

Preferred Networks is a leader in research and development of deep learning and is actively involved in many joint projects with leading corporations in Japan. The company, which is developing computing infrastructure in-house, adopted Cisco® Nexus® 9000 Series data center switches for its large-scale cluster called MN-2.

Preferred Networks (PFN, hereinafter) was established in March 2014 with the aim of delivering state-of-the-art technologies such as deep learning to the world. The company has a vision of “Making the real world computable.” By fusing software and hardware in a sophisticated manner, they are striving to solve real-world problems. PFN’s unique strength lies in the ability to combine deep learning with profound expertise of various fields to research and develop state-of-the-art technologies. The company’s problem-solving abilities are highly regarded by many companies. PFN is actively involved in joint projects with leading manufacturers, construction companies, and medical institutions in Japan.

PFN also developed Chainer, which is the deep-learning framework, which incorporated the Define-by-Run model description method and had a significant influence on PyTorch that was released later.

Challenges

- MN-1 and MN-1b, the computing infrastructure designed for deep learning, adopted 10-Gbps Ethernet for storage access, which was likely to become a bottleneck.
- The bottleneck had to be removed in PFN’s next model MN-2, but PFN wanted to avoid an overlapping investment of InfiniBand, which was used for node-to-node communication, and high-speed Ethernet.
- PFN wanted to visualize burst traffic that should occur between nodes for a short period in multi-node computing infrastructure designed for deep learning.

Solutions

- PFN integrated their network into Ethernet by adopting RoCEv2 to implement RDMA, required for multi-node deep learning on Ethernet.
- The company secured sufficient bandwidth by configuring a leaf-spine network with Cisco Nexus 9000 Series switches.
- The hardware-based streaming telemetry embedded to the Network Processing Unit of Cisco Nexus 9000 Series switches enabled PFN to visualize short burst traffic.

Results and the future

- Integrated network eliminated bottlenecks while avoiding overlapping investments.
- High stability of Cisco products and the network engineers’ familiarity with these products enabled prompt troubleshooting.
- Visualizing burst traffic helped to identify a bottleneck and made it easier to decide on a next investment.

Challenge

Developed one of the most powerful computing infrastructures in-house

One of the core strengths of PFN is that they develop hardware specifically for deep learning in-house. In September 2017, PFN started operating their first large-scale parallel computer called MN-1. Powered by 1024 units of NVIDIA Tesla P100 GPU, this multi-node computer was one of the most powerful computing infrastructures in Japan's private sector.

In the following November, the company ran the distributed deep-learning package called ChainerMN on MN-1. It took only 15 minutes for MN-1 to complete training on ImageNet, an image classification data set commonly used as a benchmark, achieving the fastest record in the world at that time. Around the same time, MN-1 recorded LINPACK performance of approximately 1.39 PetaFLOPS, making it the world's top 12th and the 1st in Japan among industrial supercomputers ranked in the TOP 500 of the most powerful supercomputers as of November 2017.

Subsequently, in July 2018, PFN adopted NVIDIA Tesla V100 32GB GPU to launch the operation of MN-1b, an extended version of MN-1. In July 2019, the company started the operation of MN-2, the latest multi-node GPGPU (General-Purpose Computing on GPU) computing infrastructure with the NVIDIA V100 Tensor Core GPUs.

Eliminated communication bottlenecks with 2nd-gen computing infrastructure

It was around the summer of 2016 when the project of the latest computing infrastructure—MN-2—kicked off.

“At first, we started discussing if we should develop the platform in-house or use an external service. In the fall of 2018, we decided to build the platform on our own,” says Yusuke Doi at PFN.

PFN then started working on detailed specifications. After deciding on the GPUs to be adopted, the company considered the methods to implement CPUs and nodes and the heat discharge. One of the important themes are the networks connecting nodes and connecting a node with storage.

The previous models MN-1 and MN-1b adopted InfiniBand for the network connecting nodes, and onboard 10-Gbps Ethernet for the network connecting a node and storage. The reason why it was not possible to use InfiniBand for storage access is that InfiniBand was not compatible with HDFS (Hadoop Distributed File System) that was adopted as a distributed file system. The reason for using onboard Ethernet came from constraints of PCI slots provided with each node. “While node-to-node communication was fast enough, storage access could easily become a bottleneck,” explains Hirochika Asai at PFN.

High reliability, quick response to the latest protocol, and hardware-based streaming telemetry are the reasons why PFN chose Cisco

“What we needed the most for the MN-2 construction was to eliminate network bottlenecks. But, we also wanted to avoid overlapping investment of InfiniBand and high-speed Ethernet.”

Yusuke Doi

Corporate Officer, VP of Computing Infrastructure, Preferred Networks, Inc.

Solution

Integrated network into ethernet to avoid overlapping investment

The bottleneck of storage access was not the only problem with the network for MN-2. Another was investment related. Doi says, “We wanted to avoid overlapping investment which should occur by adopting high-speed Ethernet to eliminate the bottleneck in addition to InfiniBand.”

What is better? To integrate storage access by using InfiniBand, or to integrate node-to-node communication by using Ethernet? PFN ultimately chose to integrate the communication that used to be interconnected by InfiniBand into high-speed Ethernet.

The core technology of InfiniBand is RDMA (Remote Direct Memory Access), which writes data directly to a memory of a destination node. The company decided to adopt RoCE (RDMA over Converged Ethernet) v2, where RDMA is applied on top of the Ethernet link layer and general IP/UDP layer.

“When we attended a supercomputer conference in October 2018, we participated in a tutorial comparing RoCEv2 and InfiniBand. The data from the tutorial demonstrated that network integration with RoCEv2 would be possible, which convinced us to go with that technology,” explains Arai.

Why partnered with Cisco—three reasons

Once PFN decided to go with Ethernet for MN-2, they contacted Cisco.

Shortly after that, the company borrowed a Cisco Nexus Series tester and validated its quality. Once they received satisfactory results, they determined the overall specifications of MN-2. PFN decided to adopt leaf-spine architecture where installing four 100-Gbps NICs to a node with eight GPU units and connect them to Cisco Nexus 9364C switch through Cisco Nexus 9336C-FX2 switch.

Asai explains the reasons for adopting Cisco products as follows.

“First of all, we valued high reliability. From my long years of experience in a network field, I think highly of NX-OS installed in Cisco Nexus for its extremely high level of stability with an extensive track record. The second reason is Cisco’s prompt response to the latest protocol. We expected that Cisco would be capable of supporting cutting-edge technology like RoCEv2. The third reason is that we can use the hardware-based streaming telemetry embedded in the Network Processing Unit (NPU). In a multi-node deep-learning computing infrastructure, communication for distributed computation called All-Reduce is required in a certain interval, and short burst traffic should occur between nodes. Hardware-based streaming telemetry will allow us to visualize such traffic more easily.”

PFN finished building MN-2 in June 2019. They started the operation in the following month. The processing capacity of MN-2 combined with MN-1 and MN-1b has reached approximately 200 PetaFLOPS.



Yusuke Doi

Corporate Officer, VP of Computing Infrastructure
Preferred Networks, Inc.



Hirochika Asai

Ph.D. in Information Science and Technology Researcher
Preferred Networks, Inc.

Results and the future

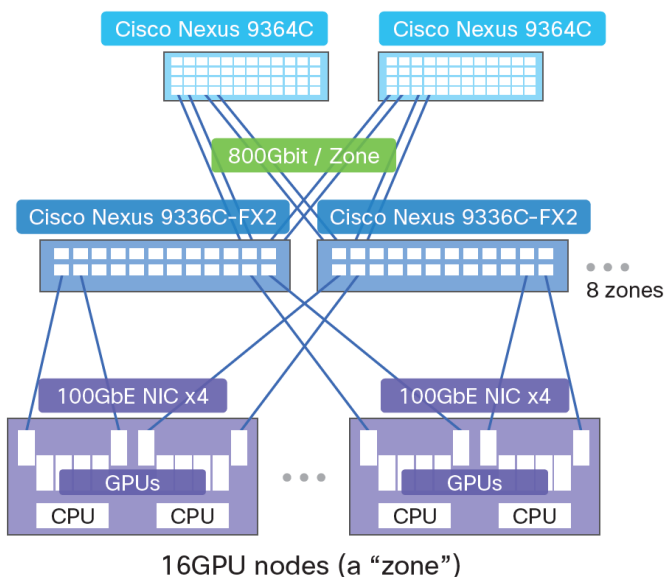
Constructed high-speed network without overlapping investment

The integration of node-to-node communication and storage access into Ethernet allowed them to avoid overlapping investment in the network, improving investment efficiency of MN-2 compared to that of MN-1 and MN-1b. “In this configuration, we installed four NICs to each node to expand bandwidth. This was achieved because of improved investment efficiency,” says Asai.

Cabling became simple as well. Compared to a configuration using InfiniBand together with Ethernet, the number of cables can be reduced to half. “The performance as a whole has also been optimized,” continues Asai.

In deep learning, depending on training data, node-to-node traffic could be heavier than storage access traffic, or vice versa. If both are independent networks, a bottleneck of either of them could limit the intended performance. However, when the network is integrated, the bandwidth can be flexibly allocated, making it easier to avoid such bottleneck. “By concentrating our investment, we can easily expand as the network scales up in the future,” says Doi.

Stable network enabling secure restart on legal inspection



Ethernet works well with HDFS, and it is easier to identify the cause and troubleshoot with Ethernet in case a problem occurs, which is evaluated as a great advantage.

“We have experienced little network trouble so far, but in case we do, Cisco products would be the best choice. We have many Linux engineers, so we initially thought of configuring the network using White Box. But not all Linux engineers are used to working with White Box as a network device. On the other hand, most network engineers can work with Cisco products. Adopting Cisco products allows us to improve our troubleshooting capabilities while minimizing education costs,” says Asai.

All facilities with MN-2 installed must have a legal inspection once a year, and all the power supplies must be shut off at that time. PFN would not have to worry during this time because they are using Cisco products for the network. The power can be turned off without typing a shutdown command, and all the devices will start working fine after restart. “We will not get this peace-of-mind feeling with White Box,” says Asai.

Visualizing burst traffic simplifies investment decision-making

Hardware-based streaming telemetry also made

“The All-Reduce traffic is likely to be a bottleneck when we try to speed up deep-learning computation. Without sufficient bandwidth, the period in which this traffic occurs becomes longer; no matter how fast each node can perform calculation, the overall performance will not be improved.

Case study
Cisco Public

Since the All-Reduce traffic is burst traffic, software-based telemetry showing an average cannot capture all of it. On the contrary, hardware-based streaming telemetry allows for visualization of burst traffic for even a dozen milliseconds and identifying what has become a bottleneck. Seeing such information has made it easier for us to determine where to invest next,” says Doi.

Participation of network expert helped PFN succeed

Upon launching the operation of MN-2, PFN’s businesses are growing at an even faster pace in a wide range of areas, including projects related to medical or resources, as well as research fields. The participation of Cisco into the MN-2 project turned out to be such a notable achievement for PFN, because they are now able to have a more in-depth discussion on a network.

“For MN-1 and MN-1b, after carefully considering feasibility-related safety, we had to design them conservatively. But for MN-2, we were able to adopt a more challenging design thanks to Cisco’s knowledge. We look forward to Cisco’s Network Processing Unit and product development in the future,” says Doi.

In fact, PFN is developing MN-Core, a dedicated chip optimized for matrix operation, which is an essential characteristic deep learning, to further enhance computing infrastructure. The company is continuing to take on new challenges, such as launching the operation of MN-3 with the new chip installed in May 2020. The network for this latest infrastructure will be configured in the same way as MN-2. The stable network that can easily circumvent bottlenecks is helping PFN to take on new challenges and make a leap forward.

Product and services

- Cisco Nexus 9364C
- Cisco Nexus 9336C-FX2



Preferred Networks, Inc.

Preferred Networks is a startup company founded in 2014, which is engaged in research and development focusing on deep learning. The company aims to leverage deep learning and robotics technology for businesses and promotes innovation in various fields. PFN’s strength is not limited to software development—they develop processors designed for deep learning and large-scale computing infrastructure in-house.

With innovative and essential technologies, the company is building by fusing software and hardware, they continue to venture into the unknown.