

# **EOS archive tool**

**2024**

**Elvin Sindrilaru**

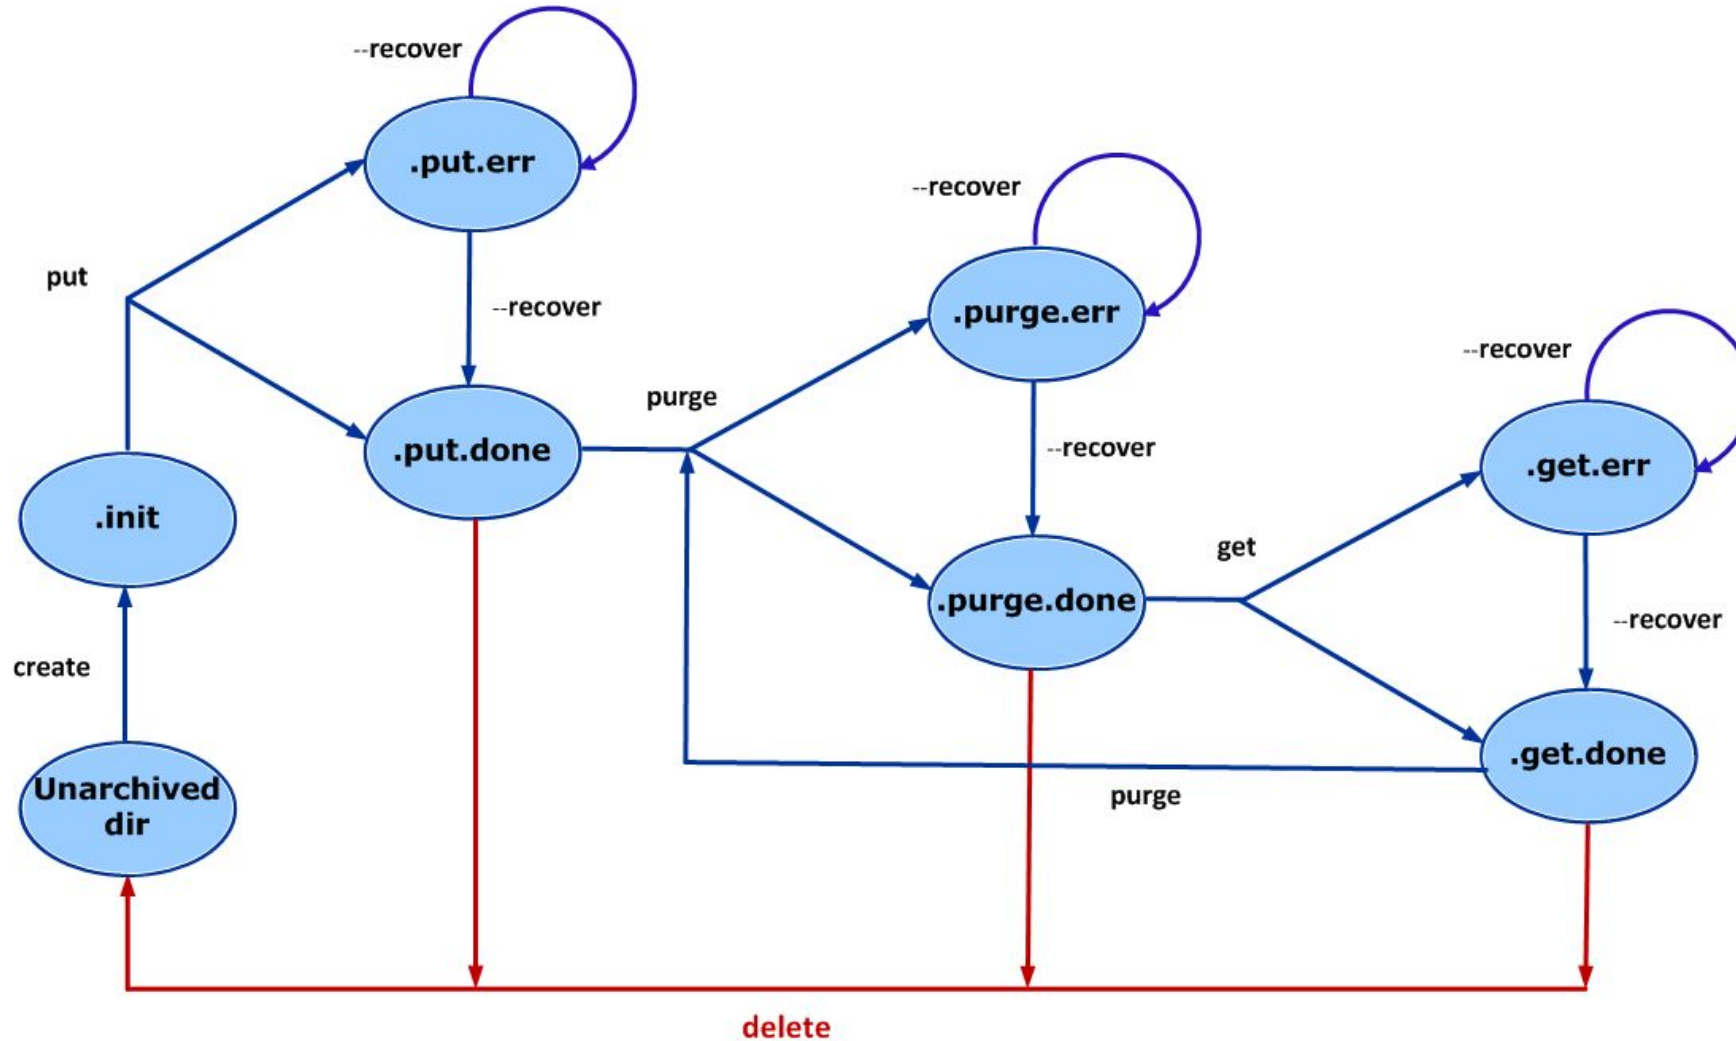
# Why an archive tool?

- **EOS quota limits** - free up storage space in EOS for users' online activities
- Spare users from developing *ad hoc* archiving solutions
- Manage **efficiently** the movement of data between disk and archive storage
- Have control over the archiving rules and easily tune the configuration

# EOS archive CLI

- **eos archive** <subcmd>
  - **create** <eos\_path> <archive\_location>
    - Create archive file for <eos\_path> directory whose contents will be archived at: <archive\_location>
  - **put** [--recover] <eos\_path>
    - Copy the contents of the EOS directory to the archive destination.
  - **get** [--recover] <eos\_path>
    - Copy the contents of the archive back to EOS.
  - **purge** [--recover] <eos\_path>
    - Delete all disk resident data keeping the archive file and the archived data.
  - **list** [all|put|get|job\_uuid]
    - List ongoing transfers. By default, all transfers are listed.
  - **delete** <eos\_path>
    - Delete all archived data including the archive file - **Admin command**.

# Archiving workflow



# Archive and log files

- **Archive file** – contains entries in JSON format
  - Header (source, destination, archive size etc.)
  - Directory/file EOS metadata information in JSON format
  - Never modified during the lifetime of an archive
  - Can be used in the future to get information about the contents of the archive
  - Archive “get” restores files in the original layout (2 replicas, RAIN etc.)
- **Archive log file – archive.log**
  - Summary of the last executed transfer
  - Hints to why a transfer has failed
  - Users can/should access it in case of errors

# Archiving internals

- **Archive transfer operation**

- Separate process handling in a multi-threaded way the **third-party copy operations, integrity checks, directory structure creation** etc.
- Only users in the **ACLs with the a - archive flag** can perform archive operations on the given path
- **Wait on tape** - jobs waits indefinitely for successfully transferred file to be migrated to tape
  - Helped us in the past to spot issues on the CTA side
- **Limitations** - user gets an error when trying to create the archive
  - can not deal with 0 size file (CTA)
  - EOS special internal files i.e. atomic/version files

- **Archive configuration**

- Control the level of parallelism of the archive transfers from the `/etc/eosarchiverd.conf`
  - `MAX_TRANSFERS=10` (default) - max number of transfers (process) to run in parallel, the other are queued
  - `MAX_THREADS=5` (default) - max number of threads per process doing work
  - `BATCH_SIZE=10` (default) - max batch size of ops to be performed per thread
  - `MAX_RETRIES=5` (default) - max number of retries when a batch fails - protect against transient failures

# Archiving in practice

- What is left on EOS after a successful archive?

```
[root@eoscms-ns-ip563 (mgm:master mq:master) ~]$ eos archive list | grep SMP_20
|Thu Feb 24 16:28:37 2022      |/eos/cms/store/group/phys_smp/archival/2022-02-16-SMP_20_005/
|purge done                |
[root@eoscms-ns-ip563 (mgm:master mq:master) ~]$ eos ls -lrta /eos/cms/store/group/phys_smp/archival/2022-02-16-SMP_20_005/
drwxr-xr--+ 1 pgras zh          18077344 Feb 24 2022 .
drwxr-xr--+ 1 root  zh          524708304 Oct 12 2022 ..
-rw-r--r--  2 pgras zh          9686603 Feb 24 2022 .archive.log
-rw-r--r--  2 root  root        8390741 Feb 16 2022 .archive.purge.done
[root@eoscms-ns-ip563 (mgm:master mq:master) ~]$ head -4 /var/tmp/archive.info
{"src": "root://eoscms.cern.ch/eos/cms/store/group/phys_smp/archival/2022-02-16-SMP_20_005/", "dst": "root://eosctacms.cern.ch/eos/ctacms/archive/archive/cms/cb2c697cb0f42cda04ee116576fe1b928ea7ba067f2e28d508a07a125b67aa4c/", "svc_class": "archive", "dir_meta": [{"uid", "gid", "mode", "attr"}, {"file_meta": [{"size", "mtime", "ctime", "uid", "gid", "mode", "xstype", "xs"}], "uid": "2814", "gid": "1399", "timestamp": 1645005866, "num_dirs": 1040, "num_files": 31696}
["d", "/", "2814", "1399", "40755", {"sys.acl": "g:1399:!d,egroup:cms-eos-pag-smp-cleaners:rwa+d,z:i", "sys.eos.btime": "1645005655.650981459", "sys.forced.blockchecksum": "crc32c", "sys.forced.blocksize": "4k", "sys.forced.checksum": "adler", "sys.forced.layout": "replica", "sys.forced.nstrips": "2", "sys.forced.space": "default", "sys.recycle": "/eos/cms/proc/recycle/", "user.acl": ""}]
["d", "SMP_20_005_archive/", "51667", "1399", "42755", {"sys.accounting.vos.0": "cms", "sys.acl": "u:22014:rw,u:31275:rw,u:5410:rw,g:1399:!d,egroup:cms-eos-pag-smp-cleaners:rw+d,z:i", "sys.eos.btime": "1644923408.890932465", "sys.forced.blockchecksum": "crc32c", "sys.forced.blocksize": "4k", "sys.forced.checksum": "adler", "sys.forced.layout": "replica", "sys.forced.nstrips": "2", "sys.forced.space": "default", "sys.recycle": "/eos/cms/proc/recycle/", "user.acl": ""}]
["d", "SMP_20_005_archive/wgamma-miniaod-genonly-2016/", "51667", "1399", "42755", {"sys.accounting.vos.0": "cms", "sys.acl": "u:22014:rw,u:31275:rw,u:5410:rw,g:1399:!d,egroup:cms-eos-pag-smp-cleaners:rw+d,z:i", "sys.eos.btime": "1585155291.964649406", "sys.forced.blockchecksum": "crc32c", "sys.forced.blocksize": "4k", "sys.forced.checksum": "adler", "sys.forced.layout": "replica", "sys.forced.nstrips": "2", "sys.forced.space": "default", "sys.recycle": "/eos/cms/proc/recycle/"}
```

- What is left on EOS after a failure?
  - All the initial files plus the archive file and logs.



[home.cern](https://home.cern)