

dCache Troubleshooting

Tigran Mkrtchyan
for dCache Team

This talk compiled on e-mails to:
support@dcache.org
user-forum@dcache.org

There is two major category of Questions:

Troubleshooting and **Tuning**

1. Do NOT customize startup scripts
2. Do NOT copy startup scripts into `/etc/init.d`
(use wrappers or sym-links if you really want to do so)
3. DO NOT customize `.batch` files
(use `dCacheSetup` instead)

Unless you have
been told by us !



Q:

After update 1.7.0-n to 1.7.0-m SRM does not work any more.

A:

While our RPMs does not have 'active' scripts inside, you have to run `install.sh` to finish update.

Q:


After successful installation I can write a file with srmcp/g-u-c, but can not read it back.

A:

The main difference between read and write, that write is an passive transfer, e.g. on write client connects to the door for a data transfer, while on read pool connects to. Please check your firewall configuration.

FTP write

dCache.ORG



Q:

After successful installation I can write a file, but can not read it back.

A:

Please check `/var/log/pnfsDomain.log` for errors. Check that '*companion*' database populated.

```
companion=# select * from cacheinfo where
```

```
pnfsid='003C000000000000001895498';
```

```
pnfsid | pool | ctime
```

```
-----+-----+-----  
003C000000000000001895498 | dcache21_1 | 2006-02-22 01:35:21.839
```

```
(1 row)
```

Q:

TURL returns internal interface.

A:

By default, doors bind to all available interfaces and SRM chooses one of them. Please check 'Complex Network Configuration' chapter of 'The Book'.

Q:

TURL returns 127.0.0.x address.

A:

Please check that /etc/hosts does not contains something like:

```
# /etc/hosts
```

```
127.0.0.1 localhost
```

```
127.0.0.2 myhost # bad entry
```

Q:

I am unable to force srmcp to use dcap, nevertheless, plain dcap works fine.

A:

To avoid dcap on wide area transfers dcap doors are not published to SRM.

Q:

I would like to write into a directory, but getting an error: “No storageInfo available”

A:

Please check that in the destination directory tags 'sGroup' and 'OSMTemplate' exists and not empty.

```
# cat '.(tag)(sGroup)'  
group  
# cat '.(tag)(OSMTemplate)'  
StoreName store  
#
```

store:group@osm

Q:

I have no HSM, but some transfers hangs in 'Tape Restore Queue'

A:

File located on a pool which is off-line or was off-line at request time. Check that file reregistered by *companion*. Use *retry* in PoolManager after pool is back.

Q:

Is there a way to control log files?

A:

we successfully using *logrotate* utility.
In addition, version 1.7.0 has a log4j support.

Tuning (who is who)

dCache.ORG

LM (LocationManager)	lmDomain	Knows locations of all dCache components (cells)
PoolManager	dCacheDomain	Knows all pools, selects a pool for a transfer according static configuration for PoolManager.conf and dynamic costs information from pools.
PnfsManager	pnfsDomain	Responsible for all operations related with the namespace. Provides other components with informations like file attributes, storageInfo, cacheInfo. Uses pnfs as the information source.
companion	part of pnfsManager	Addon to PnfsManager, which allows to store cacheInfo in dedicated database instead of inside pnfs.
Cleaner	pnfsDomain	Cleans pools from files removed from namespace
gPlazma	gPlazmaDomain	Authentication component. Maps a user DN[+VO] to the local account
SRM	SrmDomain since 1.7.0 works in tomcat	WS-based door for SRM protocol
LoginBroker	httpdDomain	Keeps the list of all active doors. Used by SRM and httpd.
PinManager	utilityDomain	Pins files in pools to keep them on disk. Used by SRM.
ReplicaManager	replicaDomain	Resilien manager. Keeps N copy of files located in resilienPoolGroup , but not
RemoteGsiftTransferManager	srmDomain	keep track of transfers, when pools interacts as an ftpClient

Q:

Which file system we have to use on the pools?

A:

XFS for a linux box and SUN's ZFS for Solaris it the best choice for now. They easily provide 30–50 MB/s on single stream on a good quality hardware.

Q:

howto move *data disk* from one pool to another?

A:

after attaching the disk to the new hosts:

- adjust x.poollist file;
- start the pools;
- run *pnfs register* in the new pool or

```
# UPDATE cacheinfo SET pool='<new pool name>'
WHERE pool='<old pool name>';
```


Q:

howto find out which files was located on crushed disk ?

A:

in companion database:

```
# SELECT * FROM cacheinfo WHERE pool='<dead pool>';
```

Q:

howto switch *companion* on?

A:

- set *cacheinfo=companion* in dCacheSetup file;

- create a database called companion:

```
psql -U postgres -O srmdache companion
```

- initialize companion:

```
psql -U srmdache <  
/opt/d-cache/etc/psql_install_companion.sql
```

- run *pnfs register* on all pools;

- restart pnfsManager

Q:

howto move the *PNFS-gdbm* to a different host ?

A:

- install the pnfs on the new host
- shutdown old pnfs
- copy /opt/pnfsdb to the new host
- start the new pnfs

We recommend to migrate to
PNFS-postgres

Q:

howto move the *PNFS-postgres* to a different host ?

A:

- install the pnfs on the new host
- shutdown old pnfs
- dump postgres databases with *pg_dumpall > pnfs.dump*
- copy `/opt/pnfsdb` to the new host
- restore postgres database with *psql -U postgres template1 < pnfs.dump*
- start the new pnfs

Q:

howto move dCache component to a different host ?

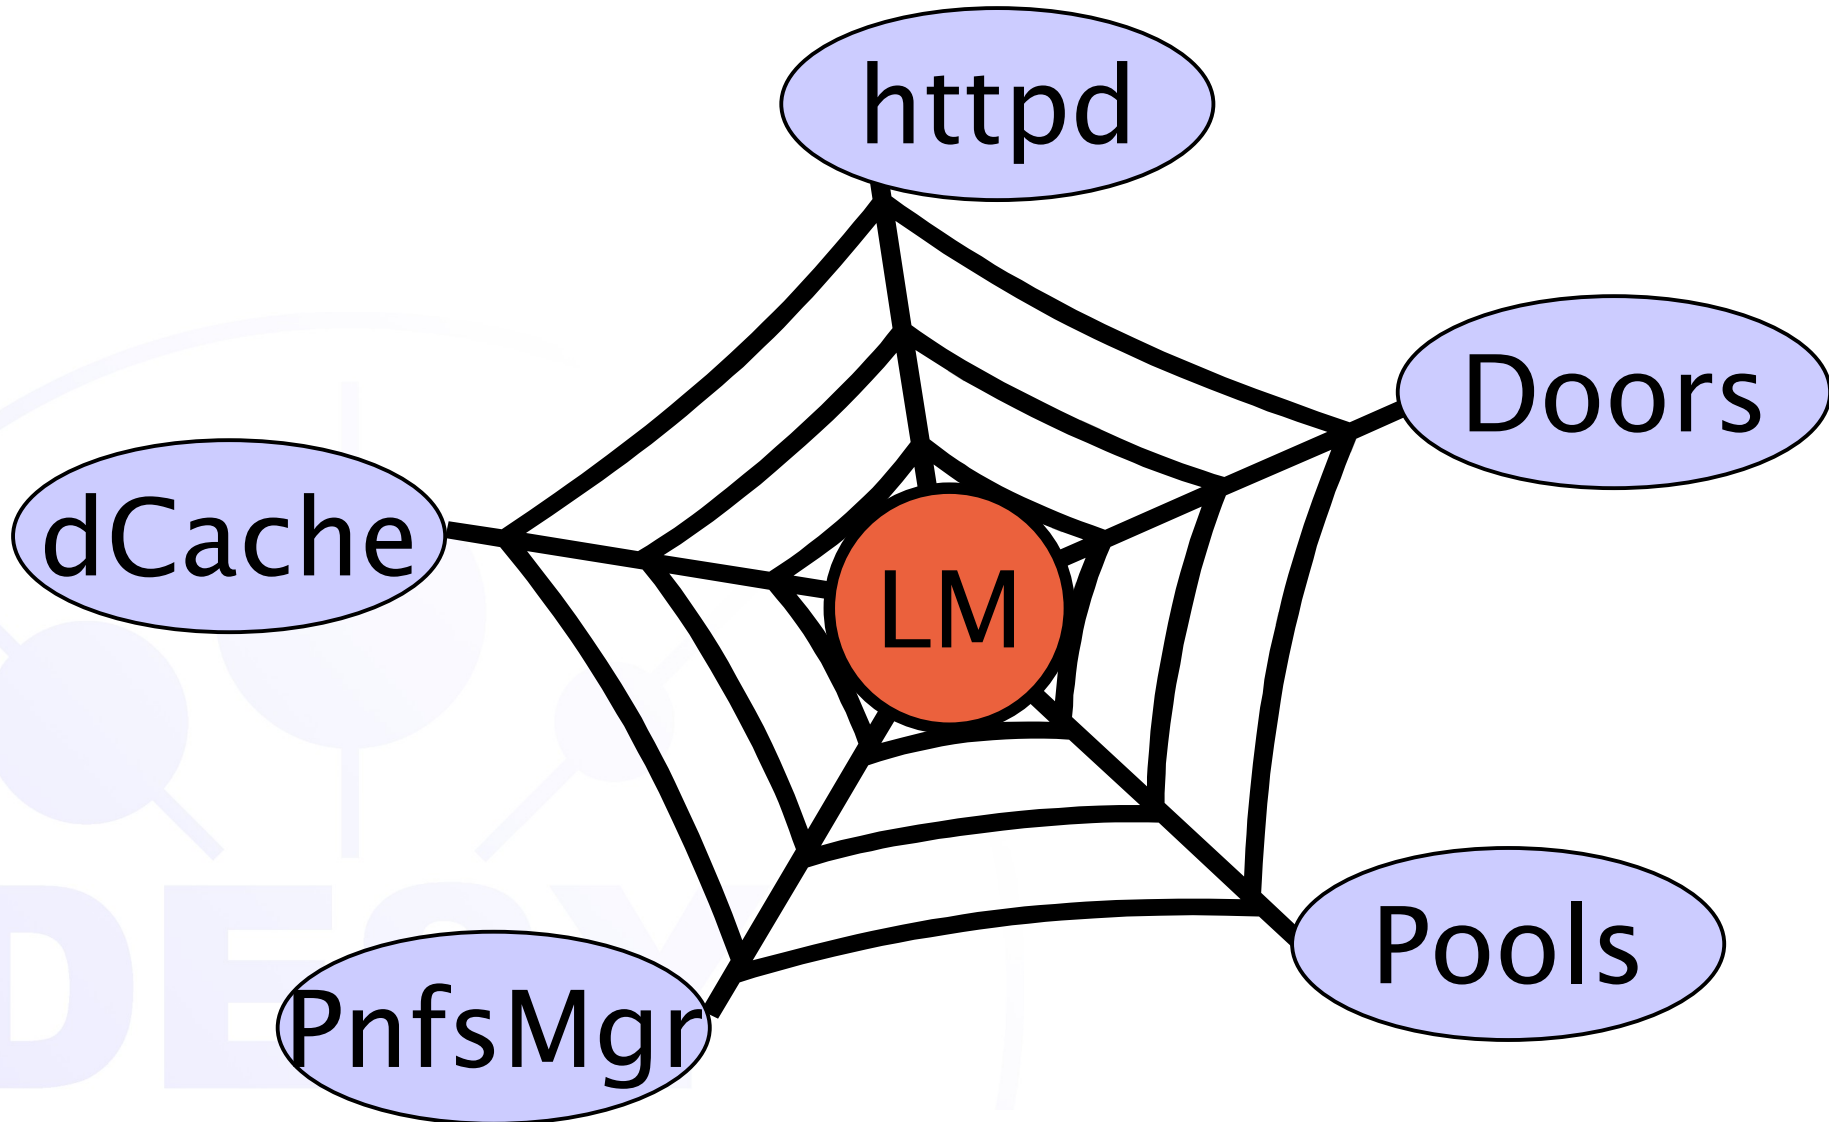
A:

1.7.0-x supports a new node type '*custom*'. Please be sure that set of all services on all custom nodes have to contain:

lm, dCache, adminDoor, pnfs, httpd, utility
Service locator should point to lm host

Tuning (howto)

dCache.ORG



Q:

do I need to mount pnfs on all host ?

A:

pnfsDomain, dirDomain and gridftp-doors requires pnfs to be mounted. In addition, all HSM backhanded write pools have to mount pnfs as well.

Tuning (howto)

dCache.ORG

Q:

How can I do some action on poolGroup

A:

GUI is the best place to do so



Tuning (howto)

dCache.ORG

The screenshot shows the dCache-HH Pool Commander interface. The window title is "Cell Login". The menu bar includes "Session", "Windows", "Specials", and "Help". The main title is "dCache-HH". Below the title are tabs for "Commander", "Restore", "Transfer", "Pools", and "CostModule". The "Pools" tab is active, showing the "Pool Commander" window. On the left, there is a "Pool Groups" list with a search box containing "*". The list includes groups like "hera-b-mc4", "hera-b-pools", "hera-b-rp-00...", "hera-b-write-...", "herab-pools", "hermes-grau-...", "hermes-pools", "hermes-write-...", "large-file-store", "raw-flush", "write-pools", "zeus-dcache-...", "zeus-grau-po...", "zeus-pools", and "zeus-write-po...". Below the list are buttons for "Update Groups" and "Toggle Source". The main area contains a table with columns "Pool", "Time", and "Result". The table lists several pools, with the "Result" column highlighted in green. At the bottom right, there is a "To Your Selection" button.

Pool	Time	Result
it-dcache15-0		
it-dcache3-0		
it-dcache3-1		
it-dcache4-0		
it-dcache4-1		
it-dcache5-0		
it-dcache5-1		
it-dementor2-0		
it-dementor2-1		

Q:

If a file in the *SUSPENDED* mode is there a way to distinguish between:

- Pool is down
- File not in dCache

A:

Not really, but we can add more logic behind for some obvious cases.

Q:

How to change default permission mask with SRM?

A:

There is no way to set something like umask in all URL-based protocols, but SRMv2.2 will have some kind of *chmod methods*.

Q:

which TCP ports used by dCache

A:

50000:52000 => gridftp in pools
33115:33145 => passive dcap & xrootd
8443 => SRM
22223 => ssh
22125 => dcap
22128 => gsidcap
1049 => xrootd
2811 => gridftp

Q:

I am not a PostgreSQL expert....

A:

- version: $\geq 8.1.0$
- disk: as more spindle as possible, RAID 5+1
- keep `<PGDATA>/data` and `<PGDATA>/data/pg_xlog` on different disks:

```
mv pg_xlog /xLog/pg_xlog
```

```
ln -s /xLog/pg_xlog pg_xlog
```

Tuning (howto)

dCache.ORG

in postgres.conf:

```
stats_start_collector = on  
stats_command_string = on  
stats_row_level = on
```

```
autovacuum = on  
autovacuum_naptime = 600  
autovacuum_vacuum_threshold = 1000  
autovacuum_analyze_threshold = 500
```

Linux, PostgreSQL Server with 4GB of RAM

```
shared_buffers = 16384  
sort_mem = 4096  
vacuum_mem = 32168  
work_mem = 8192  
max_fsm_pages = 1000000  
effective_cache_size = 262144  
random_page_cost = 2
```

Change the kernel parameter for maximum shared memory segment size in /etc/sysctl.conf :

```
kernel.shmmax = 288435456
```