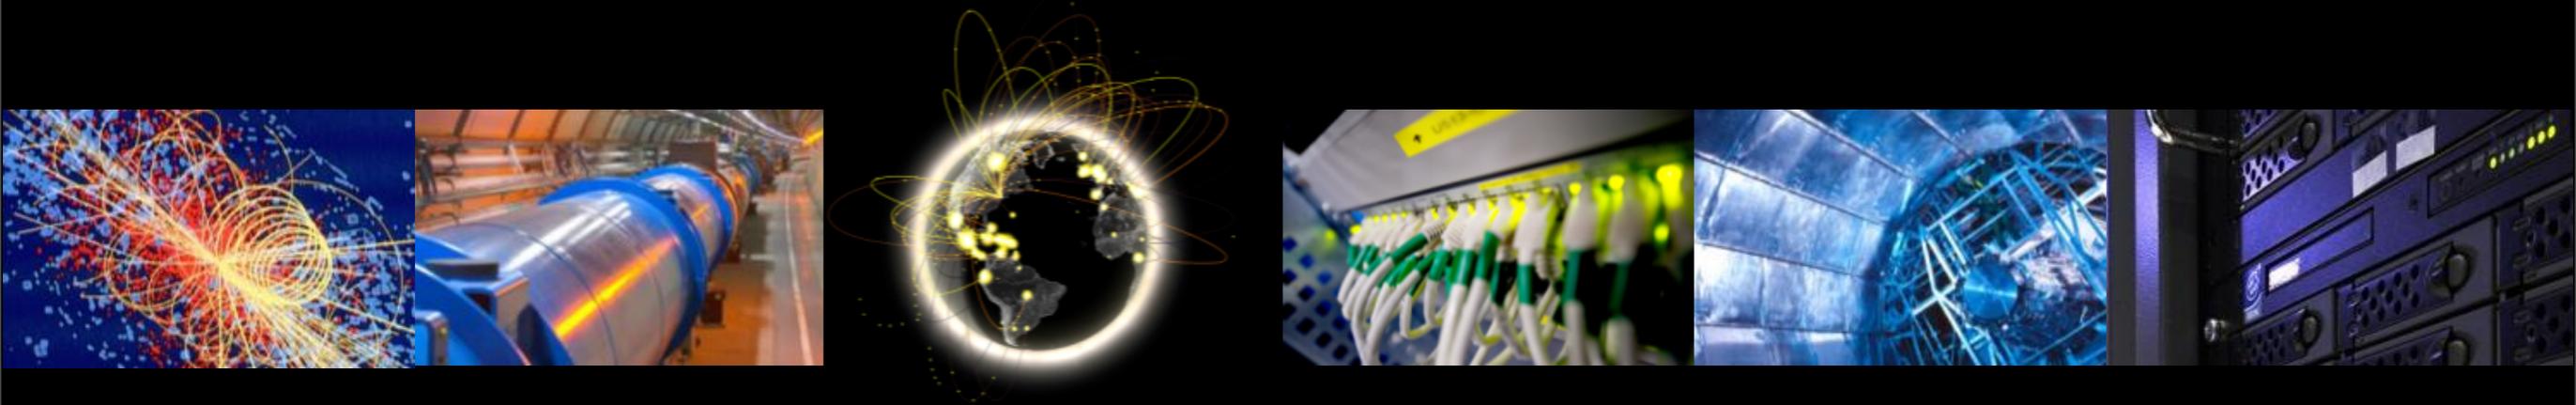


XRootD Deployment Task Force

Domenico Giordano, Simone Campana
(CERN IT-SDC)

GDB 8 May 2013



XRootD deployment TF

Main goals of this task force are

- ▶ Provide support to the XRootD deployment
 - driven internally by each project: AAA, FAX
- ▶ Liaise the monitoring efforts , collect the monitoring requirements
- ▶ Identify common needs among CMS and ATLAS to be addressed uniformly

Currently followed topics

- ▶ Monitoring
 - Workflow, plugins, configuration for dCache sites, DPM sites
- ▶ Repository for common third party plugins
- ▶ Privacy aspects of the monitoring infrastructure
- ▶ Monitor the XRootD service availability at sites

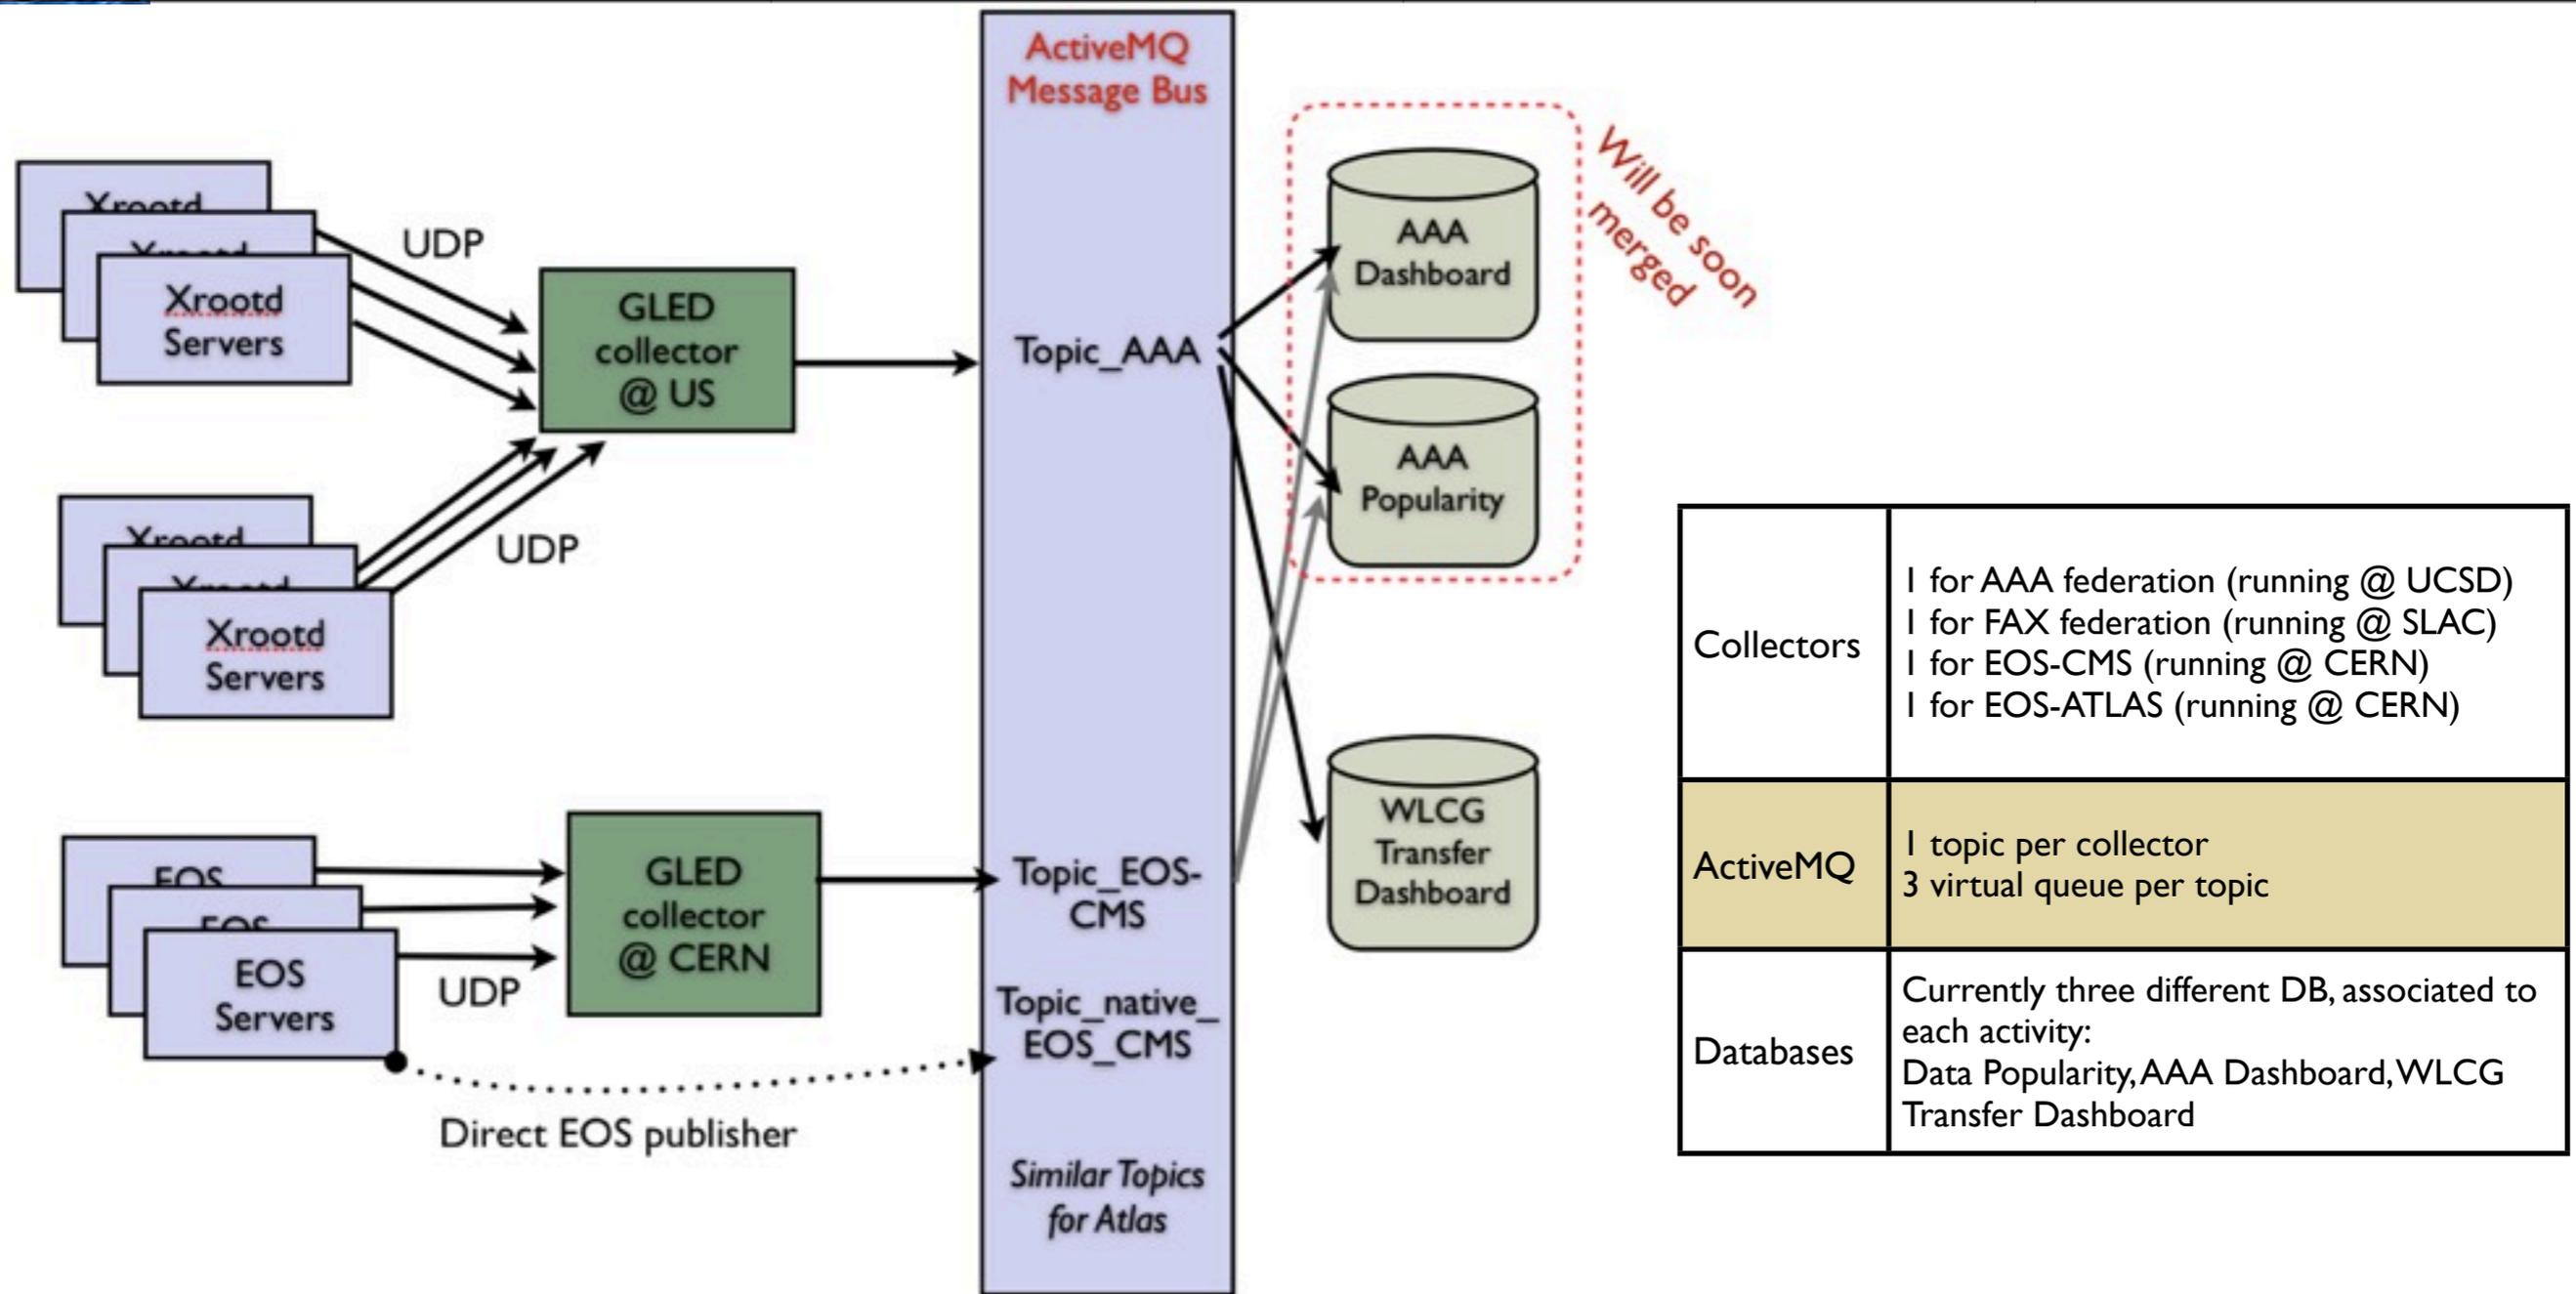
XRootD monitoring universe

XRootD monitoring infrastructure is implemented at different levels

- ▶ Two UDP streams of monitoring from XRootD
 - **Summary:** high-level statistics exposed in MonALISA
 - * # connections, #authentications, failures, in/out network traffic
 - **Detailed** (and f-stream): trace the user's activity at the level of each file open
 - * Accessed File LFN, user DN and VO, read/write operations, timestamps, client/server host & domain
 - * Requires a dedicated collector for the aggregation of single UDP packets: GLED collector (a.k.a. UCSD collector)

- ▶ In addition
 - External server health checks based on Nagios
 - SAM tests of fallback and redirections (for CMS)
 - Dedicated tools for testing, deployment tracing, performance studies

XRootD Detailed Monitoring Workflow

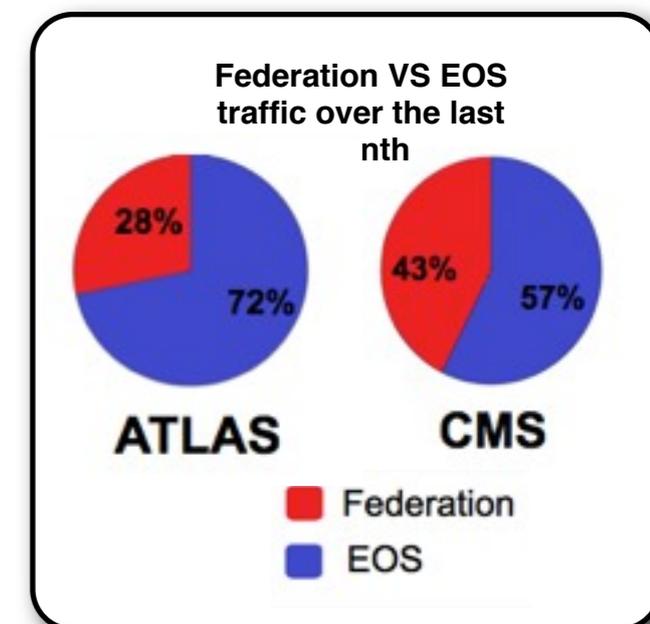


The detailed monitoring system and its complexity is growing

- ▶ Effort to unify the monitoring workflow for Dashboard and Data Popularity
 - unify the database schemas and possibly the Web UI
 - monitoring data will be kept for a period long enough to assess popularity metrics

FAX and AAA monitoring Dashboard

- ▶ Common monitoring solution for ATLAS and CMS
- ▶ Production quality data
 - Include federation and EOS traffic
- ▶ System usage growing in parallel with the federation deployment

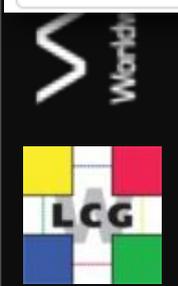
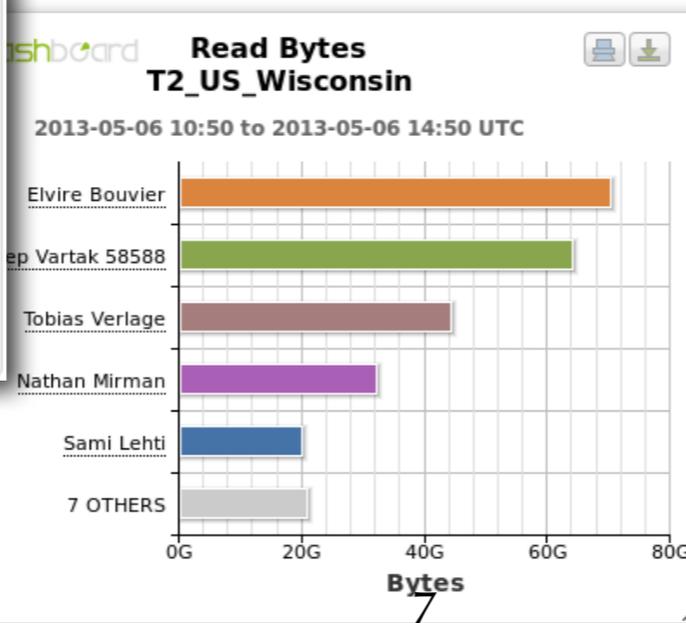
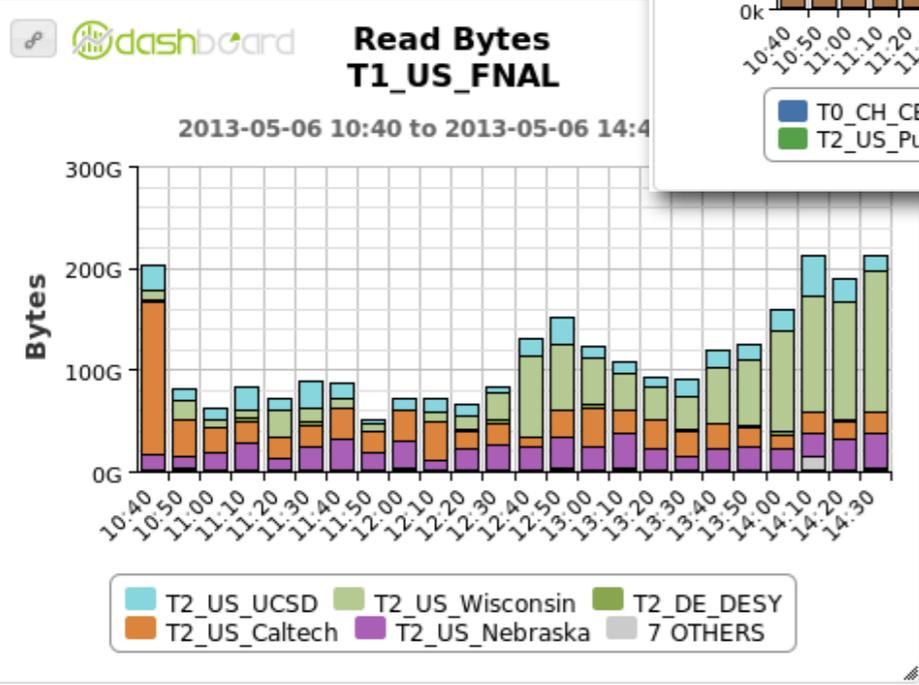
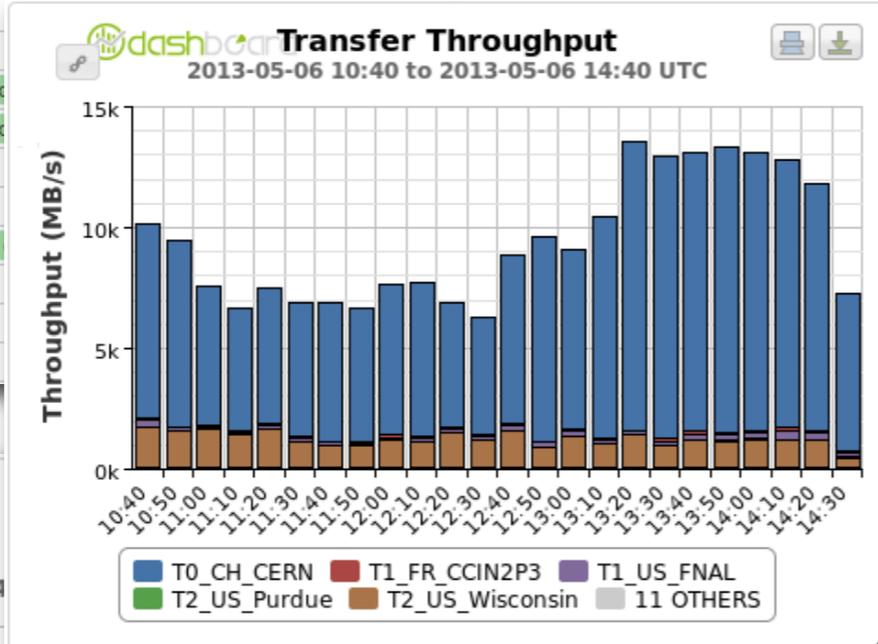


Monitoring metrics

Several levels of aggregation

- ▶ Traffic between sites
- ▶ Aggregated throughput for a given site
 - Time bin granularity from 10 mins up to day
- ▶ Server access patterns
 - breakdown by clients
- ▶ User access patterns
 - volume of data read/user
 - File lfn and fraction of read files

TOTAL	T0_CH_CERN	T1_FR_CCIN2P3	T1_IT_CNAF	T1_US_FNAL	T2_CH_CSCS	T2_IT_Bari	T2_IT_Legnaro	T2_IT_Pisa	T2_IT_Rome	
TOTAL	135 TB	114 TB	1 TB	10 GB	3 TB	47 GB	69 GB	229 MB	30 GB	47 GB
NYSGRID_CORNELL_NYS1	219 MB									
T0_CH_CERN	114 TB	114 TB								
T1_ES_PIC	496 MB					496 MB				
T1_FR_CCIN2P3	1 TB		1 TB		7 GB	550 MB		2 GB		
T1_IT_CNAF	49 GB			2 GB						
T1_UK_RAL	565 MB									
T1_US_FNAL	5 GB			2 GB						
T2_BR_SPRACE	3 GB			1 GB						
T2_CH_CSCS	612 MB									
T2_CN_Beijing	175 MB									
T2_DE_DESY	29 GB			3 GB	21 GB					
T2_DE_RWTH	295 MB									
T2_EE_Estonia	3 GB			2 GB						
T2_ES_IFCA	134 MB									



Future monitoring objectives

Provide through Dashboard a single entry point for all information relevant to XRootD federation monitor

- ▶ Integration of the summary flow in the monitoring data
 - Connections #, authentications #, authentication failures, ...
 - Expose through Dashboard the information provided by ML-based monitor

- ▶ Link to XRootD server status monitoring

- ▶ Add runtime view of the single file accesses
 - Useful for debugging and specific studies

- ▶ Adding new functionality
 - Interactive view, geographical presentation of the accesses
 - Automatic validation of Dashboard information

GLLED collectors

GLLED Collectors are a key point of the detailed monitoring workflow

- ▶ New features included in the past months (by M. Tadel)
 - Publishing messages directly in ActiveMQ
 - Accepting other UDP streams from XRootD disk servers (f-stream)

- ▶ Deep validation campaigns allowed to certify the application
 - Comparison among Summary and Detailed Monitoring
 - Comparison Vs a native collector developed by IT-DSS for EOS

Need to guarantee the steady operation of the deployed collector instances

- ▶ To be considered a critical service, up and running 24/7

WLCG Repository for third party plugins

Organize distribution and deployment of external plugins

- ▶ WLCG repository will host some packages connected to the XRootD deployment and not already associated to other repositories
 - Guarantees traceability of the packages
- ▶ Repository available at <http://linuxsoft.cern.ch/wlcg/>

Packages already identified

- ▶ GLED UDP Collector for detailed monitoring (developed by M. Tavel)
- ▶ VOMS-XRootD plugin (developed by G. Ganis)
 - Extract and validate the VOMS attributes from a proxy
 - It is meant as an add-on to the libXrdSecgsi authentication plugin
- ▶ dCache-XRootD monitoring plugin for f-stream (developed by I. Vukotic)
 - dCache sites require a dedicated monitoring plugin in order to be monitored
 - Plugin developed in the Atlas-FAX context, now adopted also by CMS dCache sites
 - * Will be improved, adding missing functionalities such as user DN identification

Open Issue: Privacy & Detailed Monitoring

Detailed monitoring allows to trace the user activity at the level of single open file

- ▶ Is there any potential concern about the user privacy?
 - Does this hold even if the user information is just stored in the monitoring DB and not exposed?

Two domains

- ▶ LHC VOs (ATLAS, CMS) rules allow monitoring of the user activities
 - Indeed it's a functionality requested by the experiments
- ▶ WLCG sites serving multi-VOs paying closer attention to the privacy
 - If disk servers are shared among multi-VOs, XRootD UDP streams will trace file accesses for all served VOs
 - In this case skimming the monitoring information from other VOs can be implemented at different levels
 - * Solution A: filter on the user_VO at the level of GLED collector
 - * Solution B: filter on the user_VO (or basepath, ..) at the site level (more work needed)

Open to discussion/feedback

Publication in GOCDDB/OIM

Two service types available in GOCDDB

- ▶ “XRootD”: XRootD entry point for the SE (including site local redirector)
- ▶ “XRootD.redirector”: regional/global redirector

Sites should publish both (if available) in order to

- ▶ Flag service downtimes to be used by VO’s automated actions
- ▶ Implement service specific SAM tests (analogous to existing SRM tests)

Summary

XRootD deployment Task Force is an initiative transversal among AAA, FAX, WLCG Operation coordination, Dashboard team

- ▶ contributes in the areas where effort can be shared
 - Monitoring, third party plugins, SAM tests, deployment instructions

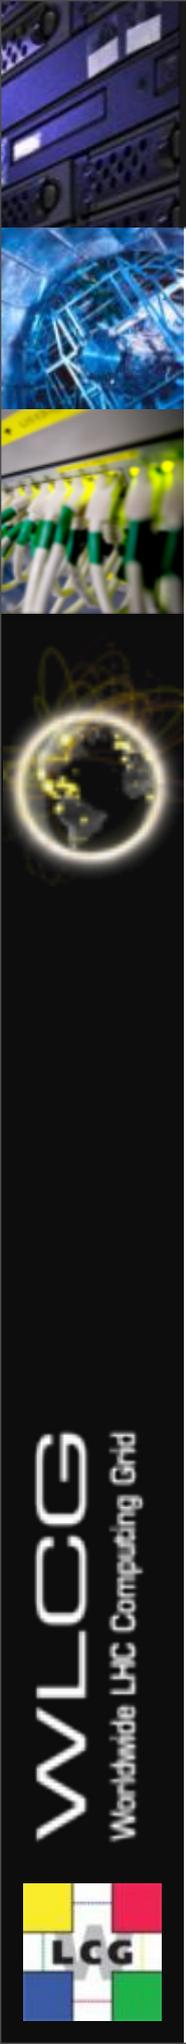
Monitoring is a crucial part of the XRootD federation projects

- ▶ Intensive WAN access is a new scenario: monitoring helps to understand and improve it
- ▶ Monitoring is in quite advanced status
 - Several level of monitoring: Service availability, Summary and Detailed XRootD streams

Confidentiality of the monitored accesses is an open question, mainly for sites serving multi-VOs

As all WLCG services, XRootD federation services need to be instrumented with dedicated service availability tests

wlcg-ops-coord-tf-xrootd@cern.ch



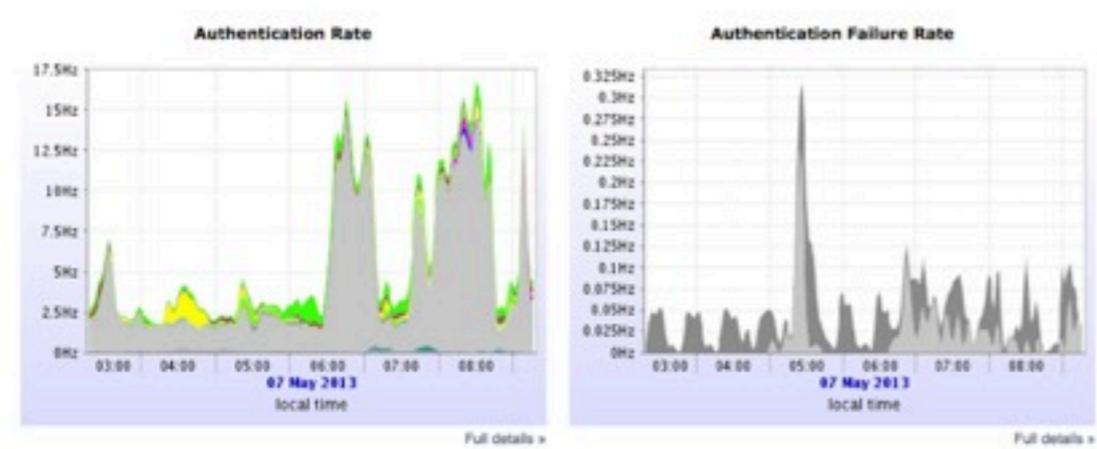
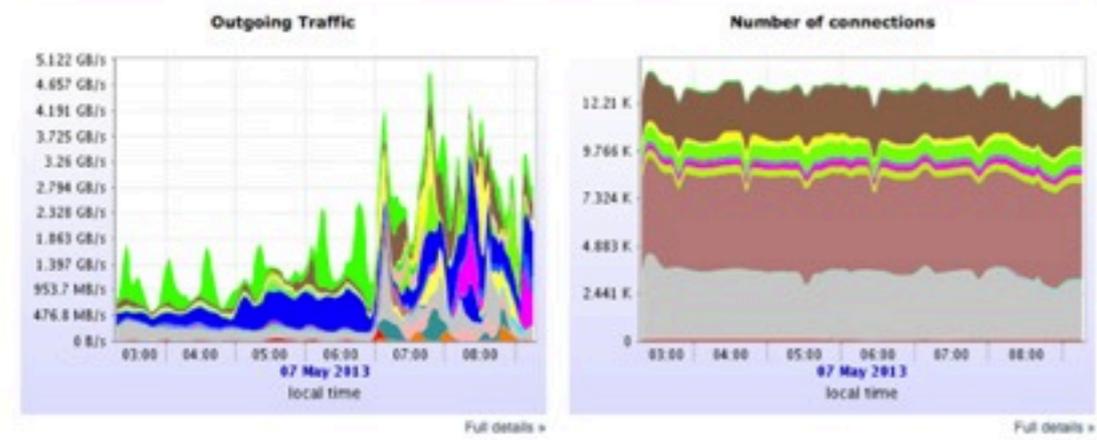
WLCG
Worldwide LHC Computing Grid



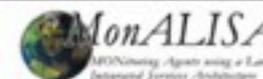
XRootD Summary monitoring

ATLAS XRootD Federation Monitoring

1 hour 3 hours 6 hours 12 hours 1 day 3 days 1 week 1 month 3 months 6 months 1 year

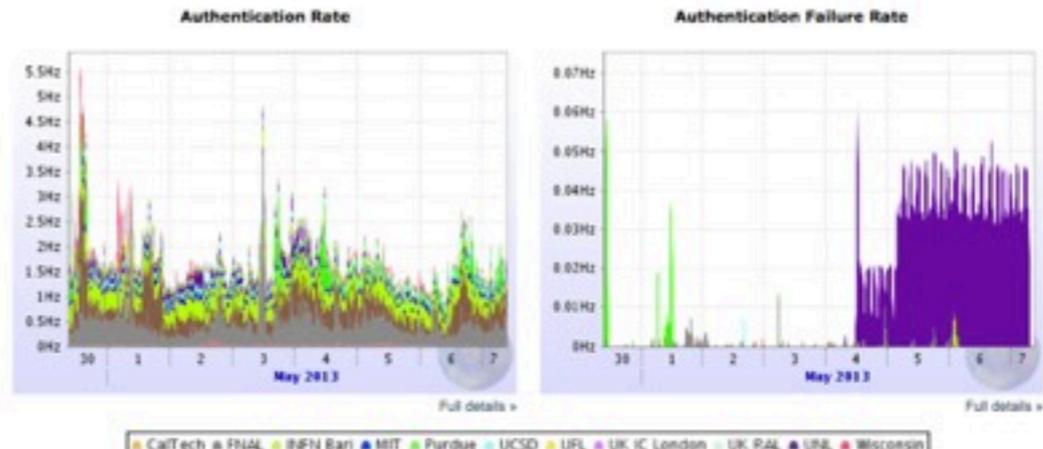
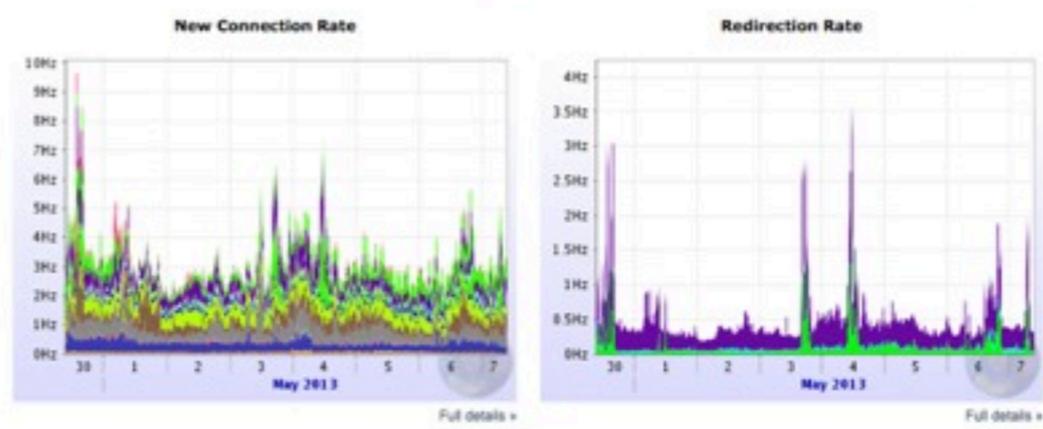
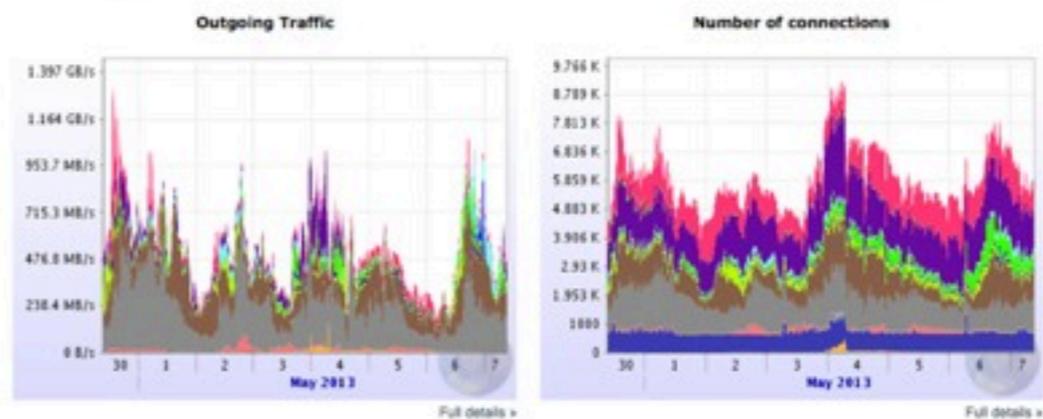


USCMS XRootD Federation Monitoring



1 hour 3 hours 6 hours 12 hours 1 day 3 days 1 week 1 month 3 months 6 months 1 year

CalTech a FNAL b INFN Bari c MIT d Purdue e UCSD f UFL g UK IC London h UK RAL i UNL j Wisconsin



CalTech a FNAL b INFN Bari c MIT d Purdue e UCSD f UFL g UK IC London h UK RAL i UNL j Wisconsin

Running @ SLAC
D. Giordano

Running @ UCSD

