Institute of SCIENCE TOKYO

Center for Information Infrastructure

TSUBAME4.0 Supercomputer

More of Everyone's Supercomputer, In Operation Since April 2024

System Overview



Components of TSUBAME4

Compute Node Racks
- HPE Cray XD665
- Direct Liquid Cooling + Rear-Door Cooling (20℃)

Network
- InfiniBand NDR200, 4 ports / node
- Full-bisection Fat Tree Topology

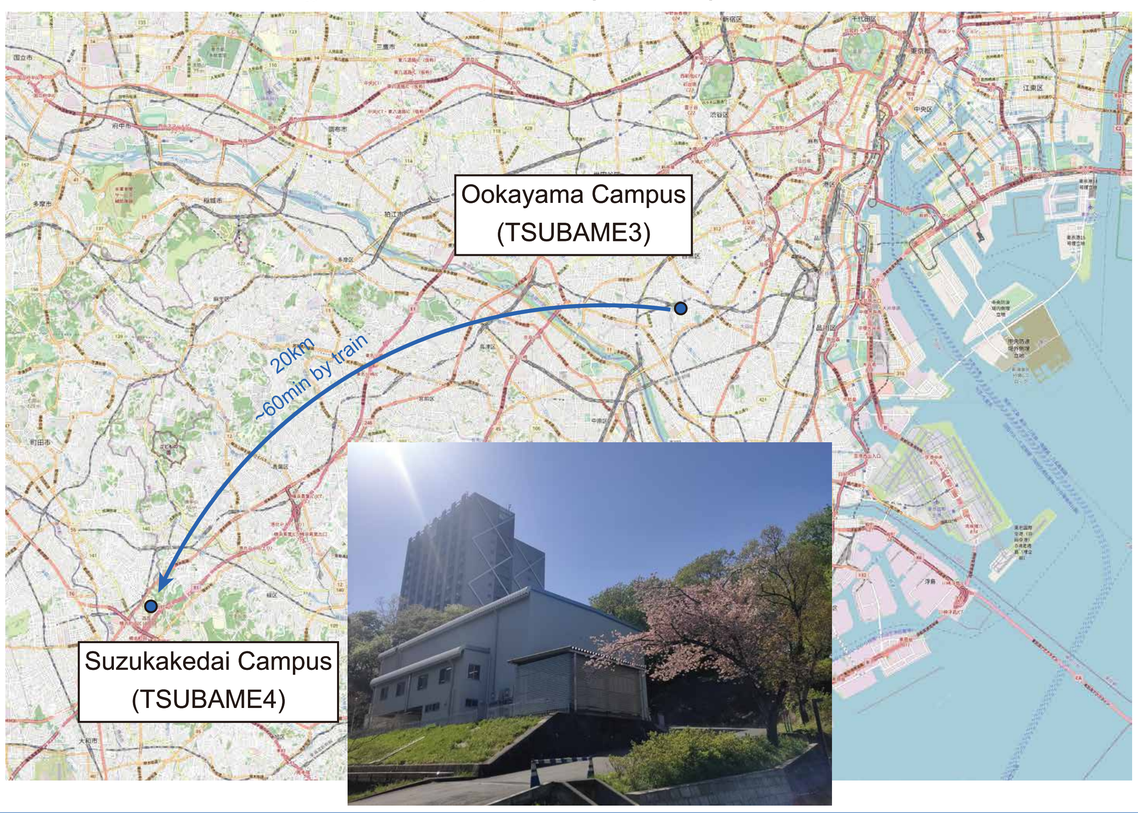
Storage
- HPE ClusterStor E1000
  - HDD 44.2PB, 280GB/s
  - SSD 327TB, 325GB/s
- Node-Local SSD: 460.8TB (in total), Read 1,632GB/s, Write 648GB/s

Comparison with previous generation

	TSUBAME3.0(2017-)	TSUBAME4.0(Apr 2024-)
Computational Performance		
• FP64 Matrix	12PFlops	**66.8PFlops (5.5x)**
• FP64 Vector		**34.7PFlops (2.8x)**
• Deep Learning (FP16 Matrix)	47PFlops	**952PFlops (20x)**
GPU Memory Bandwidth	1.56 PB/s	**3.07 PB/s (1.97x)**
Number of Nodes	540 Nodes (homogeneous config)	**240 nodes (homogeneous config)**
GPUs	2,160 NVIDIA P100	**960 NVIDIA H100**
Cooling / Inlet Water Temperature	Free Cooling with Cooling Tower, 32°C	**Chiller, 20°C**
Power Consumption (incl. cooling)	1,080kW (Spec. value) 400~600kW(Operation)	**1,820kW (Spec. value) 650~850kW(Operation)**

Location and Data Center Facilities

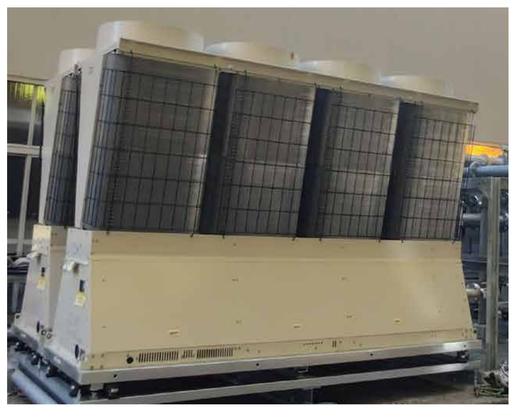
Unlike past TSUBAME series, TSUBAME4.0 is hosted at Suzukakedai campus, another campus of Science Tokyo.
In order to accommodate heavy water-cooled supercomputers, we renovated a building exclusively for supercomputers.



The new supercomputer room is 200m² in area, 30kN/m² load capacity, and 4m ceiling height.
The facility provides three-phase 415 V power with a capacity of 2,000 kVA.



All TSUBAME4.0 components are cooled by a single cooling system consisting of eight chillers. Each chiller has a cooling capacity of 200 kW and supplies 20 °C water in a 30 °C environment.
Processors in compute nodes are directly cooled. Other components, including storage and management nodes, are indirectly cooled by rear-door heat exchangers.





https://www.isct.ac.jp/