



Grid Status

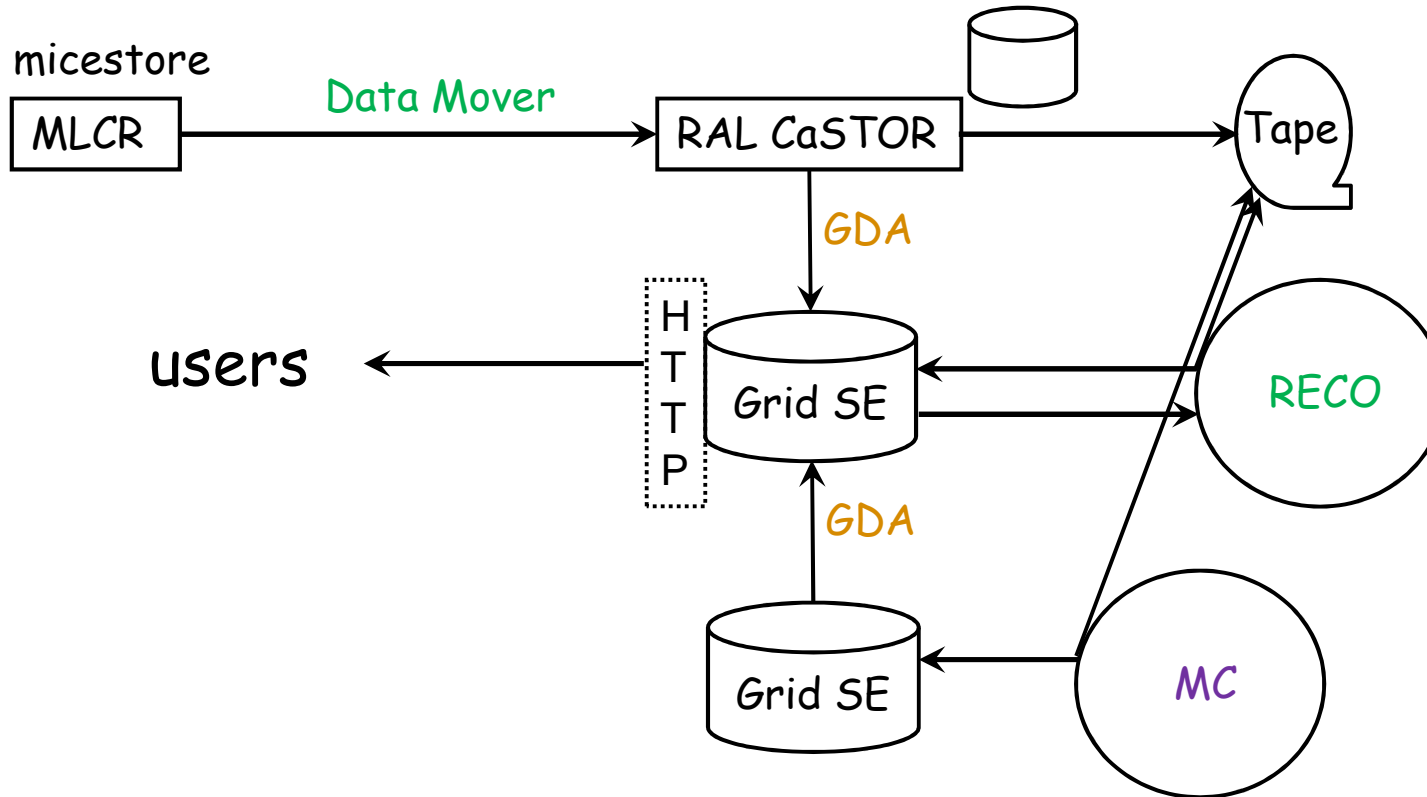


Henry Nebrensky

Brunel University



Tears of MICE



- Tear
- 0 RAL T1
 - 1 Imperial
Brunel
RAL PPD
Glasgow
 - 2

Data flow has a similar architecture to the LHC experiments
 The Data Mover and Batch Reprocessing frameworks are working.
 The GDA tool exists, but we don't yet know how to manage the process.
 MonteCarlo Production has the process defined but implementation in development



Data Mover - 1

Permanent storage of data from the MICE DAQ.

- Initial “File Compactor” automated within Run Control, still need Data Quality flags and ALH
- Manual moveFiles script still supported
- Autonomous process - “Data Mover” - then makes copy at RAL Tier 1 for tape archival
 - ◆ Proxy renewal automated from hardware token (CHEP 2016 poster accepted)
- Separate agents then make copies at Imperial and Brunel

RAW data runs identified for permanent storage *by* the act of running the File Compactor on them (auto or manually).



MICE Raw Data

Raw data on Grid, to run 8178 (as of writing)...

- Data validation up to run 8147:
- All data tarballs valid (internal checksums OK)
- All data tarballs have two copies on tape at Tier 1
- All data tarballs have overall checksums which match between Castor and MLCR copy
- Copies of RAW data at Imperial College (web) and Brunel (Grid/WebDAV)

Gathered 634 GB of data so far this User Cycle.



MICE Raw Data - 2

I have had a DOI assigned to the RAW Data, via the Brunel University Figshare account, allowing it to be cited in papers:

- Author: The MICE Collaboration
- Title: MICE Raw Data
- DOI: 10.17633/rd.brunel.3179644.v1
- <https://dx.doi.org/10.17633/rd.brunel.3179644.v1>



Data Mover - 2

Gathered 634 GB* of data so far this User Cycle, at up to 11 GB/hr stored and replicated.

CaSTOR recently tweaked for MICE, now should not wait more than an hour before actual write to tape.

LFC (Grid file catalogue service) moving to new hardware

Occasional drop-outs of the hardware dongle, possibly caused by EMI. This is (eventually) picked up by Nagios monitoring but notification depends on email (no alarm) and no on-call fix person (needs hardware intervention in the MLCR). CROC?

(* 634 GB = 171 potatoes)



CVMFS

CVMFS is

- a read-only filesystem based on HTTP; uses caching to give (usable) global coverage
- the master copy, Stratum-0, is at the RAL Tier1
- installed on Grid clusters at Brunel, Imperial and RAL PPD and Tier 1
- MAUS was compiled and built at Imperial, but 3rd-party library problems. Now build in MLCR, then the binaries moved to the Stratum-0 and replicated across the Grid
- “Recently” installed MAUS-v2.4.0, 2.5.0, 2.5.1



Offline Reco

The Offline Reco used to be running on the Tier 1 via the Grid .

Since been moved to a dedicated machine in the MLCR to allow faster MAUS updates.

- automated process waiting for new raw data files
- when a new data file appears, makes Offline Reco using latest approved MAUS and creates output tarballs
- saves them to Grid disk and tape by dedicated variant of Data Mover, also with a proxy derived from a hardware token

Has been in production during the present User Cycle.



Batch Reprocessing

A manually-triggered process, in response to a particular MAUS release becoming available via CVMFS:

- makes Offline Reco and run-specific Monte-Carlo job using that MAUS release for every run in a specified Step.
- uses Tier2 sites across the UK
- automatically runs MAUS jobs, creates output tarballs, saves them to Grid disk and tape :)
- Demonstrated ages ago, revealed issues with data unpacking in Step1 (unpacking solved, #1702)

Will need to resurrect this now we have more & bigger data.



MC Production

The procedure is conceptually similar to the reprocessing:

- Create ConfigDB entries indicating something to do
- An agent submits corresponding jobs to the Grid
- Grid jobs run a MAUS MC Production script, filling out details from ConfigDB entries, and store data locally at Grid site
- A Transfer Controller copies output to RAL Tape and Imperial SE (web access)

Taken over by Dimitrije Maletic (Belgrade) - see next talk. Production runs going through.



Miscellaneous Data

Longstanding drive to archive a variety of data from other activities in MICE, e.g.:

- Testbeam and cosmic data
- Field Maps
- Geometry and surveys
- Tracker calibrations
- Muon Beams library
- EPICS Archiver archive

(Data curation seems to be becoming a topic with the funding agencies)

The Grid storage is my responsibility (wearing my Archivist hat) but *preparation, indexing and making incoming data available rests with its creator!*

<https://micewww.pp.rl.ac.uk/projects/computing-software/wiki/GridDataStorage>₁



Data Volume

For data to continue to be available into the future it must be properly archived, not just left lying around.

Long term storage is provided on disk and tape at various Grid sites. It doesn't appear by magic...

... we do need realistic estimates of how much data different activities will produce, and how reliably we want it preserved.

We have a 75 TB tape quota at RAL expecting to grow to 300 TB next year. Are using 9.

Expect 30 TB of RAW & RECO from Step IV, with the rest being Monte Carlo. Needs clarifying #1618



Outstanding tasks

Brief list:

- RECO data mover (done)
- data distribution to Brunel, RAL PPD and Glasgow
I need to deploy GDA instances at Brunel to organise file transfers to other Tier 2 sites
- Simulation Muon Beams tarball
- Various monitoring and alarms...

More lurking in Grid meeting minutes.