



XRootD monitoring status and next steps

WLCG Monitoring Task Force

Alessandra Forti, Borja Garrido, Derek Weitzel, Fabio Andrijauskas, Julia Andreeva, Katy Ellis, Shawn McKee

WLCG Operations Coordination

June 13, 2024



Index

- Introduction
- Status
 - XRootD servers (NG)
 - dCache with XRootD door
- Next steps
- Summary



Introduction

Goal: Provide reliable monitoring for XRootD transfers under the WLCG umbrella

Evolution of current GLED infrastructure (not reliable) into NG (Next generation)

- New components: Shoveler, to be run (desirably) at sites to overcome UDP unreliability

Monitoring of dCache with XRootD door

- New components: WLCG converter to forward dCache monitoring to WLCG



Status: XRootD NG

- Deployed in few selected/volunteered sites for DC2024
 - Main DC2024 goal was to cover CERN
- Running on a “production” environment
- Identified issues:
 - Shoveler instabilities: Identified ones fixed already
 - Enrichment mismatches



Status: dCache with XRootD door

- Deployed in few selected/volunteered sites for DC2024
 - Main goal for DC2024 was to cover FNAL (CMS pileup)
- Sites need dCache 9.2+ and to run the [WLCG converter](#)
 - Special thanks to Sandro Grizzo
- Identified issues:
 - Enrichment mismatches



Status

Is this data reliable?

OSG did a first validation on a controlled environment (NG)

- [Validation](#) and [Scale validation](#)

For WLCG it was required to adapt slightly the components

- Initial validation against some specific CERN EOS servers
- Missing some real validation against the “*truth*”



Next steps: Validation (NG and dCache)

- Need to select few key sites to check data against the “*truth*”
 - Network monitoring, controlled environment, XRootD cinfo files...
 - For XRootDNG we can also use the ALICE data in MonALISA
- Data mismatch could be due to several reasons:
 - Missing transfer streams (rare)
 - Missing extra information (to be fixed upon identification)



Next steps: Fixing enrichment mismatches

Wrong enrichment = Wrong plotting

Enrichment on NG flow

- Collector translates IPs/hostnames to domain
- MONIT matches domains against CRIC data

Enrichment on dCache flow

- MONIT resolves IP against CRIC to infer site
- MONIT marches sites against CRIC data

Issues so far: Internal IPs (10.X.X.X), Not resolvable IPs, Missing info in CRIC...



Next steps: GLED retirement

- Currently running on CC7 machines (EoL 30th June)
- Data not reliable/being used (little point on migration effort)
- **Proposal:** Stop the service and focus on the new flows
- **Impact:**
 - No data for GLED flows
 - Site managers to reconfigure their XRootD servers
 - Stop sending data to old aliases (mainly for cleanliness but nothing should break if not done)



Summary

- We need to validate new flows before doing a full deployment campaign
 - One/Two sites (per VO), iterative approach with issues identified
- Things should be ready (hopefully well in advance) for DC2026
 - Will current manpower be enough?
- GLED retirement by the 30th of June
 - Little point on keep running it knowing current usefulness
 - Migration to newer OS will require effort
 - Already SLC6 to CC7 was not trivial

Q & A