

CASE STUDY

High Performance Computing (HPC)
2nd Generation Intel® Xeon® Scalable Processor
Intel® Omni-Path Architecture



Shanghai Jiao Tong University's π 2.0 Supercomputer Built for Convergence of HPC Applications

2nd Generation Intel® Xeon® Scalable Processors and Intel® Omni-Path Architecture to power traditional HPC, Big Data, and Artificial Intelligence research

π 2.0 System Highlights

- Inspur-built system with 658 nodes of 2nd Generation Intel® Xeon® Gold 6248 processors
- 26,320 total compute cores (52,640 threads)
- Intel® Omni-Path Architecture fabric
- Intel® SSD DC Series for NVMe in Lustre* parallel file system



Executive Summary

[Shanghai Jiao Tong University](#) is one of the most prestigious universities in China. The university's High Performance Computing (HPC) center installed its last supercomputer in 2013. For computational researchers from the organization's 60,000 students and 6,000 faculty, the resource lacks capacity to effectively support ongoing work. The University turned to Inspur* for their next-generation cluster built on 2nd Generation Intel® Xeon® Scalable processors and Intel® Omni-Path Architecture (Intel® OPA).

Challenge

Established in 1896, Shanghai Jiao Tong University is one of the oldest universities in China. Its 28 departments and 15 hospitals educate 30,000 undergraduate and 30,000 graduate students in a wide range of disciplines. Many of those departments require supercomputing resources for discovery and insight in materials, astrophysics, aeronautics, computational genomics, and other traditional sciences. Over the last few years, research has expanded in new areas, including big data and machine learning.

The university's existing HPC resource, named π, was built in 2013. It is a heterogeneous, 260-teraFLOPS cluster built on Intel® Xeon® processors E5 and NVIDIA* GPUs with InfiniBand* Architecture interconnect.

"Things have changed a lot over the last six years," stated Dr. James Lin, Vice Director of the HPC Center. "As research at the university has addressed ever more complex and deeper problems and included new fields in machine learning and big data, more students need computing cycles. The queues on π for researchers' jobs have gotten longer and longer, delaying important research work."

Besides not having enough capacity in π for current work, researchers want to take advantage of more scalable codes that can run their jobs faster on more processor cores. In 2018, the university contacted OEM Inspur to build a new, approximately two-petaFLOPS homogeneous system named π 2.0.

Solution

"We support research that runs commercial applications, open source codes for traditional CFD modeling and other science, and in-house high-scalability codes," explained Stephen Wang, head of Technical Support. "We provide researchers with help in optimizing and porting their codes for scalability on parallel systems."

π 2 will be a 658-node system with two-socket Inspur servers running 2nd Gen Intel® Xeon® Gold 6248 processors with 20 cores each, totaling 26,320 compute cores with around 2 PetaFlops peak performance. It will be the 3rd largest super-

computer among China's universities. The compute nodes will be connected by an Intel® Omni-Path Architecture (Intel® OPA) fabric and supported by a Lustre* scalable, parallel file system using Intel® SSD Data Center series drives for NVMe.

"Most supercomputers on the Top500 list are built on Intel® architecture (IA)," added James. "And we have a lot of experience with Intel architecture, including modernizing codes from GPUs to IA. So, we chose next-generation Intel Xeon Scalable processors for our new cluster."

Selecting Intel OPA for the interconnect was a little more involved.

"We visited the two leading HPC centers in Japan: The Joint Center for Advanced High Performance Computing (JCAHPC) by the University of Tokyo and the University of Tsukuba, and Global Scientific Information and Computing Center (GSIC) of Tokyo Institute of Technology," explained James. "JCAHPC hosts Oakforest-PACS, a very large supercomputer with Intel OPA and the largest Intel OPA deployment in the world. Tsubame3 GSIC is also a very large cluster that uses Intel OPA. We chose Intel OPA based on our research with those centers and visits in China to other Intel OPA customers."

Power was a critical concern for the HPC center. With 26,320 cores, π 2.0 will be seven to eight times larger than π .

"We are required to support a power usage effectiveness (PUE) of 1.3," commented Stephen. "With the more efficient technology of the latest Intel Xeon Scalable processors, π 2.0 power demand will be only between two and three times that of π even though the system is nearly eight times larger."

A key area of concern was the Lustre file system. With π and an increasing number of computational genomics jobs (as many as 1000 at a time), Lustre was becoming a bottleneck because genomics runs many small jobs. π 2.0's Lustre file system will include Intel® SSD Data Center series drives to accelerate IO across the storage cluster.

Results

While the system is still in the construction phase, researchers are well aware of the new capacity and technologies available on π 2.0. Stephen's technical support department fields questions from users about scaling their codes.

"We are focused on developing scalable codes," said Stephen, "offering help with methodologies, such as OpenMP* and MPI. We also have interns who will actually help researchers port their codes. Since we have experience with modernizing GPU codes, we can help them target their machine learning applications for the 2nd Gen Intel® Xeon® processors with features, such as Vector Neural Network Instructions (VNNI)."

The new supercomputer's first customers will be the biggest users at the university, running their inhouse high-scalability codes optimized for 2nd Gen Intel Xeon Scalable processors. Other early projects will include machine learning jobs.

"Users are very excited about getting access to the system," added James.

Solution Summary

- Inspur-built system with 658 nodes of 2nd Gen Intel Xeon Gold 6248 processors
- 26,320 total compute cores (52,640 threads)
- Intel SSD DC Series for NVMe for fast-response Lustre parallel file system
- Intel Omni-Path Architecture for fast communications across compute nodes
- Frameworks for machine learning, including TensorFlow* and Pytorch* using Intel® Optimizations for TensorFlow and Intel® Distribution of Python*


Where to Get More Information

Learn more about Intel Xeon Scalable Processors at <https://www.intel.com/content/www/us/en/products/processors/xeon/scalable.html>.

Find out more about Intel Omni-Path Architecture at <https://www.intel.com/content/www/us/en/high-performance-computing-fabrics/omni-path-architecture-fabric-overview.html>.

Read about Intel SSDs at <https://www.intel.com/content/www/us/en/products/memory-storage/solid-state-drives.html>.

Learn more about HPC at Shanghai Jiao Tong University at <https://hpc.sjtu.edu.cn/>.



The new π 2.0 supercomputer at Shanghai Jiao Tong University will support research for commercial applications, open source codes for traditional CFD modeling and other science, and in-house high-scalability codes.



Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at <https://www.intel.com/content/www/us/en/high-performance-computing-fabrics/omni-path-architecture-fabric-overview.html>

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

Intel, the Intel logo, Intel Xeon and Intel Open are trademarks of Intel Corporation in the U.S. and/or other countries. *Other names and brands may be claimed as the property of others.