

Handling failed requests


Volodymyr Yurchenko, IT-SD-TAB

What are failed requests?

- Normal operations workflow is described in <https://indico.cern.ch/event/1227241/contributions/5348164/>
- Requests end up in QueueFailed after
 - 2 in-mount retries for archive jobs
 - 3 in-mount retries × 2 mounts for retrieve jobs
 - no retries for repack
- Remain in the archive buffer
- May indicate a problem in the system (hardware issue, critical bug)
- Can be listed by ``cta-admin fr ls``

CTA incident 17-18 July 2021 on EOSCTAPUBLIC

EOSCTA archive buffer





File transfer



Tape server




Orchestration




EOS headnode

Orchestration



CTA incident 17-18 July 2021 on EOSCTAPUBLIC


EOSCTA archive buffer



File transfer



Tape server







EOS headnode

20-07-2021 11:20:34
Related ticket: [INC2859801](#)
Dear EOSCTA Team,
I have 3 defective DIMM on your node EOSCTAFST0114 as you can see below

CTA incident 17-18 July 2021 on EOSCTAPUBLIC

EOSCTA archive buffer

Tape server



EOS headnode

How to recover?

- **Ask user** to re-transfer
- **Reinject**: send CLOSEW event to CTA to trigger archiving of a disk file
 - enabled by EOS WorkFlow Engine
 - more details on EOS+CTA workflows:
<https://indico.cern.ch/event/985953/contributions/4238328/>
- Or **ignore** (e.g. repack failed requests)

Reasons for reinjection

- ① Reduce the number of transfers made by user
- ② Free up archive buffer space on shared instances
- ③ Quickly recover from the incidents

Scripts to handle failed requests



classify.sh



reinject.sh

1. Dump requests to a json file
2. Filter out files that:
 - were deleted by user
 - were overwritten by user
 - have wrong fileID (but same path and not on tape)
 - successfully written to tape
 - have valid request still ongoing
 - request the second copy
 - have no archiveID - to be reinjected!

Scripts to handle failed requests



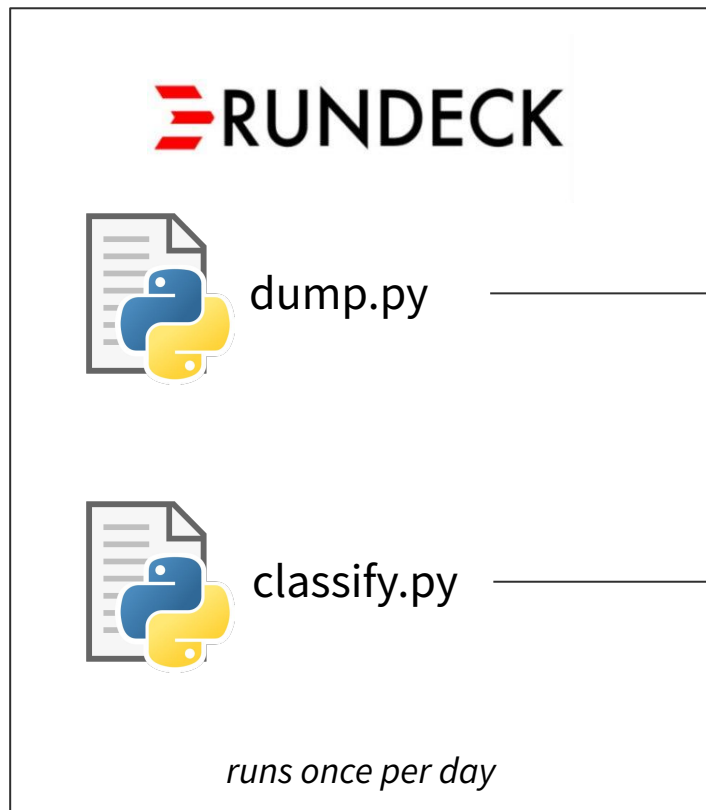
classify.sh



reinject.sh

1. Sanity checks
2. Delete
`'sys.cta.archive.objectstore.id'`
extended attribute
3. Send CLOSEW event

Automated classification



 RUNDECK



`dump.py`

Dump failed requests to a file and upload it to the CTA project space in EOS



`classify.py`

Filter out files for reinjection + upload the list to CTA project space in EOS

- TODO: upload directly to the monitoring DB

runs once per day

Next steps

- Integrate the classification scripts with the monitoring system
- Make them publicly available

