

CMS experience of running glideinWMS in High Availability mode

by
 I Sfiligoi¹, J Letts¹, S Belforte², A Mc Crea¹, K Larson³, M Zvada⁴, B Holzman³, P Mhashilkar³, D C Bradley⁵,
 M D Saiz Santos¹, F Fanzago⁶, O Gutsche³, T Martin¹ and F Würthwein¹
 for the CMS computing group

¹University of California San Diego, La Jolla, CA 92093, USA

²Universita e INFN di Trieste, 34127 Trieste, Italy

³Fermi National Laboratory, Batavia, IL 60510, USA

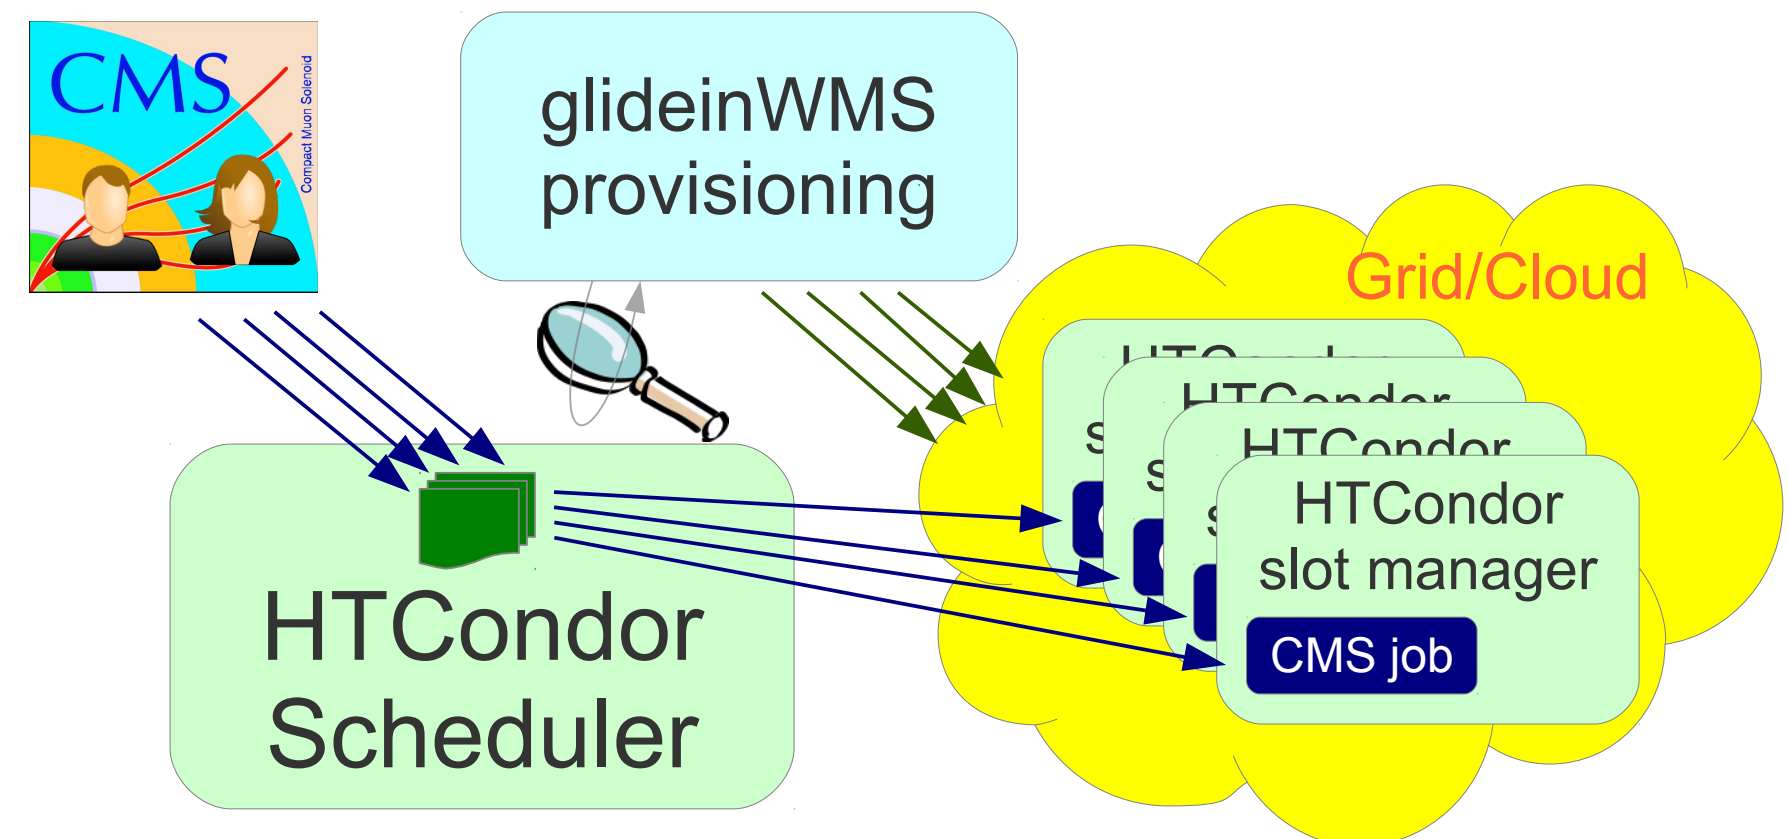
⁴Karlsruhe Institute of Technology, 76131 Karlsruhe, Germany

⁵University of Wisconsin – Madison, Madison, WI 53706, USA

⁶Universita e INFN di Padova, 35131 Padova, Italy

CMS use of glideinWMS

The CMS experiment at the Large Hadron Collider is relying on the HTCondor-based glideinWMS batch system to handle most of its distributed computing needs.



glideinWMS provisioning

VO Frontend
 Implements provisioning policy

Abstraction layer toward Grid/Cloud

Glidein Factory

HTCondor slot manager a.k.a. glidein
 Manage compute resources

Startd

Keeps list of all HTCondor processes

Collector

Acts also as network router (CCB)

Keep list of all user jobs

Schedd

Negotiator

Implements scheduling policy

glideinWMS HA capabilities

Glidein Factory

- No decision logic
- Slave to VO frontend
- Can run several independent instances in parallel, both for
 - Load balancing and
 - High Availability
- Loss of an instance results only in loss of pilot monitoring

VO Frontend

- No persistent state
- Could run several instances for HA
 - May result in slight over-provisioning
- Not in critical path
 - Hour long downtime acceptable
 - Cold spare good enough

Collector

- No persistent state
- Can run several independent instances
 - All receiving all data
 - All potentially routing traffic for all glideins
 - i.e. no load balancing
- As long as one instance is functioning, no loss

Negotiator

- Has persistent state
 - Only one can be active at any point in time
- HTCondor provides triage and sync service
- Can run several instances in hot spare mode

No good HA solution for the **Schedds**, due to user jobs' sandboxes creating large persistent footprint.

- Round-robin workflow partitioning can be used to limit damage in case of hardware failure.

CMS glideinWMS HA experience

(AnaOps only, production pool similar – See poster #112 for more details)

glideinWMS prov.

- Four factories
 - 3 in USA, 3 states
 - 1 at CERN
- One active, one cold spare Frontend
- both in USA, separate availability zones

HTCondor daemons

- Two collector/negotiator pairs
 - Both in USA, separate availability zones
- Five schedd
 - 3 in USA, in 3 different availability zones
 - 2 at CERN

Switched from single collector to HA mode live

- With O(30k) glideins running

Maintenance done of one of the Collectors

- No hiccups with O(30k) glideins running

Requires HTCondor 8.0.1+

- Found and fixed a bug on test system

