HPCI Computing Resource Handbook 2026

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

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 a shared computational environment.
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
- Center for Information Infrastructure, Institute of Science Tokyo
- Information Technology Center, Nagoya University
- Academic Center for Computing and Media Studies, Kyoto University
- D3 Center, The University of Osaka (UOsaka)
- 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(ISM)
- Information Technology and Human Factors, National Institute of Advanced Industrial Science and Technology (AIST)
- RIKEN Information R&D and Strategy Headquarters (R-IH)

CPU Architectures

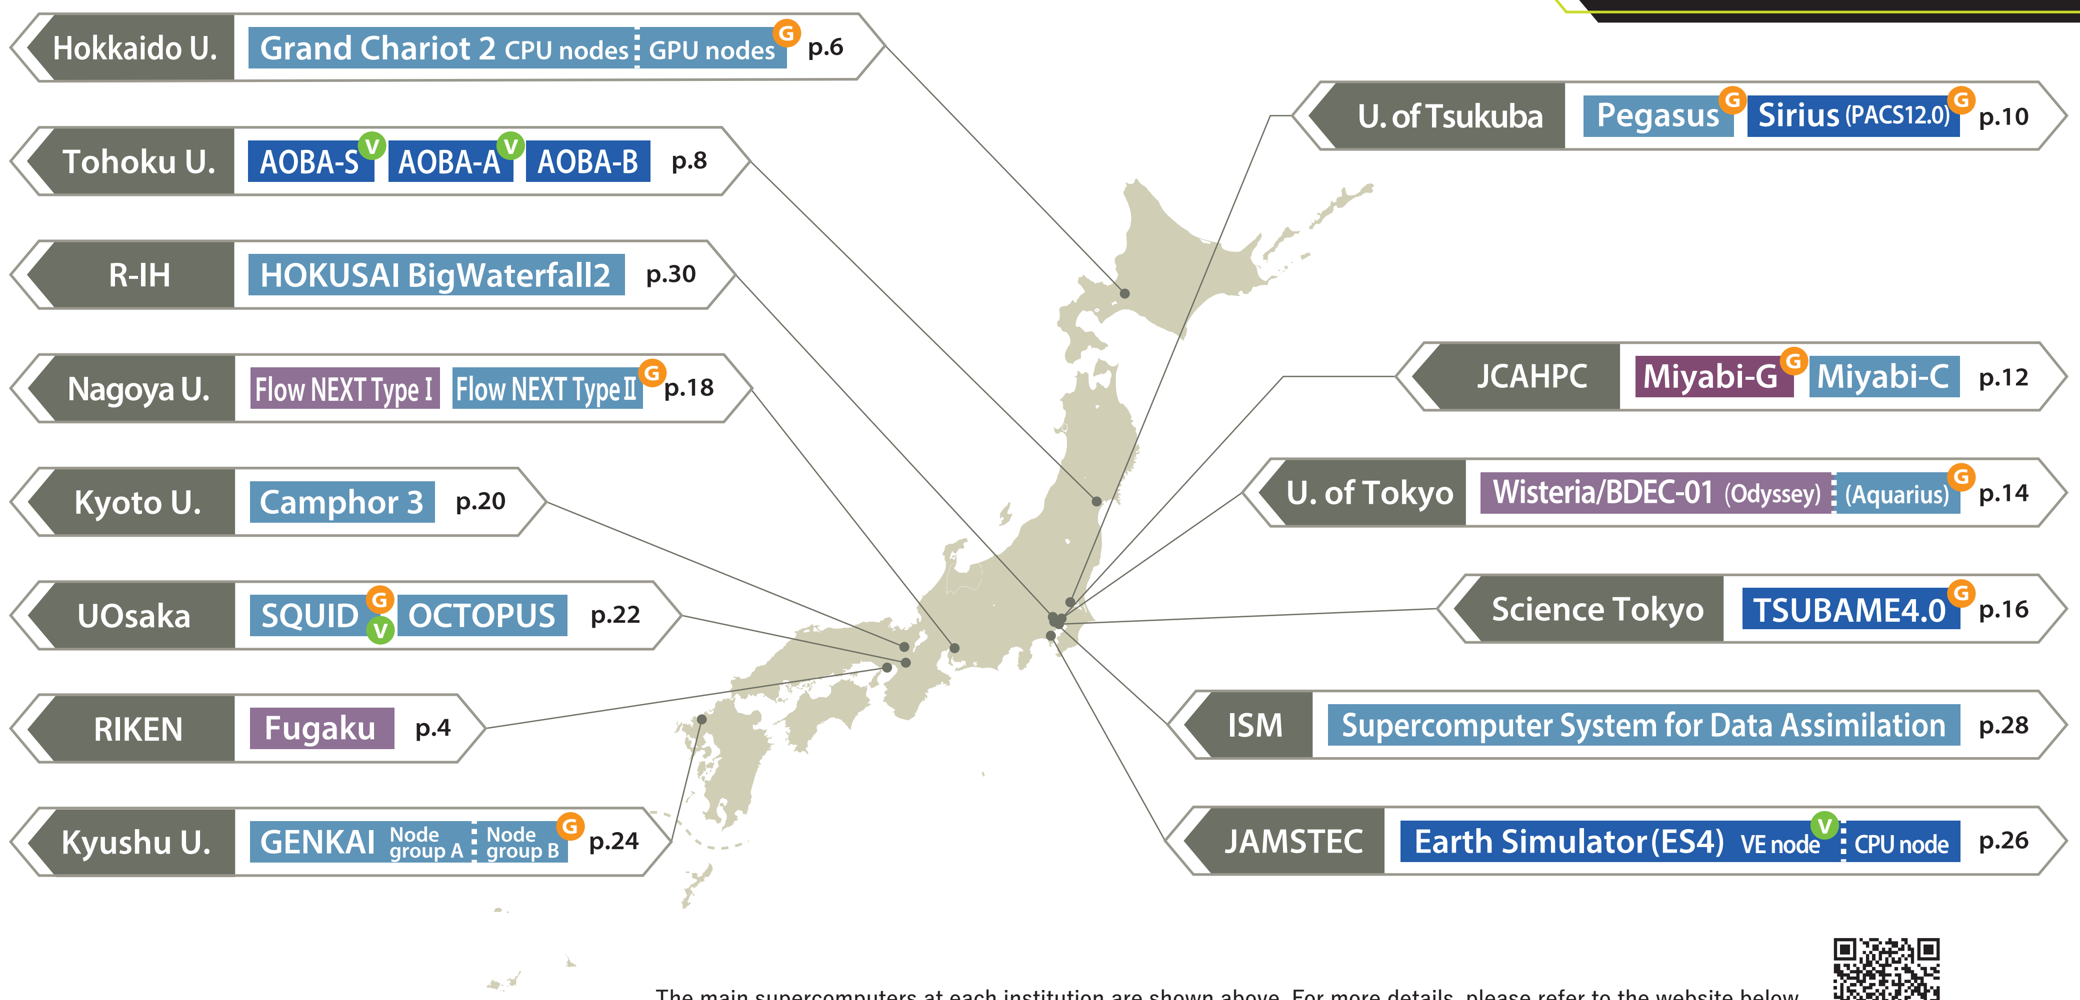

The main supercomputers at each institution are shown above. For more details, please refer to the website below.
https://www.hpci-office.jp/en/using_hpci/hardware_software_resource

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
GRACE CPU	Armv9 and SVE2 compatible processors equipped with NVIDIA GH200 superchips.
G GPU	NVIDIA/ADM's GPUs for acceleration
V Vector	x86-64 host processors with NEC's Vector Engine for acceleration

Supercomputing Resources by Software

Data is current as of January 2026, but subject to change.
Some institutions have restrictions on the use of commercial software.
Please contact our help desk for the latest information

E-mail: 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 JCAHPC		p.14 U. of Tokyo		p.16 Science Tokyo	p.18 Nagoya U.		p.20 Kyoto U.	p.22 UOsaka		p.24 Kyushu U.		p.26 JAMSTEC	p.28 ISM	p.30 R-IH
		Fugaku	Grand Chariot 2	AOBA-S (SX)	AOBA-A (SX)	AOBA-B (LX)	Pegasus	Sirius (PACS12.0)	Miyabi-G	Miyabi-C	Wisteria (Odyssey)	Wisteria (Aquarius)	TSUBAME 4.0	Flow NEXT Type I	Flow NEXT Type II	Camphor 3	SQUID	OCTOPUS	GENKAI Nodegroup A	GENKAI Nodegroup B	Earth Simulator	Supercomputer System for Data Assimilation	HOKUSAI BigWaterfall2
Molecular Dynamics	AMBER	○											○		○					○	○		○
	CHARMM													○	○								
	DeePMD-kit												○		○								
	ERmod	○																					
	GENESIS	○	○			○	○		○	○	○	○	○	○	○	○	○	○	○	○	○		
	GROMACS	○	○			○			○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	LAMMPS	○	○			○			○	○	○	○	○	○	○	○	○	○	○	○	○	○	
	MODYLAS	○	○			○			○	○	○	○	○			○	○			○	○		
	MyPresto		○																				
	N2P2	○																					
	NAMD	○	○										○	○	○	○							○
	OCTA	○																					
	Tinker												○										
Quantum Chemistry	ABINIT-MP	○	○	○	○	○			○	○	○	○	○	○	○	○	○	○	○	○	○		
	GAMESS		○										○	○	○	○	○	○	○	○	○		○
	Gaussian	○	○			○							○		○	○	○	○	○	○	○		○
	GRRM					○																	
	Molpro																			○	○		
	NTChem	○	○			○			○	○	○	○	○			○	○			○	○		
	NWChem	○							○	○	○	○											
	SMASH	○	○			○			○	○	○	○	○			○	○	○	○	○	○		
Condensed Matter Physics	ABINIT	○																					
	AkaiKKR	○	○			○	○		○	○	○	○	○			○	○			○	○		
	ALAMODE	○	○			○			○	○	○	○	○			○	○			○	○		
	CP2K	○							○	○	○	○	○	○	○					○	○		
	CPMD	○																					
	HΦ	○	○		○	○			○	○	○	○	○			○	○			○	○		
	mVMC	○				○			○	○	○	○	○			○	○			○	○		
	OpenMX	○	○			○			○	○	○	○	○			○	○			○	○		
	PHASE/0	○	○	○	○	○			○	○	○	○	○			○	○			○	○		
	Phonopy	○	○			○			○	○	○	○	○			○	○			○	○		
	Quantum ESPRESSO	○	○	○	○	○			○	○	○	○	○	○	○	○	○	○	○	○	○		
	SALMON	○	○	○		○	○		○	○	○	○	○			○	○			○	○		
	SIESTA	○																					
	VASP	○											○						○	○	○		
Computational Biology	Alphafold												○		○								

Data is current as of January 2026, but subject to change.
Some institutions have restrictions on the use of commercial software.
Please contact our help desk for the latest information

E-mail: 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 JCAHPC		p.14 U. of Tokyo		p.16 Science Tokyo	p.18 Nagoya U.		p.20 Kyoto U.	p.22 UOsaka		p.24 Kyushu U.		p.26 JAMSTEC	p.28 ISM	p.30 R-IH
		Fugaku	Grand Chariot 2	AOBA-S (SX)	AOBA-A (SX)	AOBA-B (LX)	Pegasus	Sirius (PACS12.0)	Miyabi-G	Miyabi-C	Wisteria (Odyssey)	Wisteria (Aquarius)	TSUBAME 4.0	Flow NEXT Type I	Flow NEXT Type II	Camphor 3	SQUID	OCTOPUS	GENKAI Nodegroup A	GENKAI Nodegroup B	Earth Simulator	Supercomputer System for Data Assimilation	HOKUSAI BigWaterfall2
Computational Biology	AutoDock Vina	○																					
	Colabfold														○								
	cryoSPARC														○								
	rDock	○																					
	Parabricks														○								
	Relion														○		○	○	○				
Fluid Analysis	ANSYS Fluent	○																					
	CONVERGE	○																					
	Cradle CFD scFLOW	○																					
	EXAPARTICLES/FLOW	○																					
	FDS	○																					
	FFVHC-ACE	○	○			○			○	○	○	○	○			○	○			○	○		
	FFX	○	○	○	○	○			○	○	○	○	○			○	○						
	FrontFlow/blue	○	○	○	○	○			○	○	○	○	○	○	○	○	○			○	○		
	FrontFlow/red		○											○	○								
	OpenFOAM	○	○			○			○	○	○	○	○	○	○	○	○	○	○	○	○	○	
	STAR-CCM+	○																					
	V-FaSTAR		○																				
Structural / Collision Analysis	ANSYS Mechanical												○										
	FrontISTR	○	○	○	○	○	○		○	○	○	○	○	○	○	○	○			○	○	○	
	LS-DYNA	○																					
	Marc																			○	○		
	MSC Nastran																			○	○		
	Virtual Performance Solution (VPS)	○																					
Electromagnetic Field Analysis	EXAMAG LLG シミュレータ	○																					
	Meep		○																				
	OpenFDTD	○																					
	Poynting	○																					
Multi-physics	Freefem++																	○	○				
Particle Systems	GEANT4								○	○	○	○		○	○								
Weather / Climate	NEMO	○																					
	SCALE	○																					
	WRF	○	○																				
Quantum Computing	cuQuantum														○								
	Qulacs	○																					
	RIKEN-braket	○																					

RIKEN Center for Computational Science

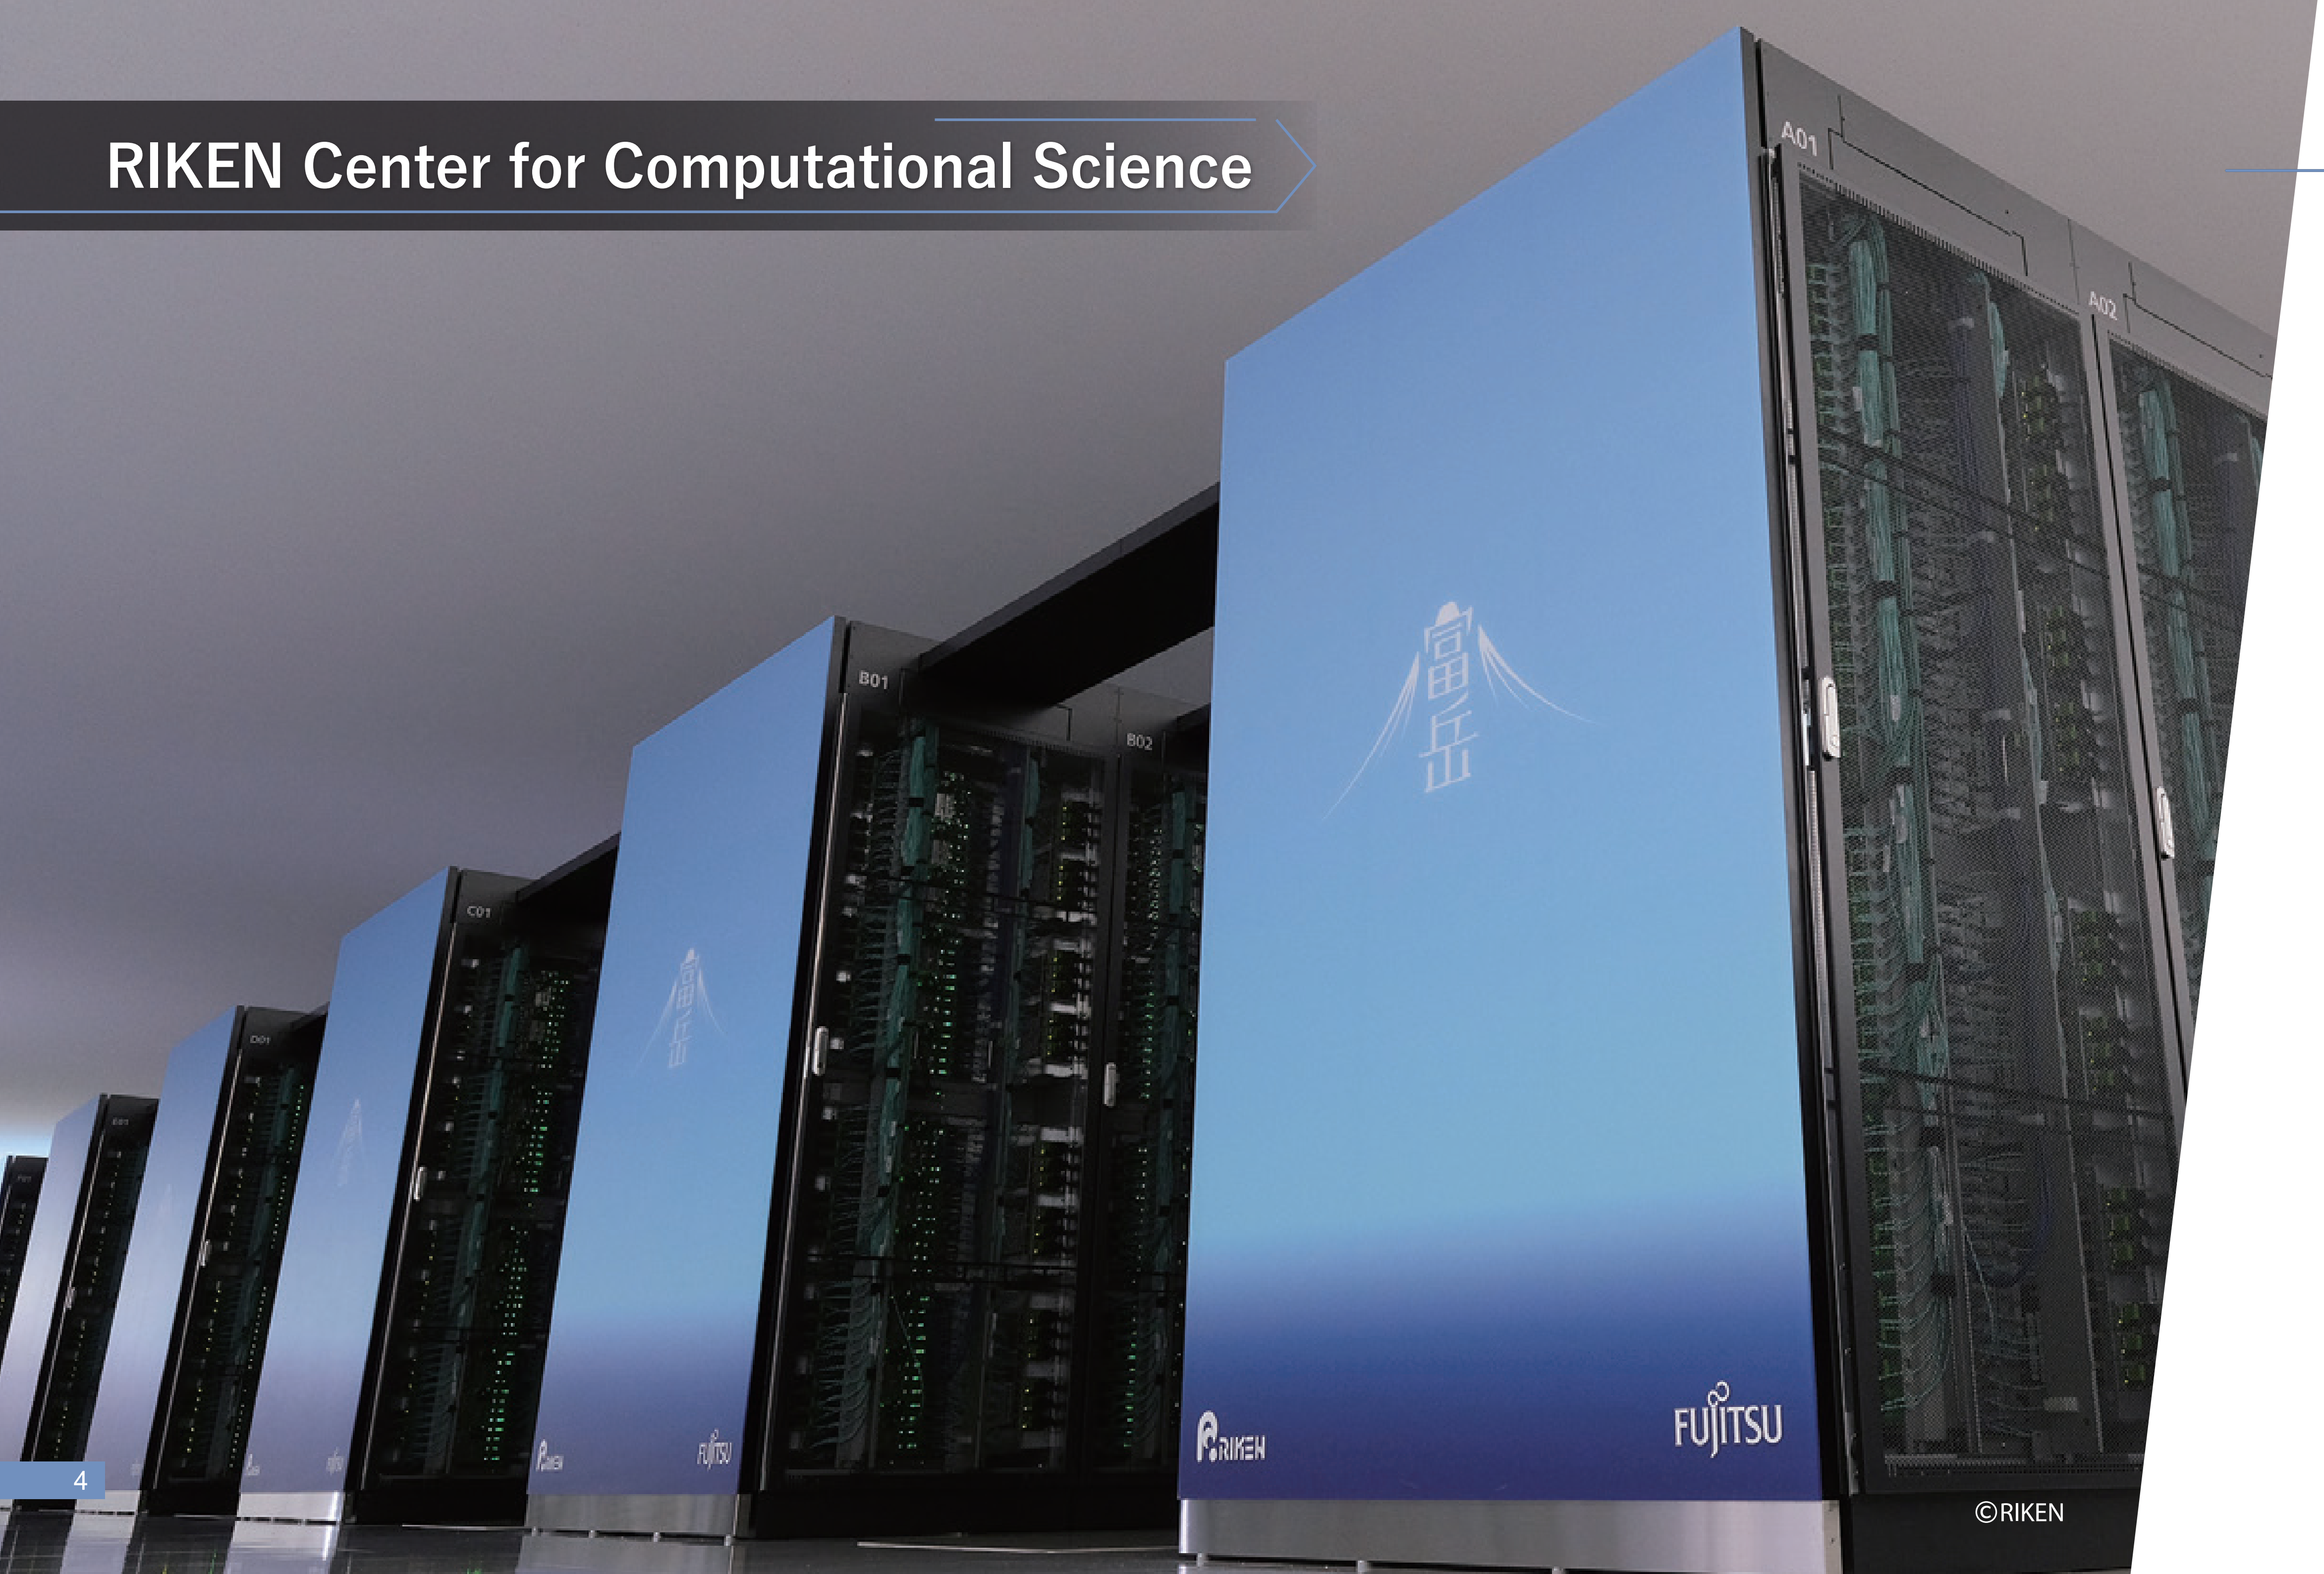

©RIKEN

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 I/O nodes (storage, I/O, boot) linked by interconnects known as “TofuD”. Each set of 16 compute nodes is equipped with a compute/storage node with about 1.6 TB of SSD storage. These compute/storage nodes constitute the first-tier storage layer. The first-tier storage is used as a cache for the second-tier storage, as a local file system for compute nodes, and as a shared file system for jobs. The second-tier storage provides a total of 6 volumes with a Lustre-based shared file system with a total capacity of about 150 PB. The third-tier storage provides servers for external cloud services.

For greater convenience, Fugaku is also expanding its cloud capabilities to provide Web-based access environments such as Open OnDemand, and increasing compatibility through a commercial cloud environment called “Virtual Fugaku”. Additional efforts are underway to expand programs for beginners, including an introductory “First-Touch” option, a workshop option, and an application services program.

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

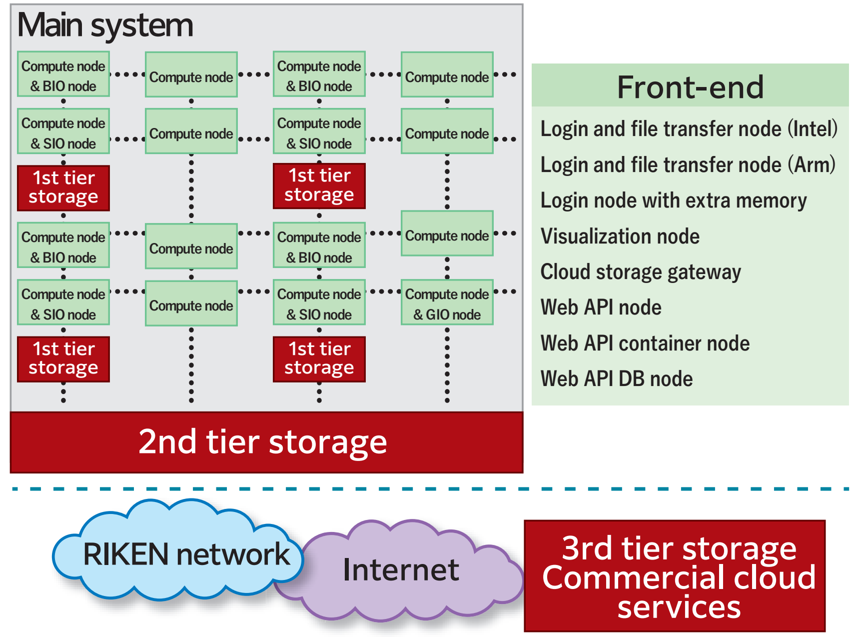

Hardware configuration

Information Initiative Center, Hokkaido University

About the Interdisciplinary Large-scale Computing System

Takayuki Umeda

The Hokkaido University Information Initiative Center updated the Interdisciplinary Large-scale Computing System, which consists of a supercomputer system and a cloud system, and began providing services for the new system in July 2025, with Fujitsu's PRIMERGY series as its core.

The supercomputer system comprises a computing system with a total theoretical peak performance of 9 PFlops (Grand Chariot 2) and a storage system with 16.95 PB of physical storage capacity, both connected via high-efficiency networking. The computing system consists of 504 computing nodes equipped with two 5th Gen Intel Xeon processors, with 24 of these nodes additionally equipped with four NVIDIA H100 GPU computing cards per node. The storage system is an all-flash configuration, with the storage medium consisting solely of flash memory (SSD). The supercomputer system offers two types of services: an "exclusive course," where users can use solely the computing resources they apply for at a flat rate, and a "shared course," where users access resources with tokens (operation hours) purchased on a pay-per-use basis. A portion of the shared course is provided as computing resources for HPCI, JHPCN, and the Center's joint research programs. Building on the design concept of the former system (Grand Chariot), the system uses multicore CPUs based on x64 architecture and the Linux OS. By enhancing node performance and shared memory capacity, the CPU node group places greater emphasis on ease of use, supporting a wide range of software stacks, including open-source software, and enabling research activities for more researchers, including those transitioning from the former system. In addition, the newly added GPU nodes address the needs in AI and machine learning.

The research cloud system adopts PRIMERGY CDI (Composable Disaggregated Infrastructure), a disaggregated computing technology. Twenty PC servers equipped with two 5th Gen Intel Xeon processors, four NVIDIA H100 GPU computing cards, and 30.72 TB of NVMe SSDs, housed in PCIe BOXs, are interconnected via PCIe fabric, enabling flexible resource assignment and configuration changes. The system uses SUSE Rancher Prime software, shifting from a conventional virtual environment to a container-based cloud environment via Kubernetes. Building on these foundational technologies, the cloud system provides interactive computing via a Jupyter environment and application deployment through shared or dedicated Kubernetes clusters.

Furthermore, establishing a tape-based remote archiving system at the Kitami Institute of Technology to make periodic backups of critical data ensures the security of crucial research data in the event of a disaster.

For system details, please refer to the Hokkaido University Information Initiative Center website for the Interdisciplinary Large-scale Computing System:

https://www.hucc.hokudai.ac.jp/en

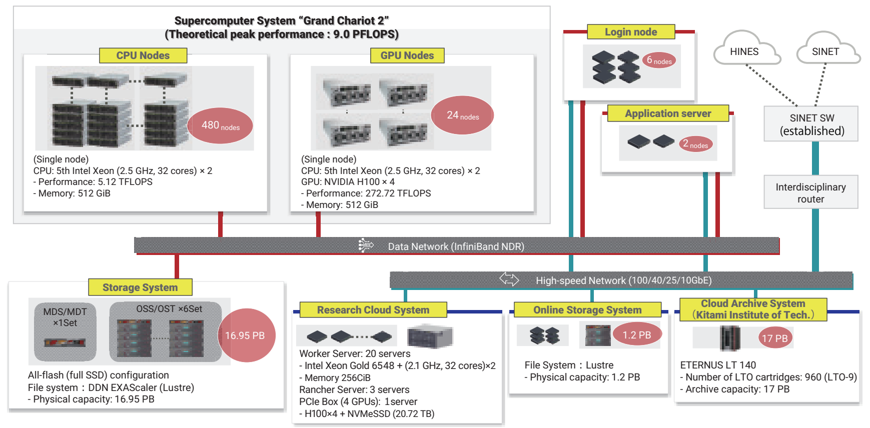

Diagram of the new Interdisciplinary Large-Scale Computing System

Supercomputer AOBA

Hiroyuki Takizawa

Tohoku University's Cyberscience Center has operated the Supercomputer AOBA since October 2020. In August 2023, this system underwent major enhancements. AOBA consists of three subsystems. In addition to the AOBA-A and AOBA-B subsystems adopting NEC SX-Aurora TSUBASA B401-8 and NEC LX406Rz-2, respectively, the latest subsystem of AOBA-S adopts NEC SX-Aurora TSUBASA C401-8. AOBA-A and AOBA-S are vector supercomputers that offer a good balance of computing power and memory performance. They are particularly effective for scientific and technical calculations, which are often memory intensive. For this reason, AOBA-A and AOBA-S are expected to mainly run user-developed code. AOBA-B is an x86 server using AMD's EPYC processors. It mainly runs open-source software and commercial applications. The nodes of AOBA-A and AOBA-B are connected via a high-speed InfiniBand HDR network, sharing a file system with a total capacity of 2 PB. The nodes of AOBA-S are connected via a high-speed InfiniBand NDR200 network, sharing a file system with a total capacity of 4.5 PB.

Since 1997, Tohoku University's Cyberscience Center has worked closely with users and computer vendors to optimize user-developed programs. This effort resulted in many valuable user-developed applications for the NEC SX-ACE system, which was the previous generation system. Because the new AOBA system architecture is significantly different, we have been working with users from the outset to help them 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 its high memory bandwidth) for scientific computing. We also help accelerate new users' computer programs.

In the event of a magnitude 7.0+ earthquake, AOBA can run emergency simulations to predict potential damage from ensuing tsunamis. These predictions are immediately forwarded to the Japanese government and other agencies to enable prompt and accurate response. Thus, AOBA is not only for academic research. It also plays a critical and prominent role in mitigating damage to societal infrastructure during emergencies.

For more details, please refer to the Cyberscience Center page on Tohoku University's website below.

https://www.tohoku.ac.jp/en/news/university_news/tohoku_university_unveils_new_supercomputer_aoba.html

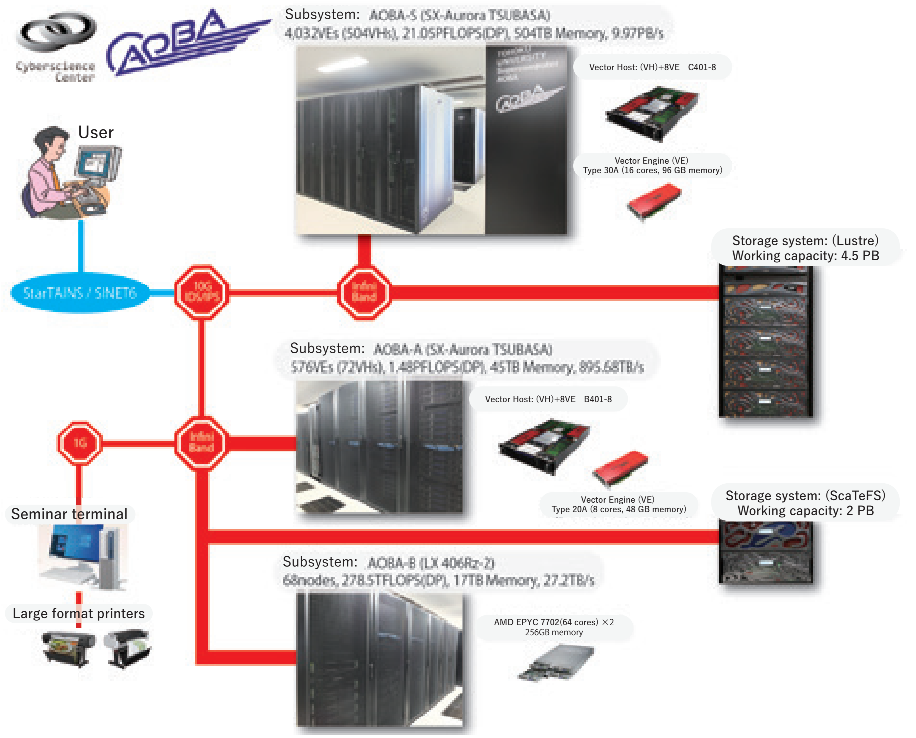

Center for Computational Sciences, University of Tsukuba

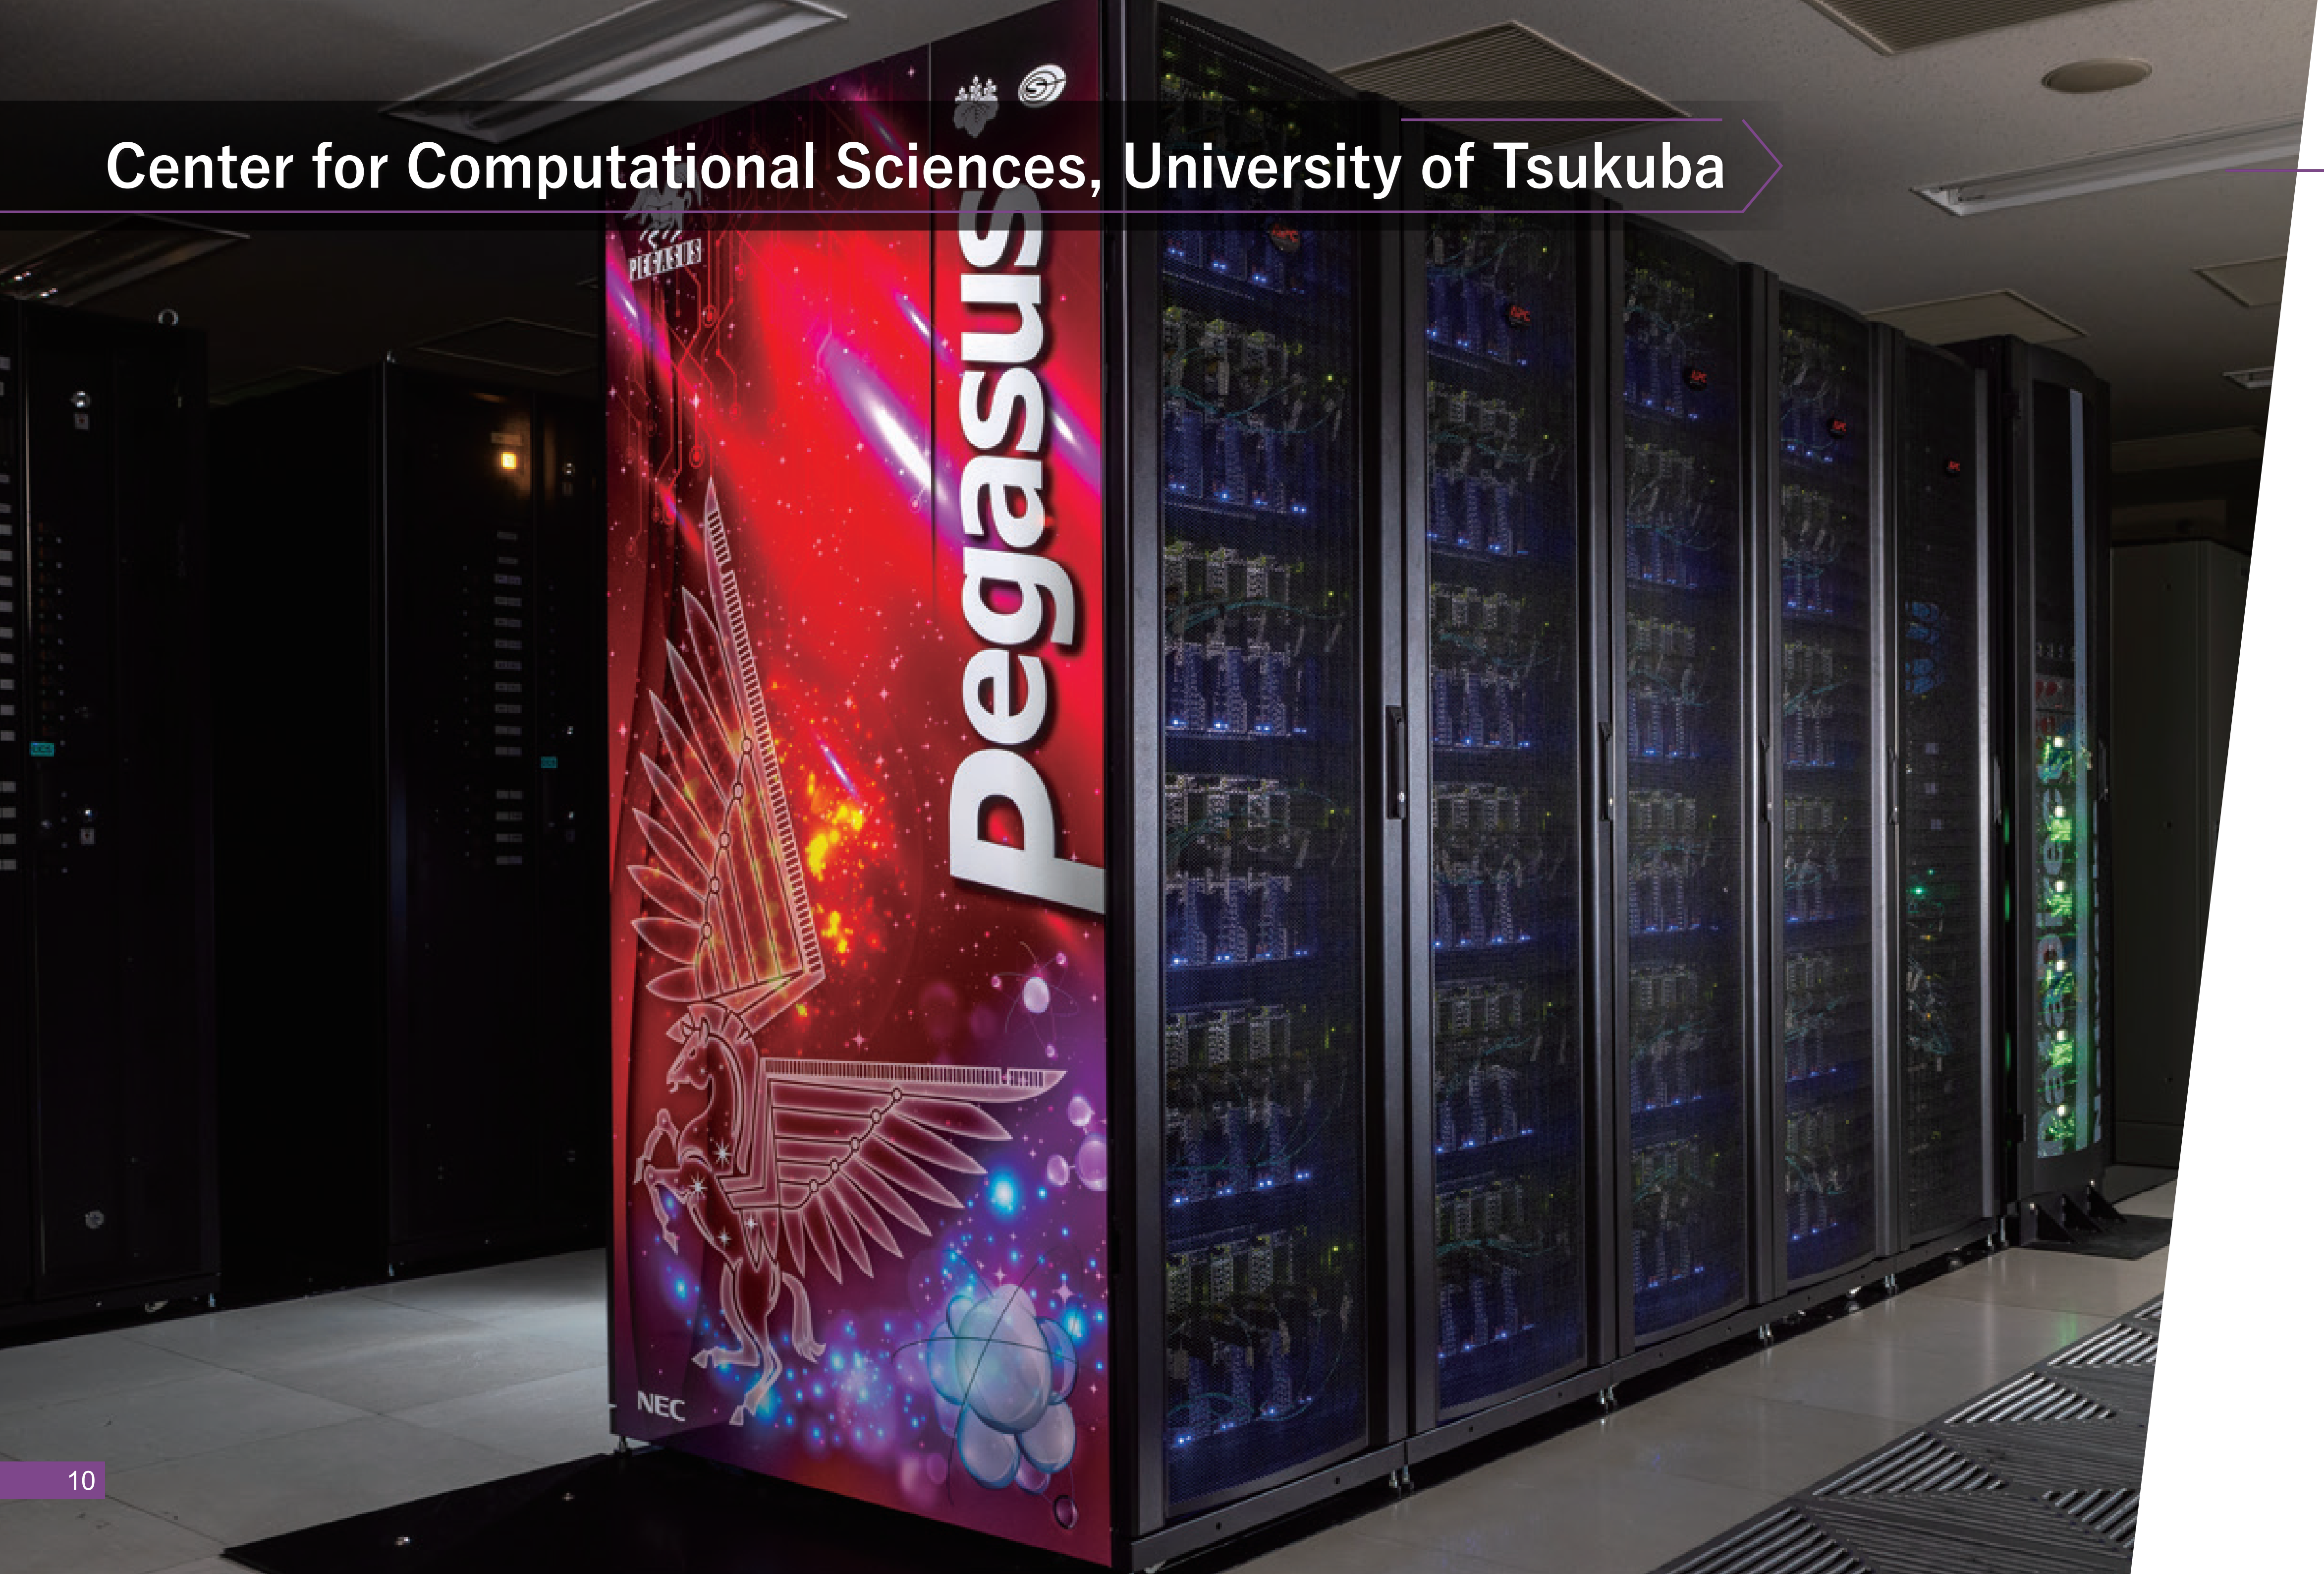

Big-Memory Supercomputer Pegasus and Unified-Memory Supercomputer Sirius (PACS12.0)

Osamu Tatebe

This year, the Center for Computational Sciences (CCS) at the University of Tsukuba began operations of the unified-memory supercomputer Sirius (PACS12.0), in addition to the big-memory supercomputer Pegasus.

Pegasus computing nodes are equipped with one fourth-generation Intel Xeon CPU and one NVIDIA H100 PCIe GPU. The system comprises 150 nodes, delivering a theoretical peak performance of 8.1 PFlops in double-precision floating-point operations. Each node is equipped with 128 GiB of DDR5 memory, 2 TiB of persistent memory, and a 6.4 TB NVMe SSD. Persistent memory can be used either as high-capacity memory or ultra-high-speed storage. Nodes are also interconnected via full-bisection InfiniBand NDR200 at 200 Gbps. A 7.1 PB parallel file system enables high-speed access to large-scale datasets. With significantly enhanced computing performance, memory bandwidth, and memory capacity, Pegasus can strongly promote not only computational sciences but also big data analysis and ultra-large-scale AI workloads.

Sirius (PACS12.0) computing nodes are comprised of four AMD MI300A APUs. Each APU integrates a 24-core EPYC Zen 4 CPU, a CDNA3 GPU, and 128 GB of HBM3 high-bandwidth memory. For four APUs, each node achieves a theoretical peak performance of 496 TFLOPS in double-precision calculations and offers 512 GB HBM3 memory. Each node is also equipped with four sets of 3.84 TB PCIe Gen5 NVMe SSDs, allowing it to meet the demands of large-scale data processing, HPC, and AI workloads. Sirius consists of 24 nodes, offering a total theoretical peak performance of 11.9 PFlops. Each computational node is connected via four InfiniBand NDR links (400 Gbps). 5.2 PB parallel file system enables it to access massive datasets at high speed. The parallel file systems of Pegasus and Sirius are interconnected via InfiniBand, enabling high-speed cross-system access.

To contribute to the advancement of nationwide interdisciplinary computational sciences, the CCS at the University of Tsukuba promotes the Multidisciplinary Cooperative Research Program (https://www.ccs.tsukuba.ac.jp/eng/use-computer/mcrp/). The CCS contributes to the development of computational and computer sciences by providing access to Pegasus, Sirius, and Miyabi (JCAHPC) as freely available computer resources in this program. The CCS was also selected under the MEXT 2023 Promotion of Development of a Joint Usage/Research System Project: Coalition of Universities for Research Excellence Program (CURE) as “the interdisciplinary hub for the social implementation of computational science in the AI era.” In resolving issues using supercomputers and computational science methods, the CCS aims to uncover new needs through collaboration among industry, government, and academia, and to expand the adoption of these methods within the private sector. As part of this project, we conduct supercomputer trial programs at the interdisciplinary hub, and companies that conduct continuous, joint research with the CCS can use its resources free of charge. In addition, fee-based research use is also available in the general access program and the industrial access program . For system details, please visit the following site:

https://www.ccs.tsukuba.ac.jp/eng/supercomputers/

Sirius (PACS12.0) Unified Memory Supercomputer

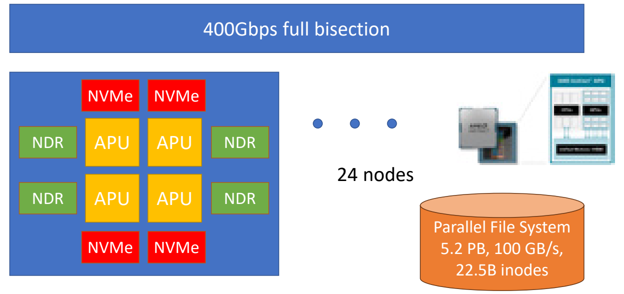

Joint Center for Advanced High Performance Computing (JCAHPC)

Supercomputer Miyabi

Shigeru Chiba, Yasuteru Shigeta

In March 2013, the University of Tsukuba and the University of Tokyo entered into an Agreement on Collaboration and Cooperation for the Advancement and Promotion of Computational Science and Engineering. Under this agreement, the Center for Computational Sciences (CCS) at the University of Tsukuba and the Information Technology Center (ITC) at the University of Tokyo established the Joint Center for Advanced High Performance Computing (JCAHPC). The JCAHPC, led mainly by staff from both institutions, is an organization that designs, installs, and operates large-scale leading-edge supercomputer systems as high-performance computing platform infrastructure within the ITC on the University of Tokyo's Kashiwa Campus. The collaborative work done by this facility helps to promote state-of-the-art computational science and contribute to the advancement of science and technology in Japan.

On January 14, 2025, the JCAHPC began operation of its next-generation supercomputer system, named "Miyabi". Miyabi is a large-scale cluster system of 80.1 PFlops by double-precision floating-point number. It consists of 1,120 "Miyabi-G" compute nodes featuring NVIDIA GH200 Grace-Hopper Superchip with NVLink-C2C ultra-high-speed links between CPU and GPU, and 190 "Miyabi-C" compute nodes each equipped with two Intel Xeon Max 9480 CPUs. All these nodes are interconnected by InfiniBand NDR200. This is Japan's first general-purpose large-scale system to use the GH200 Superchips. The Miyabi supercomputer also has a parallel file system with 11.3 PB of storage using NVMe SSDs exclusively.

Over the past decade, supercomputing has evolved considerably. Along with the traditional simulations, supercomputers are now driving diverse applications such as data analysis, data assimilation, machine learning, and AI. The University of Tsukuba and the University of Tokyo have lead the way in addressing these diversifying needs. They have pursued innovative R&D in supercomputer systems, software, and applications, integrate "simulation, data, and learning". They have also contributed to the human-centric society envisioned by the Japanese government's "Society 5.0" initiative. Building on the foundation of the JCAHPC's first-generation Oakforest-PACS supercomputer system, the Miyabi system furthers "simulation, data, and learning". Its particular focus is science and engineering driven by computer simulations. It also aims to revolutionize scientific research through a new "AI for Science" approach based on generative AI. Also, the University of Tsukuba and the University of Tokyo are participating in the JHPC-Quantum project (https://jhpc-quantum.org/en/). This project conducts R&D on quantum-HPC hybrid collaborative environments, contributing to the spread of quantum computing and the development of new scientific frontiers. Miyabi is expected to connect to multiple actual quantum computers beginning in FY2026.

For more details about the Miyabi supercomputer, please refer to the website below.

https://www.jcahpc.jp/eng/

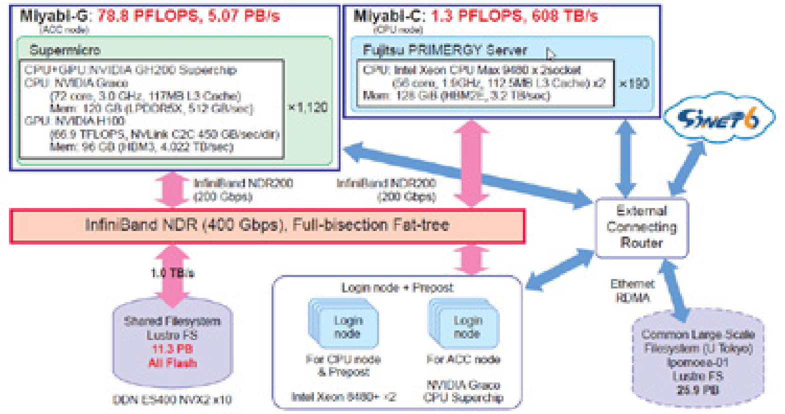

Information Technology Center, The University of Tokyo

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

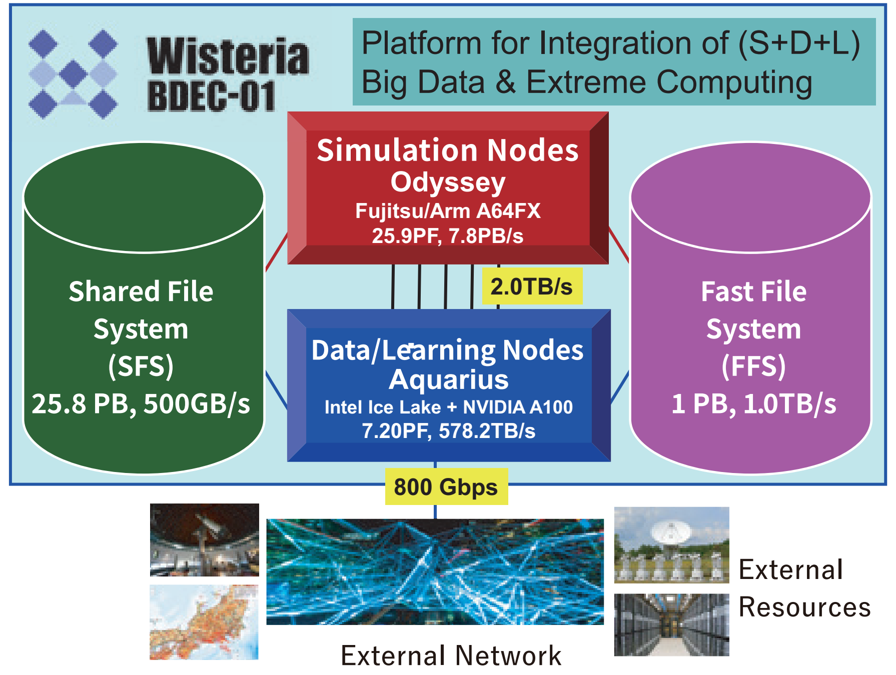

Center for Information Infrastructure, Institute of Science Tokyo

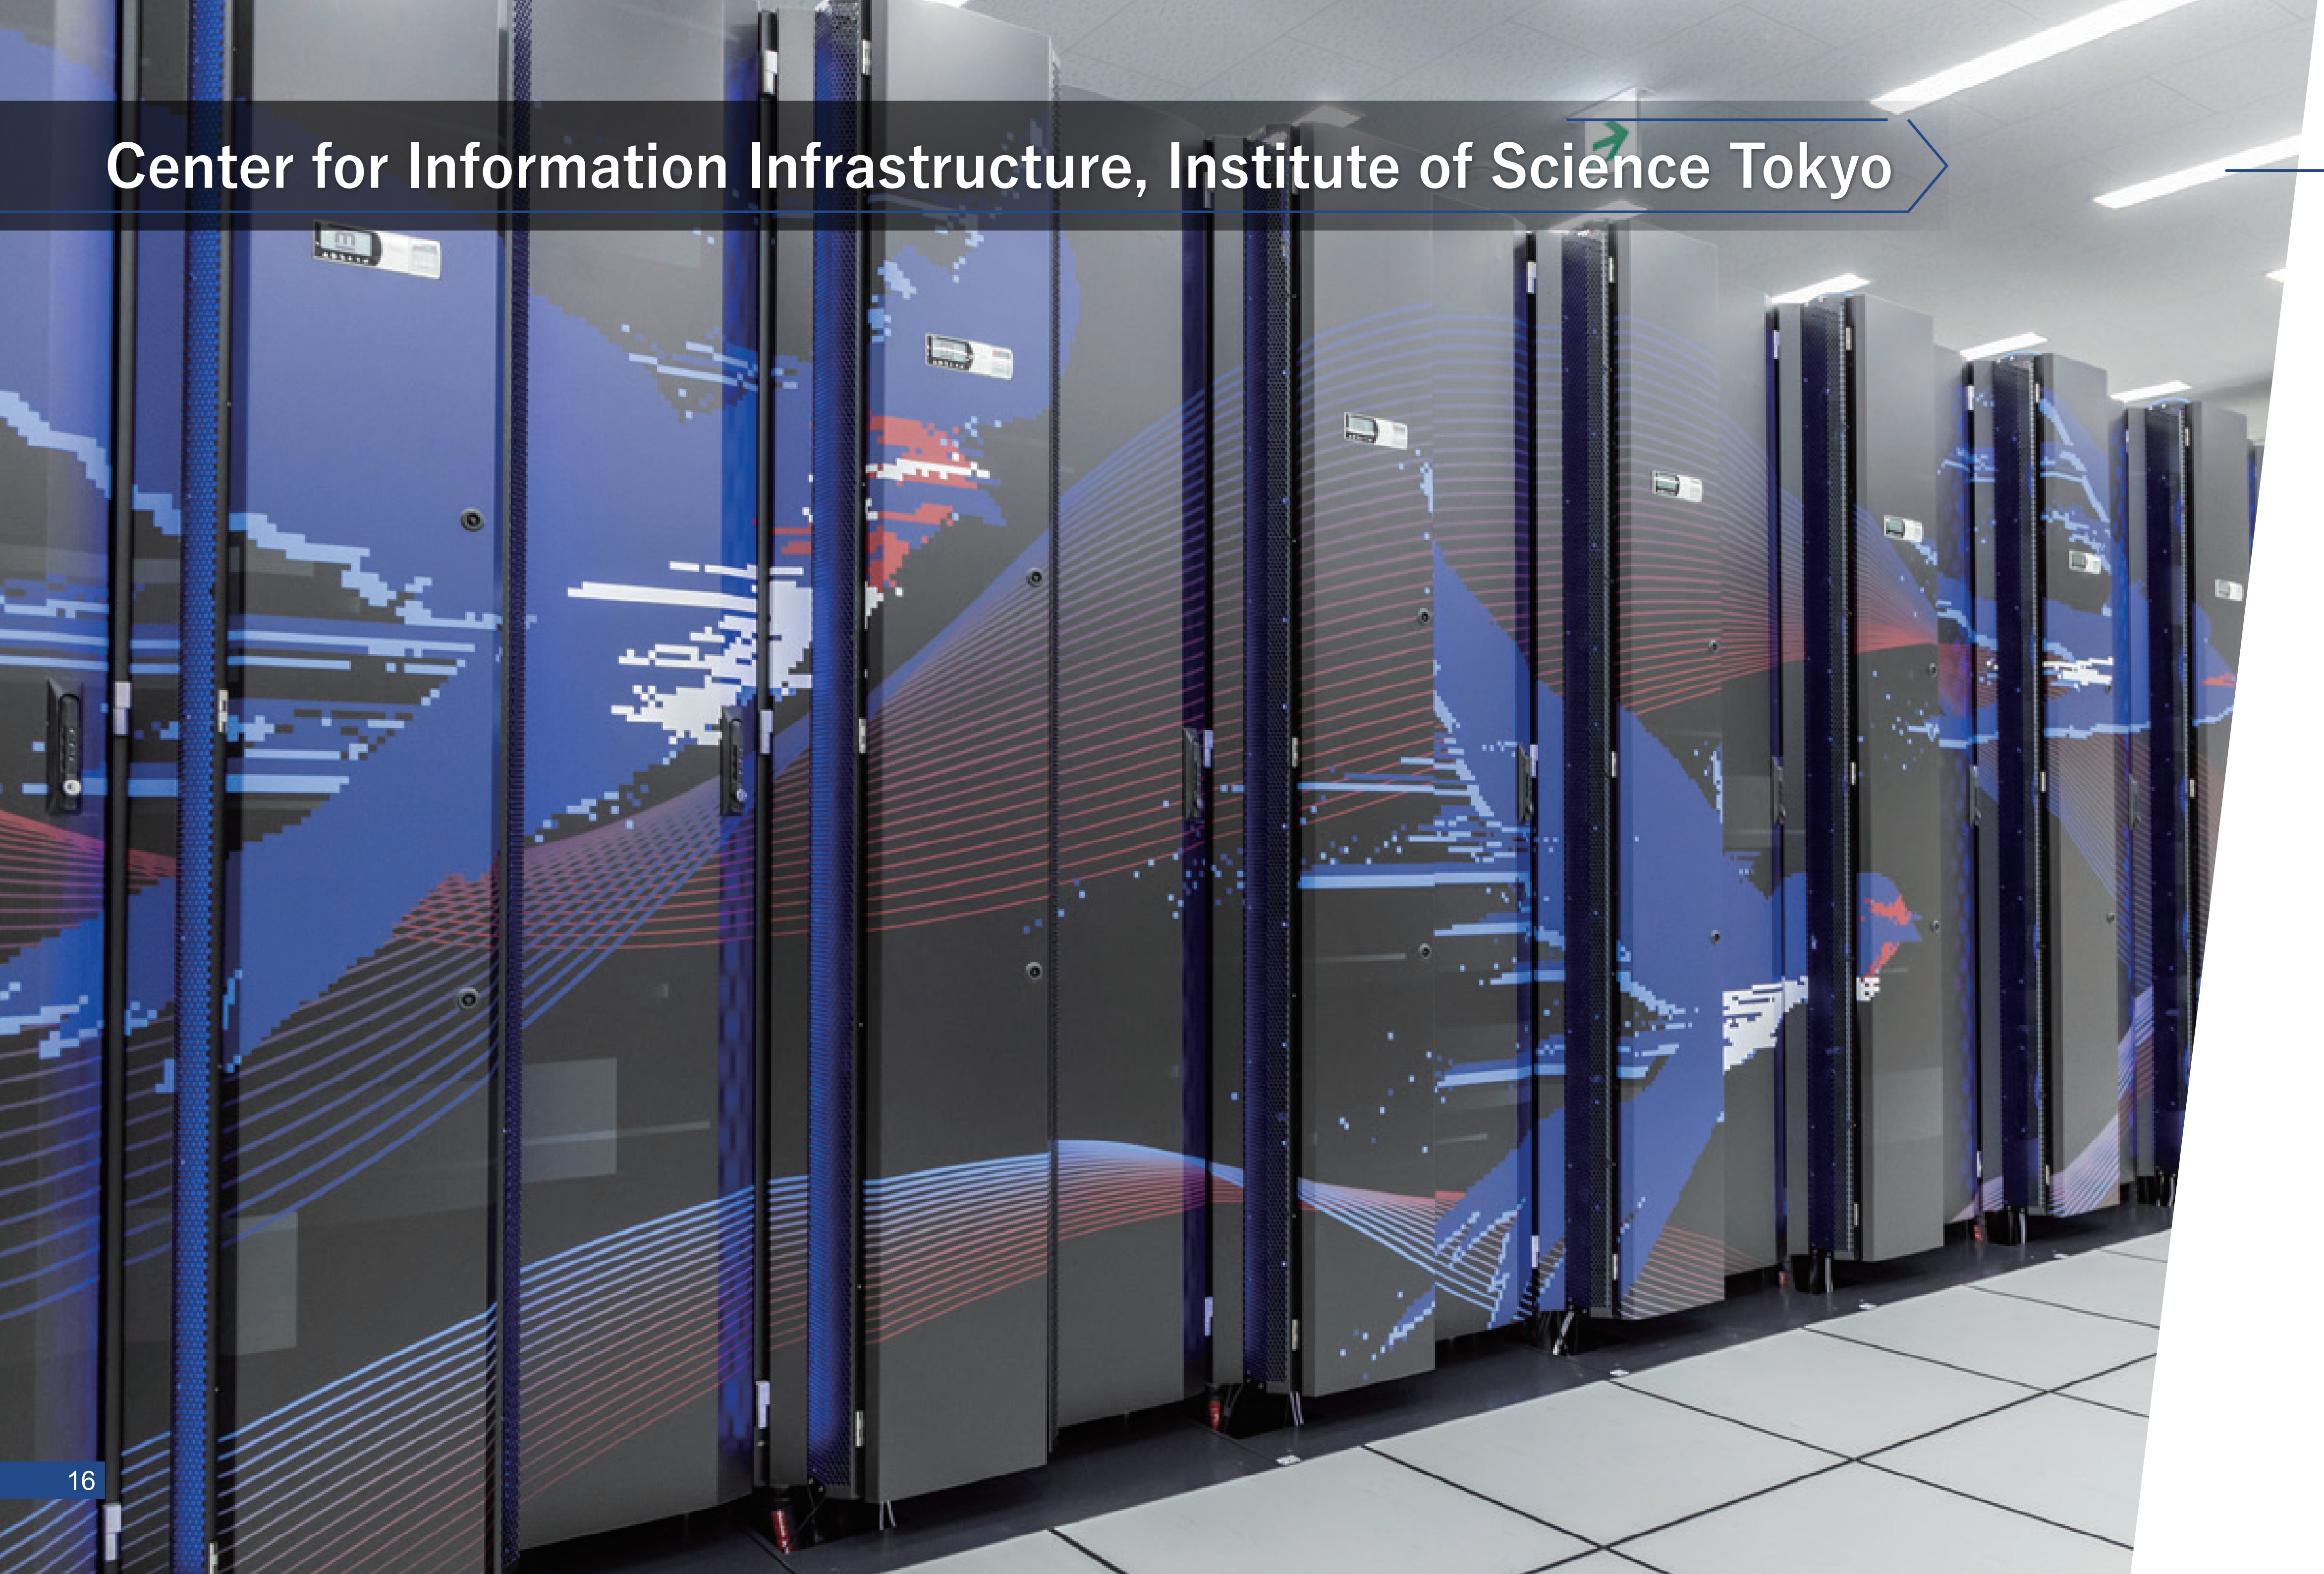

TSUBAME 4.0: A More Accessible Supercomputer for Everyone

Masahiko Tomoishi

Since 2006, the “TSUBAME” supercomputer series has been in use at the Global Scientific Information and Computing Center at the Tokyo Institute of Technology (currently, Institute of Science Tokyo). TSUBAME 1.2 achieved a milestone in 2008, becoming among the first supercomputers in the world to adopt GPUs on a large scale and provide GPU-based supercomputing services. In April 2024, TSUBAME 4.0 replaced TSUBAME 3.0, in service since 2017.

TSUBAME 4.0 delivers a total theoretical performance of 66.8 PFlops for the 64-bit double-precision matrix operations commonly used in scientific computing. It also boasts an impressive 952 PFlops for the 16-bit half-precision operations frequently used in artificial intelligence. The new system consists of 240 compute nodes connected via four 200 Gbps ports of high-speed networking. Each compute node is equipped with two AMD EPYC 9654 CPUs, four NVIDIA H100 Tensor Core GPUs, 768 GiB of main memory, and 1.92 TB of SSD capacity per node. It also provides 44.2 PB of HDD-based shared storage and 327 TB of SSD-based shared storage. In November 2024, TSUBAME 4.0 took 6th place in the HPL-MxP benchmark global ranking, which measures the performance of the newest low-precision computing algorithms running on the latest hardware.

Expanding on the TSUBAME series concept of a "supercomputer for everyone", TSUBAME 4.0 allows for innovative new use cases, including access through web applications via Open OnDemand. These user-friendly enhancements make TSUBAME 4.0 increasingly suitable for regular use by students and researchers across a range of fields. Equipped with 960 of the latest GPUs, by leveraging its GPU-level logical partitioning and Linux resource partitioning, it provides greater performance and usability than previous-generation systems.

From the outset, TSUBAME 4.0 has been made available as a supercomputing resource through the HPCI initiative. Additionally, to encourage large-scale computations producing results that are only possible by using the entire system at once, the center supports a Grand Challenge Large-Scale Computation program. The center’s other initiatives to broaden the base of supercomputer users and promote HPC in Japan include a seed funding program for young and female researchers, and proactive outreach to industry.

For more details, please refer to the TSUBAME Computing Services website below.
https://www.t4.cii.isct.ac.jp/en

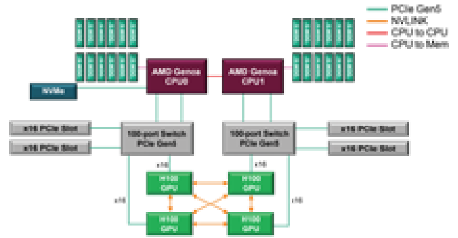

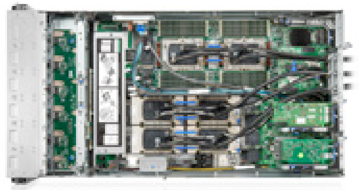

HPE Cray XD665 Server × 240	
CPU	AMD EPYC 9654 (96 cores, 2.4GHz) × 2 Socket
GPU	NVIDIA H100 SXM5 × 4 FP64 33.5 TFlops, FP64 Tensor 66.9 TFlops, FP32 66.9 TFlops TF32 Tensor 494.7 TFlops, FP16/BF16 Tensor 989.4 TFlops, INT8 Tensor 1978.9 Tops Memory 94 GB HBM2e 2395.87 GB/s
Memory	768GiB (DDR5-4800)
Local storage	1.92TB NVMe U.2 SSD
Network	InfiniBand NDR200 × 4

TSUBAME 4.0 compute node configuration

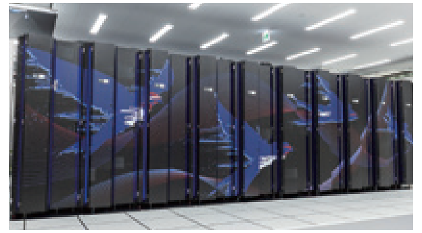

The TSUBAME 4.0 supercomputer

The new TSUBAME 4.0 data center on the Institute of Science Tokyo’s Suzukakedai Campus

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/

The supercomputer "Flow" will conclude operations in March 2026, with a new system scheduled to begin operation in October 2026.

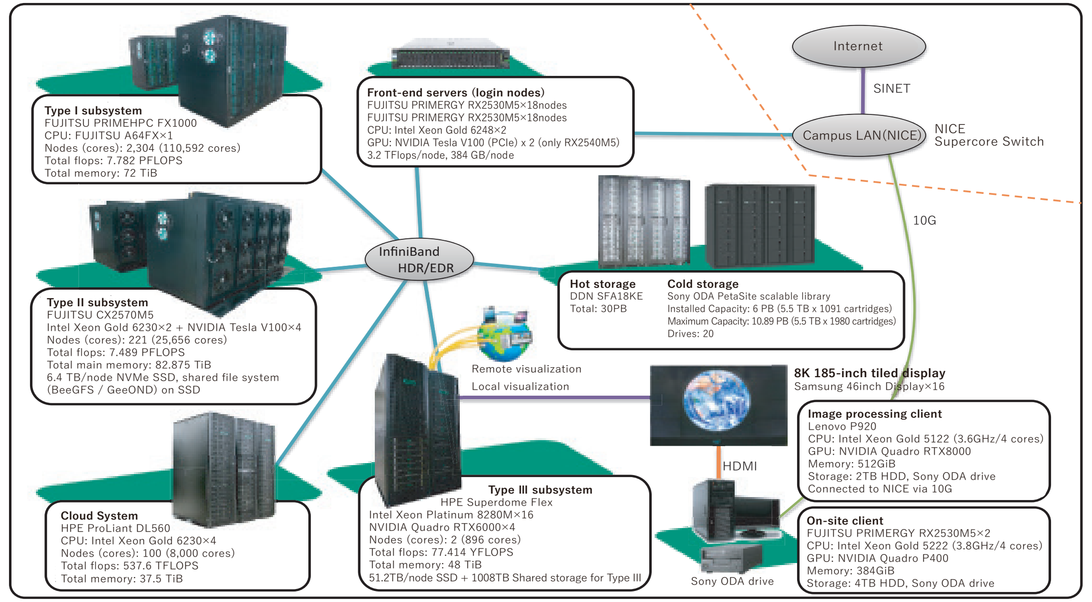

Academic Center for Computing and Media Studies, Kyoto University

Kyoto University Supercomputer Systems (Camphor / Laurel / Cinnamon / Gardenia)

Keiichiro Fukazawa

Kyoto University's Academic Center for Computing and Media Studies runs four supercomputer systems: Operating since May 2023, the Laurel 3, Cinnamon 3, and Gardenia systems were joined in October 2023 by Camphor 3.

- Laurel 3 is a general-purpose supercomputer with Dell PowerEdge C6620 rack servers with Intel Xeon Platinum 8480+ processors.
- Cinnamon 3 is a memory-intensive use Laurel 3 configuration supercomputer with two terabytes of memory per node.
- Gardenia is an AI and machine learning system with DELL PowerEdge XE8545 rack servers equipped with NVIDIA A100 Tensor Core GPUs.
- Camphor 3 is a computational performance-oriented supercomputer with Dell PowerEdge C6620 rack servers equipped with Intel Xeon CPU Max 9480 processors with HBM2e high-bandwidth memory.

The Camphor 3's 1,120 nodes yield a total theoretical performance of 7.63 PFlops, while the Laurel 3's 370 nodes yield 2.65 PFlops. The Cinnamon 3 has only 16 nodes, but with four times as much memory as the Laurel 3 configuration. Gardenia also has 16 nodes, but each node is equipped with four A100 80GB GPUs. The different features of these four supercomputers are intended to provide an environment that can meet the diverse computational needs of each user. A portion of these supercomputing resources support 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/

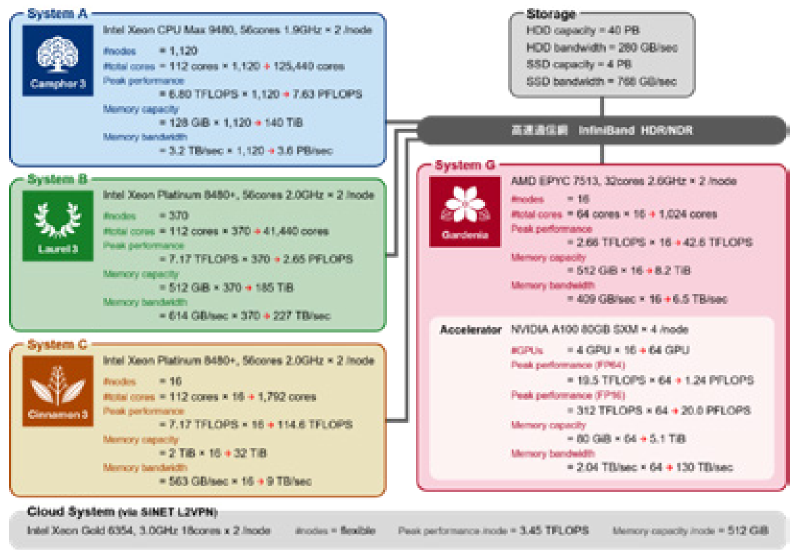

D3 Center, The University of Osaka

Introduction to the University of Osaka D3 Center High-Performance Computing Systems

Keiichi Takahashi

The University of Osaka D3 Center offers high-performance computing systems to enhance academic research and promote open science: the supercomputer SQUID (picture) and the next-generation computing and storage infrastructure OCTOPUS.

The supercomputer SQUID (Supercomputer for Quest to Unsolved Interdisciplinary Datascience), which began operation in May 2021, is a hybrid cluster system with a total theoretical computing performance of 16.591 PFlops. It comprises general-purpose CPU nodes, GPU nodes, and Vector nodes. A Lustre file system, consisting of HDDs (20.0 PB) and SSDs (1.2 PB), is available through DDN's EXAScaler. Processors and accelerators across all node groups are cooled using a direct liquid-cooling system and are designed to ensure stable, high-performance operation over the long term. SQUID enables the use of a variety of computing nodes with different architectures in the same environment, flexibly supporting computing needs across a wide range of research fields.

OCTOPUS (Osaka university Compute & sTOrage Platform Urging open Science), which began operations in September 2025, is a new computing and data platform built on the concept of promoting open science. It is composed of general-purpose CPU nodes equipped with two sixth-generation Intel Xeon Scalable processors and has a total theoretical computing performance of 2.293 PFlops. OCTOPUS is also equipped SCUP-HPC (System for Constructing and Utilizing Provenance on High-Performance Computing system), a provenance management system, developed jointly by D3 Center and NEC. This system automatically records and manages the computational and data generation processes, thereby ensuring the integrity of academic research and supporting the advancement of open science.

These systems have been built to support a variety of researchers' needs across both high-performance computing (HPC) and high-performance data analysis (HPDA). Our center's faculty and technical staff are fully committed to supporting the users of our high-performance computing systems. We welcome you to explore our services.

For more details, please refer to the University of Osaka website below.

https://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

SQUID system configuration

Theoretical Computing Speed	2.293 PFlops	
Nodes	General-purpose CPU nodes 140 nodes	CPU: Intel Xeon 6980P (Granite Rapids/2.0 GHz 128 cores) x 2 Memory : 768 GB
Storage	DDN EXAScaler (Lustre)	HDD: 3.58 PB

OCTOPUS system configuration

Research Institute for Information Technology, Kyushu University

The Genkai Supercomputer

Kazuki Yoshizoe

From July 2024, the Research Institute for Information Technology at Kyushu University has made available its “Genkai” supercomputer system, based on Fujitsu’s PRIMERGY CX2550 M7 server series. Equipped with Intel “Sapphire Rapids” CPUs and NVIDIA H100 GPUs, Genkai boasts a total theoretical performance of about 13 PFlops, making it one of the most powerful systems in Japan.

The Genkai supercomputer meets the needs of emerging research fields such as Large Language Models and contributes to the data-driven research and open science promoted under Japan’s 6th Basic Plan for Science, Technology, and Innovation, while also making accommodations for possible further computing demands in the future. As with its predecessor the “ITO” supercomputer in operation until February 2024, Genkai consists of high-performance backend compute node clusters suitable for large-scale simulations and machine learning, all connected to a high-speed file system. This system also supports Amazon S3 APIs and NextCloud-compatible cloud storage interfaces as well as Open OnDemand for browser-based access, thereby providing the high availability required for “open science”.

Kyushu University will use the Genkai supercomputer system as part of its efforts to become a “university that drives social change with integrative knowledge”, as outlined in the Kyushu University VISION 2030 initiative. It will also make this supercomputing resource available to the Japan High Performance Computing and Networking initiative (JHPCN), the HPCI initiative, and various programs run by its Research Institute for Information Technology, thereby providing a foundation for academic research in Japan and advancing new research endeavors.

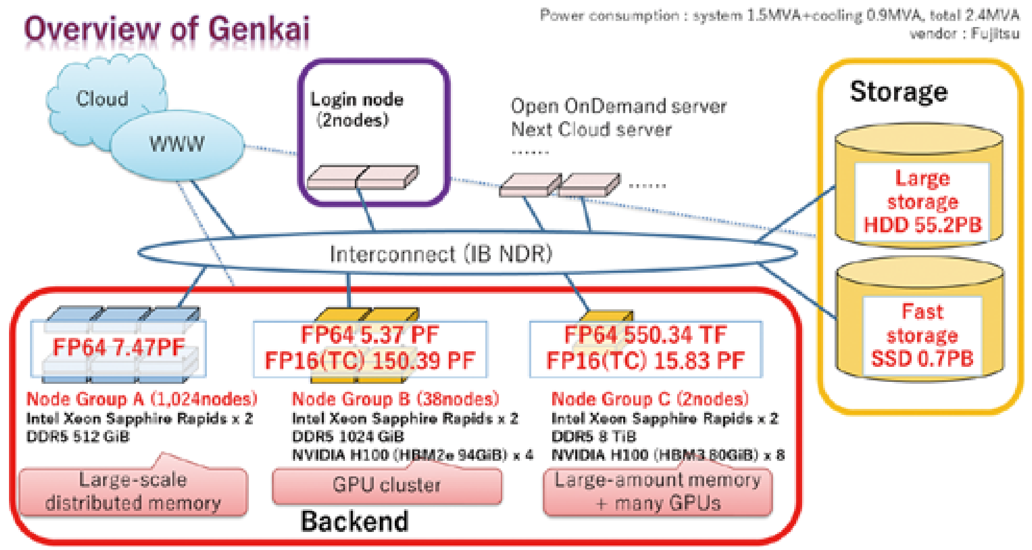

Center for Earth Information Science and Technology, Japan Agency for Marine-Earth Science and Technology

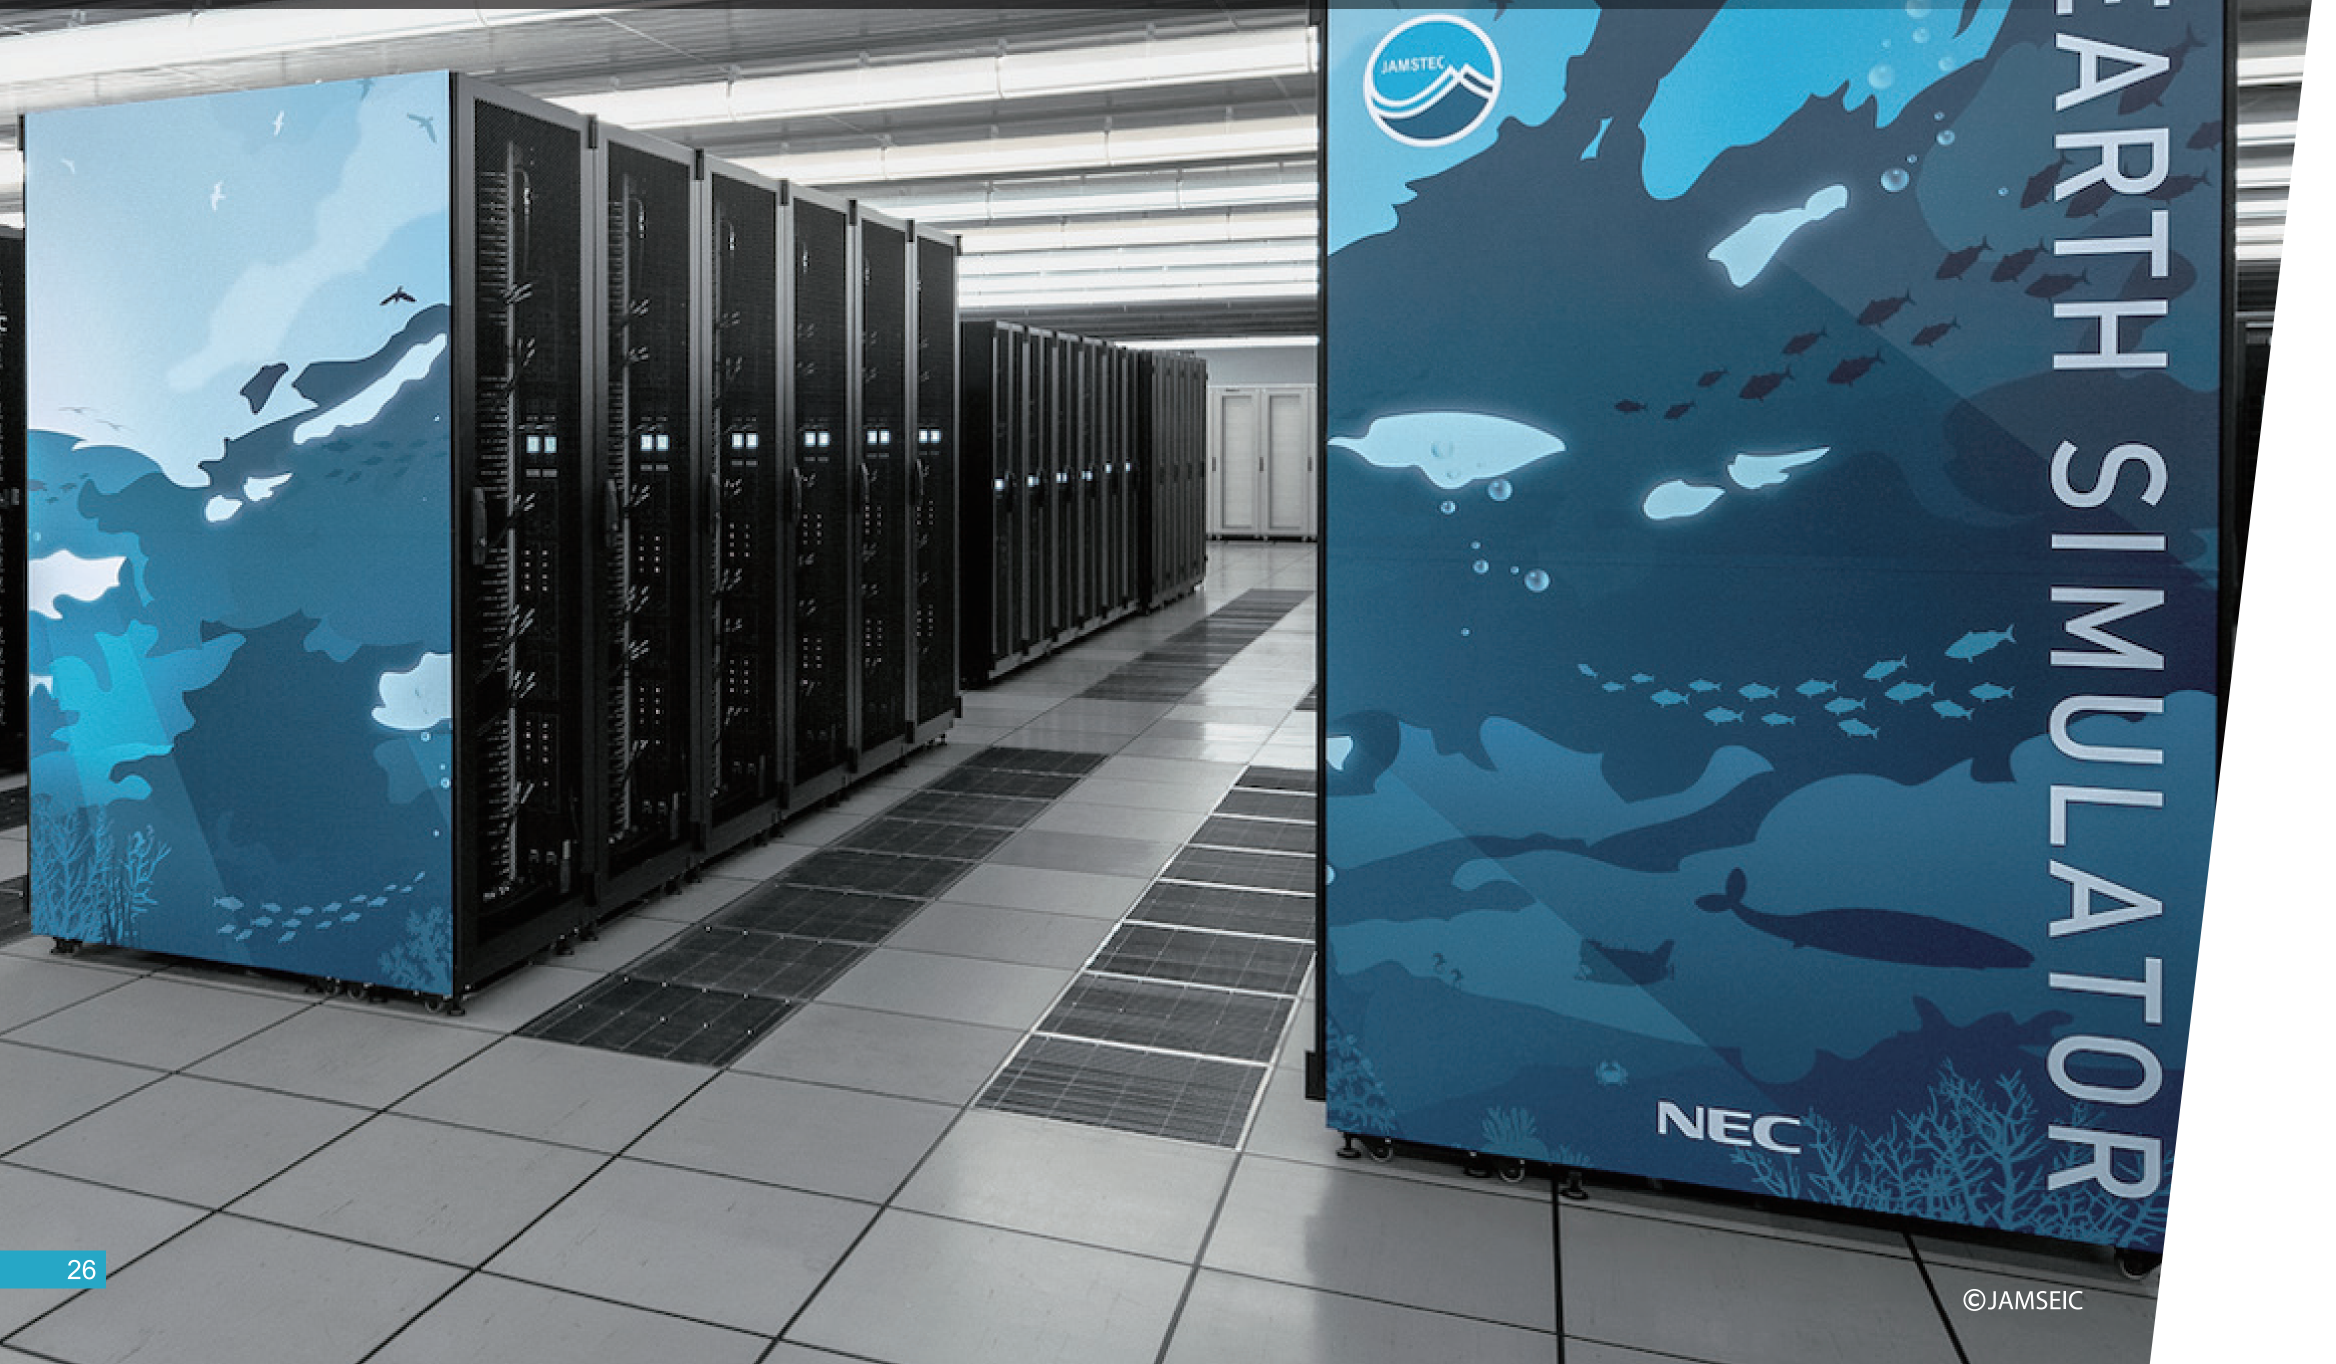

©JAMSEIC

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, 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/

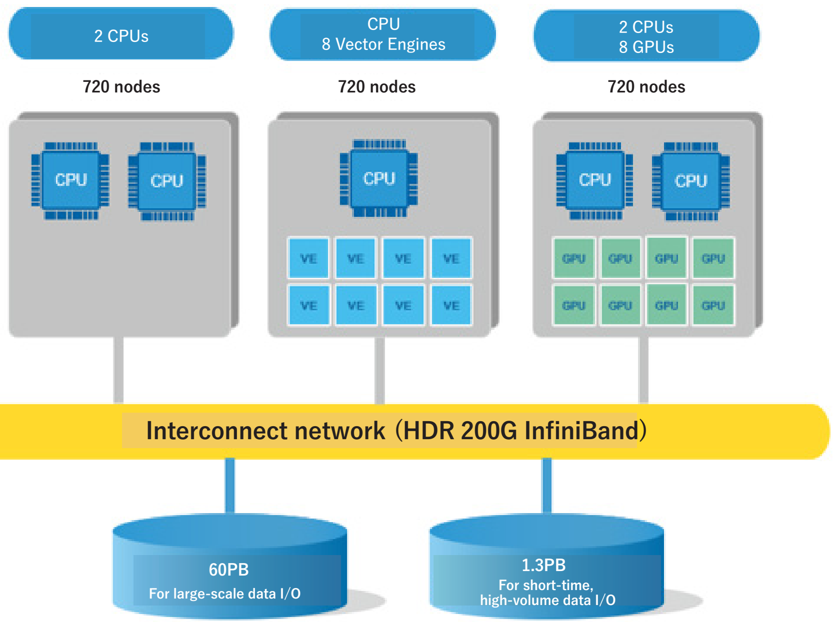

Center for Engineering and Technical Support, The Institute of Statistical Mathematics (ISM)

Supercomputer System for Data Assimilation

Genta Ueno

Data has ballooned in recent years. Obtaining useful insights requires analysis of massive amounts of data from both real-world measurements and computer simulations. Moreover, the integration of data from those measurements and simulations, known as “data assimilation”, tends to compound this problem, producing ensemble data and reanalysis data that then have to be analyzed.

Working with such massive amounts of data is not easy, largely because most supercomputers rely on distributed memory allocation and can’t fit all this massive data inside a single memory space. Analyzing big data on these distributed memory systems requires explicitly programming “parallel processing” software that can separate, compute, and then reintegrate the data. But, coding parallelization programs is rote and time-consuming work that doesn’t change the final results of the analysis. This work is a hindrance that should be avoided if possible. Doing so requires a supercomputer equipped with massive amounts of memory that can be used by any of its CPUs.

In March 2023, the Center for Engineering and Technical Support at the Institute of Statistical Mathematics launched its “Supercomputer System for Data Assimilation” to advance the analysis of large-scale data without parallel programming. This is a distributed shared-memory computer with a large memory space that can be accessed from any CPU. This system is equipped with two HPE Superdome Flex computing nodes and has a total theoretical computing performance of 154.8 TFlops. Each node is equipped with 32 28-core CPUs (Intel Xeon Platinum 8280L) with a main memory of 48 TB, and an SSD with 880 TB of usable capacity. Since October 2020, one node of this system has been made available for use as a computing resource under Japan’s HPCI initiative.

https://www.hpci-office.jp/en/using_hpci/hardware_software_resource/2025/ism_2025-1

RIKEN Information R&D and Strategy Headquarters

HOKUSAI BigWaterfall2

Motoyoshi Kurokawa

In December 2023, the RIKEN Information R&D and Strategy Headquarters began operations of the supercomputer HOKUSAI BigWaterfall2 (HBW2). To encourage internal research and development work, RIKEN has operated supercomputer systems since the 1960s. The new HOKUSAI BigWaterfall2 consists of a massively parallel computing system with a total theoretical performance of 2.12 PFlops on 312 nodes equipped with two Intel Xeon Max (Sapphire Rapids) and HBM2e (128GB), a two-node large-capacity memory processing server with 3TB of memory, and 9.8 PB of Lustre file system shared storage. They are connected by a 400Gb/s InfiniBand NDR network. The biggest advantage is that the compute nodes are equipped with HBM2e, which feature a very high memory bandwidth of 3,260Gb/s.

The supercomputer systems operated by RIKEN's Information R&D and Strategy Headquarters have historically been used mainly to advance the research of scientists and engineers within the institute. However, given the increasing diversity and depth of the research, the computing resource requirements are similarly diversifying, making it difficult for any single organization to maintain all of them on its own. Japan's High-Performance Computing Infrastructure (HPCI) initiative, including supercomputer Fugaku, provides researchers with large-scale, diverse high-performance computing platforms. RIKEN's supercomputer HOKUSAI BigWaterfall 2 is provided as a high-performance computing resource, and from 2024 it has been made available through the HPCI initiative, allowing RIKEN scientists to learn more about HPCI itself and more easily choose the computing resources best suited to their research.

For more details, please refer to the RIKEN Information R&D and Strategy Headquarters website below.

https://i.riken.jp/en/supercom/

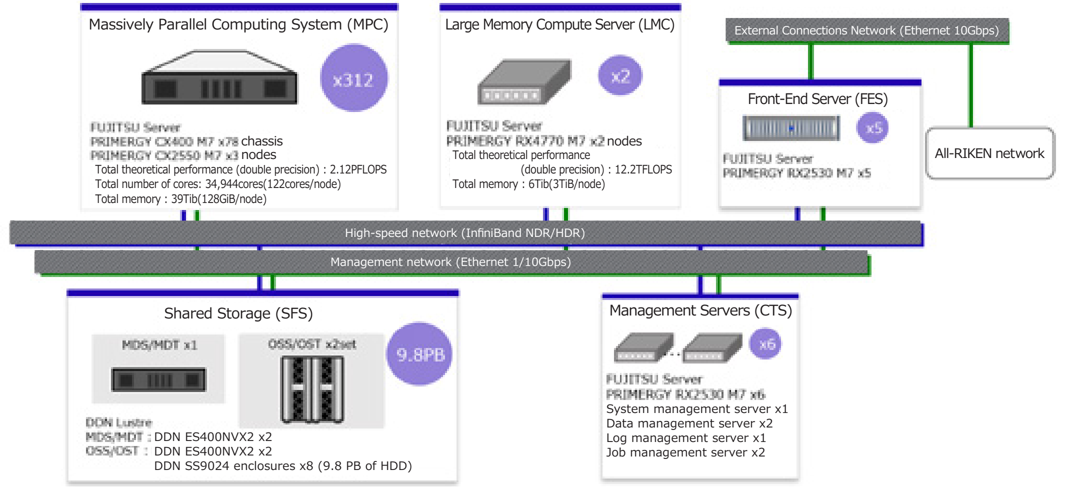

If you are considering using HPCI computing resources or would like to learn more about Japan's HPCI initiative, please refer to the HPCI portal site below, or contact our help desk.

HPCI Portal Site https://www.hpci-office.jp/en

Help Desk helpdesk@hpci-office.jp

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

The Kobe Center of the Research Organization for Information Science and Technology (RIST) 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
English version
Issued: January 2026

Research Organization for
Information Science and Technology
Kobe Center

https://www.hpci-office.jp/ristkobe/en/