Evolution of HammerCloud to commission CERN Compute resources

Jaroslava Schovancová, Alessandro Di Girolamo, Aristeidis Fkiaras, Valentina Mancinelli
(CERN)

HammerCloud at a glance

- Functional and stress tests of WLCG resources: ATLAS, CMS; Batch
  - **Functional:** steady flow of test jobs
  - **Stress:** on demand tests, configure load intensity
- Part of automation suite of the Experiments
- Testing the **full chain** of an Experiment job
  - Same environment as the “real” analysis/production jobs
- Utilization
  - ATLAS: 80k jobs/day, ~30 tests/day
  - CMS: 39k jobs/day, 36 tests/day
  - Batch: 150-750 jobs/day, ~1-2 tests/day

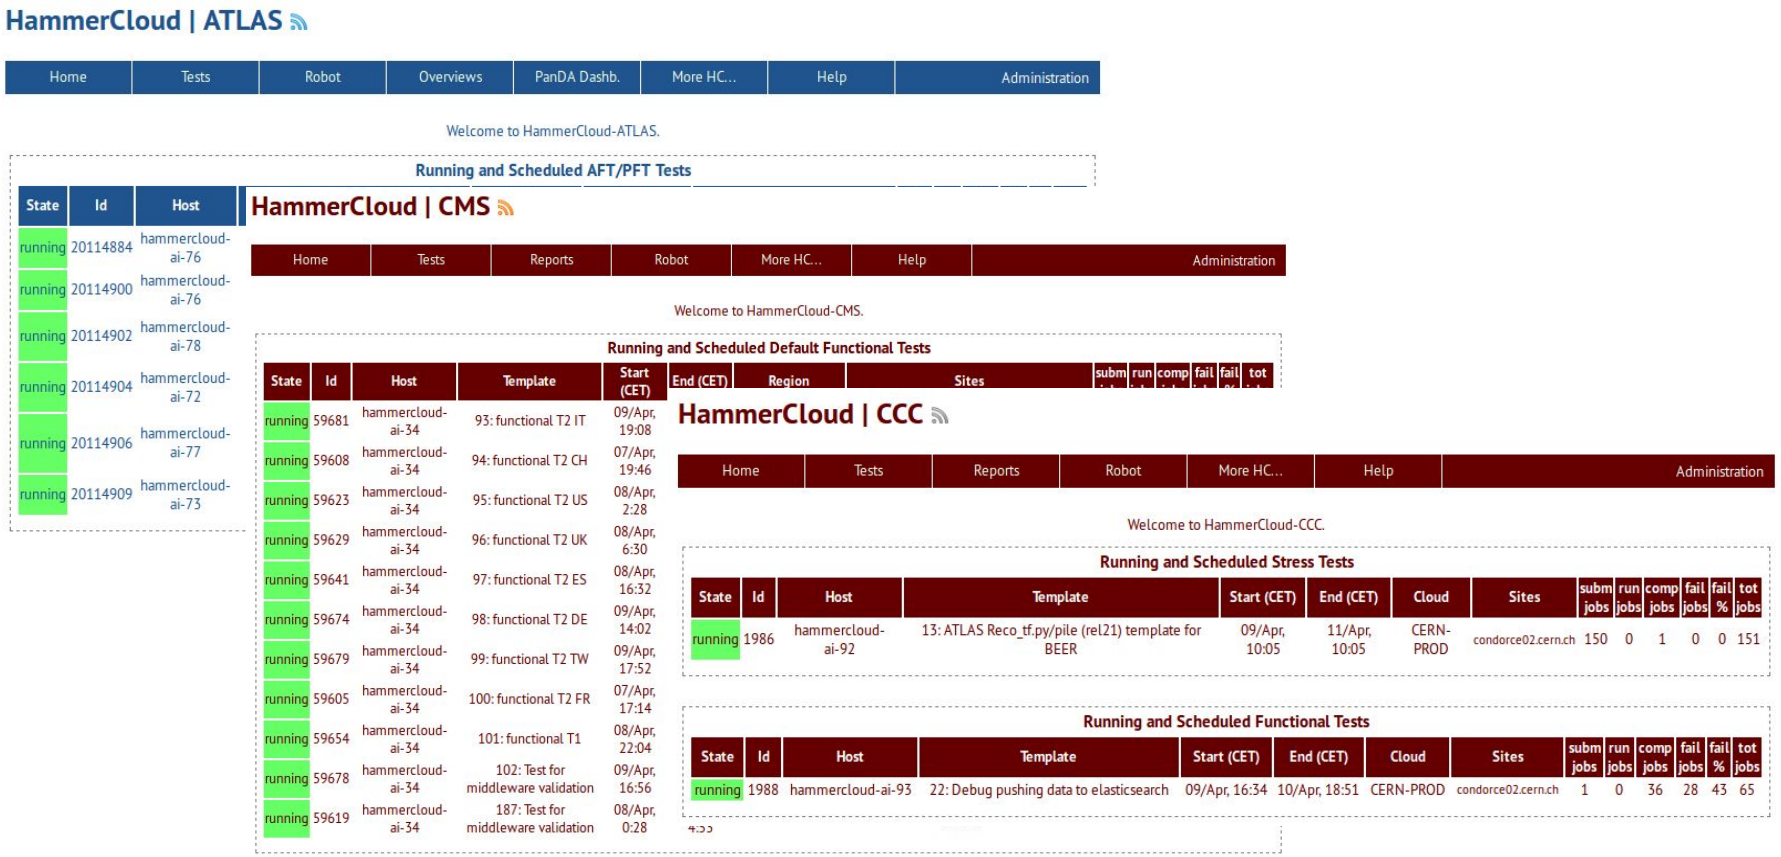
HammerCloud | ATLAS
Home
Tests
Robot
Overviews
PanDA Dashb.
More HC...
Help
Administration
Welcome to HammerCloud-ATLAS.
Running and Scheduled AFT/PFT Tests
State
Id
Host
running 20114884 hammercloud-ai-76
running 20114900 hammercloud-ai-76
running 20114902 hammercloud-ai-78
running 20114904 hammercloud-ai-72
running 20114906 hammercloud-ai-77
running 20114909 hammercloud-ai-73
HammerCloud | CMS
Home
Tests
Reports
Robot
More HC...
Help
Administration
Welcome to HammerCloud-CMS.
Running and Scheduled Default Functional Tests
State
Id
Host
Template
Start (CET)
End (CET)
Region
Sites
subm jobs
run jobs
comp jobs
fail jobs
fail %
tot jobs
running 59681 hammercloud-ai-34 93: functional T2 IT 09/Apr, 19:08
running 59608 hammercloud-ai-34 94: functional T2 CH 07/Apr, 19:46
running 59623 hammercloud-ai-34 95: functional T2 US 08/Apr, 2:28
running 59629 hammercloud-ai-34 96: functional T2 UK 08/Apr, 6:30
running 59641 hammercloud-ai-34 97: functional T2 ES 08/Apr, 16:32
running 59674 hammercloud-ai-34 98: functional T2 DE 09/Apr, 14:02
running 59679 hammercloud-ai-34 99: functional T2 TW 09/Apr, 17:52
running 59605 hammercloud-ai-34 100: functional T2 FR 07/Apr, 17:14
running 59654 hammercloud-ai-34 101: functional T1 08/Apr, 22:04
running 59678 hammercloud-ai-34 102: Test for middleware validation 09/Apr, 16:56
running 59619 hammercloud-ai-34 187: Test for middleware validation 08/Apr, 0:28
HammerCloud | CCC
Home
Tests
Reports
Robot
More HC...
Help
Administration
Welcome to HammerCloud-CCC.
Running and Scheduled Stress Tests
State
Id
Host
Template
Start (CET)
End (CET)
Cloud
Sites
subm jobs
run jobs
comp jobs
fail jobs
fail %
tot jobs
running 1986 hammercloud-ai-92 13: ATLAS Reco_tf.py/pile (rel21) template for BEER 09/Apr, 10:05 11/Apr, 10:05 CERN-PROD condorce02.cern.ch 150 0 1 0 0 151
Running and Scheduled Functional Tests
State
Id
Host
Template
Start (CET)
End (CET)
Cloud
Sites
subm jobs
run jobs
comp jobs
fail jobs
fail %
tot jobs
running 1988 hammercloud-ai-93 22: Debug pushing data to elasticsearch 09/Apr, 16:34 10/Apr, 18:51 CERN-PROD condorce02.cern.ch 1 0 36 28 43 65

HammerCloud activities

ATLAS: functional testing & auto-exclusion of resources; ESblacklist; commissioning of new resources; commissioning of new components of distributed computing systems (Pilot, Rucio, new data access protocols, …); FT of services (ObjectStore testing)

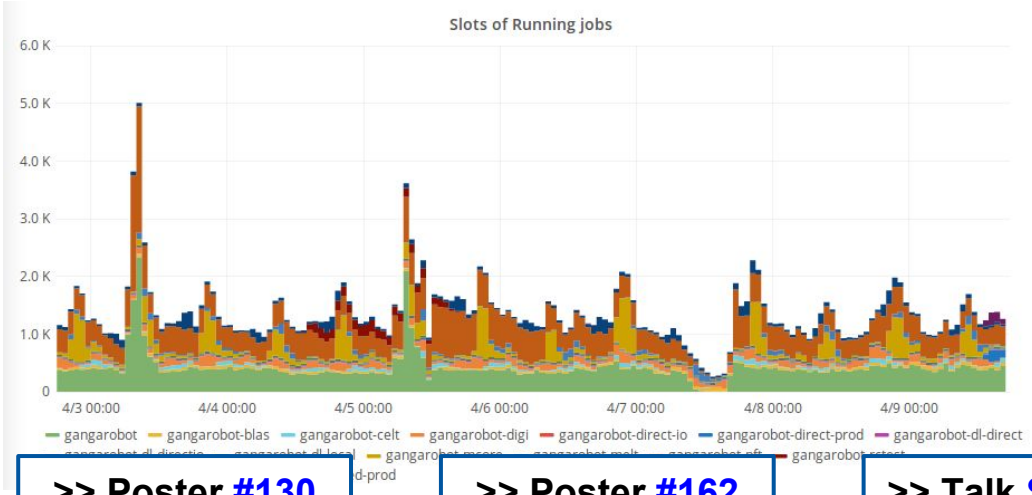

CMS: functional testing; commissioning of new resources; commissioning of new components of distributed computing systems

Batch: BEER, external cloud, CI/CD, containers usability, ...

>> Poster #130

>> Poster #162

>> Talk Sharing server nodes for storage and compute

ATLAS HammerCloud auto-exclusion

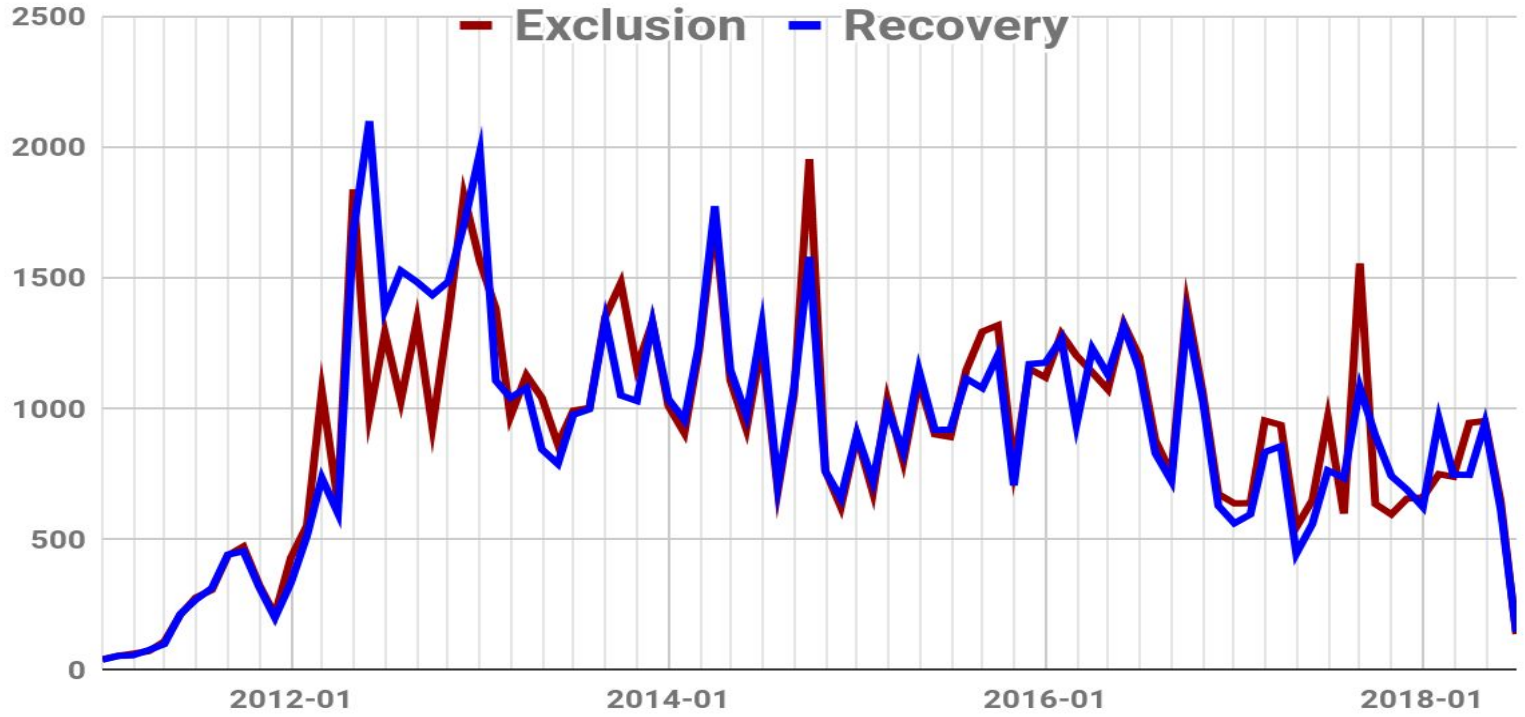

HammerCloud from far away

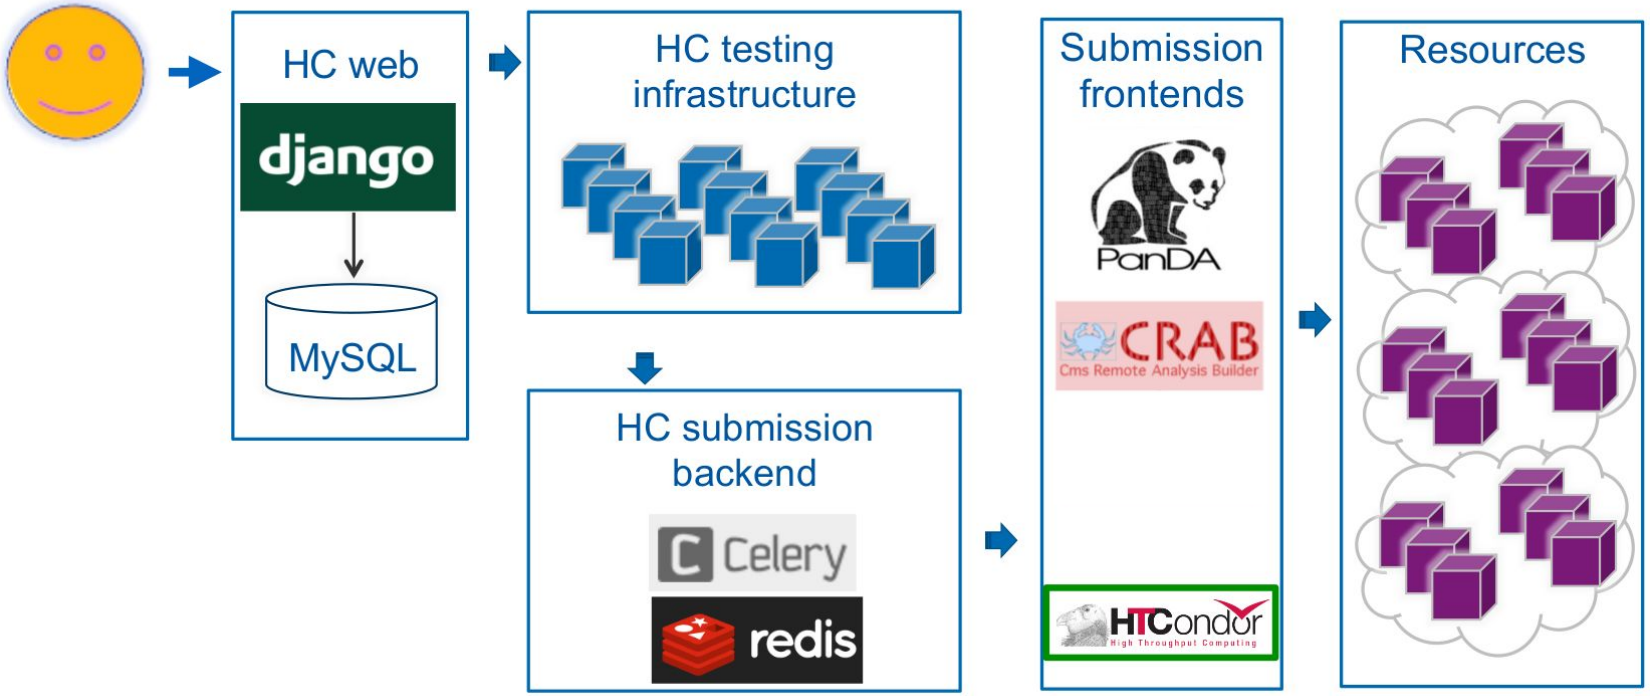

Why new submission backend(s)?

Adding compute resources: get as many CPU cycles as possible ASAP

… a complex task!

distributed, planet-scale computing

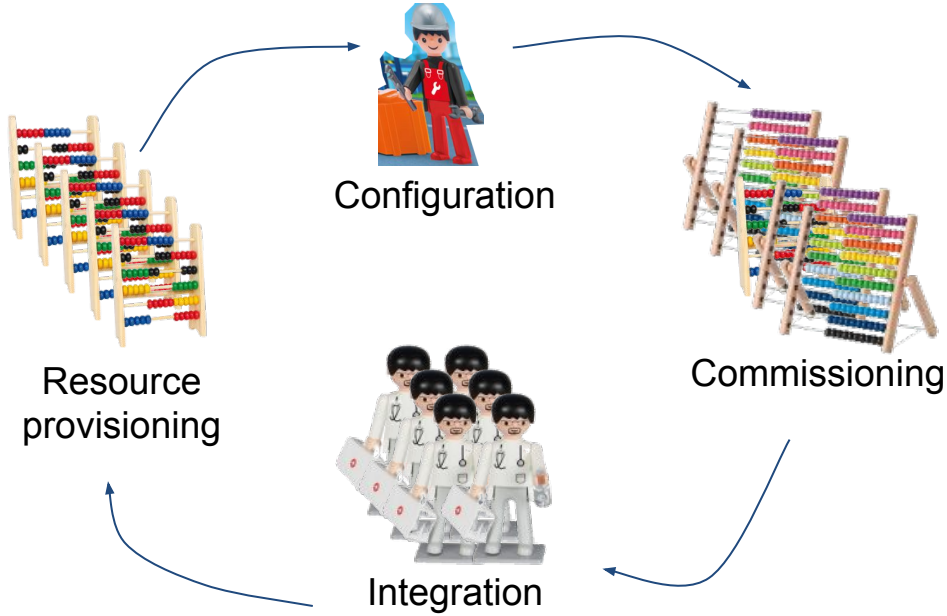

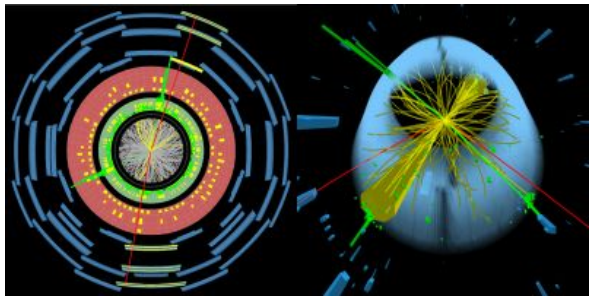

Data analysis

Common infrastructure issues

Complex distributed systems

… challenging to spot, debug & address issues

- Networking
  - DNS, firewall, bandwidth, ...
- Access to services essential to run a job
  - Squids & caches, frontier, …
- Experiments jobs suffering in various ways

“Pre-commission” the resources

How?

- we have **experience** with testing with **full-chain jobs**
  - fail early, in a controlled environment

- swap the submission backend from WMS to a batch system
  - but use the same batch resources & environment configuration

    ⇒ test resources and services at a site without the need for the full integration with the Experiment distributed computing systems

⇒ spot infrastructure issues early, and address them early

⇒ happy customers :)

Not only batch

“Pre-commission” the resources

… any resource or service essential to operate Compute activities

- Batch resources: in-house, external (cloud); test VM image readiness
- Containers: commission images and environment configuration
- ObjectStores: service functional testing
- Issuing load of any kind: front-end load testing, DB load testing, ...
- Commission components of complex distributed systems
  - a DDM client component commissioning

- Majority of Batch resources at CERN are available via HTCondor
  - Happy to collaborate to plug in other batch systems

Batch CI/CD

… continuous integration / continuous deployment of Batch VMs

Step 1. Create a Batch VM image with HashiCorp Packer

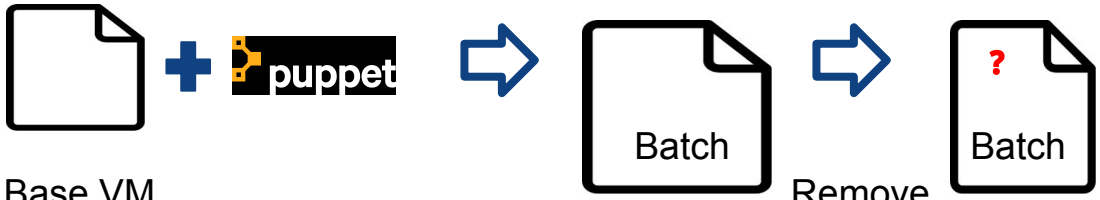

Step 2. Instantiate VM(s) on a test openstack **cluster**

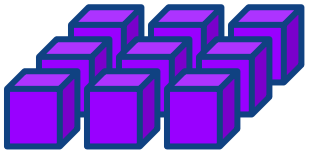

Step 3. Plug the cluster to our HTCondor High Throughput Computing **pool**

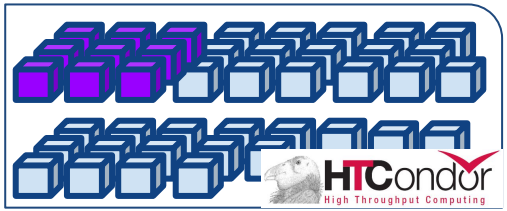

Step 4. Smoke-test the cluster with HammerCloud

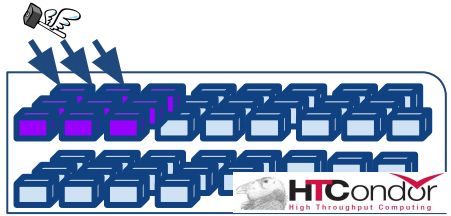

Step 5. Sign-off the new Batch VM image & deploy

Evolution of HammerCloud to commission CERN Compute resources

… to support even more Compute activities:

- “pre-commission” the resources,
- smoke-test builds & images,
- happy to collaborate!
  - Contact: <mailto:Jaroslava.Schovancova@cern.ch>

Jaroslava Schovancová, Alessandro Di Girolamo, Aristeidis Fkiaras, Valentina Mancinelli
(CERN)


CERN