



dCache 8.2

Dmitry Litvintsev

FIFE meeting

03/17/23

dCache releases

- We have been running dCache 7.2 since Oct 2021.
- Current golden release (2 year support) is 8.2.
 - Introduces implementation of new WLCG TAPE REST API.
 - Support of multiple authn/authz schemes in the same xrootd door instance (no need to run separate doors for GSI and tokens).
 - Scaled up and revamped bulk request service.
 - A new TAPE API has been made available for the purposes of bulk staging
 - Addressed issue with group destination directory permissions when using tokens.
- Upgraded to 8.2 on 02/15/2023

Post upgrade issues - xrootd

- In anticipation of token roll-out the xrootd door has been setup with GSI and token support and TLS required (for token protection).
 - All tests have been made by us using most recent xrootd-client. package available from epel repository (5.5.1)
 - Further tests were done by IFDH folks who seemed to have used similar approach.
 - “Tested for success”. All worked well.
- BUT.... Experiments use older version of xrootd-client:
 - Versions older than 4 cannot work with TLS.
 - Versions older than 5.1 cannot work with tokens.
 - Locally built xrootd-client libraries of 5.1 or newer that a part of experiments’ software stacks (particularly root) have been built without token support.
- FIX:
 - Reconfigured xrootd door to not enforce TLS. This requires clients who want to use tokens to use (x)roots scheme. That is “s”, for secure. Which implies changing many scripts.
 - Since all experiments’ software does not work with xrootd using tokens this is considered a reasonable compromise (comparing going back to running different configurations on different ports on the same door hosts).

Post upgrade issues, xrootd, talking to each other

- We have had a couple of very productive meetings with scisoft team.
- Difficulties of rebuilding frozen releases have been highlighted.
- An understanding has been reached on scope and timeline for token supporting root releases:
 - Releases depending on xrootd-client libraries older than version 5 need to be upgraded.
 - Most recent releases using xrootd-client libraries of version 5.5.1 need to be rebuilt with token support libraries included.
 - Timeline seems to be defined as 1 year.

Xrootd client version roster

- Just at a random time on a random day

64339 xrd.rn=v4.8.0

27136 xrd.rn=v5.1.0

4115 xrd.rn=v4.5.0

1144 xrd.rn=v4.8.5

952 xrd.rn=v4.12.3

79 xrd.rn=v5.4.2

41 xrd.rn=v5.5.1

9 xrd.rn=v5.5.3

2 xrd.rn=v5.4.0

1 xrd.rn=v5.5.0

Architectural changes visible to user

- Not related to 8.2 upgrade but an important note:
 - For sometime now we have been running HA door endpoint – `fndcador.fnal.gov` which is HAProxy™ setup with 3 doors – `fndcador0[1-3].fnal.gov`
 - We advise users to switch from using `fndca1.fnal.gov` to `fndcador.fnal.gov` for better scalability, load balancing and fault tolerance.
 - Also allows rolling updates with little service disruption.
 - We have no immediate plans to shutdown `fndca1.fnal.gov` and envision to utilize network alias for backward compatibility.
 - It is my understanding that recent `ifdh_config` has switched for using `fndcador.fnal.gov` when setting up data transfers.

WLCG tape API: SRM

- Unlike IF experiment, the LHC experiments have been heavily using SRM to interact with Storage Elements (SE).
- In the nutshell SRM implementations provide:
 - Space management functions.
 - Directory functions.
 - Transfer functions.
 - Elements of data lifecycle (staging, pinning, expunging).
- And it could do above in bulk.
- SRM has been designed and developed in 2003-2006.
 - Implemented as SOAP exchange based Web Service.
 - Requires special client.
 - Non-trivial amount of effort went into ensuring interoperability and protocol clarification.
- Protocol specification is somewhat ambiguous and complex.
- Time to move on having learned the lesson.

WLCG tape API

- Work on providing “storage portals” have been ongoing in dCache for several years.
- What is storage portal? A web server providing REST API to perform storage management functions.
 - Standard off-the-shelf client – curl.
 - Standard authentication/authorization mechanisms.
- dCache API is described here <https://fndcador.fnal.gov:3880/api/v1/>
- In 7.2 added tape API and bulk request support:
 - <https://docs.google.com/document/d/14sdrRmJts5JYBFKSvedKCxT1tcrWtWchR-PJhxdunT8/edit#heading=h.lassqf5agt8x>
- We think that this API will eventually supersede SRM.
- We got CERN WLCG interested in this and within DOMA BDT working group the dCache API has been used as a basis for developing WLCG Tape API Specification:
 - https://docs.google.com/document/d/1Zx_H5dRkQRfju3xIYZ2WgjKoOvmLtsafP2pKGpHqcfY/edit#heading=h.ozszs1lr7q93
- Fully implemented in dCache 8.2

WLCG tape API: Discovery

```
curl -k https://fndcador.fnl.gov:3880/.well-known/wlcg-tape-rest-api
{
  "endpoints": [{
    "metadata": {},
    "uri": "https://fndcador.fnl.gov:3880/api/v1/tape",
    "version": "v1"
  }],
  "sitename": "fnal-wlcg-tape-endpoint",
  "description": "Fermilab public dCache WLCG tape REST API"
}
```

WLCG tape API: Stage Example

```
curl -L --capath /etc/grid-security/certificates
      --cert /tmp/x509up_u`id -u`
      --cacert /tmp/x509up_u`id -u`
      --key /tmp/x509up_u`id -u`
      -X POST https://fndcador.fnl.gov:3880/api/v1/tape/stage
      -H "Content-Type: application/json"
      -d '{"files":
          [{"diskLifetime":"P1D",
            "path":
"/pnfs/fnl.gov/usr/test/litvinse/atom/disk01/g2/dmitri/xist_ttt/xi_1.root"},
          {"diskLifetime":"P1D",
            "path":
"/pnfs/fnl.gov/usr/test/litvinse/atom/disk01/g2/dmitri/xist_ttt/xi_2.root"}]}'
```

Response

```
{
  "requestId" : "66f1814c-0059-448d-b5a6-730e0927a6d8"
}
```

WLCG tape API: Check status

```
curl -L --capath /etc/grid-security/certificates
      --cert /tmp/x509up_u`id -u`
      --cacert /tmp/x509up_u`id -u`
      --key /tmp/x509up_u`id -u`
      -X GET "https://fndcador.fnl.gov:3880/api/v1/tape/stage/66f1814c-0059-448d-b5a6-730e0927a6d8"
{
  "id" : "66f1814c-0059-448d-b5a6-730e0927a6d8",
  "createdAt" : 1679024479188,
  "startedAt" : 1679024479216,
  "files" : [ {
    "path" : "/pnfs/fnl.gov/usr/test/litvinse/atom/disk01/g2/dmitri/xist_ttt/xi_1.root",
    "startedAt" : 1679024479278,
    "state" : "STARTED"
  }, {
    "path" : "/pnfs/fnl.gov/usr/test/litvinse/atom/disk01/g2/dmitri/xist_ttt/xi_2.root",
    "startedAt" : 1679024479301,
    "state" : "STARTED"
  } ]
}
```

WLCG tape API: Check Request Status

```
curl -L --capath /etc/grid-security/certificates
      --cert /tmp/x509up_u`id -u`
      --cacert /tmp/x509up_u`id -u`
      --key /tmp/x509up_u`id -u`
      -X GET "https://fndcdoor.fnal.gov:3880/api/v1/tape/stage/66f1814c-0059-448d-b5a6-730e0927a6d8"
{
  "id" : "66f1814c-0059-448d-b5a6-730e0927a6d8",
  "createdAt" : 1679024479188,
  "startedAt" : 1679024479216,
  "completedAt" : 1679024840827,
  "files" : [ {
    "path" : "/pnfs/fnal.gov/usr/test/litvinse/atom/disk01/g2/dmitri/xist_ttt/xi_1.root",
    "finishedAt" : 1679024840822,
    "startedAt" : 1679024479278,
    "state" : "COMPLETED"
  }, {
    "path" : "/pnfs/fnal.gov/usr/test/litvinse/atom/disk01/g2/dmitri/xist_ttt/xi_2.root",
    "finishedAt" : 1679024823790,
    "startedAt" : 1679024479301,
    "state" : "COMPLETED"
  } ]
}
```

WLCG tape API: Check Archive Status

```
curl -s -L --capath /etc/grid-security/certificates
      --cert /tmp/x509up_u`id -u`
      --cacert /tmp/x509up_u`id -u`
      --key /tmp/x509up_u`id -u`
      -X POST https://fndcador.fnl.gov:3880/api/v1/tape/archiveinfo
      -H "Content-Type: application/json"
      -d '{"paths":
["/pnfs/fnl.gov/usr/test/litvinse/atom/disk01/g2/dmitri/xist_ttt/xi_1.root",
"/pnfs/fnl.gov/usr/test/litvinse/atom/disk01/g2/dmitri/xist_ttt/xi_2.root"]}'
| jq .
```

Response

```
[
  {
    "path": "/pnfs/fnl.gov/usr/test/litvinse/atom/disk01/g2/dmitri/xist_ttt/xi_2.root",
    "locality": "DISK_AND_TAPE"
  },
  {
    "path": "/pnfs/fnl.gov/usr/test/litvinse/atom/disk01/g2/dmitri/xist_ttt/xi_1.root",
    "locality": "DISK_AND_TAPE"
  }
]
```

dCache 8.2 Concluding remarks

- Rolled out WLCG Tape API – try it.
- Have been running stable for a month, looks like a solid release