Science & Technology
Facilities Council

Tier 1 CASTOR Consolidation

What goes up, must come down…

Rob Appleyard

Part 1

Introduction

CASTOR

- Currently 13 PB disk storage
  - D1T0 for ATLAS, LHCb, CMS
- 36 PB tape storage
- CERN:
  - 240PB of data on tape
  - 16PB of disk, all cache
- CERN (CASTOR developers) now only use CASTOR for tape-backed storage

*http://castor.web.cern.ch/

Science & Technology Facilities Council

Echo

- Replacing CASTOR D1T0
- Currently 9PB of usable disk storage.
- ATLAS migration well underway
- CMS migration starting
- LHCb a bit later

Science & Technology Facilities Council

Tier 1 Data Flow Now

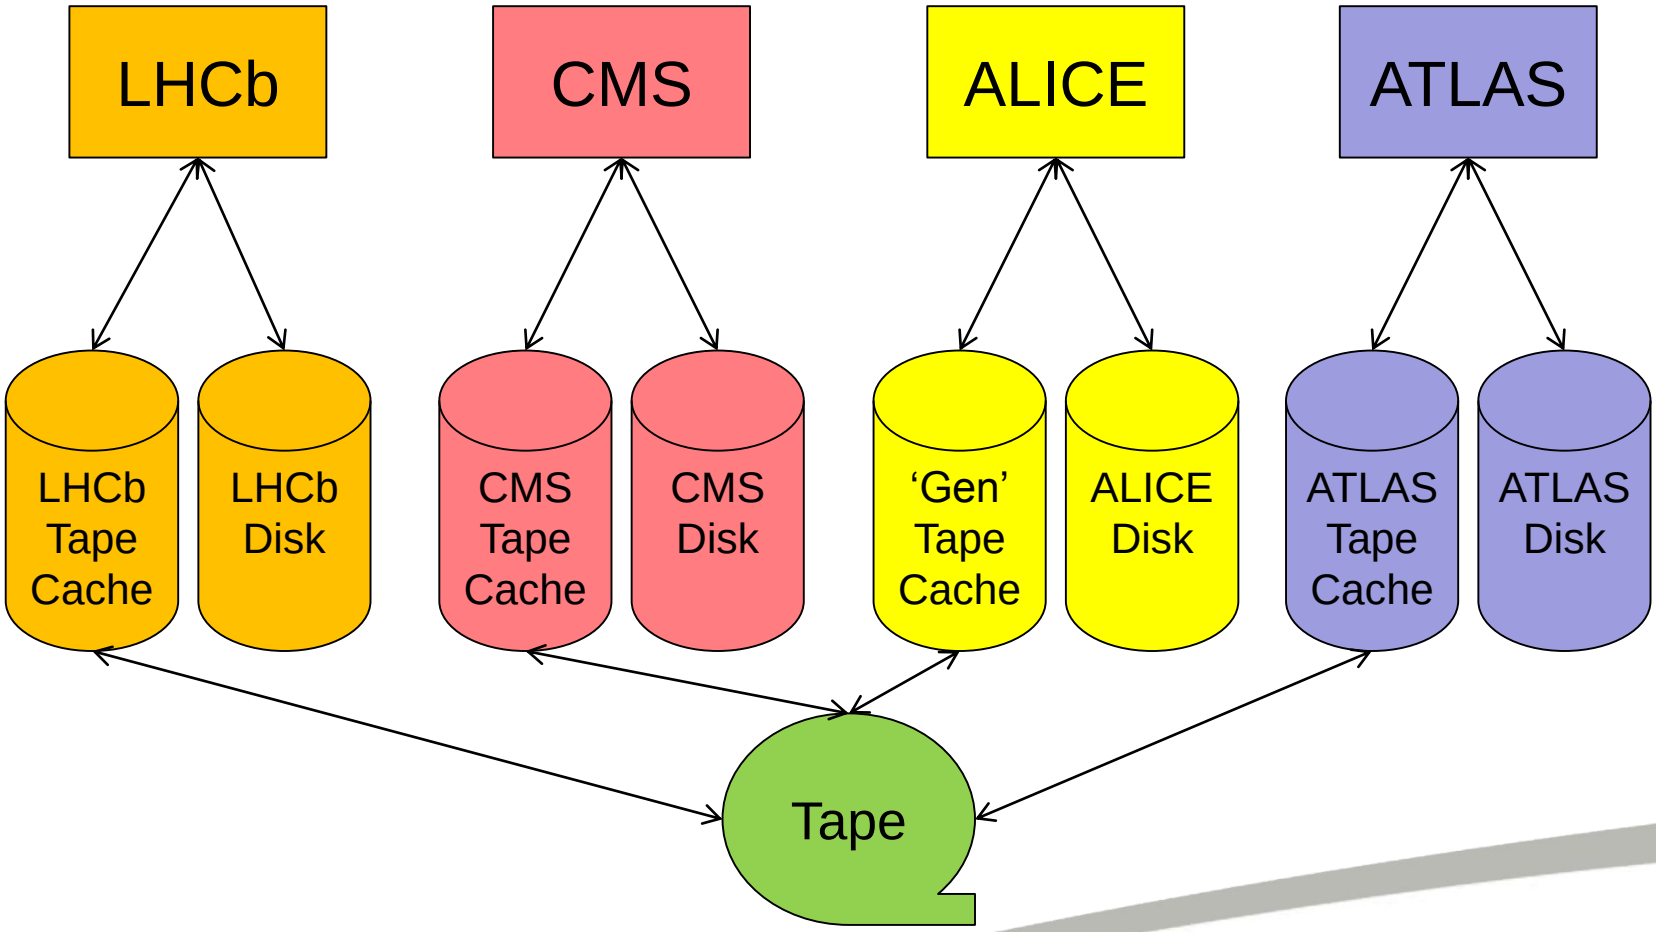

Science & Technology Facilities Council

CASTOR Databases

- *Everything* in CASTOR is based on Oracle DBs.
  - Physical data location
  - Transaction information
  - Namespace mapping
  - Tape drive state
  - Transfer scheduling

Database Groupings

- ‘Central services’ DB
  - One DB instance for all WLCG users
  - Manages namespace
  - Manages tape interface & contents of tapes
- ‘Stager’ DB
  - Manages data residing on disk
  - One DB instance per major user community
  - ‘Instance’ – one stager DB schema
- SRM DB
  - Collocated with stager

Part 2

What we have now

Science & Technology
Facilities Council

A Picture of CASTOR – GridFTP (ATLAS)

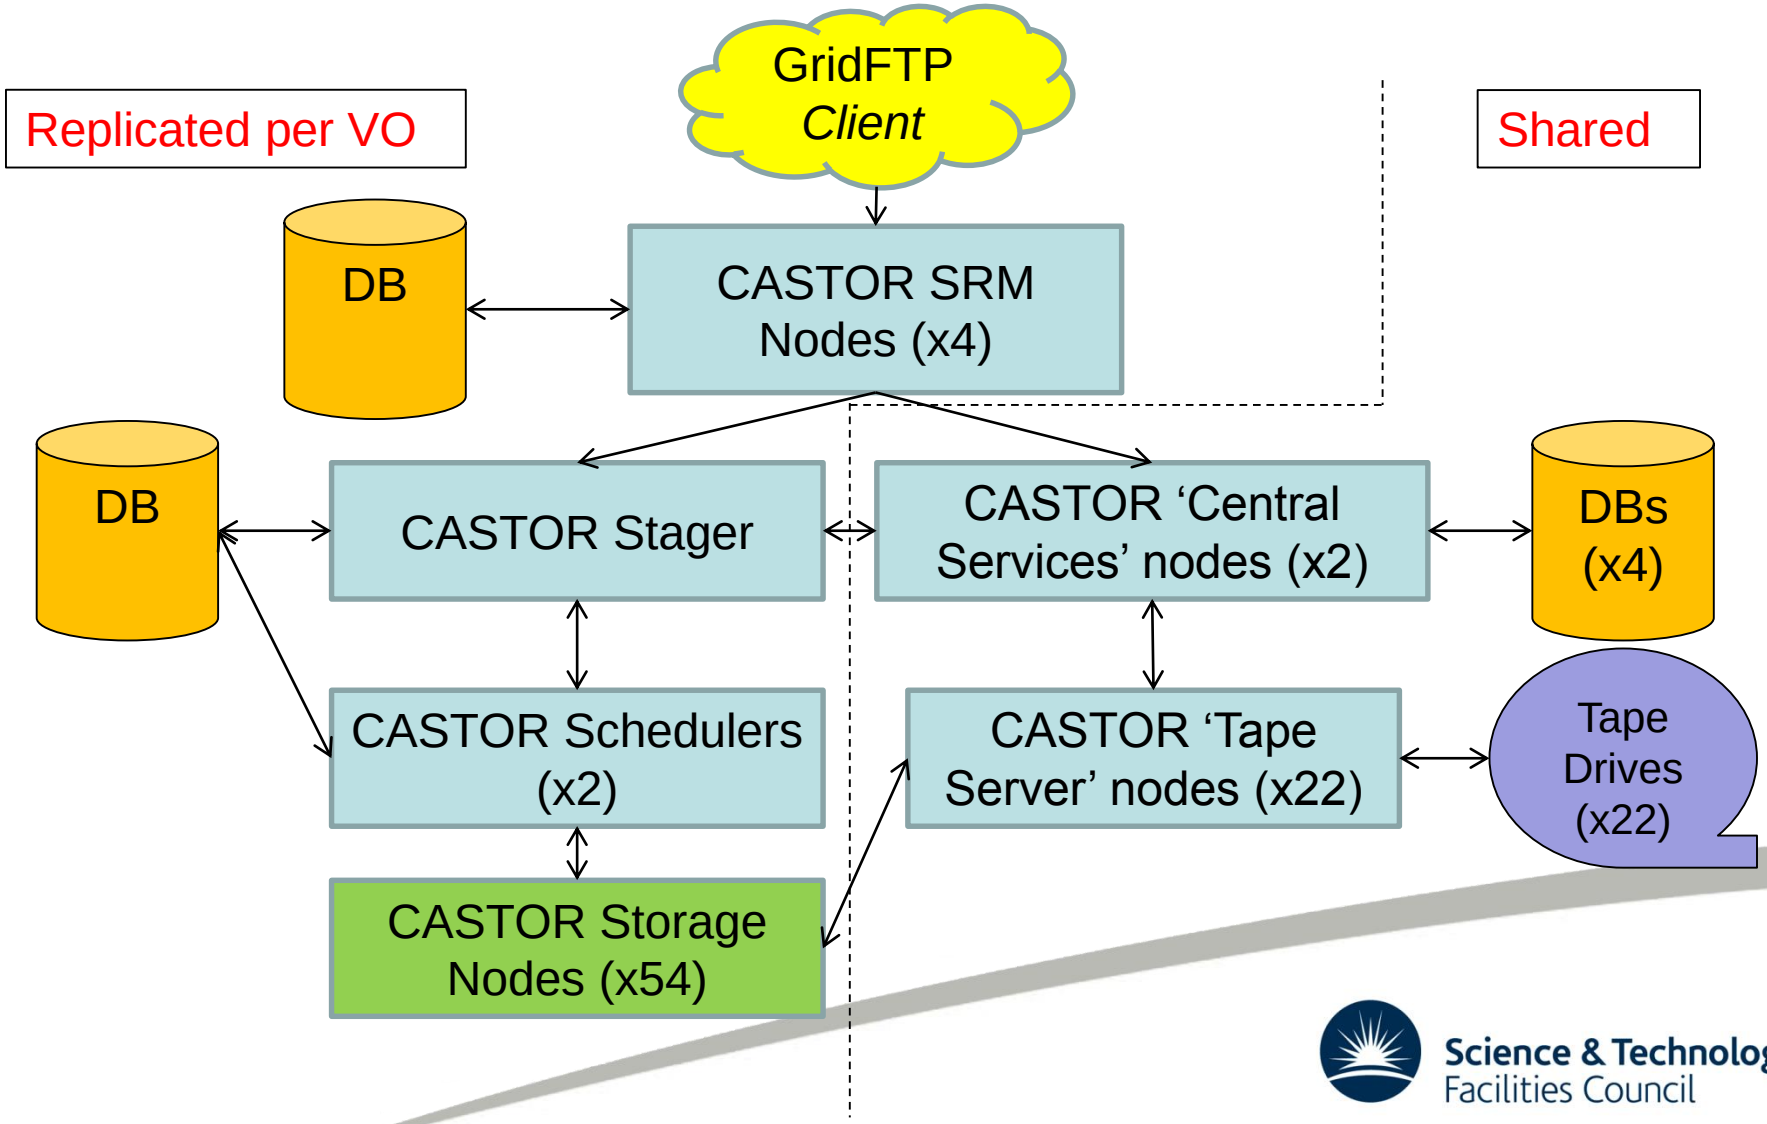

Science & Technology Facilities Council

CASTOR Current State: Databases

- Two Oracle RACs are used to support CASTOR operations
  - One hosts ATLAS and our ALICE/general-use instance
  - The other hosts CMS, LHCb, and the central services
- Transaction rate: 390hz/RAC
- Load is strongly driven by disk-only operations

CASTOR Current State: Management Nodes

- AKA ‘headnodes’
- Each instance has 3 dedicated management nodes headnodes, and 2-4 dedicated SRM interface nodes
  - Interface nodes handle control traffic only
- Plus two ‘name servers’ for the central services
  - Including tape system
- Grand total of 25 core management nodes
- Management nodes are currently ‘pets’, not ‘cattle’
  - One management node failure -> service offline ☹

CASTOR Current State: Storage nodes (1)

- Aka ‘disk servers’
- 137 nodes
- Each node is 60-120 TB
  - One big RAID 6 array
- 10Gb networking
- Peak i/o performance typically ~3Gb/s/node
  - Constrained by disk i/o

Science & Technology Facilities Council

CASTOR Current State: Storage nodes (2)

- 29 of those storage nodes are used only for tape-backed storage
  - Caching data on its way to/from tape
- Remaining 108 are disk-only
  - Will be retired when migration to Echo is complete

What if we do nothing?

- Disk server count drops from 137 to ~30
  - Transaction rate drops to ~5% of current (or lower)
- But we still have…
  - 29 management nodes
  - 2 RACs
  - Management nodes outnumber storage!
- Unacceptable management overhead

Part 3

What we are going to do

Project Objectives

- Reduce node count
- Reduce management overhead
- Improve service quality
- Don’t lose any data!

User migration

- Users responsible for their own data management
  - LHC VOs well aware of need to migrate
- ATLAS: good progress at drawing down CASTOR disk
  - Production use of Echo
- CMS also using Echo in early stage production
- LHCb lagging a bit, work ongoing
- Once user says ‘all clear from CASTOR’, we can clean up any remaining data
  - There is always some

‘Tier1Tape’

- Plan: create a fresh CASTOR instance using new DB hardware
- Once a user has nothing on disk…
  - Suspend new tape writes
  - Flush remaining migration candidates to tape
  - Repoint aliases to new instance
  - Restart
- Downtime for tape system ~ few hours
- Users can be migrated one-by-one

Post-Echo CASTOR Data Flow

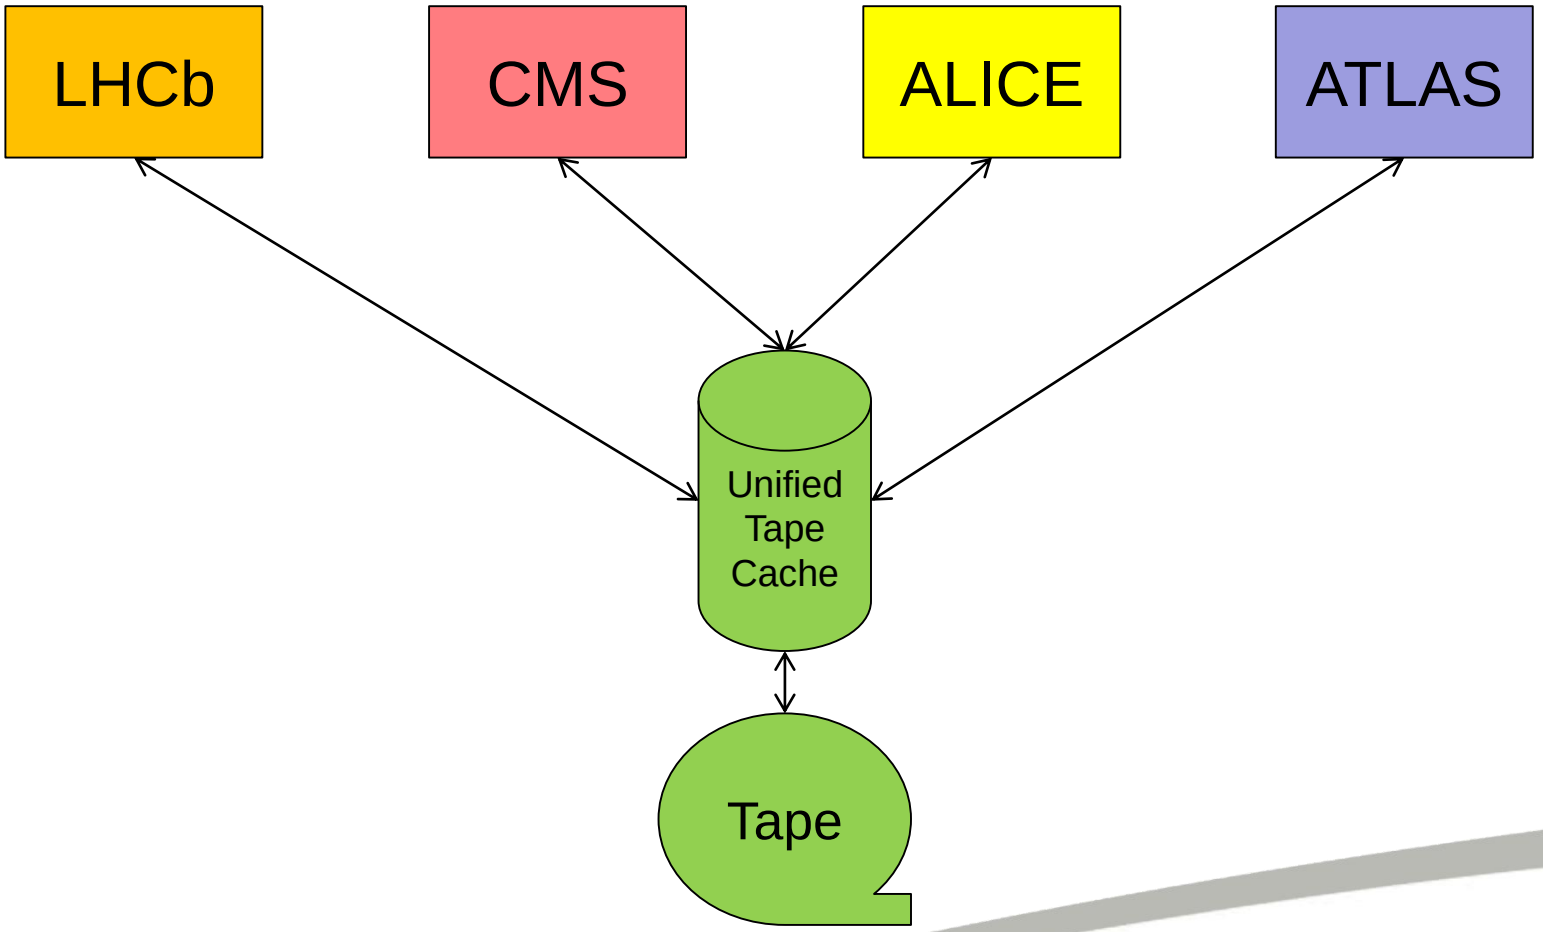

Science & Technology Facilities Council

Issues: Contention

- Potential for contention between users introduced into system
  - Disk cache needs to be relatively big to mitigate this
- Issue already in play for other system elements
  - Tape drives are a shared resource for all users
- Partitioning of cache is possible…
  - …but not desirable

Issues: Scheduling Interventions

- Advantage of separate infrastructure for each user community: easy intervention scheduling
  - Not present when everyone shares ☹
    - Need to find a date that suits everyone
- Difficult to mitigate
  - Saving grace: WLCG Tape access is usually orderly and planned
  - Able to plan with experiment data admins

Other Improvements: Management Nodes

- Change of structure is an opportunity to address other issues
- CERN CASTOR implementation uses ‘cattle headnodes’
  - All management processes run on a set of identical nodes
  - HAProxy for failure-tolerence
  - George will be replicating these
- Shift to from physical to virtualized infrastructure
- Also push to SL7

Other Improvements: RAO

- New available feature in CASTOR 2.1.17: Recommended Access Order for recalls
- Tape drive gets many recall requests, figures out sensible order to minimise seek time
- 40%-60% improvement in seek time for large reads (says vendor)
  - Seen even better numbers than this from CERN…

CASTOR's Long-Term Future

- CERN CASTOR service is scheduled to be discontinued ~ mid 2019
  - New product: 'CTA[1]'
  - No more development effort from CERN
- …but…
  - Migrating away from CASTOR will take time
  - Improvements have time to bear fruit
- See other talks/discussions for more detail on our plans

1: **An efficient, modular and simple tape archiving solution for LHC Run-3, S Murray, et al** http://iopscience.iop.org/article/10.1088/1742-6596/898/6/062013/pdf

Science & Technology Facilities Council

Any Questions?

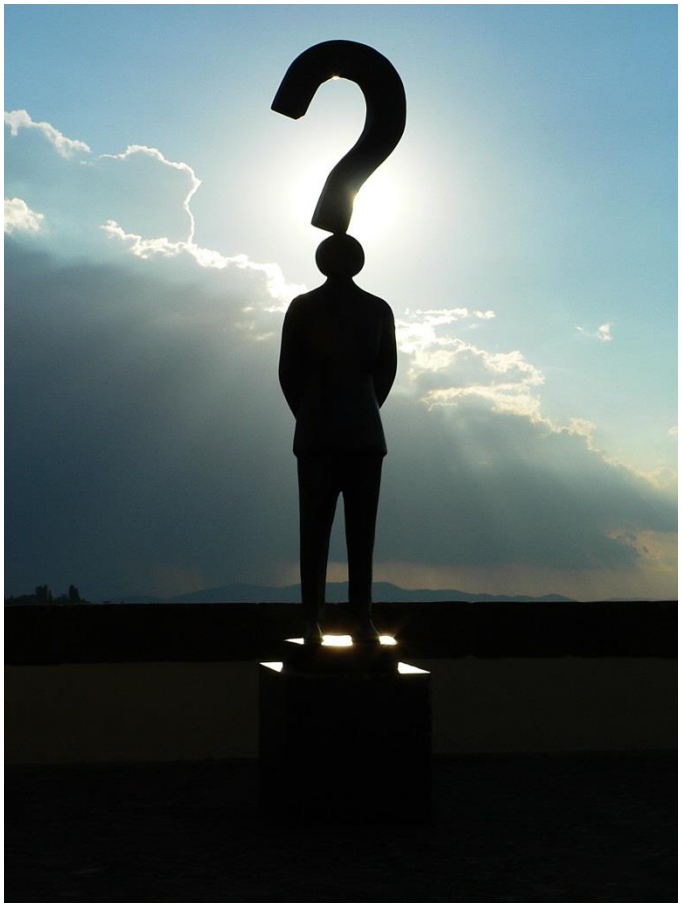

Image by Marco Belluci, distributed under CC 2.0 license.
https://www.flickr.com/photos/marcobellucci/3534516458

Science & Technology Facilities Council