

Lightweight sites

- Overview. T2 effort target areas.
- Summary of GDB discussion on approaches and ideas.

Storage

- Churn. Low reuse. T2 storage is cache-like.
- Easier to manage without reducing capacity.
- **What do we do with storage that will reduce?** Slow degrade or physical co-location?
Reduce manpower/maintenance/reliability requirements/features -> Networking – Caching – Barebones – Egalitarian.
- Eevee – Pokemon applies to everything.
- Olympics.
- Data CVMFS (caching/read only) **How many does this impact?** RAL experience?
- T2C caching.
- ARC caching. No visible Rucio integration testing.
- **Is caching efficient? These are open questions.**
- **Pure network model – Alastair** testing. Hammercloud – work goes to data.
- **Is it workable to run sites without storage?** Bristol. 2.5MB/s. US CMS T3s. Not storage. Network capacity at serving sites. UCL-QMUL.
- Network close option.
- Understand efficiencies of different style of jobs.
- Simpler systems – ZFS software ‘RAID’.
- SE “retirement”. Transition is hard. DPM multiple pools. Workload high. SRM retirement – spacetokens still an issue.
- Single protocol – xroot/xrootd. Supports http(s). SE-> xrootd.
- How will you know what storage is available? Xroot redirector.
Worldwide view of storage – multiple levels. Not in standard redirector. Consistent namespace coming.
- Interoperability. **DIRAC and idea of global redirector?** Xroot supported. If not making where data is visible then DIRAC cannot place jobs.
- Experiment view on things – **use this end point to view storage?** Needs thought.
- Funding signals? **How will cache/non-reliable storage count?** Can we count it?
- Policy. GridPP does what the experiments want. Discussion between tech suppliers and experiments. We then implement policy. Well run sites... reliable with good network access. Better access. ADC want fewer sites. 2000 sites tension. Centrally there is limited resource. Push things down.
- Xrootd federation pilot. Setup to support many VOs. Single name space – multiple VOs. Authn.

Vac/Vcycle.

- ramping up

- Multiprocessor VM support using “superslots”
- Needs “machinegroups”. VM definitions.
- How do you impose workloads? Experiment defines VM... at moment gives lifetime. Will also need to specify constraints. Machines jobs features.
- **CMS Glidein algorithms rubbish**
- Volunteers through WG may see more interest outside the UK. Talked with ML about ALICE approach.

GridPP Technical WG

- Clouds directly. No easier than traditional grid. IC seeing this now with manpower changes.
- Work within experiments to **access other generic site resources**. Uses microCERN VM.
- Cloud WG. RCUK. Seed academic community. Going for a year now. Led by Kershaw at RAL. Model is one where people **perform science on public and private clouds** trial. Working with 3 providers. A Lahiff containers with increasingly complicated workflows – AWS; Azure and Google.
- Too much effort to run clouds for just WLCG. Will benefit if University running a cloud service then use Vcycle to access.
- Evolve to using **cloud if someone can run it for you..** predictable requirements mean Grid is not bad. HLT needs to turn quickly for various use-cases.
- New site – OpenStack vs Grid without support?
- OpenNebula easier than OpenStack. -> Relgion!

Service Evolution Ideas

- Can some larger sites run services for smaller sites? Exploring ideas not making a decision.
- Distributed site idea is not new
- Not everyone wants to run jobs in VMs.
- There are **several alternative ways to run services across sites**
- Some example services to pull out: PhEEx; Site BDII; CE; Monitoring & logging
- Experience running remotely (Bristol)... policy issues.
- **How does this fit with our funding model?** Issues spanning sites hard to resolve. Security concerns.
- Sussex – no HEP sysadmin. Model 2 better for Sussex. Quite a lot of opportunistic resource.
- Concern grid services are not built to be distributed.