

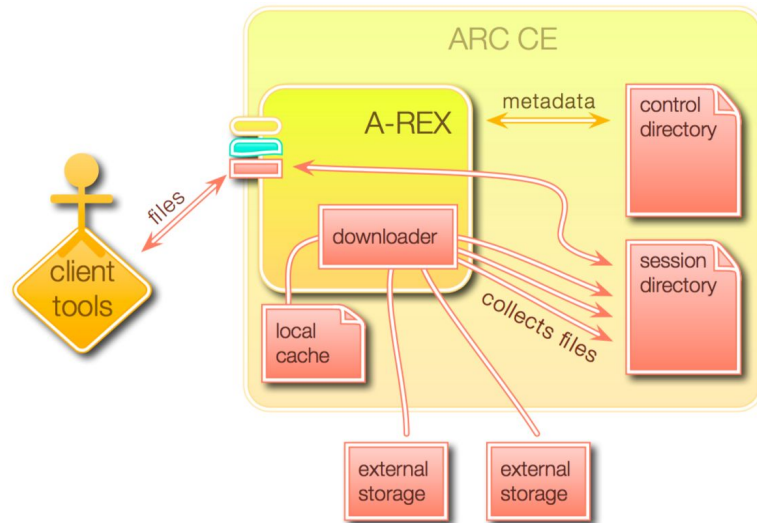
ARC Cache



David Cameron
University of Oslo

ARC CE

- ARC CE can be used in three ways:
 - Simple pilot gateway (pull mode, APF)
 - True pilot (push mode, aCT)
 - Real payload is sent with the job, but everything else is a “normal” pilot job
 - NorduGrid mode (push mode, aCT)
 - Separating Grid interactions from the job
 - aCT handles all Panda interaction
 - Data transfer done by the CE
 - Local caching of input files
 - “mv” pilot mover
 - Used in NorduGrid, HPCs, BOINC, etc



ARC Cache

- ARC CE maintains an internal cache of input files
 - Normally sits on a shared filesystem, files are linked to job working directory
 - Also possible without shared filesystem, files are copied to job working directory
- In push mode input files are known before job is submitted to the CE
- Input files (Rucio URLs) are downloaded to the cache and symlinked/copied to the job's working directory
 - The download is skipped if the file is already in the cache
 - The source SE can be selected based on closeness, or forced to particular SE
 - Full inputs are downloaded as opposed to block caching in Xcache
- Cache space is managed using LRU and watermarks in CE configuration

ARC Cache in Pull mode

- The cache could also be used in the pull model where pilots do data transfer (with or without ARC CE)
 - Replace lcg-cp/gfal-copy/xrdcp with arccp -y
 - `arccp -y /shared/cache rucio:///... localfile`
 - (Caveat: requires all local users having write access to cache and requires implementing own cache cleaning)
 - needs a new Rucio mover
- To be tested

ARC Cache Integration with Rucio

- Cache data can be registered in Rucio on a volatile RSE
 - i.e. Rucio does not manage the data on the RSE but can index it
- Useful for brokering jobs to where data is already cached
 - It is not guaranteed that the data is still in the cache when the job gets there, but not a problem
 - ARC will download it again
- A probe on the CE (the “whistle-blower”) periodically sends lists of files added to and deleted from the cache
 - ARC CE is configured to periodically dump the cache content to files
 - The whistle blower compares the dumps and looks at the difference
 - ActiveMQ messages with add/delete replica are sent to the Rucio message brokers
- The mechanism is not ARC-specific, can be used for any cache or non-Rucio-managed storage
- Current status:
 - Whistle-blower publication was tested on a couple of sites
 - Main sticking point is how to handle in Panda brokering

Use cases for ARC cache

- Distributed storage (NDGF-T1)
- No local Grid storage but shared file system (HPC)
 - For running all workloads needs to be big and performant
- Lightweight site which doesn't want to run an SE and has shared filesystem (tested at Durham)
- Very lightweight site (or cloud, vac?) without SE or CE (arccp -y, untested)

Probably not useful for

- Sites with local SE
- Sites mainly running workloads where small fractions of files are used, (eg only analysis)
 - Better to use xcache and/or direct I/O