

Benchmarking of XRootD Disk-Caching-On-The-Fly Status Update

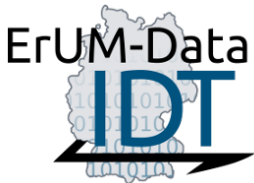
Michael Böhler, Anton J. Gamel, Stefan Kroboth,
Dirk Sammel, Markus Schumacher

26.11.2020
IDT-UM Storage Meeting

Albert-Ludwigs-Universität Freiburg



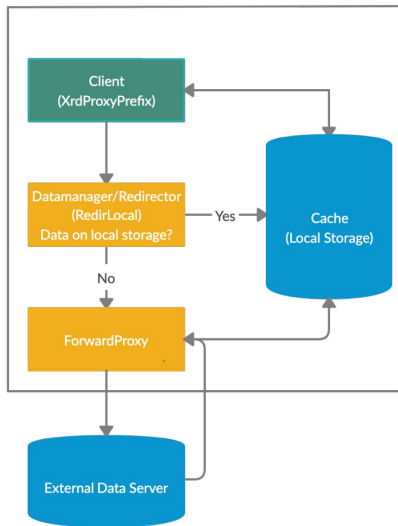
UNI
FREIBURG



Goals:

- ▶ Benchmarks
 - ▶ XrootD disk-caching-on-the-fly (DCOTF) vs direct access
 - ▶ Different underlying file systems for cache location
- ▶ Functional tests for required components of DCOTF

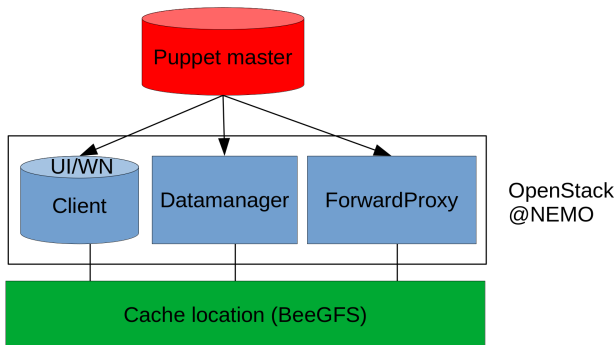
Architecture of DCOTF



by Serhat Atay

- ▶ Implementation with OpenStack and Puppet
- ▶ 3 VMs running on NEMO cluster: Client, Datamanager, ForwardProxy
- ▶ Client: typical worker node/ user interface + XrdProxyPrefix plugin
"typical" → same as machines in production
- ▶ Cache location: currently BeeGFS
 - ▶ Mounted on Client, Datamanager and ForwardProxy

Configuration management tool



- ▶ Python script: pyROOT
- ▶ "Pseudo analysis", close to typical user analysis:
open file, loop over events, read some branches of DAOD, loop over all electrons in event
- ▶ Output: information in JSON format → visualize with pandas
- ▶ Things to be tested:
 - ▶ Caching vs. direct access
→ access to external sites important (VOMS issue, see next slide)
 - ▶ File size, number of events
 - ▶ Cache location: BeeGFS vs. SSDs

- ▶ Caching of files from external sites:
 - ▶ Client credentials (VOMS) are not forwarded to external site
→ authentication fails
 - ▶ Suggestion from KIT in last meeting
 - ▶ Postponed: currently only using Freiburg dCache
- ▶ Reading of cached files by a given user:
 - ▶ Files are written by user 'xrootd'
 - ▶ Default permissions are too strict (600)
 - ▶ Known issue: <https://github.com/xrootd/xrootd/issues/649>
 - ▶ Possible solution: [Multi-user plugin](#)
→ "works", but plugin can't retrieve username (empty), current workaround: using a predefined service user
- ▶ DCOTF with ROOT: files are cached but can't be read:
TNetXNGFile::ReadBuffers ERROR [ERROR] Operation is not implemented → in contact with Serhat Atay, who can reproduce this issue

- ▶ Open issues: VOMS, DCOTF with ROOT, (Multi-user plugin)
→ required for benchmarks
- ▶ Benchmarks
 - ▶ Compare caching setup with direct access to external sites
 - ▶ Compare different options for local storage: BeeGFS, SSDs, ...
 - ▶ Test complete example user physics analysis (different samples, ...)
- ▶ Implement component tests
 - ▶ Check functionality of every component: is service running, is local storage available, ...