

# Disk-Caching-On-The-Fly vs. XCache

Volker Lindenstruth, Kilian Schwarz, Paul-Niklas  
Kramp, Dirk Hutter, **Serhat Atay**

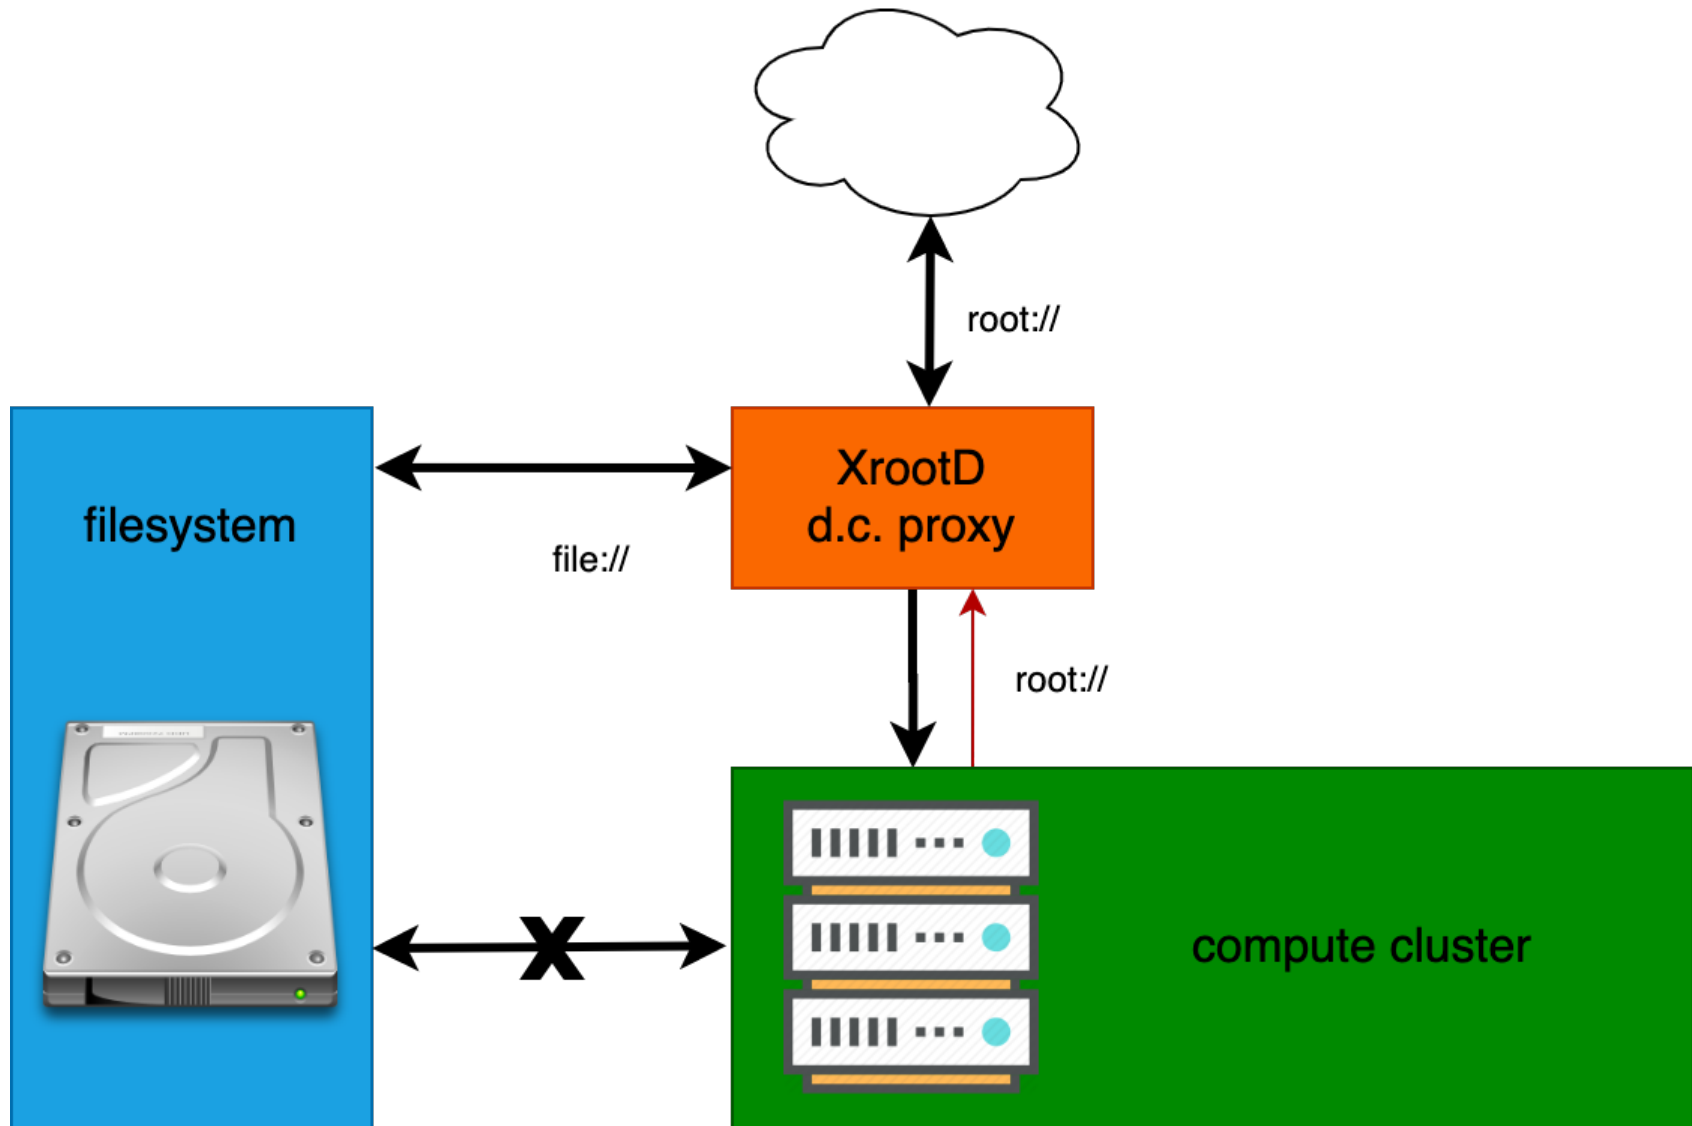
Goethe University, Frankfurt

GSI, Darmstadt

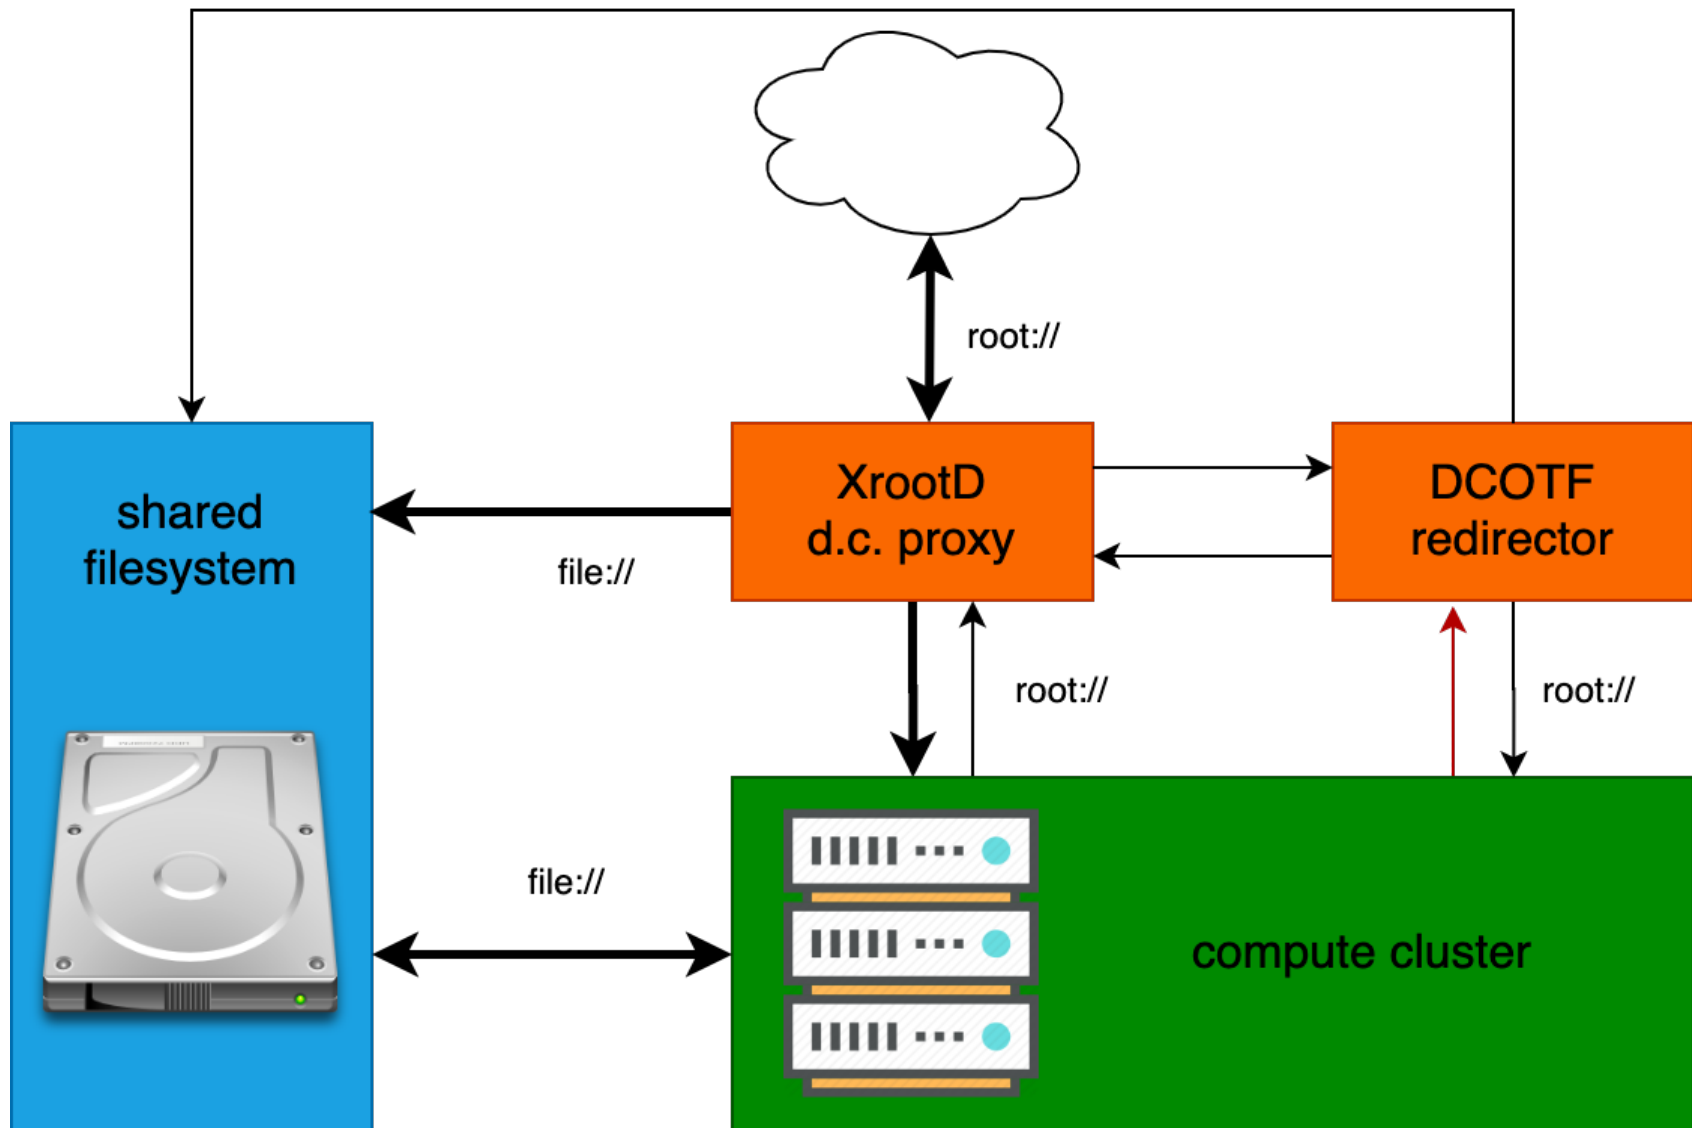
24 July 2020



# XCache (Standalone)



# Disk-caching-on-the-fly (DCOTF)



# Client Plugin by GSI

- Employed by the client
- Addition of the redirector address
  - `export XRD_PLUGIN="/path/to/client_plugin.so"`
  - `export XRD_DEFAULT_PLUGIN_CONF="/path/to/plugin_config.conf"`
    - Contents of `plugin_config.conf`
      - `proxyPrefix=://redirector_ip:redirector_port///x`
      - `enable = true`
- Command by the client
  - `xrdcp root://extDataserver_ip:extDataserver_port//path/to/file ./tfile`
  - ExtDataserver as data source (e.g. GSI via WAN access)
- After the client plugin
  - `xrdcp root://redirector_ip:redirector_port///xroot://extDataserver_ip:extDataserver_port//path/to/file ./tfile`

# Local Redirect Plugin by GSI

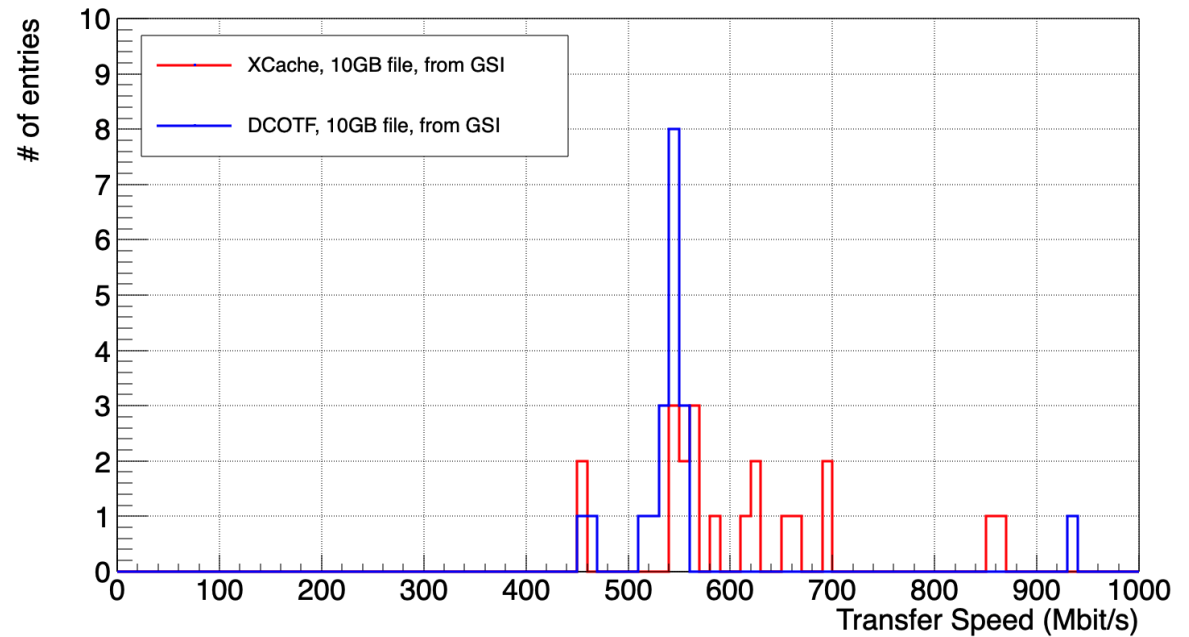
- Employed by the redirector
  - ofs.cmslib `/path/to/local_redirect_plugin.so`
  - Redir.prefix `caching_proxy_ip:caching_proxy_port`
- Access to the cache location
  - oss.localroot `/path/to/cache`
- Command from the client plugin
  - root://redirector\_ip:redirector\_port///xroot://extDatasever\_ip:extDatasever\_port//`path/to/file` ./file
  - ExtDatasever as data source (e.g. GSI via WAN access)
- Checking existence of file
  - Cut “`/path/to/file`” from the address
  - Add “`/path/to/file`” to “`/path/to/cache`”
  - Check by “access” method if “`/path/to/cache/path/to/file`” exists
  - If found, SFS\_REDIRECT is returned with port -1 and `/path/to/cache/path/to/file`
  - If not found, SFS\_REDIRECT is returned with `caching_proxy_port` and `caching_proxy_ip`
- LocalFileHandler (SFS\_REDIRECT) (XrootD version >4.8.0)
  - with port -1 and “path” → switches to file:// protocol for POSIX operation as “`file://path/to/cache/path/to/file`”
  - with valid port address and ip → keeps root:// protocol adding the port and the ip address as “`root://caching_proxy_ip:caching_proxy_port///xroot://extDatasever_ip:extDatasever_port//path/to/file`”

# First Tests

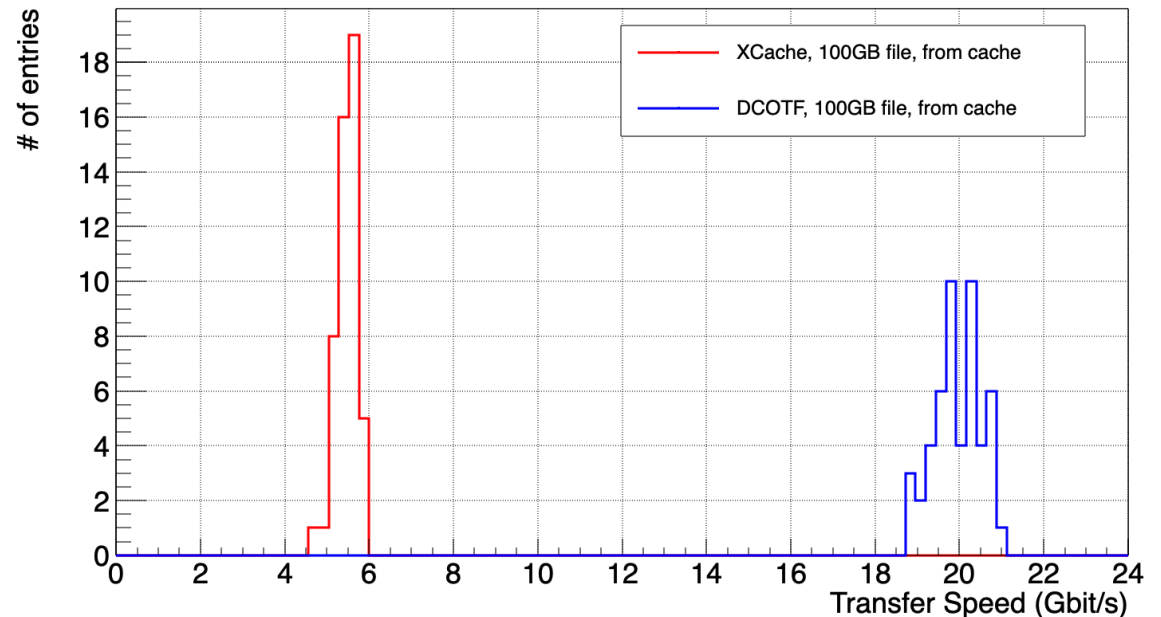
- Establishment of connection from Goethe-HLR to GSI
  - 1Gbit/s to GSI via WAN (shared)
  - 120Gbit/s via private link (to be reestablished)
- QDR links between the client node and the cluster filesystem at Goethe-HLR
  - Theoretical max bandwidth 32Gbit/s (shared)
- Sample data to test connections and the caching functionality
  - 100GB (from cache) and 10GB (from GSI)
- Transfer rates of data transfers
  - Data transfer from GSI with disk caching (10GB file, x20 serially)
  - Data transfer from the cache (100GB file, x50 serially)

# Transfer Speeds

- From GSI with caching (upper plot)
  - Max bandwidth: 1Gbit/s
  - Shared resource



- From cache (lower plot)
  - Max bandwidth: 32Gbit/s
  - Shared resource
  - Difference in transfer rates is under investigation (limited as a non-root user)



# Future plans

- Setting up a new server to be added to the Goethe-HLR in Frankfurt with root privileges
  - 120Gbit/s private link to GSI (to be reestablished)
  - 100Gbit/s link to the cluster filesystem (unshared by other users)
  - Few TBs of NVMe SSD as a local filesystem in case the cluster filesystem is slow. (The cluster filesystem is always shared)
- Improvement and debugging of XrootD configuration files with the feedback from performance tests
- Development on the local redirect plugin to make use of XCache features, i.e.:
  - Caching blocks of files instead of complete files -not implemented in the local redirect plugin
  - Status of data in cache (parsing .cinfo files) -not implemented in the local redirect plugin
- Submission of local redirect plugin to XrootD core package after debugging and aforementioned improvements

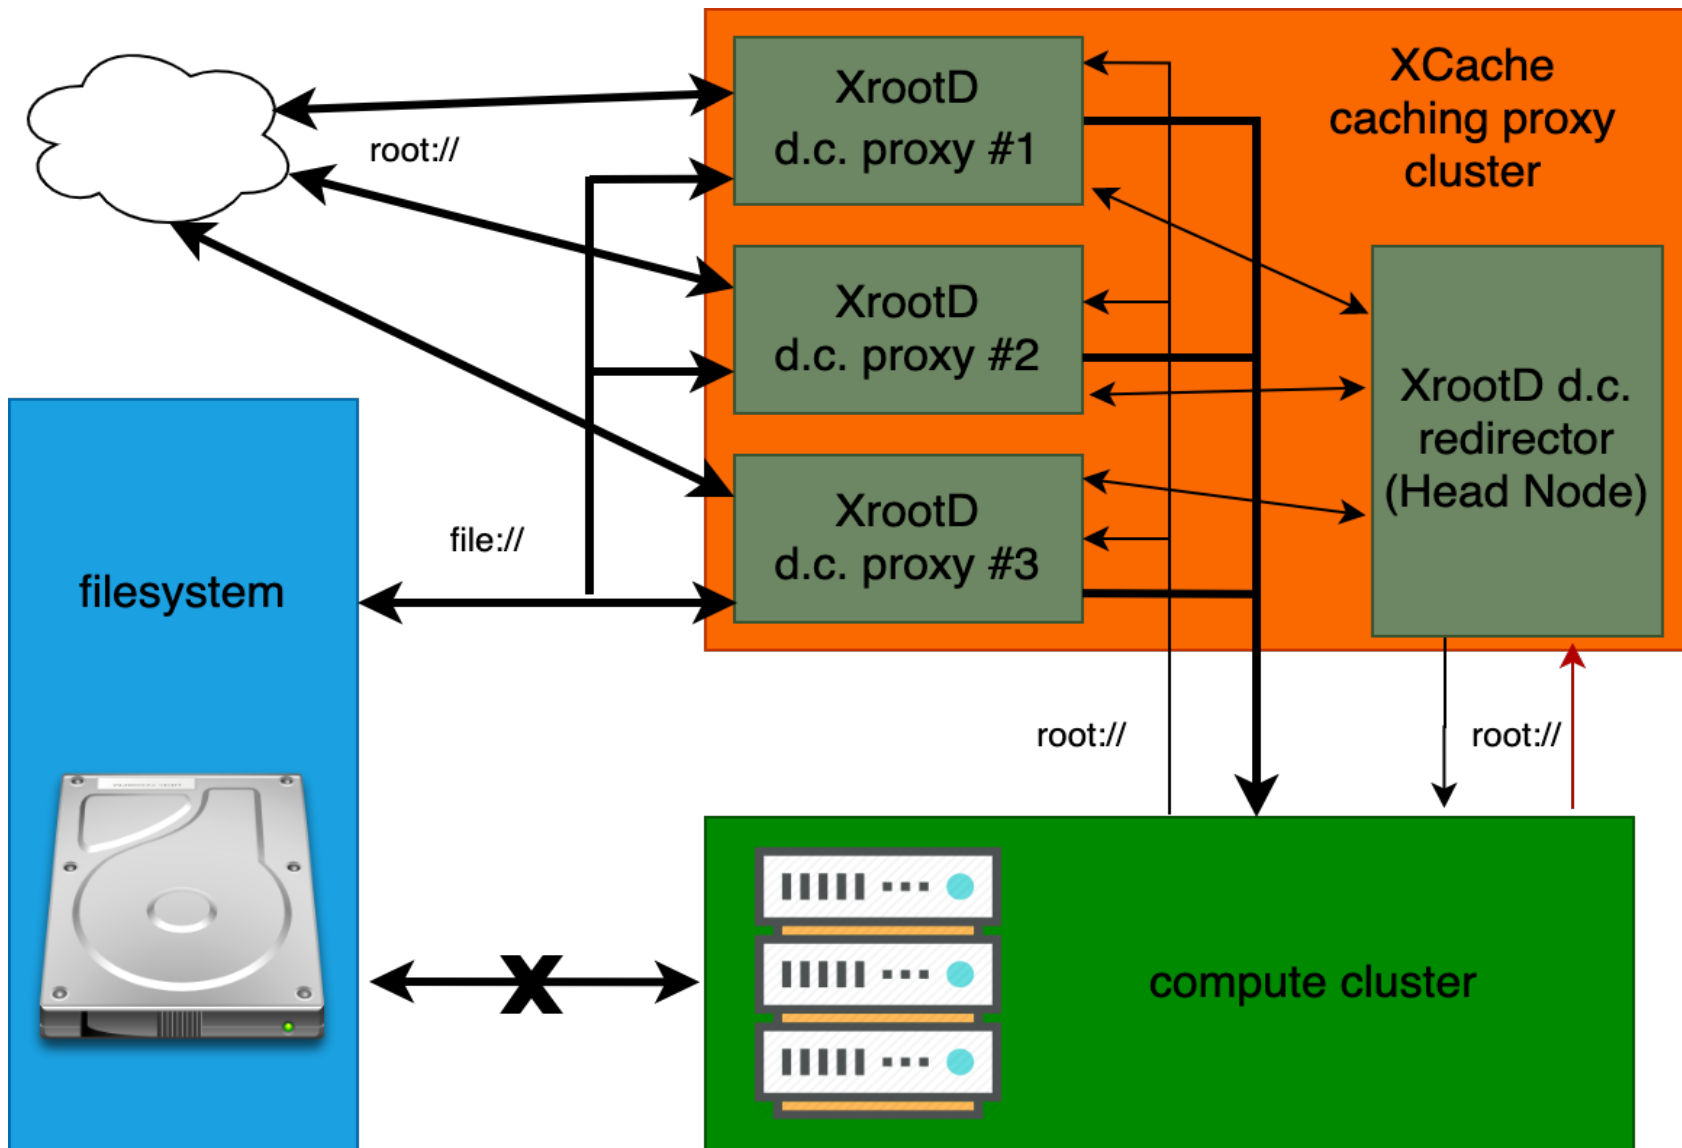


# Summary

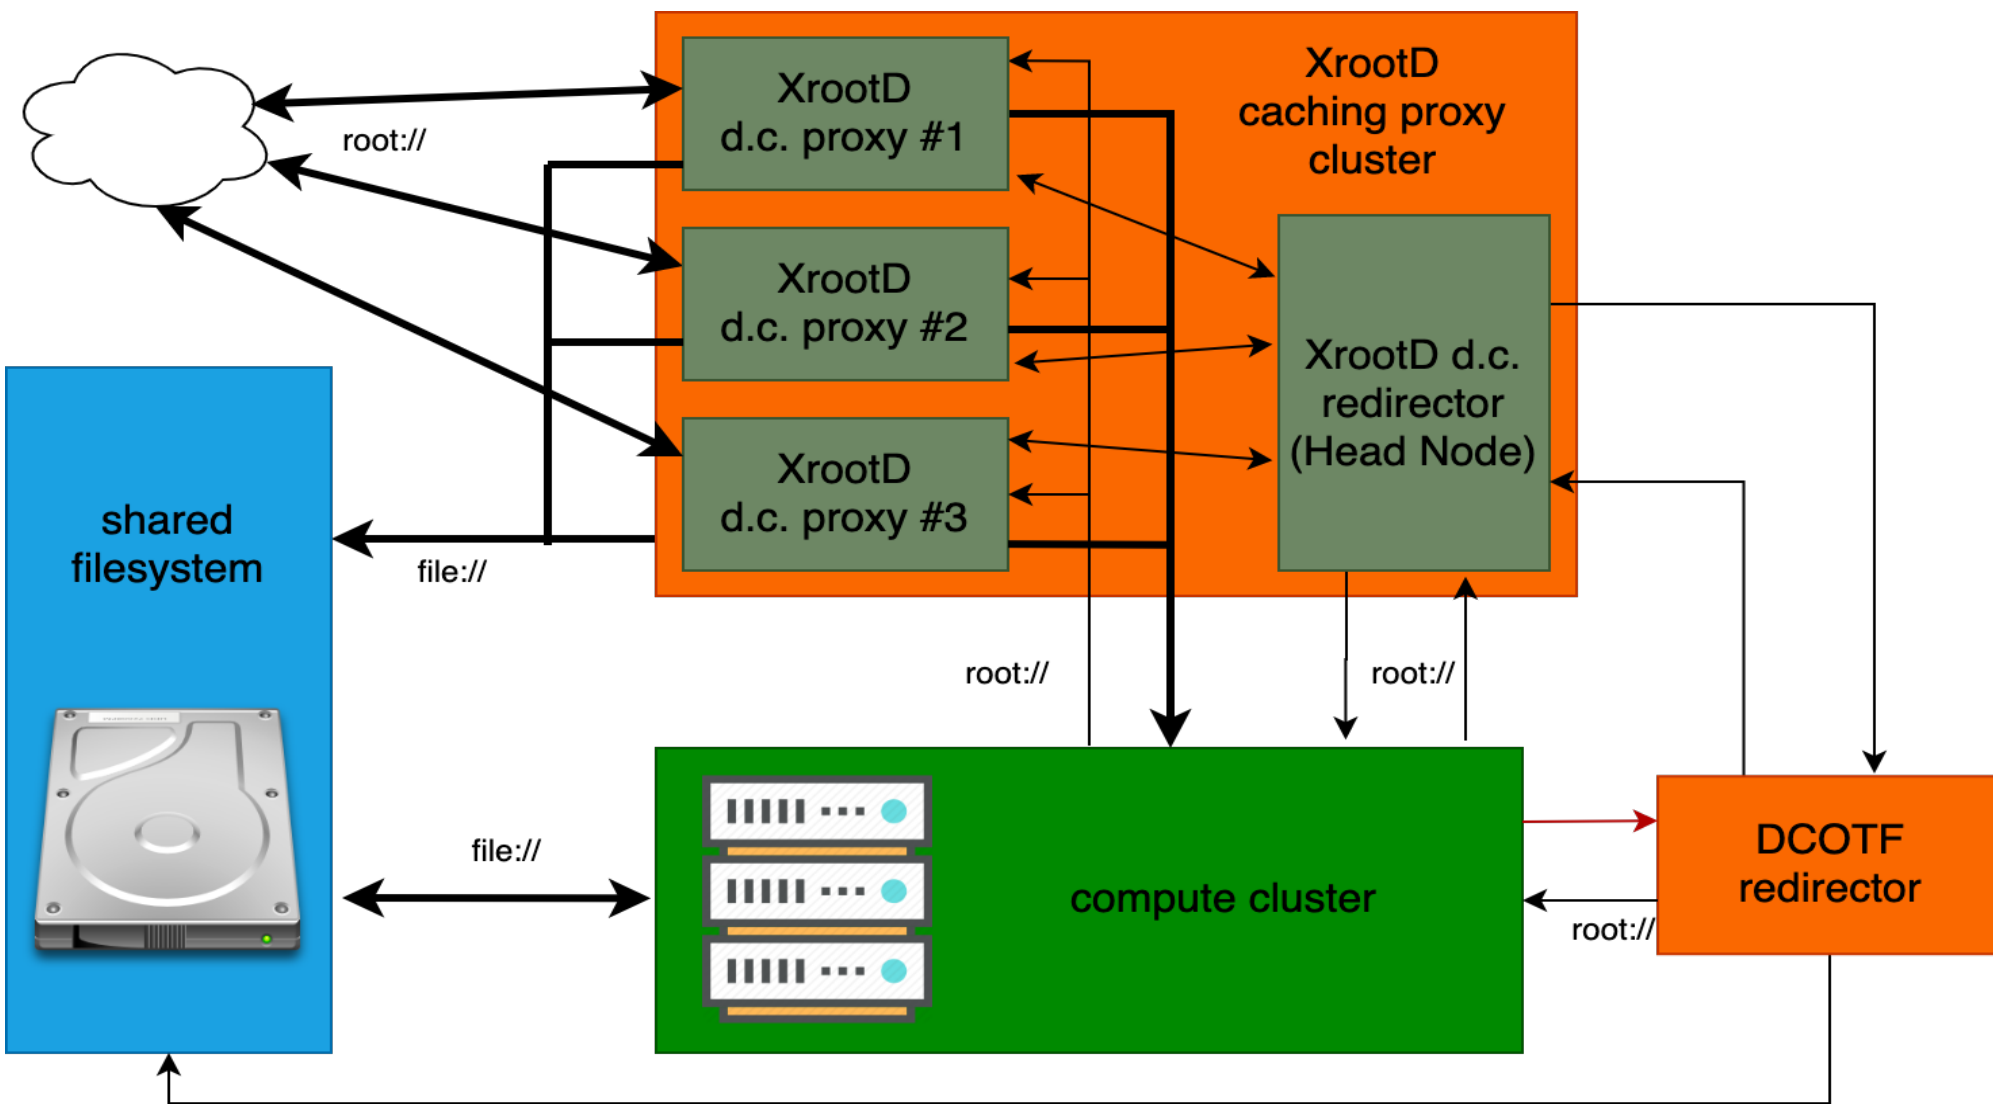
- Added local redirection feature to XCache for the cached data by DCOTF
  - Reduction of caching proxy load and network traffic
  - Possibility to speed up transfer rates for the successive read operations
- Client plugin to contact the DCOTF redirector first
- Local redirect plugin to manage POSIX operation within shared filesystem
- First tests of the DCOTF and XCache at Goethe-HLR
- New server setup with private links to GSI and the cluster filesystem at Goethe-HLR
- Development of the local redirect plugin to enable more XCache features

Thanks

# XCache (Clustered)



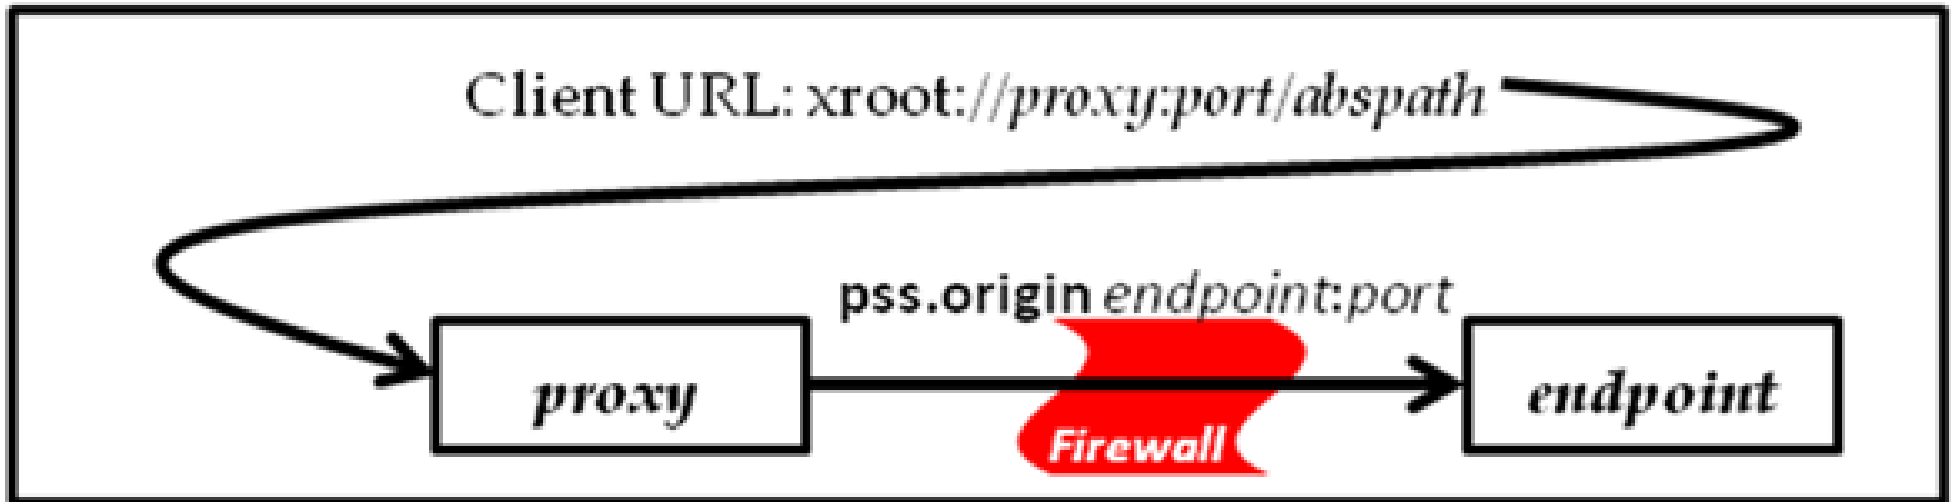
# Disk-caching-on-the-fly (clustered)



# XrootD d.c. proxy

- Employment of two XrootD libraries
  - **LibXrdPss** -proxy server library
  - **LibXrdFileCache** -caching library (named libXrdPfc from XrootD v5.0.x (2020))
- Access to the cache location
  - `oss.localroot path/to/cache`

# direct mode proxy



# forwarding proxy

