SRM2.2 deployment workshop - dCache tutorial session

Greig Cowan and Timur Perelmutov

14th November 2007

1 Logging on

Use an ssh client on your laptop to log on to your dCache machine.

- Username: `labuser`
- Hostname: `lab-XX.nesc.ed.ac.uk`
- Password: `fruit48BAT`
- Root access via sudo: `sudo bash`

2 SRM2.2 configuration

The developers will talk through the necessary commands that must be issued to configure SRM2.2 spaces. Use PPS Systems Requirements as a basis

2.1 Atlas Requirements

We want the following users with these VOMs extensions to be able to write to reserve space in the system. We include the dteam001 user so that we are able to test the system with the proxy certificate that is on the system (see Section on Client tools).

```
/atlas/Role=production
/atlas/soft-valid/Role=production
dteam001/Role=*
```

2.1.1 Space Tokens with Descriptions

We will use the following space token descriptions for the storage classes. The user certificate maps us to the dteam001 user, hence the use of the path below which is in the dteam part of the PNFS namespace. This is accecptable since thisnis only a test.

- `TAPE` for the CUSTODIAL-NEARLINE storage class
  - `/pnfs/nesc.ed.ac.uk/data/dteam/atlas/tape`
- `DISK` for the REPLICA-ONLINE storage class Paths
  - `/pnfs/nesc.ed.ac.uk/data/dteam/atlas/disk`
- *https://twiki.cern.ch/twiki/bin/view/LCG/GSSDATLASPPS*

2.2 LHCb requirements

We want the following users with these VOMs extensions to be able to write to reserve space in the system. We include the dteam001 user so that we are able to test the system with the proxy certificate that is on the system (see Section on Client tools).

```
/lhcb/lcgprod
dteam001/Role=*
```

2.2.1 Space Tokens with Descriptions

We will use the following space token descriptions for the storage classes. The user certificate maps us to the dteam001 user, hence the use of the path below which is in the dteam part of the PNFS namespace. This is accecptable since this is only a test system.

- `LHCb_RAW` and `LHCb_RDST` for the CUSTODIAL-NEARLINE storage class.
- *https://twiki.cern.ch/twiki/bin/view/LCG/GSSDLHCBPPS*

2.3 PoolManager.conf

You need to make the following changes to `/opt/d-cache/config/PoolManager.conf`. This is where we tell the system about the pool configuration and the properties of the spaces.

```
pools and links psu create pool lab-XX_1 psu create pool lab-XX_2
psu create pgroup dteam-disk
psu addto pgroup dteam-disk lab-XX_1
psu create pgroup dteam-tape
psu addto pgroup dteam-tape lab-XX_2
psu create link dteam-link-disk world-net dteam-groups
psu set link dteam-link-disk -readpref=20 -writepref=20 -cachepref=20
psu add link dteam-link-disk dteam-disk

psu create link dteam-link-tape world-net dteam-groups
psu set link dteam-link-tape -readpref=20 -writepref=20 -cachepref=20
psu add link dteam-link-tape dteam-tape

PoolManager.conf link groups
psu create linkGroup dteam-tape-link-group
psu addto linkGroup dteam-tape-link-group dteam-tape-link
psu set linkGroup custodialAllowed dteam-tape-link-group true
psu set linkGroup outputAllowed dteam-tape-link-group false
psu set linkGroup replicaAllowed dteam-tape-link-group false
psu set linkGroup onlineAllowed dteam-tape-link-group false
psu set linkGroup nearlineAllowed dteam-tape-link-group true

psu create linkGroup dteam-disk-link-group
psu addto linkGroup dteam-disk-link-group dteam-disk-link
psu set linkGroup custodialAllowed dteam-disk-link-group false
psu set linkGroup outputAllowed dteam-disk-link-group true
psu set linkGroup replicaAllowed dteam-disk-link-group true
psu set linkGroup onlineAllowed dteam-disk-link-group true
psu set linkGroup nearlineAllowed dteam-disk-link-group true
```

2.4 Reload PoolManager.conf

We must tell the dCache about the new configuration so we should login to the administration interface (use the command line or the GUI from your laptop). This is an important message from this session: **if you are not already, then you really should become very familiar with the use of the admin interface for controlling and understanding the dCache system. Most problems can be fixed by inspecting cells and the info that they produce.**

```
# ssh adminlocalhost -c blowfish -p 22223
admin@localhost's password:  dickerelch

   dCache Admin (VII) (user=admin)

(local) admin > cd PoolManager
(PoolManager) admin > reload -yes
```

Now have a look at the new configuration by running some standard commands.

```
   (PoolManager) admin > psu ls -l link
(PoolManager) admin > psu ls -l linkGroup
```

2.5 dCacheSetup

Now that the pool configuration is complete, we need to make the relevant changes to the main system configuration file `/opt/d-cache/config/dCacheSetup`. Open up the file in your favourite editor (well, `vi`) and ensure the following options are set. Make sure to ask if you do not know what any of them mean.

```
srmSpaceManagerEnabled=yes
srmImplicitSpaceManagerEnabled=yes
SpaceManagerDefaultRetentionPolicy=CUSTODIAL
SpaceManagerDefaultAccessLatency=NEARLINE
SpaceManagerLinkGroupAuthorizationFileName=/opt/d-cache/etc/LinkGroupAuthorization.conf
```

2.6 LinkGroupAuthorization.conf

`/opt/d-cache/etc/LinkGroupAuthorization.conf` is a new configuration file that determines precisely which users can make space reservations in the link groups that you defined in the PoolManager.conf file. Again, using `vi` open the file.

```
LinkGroup dteam-tape-link-group
/atlas/Role=production
/atlas/soft-valid/Role=production
dteam001/Role=*

LinkGroup dteam-disk-link-group
/lhcb/lcgprod
dteam001/Role=*
```

2.7 Restart dCache, SRM

You should now force dCache to pick up all of the changes that you have made by restarting.

```
/opt/d-cache/bin/dcache-core stop
/opt/d-cache/bin/dcache-core start
```

3 Testing with the client tools

If you want to test out your SRM2.2 configuration you will find the latest (not even in production!) versions of the lcg-* and srm* client tools available. Be aware that the lcg-* commands come from EGEE and provided by the `lcg_util` package which depends on GFAL (Grid File Access Library). The srm* clients are written in Java by FNAL developers. They are distributed as part of dCache.

```
$ rpm -qa|grep lcg_util
lcg_util-1.6.3-1.slc4
$ rpm -qa|grep GFAL
GFAL-client-1.10.4-1.slc4
$ rpm -qa|grep dcache-client
dcache-client-1.8.0-0
```

Older versions of `lcg_util/GFAL` always required a lookup of the BDII in order to work out which storage element (SE) endpoint to use. The latest version of the client tools have a option (`-b`) which switches off the calls to the BDII. Will try both methods in the examples below. The BDII in use is defined by:

```
$ echo $LCG_GFAL_INFOSYS
wn3.epcc.ed.ac.uk:2170
```

The srm* tools never interact with the BDII or LFC.

3.1 Procedure

Before running the clients you should switch to the `labuser` account.

- `$ su - labuser`
- `$ grid-proxy-init`
- Password: `SRMWorkshop`

The SRM only starts once you make an initial connection to it. The easiest way of doing this is to srmPing the system.

```
$ srmping -2 -debug srm://lab-22.nesc.ed.ac.uk:8443/pnfs/nesc.ed.ac.uk
```

Now check the admin interface again.

```
ssh localhost -l admin -c blowfish -p 22223
(local) admin > cd SrmSpaceManager
(SrmSpaceManager) admin > ls
Reservations:
total number of reservations:  0
total number of bytes reserved:  0
LinkGroups:
0 Name:dteam-disk-link-group FreeSpace:2147483648...
1 Name:dteam-tape-link-group FreeSpace:2147483648...
total number of linkGroups:  2
total number of bytes reservable:  4294967296
last time all link groups were updated:  1194888073861
```

3.2 Reserve Space

We are now ready to reserve some space in the link groups that we setup. Use the client tool to do this easily.

```
$ srm-reserve-space -space_desc=TAPE -retention_policy=CUSTODIAL -access_latency=NEARLINE
-desired_size=1000 -guaranteed_size=1000 -lifetime=-1 srm://lab-22.nesc.ed.ac.uk:8443/
Space token =2
```

Once the space is reserved, we can get detailed information about it.

```
$ srm-get-space-metadata srm://lab-22.nesc.ed.ac.uk:8443/ 2Space Reservation with token=2
owner:VoGroup=dteam001 VoRole=
totalSize:1000
guaranteedSize:1000
unusedSize:1000
lifetimeAssigned:-1
lifetimeLeft:-1
accessLatency:NEARLINE
retentionPolicy:CUSTODIAL
```

3.3 Writing into the space

Now use the standard srmcp to try copying a file into the new space reservation. Note that you must specify the space token that was returned from the metadata command. It would be better if the human readable space token description could be used, rather than the space token itself (the number).

```
$ srmcp -debug=true -space_token=2 file:////etc/group
srm://lab-22.nesc.ed.ac.uk:8443/pnfs/nesc.ed.ac.uk/data/dteam/testspace
```

We can then list the properties of the file.

```
$ srmls -l srm://lab-22.nesc.ed.ac.uk:8443/pnfs/nesc.ed.ac.uk/data/dteam/testspace
```

```
2505 /pnfs/nesc.ed.ac.uk/data/dteam/testspace
storage type:PERMANENT
retention policy:CUSTODIAL
access latency:NEARLINE
...
```

3.4 Useful links

- *http://trac.dcache.org/trac.cgi/wiki/manuals/SRM_2.2_Setup*
- *http://trac.dcache.org/trac.cgi/wiki/manuals/dCache_clients/client_dcache_srm*
- *http://trac.dcache.org/trac.cgi/wiki/manuals/dCache_clients/client_list*

4 Additional examples of client commands

4.1 srmPing

Make sure that the SRM endpoint is up and responding by using the srmping client. Remember that the dCache SRM2.2 endpoint accessed (by default) on port 8443.

```
$ srmping -debug srm://lab-XX.nesc.ed.ac.uk:8443
```

If the ping is successful (i.e. the SRM responds) and the `-debug` option is not used then there is no output and the client exits with code `0`.

4.2 srmLs

In SRM2.2 there is a method which can be used to list the contents of directories in the namespace.

4.2.1 lcg-* client with BDII

```
$ lcg-cp -D srmv2 -v --vo dteam file:/etc/group
srm://lab-XX.nesc.ed.ac.uk:8446/pnfs/nesc.ed.ac.uk/data/dteam/
```

4.2.2 lcg-* client without BDII

In this case, it is necessary to specify the full endpoint, which includes the `/srm/managerv2?SFN=` string and the `-D` option which defines the type of the destion SRM.

```
$ lcg-ls -D srmv2 -b -v
srm://lab-XX.nesc.ed.ac.uk:8443/srm/managerv2?SFN=/pnfs
```

4.2.3 srm* client

```
$ srmls srm://lab-XX.nesc.ed.ac.uk:8443/pnfs/nesc.ed.ac.uk/data/dteam
```

4.3 srmPut

We can copy a file into the system in the following way:

4.3.1 lcg-* client with BDII

```
$ lcg-cp -D srmv2 -v --vo dteam file:/etc/group
srm://lab-XX.nesc.ed.ac.uk:8443/pnfs/nesc.ed.ac.uk/data/dteam/‘date +%s‘
```

4.3.2 lcg-* client without BDII

```
$ lcg-cp -D srmv2 -b -v --vo dteam file:/etc/group
srm://lab-XX.nesc.ed.ac.uk:8443/srm/managerv2?SFN=/pnfs/nesc.ed.ac.uk/data/dteam/‘date +%s‘
```

4.3.3 srmcp client

```
$ srmcp -debug file:////etc/group
srm://lab-XX.nesc.ed.ac.uk:8443/pnfs/nesc.ed.ac.uk/data/dteam/‘date +%s‘
```

4.4 srmGet

4.5 srmCopy

srmCopy is not supported by the lcg-* client tools. Rather, to copy a file from one SRM to another it performs a third party copy.

5 Note about DPM

For DPM SRM endpoints, note that the SRM1 and SRM2.2 services are accessed via different ports: SRM1 ↔ 8443 and SRM2.2 ↔ 8446 (by default). You do not have to worry about this if you are using the lcg-* clients without the `-b` option (since the relevant port information will be obtained from the information system. Remember that the DPM namespace is typically structured like this:

```
/dpm/domain/home/voName/
```

However, this is not always the case. Again, this information is published by the SE into the information system. If the lcg-* client performs a BDII lookup, then it will obtain the correct path.

6 Conclusion

So you can now all configure dCache SRM2.2 space management and perform space reservation. Congratulations! Note that we will put this material online, probably somewhere in the GridPP or dCache wiki. We will let you know when this happens.