

Institute for Advanced Computational Science (IACS) at Stony Brook University

The Institute for Advanced Computational Science at Stony Brook University was established in 2012 with a \$20M private endowment. IACS Core faculty expertise spans chemistry, materials, astrophysics, atmospheric science, linguistics, ecology, sociology, applied mathematics, computer science, artificial intelligence and the arts. The Institute seeks to make sustained advances in the fundamental techniques of data, computation and in high-impact applications.

Our state-of-the-art facility is located next to the Laufer Center for Physical and Quantitative Biology. This space houses the Institute's faculty, about 45 students/postdocs, and technical/administrative staff, and includes conference space and small meeting rooms to accommodate in-person and hybrid/virtual meetings. There is 10 Gigabit networking throughout the building (optical fiber within the building to major hubs across campus). The Institute provides to its members office space, office materials and supplies, administrative support, commonly used software packages including select compilers and libraries, technical assistance with computer administration and software installation, and use of the Institute's computing clusters.

Networking at SBU is provided by a redundant 100 Gigabit connection to Internet2, connecting to ESNET, as well as Amazon, Microsoft, & Google, at 100 Gigabits/sec at 32 Avenue of the Americas in New York City. The Campus Data Center and the campus network operations center are connected via single-mode fiber, with 2 links providing redundant connectivity at not less than 40 Gbit/s with 100 Gbit/s in the campus network plan. Cybersecurity is provided through a layered approach, including a high-availability pair of firewalls that can support up to 40 Gigabits/sec.

Stony Brook University Research Computing Clusters:

SeaWulf is available for use by all Stony Brook University faculty and students. This cluster includes **405** nodes and **23,372 cores**, with a peak performance of **~1.8 PFLOPS** available for research computation. Available to all nodes is a 3 Petabyte shared storage array running the GPFS file system. Recent upgrades include:

2023: "SeaWulf II" with support from NSF Major Research Instrumentation award 2215987, matching funds from the New York State Department of Economic Development contract C210148, plus additional funding from SBU's President, Provost, Vice President for Research, CIO, Professor Dilip Gersappe and the Deans of CAS, CEAS, and SoMAS. This latest addition, a set of servers from HPE, is composed of 94 compute nodes and a total of 9024 cores, offering 606 TFLOPS. Each node uses Intel Xeon Max 9468 processors, with 96 cores across both sockets, running at a base clock speed of 2.1 GHz, and 256GB of DDR5-4800 RAM. What is notable about these nodes is the presence of 128 GB of HBM2e RAM offering 2 TB/s of aggregate memory bandwidth. These servers are interconnected via an NDR400 InfiniBand fabric at a cutting-edge 400 Gbps. The login node uses dual Intel Xeon Platinum 8468 Processors, with a total of 96 cores running at 2.1 Ghz and 512GB of DDR5-4800 of memory.

2022: IACS expanded the SeaWulf cluster thanks to generous funding by Professors Dilip Gersappe, Benjamin Levine and Robert Rizzo. The first project in this expansion was 48 nodes, totaling 4608 cores with a peak performance of 339 TFLOPS. Each of those compute nodes have AMD EYPC 7643 CPU's with a total of 96 cores with a base clock speed of 2.3 GHz and 256 GB of DDR4-3200 of RAM plus an HDR100 InfiniBand adapter providing 100 Gbps of connectivity. There are also 2 login nodes providing high availability access, configured identically to the compute nodes but with 512GB each instead of 256GB. The second project in this expansion is a GPU-focused set of servers, with eleven nodes and an aggregate of 352 cores and 44 GPUs providing 449 TFLOPS of performance. These GPU-enabled nodes each include 4x Nvidia A100 GPUs with 80GB of video memory alongside two Intel Xeon Gold 6338 Processors (64 cores per node) running at a base speed of 2.3 GHz with 256 GB of DDR4-3200 RAM, and an HDR100 InfiniBand adapter. For memory intensive tasks, the system offers an HPE server containing 3 TB of DDR4-2933 RAM and 4 Intel Xeon Platinum 8360H Processors running at 3 GHz and a total of 96 cores.

2019: SeaWulf was expanded by 64 compute nodes, each with Intel Xeon Gold 6148 CPUs with 40 cores that operate at a base speed of 2.4 GHz and have 256 GB of RAM. This expansion, in aggregate, offers up to 111 TFLOPS and 2,560 cores. In addition, there is a large memory node, with 3 TB of DDR4 RAM and Intel E7-8870v3 processors with 72 cores operating at 2.1 Gigahertz, for a total of 72 cores and 144 threads (via Hyper-Threading) and also a Nvidia V100 16GB GPU.

NVwulf is Stony Brook University's newest high performance computing cluster, designed to accelerate artificial intelligence, machine learning, and data intensive scientific research. Functioning as a sister system to the general purpose SeaWulf cluster, NVwulf represents a major expansion of campus GPU computing resources. NVwulf was developed through collaboration between the Research Computing and Innovation team within the Division of Information Technology, multiple academic departments, and leadership across Stony Brook's East and West campuses. It is part of a larger multiyear initiative supported by a five million dollar New York State grant to establish Stony Brook as a national hub for interdisciplinary AI research.

Launched in June 2025 and expanded in Q1 2026, NVwulf features 32 NVIDIA H200 NVL GPUs, capable of delivering up to 107 petaFLOPS of FP8 performance for AI and machine learning workloads and 1 petaFLOP of FP64 performance for scientific computing. In addition, 36 NVIDIA RTX PRO 6000 Blackwell ("B40") GPUs are present, capable of delivering up to 36 petaFLOPS of FP8 and 73 petaFLOPS of FP4 performance for AI and machine learning workloads.

Each NVIDIA H200 compute node has 4-8 NVLink GPUs with 141 GB VRAM each, and either two AMD EPYC 9355 32-core processors or two Xeon 6 6530P ("Granite Rapids") 32-core processors (64 cores per node), and 768 GB - 1.5 TB of DDR5 memory. Each NVIDIA B40 node is equipped with 4 GPUs with 96 GB VRAM, two Xeon 6 6530P ("Granite Rapids") 32-core processors (64 cores per node), and 512 GB DDR5 memory.