


CASE STUDY

High Performance Computing (HPC)
2nd Generation Intel® Xeon® Scalable Processors
Intel® Optane™ Persistent Memory



Preferred Networks Launches Next-Generation Deep Learning Supercomputer, MN-3, with Custom Training Accelerator

2nd Generation Intel Xeon Scalable processors and Intel Optane persistent memory enable up to 3.5X faster data pipeline¹

MN-3 HPC cluster highlights:

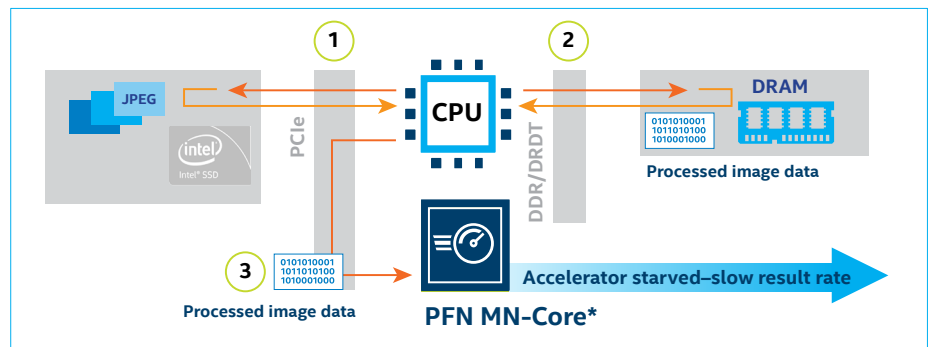
- Supermicro 50-node deep learning training cluster with custom accelerator
- Two 24-core Intel Xeon 8260M processors per node
- 3 TB of Intel Optane persistent memory per node (153.6 PB total)



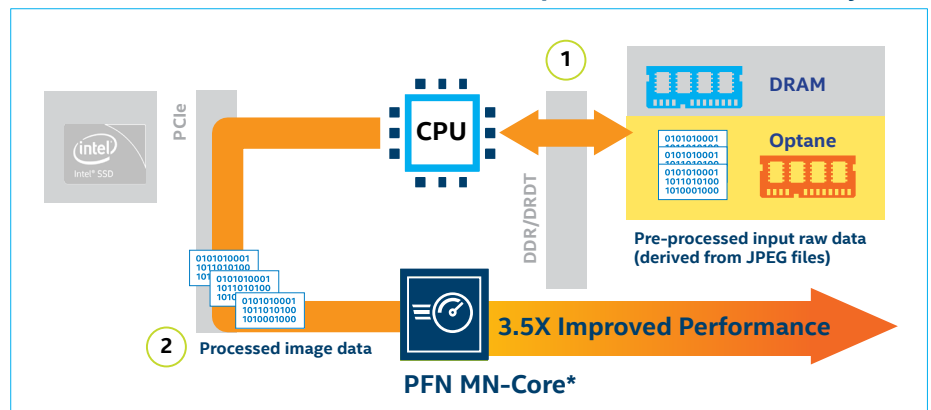
Executive Summary

Preferred Networks (PFN) uses 2nd Gen Intel Xeon 8260M processors and Intel Optane persistent memory to create a high-performance data pipeline to keep their custom, high-performance deep learning training accelerator busy in their new MN-3 HPC system. Located in Tokyo, Preferred Networks is a deep learning company, deploying High Performance Computing (HPC) clusters to build and train algorithms used in domestic and industrial applications. Their latest system, MN-3, integrates a custom-designed deep learning accelerator they engineered. But it is too fast for a traditional storage approach on the nodes.

Workflow with Preferred Networks (PFN) Accelerator



Workflow with PFN Accelerator + Intel Optane Persistent Memory



Traditional SSDs could not meet the I/O throughput requirements of the new architecture, so Preferred Networks turned to Intel Xeon 8260M processors and Intel Optane persistent memory to enable a balanced node with fast access and high capacity for training data.

Challenge

Preferred Networks develops artificial intelligence solutions for industrial and domestic robotics, Industrial Internet of Things (IIoT), manufacturing systems, and other industries. It is a leader in the robotics revolution.² Among its projects, the company is collaborating with industrial systems developer, [FANUC](#), using their unique deep learning capabilities to engineer domestic robots for the home.

The company's Research and Development (R&D) team uses HPC systems designed specifically to create and train algorithms for automated functions, such as:

- Predictive analytics of industrial machines to optimize the use and maintenance of them for increased productivity
- Controlling a robot to easily navigate in a home, recognize objects out of place, pick them up, and put them where they belong
- Other autonomous operations based on vision computing

Preferred Network's largest R&D supercomputers, MN-1 and MN-2, include more than 2500 GPUs total. Yet Preferred Networks needed to accelerate computations to support the many projects the engineering team is working on.

Solution

"We believe more computational power makes our engineers and researchers more effective," vice president of Computing Infrastructure Yusuke Doi said. "By keeping a leadership position in our computational capabilities, we can better compete in our industry and provide advanced solutions to our customers."

So, Preferred Networks designed a unique custom accelerator called MN-Core.³ MN-Core is a custom processor based on a four-die multi-chip package designed specifically for PFN's own R&D projects. The quadruple-chip package—specialized for deep learning training tasks—is at the center of a design for a new supercomputing cluster, MN-3. However, due to the dramatic increase in computing performance, they ran into I/O bottlenecks when they began to design and evaluate the data loading path for the training system.

Many of Preferred Networks' projects are computer vision problems. The training data set, consisting of millions of JPEG image files, is archived on a large external storage system. It is not practical to store the entire data set directly in system memory to take advantage of the faster access. For training, the data is first copied to the nodes into high-performance NVMe SSD drives.

"We first benchmarked node performance with the Intel Xeon 8260M processors," engineer Tianqi Xu of Preferred Networks explained. "During the I/O phase, the processor must get the JPEG files out of block storage and into memory, decode them, and then perform model-specific augmentations. With the 2nd Generation Intel Xeon Scalable processors and current GPUs, the node was well balanced for I/O, computing, and storage."

But with terabytes of data to move during training and the I/O challenges discovered in the data path, traditional storage hierarchy with SSDs would not be able to keep up with the custom accelerator. The accelerator would be starved for data. Preferred Networks needed high capacity storage at DIMM-like speeds in the node. Engineers worked directly with Intel to

understand how the high memory bandwidth of 2nd Gen Intel Xeon Scalable processors and support for high-capacity Intel Optane persistent memory could create a very fast and very large data pipeline.

Once Preferred Networks became aware of Intel Optane persistent memory's capability to speed up their AI pipeline they initiated a proof of concept to verify that the design would support high capacity storage. Intel continues to advise the company as it moves forward with their AI technology efforts.

Leveraging a New Hierarchy of Storage with Intel Optane Technology

Intel Optane persistent memory is a high-density, byte-addressable, 3D memory technology in a DIMM form factor that delivers a unique combination of large capacity, low latency, low power, and data persistence. The persistent memory modules integrate a new layer into the memory/storage hierarchy of an HPC system, offering DIMM-like speeds of byte-addressable data access with terabytes of capacity on the memory bus. Most 2nd Generation Intel Xeon Scalable processors support Intel Optane persistent memory modules. A node with the Intel Xeon 8260M processors can support up to 3 TB of Intel Optane persistent memory.

Intel Optane persistent memory can operate in different modes (memory, app-direct, and storage over app-direct). In memory mode, the CPU uses the Intel Optane persistent memory as system memory and uses the system memory (DIMMs) as a cache. In app-direct mode, software is made aware of both types of memory and is configured to direct data reads and writes based on suitability for DRAM or Intel Optane persistent memory. This offers larger capacity and higher performance to Preferred Networks' training processes.

"In memory mode, the entire memory domain would reside in the persistent memory," Xu added, "which means we wouldn't get optimal use of the entire three terabytes. Additionally, deep learning data access patterns are very random. DRAM as cache doesn't work effectively for those accesses. We needed direct control over the persistent memory, so we developed custom code to control it in app-direct mode."

In addition to their own code, Preferred Networks developed a custom library to take advantage of the large capacity, low latency, and byte-addressable features of Intel Optane persistent memory. To optimize performance for the entire data pipeline and custom accelerator, they included a staging phase to pre-process the JPEG images by converting them to raw pixel data and loading the data set into Intel Optane persistent memory.

Result

The company is manufacturing its accelerator and launching MN-3 with the accelerator. MN-3 is a cluster with up to 48 nodes initially. The company will expand MN-3 into a half-precision exascale supercomputer. The Intel Xeon 8260M processors will allow MN-3 to optimize pre-processing performance to stage the data set and effectively handle the post-processing phase to manage the results.

Early benchmarking of the data pipeline with Preferred Networks' accelerator MN-Core, Intel Xeon 8260M processors, and Intel Optane persistent memory is returning up to 3.5X faster data throughput compared to the system

with NVMe SSDs.⁴ Preferred Networks expects to grow the system over five years as much as 20X to exascale performance for deep learning training.

Solution Summary

Preferred Networks has been using HPC clusters for deep learning training to support their customers. They needed more performance, so they built their own deep learning accelerator and the first stage of a new cluster around it named MN-3. Traditional SSDs could not meet the I/O throughput requirements of the new architecture, so Preferred Networks turned to Intel Xeon 8260M processors and Intel Optane persistent memory to enable a balanced node with fast access and high capacity for training data. The new system design is expected to deliver up to 3.5X faster performance, according to Preferred Networks.

Where to Get More Information

Learn more about [Preferred Networks](#).

Explore the capabilities of the [2nd Generation Intel Xeon Scalable processors](#) with integrated Intel® Deep Learning Boost capabilities for accelerated AI inferencing.

Learn more about [Intel Optane persistent memory](#).

Get insights on [Supermicro's SuperServer platform](#).

Solution Ingredients

- 48-node deep learning training cluster with custom accelerator
- Two 24-core Intel Xeon 8260M processors per node
- 3 TB of Intel Optane persistent memory per node (153.6 PB total)



Spotlight on Supermicro

Supermicro's SuperServer hardware was deployed at Preferred Networks. The SuperServer platform offers high levels of performance, efficiency, and supports 2nd Gen Intel Xeon Scalable processors.

Supermicro (Nasdaq: SMCI), is a leading innovator in high-performance, high-efficiency server technology, is a premier provider of advanced Server Building Block Solutions for Data Center, Cloud Computing, Enterprise IT, Hadoop/Big Data, HPC and Embedded Systems worldwide.

¹ Benchmark information provided by [Preferred Networks](#).

² <https://www.rocketinspace.com/corporate-innovation/10-of-japans-top-robotics-companies-leading-the-world-into-2020>

³ <https://projects.preferred.jp/mn-core/en/>

⁴ Benchmark Information provided by Preferred Networks who measured throughput with the following steps: data read (from ndarray format), ImageNet augmentation (crop, resize, flip), and memory layout for the Preferred Networks accelerator (e.g. data copy).

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.

Performance results are based on testing as of January 2020 and may not reflect all publicly available security updates.

No product or component can be absolutely secure.

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

Results have been estimated or simulated.

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.