

Highlights from CHEP 2018



Alastair Dewhurst

Introduction

- In April I moved to my new jobs as Tier-I manager.
 - I therefore focus more on the infrastructure side of things.
- Highlights:
 - Tape Evolution
 - Commercial Clouds
 - DynaFed + Large Scale CVMFS (My Talk!)



RAL Tape Service



Hardware



1 x SL8500
Tape Library



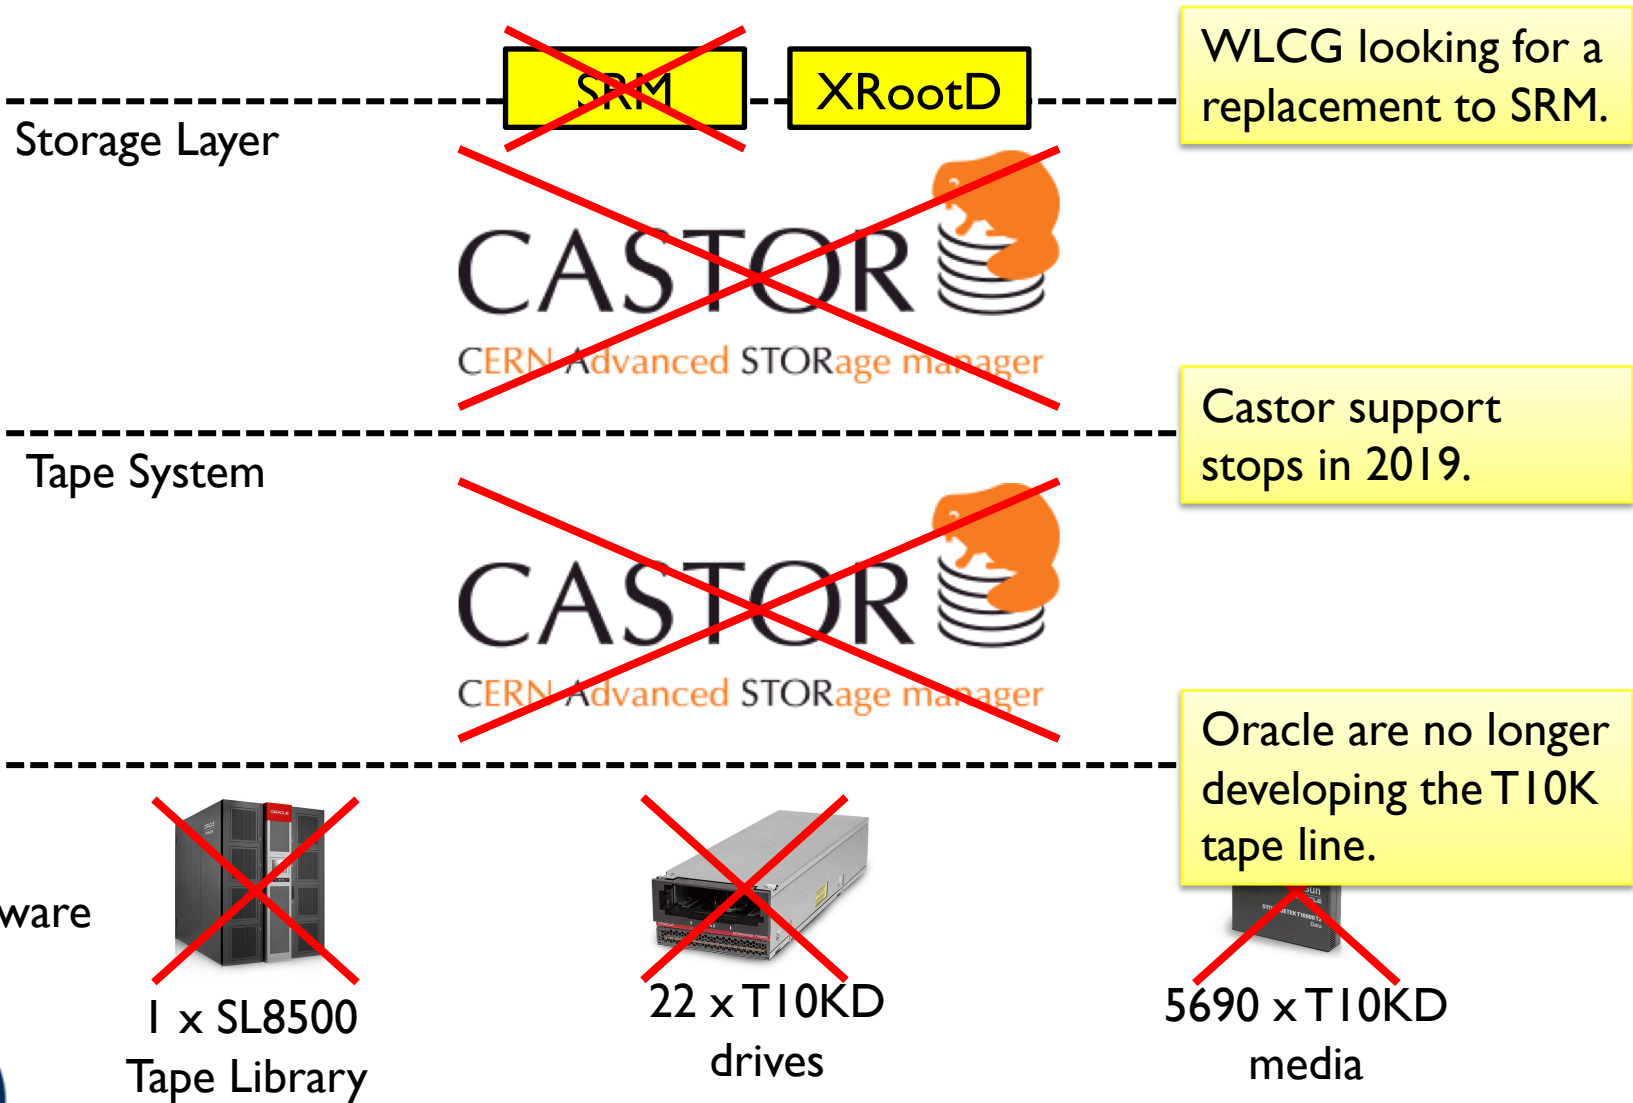
22 x T10KD
drives



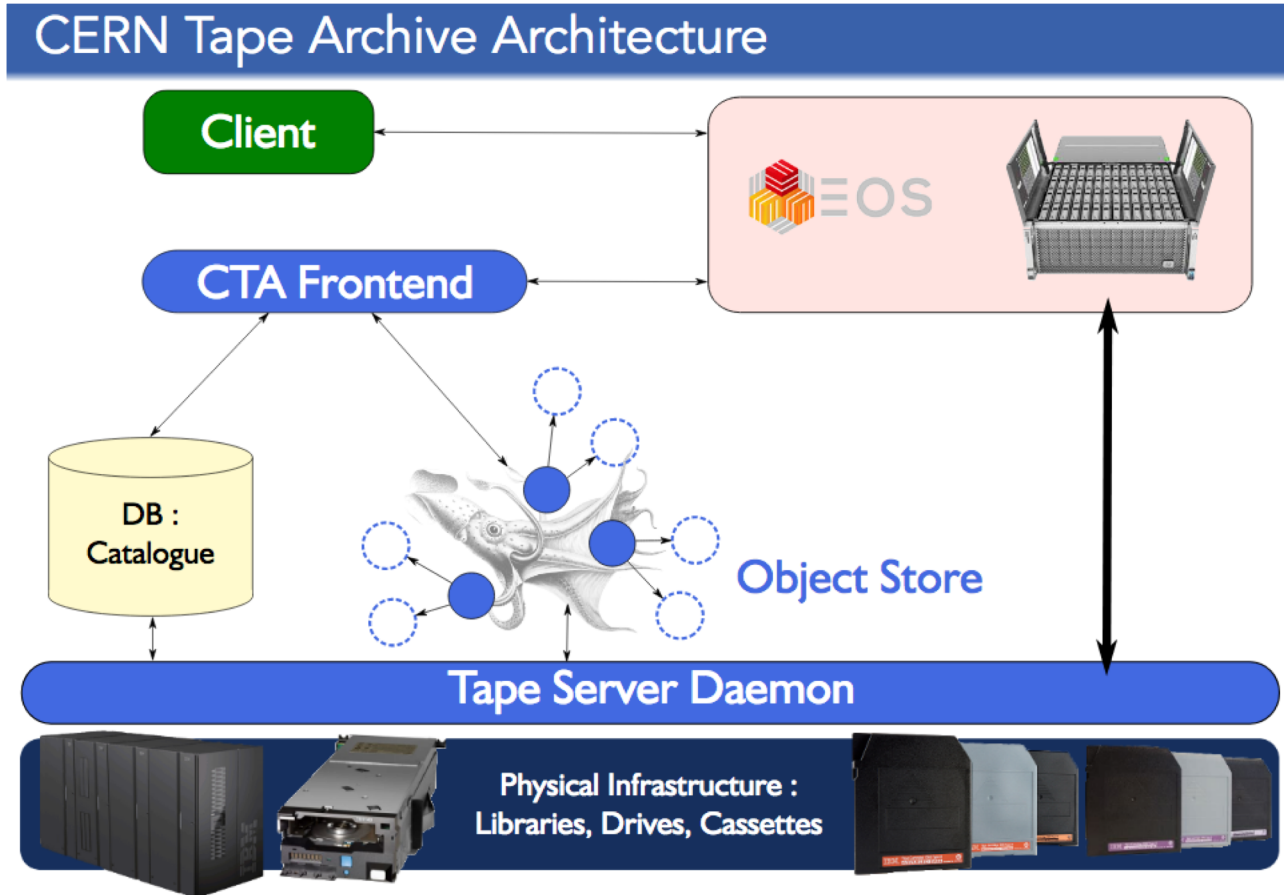
5690 x T10KD
media



Tape Outlook



CTA



KIT Tape environment

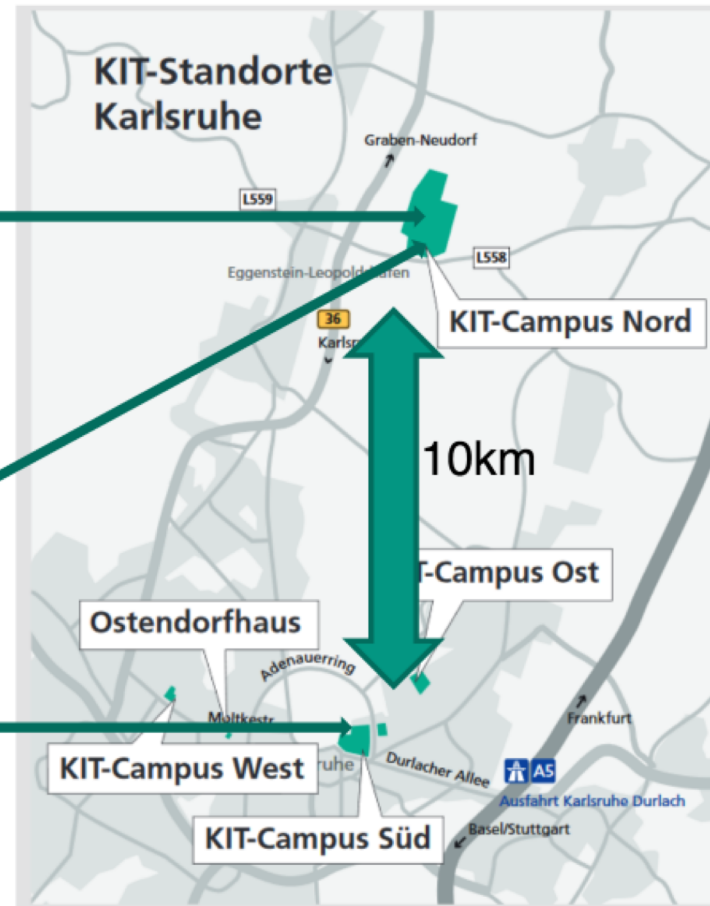
3 SL8500 Libraries
SCC building



2 TS3500 Libraries
OKD building

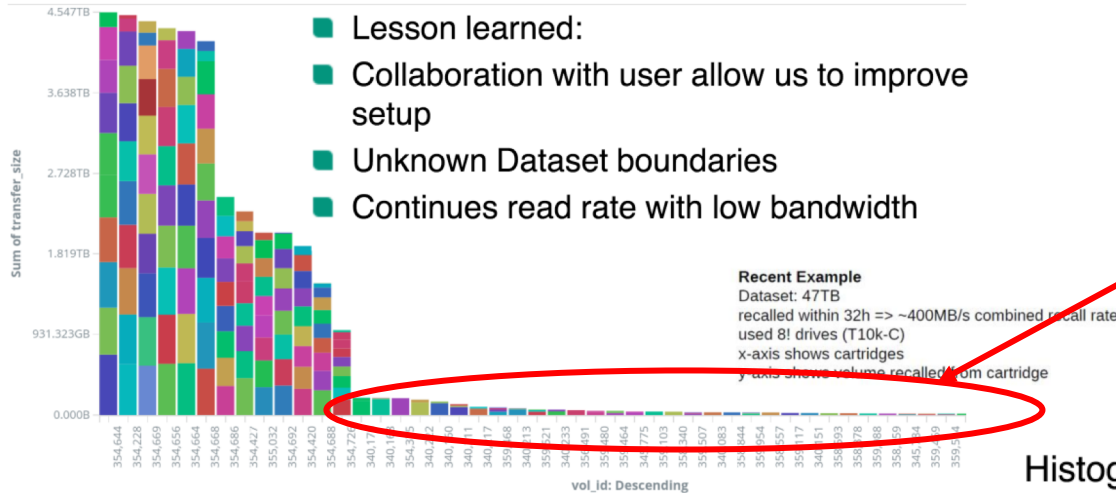


1 TS3500 Library
KIT Campus South



KIT improvements

ATLAS recall example of one dataset



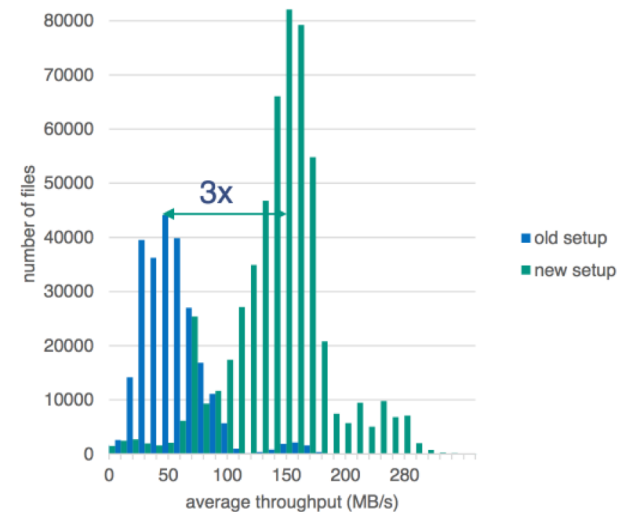
Long tail of recalls that require only a few files from each tape.

Tape families are important

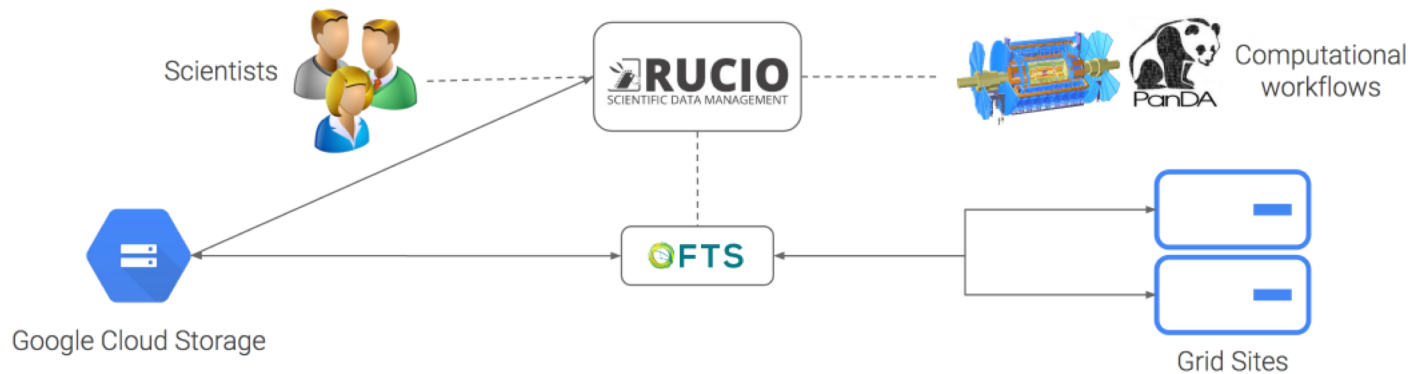
KIT added in proxy nodes between their tape drives and storage services.

Dramatic improvement in both stability and performance!

Histogram of Recall Rates per File




Google Data Ocean




- Terabyte transfer test

- Created rules to transfer 1 TB of user analysis data
- 1TB each to both US and EU Google Data Centres
- Worked off at 0.6TB/h aggregate
- Maxing out FTS intermediate stream



Google Compute

- Operated a 120 core cluster running standard **simulation jobs** for 1.5 months
 - I/O to CERN storage
 - Excellent success rate (<<5% errors) using normal VMs
- Preemptible VMs
 - Significantly higher error rate (20-30%)
 - Still gain on a \$/event basis
- **Analysis queue** ramping up
 - I/O intensive workloads reading from GCS



Efficiency of preemptible VMs can be optimized through usage of Event Service.



CVMFS

- CVMFS has been primarily developed for distributing large software stacks.
 - Provides read only POSIX access.
 - Extremely effective and popular.
- Large Scale CVMFS is an extension to the base software which allows it to distribute large, non-public datasets.
 - In use for LIGO since 2016.
 - I have been trying to improve it by integrating it with DynaFed.



CVMFS






A VO admin uploads files to the Stratum 0.

The SHA1 checksum is calculated and added to the catalog file.

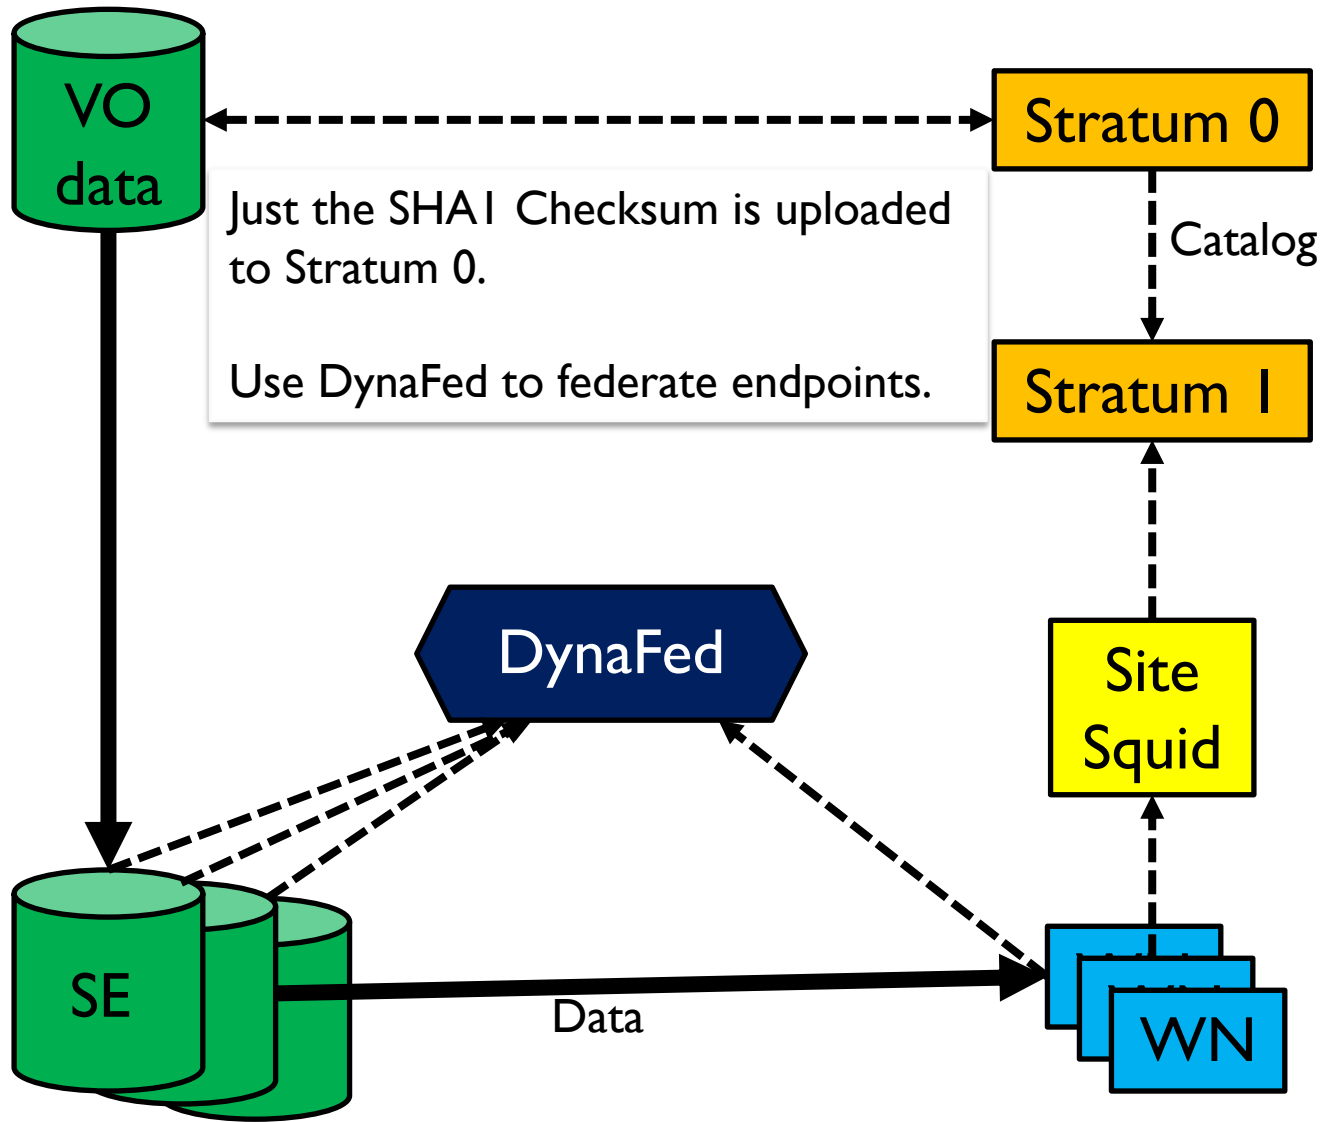
A copy of all the files are replicated to all the Stratum 1s

There is a cache on each WN as well as site squids.

The CVMFS client can be provided with a list of squids which it will try in turn to connect to if the file is not cached locally.



Large Scale CVMFS



Questions?



References

- Tape:
 - <https://indico.cern.ch/event/587955/contributions/2936881/>
 - <https://indico.cern.ch/event/587955/contributions/2936822/>
- Google Data Ocean:
 - <https://indico.cern.ch/event/587955/contributions/2947395/>
- LS CVMFS:
 - <https://indico.cern.ch/event/587955/contributions/2936825/>

