

GDB- CERN- 18/4

Agenda: <https://indico.cern.ch/conferenceDisplay.py?confId=155067>

Table of Contents

Introduction- M. Jouvin	2
TEG Future	2
Data and Storage Management- D. Duellmann	2
Operation Tools- M. Girone	3
Workload Management- D. Salomoni.....	3
Security- R. Wartel.....	4
Databases- D. Dystra	5
perfSONAR-PS Experience- S. McKee	6
Middleware	6
EMI- C. Aiftimei.....	6
UMD- M. Drescher	7
Conclusion	8
SHA-2 and RFC proxies- M. Litmaath.....	8
glExec update- M. Litmaath.....	8
OSG SW Update- b. Bockelmann.....	9
Conclusions – Actions.....	10

Introduction- M. Jouvin

See

<https://indico.cern.ch/getFile.py/access?contribId=1&sessionId=0&resId=0&materialId=slides&confId=155067>.

Last experiment session in March: next in June or July

TEG day before May GDB: probably too early, in particular as there will be a TEG day in NY.

- Probably a TEG day in June as a follow-up for WCLG workshop

Need volunteers to act as minute makers in next meetings

- Easiest would be a team of rotating volunteers

Agenda for next meeting ready soon after the current meeting

TEG Future

Data and Storage Management- D. Duellmann

Data placement

- Users interacting with the disk layer, managed by the placement layer
 - o Originally managed by experiment pushes
 - o Transparent HSM proved not to be optimal: SRM BringOnline (experiment decides what to move to cache)
- Experiment workflows interact with archive layer hosted in dependable sites
 - o Access forbidden to users
 - o T2s forbidden: insufficient service level
- Federation layer added by some experiments, tested by others
 - o Opportunistic end user pull: placement triggered by user jobs
 - o May also be used for transparent repair
 - o The client protocol has to be the federation protocol
- Tendency to use the intersite transfer protocol infrastructure (e.g. FTS) to do the archive to disk cache migration

Storage element components

- Raw media: data storage
- File cluster: expose a consistent namespace
- Client protocols for random I/O (jobs), sequential I/O (FTS), storage management (e.g. SRM)
 - o Ideally should share the user authorization/mapping layer but not always true

- Cloud storage: may be integrated as another "raw media" layer or as a media+cluster replacement
 - o Second approach will have more dramatic impacts, including on client protocols (e.g. S3)

Recommendations

- Many recommendations see slides for details.
- Need more work to identify what are really the actions requiring a working group and what are the recommendations that are just a statement and those who require a follow-up action.

Operation Tools- M. Girone

Not many things new since February report

- Got feedback from OSG
- TEG has the appropriate membership to guide the implementation of the recommendations that will receive the green light.

Goal or where we would like to go/be

- A small number of well-defined services
- Installing, configuring, upgrading should be trivial
- Services resilient to glitches and problems

See Twiki for the full list of recommendations

- Several proposed for implementation during LS1 with several criteria of prioritization
- Some considered medium-term and require a WG
- Long term: require coordination and communication

Workload Management- D. Salomoni

Recommendations

- Many recommendations see slides for details.

Future work

- Complete the report editorial process
- Use next pre-GDB day for a F2F meeting
- Groups, priorities and activities by CHEP

Dependencies on some external providers, including other TEGs.

Security- R. Wartel

Steffen Schreiner writing his thesis, resigned as a co-chair

- Was leading the WN security subgroup

Current status: document complete, waiting for MB feedback

In progress: prepare recommendations and their implementation by the end of the year

Risk analysis: 11 risks identified and scored.

- 3 top risks; misused identity, attack propagation, OS vulnerability exploit

AAI on WN: fine grained traceability and physical identity switching required

- But some important stakeholders absent from the discussion
- Several topics need further discussions: see slides. Hope to be able to make progress on these topics by the end of the year.
- Common meetings during pre-GDB days with other TEGs may help to bring all experts together
- May pre-GDB may not be convenient as there is an EMI AH meeting in DESY that glexec and Argus developers will probably attend...

AAI on storage systems

- Work in progress
- Main conclusions: traceability needs to be improved, effort required to improve data protection, data ownership
- Recommendations expected by June

Other subtasks in progress

- Identity federaton: hope to have recommendations ready by end of June
- Usability vs. security: interim report ready, final in June

Very important: contribution from experts of other TEGs

- Fold the security discussions into the appropriate TEGs?
- Use pre-GDB days to organize common meetings

Databases- D. Dystra

Conditions DB: no experiment plan to change their current condition DB model or support tools

- ATLAS, CMS and LHCb expect CORAL and COOL will continue to be supported by CERN

FrontTier service: already used by ATLAS & CMS, LHCb planning its use

- CMS planning to add a second server
- ATLAS has 6 server sites, may evaluate the possibility to reduce this next year
- CMS planning failover Squids

Squid

- Site Squid proxy closely monitored by CMS
- Recommend each site use one set of Squids for all production, approved apps
- Recommend WLCG decide a standard way of discovering Squid

Oracle: no major changes planned, scale-up gradually

- Move from Streams to Active Data Guard where possible

Online databases

- All experiments using PVSS: major upgrade coming, may conflict with timing of LS1
- CMS & ATLAS using different methods of loading condition data into HLT: consolidation possible but no plans to do so

NoSQL

- CMS & ATLAS had done a good job evaluating them with different implementations in production
- Hadoop MapReduce identified as potentially more disruptive but attractive: recommend that CERN IT set up a cluster usable for evaluation
- Build LHC-wide communities around these tools

No plan to continuation work but may contribute to drive to building communities around NoSQL tools

perfSONAR-PS Experience- S. McKee

See <https://indico.cern.ch/getFile.py/access?contribId=6&resId=1&materialId=slides&confId=155067> slides.

Action: get T2s from VOs other than Atlas to join the existing infrastructure

Middleware

Issues have been observed that affect WLCG and others

- Not a criticism against anybody
- Better to know the solution

Quality of SW reaching production

- At the origin, EMI's responsibility but made through UMD...
- WN content: Hydra client missing into the EMI WN because it has not been released in EMI yet
- EGI should have a better view of what is deployed
- Sites requiring many YUM repositories because some EMI components failed to be included by UMD (for good reasons...) despite they are required...

EMI- C. Aiftimei

EMI 1

- Update 13 (Feb.) and 14(March)
- Mainly BDII
- Preparing update 15: a bit delayed
 - o BLAH job registry corruption under heavy load
 - o WMS: performance/scalability issue affecting all GRAM-enabled CEs
 - o DPM 1.8.3
 - o GFAL/lcg_util 1.12.0: Glue 2.0 support in clients
 - o Proxyrenewal: fixing problems with expired VOMS extensions
- Future updates
 - o VOMS Oracle plugin (fixes for crashes)
 - o VOMS-admin 2.7.0 with several new features
 - o WMS 3.3.6: propagation of job exit code

UMD 1.6.0 released beginning of April : EMI products missing

- dCache 1.9.12-13 (unverified)
- Not planned for inclusion: AMGA, LFC/Oracle, VOMS/Oracle

gLite security updates until EMI-2 release (30/4)

- No update foreseen so far
- In fact security updates extended until end of Sept. for WN and UI only

EMI-2: 4 RC since Dec. 11

- After RC4: 98% successful build rate on SL5, 95% on SL6
 - o Problematic products: Hydra, WMS
- Deployment on EMI testbed to be completed today
 - o Most products successfully tested
 - o Estimated release date: May 7

Note on DPM: a bug present in all versions up to 1.8.3 included is causing DPM daemon crashes

- Workaround to be advertised soon
- Real fix in the next weeks

VOMS: how to verify that VO managers configured AUP signing policy according to EGI policies

- Renewal at least once a year

UMD- M. Drescher

EGI policy: UMD will contain only products coming through releases from EMI, IGE and a few others

- Hydra has not been released as part of EMI yet
- UMD will not contain products from other Technology Providers than those EGI has relationship with
- gLite 3.2 is still maintained by EMI but is not an EMI product

Provisioning: product updates are verified and are rejected from inclusion in next UMD update if they are failing some quality criteria

- E.g. happened to WMS
- All provisioning tests are publicly documented and full disclosure of provisioning results to technology providers

EMI vs. EGI testing: not the same scope

- EMI: remove bugs
 - o Unit testing, integration testing, system testing
- EGI: continuous service delivery
 - o Acceptance testing
 - o Assess whether delivery into production is reasonable (doesn't guarantee against bugs)

Question from Ian: why UMD distribution rather than direct use of technology provider distributions

- EGI will do only the version quality assessment

Conclusion

Follow-up at June's GDB about impact of end of EMI project on MW delivery, support and maintenance

- Including impact on EGI process and UMD

SHA-2 and RFC proxies- M. Litmaath

See January GDB for initial discussion.

IGTF plans are still to allow SHA-2 certificate no later than Jan. 2013.

Closed discussion list about SHA-1 and its vulnerabilities started mid March

- WLCG well represented
- Risk assessment in progress
- Define mitigations where possible
- May give us extra time for SHA-2 migration that remains required

Conflicting requirements/contraints between different MW products

- EMI-2 products should be ready for RFC proxies (and thus SHA-2)

Proposed milestones

- Deploy RFC-ready version of EMI products by the end of the year but still use SHA-1
- Switch to RFC proxies in Jan. 2013
- Upgrade dCache and BestMan
 - o When upgraded will not handle anymore Globus proxies...
- Ready for SHA-2 in spring 2013

No plan B yet... not clear if it will be needed, not excluded

Official WLCG message transmitted by WLCG to IGTF about our plan and our need to avoid disruption of service.

glExec update- M. Litmaath

As of early April, tests are no longer submitted: each NGI in charge of submitting their own tests

Each EGI site supporting glExec needs to update the GOCDB

- Add glExec flag for each CE supporting glExec on WN
- Results are available in MyWLCG (<http://cern.ch/go/PX7p>)
 - o Currently only 69 CEs, some T1 missing

- 72 sites claiming glxexec support in BDII

Experiment plans

- CMS: enabling glxexec in glideinWMS per site
- ATLAS working on glideinWMS back-end for PanDA
- ALICE will look into implementing Security TEG proposal
 - Doesn't use user proxies: idea is to use specifically crafted proxies with critical extension that only glxexec can understand
- LHCb will check if glxexec support in DIRAC needs adjustments

Action: Need to ramp-up glxexec deployment

- Presentation at WLCG workshop with hints for each configuration strategies
- Tarball WN: specific problem as the glxexec configuration cannot be made relocatable/generic but a few recipes to help
- Quattor QWG templates: better advertize how to enable it (simple!)

OSG SW Update- b. Bockelmann

OSG3: next version after OSG 1.2, merged with VDT 2.x

Distribution methodology changes: give up with packman on RHEL clones, use RPM instead

- Additionally adhere to Fedora/EPEL
- Do not provide any package provided by RHEL or EPEL
- Goal: be net contributors to wider communities, benefit from their innovations
 - Submit patches upstream everytime this is possible

Packaging done with Koji, used by RHEL and Fedora

- Highly automated build process
- Most packages already available in Koji, last ones expected end of April

RHEL5 fully supported, RHEL6 only fr WN

- Still discussing how to do the packaging for Debian but less needs

OSG3 also contains major release of many components.

Sites generally happy with the new approach but some sites willing to keep a relocatable distribution for some components like WN: still trying to figure out how to implement this.

- Moving from obscure packaging format (Pacman) to standard RPM is still a migration!
- Was a discussable approach to have in the same release both a new packaging approach and major releases for most components

Conclusions – Actions

GDB organization

- A team of volunteers for making minutes
- Decide next “TEG days”, taking into account the need for having experts from several groups present to discuss some issues like computing or storage security (AAI) or operational tools.

TEGs

- Recommendation prioritization (MB)
- Suggested WGs: identify what is really WLCG specific and what can be handled by other (wider) groups like HEPiX

perfSONAR

- Encourage deployment at T2s of VOs other than ATLAS

Glexec

- Presentation at next WLCG workshop highlighting pointers on how to configure it for different configuration tools

Middleware

- Discussion on MW support/development/validation after EMI at June’s GDB with EMI and EGI

SHA2/RFC proxies: follow-up in June and review need for plan B