Heterogeneous Architectures Goals

This is a list of important topics for discussion at the HSF-WLCG Pre-CHEP workshop for the *Non-x86 and Heterogeneous Computing* session on Sunday.

The outcome of the workshop will be a short document on the status re. the above questions and some view of the next actions to be taken to advance our use of different architectures.

Meta-question: *When will WLCG be able to effectively use non-x86 hardware?*

- Experiment software
    - Can non-x86 CPUs be used?
        - Software can be built (CERN support?)
        - Software can be validated (CERN support?)
        - Software can be distributed (CVMFS)
        - Jobs can be steered to sites and run correctly (Production system readiness)
    - Can GPUs be used?
        - Ditto to above
        - With the caveat that pure GPU workloads are unlikely to be available for a while
            - Mixed mode workloads?
            - Can we load balance?
        - Can we adapt to the breadth of non-CPU devices?
            - Will sites have earlier useful options? (Nvidia, being the most likely, but non-Nvidia are interesting for testing)
- Common software
    - Geant4
        - Yes, it's ok for ARM
        - For GPU still highly R&D (AdePT, Celeritas) and will only be partial offload (workload balance again)
    - ROOT
        - OK for ARM
        - Nothing on GPU
    - Generators
        - MG4GPU project
            - Still only LO at the moment, but NLO in the pipeline

General maturity of aarch64 linux distributions seems good.
Are we only caring about ARM, or is PowerPC still interesting? RISC-V is getting attention (strategic from the EU).