



Science and
Technology
Facilities Council

Antares - status and plans

George Patargias on behalf of
the Antares team

GridPP49 & SWIFT-HEP05 30/03/2023

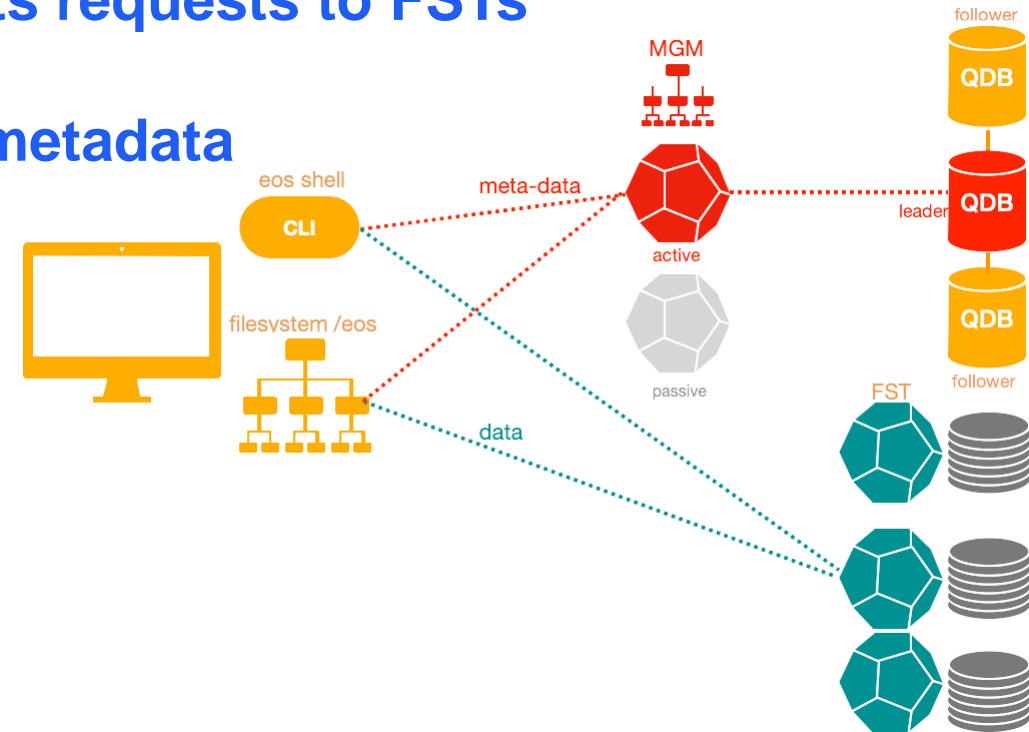
The Big Picture

- **Antares = EOS + CTA**
- **Deployed at CERN Tier-0, so an obvious choice for RAL**
- **Has replaced CASTOR WLCGTape (Tier-1)**

EOS



- Namespace management and external request handling
- Multiprotocol support: XRootD, HTTP/WebDAV (also S3, GridFTP, GRPC, FUSE)
- MGM: manages namespace metadata and re-directs requests to FSTs
- QuarkDB: High available backend for namespace metadata
- FST: Data storage server
- Provides a bandwidth oriented tape buffer
- Tier-1 EOS cluster: 13 x 1.5TB SSD nodes
- EOS production version: 4.8.88-1



CTA

The tape back-end system evolved from CASTOR

- **CTA Frontend:** Central XRootD server that accepts all tape requests from EOS and stores them in the CTA Object Store
- **CTA Object Store:** Backend tape scheduler DB implemented as a Ceph object store
- **CTA DB (aka CTA Catalogue):** Backend Oracle DB holding all tape file metadata
- **CTA Tape Servers:** CASTOR type tape servers controlling tape library drives
- **CTA production version: 4.7.14-1 (v12.0 schema)**

VOs served

LHC

VO	Data volume	Pledge*
ATLAS	42.2 PB	53 PB
CMS	14.8 PB	24.4 PB
LHCb	14.3 PB	39.9 PB
ALICE	921.4 TB	2.5 PB

* From 2023 values in CRIC

Non LHC

VO	Data volume	Pledge**
NA62	8.2 PB	7.7 PB
DIRAC	2.6 PB	3.0 PB
SOLID	1.7 PB	2.0 PB
T2K.ORG	654.8 TB	1.3 PB
SNOPLUS	481.6 TB	500 TB
ILC	216.3 TB	400 TB
MICE	62.2 TB	100 TB
GRIDPP/MIGDAL	7.9 TB	n/a
DUNE	547 GB	n/a

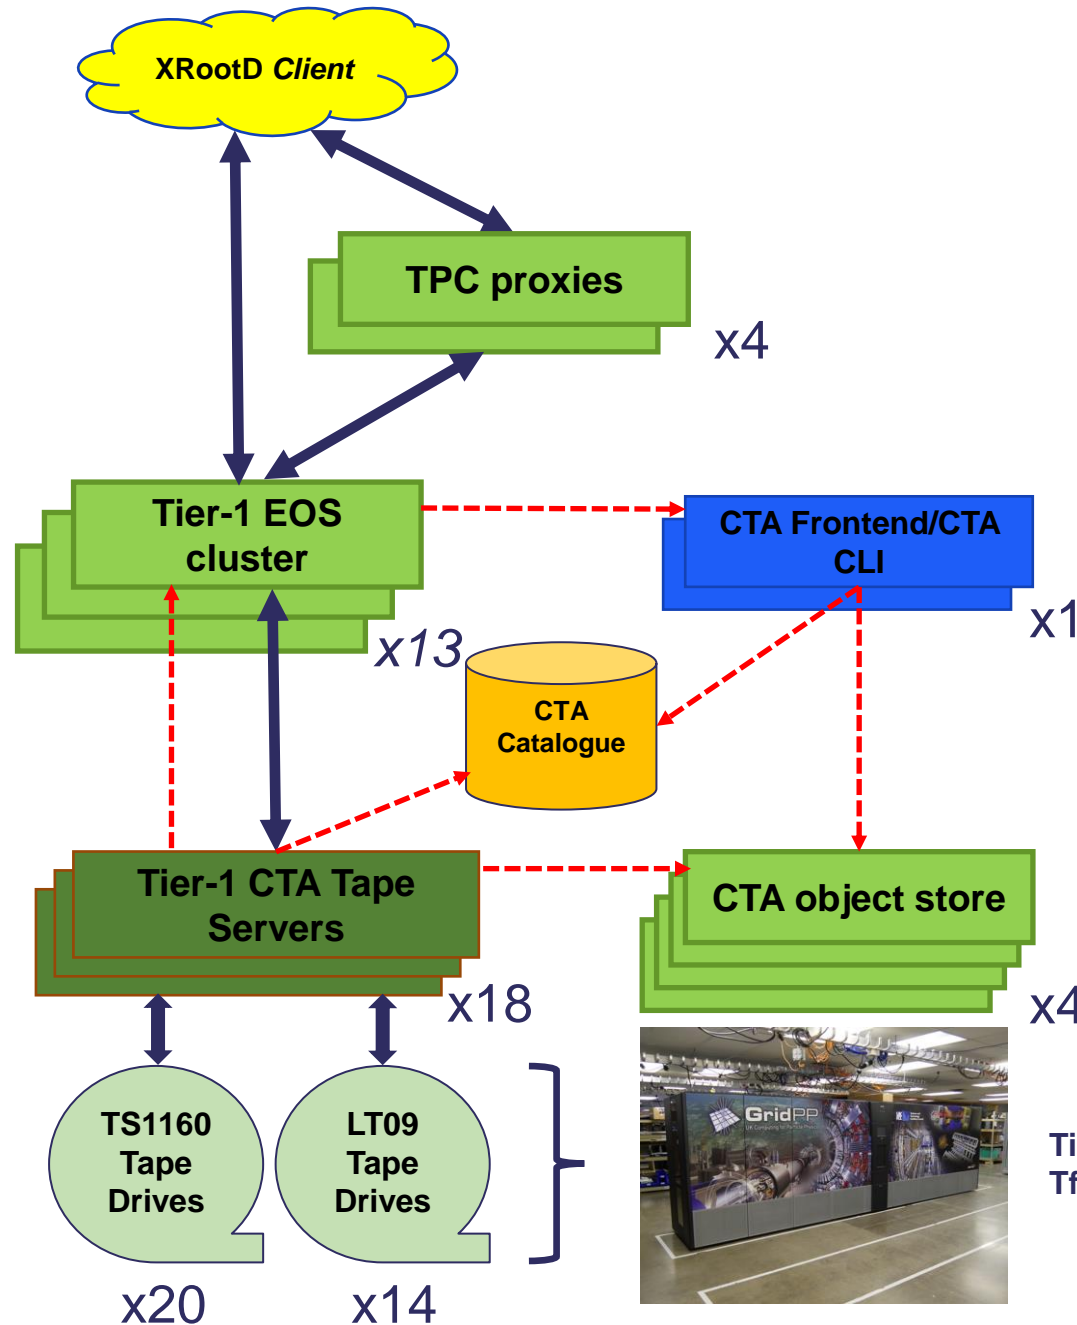
** From Oct 2022 Resources Review Meeting
(<https://indico.cern.ch/event/1208356/>)

STFC Facilities

Facility	Data volume
CLF	13.4 TB

Antares: Current setup

Tier-1 EOS cluster
against one CTA instance



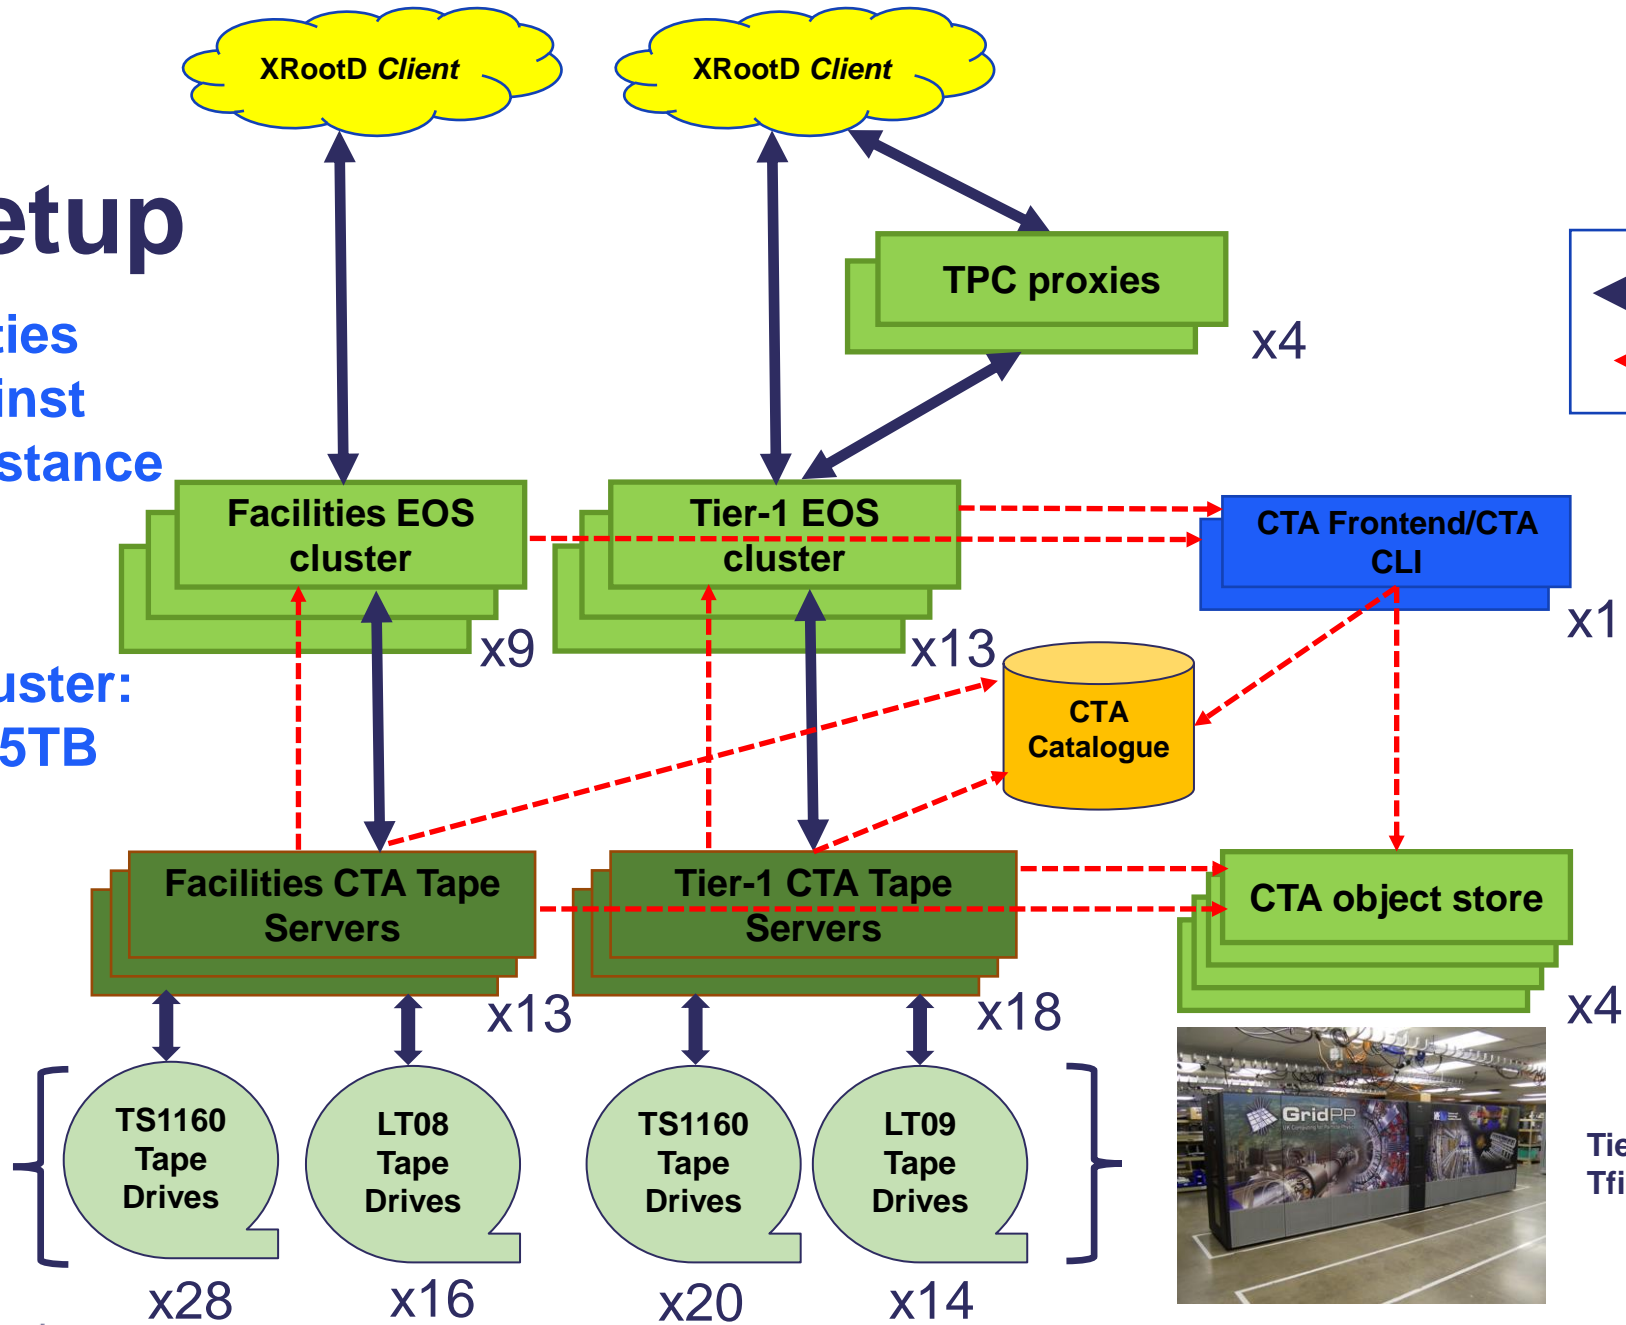
Tier-1 Spectra Logic
Tfinity tape library

Antares: Future setup

Tier-1 and Facilities
EOS cluster against
the same CTA instance

Facilities EOS cluster:
2 x 3.5TB + 7 x 1.5TB
SSD nodes


Facilities
SpectraTfinity
tape library



Tier-1 Spectra Logic
Tfinity tape library



Tier-1 tape library (Asterix)



- **Total of 15 frames - longest in Europe/UK!**
- **20 x TS1160, 16 x LTO9**
- **Capacity: 250PB**

Facilities tape library (Obelix)



- **13 frames**
- **28 x TS1160, 17 x LTO8, 6 x LTO9**
- **Capacity: 212 PB**

Antares performance

➤ Overall performance: very good

- LHC Run 3 rates are achieved (almost) always
- Number of concurrently mounted drives is the rate limiting factor
- Resolved issue with LTO9 drive performance

➤ Repeated LHCb tape challenge (Dec 2022 - Jan 2023)

- Average writing speed: 2.6 GB/s (Run3 target 2.96 GB/s)
- Average staging speed: 3 GB/s (Run3 target 1.93 GB/s)
- Much better results compared with the previous challenge

Operational issues

➤ EOS MGM becomes unresponsive

- Mitigation: a cron script restarts eos@mgm if necessary
- Manual clean-up of failed archive requests

➤ cta-taped shuts down after repeated mount failures

- Number of failed mount attempts that trigger shutdown is hard-coded
- Because we have 2 drives/tape server, both drives are brought down
- Suggested mitigation: run cta-taped in a container

Short-term plans

- **CASTOR Facilities migration**
- **Configure Tape REST API**
- **Upgrades to EOS/CTA 5**
- **Hardware**
 - Add 4 x TPC proxy nodes**
 - Move the CTA Frontend to a physical host**

Mid-term plans

➤ Accounting

- Tape accounting as per WLCG requirements

➤ Monitoring

- MGM metadata operations
- EOS functional check
- Garbage collection in the retrieve space
- Status of drives

➤ Network: Enable IPv6 traffic

THANK YOU!