



XRootD Monitoring Report

A.Beche

D.Giordano




Outlines

- Talk 1: XRootD Monitoring Dashboard
 - Context
 - Dataflow and deployment model
 - Database: storage & aggregation
 - User interface & use cases
 - Open issues & future work
 - Summary
- Talk 2: Beyond XRootD monitoring
 - HTTP/WebDAV integration
 - Integration in the WLCG Transfers Dashboard

XRootD federation monitoring

- Activity started during summer 2012
 - 4 sites for FAX, 11 for AAA

Number of sites reporting



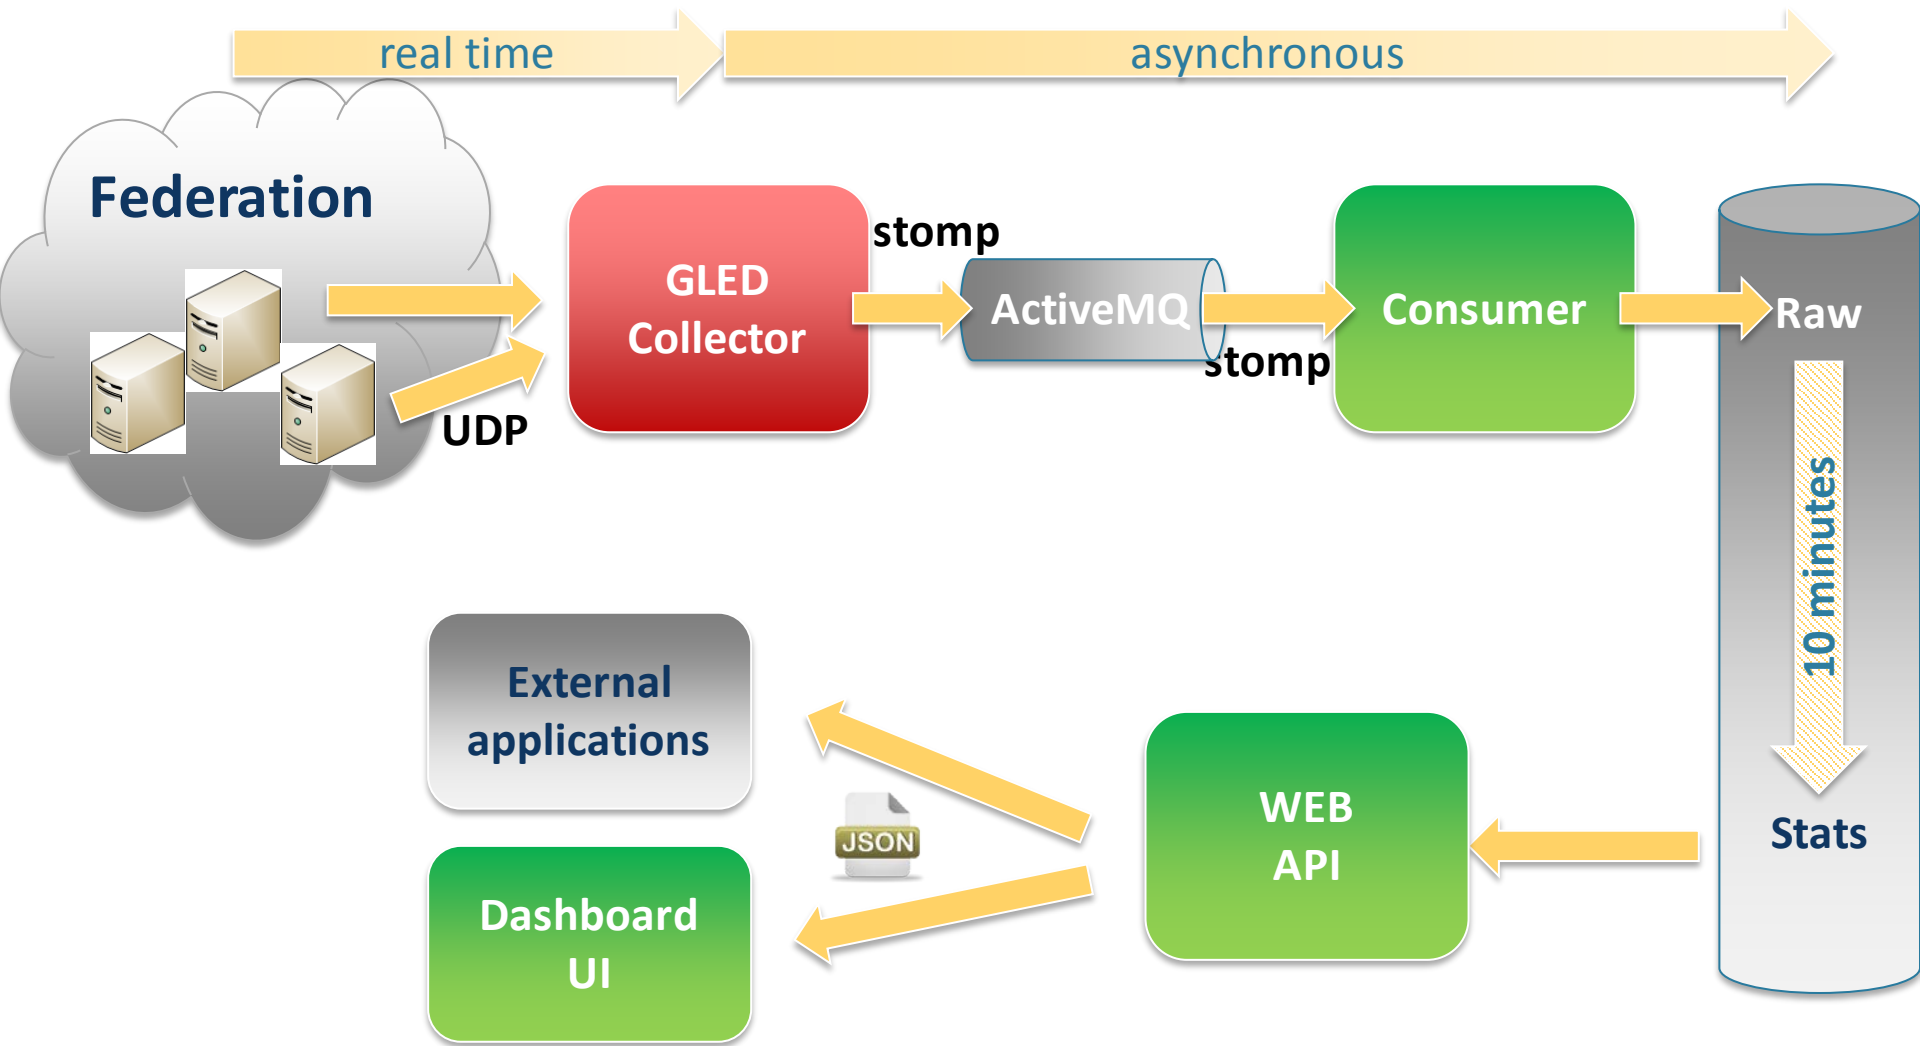
Monitoring data increased accordingly

	July 2012	March 2014
AAA	606k	43M
FAX	15k	22M


Why monitoring ?

- Understand data flows to estimate data traffic
- Provide information for efficient operations
- Identify access patterns and propose data placement strategies

XRootD monitoring dataflow



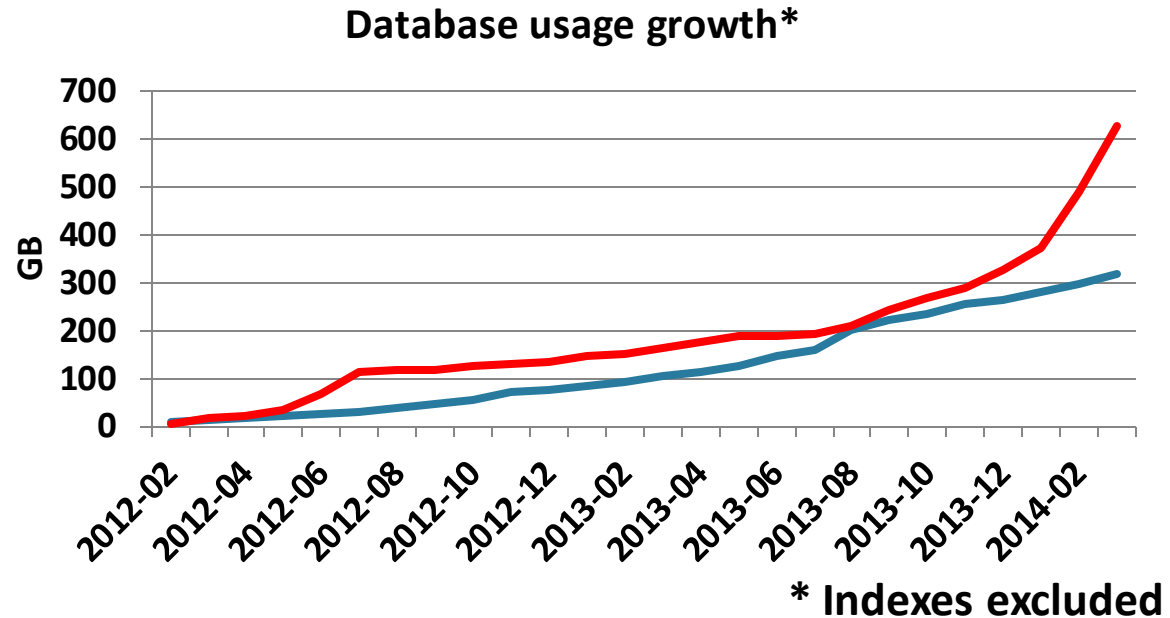

GLEED Deployment model



Consolidated dataflow

- Two usage of these raw data:
 - Dashboard monitoring
 - XRootD popularity
- Now share the same database:
 - Storage optimization
 - Consistency guaranteed

Database




- Storage
 - Raw, statistics, metadata
 - Tables daily **partitioned**, no global indexes


Database

- Raw data aggregation:
 - Done using PL/SQL procedures
 - Events are **unordered**
 - **Stateless: Full re-computation** of touched bins each time
 - Compute stats from raw data in 10 min bins
 - Aggregate 10 min stats in daily bins

Aggregation methods




Easy method



Aggregation methods

Transfers



Transfers

1

0

0

2

1

Bytes

10

0

0

15

20

Transfers

1 (1)

1 (0)

2 (0)

3 (2)

1 (1)

Bytes

8


1

14 (9+6)


15 (1+9+5)

5

Easy method



Adopted method



Visualization Interface



FAX Monitoring (RC.34)

latest statistics update: 2013-11-27T09:10:00.038192

TRANSFER PLOTS (2013-11-27 05:20 to 2013-11-27 09:20 UTC SLIDING)

PLOT: GROUPING ▾ TYPE ▾ SERIES ▾ SIZE ▾ STYLE ▾ BIN: SIZE ▾ FORMAT ▾ STEP ▾

▼ Summary

Common options

Interval
Last 4 hours

Access type
Remote access

Transfer VS Reading
Reading
Copy

Client - Server options

Server:
Grouping: SITE

Clients:
Grouping: SITE

Users:

File:

Interval





Access type

Transfer VS Reading

Sources

Destinations

Matrix Transfer Plots Access Plots Site Statistics Custom Ranking Plots Site History Access Pattern MAP (Experimental)




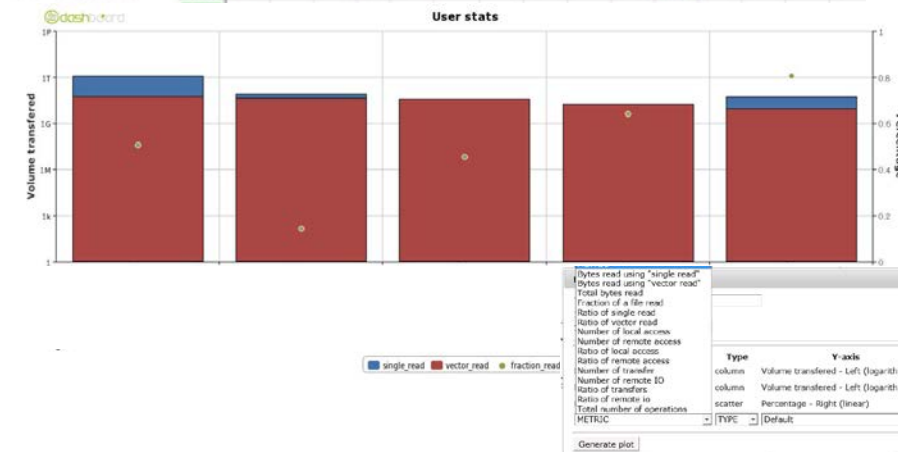

IT-SDC

A.Beche – Federated Workshop

10 – April - 14

12

Pre-defined set of views




Use case example

Understand site access patterns




Read Bytes
T1_US_FNAL

2013-09-24 04:20 to 2013-09-24 08:20 UTC



Read Bytes
T1_US_FNAL

2013-09-24 04:30 to 2013-09-24 08:30 UTC




1. Which sites are reading from FNAL
2. Zoom to a specific site to understand which users are reading
3. Understand which files are read by a user



Number of accesses
T1_US_FNAL

2013-09-24 04:30 to 2013-09-24 08:30 UTC




Data popularity

- XRootD monitoring provides information about file access patterns:
 - Including non official collections (ie: user files)
 - Contribute to simplify and make more efficient the usage of disk resources
- Popularity data analytics built on this information:
 - Adopted already for CMS-EOS
 - will be extended to full AAA

Archive recommendation for CMS-EOS

- Help to manage the disk space of EOS including user space
 - No central bookkeeping system



- Unused files: created > 4 months ago, no access in the last 3 months:
 - ~500 TB of space occupied and not used \Leftrightarrow 30% of total for these areas

Open issues

- Missing servers:
 - Dcache sites
- Server should provide their site name.
 - CMS: only 5 sites:
 - anon, BUDAPEST, Hefy-Vienna, T2_US_USCD, UKI-LT2-Brunel
 - Not coherent convention naming
 - ATLAS: GLED RPM to be deployed
- GLED Collector improvements:
 - Reliability of the service:
 - Recover time, can be long due to time difference
 - GLED should be operated as a production service
 - Scalability:
 - to be fixed with automatic reconnection soon

Future work

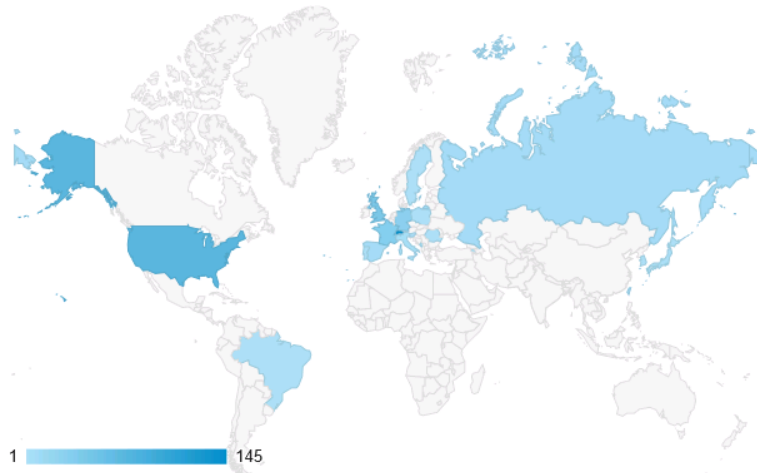

- Strong requirement from ATLAS to understand efficiency:
 - Need the concept of error / failure
 - How XRootD server could be instrumented to report it?
- European GLED collector is up and running:
 - Only 1 pilot site is reporting to it (CNAF)
 - Should we keep it?
- Data mining activity (not started yet):
 - Almost 2 years of raw data (1TB)

Data Mining

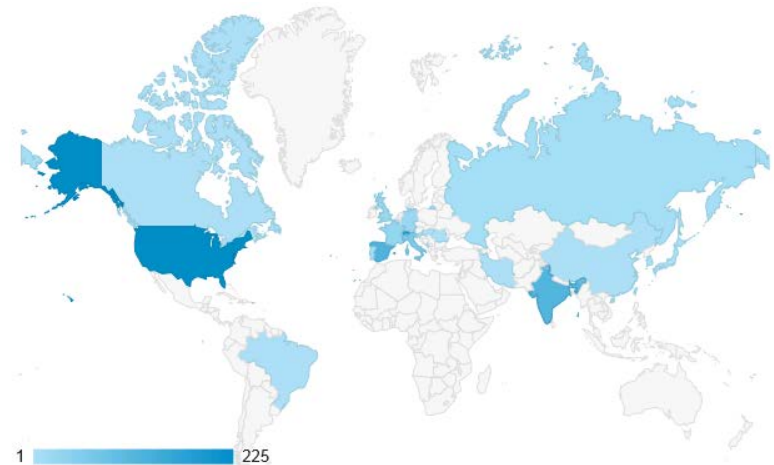

- Extract further knowledge from the data...
 - Detect inefficiencies
 - Propose deletion strategies
 - Define data placement
- ... by
 - Understand access patterns and data usage
 - Correlate data traffic and data access performance
- Possibility to automate some operations

Application usage

FAX



AAA



Summary

- Monitoring federations is a challenge
 - High rate of traffic & information
 - Challenge met by data aggregation, scalable technologies
- Dashboard is not actively used
 - Less than 10 daily users (FAX), less than 15 (AAA)
 - Is there any missing functionalities?
- Improvement work is ongoing
 - New requests are coming
- XRootD monitoring is a one piece of the entire Data transfers puzzle
 - See next talk



Beyond XRootD monitoring

A.Beche

D.Giordano



Outlines


- Talk 1: XRootD Monitoring Dashboard
 - Context
 - Dataflow and deployment model
 - Database: storage & aggregation
 - User interface & use cases
 - Open issues & future work
 - Summary
- Talk 2: Beyond XRootD monitoring
 - HTTP/WebDAV integration
 - Integration in the WLCG Transfers Dashboard

HTTP Federation is coming

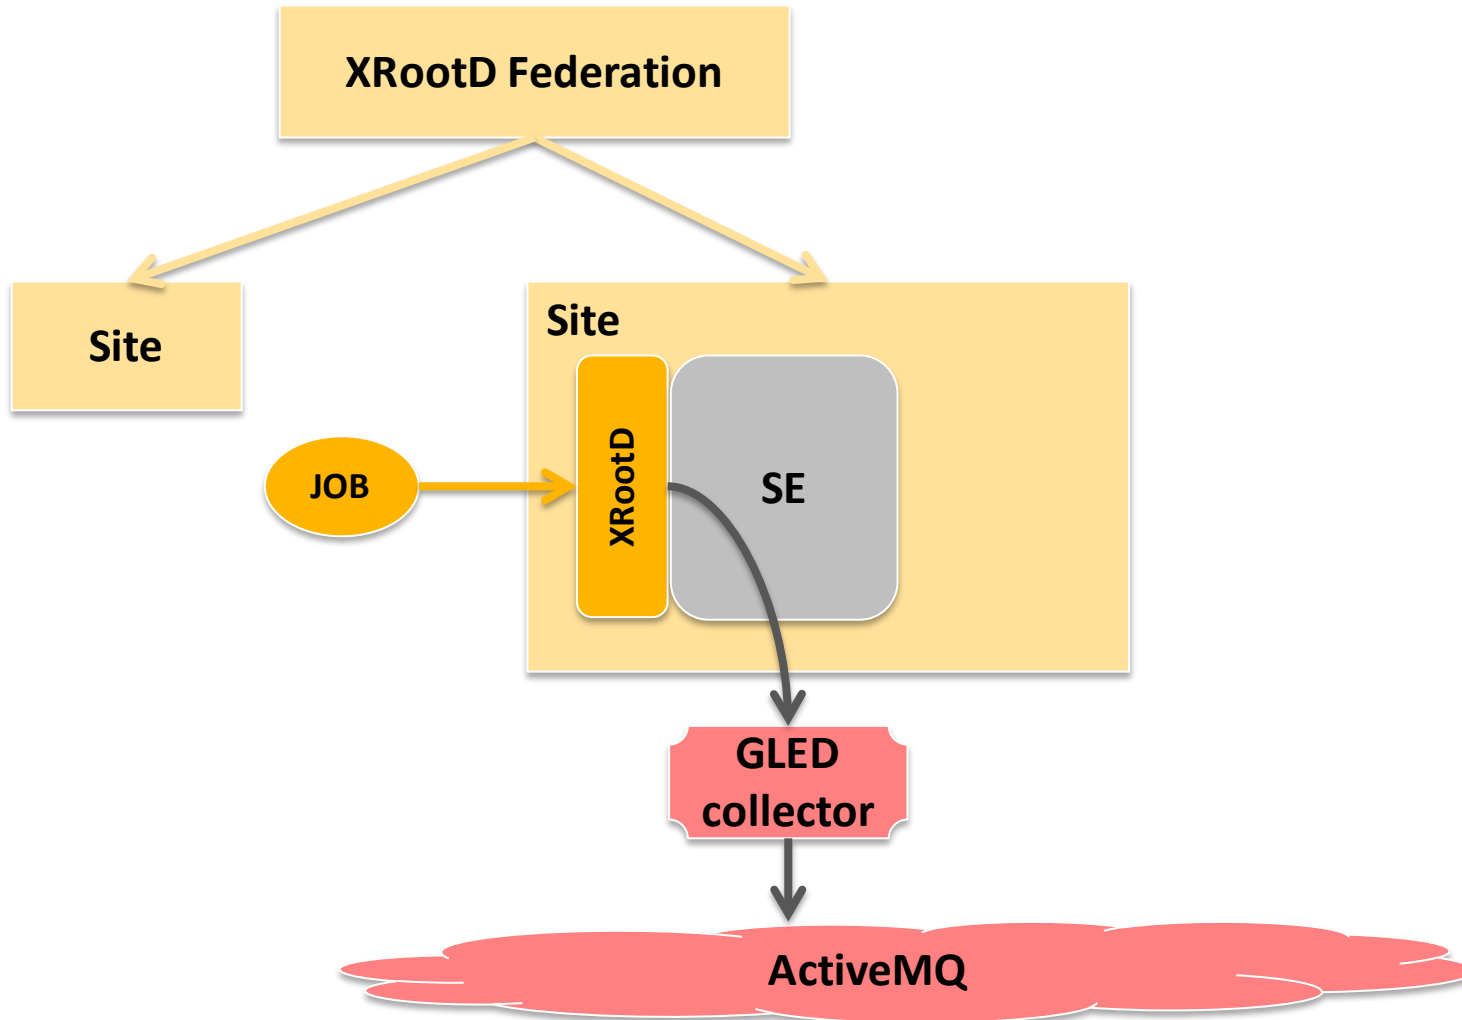
- HTTP protocol will be used in the future
 - XRootD servers can be accessed
 - See Fabrizio's presentation on xrdhttp
- Two kind of accesses:
 - Pure HTTP access (through Apache)
 - HTTP gate to XRootD server
- Can't be monitor in the same way

Monitoring XRootD access protocol

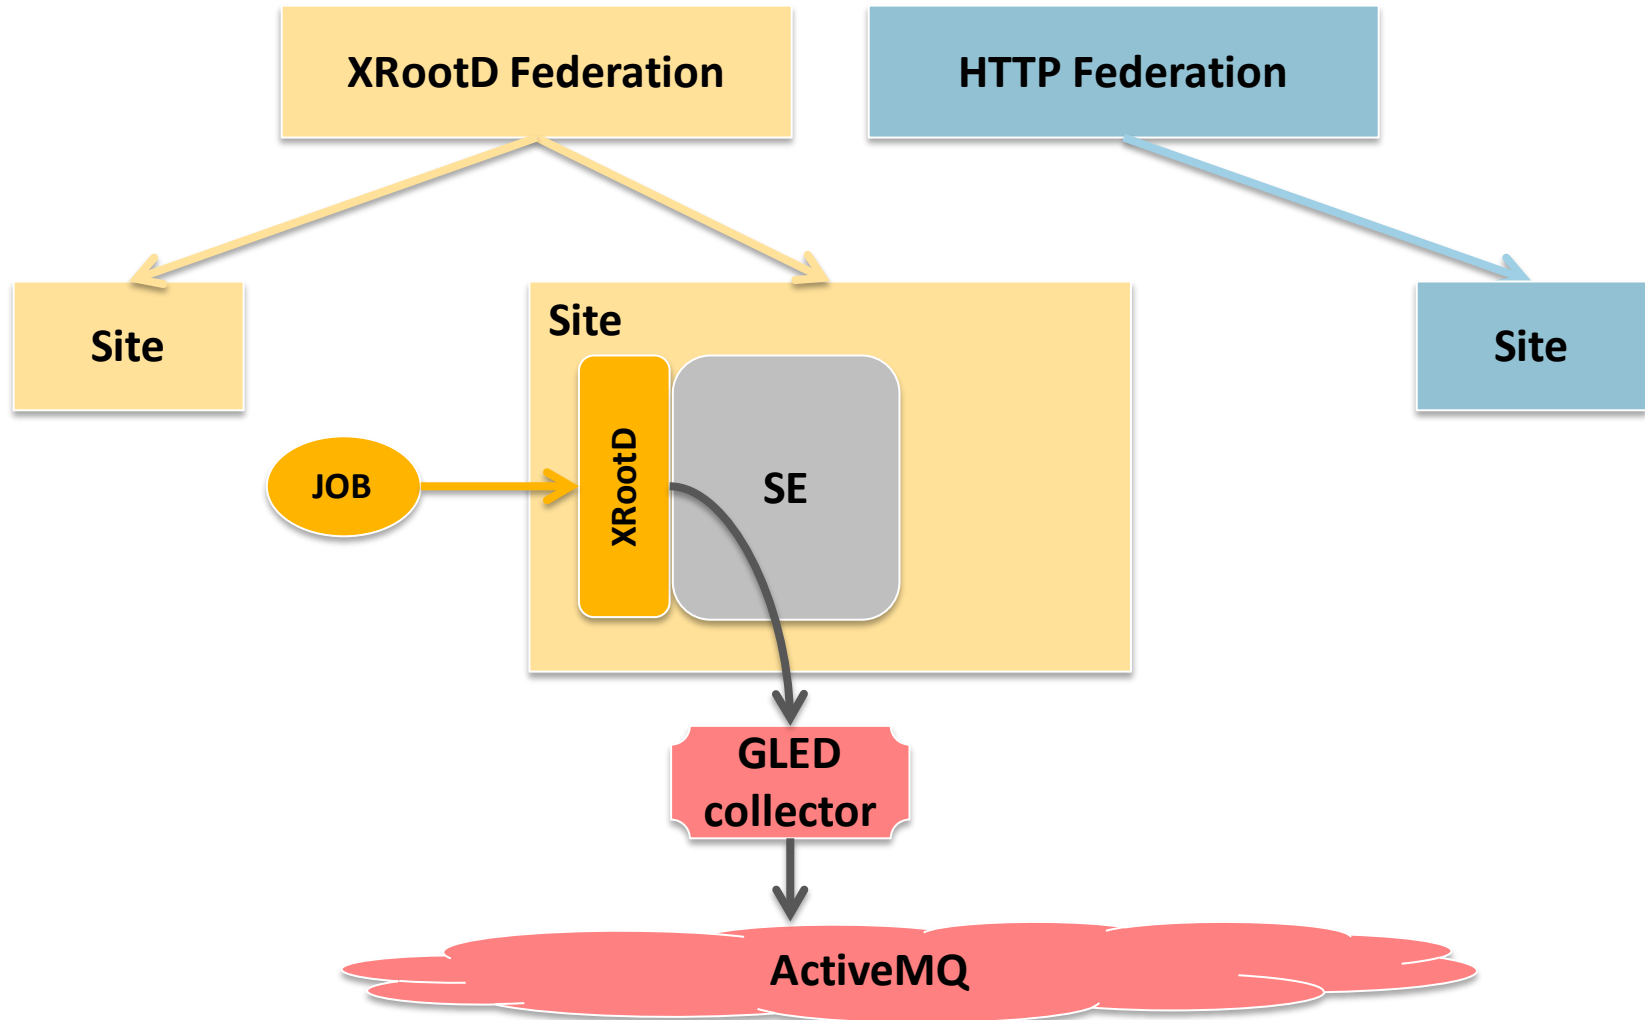
- XRootD 4 will now reports the user protocol:
 - All the monitoring chain needs to be updated
 - Dashboard DB and UI are fully ready



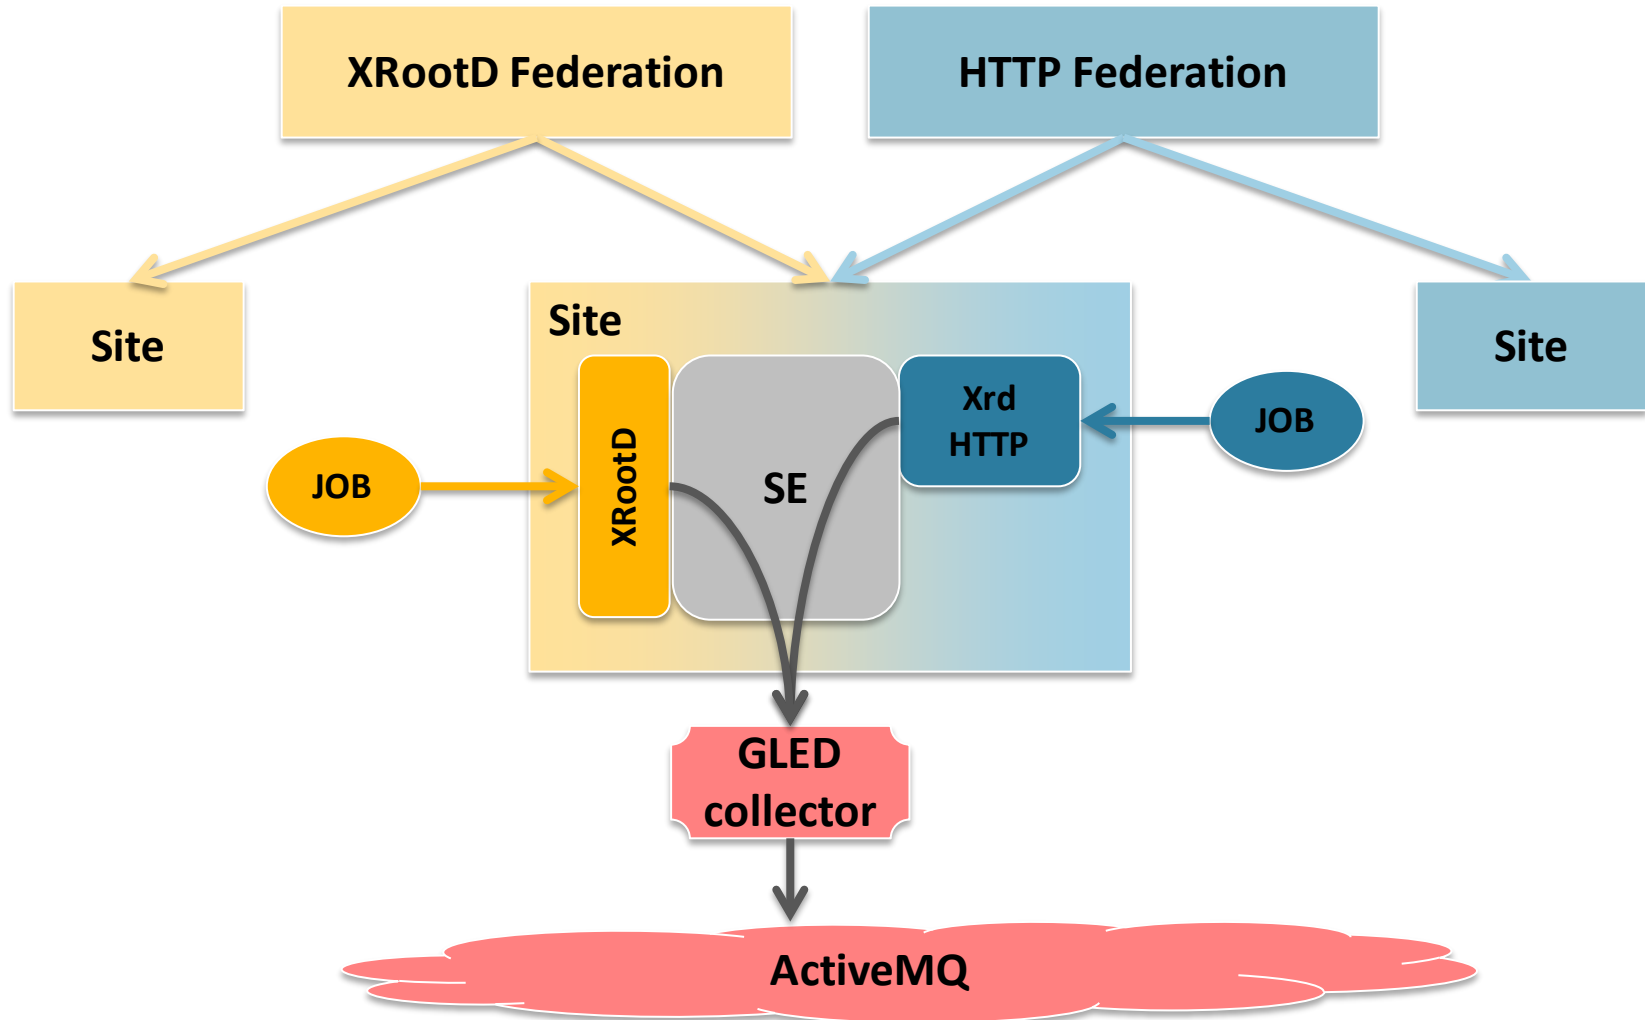
HTTP/WebDAV federation monitoring



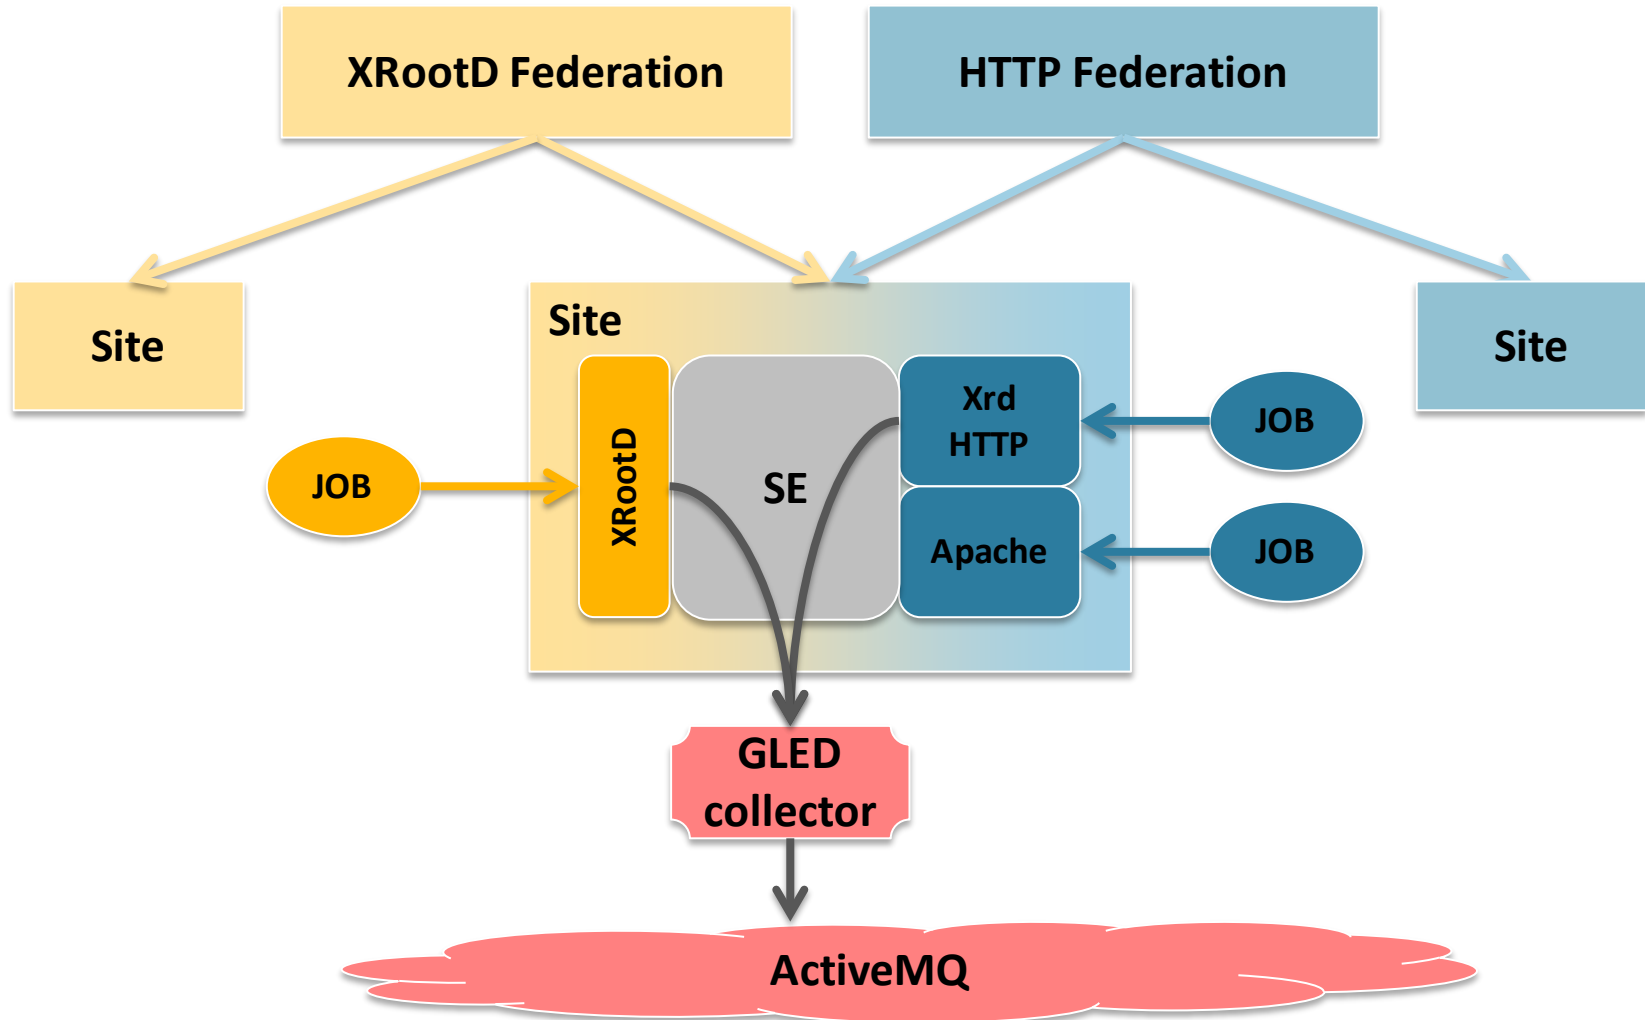
HTTP/WebDAV federation monitoring



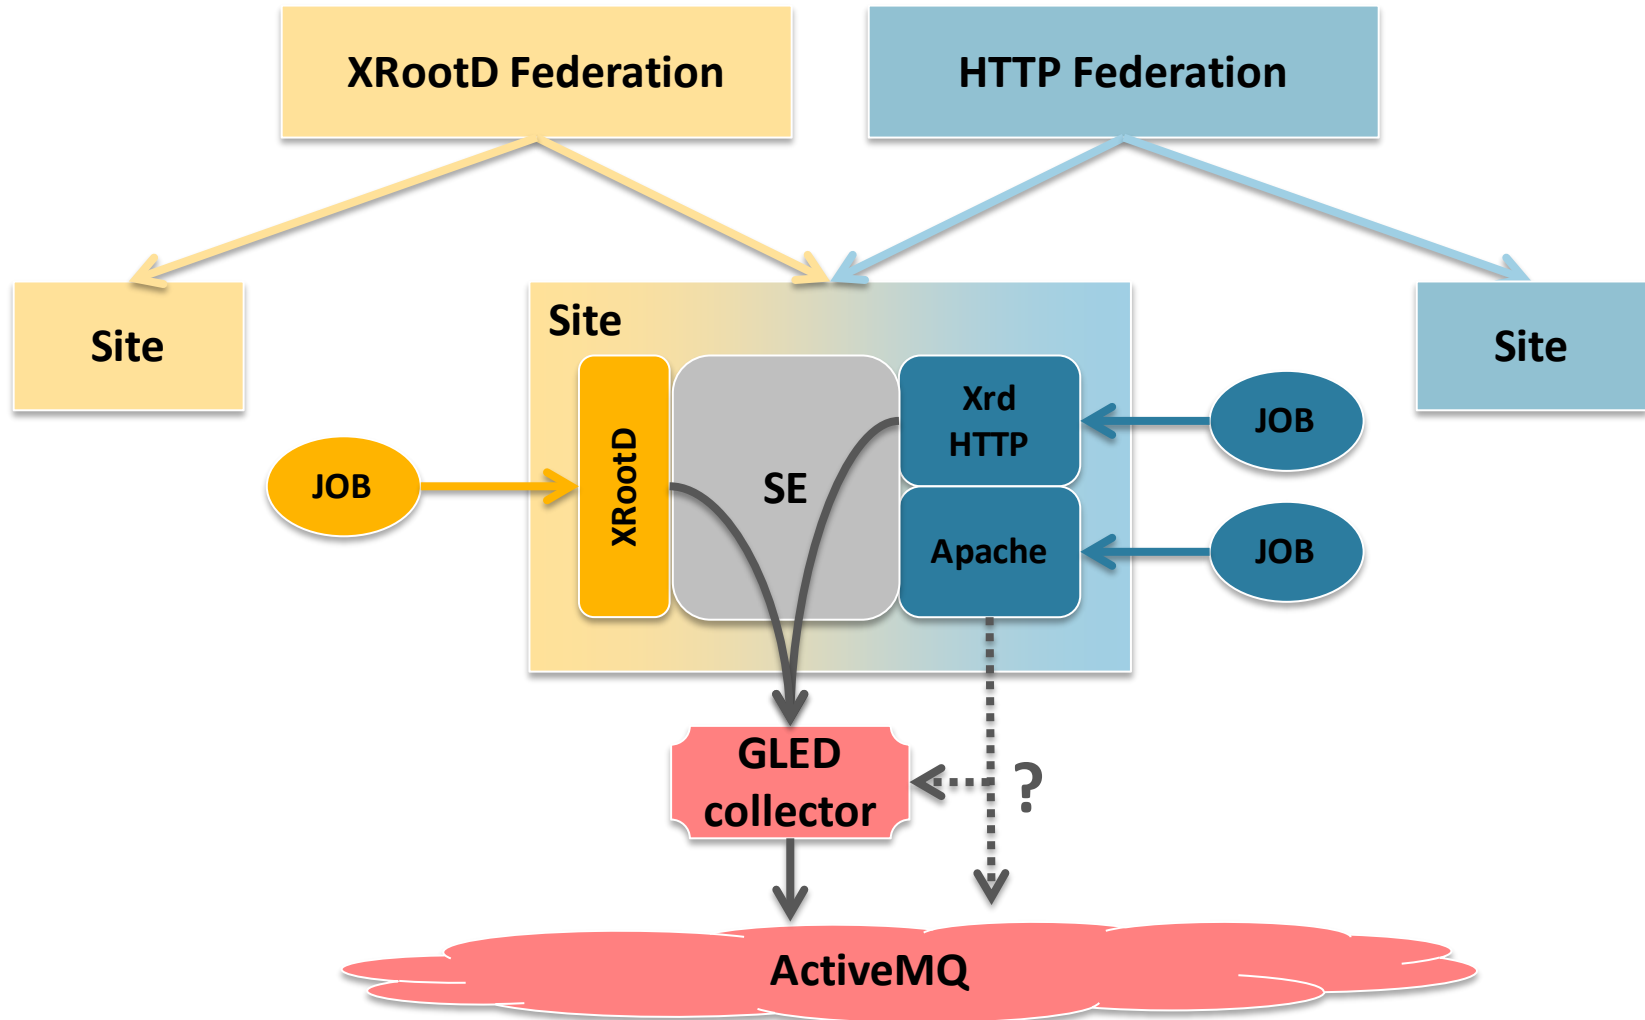
HTTP/WebDAV federation monitoring




HTTP/WebDAV federation monitoring



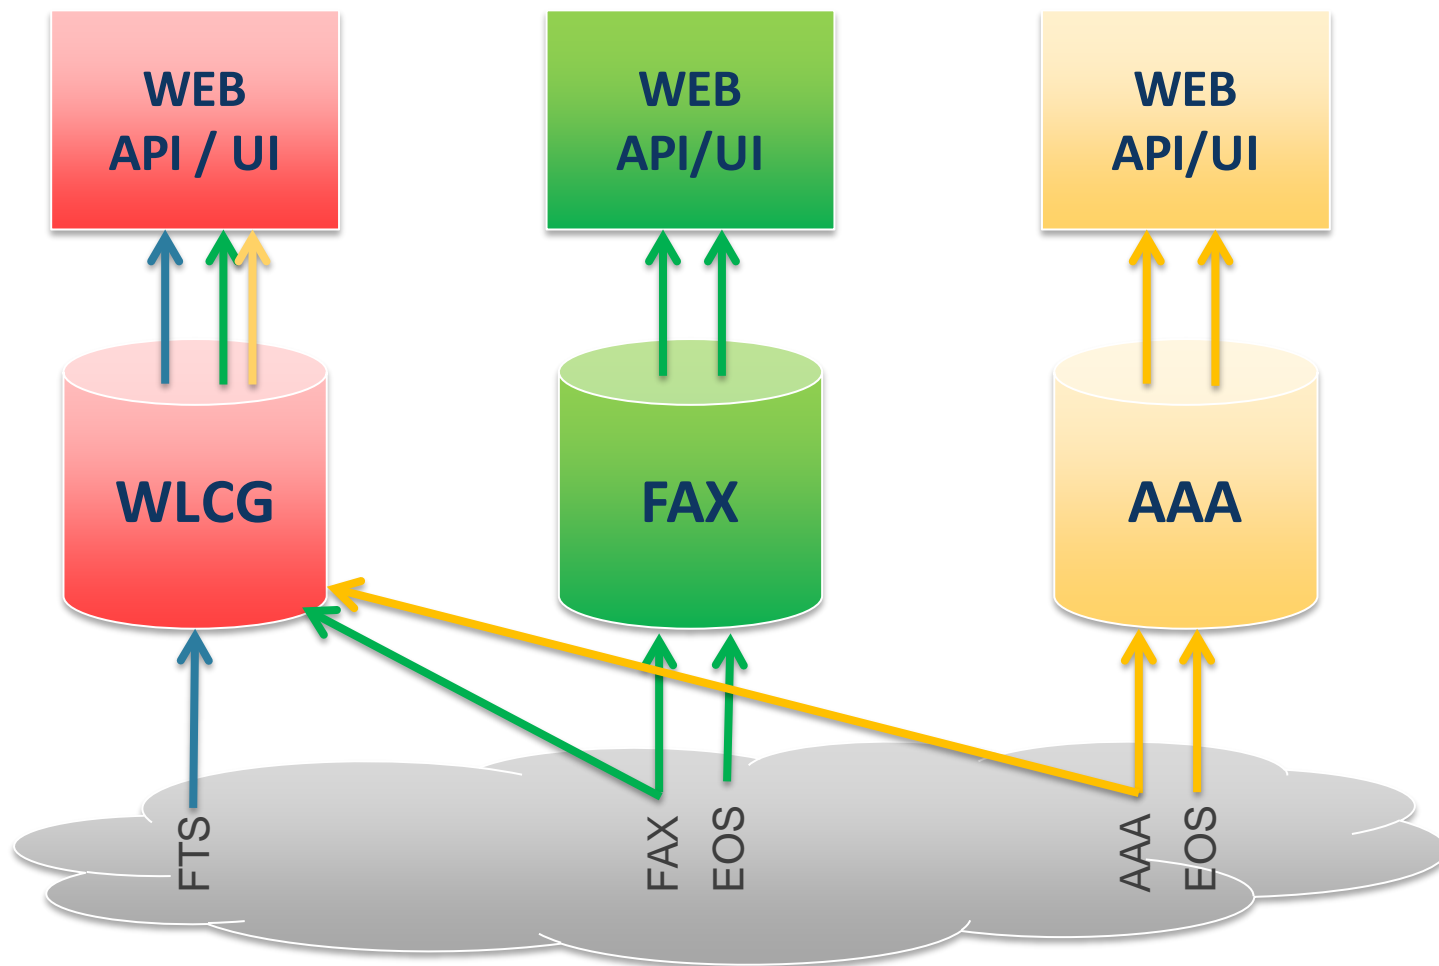
HTTP/WebDAV federation monitoring



How to compare data from different applications?

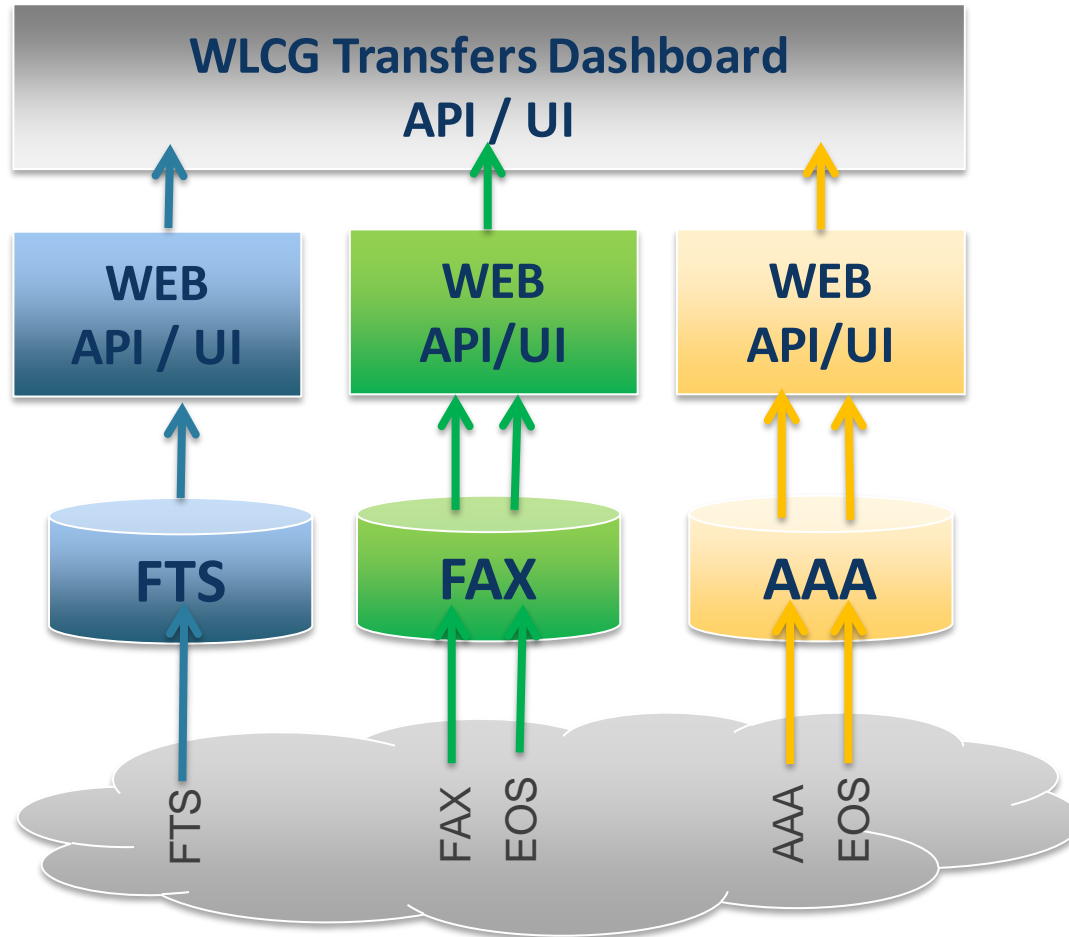


data transfers & accesses monitoring tools






WLCG Transfers Dashboard

federated approach



Some plots



Summary

- Lots of effort has been put in XRootD monitoring workflow and dashboard in the last 2 years
 - Reliable system achieved
 - Lots of use cases covered
- HTTP Monitoring already started
 - Will require a lot of effort to reach XRootD monitoring level
- New WLCG Transfers Dashboard architecture
 - Highly extensible system
 - Cross-VO or cross-technology analysis

Credits

- Andreeva Julia
- Cons Lionel
- Giordano Domenico
- Saiz Pablo
- Tadel Matevz
- Tuckett David
- Vukotic Ilija
- The AAA and FAX deployment team
-

Useful links

- AAA Dashboard
 - <http://dashb-cms-xrootd-transfers.cern.ch>
- FAX Dashboard:
 - <http://dashb-atlas-xrootd-transfers.cern.ch>
- CHEP materials
 - <https://indico.cern.ch/abstractDisplay.py?abstractId=101&confId=214784>
 - <https://indico.cern.ch/getFile.py/access?contribId=94&sessionId=6&resId=0&materialId=slides&confId=214784>
 - <https://indico.cern.ch/getFile.py/access?contribId=265&sessionId=6&resId=1&materialId=slides&confId=214784>
- Xbrowse framework:
 - <https://twiki.cern.ch/twiki/bin/view/ArdaGrid/XbrowseFramework>

Thanks for your attention

