HPCI Computing Resource Handbook

RIST
Registered Institution for Facilities Use Promotion / The Representative for HPCI Operation
Research Organization for Information Science and Technology (RIST)

Jan. 2023

High-Performance Computing Infrastructure (HPCI)

The HPCI initiative established by Japan's Ministry of Education, Culture, Sports, Science, and Technology (MEXT) promotes supercomputing resources for superlative research infrastructure. The HPCI initiative provides an innovative shared computing environment that meets diverse user needs by connecting world-class, advanced supercomputers and storage devices installed at Japanese universities and research institutions, including Fugaku, through a high-speed network.

The organizations operating HPCI systems

RIKEN Center for Computational Science
National Institute of Informatics
Information Initiative Center, Hokkaido University
Cyberscience Center, Tohoku University
Center for Computational Sciences, University of Tsukuba
Joint Center for Advanced HPC (JCAHPC)
Information Technology Center, The University of Tokyo
Global Scientific Information and Computing Center, Tokyo Institute of Technology
Information Technology Center, Nagoya University
Academic Center for Computing and Media Studies, Kyoto University
Cybermedia Center, Osaka University
Research Institute for Information Technology, Kyushu University
Center for Earth Information Science and Technology, Japan Agency for Marine-Earth Science and Technology(JAMSTEC)
Center for Engineering and Technical Support, The Institute of Statistical Mathematics
Information Technology Research Institute, National Institute of Advanced Industrial Science and Technology (AIST)

CPU Architectures

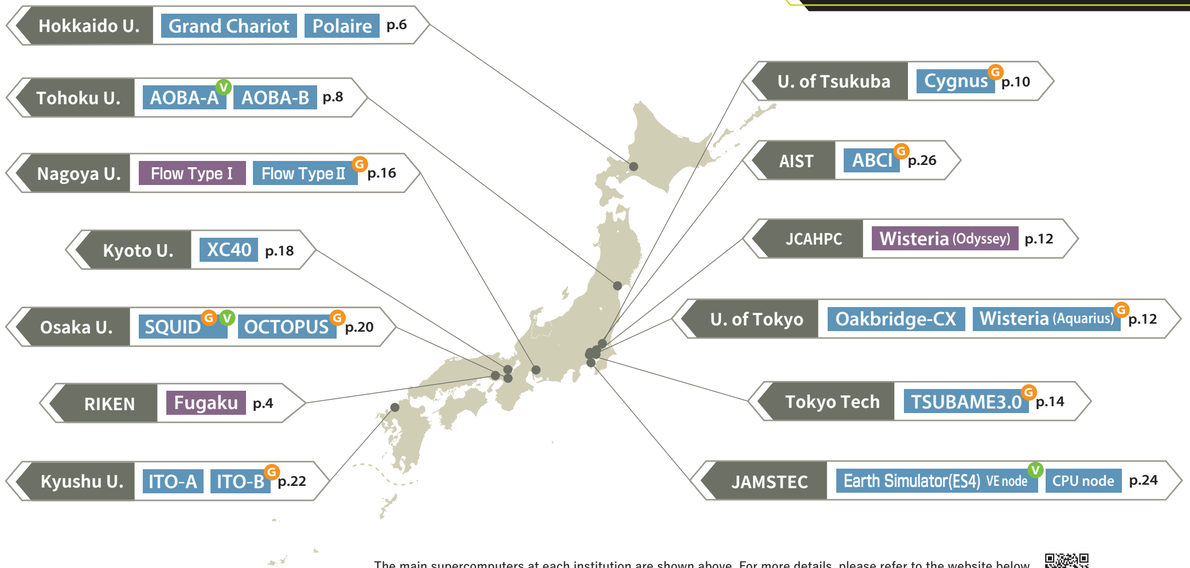

The main supercomputers at each institution are shown above. For more details, please refer to the website below.
https://www.hpci-office.jp/en/using_hpci/hardware_software_resource

CPU Architectures

Xeon(x86-64)	Processors with x86-64 instruction sets manufactured and sold by Intel for servers or workstations
EPYC(x86-64)	Processors with x86-64 instruction sets designed and developed by AMD based on the Zen microarchitecture
A64FX	Fujitsu Arm microprocessors compliant with Armv8.2-A SVE
GPU	x86-64 host processors with NVIDIA HPC GPUs for acceleration
Vector	x86-64 host processors with NEC's Vector Engine for acceleration

Supercomputing Resources by Software

Data is current as of January 2023, but subject to change.
Please contact our help desk for the latest information.
helpdesk@hpci-office.jp

	Name of Software	p.4 RIKEN	p.6 Hokkaido U.		p.8 Tohoku U.		p.10 U. of Tsukuba	p.12 U. of Tokyo JCAHPC	p.12 U. of Tokyo		p.14 Tokyo Tech	p.16 Nagoya U.		p.18 Kyoto U.	p.20 Osaka U.		p.22 Kyushu U.	
		Fugaku	Grand Chariot	Polaire	AOBA-A (SX)	AOBA-B (LX)	Cygnus	Wisteria (Odyssey)	Oakbridge-CX	Wisteria (Aquarius)	TSUBAME 3.0	Flow Type I	Flow Type II	CRAY XC40	OCTOPUS	SQUID	ITO-A	ITO-B
Molecular Dynamics	AMBER										○	○	○				○	○
	feram								○									
	GENESIS	○	○			○	○	○	○		○	○	○	○	○		○	○
	GROMACS	○	○	○		○		○	○	○	○	○	○	○	○	○	○	○
	LAMMPS	○				○		○		○	○	○	○	○	○	○	○	○
	MODYLAS	○	○			○		○	○	○	○	○	○	○	○	○	○	○
	MyPresto		○	○														
	NAMD		○	○							○	○	○	○				
	Tinker										○							
	N2P2	○																
Quantum Chemistry	ABINIT-MP	○	○		○	○		○	○	○	○	○	○	○	○	○	○	○
	BoltzTrap													○				
	GAMESS		○	○							○		○	○	○	○	○	○
	Gaussian	○	○			○					○	○	○	○	○	○	○	○
	GRRM					○												
	Molpro																○	○
	NTChem	○	○			○		○	○		○	○	○	○	○	○	○	○
	NWChem	○						○		○				○				
	Scigress																○	
	SMASH	○	○			○		○	○		○	○	○	○	○	○	○	○
Condensed Matter Physics	ABINIT	○												○				
	AkaiKKR	○	○			○	○	○	○		○	○	○		○		○	○
	ALAMODE	○	○			○		○	○		○	○	○		○		○	○
	ALPS								○									
	CP2K	○						○		○	○							
	HΦ	○	○		○	○		○	○		○	○	○	○	○	○	○	○
	mVMC	○	○			○		○	○		○	○	○		○		○	○
	OpenMX	○	○			○		○	○	○	○	○		○	○	○	○	○
	PHASE/0	○	○		○	○		○	○	○	○	○	○	○	○	○	○	○
	Phonopy	○	○			○		○	○		○	○	○		○		○	○

Data is current as of January 2023, but subject to change.
Please contact our help desk for the latest information.
helpdesk@hpci-office.jp

	Name of Software	p.4 RIKEN	p.6 Hokkaido U.		p.8 Tohoku U.		p.10 U. of Tsukuba	p.12 U. of Tokyo JCAHPC	p.12 U. of Tokyo		p.14 Tokyo Tech	p.16 Nagoya U.		p.18 Kyoto U.	p.20 Osaka U.		p.22 Kyushu U.	
		Fugaku	Grand Chariot	Polaire	AOBA-A (SX)	AOBA-B (LX)	Cygnus	Wisteria (Odyssey)	Oakbridge-CX	Wisteria (Aquarius)	TSUBAME 3.0	Flow Type I	Flow Type II	CRAY XC40	OCTOPUS	SQUID	ITO-A	ITO-B
	Quantum ESPRESSO	○	○		○	○		○	○	○	○	○		○	○	○		
	SALMON	○	○			○		○	○		○	○	○	○	○	○	○	○
	SIEATA	○																
	xTAPP								○									
Computational Biology	Alphafold										○		○					
	rDock	○																
	Relion														○	○		
Fluid Analysis	FLASH code														○			
	FrontFlow/blue		○			○		○	○	○	○	○	○	○	○		○	○
	FrontFlow/red	○	○	○								○	○					
	OpenFOAM	○	○	○		○		○	○	○	○	○	○	○	○	○	○	○
	V-FaSTAR		○															
	FDS	○																
Structural / Collision Analysis	FrontISTR	○	○			○		○	○	○	○	○	○	○	○		○	○
	LS-DYNA												○					
	Marc																○	○
	Mentat																○	○
	MSC Nastran																○	○
	Patran																○	○
Electromagnetic Field Analysis	Meep		○	○										○				
	OpenFDTD	○																
Multi-physics	Freefem++														○			
	HyperWorks												○					
Particle Systems	GEANT4							○		○		○	○					
Weather / Climate	SCALE	○																
	WRF	○	○	○										○				

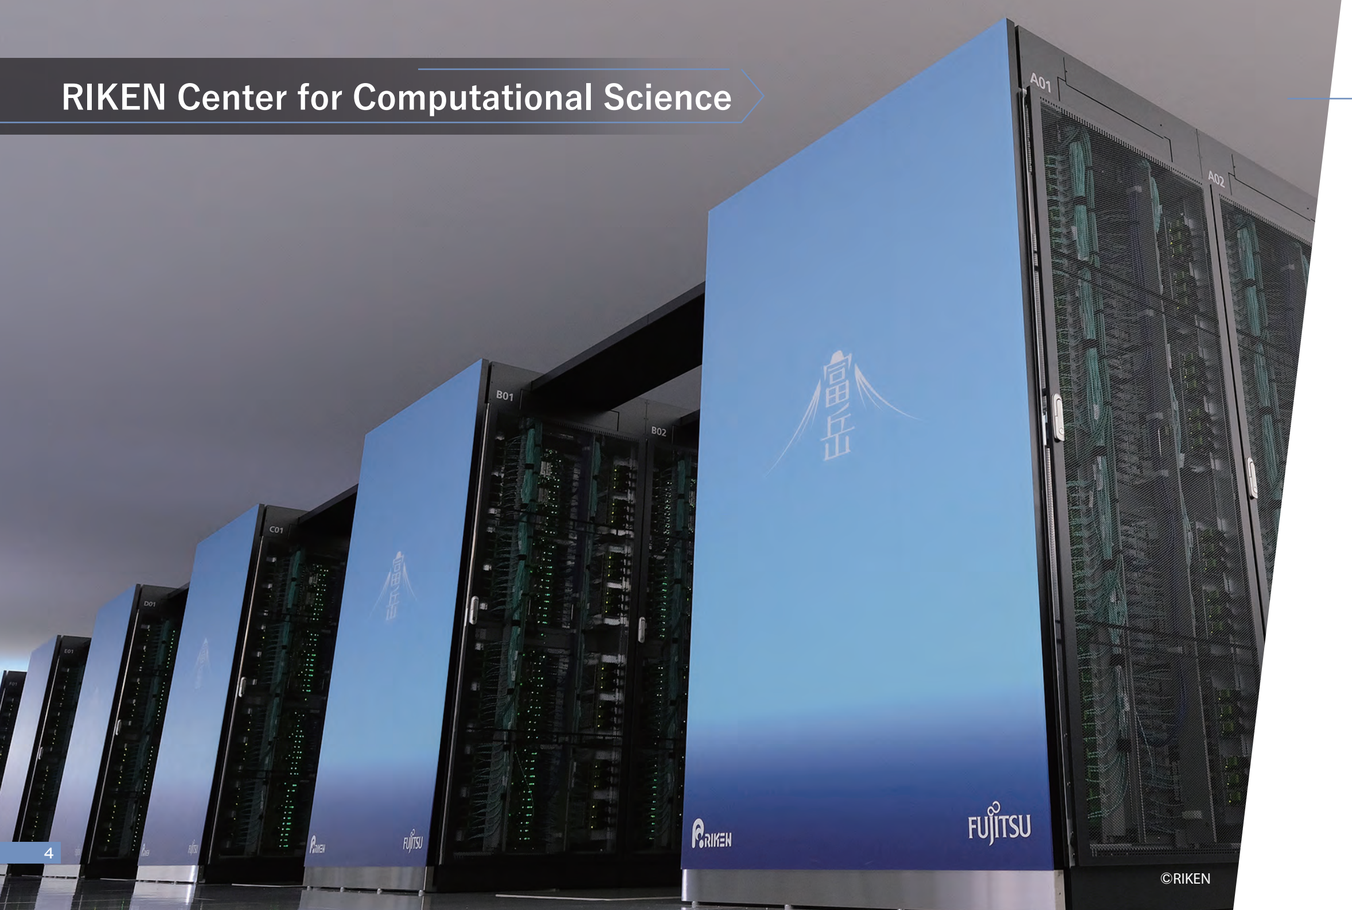

©RIKEN

Supercomputer Fugaku

Fumiyoshi Shoji

The supercomputer Fugaku officially became available for use on March 9, 2021. Fugaku features CPUs based on the Arm architecture widely used in smartphones and other devices, along with high-speed CPU-to-CPU interconnects. Taking a “co-design” approach, software and hardware engineers worked closely with each other to develop a system versatile enough to efficiently execute applications with diverse requirements across a range of fields.

The hardware configuration is shown in the figure on the right. Fugaku consists of compute nodes and IO nodes (storage, IO, boot), which are connected by an interconnect called “TofuD”. Each set of 16 compute nodes is equipped with a compute/storage node (approx. 1.6 TB of SSD storage). These compute/storage nodes constitute the first-tier storage. The first-tier storage is used as a cache for the second-tier storage, and as a local file system for compute nodes and a shared file system for jobs. The second-tier storage provides a total of 6 volumes with a Lustre-based shared file system and a total capacity of about 150 PB. The third-tier storage provides servers for external cloud storage services.

For greater convenience, Fugaku is also expanding its cloud capabilities to include REST API access, a container execution environment and orchestration tools, object storage, and high-speed connectivity to commercial clouds.

For more details, please refer to the website below.
https://www.r-ccs.riken.jp/en/fugaku/

Architecture	Armv8.2-A SVE (512 bit SIMD) +Fujitsu extensions
Core	48 cores for compute and 2/4 for OS activities
	Double-precision floating-point arithmetic : 2.7+ TF
	Single-precision floating-point arithmetic : 5.4+ TF
	Half-precision floating-point arithmetic : 10.8+ TF
Cache	L1D/core: 64 KiB, 4way, 230+ GB/s (load), 115+ GB/s (store)
	L2/CMG: 8 MiB, 16way
	L2/node: 3.6+ TB/s
	L2/core: 115+ GB/s (load), 57+ GB/s (store)
Memory	HBM2 32 GiB, 1024 GB/s
Interconnect	TofuD (28 Gbps x 2 lane x 10 port)
I/O	PCIe Gen3 x 16 lane
Technology	7nm FinFET

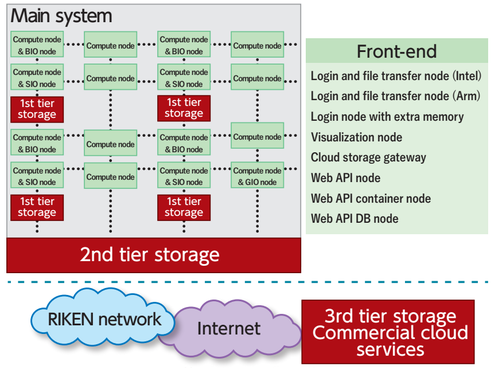

Hardware configuration

Information Initiative Center, Hokkaido University

Interdisciplinary Large-scale Computing System

Masaharu Munetomo

In December 2018, Information Initiative Center of Hokkaido University updated their “Interdisciplinary Large-scale Computing System”, which is comprised of supercomputer and cloud systems. Performance of the new supercomputer system is 20 times larger than that of the previous system. Moreover, the cloud system involves an advanced nationwide distributed computing environment, which is operated as the “Hokkaido University High-Performance Inter-Cloud”.

The supercomputer system consists of two computing subsystems named “Grand Chariot” and “Polaire”, and a storage subsystem of 16 PB. These subsystems are connected via a high-speed network based on Omni-Path architecture. The total computing performance of the supercomputer system is 4PFlops. Multi-core and many-core processors based on the x86 architecture are used for the system. The operating system of the system is based on Linux. High performance compiler, numerical libraries, and various tools for programming are available. Furthermore, various open-source software can be used for the system. These hardware and software stacks are expected to be used for cutting-edge computational science and HPCI collaborative research. Moreover, an originally designed job scheduler is used for the operation of the system, which provides the efficient use of computing resources.

The “Inter-Cloud” system consists of high-performance cloud servers, including baremetal (physical), virtual and GPU-equipped servers. These servers are managed by the OpenStack cloud computing platform, which provides the users with a cloud system environment that ensures performance and convenience.
In the “Inter-Cloud” system, an “Inter-Cloud” package is available. In this package, the user can use remote sites established at the University of Tokyo, Osaka University, and Kyushu University. Servers at the remote sites and Hokkaido University are connected via Japan’s ultra-high speed Science Information Network: SINET6. The users can simply apply and use this wide-area distributed system immediately, eliminating the need to coordinate between sites.

Moreover, a remote backup site has been set up at the Kitami Institute of Technology. Important data is regularly backed up to tape archiving equipment there, ensuring that research data is safely preserved in the event of a disaster.

For more details, please refer to the “Interdisciplinary Large-scale Computer System” page on the Hokkaido University Information Technology Center’s website below.
https://www.hucc.hokudai.ac.jp/en/overview/ilcs/

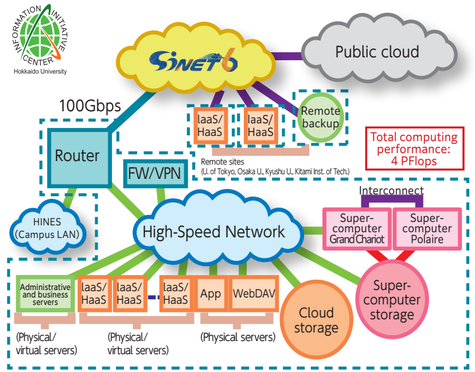

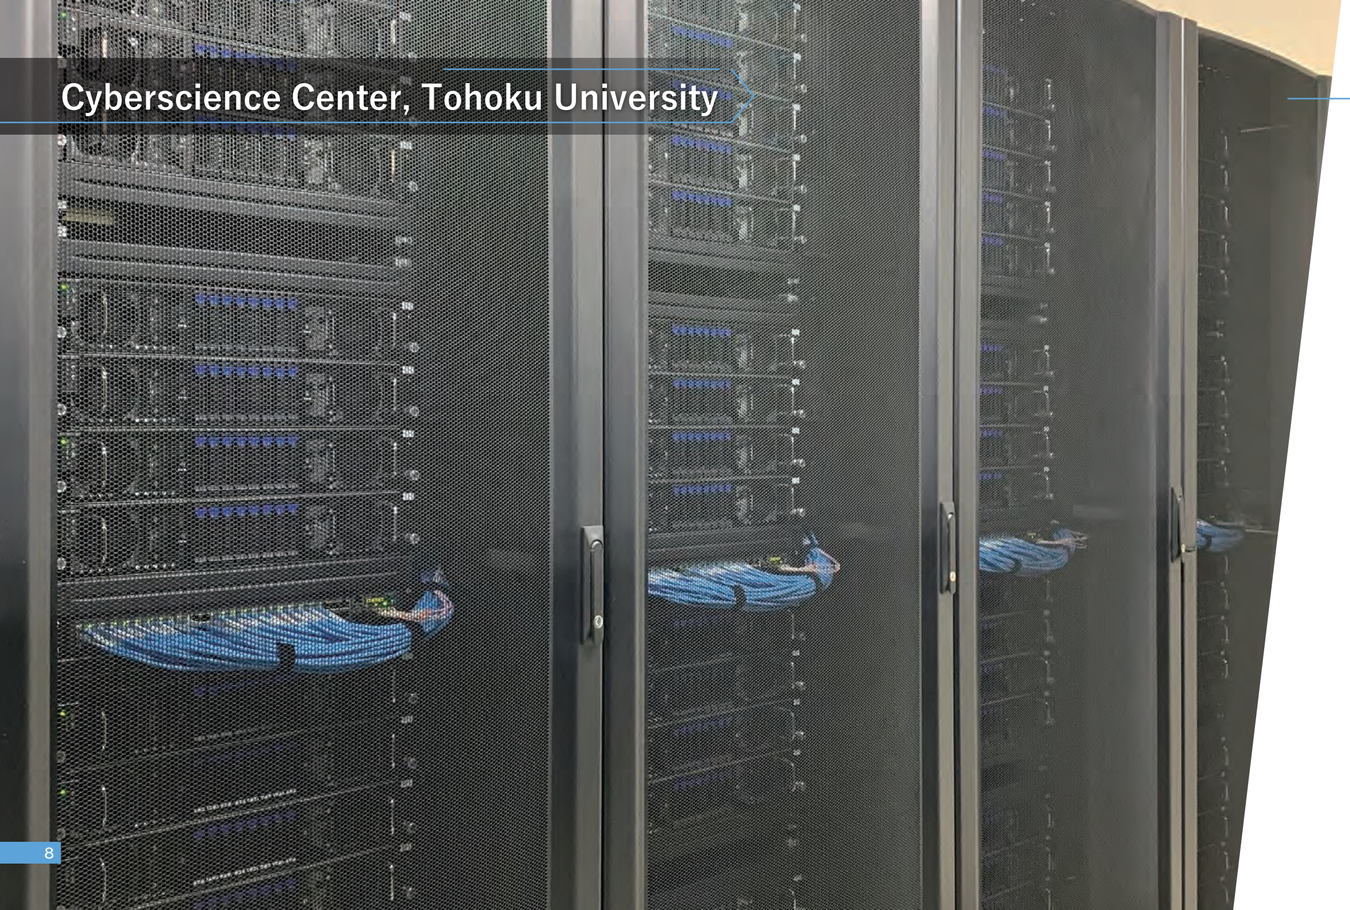

Cyberscience Center, Tohoku University

Supercomputer AOBA

Hiroyuki Takizawa

In October 2020, the Cyberscience Center at Tohoku University began operations of Supercomputer AOBA. AOBA consists of two subsystems. The AOBA-A subsystem uses NEC SX-Aurora TSUBASA B401-8 series vector processors. The AOBA-B subsystem uses NEC LX 406Rz-2 series. AOBA-A is a vector supercomputer with a total computing performance of 1.48 PFlops and a total memory bandwidth of 895 TB/s. It features a good balance between computing and memory performance. In particular, it is effective at performing scientific and technical calculations, which are often memory intensive. For this reason, AOBA-A is expected to mainly run user-developed applications. AOBA-B is an x86 server using AMD's latest EPYC processors, and used to run open-source software and commercial applications. The nodes of these subsystems are connected via a high-speed InfiniBand HDR network, sharing a file system with a total capacity of 2 PB.

Since 1997, Tohoku University's Cyberscience Center has worked closely with computer vendors as well as users to optimize their programs. This effort resulted in many valuable user-developed applications for AOBA's predecessor, an NEC SX-ACE system. Because its system architecture is significantly different from that of the SX-ACE, as soon as AOBA began operating we started actively working to help users migrate their applications to the new system. Also, the operating system is now a standard Linux environment. This makes it easier for new users to take full advantage of the performance of vector processors (especially for high memory bandwidth) for scientific computing. We can also help accelerate new users' computer programs.

It is also worth noting that AOBA is used to run emergency simulations that predict possible tsunami-related damage after an earthquake of magnitude 7.0 or greater occurs. These prediction results are immediately sent to the Japanese government's Cabinet Office to help the government and other agencies make prompt and accurate decisions in response to a tsunami. Thus, AOBA is not only used for academic research, but also plays a critical and prominent role in societal infrastructure that mitigates damage during emergencies.

For more details, please refer to the Cyberscience Center's page on Tohoku University's website below.
https://www.tohoku.ac.jp/en/news/university_news/tohoku_university_unveils_new_supercomputer_aoba.html

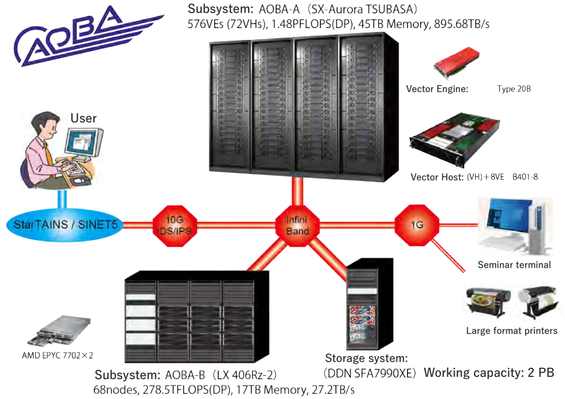

Center for Computational Sciences, University of Tsukuba

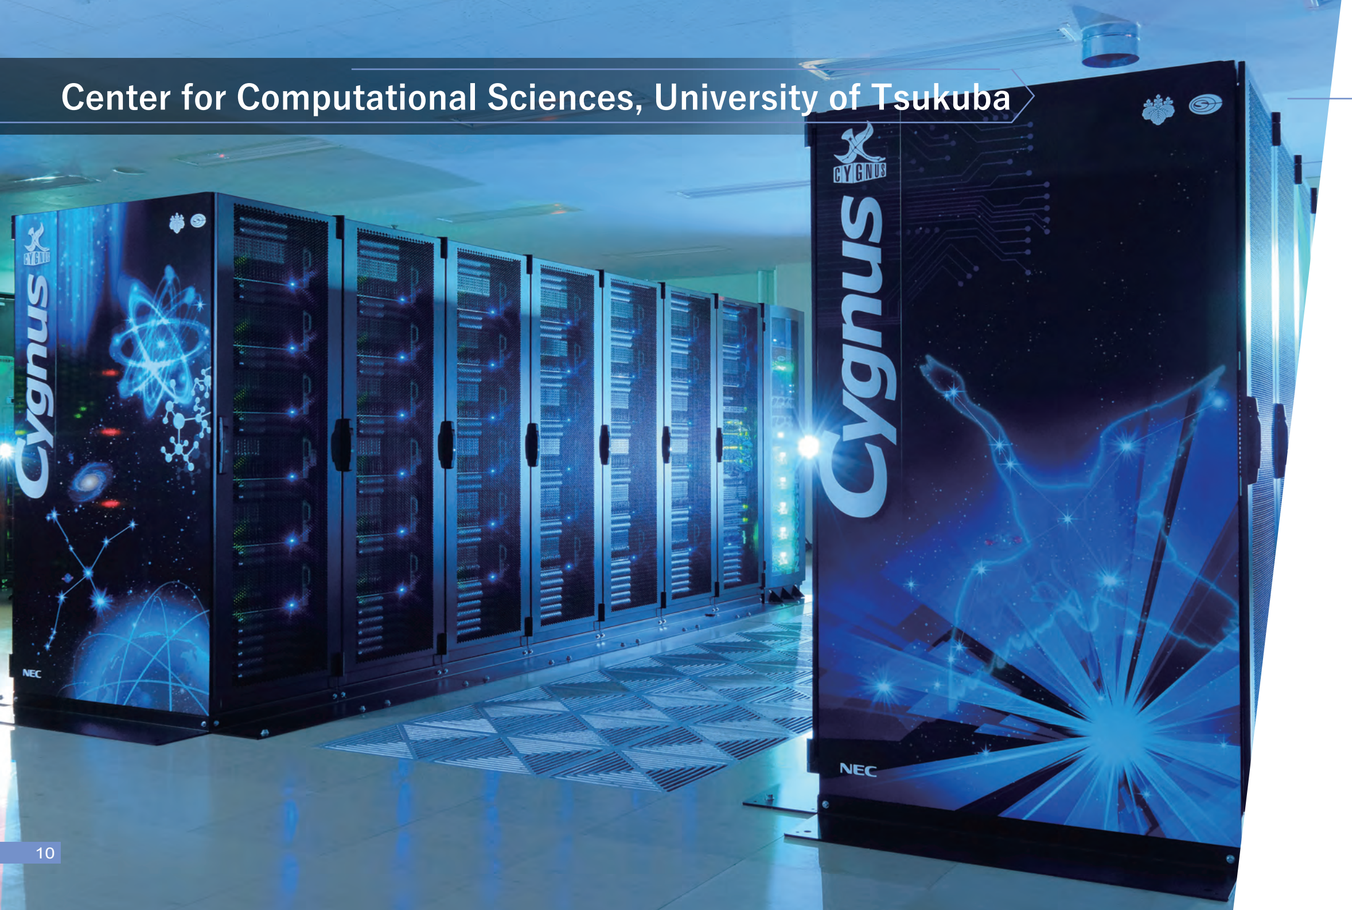

Cygnus

Taisuke Boku

In April 2019, the University of Tsukuba's Center for Computational Sciences (CCS) began operations of its Cygnus (PACS-X) supercomputer, the 10th generation system of the PACS series, which have been researched and developed at the center.

Cygnus is a massively parallel GPU/FPGA cluster system, consisting of a total of 78 nodes with 4 GPUs (NVIDIA Tesla V100) per node and a theoretical peak performance of 2.34 PFlops excluding FPGA. The compute nodes are connected by an InfiniBand HDR100 4-port (400 Gbps) "fat tree" network topology with full bisection bandwidth. In addition, 32 nodes are equipped with two Intel Stratix10 FPGAs per node, connected by 100 Gbit optical links in an 8-by-8 2D torus. The system is named after the constellation Cygnus (the Swan); its GPU nodes after Deneb, the brightest star in Cygnus; and its FPGA nodes after Albireo, the beautiful "double star" pair in Cygnus. All the compute nodes are equipped with 3.2 TB of NVMe SSD (Intel SSD DC P4610), which speeds up file input/output, and also provides substantial support for big data and AI applications.

At the University of Tsukuba's CCS, specialists in computer science and computational science have taken a "co-design" approach, collaborating to develop computers that are optimized for scientific work and propelling research. Since 2007, we have promoted a "Multidisciplinary Cooperative Research Program" to help expand multidisciplinary computational science work throughout Japan. By providing use of Cygnus free of charge under this program, we are contributing to the advancement of computer science and computational science.

For more details, please refer to the website below.
https://www.ccs.tsukuba.ac.jp/eng/supercomputers/

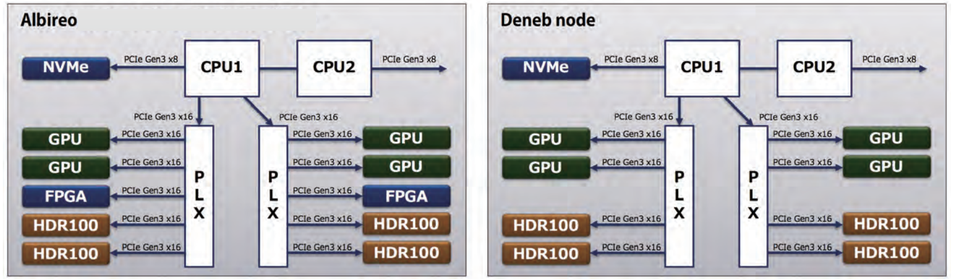

Block diagrams of nodes

Information Technology Center, The University of Tokyo

Wisteria / Big Data & Extreme Computing

Kengo Nakajima

Bringing about the dawn of new science: The “Wisteria/BDEC-01” supercomputer system, operated since May 14, 2021 by the University of Tokyo’s Information Technology Center, combines computer Simulations with big Data and machine Learning (S+D+L). This is a hybrid system with two types of compute nodes: simulation nodes (Odyssey) and data/machine learning nodes (Aquarius). With the human-centered vision of Japan’s “Society 5.0” initiative in mind, the “S+D+L” approach incorporates expertise in data science and machine learning into computational science and computational engineering.

The Simulation nodes cluster (Odyssey) is equipped with 7,680 Fujitsu A64FX processors, the same CPUs used in supercomputer Fugaku, for a peak performance of 25.9 PFlops. The Data & Learning nodes cluster (Aquarius) is equipped with 90 Intel Xeon Platinum 8360Y (Ice Lake) CPUs, and 360 NVIDIA A100 Tensor Core GPUs, for a peak performance of 7.2 PFlops. Odyssey and Aquarius are connected by an InfiniBand EDR 100 Gbps network with a bandwidth of 2 TB/s. In addition, some Aquarius nodes can directly access various external resources, including servers, storage, and sensor networks, via networks such as Japan’s SINET, and can record data in real time for analysis and simulations.

The Information Technology Center provides libraries, tools, and applications for a wide range of fields such as computational science, data science, artificial intelligence, and machine learning. The center has also created open source software to make developing high-performing S+D+L applications easier. These include “ppOpen-HPC”, an application development and execution environment with automatic tuning functionality, and “h3-Open-BDEC”, a novel software platform for S+D+L applications.

The Wisteria/BDEC-01 supercomputer is the world's first heterogeneous large-scale system that combines computer Simulations with big Data and machine Learning. It is expected to play an important role in realizing the goals of Japan’s “Society 5.0” initiative.

For more details, please refer to the website below.
https://www.cc.u-tokyo.ac.jp/en/supercomputer/wisteria/system.php

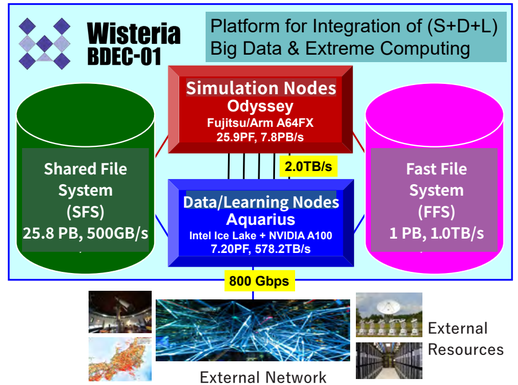

TSUBAME3.0

Toshiya Itoh

Since 2006, the Tokyo Institute of Technology's Global Scientific Information and Computing Center has developed and operated the TSUBAME series of supercomputers. In 2008, TSUBAME 1.2 was the first supercomputer in the world to introduce a large number of GPUs. Since then, we have provided computing resource services on GPU-based supercomputers. The current TSUBAME 3.0 began operations in August 2017, with 2,160 NVIDIA Tesla P100 GPUs installed in 540 compute nodes, each connected by a 400 Gbps high-speed network, providing 12.15 PFlops (double precision) of total computing performance and 16 PB of shared storage. Both TSUBAME 3.0 itself and its cooling system of are among the world's most energy-efficient. In 2017, TSUBAME 3.0 topped the Green500 list of the most energy-efficient supercomputers in the world. Its exceptionally efficient cooling system uses water at near outdoor temperatures to cool GPUs and CPUs directly. An extremely high Power Usage Effectiveness value of 1.033 shows just how efficient this cooling system really is. TSUBAME 3.0 provides a wide variety of fine-tuned computing resource options for more effective usage, and container environments that allow the user to more flexibly configure the software environment that suits their needs. High-performing GPUs have been widely embraced for simulation work in computational science. More recently, they're also being adopted in artificial intelligence and data science work, for which we provide the corresponding software.

In FY2023, 20% of TSUBAME 3.0's computing resources are being provided to Japan's HPCI initiative. We are also actively working to open our doors to industrial projects and human resource development initiatives to broaden the base of supercomputer users, such as through our Budding Researchers Support Program, which encourages work by young and female researchers, and through our Grand Challenge Large-Scale Computing Program, which aims to create results that can only be obtained by using all of TSUBAME's nodes at once.

For more details, please refer to the TSUBAME Computing Services website below.
https://www.t3.gsic.titech.ac.jp/en

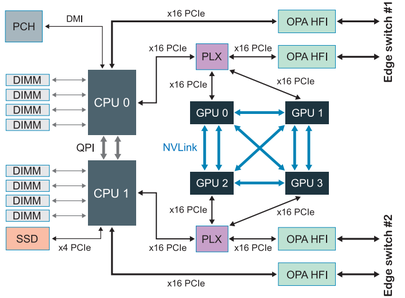

Block diagram of TSUBAME 3.0 compute node

TSUBAME 3.0 node

Free cooling tower producing cooling water at near outdoor temperatures

Information Technology Center, Nagoya University

Supercomputer Flow

Takahiro Katagiri

On July 1, 2020, Nagoya University's Information Technology Center began operations of the supercomputer "Flow", which consists of three subsystems and a cloud system.

- Subsystem I (Fujitsu) features 2,304 Fugaku-type A64FX-based nodes like those used at the RIKEN R-CCS.
- Subsystem II (NVIDIA) features 221 nodes, each equipped with 4 Tesla V100 Volta GPUs.
- Subsystem III (Hewlett Packard Enterprise) features a sizeable 48 TB of memory.
- The cloud system (Intel) features 100 nodes with 4-socket Xeon Gold 6230 CPUs.

With a theoretical performance of 15.88 PFlops, it is one of Japan's top supercomputers for numerical computations and data science. It also offers 6 PB of "cold storage" optical disc archiving in operation for the first time.

The Flow supercomputer marked the world's first use of a Fugaku-type node. It is available to any qualified user through an application process. It is particularly useful for preliminary development of Fugaku-targeted software, allowing for a seamless transition to Fugaku itself. Flow is also well-suited for the rapidly growing field of data science. Subsystem II, with its powerful GPUs for machine learning and a massive 30 PB of "hot storage" is particularly effective for data science work. Each node of Subsystem II is equipped with 6.4 TB of NVMe SSD storage (1.4 PB in total), and another 50 nodes (up to 320 TB) can be used to create a shared file system using BeeGFS. These features allow the high-speed file access needed for machine learning.

For more details about Flow, please refer to the website below.
https://icts.nagoya-u.ac.jp/en/sc/

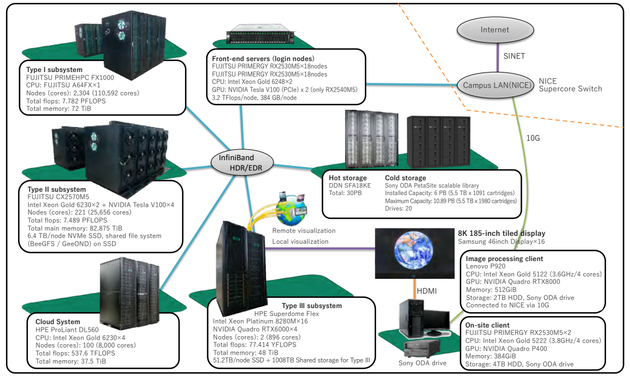

Academic Center for Computing and Media Studies, Kyoto University

Supercomputer Systems (Camphor / Laurel / Cinnamon)

Keiichiro Fukazawa

Kyoto University's Academic Center for Computing and Media Studies runs three supercomputer systems: Camphor 2, Laurel 2, and Cinnamon 2.

- Camphor 2, introduced in October 2016, is a computational performance-oriented Cray XC40 series supercomputer with Intel Xeon Phi (KNL) processors.
- Laurel 2, introduced in December 2016, is a general-purpose Cray CS400 2820XT supercomputer with Intel Xeon Broadwell processors.
- Cinnamon 2, also introduced in December 2016, is a memory-intensive use Cray CS400 4840X supercomputer with 4-socket Intel Xeon Haswell processors.

The Camphor 2's 1,800 nodes yield a total theoretical performance of 5.48 PFlops, while the Laurel 2's 850 nodes yield 1.03 PFlops. The Cinnamon 2 has only 16 nodes, but with 3 TB of memory per node. The different features of these three supercomputers are intended to provide an environment that can meet the diverse computational needs of each user. A portion of these supercomputing resources are supported by Japan's HPCI and JHPCN initiatives.

The Academic Center for Computing and Media Studies also has its own supercomputer joint research system, under which three types of research incentive programs are offered.

- The first program covers all or part of the usage fees for researchers under 40 years of age or female researchers regardless of age.
- The second program covers part of the usage fees for research groups with certain large jobs.
- And the third program supports improvements and refinements to computer programs.

For more details about our supercomputer resources and joint research programs, please refer to the Kyoto University Academic Center for Computing and Media Studies website below.

https://www.media.kyoto-u.ac.jp/accms_web/en/

Please note: As of July 2022, these supercomputer resources will have been taken offline to allow for a new supercomputer system that will become fully operational from October 2023.

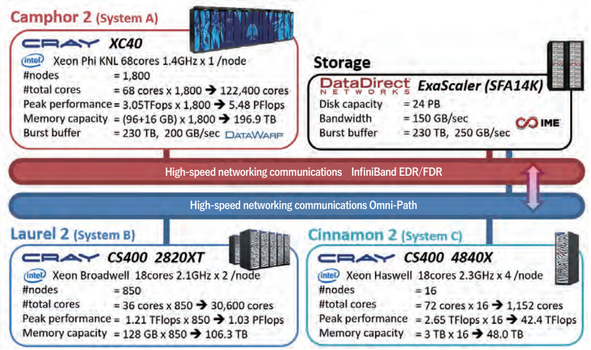

Cybermedia Center, Osaka University

Supercomputer SQUID

Susumu Date

Osaka University's Cybermedia Center offers the use of two supercomputer systems: the "SQUID" for cloud-linked high-performance computing and data analysis, and the petaflops-class hybrid "OCTOPUS". This introduction will focus on the newer "SQUID" supercomputer, made available in May 2021.

SQUID, an acronym from "Supercomputer for Quest to Unsolved Interdisciplinary Datascience", is a hybrid cluster system with general-purpose CPU nodes, GPU nodes, and vector nodes, for a total computing performance of 16.591 PFlops. SQUID's Lustre parallel file system with 20 PB of HDD and 1.2 PB of SSD storage can be accessed using DDN's EXAScaler. The processors and accelerators in each node cluster feature direct liquid cooling and are designed and built for reliable high-performance. SQUID can provide high-performance compute nodes with different processors, accelerators, and architectures within a single computing environment, making SQUID distinctively capable of accommodating each user's diverse computational needs.

Our faculty and staff are ready to support you in using Osaka University supercomputers. We hope you'll give SQUID a try!

For more details, please refer to the Osaka University Cybermedia Center website below.
http://www.hpc.cmc.osaka-u.ac.jp/en/

Theoretical Computing Speed	16.591 PFlops	
Nodes	General-purpose CPU nodes 1,520 nodes (8.871 PFlops)	CPU: Intel Xeon Platinum 8368 (Ice Lake / 2.4 GHz 38 cores) x 2 Memory: 256 GB
	GPU nodes 42 nodes (6.797 PFlops)	CPU: Intel Xeon Platinum 8368 (Ice Lake / 2.40 GHz 38 cores) x 2 Memory: 512 GB GPU: NVIDIA Delta HGX A100 8 GPU board
	Vector nodes 36 nodes (0.922 PFlops)	CPU: AMD EPYC 7402P (Rome / 2.8 GHz 24 cores) x 1 Memory: 128 GB Vector Engine: NEC SX-Aurora TSUBASA Type20A x 8
Storage	DDN EXAScaler (Lustre)	HDD: 20.0 PB NVMe: 1.2 PB
Interconnect	Mellanox InfiniBand HDR (200 Gbps)	

SQUID system configuration

Cybermedia Center supercomputers SQUID and OCTOPUS

Research Institute for Information Technology, Kyushu University

Research Computer System ITO

Kenji Ono

On January 9, 2018, Kyushu University's Research Institute for Information Technology made available the Fujitsu-based “ITO” supercomputer system. This system is equipped with Intel Skylake Scalable Performance CPUs and NVIDIA Tesla P100 GPUs, for a total theoretical computing performance of approximately 10 PFlops, making it one of the most powerful systems in Japan.

The ITO system specifications were chosen for flexibility, to help achieve the super smart “Society 5.0” envisioned in the Japanese government’s 5th Science and Technology Basic Plan, and to provide a research platform for artificial intelligence, machine learning, big data, data science, and more. This is the first supercomputer in Japan configured to connect, via a high-speed file system, the type of large private front-end cloud computing environment needed for interactive data science and multi-system collaborative work with the type of high-performance back-end computing node clusters needed for large simulations and machine learning. In addition to the standard front-end specifications, large memory types can be used, and virtual and bare metal servers can be reserved in advance through a web-based system. ITO also provides a platform for new classes of users and research projects, implementing a full-fledged collaborative interface for interacting with public clouds, and providing support for supercomputing that takes advantage of open data found on the Internet. We are also working on new technologies to allow intelligent power-saving operations. The ITO system introduces new detailed power monitoring and power-limited job scheduling functions.

Through the use of the ITO system, for various programs provided by our center, providing it broadly to researchers inside and outside the university, and as a computing resource for Japan’s JHPCN & HPCI initiatives, the Kyushu University Research Institute for Information Technology contributes in developing new academic research and in strengthening the foundations for academic research in Japan.

For more details, please refer to the Kyushu University Research Institute for Information Technology website below.
https://www.cc.kyushu-u.ac.jp/scp/eng/system/ITO/01_intro.html

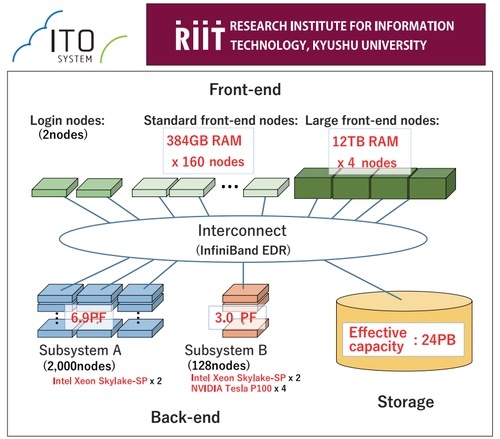

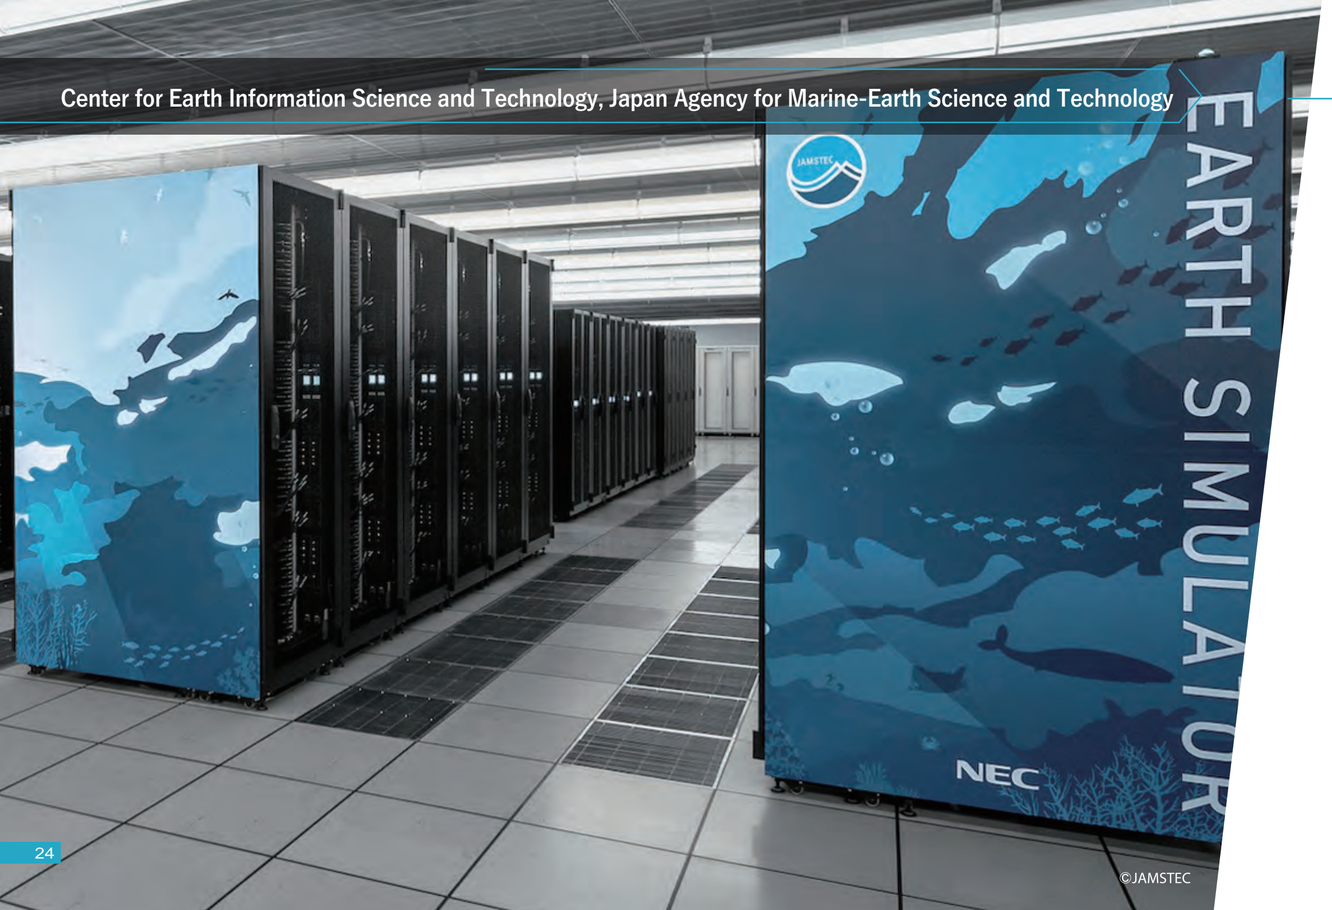

©JAMSTEC

Earth Simulator 4

Hitoshi Uehara

The Japan Agency for Marine-Earth Science and Technology (JAMSTEC) Research Institute for Value-Added Information Generation (VAiG) Center for Earth Information Science and Technology (CEIST) has updated its “Earth Simulator” and made these supercomputing resources available for Japan’s HPCI initiative from June 2021.

The upgraded Earth Simulator 4 supercomputer consists of CPU nodes with AMD EPYC Rome processors on HPE Apollo, GPU nodes with NVIDIA A100 GPU, Vector Engine nodes with NEC SX-Aurora TSUBASA processors, high-capacity DataDirect Networks storage, and a high-speed InfiniBand network. The Earth Simulator 4’s CPU and Vector Engine resources are available for Japan’s HPCI initiative.

Its 684 vector engine nodes provide a total computing power of 14.97 PFlops with a total memory bandwidth of 8.5 PB/s, making this a powerful tool for research that uses vector computing. Its 720 CPU nodes, based on the widely-used and versatile x86 architecture, provide a total computing power of 3.3 PFlops and 180 TB of total memory, making it suitable for a wide variety of research projects. These nodes can be used for one-off batch jobs to meet the particular computational needs of various policy, industrial, and academic projects.

The data storage is a Lustre-based shared file system with 60 PB or 1.3 PB capacities, composed of HDDs or SSDs.

The storage is directly accessible from all nodes and front-end servers, as are large-scale shared memory servers with 9 TB memory for pre-post processing.

The Center for Earth Information Science and Technology provides seminars on how to use the Earth Simulator, as well as robust support for porting and optimizing programs. Please consider using the Earth Simulator.

For more details, please refer to the Earth Simulator website below.
https://www.jamstec.go.jp/es/en/

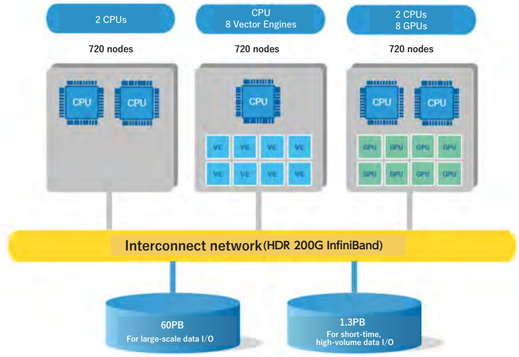

AI-Bridging Cloud Infrastructure

Yusuke Tanimura

To advance Japanese R&D in artificial intelligence and accelerate AI's use in society, Japan's National Institute of Advanced Industrial Science and Technology (AIST) has made its AI Bridging Cloud Infrastructure (ABCI) available since August 2018. An upgrade in May 2021 gave the integrated "ABCI 2.0" system new energy-efficient high-performance GPU-based nodes and enhanced storage.

The upgraded ABCI is now a massively parallel cluster supercomputer with 1088 "V-type" compute nodes, each consisting of 2 Intel Xeon Skylake CPUs and 4 NVIDIA V100 GPUs, and 120 "A-type" compute nodes, each consisting of 2 Intel Xeon Ice Lake CPUs and 8 NVIDIA A100 GPUs, connected by a high-speed InfiniBand network. ABCI's peak performance is 56.6 PFlops at double precision, and 851.5 PFlops at half precision. It has 573.5 TB total memory, and 2.22 PB total NVMe SSD storage. The system also includes a shared file system with an effective capacity of 35 PB, and 17 PB of Amazon S3-compatible storage.

The system is installed at the AIST Kashiwa facility's AI Data Center building, an ultra-high density, ultra-low power data center constructed prior to the installation. The AI Data Center employs a "free cooling" system that produces chilled water using only a cooling tower. Compute nodes are cooled with a combination of this chilled water and air from Fan Coil Units. This results in an average annual Power Usage Effectiveness ratio of 1.1.

Since it began operating, ABCI has consistently held high positions on the TOP500 ranking. In July 2019, it set a world record in the Image Classification category on the MLPerf Training v0.6 benchmark for deep learning, taking only 70 seconds training time to recognize ImageNet images on a ResNet-50 convolutional neural network. In November 2020, it also achieved the highest speed on the MLPerf HPC v0.7 machine learning benchmark.

To meet the widely varying demands for AI computing resources for R&D and community uses, ABCI does not draw a distinction between academic and industrial uses. Anyone affiliated with a company, university, or research institute in Japan can apply and use ABCI for a flat rate. ABCI is also available as an HPCI resource, so please consider it for your next project.

For more details, please refer to the ABCI website below.
https://abci.ai/

If you are considering using HPCI computing resources or would like to learn more about Japan' s HPCI initiative, please refer to the HPCI Help Desk website below.

Help Desk helpdesk@hpci-office.jp

About the Research Organization for Information Science and Technology's Kobe Center

About the Research Organization for Information Science and Technology' s Kobe Center

The Research Organization for Information Science and Technology' s Kobe Center is responsible for promoting and supporting the use of Fugaku and other Japanese supercomputers.

Supercomputer simulations across a wide range of fields are making major contributions toward a more safe and secure society. These include: elucidating the fundamental laws of matter and the evolution of the universe, realizing new sources of energy, analyzing genomic and intracellular dynamics, finding physical materials with new properties and capabilities, making highly accurate predictions of typhoons and tsunamis, efficiently designing new drugs, and allowing highly reliable design and manufacturing research that alleviates the need for physical prototyping and experiments.

The RIST Kobe Center strives to make its world-class supercomputers available to researchers and engineers in a wide range of fields in a fair and efficient manner so they can produce many fruitful research results.

HPCI Computing Resource Handbook
January 2023 English version

Research Organization for Information Science and Technology
Kobe Center

https://www.hpci-office.jp

*Please refrain from copying or reprinting the contents of this booklet without permission.

© 2023 Research Organization for Information Science and Technology