



WLCG Monitoring Task Force

Borja Garrido

11.05.2022

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
- **Focused on three main areas**
 - XRootD monitoring improvements
 - WLCG transfers harmonization
 - Site network monitoring integration

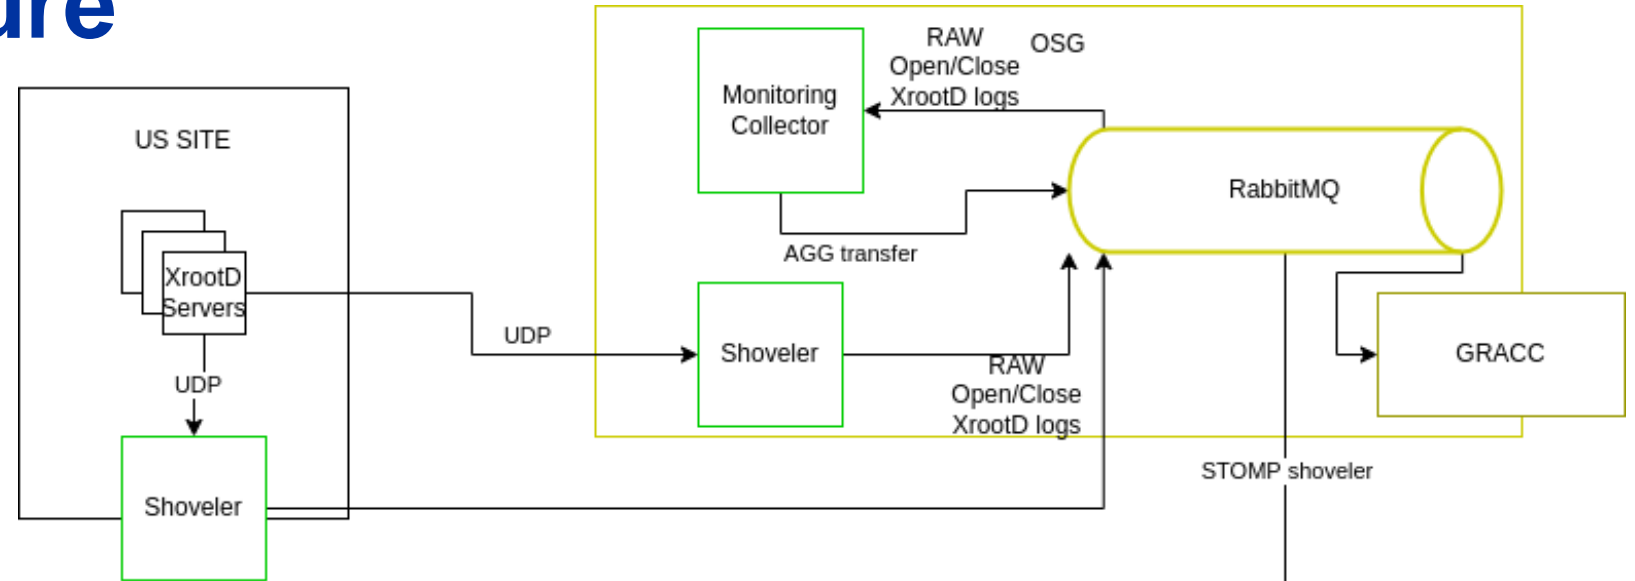
XRootD monitoring improvements

Main goals

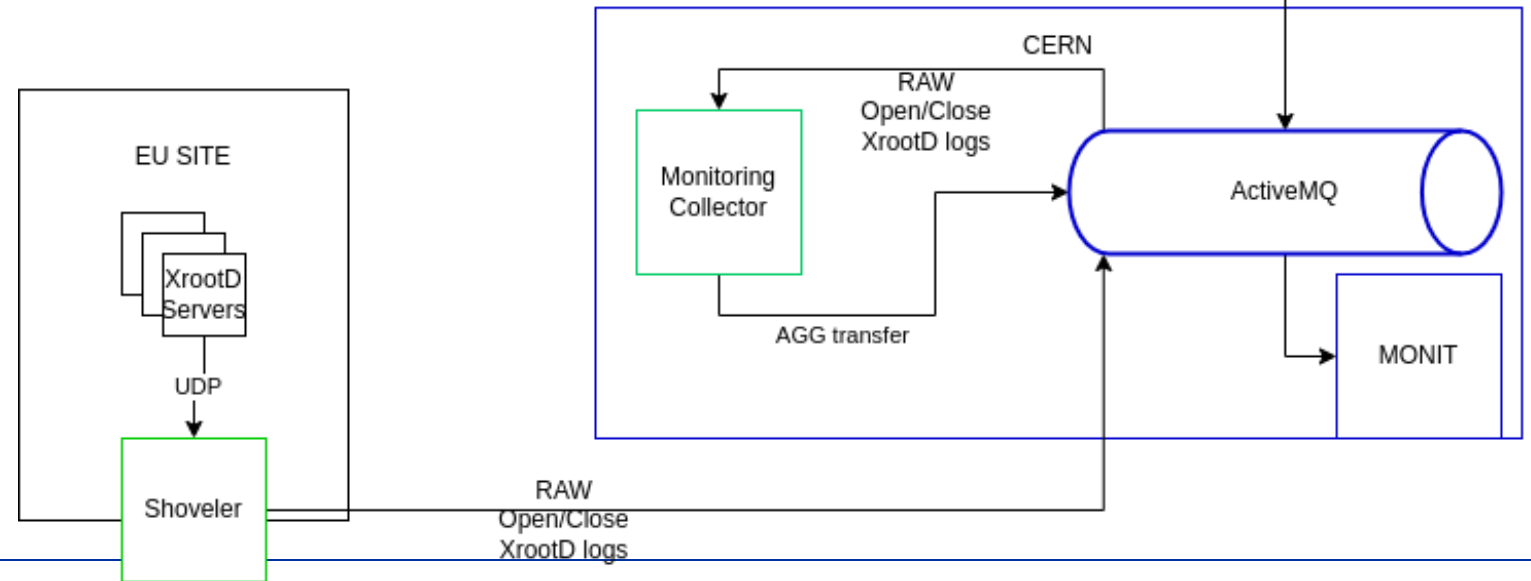
- **Redesign current implementation based on XRootD server reports**
 - Relying on the UDP protocol
 - Work implies collaboration between MONIT and OSG developers
- **Coordinate with dCache developers to enable monitoring flow**
 - For the use case dCache+XRootD port
- **Make sure that XRootD is properly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

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)




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)...
 - Will allow us to compare/validate numbers with Monalisa
- **Some UK sites contacted and volunteered for testing the new flows (Thanks!)**
 - T1 (CMS, ATLAS) T2 (Edinburgh, Lancaster, Manchester)

Converge OSG and WLCG



Once WLCG flow is fully tested/validated



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**






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 being composed and will be circulated when ready



- **Harmonise the DC dashboard cross-experiment functionality with the agreed common schema**

- **Gather feedback from experiments and dashboard users for improvements**

Site network monitoring

Shawn McKee / University of Michigan

Main goals

- As identified in WLCG Network Data Challenge 1, we need **site specific** network traffic information to better understand our sites and infrastructure.
- CRIC already has locations to define information and monitoring links for networks.
- **GOAL-1: Collect site network information in a single place**
 - Network topology information
 - Monitoring links
- **GOAL-2: Make site level network monitoring available for experiment and central operations**
 - Provide a central network monitoring dashboard for the different sites

Site Network Monitoring Components

To better manage, monitor and diagnose site's networks we need to gather some information.

To keep this data up-to-date, we need to make sure site's control their own information.

There are two main things we would like sites to provide

1. **Site network information:** A human readable description of the relevant site network details including optional details about hardware, equipment models, diagrams, etc.
2. **An updating source of network traffic IN/OUT of the site as defined at <https://gitlab.cern.ch/wlwg-doma/site-network-information/-/tree/master/WLCG-site-snmp>**

These two items need to have associated URLs that can be added to CRIC in the network Information and Monitoring URLs.

Site Expectations for Network Monitoring

We have set up a CERN Gitlab to host the documentation and associated example code at <https://gitlab.cern.ch/wlwg-doma/site-network-information>

Our target is to have at least the largest sites summarize their network information by:

- Copying the [template](#) to a new document named as <SITE>.md in [SitePages](#)
- Editing the template to provide the requested (and optional) data

These sites will also need to set up a simple python service that uses SNMP to query their network border interfaces to extract the IN and OUT traffic in Bytes / second.

- The details are available at <https://gitlab.cern.ch/wlwg-doma/site-network-information/-/tree/master/WLCG-site-snmp>
- Sites will need to find a server to host this python script and properly configure it.
- Example for AGLT2 at <https://head01.aglt2.org/aglt2-netmon.json>

Site Network Details


- **We need to understand what information is REQUIRED vs OPTIONAL**
- **We may want to discuss limiting access to some of the information site's might consider sensitive, e.g., detailed network diagrams or network operating system versions**
- **The SNMP query to JSON output for gathering each site's IN/OUT is a possible mechanism to gather this information but we may wish to evolve both what is gathered and how it is gathered.**
- **The site traffic metrics need to end up in CERN Monit so we can enable appropriate dashboards and have the data available for future testing, diagnosis and data challenges.**

TF short-term to-do list

- Network information template and draft of documentation for sites is ready



-
- **We need to get a few big sites to fill out the Site Network Template**
 - Part of this work is to deploy site traffic monitoring (SNMP to JSON)
 - These sites can help identify any issues and provide feedback
 - **We also need to determine if JSON format and a web URL are the best method to make site IN / OUT data usable for WLCG**
 - **We are very interested in getting any feedback on the network information template and the site traffic monitoring plans.**



Start Deployment Campaign

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