

Archiving Tools for

EOS

Disk Storage @ CERN

Elvin-Alin Sindrilaru

Presenter: Andreas-Joachim Peters


IT Data & Storage Service Group

CHEP Okinawa 13.-17.4.2015

OUTLINE



- Why an archiving tool?
- Architecture
- CLI & archive file structure
- Summary

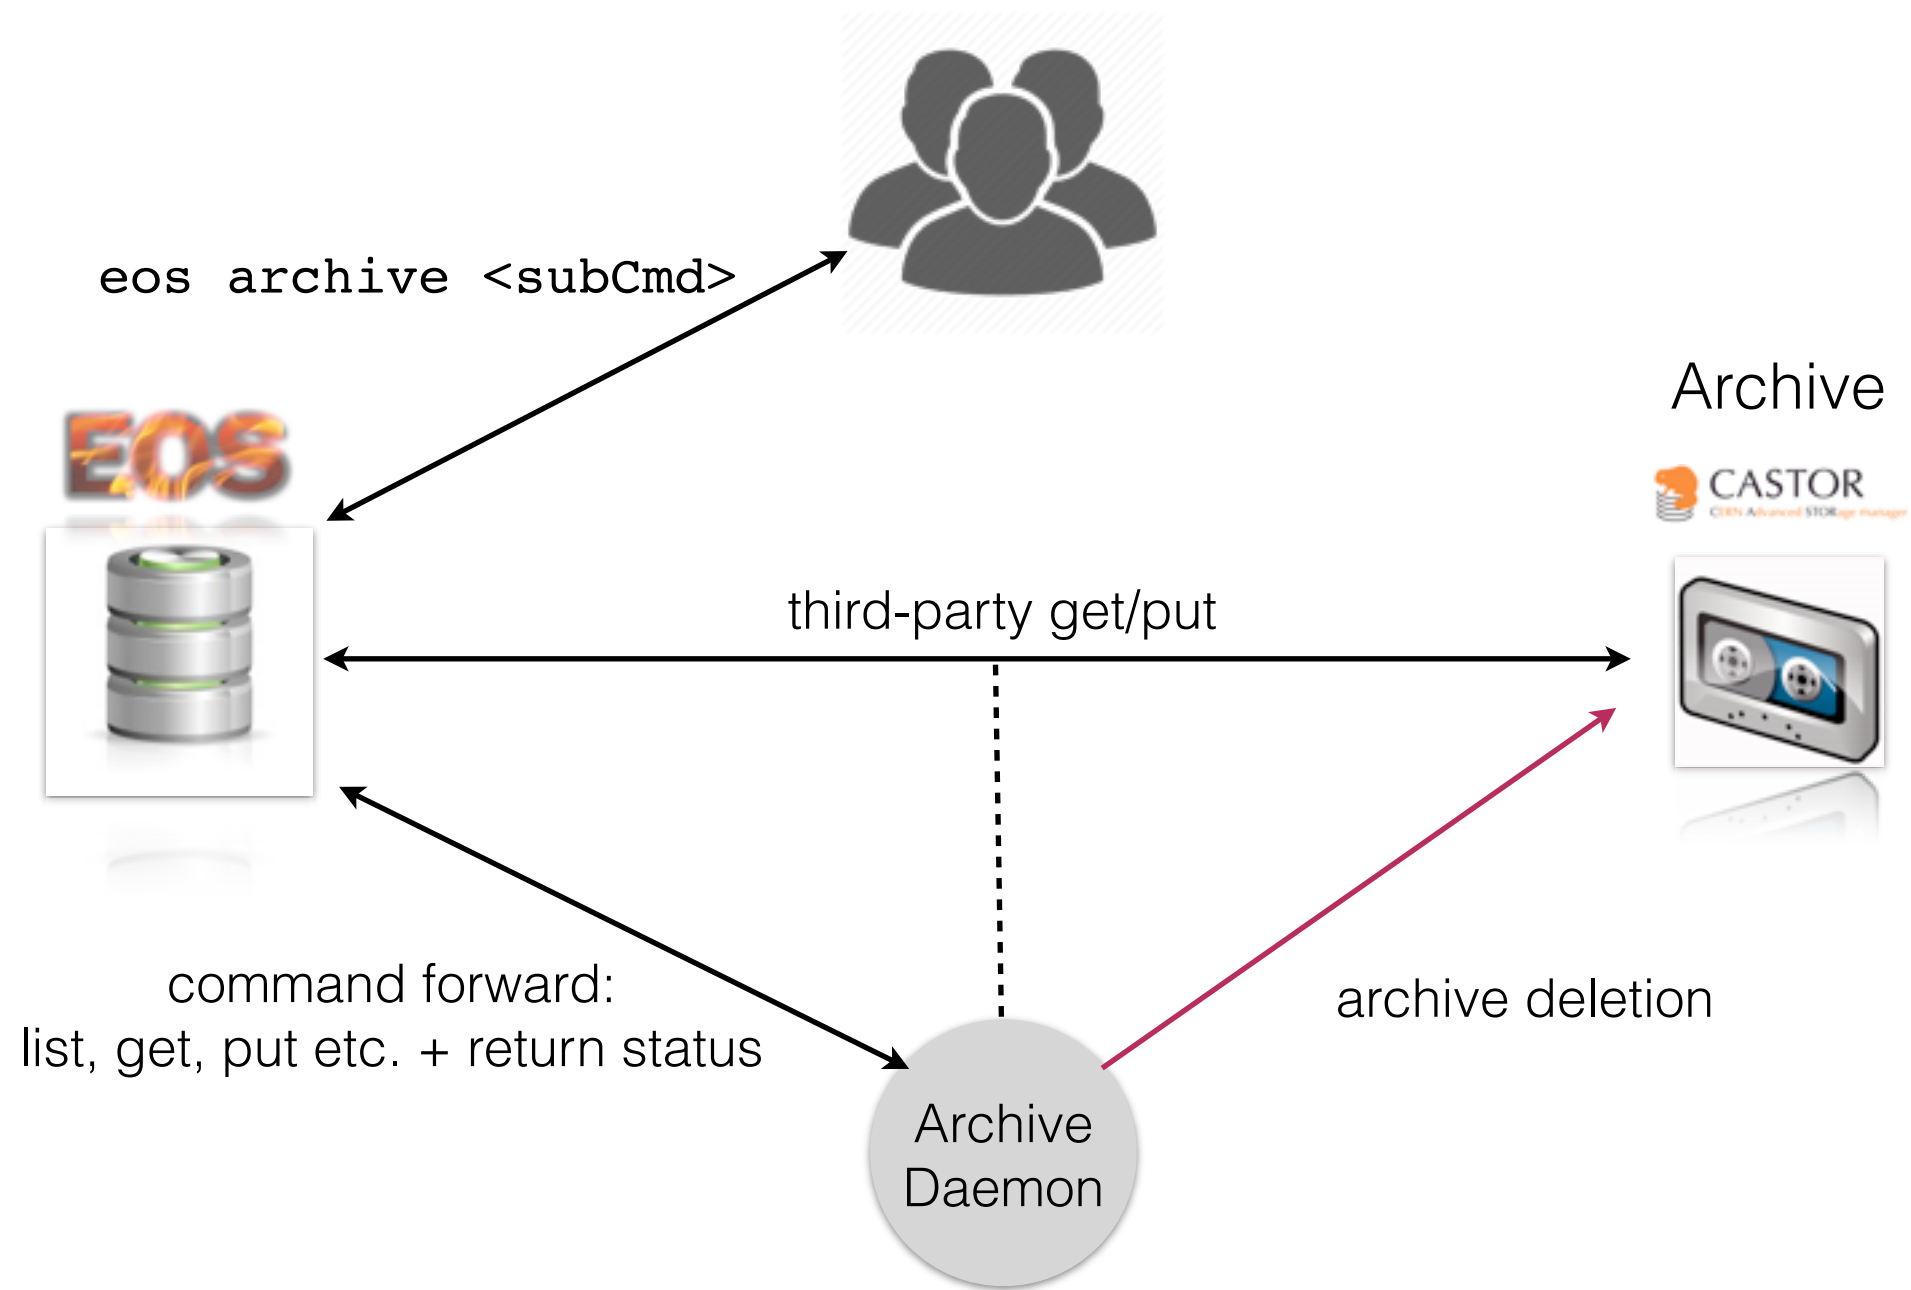


WHY AN ARCHIVING TOOL?



- **users/group quota space is limited**
 - want to free cold storage space in EOS for users' online activities
- avoid user-made ad hoc archiving solutions
- **manage efficiently** the movement of data between disk and archive storage
 - ▶ no client side copy
- easily implement archiving rules and best practices

ARCHITECTURE

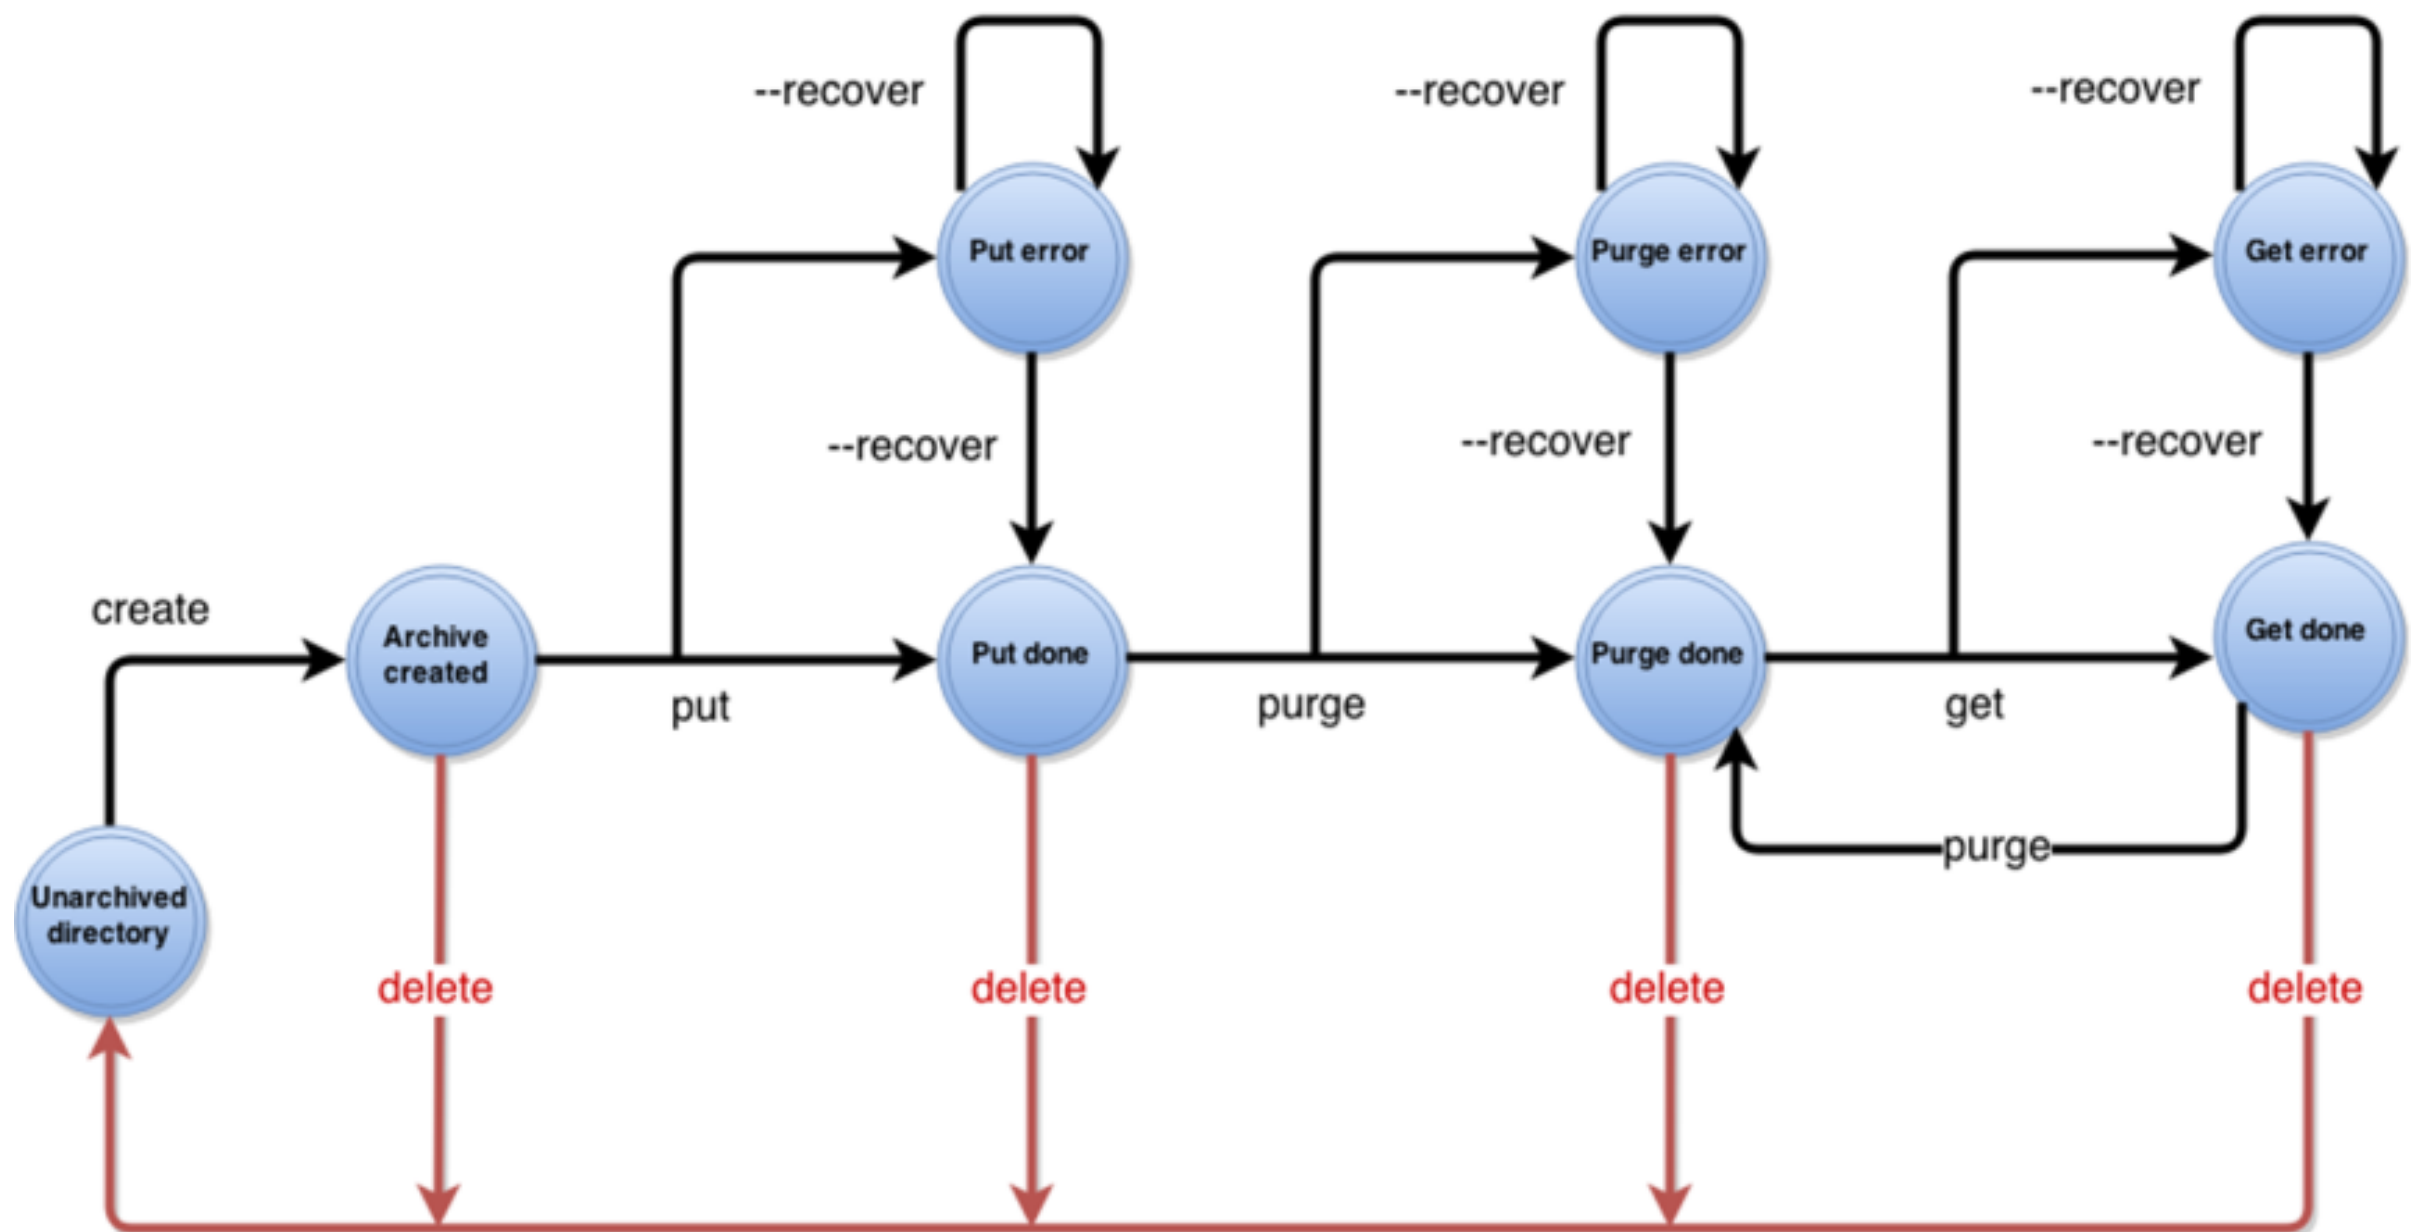


EOS ARCHIVE CLI



- **eos archive** <subcmd>
 - ▶ **create** <eos_path>
 - create archive file for <eos_path> directory
 - ▶ **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** [<eos_path>]
 - show status of archived directories in the <eos_path> subtree
 - ▶ **transfers** [all|put|get|purge|job_uuid]
 - display ongoing transfers. By default, all transfers are displayed
 - ▶ **delete** <eos_path>
 - delete all archived data including the archive file - **Admin command**

ARCHIVE WORKFLOW



ARCHIVE & LOG FILES



- **archive definition 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
 - ▶ used to get information about the contents of the archive
 - ▶ archive “*get*” restores files in their original layout (2 replicas, RAIN etc.)

```
{ "src": "root://p05508916e55509.cern.ch:1094//eos/cmsarchive/45964a1c85413c4d6c5346bf82c23b635/", "svc_class": "archive", "uid": "2672", "gid": "1028", "timestamp": 1427447120, ["d", ".", "2672", "1028", "42775", {"sys.acl": "u:lmascet", rced.blocksize": "4k", "sys.forced.checksum": "adler", "sys.cle/"}] ["d", "test/", "2672", "1028", "42775", {"sys.acl": "u:lmascet", rced.blocksize": "4k", "sys.forced.checksum": "adler", "sys.cle/"}] ["d", "test/test1/", "2672", "1028", "42775", {"sys.acl": "u:lmascet", rced.blocksize": "4k", "sys.forced.checksum": "adler", "sys.cle/"}] ["d", "test/test1/test1", "2672", "1028", "42775", {"sys.acl": "u:lmascet", rced.blocksize": "4k", "sys.forced.checksum": "adler", "sys.cle/"}] }
```

```
C\LECLC\G\}}  
2λ2*40LC6Q*P\OCk2T5G: "qK" "2λ2*40LC6Q*CM6Ck2nW: "9qf6L"  
["q" "622f\622fT" "5015" "7058" "45112" {"2λ2*8Cf: "f  
2672:1028
```

- **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

```
10:17:21,674 transfer[118620] transfer.py:91 LVL=INFO Starting the status thread  
10:17:21,674 transfer[118620] utils.py:43 LVL=DEBUG Execute: root://p05508916e55509.cern.ch:1094//proc/user/?mgm.cmd=file&mgm.subcmd=rename&mgm.path=/eos/cmsarchive/arch  
7/.archive.purge.done&mgm.file.source=/eos/cmsarchive/archive/20150327/.archive.purge.done&mgm.file.target=/eos/cmsarchive/archive/20150327/.archive.get.err  
10:17:21,706 transfer[118620] archivefile.py:84 LVL=DEBUG fseek_dir=468, fseek_file=37764  
10:17:21,707 transfer[118620] utils.py:43 LVL=DEBUG Execute: root://p05508916e55509.cern.ch:1094//proc/user/?mgm.cmd=find&mgm.path=/eos/cmsarchive/archive/20150327/&mgm.  
10:17:21,708 transfer[118620] utils.py:43 LVL=DEBUG Execute: root://p05508916e55509.cern.ch:1094//proc/user/?mgm.cmd=chown&mgm.path=/eos/cmsarchive/archive/20150327/&mgm  
r=2672:1028  
10:17:21,709 transfer[118620] utils.py:43 LVL=DEBUG Execute: root://p05508916e55509.cern.ch:1094//proc/user/?mgm.cmd=chmod&mgm.path=/eos/cmsarchive/archive/20150327/&mgm  
=42775  
=45112  
T0:T1:5T*100 fL9u2f6L[JT0250] nETf2*bλ:43 ΓΛΓ=DEBNC EX6CnΓ6: L00Γ:\\b022080J0622200*CGLU*CM:T00d\\bLOC\p26L\jwδw*CWQ=CMW0Q2Wδw*b9Γμ=\\602\CW29LCμTΛ6\9LCμTΛ6\50T20351\QWδw  
L=5015:T050  
T0:T1:5T*100 fL9u2f6L[JT0250] nETf2*bλ:43 ΓΛΓ=DEBNC EX6CnΓ6: L00Γ:\\b022080J0622200*CGLU*CM:T00d\\bLOC\p26L\jwδw*CWQ=CMW0Q2Wδw*b9Γμ=\\602\CW29LCμTΛ6\9LCμTΛ6\50T20351\QWδw  
L=5015:T050
```




SUMMARY

- ▶ archive creation leads to **sub-tree freeze** in EOS via ACLs
 - ▶ immutable flag: `sys.acl=z:i`
 - ▶ **no updates/writes possible** to an archived sub-tree
- ▶ XRootD **Third-Party Copy** transfers are done sequentially
 - ▶ starting with XRootD 4.1 will run in parallel to boost performance
- ▶ EOS keeps track of any archived sub-tree to facilitate quick access/search operation

```
EOS Console [root://localhost] |eos/> archive list
```

Creation date	Path	Status
Tue Mar 3 11:20:50 2015	/eos/esindril/raiddp/	get done
Sun Feb 22 18:03:55 2015	/eos/esindril/raid6/	created
Wed Nov 26 15:27:00 2014	/eos/cmsarchive/store/user/lmascett/archive001/	put done
Fri Feb 13 17:58:09 2015	/eos/cmsarchive/store/group/comm_ecal/upgrade/GluGluToHToGanGan_M-125_14TeV-powheg-pythia6_SLHC17/	put failed
Mon Nov 3 15:12:19 2014	/eos/cmsarchive/store/group/comm_ecal/matbudget/	get failed
Fri Mar 27 10:04:56 2015	/eos/cmsarchive/store/group/comm_ecal/upgrade/RelValGamma_pT35_SLHC15/	purge done

- ▶ **enables** users and working groups to **extend** their virtual available **storage space**
 - ▶ archiving rights are granted via the 'a' ACL flag : `sys.acl="u:xavierc3:rwxcqa"`
- ▶ has been extended as **full/incremental backup** tool
 -  use case

Thank you!



Questions or Comments?