HPCI Computing Resource Handbook

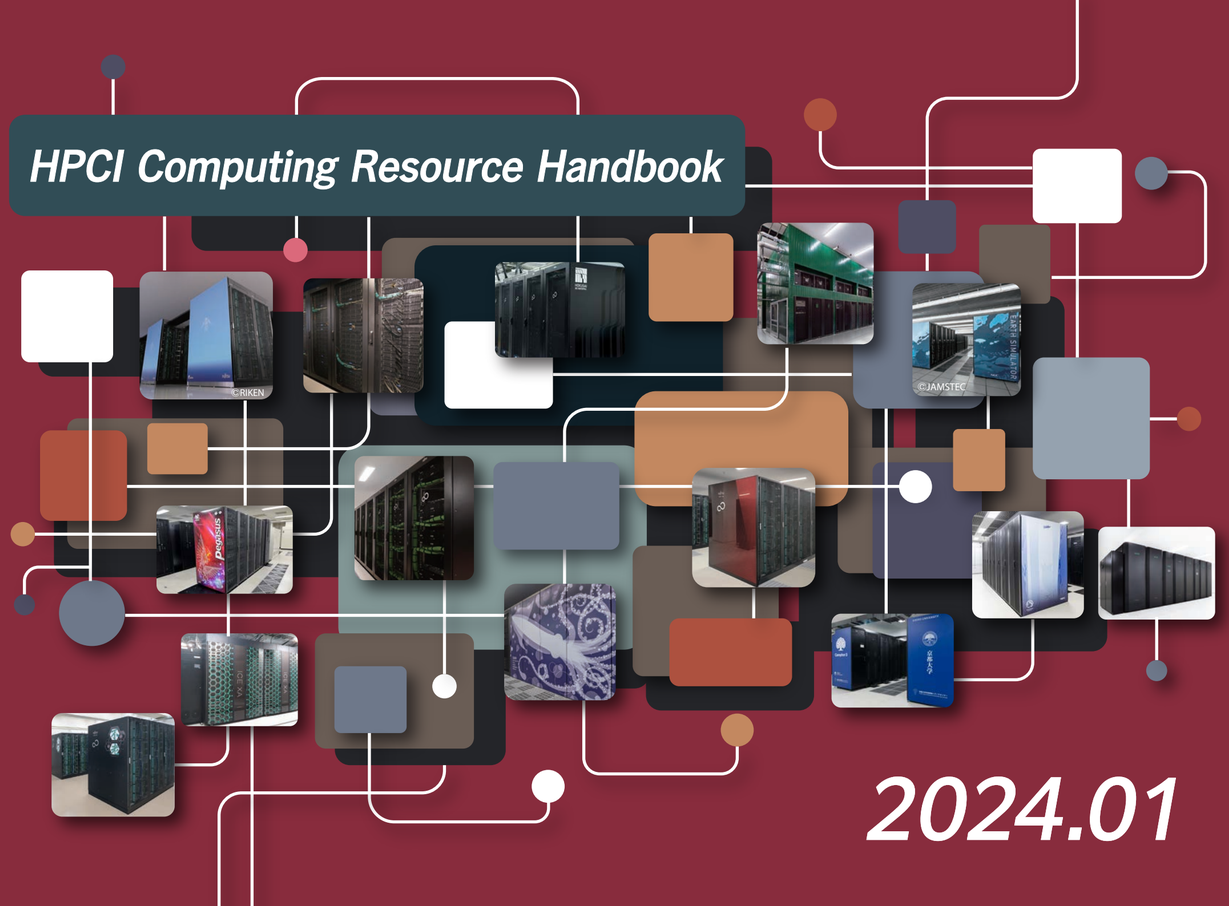

2024.01

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
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(ISM)
Information Technology Research Institute, National Institute of Advanced Industrial Science and Technology (AIST)
RIKEN Information R&D and Strategy Headquarters (R-IH)

CPU Architectures

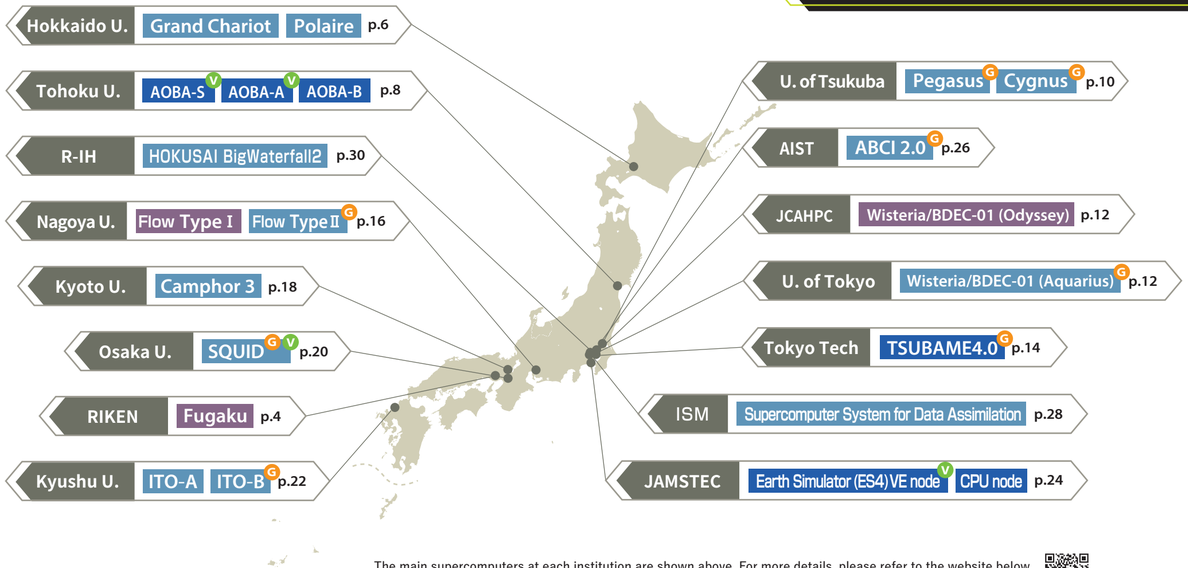

The main supercomputers at each institution are shown above. For more details, please refer to the website below.
https://www.hpci-office.jp/en/using_hpci/hardware_software_resource

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
G GPU	x86-64 host processors with NVIDIA HPC 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 2024, but subject to change.
Some institutions have restrictions on the use of commercial software.
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.		p.24 JAMSTEC	p.26 AIST	p.28 ISM	p.30 R-IH
		Fugaku	Grand Chariot	Polaire	AOBA-S (SX)	AOBA-A (SX)	AOBA-B (LX)	Pegasus	Cygnus	Wisteria (Odyssey)	Wisteria (Aquarius)	TSUBAME 4.0	Flow Type I	Flow Type II	Camphor 3	SQUID	ITO-A	ITO-B	Earth Simulator	ABCI 2.0	Supercomputer System for Data Assimilation	HOKUSAI BigWaterfall2
Molecular Dynamics	AMBER	○										○	○	○			○	○				○
	GENESIS	○	○				○		○	○	○		○	○	○	○	○	○				
	GROMACS	○	○	○			○			○	○	○	○	○	○	○	○	○			○	○
	LAMMPS	○					○			○	○	○	○	○	○	○	○	○			○	
	MODYLAS	○	○				○			○	○		○	○	○	○	○	○				
	MyPresto		○	○																		
	N2P2	○																				
	NAMD	○	○	○								○	○	○	○							○
	Tinker											○										
Quantum Chemistry	ABINIT-MP	○	○			○	○			○	○		○	○		○	○	○	○			
	GAMESS		○	○								○	○	○	○	○	○	○				○
	Gaussian	○	○				○					○	○	○	○	○	○	○				○
	GRRM						○															
	Molpro																○	○				
	NTChem	○	○				○			○	○		○	○	○	○	○	○				
	NWChem	○								○	○											
	Scigress																○					
	SMASH	○	○				○			○	○		○	○	○	○	○	○				
Condensed Matter Physics	ABINIT	○																				
	AkaiKKR	○	○				○		○	○	○		○	○	○	○	○					
	ALAMODE	○	○				○			○	○		○	○	○	○	○					
	CP2K	○									○	○										
	CPMD	○																				
	HΦ	○	○			○	○			○	○		○	○	○	○	○	○				
	mVMC	○	○				○			○	○		○	○	○	○	○					
	OpenMX	○	○				○		○	○	○		○	○	○	○	○	○				
	PHASE/0	○	○			○	○		○	○	○		○	○	○	○	○	○	○			
	Phonopy	○	○				○			○	○		○	○	○	○	○					
	Quantum ESPRESSO	○	○		○	○	○			○	○	○	○	○	○	○			○			
	SALMON	○	○				○			○			○	○	○	○	○	○				

Data is current as of January 2024, but subject to change.
Some institutions have restrictions on the use of commercial software.
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.		p.24 JAMSTEC	p.26 AIST	p.28 ISM	p.30 R-IH
		Fugaku	Grand Chariot	Polaire	AOBA-S (SX)	AOBA-A (SX)	AOBA-B (LX)	Pegasus	Cygnus	Wisteria (Odyssey)	Wisteria (Aquarius)	TSUBAME 4.0	Flow Type I	Flow Type II	Camphor 3	SQUID	ITO-A	ITO-B	Earth Simulator	ABCI 2.0	Supercomputer System for Data Assimilation	HOKUSAI BigWaterfall2
	SIESTA	○																				
	VASP	○										○					○	○				
Computational Biology	Alphafold											○		○								
	rDock	○																				
	Relion															○						
Fluid Analysis	COLMINA CAE (※1)	○																				
	CONVERGE	○																				
	Cradle CFD scFLOW	○																				
	FDS	○																				
	FFVHC-ACE	○					○			○	○		○	○	○	○						
	FFX	○					○			○	○		○	○	○	○						
	FrontFlow/blue	○	○			○	○			○	○		○	○	○	○	○	○	○			
	FrontFlow/red		○	○									○	○								
	OpenFOAM	○	○	○			○			○	○	○	○	○	○	○	○	○			○	
	STAR-CCM+	○																				
	V-FaSTAR		○																			
Structural / Collision Analysis	FrontISTR	○	○			○	○		○	○	○		○	○	○	○	○	○	○		○	
	LS-DYNA	○												○								
	Marc																○	○				
	Mentat																○	○				
	MSC Nastran																○	○				
	Patran																○	○				
Electromagnetic Field Analysis	COLMINA CAE (※2)	○																				
	Meep		○	○																		
	OpenFDTD	○																				
	Poynting	○																				
Multi-physics	HyperWorks										○			○								
Particle Systems	GEANT4									○	○		○	○								
Weather / Climate	SCALE	○																				
	WRF	○	○	○																		

※1…Particle-based Casting Simulator　※2…Magnetic Simulator

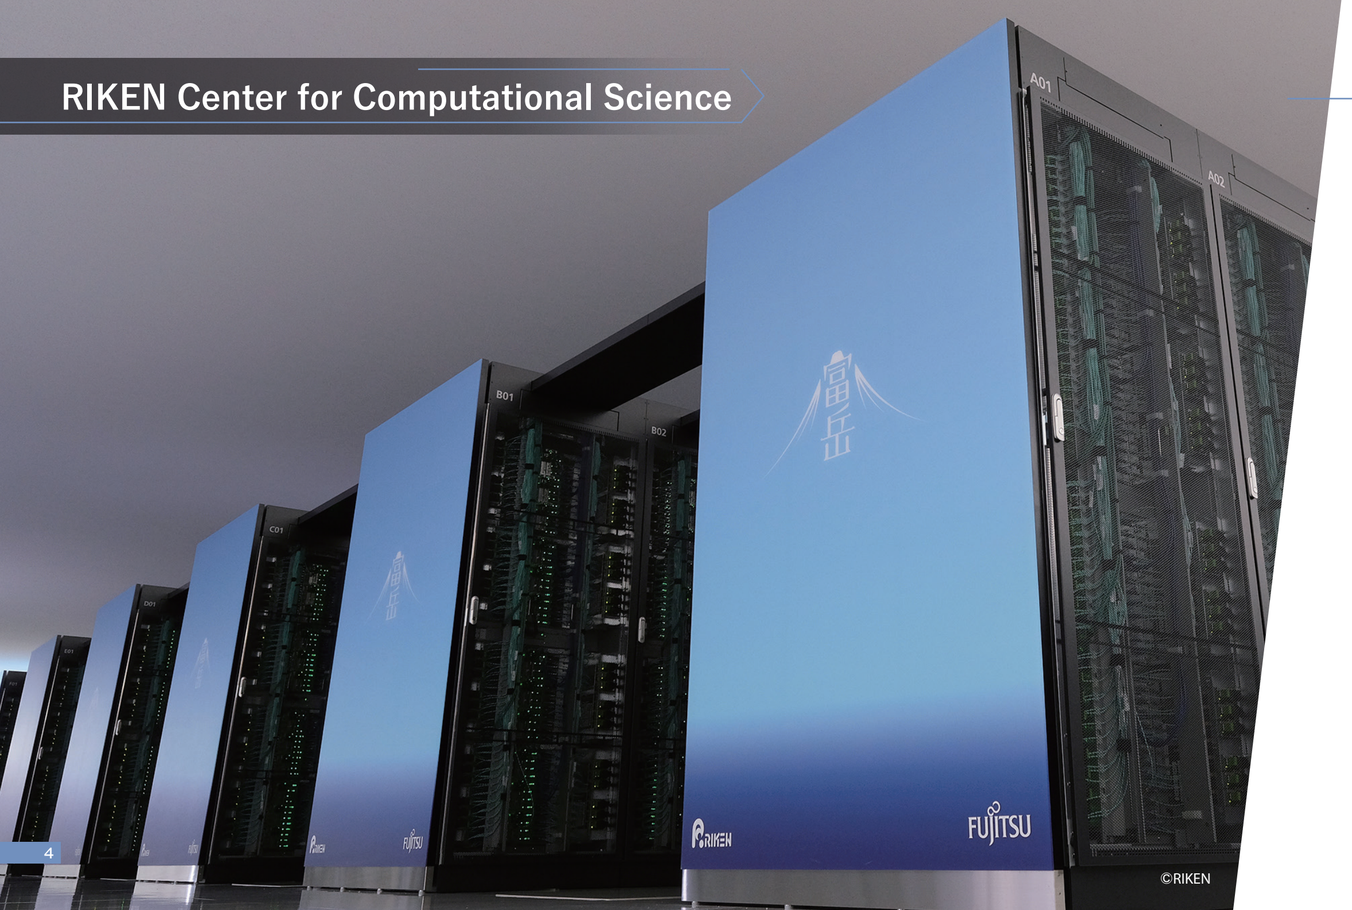

Supercomputer Fugaku

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

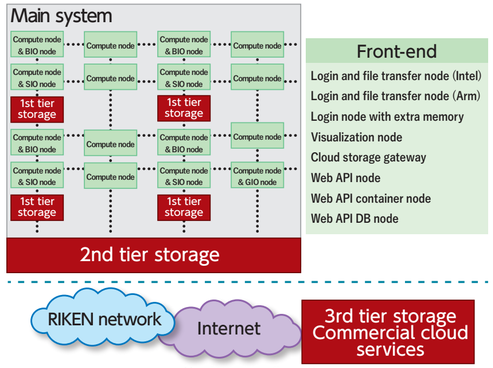

Hardware configuration

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 Initiative Center’s website below.
https://www.hucc.hokudai.ac.jp/en/overview/ilcs/

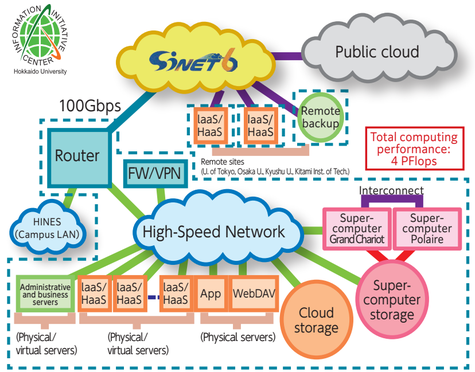

Cyberscience Center, Tohoku University

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

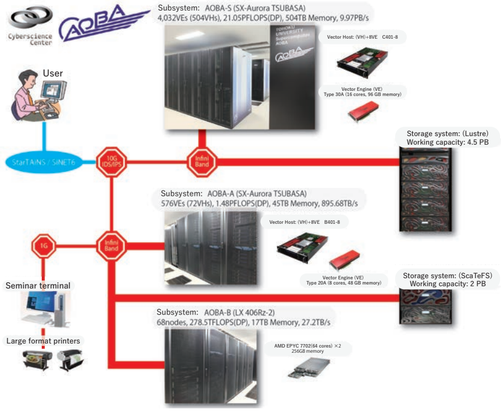

Center for Computational Sciences, University of Tsukuba

Pegasus & Cygnus

Taisuke Boku

The University of Tsukuba's Center for Computational Sciences (CCS) began operating the “big memory” supercomputer Pegasus in April 2023. Pegasus consists of 120 nodes with one NVIDIA H100 PCIe GPU per node, with an overall theoretical peak performance of 6.5 PFlops. Equipped with DDR5 memory and non-volatile memory, it can be used as large capacity memory or ultra-high speed storage. Its greatly improved computational performance, memory size, and memory bandwidth are major advances for computational science and the fields of big data analysis and extremely large-scale AI.

We also operate Cygnus, a massively parallel GPU/FPGA cluster system. Cygnus consists of a total of 78 nodes with 4 GPUs (NVIDIA Tesla V100) per node. The theoretical peak performance excluding FPGA is 2.34 PFlops. 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.

At the University of Tsukuba's CCS, we promote a “Multidisciplinary Cooperative Research Program” (https://www.ccs.tsukuba.ac.jp/eng/use-computer/mcrp/) to help expand interdisciplinary computational science work throughout Japan. By providing Pegasus and Cygnus as computing resources that can be used 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/

Supercomputer Cygnus(PACS-X)

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

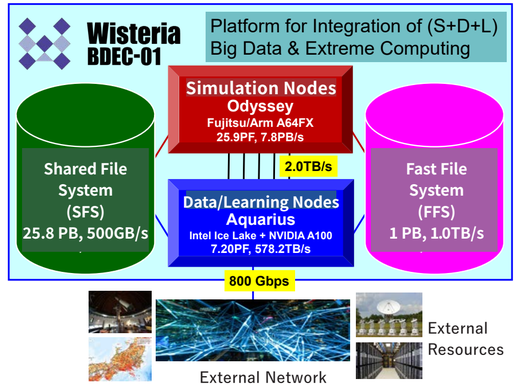

TSUBAME 4.0 (from April 2024)

Toshiya Itoh

Since 2006, the “TSUBAME” supercomputer series has been in use at the Global Scientific Information and Computing Center, Tokyo Institute of Technology. TSUBAME 1.2 achieved a milestone in 2008, becoming a globally pioneering supercomputer with integrated GPUs, and then providing GPU-based supercomputer resources as a service. In April 2024, the upcoming TSUBAME 4.0 will replace the current TSUBAME 3.0, operational since 2017.

TSUBAME 4.0 has a theoretical performance of 66.8 PFlops (64-bit double-precision) used in scientific and technical calculations. It also boasts an impressive theoretical performance of 952 PFlops (16-bit half-precision), particularly relevant in artificial intelligence applications. The system consists of 240 interconnected compute nodes linked by a 200 Gbps x 4 port high-speed network. Each compute node is equipped with two AMD EPYC 9654 processors, four NVIDIA H100 Tensor Core GPUs, 768 GB of main memory, and 1.92 TB of local SSD storage. Collectively, these nodes share 44.2 PB of HDD storage and 327 TB of SSD storage.

Expanding on the TSUBAME concept of a "Supercomputer for Everyone", TSUBAME 4.0 allows for innovative use cases, including access through web applications. 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, using its GPU logical partitioning and Linux resource partitioning contributes to improved performance and usability compared to its predecessor.

From April 2024, TSUBAME 4.0 computing resources will be provided to HPCI initiative. We are also accepting industrial and apprentice users to broaden the scope of supercomputer users. We also support challenging projects, such as through our Budding Researchers Support program, which encourages young and female researchers to use TSUBAME, and through our Grand Challenge Large-Scale Computing program, which encourages research results that can only be obtained by using all of TSUBAME’s nodes at once, thereby broadening the range of supercomputer users.

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

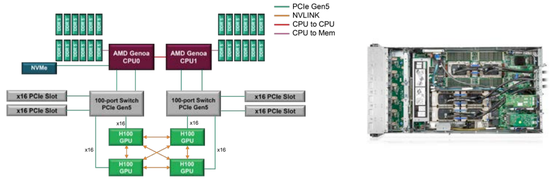

HPE Cray XD665 Server × 240	
CPU	AMD EPYC 9654 (96 cores, 2.4GHz) × 2 Socket
GPU	NVIDIA H100 SXM5 × 4
	FP64 33.5TFlops, FP64 Tensor 66.9TFlops, FP32 66.9TFlops
	TF32 Tensor 494.7TFlops, FP16/BF16 Tensor 989.4TFlops, INT8 Tensor 1978.9Tops
	Memory 94GB HBM2e 2395.87GB/s
Memory	768GiB (DDR5-4800)
Local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

TSUBAME4.0 compute node configuration

Rendering of the TSUBAME 4.0 supercomputer

New data center on the Tokyo Tech Suzukakedai Campus exclusively for TSUBAME 4.0

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/

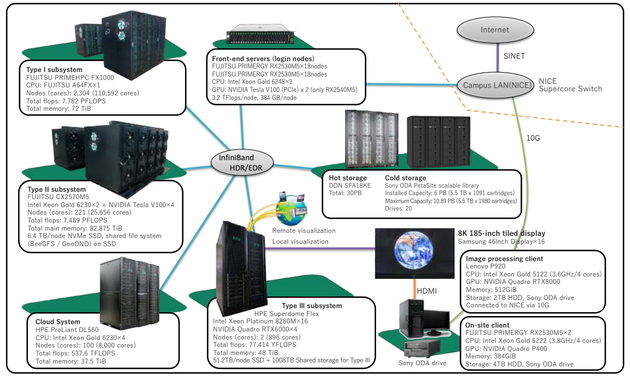

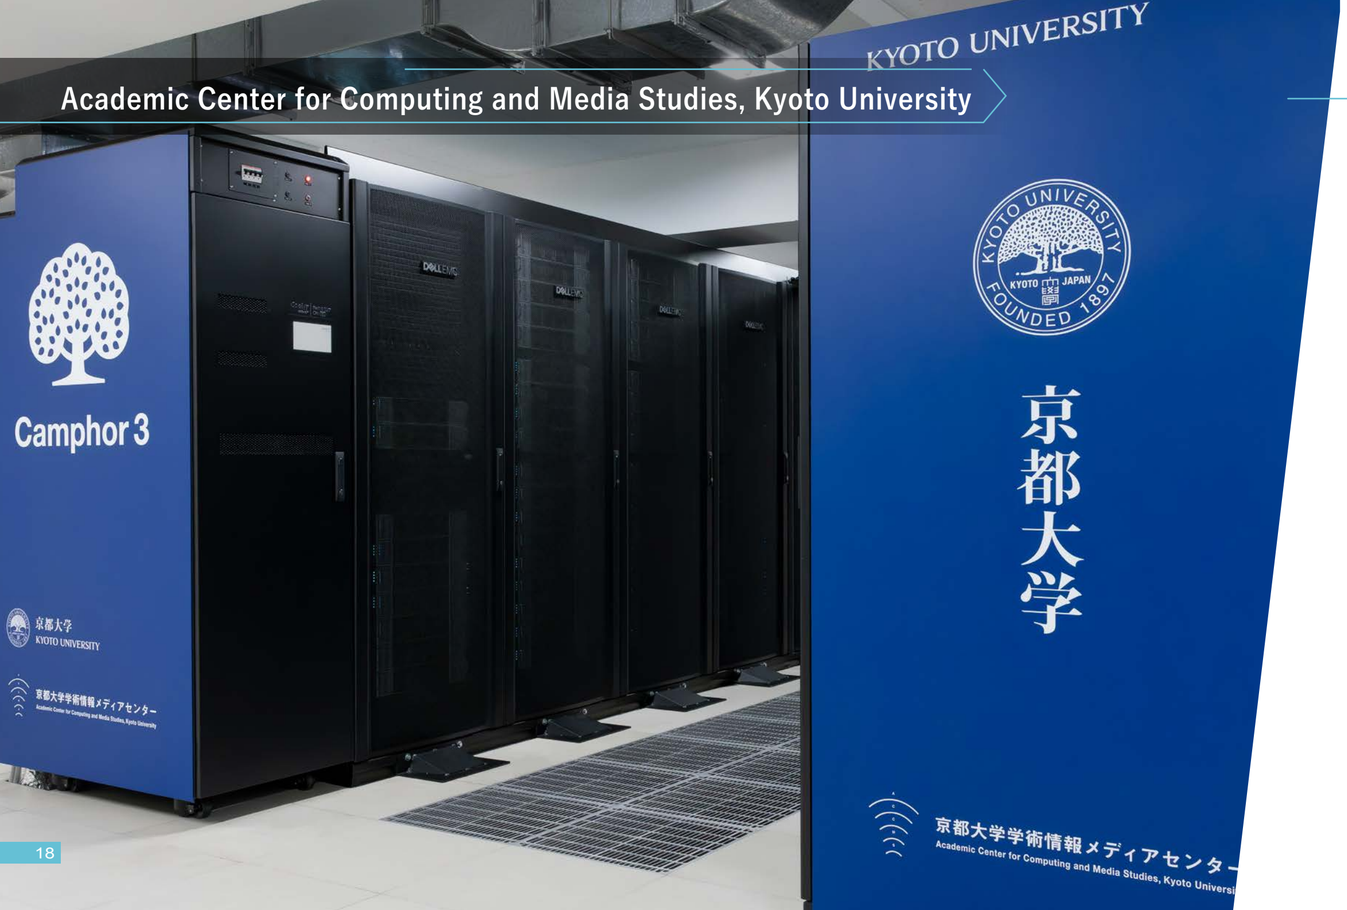

Academic Center for Computing and Media Studies, Kyoto University

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/

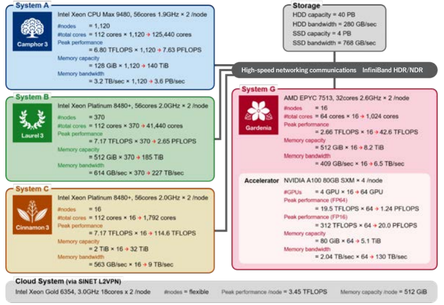

Cybermedia Center, Osaka University

Supercomputer SQUID

Susumu Date

Osaka University's Cybermedia Center offers the use of supercomputer system "SQUID" for cloud-linked high-performance computing and data analytics. This system was first 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 supercomputer. 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

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.

This system will cease operations at the end of February 2024. The next system going forward is planned to begin service from July 2024.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

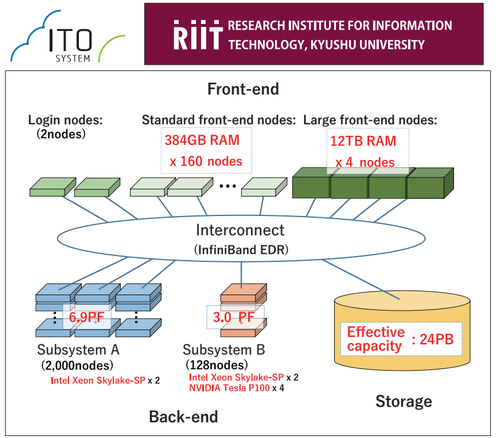

©JAMSTEC

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, 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/

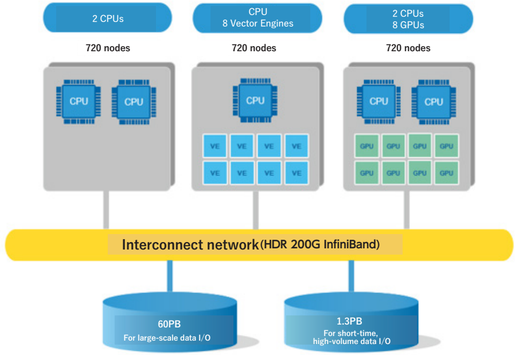

Information Technology Research Institute, National Institute of Advanced Industrial Science and Technology

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/

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.

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.

Until now, supercomputers at the RIKEN Information R&D and Strategy Headquarters have been mainly used to drive research work by researchers and engineers within the institution. But, it is getting difficult for a single institution to provide and maintain the increasingly varied computing resources needed for ever more diverse and specialized types of research. To address this, Japan's collaborative High-Performance Computing Infrastructure (HPCI) initiative provides researchers with access to Fugaku and other varied supercomputing resources. In 2024, we plan to make a portion of HOKUSAI BigWaterfall2 availale as an HPCI resource. This will make it easier for researchers at RIKEN to learn about HPCI resources and use them as appropriate.

For more details, please refer to the RIKEN Information R&D and Strategy Headquarters website below.
https://i.riken.jp/en/supercom/

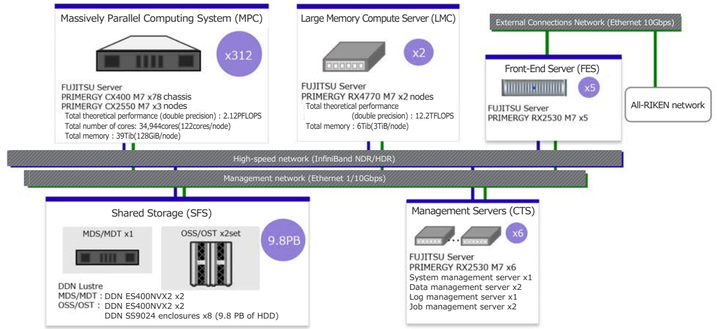

If you are considering using HPCI computing resources or would like to learn more about Japan' s HPCI initiative, please refer to the HPCI portal site below, or contact our help desk.

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

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 2024 English version

Research Organization for Information Science and Technology
Kobe Center

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