

Table of Contents

Summary of GDB meeting, February 14, 2015 (CERN).....	1
Agenda.....	2
Introduction - M. Jouvin.....	3
Low Power System-on-Chip Evaluation @CNAF - D. Cesini.....	4
Actions in Progress.....	6
Report - J. Flix.....	6
LHCOPN/ONE Workshop Summary - J. Coles.....	7
Summary of pre-GDB on Cloud Traceability - I. Collier.....	8
HEP SW Foundation SLAC Workshop Summary - M. Jouvin.....	9
INDIGO Project - G. Donvito.....	10

Summary of GDB meeting, February 14, 2015 (CERN)

Agenda

<https://indico.cern.ch/event/319744/>

Introduction - M. Jouvin

2014-2015 Meetings

- March meeting is in Amsterdam and includes pre-GDB. Will close 3:30pm
 - ◆ March pre-GDB will cover other sciences discussion in morning and cloud issues in afternoon.
- April meeting cancelled
- Consideration of October GDB being at BNL alongside HEPiX. Please give feedback.

Upcoming pre-GDBs thereafter: Batch system discussions; volunteer computing and accounting.

WLCG workshop will take place April 11-12.

- Agenda <https://indico.cern.ch/event/345619/other-view?view=standard>
- Day 1: status & short term
- Day 2: Toward run 3.
- Early bird registration for CHEP is the end of the week!

Actions in progress

- Glxec in production for a few months now within CMS. About to be in ATLAS.
- Machine/job feature adoption: **looking for early adopter sites**
 - ◆ Documentation issue identified and being worked on
- Multicore accounting: good progress but still 20% sites not yet publishing multicore accounting
 - ◆ Catherine: efficiency in APEL is still crazy at French sites
 - ◆ John: to be followed offline
- perfSONAR reinstallation deadline is 16th February.
 - ◆ Quite a lot not yet properly configured.
- Little feedback on ginfo (GLUE2 client)

Upcoming meetings:

- ISGC (15-20 March);
- HEPiX 23-27 March;
- WLCG workshop (11-12 April) + CHEP (13-17 April).

Low Power System-on-Chip Evaluation @CNAF - D. Cesini

System-on-Chip (SoC) originally targeted for mobile and embedded systems but now users for more general purpose servers

- Everything except connectors on the chip (no chipset or memory)
- Strength: huge market compared to general purpose processors
- Processing power now closer to general purpose processors

Now a lot of small boards available

- Basically all implementing the Heterogeneous Multi Processors (HMP): 2 types of ARM processors
- All have an associated GPU
- Generally TDP \leq 15W per board
- Theoretical performance could be 1/2 an E5 + K40: worth testing!

HW limitations

- 32-bit
- Small memory size
- Low performance Ethernet: 10/100 Mb/s
 - ◆ Often through USB

Programming constraints

- GCC + OpenMP available for ARM
- GPU: OpenCL on most boards, CUDA on Jetson K1
- Cross compilation remains difficult
- Ongoing project: gcc v5 + OpenMP v4
 - ◆ No results yet

Tests

- Very simple algorithm like Pi computation: up to 20x perf/W ratio
- Non trivial algorithms (like prime numbers, FFT) shows that Xeon remains significantly more performant and that the perf/W ratio is around \sim 5 time better with SoC
- GROMACS (Molecular Dynamics)
 - ◆ Easy recompilation
 - ◆ Execution time on Jetson-K1 (CPU or GPU) is about 10x slower with a 20x better perf/W but these apps being long to run, cannot wait 10x more...
- Filtered Backprojection (Tomography): performance pretty similar between Xeon and Jetson K1 with a 6x perf/W improvement
- Lattice Boltzmann: very bad perf observed, still under investigation (bandwidth problems?)

SoCs also available from Intel (Atom)

- 64-bit support
- No GPU
- Avoton: one of the promising SoC
 - ◆ See presentation on HP Moonshot by A. Chierici at HEPiX
- Not as low cost as ARM-based SoCs
- Tend to maintain better perfs than ARM-based under higher concurrent load

Conclusion

- ARM-based SoCs interesting for selected apps: imaging, no high mem requirements
 - ◆ Still many limitations: 32-bit, no ECC
- Intel SoC looks promising

Discussion

- Maarten: Has anyone verified run on new infrastructure is scientifically okay?
 - ◆ DC: Yes.
- S. Lin: with 64-bit version will the overall performance improve?
 - ◆ DC: server version tests did not show great performance power consumption was very high. ARM 64-bit rather disappointing so far...
- Michel: is there still an ongoing activity in WLCG experiments?
 - ◆ Claudio: as far as known, CMS tests continuing but no new results have been shared.
- MJ: CNAF mentioned that they may purchase some for HP Moonshot in a future procurement, is it still planned?
 - ◆ DC: It was postponed. Performance interesting but price too high and not really cost effective when everything is put together (in particular impact in term of footprint, network ports...). Also when putting storage with servers, density no so much attractive.
- MJ: to be followed up. Would be interesting to hear from other sites who are doing some similar tests...

Actions in Progress

Report - J. Flix

General news

- N. Magini moving to something new: thanks for the work done!
- WLCG survey closed, answers analysed, first report expected at March GDB
 - ◆ ~100 answers received: pretty good participation, thanks to all sites who took the time to answer
- WLCG workshop: please register asap
 - ◆ An updated agenda will be produced soon (in fact was produced right after the GDB!)

Baseline changes

- New FTS: upgraded at RAL, BNL and CERN
- New Frontier/Squid: 2.7.STABLE9-22
 - ◆ DOS on squid3, not high priority but recommended
- dCache 2.6 end of support in June 2015
- MW Readiness: StoRM 1.11.5/6 under verification
 - ◆ CREAM 1.16.4 verified

MW news

- issue identified in ARGUS with latest Java (SSLv3 disabled by default), workaround available
- Also a ghost vulnerability identified and rated "high risk"

T0

- VOMRS will be decommissioned next week
- AFS UI closed

T2s:

- Request to get CERN accounts for T2 sysadmins: to be studied

Experiments

- ALICE: high activity, large data loss at SARA have been recovered
- ATLAS: Prodsys2 + Rucio now validated and stable, high activity with some hiccups
 - ◆ Production now runs mostly multicore
- CMS: large MC production on going for Runé preparation, cleanup of disk areas in T1s in progress, 50% of multicore resources requested at T1s
 - ◆ Also analysis and production Condor pool being merged: pilot will not have the production role anymore. May require batch system reconfiguration at sites
- LHCb: affected by data loss at SARA, several user files with a single replica

Task forces

- glexec: Panda testing campaign ramping up, already 54 sites
- Machine/job features: asking for volunteer sites
- IPv6 validation and deployment: by April, all T1 PS instances must be dual-stacked

- MW Readiness: good participation, sites participating to the TF asked to install WLCG Package reporter
 - ◆ Package reporter scrutinized and compliant with EGI security requirements
- Squid monitoring and http discovery
 - ◆ 154/178 squid registered in GOC/OIM
 - ◆ When restricting Stratum 1 instances reachable, need to add 2 new ones (see slides)
- Network and Transfer Metrics: perfSonar reinstallation deadline is Feb. 16, experiments encouraged to test the API to access PS data

Next meeting is February 19.

- T1s and T2s are encouraged to participate to OpsCoord meetings
 - ◆ T1s in particular should all participate

LHCOPN/ONE Workshop Summary - J. Coles

Last meeting in Cambridge beginning of this week, good attendance (46)

LHCOPN

- Most T1-T1 connections moving to LHCONE
- Increased bandwidth to T0

LHCONE

- Belle2 (PSNC) has joined LHCONE and happy with it
 - ◆ Also started to use FTS3
 - ◆ KEK not yet connected
- Brasil expressed interest in joining LHCONE
 - ◆ 4 or 5 WLCG sites
- L3VPN very stable, BGP filtering successful
- New Transatlantic links: 3x 100G links
- AUP still being discussed, several refinements
 - ◆ Open question about use of LHCONE by ATLAS/CMS T3 sites who didn't sign the WLCG MoU
- P2P service: some progress (thanks to collaboration with AutoGOLE) but still several issues, in particular with routing

Brocade set a special deal for all WLCG sites with special discounts accessible through revendors

Summary of pre-GDB on Cloud Traceability - I. Collier

Traceability goal: contain and limit the impact of security issues

- Preserve our reputation
- Ensure resources are used for the purposes intended

Answer the question: who did what? when? where?

Good pre-GDB participation with a good mix of sites, security teams and experiments

- 12 local + 8 remote
- Meeting notes: <https://twiki.cern.ch/twiki/bin/view/LCG/20150210PreGDB>

Issues discussed

- Just use syslog to collect info about what is done in the VM
 - ◆ Still discussion on machine/job features vs. contextualization to pass syslog info
- Hypervisor and netflow logging: externally observable behaviour, complementary to syslog
 - ◆ A survey of actual experience and possible recommendations will be done: work led by Raul Lops and David Crooks
- Giving sites root access: general agreement to make it possible in the WLCG context
- Quarantining VMs if possible
 - ◆ Useful for VM forensics and also to troubleshoot strange behaviours or other non security related issues (e.g. misconfigured CVMFS)
 - ◆ Currently built-in in StratusLab only, will investigate for other cloud MW
 - ◆ Has a cost in storage: has to find the right balance
- Classification of VMs for incident response
- Policy evolution: recognize the increased role of VOs
 - ◆ Exact work will depend of the outcome of other tasks
- VO logging gap: what is missing in VO logs to make them useful for security traceability
 - ◆ Proposal of a traceability challenge: NIKHEF agreed to lead

Logging remarks

- Cross checking multiple source is crucial
- VM images well controlled by VOs: allow more trust on what is done in the VM
 - ◆ In particular user payload not run as root and under a different user as the 'supervisor' (pilot) user
- Use site syslog rather than a central syslog: less policy implication, do not put load on VOs

Contact Ian Collier if you are interested to participate in this work

Discussion

- Helge: are classes based on VM lifetime. If services running on machines may be more vulnerable, whatever the lifetime.
 - ◆ IC: classes based on what is run, lifetime almost irrelevant...
 - ◆ Dave K.: the risks and threats are different. Classifications may be important for policy.
 - ◆ IC: classification is quite a distinct issue from traceability. More related to incident handling
 - ◆ Michel: general agreement that VM hosting services should be treated the same way a bare metal machine running the same service as far as incident handling is concerned.

HEP SW Foundation SLAC Workshop Summary - M. Jouvin

Reminder of goals: facilitate coordination and common effort in HEP software and computing.

SLAC workshop:

- 100 people attended (80 local + 20 remote)
- Several non pure HEP experiments represented
- Agenda: <http://indico.cern.ch/event/357737/other-view?view=standard>

Main sessions

- Learning from others
 - ◆ Apache Software Foundation similar umbrellala organization; but it started before projects; do-craacy (active people have say); Darwinian approach; ASF incubates. Transparency essential.
 - ◆ D. Katz: nice summary on lessons learnt from other projects. Try-and-fail most productive approach. Give credit. Flat governance model. Get people involved (to avoid reinvention).
 - ◆ Software Sustainability Institute (UK). Same message as D. Katz don t try to perfect HSF. Training focus. Lobbying/communication for Software Engineers.
- Community & project views: Agreement HSF could help. No real conflicts of view but different focus areas.
- Non-topic: Governance it should be a light structure without too much formal management.
- Next steps
 - ◆ guinea pig projects. Try incubator idea.
 - ◆ Technical forum: started as a Google group. Split into smaller, focused, groups if/when needed. Publish "technical notes" as a way of sharing expertise.
 - ◆ training. Consensus this should be an initial HSF focus.
 - ◆ Services: software knowledge base. Prototype now exists. Register your favourite software.
- Many open questions: Licensing; Consultancy SWAT teams ; Access to scientific journals.
- Conclusion: Productive meeting. Next milestone CHEP face-to-face meeting. Encourage people to join!
 - ◆ Membership is mainly based on the mailing list membership: see <http://hepsoftwarefoundation.org> for details.

INDIGO Project - G. Donvito

INDIGO: INtegrating Data Infrastructures for Global Operations

- INDIGO DataCloud funded as part of EINFRA-1-2014 (item 4 and 5)
- 5 private companies among the partners
- Lead by INFN
- Several scientific communities involved
 - ◆ Biological, medical, social science, arts and humanities, environmental and earth, physical sciences

Gap analysis / main goals

- Federated identity support
- Performance limitations limiting adoption of clouds in large data centers
- Orchestration and federation of cloud, grid and HPC resources
- Non interoperable interfaces preventing adoption of PaaS
 - ◆ Particularly of true for storage access
- Lack of flexible data sharing between group members
- Static allocation and partitioning of both storage and computing resources
- Integration of specialized HW like GPUs
- Inflexible ways of distributing/deploying applications
- Service discovery and monitoring
- Support for containers in addition to VMs
- Enhance cloud schedulers to allow handling of different load priorities and achieve "fairshare"
- Network virtualization: local virtual networks, SDN
- Authz: role/group management, policy composition
- Science gateways

INDIGO will deliver SW based on open-source solutions

- Adopt well established solutions when available
- Prefer extending existing ones rather than starting new solutions
 - ◆ Look at presentation for details: many components already available and only requiring an extension

Project facts

- Duration: 30 months
- Budget: 11.1 Meuros.
 - ◆ 1500 personxmonths
- Likely start April 2015

Discussion

- Michel: Are you aware of CYCLONE project? The project has overlaps with some of the packages mentioned, in particular federated identity support and network virtualization
 - ◆ GD: no. Happy to see how we can collaborate

-- MichelJouvin - 2015-02-11

This topic: LCG > GDBMeetingNotes20150211
Topic revision: r1 - 2015-02-11 - MichelJouvin



Copyright &© 2008-2019 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.
Ideas, requests, problems regarding TWiki? Send feedback