KTH ROYAL INSTITUTE OF TECHNOLOGY

Investments in e-infrastructure in Sweden: PDC

Dirk Pleiter

SeRC
Swedish e-Science Research Centre

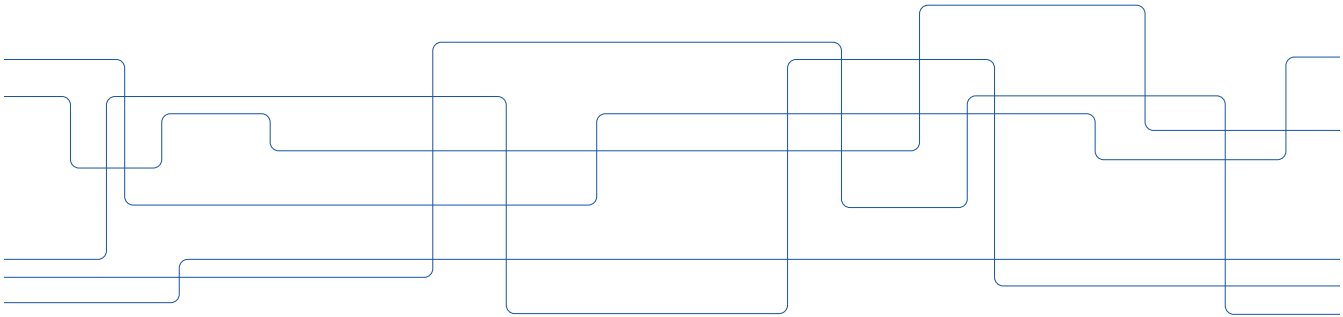

Next System at PDC: Cray XE System Dardel

Key performance features

- More than 13.5 PFlops/s
  - CPU partition with about 110,000 cores (~70% procured by SNIC)
  - Most performance delivered by a 56 GPU-accelerated nodes
- Lustre-based parallel file system
  - 8 PByte usable capacity
  - 180 GByte/s bandwidth

Key Architectural Features

AMD EPYC CPUs with large number of cores per socket

- 64 cores per socket → 128 cores per node
- Less memory bandwidth and (on most nodes) less memory capacity per core compared to Beskow

Future generation of AMD Instinct GPUs

- 4 GPUs per node interconnected with high-bandwidth links
- Significant increase in throughput of Flops per GPU
- Very high memory bandwidth and significant increase of memory capacity per GPU

Cray Slingshot network technology

- Based on Ethernet, but optimised for HPC
- Very high link bandwidth: 200 Gbit/s
  - 1-4 network ports per node
- Dragonfly topology with dynamic routing
  - Moderate to low bi-section bandwidth between switch-groups

GPU-Accelerated High-Performance Computing

Why using compute accelerators?

- GPUs allow to provide much higher throughput of Flops within a given power envelope
  - Green500 (November 2020)
    - > *#1: NVIDIA DGX SuperPod at 26 GFlop/s/W*
    - > *#6: A64FX Prototype at 24 GFlop/s/W*
    - > *#44: Bell (Dell, AMD EPYC) at 5.2 GFlop/s/W*

Why using AMD instead of NVIDIA GPUs?

- Better price/performance in this round
- AMD entering the HPC market will create new completion
- BUT:
  - NVIDIA software ecosystem much better compared to AMD
  - Less experience available on how to optimise for AMD architectures

→ **Computational science community needs stronger focus on performance portability**

  - Experience: good investment improving sustainability of software

Timeline and Support

Planned start of use access

- Phase 1: September 2021
  - Likely earlier access for selected users
- Phase 2: January 2022

Support

- Standard support models via SNIC and PDC
- Dedicated training events together with HPE and AMD
  - Current focus on trainer-of-trainers
- Annual hackathons
- Support services from HPE

Thank you!