



CTA Outlook

Oliver Keeble on behalf of the CTA team

Pre-GDB 9th Feb 2021

The CERN Tape Archive (“CTA”)

- **CTA is CERN’s response to the challenge of archiving LHC data during Run 3 and beyond**
 - 150PB/yr, 40GB/s
 - O(1EB) data, O(100k) tapes, O(1 billion) files, ~200 tape drives
- **Combines the best of Castor tape with CERN’s disk system, EOS.**



The CERN Tape Archive (“CTA”)

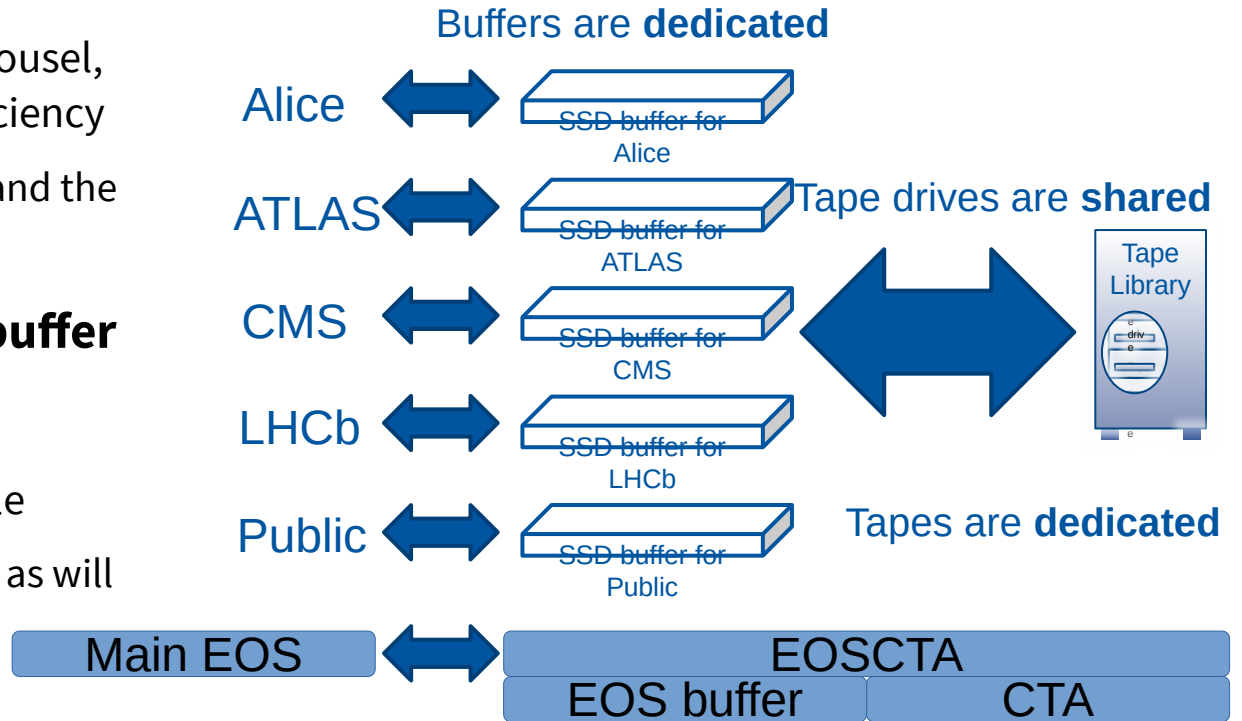


- **Designed to exploit modern tape hardware to the max**

- Anticipated use cases, e.g. Data Carousel, require greater throughput and efficiency
 - This is handled in both the design and the deployment model

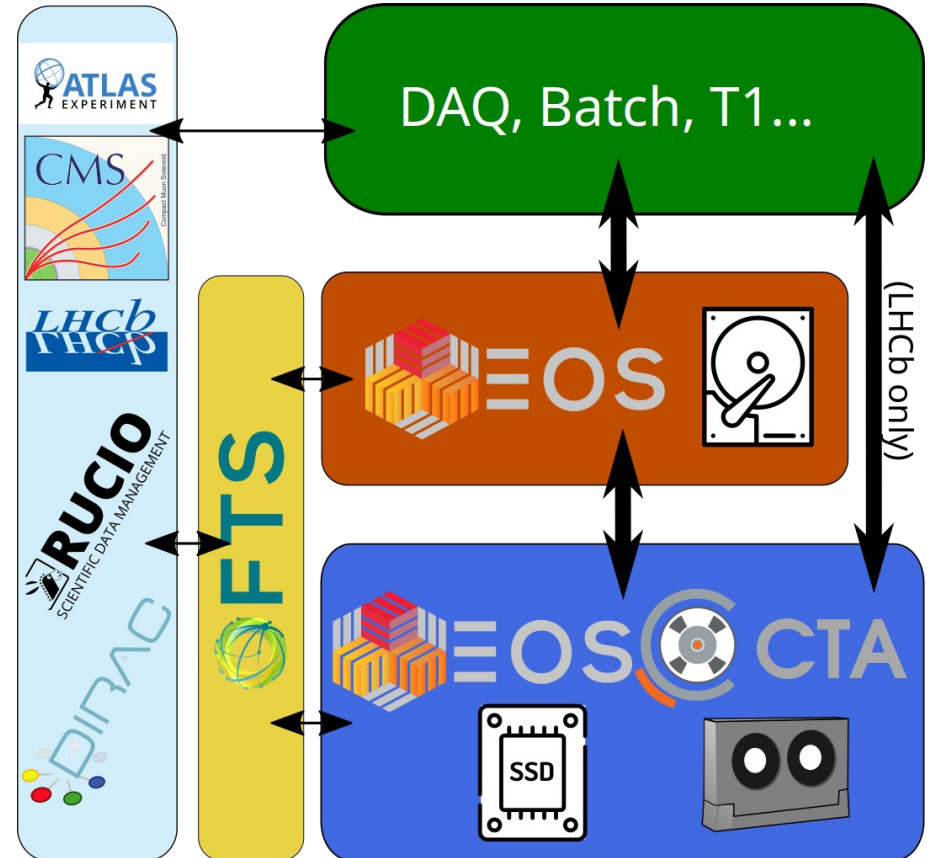
- **CTA’s SSDs are conceived as a buffer**

- Data residency is transitional
- A read cache is nevertheless possible
 - Alice is running with a recall cache, as will some smaller experiments



How is CTA deployed and used?

- **CTA is a tape system and needs a disk management system**
 - CERN will use EOS → “EOSCTA”
 - EOS provides the interface, file operations and namespace
 - CTA provides high performance tape operations based experience from Castor
 - Other disk systems can in principle be used
- **“EOSCTA” at CERN is part of a broader integrated stack which includes**
 - Main EOS disk instance
 - FTS
 - FTS clears recall buffer after transfers
 - Experiment data orchestration



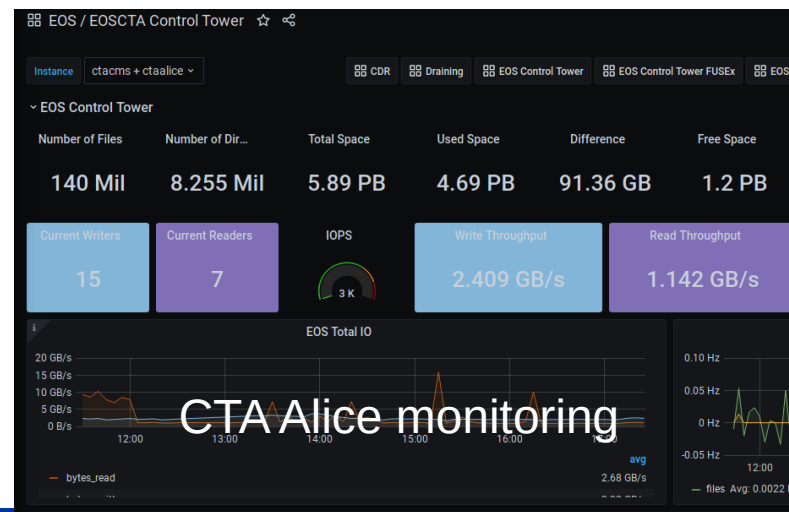
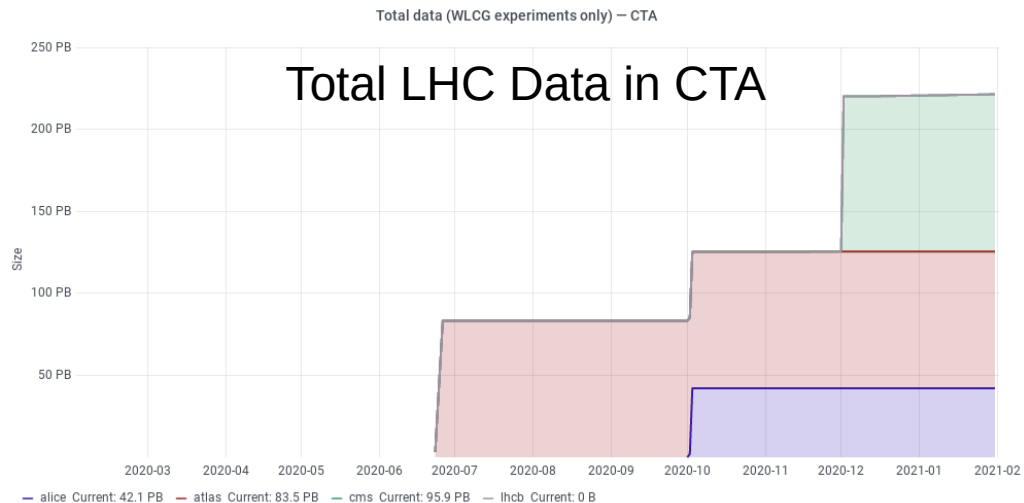
HDD icon: <https://commons.wikimedia.org/wiki/File:Hard-drive.svg>

SSD icon: <https://commons.wikimedia.org/wiki/File:Ssd.svg>

Tape icon: https://commons.wikimedia.org/wiki/File:Tape_cinta_cassette_backup.svg

CTA deployment status

- **CTA has deployed production instances for Alice, Atlas and CMS**
 - All their Castor data is now managed by CTA
 - The Castor endpoints have been retired
- **One other experiment (CAST) is data-taking**
- **Babar data preservation**



Some CTA highlights

- **Postgres**

- CTA has entered production using an Oracle db
- Both Postgres and MySQL backends also exist
 - We're not using Oracle tricks, PL/SQL etc
 - Longer term, CERN is targeting Postgres
- We recommend all adopters of CTA to use Postgres if possible

- **Unified scheduling (cf Castor)**

- **RAO**

- LTO does not have “recommended access order”
 - An optimisation for “random” access of files on a tape
- Support in software has been implemented and deployed in the CTA tape server
- Substantial improvements
 - → Vladimir Bahyl's talk

Pre-Run 3 outlook

- **Migrate LHCb from Castor**
- **Substantial progress on migration of smaller experiments**
 - NA62 leading the way
 - NA61, Compass, AMS ...
- **HTTP**
 - Deploy xrd-http
 - HTTP access
 - Third Party Copy
 - Token auth
- **XrootD v5**
- **Deploy Run3 hardware**
 - Pack 4 libraries
 - 3xIBM + 1xSpectraLogic
 - New tapeservers
 - New drives 1160 and LTO-9
 - Full SSD buffer
 - → Vladimir Bahyl will give more details at HEPiX
- **Ramp up infrastructure and perform Run-3 stress tests & data challenges**
 - Concentrate on performance and reliability engineering this year
 - Feed ops experience back into the tools

Community

- **CTA is in the early stages of establishing a community – we welcome all potential collaborations**
- **Software collaborations**
 - d-Cache
 - IHEP
- **Sites and infrastructures**
 - RAL
 - AARNET
 - IHEP
- **WLCG**
 - RESTful interface for tape systems
 - HTTP bringonline
 - “Making the most of tape”
 - Working with the experiments to boost efficiency
 - Rucio

Outlook

- **CTA is feature-frozen in preparation for Run 3**
- **We expect consolidation during the start of the run**
 - Performance, reliability & operability
- **Scheduler optimisation**
 - Both experiment-driven and internal (repack)
- **CERN move to Postgres**
- **Community evolution**
 - Integrations
 - Deployments and distribution
 - Interfaces
- **Beyond LHC**
 - Finish Castor migration and support new experiments
 - <https://home.cern/science/experiments>
 - Data preservation
 - Babar, LEP, ...
 - Other targeted archival duties
 - Backup use-cases
- **Longer term**
 - <insert prediction here>

References and further info

- **CTA website** - <http://cta.web.cern.ch/cta>
- **CTA community** - <https://cta-community.web.cern.ch/>
- **CTA docs** - <https://eoscta.docs.cern.ch/>
- **CTA code** - <https://gitlab.cern.ch/cta/CTA>
- **cta-support@cern.ch**
- **Find us at**
 - **The EOS workshop**, 1st-5th March - <https://indico.cern.ch/event/985953/>
 - **HEPiX**, 15th-19th March - <https://indico.cern.ch/event/995485/>