WLCG DC Monitoring follow-ups

WLCG Monitoring TaskForce

07.11.2022

WLCG Monitoring Task Force

- **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

Introduction

- **WLCG DC 2021**
  - Successfully performed
  - Last presentation provided May 2022
  - Monitoring recommendations were done
- **WLCG Monitoring TF**
  - Focused on three main areas:
    - XrootD Monitoring re-structure
    - Transfer data harmonization
      - Common Schema
      - Common DC Dashboards
    - Site monitoring
- **Last update provided during GDB**

Monitoring

- DC Dashboard generally considered useful
  - Used for both DC and TC I & II
    - Required expertise in the data structures to plot consistent information
    - Many plots could only be static selections
    - xrootd likely under estimated
- WLCG monitoring TF following DC recommendations
  - Re-structure xrootd monitoring infrastructure
  - Agree a common schema between the experiments, FTS & xrootd
  - Refactor DC dashboard to use the new schema
  - Add site monitoring

Tape only traffic - static selection

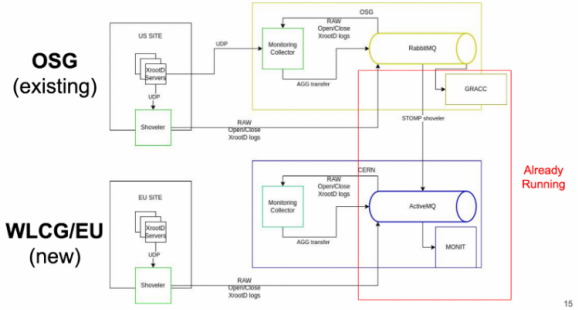

New xrootd monitoring infrastructure

14

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 specificalities for experiments

General Plots

Attempted Transfers

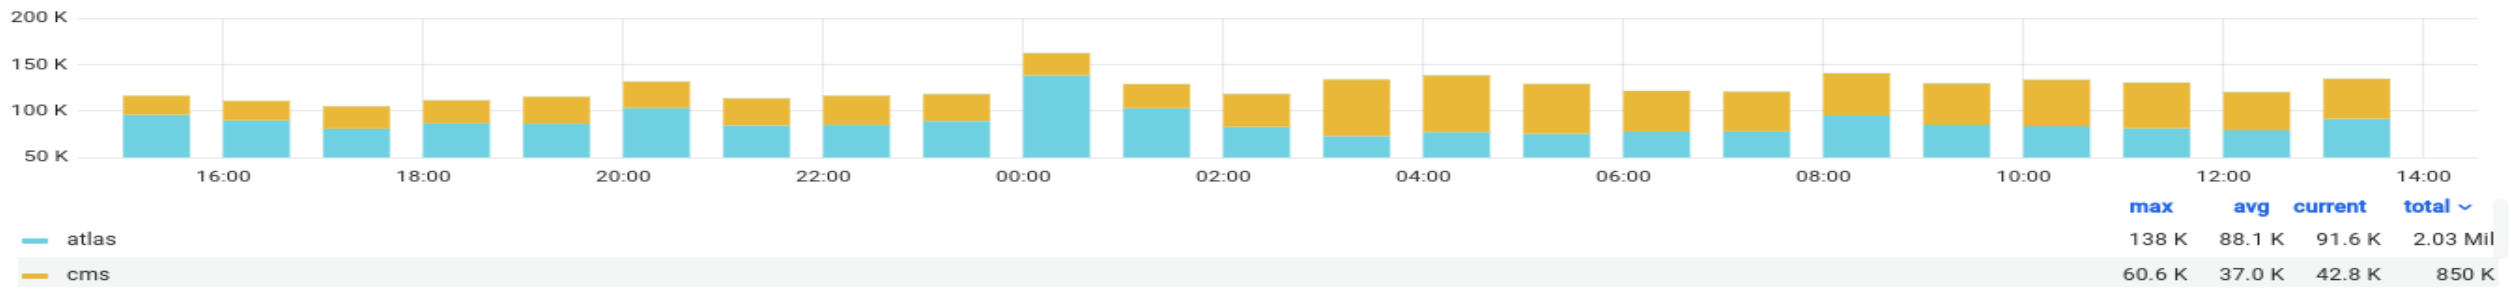

	max	avg	current	total
atlas	138 K	88.1 K	91.6 K	2.03 Mil
cms	60.6 K	37.0 K	42.8 K	850 K

Successful Transfers (%)

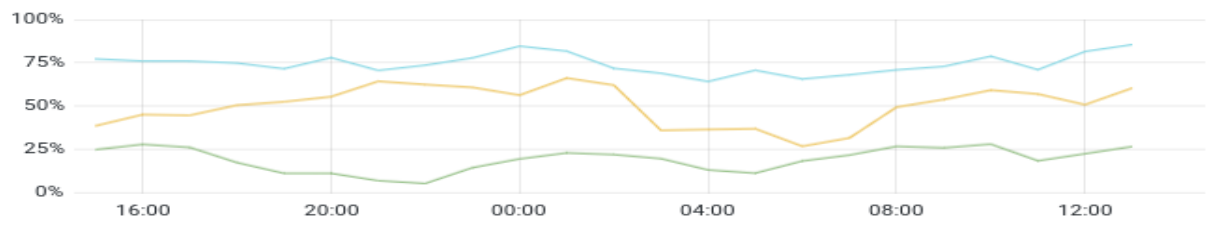

	avg	current
atlas	74.5%	85.5%
cms	50.5%	60.6%

Transfers Throughput

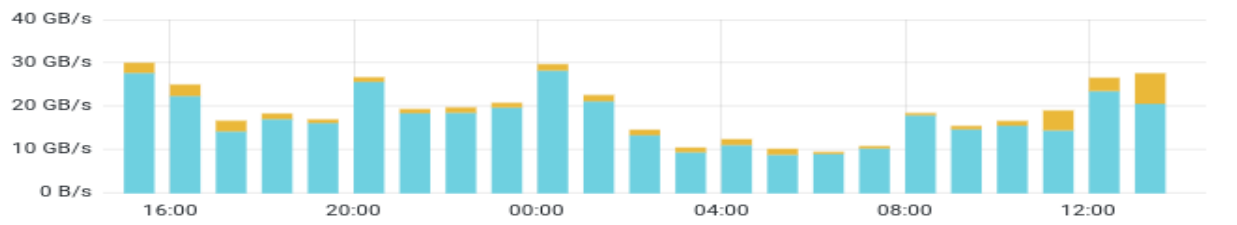

	max	avg	current
atlas	28.3 GB/s	17.3 GB/s	20.6 GB/s
cms	7.05 GB/s	1.75 GB/s	7.05 GB/s

Successful Transfers

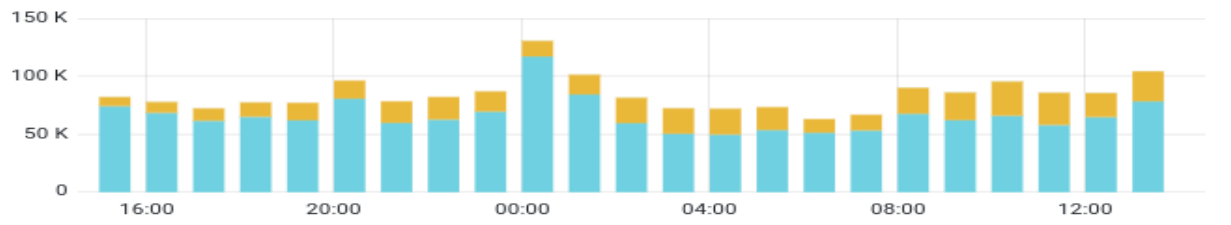

	max	avg	current	total
atlas	117 K	66.1 K	78.3 K	1.52 Mil
cms	29.5 K	18.2 K	25.9 K	419 K

Failed Transfers

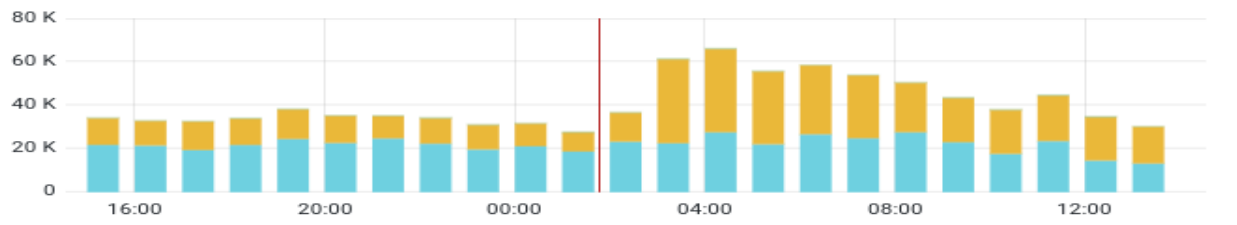

	max	avg	current	total
atlas	27.6 K	22.0 K	13.2 K	505 K
cms	38.6 K	18.8 K	16.9 K	432 K

Status update

- **First dashboard version was created**
  - With striped VO specificalities and current available schema
- **Draft document was created and circulated**
  - Proposing the minimum required fields that the tools will need to provide
- **Feedback was collected and discussed**
- **Agreement reached with experiments on the fields to aim for**
- **Meet developers (XRootD, FTS, dCache)**
  - To understand need for changes to fit the schema
- **Adapt dashboard to make full usage of the new schema**

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

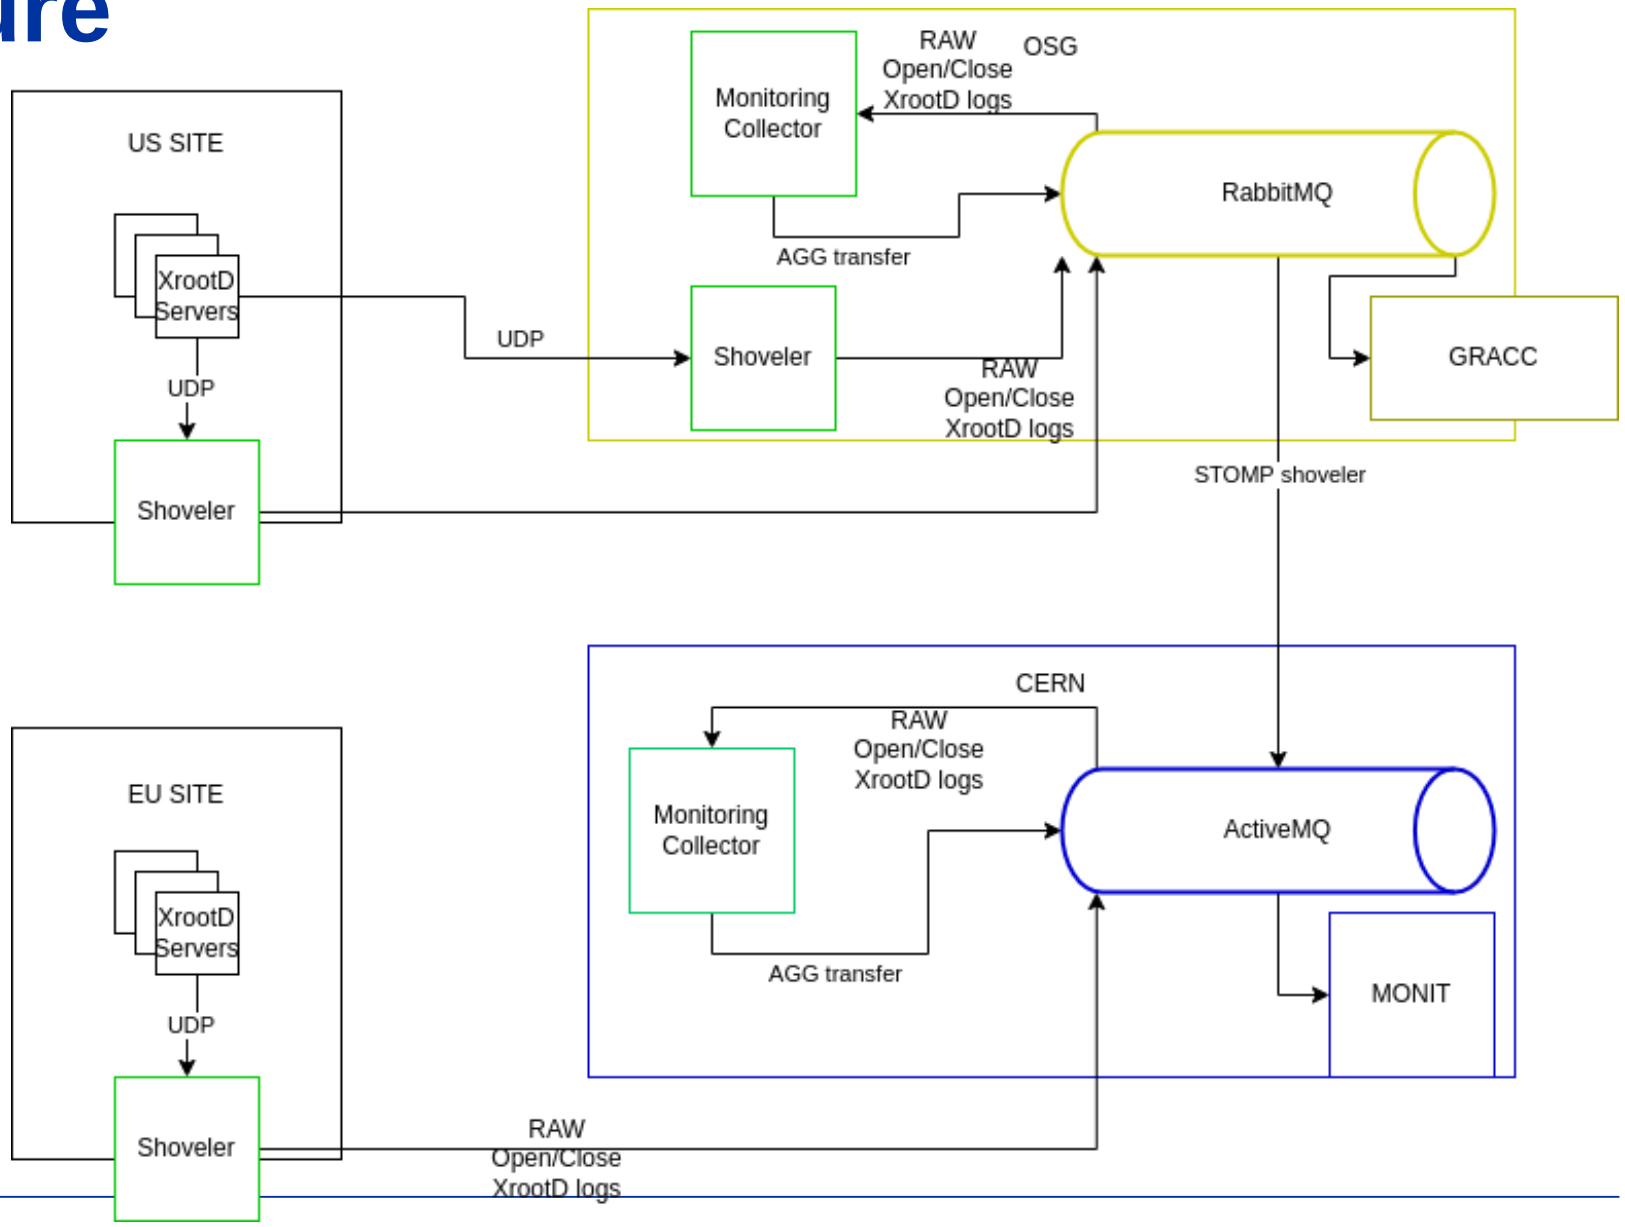

Architecture Components

- **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

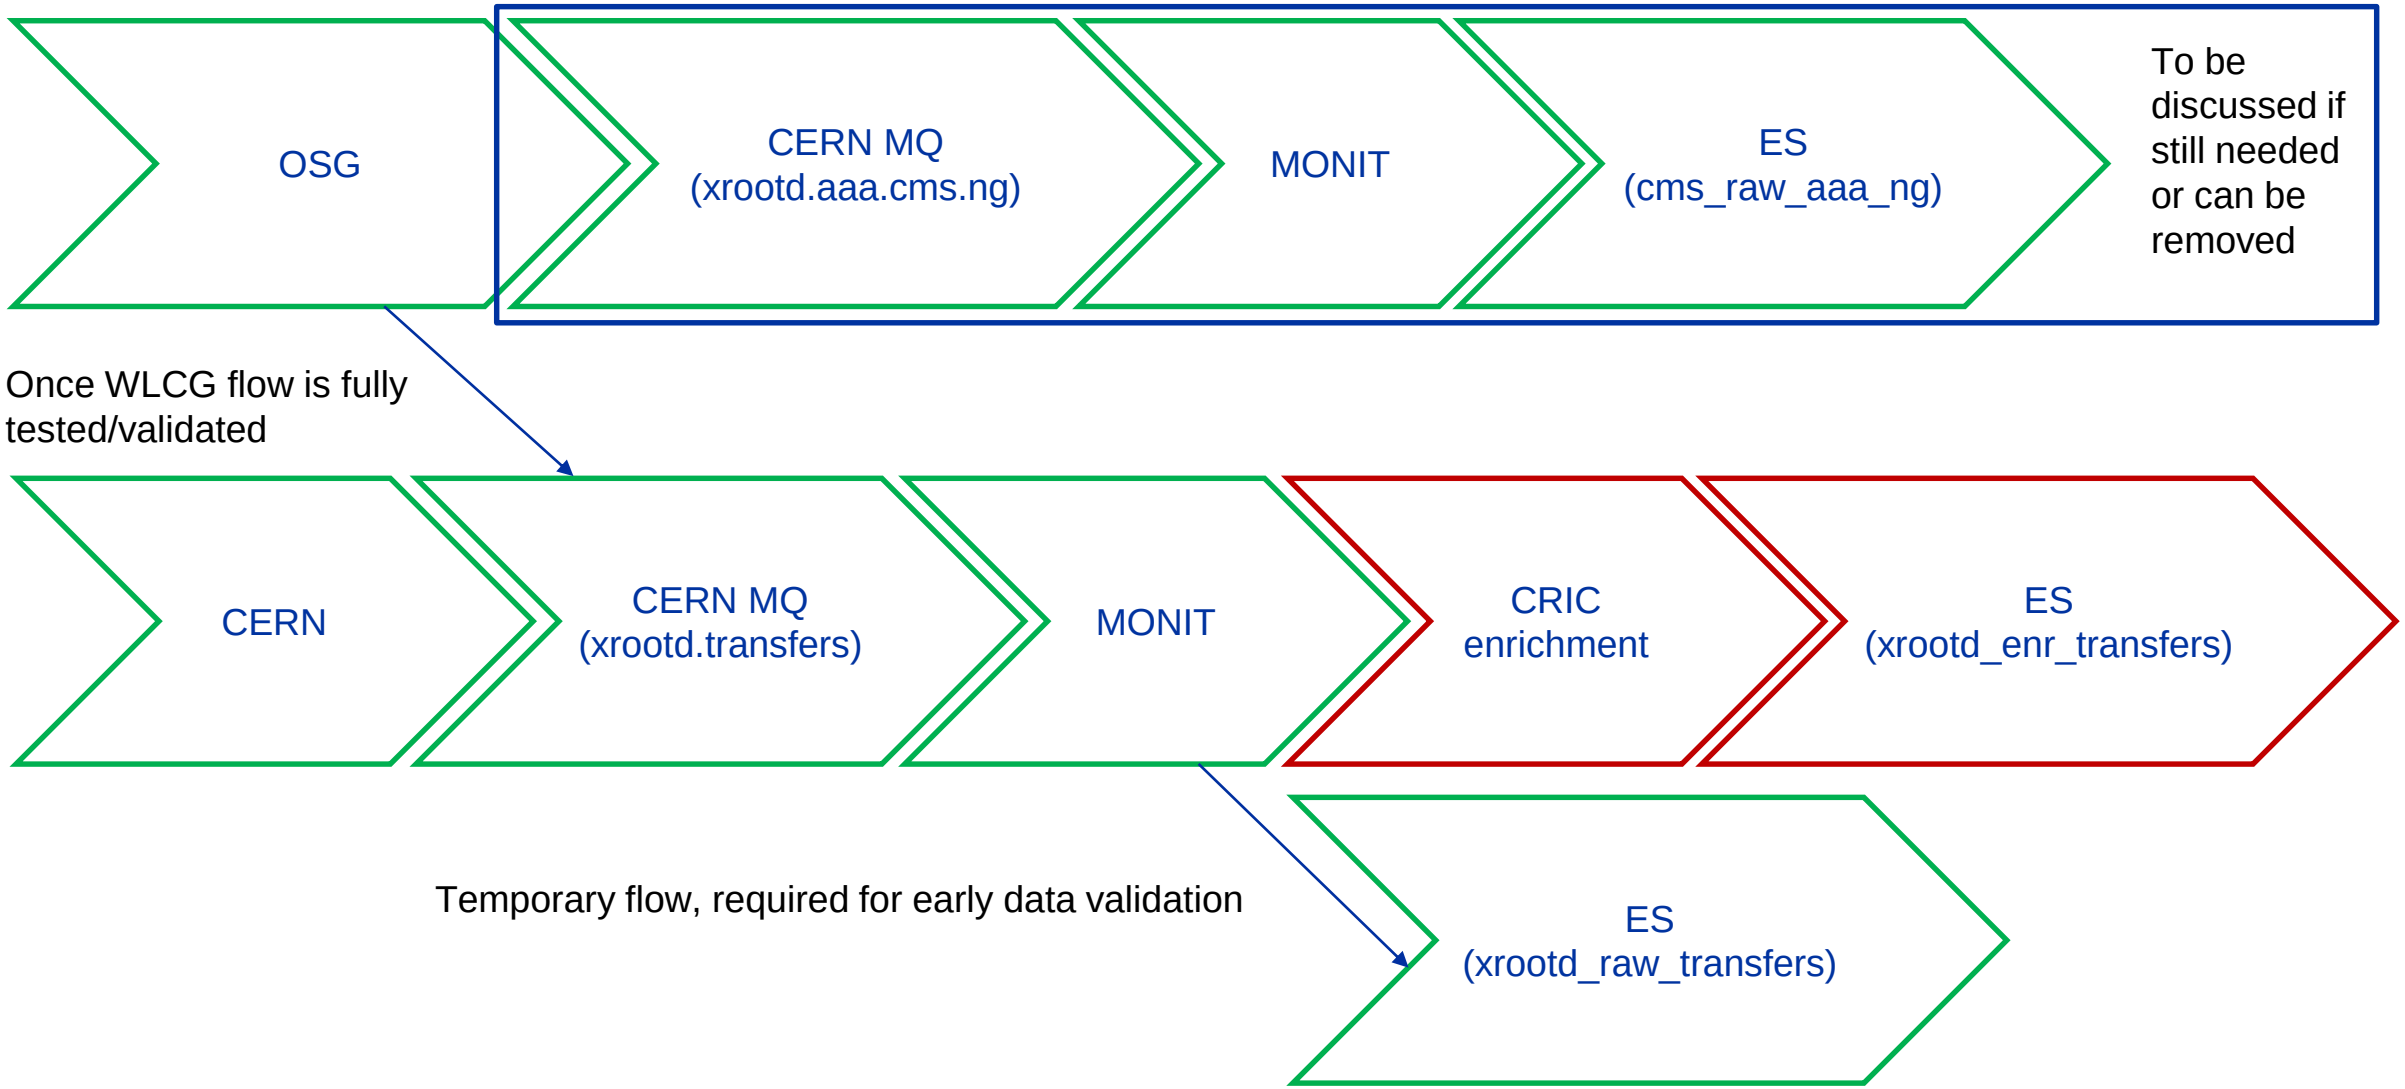

Current situation (OSG)

- **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

Last reported 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

Current situation (WLCG)

- **Started several tasks to tackle issues found on first phases**
- **Numbers validation**
  - Validation of numbers between new flow and Monalisa
  - Set a new test flow between OSG CERN to compare numbers between both deployments
  - Re-do numbers validation for RAL with experts to understand discrepancies
- **Lack of fields**
  - Deemed to be produced by specific server configurations
  - The idea is to start deploying the available components to more sites to assess situation better
    - Missing a small development for the shoveler to work with certificates (CERN MQ requirement)

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

Network site monitoring

Motivation

- **Lack of information about network traffic for sites During DC1 in October 2021**
- **Total traffic to/from our sites needed to identify issues/bottlenecks**
- **Find a way to gather the minimal amount of information that will help us understand and improve how our sites work across the WAN.**

Main goals

- **Provide human readable information about a sites network**
  - Template created for sites to clone and fill our following the instructions
  - Examples can be found in Gitlab
  - File needs to be converted to HTML, uploaded to a Site webserver and linked in WLCG CRIC
- **Provide site's IN and OUT network traffic (total)**
  - Updated each minute, in a publicly accessible URL in JSON format
  - Example Python3 script implemented
    - Queries one or more interfaces representing site's boundary

Status update

- Preparation of Network template for sites to fill in
- Add needed fields in CRIC for sites to define Monitoring links
- Develop script to gather and expose minimum set of network metrics
- Integrate metrics in MONIT for visualization
- Tune-up script as required for different sites
- Deploy script in T1s initially so a dashboard can be created

Site’s homework

Sites will be expected to:

- **Document their network at a high-level, with options to provide helpful details**
  - Information about site topology, peering, hardware and capacity can allow WLCG network experts to better support, diagnose and fix network problems
- **Gather (via snmp or other data source) the IN/OUT traffic of their whole site**
  - This requires identifying the interface(s) that represent the “border” of the site and a mechanism to gather the interface(s) traffic
- **Provide URLs to access the network information and monitoring files**
- **Maintain the documentation and update monitoring as networking at the site evolves**

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