XRootD
2018/19 Releases & Plans

US ATLAS Computing Facility Meeting

ANL

December 3 - 5, 2018

Andrew Hanushevsky, SLAC

http://xrootd.org

Introduction

- **XRootD** has an aggressive release schedule
  - Perhaps too aggressive for available FTE's
    - Release delays occur
- Latest patch/enhancement release is 4.8.5
- Upcoming feature release is 4.9.0
  - Currently available as a release candidate
- Next planned major release is 5.0.0
  - Target is 1Q19 barring 4.9.0 patch releases

XRootD 4.8.5 Part 1

- Issues addressed
  - XrdCrypto certificate parsing
  - xrdcp recursive copy failures
  - errno assignment for proper error recovery
  - HTTP performance additions
  - HTTP error status code corrections
  - HTTP header parsing
  - HTTP error handling

XRootD 4.8.5 Part 2

- Issues addressed
  - Avoid double I/O for local checksum calc
  - Expose kXR_cancel for prepare requests
  - Zip file extraction via CGI (enhancement)
  - SSS keytab selection via CGI (enhancement)
  - Pass proper endorsement for sss security protocol (enhancement)
  - Various client edge cases to avoid a SEGV

XRootD 4.9.0 Part 1

- Features
  - Vector write (kXR_writev)
  - Support xrdcp stream specification via TPC
  - Support Xcache ingest via HTTP
  - xrdcp cross protocol copies (xroot <-> HTTP)
  - Allow client to force server disconnect
  - Enhancements to ease containerization
  - Pipelining dataflow API in the client

XRootD 4.9.0 Part 2

- Features
  - Redirect trace back
  - HTTP Macaroon support
  - Deferred close requests
  - Fix Xcache SEGV when reading a file from multiple sources.
  - Subject Alternative Name (SAN) support
    - X509 and RFC 2818 compliance

XRootD 4.9.0 Part 2

- Features
  - Full delegated proxy support
    - Congruent with RFC 2818
      - A rather convoluted interaction with DNS usage
        - Largely on how sites register DNS names
    - Proper security requires reissuance of certs
    - Currently a very contentious feature
  - TPC support for delegated proxies

XRootD 5.0.0 Part 1

- Features
  - User settable file extended attributes
    - Restricted to user namespace
  - Full TLS support (i.e. xroots)
  - Extended stat information
    - Support for uid/gid tracking
  - Trivialize ofs plug-in wrapping
    - Avoids disabling XRootD performance features

XRootD 5.0.0 Part 2

- Features
  - Allow checksum check on close()
    - New request code: kXR_closev
  - State full redirects (i.e. on read)
    - Redirect to local data source
      - Implements HPC RDMA data access
  - Remove old client
    - At least no longer compiled
      - This impacts ALICE

XRootD 5.0.0 Part 3

- Features
  - Integrated monitoring interface for **Xcache**
  - Prepare for TPC version 3
    - Add getFile() & putFile() ofs APIs
      - Set stage for safe SciToken support
      - Compliance with **XRootD** security architecture

Beyond XRootD 5 Chalkboard I

- Additions being contemplated
  - Recursive delete (mostly for HTTP)
  - Response data streaming (improved copy)
  - Getfile and Putfile functionality for new TPC
  - Apply/Map operation for data pipelining
    - Sometimes known as request bundling
  - Erasure encoding plug-in
  - Native data striping across partitions

Beyond XRootD 5 Chalkboard II

- Additions being contemplated
  - RDMA support
  - Erasure encoding plug-in
  - The uid/gid tracking for files/directories
  - Allow appends to a zip archive
  - Dynamic data source selection in the client
  - Enable mock testing of client
  - Docker based distribution
  - Implement package config functionality

Functionality Drivers

- ATLAS (majority) and CMS requests
  - ALICE seems happy with current state
- External service providers
  - CTA, dCache, DPM, and EOS
- OSG and WLCG
- HPC usability (it's increasing)
- Exploratory endeavors

Conclusion

- **XRootD** has a plethora of requests
  - Constant struggle to prioritize these
    - **XRootD** is now embedded in practically every HEP data delivery system
      - EOS, DPM, CTA, dCache (Java version), native …
    - New experiments rely on the **XRootD** as well
      - E.g. Dune and LSST
  - Other data access modes ratchet it up as well
    - **Xcache**, StashCache, CERNVMFS, token based authorization, etc. (**Xcache** is based on **XRootD**)