



Science and
Technology
Facilities Council

RAL site update

George Patargias on behalf of
the Antares team

pre-GDB Tape Evolution

07/11/2023

Introduction

- **Two main storage endpoints at RAL**
 - ✓ **Echo: Ceph cluster for disk storage**
 - ✓ **Antares: EOS/CTA instance for tape storage**
- **Antares main features**
 - ✓ **Tier-1 EOS cluster: 13 x 1.5TB SSD nodes**
 - ✓ **Facilities EOS cluster: 2 x 3.5TB + 7 x 1.5TB SSD nodes**
 - ✓ **Tier-1 tape library: 20 x TS1160, 14 x LTO9 drives**
 - ✓ **Facilities tape library: 28 x TS1160, 12 x LTO8, 6 x LTO9 drives**
 - ✓ **EOS: 4.8.105-1, CTA: 4.8.7-1**

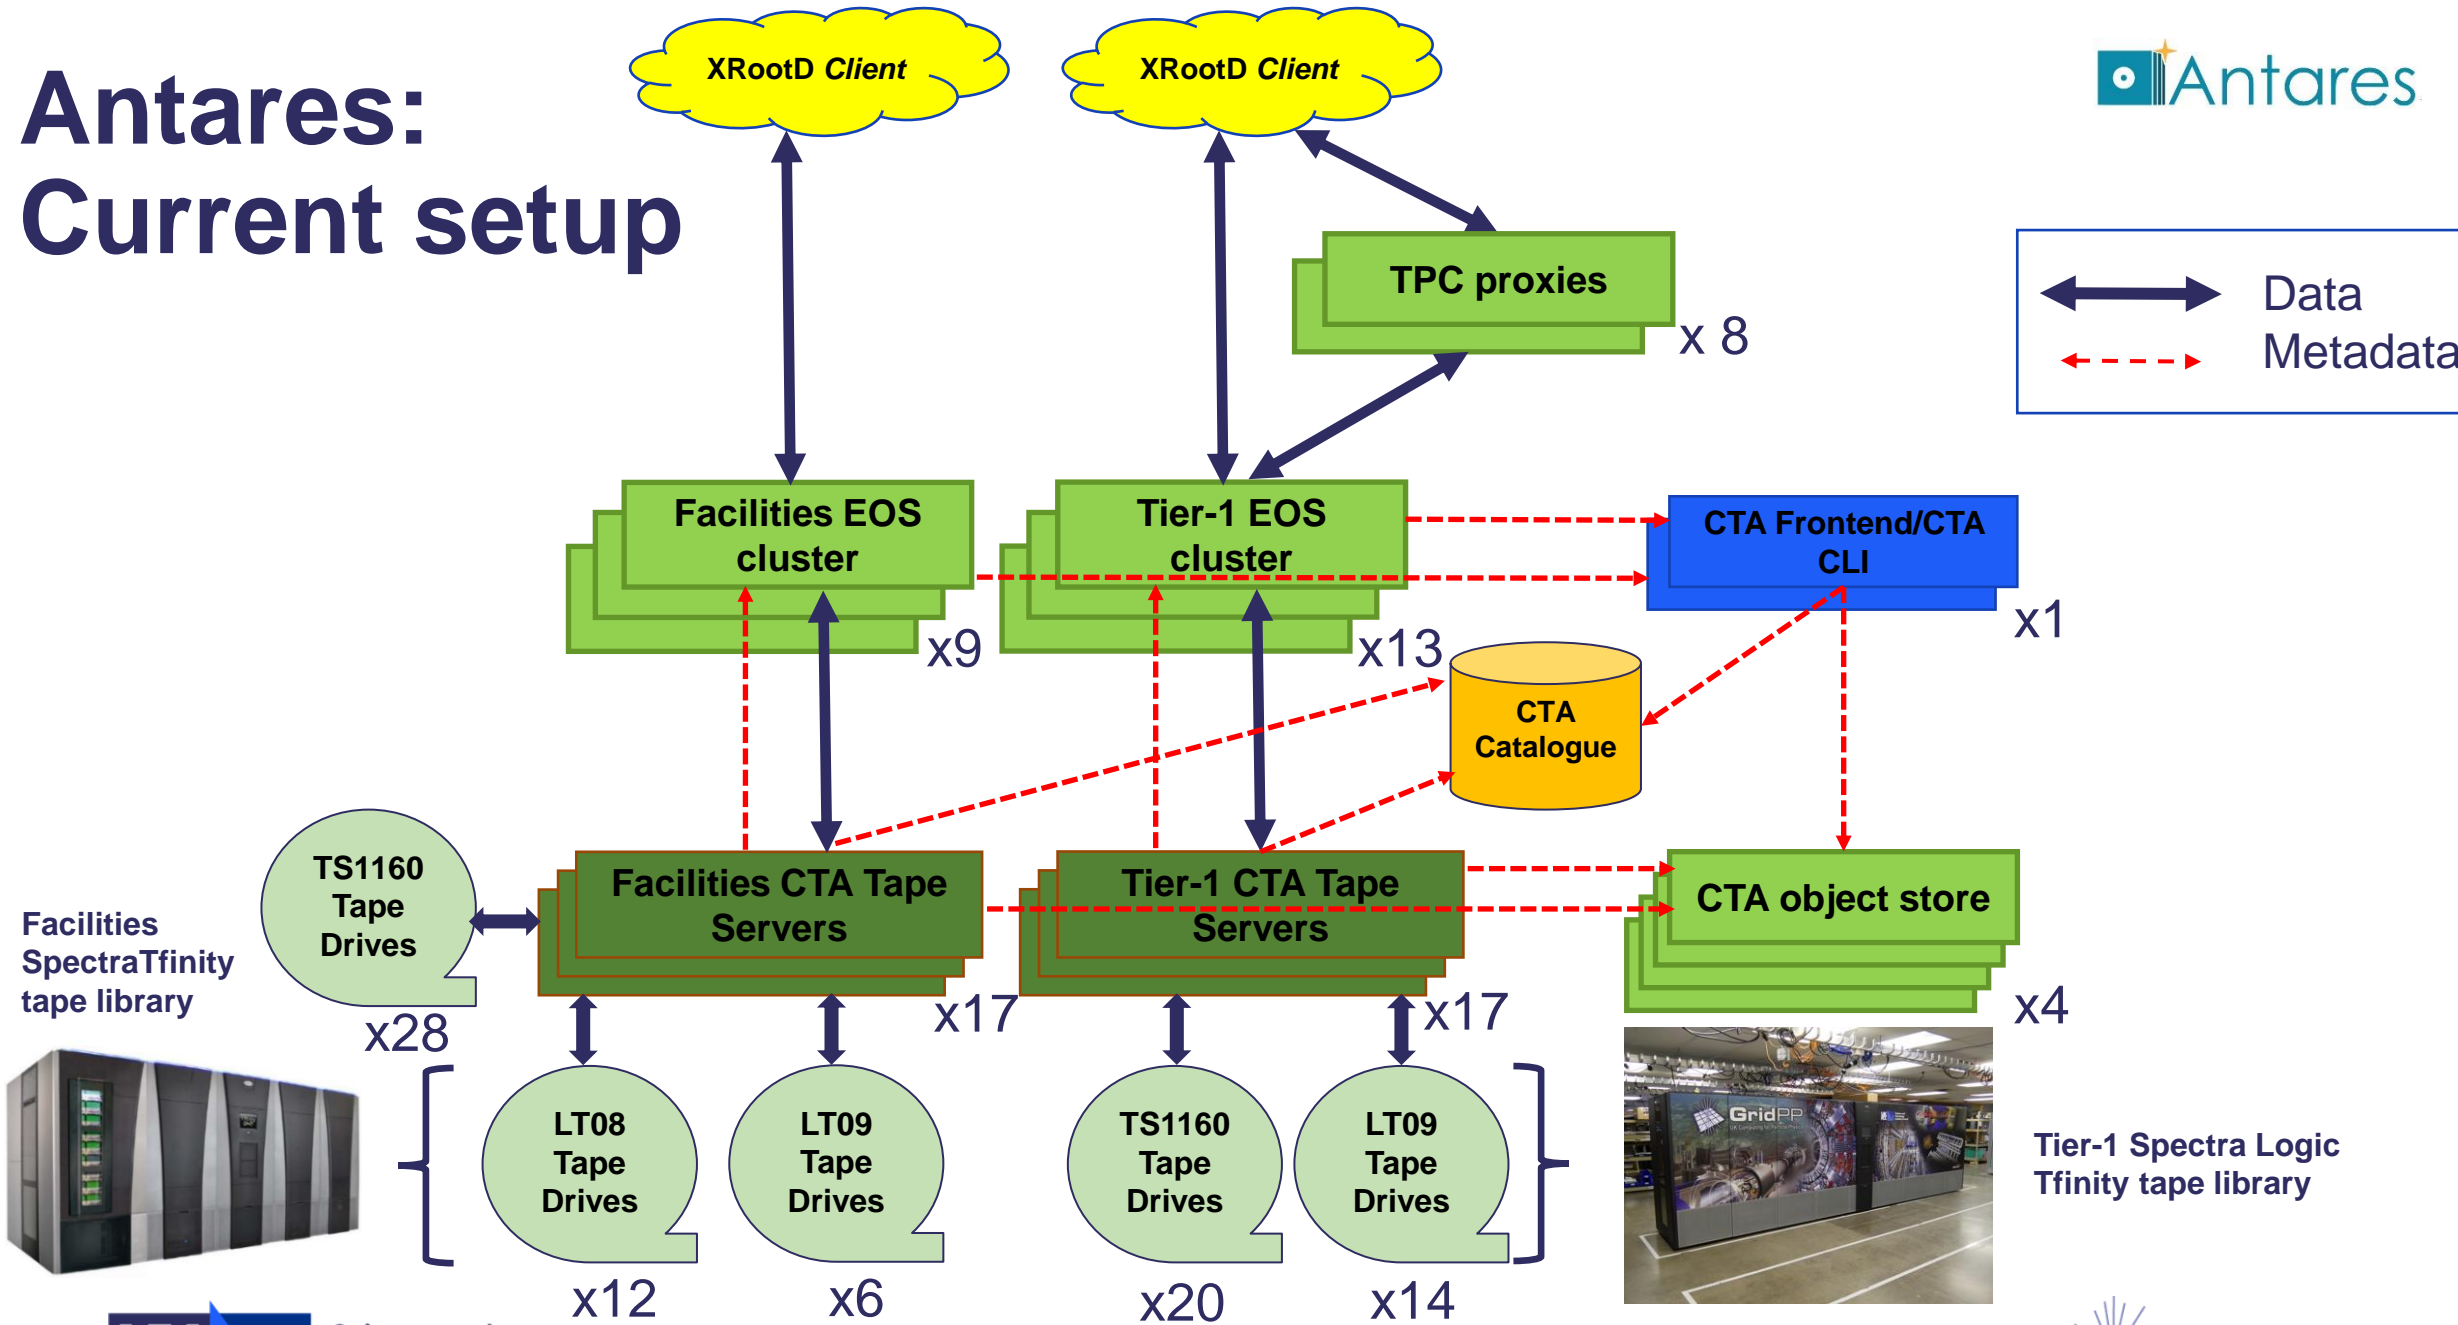
Recent developments

- May 2023: Minor version upgrade to EOS 4.8.98 and CTA 4.8.7 and enabled HTTP Tape REST API
- June 2023: CASTOR Facilities migration to Antares completed
- August 2023: CASTOR decommissioned at RAL!

CASTOR Facilities migration

- The last CASTOR to CTA migration
- Data migrated: ~13.3 million files (>100PB)
- More complex than the Tier-1 data migration
 - ✓ Integrate Facilities client code with EOS-CTA
 - ✓ Remove FileID clashes with ~100,000 Tier-1 files
 - ✓ Care for the dual-copy tape pools

Antares: Current setup



RAL site update

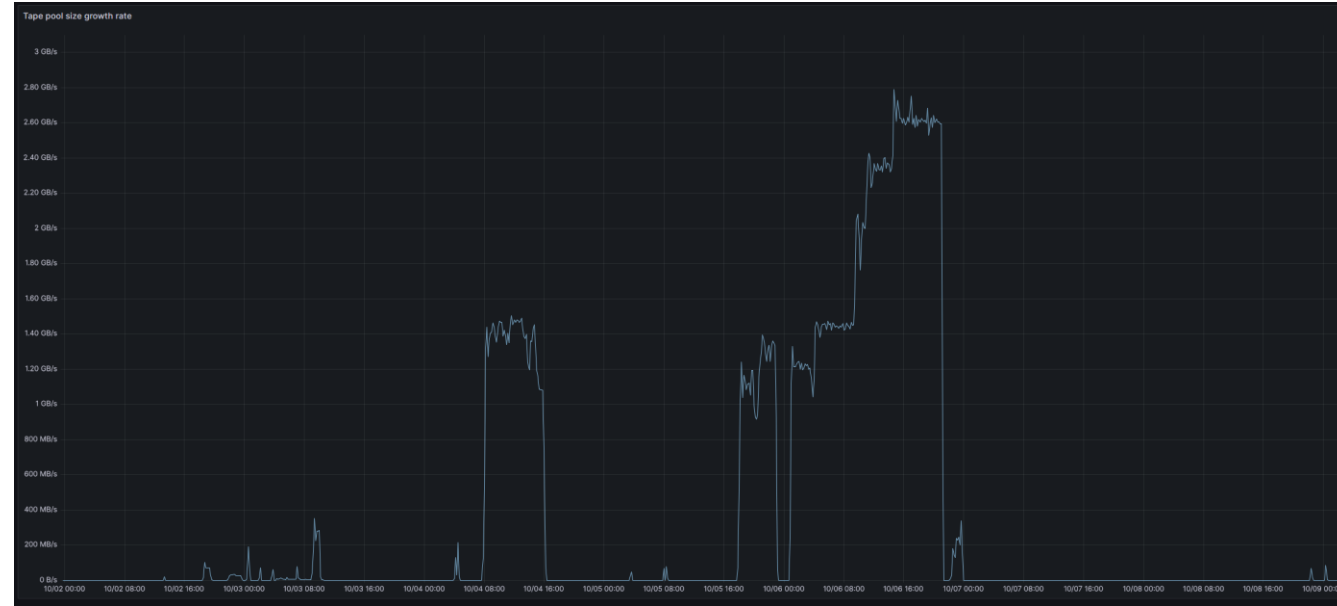
Experiment file sizes

- For LHC, the max file size is limited to 128GB
 - For local users 1.1TB
- Small SSD buffer that fixes the files in place until archived/recalled
- Size of the largest available EOS FS determines CTA max file size
- CTA max file size may need change due to some uncertainty in the size of the file aggregates that Facilities ingest servers send to Antares
- What is CTA max file used by other sites?

Antares Tier-1

- LHC Run-3 ingest rates are achieved

atl23 pool
growth rate



- Initially, only XRootD access was offered to LHC users
 - Use TPC proxies for external transfers
 - Combine with WebDAV by doing multi-hop transfers
- WebDAV fully enabled
 - Used by ATLAS and LHCb
 - Deployed fix (EOS 4.105) for the bug that disrupted ATLAS recalls

CTA Databases

- Until August, the Antares DB was a 3 node RAC
- The DBA team decided that a third node was not needed – DB load very low
- CTA Oracle service is running on a single node anyway
- The two Oracle nodes removed from CTA were used to create a new OL8 DB

Future plans

- Upgrade to EOS5 following CERN
- Run cta-taped in a container
- Plan and carry out upgrade to Rocky 9
- Transition to LTO media

Tape storage for SKA



- **Construction phase expected to finish 2028**
 - Marking start of full operations
- **Keep recent (and popular) data on disk**
- **Custodial data archived to nearline storage**
- **Current estimates suggest ~ 1 Exabyte of data to tape (or equivalent nearline storage) (including replicas for resilience).**
- **Increasing linearly per year.**
- **UK will represent up to O(15-20)% of SKA data.**