

# Operation Coordination and Commissioning [WLCG]

Josep Flix (PIC/CIEMAT)  
On behalf of the WLCG Operations Coordination Team

GDB – CERN  
[15<sup>th</sup> January 2014]



A 3D rendering of the words "Happy New Year" in a bold, red, sans-serif font. The letters are thick and have a slight shadow, giving them a three-dimensional appearance. They are set against a white background with a light gray floor. Numerous small, red, diamond-shaped confetti pieces are scattered throughout the scene, some floating in the air and some on the floor. The overall effect is festive and celebratory.

Happy  
New  
Year

*Thanks to sites for the best efforts made  
during the Xmas break*

**Much appreciated!**

# Outline

- ▶ Next meetings [**schedule**]
- ▶ News
- ▶ (some) Task Force Updates
- ▶ Conclusions

# Next Meetings (schedule)

- ▶ Meetings schedule until March 2014:
  - ▶ January 16<sup>th</sup> [VIRTUAL]
    - ▶ Experiments, sites and task forces should add their reports to the twiki as usual (if there is anything to report) but **there will be no meeting**  
<https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOpsMinutes140116>
  - ▶ January 30<sup>th</sup>?
  - ▶ February 11<sup>th</sup>: WLCG OpsCoord **F2F@CERN** [pre-GDB]
  - ▶ February 20<sup>th</sup>
  - ▶ March 6<sup>th</sup>

- ▶ Today we have a separate talk on the “***openSSL issue***”
  - ▶ Detailed report provided by Maarten L.
- ▶ **SAM tests:** Since Jan. 2014, VO-based reports replace the OPS-based reports for the official WLCG Availability & Reliability Reporting

<https://tomtools.cern.ch/confluence/display/SAMDOC/New+Joint+T1-T2+reporting>

# Task Forces Updates (1/12)

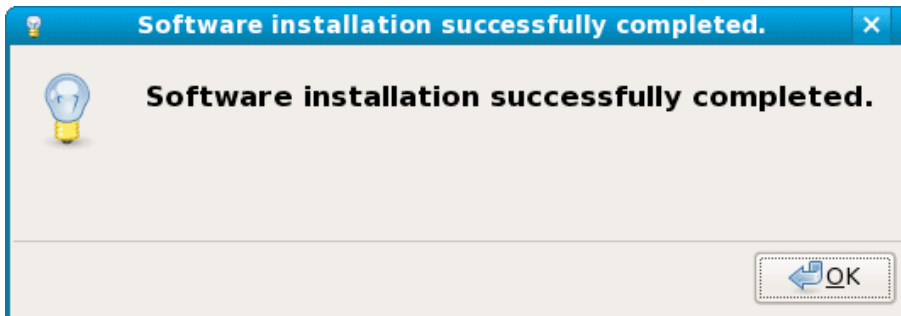
## ► Middleware deployment (CVMFS)

<https://twiki.cern.ch/twiki/bin/view/LCG/WLCGBaselineVersions>

- Deadline for sites to upgrade CVMFS from v2.0.x to v2.1.15 or higher is March 1<sup>st</sup>
- Then, Stratum-0 service will be upgraded to v2.1
- Experiments might add CVMFS tests to their **critical SAM profiles** at some point

## ► SL6 Task Force (Concluded)

<https://twiki.cern.ch/twiki/bin/view/LCG/SL6Migration>



*11<sup>th</sup> Dec: 92.6% of resources moved*  
*Today: >95% of resources moved*

Sites that still have a ticket open → <http://tinyurl.com/pabsftu>

# Task Forces Updates (2/12)

## ▶ gLExec deployment

<https://twiki.cern.ch/twiki/bin/view/LCG/GlexecDeploymentTracking>

- ▶ # of closed tickets is 68
- ▶ # of open tickets is 27
- ▶ ALICE & ATLAS still need some development
- ▶ LHCb is testing DIRAC pilot with gLExec
- ▶ CMS almost done:
  - ▶ will make gLExec **SAM test critical** in January

# Task Forces Updates (3/12)

## ▶ PerfSONAR deployment

<https://twiki.cern.ch/twiki/bin/view/LCG/PerfsonarDeployment>

▶ All sites must deploy **pS 3.3.1** (or later) by **April 1<sup>st</sup>**

▶ The TF provides all the needed support

▶ Sites not yet ready are now tracked by tickets


⚠ The **perfSONAR dashboard** is not developed/supported anymore



We need to find **some volunteer** to take responsibility of the code



**THIS network monitoring is the THING everyone wants to have!**





# Task Forces Updates (4/12)

## ▶ **SHA-2 (1/2)**

<https://twiki.cern.ch/twiki/bin/view/LCG/SHA2readinessTesting>

- ▶ The **EOS SRM is OK** now for **ATLAS and CMS**
  - ▶ for CMS a Java *config* change was needed to allow proxies with 512-bit keys; that change may also be needed for the others
- ▶ The **EOS SRM** for **LHCb is not OK** yet
  - ▶ a patch will need to be ported to let the new SRM support the "*root*" protocol expected by LHCb jobs
- ▶ The **OSG CA** has delayed switching to SHA-2 from Jan 14<sup>th</sup> to Feb 11<sup>th</sup> on request of CMS
- ▶ The **CERN CA** has delayed making SHA-2 the default to Feb 11<sup>th</sup> as well (SHA-1 remains available until June)
- ▶ The **UK CA** intends to switch mid March
- ▶ The **French CA** should switch on Jan 14<sup>th</sup> (the day before this GDB)

# Task Forces Updates (5/12)

## ▶ **SHA-2 (2/2)**

- ▶ The **VOMS-Admin instability has been fixed** (Java 7→6)
  - ▶ many thanks to Valerio Venturi [**RIP**] for debugging this!
- ▶ The **VOMS-Admin test cluster** is being reinstalled and should become available in the coming days
- ▶ **SHA-2 timeline:** by mid Feb the WLCG infrastructure is expected to be essentially ready
  - ▶ But, the migration from VOMRS to VOMS-Admin may take longer...

# Task Forces Updates (6/12)

## ▶ IPv6 validation and deployment

<https://twiki.cern.ch/twiki/bin/view/LCG/WlcgIpv6>

- ▶ Job submission to a CE via **gLite WMS**, both using IPv6 to communicate, was **successfully** tested to work
- ▶ A similar test via **Condor-G** **failed**
  - ▶ bug in HTCondor, which uses a rather old version of CREAM (condor-8.1.3/externals/bundles/cream/1.12.1\_14)
  - ▶ Waiting for feedback from the developers
- ▶ **CMS tests** (next steps):
  - ▶ to get a BeSTMAN SE and an OSG CE to run data transfer and job submission tests respectively
  - ▶ If successful, will proceed to test glideinWMS in dual stack submitting jobs to IPv6-only CEs and WNs
  - ▶ Most central services do not need to be tested on IPv6 for the time being (typically those services that do not directly interact with site resources)

# Task Forces Updates (7/12)

## ▶ WMS decommissioning

### ▶ Experiment usage:

- ▶ Still used for a small fraction of CMS analysis
  - ▶ CMS users strongly encouraged to move to *glidein* submission
  - ▶ Usage is indeed decreasing
- ▶ VOs still using the shared WMS instances will be contacted in the coming days, in particular LHCb

## ▶ Tracking tools evolution

- ▶ Transition from Savannah to JIRA: slow progress
- ▶ One tracker has been used as a prototype, and there are still several open issues like broken cross references between tickets and missing version names

# Task Forces Updates (8/12)

## ▶ Machine/job features (1/2)

<https://twiki.cern.ch/twiki/bin/view/LCG/MachineJobFeatures>

- ▶ Machine features **available** as of today at CERN:
  - ▶ "*mjf.py*" [1] is deployed via an rpm from the wlcg-repository [2] on SLC5 (bare metal) and SLC6 (OpenStack virtualised) WNs
  - ▶ Ulrich is taking care of providing the features to each worker node and via *mjf* they can be picked up
    - ▶ *mjf* makes use of the OpenStack meta data service. If this is done too often (from too many machines), it might not scale... (moderate testing going on)
  - ▶ The source code of the client is available at Usage:

```
from mjf import mjf
m = mjf()
m.collect()
print m.features()
```

[1] <http://cern.ch/go/BKp7>

[2] <http://linuxsoft.cern.ch/wlcg/>

# Task Forces Updates (9/12)

## ▶ Machine/job features (2/2)

### ▶ What is left to do:

- ▶ make job/machine features available on the OpenStack *Cloud environment*
- ▶ provide the *mjf* client also in the Applications Area as a "**middleware package**" such that it can be picked up by experiments taking their software from there

### ▶ Proposal:

- ▶ Give experiments 1 month time to play with this and give feedback for this prototype
- ▶ If all OK we can extend machine/job features to other sites (we do have first implementations for all batch system types including condor via Igor Sfiligoi)
- ▶ At the same time look into other IaaS infrastructures than OpenStack to see how we can implement features on those

# Task Forces Updates (10/12)

## ▶ **Middleware readiness**

<https://twiki.cern.ch/twiki/bin/view/LCG/MiddlewareReadiness>

- ▶ **1<sup>st</sup> kick-off meeting** occurred last 12<sup>th</sup> December:
  - ▶ Minutes: <https://indico.cern.ch/getFile.py/access?resId=1&materialId=minutes&confId=285681>
  - ▶ There **wasn't enough participation from experiments and sites** so the discussion concentrated on software repositories and the criteria for judging a product deployment process successful
- ▶ How to involve experiments and sites should be discussed in the next meeting, which is planned for Feb 6 at 16:00 CET
  - ▶ Please try to get involved!

# Task Forces Updates (11/12)

## ► Multicore deployment (1/2)

<https://twiki.cern.ch/twiki/bin/view/LCG/DeployMultiCore>

wlwg-ops-coord-tf-mc e-group → **subscribe!**

- **Mandate**
  - Find out how experiments plan to **request and use** multicore resources
  - **Agree** on practical multicore resource provisioning **models** for WLCG and **evaluate** them at selected sites with real workflows
  - Define a **deployment strategy**
  - Drive the **deployment process**
- The scope is **Grid resources**, not cloud
  - But interaction with the Cloud WG is highly desirable
- Coordinate also with the Machine/Job features TF
- To be led by **Alessandra Forti** and **Antonio Perez-Calero**



# Task Forces Updates (12/12)

- ▶ **Multicore deployment (2/2)**
  - ▶ Activities going on in the context of experiments, at least ATLAS and CMS
    - ▶ recent ATLAS multicore production
    - ▶ CMS scheduling models
  - ▶ TF will soon call for a first meeting to introduce the TF and start collecting feedback and start providing guidance
    - ▶ Doodle: <http://www.doodle.com/tnxn9zxhcqdtv758>

# Conclusions

- ▶ All of the **TFs are progressing well!**
- ▶ In particular, WMS decommissioning, Multicore deployment, and Middleware readiness Task Forces have started their activities:
- ▶ **Sites and experiments** are **encouraged to actively participate** on the discussions and the TFs!

# Links

- ▶ WLCG Operations coordination twiki:  
<https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOpsCoordination>
- ▶ Mailing list:  
`wl cg - ops - coord @ cern . ch`