


Large-scale Computing Systems at the Cybermedia Center

Overview of High-Performance Computing Environment at the CMC



Large-scale computing systems (OCTOPUS and SQUID) and data aggregation infrastructure (ONION) are deployed on CMC-Supercomputer network, a.k.a CMC-SCinet, a low-latency and wide-bandwidth network. This architectural design allows users to access to large-scale storage systems, perform large-scale high-performance computation and analysis on our large-scale computing systems.

Large-scale Computing Systems

OCTOPUS



OCTOPUS is short for **O**saka university **Cybermedia cenTer** **O**ver-Petascale **U**niversal **S**upercomputer. OCTOPUS is a cluster system being operated since December 2017. This system is composed of General purpose CPU nodes, GPU nodes, Large-scale shared-memory nodes, and Xeon Phi nodes, total 319 nodes. These nodes and large-scale storage EXAScaler (Lustre 3.1 PB) are interconnected on InfiniBand EDR (100 Gbps) and form a cluster.

Table 1 Data Sheet of OCTOPUS

Type of nodes	General purpose CPU	GPU	Large-scale shared-memory	Xeon Phi
CPU	Intel Xeon Skylake (2.6 GHz, 12 cores) x 2	Intel Xeon Skylake (2.0 GHz, 16 cores) x 8		Intel Xeon Phi KNL (1.3 GHz, 64 cores)
OS	RHEL 7.6			
# of nodes (total)	236	37	2	44
# of cores (total)	5,664	888	256	2,816
# of memories (total)	45 TB	7 TB	12 TB	8 TB
Peak performance	471.2 TFLOPS	858.3 TFLOPS	16.4 TFLOPS	117.1 TFLOPS
Accelerator	NVIDIA Tesla P100 (NVLink) x 148			

SQUID




SQUID is short for **S**upercomputer for **Q**uest to **U**nsolved **I**nterdisciplinary **D**atascience. 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) and form a cluster.

Table 2 Data Sheet of SQUID

Type of nodes	General purpose CPU	GPU	Vector
CPU	Intel Xeon IceLake (2.4 GHz, 38 cores) x 2		AMD EPYC Rome (2.8 GHz, 24 cores)
OS	Rocky 8.6		
# of nodes (total)	1,520	42	36
# of cores (total)	115,520	3,192	864
# of memories (total)	389 TB	22 TB	5 TB
Peak performance	8,871 PFLOPS	6,797 PFLOPS	0.922 PFLOPS
Accelerator	NVIDIA HGX A100 8-GPU x 42 NEC SX-Aurora TSUBASA Type20A x 288		

Data Aggregation Infrastructure

ONION



ONION stands for **O**saka university **N**ext-generation **I**nfrastructure for **O**pen research and open innovation. ONION is a new data aggregation infrastructure that is linked to SQUID. ONION consists of ONION-object (AWS S3 compatible object storage), ONION-file (storage service using Nextcloud), and EXAScaler (a parallel file system based on Lustre).

ONION makes it easy for users to transfer data between your PC and large-scale computing system. In addition, ONION can be used in a variety of ways, such as immediate sharing of calculation results with those who do not have a SQUID or OCTOPUS account and manipulating data from a smartphone. Of course, it can also be used to store and share research data in the laboratory.

Table 3 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

Table 4 ONION-object

Effective capacity	950 TiB * We plan to expand sequentially
Data protection method	Erasure Coding (Data chunk:4 + Parity chunk:2)