

# EOSCTA File Restoring

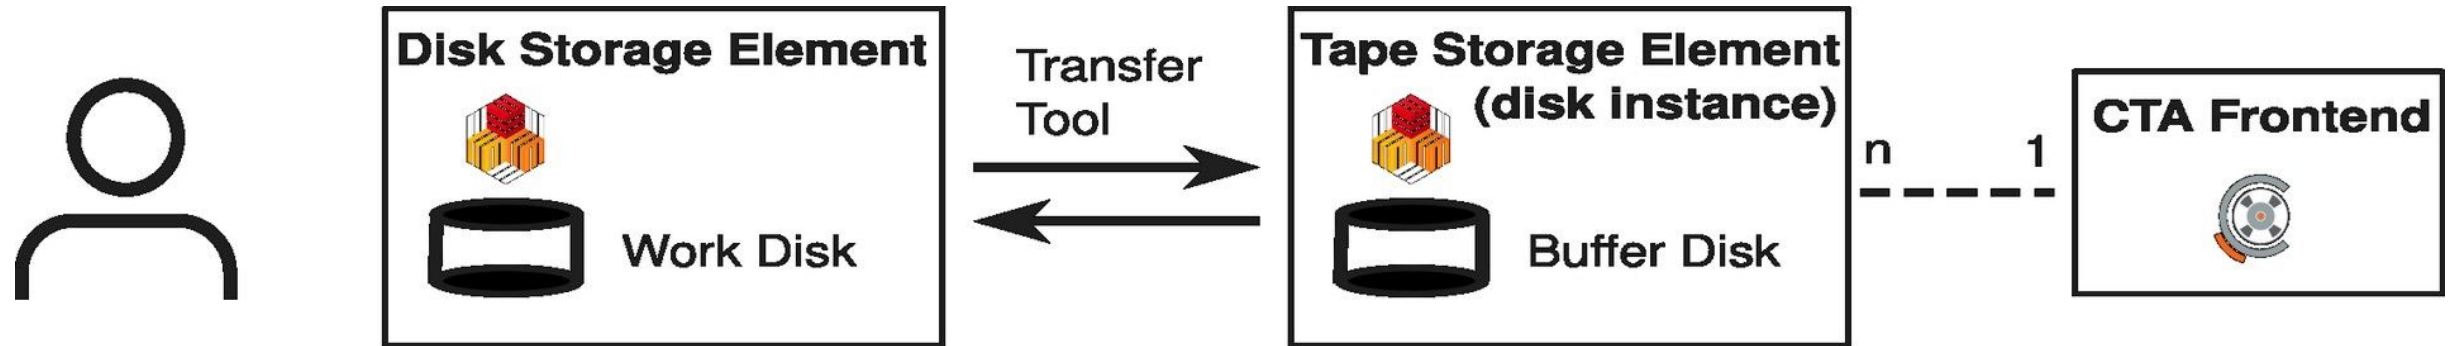
Miguel Barros - IT-ST-TAB

Wednesday, March 9th 2022



CERN  
Tape Archive

# EOSCTA Architecture Overview



# EOSCTA File Metadata

```
[root@ctacli /]# cta-admin tf ls --vid V01007
archive id copy no      vid fseq block id instance disk fxid size checksum type checksum value  storage
class owner group      creation time path
      10062      1 V01007 10006 100051 ctaeos      271f 407      ADLER32      20ce7e60 ctaStorag
eClass 11001 1100 2022-03-04 09:18 -
      10063      1 V01007 10007 100061 ctaeos      272b 407      ADLER32      20ce7e60 ctaStorag
eClass 11001 1100 2022-03-04 09:19 -
      9940      1 V01007 10008 100071 ctaeos      272d 407      ADLER32      20ce7e60 ctaStorag
eClass 11001 1100 2022-03-04 09:20 -
      9957      1 V01007 10009 100081 ctaeos      2730 407      ADLER32      20ce7e60 ctaStorag
eClass 11001 1100 2022-03-04 09:21 -
```

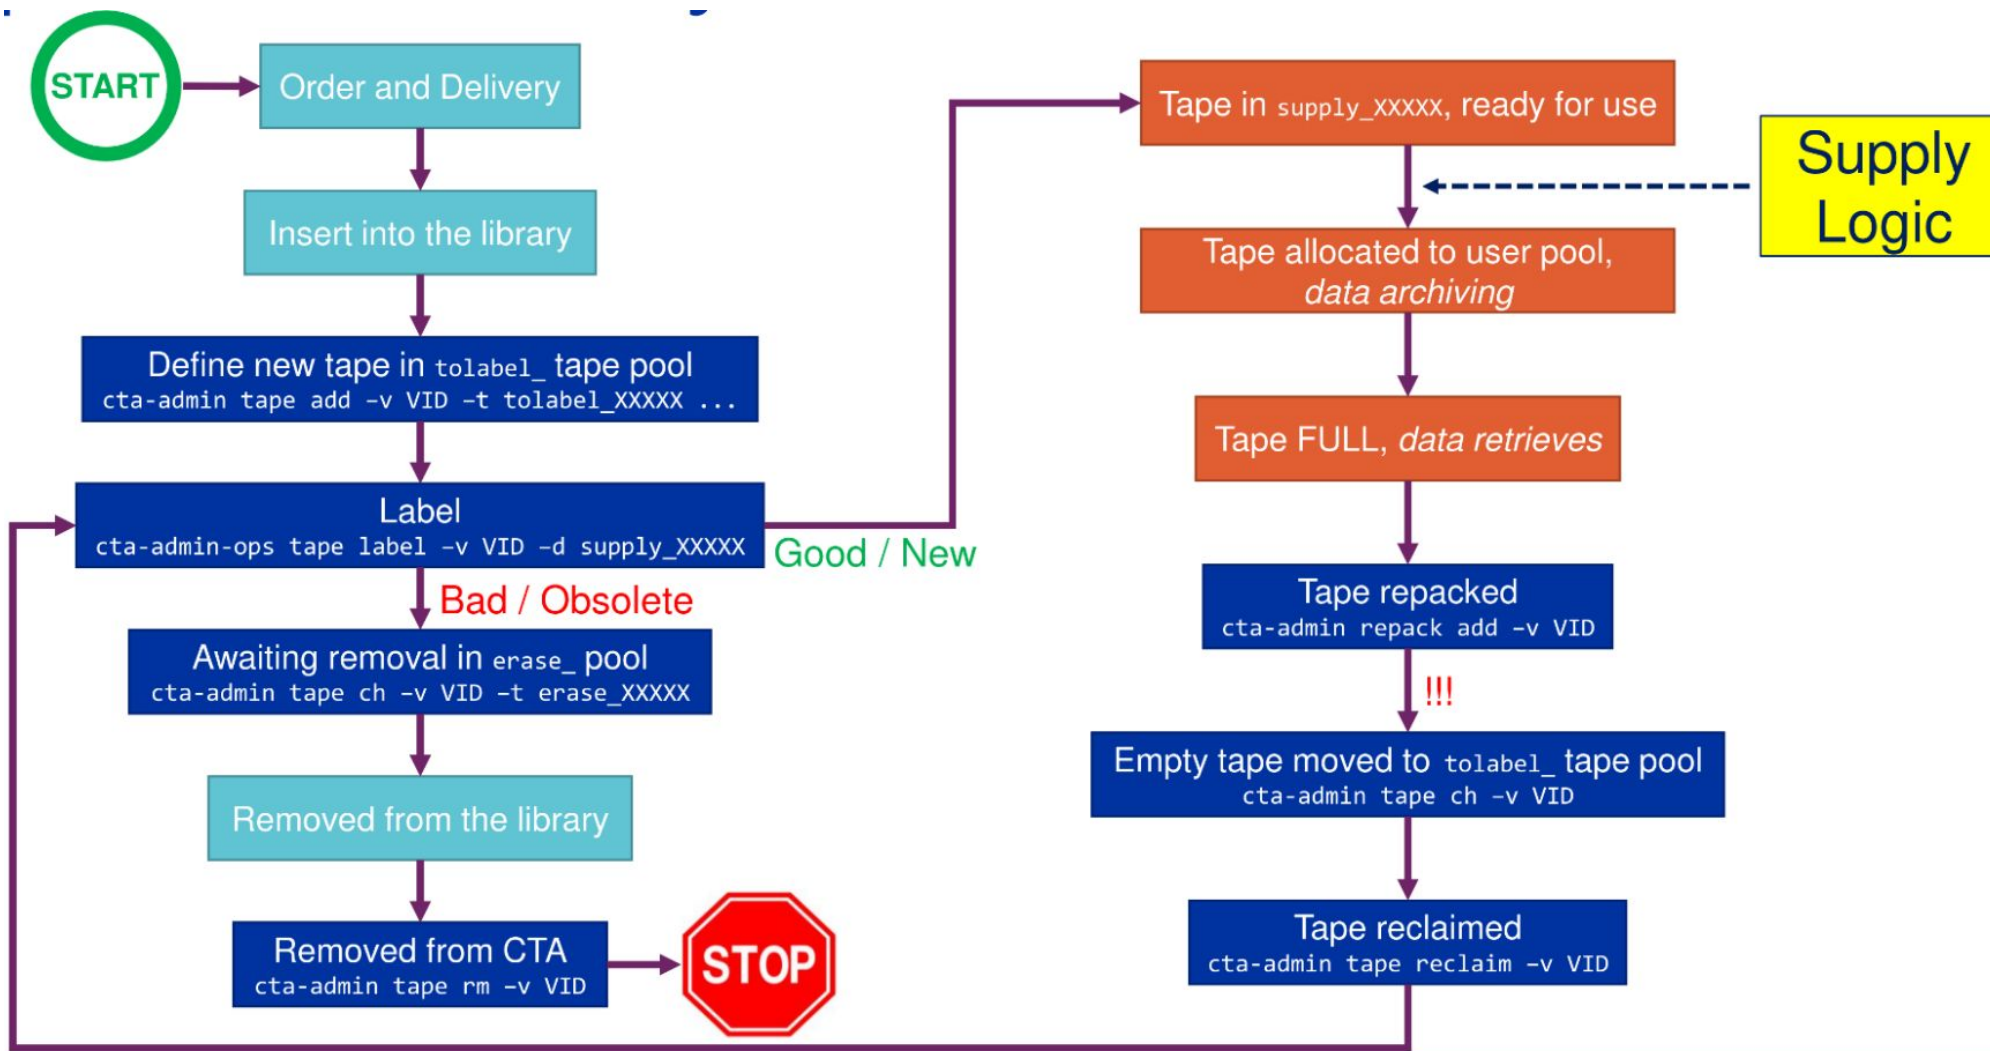
```
[root@ctaeos /]# eos ls -ly eos/ctaeos/cta/
d0::t1 -rw-r--r-- 1 user1 eosusers      407 Mar  4 09:19 06df4433-bd5d-4643-b131-25090f6943d
4
d0::t1 -rw-r--r-- 1 user1 eosusers      407 Mar  4 09:20 3a7f0aa2-1376-4347-9da2-293f174a38f
8
d1::t0 dr-xr-xr--+ 1 user1 eosusers           0 Mar  4 09:22 48d8a760-a240-467a-b17e-c464a6dfa63
2
d1::t1 -rw-r--r-- 2 user1 eosusers      407 Mar  4 09:18 67ba71b5-abe3-4b5b-9472-9e7040755b8
5
d0::t1 -rw-r--r-- 1 user1 eosusers      407 Mar  4 09:20 6eac6cda-6543-4a3c-80df-4a6439ba3c6
8
d0::t1 -rw-r--r-- 1 user1 eosusers      407 Mar  4 09:21 d38c7f9b-cc48-4227-8045-b7d249c3f08
d
d0::t1 -rw-r--r-- 1 user1 eosusers      407 Mar  4 09:19 ddeed1ea-da44-4b39-97a1-8ef41650d66
c
```

# EOSCTA File Metadata

```
[root@ctacli /]# cta-admin rtf ls --vid V01007 | head -n 5
```

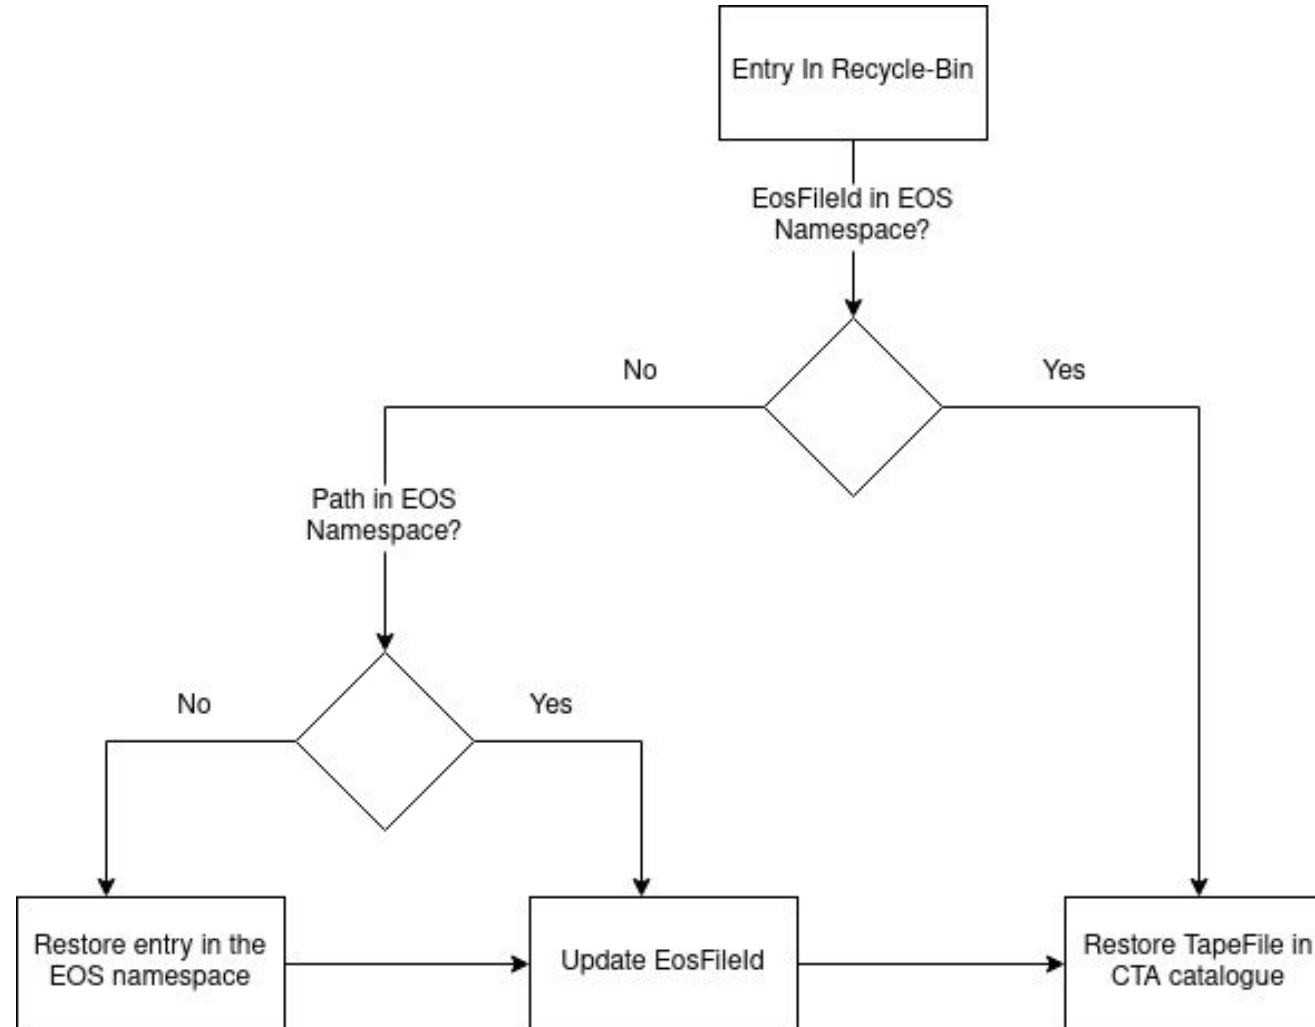
archive	id	copy	no	vid	fseq	block	id	instance	disk	fxid	size	checksum	type	checksum	value	storage
class	owner	group		deletion	time											path when del
eted	reason															
1	1	V01007	1	0	ctaeos	7	407					ADLER32		20ce7e60	ctaStorag	
eClass 11001	1100	2022-03-04 09:08														/eos/ctaeos/cta/2e6871ee-a3f4-4413-9b1d-ed7f9b3e
4463	File deleted by user1 from the ctaeos instance															
4	1	V01007	2	11	ctaeos	b	15.4K					ADLER32		977bffd4	ctaStorag	
eClass 11001	1100	2022-03-04 09:17														/eos/ctaeos/preprod/e2384e2a-670d-44c1-8322-db53ea424075/0/test0000
0000	File deleted by poweruser1 from the ctaeos instance															
7	1	V01007	3	21	ctaeos	e	15.4K					ADLER32		a832fee9	ctaStorag	
eClass 11001	1100	2022-03-04 09:17														/eos/ctaeos/preprod/e2384e2a-670d-44c1-8322-db53ea424075/0/test0000
0002	File deleted by poweruser1 from the ctaeos instance															
5	1	V01007	4	31	ctaeos	c	15.4K					ADLER32		80b0ff7f	ctaStorag	
eClass 11001	1100	2022-03-04 09:17														/eos/ctaeos/preprod/e2384e2a-670d-44c1-8322-db53ea424075/0/test0000
0001	File deleted by poweruser1 from the ctaeos instance															

# EOSCTA Tape Lifecycle



# EOSCTA File Restoring Procedure

```
[root@ctafrontend /]# cta-restore-deleted-files V01007 -I 1
```



# EOSCTA File Restoring Limitations

- **Lack of atomicity**
  - File may be restored in the EOS namespace but not in CTA.
  - Inconsistencies between the catalogue and the namespace can be caught in the namespace reconciliation.
- **Slow**
  - Each file restored implies multiple network calls to the CTA and EOS frontend.
  - Restoration done a file at a time to minimize possible inconsistencies between the namespace and the catalogue.
- **Setup required**
  - Create EOS access token with right permissions and delete it afterwards.
  - Wrapper script can automate token management.



# Conclusion

- CTA allows restoring files that were deleted previously.
- File restoration is a last resort tool, there are limitations, it is separate from normal workflows.
- Documentation on file restoration available [here](#).
  - Available outside CERN!

