



WLCG Monitoring Task Force

Borja Garrido

14.09.2022

Introduction




- **WLCG Monitoring TaskForce was presented on December 2021**
 - During WLCG Operations coordination [meeting](#)
 - Real activities started January 2022: meetings, JIRA project...
- **Core team of ~6 people in "best effort"**
 - Alessandra Forti, Borja Garrido, Derek Weitzel, Julia Andreeva, ~~Rizart Dona~~, Shawn McKee
 - Meeting every 2 weeks for checkpointing and planification
 - Special thanks to Katy Ellis and Robert Currie for their contributions in XRootD Improvements
- **Focused on three main areas**
 - WLCG transfers harmonization
 - XRootD monitoring improvements
 - Site network monitoring integration

WLCG transfers data harmonization

Main goals

- **Consolidate schema between FTS/XRootD transfer documents**
 - Agree on a minimum required schema for both flows
- **Adapt WLCG transfers dashboards to new common schema**
 - Provide a set of useful dashboards under the WLCG umbrella
 - Avoid specificities for experiments

General Plots



TF short-term to-do list

- **Discuss minimum required schema with different data producers**
 - FTS schema already provides all initially needed fields for ATLAS and CMS
 - Meetings hold with XRootD and dCache developers, ALICE and FTS
- **Make minimum schema public and get feedback**
 - Draft [document](#) circulated and feedback gathering in progress (**Deadline 30th of September!**)
- **Harmonise the DC dashboard cross-experiment functionality with the agreed common schema**
- **Gather feedback from experiments and dashboard users for improvements**



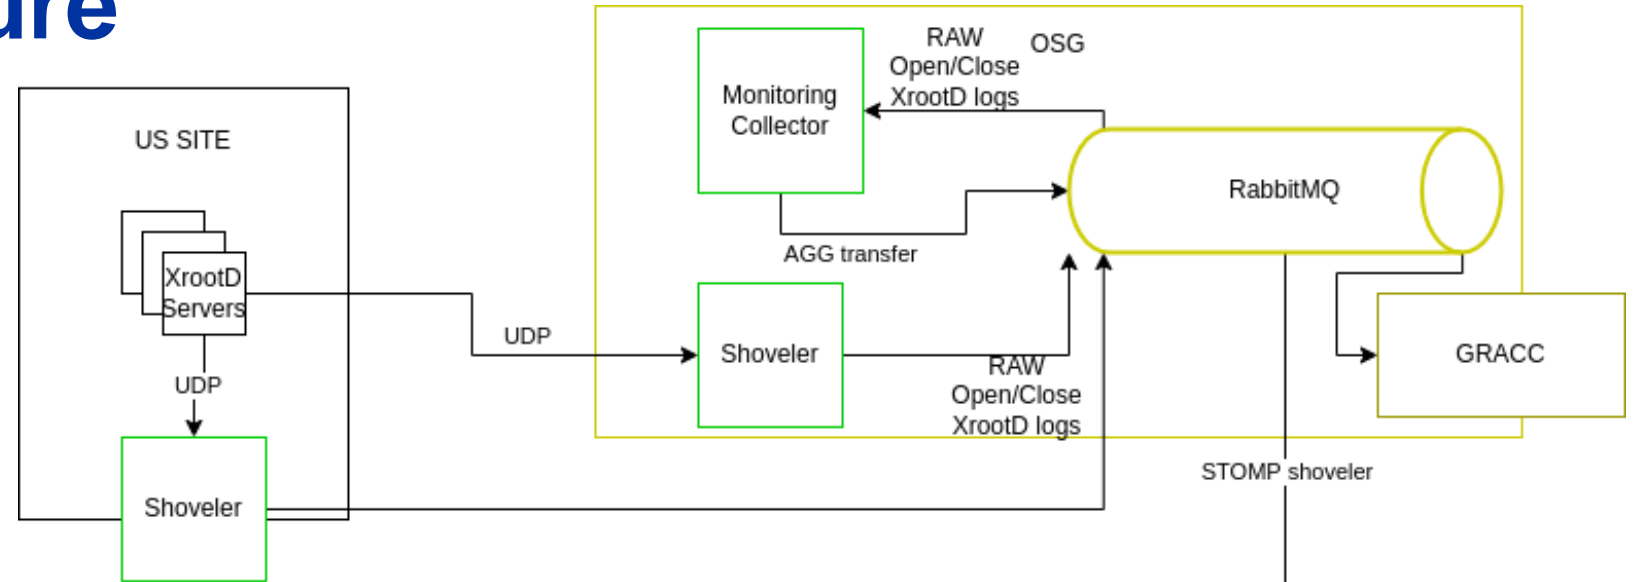
XRootD monitoring improvements

Main goals


- **Redesign current implementation based on XRootD server reports**
 - Relying on the UDP protocol
 - Work implies collaboration between MONIT and OSG developers
- **Coordinate with dCache developers to enable monitoring flow**
 - For the use case dCache+XRootD port
- **Make sure that XRootD is properly integrated in the WLCG transfer monitor**
 - Including also ALICE XRootD monitoring flow

"New" Architecture

OSG



WLCG



Architecture Components

- **Two new components XRootD shoveler and XRootD collector**
 - Already developed and deployed for OSG when WLCG work started
- **XRootD Shoveler**
 - New component that ships XRootD monitoring streams to a message queue
 - Main goal is to deploy it as close as possible to the XRootD server
 - Motivation is to reduce the chance of losing UDP packets
- **XRootD Collector**
 - Similar component to the previous GLED collectors
 - Receives and aggregates XRootD monitoring streams into a "transfer" document

Architecture Components (Missing development)


- **XRootD Shoveler**

- Currently using non-TLS connection with basic auth
 - Request to allow TLS (will require the usage of robot certificates)
- Improve shoveler installation docs for “non-docker” deployments


- **XRootD Collector**

- Possible improvements driven by few issues found during first validation phase
 - Will be mentioned later on another slide

Current situation OSG + WLCG



Once WLCG flow is fully tested/validated



Current situation (OSG)

- **Running deployment for several sites already**
 - All Open Science Data Federation Caches - 9 shovelers across the U.S. and Europe (Amsterdam, Cardiff), Purdue, Florida, Nebraska, UCSD, Caltech, MIT T2 and T3
- **Sending monitoring data to CERN**
 - On a CMS specific flow (so not integrated with WLCG XRootD as for now)
- **Completed validation of the new flow**
 - Correctness and Scale

Current situation (WLCG)


- **Test bed deployment running on a Kubernetes cluster**
 - Currently we are running a battery of shovelers and a collector
 - Shovelers should not be run centrally, but required for testing phase
- **Integrating EOS ALICE servers at CERN**
 - Closer to the development team, within CERN network (so less UDP loss risk)...
 - Few servers reconfigured temporarily for sending data to the new flow
 - Data was flowing as expected, but we observed a **lack of VO information**
- **Shoveler deployed in a few sites for testing purposes**
 - Manchester (ATLAS), RAL (CMS)
 - Data integrated shows as well a **lack of VO information** in some of the cases
 - First numbers comparison with internal monitoring from RAL don't seem to match

Other producers of XRootD data

- **ALICE Monalisa**
 - Current aim is to converge with this new flow
 - XRootD servers will report in parallel to Monalisa and new shovelers
 - WLCG Monitoring information will be based on the shovelers flow
- **xCache**
 - OSG already monitors their XCache instances with this new flow
 - The same will be applied for WLCG
- **dCache**
 - Data will need to be integrated in a separate flow in MONIT
 - Schema of the "final" data checked to be compatible

TF short-term to-do list

- **Adapt the new components to use the message queue implementation at CERN**
 - Shoveler implementation is finished and merged
 - Collector implementation is in testing phase
- **Integrate new flow in monitoring**
 - First testing subject will be CERN internal
- **Validate numbers reported**
- **Converge OSG and WLCG flows within MONIT**
- **Integrate more non US sites**




Questions & Answers

Contact: wlcgmon-tf@cern.ch
Shoveler installation [docs](#)

Acknowledgments

- **OSG contributed XRootD and Site Monitoring improvements:**
This project is supported by the National Science Foundation under Cooperative Agreements OAC-1836650, MPS-1148698 and OAC-2030508. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.



home.cern