




Michal Simon

XRootD client status update: latest developments and plans



Outline

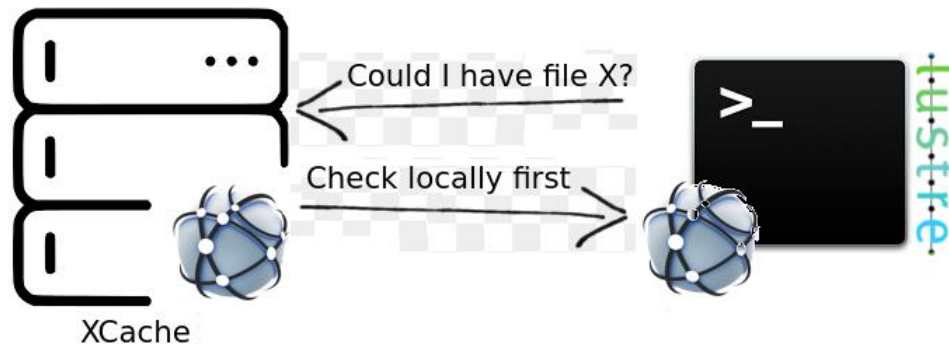
- 4.8.0 & 4.9.x (recent developments)
- 4.10.0 (about now)
- 5.0.0 (on the horizon)
- 5.x.x (plans)



Redirects / Retries: local file redirects

Release 4.8.0

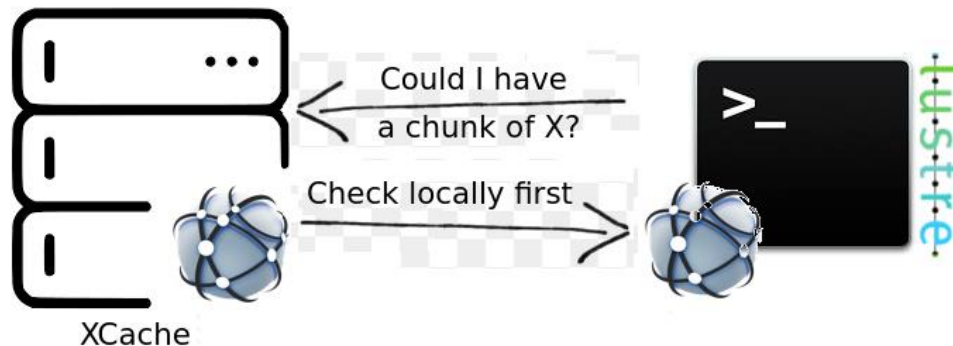
- **Redirection to local file** (and local file support)
 - Meant for HPC clusters with local Lustre mounts



Redirects / Retries: state redirect

Release 4.9.0

- **Redirection on stateful operation** (e.g. read)
 - Also, meant for HPC clusters with local Lustre mounts



Redirects / Retries: redirect trace-back

Release 4.9.0

- **Redirect trace-back** - printed to logs with level:
 - **Warning** for `errNotFound` and `errRedirectLimit`
 - **Debug** otherwise

```
[Debug][XRootD] Redirect trace-back:
```

```
[Debug][XRootD] 0. Redirected from: root://metaman:1094/ to: root://man1:1094/
```

```
[Debug][XRootD] 1. Redirected from: root://man1:1094/ to: root://srv1:1094/
```

```
[Debug][XRootD] 2. Retrying: root://metaman:1094/
```

```
[Debug][XRootD] 3. Redirected from: root://metaman:1094/ to: root://man2:1094/
```

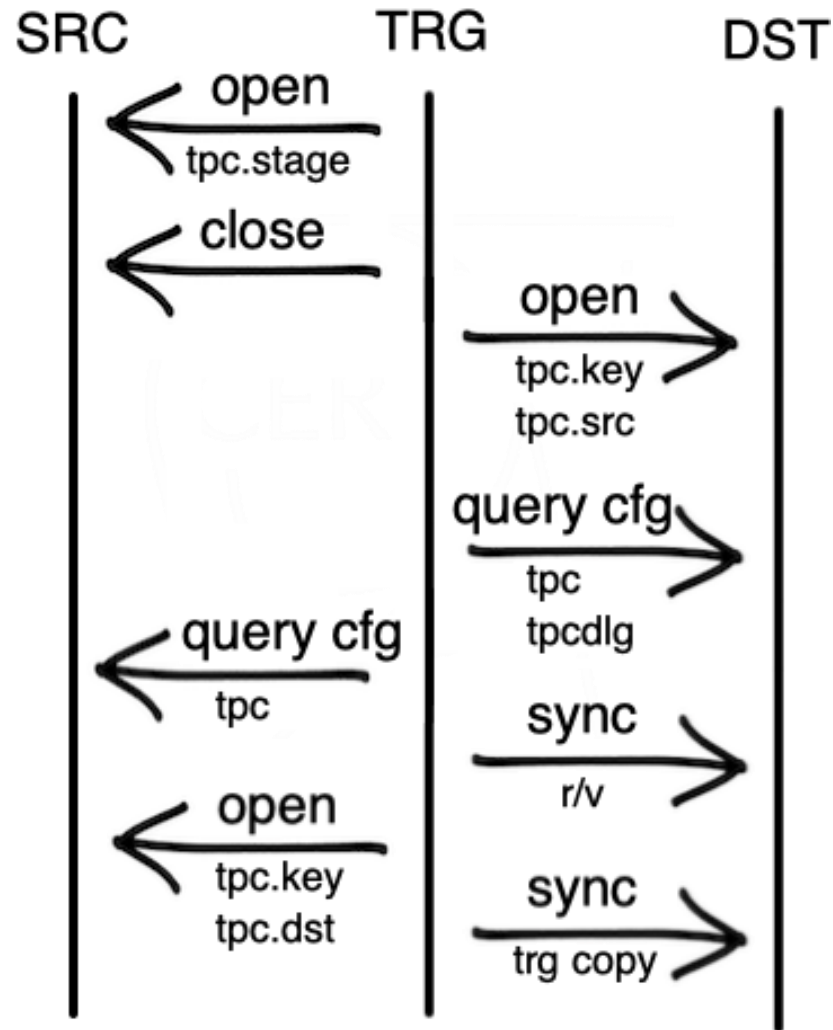
```
[Debug][XRootD] 4. Redirected from: root://man2:1094/ to: root://srv3:1094/
```

Third party copy


Release 4.9.0

- Motivation: enable xroot TPC in Atlas workflows and address destination vs source authentication problem
- Forward number of streams
- Enable **proxy certificate delegation**

Third party copy: vanilla TPC




Third party copy: TPC lite



New features: vector writes

Release 4.9.0

- Write scattered data chunks in one operation (within a single file)



- The protocol also supports vector writing to multiple files (in case this is would be needed in the future)

New features: protocol-level plug-in

Release 4.9.0

- For SFT xrdcl-http plug-in
- Client plug-in configuration:

```
url=http://*  
lib=/usr/lib64/libAwesomeHttpPlugin.so  
enabled=true
```

New features: force disconnect

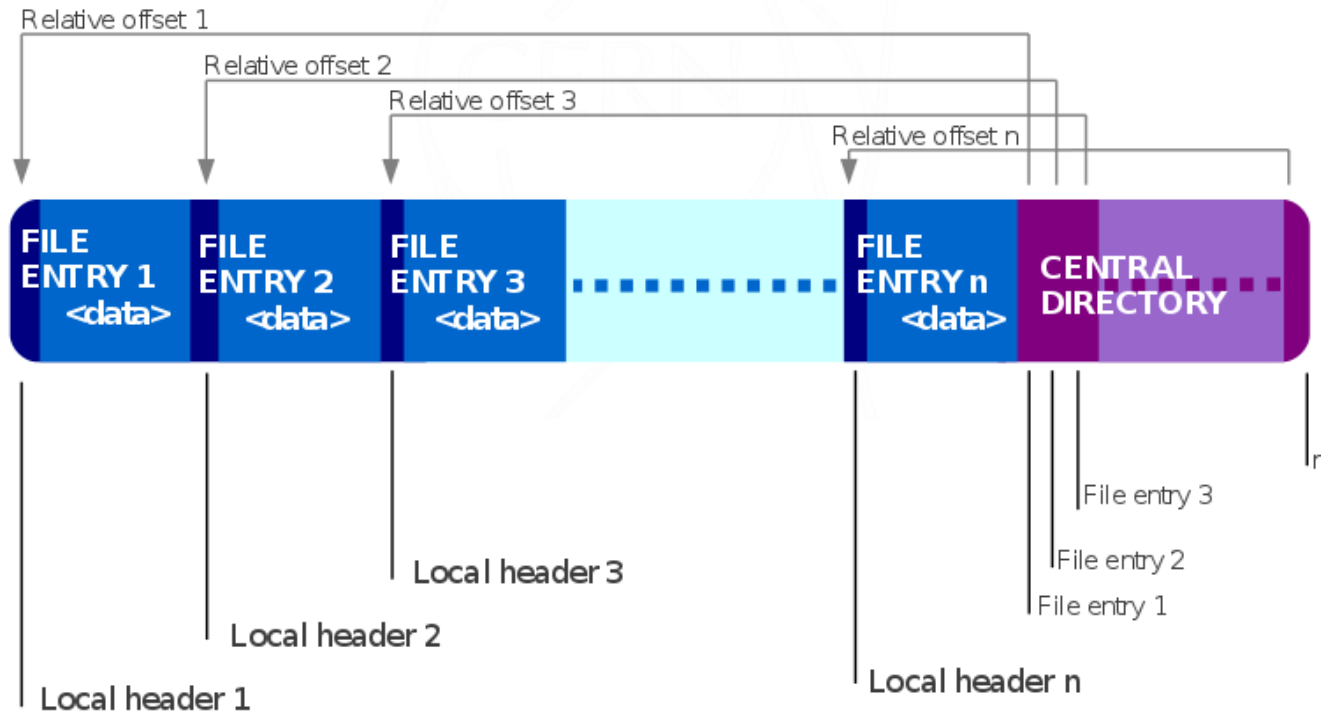
Release 4.9.0

- New API:
`Status PostMaster::ForceDisconnect(const URL &url);`
- Previously only way to influence disconnection were TTL timeouts
- Use case: **force authentication with new credential, force proxy delegation**

ZIP support: extraction

Release 4.9.1 (last bugs fixed)


- Extract data from root files (**no decompression**)



ZIP support: extraction & listing

Release 4.9.1 (last bugs fixed)

- `xrdcp -z file1 root://host//path/file.zip .`
- `xrdcp root://host//path/file.zip?xrdcl.unzip=file1 .`
- `xrdfs ls /path/to/file.zip`
- `xrdfs ls -Z /path/to/zip.file`



Redirects / Retries: retry policy changes

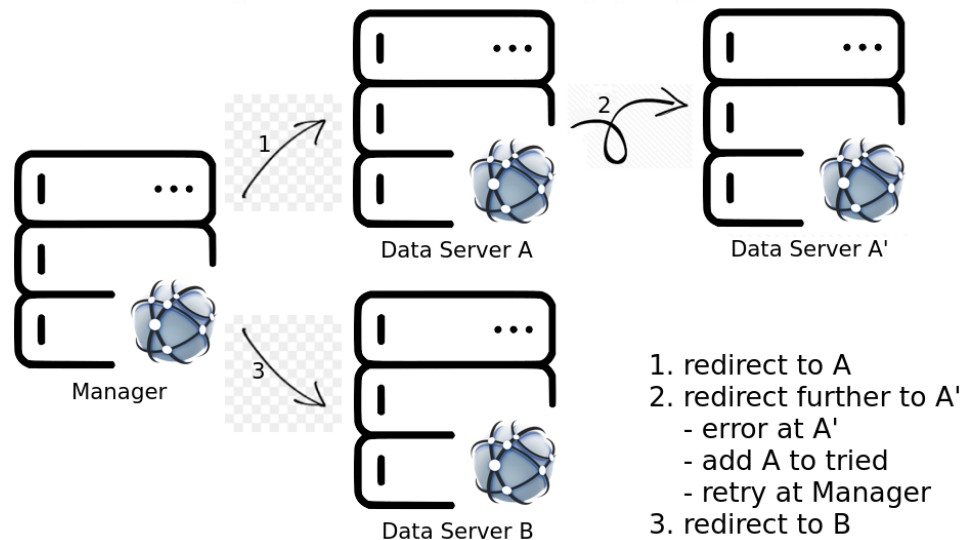
Release 4.10.0

- All virtual redirectors (**metalinks**) are treated as **meta-managers**
- For meta-managers following errors are treated as **retriable**: `kXR_NotAuthorized`, `kXR_FileLocked`, `kXR_Unsupported`
- For `kXR_NotAuthorized` recovery wont be attempted by default more than 3 times (configurable through `XRD_NOTAUTHORIZEDRETRYLIMIT`)

Redirects / Retries: retry policy changes

Release 4.10.0

- In case a data server responds with a `kXR_redirect` and we fail at the node where we were redirected to, the original data server is included in the `tried` and `triedrc` cgi.



New features: bring online interface

Release 4.10.0

- Use case: CTA and FTS bring-online utility
- `xrdfs prepare -s [paths]`
 - returns a staging request ID
- `xrdfs prepare -a requestid [paths]`
 - cancel whole staging request or just some “paths” within a request
- `xrdfs query prepare requestid`
 - query status of a staging request
- Also, available through `FileSystem API`


New features: chunked response

Release 4.10.0

- XRootDMsgHandler **supports chunked response** handling
- Currently available only for `kXR_dirlist`, enabled through `DirListFlags::Chunked`
- Could be enabled for others (except `kXR_read` which requires additional work)
- Use case: e.g. streaming big directory listing

New features: chunked response

Release 4.10.0



New features: chunked response

Release 4.10.0

- Response status:
 - **suContinue**: more data will follow (**mustn't destroy handler**)
 - **suDone**: this is the last chunk
 - **suPartial**: all chunks were received, but not all were successful

Python bindings

4.9.x & 4.10.0

- Use cases: experiment frameworks, but also things like Invenio
- **Fix deadlock on Python interpreter exit**
- New functionalities exposed in Python: parallel copy jobs, extreme copy, chunked directory listing, etc.
- **PyPI release**
- Globbing

On the horizon: 5.0.0

5.0.0 major release

- **Extended attributes** (completely implemented)
- **Encryption** (core mostly implemented)
- Review **ABI** (pimpl idiom) and public headers
- Remove old client



ComputerHope.com

On the horizon: 5.0.0

5.0.0 major release

- Channel level plug-ins
 - A plug-in that resides on the **same level as a virtual redirector** (Metalink)
 - Would allow for **redirecting requests to protocol plug-ins** (e.g. HTTP)

Plans: post R5

5.x.x series

Bundled requests:

- mapping file ID on subsequent requests
- allowing for close in presence of outstanding request



Plans: post R5

5.x.x series

Delete on close

- Protocol extension needed: new request type for verified close
- Dedicated server side support
- Adjustments on client side + compatibility layer
- Re-factoring of xrdcp logic so it can take advantage of delete on close

Plans: post R5

5.x.x series

Extended stat:

- Probably there's no need to extend protocol
- Server will just send additional data and the client needs to make sure all of it gets parsed
- Can be combined with bulk listing of extended attributes

Plans: post R5

5.x.x series

Partial response handling for reads:

- Use cases: **streaming copy**
- A protocol enhancement is needed: data offset is missing in read response (probably new type of request to ensure protocol level compatibility)
- New client side logic that would handle callbacks with `suContinue` status

Plans: post R5

5.x.x series

New TPC:

- One **request to handle whole TPC** operation
- Protocol extensions: **kXR_getfile / kXR_putfile**
- New authentication model for TPC

Plans: post R5

5.x.x series

Client:

- Extending testing infrastructure: **mock event-loop**
- Extending ZIP support (**append** – difficult: requires server side support, **decompression** – easy: funzip)
- Modernize code base (C++11, pointer ownership)

Packaging:

- Docker, cmake/pkg-config module

Question?

XrdCI docs: <http://xrootd.org/doc/xrdcl-docs/xrdcldocs.pdf>

