PIC
port d'informació
científica

CERN
Tape Archive

CERN Tape Archive @ PIC

Site report

Elena Planas, Elisabet Carrasco (ex PIC, now CERN), Esther Acción, Jordi Casals
CTA Team @ PIC

IFAE 25 years 1991-2016 Institut de Física d'Altes Energies

EXCELENCIA SEVERO OCHOA

BIST Barcelona Institute of Science and Technology

GOBIERNO DE ESPAÑA MINISTERIO DE ECONOMÍA Y COMPETITIVIDAD

Ciemat Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas

EXCELENCIA MARÍA DE MAEZTU

cfp CIEMAT física de partículas

PIC
port d'informació
científica

PIC as a data center

- Spanish WLCG Tier-1 centre → Provides ~4% of LHC Tier-1 data processing
- Support to projects from different scientific disciplines: astronomy, cosmology, genomics, superconductivity, bioimaging…
- 2x100 Gbps to Academic Network → One of the largest data mover in Spain
- Member of the Spanish Supercomputing Network (RES)
- 24 employees, 8 at the Operations Team

CPU: +11k cores on HTCondor v.9.0.17
Disk: ~20PB - running on dCache v9.2.11-1
Tape: ~63PB - running on Enstore v6.3.4-14

PIC
port d'informació
científica

CTA Team at PIC

- Person in charge of the CTA implementation at PIC left
  - She is now at CERN
  - Waiting on opening a position for future CTA administrator
- Enstore admin
  - Performing tests related to the migration of Enstore to CTA
- dCache admin
  - Initially helped with the functional testing and now setting up the connection between dCache and CTA
- Monitoring admin
  - Previously coordinating the implementation of CTA and integration of CTA with the monitoring systems at PIC

PIC
port d'informació
científica

Our dCache + CTA instance

$pool.pic.es

dCache

dcache-cta
nearline driver

ctafrontend.pic.es

cta-frontend

cta-frontend-grpc

ctadb.pic.es

Catalogue
(postgreSQL)

ObjectStore
(NFS)

dCache 9.2.11-1
dcache-cta 0.11.0-1
CTA 5.8.7-1
PostgreSQL 14.6
IBM TS4500
1xLTO8

ctalogs.pic.es

Centralized
Log server

elastic

ctamon.pic.es

cta-cli

monitoring
scripts

graphite

Grafana

ctatps001.pic.es

ctatps002.pic.es

cta-taped

cta-rmcd

cta-frontend

Our tape infrastructure

IBM TS4500

IBM TS4500

- ➔ 5 LTO frames (L55+D55+3xS55)
- ➔ 10 L8 drives + 11 L9 drives - x4 on each tapeserver
- ➔ 6 tapeservers
- ➔ 5500 tape volumes (only LTO8)

Enstore is our tape storage system since 2007 + **dCache** for disk

SL8500

Oracle SL8500

Finally decommissioned and removed from our datacenter

Next challenge!

Build a separate room to accommodate a new library due to humidity problems appeared in newer tape technologies

- ➔ still don't know what manufacturer, technology or final budget

LTO9 specs change and humidity problems

- Talks with IBM and Fujifilm confirmed that the relative humidity must be between 20%RH - 50%RH from LTO9 and less than a ±5%RH diff in one hour
- We don’t use LTO9 tapes, only drives
  - All tapes are LTO8
- CTA test environment initially with two LTO9 drives
  - They have humidity and temperature sensors (LTO8 only temp)
  - Later switched to LTO8 drives to continue testing
- CTA error appeared when humidity increased
  - In summer, the (usual) heat wave increases humidity over 50%RH
  - Drives detect an error are are not able to mount tapes

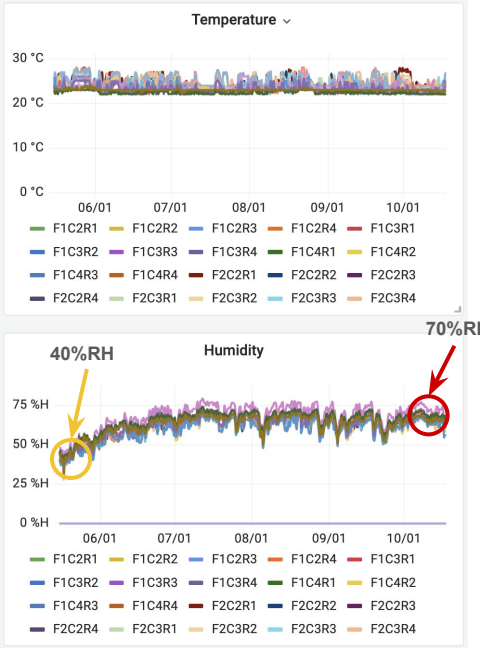

Progress made so far

- Setting up infrastructure with multiple components in different servers
  - Virtual + physical, tested with CentOS 7
- dCache connection to CTA using DESYs provided plugin
- Testing things for general understanding of CTA
  - Finding similarities between concepts from Enstore and CTA
    - Use storage classes+tapepools as Enstore File Families
  - Daily operations: read, writes, double copies, file deletion, tape repacks, etc
  - Understanding mount policies, criteria, rules…
- Integration with monitoring tools
  - Icinga for active checks and alerts on components status, Graphite as metric backend (collectd + custom scripts), Elasticsearch for log storage and search and Grafana for plotting
- Automated all our configuration with puppet

Migration tests from Enstore to CTA

- Previous steps
  - Clear existing CTA database
  - Configuring dCache test with both HSM instances
- Testing migration with script provided by the team at FNAL
  - Understanding organization from CTA, relation with Enstore concepts
  - Only active files are migrated
  - If necessary, deleted files can always be recovered by maintaining a copy of Enstore DB
- Refined and decided the tapepool organization (file family concept)
  - With help also from Dmitry from FNAL, modifying the migration script
- Tested migration for files with double copies
- Read tests for the 2 formats supported by Enstore (cern and cpio)
- Testing migration for LTO8 and LTO7M8 technologies
- Defining a migration procedure

PIC
port d'informació
científica

Plans for 2023

- Test tapeservers efficiency with 4 drives
  - Not tested yet. We'll probably go for 2 drives per tapeserver
- Test with more realistic loads to test performance & bandwidth
  - Tested some loads, but not the big ones we thought last year
- Keep refining our monitoring system
  - Finished integration with Elasticsearch, Graphite, Grafana. Monitoring with Icinga (Nagios)
- Discuss the use of compression - never enabled with Enstore
  - We didn't have it in Enstore and, for accounting purposes, we won't have it in CTA
- Start using CTA for internal backups
  - Built a simple system sending dCache and Enstore database backups regularly to CTA
- Define how the data migration from Enstore will be done
  - FNAL team helped us with a migration script and constant communication - THANK YOU!
- No date to switch to production yet!
  - We have now an approximate date - mid 2025

Plans for 2024 and 2025

- Waiting for the EL9 version of CTA to update our system (Q2-2024?)
- Testing migration with production database replica (Q2-2024)
- Continuing functional tests based on our needs (Q2,Q3-2024)
  - Adapting from Enstore to CTA
- Acquiring new hardware for the final migration (Q3,Q4-2024)
- Start migrating by experiment (Q1-2025)
- Planned final migration and going into production (Q2-2025)
  - Finishing new room construction and new tape library arrival

PIC
port d'informació
científica

CERN
Tape Archive

Thanks!

But also thanks to …

The **Enstore team** for the meetings and the migration help,
the **dCache team** for the plugin integration and the initial talks,
the **CTA Community Forum** for answering our amateur questions

IFAE
25 years 1991-2016
Institut de Física d'Altes Energies

EXCELENCIA SEVERO OCHOA

BIST Barcelona Institute of Science and Technology

GOBIERNO DE ESPAÑA
MINISTERIO DE ECONOMÍA Y COMPETITIVIDAD

Ciemat
Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas

EXCELENCIA MARÍA DE MAEZTU

cfp CIEMAT
física de partículas

LTO Specifications

Main specifications of LTO cartridges

		LTO CL※	LTO G4/G4 WORM	LTO G5/G5 WORM	LTO G6/G6 WORM	LTO G7/G7 WORM	LTO G8/G8 WORM	LTO G9/G9 WORM
Basic specifications	Capacity (Max compression)	—	800GB(1.6TB)	1.5TB(3.0TB)	2.5TB(6.25TB)	6TB(15TB)	12TB(30TB)	18TB(45TB)
	Data transfer rate (Max compressed) *Depends on the drive interface.	—	120M/s (240MB/s)	140MB/s (280MB/s)	160MB/s (400MB/s)	300MB/s (750MB/s)	360MB/s (750MB/s)	400MB/s (1,000MB/s)
	Number of tracks	—	896(16 track head serpentine)	1,280(16 track head serpentine)	2,176(16 track head serpentine)	3,584(32 track head serpentine)	6,656(32 track head serpentine)	8,960(32 track head serpentine)
	Servo method	—	Timing-based					
	Cartridge memory	32,786 bits (4,096 bytes); internal EEPROM with electromagnetic induction antenna	65,280 bits (8,160 bytes); internal EEPROM with electromagnetic induction antenna		130,816 bits (16,352 bytes); internal EEPROM with electromagnetic induction antenna			261,888 bits (32,736 bytes); internal EEPROM with electromagnetic induction antenna
Physical specifications	Tape with	12.65mm						
	Tape thickness	8.9μm	6.6μm	6.4μm	6.1μm	5.6μm		5.2μm
	Tape length	319m	820m	846m		960m		1,035m
	Cartridge dimensions	H102.0 x W105.4 x D21.5mm						
Recommended operation environment	Temperature	10~45℃						15~25℃
	Humidity	10~80%RH						20~50%RH
	Max wet-bulb/dew-point temperature	26℃ (Max wet-bulb)						22℃ (Max dew-point)
Recommended storage environment	Temperature	16~35℃ (short-term) / 16~25℃ (long-term)						15~25℃
	Humidity(short/long-term)	20~80%RH (short-term) / 20~50%RH (long-term)						20~50%RH
	Max wet-bulb/dew-point temperature	26℃ (Max wet-bulb)						22℃ (Max dew-point)
Allowable operation environment	Temperature	10~45℃						15~35℃
	Humidity	10~80%RH						20~80%RH
	Max wet-bulb/dew-point temperature	26℃ (Max wet-bulb)						22℃ (Max dew-point)
Allowable storage environment	Temperature	16~35℃ (short-term) / 16~25℃ (long-term)						15~35℃
	Humidity(short/long-term)	20~80%RH (short-term) / 20~50%RH (long-term)						20~80%RH
	Max wet-bulb/dew-point temperature	26℃ (Max wet-bulb)						22℃ (Max dew-point)
Physical specifications	Encryption support	×	○	○	○	○	○	○
	LTFS support	×	×	○	○	○	○	○