

# CTA Operator Tools News

New tools and features for the CTA administrator


Richard Bachmann,  
on behalf of the CTA team



# What are the CTA Operations Utilities?

- Scripts to make tape operations easier
  - Automation
  - Monitoring
- Python code, pip packages
- FOSS


# General features



## Configuration

- One file
  - keytab location
  - cta-cli file
  - Library/tapepool setup
  - External binary selection for command execution
  - Configurable tabulation

# General features



## Configuration

- One file
  - keytab location
  - cta-cli file
  - Library/tapepool setup
  - External binary selection for command execution
  - Configurable tabulation

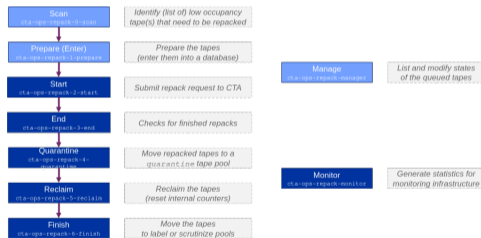
## Misc

- Shared libraries
  - Consistent behavior
- Translate `cta-admin` error messages
- Consistent logging and built-in rotation

# The Tools — ATRESYS

*Moving data between tapes at scale is cumbersome*

- **Automated Tape Repacking SYSTEM**
- Manages media transfer and re-use workflow
- Previously presented at [EOS 2023 Workshop](#)



# The Tools — CTA Ops Admin

*I wish I could easily extend the cta-admin command*

\$ `cta-ops-admin`

- Customisable wrapper for `cta-admin`
- Everything `cta-admin` can do
- Is 60% `cta-admin` json output + formatting

# The Tools — CTA Ops Admin

*I wish I could easily extend the cta-admin command*

\$ `cta-ops-admin`

- Customisable wrapper for `cta-admin`
- Everything `cta-admin` can do
- Is 60% `cta-admin` json output + formatting

## Queues

- \$ `cta-ops-admin showqueues`  
—Quick queue summaries

# The Tools — CTA Ops Admin

*I wish I could easily extend the cta-admin command*

## \$ cta-ops-admin

- Customisable wrapper for `cta-admin`
- Everything `cta-admin` can do
- Is 60% `cta-admin json` output + formatting

## Queues

- \$ `cta-ops-admin showqueues`  
—Quick queue summaries

## Tapes

- \$ `cta-ops-admin tape mount/unmount`
- \$ `cta-ops-admin tape mediacheck` —Check tapes for problems
- \$ `cta-ops-admin tape label` —Bulk labeling



# The Tools — CTA Ops Admin

*I wish I could easily extend the cta-admin command*

## \$ cta-ops-admin

- Customisable wrapper for `cta-admin`
- Everything `cta-admin` can do
- Is 60% `cta-admin` json output + formatting

## Queues

- \$ `cta-ops-admin showqueues`  
—Quick queue summaries

## Tapes

- \$ `cta-ops-admin tape mount/unmount`
- \$ `cta-ops-admin tape mediacheck` —Check tapes for problems
- \$ `cta-ops-admin tape label` —Bulk labeling

## Drives

- \$ `cta-ops-admin drive test` - Check drive health

# The Tools — CTA Ops Admin ShowQueues

*I need a summary of archive/retrieve activity*

## Example

```
$ cta-ops-admin showqueues archive
```

| TAPE_POOL    | PRIORITY | Qd_BYTES | Qd_FILES | MOUNTS | CURRENT_FILES | CURRENT_BYTES | MIN_AGE | OLDEST | TAPES_BYTES | TAPES_CAPACITY | TAPES_FILES |
|--------------|----------|----------|----------|--------|---------------|---------------|---------|--------|-------------|----------------|-------------|
| r_cms_mc     | 50       | 1.5T     | 487      | 1      | 497           | 1.8T          | 14400   | 3841   | 39.0P       | 38.0P          | 11633169    |
| r_alice_raw  | 50       | 1.3T     | 1621     | 1      | 12251         | 11.6T         | 14400   | 2793   | 141.3P      | 123.8P         | 65893108    |
| r_na62       | 2 50     | 1.2T     | 2277     | 6      | 184880        | 116.3T        | 14400   | 934    | 11.6P       | 11.0P          | 23274200    |
| r_atlas_prod | 2 50     | 345.5G   | 829      | 0      | 0             | 0B            | 14400   | 10218  | 34.6P       | 32.9P          | 18355362    |
| r_lhcb       | 50       | 10.7G    | 2        | 1      | 7             | 53.7G         | 14400   | 194    | 30.8P       | 26.3P          | 10701351    |
| ...          |          |          |          |        |               |               |         |        |             |                |             |

# The Tools — CTA Ops Admin Mediacheck

*Is this tape cartridge in good condition?*

## Example

```
$ cta-ops-admin tape mediacheck --drive I3601424 --vid L94871
```

```
Drive I1L90624 is DOWN and has a reason (<operator> 20240301 - Mediacheck on tape L94871) - OK
Source pool checked: OK
Checking if tape [L94871] is empty...
Tape [L94871] is empty: OK
Cartridge L94871 has capacity of 18000 GB and it will be filled with 3960 files(each with size of 5368709120 bytes)
Executing shell command: LANG=C /usr/bin/dd if=/dev/./.<random_file>.bin of=/dev/nst0 bs=256K 2>&1
WRITE finished - EOT hit - written 3464 files
Starting to READ 3463 files from tape L94871 ...
Executing shell command: LANG=C /usr/bin/dd if=/dev/nst0 of=/dev/./.<random_file>.bin-1 bs=256K 2>&1
Checksums are equal
Read back all 3463 files from tape, all checksums match.
Media on tape L94871 checked, EVERYTHING WENT FINE - MEDIA LOOKS GOOD.
----- TAPE MEDIACHECK SUMMARY -----
1 [L94871] - SUCCESS
```

# The Tools — CTA Ops Admin Mediacheck

*Is this tape cartridge in good condition?*

## Example

```
$ cta-ops-admin tape mediacheck --drive I3601424 --vid L94871
```

```
Drive I1L90624 is DOWN and has a reason (<operator> 20240301 - Mediacheck on tape L94871) - OK
Source pool checked: OK
Checking if tape [L94871] is empty...
Tape [L94871] is empty: OK
Cartridge L94871 has capacity of 18000 GB and it will be filled with 3960 files(each with size of 5368709120 bytes)
Executing shell command: LANG=C /usr/bin/dd if=/dev/./.<random_file>.bin of=/dev/nst0 bs=256K 2>&1
WRITE finished - EOT hit - written 3464 files
Starting to READ 3463 files from tape L94871 ...
Executing shell command: LANG=C /usr/bin/dd if=/dev/nst0 of=/dev/./.<random_file>.bin-1 bs=256K 2>&1
Checksums are equal
Read back all 3463 files from tape, all checksums match.
Media on tape L94871 checked, EVERYTHING WENT FINE - MEDIA LOOKS GOOD.
----- TAPE MEDIACHECK SUMMARY -----
1 [L94871] - SUCCESS
```

# The Tools — CTA Ops Admin Mediacheck

*Is this tape cartridge in good condition?*

## Example

```
$ cta-ops-admin tape mediacheck --drive I3601424 --vid L94871
```

```
Drive I1L90624 is DOWN and has a reason (<operator> 20240301 - Mediacheck on tape L94871) - OK
Source pool checked: OK
Checking if tape [L94871] is empty...
Tape [L94871] is empty: OK
Cartridge L94871 has capacity of 18000 GB and it will be filled with 3960 files(each with size of 5368709120 bytes)
Executing shell command: LANG=C /usr/bin/dd if=/dev/./.<random_file>.bin of=/dev/nst0 bs=256K 2>&1
WRITE finished - EOT hit - written 3464 files
Starting to READ 3463 files from tape L94871 ...
Executing shell command: LANG=C /usr/bin/dd if=/dev/nst0 of=/dev/./.<random_file>.bin-1 bs=256K 2>&1
Checksums are equal
Read back all 3463 files from tape, all checksums match.
Media on tape L94871 checked, EVERYTHING WENT FINE - MEDIA LOOKS GOOD.
----- TAPE MEDIACHECK SUMMARY -----
1 [L94871] - SUCCESS
```

# The Tools — CTA Ops Admin Mediacheck

*Is this tape cartridge in good condition?*

## Example

```
$ cta-ops-admin tape mediacheck --drive I3601424 --vid L94871
```

```
Drive I1L90624 is DOWN and has a reason (<operator> 20240301 - Mediacheck on tape L94871) - OK
Source pool checked: OK
Checking if tape [L94871] is empty...
Tape [L94871] is empty: OK
Cartridge L94871 has capacity of 18000 GB and it will be filled with 3960 files(each with size of 5368709120 bytes)
Executing shell command: LANG=C /usr/bin/dd if=/dev/./.<random_file>.bin of=/dev/nst0 bs=256K 2>&1
WRITE finished - EOT hit - written 3464 files
Starting to READ 3463 files from tape L94871 ...
Executing shell command: LANG=C /usr/bin/dd if=/dev/nst0 of=/dev/./.<random_file>.bin-1 bs=256K 2>&1
Checksums are equal
Read back all 3463 files from tape, all checksums match.
Media on tape L94871 checked, EVERYTHING WENT FINE - MEDIA LOOKS GOOD.
----- TAPE MEDIACHECK SUMMARY -----
1 [L94871] - SUCCESS
```

# The Tools — CTA Ops Admin Mediacheck

*Is this tape cartridge in good condition?*

## Example

```
$ cta-ops-admin tape mediacheck --drive I3601424 --vid L94871
```


```
Drive I1L90624 is DOWN and has a reason (<operator> 20240301 - Mediacheck on tape L94871) - OK
Source pool checked: OK
Checking if tape [L94871] is empty...
Tape [L94871] is empty: OK
Cartridge L94871 has capacity of 18000 GB and it will be filled with 3960 files(each with size of 5368709120 bytes)
Executing shell command: LANG=C /usr/bin/dd if=/dev/./.<random_file>.bin of=/dev/nst0 bs=256K 2>&1
WRITE finished - EOT hit - written 3464 files
Starting to READ 3463 files from tape L94871 ...
Executing shell command: LANG=C /usr/bin/dd if=/dev/nst0 of=/dev/./.<random_file>.bin-1 bs=256K 2>&1
Checksums are equal
Read back all 3463 files from tape, all checksums match.
Media on tape L94871 checked, EVERYTHING WENT FINE - MEDIA LOOKS GOOD.
----- TAPE MEDIACHECK SUMMARY -----
1 [L94871] - SUCCESS
```

# The Tools — Pool Supply

*I want to write data such we keep read back times high*


\$ `cta-ops-pool-supply`

- Gradually feeds empty tapes to tape pools used by drives
- Avoids excessive spread of files across media



XX:15

XX:30





# The Tools — Pool Supply

*I want to write data such we keep read-back times high*

## Example

```
$ cta-ops-pool-supply
```

```
Starting processing all tape pools at: 2024-03-01 08:48:02  
Tape pool: r_backup_hadoop which should have at least: 10 eligible partial tape(s) is supplied from: supply_alltypes  
Tape pool: r_backup_hadoop only has: 9 eligible partial tape(s) available, re-filling  
Identified: 50 supply tapes, moving random 1 to the pool: r_backup_hadoop  
Tape: L62683 moved to the pool: r_backup_hadoop  
...
```

# The Tools — Pool Supply

*I want to write data such we keep read-back times high*

## Example

```
$ cta-ops-pool-supply
```

```
Starting processing all tape pools at: 2024-03-01 08:48:02
```

```
Tape pool: r_backup_hadoop which should have at least: 10 eligible partial tape(s) is supplied from: supply_alltypes
```

```
Tape pool: r_backup_hadoop only has: 9 eligible partial tape(s) available, re-filling
```

```
Identified: 50 supply tapes, moving random 1 to the pool: r_backup_hadoop
```

```
Tape: L62683 moved to the pool: r_backup_hadoop
```

```
...
```

# The Tools — Pool Supply

*I want to write data such we keep read-back times high*

## Example

```
$ cta-ops-pool-supply
```

```
Starting processing all tape pools at: 2024-03-01 08:48:02  
Tape pool: r_backup_hadoop which should have at least: 10 eligible partial tape(s) is supplied from: supply_alltypes  
Tape pool: r_backup_hadoop only has: 9 eligible partial tape(s) available, re-filling  
Identified: 50 supply tapes, moving random 1 to the pool: r_backup_hadoop  
Tape: L62683 moved to the pool: r_backup_hadoop  
...
```

# The Tools — Pool Supply

*I want to write data such we keep read-back times high*

## Example

```
$ cta-ops-pool-supply
```

```
Starting processing all tape pools at: 2024-03-01 08:48:02  
Tape pool: r_backup_hadoop which should have at least: 10 eligible partial tape(s) is supplied from: supply_alltypes  
Tape pool: r_backup_hadoop only has: 9 eligible partial tape(s) available, re-filling  
Identified: 50 supply tapes, moving random 1 to the pool: r_backup_hadoop  
Tape: L62683 moved to the pool: r_backup_hadoop  
...
```


# The Tools — Tape Verification

*Are my files ok? Let's check in an automated fashion*

```
$ cta-ops-verify-tape
```

```
$ cta-ops-verification-feeder
```

- Picks a selection of files and retrieves them
- Sample-based error detection



# The Tools — Tape Verification

*Are my files ok? Let's check in an automated fashion*

## Example

```
$ cta-ops-verification-feeder ...
```

```
Verification for tape 197064 has finished since the last cta-ops-verification-feeder run
```

```
Currently running verification on tapes:
```

```
0 tapes are currently being verified, target is 20
```

```
42298 tapes are eligible for verification
```

```
--min_data_on_tape specified, selecting only tapes with at least 1000000000000 bytes written
```

```
After selecting only tapes with at least 1000000000000 bytes written, 42242 are eligible based on policy: random,  
the following new tapes have been selected for verification: L95548, I53420, L73747, ...
```

```
Submitting verification for tape 195548 (media type: lto9, logical library: ibm119, tape pool: r_alice_raw, total files
```

```
Tape 195548 successfully submitted for verification.
```

```
Waiting 120 seconds to start the next verification job
```

```
...
```

# The Tools — Tape Verification

*Are my files ok? Let's check in an automated fashion*

## Example

```
$ cta-ops-verification-feeder ...
```

```
Verification for tape 197064 has finished since the last cta-ops-verification-feeder run  
Currently running verification on tapes:
```

```
0 tapes are currently being verified, target is 20
```

```
42298 tapes are eligible for verification
```

```
--min_data_on_tape specified, selecting only tapes with at least 1000000000000 bytes written
```

```
After selecting only tapes with at least 1000000000000 bytes written, 42242 are eligible based on policy: random,  
the following new tapes have been selected for verification: L95548, I53420, L73747, ...
```

```
Submitting verification for tape 195548 (media type: lto9, logical library: ibm119, tape pool: r_alice_raw, total files
```

```
Tape 195548 successfully submitted for verification.
```

```
Waiting 120 seconds to start the next verification job
```

```
...
```

# The Tools — Tape Verification

*Are my files ok? Let's check in an automated fashion*

## Example

```
$ cta-ops-verification-feeder ...
```

```
Verification for tape 197064 has finished since the last cta-ops-verification-feeder run
```

```
Currently running verification on tapes:
```

```
0 tapes are currently being verified, target is 20
```

```
42298 tapes are eligible for verification
```

```
--min_data_on_tape specified, selecting only tapes with at least 1000000000000 bytes written
```

```
After selecting only tapes with at least 1000000000000 bytes written, 42242 are eligible based on policy: random,  
the following new tapes have been selected for verification: L95548, I53420, L73747, ...
```

```
Submitting verification for tape 195548 (media type: lto9, logical library: ibm119, tape pool: r_alice_raw, total files
```

```
Tape 195548 successfully submitted for verification.
```

```
Waiting 120 seconds to start the next verification job
```

```
...
```



# The Tools — Tape Verification

*Are my files ok? Let's check in an automated fashion*

## Example

```
$ cta-ops-verification-feeder ...
```

```
Verification for tape 197064 has finished since the last cta-ops-verification-feeder run
```

```
Currently running verification on tapes:
```

```
0 tapes are currently being verified, target is 20
```

```
42298 tapes are eligible for verification
```

```
--min_data_on_tape specified, selecting only tapes with at least 1000000000000 bytes written
```

```
After selecting only tapes with at least 1000000000000 bytes written, 42242 are eligible based on policy: random,
```

```
the following new tapes have been selected for verification: L95548, I53420, L73747, ...
```

```
Submitting verification for tape 195548 (media type: lto9, logical library: ibm119, tape pool: r_alice_raw, total files
```

```
Tape 195548 successfully submitted for verification.
```

```
Waiting 120 seconds to start the next verification job
```

```
...
```

# The Tools — Tape Drive Config Generator

*I need a solid device naming convention based on library data*



```
$ cta-ops-drive-config-generate
```

- Extracts drive information from library output files
- Encapsulates drive naming convention, generates unique names

# The Tools — Tape Drive Config Generator

*I need a solid device naming convention based on library data*



```
$ cta-ops-drive-config-generate
```

- Extracts drive information from library output files
- Encapsulates drive naming convention, generates unique names

```
$ cta-ops-drive-config-generate
```

```
{  
  "tape_drives": [  
    {  
      "DriveLogicalLibrary": "IBMLIB4-TS1160",  
      "DriveName": "IBMLIB4-TS1160-F2C1R4",  
      "DriveDevice": "/dev/nst0",  
      "DriveControlPath": "smc3"  
    }  
  ]  
}
```

# The Tools — CTA Ops EOS

**EXPERIMENTAL**

*I have complex EOS instance + CTA operations to perform*

- EOSCTA metadata operations
- Help manage *little EOS* disk buffers
- `$ cta-ops-get-path-from-eos`
  - Fetch EOS path of given file ID
- `$ cta-ops-change-storage-class`
  - Bulk change storage class of files in an EOS directory

# The Tools — Tape Alerting System (TAS)

*I need alerts and instant actions when something goes wrong*

\$ `cta-ops-tape-alerting-system`

- Automatic response to drive or media related issues
- Prevent damage, alert operators

**EXPERIMENTAL**

Dedicated talk:  
CTA Workshop 2024

# Summary and Q&A

*Upcoming release: 2.0*

- Repository: <https://gitlab.cern.ch/cta/cta-operations-utilities>
- Documentation: <https://cta.web.cern.ch/cta/pages/Documentation.html>

## Installation:

```
python3 -m pip install --extra-index-url  
https://cta-public-repo.web.cern.ch/cta-operations/pip/simple/ --requirement  
requirements.txt
```



# Option 1

VO -wmd=5, TP #p=3

