

Mid-SLA jobs and accounting

Gavin McCance [CERN]
2 March 2017
WLCG Ops



Extra capacity

- Significant potential extra capacity, if we are able to relax the normal HEP SLA a bit
- Examples:
 - disk-server CPU
 - spare service “headroom”, choppy cloud compute capacity, external cloud spot
 - HPC backfill, pre-empt by prompt work (T0/CAF),
 - intervention draining

Current levels

- **Standard production:** usual kit, usual networking, access to storage, (careful) draining
- **Pure opportunistic:** BOINC, random volunteer kit, data bridges, no direct access to storage. Jobs potentially shot with no notice

Current levels

- **Standard production:** usual kit, usual networking, access to storage, (careful) draining
- **Mid-SLA:** usual kit, usual networking, access to storage. Potential resource squeezed (e.g. local IOPS). Jobs potentially shot with no notice.
- **Pure opportunistic:** BOINC, random volunteer kit, data bridges, no direct access to storage. Jobs potentially shot with no notice

Mid-SLA notice

- Quite a bit more capacity could be made available if we abandon “won’t kill” or even abandon “give 30 minutes notice”
- Fine to use someone else’s stuff if you give it back immediately
 1. Data servers will kill the job the very moment they think they’re in trouble (data server shoots container running job)
 2. Cloud ramp up (Openstack shoots VM running job). Spot resources.
 3. Guaranteed prompt capacity that we’re “borrowing” (T0 / HPC). Potential batch system pre-empt.

Mid-SLA notice

- Batch system pre-emption likely doable for case 3 (if useful?)
 - Would use pre-empt kill (rather than SIGSTOP)
 - Many batchsys (LSF, HTCondor @CERN) can send a signal to job on kill pre-emption. Not always tuneable (it is in HTCondor)
 - You have 10-30 seconds before SIGKILL hits you
 - compatible with the guarantee we give the real owner
 - probably enough to notify book-keeping
 - not enough to save the data
- Hard to reliably provide any notice for case 1 and 2 without hackery

Cases where we know

- Some cases we do (or can) know the end time
 - Burn-in, draining for capacity shuffles, end of life
- In reality, we often just run standard production on this - we just don't accept that last 12-hour job

We acknowledge MJF

- /etc/machinefeatures/shutdowntime update 30 mins notice
- We think most of the potential extra capacity CERN could provide doesn't fit with this
 - It's seconds notice, if any
 - Or it's 24 hours and we run standard production on it anyway

Accounting

- Would propose explicit above-pledge accounting line in APEL for “mid-SLA”

Mid-SLA

- Would propose explicit above-pledge accounting line in APEL for “mid-SLA”
- Tied to SLA:
 - Standard kit, standard access to storage
 - Potential resource squeeze (notably IOPs)
 - No notice termination
 - best effort, not guaranteed SIGINT/TERM/XCPU (tbd) pre-empt with seconds, to notify book-keeping (if useful?)

Mid-SLA plans

- CERN will be making capacity available with this SLA (via distinct CE / route) this year
- Other WLCG sites likely have similar capacity that would suit this SLA (e.g. disk-servers?)
- Open questions:
 - how to let experiments know how much we have (in case of variable)
 - Best effort pre-empt useful for experiments? (vs. nothing)
 - Can we agree a set of signals for this?