

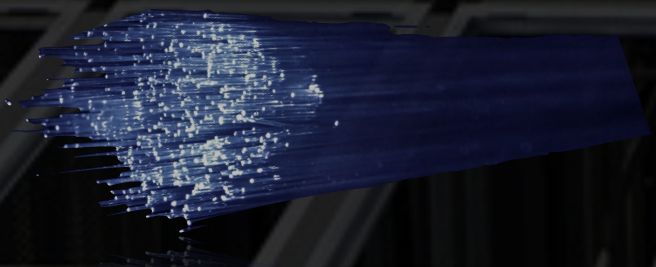
8th March 2022

CERN

CEBN

EOS migration tools

motivated by AMS experiment migration

A bundle of fiber optic cables is shown, with light reflecting off the ends, creating a starburst effect.

Dr. Jaroslav Guenther
on behalf of EOS PDS team


(CERN IT-ST-PDS)

(CEBN IT-ST-PDS)

Objective and Tools

EOS space migration:

- > data: directories, files, links
- > metadata: permissions, attributes
- > instance and space settings (quotas, access rights etc.)
- > sync, freeze source and redirect all clients to use new location



Low level tools:

- > eos cp, xrdcp, rsync, fusex, eos-ns-inspect
- > Incomplete/no attribute or permission transfer
- > requires special care per migration item type
- > need for higher level tool automating the migration process

Procedure pitfalls:

- > **very diverse file occupancy among directories**
(e.g. one can find $O(1M)$ very small files in 1 dir)
- > **user permission inconsistencies (posix/ACL)**
(users not allowed to browse or write to the directory where they own files)
- > **interference with durability tools (fsck)**
(problems during migration process may be difficult to spot)
- > **random data/metadata corruption at source and destination**
(HW/SW problems)
- > **human factor**
Do you really want to run this from master headnode ?

Some of the automation pitfalls:

- > **scale** (e.g. metadata volume to keep in sync)
- > **sync time** (in case of no source freeze)
(AMS 60M files, 400k directories, 13PB)
- > **problem catching, reporting and fixing**
(many migration jobs)
- > **low level tool issues**
(e.g. string byte decoding in python)



Existing tools



Per username migration tools:

- > repository
- > migration item = user
- > `/eos/user/<letter>/<username>`
- > no long term automatic sync functionality
- > philosophy: migrate and switch “in one go”

General migration tools:

- > different approaches being developed
- > eos-instance-migration tool (docs)
 - * migrating any path/dir/file/link/quota
 - * philosophy:
 - migrate
 - keep in sync
 - switch whenever convenient
- > EOS fusex could provide a general solution as well (extended attribute support in development)

The screenshot shows a web page titled "EOS Instance Migration Documentation" with a search bar. The main heading is "EOS Migration Tools". The text explains that several tools have been developed for various use cases, including archiving and usage targeting. It mentions the AMS02 migration from EOSPUBLIC to EOSAMS02 (2021-2022) and the development of a new migration tool. The tool is written in Python and provides options for quota, directory tree, or file migration. The page also lists the tool's components, including the main executable `migration_daemon.py`.



New migration tool

Migration steps:

- 1.migrate all directories
- 2.migrate eos quota settings
- 3.directory sync and file sync
- 4.automatic sync replays

Tool Configuration:

- > run options (emailing, file names)
- > access to EOS instances
- > attribute check settings
- > dir/file/link removal behaviour
- > quota node list
- > sync optimisation options
- > Thread settings
- > Timeouts
- > check the docs



The screenshot shows a GitHub README page for the 'EOS instance migration tool'. At the top, there is a dark blue header with a GitHub logo and a 'Menu' button. Below the header, the file name 'README.md' is displayed. The main content area features the EOS logo and the title 'EOS instance migration tool'. A line of text says 'Check the [documentation](#) for more info.' Below this, there are two sections: 'Dependencies' and 'Repository structure'. The 'Dependencies' section lists: 'python (v3.7+)', 'eos-ns-inspect tool (part of EOS v4.8.57 and later only is optimised for this use case)', and 'optional: pip3 instal ijson'. The 'Repository structure' section lists: 'docs: content for [documentation](#) for more info.', 'code:', and a list of Python files with their functions: 'migration_daemon.py: main executable', 'migrate.py: prepare and run migration tasks', 'namespace.py: query namespaces', 'xcheck.py: miration item cross checking functionality between source and destination', 'transfer.py: data transfer functionality (for directories/files/links)', 'utils.py: utility functions', 'quota.py: quota migration methods', 'decode_json_lines.py: if namespace output requested as json, this module will be called, requires pip3 instal ijson (allowing to read big json line by line)', 'accounting.py: general accounting functionality, consolidation of job reports', 'logsetup.py: helper methods for logging', and 'config.py: configure the migration tool (expected to be modified according to use-case)'. At the bottom, there is a 'tests: testing scripts (unit tests supposed to come here)'.

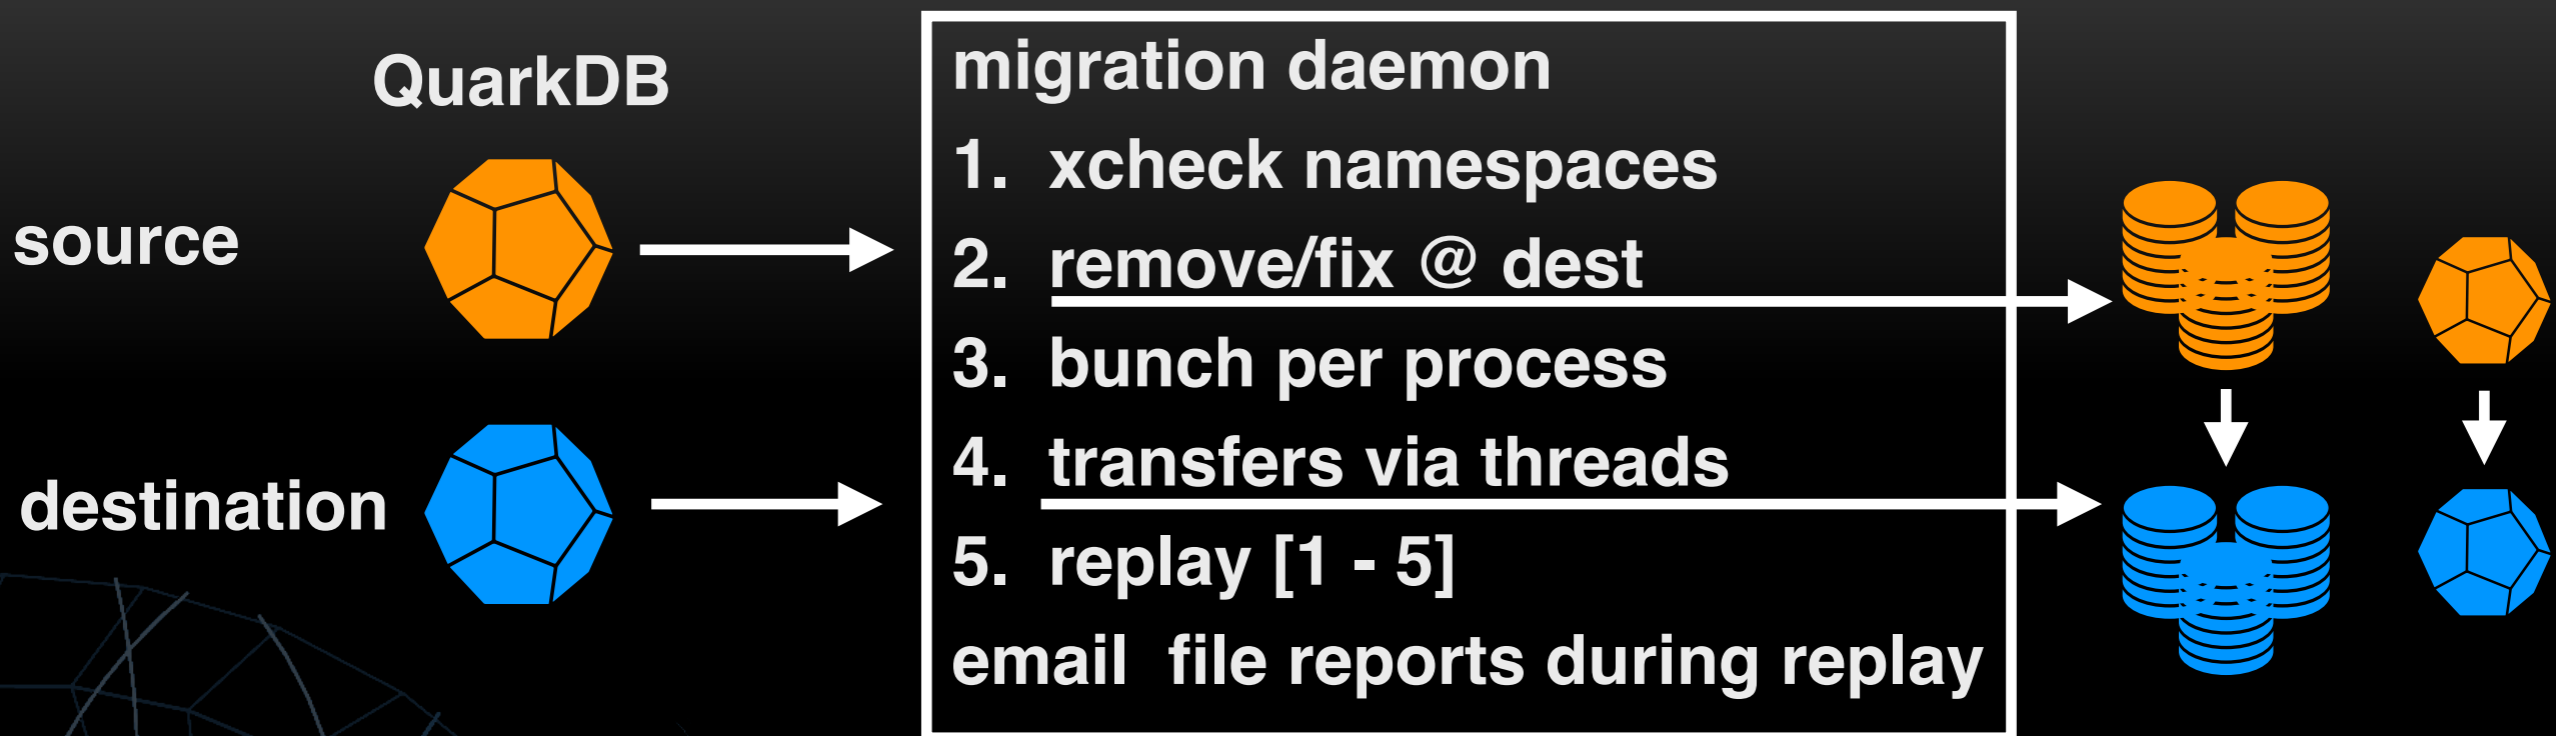


Functionality



Operator Input:

- > configuration file
- > root directory to migrate
- > migration script output path



Error reporting:

- > in categorised files
 1. namespace check and filter categories
 2. transfer crash categories
- > check the docs



Directory migration (hands-on)



- > Migrating from EOSPILOT to EOSAMS02
- > Instances setup for migration (TPC, keytab files, access)
- > directory order is important - processing sequential processing for the moment
- > 1 process for namespace xcheck and 1 for directory processing

```
[root@eosams02-ip554 (mgm:slave mq:slave) migtest]$ workdir=/var/migtest/
[root@eosams02-ip554 (mgm:slave mq:slave) migtest]$ codedir=/var/migtest/code/
[root@eosams02-ip554 (mgm:slave mq:slave) migtest]$ rundir=/var/migtest/run/
[root@eosams02-ip554 (mgm:slave mq:slave) migtest]$ migrate=$codedir/eos-instance-migration/code/
migration_daemon.py
[root@eosams02-ip554 (mgm:slave mq:slave) migtest]$ cd run

[root@eosams02-ip554 (mgm:slave mq:slave) run]$ rootpth=/eos/pilot/cms/tier0/store/backfill/1/data/;
[root@eosams02-ip554 (mgm:slave mq:slave) run]$ python3.7 ${migrate} --root-nsdir $rootpth --sync-dirs-only
--output-path ${rundir} > ${rundir}/ams.log 2>&1 &

[root@eosams02-ip554 (mgm:slave mq:slave) run]$ less ams.log
```



eosams02-ip554-migration@cern.ch

/eos/pilot/cms/tier0/store/backfill/1/data/: DIR migration finished

To: Undisclosed recipients;

1 directories were queued for processing, from these:
- 0 initially matched the source
(no jobs ran for them)
- 1 were (at least partially) migrated
-- 1 were completely migrated
0 directories have a timestamp of complete *file* migration


Subdirectories of the directories above were migrated,
their counts were:
- 244 queued for processing
-- 244 migrated successfully
--- 244 part of completely migrated directories
-- 0 crashed
-- 0 left to be processed (from the queued ones at the moment)

```
[root@eosams02-ip554 (mgm:slave mq:slave) code]$ ll ../run/
total 232
-rw-r--r--. 1 root root 2227 Mar 6 05:59 ams.log
drwxr-xr-x. 6 root root 4096 Mar 6 04:48 dir_bunches
-rw-r--r--. 1 root root 3400 Mar 6 04:48 DIR_migration_summary_report.json
drwxr-xr-x. 2 root root 4096 Mar 6 03:56 dir_reports
-rw-r--r--. 1 root root 2 Mar 6 03:55 dir_sync_timestamps.json
-rw-r--r--. 1 root root 216542 Mar 6 04:48 file_processing_dirs.json
```



Directory migration (hands-on)



 eosams02-ip554-migration@cern.ch
/eos/pilot/cms/tier0/store/backfill/1/data/: DIR migration summary
To: Undisclosed recipients;

DIR migration summary for 244 entries from
/eos/pilot/cms/tier0/store/backfill/1/data/

Last job attempt summary:

- 244 entries seen in bunches for processing
- 244 should equal the above (sanity check)
- 0 crashed
- 244 were processed,
- 244 should equal the above (sanity check)
- 0 remain to process (job omissions)
- 0 should equal the above (sanity check)
- 0 tpc remote read failed
- 0 crashed due to entries existing at destination
- 0 crashed due to i/o error at source
- 0 crashed due to transfer operation expired
- 0 crashed due to entries disappearing from the source namespace during migration
- 0 crashed due to entries being broken at source
- 0 crashed due to insufficient permissions at source (/entries could be broken)
- 0 crashed due to network being unreachable
- 0 crashed in setting attributes
- 0 crashed in setting mode
- 0 crashed in setting ownership
- 0 crashes of unknown type
- 0 unaccounted crashes (xcheck jobs vs category summary)

Processing of: 0.0 GB (retries not accounted) took (retries included) 0.04215920454925961 hours.
The transfers took: 0.030555555555555555 hours (correct job average)
Average file size: 0.0 MB.
Average transfer speed: 0.0 MB/s (correct job average)
Job report can be found in:
/tmp/migttest/run/DIR_migration_summary_report.json.

Initial namespace exploration found the following items:

- 244/0 [src/dest] totals seen
- 0/0 [src/dest] items from directories explicitly excluded by operator
- 0/0 [src/dest] failed prefiltering (see below [1])
- 0 [dest] failed prefiltering and were removed
(not on storage, ignored, decode errors, zero size/xs corrupted)
- 0 [dest] remained failed (removal omissions)
- 244/0 [src/dest] items passed the initial filter (see below [1])
- 0/0 [src/dest] items with matching all MD between the source and destination
- 0 [dest] had bad metadata (attribute, ownership, posix perm) crosscheck (size and checksum were O.K.)
- 0 [dest] removed due to bad size or xs
- 0 [dest] remained bad (removal omissions)
- 0 [dest] not present at source:
- 0 [dest] removed
- 0 [dest](removal omissions)
- 0 [src] were skipped (fully migrated already)
- 0/0 [src/dest] unaccounted items (sanity check, should be 0)
- 0/0 [src/dest] items with uid without a username (included in migration)
- 244 [src] were selected for processing input

[1]

Failed cases at source:

```
{
  "ignored": 0,
  "decode_errors": 0,
  "zero_size_xs_corrupted": 0,
  "no_storage_only_in_namespace": 0
}
```

Failed cases at destination:

```
{
  "ignored": 0,
  "decode_errors": 0,
  "zero_size_xs_corrupted": 0,
  "no_storage_only_in_namespace": 0
}
```

End of migration summary.

```
[root@eosams02-ip554 (mgm:slave mq:slave) ~]$ eos ls -l /eos/pilot/cms/tier0/store/
backfill/1/data/tier0_REPLAY_2021/HLTPhysics
drwxr-x--- 1 cmst0 zh 0 Mar 6 04:46 ALCARECO
drwxr-x--- 1 cmst0 zh 0 Mar 6 04:46 AOD
drwxr-x--- 1 cmst0 zh 0 Mar 6 04:46 DQMIO
drwxr-x--- 1 cmst0 zh 0 Mar 6 04:46 MINIAOD
drwxr-x--- 1 cmst0 zh 0 Mar 6 04:46 RAW
drwxr-x--- 1 cmst0 zh 0 Mar 6 04:46 RAW-RECO
drwxr-x--- 1 cmst0 zh 0 Mar 6 04:46 USER
```

```
[root@eospiilot-ns-ip700 (mgm:master mq:slave) ~]$ eos ls -lh /eos/pilot/cms/tier0/store
/backfill/1/data/tier0_REPLAY_2021/HLTPhysics
drwxr-x--- 1 cmst0 zh 7.03 M Nov 22 16:15 ALCARECO
drwxr-x--- 1 cmst0 zh 8.55 G Nov 22 16:26 AOD
drwxr-x--- 1 cmst0 zh 1.42 G Nov 22 16:15 DQMIO
drwxr-x--- 1 cmst0 zh 1.93 G Nov 22 16:29 MINIAOD
drwxr-x--- 1 cmst0 zh 109.39 G Nov 16 12:35 RAW
drwxr-x--- 1 cmst0 zh 35.56 G Nov 16 14:44 RAW-RECO
drwxr-x--- 1 cmst0 zh 44.89 M Nov 22 16:07 USER
```

```
[root@eosams02-ip554 (mgm:slave mq:slave) ~]$ eos attr ls /eos/pilot/cms/tier0/store/
backfill/1/data/tier0_REPLAY_2021/HLTPhysics
sys.acl="g:1399:rx,u:103031:rx,u:8959:rw,u:22014:rw"
sys.eos.btime="1646538352.279470950"
sys.forced.atomic="1"
sys.forced.blocksize="4k"
sys.forced.checksum="adler"
sys.forced.layout="replica"
sys.forced.nstrips="2"
sys.forced.space="default"
sys.lru.expire.match="*.root:15d"
```

```
[root@eospiilot-ns-ip700 (mgm:master mq:slave) ~]$ eos attr ls /eos/pilot/cms/tier0/
store/backfill/1/data/Tier0_REPLAY_2021/HLTPhysics
sys.acl="g:1399:rx,u:103031:rx,u:8959:rw,u:22014:rw"
sys.eos.btime="1637062517.44885956"
sys.forced.atomic="1"
sys.forced.blocksize="4k"
sys.forced.checksum="adler"
sys.forced.layout="replica"
sys.forced.nstrips="2"
sys.forced.space="default"
sys.lru.expire.match="*.root:15d"
```



Quota migration (hands-on)



```
[root@eosams02-ip554 (mgm:slave mq:slave) run]$ python3.7 ${migrate} --root-nsdir $rootpth --set-quota --output-path ${rundir} > ${rundir}/ams.log 2>&1 &
```

```
[root@eospiilot-ns-ip700 (mgm:master mq:slave) ~]$ eos quota ls /eos/pilot/cms/tier0/
```

```
➤ Quota Node: /eos/pilot/cms/tier0/
```

user	used bytes	logi bytes	used files	aval bytes	aval logib	aval files	filled[%]	vol-status	ino-status
cmst0	111.16 TB	55.58 TB	233.20 K	0 B	0 B	0	100.00 %	ignored	ignored
nobody	0 B	0 B	0	0 B	0 B	0	100.00 %	ignored	ignored

group	used bytes	logi bytes	used files	aval bytes	aval logib	aval files	filled[%]	vol-status	ino-status
nobody	0 B	0 B	0	0 B	0 B	0	100.00 %	ignored	ignored
zh	111.16 TB	55.58 TB	233.20 K	300.00 TB	150.00 TB	10.00 M	37.05 %	ok	ok

summary	used bytes	logi bytes	used files	aval bytes	aval logib	aval files	filled[%]	vol-status	ino-status
All users	111.16 TB	55.58 TB	233.20 K	0 B	0 B	0	100.00 %	ignored	ignored
All groups	111.16 TB	55.58 TB	233.20 K	300.00 TB	150.00 TB	10.00 M	37.05 %	ok	ok

```
[root@eosams02-ip700 (mgm:master mq:slave) ~]$ eos quota ls /eos/pilot/cms/tier0/
```

```
➤ Quota Node: /eos/pilot/cms/tier0/
```

group	used bytes	logi bytes	used files	aval bytes	aval logib	aval files	filled[%]	vol-status	ino-status
nobody	0 B	0 B	0	0 B	0 B	0	100.00 %	ignored	ignored
zh	0 B	0 B	0	300.00 TB	150.00 TB	10.00 M	0.00 %	ok	ok

summary	used bytes	logi bytes	used files	aval bytes	aval logib	aval files	filled[%]	vol-status	ino-status
All users	0 B	0 B	0	0 B	0 B	0	100.00 %	ignored	ignored
All groups	0 B	0 B	0	300.00 TB	150.00 TB	10.00 M	0.00 %	ok	ok



File migration



- > directory loop as before to sync with source if changes were made
- > prepares list of dirs for file processing
(to be used if `--limit-depth` is used, otherwise namespace queried recursively with top dir)
- > Tool configured with:
 - 1 process running namespace exploration with
 - * 5+5 threads parallel queries to src and dest
 - * 5 threads for xchecks
 - 5 migration processes with 10 threads each (50 transfers in parallel)

```
[root@eosams02-ip554 (mgm:slave mq:slave) run]$ python3.7 ${migrate} --root-nsdir $rootpth --limit-depth --max-job-files 20 --max-retries 1
--output-path ${rundir} > ${rundir}/ams.log 2>&1 &

watch tail -n 40 ${rundir}/ams.log

[...]
[INFO] 2022-03-06 13:31:37,106 [migrate.py:200] 114864 (MainThread) FILE: PID 0 RETRY 0: Preparing list of entries for 5 paths [ /eos/pilot/
cms/tier0/store/backfill/1/data/Tier0_REPLAY_2021/HLTPhysics/, '/eos/pilot/cms/
tier0/store/backfill/1/data/Tier0_REPLAY_2021/HLTPhysics/ALCARECO/', '/eos/pilot/cms/tier0/store/backfill/1/data/Tier0_REPLAY_2021/
HLTPhysics/ALCARECO/TkAlMinBias-PromptReco-v1/', '/eos/pilot/cms/tier0/store/backfill/1/d
ata/Tier0_REPLAY_2021/HLTPhysics/ALCARECO/TkAlMinBias-PromptReco-v1/000/', '/eos/pilot/cms/tier0/store/backfill/1/data/Tier0_REPLAY_2021/
HLTPhysics/ALCARECO/TkAlMinBias-PromptReco-v1/000/344/' ]
[...]
[INFO] 2022-03-06 13:31:37,106 [migrate.py:109] 114864 (MainThread) NS FILE threading PID 0 Starting thread processing.
[INFO] 2022-03-06 13:31:37,106 [namespace.py:309] 114864 (NS FILE threading PID 0_0) ILE: PID 0: Setting max depth to descend to 10; cmd:
[['eos-ns-inspect', 'scan', '--members',
'eosams02-ip700.cern.ch:7777', '--password-file', '/etc/eos.keytab', '--path', '/eos/pilot/cms/tier0/store/backfill/1/data/
Tier0_REPLAY_2021/HLTPhysics/', '--no-dirs', '--maxdept
h', '10']]
[INFO] 2022-03-06 13:31:37,106 [namespace.py:309] 114864 (NS FILE threading PID 0_1) ILE: PID 0: Setting max depth to descend to 11; cmd:
[['eos-ns-inspect', 'scan', '--members',
'eosams02-ip700.cern.ch:7777', '--password-file', '/etc/eos.keytab', '--path', '/eos/pilot/cms/tier0/store/backfill/1/data/
Tier0_REPLAY_2021/HLTPhysics/ALCARECO/', '--no-dirs', '
--maxdepth', '11']]
[...]
[INFO] 2022-03-06 13:33:26,571 [migrate.py:200] 114864 (MainThread) FILE: PID 0 RETRY 0: Preparing list of entries for 2 paths [ /eos/pilot/
cms/tier0/store/backfill/1/data/Tier0_REPLAY_2021/HLTPhysics/RAW/v1/000/344/063/00000/, '/eos/pilot/cms/tier0/store/backfill/1/data/
Tier0_REPLAY_2021/HLTPhysics/USER/LogErrorMonitor-PromptReco-v1/000/344/063/00000/' ]
[...]
```



- > namespace file entry xcheck → entries to process bunched to files
- > main process picks up prepared bunches and dispatches them

```
[...]
[INFO] 2022-03-06 13:34:40,644 [migrate.py:745] 114864 (MainThread) FILE: PID 1: Migrating 13 entries to destination.
[INFO] 2022-03-06 13:34:40,645 [migrate.py:52] 114864 (MainThread) FILE: PID 1 RETRY 0: Bunching thread input.
[INFO] 2022-03-06 13:34:40,645 [migrate.py:109] 114864 (MainThread) MIG FILE threading PID 1 Starting thread processing.
[INFO] 2022-03-06 13:34:40,646 [migrate.py:131] 114864 (MainThread) MIG FILE threading PID 1 Waiting for threads to finish.
[...]
[INFO] 2022-03-06 13:35:11,353 [migrate.py:149] 114864 (MainThread) MIG FILE threading PID 1 MIG threads finished processing.
[WARNING] 2022-03-06 13:35:11,354 [migrate.py:790] 114864 (MainThread) FILE: PID 1: Crashed processing entries:
{'unknown': 0, 'file_exists': 0, 'not_in_src_ns': 0, 'set_attr': 0, 'chmod': 0, 'chown': 0, 'broken_source': 0, 'not_permitted': 0,
 'remote_read_failed': 0, 'io_error_at_source': 0, 'network_unreachable': 0, 'operation_expired': 0}
[INFO] 2022-03-06 13:35:11,354 [migrate.py:813] 114864 (MainThread) FILE: PID 1: returning.
[INFO] 2022-03-06 13:35:11,354 [migration_daemon.py:150] 122856 (Thread-7 ) FILE PID 1: Job job FILE ID 1 RETRY 0 107298
ended processing. 'processing summary: {'processed': 13, 'remainder_job': 0, 'to_process': 13, 'total_size': 10742603741,
'total_time': 30, 'total_crashed': 0, 'transfer_rate_[MB/s]': 358.08679136666666, 'avg_file_size_[MB]': 826.35413392
30769, 'exception_thrown': 0}, crashes: {'unknown': 0, 'file_exists': 0, 'not_in_src_ns': 0, 'set_attr': 0, 'chmod': 0,
'chown': 0, 'broken_source': 0, 'not_permitted': 0, 'remote_read_failed': 0, 'io_error_at_source': 0, 'network_unreachable': 0,
'operation_expired': 0}
[INFO] 2022-03-06 13:35:12,786 [utils.py:174] 122856 (MainThread) FILE: PID 1: Starting file consolidation for files ['bunch_src_id_1.json'].
[INFO] 2022-03-06 13:35:12,786 [utils.py:224] 122856 (MainThread) FILE: PID 1: Removing job dir /var/migttest/run/file_bunches/0001/
[INFO] 2022-03-06 13:35:12,787 [migration_daemon.py:353] 122856 (MainThread) FILE: PID 1: Finished job PID 1 with result [ CUT ]
[INFO] 2022-03-06 13:35:12,787 [migration_daemon.py:354] 122856 (MainThread) MP: Processed in total 28 entries.
[INFO] 2022-03-06 13:35:12,787 [migration_daemon.py:368] 122856 (MainThread) MP: Finished 5 dir/file processing jobs and 0 bunching jobs.;
5 Active process ids from maximum 6: [0, 2, 4, 5, 3]
```

Sync replays (hands-on)



> `--replays <n>` to run `<n>` sync loops

> timestamps after full directory migration recorded to json

> directories changes at later date are taken for md sync and/or file migration

```
[root@eosams02-ip554 (mgm:slave mq:slave) run]$ python3.7 ${migrate} --root-nsdir $rootptl --replays 2 --only-dir-changes --max-job-files 20 --max-retries 1 --output-path ${rundir} > ${rundir}/ams.log 2>&1 &
```


```
watch tail -n 40 ${rundir}/ams.log
```

```
[...]
[INFO] 2022-03-06 14:13:38,916 [migration_daemon.py:528] 122856 (MainThread) Starting migration replay 0
[INFO] 2022-03-06 14:13:39,002 [migration_daemon.py:227] 122856 (MainThread) MP DIR: queued namespace MD sync and item bunching job (id 0).
[...]
[INFO] 2022-03-06 14:13:39,180 [migrate.py:416] 122867 (MainThread) DIR: PID 0 RETRY 0: Accounting total items to process after prefilter:43
[INFO] 2022-03-06 14:13:39,180 [migrate.py:420] 122867 (MainThread) DIR: PID 0 RETRY 0: 0 items skipped
[INFO] 2022-03-06 14:13:39,180 [migrate.py:440] 122867 (MainThread) DIR: PID 0 RETRY 0: 43 directories collected for file migration
processing.
[...]
[INFO] 2022-03-06 14:13:59,568 [migration_daemon.py:412] 122856 (MainThread) MP: Finished all processing jobs, saving stat output and
exiting unless more migration replays will be run.
[INFO] 2022-03-06 14:16:05,130 [migration_daemon.py:528] 122856 (MainThread) Starting migration replay 1
[INFO] 2022-03-06 14:16:05,213 [migration_daemon.py:227] 122856 (MainThread) MP DIR: queued namespace MD sync and item bunching job (id 0).
[...]
[INFO] 2022-03-06 14:16:05,384 [migrate.py:416] 135763 (MainThread) DIR: PID 0 RETRY 0: Accounting total items to process after prefilter:0
[INFO] 2022-03-06 14:16:05,384 [migrate.py:420] 135763 (MainThread) DIR: PID 0 RETRY 0: 43 items skipped
[INFO] 2022-03-06 14:16:05,384 [migrate.py:440] 135763 (MainThread) DIR: PID 0 RETRY 0: 7 directories collected for file migration
processing.
[...]
[INFO] 2022-03-06 14:16:07,016 [migrate.py:200] 135808 (MainThread) FILE: PID 0 RETRY 0: Preparing list of entries for 5 paths ['/eos/pilot/
cms/tier0/store/backfill/1/data/Tier0_REPLAY_2021/HLTPhysics/ALCARECO/TKA/MinBias-PromptReco-v1/000/344/063/00000/', ...CUT ...]
[...]
[INFO] 2022-03-06 14:16:07,948 [migrate.py:200] 135808 (MainThread) FILE: PID 0 RETRY 0: Preparing list of entries for 2 paths ['/eos/pilot/
cms/tier0/store/backfill/1/data/Tier0_REPLAY_2021/HLTPhysics/RAW/v1/000/344/063/00000/', '/eos/pilot/cms/tier0/store/backfill/1/data/
Tier0_REPLAY_2021/HLTPhysics/USER/LogErrorMonitor-PromptReco-v1/000/344/063/00000/']
[...]
[INFO] 2022-03-06 14:16:08,026 [migrate.py:416] 135808 (MainThread) FILE: PID 0 RETRY 0: Accounting total items to process after prefilter:0
[INFO] 2022-03-06 14:16:08,026 [migrate.py:475] 135808 (MainThread) FILE: PID 0 RETRY 0: Finished xchecks and filtering out entries which
exist at source and destination.
[...]
[INFO] 2022-03-06 14:16:09,930 [migration_daemon.py:557] 122856 (MainThread) Migration replays finished.
[...]
```

```
[root@eosams02-ip554 (mgm:slave mq:slave) run]$ wc -l dir_sync_timestamps.json
44 dir_sync_timestamps.json
```


Migration done (hands-on)

```
[root@eosams02-ip554 (mgm:slave mq:slave) ~]$ ll /var/migtest/run/
total 356
-rw-r--r--. 1 root root 314214 Mar  6 14:16 ams.log
drwxr-xr-x. 2 root root  4096 Mar  6 14:16 dir_bunches
-rw-r--r--. 1 root root  1662 Mar  6 14:13 DIR_migration_summary_report.json
drwxr-xr-x. 2 root root  4096 Mar  6 14:13 dir_reports
-rw-r--r--. 1 root root  5572 Mar  6 14:16 dir_sync_timestamps.json
drwxr-xr-x. 2 root root  4096 Mar  6 14:16 file_bunches
-rw-r--r--. 1 root root 10596 Mar  6 14:16 FILE_migration_summary_report.json
-rw-r--r--. 1 root root  5978 Mar  6 14:16 file_processing_dirs.json
drwxr-xr-x. 2 root root  4096 Mar  6 14:14 file_reports
```

 eosams02-ip554-migration@cern.ch
/eos/pilot/cms/tier0/store/backfill/1/data/Tier0_REPLAY_2021/HLTPHy
To: Undisclosed recipients;

43 directories were queued for processing, from these:


- 36 initially matched the source (no jobs ran for them)
- 7 were (at least partially) migrated
- 7 were completely migrated

replay 1

36 directories have a timestamp of complete *file* migration

Files of the directories above were migrated, their counts were:

- 105 queued for processing
- 105 migrated successfully
- 105 part of completely migrated
- 0 crashed
- 0 left to be processed

 eosams02-ip554-migration@cern.ch INSTANCE-MIGR
/eos/pilot/cms/tier0/store/backfill/1/data/Tier0_REPLAY_2021/HLTPHy: FILE migration finished
To: Undisclosed recipients;

7 directories were queued for processing, from these:


- 7 initially matched the source (no jobs ran for them)
- 0 were (at least partially) migrated
- 0 were completely migrated

43 directories have a timestamp of complete *file* migration

Files of the directories above were migrated, their counts were:

- 0 queued for processing
- 0 migrated successfully
- 0 part of completely migrated directories
- 0 crashed
- 0 left to be processed (from the queued ones at the moment)

replay 2

 eosams02-ip554-migration@cern.ch
/eos/pilot/cms/tier0/store/backfill/1/data/Tier0_REPLAY_2021/HLTPHy: FILE migration summary
To: Undisclosed recipients;

FILE migration summary for 105 entries from
/eos/pilot/cms/tier0/store/backfill/1/data/Tier0_REPLAY_2021/HLTPHy/

Last job attempt summary:

- 105 entries seen in bunches for processing
- 105 should equal the above (sanity check)
- 0 crashed
- 105 were processed,
- 105 should equal the above (sanity check)
- 0 remain to process (job omissions)
- 0 should equal the above (sanity check)
- 0 tpc remote read failed
- 0 crashed due to entries existing at destination
- 0 crashed due to i/o error at source
- 0 crashed due to transfer operation expired
- 0 crashed due to entries disappearing from the source namespace during migration
- 0 crashed due to entries being broken at source
- 0 crashed due to insufficient permissions at source (/entries could be broken)
- 0 crashed due to network being unreachable
- 0 crashed in setting attributes
- 0 crashed in setting mode
- 0 crashed in setting ownership
- 0 crashes of unknown type
- 0 unaccounted crashes (xcheck jobs vs category summary)

replay 1

Processing of: 156.902057556 GB (retries not accounted) took (retries included) 0.05676464325851864 hours.

The transfers took: 0.09083333333333334 hours (correct job average)

Average file size: 1494.305310057143 MB.

Average transfer speed: 479.82280598165136 MB/s (correct job average)

Job report can be found in:

/tmp/migtest/run/FILE_migration_summary_report.json.

Initial namespace exploration found the following items:

- 105/0 [src/dest] totals seen
- 0/0 [src/dest] items from directories explicitly excluded by operator
- 0/0 [src/dest] failed prefiltering (see below [1])
- 0 [dest] failed prefiltering and were removed (not on storage, ignored, decode errors, zero size/xs corrupted)
- 0 [dest] remained failed (removal omissions)
- 105/0 [src/dest] items passed the initial filter (see below [1])
- 0/0 [src/dest] items with matching all MD between the source and destination
- 0 [dest] had bad metadata (attribute, ownership, posix perm) crosscheck (size and checksum were O.K.)
- 0 [dest] removed due to bad size or xs
- 0 [dest] remained bad (removal omissions)
- 0 [dest] not present at source:
- 0 [dest] removed
- 0 [dest](removal omissions)
- 0 [src] were skipped (fully migrated already)
- 0/0 [src/dest] unaccounted items (sanity check, should be 0)
- 0/0 [src/dest] items with uid without a username (included in migration)
- 105 [src] were selected for processing input

[1]

Failed cases at source:

```
{
  "ignored": 0,
  "decode_errors": 0,
  "zero_size_xs_corrupted": 0,
  "no_storage_only_in_namespace": 0
}
```

Failed cases at destination:

```
{
  "ignored": 0,
  "decode_errors": 0,
  "zero_size_xs_corrupted": 0,
  "no_storage_only_in_namespace": 0
}
```

End of migration summary.



Migration done (hands-on)



```
[root@eosams02-ip554 (mgm:slave mq:slave) ~]$ eos ls -lh /eos/pilot/cms/tier0/store/backfill/1/data/Tier0_REPLAY_2021/HLTPhysics/
drwxr-x--- 1 cmst0 zh          7.03 M Mar  6 13:31 ALCARECO
drwxr-x--- 1 cmst0 zh          8.55 G Mar  6 13:31 AOD
drwxr-x--- 1 cmst0 zh          1.42 G Mar  6 13:31 DQMIO
drwxr-x--- 1 cmst0 zh          1.93 G Mar  6 13:31 MINIAOD
drwxr-x--- 1 cmst0 zh        109.39 G Mar  6 13:31 RAW
drwxr-x--- 1 cmst0 zh         35.56 G Mar  6 13:31 RAW-RECO
drwxr-x--- 1 cmst0 zh         44.89 M Mar  6 13:31 USER

[root@eosipilot-ns-ip700 (mgm:master mq:slave) ~]$ eos ls -lh /eos/pilot/cms/tier0/store/backfill/1/data/Tier0_REPLAY_2021/HLTPhysics/
drwxr-x--- 1 cmst0 zh          7.03 M Nov 22 16:15 ALCARECO
drwxr-x--- 1 cmst0 zh          8.55 G Nov 22 16:26 AOD
drwxr-x--- 1 cmst0 zh          1.42 G Nov 22 16:15 DQMIO
drwxr-x--- 1 cmst0 zh          1.93 G Nov 22 16:29 MINIAOD
drwxr-x--- 1 cmst0 zh        109.39 G Nov 16 12:35 RAW
drwxr-x--- 1 cmst0 zh         35.56 G Nov 16 14:44 RAW-RECO
drwxr-x--- 1 cmst0 zh         44.89 M Nov 22 16:07 USER
```

```
[root@eosams02-ip554 (mgm:slave mq:slave) ~]$ eos file info /eos/pilot/cms/tier0/store/backfill/1/data/Tier0_REPLAY_2021/HLTPhysics/USER/
LogErrorMonitor-PromptReco-v1/000/344/063/00000/070495ac-14e3-433f-aa3f-beb6ae769e21.root
File: '/eos/pilot/cms/tier0/store/backfill/1/data/Tier0_REPLAY_2021/HLTPhysics/USER/LogErrorMonitor-PromptReco-v1/000/344/063/00000/
070495ac-14e3-433f-aa3f-beb6ae769e21.root'  Flags: 0644
Size: 15002478
Modify: Mon Dec 13 08:32:03 2021 Timestamp: 1639380723.458453749
Change: Sun Mar  6 14:15:20 2022 Timestamp: 1646572520.349817905
Birth: Mon Dec 13 08:32:03 2021 Timestamp: 1639380723.458125257
CUid: 103031 CGid: 1399 Fxid: 0f4ce749 Fid: 256698185 Pid: 838735 Pxid: 000cccc4f
XStype: adler  XS: 8e 67 69 1c  ETAGs: "68906894344847360:8e67691c"
Layout: replica Stripes: 2 Blocksize: 4k LayoutId: 00100112 Redundancy: d2::t0
#Rep: 2
```

no.	fs-id	host	schedgroup	path	boot	configstatus	drain	active	geotag
0	103	st-096-dd906d68.cern.ch	default.72	/data29	booted	rw	nodrain	online	740376cb::0513::S::0034::SM05
1	894	st-096-ee809e30.cern.ch	default.72	/data91	booted	rw	nodrain	online	853a68f0::0513::S::0034::SE10

```
[root@eosipilot-ns-ip700 (mgm:master mq:slave) ~]$ eos file info /eos/pilot/cms/tier0/store/backfill/1/data/Tier0_REPLAY_2021/HLTPhysics/USER/
LogErrorMonitor-PromptReco-v1/000/344/063/00000/070495ac-14e3-433f-aa3f-beb6ae769e21.root
File: '/eos/pilot/cms/tier0/store/backfill/1/data/Tier0_REPLAY_2021/HLTPhysics/USER/LogErrorMonitor-PromptReco-v1/000/344/063/00000/
070495ac-14e3-433f-aa3f-beb6ae769e21.root'  Flags: 0644
Size: 15002478
Modify: Mon Dec 13 08:32:03 2021 Timestamp: 1639380723.458453749
Change: Mon Dec 13 08:32:03 2021 Timestamp: 1639380723.458125257
Birth: Mon Dec 13 08:32:03 2021 Timestamp: 1639380723.458125257
CUid: 103031 CGid: 1399 Fxid: 00266012 Fid: 2514962 Pid: 8685 Pxid: 000021ed
XStype: adler  XS: 8e 67 69 1c  ETAGs: "675104971292672:8e67691c"
Layout: replica Stripes: 2 Blocksize: 4k LayoutId: 00100112 Redundancy: d2::t0
#Rep: 2
```

no.	fs-id	host	schedgroup	path	boot	configstatus	drain	active	geotag
0	275	p06636710a25939.cern.ch	default.9	/data40	booted	rw	nodrain	online	71b9bb38::0513::R::0050::CH09
1	370	p06636710a08535.cern.ch	default.9	/data37	booted	rw	nodrain	online	311bd43f::0513::R::0050::CI22



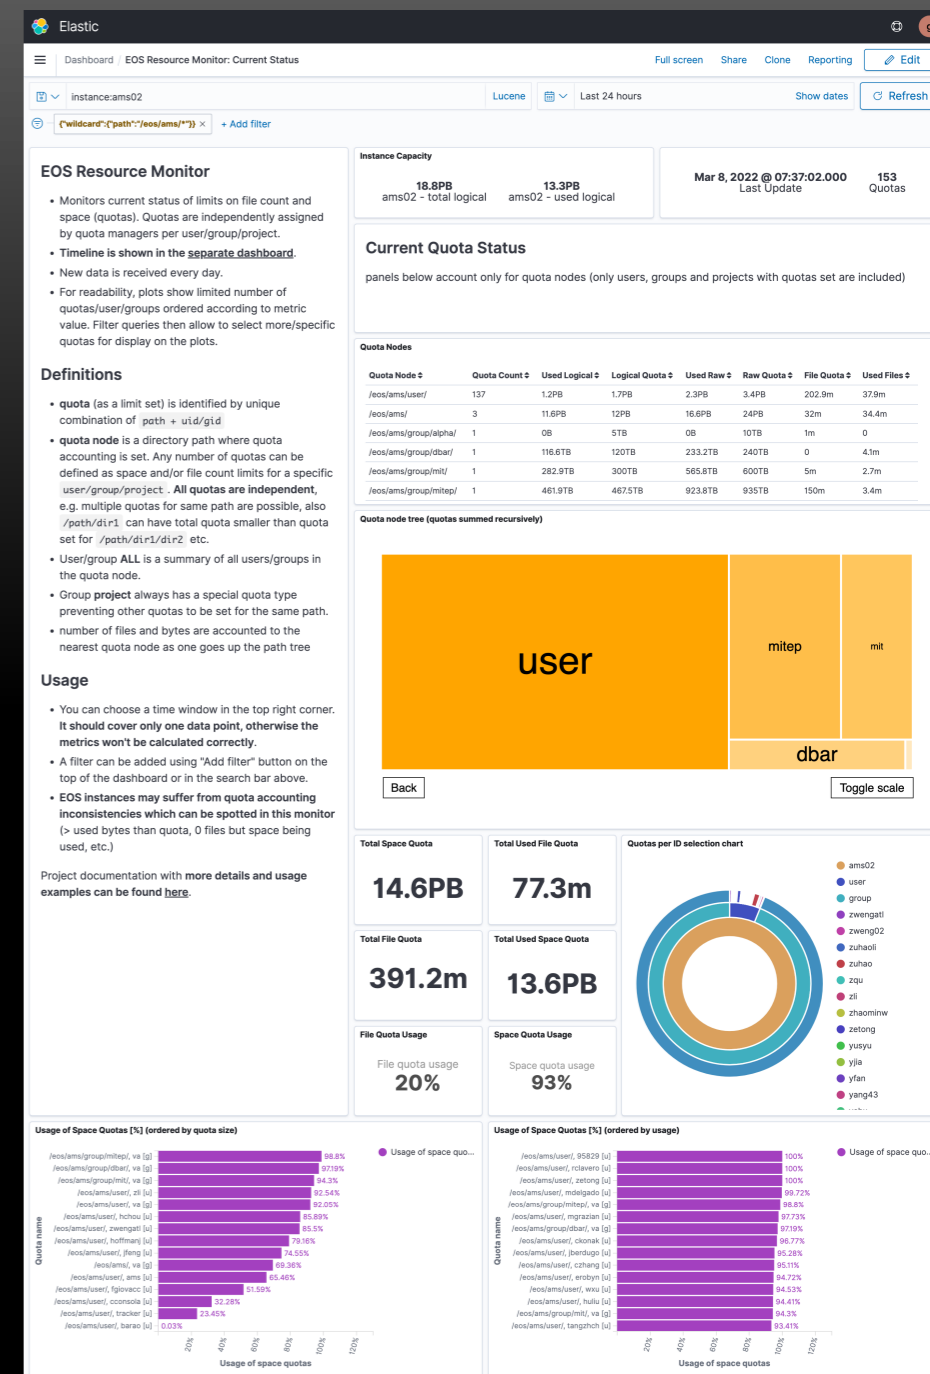
Sync of AMS at scale



> every 30 min sync between EOSPUBLIC with EOSAMS02

```
[root@eospublic-ns-ip563 (mgm:master mq:master) ~]$ eos ls -lh /eos/ams
drwxr----+ 1 ams      va      3.82 P Aug 29 2021 Data
drwxr----+ 1 ams      va      7.88 P Jan 24 2014 MC
drwxr----+ 1 ams      va     514.92 M Mar  3 2021 aegorov
drwxr----+ 1 ams      va      0 Oct 23 2018 ams
drwxr----+ 1 ams      va     861.46 T Dec 18 2020 group
drwxr----+ 1 root     root    1.39 G Apr 30 2021 opstest
lrwxrwxrwx 1 ams      va      19 Dec 18 2020 pocc -> /eos/ams/group/pocc
drwxr----+ 1 ams      va     45.63 M Aug 21 2019 qqq
drwxr----+ 1 ams      va     16.04 G Aug  5 2021 scratch
drwxr-x---+ 1 ams      va    741.04 T Apr 26 2013 user

[root@eosams02-ip700 (mgm:master mq:slave) ~]$ eos ls -lh /eos/ams
drwxr----+ 1 ams      va      3.82 P Jan 26 08:44 Data
drwxr----+ 1 ams      va      7.88 P Jun  2 2021 MC
drwxr----+ 1 ams      va     514.92 M Jul  2 2021 aegorov
drwxr----+ 1 ams      va      0 Mar  4 2021 ams
drwxr----+ 1 ams      va     861.46 T Apr  6 2021 group
drwxr----+ 1 root     root    0 Jul  2 2021 opstest
lrwxrwxrwx 1 ams      va     19 Nov 17 15:32 pocc -> /eos/ams/group/pocc
drwxr----+ 1 ams      va     45.63 M Jul  2 2021 qqq
drwxr----+ 1 ams      va     16.04 G Nov  5 05:04 scratch
drwxr-x---+ 1 ams      va    739.83 T Jan 26 06:50 user
```



> we also have in place EOS Quota monitoring for IT-ST-PDS instances (showing here that EOSAMS02 need quota sync ;)



Thank you



Questions ?

