

Case Study



High Performance Cloud Computing
3rd Generation Intel® Xeon® Scalable Processors
Intel® DL Boost

Osaka University CMC SQUID Supercomputer Enables Large-scale Interdisciplinary Research

A heterogeneous architecture with CPUs, GPUs, vector processors, and a data aggregation and sharing system, SQUID's 16+ petaFLOPS opens new doors to national and global scientific collaboration.

Solution Summary:

- 1,520 nodes with Intel Xeon Platinum 8368 processors
- 42 nodes with Intel Xeon Platinum 8368 processors and eight GPUs per node
- 36 nodes with NEC SX-Aurora TSUBASA Type20A vector accelerators
- DDN EXAScaler storage appliances

Executive Summary

[Osaka University](#) is a national university corporation in Japan. It supports researchers in academia and industry throughout the country. Osaka University [Cybermedia Center](#) (CMC) provides supercomputing resources for a wide range of science, from physical to life sciences and more. In 2017, The CMC deployed OCTOPUS, 1.463 petaFLOPS, world-class, heterogeneous cluster with 1st Gen Intel® Xeon® Scalable processors targeted at scientific computing for a variety of workloads using different architectures. OCTOPUS has allowed new levels of discovery. To continue the university's leading position in scientific research, CMC deployed SQUID in 2021. The new cluster, built by NEC with 3rd Gen Intel® Xeon® Scalable processors, is over 11 times faster than CMC's previous system with a peak performance of over 16 petaFLOPS.¹ It will allow Osaka University to support new initiatives and interdisciplinary research across the sciences, using shared data and expanded capacity and capabilities.

Challenge

Prior to 2017, Osaka University CMC resources were used by or both general purpose and scientific computing. OCTOPUS was designed for computational science alone in traditional simulation and modeling and emerging work at the time in Artificial Intelligence (AI) and machine learning (ML). Its heterogeneous



SQUID is over 11 times faster than CMC's previous system with a peak performance of over 16 petaFLOPS

Case Study | Osaka University CMC SQUID Supercomputer Enables Large-scale Interdisciplinary Research

architecture included Intel® Xeon® Scalable processors and NVIDIA GPUs. After deployment, its usage increased rapidly. It has been a key resource in supporting the new achievements of Osaka University researchers and students.

“OCTOPUS is still an important component in researchers’ tools,” Dr. Susumu Date, Associate Professor, Osaka University CMC stated. “But, by 2021, it saw 90 percent average utilization, with many users waiting days in the queue for their work to start.”

Today, as a national university corporation, Osaka University and the CMC support a greater array of researchers in academia and industry from around the country, plus students working on projects. Additionally, AI, Internet of Things (IoT), High Performance Data Analytics (HPDA), and use of shared data are increasingly crucial to expanding understanding and breakthroughs in science. To support expanding research, the CMC will enable innovation and greater interdisciplinary work across the sciences through shared data in a safe and responsible manner. This capability plus the need for more capacity and performance and user scalability led to the design and deployment in 2021 of SQUID (Supercomputer for Quest to Unsolved Interdisciplinary Datascience).

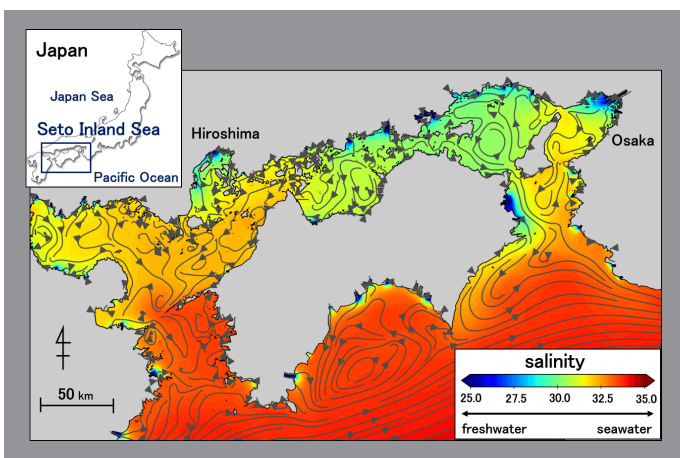
Solution

SQUID was designed to explore unsolved data science problems using the latest techniques and methods of computational science. For such a vision, SQUID, like OCTOPUS, required multiple compute architectures.

“Some users will use different types of compute nodes in a combinational way,” Date added. “Others will compare them. SQUID, like OCTOPUS, was designed with heterogeneity to accommodate the needs of users.”

Built by NEC, SQUID comprises three different groups of compute nodes, totaling 1598 servers:

- 1520 general-purpose HPC compute nodes, each with dual-socket Intel® Xeon® Platinum 8368 processors with Intel® Deep Learning Boost (Intel® DL Boost) for AI inference acceleration



This 3-D high-resolution hydrodynamic simulation of coastal waters (spatial distribution of currents and salinity) represents one of the projects that utilize the Osaka University's supercomputer. Image courtesy Associate Professor Yusuke Nakatani, Osaka University)

- 42 GPU nodes with dual-socket Intel Xeon Platinum 8368 processors, each with eight NVIDIA A100 accelerators
- 36 vector nodes, each with NEC SX-Aurora TSUBASA Type 20A accelerators with high-bandwidth memory

But additionally, SQUID needed much larger data capacity and management capabilities, extreme security, many more petaFLOPS, and the ability to easily support more users.

Five Key Challenges Addressed

“Five challenges were explored in deploying SQUID: HPC and HPDA integration, cloud bursting, a secure computing environment, tailor-made computing, and data aggregation,” explained Date. “SQUID was designed around these five criteria.”

HPDA integration: Users today have opportunity to use many types of computation for different purposes and use them in different ways, whether for simulation or analytics. HPDA has emerged as an important tool for revealing insights in research, so it was important to integrate both traditional HPC and HPDA in the design of SQUID, according to Date.

Cloud bursting: Even with a much larger cluster, Osaka University CMC needed the ability to scale quickly to meet the needs of various users in order to avoid long wait times on the system as the user base grows. The answer was to build in the ability to burst some workloads to the cloud when needed. Users have the choice to run either only on SQUID or to burst to the cloud if necessary. A sophisticated NEC job scheduler can push jobs to either Oracle Cloud Infrastructure or Microsoft Azure cloud to accommodate the needs of users.

Secure computing: Users have access to more data in a very secure environment, achieved through collaborative development between NEC and Osaka University CMC. The environment provides dynamic partitioning to isolate compute and networking for a particular group to protect data and computation. Additionally, an experimental program is exploring how to use sensitive, confidential data in on-premises repositories without moving data from storage.

Tailor-made computing: Osaka University CMC supports Singularity containers to allow users to create and run their projects in a tailor-made workspace. Users can build their projects on a local desktop or laptop and transport the container file to SQUID where it will run using the needed resources.

Data-aggregation: Modern global research largely shares data generated by supercomputing systems. Data generated by one project can be important to another endeavor. Thus, SQUID was designed with the ability to aggregate and share data between researchers across the planet.

“We designed a data aggregation infrastructure named ONION (Osaka university Next-generation Infrastructure for Open research and open innovatiON),” Date added. “It allows researchers to share computing results with other researchers immediately after the computation completes through a smartphone or local computing environment.”

ONION works in conjunction with the Cloudian* Object Storage HyperStore platform and accommodates a variety of data access protocols to improve storage flexibility. For

example, S3-compatible IoT devices can be configured to aggregate their data onto the SQUID parallel file system, so users can use that data in simulations.

The data aggregation infrastructure is built on Data Direct Networks (DDN) EXAScaler appliances, providing 20 petabytes of hard disk storage and 1.2 petabytes of fast, NVMe storage in a parallel file system.

Designed with these capabilities, SQUID now allows researchers across many fields to run their jobs and share their data using one of the fastest clusters in the country.

Result

In addition to supporting university researchers, CMC provides SQUID resources to national research projects through two programs. These projects are approved by the [High Performance Computing Infrastructure office of Japan](#) and the [Joint Usage/Research Center for Interdisciplinary Large-scale Information Infrastructures](#).

“Last year, 17 projects accepted by these two institutions were run on Osaka University CMC resources,” Date said. “Eleven of those were completed on SQUID. These were related to quantum chromodynamics (QCD), molecular dynamics, Covid-19, astrophysics, and others.”

One of the research groups in the university is using SQUID to explore the mixture of queues, according to Date. “The group uses different compute nodes in a combinational way to take advantage of different characteristics of the processors and accelerators in the nodes. The effort is examining how to use heterogeneous compute nodes more effectively,” he concluded.

Addressing the five challenges considered in designing SQUID was in response to how research has turned more global. Scientists work in greater collaboration to achieve new insights and breakthrough discoveries. SQUID supports the global research community with greater capacity and a data aggregation/sharing infrastructure.

Solution Summary

Osaka University CMC needed to augment the resources of OCTOPUS, deployed in 2017, with higher performance, greater capacity, and the ability to meet the needs of a growing research community. NEC built the heterogeneous architecture cluster with 3rd Generation Intel Xeon Scalable processors, GPUs, and vector accelerators, achieving over 16 petaFLOPS. Its data aggregation infrastructure, built on DDN EXAScaler appliances and Cloudian Object Storage HyperStore platform allows scientists to run their calculations and share data immediately with others around the world. SQUID is a key resource for researchers in academia and industry in Japan, enabling discovery and insight across multiple scientific disciplines.

Where to Get More Information

Learn more about Osaka University CMC at <http://www.hpc.cmc.osaka-u.ac.jp/en/>.

Explore the capabilities of the [3rd Generation Intel Xeon Scalable processors](#) with integrated Intel® Deep Learning Boost for accelerated AI inferencing.

Solution Ingredients

- 1,520 nodes with Intel Xeon Platinum 8368 processors
- 42 nodes with Intel Xeon Platinum 8368 processors and eight GPUs per node
- 36 nodes with NEC SX-Aurora TSUBASA Type20A vector accelerators
- DDN EXAScaler storage appliances



¹ Comparing SQUID specifications at <http://www.hpc.cmc.osaka-u.ac.jp/en/squid/> to Octopus specifications at <http://www.hpc.cmc.osaka-u.ac.jp/en/octopus/>

Performance varies by use, configuration, and other factors. Learn more at www.Intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

For workloads and configurations visit www.Intel.com/PerformanceIndex. Results may vary.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Your costs and results may vary.

Intel technologies may require enabled hardware, software, or service activation.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

092022/RJM/J/PA/PDF ♻️ Please Recycle 352527-001US