

Maintaining consistency in an EOSCTA system

EOSCTA Namespace Reconciliation

Richard Bachmann


CERN IT-ST-TAB
EOS Workshop, March 2022

09.03.2022




Motivation

EOSCTA System:



Motivation

EOSCTA System:



Accidents happen.

- 1 "Line of defence" after transfer tool feedback
- 2 Know before user if archival has gone wrong
- 3 Give proof that data is safe and sound after transfer
 - Regularly updated "table of contents"



EOS Namespace and CTA Catalogue

EOS:

- Source:
 - The MGM's QuarkDB
- Exported using
'eos-ns-inspect scan'

Critical Metadata

- *archive_file_id*
- *path*

CTA:

- Source:
 - Catalogue DB
- Exported using DB query

Critical Metadata


- *archive_file_id*
- *storage_class_name*



Steps

For each Virtual Org.

- 1 Create raw dump
 - Store raw archives on EOS
- 2 Merge
 - On *archive_file_id*
- 3 Validate
 - Run validator scripts
- 4 Produce export dump
 - Trimmed version of merged dumps
- 5 Upload JSON logs
 - Upload with Fluentd
 - Store in InfluxDB



Validation Process

Presently we check

- 1 Presence in namespace
 - EOS only?
 - CTA only?
 - Incorrectly marked as archived?
 - CASTOR import?
- 2 Metadata properties
 - Checksum, file size, file IDs, ...
- 3 Number of copies
 - Do we have the promised number of copies?
- 4 Broken tapes
 - Is the file completely unavailable due to broken tapes?



Run steps in containers

Advantages

- Encapsulates dependencies
- Per-VO isolation
- 'root'-less with Podman




podman



Orchestrate with Rundeck

- Scheduled workflows
 - Weekly for now
- Version controlled with git
 - Schedule
 - Config

RUNDECK



EOSCTA Namespace Dump Archives

- LZMA-compressed JSON archive of valid merged file records

v1.0 format

```
{ "eoscta_namespace_export": { "disk_instance": "eosctaalice" ...  
{ "adler32": "8ac7845f", "on_tape": true, ...  
{ "adler32": "2ab4ac13", "on_tape": true, ...
```

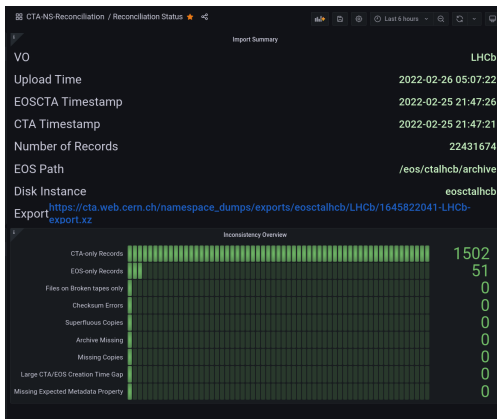
v1.0 Single record

```
{  
  "adler32": "8ac7845f",  
  "on_tape": true,  
  "path": "/eos/ctaalice/archive/some_file.txt",  
  "size": 47120561  
}
```

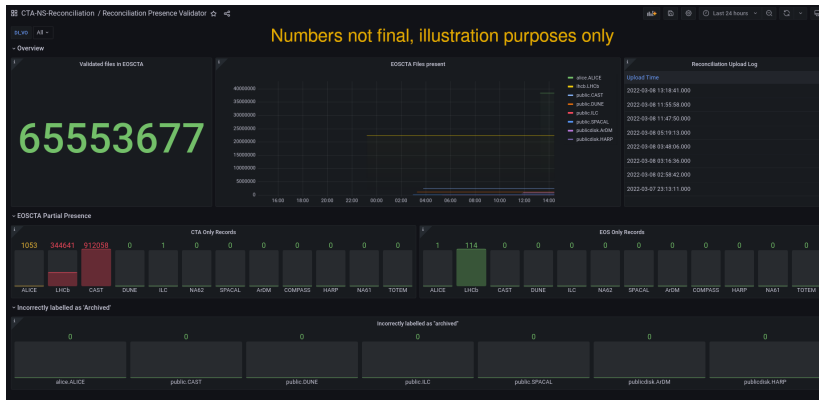


Grafana VO Dashboard

■ View for operators



Grafana Presence Dashboard



Comments & Questions

- Exported namespace dump article: [KB0007876](#)

We have

- Set of validation scripts
- Containers to run them
- Rundeck workflow
- Regularly exported merged dumps
- Grafana monitoring (work in progress)

