

ARC Brainstorming Camp Summary

Andrew Washbrook

University of Edinburgh

GridPP UK Operations Meeting

13th December 2016

ARC Brainstorming Camp

- Meeting of the **ARC4eInfrastructures** project
- Workshop held in Monor, Hungary (28 November - 2 December)
- 18 participants
 - Mostly Nordugrid members
 - 5 participants affiliated to ATLAS

Brainstorming

- Long term development ideas (even crazy ones)
- Sources of funding and potential collaborations
- Long term feedback from users

R&D

- **Site and users feedback**
 - Breaking down brainstorming ideas to specific R&D topics
 - ARC hardening
 - ARC maintenance
 - Next major release planning
-

Feedback Call

Setup and Configuration

- What release are you currently running?
- What was your experience with initially setting up and configuring an ARC CE at your site?
Were there any factors that made set up difficult?
- Was the documentation sufficient?
- How did you define the ARC configuration file (arc.conf) for your site? Did you use the provided examples as a starting point or did you modify a copy from another location?
- Do you make use of the cache and web service?

Troubleshooting

- Is the service logging information sufficient enough for you to troubleshoot common issues?
- Is there any way that issue resolution could be made easier?
- What resources have you found most useful when encountering problems with the service?
- What ARC-specific bugs affect your site operations that have still not been addressed in later releases?

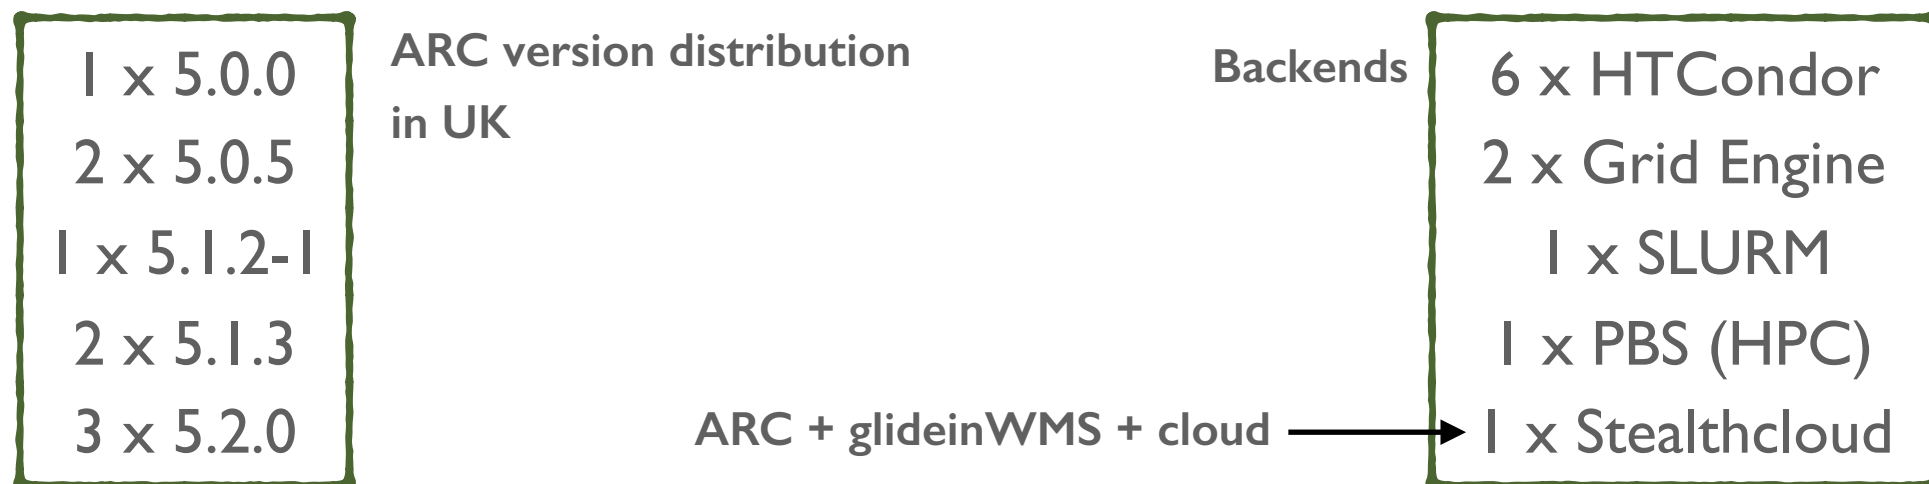
Development

- Have you added any site-specific changes to the ARC software to cover missing functionality or to override default behaviour?
- What features are missing that you would like to see considered?

A big thank you to everyone who took the time to participate!

Use of ARC in the UK

- Is there an incentive for sites to move to the latest stable version?
- Formal support expiry on earlier versions?
- Fast track for critical (security) bugs on previous releases?
- Formal support and development statement for less well used backends? For emerging backends?



- Sites should update to the latest version for important bug fixes rather than expect fixes to earlier versions (but appreciate this is not trivial)
- Discussion on which distribution channel (EPEL, UMD, Nordugrid) sites tend to use
- Change the "organic" release cycle?

Setup and Configuration (I)

“*Setting up was fine for the most part using the System Admin document Nordugrid produce, however I got totally stuck with the runtime environment scripts as it wasn't clear what they were for.*

Setting this up was relatively easy bar some oddities with access [and] runtime environment scripts.

The documentation is OK, but without an example it would probably be next to impossible for someone new to ARC.

Setting up the configuration for the information system was not easy.

We had to document the specific use case (i.e. with htcondor) ourselves to include the integration tips from many sources

A lot of integration was needed by [other sites] to make it operate with htcondor.

Made use of the provided examples for arc.conf and also bugged [other sites] many times

I'm installing it using puppet modules, but if I wasn't using these I'd certainly prefer to start from an arc.conf example rather than creating it from scratch.

”

Setup and Configuration (2)

- Configuration step was recognised as troublesome
 - Site customisation is even more troublesome especially for integration with other services (e.g. authorisation, job accounting)
 - Discussion on provenance of configuration templates
 - Sites pass around magic configuration snippets (e.g. ARGUS)
 - Official templates vs working configuration from other sites
 - Suggestion to gather copies of working arc.conf to improve provided templates
 - Proposed improving current configuration scheme as part of a new major release
 - In future distribute RTEs (like ENV/PROXY) out of box
 - Collect site configuration templates and publish them on the site and in the distribution before the next minor release
-

Troubleshooting

“ I find debugging anything going wrong quite hard on ARC and there seem to be a disproportional number of problems which go away when the machine is restarted.

Logging is fine, we haven't come across any issues we weren't able to diagnose other than the information system instability.

Logging for the most part for job submission and the gridftp interface is fine but we've had a few odd errors related to proxy delegation where we had to add our own debugging to the code to get the information we need.

Authentication errors are not easy to debug: not much info is provided in gridftpd.log. Tools like /usr/libexec/arc/arc-lcmads sometimes provide much more useful information, but it can be difficult to get access to a proxy in order to do this.

Sometimes a-rex or grid-manager locks up. We have to detect when the gm-heartbeat file is stale, then restart by hand.

There seems to be an issue with starting services, almost certainly another race condition. ”

- Starting services issue (and startup ordering) also reported by non-UK sites, will be investigated
 - Infosys and authorisation failure modes are particularly difficult to troubleshoot
-

Information Publishing

“ We’ve added the same per-VO job information to the glue-generator that other sites are using, slightly adapted for SLURM.

We're using a modified glue-generator.pl script in order to provide per-VO job information & different local SE for each VO and to publish walltime limits.

I had to hack glue-generator.pl for LHCb to given accounts per VO of jobs running, waiting, etc. I guess everybody else has done it.

[I suggest] add to the core to ARC per-VO job information.

Clarification on how ARC is going to deal with the demise of the BDII. ”

- Hot topic: these issues are a pain point for service performance and data consistency
- For the LHCb GLUEI issue each site hacks it in their own way
 - Difficult to merge to a single solution
- Raised issue on info publishing blocking job submission upon an a-rex start
 - GLUEI data has to go through 3 iterations of collection before publishing
- Information system integration will be reviewed
 - No easy way forward. What is reasonably expected and actually used?
- Proposal to deal with the infosystem (including investigation what is needed) by the NG 2017 conference
- Agreed to create testbed for performance measurement

Job Accounting

“ *The StartTime in the ARC usage records does not refer to when the job started running on the batch system. It refers to when it was queued..the convention has been that Starttime and Endtime are raw wall clock, while durations are (potentially) scaled to a publishing benchmark. With the current ARC scheme, it is not possible to get the true raw wallclock unless one uses the scale factor of a particular node to reverse the conversion.*

”

Write some custom scripts to archive the APEL directory.

Very poor backup procedure for accounting. I do backup of condor history.

Better mechanism for archiving APEL records as I record per job fills up the archive directory quickly and we've had to write some custom scripts to tar them up

- Some discussion on wallclock point above and will be looked at by developers
- JURA re-publishing recognised as being difficult with no helper tools provided and no default record back up
 - Agreed to be worked on in the short term

Proxy Renewal

“ There doesn't appear to be support for renewing proxies for jobs running in SLURM. We have added a section to scan-SLURM-job that pushes the arcnew'd proxies to the node running directory. ”

- Proxies are moved to WNs by RTE only in the beginning of the job and cannot be renewed (needed by pilot jobs) with site-level hacks needed
 - HTCCondor it is somehow easier
- Developers would like to get a copy of scan-SLURM-job changes to add upstream

Future Development

“ “ *Any docker support envisioned?*

I'm curious what the status of the Kubernetes backend is.

Support for shared infrastructures (University clusters, HPC).

” ”

Shared Infrastructures

- Suggested using Python LRMS tools provided in the latest release
 - Used by Chinese HPC sites
 - SLURM, SCIAPI, LSF are supported

Data Management and Caching

- Lightweight site support needs a high-performing cache
- New services: Cache and deliver your pilot on-demand data (Candypond)
 - Remove data handling from jobs, pre-initialise jobs; pre-place data on worker nodes
- Investigation of making ARC Cache work with CEPH

Containers

- Start providing Docker RTEs (would need worker node components)
 - Investigate how to use Kubernetes
 - Common approaches to using *Singularity* containers
-

Support and Feature Requests

“

Usually seem to deal with bugs in a reasonable time frame (especially with site debugging)

Minor feature request hangs in the queue for months on end

”

- Developers suggested they upgrade their Bugzilla service to hold requests and allow upvoting
 - Floated idea of half day developer/operations workshop
 - Perhaps aligned with an ATLAS software week
 - Main mailing list getting noisy, considering options on how to streamline
 - Would a dynamic troubleshooting knowledge base be useful for sites?
 - Release announcements not getting through to all sites
 - Targeted communication for ongoing site feedback
 - Occasional surgery in UK Ops meeting?
 - Site representation role at regional level
-

Conclusions

- Productive meeting for discussing long term ideas and communicating operational issues
- Developer community appreciated the good quality feedback from UK sites
 - Mostly aware of the main issues but consequent discussions helped prioritisation
- No simple solution for configuration issues affecting most sites and site-level customisation
- Work underway on configuration format and better integration support (e.g job accounting)
- Their long term development plans appear to align with T2 evolution ideas
- Suggested we are willing to evaluate new development areas (and are doing so already)

Potential Actions

- Post site configuration and customisation into GridPP github?
 - Any sensitive information we should redact?
 - Put common feature requests into new workflow
 - Any suggestions for R&D that would be useful for future UK site operations?
-