




Science and
Technology
Facilities Council


Antares migration

Tom Byrne, George Patargias, Alison Packer,
Alastair Dewhurst

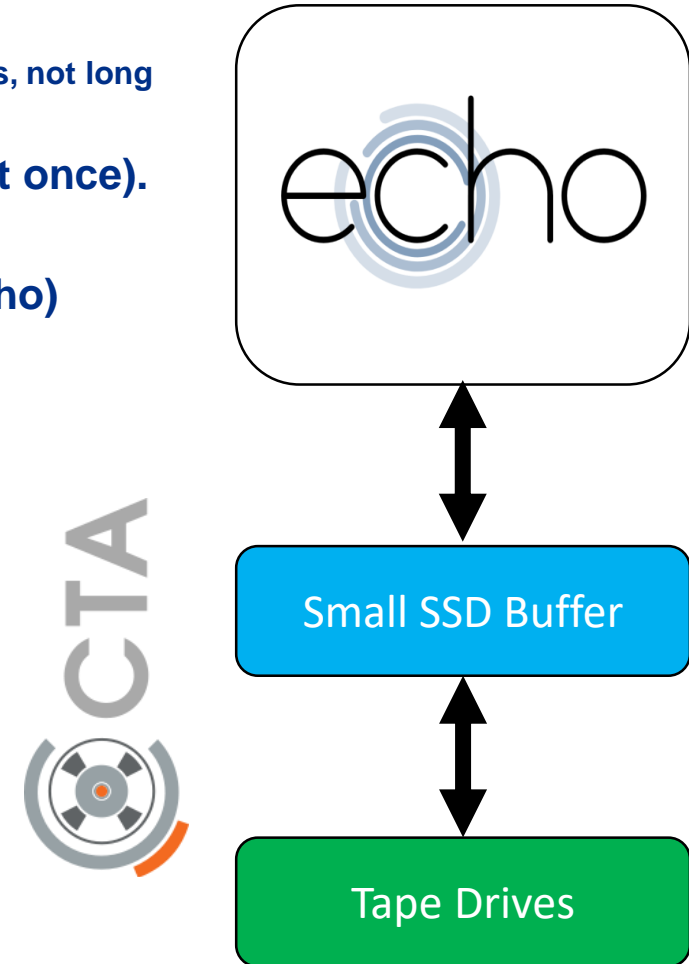


Castor vs Antares

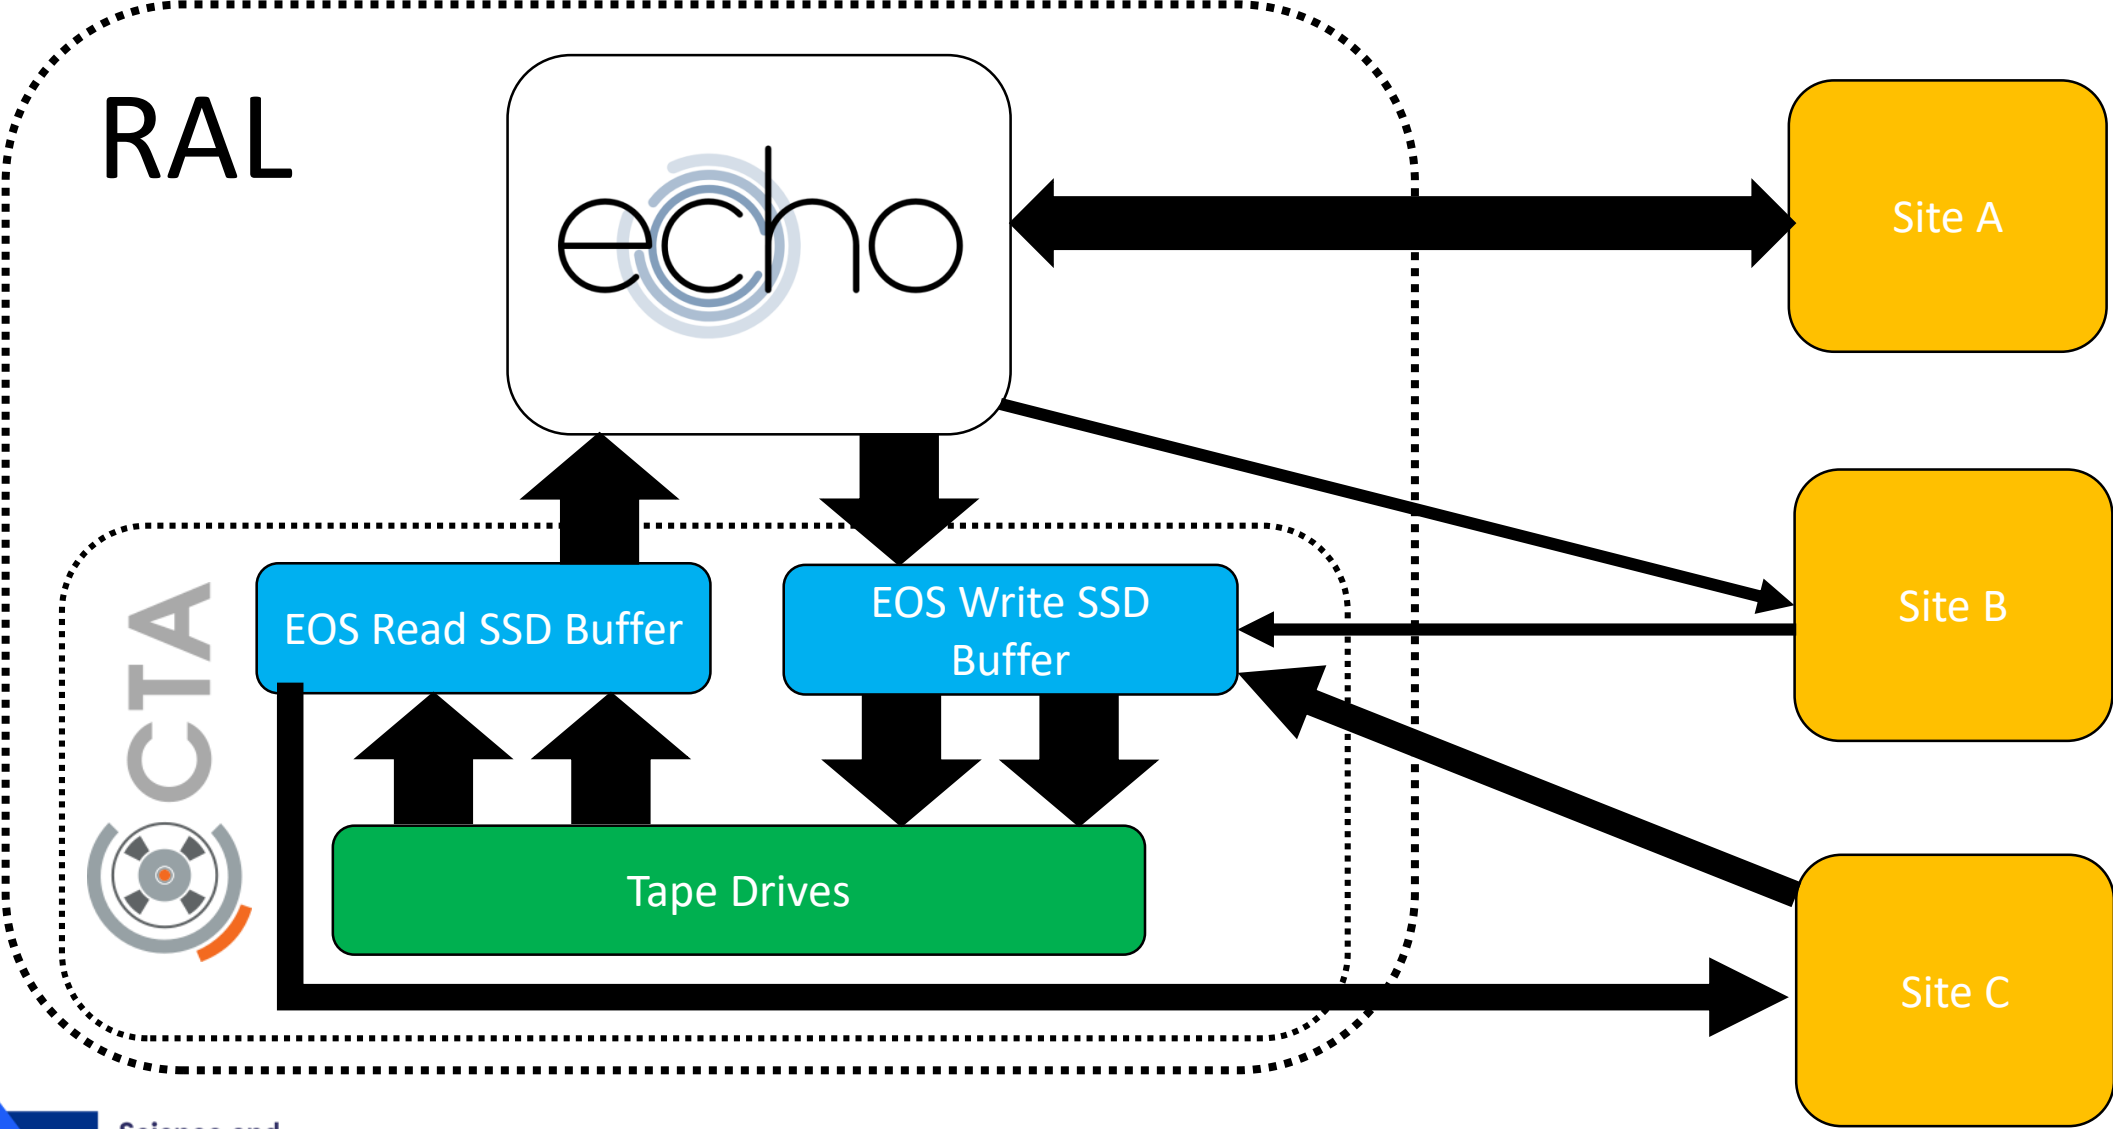
- **Castor had a large HDD based buffer in front of the Tape Drives**
 - Used 5% of VO disk pledge
 - Data not controlled by VO, it would stay until space needed, long enough for transfers, not long enough to reliably re-process data.
- **Tape drives are single stream (only one file can be written to a tape at once).**
 - Difficult to maintain high throughput if two files happen to be on the same HDD.
- **CTA puts the disk space back in the VO control (by providing it in Echo)**
 - In most cases VO still needs to use the tape buffer.



SSD buffer is designed to move data to and from tape drives to local disk rapidly.



Access to Antares




Direct Access to Antares



For non-LHC VOs we don't expect the amount of data cause any problems.

Reading / Writing via Webdav / XrootD will be supported (direct or TPC).

Multi-hop access



If you can only access via GridFTP then a multi-hop via Echo would still be required.

Migration

We will be migrating all VOs simultaneously during a 2 – 3 day outage on the week beginning the 28th February.

The migration from Castor to Antares only requires the migration of metadata between the systems and will not require any data to be moved between tapes. During the tape migration all other services at RAL are expected to be running normally.

The upgrade can be split into 4 stages:

1. Shutting down Castor, ensuring all data is safely migrated to tape and taking backups. The time this takes is highly dependent on the activity of the VOs before the downtime. If VOs are able to reduce any writing in the days before the outage then it can be quite quick. We expect for the official downtime to start some time on Monday the 28th February. Individual VOs may wish to stop scheduling writes to Castor a few days before.
2. Migrating the metadata from Castor to Antares. The actual migration will migrate the Castor Oracle database to the Antares Oracle database and EOS namespace. In our tests, the migration proceeds at slightly above 100,000 files a minute, which means we expect the total migration to take up to 8 hours (on Tuesday 1st March). There will be some contingency in case part of the migration fails so we can try again. The advantage of migrating everyone at the same time, means that we can be much more aggressive with the database migration as the Antares database is effectively a blank slate.
3. Checking the migration has been successful and reconfiguring VO endpoints. Once the migration has been completed RAL will internally check that the data has been migrated successfully. However given the scale of the change and the fact that both the surls and access method will have also changed (e.g. `srm://srm-cms.gridpp.rl.ac.uk/...` to `root://antares...`) we would request a more detailed test in collaboration with VOs before allowing normal production work to restart. E.g. full end to end write/read test with non-critical data.
4. Migration of tape drives and servers to Antares. This will be transparent to the VOs and happen a few days after (but hopefully before the tape challenge) the migration is complete. Currently we have half our tape hardware assigned to Castor and half to Antares. Once we are successfully using Antares in production, the tape hardware will be re-assigned exclusively to Antares doubling our maximum bandwidth.

Non-LHC VOs

The non-LHC VOs can be divided into two categories: Active and dormant

The dormant VOs haven't accessed their data in at least a year. We will therefore be providing them with an updated file list and a recommended way of accessing it should they need to. But the assumption is, if they ever do actually need to recall their data we will probably need to provide them with help.

The Active VOs are:

1. Solid
2. SNO+
3. NA62

Backup