




TSUBAME3.0 Compute Node

High Density Using Water Cooling

Compute Node



TSUBAME3.0 compute nodes are designed for HPC, BigData, and AI workloads which require all of computing power, memory bandwidth, network interconnect, and local storage. All of these components are densely installed in a compact blade using water cooling.




Tesla S1070 (Tesla GT200) on Tsubame1.2



Tesla M2050 (Fermi) on Tsubame2.0



Tesla K20X (Kepler) on Tsubame2.5



Tesla P100 (Pascal) on Tsubame3.0


TSUBAME3.0 continues accelerated computing using GPU.

Implementation

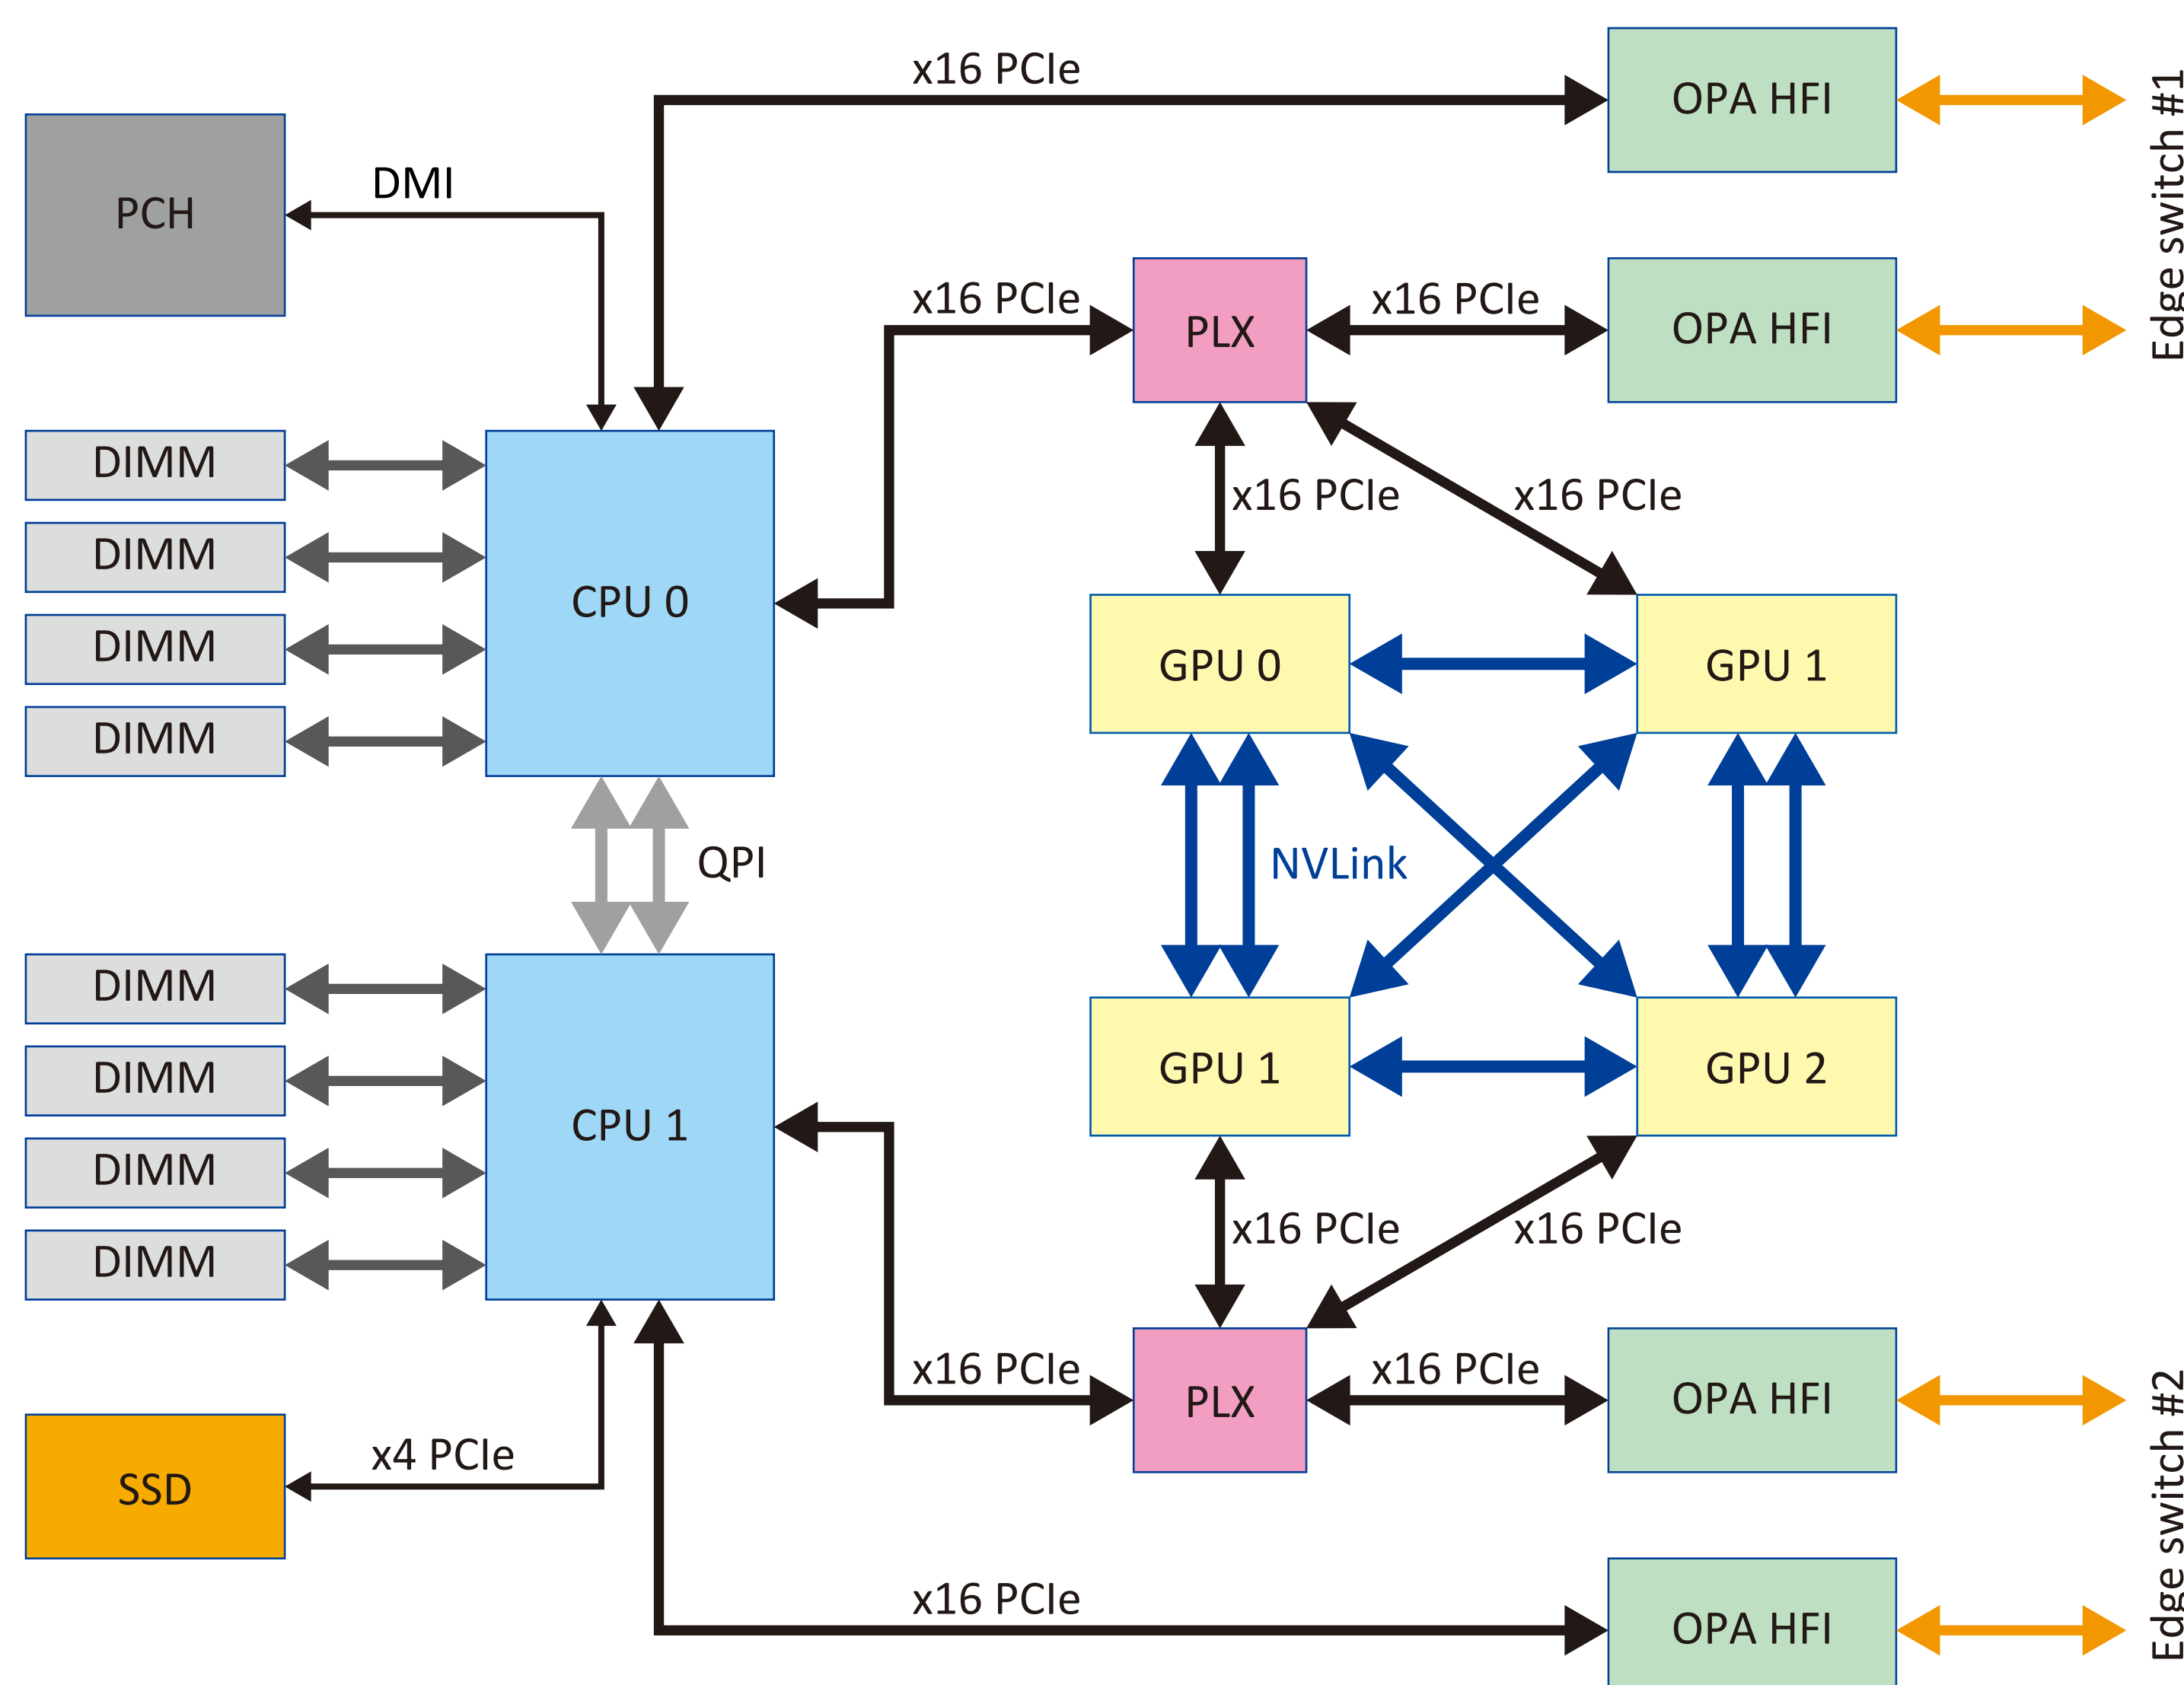
Specifications

SGI ICE-XA (HPE SGI 8600) IP139-SXM2

- CPU: Intel Xeon E5-2680 V4 (Broadwell-EP, 2.4GHz) x2 sockets
14 cores per socket, total 28 cores per node.
- GPU: NVIDIA TESLA P100 for NVLink-Optimized servers x4
Pascal GPU. 5.3TFlops, 16GB HBM2 @ 732GB/s.
- Mem: 256GB (DDR4-2400 32GB module x8)
- SSD: Intel DC P3500 2TB (NVMe, PCI-E 3.0 x4)
sequential read 2700MB/s, sequential write 1800MB/s.
- Net: Intel Omni-Path Architecture HFI (100Gbps) x4



Block Diagram



Compute node consists of four layers: CPU and Memory on 1st, OPA HFI (connected to CPU) and GPU board on 2nd, OPA HFI (via PLX) on 3rd, and Power board on 4th.

CPU and GPUs are directly water cooled. The other components are air-cooled indirectly using water.

We will operate Tsubame3.0 until March 2024.
We are designing its successor, Tsubame4.0.
The specification will be announced very soon.

