Latest Procesor-related Updates and COLA Workshop Report

David Abdurachmanov (FNAL)

March 14, 2016

The presentation was compiled on 5” Nexus 5 smartphone!

Architecture Updates #1

- **PRoot + QEMU + rootfs** (Fedora 22 or CentOS 7.2) solution is now used for installing CMSSW IBs and releases into CVMFS (**cms-ib.cern.ch** or **cms.cern.ch**)
  - SCRAM project relocation memory consumption issue is resolved (from **~12.9G** down to **~20MB** of **RSS**)
  - Giulio from ALICE also picked it up and made Docker arm64 builder image (to be improved not to use PRoot)
- CMSSW_8_0_0, CMSSW_8_0_1, CMSSW_8_0_2 should be available for aarch64 and ppc64le on cms.cern.ch CVMFS repository.
  - They there used to validate memory consumption fix, Jenkins job setup and Bockjoo installation scripts
- CMSSW_8_1_0_pre1 is the first build to be fully automatically built and installed
- My yum repository with CVMFS pre-2.3.0 build (Feb 14) for CentOS 7.2 aarch64 is featured on CVMFS download page

Architecture Updates #2

- CVMFS team ordered Geekbox (to be supported by Fedora 24) as private CI machine
- We have another m400 machine prepared to be used for HLT workloads testing and investigation
- Experiments are already asking for LCG builds for aarch64 at meeting with CERN IT (most likely only available via CVMFS, no AFS)
- GCC 6.1.0 is one month away and we had first CMSSW build (private) on x86_64 done a week ago:
  - One P2 confirmed and possible resolved issue in GCC; one issue in Boost, but initial “fix” is committed, but is pending my futher testing
  - Reminder: this makes C++14 and C11 as default languages, but still using old C++ ABI; some progress was done to support new ABI in Clang 3.9 (6 months away)

Architecture Updates #3

- Meeting with Intel in early April to discuss next steps for CMSSW + ICC builds
  - One is to get newer CMSSW version delivered to Intel for ICC internal testing
- Similar could be done with ARM HPC group in future; there are interested people
- There are some discussions between CERN Techlab, CMS and CentOS for providing access to some hardware are CERN to speed up upstreaming efforts and getting CentOS on new hardware
  - Suggestion was to start this at Linaro Connect as hardware and people are available at the event
- Visit DOME project in coming months in Geneva to learn and test their super-high-density hardware prototypes for SKA big data challenge
- Some new NDA and NDA-free aarch64 hardware will be available at CERN in near future
- Initial near complete rebuild of OSG repository for aarch64 was done (prep work)

COLA Workshop

Workshop on **Computing on Low-Power Architectures** where **power**, **cost** and **space** are main **limiting factors**

Notes #1

- Main metrics are Energy-to-Solution (mJ or J), Time-to-Solution (s.), CPU and memory frequency (MHz)
- CPU freq governors and C-States are also being tested
- The broad picture is to allow tuning hardware for your workflow, i.e. deeper integration
    - Generating Power-:Performance profile (predict run time and avg power based on CPU freq) to select CPU frequency for workload. E.g., default is 2.3GHz and you get to run at 2.7GHz is there is >12% increase in performance
- No common solution thus in some cases people are willing to spend 6 to 12 months rewriting their software to new architectures and in some cases it's not possible (lack of people, validation, other resources?)

Notes #2

- Montblanc project provides ability to query power profile for your executed job. Similar judge system (think top coder like systems) are developed by Finish. You can exploit such information to figure out which method is more efficient OpenCL, OpenACC, OpenMP, etc for specific tasks.
- People like ability to monitor power consumption in different domains, but also also require measuring full box. At the end full box is what matters.
- IBM has patented a energy aware scheduling idea/technology thus details were not shared (be careful developing such systems).
- There is interested in energy aware frameworks (e.g. turn off cores if doing data transfers or lower the frequency)
    - Can you application do better CPU control than OS kernel utilising workflow specific knowledge?

Notes #3

- In some cases people are looking for specific hardware, e.g. single low-power SoC which can drive 8-16 GPUs without neededing to transfer data via InfiniBand between cards.
- Low-power is not about just low-power general purpose cores, but about having highly efficient accelerators (GPUs, FPGAs, custom ASIC blocks, etc).
- We need silicon manufacturers to look into our workloads and there was a suggestion to form a group around this idea.
- There will be another ARM HPC User Group meeting in summer, ISC.

DOME (SKA):
It’s all about SoC and packaging!

Q & A