CERN
Tape Archive

Tape Drive Status Lifecycle

Jorge Camarero Vera

CERN, IT Department, Storage Group (IT-ST-TAB)

Tape Drives monitoring and Infrastructure

- For CTA operation and monitoring is important to know the states that a tape drive can reach during its lifecycle.

```
[root@tpsrv449 ~]# cta-admin dr ls
library    drive     host desired          request   status  since    vid
IBM1L8 I1L80913 tpsrv432      Up          RETRIEVE Transfer  20253 L76268
IBM1L8 I1L80914 tpsrv433    Down          RETRIEVE Transfer  46275 L75902
IBM1L8 I1L80933 tpsrv318      Up          RETRIEVE Transfer  32012 L75986
IBM1L8 I1L80934 tpsrv319      Up          NO_MOUNT     Free   5576      -
IBM1L8 I1L80942 tpsrv313      Up          RETRIEVE Transfer  19147 L76627
IBM1L8 I1L80943 tpsrv314      Up          NO_MOUNT     Free   9007      -
IBM1L8 I1L80944 tpsrv315      Up          RETRIEVE Transfer  17653 L76833
IBM1L9 I1L90611 tpsrv452      Up          NO_MOUNT     Free  97370      -
IBM1L9 I1L90612 tpsrv453      Up          NO_MOUNT     Free  16617      -
IBM1L9 I1L90613 tpsrv458      Up          NO_MOUNT     Free     44      -
```

- **Tape drives** are controlled and managed by **tape servers**.
- A **tape server** could manage multiple **tape drives**.
- At CERN, CTA is configured to use one **tape server** for each **tape drive.**
- Specific group of tape drives can access only to a certain **pool of tapes**.

Exporting 'cta-admin dr ls' as json

- We can call `**cta-admin dr ls**` as `**cta-admin –json dr ls**` to export all the information of a tape Drive in json format. This allows us to build interactive dashboards to monitor the tape drives more efficiently.

```
[root@tpsrv458 ~]# cta-admin --json dr ls | jq
[
  {
    "logicalLibrary": "IBM1L8",
    "driveName": "I1L80913",
    "host": "tpsrv432",
    "desiredDriveState": "UP",
    "mountType": "NO_MOUNT",
    "driveStatus": "UP",
    "driveStatusSince": "17909",
    "vid": "",
    "tapepool": "",
    "filesTransferredInSession": "0",
    "bytesTransferredInSession": "0",
    "latestBandwidth": "0",
    "sessionId": "0",
    "timeSinceLastUpdate": "11",
    "currentPriority": "0",
    "currentActivity": "",
    "ctaVersion": "4.6.0-1",
    "devFileName": "/dev/nst0",
```

cta-admin dr ls dashboard

This dashboard shows the status of all the tape drives for the various cta instances:

- *ctaproduction* is THE production tape backend of eosctaatlas, eosctaalice, eosctacms, eosctapublic...
- *preproduction* is the tape backend of all the preproduction instances: eosctacmspps, eosctapublicpps, eosctaalicepps...

Time	driveName	driveStatus	host	logicalLibrary	last	ds
2022-03-03 14:10:16	I3550944	CLEANING_UP	tpsrv082	IBM355	55	1
2022-03-03 14:10:16	I3601411	DOWN	tpsrv029	IBM360	2	0
2022-03-03 14:10:16	I4600244	DOWN	tpsrv478	IBM460	5	0
2022-03-03 14:10:16	I4601421	DOWN	tpsrv439	IBM460	89649	0
2022-03-03 14:10:16	S1L90514	DOWN	tpsrv330	SPC1L9	1	0
2022-03-03 14:10:16	I3550943	MOUNTING	tpsrv081	IBM355	9	1
2022-03-03 14:10:16	I1L80913	TRANSFERRING	tpsrv432	IBM1L8	11	1
2022-03-03 14:10:16	I1L80914	TRANSFERRING	tpsrv433	IBM1L8	4	1
2022-03-03 14:10:16	I1L80933	TRANSFERRING	tpsrv318	IBM1L8	5	1
2022-03-03 14:10:16	I1L80934	UP	tpsrv319	IBM1L8	9	3

CTA Tape Drives Status

STATUS OF A TAPE DRIVE

- **UP/FREE**: The tape drive is waiting for a data transfer session.
- **DOWN**: The tape drive isn't operative.
- **PROBING**: Checking if the tape drive is empty.
- **STARTING**: It marks the beginning of a Data Transfer Session
- **MOUNTING**: A tape is inserted in the tape drive
- **TRANSFERRING**: Data writing or reading between tape and disk.
- **CLEANING_UP**: It marks the process to restore the tape drive to an idle state.
- **UNLOADING**: It stops the actions of the tape drive on the tape
- **UNMOUNTING**: Extraction of the tape from the Tape Drive
- **DRAINING_TO_DISK**: Wait until all the data are transferred to disk.

Starting A Data Transfer Session

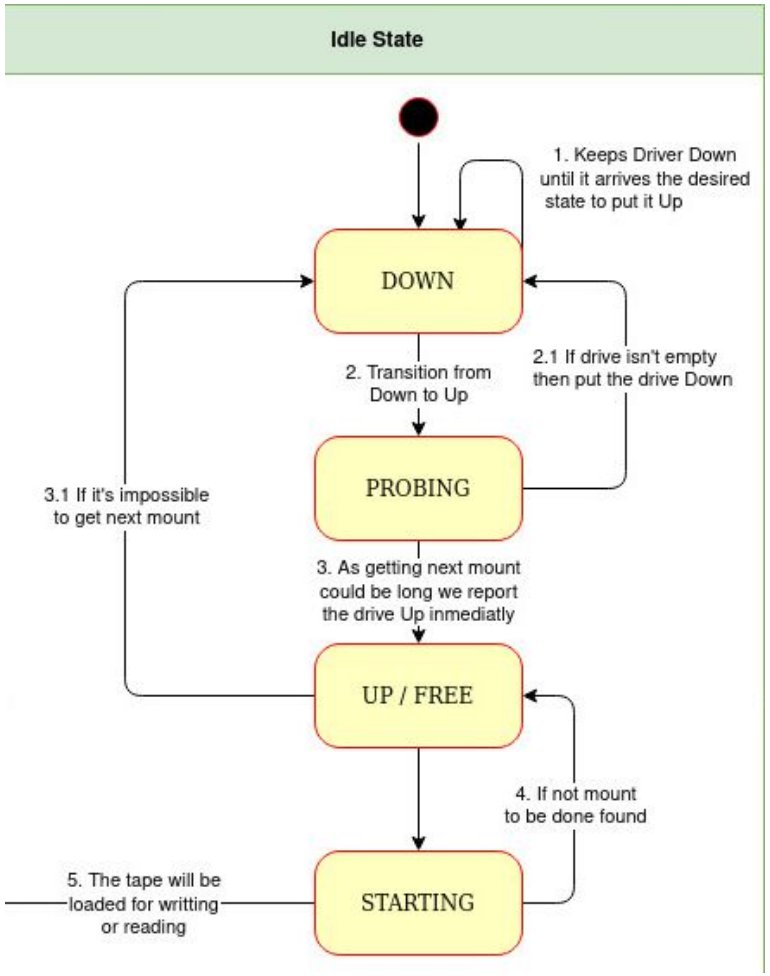

- When the Data Transfer Session begins the Tape Drive can be in two states: **DOWN** or **UP**.
- If it's in **DOWN** state then will wait until the drive receives a desired state to put it **UP**, passing first through the **PROBING** state.
  - If it fails to put the drive **UP** the system will put back the drive in **DOWN** state.
- If the cta scheduler finds a mount for a tape, then the data transfer session will pass to the **STARTING** state.

Data Writing Session

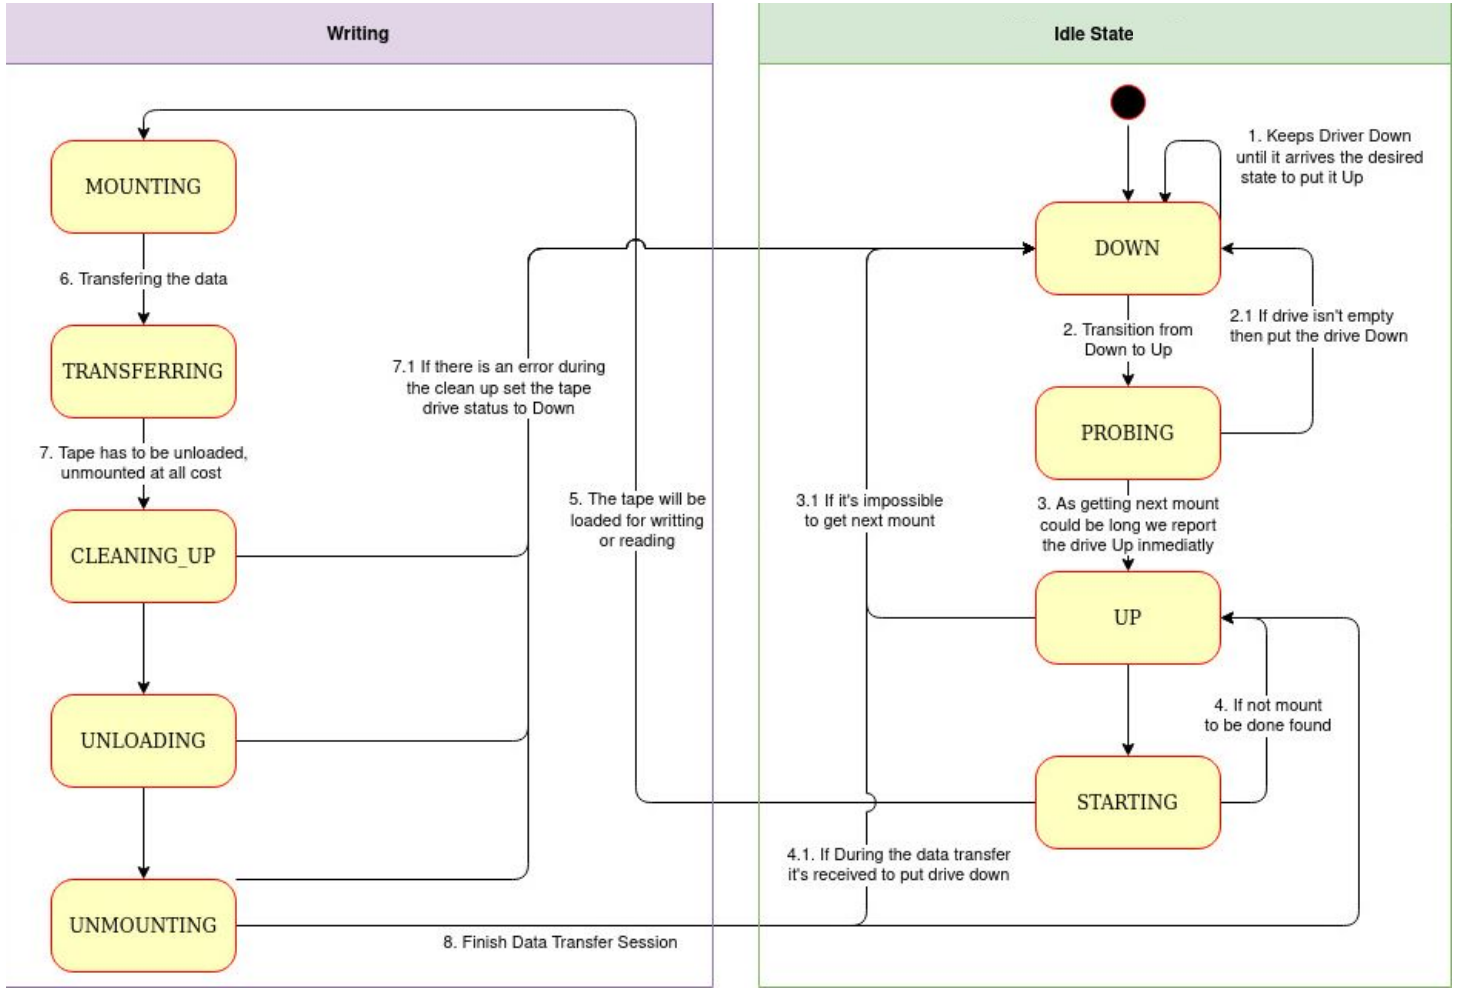

- It starts the **MOUNTING** of a tape and the **TRANSFERRING** of the data will start.
- After the data **TRANSFERRING** is finished the drive will start a **CLEANING_UP**, for that process it has to be done an **UNLOADING** and an **UNMOUNTING** of the tape
- After unmounting the tape the system will pass to **UP** or **DOWN** if during the data transfer session a signal to put the drive down is received

Data Reading Session

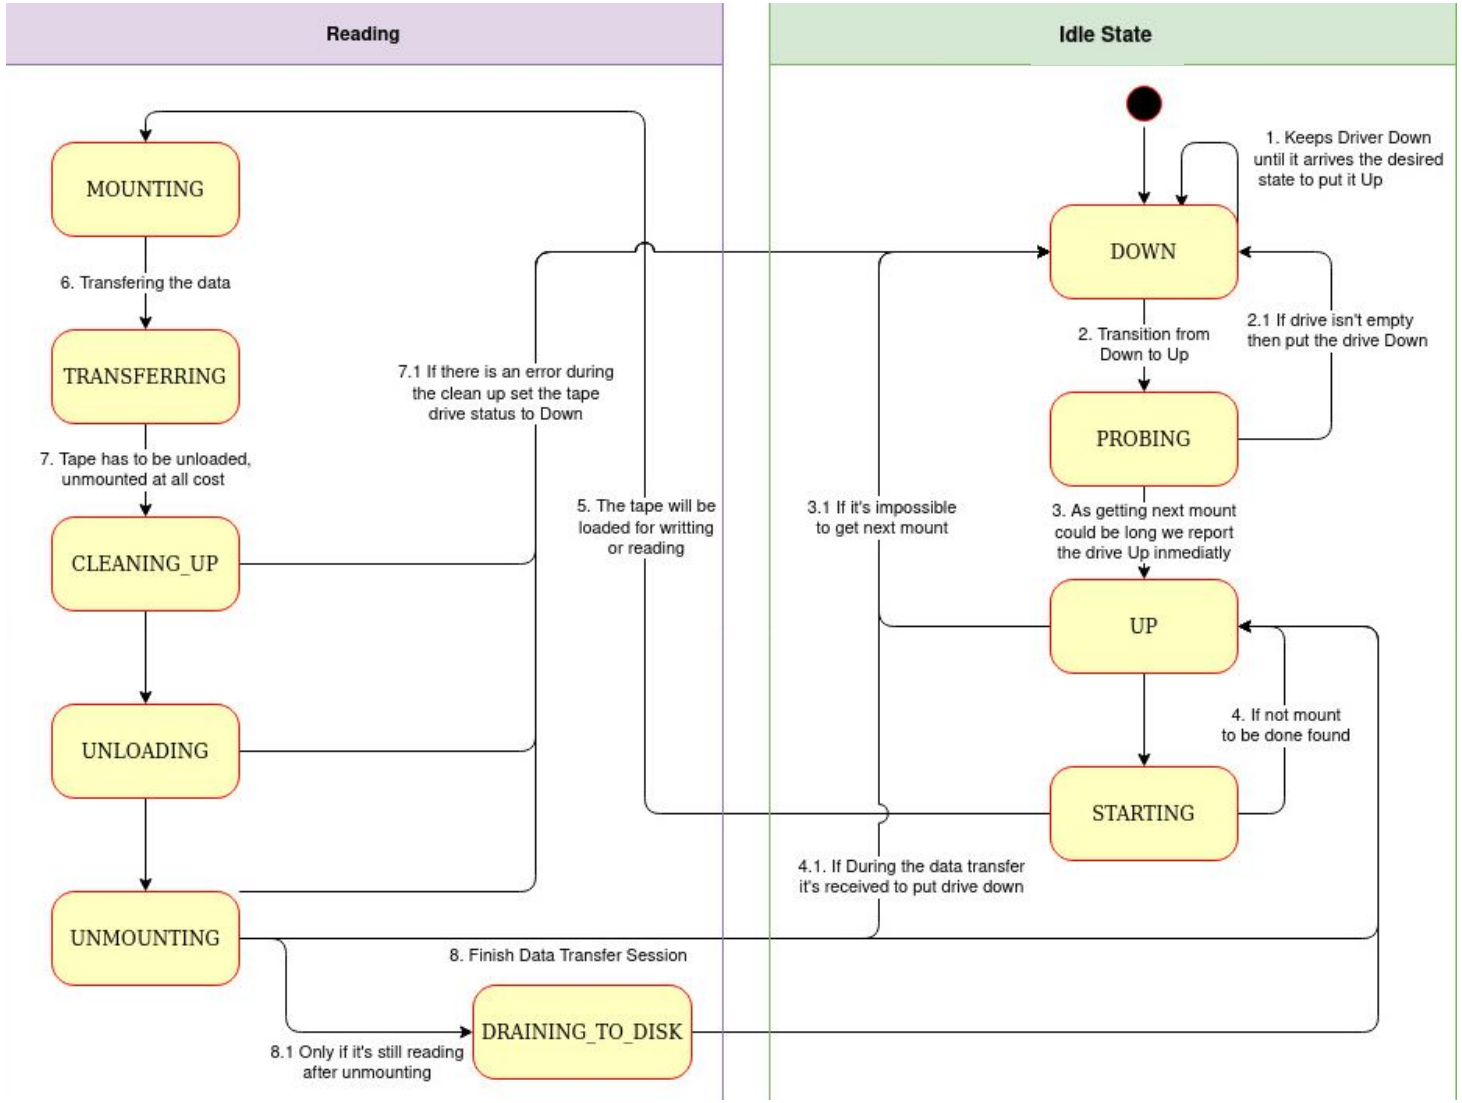

- A data reading session has a similar lifecycle as a writing one.
- The main difference is if after **UNMOUNTING** the tape the drive is still reading, the system will trigger a **DRAINING_TO_DISK** status.
- As a writing session when it ends the system will pass to **UP** or **DOWN**.

Summary

- Tape drives are the elements of CTA that allows to write and read the data on tapes.
- The command '**cta-admin dr ls**' allows us to monitor the tape drives.
- **Tape drives** are controlled and managed by **tape servers**.
- It is important to know the states that the tape drives can reach during a **data transfer session**.
- Knowing the lifecycle of a tape drive can help an operator to **solve problems** related with the tape drives and to do a better **monitoring**.

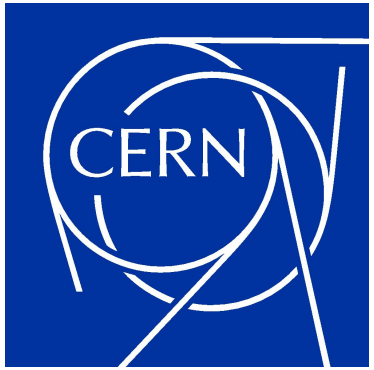

home.cern