

Case Study

High Performance Computing
Intel® Xeon® CPU Max Series



Stony Brook University's Seawulf is First U.S. Academic Supercomputer with Intel® Xeon® CPU Max Series

Learning from their Ookami supercomputer, Stony Brook's Institute for Advanced Computational Science chose Intel Xeon CPU Max Series with High-bandwidth Memory (HBM) to drive performance up to 9x compared to previous generation Intel CPU

Solution Summary

- First academic supercomputer in U.S. with Intel® Xeon® CPU Max Series
- Built on HPE ProLiant DL360 Gen11 servers
- 94 nodes with Intel Xeon Max CPU 9468 (48 cores each/9,024 cores)
- 128GB HBM2 per node

Executive Summary

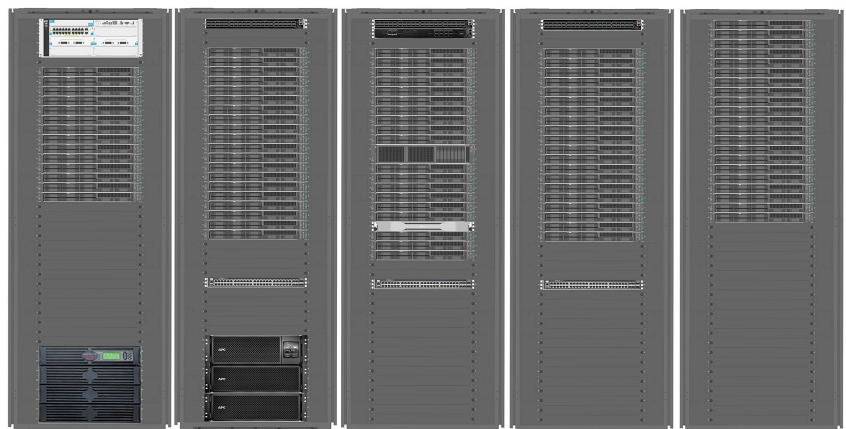
Stony Brook University recently upgraded its Seawulf supercomputer, a heterogeneous system comprising CPUs and GPUs, delivering up to 1.86 petaFLOPS. The most recent addition was built on Intel® Xeon® CPU Max Series with High-bandwidth Memory (HBM). Stony Brook researchers and computational scientists have experience with HBM through the university's Ookami cluster, built on the Fujitsu A64FX processor with HBM.

The Intel Xeon CPU Max Series with HBM accelerates execution of memory bandwidth bound applications, while the new Intel Xeon processor architecture increases overall performance compared to previous generations of Intel® Xeon® Scalable processors.

Stony Brook researchers benchmarked the new Seawulf partition, which are presented below.

Challenge

[Stony Brook University's Seawulf](#) cluster is a heterogeneous platform that has greatly evolved over the last 10 years. Seawulf is a mix of multi-generational x86 CPUs and GPUs interconnected by multi-generational InfiniBand fabrics and a GPFS storage array with both spinning disks and SSDs.



Stony Brook's Institute for Advanced Computational Science chose Intel® Xeon® CPU Max Series with High-bandwidth Memory (HBM) for their recent addition to their Seawulf cluster.

Stony Brook is unique among universities, in that it is one of a very few institutions² that also runs a Fugaku-like supercomputer—named [Ookami](#)—built on the Fujitsu A64FX CPU with high-bandwidth memory (HBM). Fugaku, at the Riken Center For Computational Science in Japan, has ranked in the top four of the Top500.org's list for the last four years. It held first place in 2020 and 2021. Its performance benefited from HBM technology. Students of Stony Brook's [Institute for Advanced Computational Science \(IACS\)](#) have been learning from Ookami, running a variety of scientific workloads since its deployment.

To continue to provide researchers and students of the IACS with advanced computing resources—including for machine learning (ML) and deep learning (DL)—the IACS and [Information Technology](#) departments needed to expand the cluster. In late 2022, having learned much from Ookami while also building on its experience with x86 architecture, the university began its search for the technology that would power Seawulf's new partition—HBM would be an important part of it.

Solution

Stony Brook researchers run a variety of scientific simulations on their clusters (both Seawulf and Ookami), including astrophysics, ocean modeling, molecular dynamics, and others. Some of these are memory bandwidth bound.

“The limited memory bandwidth of mainstream processor technology meant it was very hard to get good thread scaling and utilize all the cores on modern systems,” said Robert Harrison, Director of the IACS at Stony Brook. “Our very positive experiences on the National Science Foundation-sponsored Ookami cluster led us to look at the Intel Xeon CPU Max Series and believe it could deliver on performance with HBM.”

Robert and his team engaged with Intel early in the system architecture design and selection process for information needed to determine that the Intel® Xeon® CPU Max Series was the technology and architecture to move forward with. After a typical tender, in 2023, Stony Brook chose HPE to expand Seawulf.

In addition to leveraging technologies from HPE and Intel, the new HPC and AI solution was deployed by ComnetCo, HPE's HPC-focused solution provider and award-winning public sector partner.

Up to 9X Speedup

with Intel® Xeon® Max over ARM CPU⁴

The new partition comprises 94 HPE ProLiant DL360 Gen11 compute nodes, each with two Intel® Xeon® Max 9468 processors (48 cores each/9,024 cores total), 256GB of DDR5 memory, 128GB of HBM, and an InfiniBand NDR (400 Gbps) fabric. The new partition added 0.78 petaFLOPS to the cluster, bringing it to its 1.86 petaFLOPS rating.³

The system became production-ready in November 2023, after several benchmarks and application performance

tests were completed. The testing showed significant performance improvements provided by the Intel Xeon CPU Max series with HBM (described below).

Result

A group of researchers and computational scientists with the IACS ran a battery of benchmarks to evaluate the performance impact from HBM. They also ran deep learning problems and real-world applications to see how they might benefit from HBM and the new CPU architecture. Scientific applications included Gromacs molecular dynamics and OpenFOAM computational fluid dynamics. The study was documented in a reviewed research paper.¹

Additionally, the new Seawulf cluster has been the target for running astrophysics and ocean modeling workloads by Stony Brook researchers.

Testing Platforms

The testing compared the Intel Xeon CPU Max Series nodes against Seawulf nodes built on Intel® Xeon® Gold 6184 processor and AMD EPYC 7643 CPU (codenamed Milan), and on Ookami's Fujitsu A64FX-FX700 processor.¹ The testing included all these processors where appropriate; some processors do not support operations and capabilities of the Intel Xeon CPU Max Series, so those processors were not included in certain tests.

For details about the testing platforms, protocol, benchmarks, and applications, [see the paper](#).¹ The following charts and summary performance results reference that study. Additionally, the study looked at other factors beyond performance, including performance per core and power consumption (not included here).

Benchmarks

Benchmark results are shown in Figure 1 (for HBM) and Figure 2 (for DDR). The chart shows the speedup performance improvement of Intel Xeon Max 9468 processor compared to other CPUs for each of the benchmarks. Memory bandwidth tests included sustained bandwidth (STREAM TRIAD) and scaling and thread distribution (results not shown). The scaling and distribution




Figure 1. HBM benchmark results for Seawulf and Ookami clusters. Test results conducted by Stony Brook.




Figure 2. DDR benchmark results for Seawulf cluster. Test results conducted by Stony Brook.

benchmark examined memory bandwidth across threads as operations were scaled to more threads.

Sustained memory bandwidth was 3.5x higher with HBM versus just DDR on the new Intel Xeon CPU Max Series. According to the study, the memory copy benchmark revealed HBM improves memory bandwidth utilization nearly linearly up to full subscription on a node, while with DDR only, bandwidth saturates quickly.

It's worth noting that, while all benchmarks show improvement with HBM versus DDR on the new Intel CPU, the degree of benefit varies across tests. This is expected, since some codes are not memory bandwidth bound, as was also seen in application testing.

oneDNN Benchmark

For deep learning training and inference applications, functions of the Intel® oneAPI oneDNN library were evaluated.

“While other benchmarks looked at the memory bandwidth, the oneDNN benchmark was of particular interest to me,” commented Smeet Chheda, a PhD candidate in the IACS, “because of the tile multiplication unit, which was new on the system. It is helpful in accelerating deep learning applications. I benchmarked the convolution operator because it’s an expensive and widely used operation. And it has matrix multiplication as well.”

Different inputs with different configurations were tested, including different data types. Data types included fp32, BF16, and INT8, which are types used in various training and inference use cases. For fp32 data types, Intel® Xeon® CPU Max Series with HBM ran up to 1.7 to 3.5 times faster than the previous version Intel processor and the A64FX CPU, respectively. For BF16 data types, Intel Xeon CPU Max Series with HBM was 1.79 to 9.2 times faster than the fp32 benchmarks. For INT8 data types, the new Intel processor delivered from 16.1x to 17.5x speedup over the previous version Intel processor and A64FX CPU, respectively.

“Generally, for the Intel Xeon CPU Max series, I saw a 2x benefit with HBM versus DDR on the new CPU,” concluded Chheda.




Figure 3. Performance speedup of Intel Xeon Max 9468 processor with HBM running Gromacs. Test results conducted by Stony Brook.




Figure 4. Performance speedup of Intel Xeon Max 9468 processor with DDR running Gromacs. Test results conducted by Stony Brook.

Gromacs

For Gromacs, several sized problems were run. The results show that HBM does not significantly benefit performance versus DDR on the new Intel processor. However, it does significantly outperform other CPUs, both when using HBM (Figure 3) and DDR (Figure 4).

Exploding Stars

PhD candidates Catherine Feldman and Josh Martin are working on optimizing the astrophysics suite called Flash to run it on the new Seawulf partition. Flash is a Multiphysics code that simulates the explosion of a star.

“We’re studying Type Ia supernovae, which is a stellar explosion that’s ten billion times brighter than our sun,” explained Feldman. “Observational astronomers use the brightness of these explosions to determine distances across the universe. The better we understand how these explosion happen, the better we can calibrate our measuring tape, so to speak, and make a better map of the universe. So, we use Flash to simulate an explosion and thereby better understand how they happen.”

Flash has been around for 20 years, and Alan Calder, professor of astronomy and physics and deputy director of

the IACS, is a contributor to part of the code. Until now, with current computing resources at Stony Brook, researchers have been limited to 2D simulations of these explosions. But Calder expects with the new partition of Seawulf, they will be able to perform 3D simulations of explosions.

Feldman and Martin have been working with the code to run different sized 2D and 3D simulations on the Intel Xeon CPU Max Series. Their testing is currently ongoing and will be revealed in an upcoming paper.

Solution Summary

Stony Brook University's Seawulf cluster was due for an upgrade. With experience from their Fugaku-like Ookami cluster, they understood the benefits of HBM for memory bandwidth bound workloads. Thus, they looked at the new Intel Xeon CPU Max Series as the foundation for a new set of nodes on Seawulf.

HPE built the new partition with 94 nodes of HPE ProLiant DL360 Gen11 servers, each hosting two Intel Xeon Max 9468 processors with HBM. The system was put into production in 2023, and several researchers and computational scientists have been running it through its paces, resulting in one published research paper on performance of the new system compared to other IACS computational resources and another on the way revealing results of the astrophysics code, Flash, on the new system.

Overall, the new partition delivers as much as 3.5x sustained memory bandwidth and better thread scaling running with HBM versus DDR. With its new tile multiplication unit, the new Intel CPU delivers better performance on deep learning with oneDNN across multiple data types used in training and inference. And, while all workloads do not benefit from HBM—irrespective of the computing resource it runs on—the new nodes with Intel Xeon Max 9368 processor outperforms the other CPUs on nearly all workloads tested.

Where to Get More Information

Learn more about [Stony Brook University](#).

Find out more about the [Seawulf cluster](#) with Intel Xeon CPU Max Series.

Learn more about [Stony Brook University's Institute for Advanced Computational Science](#).

Explore the capabilities of the [Intel Xeon CPU Max Series](#) with integrated high-bandwidth memory (HBM).



¹ Siegmann, Eva, et al, [First Impressions of the Sapphire Rapids Processor with HBM for Scientific Workloads](#).

² [Brookhaven National Lab, Stony Brook University, and University of Buffalo have A64FX-based clusters](#).

³ <https://it.stonybrook.edu/help/kb/understanding-seawulf>: Cores per node x # nodes x base clock per core x Double Precision FLOPS per cycle + Nvidia GPUs Teraflops

⁴ See Figure 3.

Notices & Disclaimers

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.