

CTA: CERN Tape Archive

Future of archival storage at CERN

Vladimir Bahyl, Germán Cancio, Eric Cano, Michael Davis, Julien Leduc,
Cristina Moraru, Steven Murray

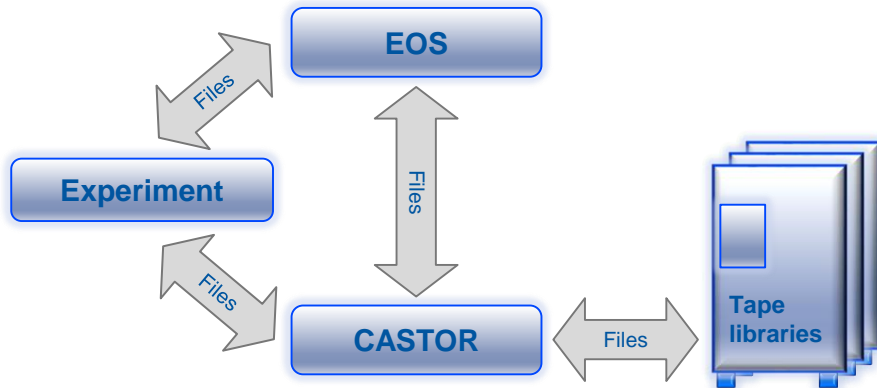
Overview

- What is EOS+CTA?
- Making the most out of tape
- Status and Timeline

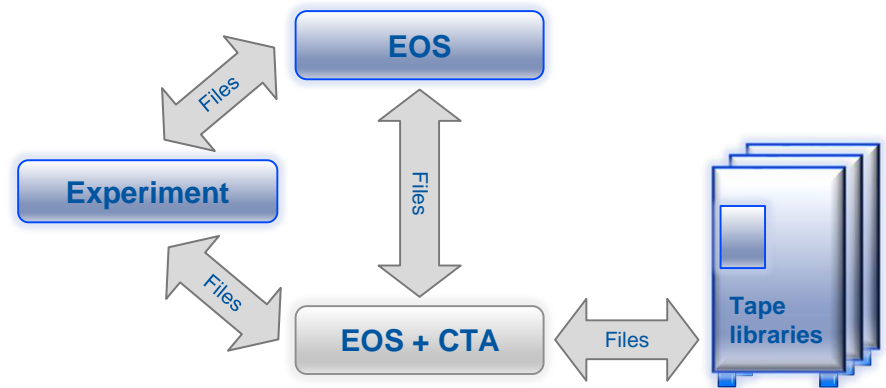
What is EOS+CTA?

- EOS plus CTA is a “drop in” replacement for CASTOR

Current deployments with CASTOR



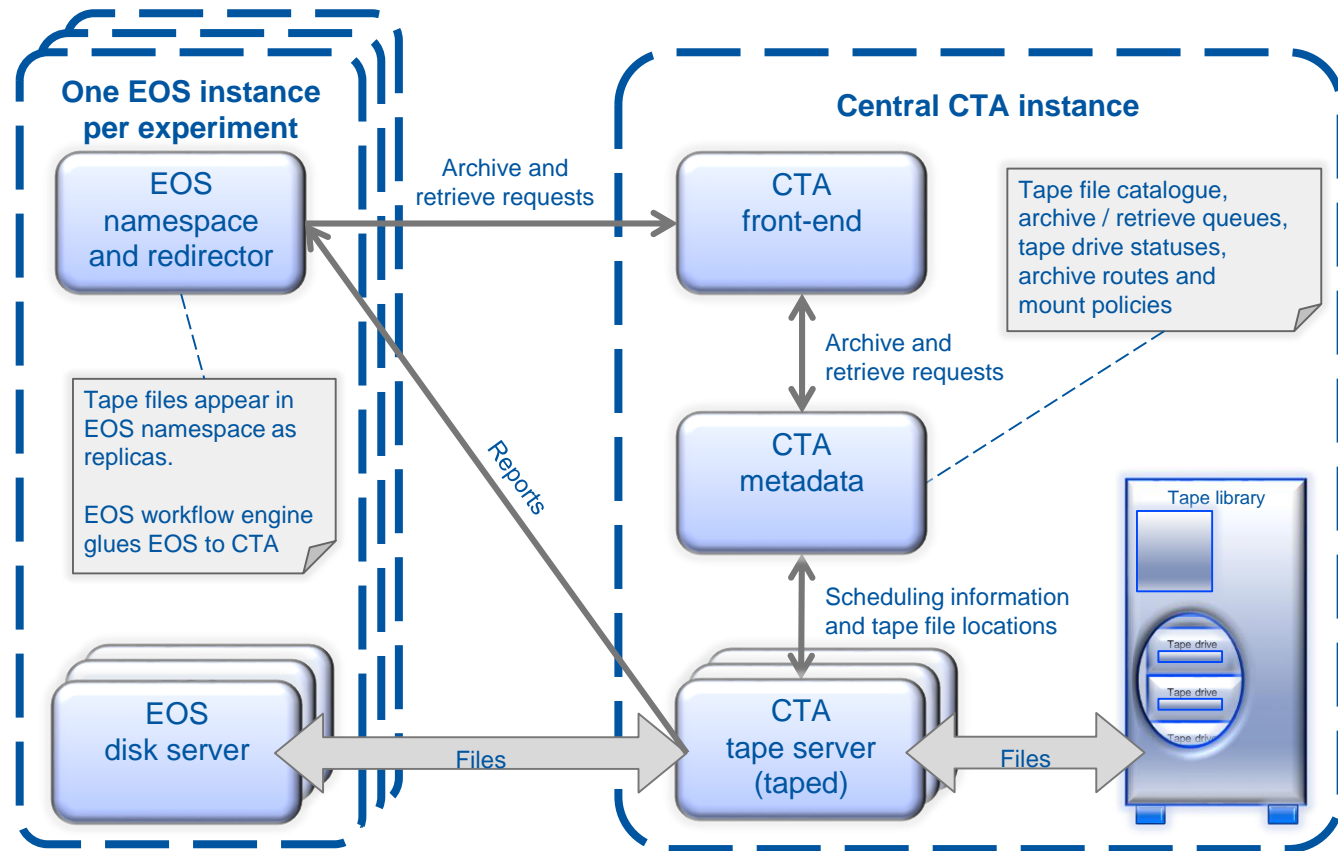
Future deployments with EOS plus CTA



EOS+CTA architecture

CTA is:

- Natural evolution of CASTOR
- A tape backend for EOS
- A preemptive tape drive scheduler
- A clean separation between disk and tape
- Xrootd based recalls (not SRM)



Making the most out of tape

- Single step mount scheduling for tape+drive
 - Avoids choosing tapes that no drive will serve when writing => improves write mount latency
- Mount preemption
 - Replace handcrafted/scripted drive dedications
 - Allow full utilization of drives all the time
 - Background repack/verify user will yield to user mounts
 - Many possible combinations (priority based)
- Recommended access order
 - Optimizes read order to minimize seek time
 - 4x – 5x speedups for a typical O(100s) files recall mount
 - Backport to CASTOR coming soon
 - ... but enterprise drives only

Currently achieved

- Preliminary EOS-CTA integration
- Archive and retrieve files to/from EOS
 - Request queueing
- Decentralized single step scheduling
- Administrative interface
 - Drive control
 - Tape management
 - File routing to tape
- Recommended access order (RAO) prototype
- All drive features from CASTOR

Currently achieved

```
[root@tpsrv600 ~]# cta drive ls
library      drive      host      desired  request  status  since  vid  tapepool  files  MBytes  MB/s  session  age
IBM1JB      I1JB0402  tpsrv033  Up       -        Free   4722  -    -         -     -      -      -       31
```

```
[root@tpsrv600 ~]# cta tapepool ls -h
name # partial tapes encrypt c.user c.host c.time m.user m.host
supply_IBM1JB 1 false vlado tpsrv600.cern.ch Mon Jun 12 11:10:40 2017 vlado tpsrv600.cern.ch Mon J
supply_tpsrv601 1 false vlado tpsrv600.cern.ch Wed Aug 30 16:14:58 2017 vlado tpsrv600.cern.ch Wed A
supply_tpsrv602 1 false vlado tpsrv600.cern.ch Wed Aug 30 16:14:58 2017 vlado tpsrv600.cern.ch Wed A
supply_tpsrv603 1 false vlado tpsrv600.cern.ch Wed Aug 30 16:14:58 2017 vlado tpsrv600.cern.ch Wed A
supply_tpsrv604 1 false vlado tpsrv600.cern.ch Wed Aug 30 16:14:58 2017 vlado tpsrv600.cern.ch Wed A
supply_tpsrv605 1 false vlado tpsrv600.cern.ch Wed Aug 30 16:14:58 2017 vlado tpsrv600.cern.ch Wed A
supply_tpsrv606 1 false vlado tpsrv600.cern.ch Wed Aug 30 16:14:58 2017 vlado tpsrv600.cern.ch Wed A
supply_tpsrv607 1 false vlado tpsrv600.cern.ch Wed Aug 30 16:14:58 2017 vlado tpsrv600.cern.ch Wed A
supply_tpsrv608 1 false vlado tpsrv600.cern.ch Wed Aug 30 16:14:59 2017 vlado tpsrv600.cern.ch Wed A
supply_tpsrv609 1 false vlado tpsrv600.cern.ch Wed Aug 30 16:14:59 2017 vlado tpsrv600.cern.ch Wed A
supply_tpsrv610 1 false vlado tpsrv600.cern.ch Wed Aug 30 16:14:59 2017 vlado tpsrv600.cern.ch Wed A
test_gcancio 1 false gcancio tpsrv600.cern.ch Tue Jun 13 13:47:31 2017 gcancio tpsrv600.cern.ch Tue J
test_jleduc 1 false jleduc tpsrv600.cern.ch Wed Jul 19 15:54:01 2017 jleduc tpsrv600.cern.ch Wed J
test_steve 1 false vlado tpsrv600.cern.ch Mon Jun 12 11:09:30 2017 vlado tpsrv600.cern.ch Mon J
test_unit_513 1 false vlado tpsrv035.cern.ch Fri May 26 16:01:40 2017 vlado tpsrv035.cern.ch Fri M
test_vlado 1 false vlado tpsrv035.cern.ch Mon May 29 16:38:07 2017 vlado tpsrv035.cern.ch Mon M
```

```
tpsrv602 L602D1 tpsrv602 Up - Free 4474 - - - - - 1
tpsrv602 L602D2 tpsrv602 Up - Free 4474 - - - - - 5
tpsrv602 L602D3 tpsrv602 Up - Free 4363 - - - - - 20
tpsrv602 L602D4 tpsrv602 Up - Free 4478 - - - - - 24
```

All drive features from CASTER

Upcoming features

- Streamlining and extension of EOS-CTA interface (from scripts to RPC)
- Object store performance optimization
- Repack & verification
- Mount preemption
- Library support rework (currently using the one of CASTOR)
- EOS-CTA catalogue synchronizations
 - Reconciliation & tape based disaster recover

Timeline

