SC19 briefing notes

J. Simone

Intel Ponte Vecchio GPU and OneAPI SW

Promotional keynote at Intel sponsored HPC Developer Conference by Raja Koduri: senior VP, chief architect, and GM of Architecture, Graphics, and Software at Intel Corporation.

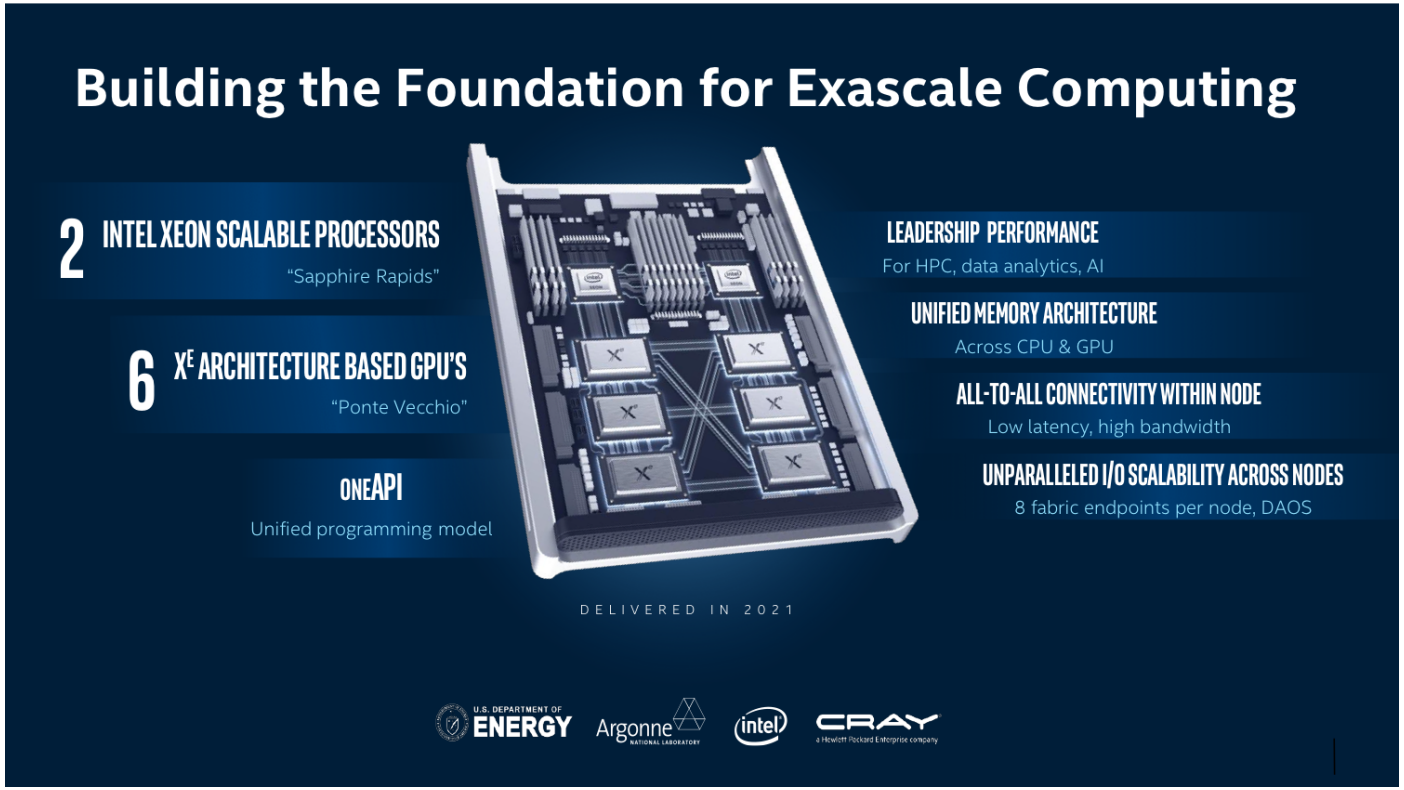

- Official announcement of the Xe Graphics Architecture spanning mobile, desktop, and HPC.
- First generation “Ponte Vecchio” product designed for HPC and AI (Aurora exascale).
- Low latency, high BW interconnect on node permits memory coherency.
- “Sapphire Rapids” CPU in 2021

OneAPI

Intel promoted software tools – open standard

Bundles many familiar Intel SW products under one umbrella

Data Parallel C++ (DPC++) an evolution of C++ based on SYCL with Intel extensions
- Platform model
- Data management
- Expressing || - kernels

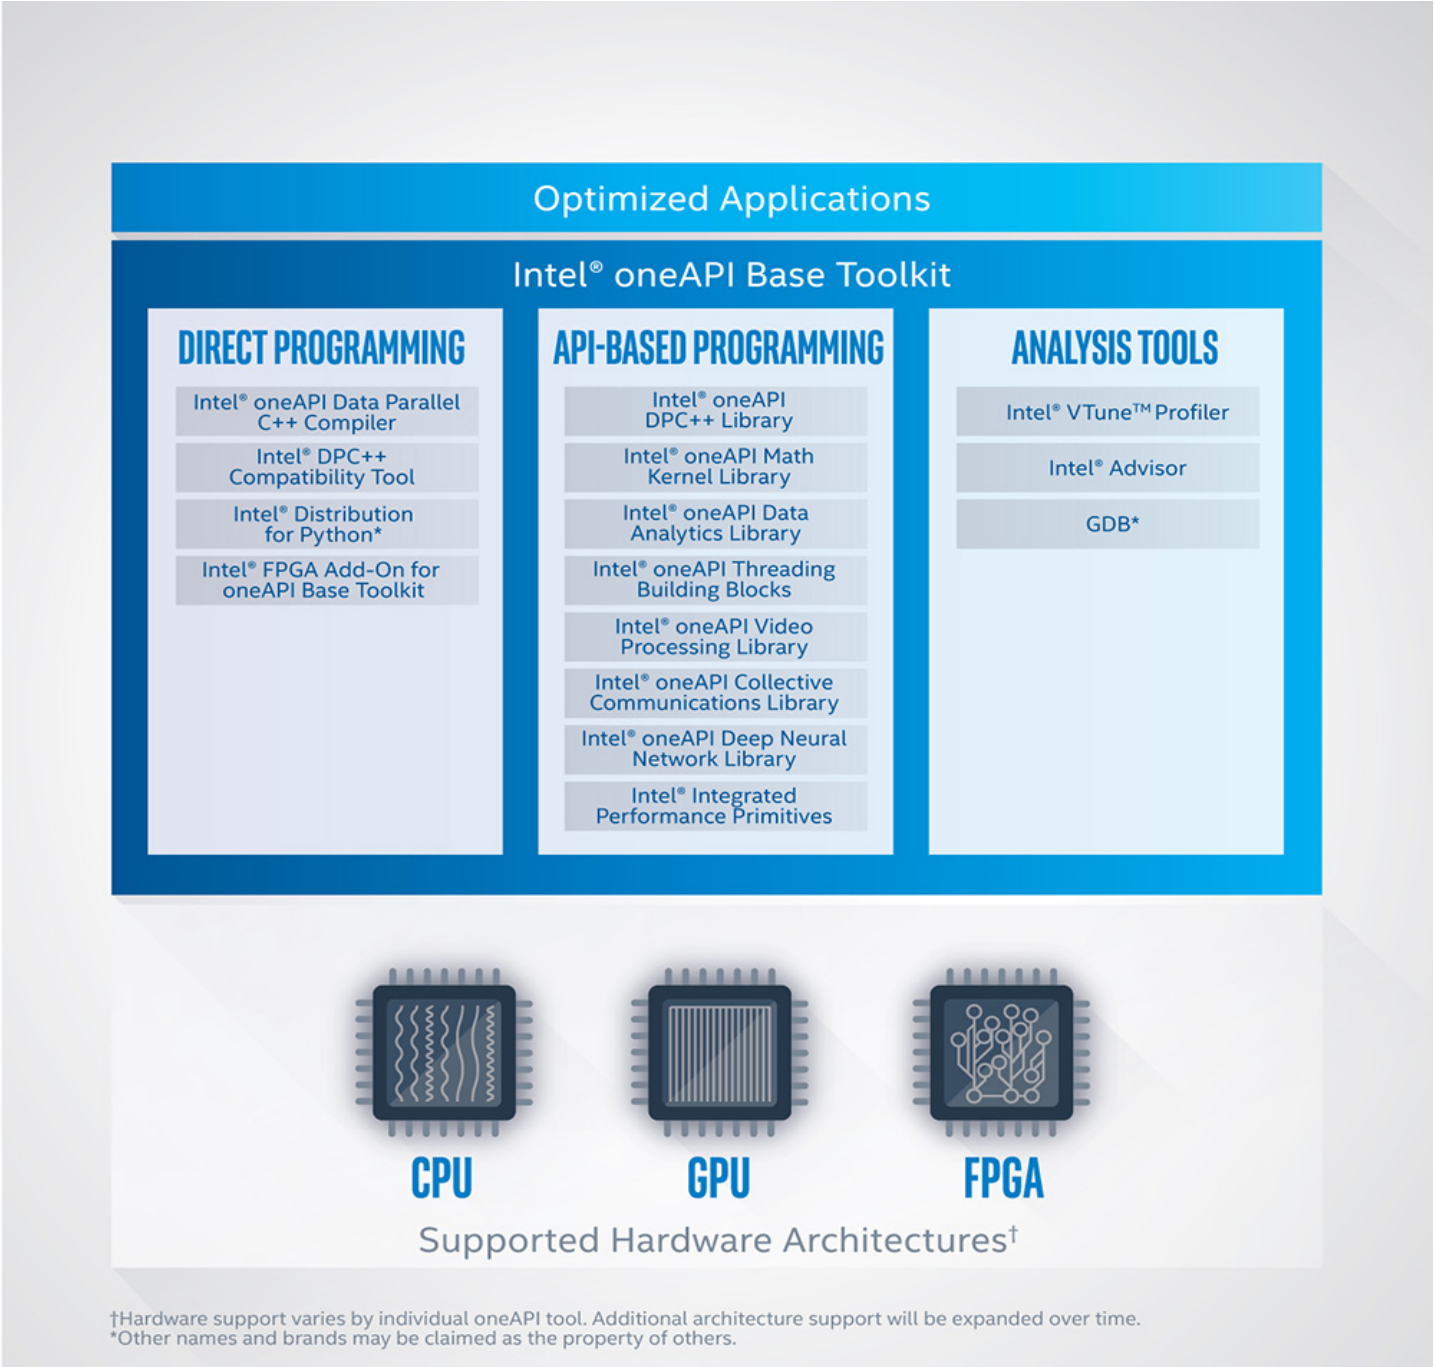

Learn DPC++ and oneAPI

Book preview and online trainining available from https://jamesreinders.com/dpcpp/

Free accounts on Intel Devcloud https://software.intel.com/en-us/devcloud/oneapi provides a pre-configured development sandbox to test code.

Hardware: Xeon (CPU+GPU) and Arria (FPGA)

Lustre BOF

- Lustre is > 20yrs old, still leading in features
- Will be used at Aurora and Perlmutter installations. NERSC will deploy a 30 PB all VME instance
- Lustre v2.10 is most used in production
- Lustre v2.12 LTS is current, supports recent kernels
- Integrated Manager is GA; simple, powerful management / monitoring tools and dashboard.
- Under development:
  - Small files stored on MDS
  - Persistent VME cache local to clients, use HSM to manage files

SLURM BOF

- I spoke with staff at SchedMD booth. I will follow up on possibility of arranging SLURM half/full day training on an hourly basis by a consultant.
- Current release 19.05
  - Plugin ‘cons_res’ will be removed in the future
  - Can move to newer, better ‘cons_tres’ (consumable trackable resource, e.g. GPUs, CPUs, memory) plugin without losing queued jobs
  - There is also a ‘nersc_cli_filter’ that runs slow computational checks on job submissions on the client side. Example: checking resource quotas requested at job submission. Users can bypass checks with some ingenuity!
- Next release is 20.02
  - REST API to slurmctl
  - No more *slurm.conf* file, slurmd processes do RPCs to server for configuration
- Only the last two tagged releases are to be supported