



Science & Technology  
Facilities Council

UK Research  
and Innovation

# XRootD for ALICE at RAL

Rob Appleyard, George Patargias, Ian  
Johnson

# Storage at RAL: An Overview

- CASTOR for WLCG tape storage
  - Re-implemented service as a dedicated tape endpoint\*
- Ceph-based Echo storage system for disk
  - Migration from CASTOR nearly complete

\*<https://indico.cern.ch/event/765497/contributions/3357957/>





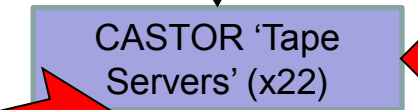

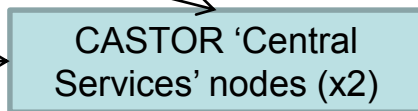


# CASTOR XrootD for ALICE

- CASTOR at RAL recently reimplemented
- Previous: 4 separated instances, one per user
  - Necessary when CASTOR serving both disk and tape
- Now: 1 tape endpoint shared between all users

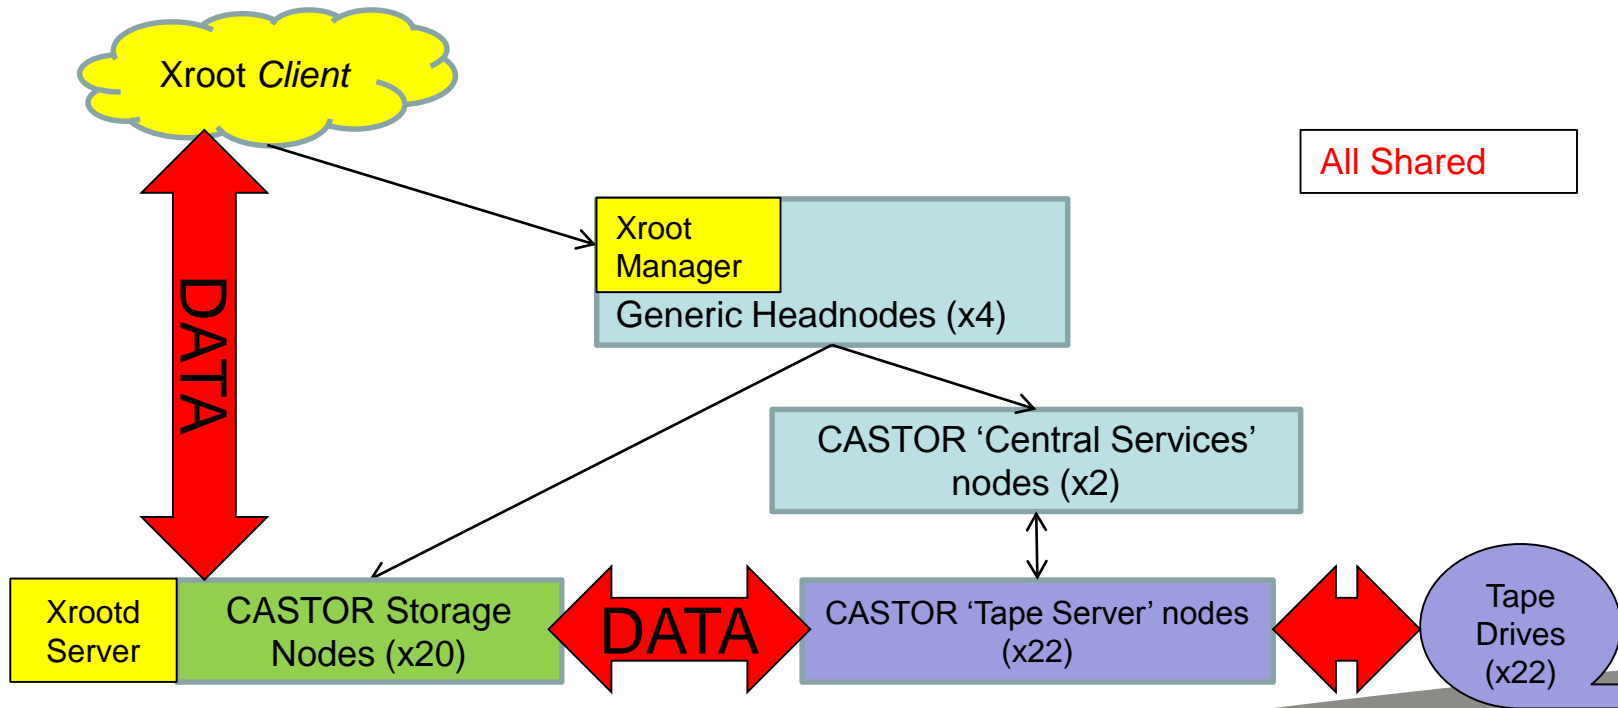
# A Picture of CASTOR pre-2018

Replicated per VO (4)

Shared



# A Picture of CASTOR 2018-



# Non-ALICE CASTOR Xrootd Config

```
xrootd.fslib /usr/lib64/libxrdxcastor2fs.so.2.1
xrootd.chksum adler32
xrootd.async off nosf
xrd.sched mint 16 maxt 512 idle 128
all.export /
all.role manager
all.manager localhost 2131
oss.fdlimit 16384 32768

xcastor2.nsmmap / /
xcastor2.stagerhost castor-stager.gridpp.rl.ac.uk
xcastor2.targetport 1095
xcastor2.capability true
xcastor2.privatekey /opt/xrootd/keys/key.pem
xcastor2.tokenlocktime 60
xcastor2.proc /var/log/xrootd/manager
xrootd.seclib /usr/lib64/libXrdSec.so

# GSI authentication
sec.protocol gsi -crl:3 -gmapopt:2 -gmapto:3600 -d:0
# Configure gsi as the default
sec.protbinding * only gsi

xcastor2.stagermap /castor/ads.rl.ac.uk/test/func/wlcgtape/ gridTape
#ATLAS mappings
xcastor2.stagermap /castor/ads.rl.ac.uk/prod/atlas/farmOut/ atlasT0Raw
xcastor2.stagermap /castor/ads.rl.ac.uk/prod/atlas/raw/ atlasT0Raw
(.....)
```

# Non-ALICE CASTOR Xrootd Config

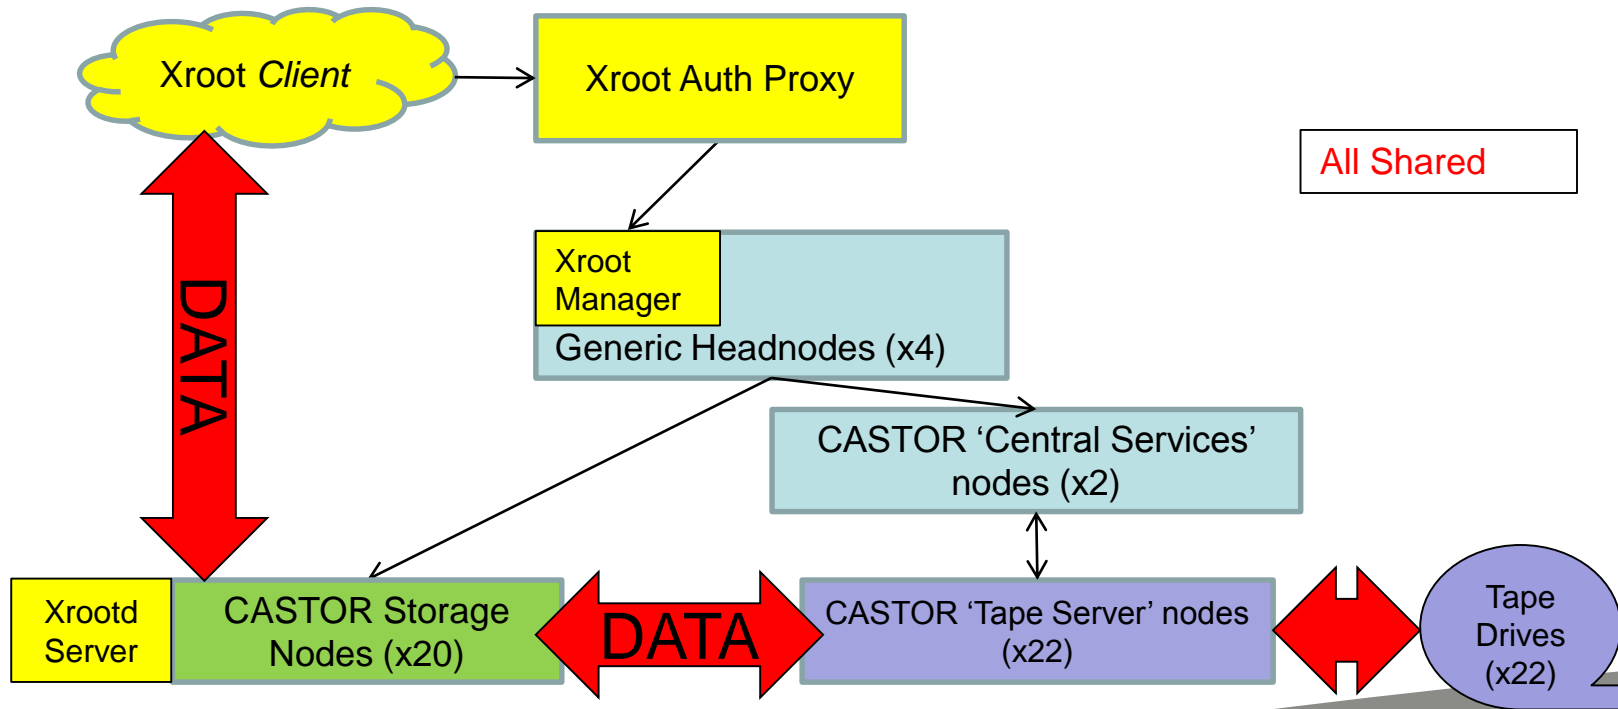
- xcastor2 is the CASTOR xrootd plugin
- GSI auth
  - Go look at a grid-mapfile
- Lots more namespace linkages

# Why can't ALICE use that?

- ALICE use their own authentication system
- Incompatible with generic config
  - Needs separate daemon
- Could run two daemons on same nodes with segregated config, but easier to host this elsewhere.
  - Also, in ~2 years we will have a new tape system, useful to have all ALICE specials in one place



# A Picture of CASTOR - ALICE



# ALICE CASTOR Xrootd Config

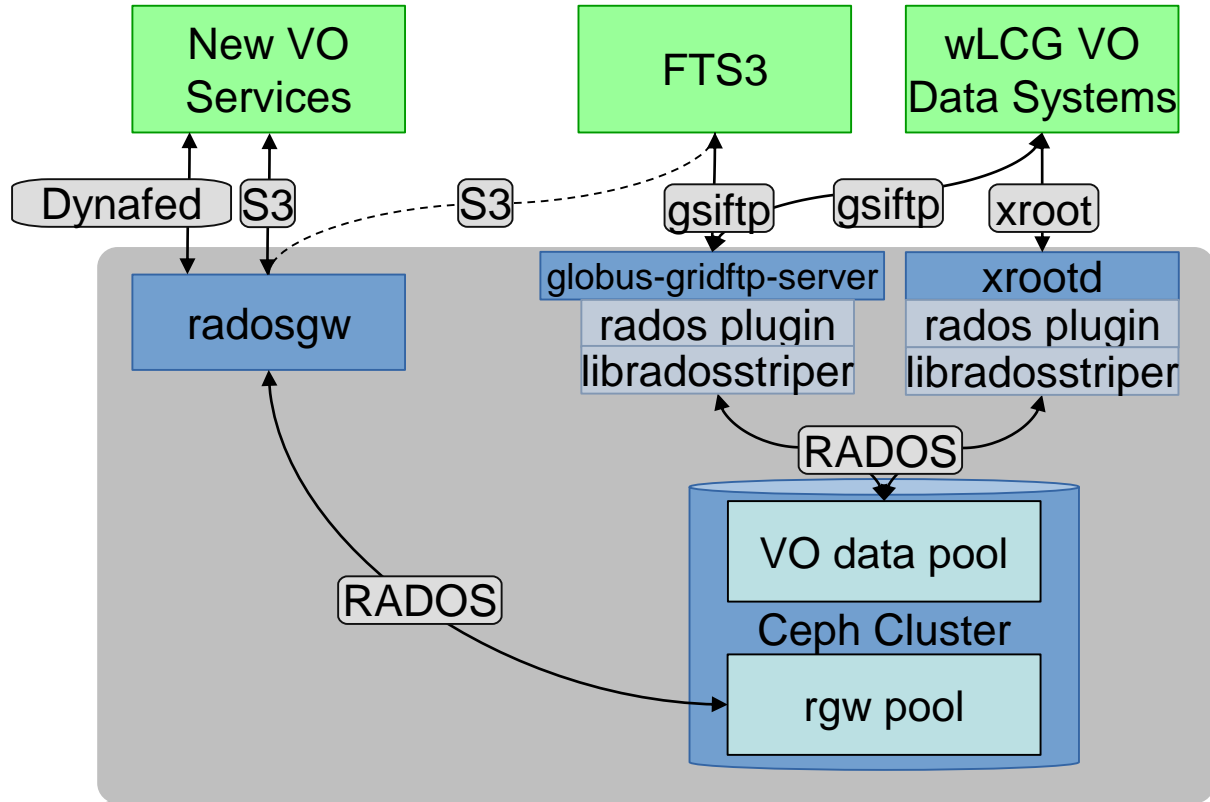
```
xrootd.fslib /usr/lib64/libxrdxcastor2fs.so.2.1
xrootd.chksum Adler32
xrootd.async off nosf
xrd.sched mint 16 maxt 512 idle 128
all.export /
all.role manager
all.manager localhost 2131
oss.fdlimit 16384 32768

xcastor2.nsmap / /
xcastor2.stagerhost castor-stager.gridpp.rl.ac.uk
xcastor2.targetport 1095
xcastor2.capability true
xcastor2.privatekey /opt/xrootd/keys/key.pem
xcastor2.tokenlocktime 60
xcastor2.proc /var/log/xrootd/manager
xrootd.seclib /usr/lib64/libXrdSec.so
# If "skipabort" option is true then we no longer abort requests for which
# the client disconnected in before getting the file. Default value is "false".
xcastor2.skipabort true
# Alice-Specific Extra Configuration
alicetokenacc.noauthzhost castor-functional-test1.gridpp.rl.ac.uk
alicetokenacc.noauthzhost lcgui03.gridpp.rl.ac.uk
alicetokenacc.noauthzhost lcgui04.gridpp.rl.ac.uk
xcastor2.authlib /usr/lib64/libXrdAliceTokenAcc.so
xcastor2.authorize 1
xcastor2.role * alicesgm
ofs.trace none
```

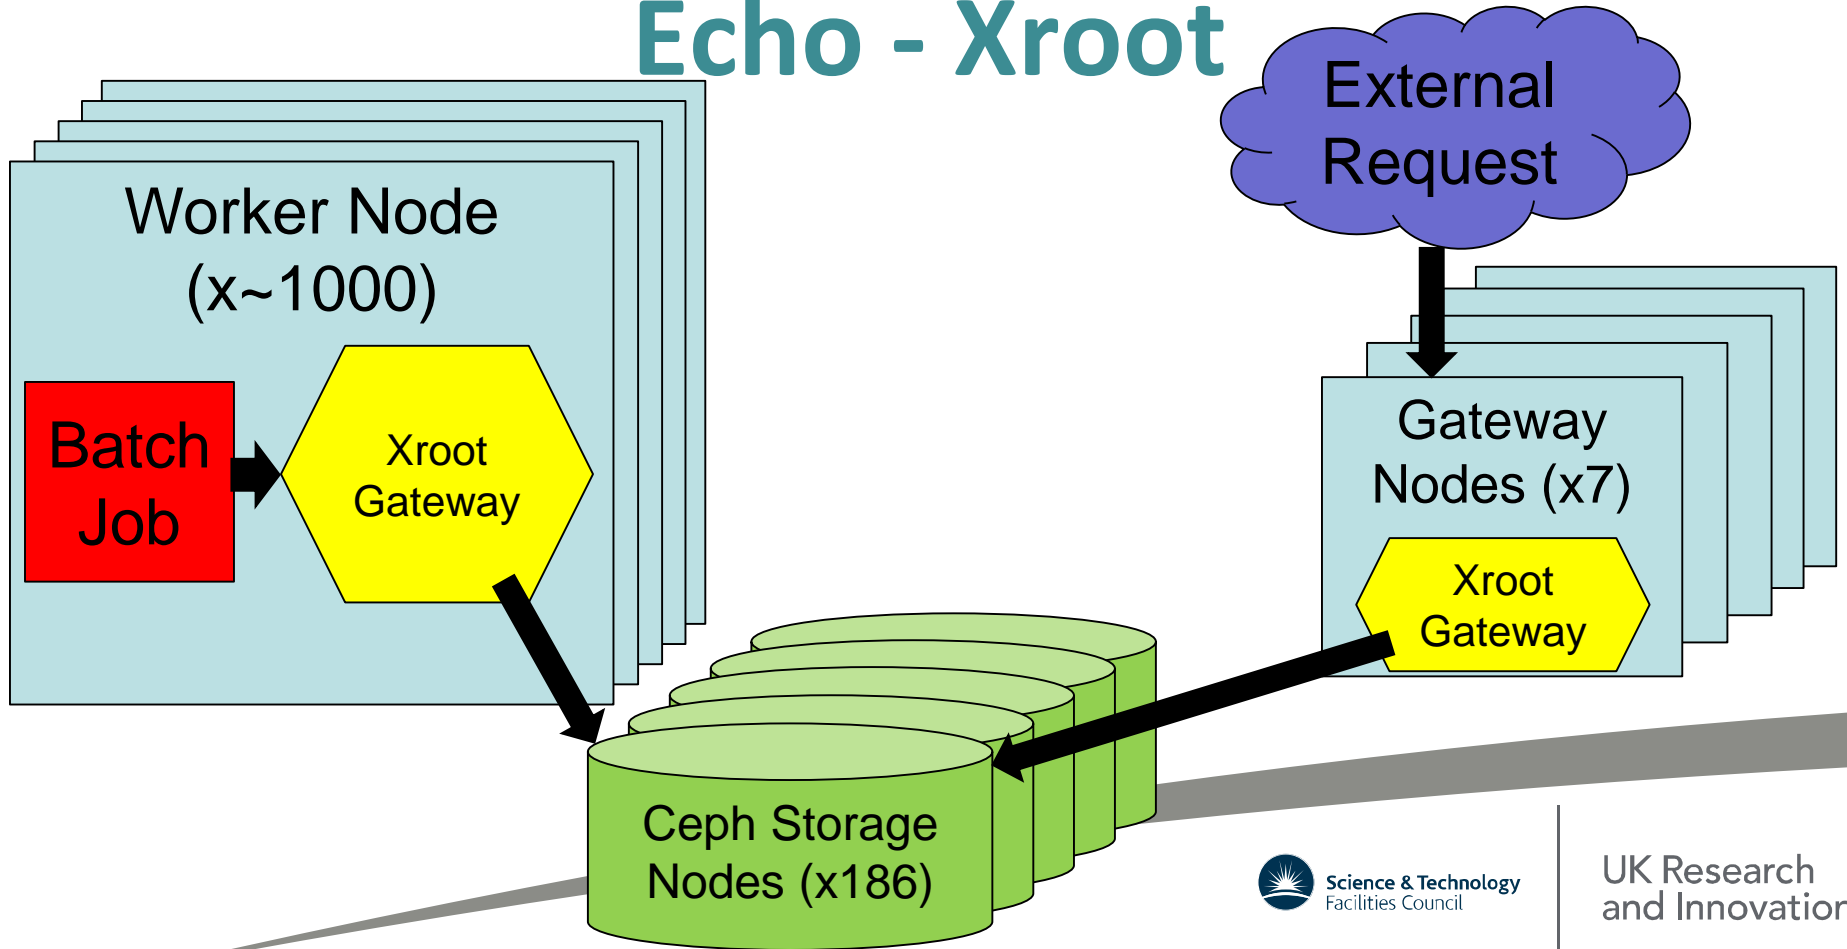
# ALICE CASTOR Xrootd Config

```
#-----  
# Security plugins  
#-----  
  
# UNIX authentication  
sec.protocol /usr/lib64/ unix  
  
# Configure ALICE  
sec.protbinding * only unix  
  
#-----  
# Mapping between paths and svcClasses. The path "/" is considered as the  
# default mapping. SvcClasses need to be separated by commas and at the end  
# there is the possibility to specify the "nohsm" option. SvcClasses are listed  
# in order of their priority. The special character "*" means that any svcClass  
# provided by the client is accepted.  
#-----  
#####  
#Alice mappings  
xcastor2.stagermap /castor/ads.rl.ac.uk/prod/alice/tape/      aliceTape  
# For Alice functional tests  
xcastor2.stagermap /castor/ads.rl.ac.uk/prod/alice/test/    aliceDisk  
xcastor2.stagermap /castor/ads.rl.ac.uk/test/              aliceDisk
```

# Echo



# Echo - Xroot



# ALICE on Echo

- ALICE use token scheme passed as parameter in the request URL
- Echo uses 'gateways' as clients to Ceph.
  - Similar to CASTOR in that generic setup not good for ALICE
- ATLAS/CMS/LHCb use gateways on WNs
- ALICE small so don't (just use a separate gateway for now)
  - Necessary ALICE layer can be containerised

# ALICE Tokens

- Echo uses an 'authdb' file to determine whether a request can proceed.
  - Cross references role making request and target path
- ALICE uses 'capability-based' scheme.
  - A request comes packaged with an affirmation from an authority to say that it is authorised.
  - Check whether sig comes from recognised authority

# Implementation

- Xrootd authenticates a request coming from a specific ALICE data-mover user.
  - One half of a shared secret.
- Xroot daemons runs as alicesgm. That user is authorised for Echo
- Rely ALICE on central authority to authorize access to alicesgm
- TPC the final hurdle



# Any questions?