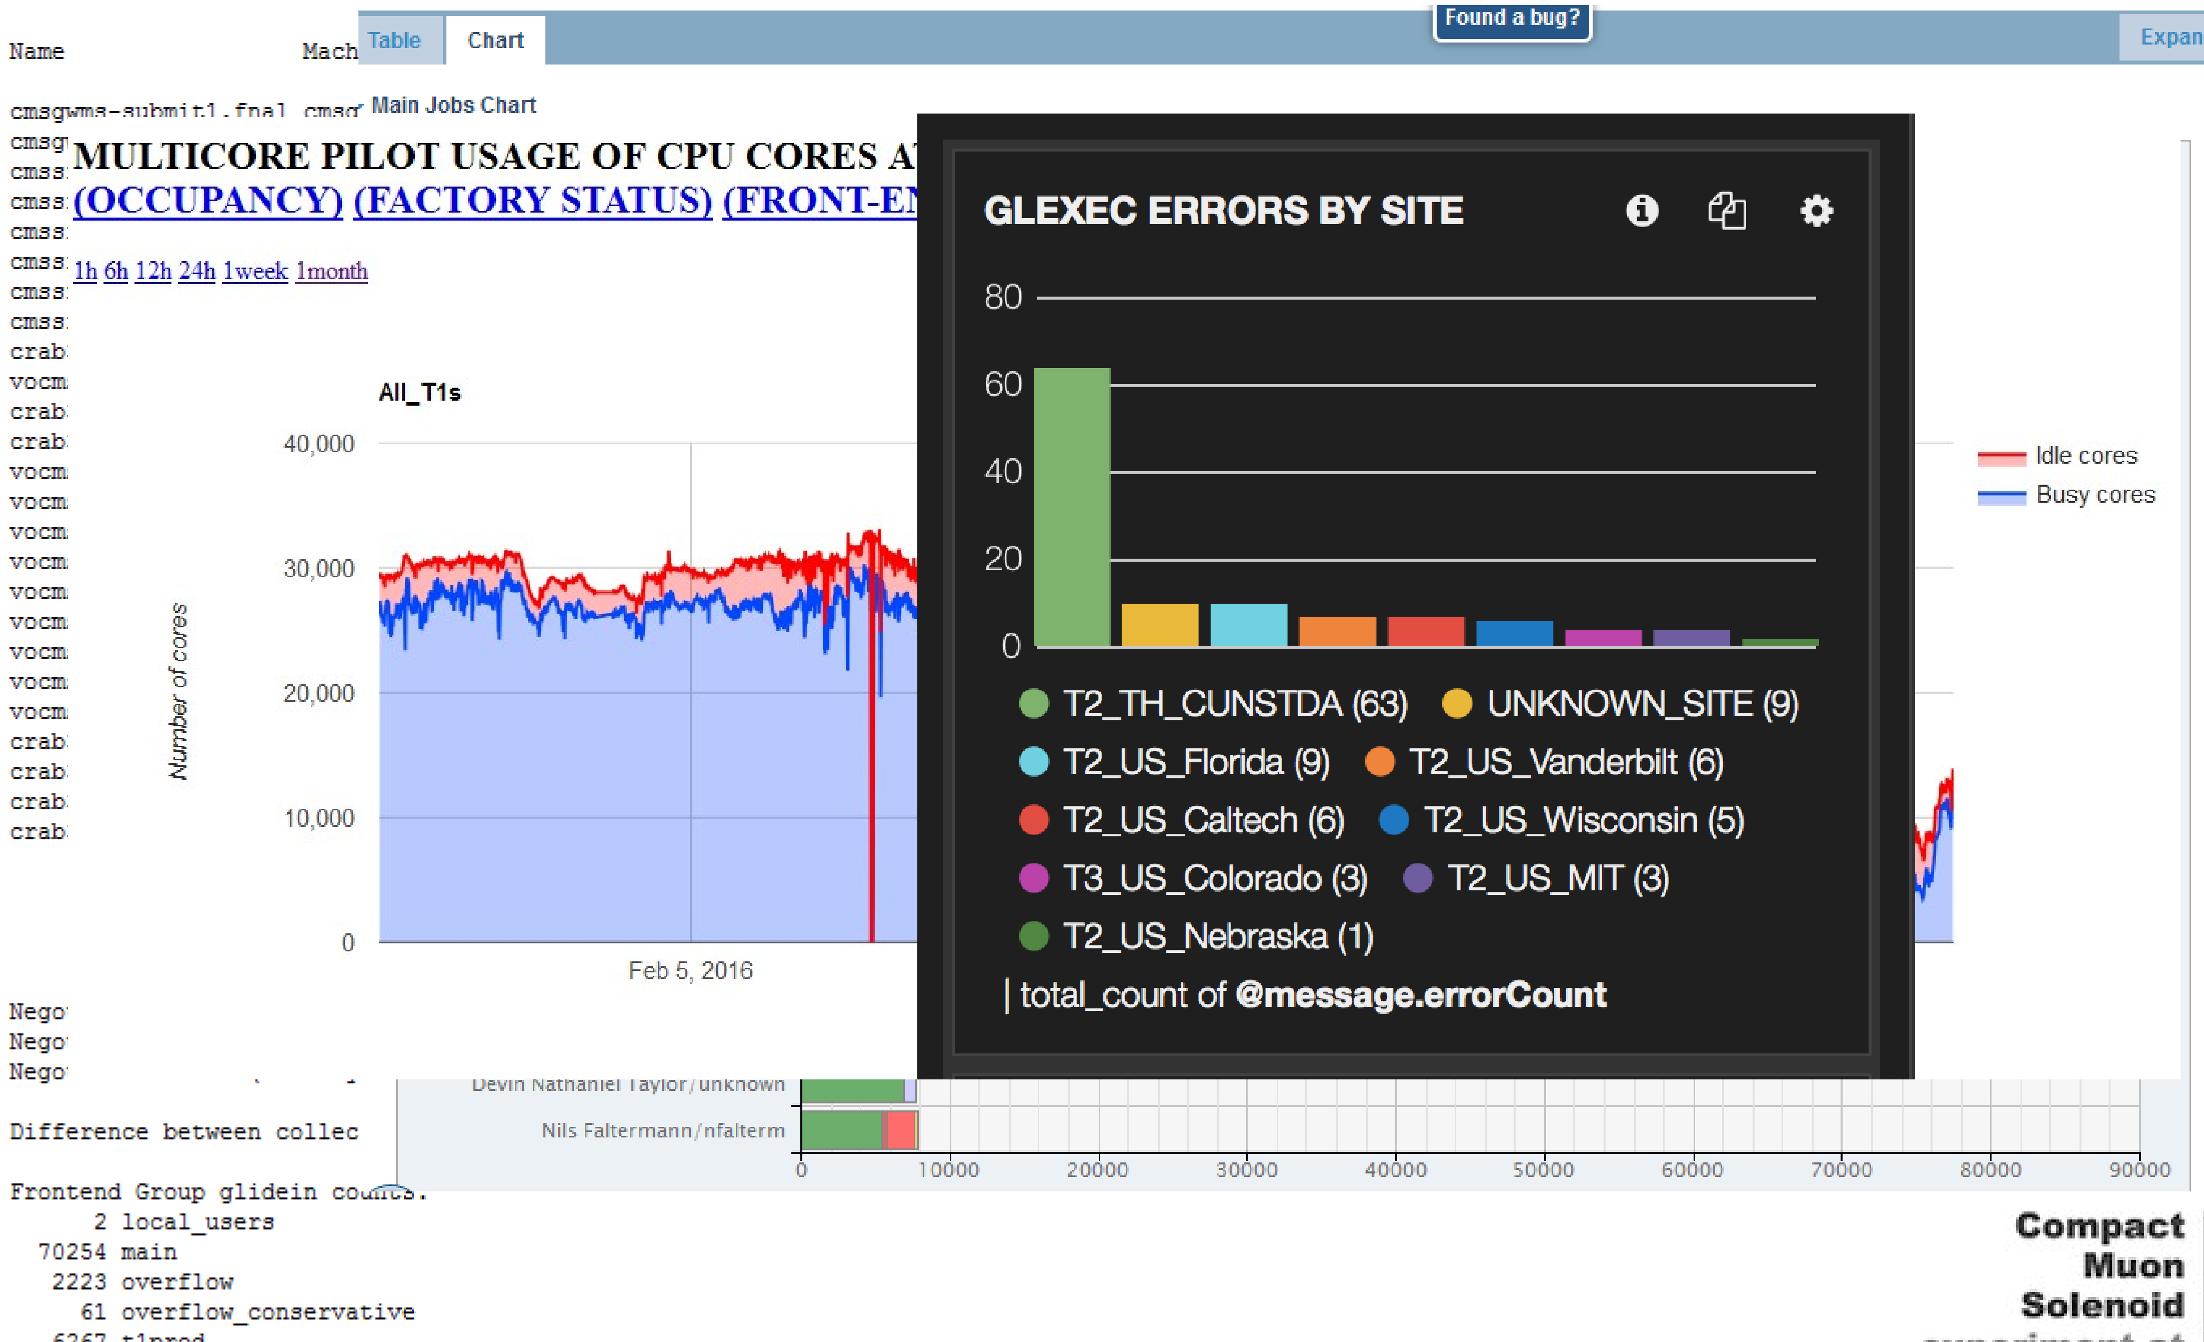

Compact Muon Solenoid experiment at CERN's LHC

CMS Global pool - Monitoring

Various different monitoring pages available to monitor the pool

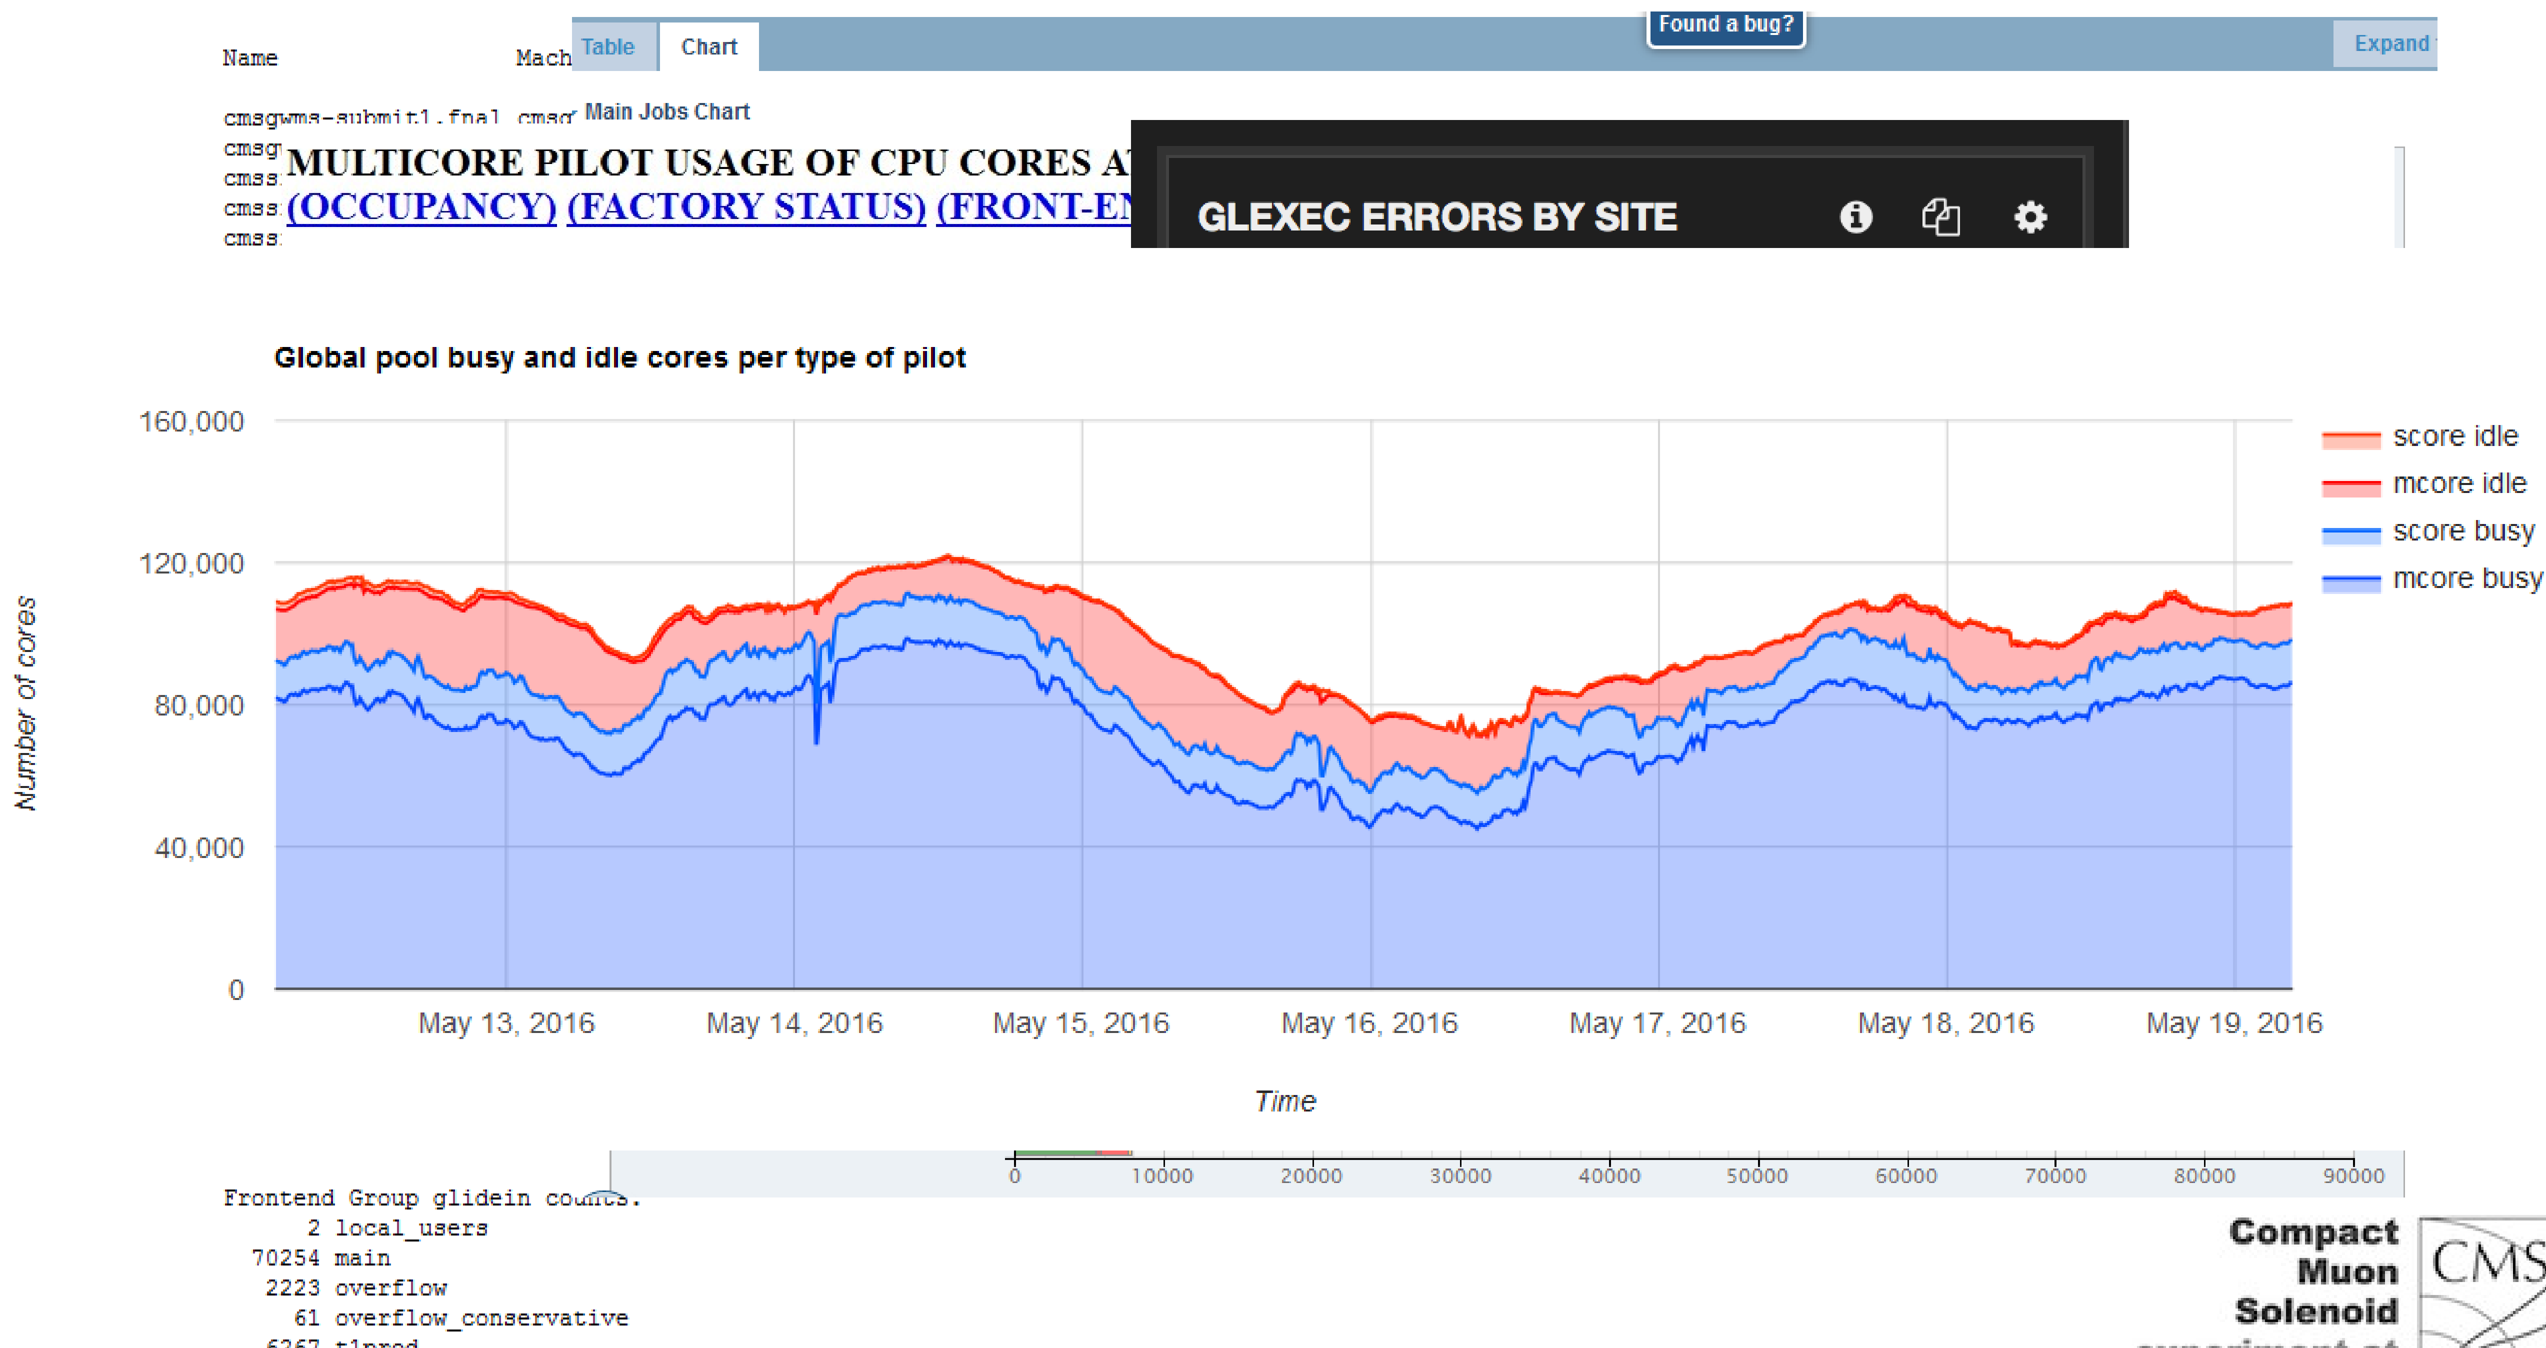

Compact Muon Solenoid experiment at CERN's LHC

CMS Global pool - Monitoring

HTCondor's innate ganglia support **VERY** helpful!

6267 tlprod

experiment at CERN's LHC

CMS Global pool - Monitoring

But the most comprehensive is cms-gwmsmon (more details in Justas' talk from EU HTCondor workshop (*))

Total Overview | Main Page | Total Overview | Production Overview | Analysis Overview | Analysis CRAB2 Overview | Pool Overview | Factory Overview

	Running	Idle	Unique Pressure	Tasks
Total	73657	76189	25497	1037
Production	20422	8611	362	114
Analysis	51488	64162	21722	895
CRAB2	1747	3416	3413	28
Update Time	2/29/2016, 8:06:40 PM			

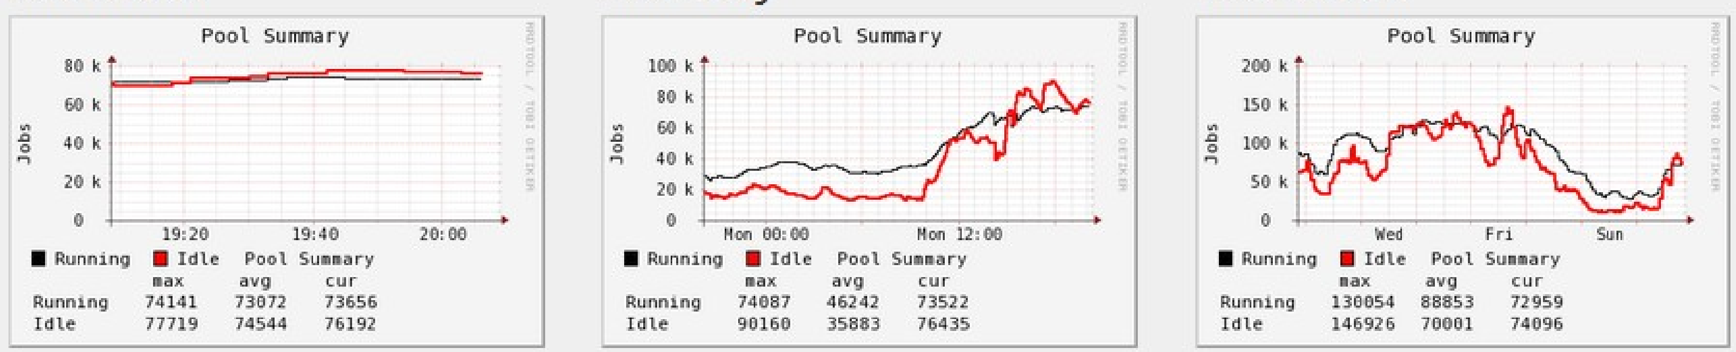

Site Summary

Site

Site	Running	Matching Idle	Unique Pressure	Max Running Achieved	Site History	Production History	CRAB3 History
T1_DE_KIT	231	2708	0				

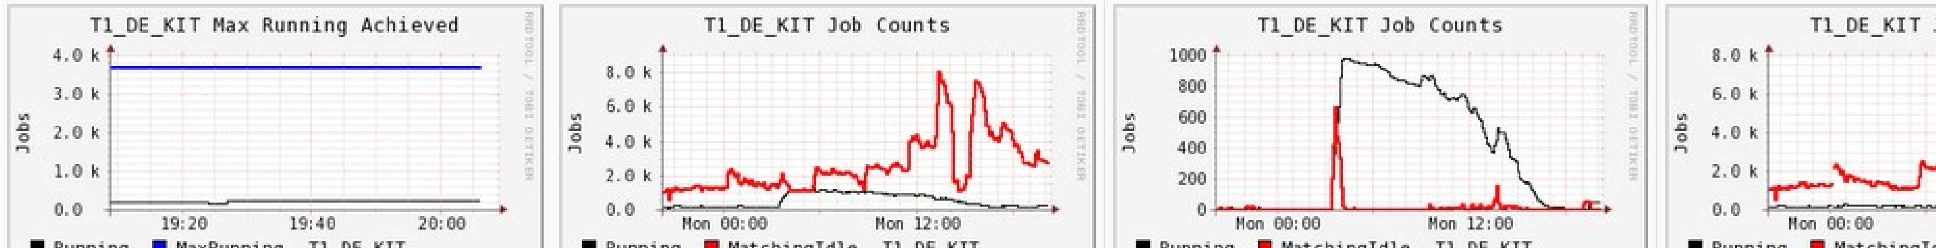

(*) http://slides.com/justasbalcas/monitoring-needs#/

Muon Solenoid experiment at CERN's LHC

CMS Global pool - Scalability

- OSG scale tests reached 200k parallel running jobs and have served as our benchmark throughout the last year or so
- We have only been able to reach 150k yet primarily due to a more chaotic job and resource mix than OSG tested
- Greater number of schedds and difference in specifications of hardware also a factor
- We have found and worked our way around various bottlenecks since the inception of the global pool thanks to the awesome (and constant) support provided by the HTCondor and glideinWMS developers
- Regular biweekly meeting

Compact Muon Solenoid experiment at CERN's LHC

CMS Global pool - Scalability

- HTCondor CCB separation
- Pre-fetching classAds for negotiation
- Dropping redundant startd updates to ease load on the collector
- Running a multi collector central manager
- Running a multi negotiator central manager
- Extra knobs to control time spent per schedd and per submitter

and many more..

HTCondor team always available for valuable advise and suggestions when we hit a wall, iterating with us until solutions are found

Compact Muon Solenoid experiment at CERN's LHC

Acknowledgements

- **Submission infrastructure group leaders:**

Antonio Perez-Calero Yzquierdo (PIC), David Alexander Mason (FNAL) and James Letts (UCSD)

- **GlideinWMS operations and development team at FNAL:**

Anthony Tiradani, Burt Holzman, Krista Larson, Marco Mambelli , Parag Mhashilkar and others

- **HTCondor developers**

- **OSG factory operations team:**

Brendan Denis, Jeffrey Dost, Martin Kandes and Vassil Verguilov

- **CRAB3 and WMAgent Operations team:**

Alan Malta, Justas Balcas, Jadir Silva, Marco Mascheroni and Stefano Belforte

- ***Brian Bockelman***

And many many others!

IMPORTANT!

In dire situations, HTCondor manual is your best friend!

Thank you!