Autogenerating Computing Resource Records for Vac and Vcycle sites

Andrew McNab

LHCb, GridPP, DUNE

Vac and Vcycle

- Lightweight systems for managing VMs that run jobs for experiments
  - either on bare metal (Vac) or OpenStack etc (Vcycle)
- Aim is to make is simple for sites to operate, without running lots of services
  - Currently just need Vac and Vcycle
  - No BDII
  - APEL done with the ssmsend agent rather than a local APEL server
- We provide a monitoring service, VacMon, to sites which provides Ganglia like functionality
  - Avoids need to run something like that locally

VacMon

- Vac or Vcycle send JSON messages to a central vacmond hosted by Manchester for GridPP
- These include
  - per hypervisor workload, broken down by type of VM and experiment
  - per hypervisor capacity, in terms of CPUs, HS06, memory
  - VM definition requirements eg max lifetime
- These are stored in an ElasticSearch database
  - For site admins, there's a web dashboard to view graphs
- So we have enough information to construct CRR files for each Vac or Vcycle "space" (~cluster, ~CE)

VacJson

- https://vacjson.gridpp.ac.uk has a prototype creating the CRRs dynamically for all Vac and Vcycle spaces reporting to VacMon
- These are taken live from the ElasticSearch database
  - Actually, they're always written to a cached copy which is still used if the ES database cannot be contacted
- Don't worry about the format of the CRR JSON
  - It will be updated to track the current draft of this TF
- Pros: site admins don't need to maintain CRR. It always reflects the truth on the ground. The CRR format can be updated centrally.
- Cons: it doesn't represent machines which are temporarily offline (the "installed capacity" issue)

Points to note

- This model assumes that JSON files for a single site can be collected from multiple sources
  - ie per CE or per batchsystem JSON files are allowed
  - A site with ARC and Vac will need this, for example
- CRR is just JSON
  - So sites could fetch the file from vacjson
  - Cache it, edit it, process it somehow?
- There are ways we could add “installed capacity” numbers including machines offline
- VacMon messages can be sent to multiple locations so a site could run it’s own vacjson instance