



WLCG Monitoring Consolidation

ALICE Physics Analysis and T1-T2 Workshop

6th March 2014

Pablo Saiz
IT/SDC

Presented by Maarten Litmaath
IT/SDC



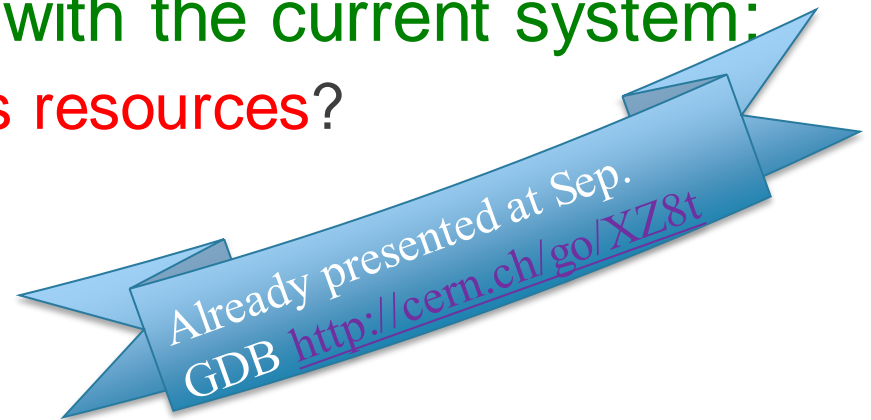
IT-SDC : Support for Distributed Computing

Table of contents

- WLCG Monitoring consolidation:
 - Definition of the project
 - How to do it
 - Current status
 - Impact on the experiments
 - Summary

WLCG monitoring consolidation group

- Main goals:
 - Reduce complexity, modular design
 - Simplify operations, support and service
 - Common development and deployment
 - Unify, where possible, components
- We know we can monitor with the current system:
 - How can we do it with **less resources**?
- <http://go.cern.ch/B6NS>



Organization of work

- **Small core team:**
 - Monitoring development teams
 - LHC experiment representatives
 - WLCG operations representative
 - Agile Infrastructure Monitoring representative
- Even smaller taskforces on **dedicated subjects**
- Mailing list: `wlcg-mon-consolidation@cern.ch`
- Fortnightly meetings with summary reports
 - <http://go.cern.ch/6XQQ>

Timeline

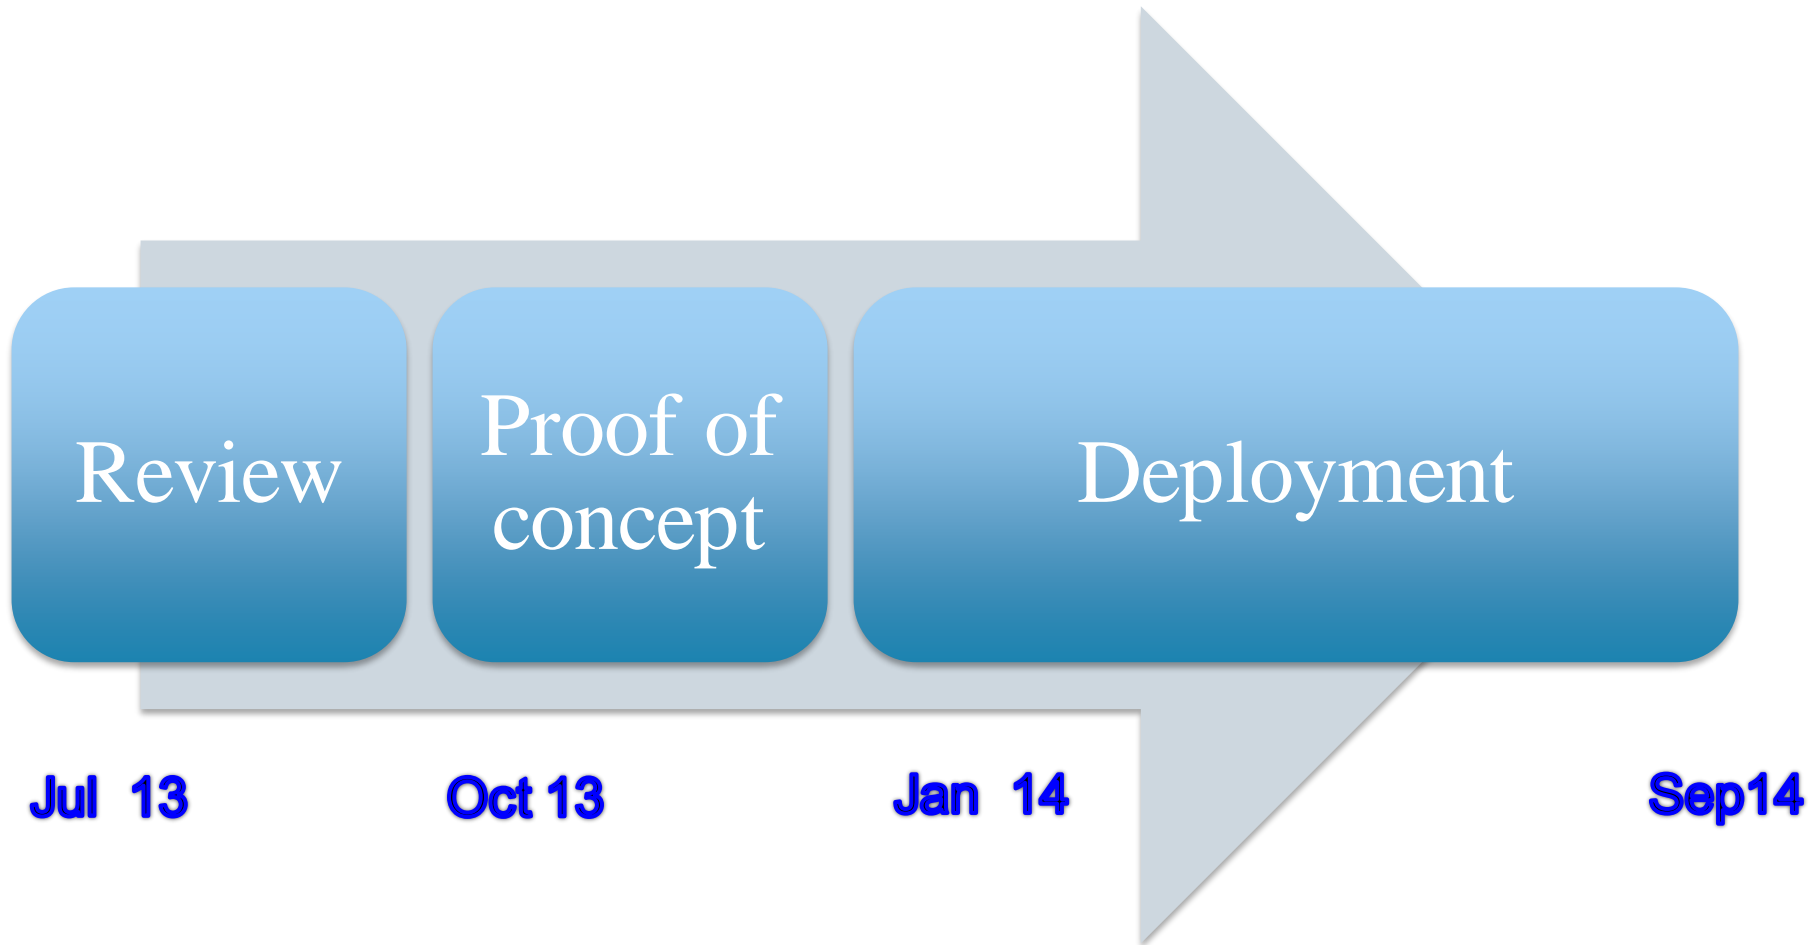


Table of contents

- WLCG Monitoring consolidation:
 - Definition of the project
 - How to do it
 - Current status
 - Impact on the experiments
 - Summary

How to reduce complexity

1. Remove applications/functionality
2. Reduce scope
3. Modular design
4. Take existing solutions
5. Merging applications

1. Identify applications no longer needed

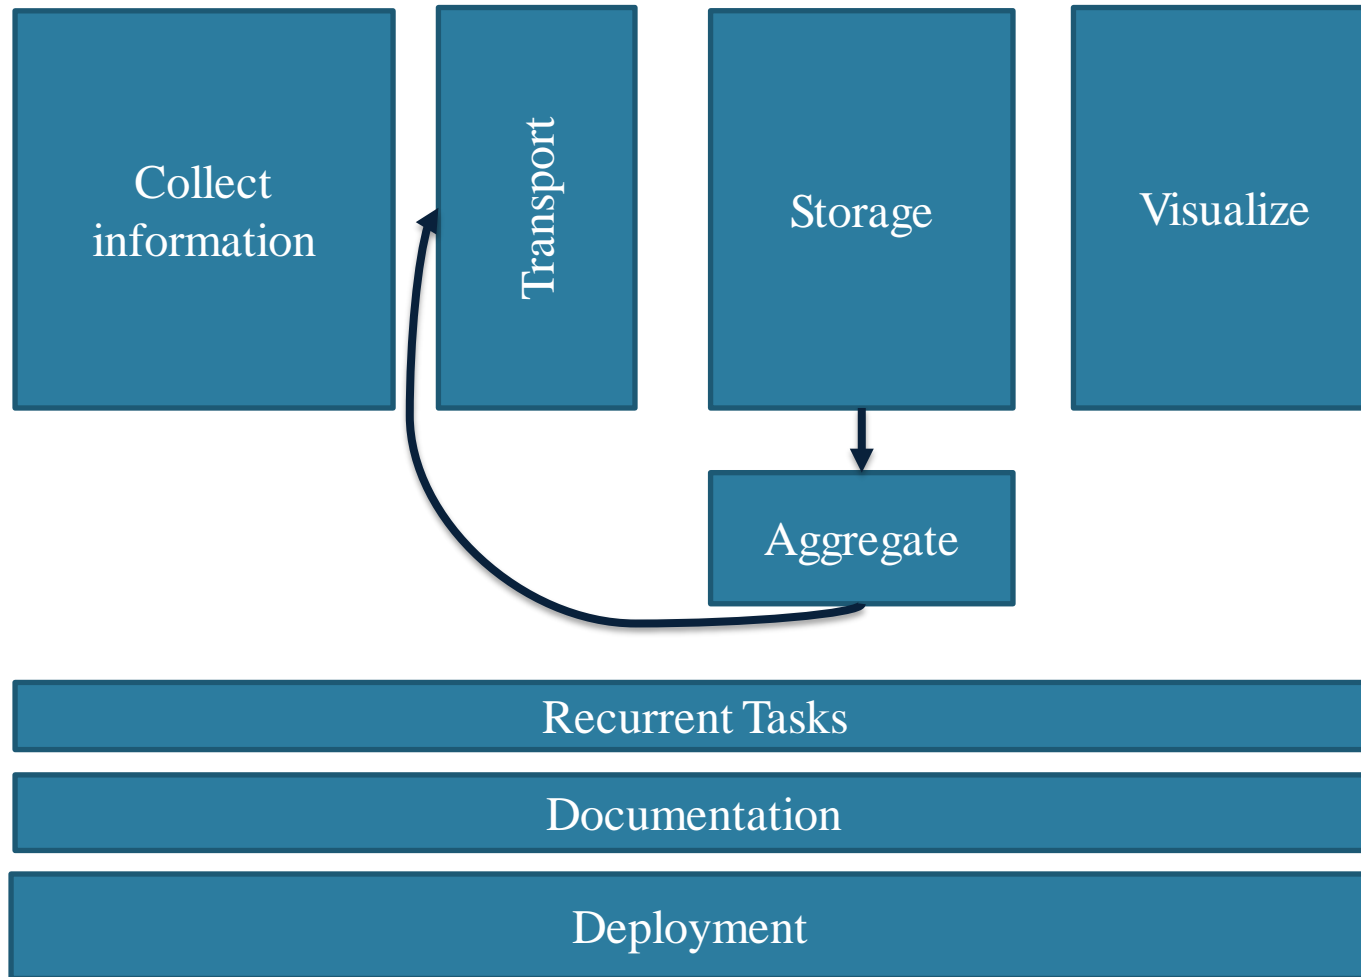
- **Used**
- **nice to have**
- **not used**

Category	ALICE	ATLAS	CMS	LHCb	Comments
Job Monitoring					
Current view		ATLAS Job View	ATLAS Job View		
Accounting		Task Monitoring	Task Monitoring	Task Monitoring on Android	
WLCG job trends		WLCG Job Trends			
Data Management					
Current view		DM	DM		
Accounting		DM Accounting			
Transfers		WLCG Transfers Dashboard			
	ATLAS Data Transfer			DM	
Site/Service Monitoring					
VO Feed	ATLAS VO Feed		ATLAS VO Feed		
SSB	ATLAS SSB	ATLAS SSB	ATLAS SSB	LHCb SSB	
SUM	ATLAS SUM	ATLAS SUM	ATLAS SUM	ATLAS SUM	
Regional/Experiments SAM/Nagios	ATLAS Nagios	ATLAS Nagios	ATLAS Nagios	ATLAS Nagios	Other 35 regional instances (NGIs, ROCs, CERN)
VO SAM/Nagios		MIDMON			Middleware Nagios, Security Nagios, former gLExec Nagios, SuperB Nagios, etc.
SAM Operational Monitoring (ops-monitor)		OPS-MONITOR			Monitors sam-atlas/alice/lhcb/cms prod and preprod SAM, SAM central, etc.
SAM central service		CMS SAM GridMon			Web API, myWLCG AR,
SAM WLCG Reports		ATLAS CMS Monthly reports			
SAM A/R Trends		AR Trends			
SAM T0/I SiteView		ATLAS SiteView			
SAM Tree Map		ATLAS			
SAM Site Nagios		SAM Nagios Installation			
Critical Services		ATLAS might be interested (the application has not been setup)	ATLAS Critical Services		
SAM Probe Development Framework and SAM probes		Probe Development Documentation			
Personalized Dashboard		Personalized dashboard			
Dissemination					
Google Earth		Web interface to Google Earth			Also installed at Globe, CC, CMS Centre Meyrin
Siteview		CMS ATLAS from site admin POV (tbody) SiteView			
Development Management					
SAM Nightly Validation Framework		ATLAS Nightly Validation Framework			Used to run sam-* dev boxes and CI

2. Reduce scope

- **Concentrate on WLCG**
 - From distributed to centralized deployment
 - Moving SAM central service to France/Greece/Croatia
- **Removing 'OPS' Virtual Organization**
 - From January, WLCG reports based on experiment metrics
 - Experiments do not depend on OPS test

3. Modular design



4a. Deployment

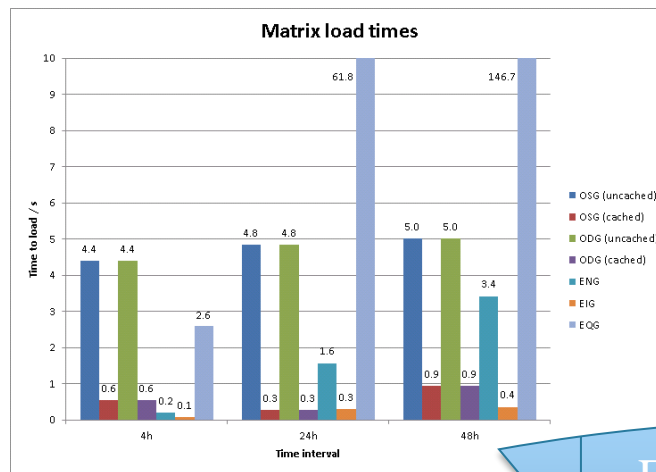
- Using **openstack**, puppet, hiera, foreman
 - Quota of 100 nodes, 240 cores
 - 28 hostgroups already created
 - Development machine (8 nodes)
 - Web servers (SSB, xrootd, WLCG transfers, Job: 40 nodes)
 - Elastic Search (6 nodes), Hadoop (4 nodes)
 - **Nagios** with info from vofeed.
- Migrating machines from quattor to AI
- **Koji** and **Bamboo** for build system and continuous integration

Documentation

Deployment

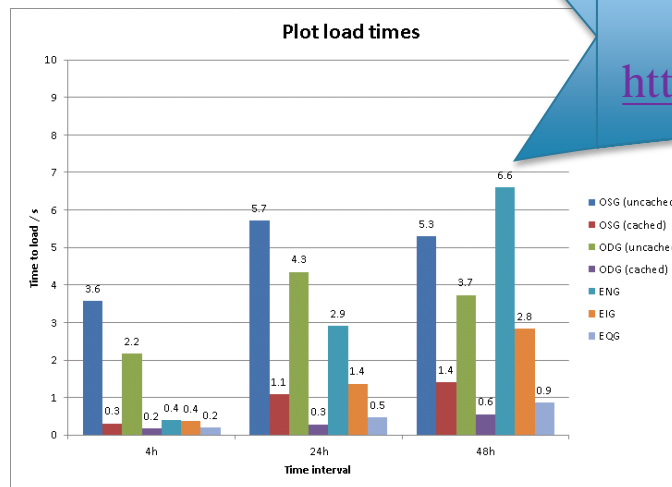
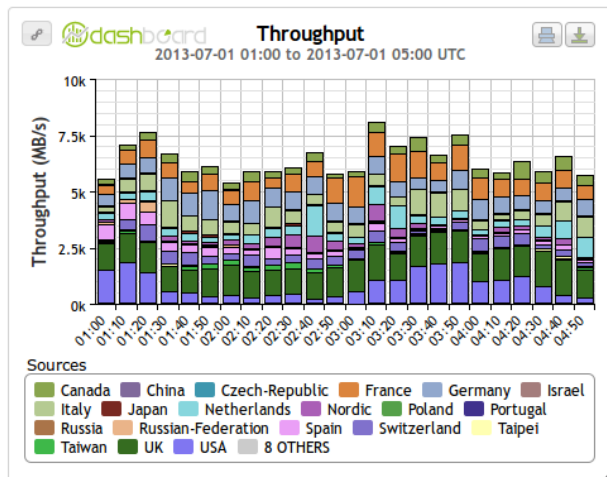
For details on other layers, see:
<http://cern.ch/go/Bp9j>

4b. Storage evaluation: Elasticsearch



For details on other layers, see:

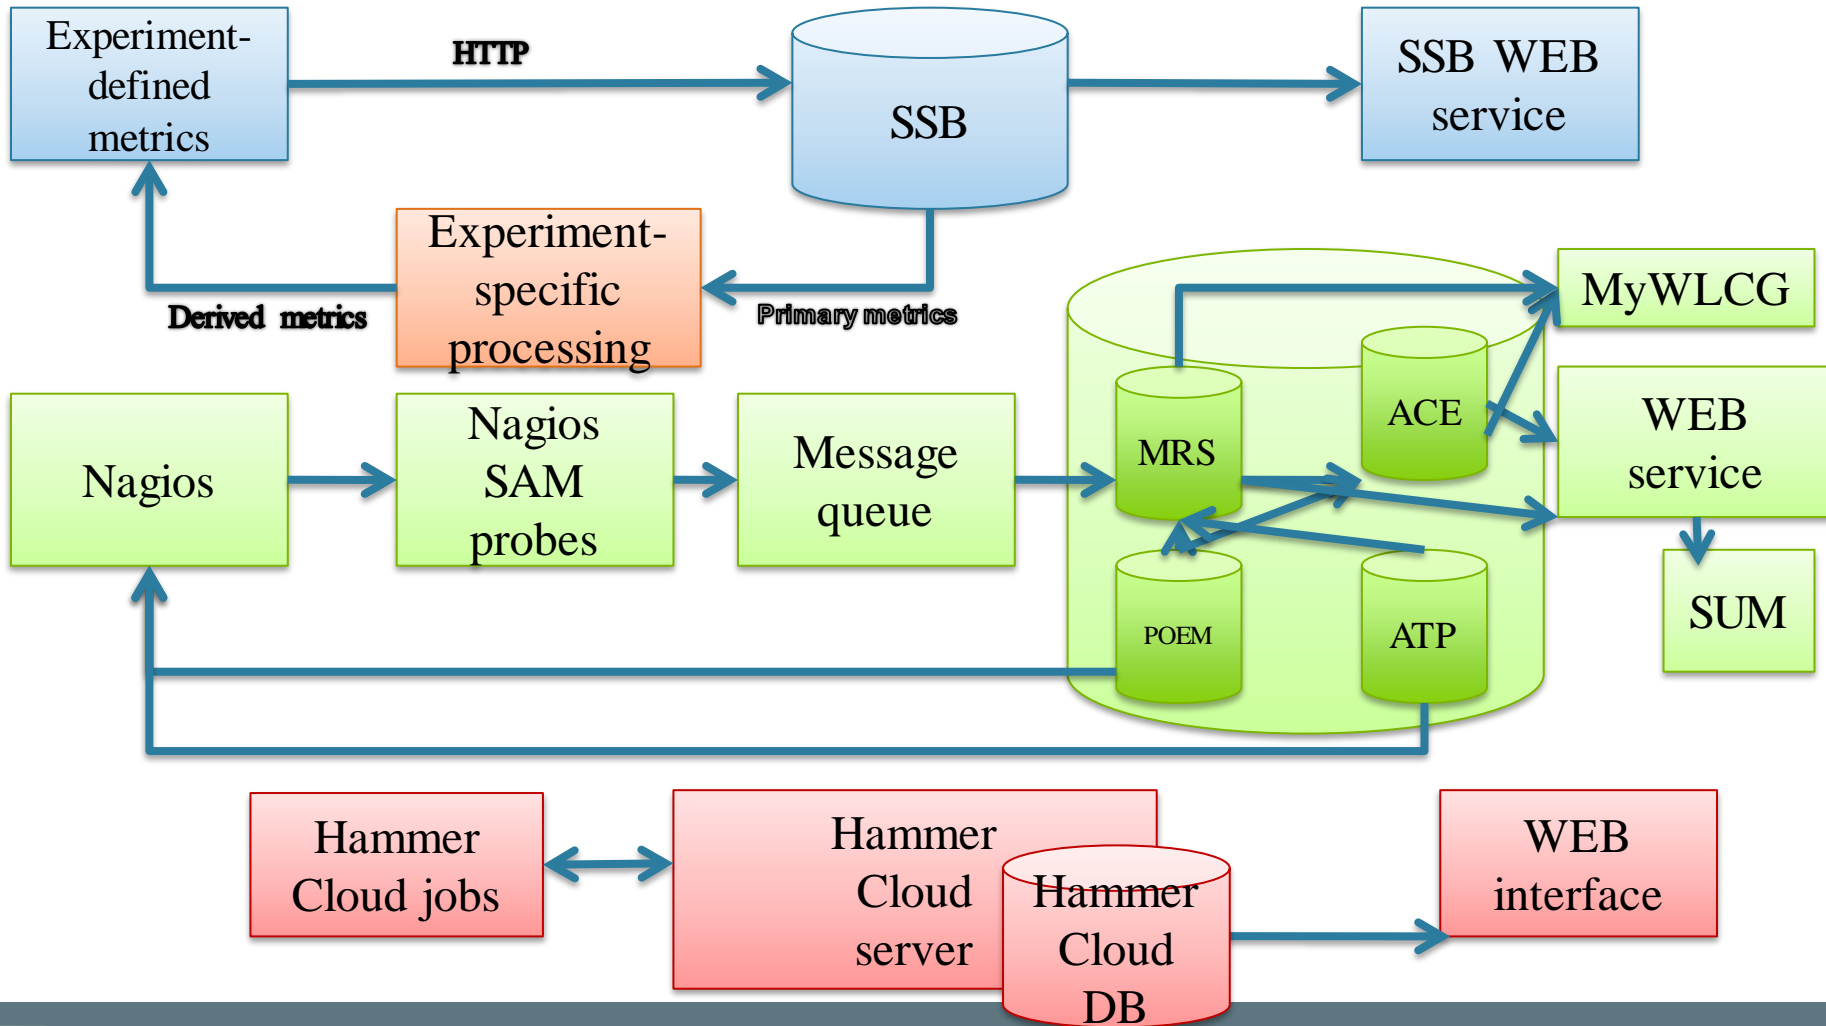
<http://cern.ch/go/pw7F>



5. Merging applications

- Categories:
 - Job Monitoring
 - Data Management /Transfers
 - Infrastructure Monitoring
 - Dissemination
- Infrastructure monitoring is currently the area with more applications
 - And where we can benefit the most from common components

Infrastructure monitoring (Current state)



Infrastructure monitoring (Goal)

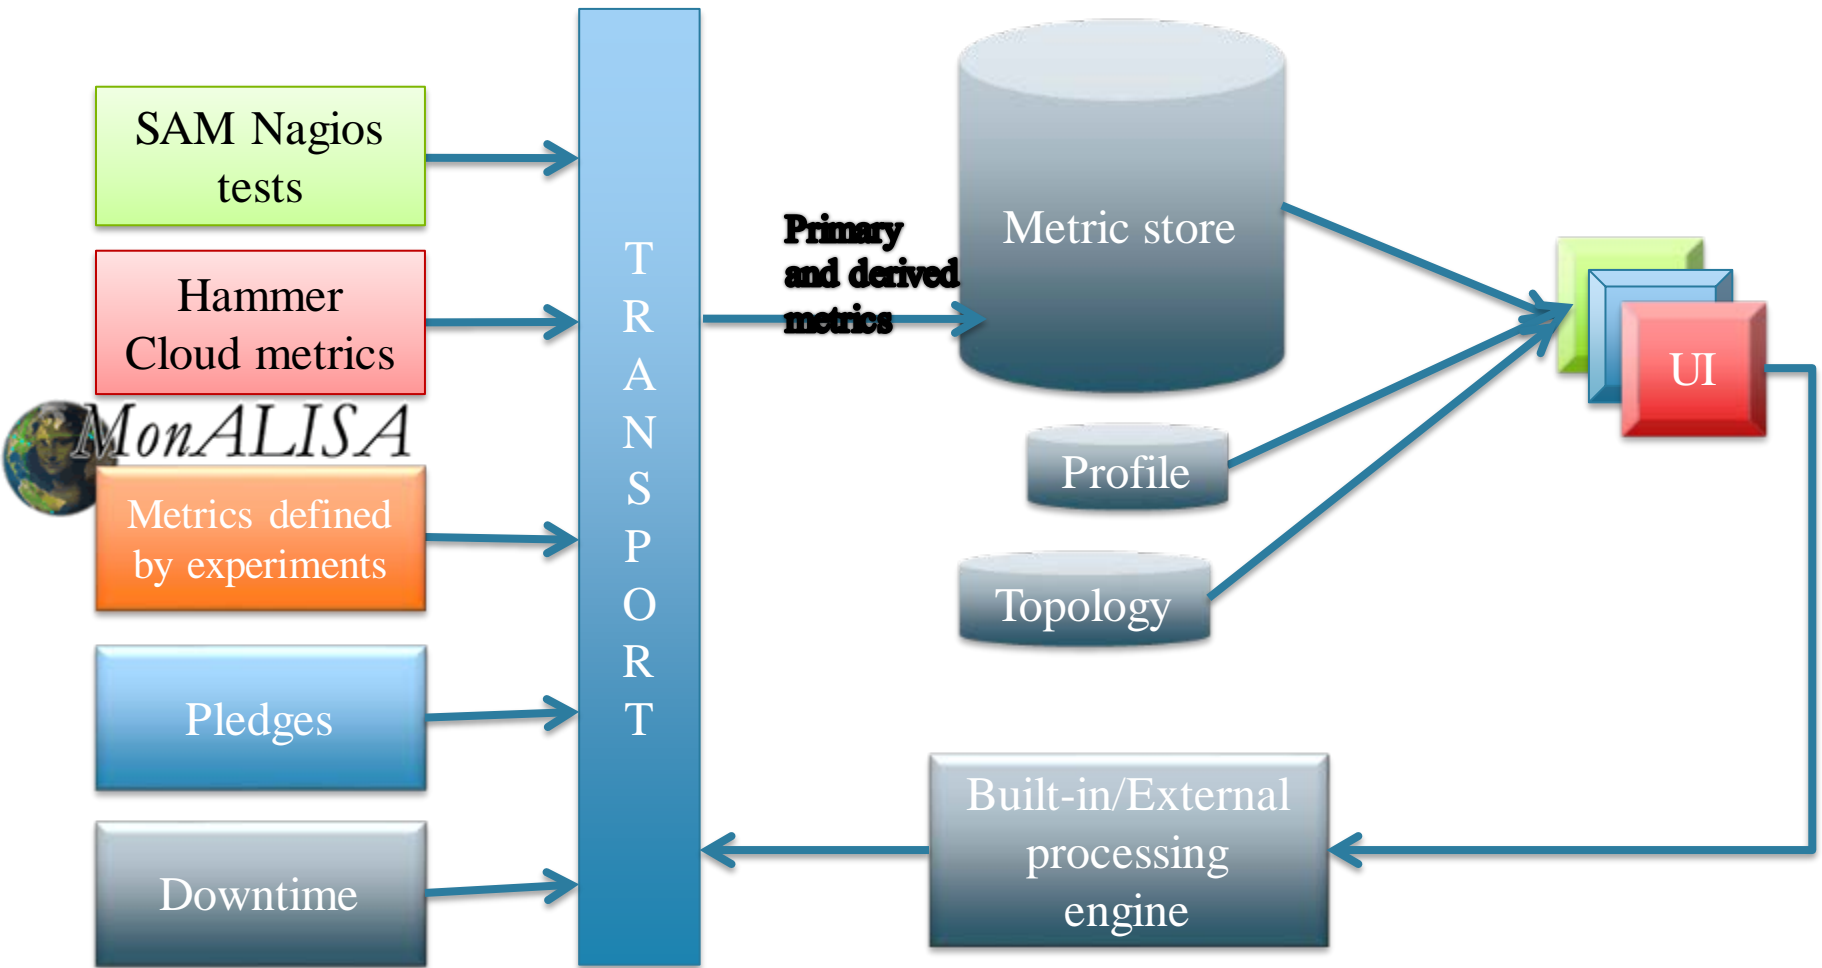


Table of contents

- WLCG Monitoring consolidation:
 - Definition of the project
 - How to do it
 - **Current status**
 - Impact on the experiments
 - Summary

Project plan

Report with the evaluation created.

Tasks identified and inserted in JIRA

WLCGMONCON plan Save as Details ★ Share Export ⚙

Monitoring Consolidation Type: All Status: All Assignee: All Contains text More 🔍 Advanced

54

Key	Summary	Assignee	P	Status	Updated	Original Estimate	Remaining Estimate	Due
VLCGMONCON-21	Ensure that the experiments can inject their own values in the transport layer	Marian Babik	↓	Done	17/Jan/14	5 weeks	0 minutes	31/Jan/14
VLCGMONCON-5	Ensure that GoogleEarth and Siteview are moved to the common framework	Edward Karavakis	↓	Done	17/Jan/14	1 week	0 minutes	31/Jan/14
VLCGMONCON-18	Stop ALICE FTD efficiency and CMS datasets	Pablo Saiz	↓	Done	20/Jan/14	2 days	0 minutes	31/Jan/14
VLCGMONCON-38	Incorporate simplevisor	Alexandre Beche	↓	Done	17/Jan/14	1 week	0 minutes	31/Jan/14
VLCGMONCON-39	Integrate Bamboo as a continuous integration system	Pablo Saiz	↓	Done	17/Jan/14	2 days	2 days	31/Jan/14
VLCGMONCON-24	Use memcache in JobMonitoring and Job accounting	Edward Karavakis	↓	Done	17/Jan/14	2 weeks	0 minutes	31/Jan/14
VLCGMONCON-49	Profile definitions should allow filtering and grouping per service and treating FQANs independently	Pablo Saiz	↓	Done	30/Jan/14	4 weeks	1 day	31/Jan/14
VLCGMONCON-26	Vofeed becomes the only source of topology description	Ivan Antoniev Dzhunov	↓	Done	17/Jan/14	2 days	2 days	28/Feb/14
VLCGMONCON-45	Ensure that WLCG Transfers applications are moved to the common framework	Alexandre Beche	↓	To Do	17/Jan/14	1 week	1 day	28/Feb/14
VLCGMONCON-14	Evaluate The Common Metric Store for REBUS	Pablo Saiz	↓	To Do	17/Jan/14	1 week	1 week	28/Feb/14
VLCGMONCON-28	Evaluate Drupal and twiki for documentation	Julia Andreeva	↓	To Do	17/Jan/14	2 days	2 days	28/Feb/14
VLCGMONCON-16	Investigate HammerCloud Schema, and check differences with JobMonitoring	Edward Karavakis	↓	To Do	17/Jan/14	1 week	1 week	28/Feb/14
VLCGMONCON-52	Integrate the downtime information to calculate the reliability	Ivan Antoniev Dzhunov	↓	To Do	30/Jan/14	2 weeks	1 week	28/Feb/14
VLCGMONCON-44	Create a prototype for fast feedback cycle	Pablo Saiz	↓	To Do	30/Jan/14	2 days	2 days	28/Feb/14
VLCGMONCON-23	Evaluate the common metric store framework	Luca Magnoni	↓	To Do	04/Feb/14	2 days	2 days	28/Feb/14
VLCGMONCON-25	Con...							31/Mar/14
VLCGMONCON-15	Ensure that...							31/Mar/14

<https://its.cern.ch/jira/browse/WLCGMON>

Different set of tasks

1. Application support

- Job, Transfer, Infrastructure, GoogleEarth

2. Running the services

- Moving to AI, EGI transition, Koji, SLC6, ...

3. Merging applications

- SSB and SAM, SSB and REBUS, HammerCloud and Nagios (still under discussion)

4. Technology evaluation

- ElasticSearch, Testing framework (see Luca's [Feb GDB presentation](#)), PostgreSQL, Django, Drupal...

1. Application support

- Support the current system
- Multiple applications identified by users/developers as **under control**
 - Already modular design
 - Hit the spot
- Introduce some new functionality
 - **Not covered by the consolidation project**
 - Example: Alexandre's [Feb GDB talk](#) 'The new WLCG Transfer dashboard'

2. Running the services

- Most of the machines moved already to openstack, puppet, SLC6
 - More will be done before end of March
 - Thanks to the Agile Infrastructure Team
- Transition of SAM EGI to consortium GRNET, SRCE and CNRS
 - Coordinated by Marian
 - See his [Feb 14 Mon Consolidation talk](#)

3. Merging applications

- Goal: reduce number of applications to maintain.
- Infrastructure monitoring:
 - Testing framework
 - Can the Nagios tests and HC Functional tests be combined
 - SAM & SSB:
 - storage and visualization already done
 - REBUS & SSB
 - Prototype ready. Still to investigate how to populate

4. Technology evaluation

- Most of the **work done** during the review & prototype phase
 - Effort still necessary to keep up to date
- Close **collaboration** with AI Monitoring

Table of contents

- WLCG Monitoring consolidation:
 - Definition of the project
 - How to do it
 - Current status
 - **Impact on the experiments**
 - Summary

Impact on the experiments

- Main impact on **Infrastructure monitoring**
 - Transition for other applications should be transparent
- **Consistent data** across all system
- Reduced number of UI

Topology

- **Vofeed** becomes single authority of topology
 - Any service, any instance
 - Also “meta” services, like the AliEn CE, CM, ...
 - No cross-check with GOCDB/BDII
- Content and format will be revised

Collecting information

- Nagios
 - Optional component
 - If experiment wants it, there will be a Nagios
 - CREAM CE probe with payload
 - CondorG probe
 - Each FQAN treated as a different metric
- Injecting metrics
 - Any metric published in the transport with the correct format will be collected
 - Dynamic definition of metrics
 - Metrics could be anything:
 - Status, # successful jobs, downtime info, # transfer errors

Metric store

- Solution already existing (SSB)
- Access to current state and historical data
- Stores only status changes
- Combining metrics into views
- Used for operations by ATLAS and CMS
 - Deployed for LHCb and ALICE
- A common metric store:
 - Simplifies current architecture
 - Allows evaluation of new technologies
 - Allows concept of usability (Site Readiness)

Aggregation of metrics

■ Horizontal

Site Name	cms topology	
	cms topology (flavour)	cms topology (site)
CREAM-CE ce10.pic.es	CREAM-CE	T1_ES_PIC
CREAM-CE ce11.pic.es	CREAM-CE	T1_ES_PIC
CREAM-CE ce2.particles.ipm.ac.ir	CREAM-CE	T3_IR_IPM

■ Vertical

Site Name
T0_CH_CERN
T1_DE_KIT
T1_ES_PIC

■ And another horizontal

CMS CRITICAL (Service CE)			
org.sam.CREAMCE-JobSubmit _cms_Role_icgadmin)	org.sam.CREAMCE-JobSubmit (_cms_Role_production)	org.cms.WN-env (_cms_Role_icgadmin)	org.cms.glexec.WN- gLExec (_cms_Role_pilot)
OK	OK	OK	OK
OK	OK	OK	OK
OK	OK	OK	WARNING

CMS_CRITICAL (CE) Expand the details
OK
OK
OK

More options for aggregation

- **Filtering:**
 - Combine metrics depending on the value of another metric:
 - Use different profile for T1 sites
 - Take values only for one service
- Other **base algorithms:**
 - Currently: AND, OR
 - Still to be developed: SUM, AVG, downtime, sliding window
- **Multiple** levels of aggregation:
 - Queue Name → CE → Site → VO

Definition of profiles

- Experiments can define multiple profiles
 - Include list of metrics
 - Filtering
 - Topology
 - Algorithm

Site readiness

- Define status of the site, defining if the site ‘can be used by the experiment’
 - Different from Site Availability
 - Issues might be experiment specific
 - Or from a different site
- Current approach of ATLAS and CMS:
 - Read data from SSB
 - Apply their own algorithm (sliding window, depending on tier structure, granularity)
 - Inject a new metric in SSB

Table of contents

- WLCG Monitoring consolidation:
 - Definition of the project
 - How to do it
 - Current status
 - Impact on the experiments
 - Summary

WLCG Monitoring Consolidation Group

- **Simplify** scope/tools/maintenance to support it with half of the resources *less than half*
- **Prototype** already deployed
 - Now, 'just' evolve it production
- Very good progress on moving to common infrastructure, and merging applications
- And plenty of work still ahead of us...