


High-Performance Computing Systems at the D3 Center

Overview of High-Performance Computing Environment at the D3 Center



High-performance computing systems (SQUID, OCTOPUS) and data aggregation infrastructure (ONION) are deployed on the D3C-Supercomputer network, also known as D3C-SCinet, a low-latency, high-bandwidth network. This architecture enables users to efficiently access high-capacity storage resources as well as perform large-scale parallel computation and analysis on our high-performance computing systems.

High-Performance Computing Systems

SQUID



SQUID is the acronym for “Supercomputer for Quest to Unsolved Interdisciplinary Datascience.” SQUID is a cluster system being operated since May 2021. This system is composed of general-purpose CPU nodes, GPU nodes, and Vector nodes, total 1,598 nodes. These nodes and large-scale storage EXAScaler (Lustre, 21.2 PB) are interconnected on InfiniBand HDR (200 Gbps). The total peak performance is 16,591 PFLOPS.

Data Sheet of SQUID

Type of nodes	General purpose CPU	GPU	Vector
CPU	Intel Xeon Platinum 8368 (Ice Lake / 2.4 GHz, 38 cores) x 2	AMD EPYC 7402P (Rome / 2.8 GHz, 24 cores)	
Accelerator	N/A	NVIDIA HGX A100 8-GPU	NEC SX-Aurora TSUBASA Type 20A x 8
OS	Rocky Linux 8.10		
# of nodes (total)	1,520	42	36
# of cores (total)	115,520	3,192	864
Memory capacity (total)	389 TB	22 TB	5 TB
Peak performance	8,871 PFLOPS	6,797 PFLOPS	0.922 PFLOPS

OCTOPUS



OCTOPUS is the acronym for “Osaka university Compute & sTOrage Platform Urging open Science.” OCTOPUS is a next-generation computing and storage infrastructure being operated since September 2025. This system consists of 140 general-purpose CPU nodes equipped with two 6th generation Intel Xeon Scalable processors (Granite Rapids) and large-capacity storage EXAScaler (Lustre, 3.58PB). The nodes and storage are interconnected on InfiniBand NDR (200 Gbps).


Data Sheet of OCTOPUS

CPU	Intel Xeon 6980P (Granite Rapids / 2.0GHz 128 cores) x 2
OS	Rocky Linux 9.4
# of nodes (total)	140
# of cores (total)	35,840
Memory capacity (total)	107.52 TB
Peak performance	2.293 PFLOPS

Data Aggregation Infrastructure

ONION

ONION is the acronym for “Osaka university Next-generation Infrastructure for Open research and open innovaTION.”



This is a data aggregation infrastructure that is linked to SQUID. ONION consists of ONION-object (AWS S3 compatible object storage), ONION-file (file hosting service using Nextcloud), and EXAScaler (a parallel file system based on Lustre).

Data Sheet: EXAScaler (on SQUID)

Effective capacity (HDD)	20 PB
Effective capacity (NVMe)	1.2 PB
Max number of inodes	Approx. 8.8 Billion
Max expected effective throughput (HDD)	Over 160 GB/s
Max expected effective throughput (NVMe)	Write : Over 160 GB/s Read : Over 180 GB/s

Data Sheet: ONION-object

Effective capacity	950 TiB
Data protection method	Erasure Coding (Data chunk:4 + Parity chunk:2)

High-Performance Cloud Infrastructure

mdx II

mdx II is a system composed of compute nodes and storage built on OpenStack. Each resource is provided through virtual machines, delivering an isolated environment for each project. This system can be used for a broader range of flexible purposes, including hosting data repositories and data platforms.

*mdx I and II are jointly operated by 9 universities and 2 research institutes in Japan. *



For more information: <https://mdx.jp/en/mdx2>



Data Sheet: mdx II

Type of nodes	CPU	GPU
CPU	Intel Xeon Platinum 8480+ (Sagehill Rapids / 2.0 GHz, 56 cores) x 2	Intel Xeon Gold 6530 (Emerald Rapids / 2.1 GHz, 32 cores) x 2
Accelerator	N/A	NVIDIA HGX H200 4-GPU
# of nodes (total)	60	7
# of cores (total)	6,720	448
Memory capacity (total)	30 TB	7,168 TB
Peak performance	430.08 TFLOPS	982.1 TFLOPS