





Running a large multi-purpose HTCondor pool at CERN

Ben Jones



Not just the LHC...




AD & AMS



LHC vs Non-LHC



WLCG



- CERN as Tier-0 of the WLCG
- Hosts tier-0 functions of LHC experiments
- 700 000 cores
- 200+ sites
- ~30 GiB/s
- ~70 PB/year

Grid vs local



CERN's batch system...


- Isn't just for LHC experiments
- Isn't just Grid
- As we'll see in a presentation later today, isn't even the only HTCondor batch system at CERN
- As the IT batch system, it is responsible for providing batch computing to experiments, groups and users associated with CERN

Worker stats

Workers					
Cores		Total Memory		Hepspec	
cernprod	cernt0	cernprod	cernt0	cernprod	cernt0
147 K	65.2 K	672 TB	254 TB	1.531 Mil	667 K

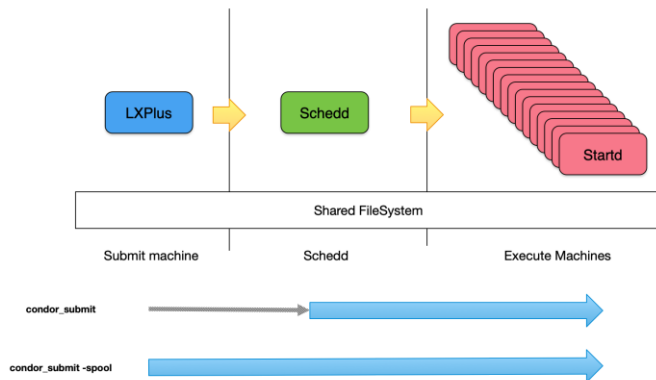
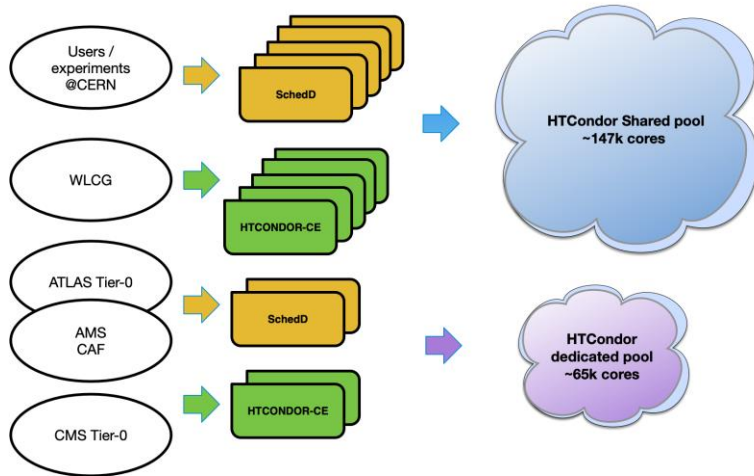
- CERN Batch system represents about 70% of the CERN Cloud
- All Batch machines are (still) SMT-on
- Majority of shared resources have 2 VMs per NUMA zone
 - VMs are scheduled by OpenStack to the physical sockets
 - (Yes, in prehistoric non-NUMA aware days we had VMs scheduled to SMT "cores")
 - 8 / 10 / 12 CPU VMs
- Older machines: just under 2Gb/core RAM
- Newer machines: just under 3Gb/core
- Significant minority of machines with non-standard configurations

Pools



- **cernprod / share**
 - CM: 2x 40 core
120Gb RAM
 - 42 sub collectors
 - 3 negotiators
 - 2 split grid, 1 for “special” resources
 - 1 CCB machine
 - 12 schedds for “local” submission
 - 15 CEs for grid
- **tzero / dedicated**
 - CM: 2x 16 core 28Gb RAM
 - 8 sub collectors
 - 1 negotiator
 - 2 schedds for “local” submission
 - 1 of which is dedicated to ATLAS T0
 - 3 CEs for grid

Local submission



- “Local” submission: remote submission to schedds, using Kerberos and passing tokens
 - Shared filesystem between remote submit machines (LXPLUS) and schedds
 - Why remote submission?
 - We don’t trust LXPLUS and nodes have short lifetime
 - We have privileged tokens on schedds
 - Kerberos token renewal via our ”ngauth” service, and HTCondor’s SEC_CREDENTIAL_* knobs
 - Mapping users to schedds
- LOCAL_CONFIG_FILE = /usr/bin/myschedd.sh |

Routing / transforming: ATLAS

StartD:

```
HostGroup = "bi/condor/tzero/atlas"
STARTD_ATTRS = $(STARTD_ATTRS) HostGroup
START = ((SendCredential =?= True) && (regexp("^group_u_ATLAS.TZERO.*", AccountingGroup)
|| regexp("^group_u_ATLAS.CAF.*", AccountingGroup)) || ((AtlasGridJob =?= True ||
EtfTestJob =?= True) && (SendCredential != True)))
```

CE
job_router:

```
[
    MaxJobs = 6000;
    MaxIdleJobs = 4000;
    TargetUniverse = 5;
    name = "AtlasT0";
    Requirements = (regexp("atlas", x509UserProxyVoName)) && (TARGET.queue =?= "AtlasT0");
    set_Requirements = (TARGET.Hostgroup =?= "bi/condor/tzero/atlas" && OpSysAndVer =?=
"CentOS7")
    set_AtlasGridJob = True;
    set_NiceUser = True;
]
```

Schedd
job transform:

```
JOB_TRANSFORM_To_Atlas_Tier0 @=end
NAME To_Atlas_Tier0
REQUIREMENTS ((regexp("group_u_ATLAS.TZERO.*", AccountingGroup) ||
regexp("group_u_ATLAS.CAF.*", AccountingGroup)) && jobUniverse =?= 5 )
COPY Requirements VanillaRequirements
SET Requirements (TARGET.Hostgroup =?= "bi/condor/tzero/atlas") && VanillaRequirements
@=end
```

Preemptible resources



- Preemptible resources both from CERN & Public Cloud
- No drain, no real warnings
- Tidy up resources in Condor afterwards

Draining

```
# cat
/usr/libexec/condor/scripts/staged
drain.sh
#!/bin/bash

FILE=/etc/shutdowntime

if [ -f $FILE ];
then
    UNIX_SHUTDOWN=`awk '{print
$0}' $FILE`
    echo "InStagedDrain = True"
    echo "ShutdownTime =
$UNIX_SHUTDOWN"
else
    echo "InStagedDrain = False"
    echo "ShutdownTime = 0"
fi
```

- We often like to drain machines for interventions.
- Jobs specify max walltime, then we accept those jobs that will finish before the intervention time.
- We kill jobs that exceed their advertised wall time
- Jobs are accepted via more START expr magic

```
((InStagedDrain =?= True &&
(time() + MaxRuntime <
ShutdownTime)) || InStagedDrain
=!= False)'
```

Healthchecks: CVMFS


- Previously: blacklisting startds if one cvmfs repo is broken.
 - Suboptimal: jobs might need only one of the healthy ones, or not require cvmfs at all.
- Need of more visibility on CVMFS led to the development of collectd-cvmfs.
 - Deployed and monitors most popular CVMFS repositories in our infrastructure.
- Extended startd crons to inject:
 - NODE_MOUNTS_CVMFS
 - CVMFS_HEALTHY_REPOS: string list with healthy repos.

CVMFS health (II)

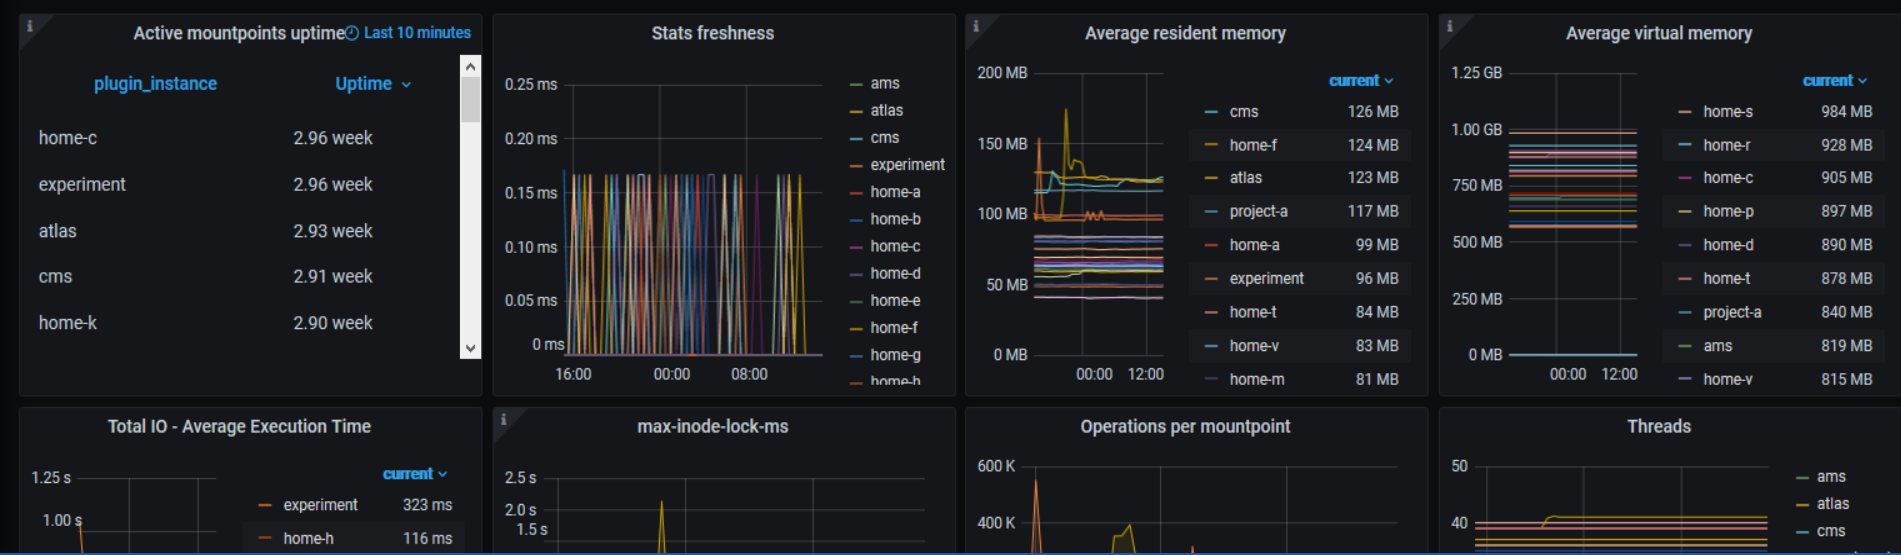
- New attributes enables us to...
 - Mapping of Grid VOs to predefined repos.
 - Allow advanced users to define their requirements.

```
+CvmfsCustomExpression = false
+CvmfsCustomExpression =
stringListMember("myrepo.cern.ch",
CVMFS_HEALTHY_REPOS)
```

Filesystems - CVMFS



Filesystems - EOS



Automated user notification

- Anticipate known issues and do early notifications to users.
 - Too many files on AFS, wrong paths, out of quota,...
- Deploy a lightweight pipeline from logs to mail notification
 - Grok_exporter, Prometheus & Alertmanager.

Automated user notification (II)

HTCondor Schedd

grok_exporter

```
metrics:  
- type: counter  
  name:  
  sched_writeuserlog_error_total  
  match:  
  WriteUserLog::initialize:  
  safe_open_wrapper\("%{AFS_AREA:  
  afs_area}%{USERNAME:username}%  
  {UNIXPATH:path}"\) failed -  
  errno %{INT:error_number}  
  paths:  
  - '/var/log/condor/SchedLog'  
  labels:  
    err_num: '{{.err_num}}'  
    username: '{{.username}}'
```

Prometheus

schedd scrape target

```
scheddxy.cern.ch:9144/metrics
```

alert rule

```
UserFullAFSFolder:  
  
delta(sched_writeuserlog_erro  
r_total{err_num ="27"}[5m]) >  
10
```

AlertManager

Notification

```
receivers:  
- name: 'email_users'  
  email_configs:  
  - to: '{{ template  
  "batch.user_notif.recipient" .  
  }}'
```

mail.tpl

```
<p>  
Dear {{.GroupLabels.username}}  
</p>
```

Never underestimate users...

- Our users often find ways to take full advantage of HTCondor's rich features in ways we did not anticipate
- Our submit requirements now becoming a historical record of various adverse reactions in the correct functioning of our plant

```
# Submit Requirement: NoJobDelay
```

```
SUBMIT_REQUIREMENT_NoJobDelay = isUndefined(NextJobStartDelay)
```


```
SUBMIT_REQUIREMENT_NoJobDelay_REASON = "Setting next_job_start_delay or  
+NextJobStartDelay in submission is not allowed:
```

```
https://batchdocs.web.cern.ch/troubleshooting/commonexceptions.html#job-delay"
```

```
# Submit Requirement: NoUsageAds
```

```
SUBMIT_REQUIREMENT_NoUsageAds = (isUndefined(RemoteWallClockTime) ||  
real(RemoteWallClockTime) =?= 0.0 ) && (isUndefined(CumulativeSlotTime) ||  
real(CumulativeSlotTime) =?= 0.0)
```

```
SUBMIT_REQUIREMENT_NoUsageAds_REASON = "The following attributes cannot be set:  
RemoteWallClockTime, CumulativeSlotTime"
```



Backup

HAGGIS Accounting Groups

Haggis 🧑

Pool ↑	Group	Quota	Role	Pledged	Charge As	Charge To
batch	group_u_ALICE	336935	n/a	✓	wall	alice
batch	group_u_CMS	68211	CMS	✓	wall	cms
batch	group_u_ATLAS	128153	ATLAS	✓	wall	atlas

Group ↑	Quota	Role	Pledged	Charge As	Charge To
group_u_ATLAS.exotics	6	ATLAS exotics	✓	wall	atlas
group_u_ATLAS.grid_ATLAS	6	ATLAS grid user	✓	wall	atlas
group_u_ATLAS.grid_ATLASPLT	34440	ATLAS grid pilot	✓	wall	atlas
group_u_ATLAS.grid_ATLASPRD	71700	ATLAS grid production	✓	wall	atlas
group_u_ATLAS.grid_ATLASSGM	200	ATLAS grid SGM	✓	wall	atlas

Rows per page: 5 ▾ 1-5 of 9



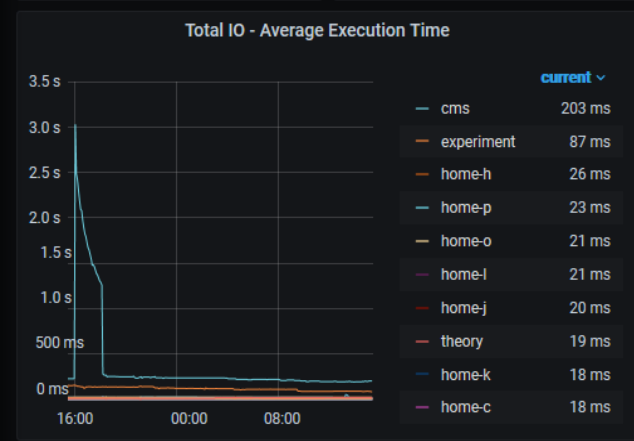

batch	group_u_LHCB	70221	LHCB	✓	wall	lhcb
-------	--------------	-------	------	---	------	------

Accounting Group transforms

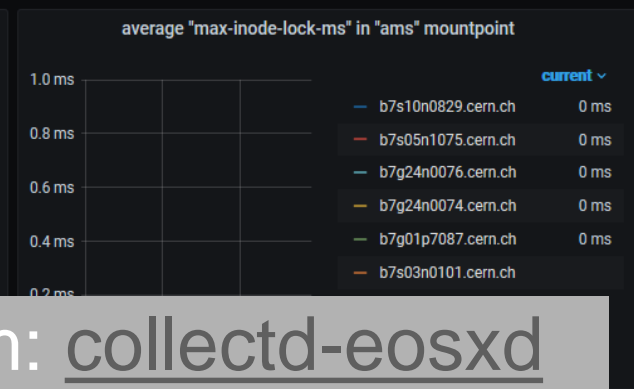

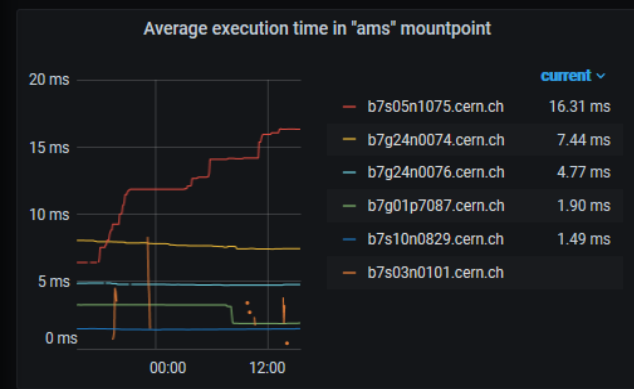
```
JOB_TRANSFORM_AccountingGroupCheck @=agc
  NAME AccountingGroupCheck
  EVALSET AccountingGroup (userMap("CERNGroups", Owner, AccountingGroup) != undefined) ?
strcat(userMap("CERNGroups", Owner, AccountingGroup), ".", Owner) : strcat("group_u_DEFAULT.all", ".", Owner)
  EVALSET MaxRuntime (userMap("JobFlavours", JobFlavour, undefined) != undefined) ?
int((userMap("JobFlavours", JobFlavour, undefined))) : (MaxRuntime =?= undefined) ? 1200 : MaxRuntime
@agc
```

- Accounting Groups map defined by HAGGIS dump
- userMap to assign to AccountingGroup, and dump to low-pri “DEFAULT” if there’s no assignment
- Also set default (low!) maxRuntime (we really should’ve called it maxWalltime)

Global metrics



Mountpoint details: ams



From: [collectd-eosxd](#)

Resident memory in "ams" mountpoint

Virtual memory in "ams" mountpoint