




TSUBAME2.0

Architecture & Services

Hardware

Storage


Parallel file system area : 5.93PB



MDS,OSS
HP DL360 G6 30nodes
Storage
DDN SFA 10000x5
(10 enclosures x5)
Lustre (5 Filesystems)
OSS:20 OST:5.9PB
MDS:10 MDT:30TB



Users' home space : 1.2PB




Storage Server
HP DL380 G6 4nodes
BlueArc Mercury 100 x2
Storage
DDN SFA10000 x1
(10 enclosures x1)



Archive Storage : 4PB

StorageTek SL8500 x2



Interconnect: Full bisection / Non-blocking

Core Switch

12switches



Voltaire Grid Director 4700 12switches
IB QDR : 324ports

Edge Switch

179switches



Voltaire Grid Director 4036 179switches
IB QDR : 36ports

Edge Switch(w/10GbE)

6switches



Voltaire Grid Director 4036E 6switches
IB QDR : 34ports
10GbE : 2ports

National Grid Federation

HPCI High Performance Computing Infrastructure

GSI, Shibboleth, Gfarm, ...

Administrative servers

Compute nodes : 2.4PF (CPU+GPU) / 224.69TF(CPU)

Dual Rail Connection

Thin compute node



1408 nodes

HP ProLiant SL390s G7

CPU : Intel Xeon E5670 x2
(2.93 Ghz, 6 cores, 76.7 GFLOPS)
GPU : NVIDIA Tesla M2050 x3
(515 GFLOPS, 3 GB VRAM)
Memory : 58 GB or 103 GB
SSD : 120 GB, 220 GB or 240 GB
OS : SUSE Linux + Windows HPC

Medium/Fat compute node



Medium : 24 nodes

Fat : 10 nodes

Medium

Fat

	Medium	Fat
Node	HP ProLiant DL580 G7	
CPU	Intel Xeon E5670 (2.0 Ghz) x4	
GPU	NVIDIA Tesla S1070 / M2070	
Memory	137 GB	274/548 GB
SSD	480GB	600 GB
OS	SUSE Linux	

Software

System Software

OS SUSE Linux Enterprise Server, Windows HPC
Job Scheduler PBS Professional
Compiler Intel Compiler / PGI CDK / gcc
MPI OpenMPI / MVAPICH2 / MS-MPI

ISV Applications (Frequently Used)

ANSYS Fluent, ABAQUS, MD Nastran, Patran, Marc & Mentat, Dytran, LS-DYNA, Gaussian, AMBER, Materials Studio, Discovery Studio, Mathematica, MATLAB, AVS Express

Green: GPU Enabled


Software for GPU Computing

Computing Architecture CUDA 4.1 / 3.2
Automatic Parallelization PGI Accerator Compiler
GPU debugger TotalView Debugger for CUDA
Linear Algebra Library for GPU CULA


TSUBAME Users

TSUBAME is open to Academia and Industries. Industrial use starts since FY2007.

TSUBAME Users



Partnership Resource Allocations for Industrial Projects



HPCI - High Performance Computing Infrastructure


National grid infrastructure for HPC research

Resources

9 supercomputers in Japan, including K & TSUBAME2.0


Services

- One-stop signup to all resources
- Single sign on to all resources using Shibboleth & GSI
- 100PB global shared storage to share data
- Distributed VM hosting service to help developing network and distributed systems (PLAN, GSIC is leading the project)



Status

- Operation starts from Sep. 2012
- Deployment starts from Jan. 2011
- 227 proposals, 62 are accepted
- 6 projects and 142 users use TSUBAME for HPCI purpose



HPCI High Performance Computing Infrastructure

<https://www.hpci-office.jp/> (Japanese Only)

<http://www.gsic.titech.ac.jp/sc12>

