


Integration of Rucio and ARC CE



David Cameron
University of Oslo

ARC CE

- ARC CE is a gateway from Grid to batch system
 - Used on many sites in Europe and beyond for LHC experiments and more
- It can be used as a simple pilot gateway or in a more advanced mode (like in NorduGrid, HPCs, etc)
 - Separating Grid interactions from the job
 - Data transfer done by the CE
 - Local caching of input files



ARC and Rucio

- ARC supports many data protocols
- Rucio was added when ATLAS switched from the LFC to Rucio in 2014
- Implemented using REST calls with native ARC HTTP client

```
> arcls -P
Protocol plugins available:
file - Regular local file
acix - ARC Cache Index
http - HTTP, HTTP over SSL (https) or DAV(s)
srm - Storage Resource Manager
s3 - Amazon S3 Store
rucio - ATLAS Data Management System
gsiftp - FTP or FTP with GSI security
ldap - Lightweight Directory Access Protocol
```

```
> arcls -L rucio://rucio-lb-prod.cern.ch/replicas/mc16_13TeV/EVNT.12714678._001433.pool.root.1
srm://srmv2.ific.uv.es:8443/srm/managerv2?SFN=/lustre/ific.uv.es/grid/atlas/atlasdatadisk/rucio/mc16_13TeV/e5/fd/EVNT.12714678._001433.pool.root.1
srm://grid002.ft.uam.es:8443/srm/managerv2?SFN=/pnfs/ft.uam.es/data/atlas/atlasdatadisk/rucio/mc16_13TeV/e5/fd/EVNT.12714678._001433.pool.root.1
srm://lapp-se01.in2p3.fr:8446/srm/managerv2?SFN=/dpm/in2p3.fr/home/atlas/atlasdatadisk/rucio/mc16_13TeV/e5/fd/EVNT.12714678._001433.pool.root.1
srm://sdrm.t1.grid.kiae.ru:8443/srm/managerv2?SFN=/t1.grid.kiae.ru/data/atlas/atlasdatadisk/rucio/mc16_13TeV/e5/fd/EVNT.12714678._001433.pool.root.1
srm://dcsrm.usatlas.bn1.gov:8443/srm/managerv2?SFN=/pnfs/usatlas.bn1.gov/BNLT0D1/rucio/mc16_13TeV/e5/fd/EVNT.12714678._001433.pool.root.1
```

ARC and ATLAS

- ATLAS submits jobs to NorduGrid in push-mode
 - An intermediate service (ARC Control Tower) picks jobs from Panda and submits the fully-defined payload to ARC CE
 - Memory, walltime, cores, and input files
- Input files are Rucio URLs
- ARC CE queries Rucio to resolve replicas
- A CE configuration parameter in arc.conf allows setting the order of replicas
 - `preferredpattern="ndgf.org$|cern.ch$"`
- The CE tries replicas in order until one succeeds
- Running in this mode for 2 years serving ~10% of ATLAS jobs

ARC Cache

- ARC CE maintains an internal cache of input files on a shared filesystem
- Input files for jobs are downloaded to the cache and symlinked to the job's working directory
 - The download is skipped if the file is already in the cache
- Cache space is managed using LRU
- The cache can also be used without ARC CE, e.g. in the pull model where pilots do data transfer
 - 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)

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

```
# python whistle-blower.py
usage: whistle-blower.py [-h] --cache-dir CACHE_DIR --rse RSE
                        [--broker BROKER] [--port PORT] [--topic TOPIC]
                        [--timeout TIMEOUT] [--chunk-size CHUNK_SIZE]
                        --username USERNAME --password PASSWORD
```

- 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

Summary

- Integrating the Rucio catalog into an existing system is rather straightforward
 - An HTTP client is all that is needed
- The volatile RSE concept allows caching systems to be integrated natively in Rucio
- Close collaboration between ARC and Rucio developers keeps the two products well connected