CERN

# Implementing a common layer for accessing HPC

Gavin McCance, Tim Bell
CERN IT Compute and Monitoring Group

# HTC Batch System at CERN

- 200k+ core HTC batch service at CERN based on HTCondor supporting
  - LHC and related experiments via Grid
  - Other CERN experiments and departments via shell

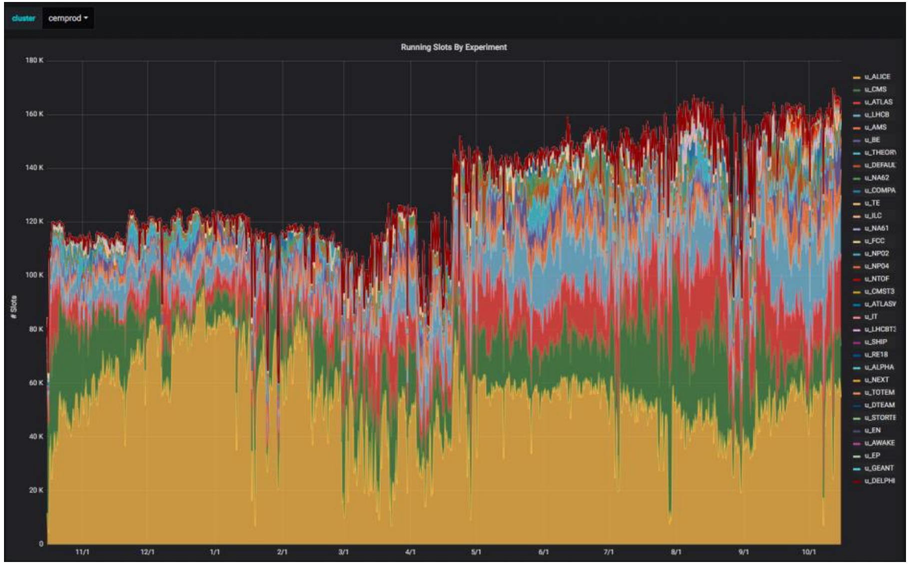

# Different Resource Types for LHC Workload

- Standard shared batch farm

- Multiple dedicated batch farms for specific activities (e.g. data recording)
  - Backfilled when not busy with primary activity

- Opportunistic resources
  - “Spare” OpenStack service capacity
  - Backfill of our SLURM resources
  - Unused CPUs of disk-servers (hyperconverged)
  - Volunteer computing via BOINC

- Externally hosted resources
  - Public cloud resources (e.g HNSciCloud and related projects)
  - PRACE centers?

# Provisioning Model

- Strategy is to hide complexity from users behind HTCondor

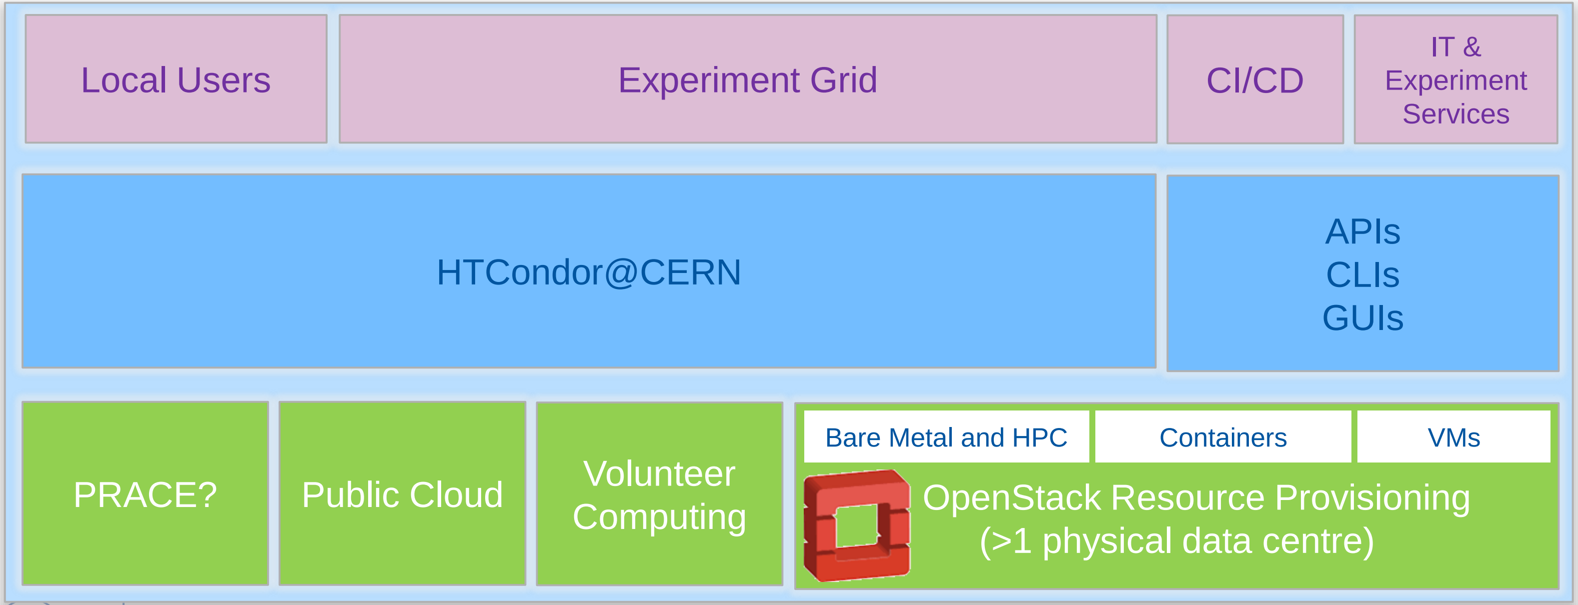


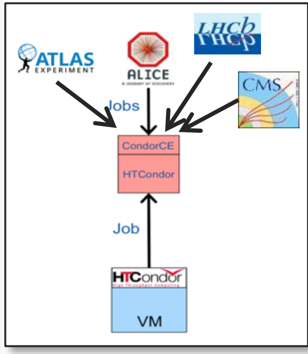

# Multiple Integration Patterns Already Used (I)

- External Cloud integration layer is based on Terraform
  - We then provision and configure machines as per our own site
  - Some recent prototypes using federated Kubernetes
  - In both cases, HTCondor daemon on box joins our cluster remotely and takes jobs

- Docker integration – start a worker node in a container
  - Docker'd-up HTCondor daemon joins our cluster remotely and takes jobs
  - This is used to integrate CERN disk-server resources (that a different team in IT run)

- Direct Job integration
  - Starts HTCondor daemon inside a (SLURM) job, joins our cluster remotely and takes one job
  - Only lasts for one job, then quits
  - We haven't tried this one, but is the basis for CMS' worldwide job system

# Multiple Integration Patterns Already Used (II)

- HTCondor GAHP integration
  - GAHP can re-submit HTCondor jobs to BOINC, SLURM and other clusters
  - Pluggable for different resource types

- This or Direct Job integration is probably the best for HPC sites, tbd

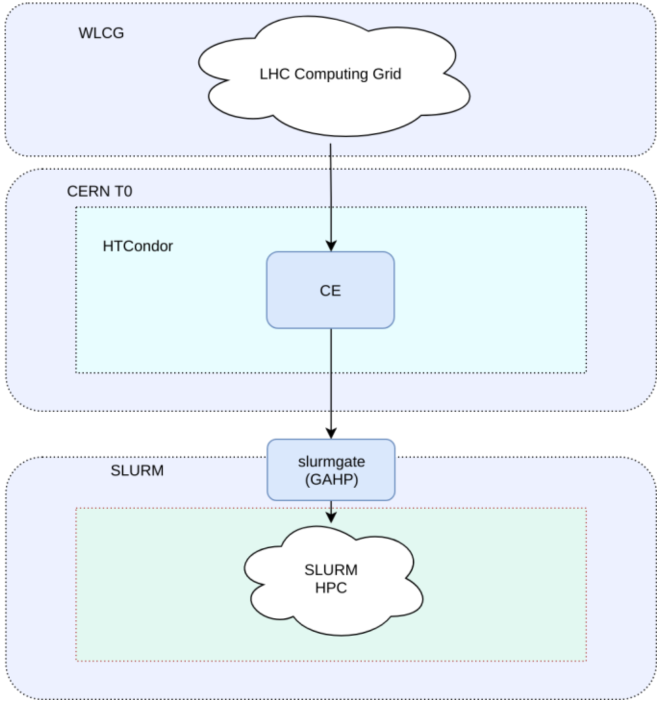

# Strawman

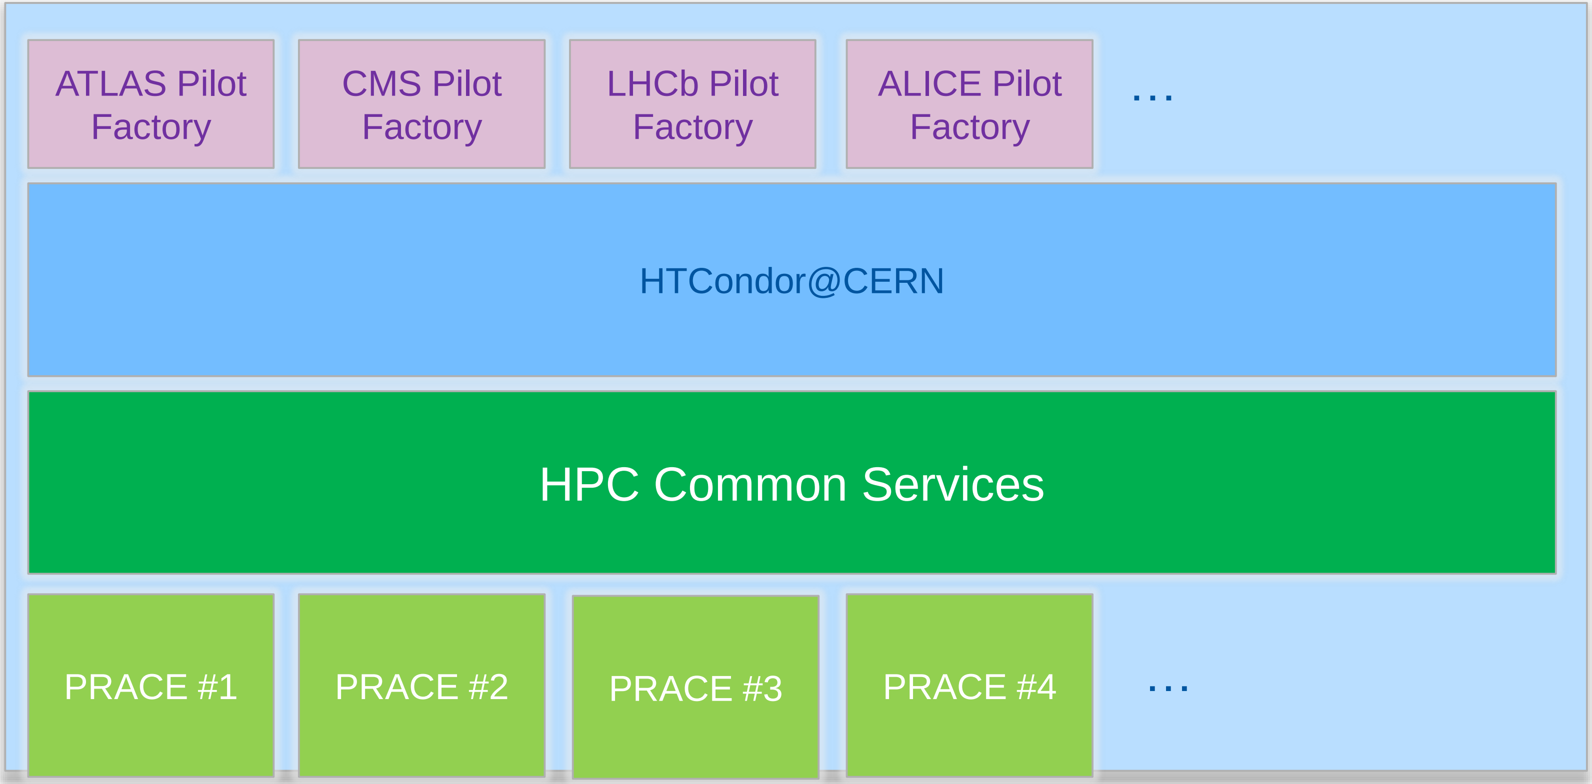
ATLAS Pilot Factory
CMS Pilot Factory
LHCb Pilot Factory
ALICE Pilot Factory
…
HTCondor@CERN
HPC Common Services
PRACE #1
PRACE #2
PRACE #3
PRACE #4
…

# Advantages

- The aim is to hide the integration complexities from users behind HTCondor
  - Single ops team for experiments and PRACE to interact with: avoids n x m PRACE site/experiment interfaces and processes
  - Reduced development effort: standard experiment pilot factories can be used
  - QoS: expose QoS differences via common properties (I/O requirements, pre-emptible?, etc) and route when resources are available
  - Special resources: specific resource properties (e.g. presence of GPUs) can be exposed to clients to match more suitable jobs, if desired
  - Backfill: we know how to operate this in backfill mode, pre-empting if a more important HPC user comes along
  - Accounting and monitoring: Existing infrastructure can be used for tracking usage

# Needs: CVMFS

- Hard to run LHC workload without CVMFS
  - Local site squid caches and per-node daemon
    - If it helps, we have containerised these
  - Non-CVMFS solutions exist .. but are expensive and typically experiment-specific
  - Workload itself can run inside Singularity, typically reading image from CVMFS

- CVMFS potentially interesting for HPC sites anyway
  - we’re looking at packaging our HPC apps with Singularity in CVMFS for efficient caching
  - Can we (or delegated owners) run edge services within the HPC sites? Examples would the Data Transfer Nodes at Argonne.

# Needs: Firewall

- Somewhat dependent on how we integrate with the site
  - e.g. HTCondorCE -> SLURM GAHP assumes we can see the site batch system over WAN to send jobs there
  - We could run HTCondorCE on the site directly (e.g. containerised) if that's better

- Outgoing connectivity from worker nodes
  - Call-us-back worker-node patterns need this
  - Most WLCG pilot jobs "call home" to get the real experiment payload

# Needs: Data Output

- We can steer specific experiment workloads via HTCondor to minimise input data needs
  - e.g. MonteCarlo-Digi-Reco
  - …but most jobs produce a fairly big output

- On external clouds, we send the output directly from the job to CERN and this works ~well
  - Assumes outgoing connectivity from workers at suitable rate but avoids the site having to run specialised storage

# Important: Input Data over WAN?

- Workload dependent (we can choose not run jobs that needs much input data)

- Work going on now inside WLCG and SKA on simple site caches to hide latency
  - Pull directly data from the data lake to avoid site having to explicitly manage data stores
  - Potential option, though typically would involve the site running some simple caching solution (containerised or otherwise)

# Conclusions

- Through CERN, we’re able run WLCG jobs on a variety of external infrastructure
- Our strategy is to hide the complexity from our users and expose everything via our HTCondor@CERN
  - And hide the complexity of HEP from the resource providers
- We have a variety of integration options already in use to achieve this
  - It would be good to agree a common solution across all the PRACE sites

# Backup

# CERN Data Centre by Numbers

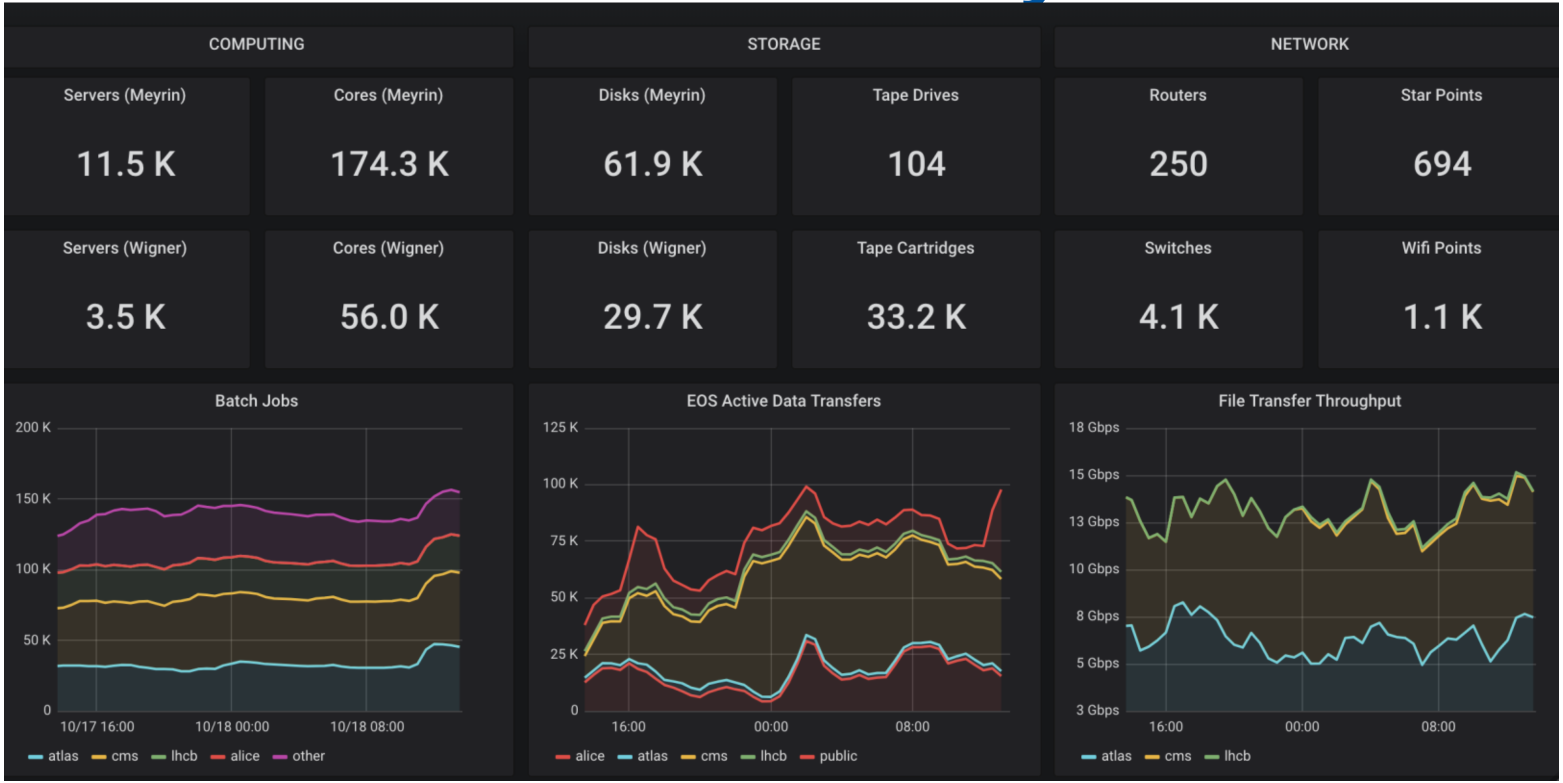


Source: http://go.web.cern.ch/go/datacentrebynumbers

# Accounting

- Since everything passes via CERN, we handle the accounting a give “credit” via WLCG report to sites that provide resources

- Mechanism already in use in HNSciCloud project where CERN runs the “WLCG share” composed of all the WLCG-cores purchased by other WLCG sites from the HNSciCloud vendors

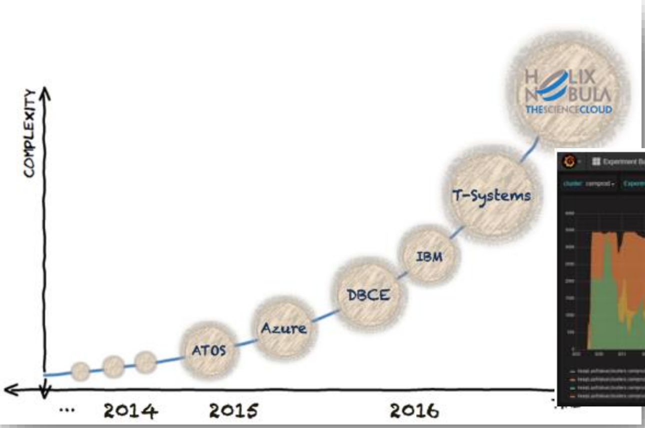



# Flat or elastic procurement of commercial cloud resources?

2nd Mar. — Atos
- *End: 31st of March 2015*
- ATLAS simulation jobs
- Single core VMs
- Up to 3k VMs for 45 days
- 1st Cloud Procurement

6th Nov. — DEUTSCHE BÖRSE CLOUD EXCHANGE
- *End: 18th of Dec. 2015*
- Target all VOs, simulation jobs
- 4-core VMs, O(1000) instances
- 2nd Cloud Procurement

1st Aug. — T-Systems
- *End: 30th of Nov. 2016*
- Provided by OTC IaaS
- 4-core VMs, O(1000) instances
- 500TB of central storage (DPM)
- 1k public IPs through GÉANT
- 3rd Cloud Procurement

*2015*

*2016*

23rd Mar. — Microsoft Azure
- Sponsored Account
- "evaluation of Azure as an IaaS"
- Any VO, any workload
- Targeting multiple DCs:
  - Iowa, Dublin and Amsterdam
- *End: 30th of Nov. 2015*

20th Nov. — SOFTLAYER an IBM Company
- Agreement between IBM and CERN
- CERN PoC to evaluate:
  - Resource provisioning
  - Network configurations
  - Compute performance
- Transparent extension of CERN's T0
- *End: 13th of May 2016*