

Table of Contents

Joining a dCache-based SE to the Xrootd service.....	1
Installation of Federation Host.....	1
Configuration.....	2
Operating xrootd.....	2
Port usage:.....	3
Testing the install.....	3
Configuring Authenticated Access.....	3
Configuring the CMS TFC Plugin in D-Cache.....	4
Federation host.....	4
dCache xrootd Door.....	5
Configuring the Monitoring Plugin.....	6
Useful Links.....	7

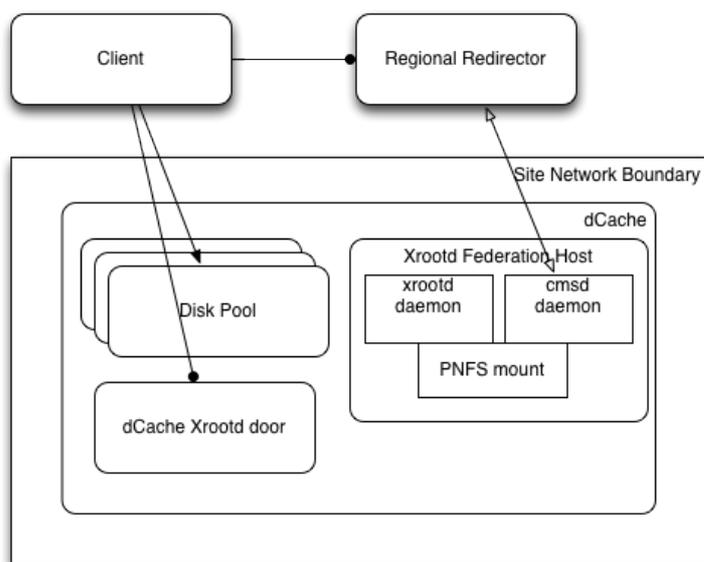
Joining a dCache-based SE to the Xrootd service.

This document covers joining a dCache-based storage element to the CMS Xrootd storage federation (AAA). This page assumes three things:

1. You are using a recent dCache version (any thing beyond 4.2).
2. All your pool nodes are on the public internet.
3. Only EL7 installations are covered (although other Linux flavors should be fine).

If you have pool nodes on a private network, you can still use this page to configure a proxy. For scalability reasons this not really recommended.

The architecture setup is diagrammed below:



This architecture uses the built-in dCache Xrootd door and adds a "federation host", which runs native xrootd components. This host integrates the dCache door with the global federation, but effectively all clients are redirected directly to the dCache xrootd door, then to the individual pools. GSI security and namespace translation are performed by dCache itself. Optionally also the xrootd federation host can be GSI enabled, to avoid exposing the namespace unprotected. At no point does data have to be "proxied", which should improve the scalability and remove complexity from the entire system.

Installation of Federation Host

For almost all configurations you need at least a few RPMs that are provided via the OSG repository. The xrootd components can be installed also via EPEL [\[1\]](#), which is likely the preferred way for non-OSG sites.

The federation host needs a Grid host certificate to authenticate itself. Procedures to obtain Grid certificates vary from country to country and are therefore not covered here.

Please also refer to the OSG admin documentation [\[2\]](#).

Since some OSG have also dependencies on EPEL, you need to install it:

```
yum install https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm
```

Install the OSG software repository.

```
yum install https://repo.opensciencegrid.org/osg/3.4/osg-3.4-el7-release-latest.rpm
```

If you want to install mostly from EPEL, set the repository to a priority lower than 98, by adding or adjusting "priority=95" in /etc/yum.repos.d/epel.repo. If you want all the components from OSG, set the priority for EPEL to "priority=99".

The host needs also some root certificates of the various certification authorities. There are several ways to obtain those.

- CVMFS grid.cern.ch
- IGTF CA Release [↗](#)
- OSG CA certificates [↗](#)

Install and enable certificate revocation lists

```
yum install fetch-crl
systemctl start fetch-crl-cron
systemctl enable fetch-crl-cron
```

Install xrootd and components...

Configuration

First, setup your dCache Xrootd door according to the instructions in the dCache book [↗](#). For the simple **unauthenticated** access it sufficient to add a proper prefix in order to make sure you set the root path so dCache will do the LFN to PFN translation. Add something according to your local setup to the layout file of the Xrootd door.

```
xrootd.root=/pnfs/example.com/data/cms
```

Configuring Authenticated Access is a bit more complex.

Next, cp /etc/xrootd/xrootd.sample.dcache.cfg /etc/xrootd/xrootd-clustered.cfg and edit the resulting config file.

```
oss.localroot /pnfs/example.com/data/cms
xrootd.redirect xrootd-door.example.com:1094 /
```

Set xrootd-door.example.com to the hostname of dCache's xrootd door and /pnfs/example.com/data/cms to match your xrootdRootPath above.

Operating xrootd

PNFS must be mounted for the xrootd federation host to function. Mount this manually, and configure /etc/fstab so this happens on boot if desired.

There are two init services, xrootd and cmsd, which must both be working for the site to participate in the xrootd service:

```
systemctl start xrootd
systemctl start cmsd
```

Everything is controlled by a proper init script (available commands are start, stop, restart, status, and condrestart). To enable these on boot, run:

```
systemctl enable xrootd
systemctl enable cmsd
```

Log files are kept in `/var/log/xrootd/{cmsd,xrootd}.log`, and are auto-rotated.

After startup, the xrootd and cmsd daemons drop privilege to the xrootd user.

Port usage:

The following information is probably needed for sites with strict firewalls:

- The xrootd server listens on TCP port 1095 (this is not the default port for Xrootd; we assume that dCache Xrootd door uses the default).
 - ◆ If you deploy the xrootd Federation host and the dCache xrootd door on different hosts, both can use the default port 1094.
- The cmsd server needs outgoing TCP port 1213 to xrootd.unl.edu.
- Usage statistics are sent to xrootd.t2.ucsd.edu on UDP ports 9931 and 9930.

Testing the install.

The newly installed server can be tested directly using:

```
xrdcp -d 1 -f xroot://local_hostname.example.com//store/foo/bar /dev/null
```

You will need a grid certificate installed in your user account for the above to work

You can then see if your server is participating properly in the xrootd service by checking:

```
xrdcp root://xrootd-itb.unl.edu//store/foo/bar /tmp/bar2
```

where `/store/foo/bar` is unique to your site

Configuring Authenticated Access

Authentication in D-Cache is (usually) done using GPLAZMA. The door for GSI enabled access needs a host certificate. The dCache xrootd door can be deployed with other doors, e.g. GridFTP. Please follow the dCache book [to configure GPLAZMA](#).

Put proper mappings and usernames in `/etc/grid-security/grid-vorolemap`. **Needs adaption to local setup!**. (Only CMS part is shown, if other VOs are needed on the Xrootd door, add them accordingly.)

```
## CMS ##
# Need mapping for each VOMS Group(!), roles only for special mapping
"*" "/cms/Role=lcgadmin" cmsusr001
"*" "/cms/Role=production" cmsprd001
"*" "/cms/Role=priorityuser" cmsana001
"*" "/cms/Role=pilot" cmsusr001
"*" "/cms/Role=hiproduction" cmsprd001
"*" "/cms/dcms/Role=cmsphedex" cmsprd001
"*" "/cms/integration" cmsusr001
"*" "/cms/becms" cmsusr001
"*" "/cms/dcms" cmsusr001
"*" "/cms/escms" cmsusr001
```

```

"*" "/cms/ptcms" cmsusr001
"*" "/cms/itcms" cmsusr001
"*" "/cms/frcms" cmsusr001
"*" "/cms/production" cmsusr001
"*" "/cms/muon" cmsusr001
"*" "/cms/twcms" cmsusr001
"*" "/cms/uscms" cmsusr001
"*" "/cms/ALARM" cmsusr001
"*" "/cms/TEAM" cmsusr001
"*" "/cms/dbs" cmsusr001
"*" "/cms/uscms/Role=cmsphedex" cmsusr001
"*" "/cms" cmsusr001

```

Setup `/etc/grid-security/storage-authzdb`. **Carefully check the usernames and UIDs GIDs, they must fit your local setup.** (Again only CMS part is shown.) Please refer once more to the dCache book [for details](#), how to set it up.

```

authorize cmsusr001 read-write 40501 4050 / / /
authorize cmsprd001 read-write 40751 4075 / / /
authorize cmsana001 read-write 40951 4060 / / /

```

You can do some first testing of the GSI enabled Xrootd door:

```
xrdcp -d 2 -f xroot://xrootd-door.mydomain.org:/store/user/<Your_HN_name>/<Your_Testfile> /dev/nu
```

Some useful debugging results are usually found in the billing logs of your D-Cache instance. The host is usually not the host you are installing the Xrootd door on.

```
/var/lib/dcache/billing/<YEAR>/
```

Configuring the CMS TFC Plugin in D-Cache

D-Cache provides a TFC Plugin such that you can send an LFN open request to the D-Cache xrootd-door and the door will resolve it to a PFN based on TFC rules. You need to resolve the LFN actually twice, on the Federation host (native xrootd) and the dCache xrootd door.

Federation host

You can also run this on the dCache xrootd door, in this case make sure that different ports are used for the Federation and the dCache xrootd door.

Install the xrootd TFC plugin. Note that this is only available from the OSG repos, it is not included in EPEL.

```
yum install xrootd-cmstfc
```

Setup the configuration `/etc/xrootd/xrootd-clustered.cfg`

```

# Port specifications; only the redirector needs to use a well-known port
# Change as needed for firewalls.
# Make sure that dCache xrootd door and Federation use different ports, when deployed on same host
xrd.port 1094

# The roles this server will play.
all.role server

# European Redirector
all.manager any xrootd-cms.infn.it+ 1213

# Site name for monitoring

```

```

all.sitename T2_XY_SiteName

# redirect to dCache xrootd door (adjust hostname and port)
xrootd.redirect dcache-xrootd-door.mysite.com:1094 /

# Allow any path to be exported
all.export / nostage

# Hosts allowed to use this xrootd cluster
cms.allow host *

### Standard directives
# Simple sites probably don't need to touch these.
# Logging verbosity
xrootd.trace msg login stall redirect
ofs.trace all
xrd.trace conn
cms.trace all

# Some tuning for disk space monitoring
# This is a pure redirector needs no real storage space
cms.space linger 0 recalc 30 min 2% 1g 5% 2g

# Integrate with CMS TFC, placed in /etc/storage.xml - protocol might be different, depends on ac
oss.namelib /usr/lib64/libXrdCmsTfc.so file:/etc/xrootd/storage.xml?protocol=direct

# Turn on authorization
ofs.authorize 1
acc.authdb /etc/xrootd/Authfile
#acc.audit deny grant

# Require GSI on Federation host
sec.protocol /usr/lib64 gsi -d:1 -crl:0 -authzfun:libXrdLcmaps.so -authzfunparms:--loglevel,1 -gm

xrootd.seclib /usr/lib64/libXrdSec.so
xrootd.fslib /usr/lib64/libXrdOfs.so
all.adminpath /var/run/xrootd
all.pidpath /var/run/xrootd

cms.delay startup 10
cms.fxhold 60s
#cms.perf int 30s pgm /usr/bin/XrdOlbMonPerf 30

if exec xrootd
# Summary monitoring configuration
  xrd.report xrootd.t2.ucsd.edu:9931 every 60s all sync
# Detailed monitoring configuration
  xrootd.monitor all fstat 60s lfn ops ssq xfr 5 ident 5m dest fstat info user CMS-AAA-EU-COLLECC
fi

```

dCache xrootd Door

For the host that runs the dCache xrootd door you need the TFC plugin. It is provided in the Download Area from dcache.org.

The RPM can be installed like this

```
rpm -ivh xrootd4j-cms-plugin-1.3.7-1.noarch.rpm
```

The following configuration parameters should be added to `/etc/dcache/dcache.conf`. The site name should be your CMS site name. Note that these settings need to be available also on all Pool nodes to generate proper monitoring messages.

```
pool.mover.xrootd.plugins=edu.uchicago.monitor
```

Federation host

```
# The following two lines are the values for EU sites
xrootd.monitor.detailed=cms-aaa-eu-collector.cern.ch:9330:60
xrootd.monitor.summary=xrootd.t2.ucsd.edu:9931:60
xrootd.monitor.vo=CMS
xrootd.monitor.site=T2_XY_MySite
```

The following should be added to the layout file of the machine(s) that host(s) the xrootd door(s), `/etc/dcache/layouts/dcache-my-xrootd-door.layout.conf` (adjust the host name). The location of the TFC file (typically named `storage.xml`) might be adjusted. The protocol might also be different for you TFC, it is just an identifier in the end.

```
[xrootd-${host.name}Domain]
[xrootd-${host.name}Domain/xrootd]
xrootd.plugins=gpplazma:gsi,authz:cms-tfc
xrootd.cms.tfc.path=/etc/dcache/storage.xml
xrootd.cms.tfc.protocol=xrootd
```

Test your setup 😊

Configuring the Monitoring Plugin

dCache can emit monitoring information similar to the SLAC Xrootd implementation. The process of enabling this is documented on the following page:

<https://twiki.cern.ch/twiki/bin/viewauth/AtlasComputing/FAXdCacheN2Nstorage>

Useful Links.

- [dCache-native TFC implementation](#)
-

This topic: [Main > DcacheXrootdRecent](#)

Topic revision: [r33](#) - 2019-11-13 - [ChristophWissing](#)



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](#)