
CERN

Containers for HTC

Luis Fernández Álvarez, Olga Datskova and Ben Jones

The Batch Service @ CERN IT

- Provides Tier-0 compute power via **HTCondor** to WLCG**.**
- Process CPU intensive workload ensuring fairshare among various user groups.
- Maximize utilisation, throughput and efficiency.
- It runs jobs from the Grid and from local CERN departments.

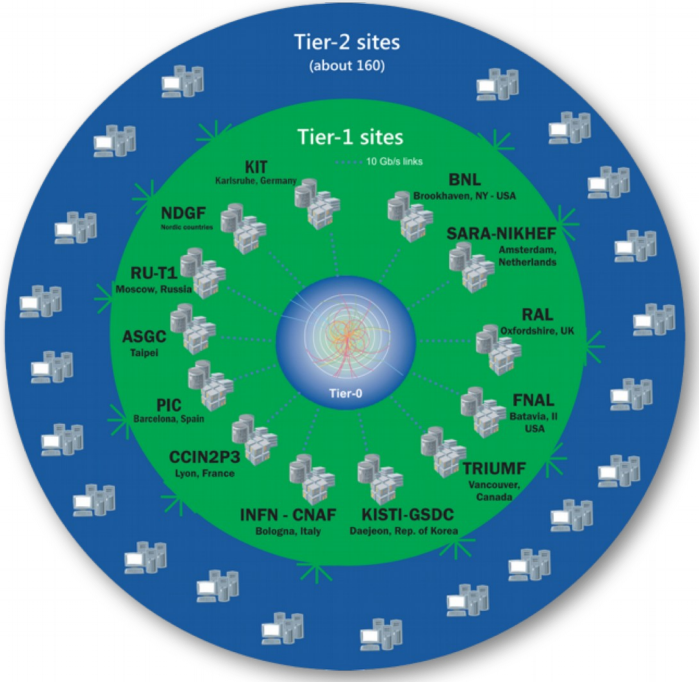

Computing Stack

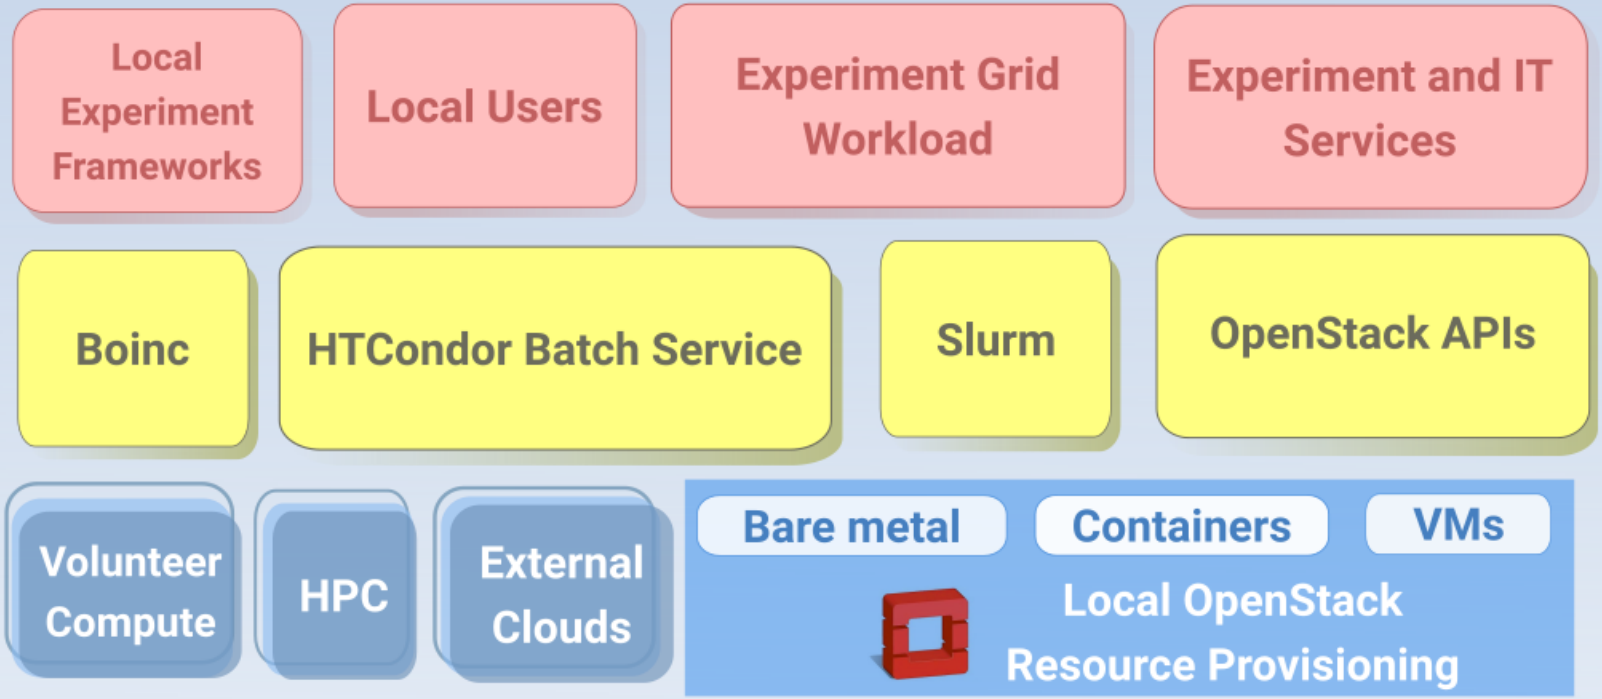

CERN Pools/Entry Points

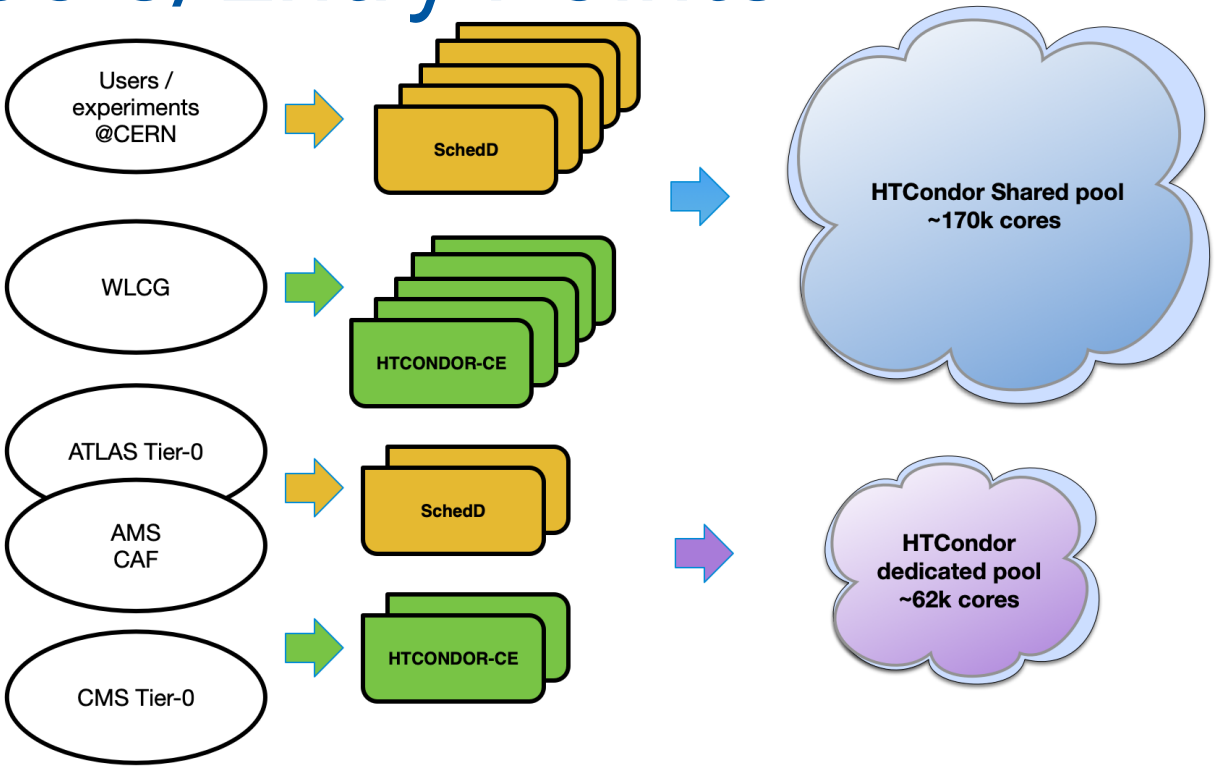

Problems to address

- Can we use spare compute capacity within CERN Computer Centre in the batch service?
- Can we get more out of the resources we already have?
- How do we best take advantage of public cloud resources?

How do containers help?

- Eliminate the need to share OS / libraries / Apps with Host machine
- Reduce the amount of things we need to configure on Host machines
- CGroups ensure that we can control resource usage of jobs

Condor + Containers on EOS = BEER

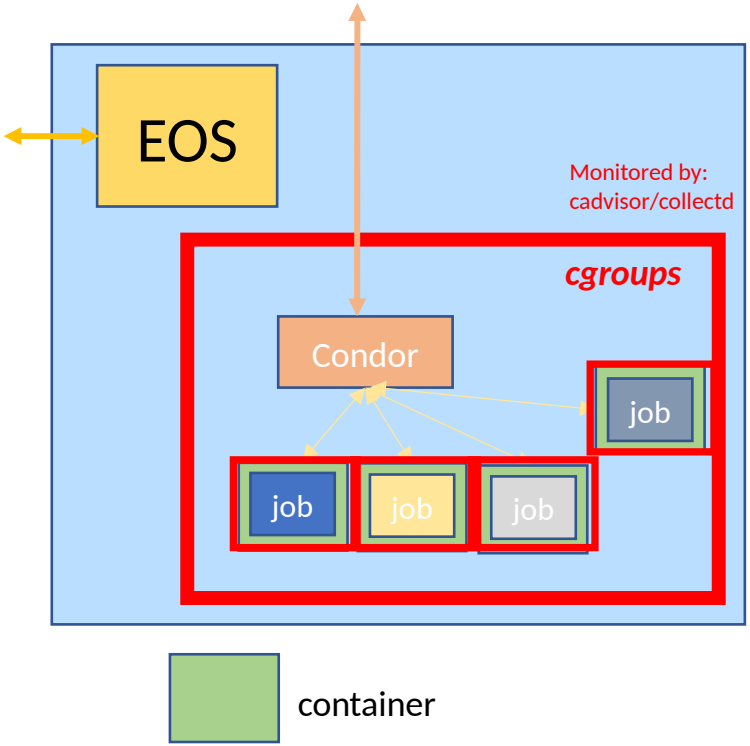

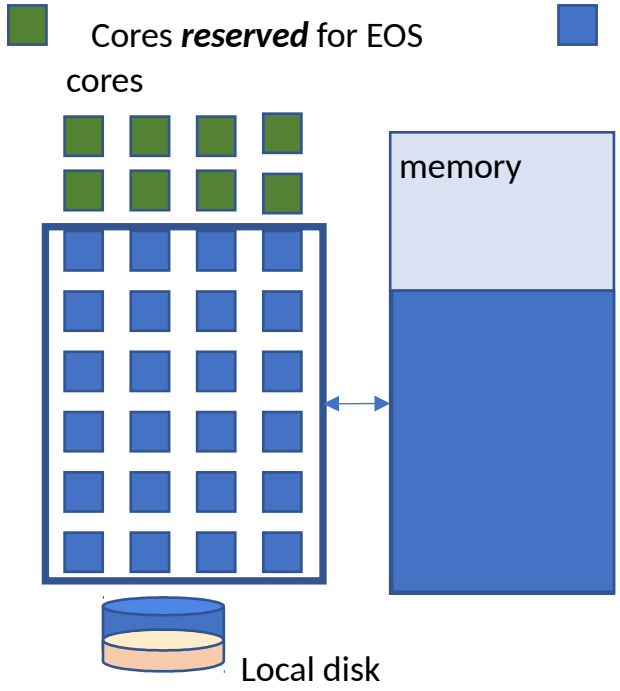

ATLAS with and without BEER

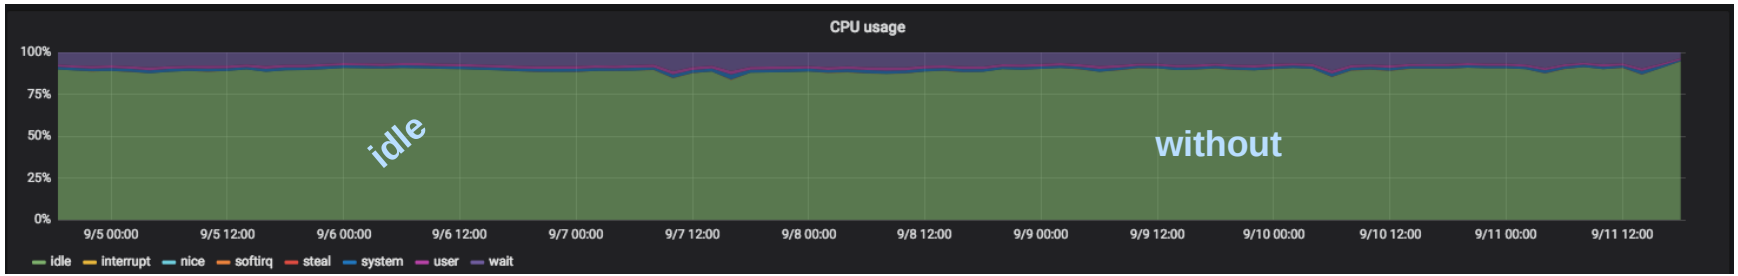

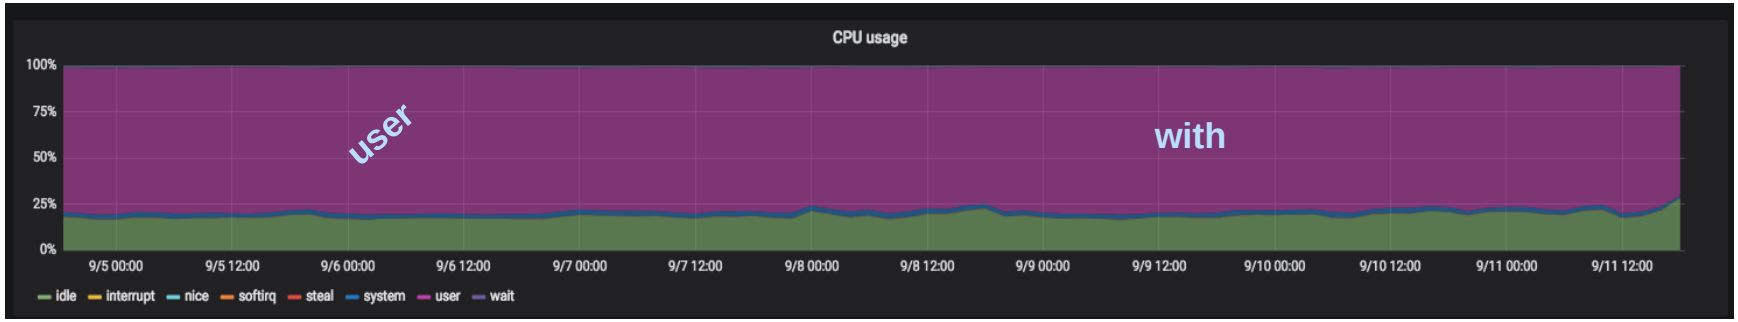

Singularity

- Business as usual.
- The **unpacked** service is being prepared to allow experiments to convert docker images into CVMFS-hosted unpacked directories suitable for use by Singularity.
- Investigating integration of Singularity CRI into Kubernetes environments.

Batch on Cloud resources

- Strategy:
  - Same provisioning & orchestration for public and local cloud.
- Dedicated to **Grid jobs** (opt-in):
  - Already designed to be location agnostic, with sophisticated job management & monitoring
- We generally have **flat capacity** & more jobs than resources
- The machines running the job live longer than the job
- Limited infrastructure required in cloud (proxies…)

Job Routes

- /etc/condor-ce/config.d/61-job-routes.config
- Defaults to set default datacentre, HEPSPEC or cores of undefined machines, Accounting Group, Max Runtime…
- Routes have helped partition public cloud whilst maintaining single point of submission

```
[
    TargetUniverse = 5;
    name = "External_Cloud";
    set_Requirements = (XBatch =?= True);
    set_WantExternalCloud = True;
    Requirements = (TARGET.WantExternalCloud =?= True) || (TARGET.queue =?=
"WantExternalCloud") || (TARGET.queue =?= "externalcloud");
  ]
```

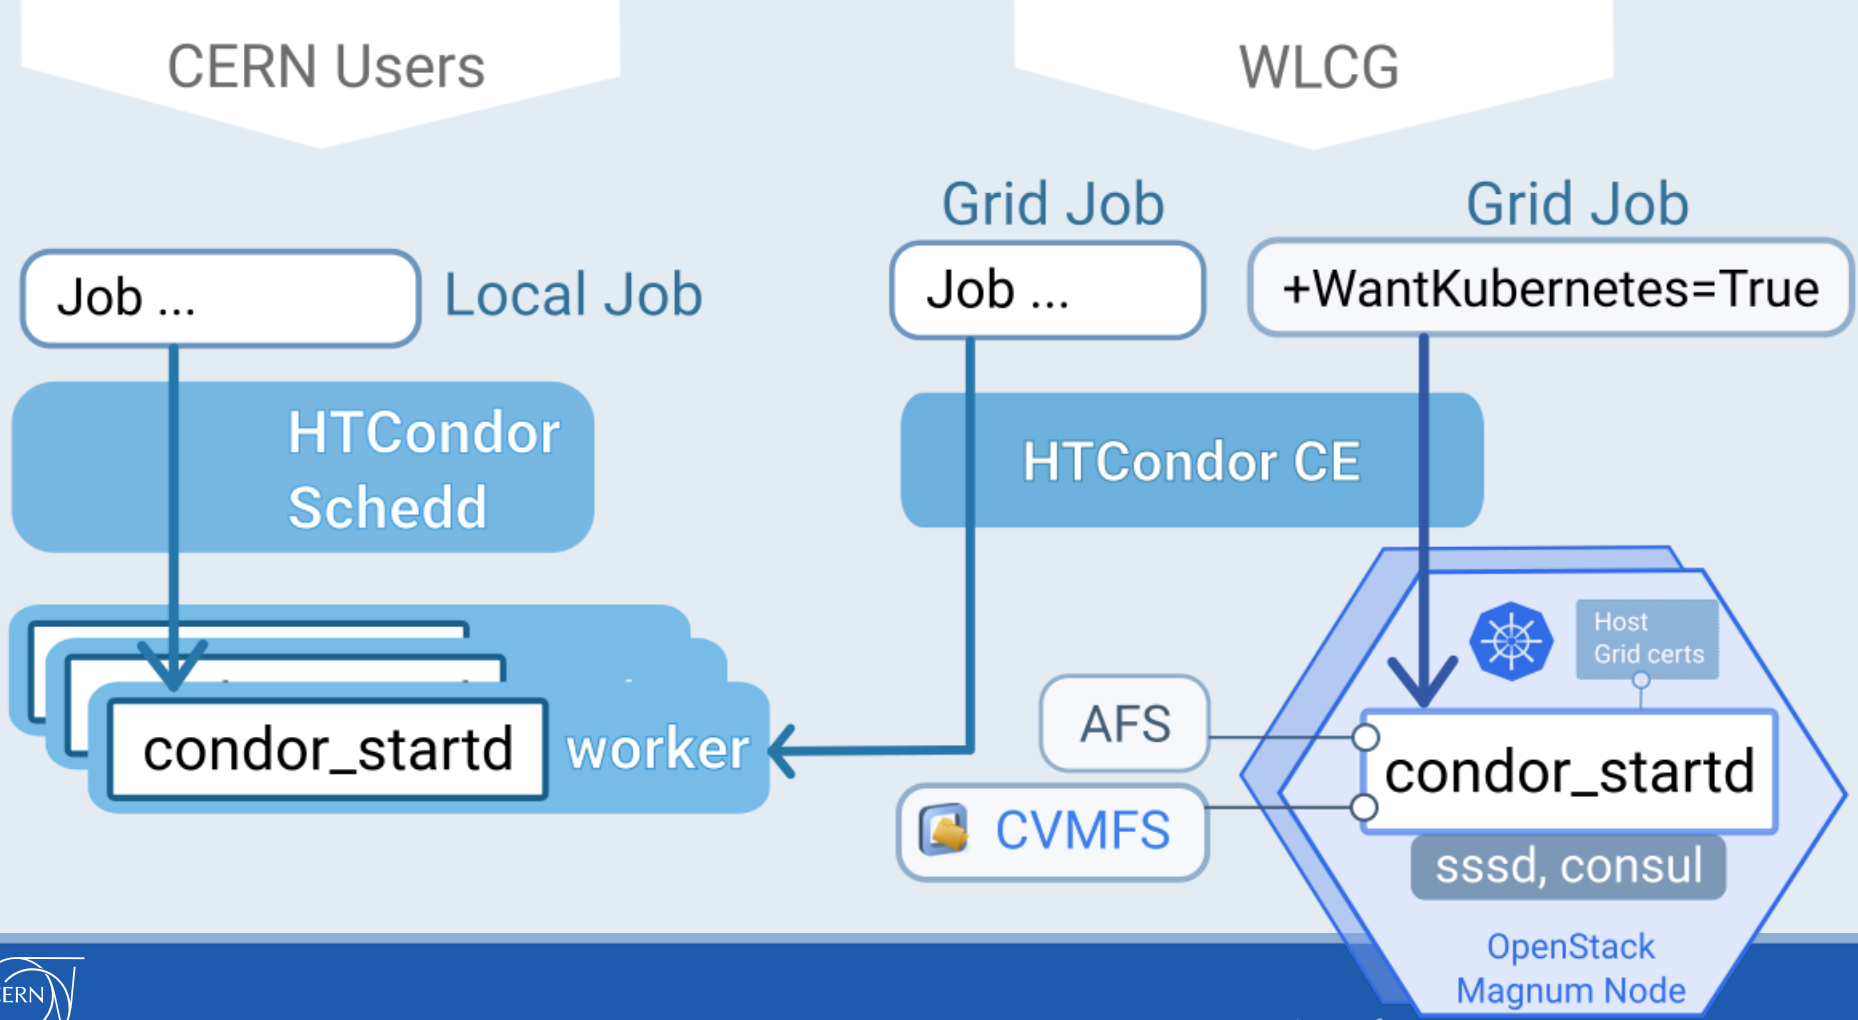
CERN Users
WLCG
Grid Job
Grid Job
Job ...
Local Job
Job ...
+WantKubernetes=True
HTCondor Schedd
HTCondor CE
condor_startd
worker
AFS
CVMFS
Host Grid certs
condor_startd
sssd, consul
OpenStack Magnum Node

Questions?