

Table of Contents

Neutrino Platform Storage at CERN.....	1
Introduction.....	2
CASTOR tape storage system.....	3
Neutrino Platform CASTOR tape area.....	4
copying single files.....	4
Long waiting time with rfcop.....	4
Reporting errors.....	5
Permissions.....	5
Removing permissions.....	5
Changing permissions.....	5
Changing ownership.....	6
How to use Personal CASTOR storage area?.....	6
Documentation.....	6
EOS storage system.....	7
How to access personal EOS space.....	7
How to access Neutrino Platform experimental EOS either from lxplus or from neutplatform computing cluster.....	7
Setting up Neutrino Platform EOS.....	8
Proposal for a new structure in NP-EOS.....	8
Proposal NP02/NP04.....	8
Proposal for DUNE.....	9
Accessing an EOS file from ROOT.....	9
Other EOS commands.....	9
How to list directories and files in EOS.....	10
Writing and reading files on EOS.....	10
Deleting files and directories on EOS.....	11
xrscp man page.....	11
EOS Knowledge Based articles and documentation.....	11

Neutrino Platform Storage at CERN

Introduction

This page describes how to access data at CERN in different locations, explaining the different commands and tools to be used and the access policy to the storage. Neutrino Platform at CERN mostly relies now on EOS [↗](#) for disk only space.

xrootd is used and works with EOS technology.

All commands and examples are thought to be issued on neutplatform nodes inside the CERN network. Most of them work anyway as well from outside CERN.

CASTOR tape storage system

CASTOR is the acronym for CERN Advanced STORage manager, a data storage system used at CERN, more on cern.ch/castor. CASTOR should not be used as a backup service, as suggested [here](#).

CASTOR is meant for storing large volumes of physics data. CASTOR is inefficient for small files, both for writing data to tape (migration) and getting data from tape (recall): the optimal file size for CASTOR is 1GB and larger.

CASTOR doesn't know about e-groups at all, just UNIX-like permissions (user::groups), so much simpler than EOS.

In order to access CASTOR service, one needs to set some variables. In bash:

```
export RFIO_USE_CASTOR_V2=YES
export STAGE_HOST=castorpublic
export STAGE_SVCCLASS=default
```

Neutrino Platform CASTOR tape area

NP base directories are:

```
/castor/cern.ch/neutplatform/protodune  
/castor/cern.ch/neutplatform/dune
```

Experimental area is granted to np-comp UNIX group, so be aware you have to belong to this UNIX computing group.

copying single files

Remember to set proper variables for CASTOR. You can use now rfcop command to copy file:

```
rfcop /castor/cern.ch/neutplatform/protodune/.../somefile.root /home/user/workspace
```

Or new XROOT [X](#) protocol (might be faster, no need to setup variables for CASTOR):

```
xrdcp root://castorpublic//castor/cern.ch/neutplatform/protodune/.../somedir/somefile.root /home/  
xrdcp /home/user/workspace/somefile.root root://castorpublic//castor/cern.ch/neutplatform/protodu
```

If some of these commands hangs - read next section of this page.

Long waiting time with rfcop

To check status of file, use stager_qry command:

```
stager_qry -M /castor/cern.ch/neutplatform/protodune/.../somefile.root
```

In most cases, as an result you can get one of following answers:

- STAGED
- STAGEIN
- Error 2/No such file

STAGED means that means "file on disk matrix (fast access)", so you should obtain it quickly with rfcop command.

```
Received 1 responses  
/castor/cern.ch/neutplatform/protodune/.../somefile.root  
312327210@castorns STAGED
```

STAGEIN means that "file is on a tape" and need to be transferred to the disk matrix, which takes some time. You need to wait for the result of rfcop command.

```
Received 1 responses  
/castor/cern.ch/neutplatform/protodune/.../somefile.root  
386766073@castorns STAGEIN
```

"Error 2/No such file" also means that "file is on a tape" and need to be transferred to the disk matrix, which takes some time. You need to wait for the result of rfcop command.

```
Error 2/No such file or directory (File /castor/cern.ch/neutplatform/protodune/.../somefile.root
```

This error is very similar to the one that you obtain when asking about not existing file. In such case (file really missing) you will see:

```
Error 2/No such file or directory (Filename /castor/cern.ch/neutplatform/protodune/.../somefile.r
```

You can check if file is existing with `nsls` command.

In second and third case you can ask stager directly to move files to the "disk matrix" by issuing following command:

```
stager_get -S default -M /castor/cern.ch/neutplatform/protodune/.../somefile.root
```

More on this topic - read this [↗](#)

Reporting errors

Switch on debugging for `rfcp` (bash/zsh):

```
export RFIO_TRACE=3
export STAGER_TRACE=3
```

or for `tcsh/csh`

```
setenv RFIO_TRACE 3
setenv STAGER_TRACE 3
```

To switch debugging for `xrdcp`, use `-d` switch (i.e. `-d 1`, `-d 2`, `-d 3`)

Permissions

To set permission we use `nssetacl` [↗](#), to get them `nsgetacl` [↗](#).

Removing permissions

Remove ACL entry:

```
nssetacl -d u:'_username_':rwx /castor/cern.ch/neutplatform/protodune/offline/test
```

Changing permissions

Modify default ACLs to 775 (will be propagated).

```
nssetacl -m u::7,g::7,o::5 /castor/cern.ch/neutplatform/protodune/offline/test
nssetacl -m d:u::7,d:g::7,d:o::5 /castor/cern.ch/neutplatform/protodune/offline/test
```

For every directory (or file) created under `/castor/cern.ch/neutplatform/protodune/offline/test` owner will be set up to the person who created it, but permissions of all directories will be set to 775. This way everybody in `np-comp` group will be able to write to any of subdirectories (whoever will create it).

Let us do the same as above, but not for group `np-comp`, but for list of users.

```
nssetacl -m d:u::7,d:g::5,d:o::5 /castor/cern.ch/neutplatform/protodune/offline/test
nssetacl -m u::7,g::5,o::5 /castor/cern.ch/neutplatform/protodune/offline/test
nssetacl -m default:group:np-comp:rx,default:other:-,default:mask:rwx /castor/cern.ch/neutplatform/protodune/offline/test
nssetacl -m default:user:_username_:rwx,default:mask:rwx /castor/cern.ch/neutplatform/protodune/offline/test
nssetacl -m user:_username_:rwx,mask:rwx /castor/cern.ch/neutplatform/protodune/offline/test
nssetacl -m d:u:_username1_:7,d:m:7 /castor/cern.ch/neutplatform/protodune/offline/test
```

Long waiting time with `rfcp`

CENFStorageAtCERN < CENF < TWiki

```
nssetacl -m u:_username1_:7,m:7 /castor/cern.ch/neutplatform/protodune/offline/test
```

For every directory (or file) created under `/castor/cern.ch/neutplatform/protodune/offline/test` owner will be set up to the person who created it, but permissions of all directories will be set to **755**. This way everybody in *np-comp* group will be able **just to read** from any of subdirectories (whoever will create it). Only users from the list can write to subdirectories.

If after some time, new user needs to be added to the list, then we shall add him permissions to write to existing files. Following command will work only with bash.

```
nsfind /castor/cern.ch/neutplatform/protodune/offline/test -type f > filechecklist.txt
cat filechecklist.txt | xargs nssetacl -m u:_username2_:rw,m:rw
nsfind /castor/cern.ch/neutplatform/protodune/offline/test -type d > dirchecklist.txt
cat dirchecklist.txt | xargs nssetacl -m d:u:_username2_:rwx,d:m:rwx
cat dirchecklist.txt | xargs nssetacl -m u:_username2_:rwx,m:rwx
```

Changing ownership

If recursive method for applying access rights was applied as it is described above, it might happen that newly created directories might have different owner as their parent directories. To change it one may use recursive version of `nschown` command, i.e.:

```
nschown -R _username3_ /castor/cern.ch/neutplatform/protodune/offline/test
```

IMPORTANT: such command might disable effective access rights. To bring them back, read from output of `nsgetacl` commands who has for example read access and grant again such access for one user. This will make effective access rights correct, i.e.:

```
nsfind /castor/cern.ch/neutplatform/protodune/offline/test -type d > dirchecklist.txt
cat dirchecklist.txt | xargs nssetacl -m d:u:_username_:rx,d:m:rx
cat dirchecklist.txt | xargs nssetacl -m u:_username_:rx,m:rx
```

How to use Personal CASTOR storage area?

Every user with CERN account has a directory in CASTOR. To list the files in your directory, type:

- `rfdir /castor/cern.ch/user/initial/username/` to list the files in your directory, where initial is the first letter of your user name.
- `nsls` - list home CASTOR directory
- `nsls /castor....`
- `nsls -l` - long listing
- `rformkdir /castor/cern.ch/user/initial/username/dir` to create directories
- `rfcop filename /castor/cern.ch/user/initial/username/dir` copy files to a CASTOR directory
- `nsrm /castor/cern.ch/user/initial/username/dir` - remove dir directory

Other useful commands:

- `xrdcp`
- `nsgetacl` , `nssetacl`
- `stager_qry`
- `stager_get`

Documentation

*CASTOR quick tutorial for beginners [☞](#) *CASTOR Knowledge Base Articles/FAQ [☞](#)

EOS storage system

Pools usually accessed directly by the users exploit the EOS technology. They can be accessed through the `eospublic.cern.ch` xrootd redirector using the xrootd tools and commands.

This section will briefly explain the easiest way to perform usual operations on EOS (i.e., listing directories, writing and reading files).

For more actions, you can open an interactive shell on eospublic by simply typing `eos`.

From within the EOS shell, type `help` for a full list of commands you can perform on EOS (most of them mimic the common unix commands, like `ls`, `mkdir`, `rm`, ...).

Similarly, from your linux shell you can also type `xrd eospublic` to open a xrootd interactive session on eospublic. Typing `help` from within the interactive shell will give you the list of xrd actions you can perform on EOS.

How to access personal EOS space

The initial and mandatory step to using your CERNBox is to create it.

Use any Web browser to go to '<https://cernbox.cern.ch>' and login via the CERN Single Sign-On page with your username and password (NICE or lxplus).

This first connection creates your CERNBox storage **space** (1 TB, up to 1 million files), hosted in

To activate the corresponding service go to:

https://resources.web.cern.ch/resources/Manage/EOS/Default.aspx?login=YOUR_UNIX_LOGIN

You must set up your environment before starting:

```
export EOS_MGM_URL=root://eosuser.cern.ch
```

and your user directories are under `/eos/user/<initial>/<userid>`

Please take into account, personal EOS space (`/eos/user/<firstletter_login_name>/<login_name>`) EOS **space** (`/eos/experiment/neutplatform/...`).

How to best access EOS from central services (LXPLUS and LXBATCH) [↗](#)

If you still have problems you can get help through the CERN Service Portal [↗](#).

If you have questions, you might also have a look at the cernbox FAQ [↗](#)

Web browser supported platforms [↗](#).

How to access Neutrino Platform experimental EOS either from *lxplus* or from *neutplatform* computing cluster

Neutrino Platform experimental EOS space is located at `/eos/experiment/neutplatform/`.

A user to have access needs to be a member of the np-comp UNIX group.

Ensure that the users who will access the space have the correct primary UNIX group.

Otherwise get them to change it in the accounts setting page

https://resources.web.cern.ch/resources/Manage/Linux/Settings.aspx?login=YOUR_UNIX_LOGIN

If **for** any reason you can't have access to the !EOS space from our cluster please take into consi

- * check **if** you are a member of the appropriate e-group
- * check **if** you have a valid kerberos certificate

- If you still have problems with the access of the space, please contact neutplatform computing support [✉](#)

Setting up Neutrino Platform EOS

- If you want to make use of most recent features mentioned in this FAQ you have to pick the 'client' version of EOS using a bash or tcsh setup script:

```
// How to setup neutplatform EOS space
// Verify first that you are connecting to 'your' experiment dedicated EOS instance
ssh neutplatform
source /afs/cern.ch/project/eos/installation/0.3.121-aquamarine.public/etc/setup.sh
mkdir <folder_name>
eosmount <folder_name>
ls <folder_name>/experiment/neutplatform

//unmount the folder
In order to unmount the folder you can either disconnect from the machine you are connected to or
eosumount <folder_name>
```

The automatic EOS endpoint setting use your group membership to sort you into an experiment. If **this** does not work **for** you, you can just specify your endpoint via an environment variable:
 export EOS_MGM_URL=root://**eospublic.cern.ch**

Using EOS

- The 'eos' CLI (Command Line Interface) supports most of the standard filesystem commands like:
 - ls, cd, mkdir, rm, rmdir, find, cp ...
 - 'eos help' shows all available commands, 'eos --help' explains each command

The 'eos' CLI can be used as an interactive shell with history

```
as a busy-box command
bash$ eos whoami

bash$ eos
EOS console [root://eospublic.cern.ch] ||> whoami
to quit from the EOS console
EOS console [root://eospublic.cern.ch] ||>.q

# check eos version
bash$ eos version
EOS_INSTANCE=eospublic
EOS_SERVER_VERSION=0.3.226 EOS_SERVER_RELEASE=1
EOS_CLIENT_VERSION=0.3.121 EOS_CLIENT_RELEASE=aquamarine

# check eos help
bash$ eos help
```

Proposal for a new structure in NP-EOS

Please see here the structure it self and the assigned rights to corresponding e-groups.

Proposal NP02/NP04

****current situation****. You could start using it and let neutplatform computing support  know.

purpose	EOS path	e-groups (ACLs)
For both protoDUNEs (np02/np04)	common path: /eos/experiment/neutplatform/protodune	

For NP02 RAW data	/eos/experiment/neutplatform/protodune/rawdata/np02	eos-experiment-cenf-np02-rea (rx)
For NP04 RAW data	/eos/experiment/neutplatform/protodune/rawdata/np04	eos-experiment-cenf-np04-rea (rx)
NP02/NP04 MCC files : detector performance, reconstruction, calibration, physics, cross-section, measurements		eos-experiment-cenf-np02-rea (rx)
		eos-experiment-cenf-np04-rea (rx)
		eos-experiment-cenf-np02-wr (rwx)
		eos-experiment-cenf-np04-wr (rwx)
NP02/NP04 physics area: output from various subgroups analyses output , ROOT files mainly		eos-experiment-cenf-np02-rea (rx)
		eos-experiment-cenf-np04-rea (rx)
		eos-experiment-cenf-np02-wr (rwx)
		eos-experiment-cenf-np04-wr (rwx)
NP04 beamgroup : studies on particle composition of the hadron beam at the entrance of the cryostat		eos-experiment-cenf-np04-rea (rx)
		eos-experiment-cenf-np04-wr (rwx)

-->

Proposal for DUNE

****NOT READY****

purpose	EOS path	e-groups (ACLs)	quotas	Notes
---------	----------	-----------------	--------	-------

Accessing an EOS file from ROOT

- You have to use URLs as file names which are built in this way:
root://eospublic.cern.ch/eos/experiment/neutplatform/... e.g.
- root://eospublic.cern.ch/eos/experiment/neutplatform/.../histo.root
- root [0]: TFile::Open("root://eospublic.cern.ch/eos/experiment/neutplatform/.../histo.root");

Other EOS commands

```
# copy a single file
eos cp /eos/experiment/neutplatform/.../histo.root /tmp/
# using xrdcp
xrdcp histo.root root://eospublic.cern.ch/eos/experiment/neutplatform/.../<destination_dir>
xrdcp root://eospublic.cern.ch/eos/experiment/neutplatform/.../<destination_dir> /tmp/
```

CENFStorageAtCERN < CENF < TWiki

```
# copy all files within a directory - no subdirectories
eos cp /eos/experiment/neutplatform/.../histodirectory/ /tmp/histodirectory

# copy recursive the complete hierarchy in a directory
eos cp -r /eos/experiment/neutplatform/.../histodirectory/ /tmp/histodirectory
# using xrdcp
xrdcp -f -R root://eospublic.cern.ch//eos/experiment/neutplatform/.../histodirectory/ /tmp/histodirectory

# copy recursive the complete hierarchy into the directory 'histodirectory' in the current local
eos cp -r /eos/experiment/neutplatform/.../histodirectory/ histodirectory
# using xrdcp
xrdcp -f -R root://eospublic.cern.ch//eos/experiment/neutplatform/.../histodirectory/ histodirectory

# copy an WEB file [ currently the reported copy size is 0 ]
eos cp http://root.cern.ch/files/<file>.root /tmp/

// If you just need to create an empty directory:
eos mkdir /eos/experiment/neutplatform/.../dir1/dir2/..
```

How to list directories and files in EOS

```
// Verify first that you are connecting to 'your' experiment dedicated EOS instance
// Creating file lists
You can run 'find' commands in EOS or XRootD storage like CASTOR using the 'eos' CLI. This command
# find all files under an EOS subdirectory (if you are an ordinary user, the file list is limited)
eos find -f /eos/experiment/neutplatform/

# find all directories
eos find -d /eos/experiment/neutplatform/

# find all files in a CASTOR directory
eos find -f root://castorpublic//castor/cern.ch/user/<login_name>/<dir_name>/

# find all files on a mounted file system
eos find -f file:/afs/cern.ch/user/<initial>/<login_name>/<dir_name>/

// Listing directories
# list files in eos
eos ls [-la] /eos/experiment/neutplatform/

# An alternative command is:
xrd eospublic dirlist /eos/experiment/neutplatform/

# list files in castor [ if you use '-la' be aware that the ownership and permissions shown are not yours ]
eos ls [-la] root://castorpublic//castor/cern.ch/user/<login_name>/<dir_name>/
```

Writing and reading files on EOS

Files on EOS can also be accessed using [xrootd](#), in particular the [xrdcp](#) command.

xrdcp can be used to copy a file to and from any EOS directory the user has reading/writing access to. The syntax is very simple and foresees a mandatory prefix that is composed of the protocol (root) and the name of the xrootd redirector (eospublic): `root://eospublic/` (notice that from outside CERN you should include the domain part in the redirector name, i.e. `root://eospublic.cern.ch/`).

So, for example, suppose you want to copy in your current directory the file which resides at

```
/eos/experiment/neutplatform/dune/mcproddisk/mcc7/kaons/1GeV/reco/13301499_1X/mu_80_20160826T1132
```

Then what you should do is simply:

```
xrdcp root://eospublic//eos/experiment/neutplatform/dune/mcproddisk/mcc7/kaons/1GeV/reco/13301499
```

Root files can also be accessed directly in ROOT using the root protocol:

Other EOS commands

```
TFile *file = TFile::Open("root://eospublic//eos/experiment/neutplatform/.../source_filename");
```

Deleting files and directories on EOS

```
// Files on !EOS can be deleted with the intuitive command :
```

```
eos rm /eos/experiment/neutplatform/.../filename=
```

```
// The rm command accepts the option -r for recursive deletion.
```

```
// This will also delete the directory it's called on:
```

```
bash$ eos rm -r /eos/experiment/neutplatform/.../directory_to_remove
```

```
// Don't use mv (move) on EOS.
```

"Moving" files or directories on EOS is NOT supported.

You can copy from one place to another, you can **delete** but you cannot move.

This is like **this** by design: moving would imply an actual recalculation of the accounting (which is based on paths in EOS) and **this** would crash the system.

xrdcp man page [↗](#)

EOS Knowledge Based articles and documentation

- EOS quick tutorial for beginners [↗](#)
- EOS Ordered FAQs List [↗](#)
- How to best access EOS from central services (LXPLUS and LXBATCH) [↗](#)
- How do I access the EOS storage underlying my CERNBox from MacOSX? [↗](#)
- EOS Documentation [↗](#)

[↩](#) Back to Neutrino Platform Computing Main Page

Major updates:

-- NectarB - 21-Sept-2016

This topic: CENF > CENFStorageAtCERN

Topic revision: r33 - 2018-09-14 - NectarB



Copyright &© 2008-2020 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback