

WLCG Batch and CE

NeIC NT1 Manager
Mattias Wadenstein
<maswan@ndgf.org>

2017-07-12
GDB

Overview

- **Note: This is September, time to make some decisions**
- Why?
- Scope
- Support statement
- Recommended batch systems
- Recommended CEs
- Stakeholder statements
- Update procedure



Why?

- Maarten's Workload Management Trends:
 - <https://indico.cern.ch/event/578984/contributions/2509205/attachments/1424665/2186386/WM-landscape-v10.pdf>
 - This work indicates that we have a (slow) shift out at sites
- Unfavorable support situation for some software like PBS and CREAM-CE
- Unnecessary diversity incurs maintenance cost, mostly on experiment side
- Some guidance might make the process easier for our sites that are looking to make a change



Scope

- Recommendations are aimed for a) new sites and b) sites want to change something
- Not intended to tell sites that they must change something
- Well-known that other communities served by the same infrastructure might have other needs (like the ones depending on “correct” Glue1 infosys data)
- To be published in WLCG web/wiki space
 - <https://twiki.cern.ch/twiki/bin/view/LCG/BatchSystemComparison>



Supported recommendations

- How strong support statement should these come with?
- The recommended solutions are supported by the software teams behind the solutions
 - WLCG can help with reproducing issues, providing feedback and patches in bug reports, testing solutions, pre-production deployment, etc.
- WLCG community takes some responsibility
 - Deployment of new features should be documented for all recommended solutions, same for installation recipes etc
 - Coordinating feature requests and developer feedback



Recommended batch systems

- HTCondor for sites with (mostly) HTC load
 - Especially at large scale in terms of number of jobs
 - Where each job is single core or a few cores up to a node
 - This is your typical WLCG site
- SLURM for sites with significant HPC load
 - Especially running multi-node MPI jobs
 - Ex: Nordic sites where pledges are on shares of HPC systems
- Local site preference of course matters
 - If site admins feel familiar with a batch system, or have local support nearby, it's probably a better choice



Recommended CEs

- HTCondor-CE
 - When connected to HTCondor batch system
 - Together with SLURM?
- ARC-CE
 - When connected to SLURM
 - Also works with HTCondor
 - Lightweight/HPC sites with data staging/cache instead of local SE
 - Some experiments say that ARC-CE gives more value to them over HTCondor-CE
- Any preference to state here?



Stakeholder input

- A couple of examples of statements:
 - ATLAS say they can work with both CEs, they prefer ARC-CE because it gives them more value
 - OSG are likely to be more of help supporting HTCondor-CE than ARC
- Statements like these are important help for making optimal choices as a site
- Rationale is important, so that they can be properly evaluated as situations change
- Suggestion: Publish these along with the rationale for the choices of software.



Update procedures

- Major changes
 - Like adding or dropping a CE/batch system to the list
 - Should probably go through a GDB meeting
- Minor changes
 - Updating stakeholder statements (“ATLAS now likes HTCondor-CE and ARC-CE equally”) or reasoning (“MPI job support now excellent in HTCondor”)
 - Just go through editors on request from relevant party?
 - Editor election?





Discussion?