

December 2012 - GDB Summary

See also: https://www.gridpp.ac.uk/wiki/GDB_12th_December_2012

Jeremy's notes

Operations Coordination Team

- Link: <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOpsCoordination>
- IGTF would like CAs to move from **SHA-1 to SHA-2** signatures ASAP, to anticipate concerns about the long-term safety of the former. Overview of status:
<https://twiki.cern.ch/twiki/bin/view/LCG/RFCproxySHA2support>
- **Middleware:** EMI-2 SL5 WNs deployment is proceeding at WLCG sites. Test queues page <https://twiki.cern.ch/twiki/bin/view/LCG/WorkerNodeTesting>. Strong warning against upgrading to SL6. It has not yet been validated by ATLAS. SL6 is OK for ALICE, CMS and LHCb.
- **Glexec:** The "ops" gLExec SAM test has been added to the EGI ROC_OPERATORS profile => failures raise alarms in the EGI operations dashboard. 1st Goal: EGI ROD teams to try and get the relevant sites in their regions to reach 75% availability (Dec-Jan). Status = <http://cern.ch/go/fgz8>.
- **Tracking tools** :E.g. savannah-GGUS bridge technical issues and savannah-JIRA migration.
- **FTS3:** Functional tests for the "FTS2-like configuration" ongoing for ATLAS and CMS (RAL, ASGC and CERN)
- **Squid monitoring:** Objective: since squids are multipurpose and multi-VO, move monitoring to
- WLCG responsibility and integrate with common WLCG operations. ATLAS/CMS (<http://frontier.cern.ch>) move to Aggregated Topology Provider (ATP).
- **CVMFS:** Used for - Software installation, conditions data, nightly build releases. 2.1.5 beta includes (last resort) NFS support. All expts. want full CVMFS deployment – target 30th April. Current status <https://twiki.cern.ch/twiki/bin/view/LCG/CvmfsDeploymentStatus>.
- **perfSONAR:** Mesh tested – script needed at the moment but native in 2013. Sites will be asked to move to the mesh configuration. MDM vs PS interoperoperation not yet clear... backend data is compatible.

EMI-2 deployment

EMI-2 deployment

- Upgrade deadline for Sept. EOL components is 17th Dec.
- Jan 31st: all DPM, LFC and WN must be upgraded to EMI or decommissioned
- From 10th Dec: ROD teams open GGUS tickets asking for a plan in 10 working days. Alarms not handled properly by ROD, and tickets assigned to unresponsive sites are escalated to NGIs and then COD.
- EMI-1 probes being developed -> alarms in March.

Requirements for Glue 2

- More sites publishing glue 2 since November. Profile document <http://go.egi.eu/glue2-profile>.
- For Jan: Integrate glue-validator in SAM. Longer term: Integrate glue-validator in the resource BDII
- Feedback for ginfo wanted.

Oracle at T1s

- ATLAS
 - Distributed databases in LS1 and Run2. Oracle + Frontier – 4 T1s continue with present system.
 - Remove HOTDISK space token from CVMFS sites. Take conditions files from CVMFS.
 - Tag evolution using Hadoop/Hbase in place of Oracle backend.
-
- CMS
 - Stable Frontier infrastructure for conditions. (Looking at backup)
 - Need Oracle for FTS only at T1s.
 - Have 4 other DBs: PhEDEx; To tracking; Organised processing; conditions.

DPM community + workshop

- A lot of recent development for DM-Lite but the old interfaces remain unchanged

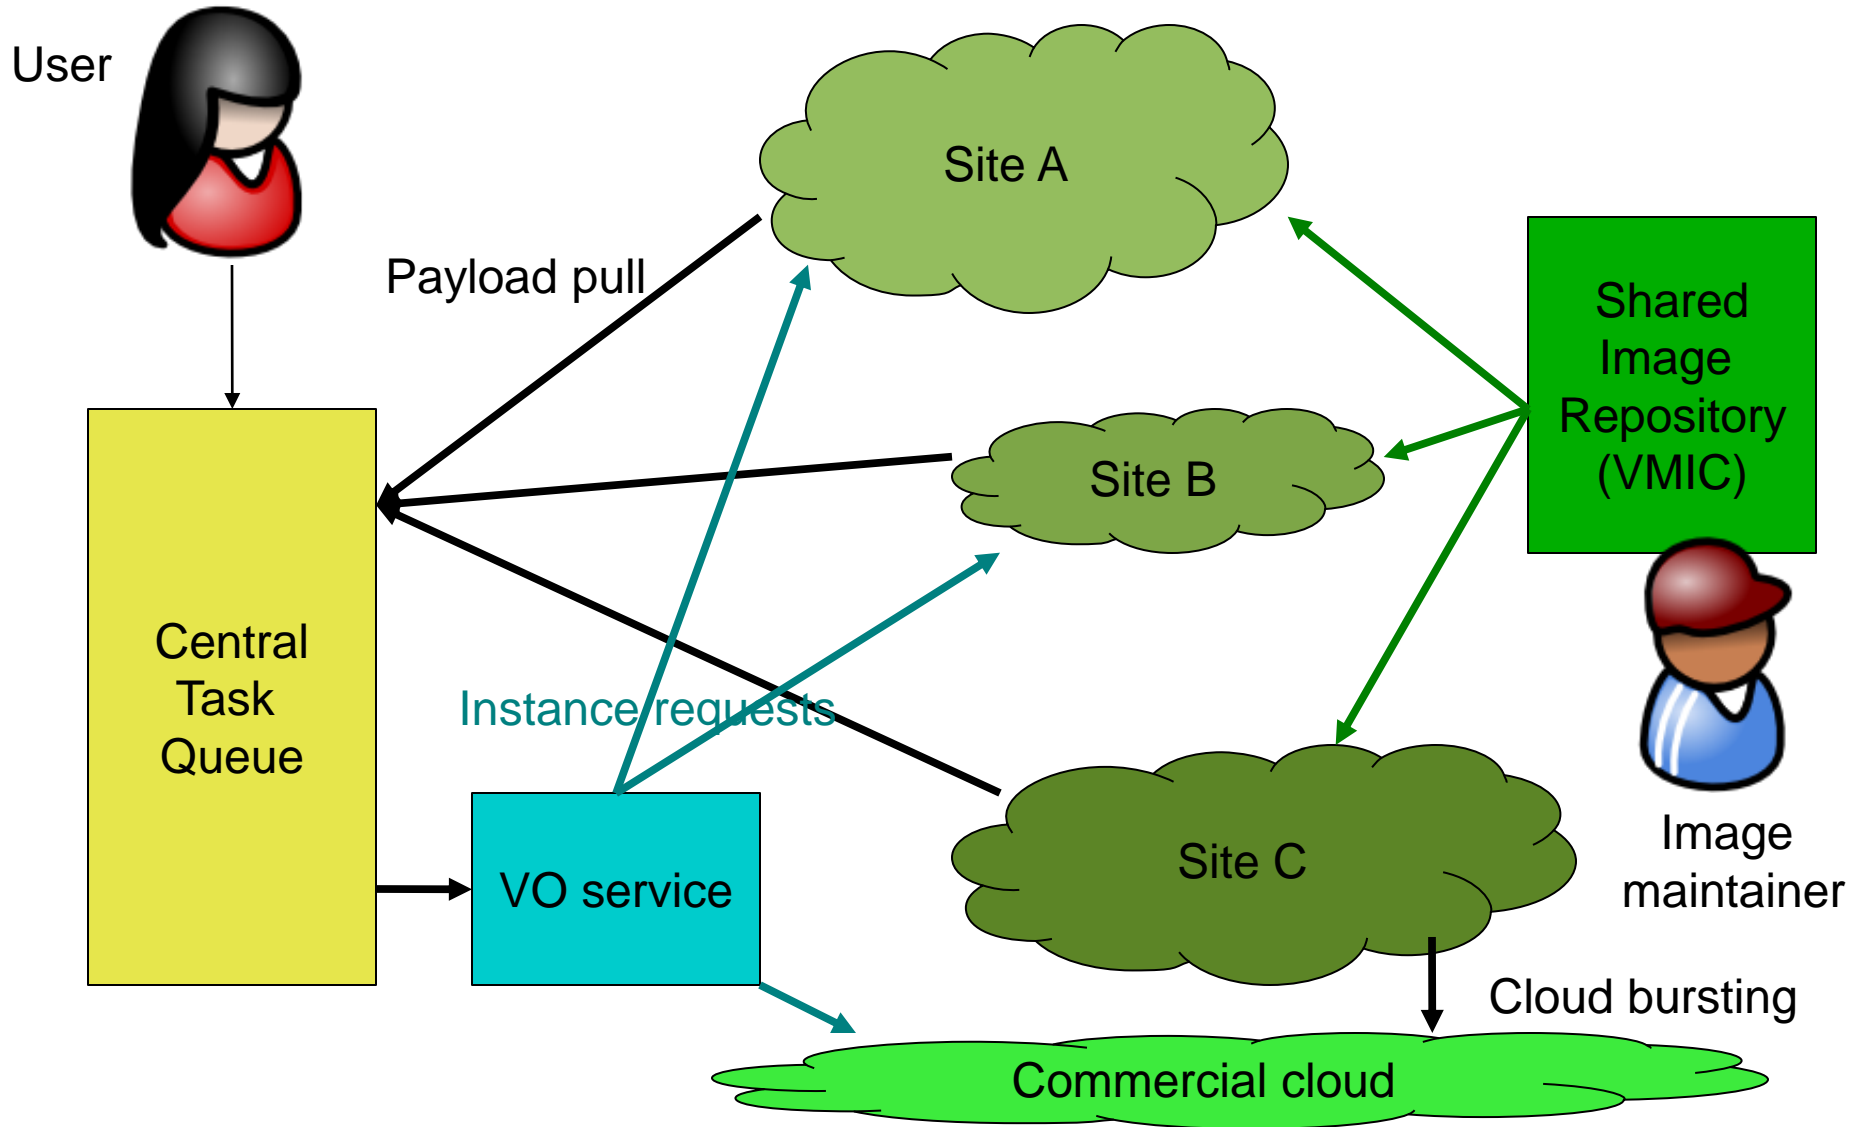
Feedback areas: I/O concurrency limits; drain performance; hot file replication; storage 'rebalancing'; metadata consistency utilities; inter-VO quotas

- Question on enabling http/xrootd. Should be a couple of hours work.
- “During a recent DPM workshop at LAL (on 5th December 2012), representatives of several countries met to discuss the future of DPM. Those present recognised the importance of DPM and agreed to establish and contribute effort to a DPM Collaboration with the express aim of driving the DPM project forwards after EMI ends on 31st March 2013. The effort indicated as available across the collaboration partners was more than adequate to maintain, develop and support DPM as an efficient storage solution. A collaboration agreement and the assignment of responsibilities will be completed in due course”

Virtualized WNs

- Virtualized WN : where we are? what next ?
- HEPiX vwg model (and S/w) endorsed by the EGI federated cloud task force.
- The HEPiX virtualisation working group was formed to facilitate the instantiation of user-generated virtual machine images at HEPiX (and WLCG) sites.
- The HEPiX VWG developed a policy that introduced the concept of image endorsers: people who would guarantee that generated images could be used safely at sites. The policy was approved.
- Now looking at implementations.

How this could be used



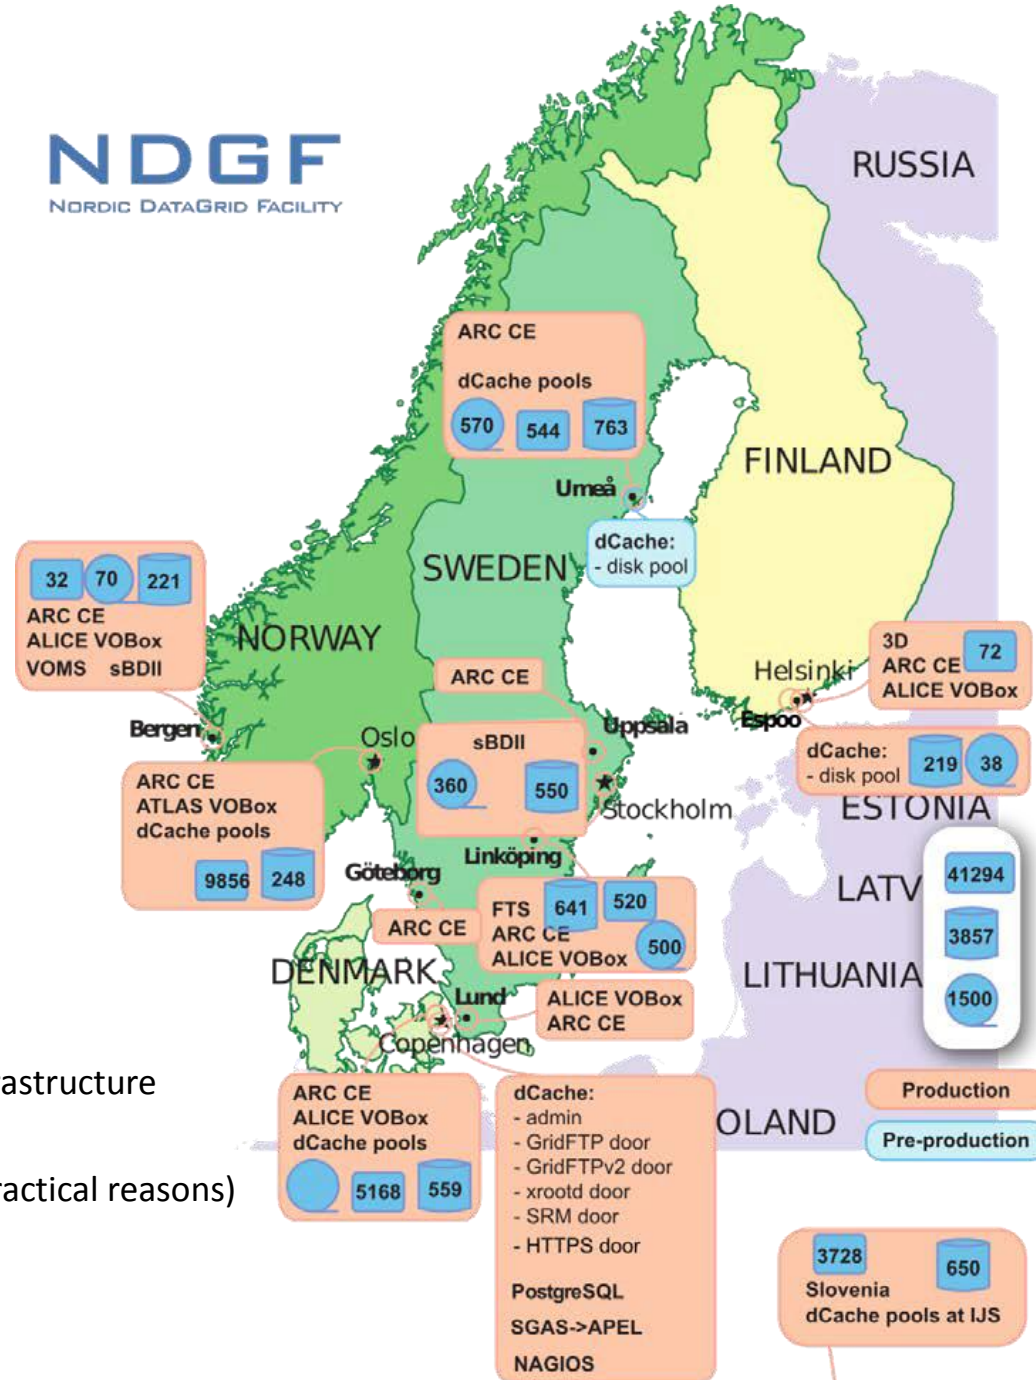
Software Defined Networks for Big-Data Science

- TCP is the underlying data transfer protocol
- Fragile; sensitive to loss; poor deployment of infrastructure
- Discussed OpenFlow

- Powerful network abstraction
 - Files / Storage
- Benefits
 - **Simplicity** for the end-site
 - Works with off-the-shelf, open-source controller
 - Topology simplification
 - **Generic** code for the network provider
 - Virtual switch can be layered over optical, routed or switched network elements
 - OpenFlow support needed on edge devices only, core stays same
 - **Programmability** for applications
 - Allows end-sites to innovate and use the WAN effectively

MW provisioning lifecycle

- Overall approach has been accepted 😊
- Most feedback was very constructive
 - very few comments reflect a significant lack of awareness of emi's status and work
- Some major and several minor issues that need clarification:
 - Role of EGI, open-ended-EMI-consortium, ScienceSoft
 - What about other sciences?
 - How to deal with the different middleware stacks?
 - Where are the gaps?
 - what can we get from EPEL/Fedora what not
 - where will testing resources come from
 - Organization of level 1,2,3 support and the role of GGUS
 - Rollback
 - Meta-Packages
 - Version management
 - Repositories (a potential can of worms)
 - OS Platforms
 - PTs' resources for communication and coordination
 - PTs' ability to run Pilot services
 - Release frequency
 - Configuration Management (YAIM or not YAIM)
- -> Some 'answers' given in Markus's talk. New proposal draft soon... for discussion in Jan.



NDGF Tier1 and NeIC – status update

Grid-centric NDGF is evolving into generic Nordic e-Infrastructure

- WLCG and EGI services are integral part of NeIC
- NDGF-T1 is kept as a name for the Nordic Tier1 (for practical reasons)
- It is business as usual, despite re-branding