



Lightweight sites: computing resources

Ops Coordination meeting
29th Sep 2016

Alessandra Forti, Maarten Litmaath

v1.0

Introduction (1)

- The May GDB afternoon had a session on lightweight sites
 - A summary is available here
- The introduction defined boundary conditions
- In particular, here we are not (directly) concerned with storage and data access
 - Those aspects fall under the Data Management Coordination initiative approved in the Sep MB
- Instead, we focus on what it takes to provide computing resources to the experiments

Introduction (2)

- In OSG every WLCG site mainly supports just a single LHC experiment
- The sites are managed in close collaboration with the US project in each experiment
 - US-ATLAS, US-CMS, US-ALICE
- Both US-ATLAS and US-CMS are already working on lighter ways to provision their resources
 - Ubiquitous Cyberinfrastructure, Virtual Clusters
 - Tier-3 in a box, Pacific Research Platform
- In EGI the situation is a lot more complex
 - Multi-experiment sites, many countries/cultures/projects/..., more MW diversity, experiments have less influence, ...
- Here we should focus on the EGI sites then
 - While learning from the OSG sites

Next

- From the GDB session summary we can get ideas on where to put effort
 - Mind that nothing comes free!
- To identify what might help sites the most, we are preparing a questionnaire
- Any question may have some context explaining the implications of a particular technology choice or operations model
 - We cannot merely shift the work from T2 sites to the experiments, CERN-IT or T1 sites
 - Instead we seek an overall reduction of complexity and operational costs, benefiting the sites in particular
- The next pages show a first draft of the questions
 - Feedback welcome, thanks!

Draft questionnaire (1)

1. **What is the name of your site?**
2. **Does your site support a single LHC experiment?**
If only one experiment is supported, experiment-specific solutions may be available instead of having to run a generic classic grid site.
3. **Must your site support other VOs on the same resources?**
Other VOs may require classic grid services, at least for the time being. The LHC experiments can also use cloud sites, as detailed further down.
4. **Does your site require classic grid services?**
There are areas to consider for simplification, as detailed next.

Draft questionnaire (2)

5. Consider reduction of the required services or phase space?

For example, the APEL service is required when CREAM is used as the CE, but not for the ARC CE, which has the functionality built-in. HTCondor can be used both as batch system and as CE. In either case the site needs to consider its other supported communities, which may need time to adjust to a new batch system or CE type.

6. Might your site benefit from a repository of OpenStack images?

Instead of having to install a grid service from scratch, a suitable image could be the starting point, already containing all the packages required for the given service and possibly even some minimal configuration.

Draft questionnaire (3)

7. Might your site benefit from a repository of Docker containers?

Containers are lighter than VMs. Also here a suitable starting image could reduce the work needed to set up a grid service.

8. Might your site benefit from a repository of Puppet modules?

In setting up a grid service, starting from a VM or container image would just be the first step: further configuration will still be needed. These days Puppet is the most popular tool for that. Some MW products already come with service-specific modules.

Draft questionnaire (4)

9. **Might your site benefit from install/configuration wizards?**
As they are well-known for OS installations and stand-alone applications, also in system administration, it is conceivable for grid services to come equipped with wizards. However, mind that the creation and support of locally customizable wizards would be a big enterprise.

10. **Could your site supply WNs dedicated to the experiment(s)?**
If the WN resources used by the experiment(s) do not have to be shared with other VOs, they could be made to join a central HTCondor pool at CERN and directly receive jobs from there. In this scenario there is no need for a local batch system, nor a CE or other auxiliary grid service, except for a local CVMFS Squid. The WNs just pop up under control of the site. They can be in the form of HW hosts, VMs (e.g. via Vac or Vcycle) or containers. Proper accounting and experiment shares per site would be implemented in the central pool.

Draft questionnaire (5)

11. **Allow remote access to your local cloud infrastructure?**
A central service would be able to instantiate VMs directly. However, this model has operational scalability drawbacks compared to the previous option. Central experts would need to engage with site admins to resolve firewall issues, cloud flavor or version incompatibilities, hypervisor overloads, and more. As always, site admins still need to monitor the state of their infrastructure and address performance issues.

12. **Allow remote access to a DMZ for the experiment(s)?**
A central team would be able to access a set of hosts on which they would install the few grid services expected at a site, in particular a CE or other interface to an existing site-managed batch system. Also this option has a scalability concern. The small team would need to engage with site admins to understand local configuration matters and operational issues between the DMZ and the site resources, in particular with the network(s).

Draft questionnaire (6)

13. Do you have comments, questions or suggestions?