




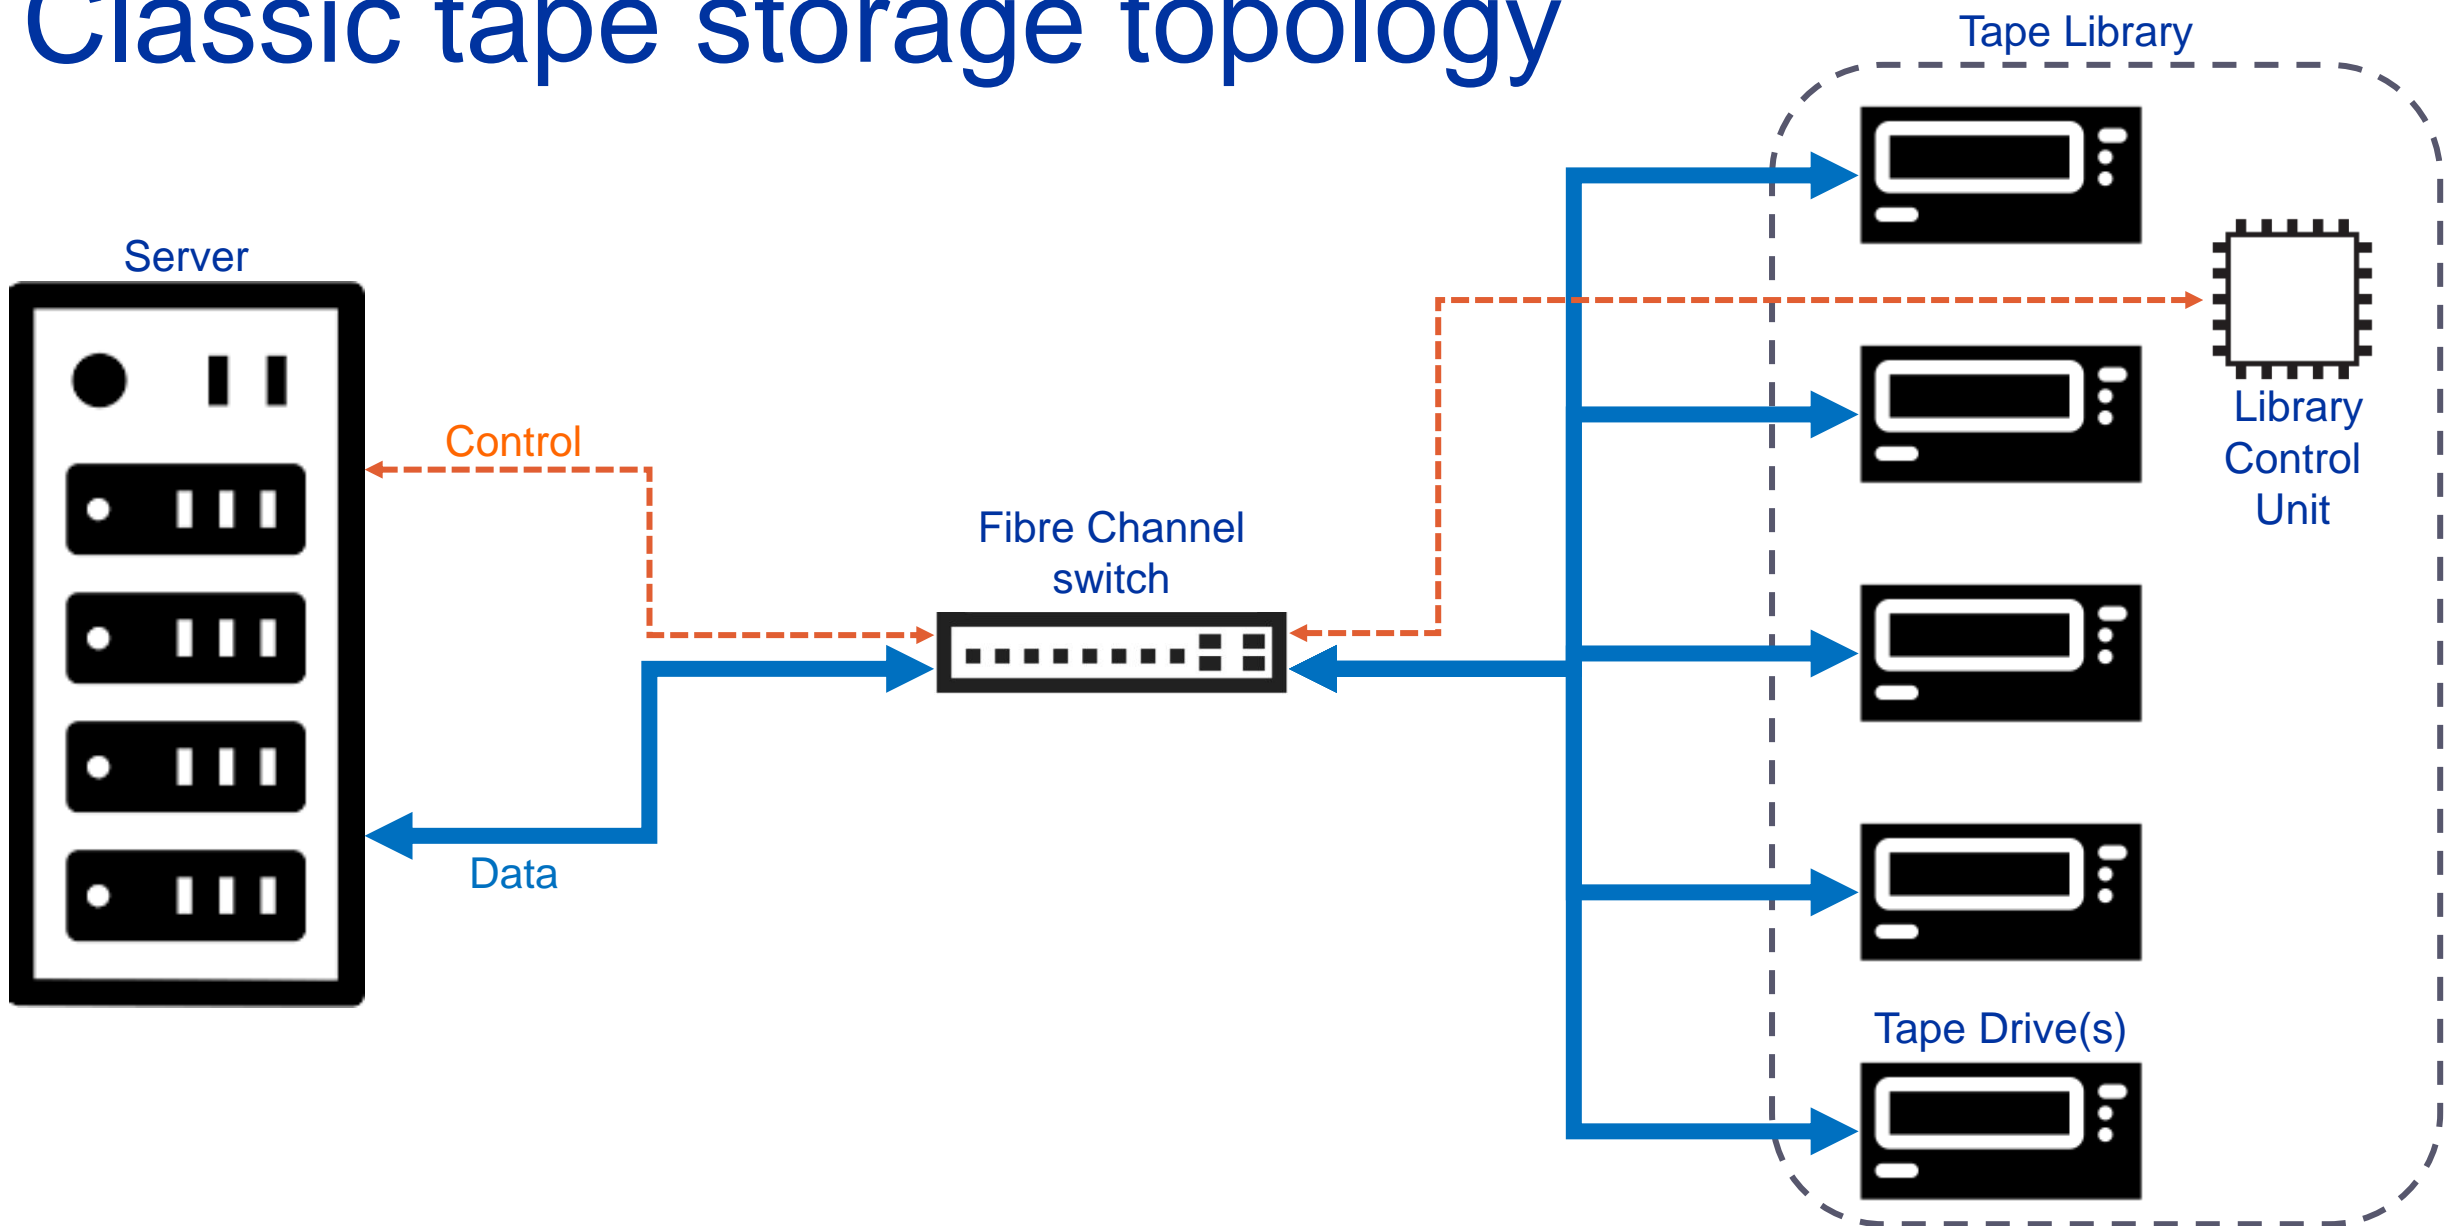


Simple and flexible drive and tape daemon configuration

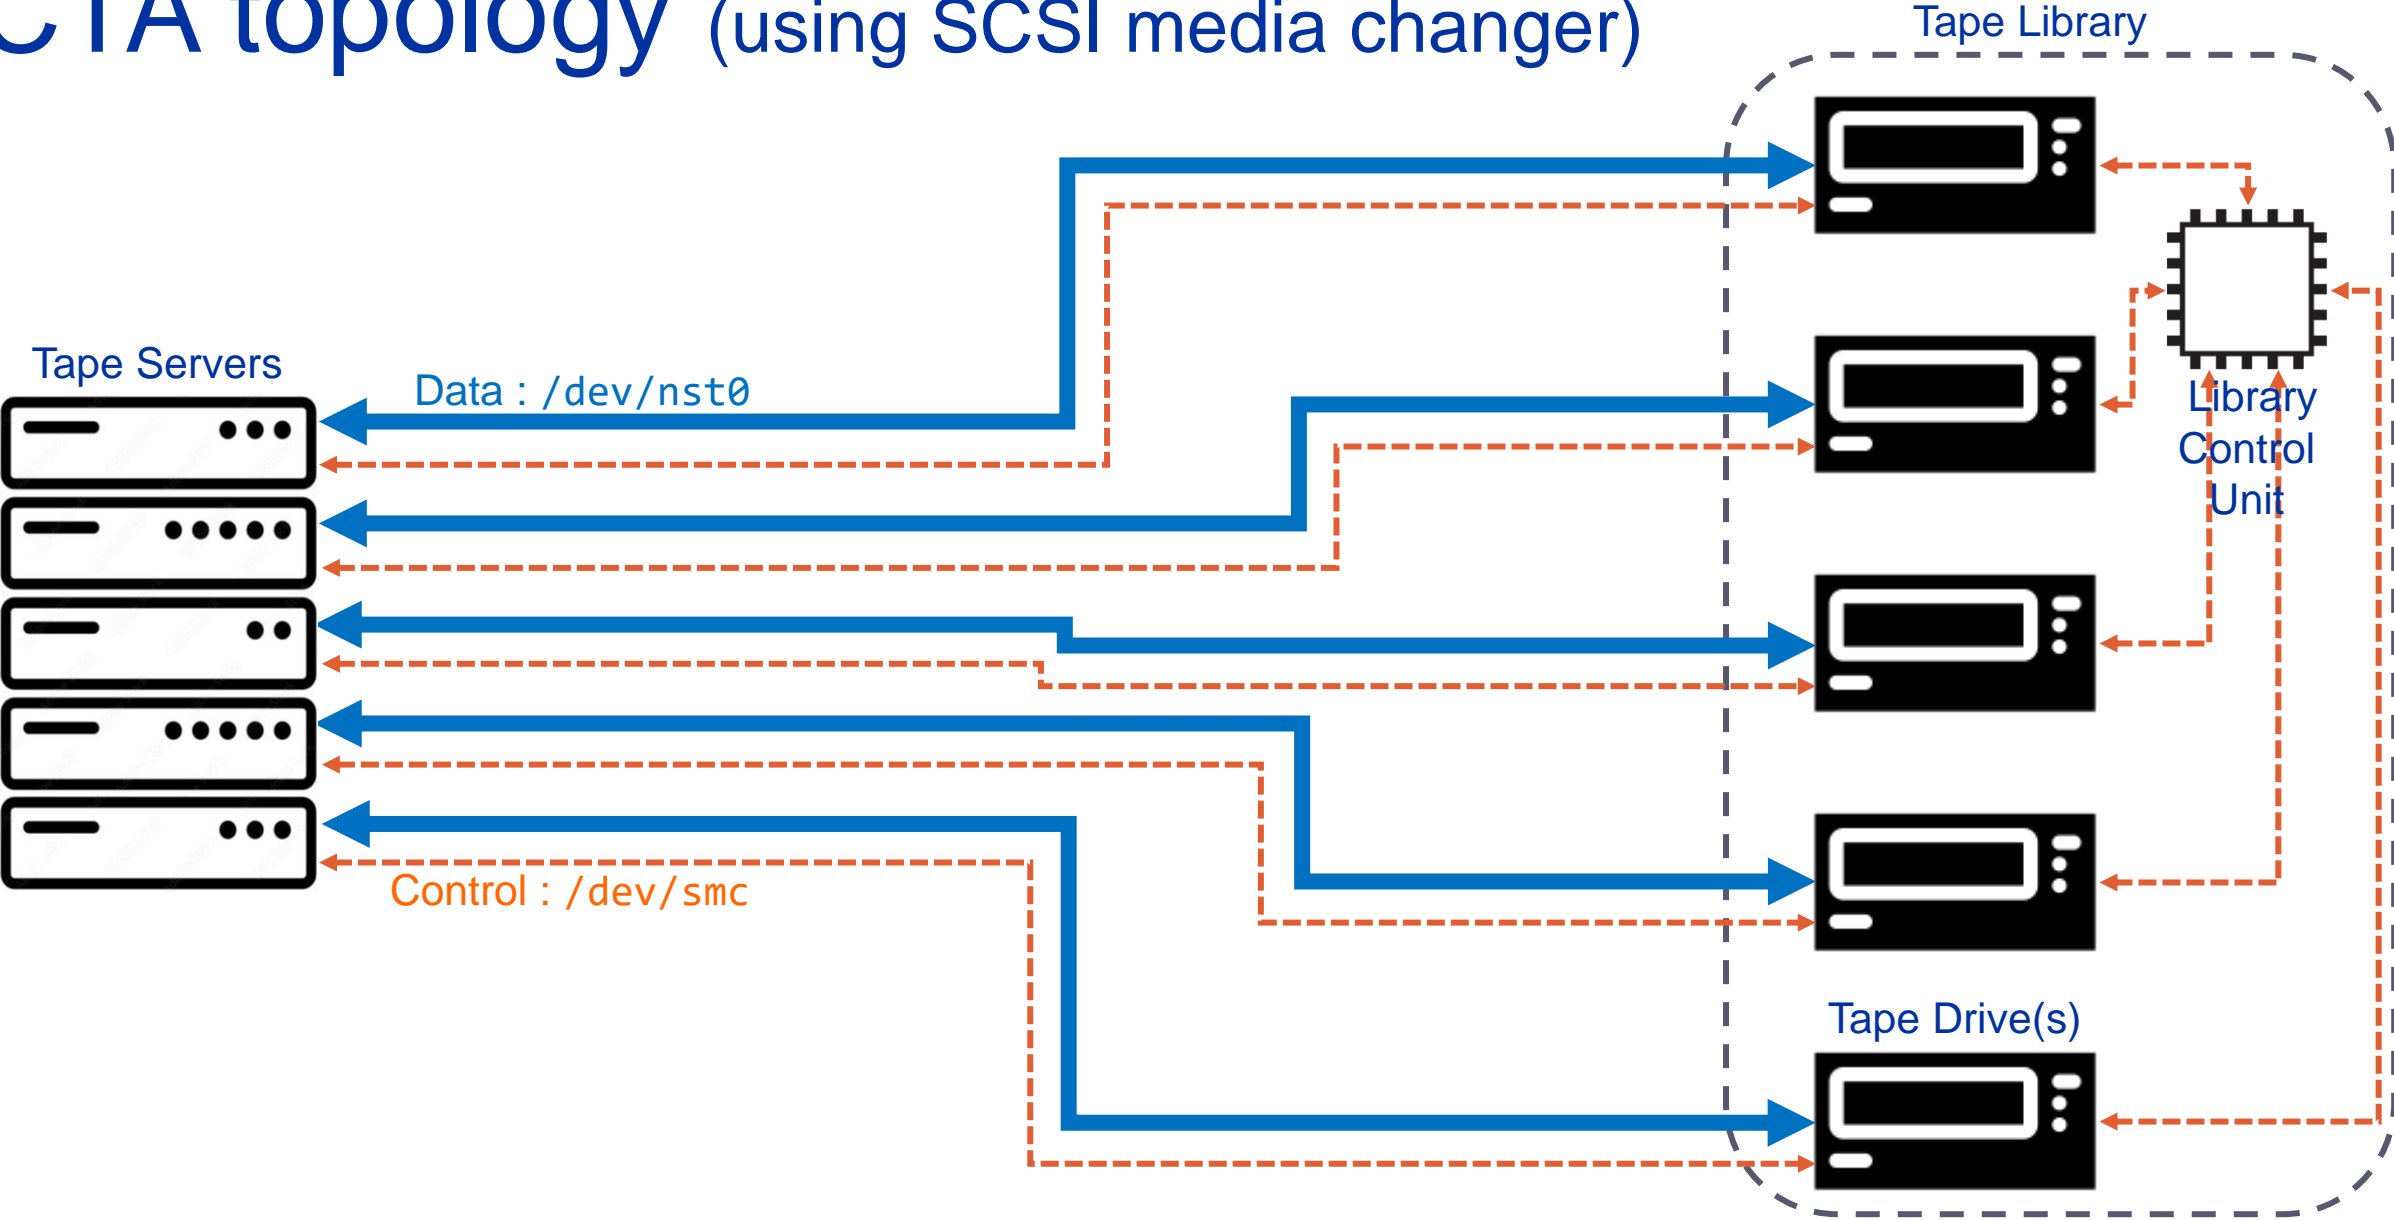
Vladimír Bahyl
Pablo Oliver Cortés
IT-SD-TAB



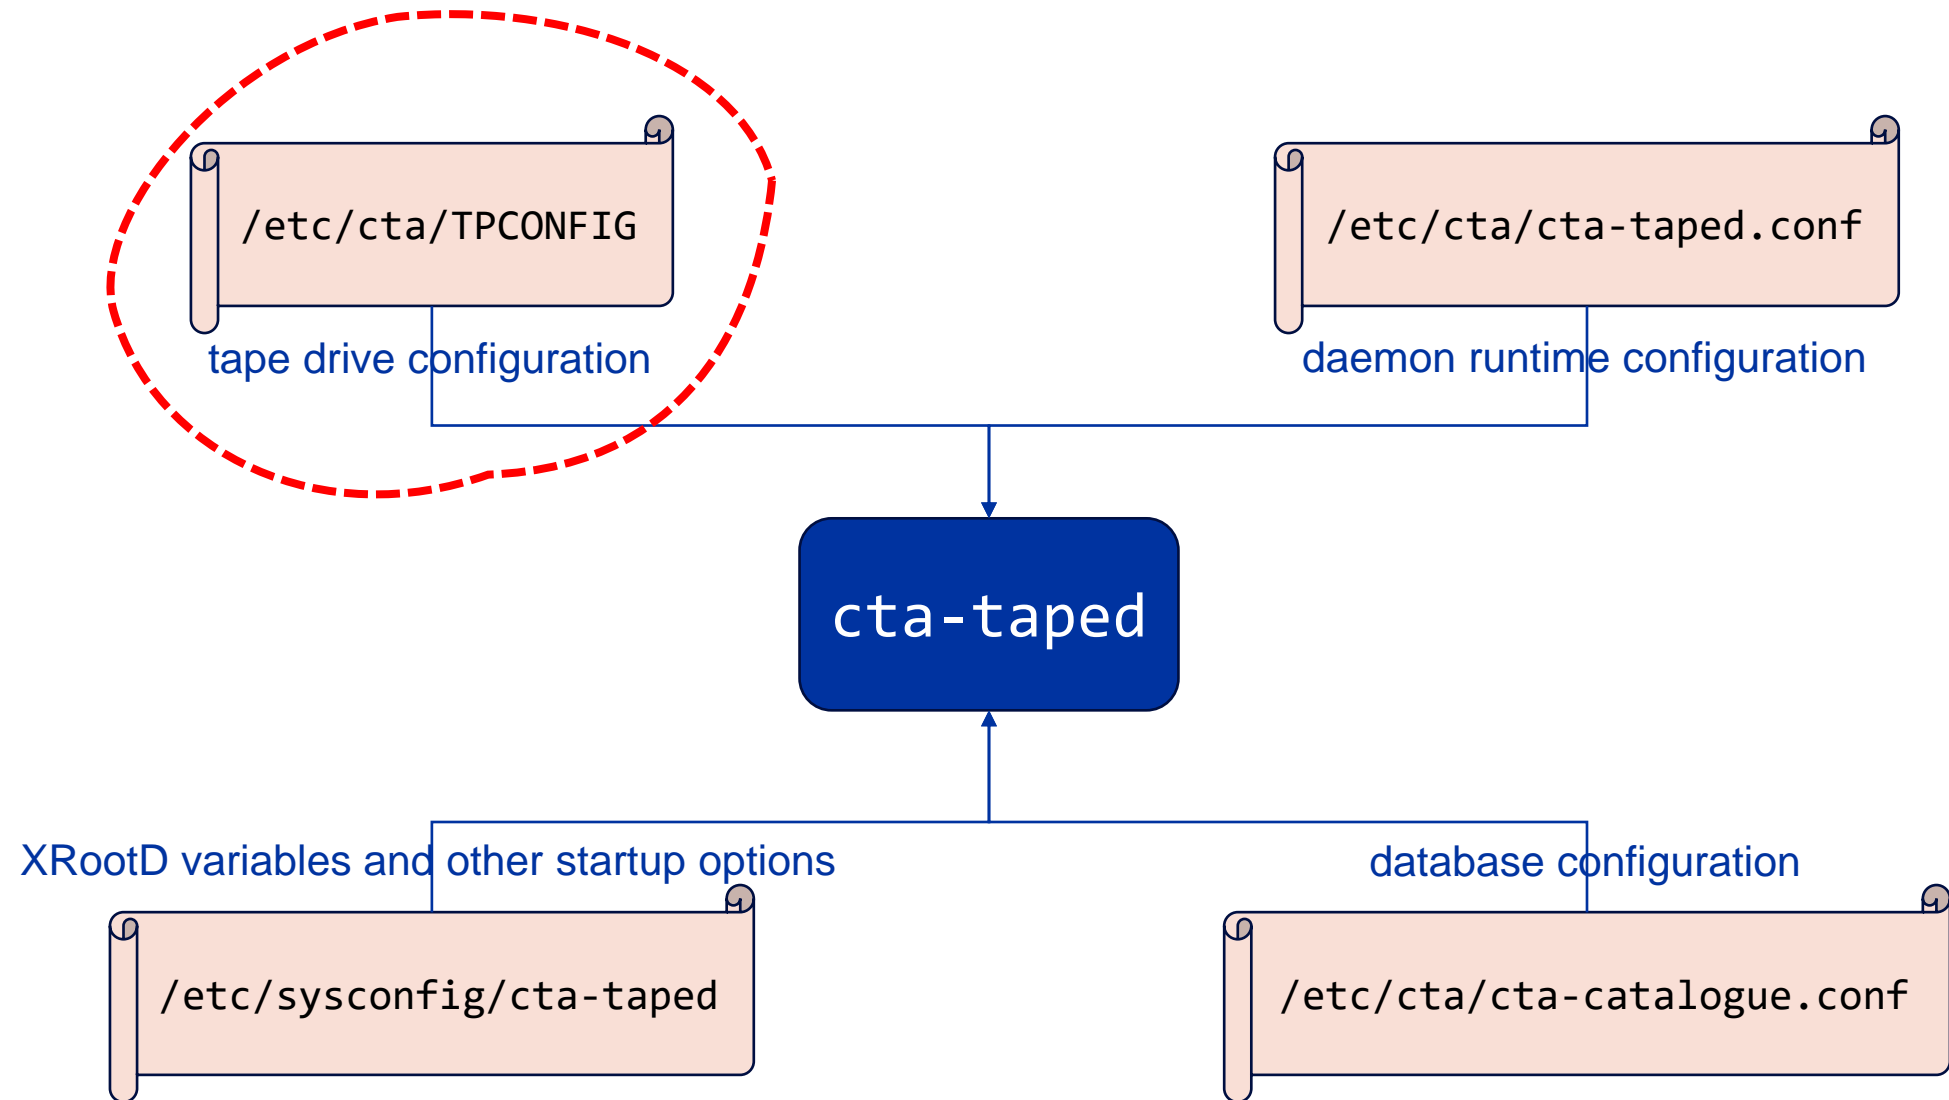
Classic tape storage topology



CTA topology (using SCSI media changer)



Current structure of cta-taped configuration files



/etc/cta/TPCONFIG particularities

- Legacy configuration file from CASTOR (CERN's predecessor of CTA)
 - At CERN it relies on manually entered configuration stored in unmaintained obsolete database
- Used by the cta-taped daemon to communicate with tape drive devices to archive and retrieve data
- The content depends on site specific details / naming conventions such as:
 - Tape libraries (logical & physical)
 - Tape drives (types)
 - Tape servers
 - Connection topology between these components

```
[root@tpsrv015 ~]# cat /etc/cta/TPCONFIG
#####
#
# CTA Tape Server Configuration file
#
# tape                library
# drive      logical  system  control
# name       library  device  path
I4550834    IBM455    /dev/nst0  smc27
#####
```

Most important field linking the drive (device) with the (position in the) library

Objective: Generate the configuration automatically

- Extract the maximum information directly from the tape drive (connected to a tape server) and use it to generate cta-taped drive configuration file
 - Problem: tape drives are independent (= they do not know about each other), only the library knows who is where (= which tape drive is defined in which library partition)

- Tape library serial number
 - Use SCSI media changer device

```
[root@tpsrv300 ~]# /usr/bin/sg_inq ` /usr/bin/lsscsi -g | /usr/bin/grep ' mediumx ' | /usr/bin/awk '{print $7}'`  
| /usr/bin/grep 'Unit serial number: '  
Unit serial number: 0000078AA6480401
```

- Tape drive type and serial number
 - Use SCSI generic (sg) device driver information

```
[root@tpsrv300 ~]# /usr/bin/sg_inq ` /usr/bin/lsscsi -g | /usr/bin/grep ' tape ' | /usr/bin/awk '{print $7}'`  
| /usr/bin/grep 'Unit serial number: '  
Unit serial number: 00000754420A
```

IBM TS4500 tape library configuration

- Accessible with `itdt` command in JSON format via REST over SCSI
 - Headers and trailers may need to be removed
- Important details of each tape drive:
 - Frame, Column, Row position
 - Serial number
 - Drive type and model
 - SCSI element address

```
[root@tpsrv300 ~]# /opt/itdt/itdt -f /dev/sg3 RoS GET /v1/drives
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 22924
[
  {
    ▶ "location": "drive_F5C4R1",
    ▶ "sn": "754420A",
      "mediaType": "3592",
      "state": "online",
      "operation": "empty",
      "accessible": "normal",
    ▶ "mtm": "3592-60F",
      "barcode": "11S02XV263YF10WT004333",
      "interface": "fibreChannel",
      "interfaceMode": "multi-mode",
      "logicalLibrary": "IBMLIB3",
      "use": "controlPath",
      "firmware": "590C",
      "encryption": "enabled",
      "wwnn": "500507604414eacc",
    ▶ "elementAddress": 269,
      "beacon": "disabled",
      "volser": null,
      "lastCleaned": "2023-06-13T11:17:43+0200"  },
    :
    :
    :
```


Example of the steps to obtain SMC value for IBM TS4500 tape library

Manual step: use IBM itdt tool to extract information about all drives installed in the library using REST over SCSI on the SCSI media changer device

Make sure the received JSON data valid

Use drive serial number to extract the element address

Get the control path with element address

```
[root@tpsrv300 ~]# /opt/itdt/itdt -f /dev/sg3 RoS GET /v1/drives > /tmp/IBMLIB3-configuration.json
[root@tpsrv300 ~]# vim /tmp/IBMLIB3-configuration.json # <<<--- REMOVING undesired headers and trailers
[root@tpsrv300 ~]# json_verify < /tmp/IBMLIB3-configuration.json
JSON is valid

[root@tpsrv300 ~]# /usr/bin/sg_inq ` /usr/bin/lsscsi -g | /usr/bin/grep ' tape ' | /usr/bin/awk '{print $7}'
| /usr/bin/grep 'Unit serial number: '
Unit serial number: 00000754420A

[root@tpsrv300 ~]# cat /tmp/IBMLIB3-configuration.json | jq -r '.[[] | select(.sn == "754420A")
| .location + " - " + (.elementAddress|toString)'
drive_F5C4R1 - 269

[root@tpsrv300 ~]# cta-smc --json -q D | jq -r '.[[] | select(.driveOrdinal == 269) | .driveOrdinal'
12

[root@tpsrv300 ~]# grep IBM3 /etc/TPCONFIG
I3600541 IBM360 /dev/nst0 smc12
```

Spectra Logic TFinity tape library configuration

- Accessible in XML format via HTTP by using 3rd party `slapi` command line tool
 - Headers and trailers may need to be removed
 - Source: <https://github.com/LLNL/slapi>

- Important details of each tape drive:

- Frame, Drive Bay, Slot position
- SCSI element address
- Drive type and model
- Serial number

```
[root@tpsrv490 ~]# /usr/bin/slapi --config /root/SPECTRA/slapi.conf
                                --server spectralib1 drivelist
:
:
:
<?xml version="1.0" ?>
<drivelist>
  <drive>
    <ID>FR5/DBA3/fLT0-DRV3</ID>
    <driveStatus>OK</driveStatus>
    <partition>SPECTRALIB1</partition>
    <partitionDriveNumber>11</partitionDriveNumber>
    <driveType>IBM Ultrium-TD9 Full Height Fibre</driveType>
    <serialNumber>10330068FC</serialNumber>
    <manufacturerSerialNumber>1097001102</manufacturerSerialNumber>
    <driveFirmware>PA60</driveFirmware>
    <dcmfirmware>8.5.39</dcmfirmware>
    <wwn>21 33 00 90 A5 00 68 FC</wwn>
    <fibreAddress>soft</fibreAddress>
    <health>Green</health>
  </drive>
  :
  :
  :
```

Example of the steps to obtain SMC value for Spectra Logic Tfinity tape library

Make sure the received XML data is valid

Use the drive serial number to extract the drive partition number

Manual step: use slapi tool to extract information about all drives installed in the library

```
root@tpsrv490 ~]# /usr/bin/slapi --config /root/SPECTRALIB1 --server spectralib1 drivelist 2> /tmp/drivelist.xml
[root@tpsrv490 ~]# vim /tmp/drivelist.xml <<<--- REMOVING undesired headers and trailers
[root@tpsrv490 ~]# /usr/bin/sg_inq ` /usr/bin/lsscsi -g | /usr/bin/grep ' tape ' | /usr/bin/awk '{print $7}'`
| /usr/bin/grep 'Unit serial number: '
Unit serial number: 10330068FC
[root@tpsrv490 ~]# xmllint --xpath '/driveList/drive[serialNumber="10330068FC"][starts-with(partition, "SPECTRALIB1")]
/partitionDriveNumber/text()' /tmp/drivelist.xml
11
[root@tpsrv490 ~]# grep SPC1 /etc/TPCONFIG
S1L90533 SPC1L9 /dev/nst0 smc10
```

Decrement by 1 (because CTA RMC numbering starts from 0) to get the control path

New semi-automated system – overview

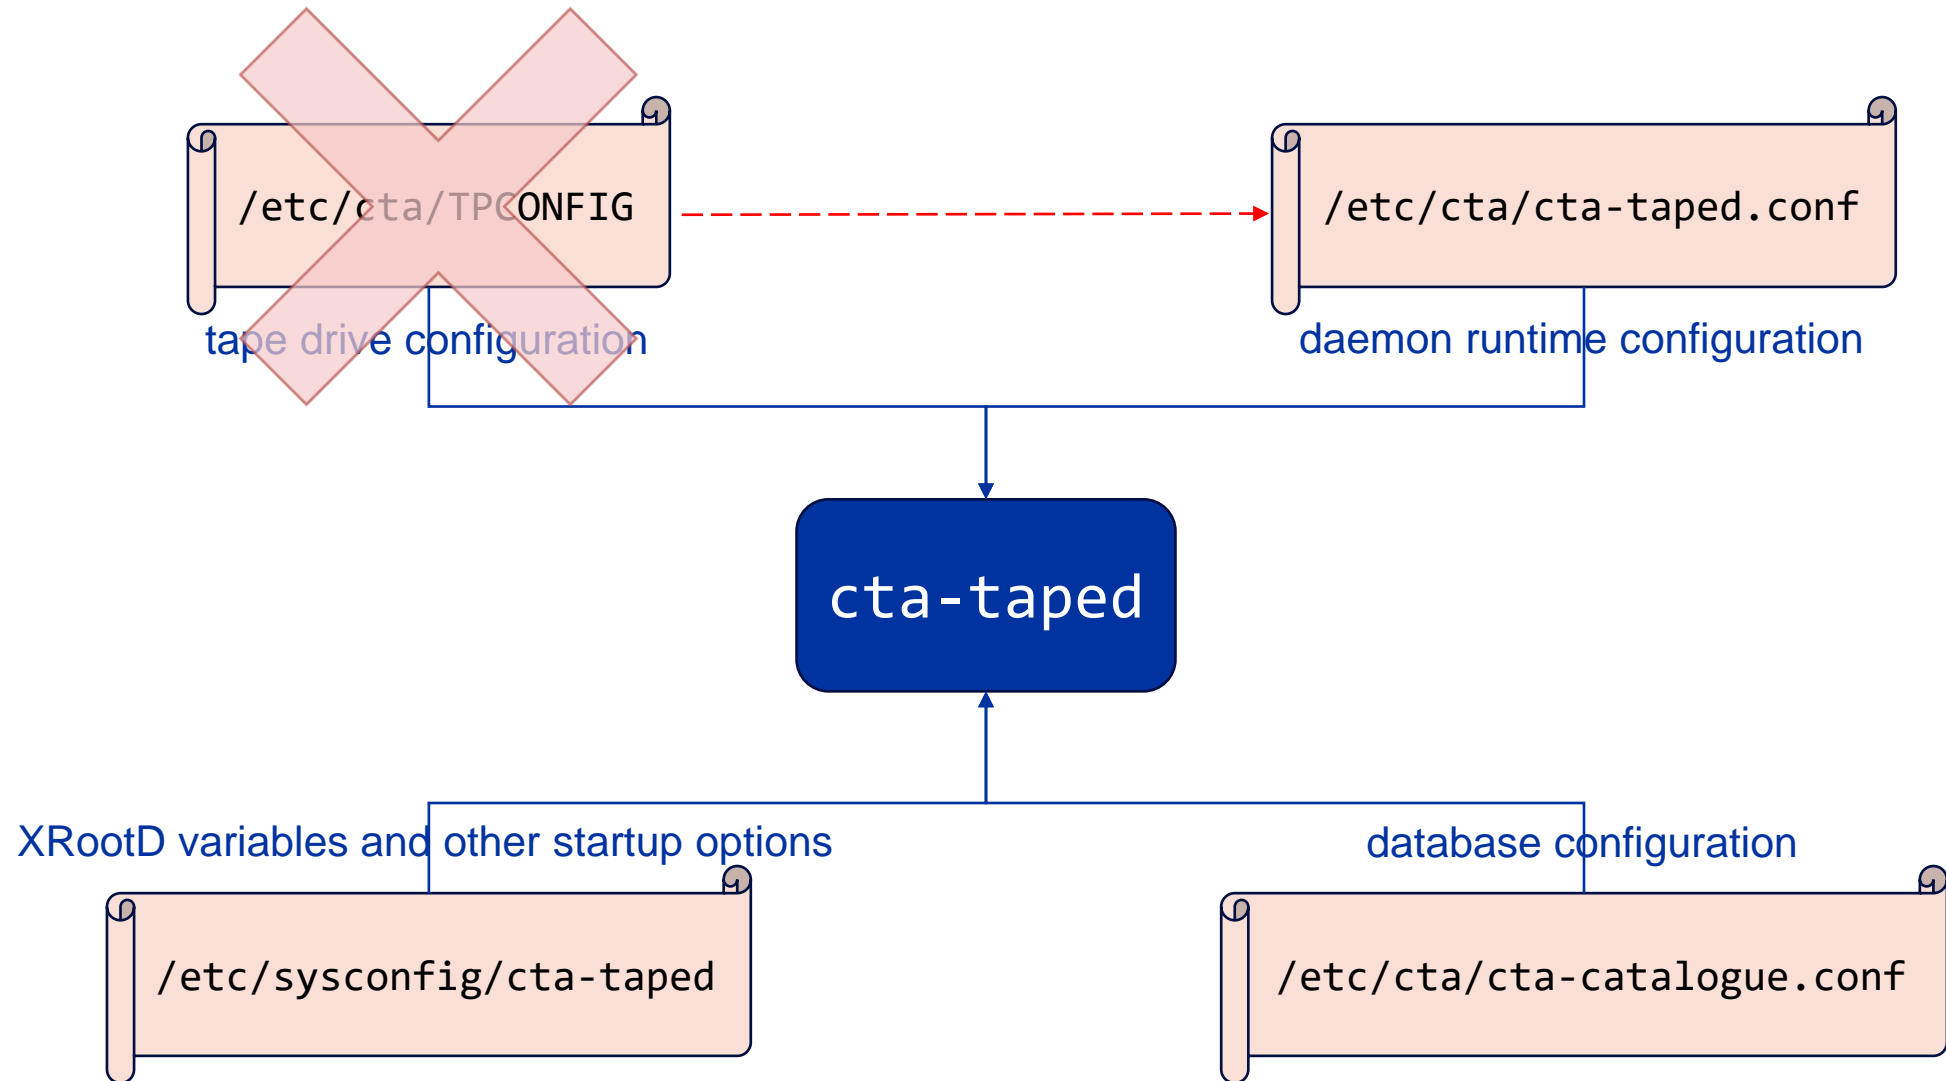
- Deploy JSON or XML library configuration files in the `/etc/cta-ops` directory on all tape servers
 - Preparation of those files is a manual step but they need to be updated very rarely (only when an existing library is modified (drives added / removed) or when a new one is installed = annually or less frequently)
- No need for separate tape drive naming conventions configuration, it is automatically generated by the new script in this format: `Library-DriveType-Position`
 - `IBMLIB4-TS1160-F14C2R1`
 - `SPECTRALIB2-LT09-F05B3S4`
- Constraint of this approach:
 - Do not connect tape drives from different libraries on the same tape server

New semi-automated system – sample script output

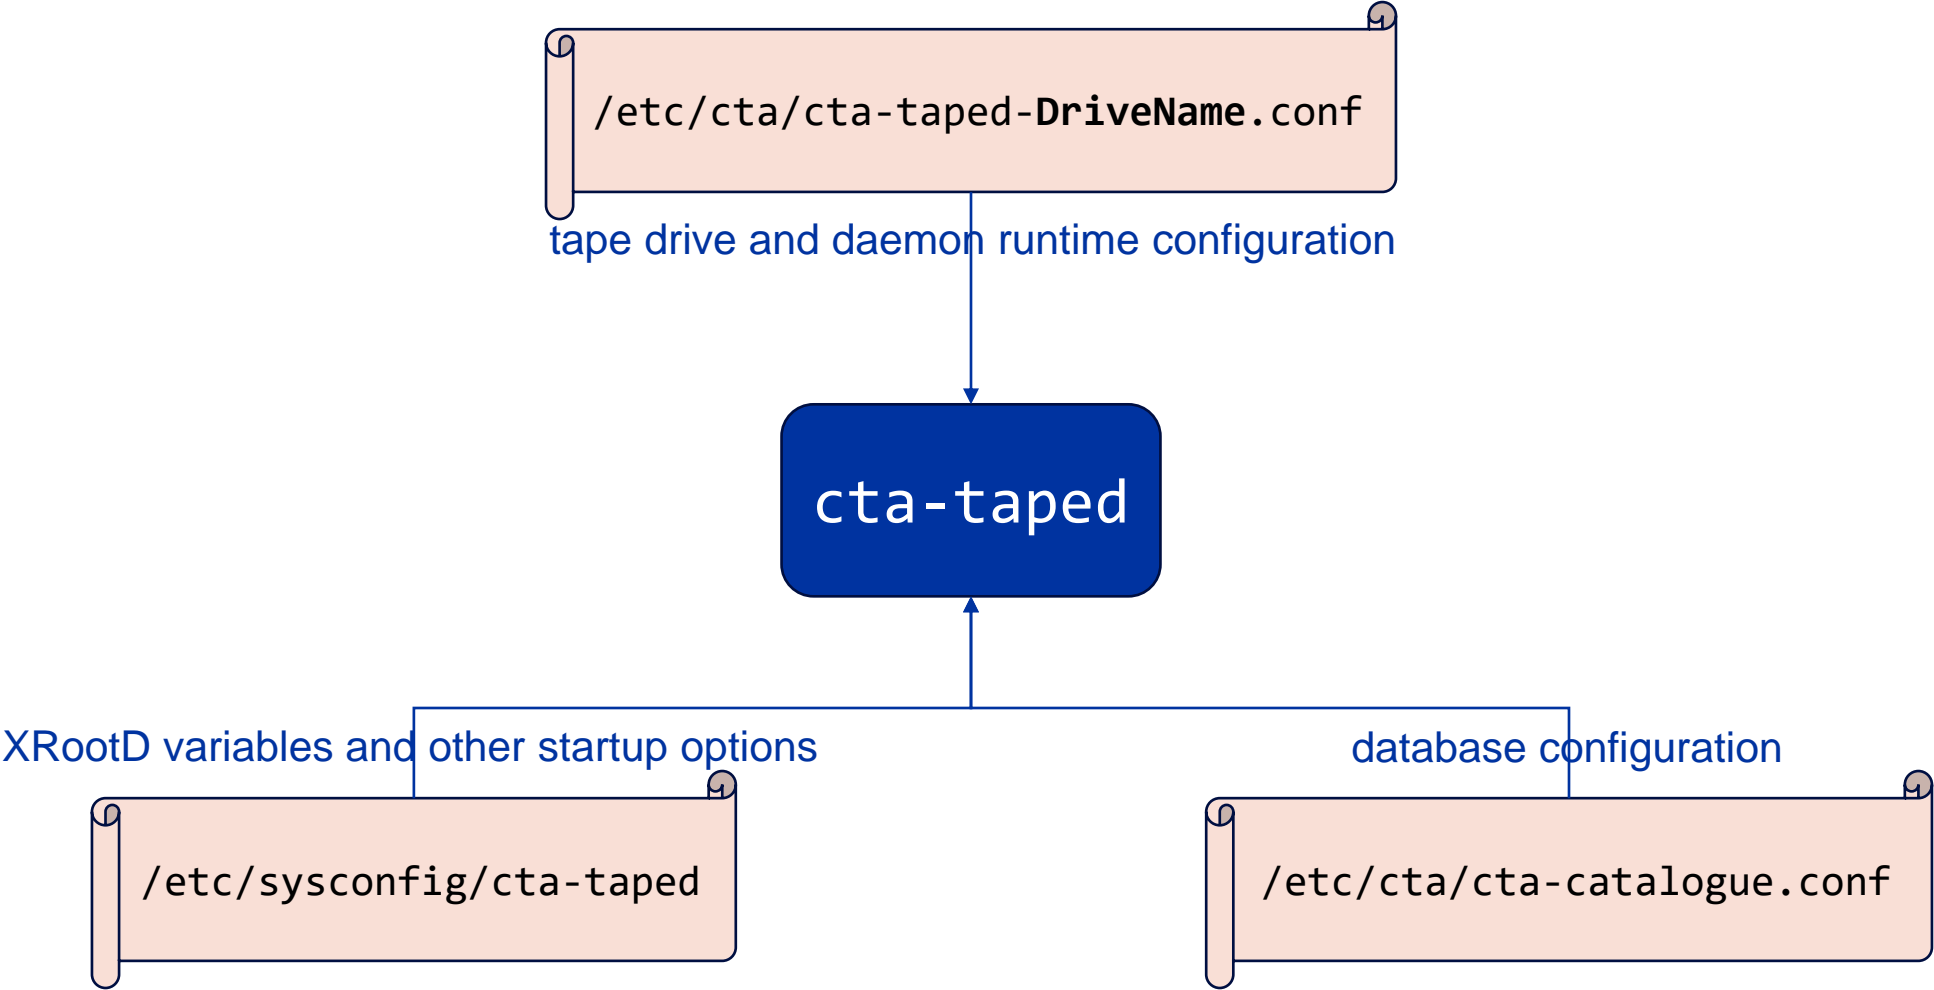
- Overview of the steps: *identify SCSI media changer device* → *detect the library serial number* → *detect all connected tape drives* → *map tape drives to st driver devices* → *combine all the data and generate the configuration*

```
[root@tpsrv454 /]# cta-ops-drive-config-generate
2024-02-19 16:48:16 [INFO] [get_scsi_media_changer] Detected first SCSI media changer device: /dev/sg1
2024-02-19 16:48:16 [INFO] [get_library_type] Detected tape library type and the serial number: IBM-03584L22-0000078AA8070401
2024-02-19 16:48:16 [INFO] [get_connected_drive_devices] Detected tape drive devices: ['/dev/sg0', '/dev/sg2']
2024-02-19 16:48:16 [INFO] [map_changer_device_to_nst] Tape drive /dev/sg0 with serial number 00000754423A is mapped to /dev/nst0
2024-02-19 16:48:16 [INFO] [map_changer_device_to_nst] Tape drive /dev/sg2 with serial number 0000075441DA is mapped to /dev/nst1
{
  "tape_drives": [
    {
      "logicalLibrary": "IBMLIB4-TS1160",
      "tapeDriveName": "IBMLIB4-TS1160-F2C1R4",
      "systemDevice": "/dev/nst0",
      "libraryControlPath": "smc3"
    },
    {
      "logicalLibrary": "IBMLIB4-TS1160",
      "tapeDriveName": "IBMLIB4-TS1160-F2C1R3",
      "systemDevice": "/dev/nst1",
      "libraryControlPath": "smc2"
    }
  ]
}
```

Current structure of cta-taped configuration files



New structure of cta-taped configuration files



New CTA's tape daemon configuration file

- Old TPCONFIG information will be moved to the cta-taped configuration file
- Each drive will require its own configuration file
- More granularity on drive configuration for multidrive setups
 - CTA drive parameters

```
[root@tpsrv454 ~]# cat /etc/cta/cta-taped-IBMLIB4-TS1160-F2C1R4.conf
[ ... ]
# DRIVE OPTIONS
#
# The name of the drive. Will be included for every line in the logs.
This can help identifying
# a problematic maintenance process of the drive after log agregation.
taped DriveName IBMLIB4-TS1160-F2C1R4

# CTA's logical library the tape drive will be linked to.
taped DriveLogicalLibrary IBMLIB4-TS1160

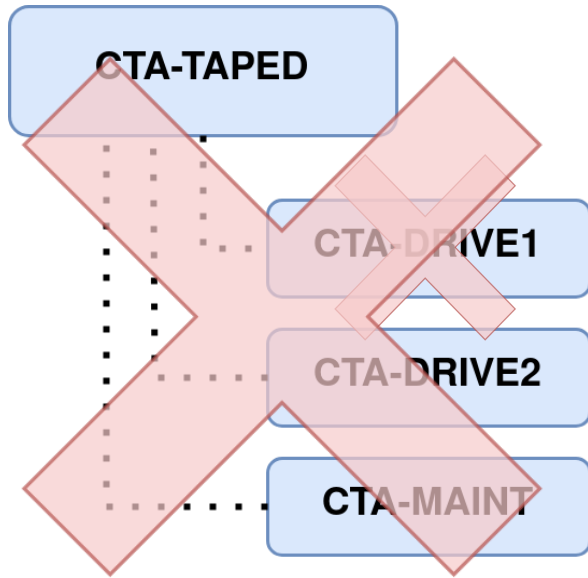
# Path to the character special device used to access the drive.
taped DriveDevice /dev/nst0

# The SCSI media changer address of the drive. This is "smc" + the
drive ordinal number of the device, which can be obtained with "cta-
smc -q D"
taped DriveControlPath smc3

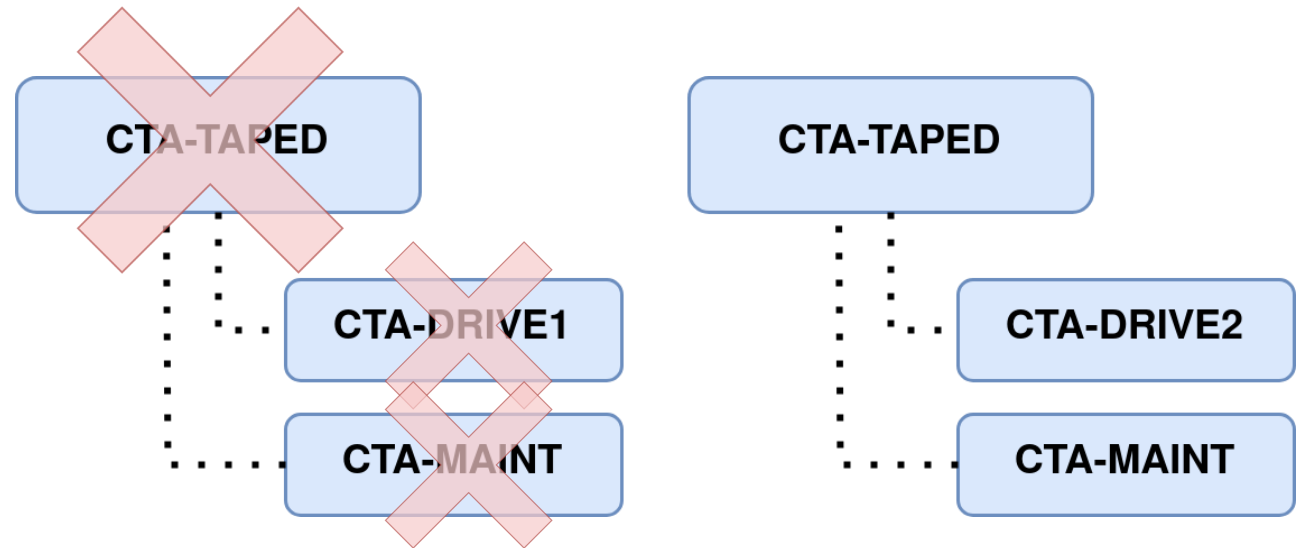
[ ... ]
```


New process topology

- Old topology



- New topology



- This new topology grants tape server independence from single drive failure
- Limit OS resources on a per process basis

New process topology – example

- Multiple drives running as independent processes

```
[root@tpsrv454 ~]# pstree -a | grep cta-
|-cta-rmcd /dev/smc
|-cta-tpd-master --log-format=json --log-to-file=/var/log/cta/cta-taped-IBMLIB4-TS1160-F2C1R3.log--config=/etc/cta/cta-taped-IBML
|   |-cta-tpd-IBMLIB4 --log-format=json --log-to-file=/var/log/cta/cta-taped-IBMLIB4-TS1160-F2C1R3.log...
|   |   `--83*[{cta-tpd-IBMLIB4}]
|   `--cta-tpd-maint --log-format=json --log-to-file=/var/log/cta/cta-taped-IBMLIB4-TS1160-F2C1R3.log...
|       `--83*[{cta-tpd-maint}]
|-cta-tpd-master --log-format=json --log-to-file=/var/log/cta/cta-taped-IBMLIB4-TS1160-F2C1R4.log--config=/etc/cta/cta-taped-IBML
|   |-cta-tpd-IBMLIB4 --log-format=json --log-to-file=/var/log/cta/cta-taped-IBMLIB4-TS1160-F2C1R4.log...
|   |   `--83*[{cta-tpd-IBMLIB4}]
|   `--cta-tpd-maint --log-format=json --log-to-file=/var/log/cta/cta-taped-IBMLIB4-TS1160-F2C1R4.log...
|       `--83*[{cta-tpd-maint}]
```

Migrating to the new setup

- Problem: tape and drive have to be in the same logical library, otherwise the tape will not get mounted
- The script to generate drive configuration has an option to map the new library names to the legacy format:
--use-legacy-mapping
- Possibility to deploy in two steps:
 1. Upgrade CTA taped while using old library names
 2. Update tape logical library field in the catalogue information and cta-taped-DriveName.conf files

```
[root@tpsrv454 ~]# cat /etc/cta-ops/cta-ops-config.yaml
[...]  
  legacy_library_mapping:  
    IBMLIB1-LT08: 'IBM1L8'  
    IBMLIB1-LT09: 'IBM1L9'  
    IBMLIB3-TS1160: 'IBM360'  
    IBMLIB3-TS1170: 'IBM370'  
    IBMLIB4-TS1160: 'IBM460'  
    IBMLIB4-TS1170: 'IBM470'  
    SPECTRALIB1-LT09: 'SPC1L9'  
    SPECTRALIB2-LT09: 'SPC2L9'  
[...]
```

Legacy library names

New naming convention


Migrating to the new setup – systemd services

- Moving from one service to pseudo-service with independent CTA taped services

```
[root@tpsrv450 ~]# cat /etc/systemd/system/cta-taped.service
[Unit]
Description=CERN Tape Archive (CTA) tape server daemon
After=syslog.target network-online.target

[Service]
EnvironmentFile=-/etc/sysconfig/cta-taped
ExecStart=/usr/bin/cta-taped
LimitCORE=infinity
OOMScoreAdjust=-1000
Type=forking
Restart=no

[Install]
WantedBy=default.target
```



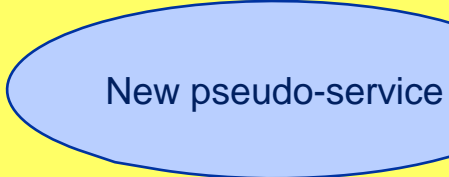
Migrating to the new setup – systemd services

- Moving from one service to pseudo-service with independent CTA taped services
 - 1 drive = 1 service. Service name: `cta-taped@DriveName`
- Each drive process will log to a different file
- Each drive service file points to its own configuration

```
[root@tpsrv454 ~]# cat /etc/systemd/system/cta-taped.service
[Unit]
Description=CERN Tape Archive (CTA) tape server daemon
After=syslog.targetnetwork-online.target

[Service]
EnvironmentFile=-/etc/sysconfig/cta-taped
ExecStart=/bin/true
RemainAfterExit=yes

[Install]
WantedBy=default.target
```



New systemd drive service example

```
[root@tpsrv454 ~]# cat /etc/systemd/system/cta-taped\@.service
[Unit]
Description=CERN Tape Archive (CTA) tape drive daemon for drive %i
PartOf=cta-taped.service
After=syslog.target network-online.target cta-taped.service

[Service]
EnvironmentFile=-/etc/sysconfig/cta-taped
ExecStart=/usr/bin/cta-taped --log-format=json --log-to-file=/var/log/cta/cta-taped-%i.log --
config=/etc/cta/cta-taped-%i.conf ${CTA_TAPED_OPTIONS}
LimitCORE=infinity
OOMScoreAdjust=-1000
Type=forking
Restart=no

[Install]
WantedBy=cta-taped.service
```

Summary

- Developed new semi-automated system to generate tape drive configuration
 - Uses information provided by the library as a source
 - Extracts the details using the tape drive serial number

- CTA's tape daemon refactored to better support configurations with multiple drives connected to a tape server

- Release status:
 - Will be available in the next CTA public release
 - Currently testing in pre-production at CERN

- Future Work:
 - Make CTA's tape daemon maintenance process independent from the CTA drive process

