

Lightweight site federation for CMS support

A. Delgado Peris on behalf of the CMS collaboration

C. Acosta-Silva^{1,2}, A. Delgado Peris³, J. Flix^{1,3}, J. M. Guerrero³, J. M. Hernández³, A. Pérez-Calero Yzquierdo^{1,3}, F. J. Rodriguez Calonge³, J. Gómez del Pulgar Ruano³

1 PIC, UAB, Bellaterra (Barcelona), Spain

2 IFAE, UAB, Bellaterra (Barcelona), Spain

3 CIEMAT, Madrid, Spain



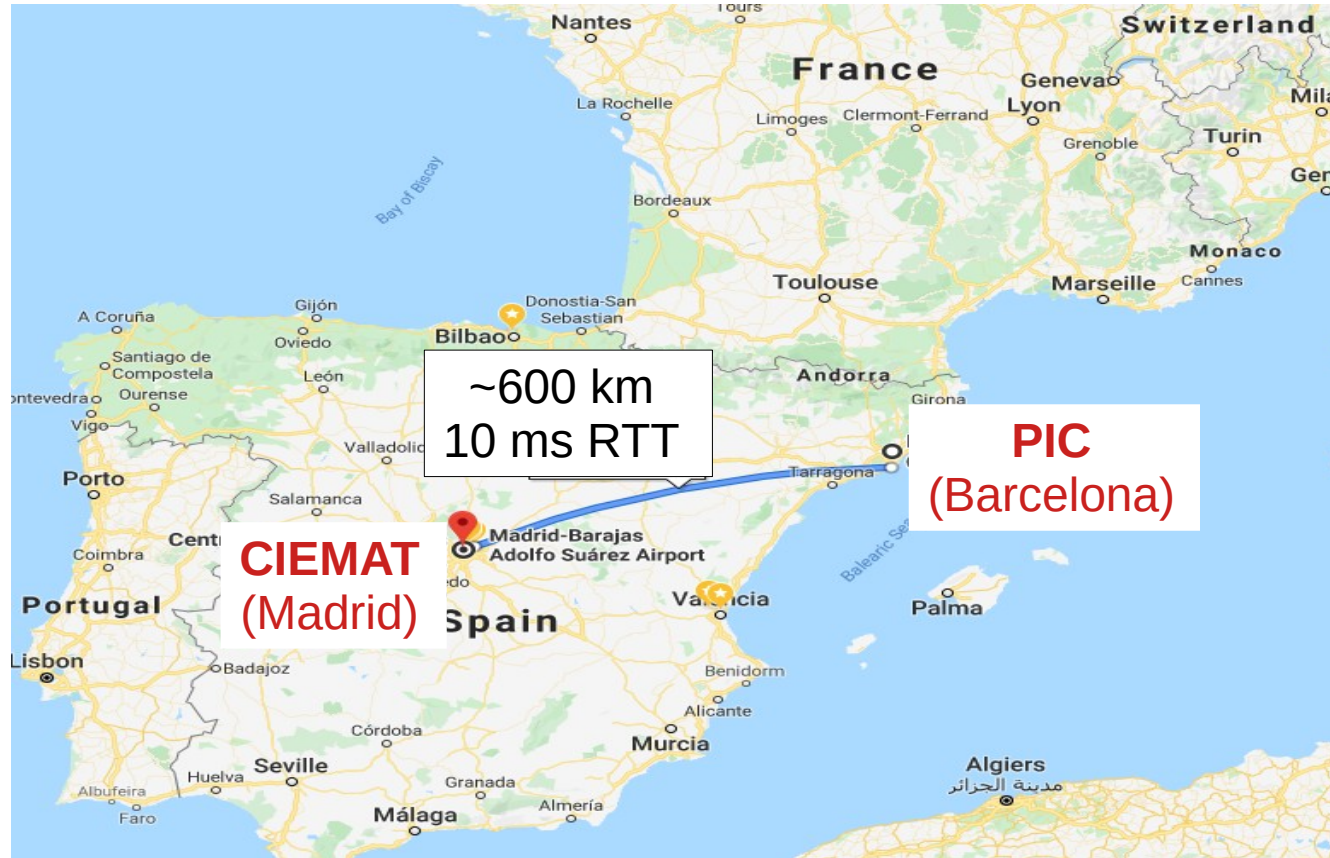
Outline

- Motivation and problem outline
- Implemented solution
- Overall results evaluation
- Results detailed view (*timewise*)
- Future work and conclusions

Federation of Computing Resources

- WLCG trend towards the federation of resources
 - Use resources/access data, of different sites, *transparently* (single access point)
 - Aim for higher simplicity, efficiency, flexibility, availability
- General solution for resource federation between independent and autonomous sites may prove difficult
- This work shows a simpler solution
 - Two close sites, shared CE/batch, one experiment (CMS)

The case of PIC and CIEMAT



The case of PIC and CIEMAT

- Relatively nearby locations
 - Good network interconnection (to be improved, midterm future)
- Shared CE/batch and SE/data-access technology
 - HTCondor and dCache/xrootd
- Long-standing close collaboration of both site teams
 - CMS Tier-1 and Tier-2 in the Spanish region
 - CIEMAT provides 50% of PIC funding



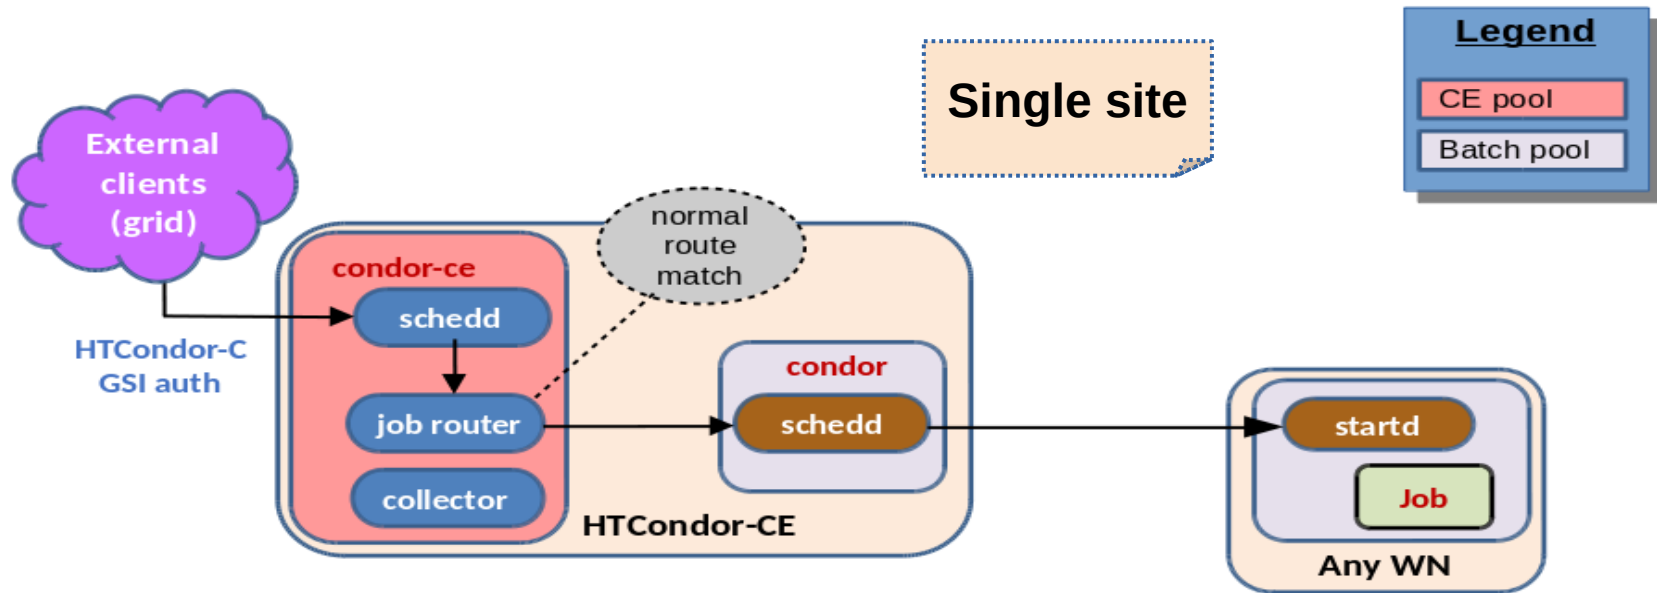
Outline

- Motivation and problem outline
- **Implemented solution**
- Overall results evaluation
- Results detailed view (timewise)
- Future work and conclusions

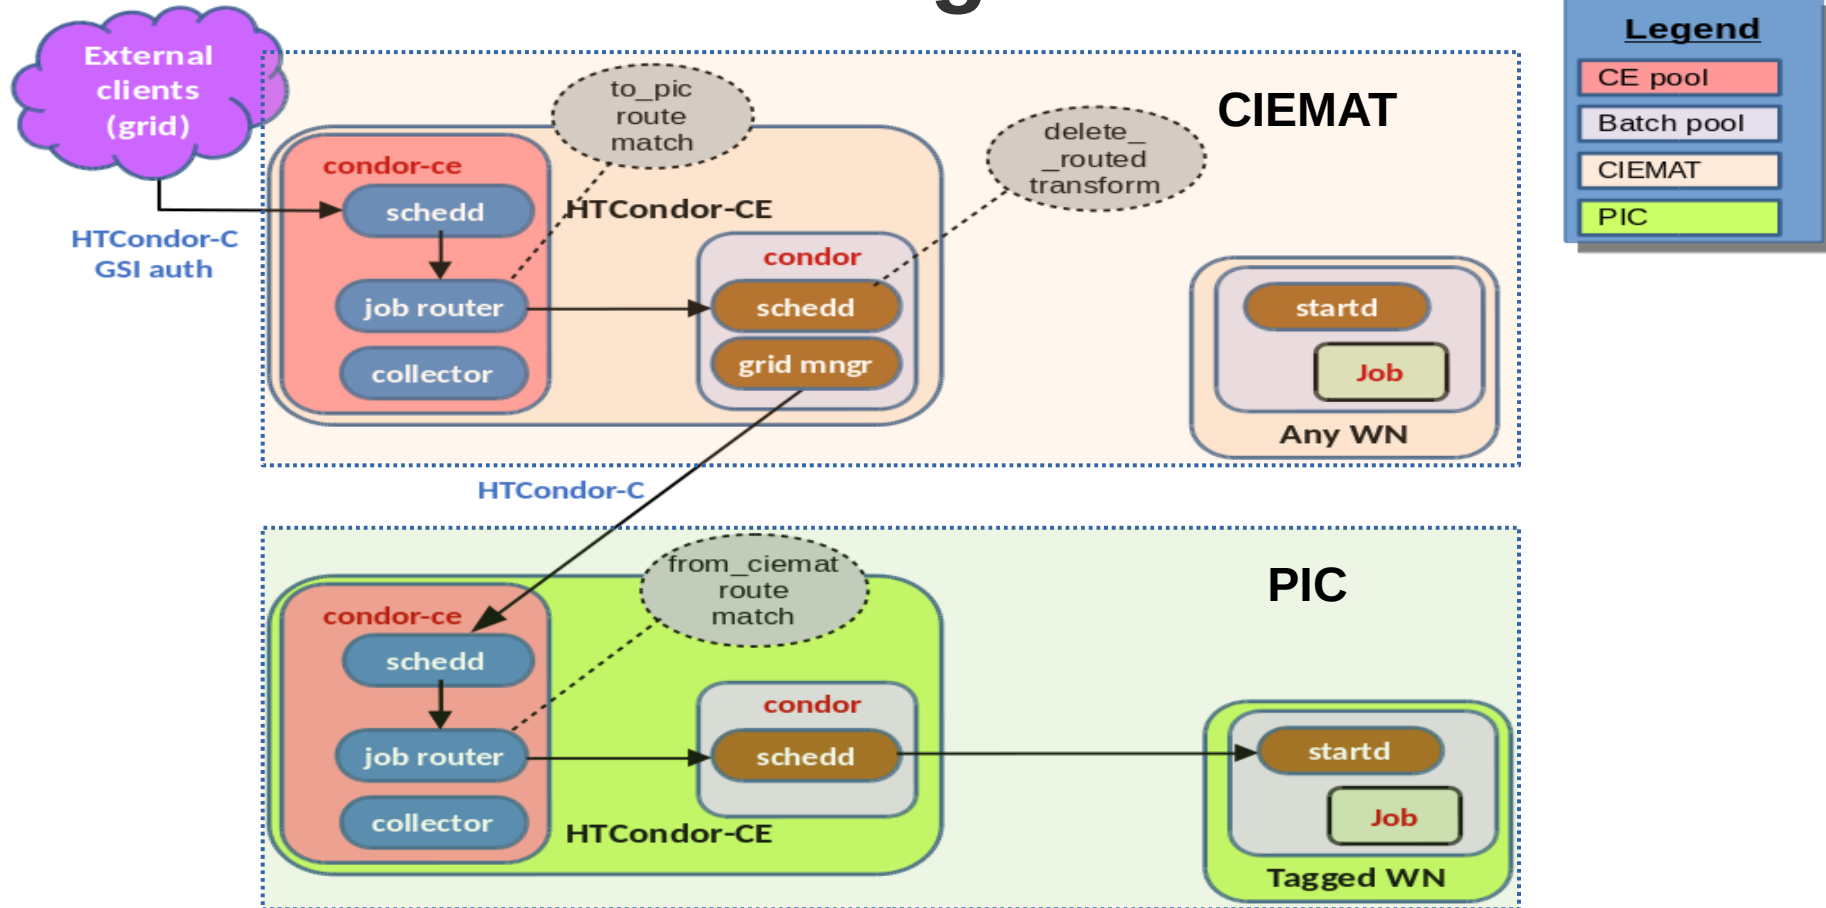
Job handling: HTCondor

- Take advantage of both sites using HTCondor as technology for CE and batch system
 - Many possible federation mechanisms
- HTCondor's *job re-routing* chosen
 - Completely functional for our test case
 - Keep independent configurations for the sites
 - Don't open networks, don't share daemon passwords
- Still, open to reassessment

HTCondor: standard deployment



HTCondor: re-routing



N. B.: we are re-routing pilots

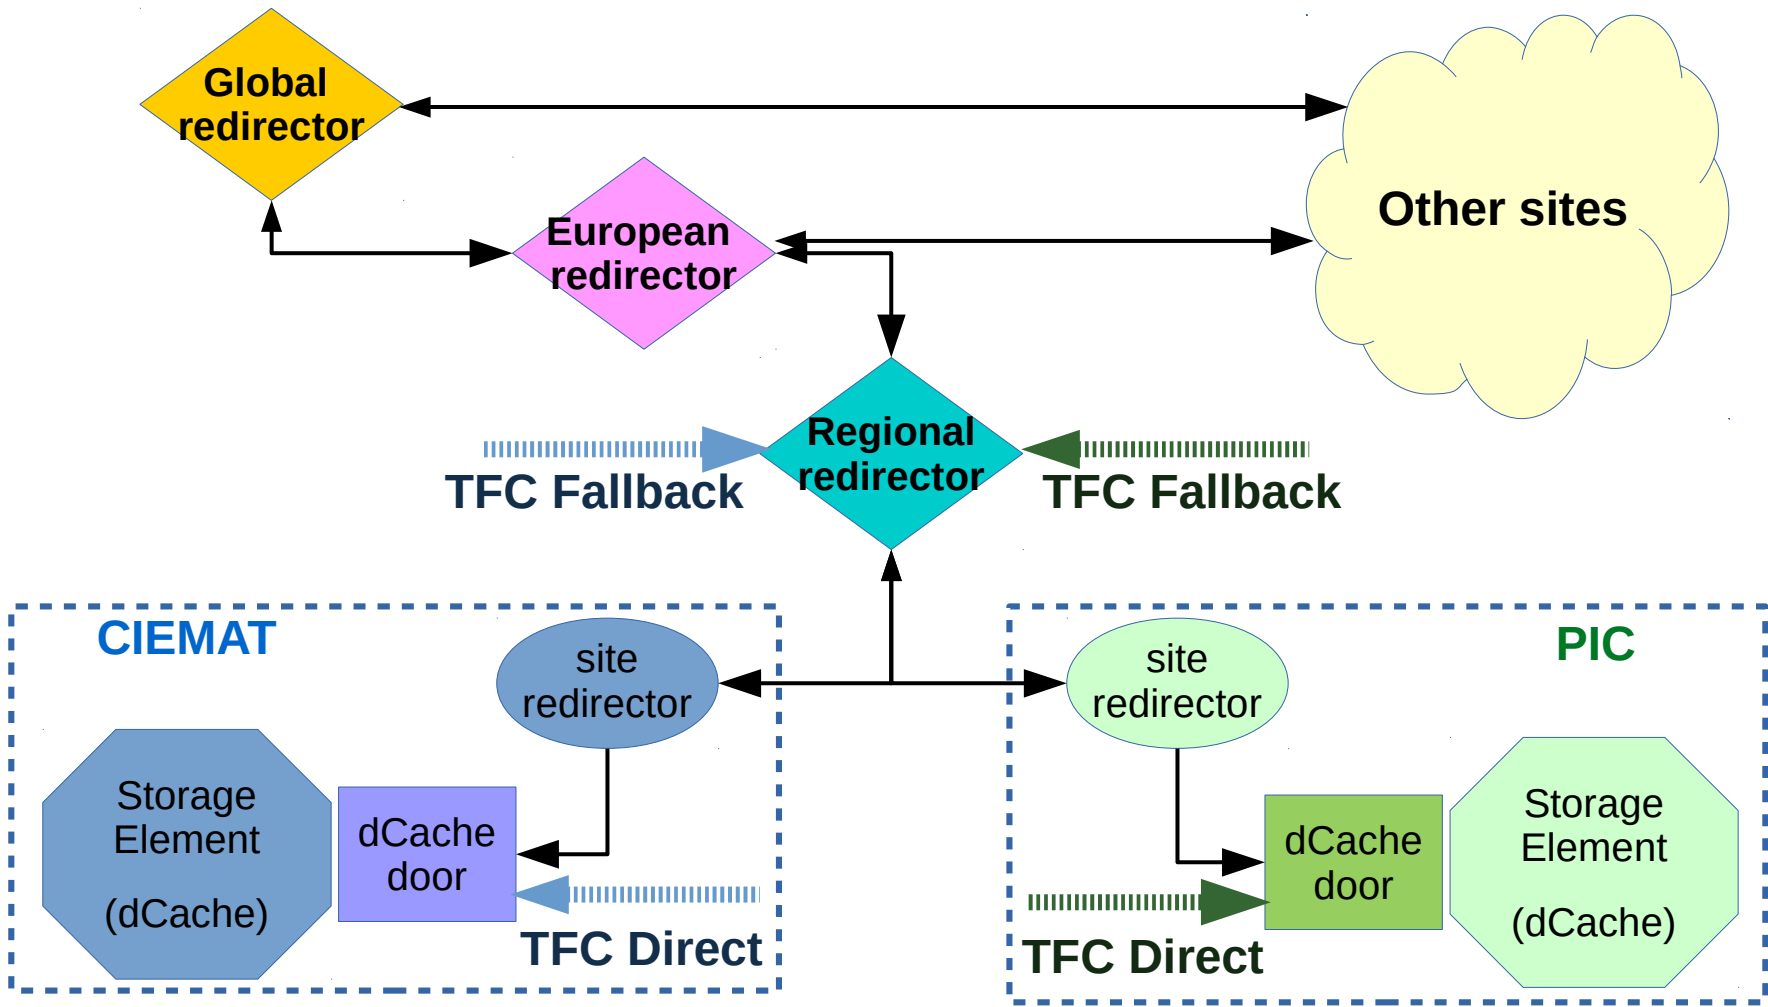
- CMS uses a late-binding scheduling model
 - Pilot jobs are sent to sites (*that's all the batch sees*)
 - Real payload jobs are matched to already-running pilots
- So: a *meant-for-CIEMAT* pilot might end up running at a PIC's worker node (and vice versa)
 - And: a *meant-for-CIEMAT* payload job might be matched by a re-routed pilot running at PIC (and vice versa)
- All pilots are essentially the same
 - They may match *any type* of payload job

HTCondor re-routing pitfalls

- Unavailability of any of the two involved sites affects the ability to re-route pilots.
 - Recovery automatic
 - WNs usable by *normal* local jobs (no loss of resources)
- Found a (possible) bug (reported)
 - `AssertionError on GridManager` caused by a single job prevents processing of further jobs
- Need to delete the `routed` attribute
 - This is known, but need to be aware

Access to data: xrootd

- CMS xrootd federation: transparent location discovery and access of data
- Added a new layer to favor access to *regional* (PIC-CIEMAT) data (rather than *elsewhere*)
 - Only for *locally running CMS jobs*
 - Only go *regional* when *data not found locally*
- Makes use of CMS *Trivial File Catalog*
 - *Direct* entry points to local dCache door
 - *Fall-back* entry points to regional redirector






Outline

- Motivation and problem outline
- Implemented solution
- **Overall results evaluation**
- Results detailed view (timewise)
- Future work and conclusions

Rates by job type: jobs meant for PIC



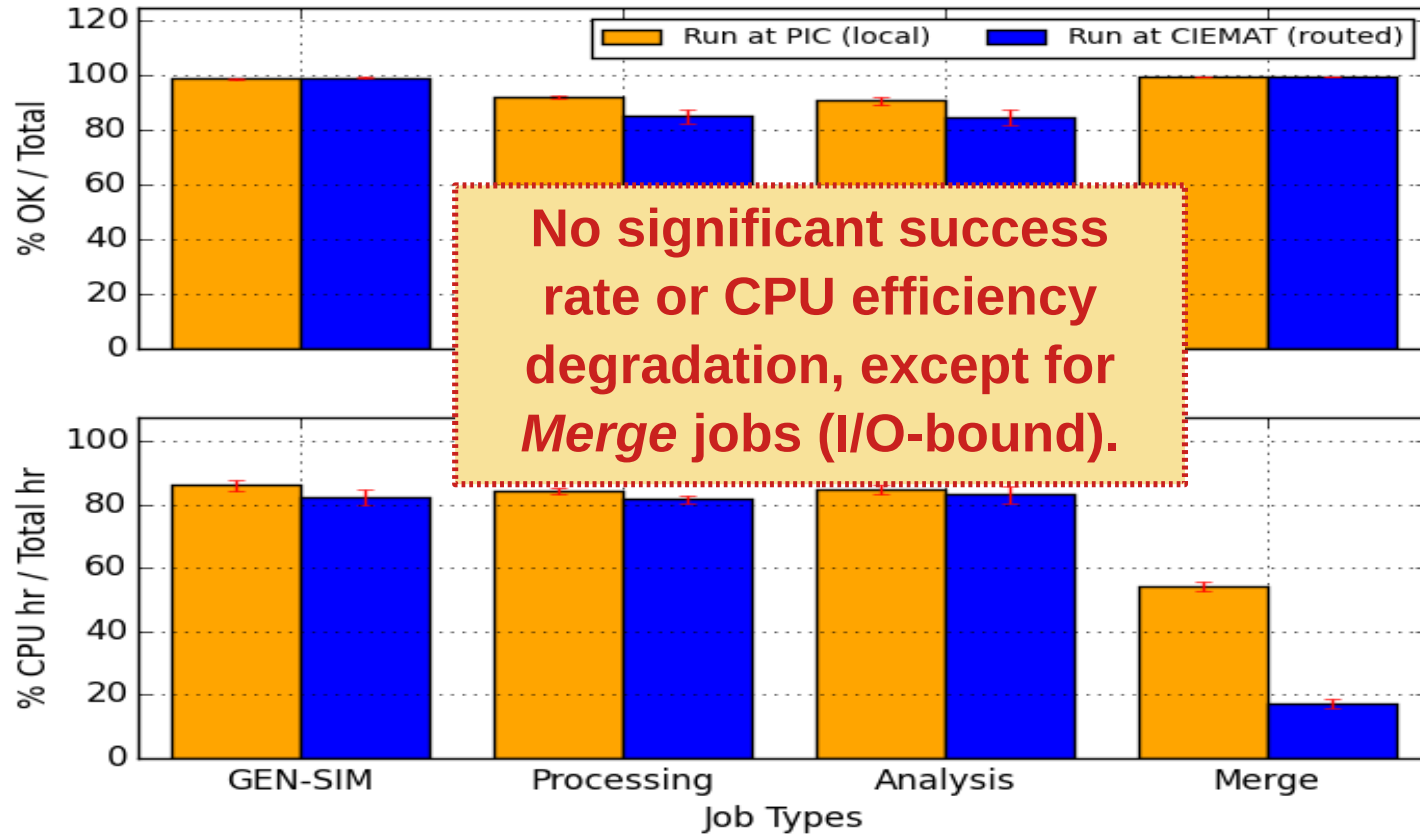
Re-routed

10
dedicated
8-core
pilot slots
per site


(up to 80
payloads)

6 months


Rates by job type: jobs meant for PIC




Rates by job type: jobs meant for CIEMAT



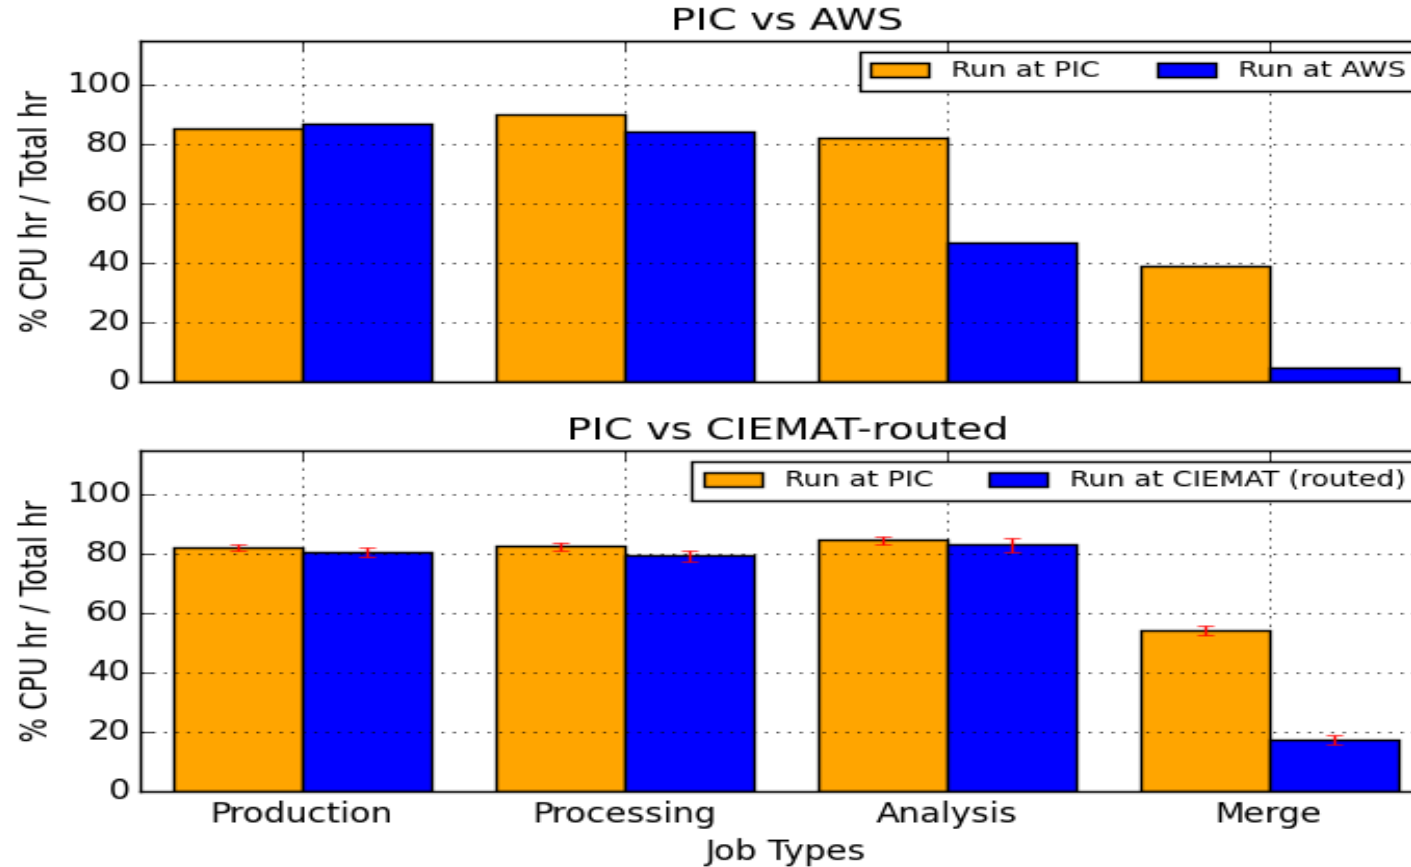
Rates by job type: jobs meant for CIEMAT



Lower success rate for CIEMAT's analysis jobs. Reasons unclear, but unrelated to re-routing.



CPU efficiency by job type: PIC vs CIEMAT vs AWS




Compare

Elastic extension of PIC farm on Amazon AWS

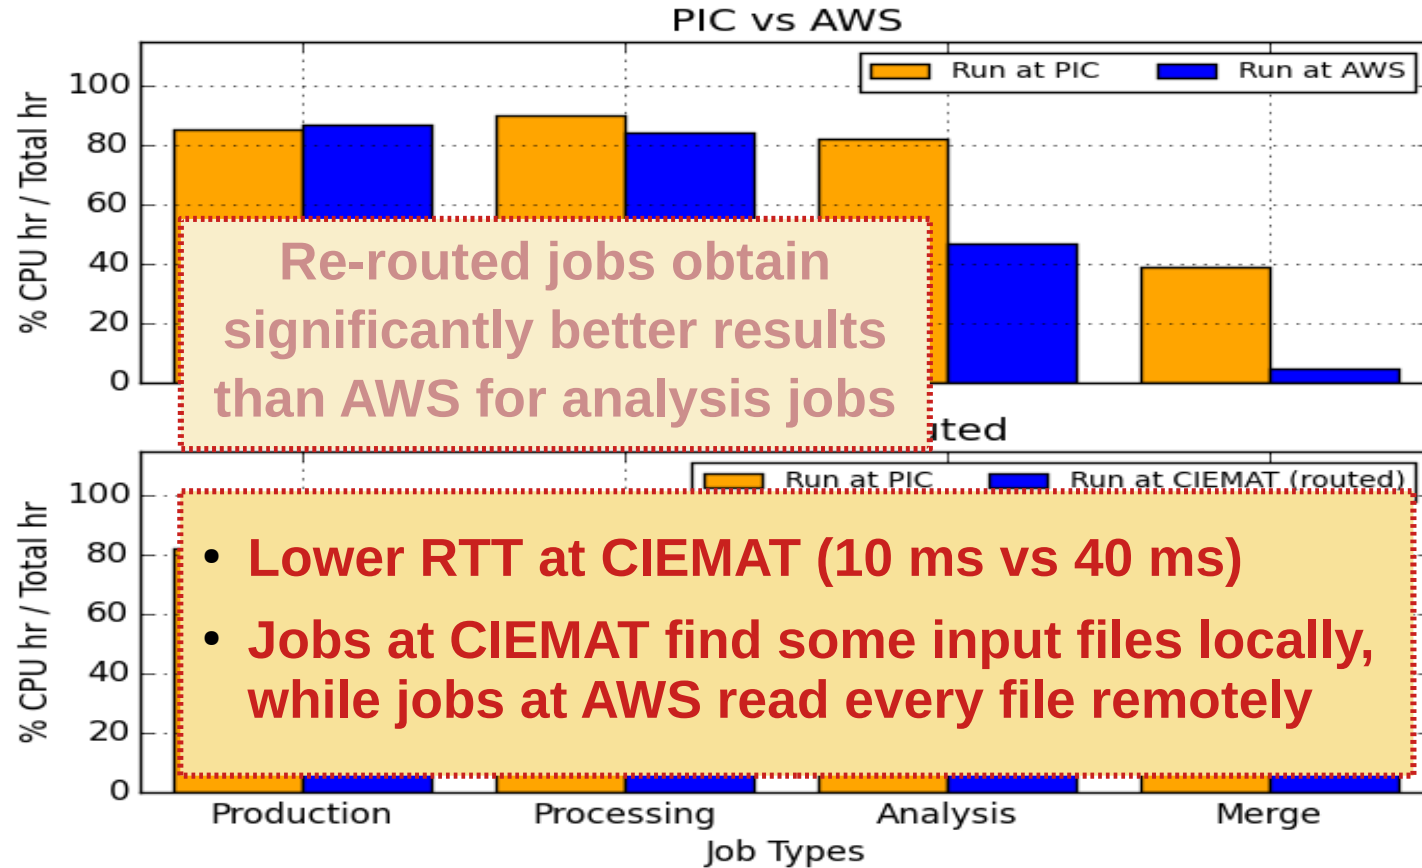
1-week test

Frankfurt center, 40 ms RTT


CPU efficiency by job type: PIC vs CIEMAT vs AWS



CPU efficiency by job type: PIC vs CIEMAT vs AWS




Rates by read origin for analysis jobs



**Closer
look**

Where
analysis
jobs really
read data
from

Rates by read origin for analysis jobs







Outline

- Motivation and problem outline
- Implemented solution
- Overall results evaluation
- **Results detailed view (*timewise*)**
- Future work and conclusions

Success/effic. for jobs meant for CIEMAT



Success/effic. for jobs meant for CIEMAT




Impact of CMS xrootd activity

- Heavy external xrootd access to CIEMAT/PIC dCache degrades performance
 - Limited number of xrootd movers per pool
 - Long-duration of connections
 - Unpredictable access *storms* (many clients)
- Sites deploy separate WAN/LAN queues
 - But regional access uses WAN queue
 - ➔ Can be affected by external xrootd activity


Analysis jobs at CIEMAT reading from PIC vs PIC xrootd

CLEANED



NOTE
 To ease visual correlation, outliers $> \text{mean} + 3\sigma$ have been removed.

Analysis jobs at PIC reading from CIEMAT vs CIEMAT xrootd **CLEANED**





Outline

- Motivation and problem outline
- Implemented solution
- Overall results evaluation
- Results detailed view (timewise)
- **Future work and conclusions**

Conclusions

- The deployed lightweight federation has proven to work as expected
 - Most jobs *just work*, with limited efficiency degradation
 - A fraction of analysis jobs suffer somewhat higher degradation, but still tolerable
- Matching of *Merge* (and auxiliary) jobs on re-routed pilots should be avoided (*see next slide*)
- For larger scales, a dedicated regional xrootd mover queue on dCache might be advisable



Future work

- Modify pilots to match only certain types of jobs
 - New CMS *site-customizable pilots*
- Data caching for re-routed pilots use (e.g. *Xcache*)
- Re-evaluate mechanisms to federate HTCondor pools
 - Final solution might be a single regional HTCondor pool
- Work on better monitoring for finer-grain studies
- Extend to other regions?
 - Where latency is tolerable



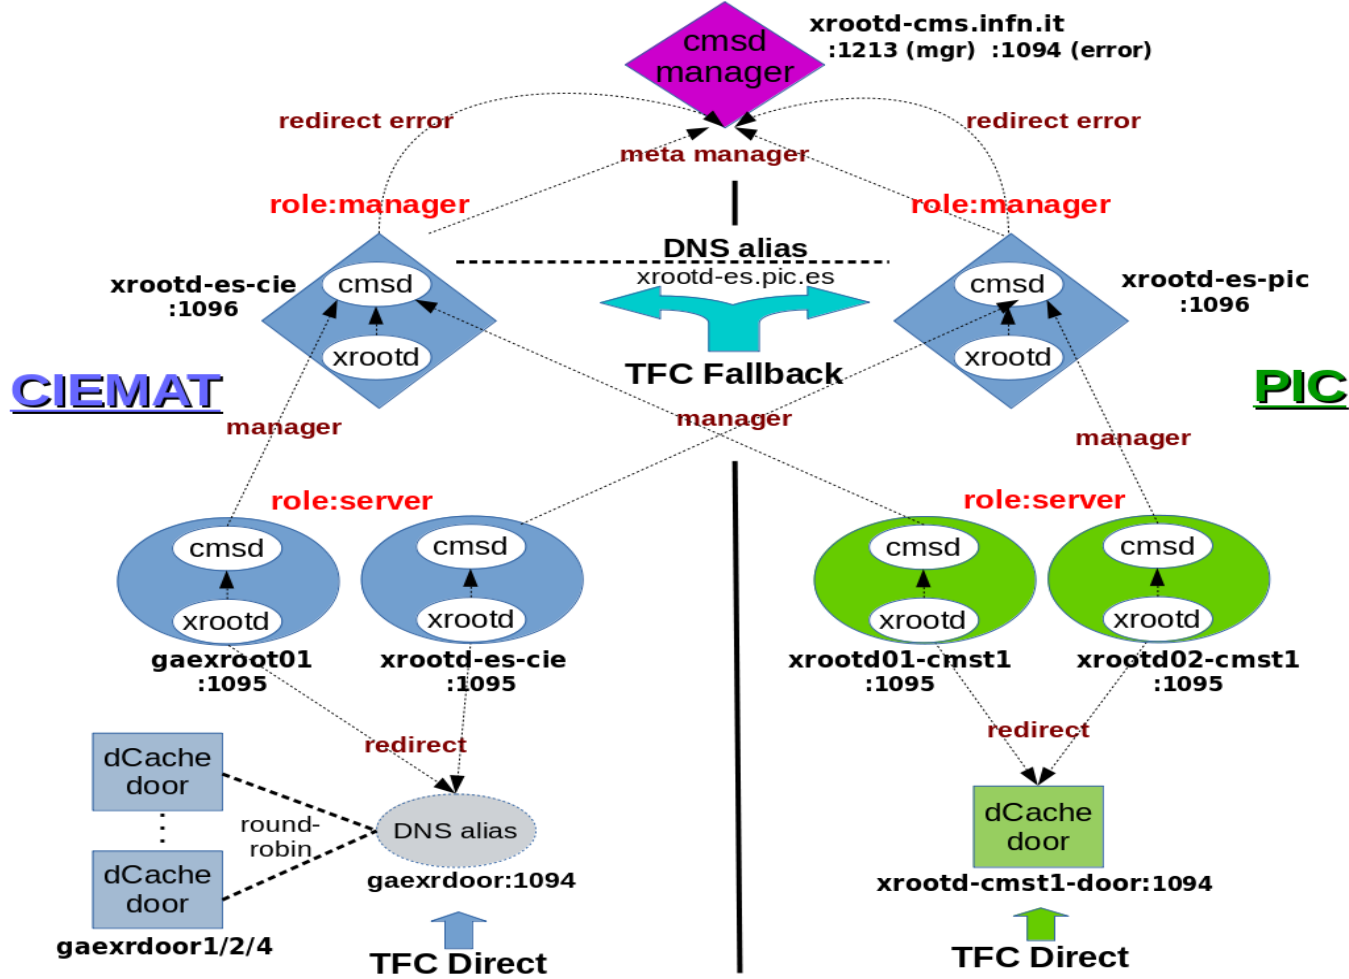
Thank you!

Questions?




Back-up Slides


European redirector




Rates by read origin for analysis jobs



Success/effic. for jobs meant for CIEMAT (analysis only)




Success/effic. for jobs meant for PIC



Success/effic. for jobs meant for PIC


Success rate



CIEMAT
dCache
problem


Single user
declaring
data at PIC
but reading
from USA

No re-routed
jobs (network
problem)

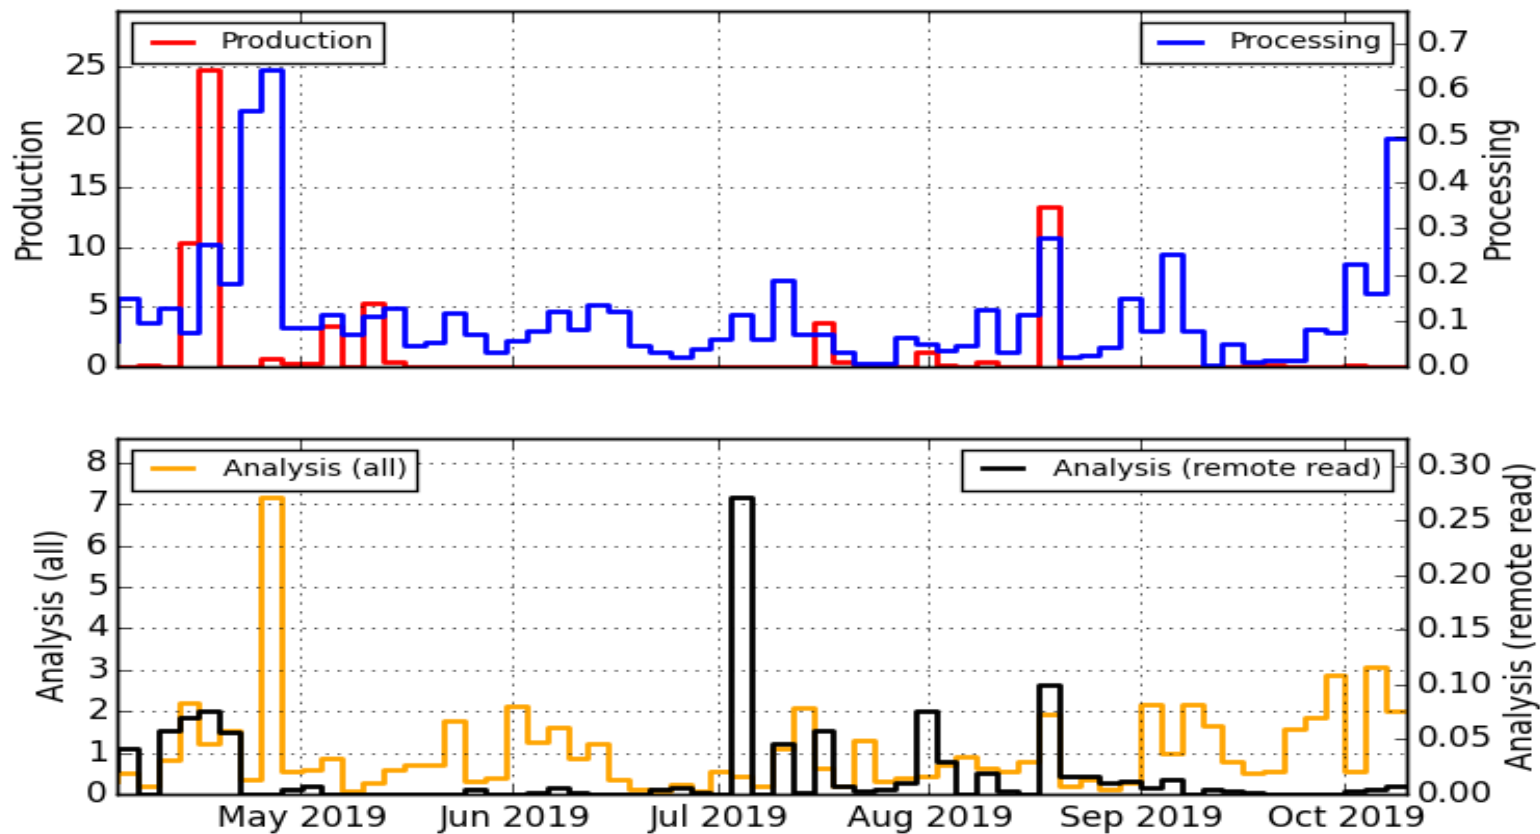


Production jobs at CIEMAT vs PIC xrootd

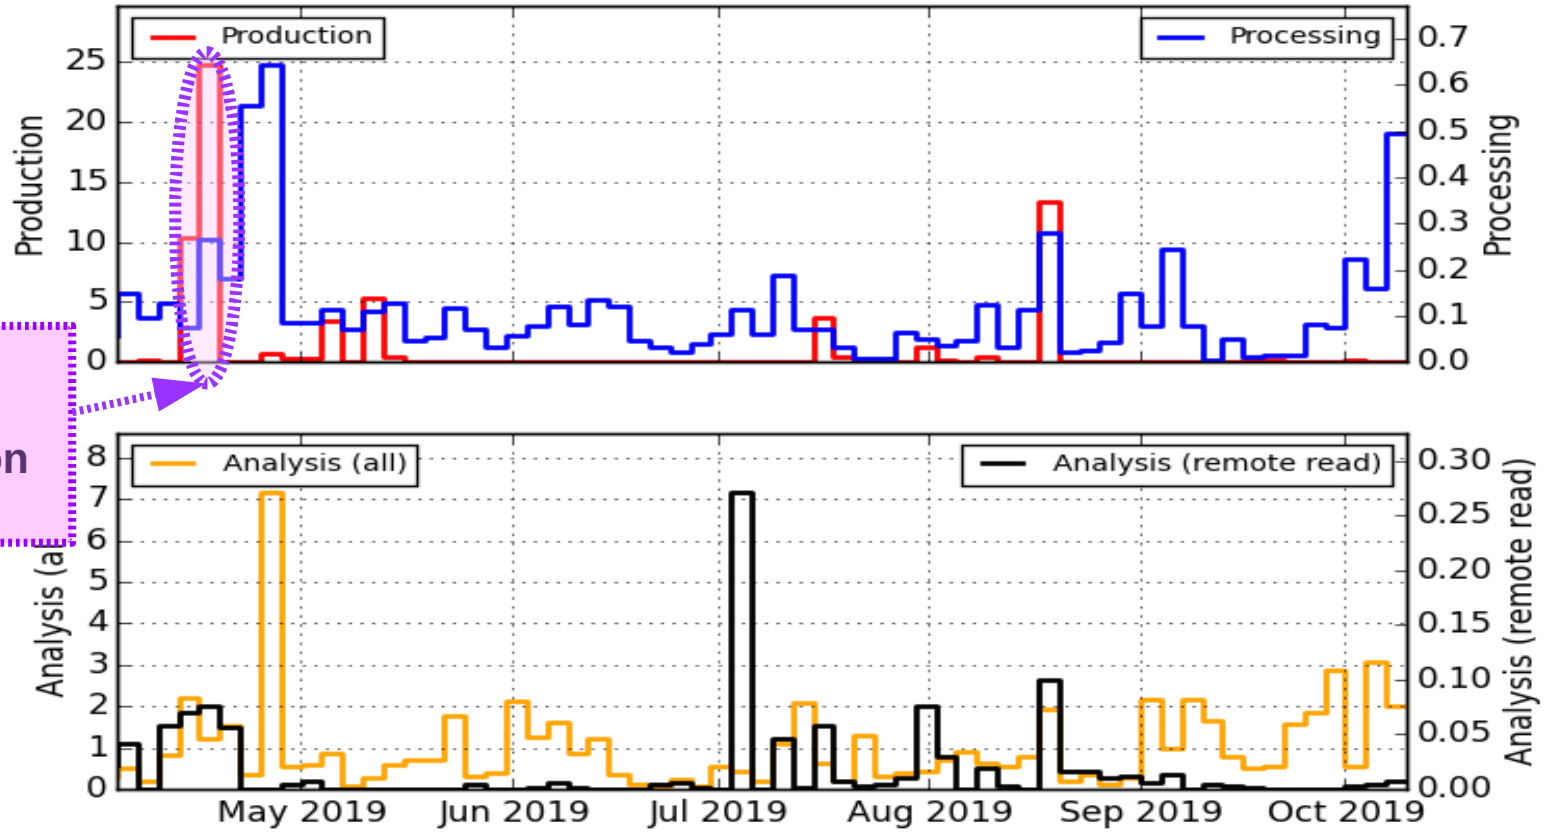
CLEANED



CIEMAT: Number of failed jobs per job type (thousands/day)




CIEMAT: Number of failed jobs per job type (thousands/day)




General CMS production problem


CIEMAT: Number of failed jobs per job type (thousands/day)




CIEMAT: Number of failed jobs per job type (thousands/day)




PIC: Number of failed jobs per job type (thousands/day)



PIC: Number of failed jobs per job type (thousands/day)




Success/effic. for jobs meant for PIC (analysis only)




Analysis jobs at CIEMAT reading from PIC vs PIC xrootd

Num failed jobs




CPU efficiency




Production jobs at CIEMAT vs PIC xrootd

Num failed jobs



CPU efficiency



Analysis jobs at PIC reading from CIEMAT vs CIEMAT xrootd

