


# Evolution of the STFC's Tape Archival Service

George Patargias, RAL Tier-1



# About RAL



# CASTOR at RAL Tier1

- Manages tape-backed data

- LHC experiments: 52 PB

- STFC Facilities: 58 PB

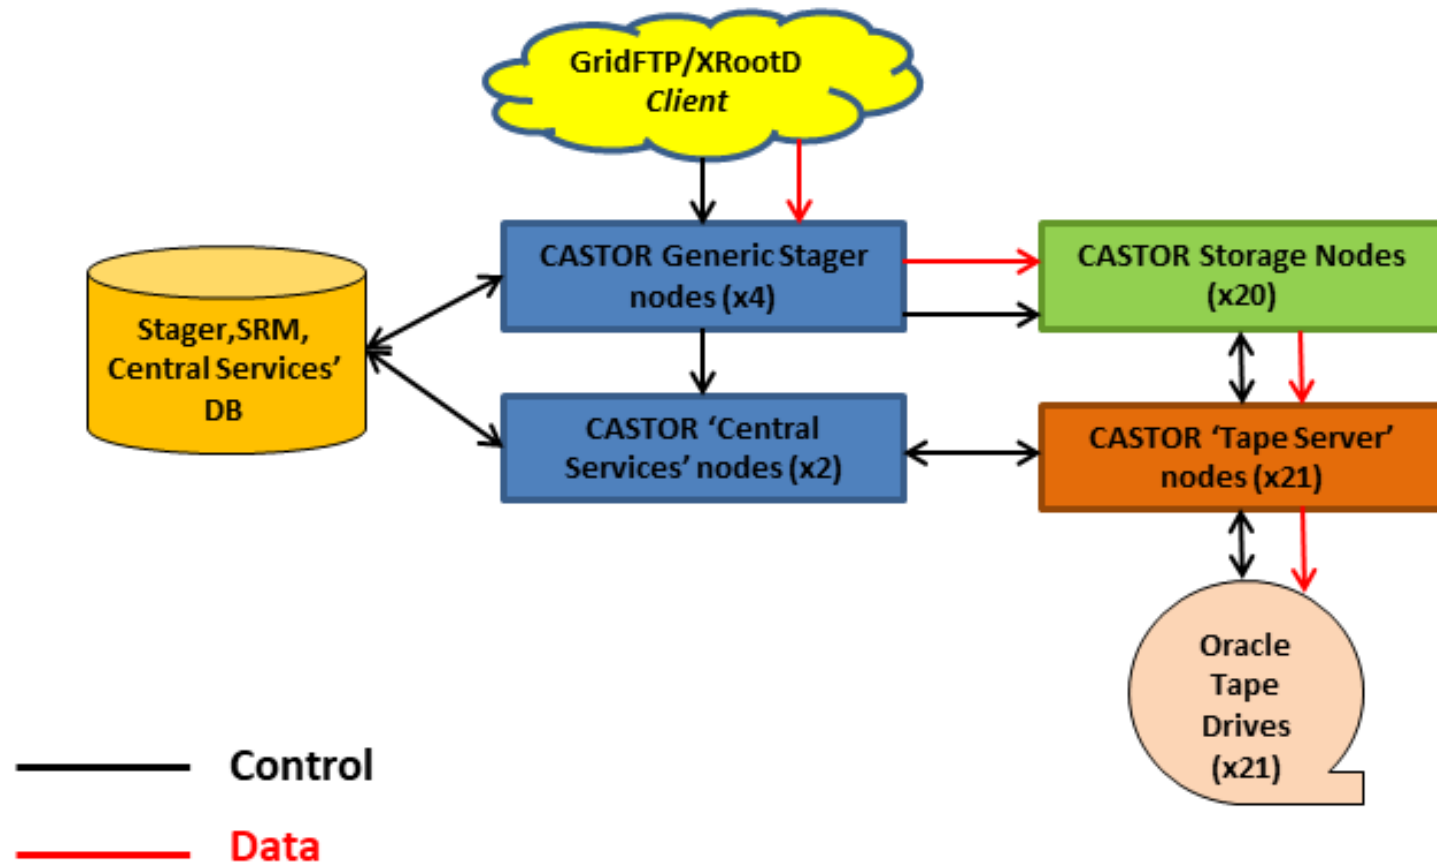


- Two major instances: CASTOR WLCGTape (Tier-1) and CASTOR Facilities (Diamond, CEDA, JASMIN)
- ALICE still have some disk-only data on CASTOR but they will migrate to WLCGTape by the end of 2019

# CASTOR WLCGTape setup

- Two sets of failure-tolerant management nodes (VMware)
  - 4 x stagers (external requests and staging of files to disk)
  - 2 x nameservers (namespace and tape services)
- A single back-end Oracle DB
- Single tape buffer of ~2PB (~~OCF2014~~/Dell2015/Dell2018)
- 21 tape servers connected to Oracle StorageTek T10KD tape drives
- All nodes run on SL7 and CASTOR 2.1.17-35
- Access through DNS round-robin aliases:  
**srm-atlas.gridpp.rl.ac.uk, ...**

# CASTOR WLCGTape (setup)



# StorageTek SL8500 tape library

Two libraries used by, **Tier-1** and **Facilities** run as a linked pair

- More than 70 drives installed: mix of T10KC/T10KD and media (5/8TB)
- 23 x T10KC/T10KD drives for Facilities
- 21 x T10KD drives for Tier-1 use (52PB)
- Total amount of data: 110 PB

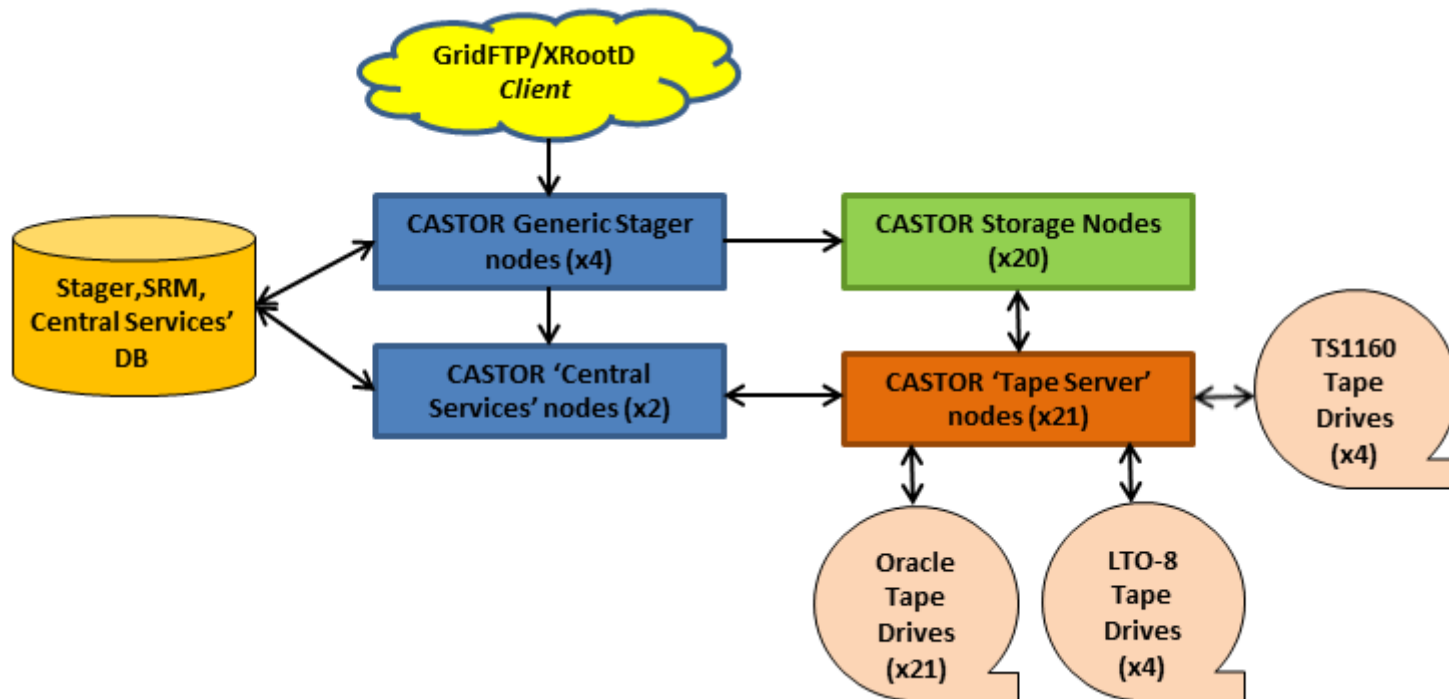


# SpectraTFinity tape library

- Mix of LTO8 and IBM TS1160 media and drives
- 33 drives installed: 17 x LTO and 16 x TS1160
- LTO media (LTOM7) hold 9TB/cartridge
- TS1160 media hold 20TB/cartridge
- 60PB of media has been installed (space for up to 24 PB)
- Tier-1 configured to use : 4 x TS1160 and 4 x LTO8 drives (DUNE)



# CASTOR WLCGTape (current set up)





# SpectraTFinity – functional testing

**Simple functional tests:** write 50 files via XRootD and GridFTP to LTO-8 and TS1160 media, recall and verify checksum

**CEDA functional tests:** write/read/verify files using rfiio and XRootD protocols

<b>CEDA functional test #1</b>	<b>Client /Tape medium</b>	<b>LTO-8</b>	<b>TS1160</b>	<b>Oracle</b>
	<b>rfio</b>	38792 (1745 MB)	38791 (1745 MB)	32264 (1452 MB)
	<b>XRootD</b>	1434 (1044 MB)	1439 (1062 MB)	1442 (1047 MB)
<b>CEDA functional test #2</b>	<b>XRootD</b>	11000 (1 TB)	2000 (240 GB)	2000 (240 GB)

# SpectraTFinity – performance testing

Four sets of performance tests (file recalls) were carried out for *LTO-8* and *TS1160* tape drives

- Files recalled in the order stored on tape
- Files were recalled in random order

**Single tape/  
40,000 files**

**Multiple tapes (4-5)/  
40-50,000 files**

# SpectraTFinity – performance testing

For every successful file transfer from tape to disk, CASTOR reports a *driveTransferSpeedMBps* value

A possible metric of tape drive performance would be the average of *driveTransferSpeedMBps* values for all files recalled from tape

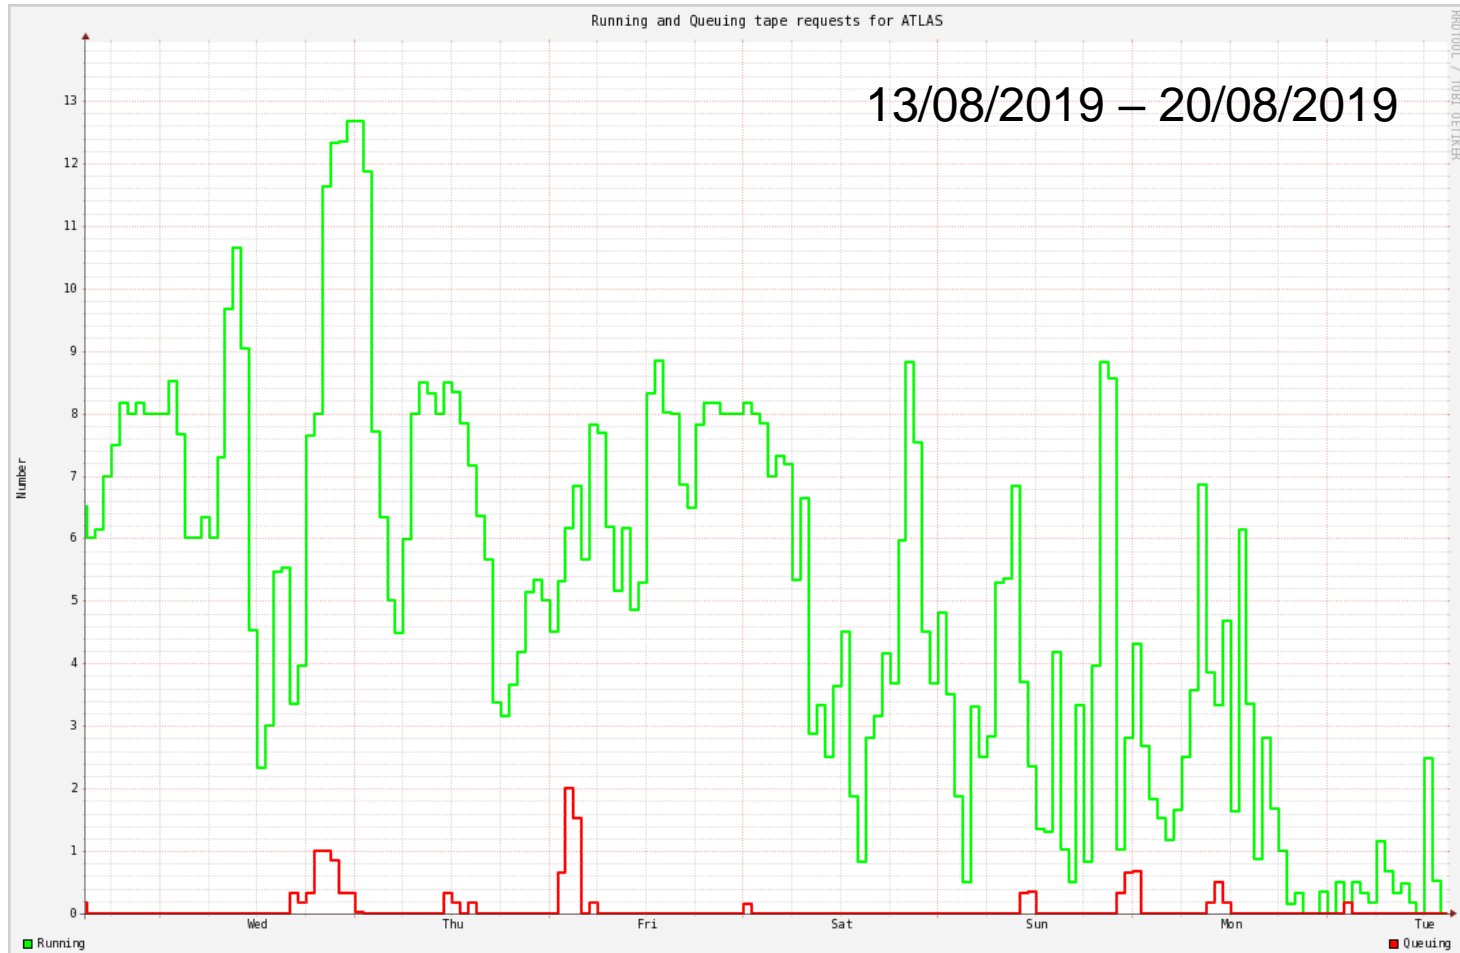
<b>Tape recall test</b>	<b><i>LTO-8</i></b>	<b><i>TS1160</i></b>
Single tape ordered read	291	309
Single tape random read	267	318
Multiple tape ordered read	258(±11)	301(±29)
Multiple tape random read	274(±16)	321(±31)

**Tape drive transfer speeds (MB/s) resulted from the Spectra TFinity recall tests**

# ATLAS tape carousel 2019

- Large (unannounced) reprocessing campaign
- Transferred 1019 TB between 08/08/2019 – 19/08/2019
  - 70% going to RAL Echo (disk-only)
  - the rest to Glasgow, Manchester, and QMUL
- Only Oracle StorageTek T10KD tape drives were used in RAL
- Initially 9 Tier-1s participated but due to an unscheduled downtime of CNAF, its share was reassigned to CERN

# ATLAS tape drive usage



# ATLAS tape carousel 2019

Tape	Total size	Avg. size per dataset	Avg. files per dataset	Avg. dataset staging time (days)	Avg. task executing time(days)
BNL-OSG2_DATATAPE	1.2 PB	26 TB	12000	15	22
FZK-LCG2_DATATAPE	0.8 PB	31 TB	14000	33	34
IN2P3-CC_DATATAPE	0.6 PB	29 TB	13000	15	27
NDGF-T1_DATATAPE	0.5 PB	30 TB	13000	21	24
PIC_DATATAPE	0.4 PB	34 TB	15000	22	27
RAL-LCG2_DATATAPE	1 PB	32 TB	14000	11	16
SARA-MATRIX_DATATAPE	0.7 PB	31 TB	14000	13	30(?)
TRIUMF-LCG2_DATATAPE	0.6 PB	32 TB	15000	18	25



# Future of CASTOR

- Plans to replace CASTOR before LHC Run 3
- Different tape storage options are considered
- Procure a new Tier-1 tape robot



Science & Technology  
Facilities Council

UK Research  
and Innovation

**Thank you!**