



# dpm-xrootd status for xroot federations

**David Smith**  
CERN

DPM Workshop

07/12/2015

# Outline

- Puppet configuration
  - How to join ATLAS and CMS xroot federations
- Upcoming release
  - Changes in processing federation requests, to improve performance and stability

# Puppet configuration

All federation specific configuration is on the head node

```
$atlas_fed = {  
  name           => 'fedredirect_atlas',  
  fed_host       => 'atlas-xrd-XX.cern.ch',  
  xrootd_port    => 1094,  
  cmsd_port      => 1098,  
  local_port     => 11000,  
  namelib_prefix => "/dpm/cern.ch/home/atlas",  
  namelib        => "XrdOucName2NameLFC.so pssorigin=localhost  
sitename=ATLAS_SITENAME",  
  paths         => [ '/atlas' ]  
}
```

Ask ATLAS for definitive name of the regional redirector (i.e. *fed\_host* variable) to join.

# Puppet configuration

```
$cms_fed = {  
  name           => 'fedredir_cms',  
  fed_host       => 'xrootd-cms.infn.it',  
  xrootd_port    => 1094,  
  cmsd_port      => 1213,  
  local_port     => 11001,  
  namelib_prefix => "/dpm/cern.ch/home/cms",  
  namelib        => "libXrdCmsTfc.so"  
file:/etc/xrootd/storage.xml?protocol=direct",  
  paths          => [ '/store' ]  
}
```

Ask CMS for definitive name of the regional redirector (i.e. *fed\_host* variable) to join.

# Puppet configuration

```
class{"dmlite::xrootd":
  nodetype           => [ 'head','disk' ],
  domain             => "${localdomain}",
  dpm_xrootd_debug   => $debug,
  dpm_xrootd_sharedkey => "${xrootd_sharedkey}",
  dpm_xrootd_fedredirs => { "atlas" => $atlas_fed, "cms" =>
  $cms_fed },
  xrd_report         => "xrootd.t2.ucsd.edu:9931,atl-
  prod05.slac.stanford.edu:9931 every 60s all sync",
  xrootd_monitor    => "all flush 30s ident 5m fstat 60 lfn
  ops ssq xfr 5 window 5s dest fstat info user redir CMS-AAA-EU-
  COLLECTOR.cern.ch:9330 dest fstat info user redir atlas-fax-eu-
  collector.cern.ch:9330"
}
```

Add to your `dmlite::xrootd` block: This connects up the `$atlas_fed` and/or `$cms_fed` variables, and it includes the `xrd_report` and `xrootd_monitor` settings for the experiments' federation monitoring. Check the puppet wiki for the monitoring settings for ATLAS alone or CMS alone.

# Other configuration (head node)

- ATLAS federation needs:
  - xrootd-server-atlas-n2n-plugin rpm from the wlcg repository (current v2.2-0)
- CMS federation needs:
  - xrootd-cmstfc from the OSG repository (current v1.5.1-10)
  - Write /etc/xrootd/storage.xml like this:

```
<lfn-to-pfn protocol="direct" destination-match=".*" path-  
match="/+store/test/xrootd/<CMS Site Name>/store/(.*)" "  
result="/dpm/<domain>/home/cms/store/$1"/>  
<lfn-to-pfn protocol="direct" destination-match=".*" path-  
match="/+store/(.*)" result="/dpm/<domain>/home/cms/store/$1"/>
```

# Upcoming release

- Next release will be dpm-xrootd 3.6.0:
  - Expected to be in epel-testing before Christmas
  - To epel stable in January.
- Changes in 3.6.0:
  - Fix an RPM dependency
  - Add option to the configuration to make the cluster ID unique
  - **Introduce statinfo library and support a new Name2Name interface: performance and stability related**



# cmsd and dpm-xrootd

- A daemon called cmsd runs on the head node for each federation joined.
- It answers queries from a regional redirector to discover if a file is available at the site.

Query: /store/thefile



*stat /store/thefile*  
(the cmsd uses an xrootd client)

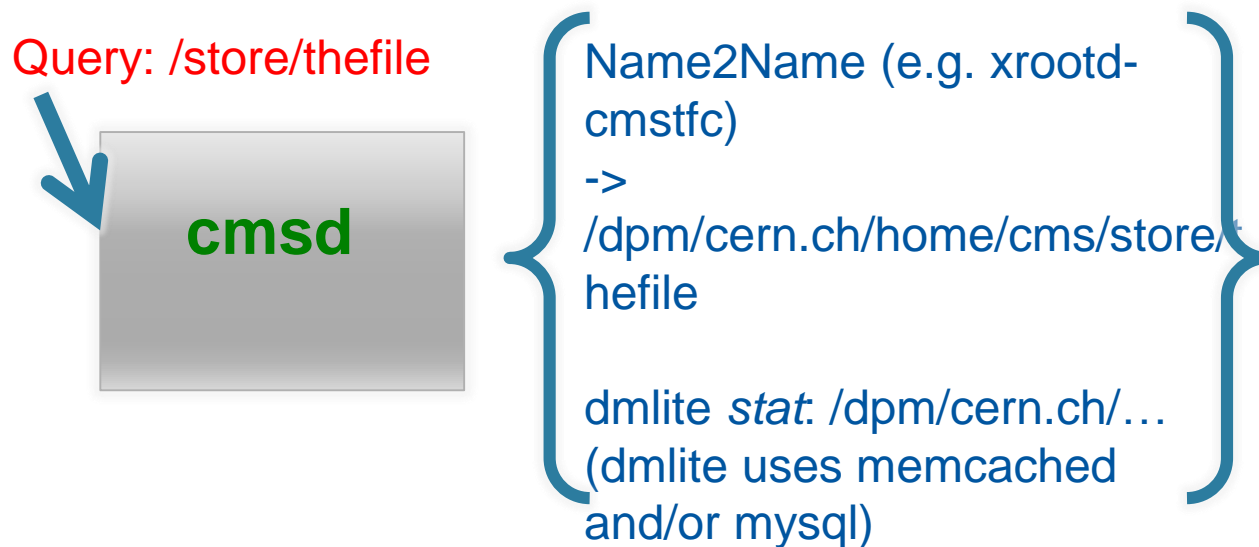
**Inside the xrootd**

```
Name2Name (e.g. xrootd-  
cmstfc)  
->  
/dpm/cern.ch/home/cms/store/  
hefile  
  
dmlite stat: /dpm/cern.ch/...
```

(memcached -> mysql)

# New plugin for the cmsd

- Avoid the use of the xroot proxy at the cmsd: use a new plugin called **statinfo**



# name2name internal queries



- Some name2name libraries, e.g. xrootd-server-atlas-n2n-plugin, do name translation differently
- A name may resolve into several possibilities: the name2name library itself makes a call to find the first one which exists

# name2name internal queries

Query (from cmsd)

or

open (from client)



Name2Name (e.g. xrootd-server-atlas-n2n-plugin)

->

/dpm/cern.ch/home/atlas/atlasdatadisk/scope/ac/45/thefile  
(check)

->

/dpm/cern.ch/home/atlas/atlasgroupdisk/scope/ac/45/thefil  
e (check)

The first one which exists:

dmlite *stat/open*: /dpm/cern.ch/...

(dmlite uses memcached and/or mysql)

dmlite *stat*: /dpm/cern.ch/home/atlas/atlasdatadisk/...

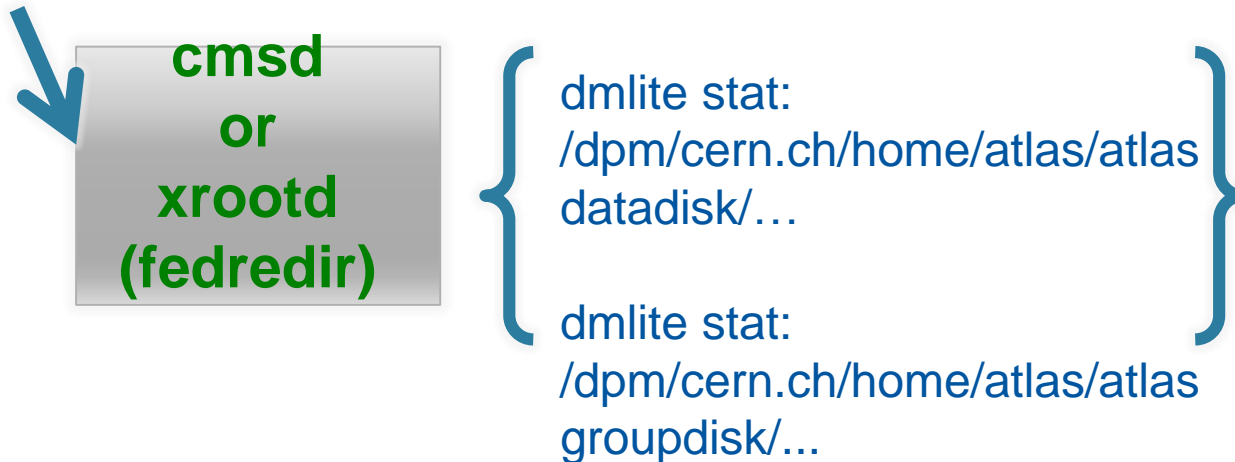
dmlite *stat*: /dpm/cern.ch/home/atlas/atlasgroupdisk/...

...

# new: name2name interface

- New class XrdOucName2NameVec
- ```
std::vector<std::string *> *n2nVec(const char *glfn);
```

Query or open:  
`/atlas/rucio/scope:thefile`



# Summary

- See puppet documentation:
- <https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm/Xroot/PuppetSetup>
- Federation performance:
  - Eliminate internal use of xroot client in the dpm-xrootd setup to make calls between cmsd/xrootd daemons on the headnode
  - One source of such calls, by the cmsd, will be removed by a new component included in dpm-xrootd 3.6.0 (will require a change in puppet module to active it)
  - The other source of calls is internal to some N2N libs: will be removed by use of a new interface, which dpm-xrootd 3.6.0 will support.