

# WLCG Monitoring

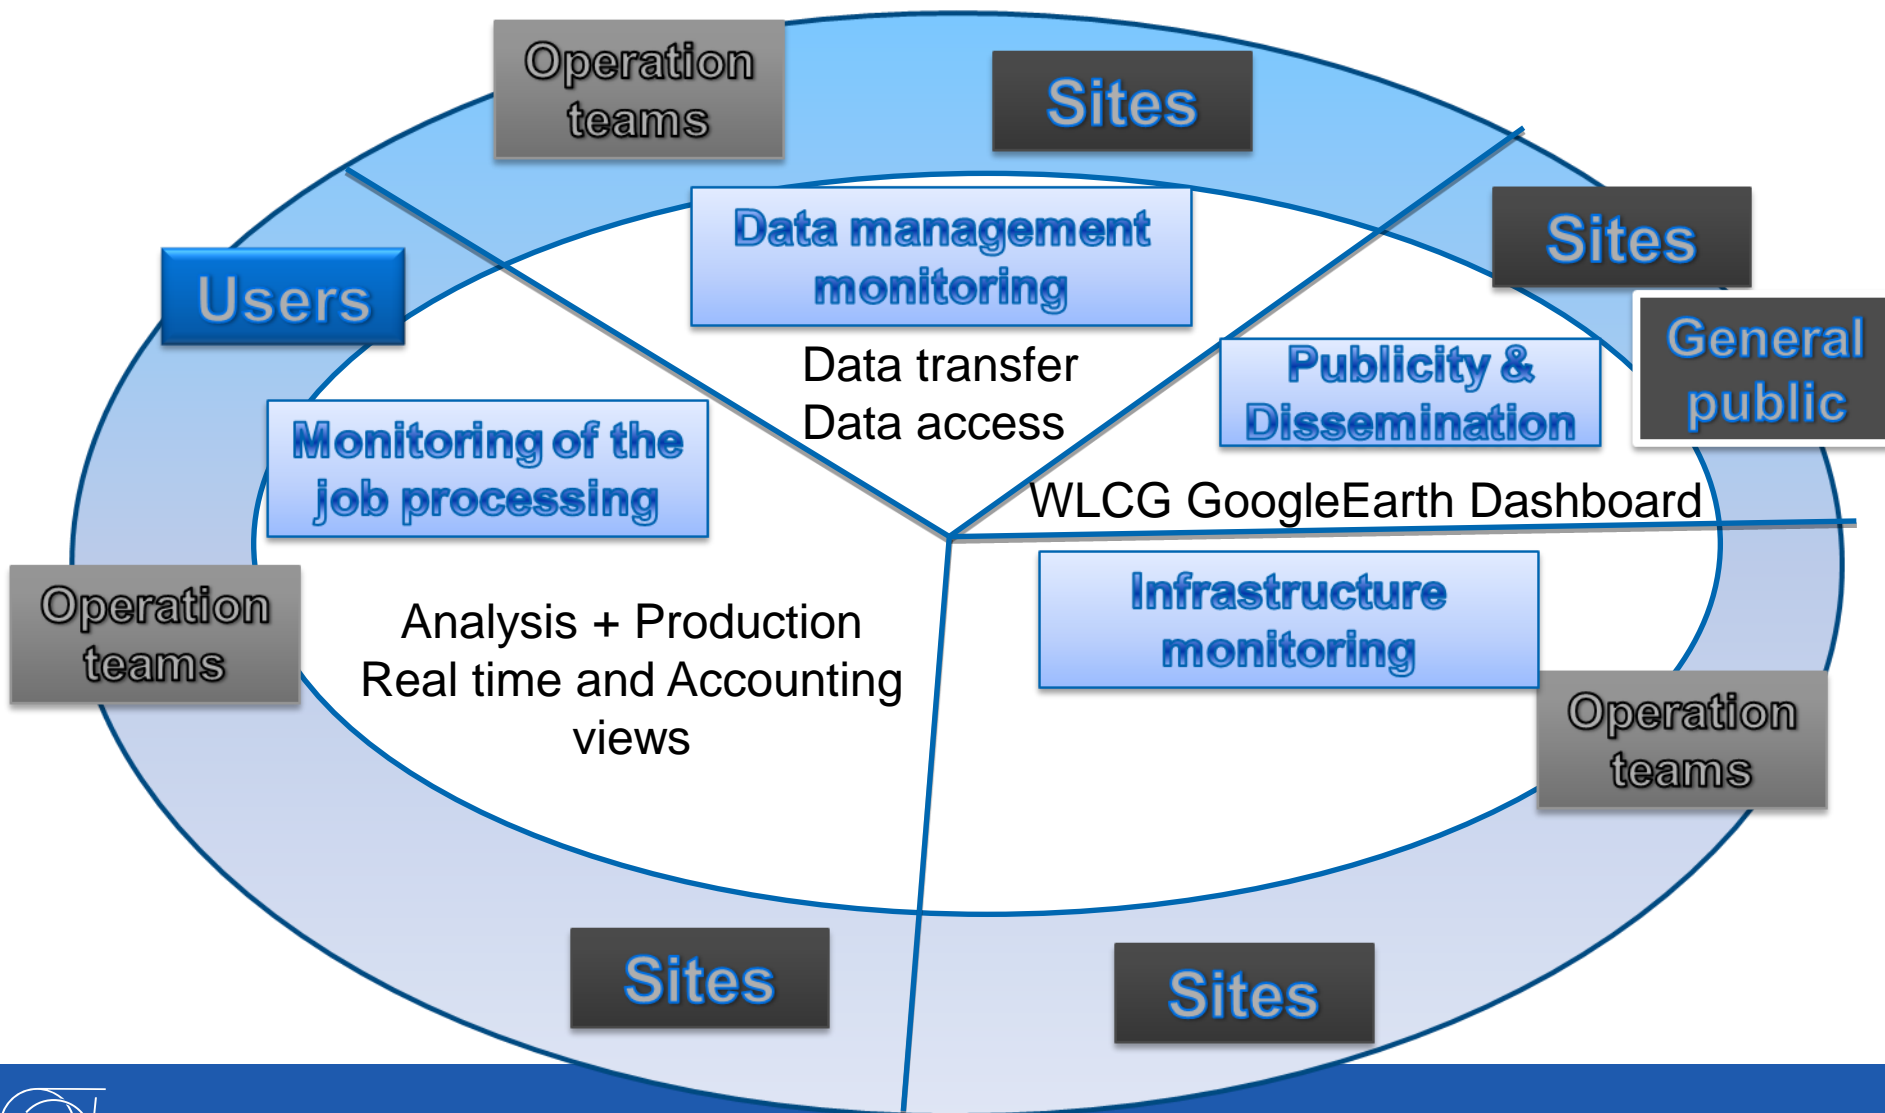
Julia Andreeva  
CERN IT



# Outline

- WLCG Monitoring Consolidation project: goals and outcome
- WLCG Monitoring: future evolution

# WLCG Monitoring landscape



# Activity monitoring vs remote testing

- Activity monitoring
  - Gather real time information about various activities running at the site : data transfer, data access, job processing
- Remote testing (Functional and stress testing)
  - Functional tests check that the site/service can perform a required task and to determine eventual failure modes for further debugging. Performed with SAM and HammerCloud
  - Stress tests assess the ability of a site or service to work reliably under regular operational load and even beyond it. Performed by HammerCloud
- Complimentary to each other
  - Activity monitoring identifies failing activities at a site
  - Remote testing helps to identify the exact reason of failure

# Goals of the WLCG Monitoring Consolidation project

Main goal was to reduce effort required for the development and support of the WLCG monitoring infrastructure

How it could be achieved

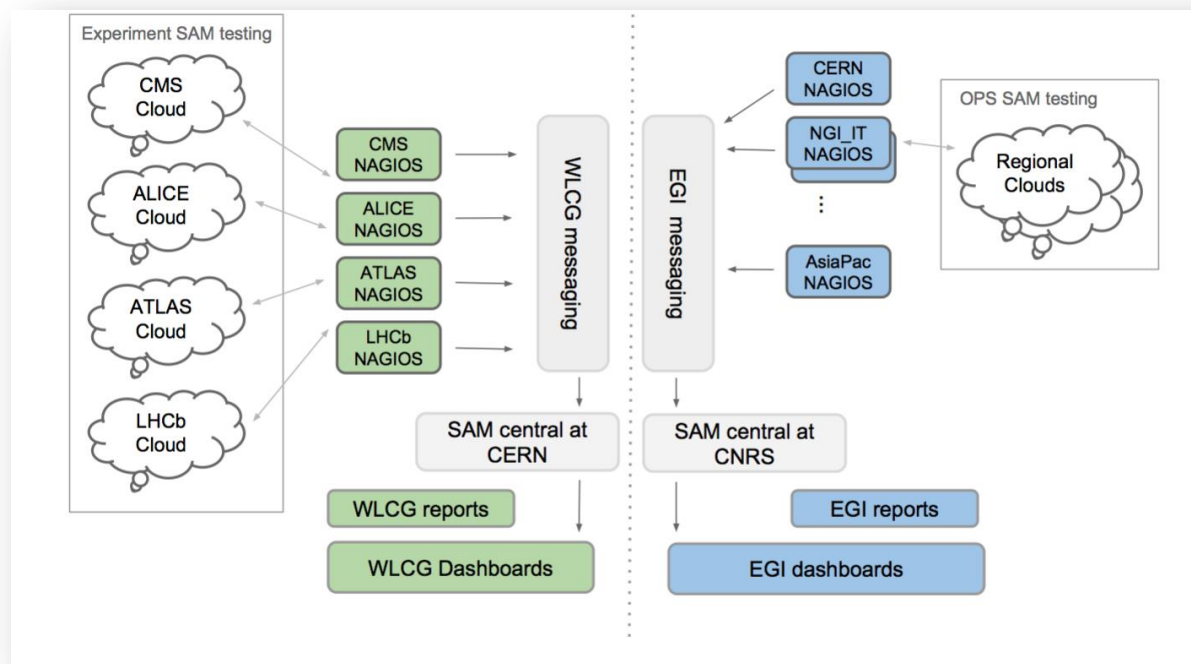
- Focus on the WLCG scope
- Reduce complexity
- Simplify operations, support and service
- Common development and deployment
- Unify where possible the components
- Enforce modular design of all monitoring systems

# WLCG Monitoring Consolidation project. Organization of work

- 3 months evaluation + 3 months design + 9 months implementation
- Working group led by Pablo Saiz consisted of the WLCG monitoring developers, representatives of sites, experiments, WLCG operations and Agile Infrastructure monitoring team
- Small taskforces on the dedicated subjects
- Project progress tracked in JIRA

# Focusing on the WLCG scope

- SAM2 was developed following the philosophy of EGI
- Highly distributed system running at different regional National Grid Initiatives (NGIs)



The change was transparent for the WLCG user community  
Helped the EGI partners in the migration of the central services  
Credits to Marian Babik and Lionel Cons

- EGI took over EGI SAM in Spring 2014
- SAM3 focuses on the WLCG scope

# Improvements to operations/deployment

- The WLCG Monitoring Cluster migrated to OpenStack. 115 hosts of various types puppetized (web portals, services, development nodes...)
- Accomplished migration to SLC6
- Migrated to a new set of release and software configuration management tools (CERN Koji, Git, JIRA)
- Reviewed, improved and documented support units workflow (using GGUS, SNOW, Roger, e-groups)

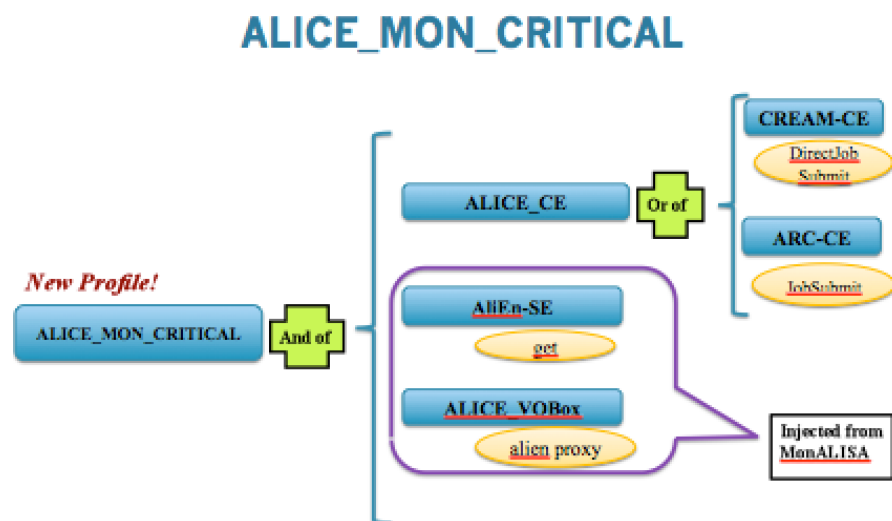


# Applying common solutions

- Extended the Site Status Board (SSB) framework to cover additional use-cases
  - SAM metric store and availability calculation was re-implemented using SSB framework
  - REBUS is being re-implemented in the SSB framework
  - SSB is actively used by the experiments, WLCG operations team. New use-cases GLUE validation, cloud accounting
  - Re-using SSB framework allowed to decrease variety of applications we need to support

# Improving functionality

- Enabled new submission methods for SAM tests (direct CREAM CE submission, Condor-G submission, ARC CE submission is under development)
- Improved flexibility for availability calculation.
- Experiments obtained more power (and responsibility) :
  - In definition of sites and services
  - Definition of profiles
  - Injecting their own metrics
  - Recalculating (overwriting) metrics
- Site admins have a possibility to integrate SAM results in their local Nagios instances.



More details can be found in the [presentation](#) of Rocio Rama at the March GDB

# Outcome of the WLCG Monitoring Consolidation project

- Decoupled WLCG and EGI SAM by migrating responsibility of SAM Central Services to EGI partners
- Considerably improved service deployment/support
- Migrated to a new architecture with common metric store and new data processing engine with variety of UIs tailored to different needs
- Improved the flexibility and functionality of the infrastructure monitoring system (SAM3)
- Enforced modular design of all monitoring components
- Support effort for the infrastructure monitoring tools dropped from 6 FTEs as of summer 2013 to ~1 FTE by the end of the project (end of 2014)
- More details can be found in the CHEP poster “WLCG Monitoring Consolidation”



# WLCG monitoring evolution

Watch what is going on

Follow trends, commissioning of the new services. Identify inefficiencies and take corrective actions

Analyze monitoring data in order to understand trends and correlations and to predict future outcomes. Therefore keep improving the infrastructure



Passive



Active




Proactive

# Active monitoring phase



- Examples of how monitoring helps to correct inefficiencies and to prove improvements:
  - ATLAS blacklisting of faulty sites
  - CMS site readiness
  - Commissioning of the new FTS functionality

## FTS optimiser commissioning



See more details in the CHEP presentation of Oliver Keeble [“FTS3 – quantitative monitoring”](#)

# Proactive monitoring

- Common effort of the LHC computing community
- The task of the WLCG monitoring team is to collect complete and reliable data and to make it available for analysis
- IT Data Analytics project stimulated this activity
- WLCG monitoring data is available for analysis in the Hadoop cluster. APIs to retrieve data from the ORACLE DB are also in place





# Other challenges ahead of us


- Need to ensure that the monitoring infrastructure can cope with the constant increase of monitoring data volume and complexity
- New areas to cover: new data transfer protocols (http), new dynamic types of resource providers – cloud computing





# Growing data volume

- Steadily growing amount of data which has to be processed, archived and served through UI/API
  - Just for the ATLAS DDM Dashboard ~ 1TB of data (size of ORACLE tables) collected per year
- Data rates are changing over time, should be ready to handle spikes
  - Up to 1000 Hz update rate for ATLAS EOS
- Time to time need to re-process big volumes of the historical data



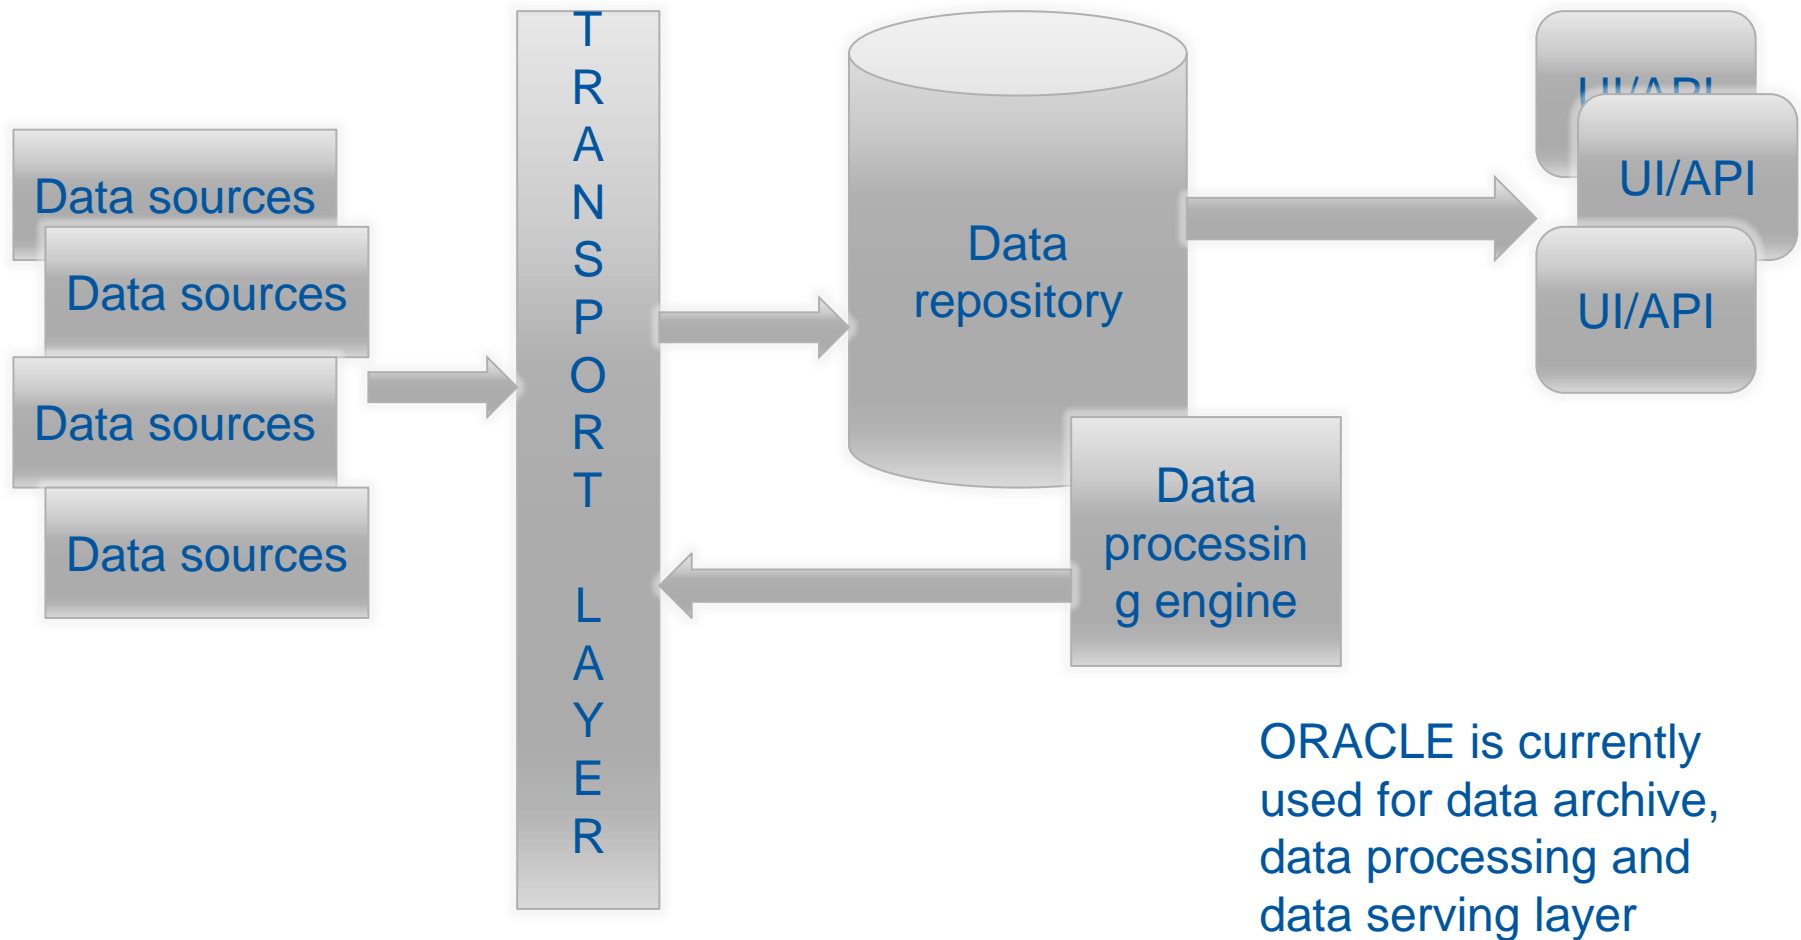
Number of jobs  
processed by ATLAS  
increased 3 times  
compared to 2010

# Are We confident our system will scale in future?

- We have a working system which can handle the load
- However, already today we experience some bottlenecks
- We better look for a new arm now, not waiting for an end of the battle



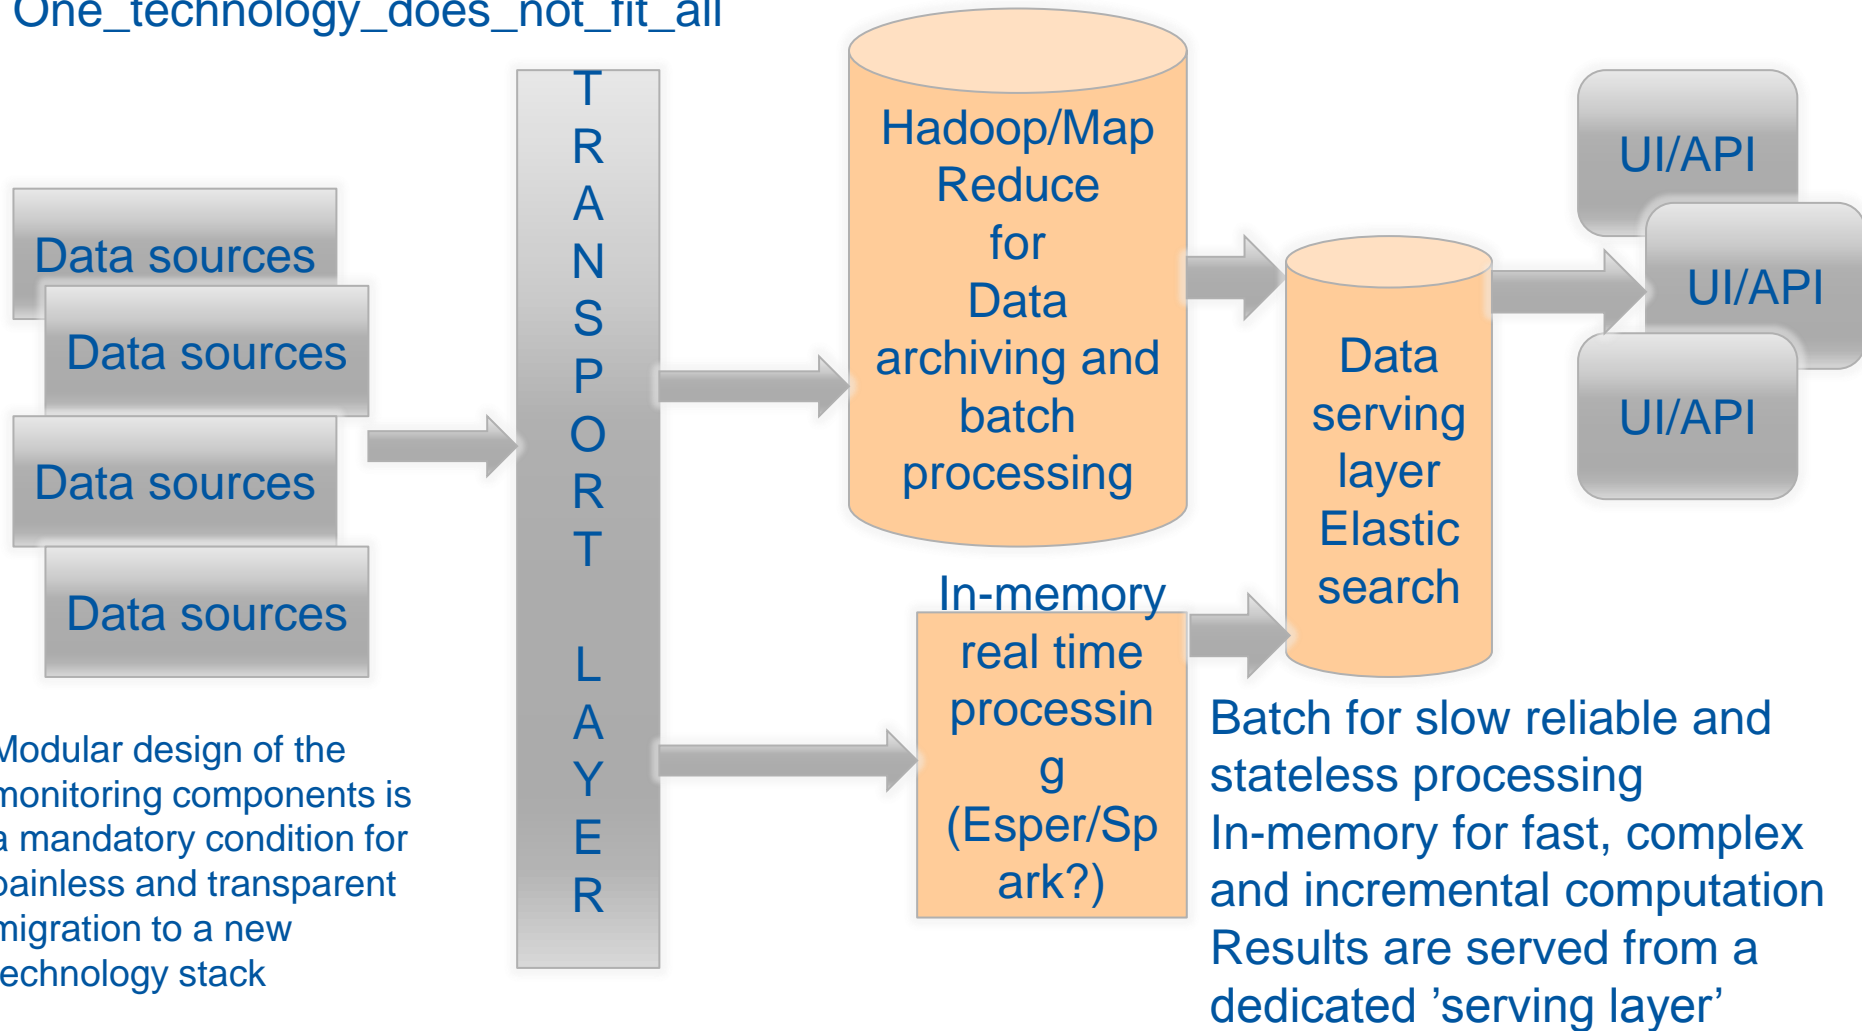
# Current architecture



# Architecture evolution


Taking inspiration from Lambda architecture

One\_technology\_does\_not\_fit\_all



# Status of migration to a new technology stack (1)

- As a pioneer application took the most challenging one – xrootd/eos data access and data traffic monitoring
- Re-implemented data processing algorithms using MapReduce (MR). Demonstrated big performance gain compared to a current sql implementation of data processing
- The prototype with a new data flow is ready and validated. Will be deployed in production by June 2015
- More details see in the CHEP presentation of Luca Magnoni “Monitoring WLCG with the *lambda* architecture”



Processing of 3 weeks of statistics takes couple of minutes

# Status of migration to a new technology stack (2)

- In parallel perform evaluation of the Elasticsearch to be used as a data serving layer
- Version of Elasticsearch released this year enabled multi-field grouping which is required for the WLCG use-cases
- Very promising results so far
- Plan to prototype a new workflow which uses Elasticsearch for serving data by the end of summer this year
- Moreover, the lambda architecture can be further simplified thanks to new tools as Apache Spark, which allows to use the same processing framework and similar APIs for both batch and streaming. Spark evaluation is on going.

# Conclusions

- Success of the WLCG Monitoring Consolidation project allowed to substantially decrease operational effort and created conditions for migration of the WLCG monitoring system to a new technology stack
- The goal of the current work is to ensure that the next generation of the monitoring infrastructure provides cost-effective, scalable and performant solution which can cope with steady growing volume and complexity of the monitoring data

