Case study

Overcoming AI Infrastructure Challenges with Intel® Ethernet 800 Series

Results at a glance:

- Compelling cost efficiency , with up to **$1.1 million in cost savings** when deployed at scale compared with the competition[3,4]
- Widespread availability with **Intel® Ethernet Network Adapters**
- **Strong performance** on parity with competition

Executive Summary

The Intel® Tiber™ Developer Cloud, an AI-as-a-service platform, faced significant delays and other challenges in scaling its AI cloud infrastructure using NVIDIA Mellanox ConnectX-5 Adapters. The high cost, proprietary ecosystem, and extended delivery times of these adapters led the Intel team to look for a high-performance alternative.

The Intel team optimized Intel® Ethernet E810 Network Adapters for distributed storage systems and conducted proof-of-concept benchmark tests that confirmed Intel Ethernet E810 Network Adapters offer performance parity with NVIDIA Mellanox ConnectX-5 Adapters.

As a result, the Intel Tiber Developer Cloud can now future-proof and steadily scale its AI infrastructure using Intel® Ethernet 800 Series Network Adapters while benefiting from the compelling performance, lower cost, faster time to deployment, and open standards these network adapters provide. This simplifies deployment of future connectivity upgrades and capacity expansions to meet evolving AI data center needs.

Challenge: High costs and long lead times for essential AI hardware

The Intel Tiber Developer Cloud provides developers, researchers, and enterprises with access to the latest hardware and software technologies for AI, machine learning (ML), high-performance computing (HPC), and data-intensive workloads. The cloud-based platform enables developers to build, test, and optimize AI and machine learning models in a scalable, flexible environment that reduces development costs while accelerating AI innovation.

"The Intel Tiber Developer Cloud is effectively AI as a service. It is a cloud service that offers AI capabilities. We have customers that come to us to kind of kick the tires on some of the AI solutions that we have, and we have other customers that use it for serious production work," explained Steven C. Miller, Intel Fellow AI Data Center.

Like many AI data centers, the Intel Tiber Developer Cloud is experiencing dramatic growth in demand for its resources and needs to rapidly build more high-performance AI clusters. The AI cloud service also faces challenges of extended delivery times, high costs, and deployment complexity that come with scaling AI resources that depend on specialized hardware from a single vendor's closed ecosystem.

In addition to thousands of GPUs, AI accelerators, and Intel® Xeon® Scalable processors, the Intel Tiber Developer Cloud relies on a large number of high-performance network adapters to rapidly transfer massive volumes of data across multiple nodes in the data center. Bottlenecks in the network can result in inefficient utilization of expensive compute resources—leading to longer training times, increased costs, and lower customer satisfaction among the users.

"In order to provide a meaningful cloud service, along with AI hardware and the actual accelerators that are in there, we have to provide storage as well. We use a very high-performance storage system built by a leading provider of HPC storage solutions. And like most high-performance storage systems, its primary network adapter was the NVIDIA Mellanox ConnectX series," said Miller. "As we were ordering for our next generation of clusters, the lead times for the adapters were going out longer and longer. That definitely put some of the builds at risk."

Intel Tiber Developer Cloud's system architects wanted ways to scale and optimize AI compute and storage resources while also eliminating the constraints, delays, and high cost related to sourcing essential, high-performance adapters.

Solution: Intel® Ethernet 800 Series Network Adapters

To address these challenges, the team sought a cost-effective, high-performance alternative to NVIDIA network adapters. They identified the Intel® Ethernet 800 Series Network Adapters as potential replacements. The Intel® Ethernet E810 Network Adapters are widely available and can provide comparable performance for AI and data-intensive workloads, often at a lower cost per unit.

The Intel Tiber Developer Cloud team collaborated with Intel to design a proof-of-concept benchmark test that compared the performance of both the Intel Ethernet E810 Network Adapter and NVIDIA Mellanox ConnectX-5 adapter in a distributed-storage data center environment. To conduct the test, however, the team first needed to address some previously identified data-performance limitations in Intel Ethernet E810 Network Adapters' software stack.

> Intel® Ethernet E810 Network Adapters offer high throughput and reduced latency, which are critical for optimizing performance with distributed storage for AI workloads.

> "A distributed storage system's performance is heavily tied into the network infrastructure's capability. We used SR-IOV in switchdev mode and DPDK running on virtual functions to guarantee its networking performance for AI data centers."
>
> **– Brian Johnson, Principal Engineer Solutions Architect, Intel**

Brian Johnson, Intel Principal Engineer Solutions Architect explained that the team used Single Root I/O Virtualization (SR-IOV) in switchdev mode and the Data Plane Development Kit (DPDK)—key features of Intel Ethernet E810 Network Adapters—to remove those limitations.

"A distributed storage system's performance is heavily tied into the network infrastructure's capability," Johnson said. "We used SR-IOV in switchdev mode and DPDK running on virtual functions to guarantee its networking performance for AI data centers."

To minimize the overhead associated with the network software stack, deploying SR-IOV in switchdev mode allows a VF to run within the DPDK container, using a single dedicated CPU core per instance. This configuration shifts the traffic routing from a software-based switch within the operating system to the hardware e-switch on the network adapter, effectively bypassing the kernel's software stack. Plus, using a DPDK application enables the use of the DPDK Poll-Mode Driver that also bypasses the kernel stack and enables zero-copy in addition to other DPDK library enhancements.

In short, by avoiding the kernel software stack, Intel® Ethernet E810 Network Adapters can now achieve higher network throughput and reduced latency, which are critical for optimizing performance with distributed storage for AI workloads.

With this technical hurdle overcome, the team was then able to conduct benchmark tests that compare the performance of NVIDIA Mellanox ConnectX-5 and Intel Ethernet E810 Network Adapters in the storage environment.

Results: Lower costs and shorter lead times

Consistent components of the test cluster included Inspur servers running Ubuntu 22.04.3 with the 5.15.0-92-generic kernel, Supermicro storage nodes, and KIOXIA storage media, each with a capacity of 13.97 TiB.

Intel® Ethernet Network Adapter E810-CQDA2.

The project team performed industry-standard FIO[1] and IO-500[2] HPC benchmark tests using either an Intel® Ethernet Network Adapter E810-CQDA2 or NVIDIA ConnectX-5 Adapter per client for network connectivity.

"In a typical distributed storage infrastructure, the networking is often the bottleneck for the overall performance. In our comparative analysis, the only variable was the type of adapter being used on the storage client," said Vincent Du, Intel System Architect. The FIO test results showed that both the Intel Ethernet Network Adapter E810-CQDA2 and the NVIDIA Mellanox ConnectX-5 Adapter can saturate the storage network bandwidth using the same workload.

The team plans to evaluate the soon-to-be-released Intel® Ethernet E830 Network Adapters, which offer up to 200 gigabits per second (Gbps) max data rate and comprehensive features for data center needs.

"We crafted the test so that the workload we're running saturates the network bandwidth for each client node," Du explained. "In this case, for the 100GbE interface of the NVIDIA Mellanox ConnectX-5 Adapter, the upper limit of the network bandwidth should be 12.5 gigabytes per second. Minus overhead, we were getting around 11 gigabytes per second. When we replaced it with the Intel Ethernet E810 Adapter and repeated the test, we achieved the same network bandwidth saturation—around 11 gigabytes per second."

The IO-500 benchmark results also showed similar comparable scores for Intel Ethernet E810 Network Adapters versus NVIDIA Mellanox ConnectX-5 Adapters. This enabled the team to conclude that Intel Ethernet E810 Network Adapters are equally as performant as NVIDIA Mellanox ConnectX-5 Adapters for Intel Tiber Developer Cloud's AI infrastructure, but at a significantly lower cost and with shorter delivery timelines.

Cost comparison

Large data centers with distributed storage like the Intel Tiber Developer Cloud depend on many thousands of high-performance network adapters to ensure high network throughput and low latency. The recommended customer price for Intel Ethernet E810 Network Adapters is $557.00-$580.00 per unit.[3] Meanwhile, NVIDIA's recommended customer price for the NVIDIA Mellanox ConnectX-5 Adapter model tested is $1,169 per unit.[4]

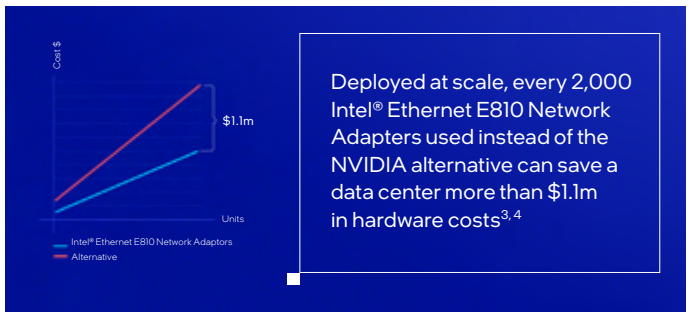

Deployed at scale, this means every 2,000 Intel Ethernet E810 Network Adapters used instead of the NVIDIA alternative, can save a data center more than $1.1 million in network hardware costs.

Intel's Steven Miller notes that eliminating the opportunity cost and the risks of lengthy delivery timelines are huge wins as well. "The biggest benefit for ITDC is that we're not supply constrained. We can move this off the supply constrained list. So that's no longer a gating item for us to build," he said.

Future plans

Having demonstrated performance parity between Intel Ethernet E810 Network Adapters and NVIDIA Mellanox ConnectX-5 Adapters, the team is evaluating the soon-to-be-released Intel(R) Ethernet E830 Network Adapters, which offer up to 200 gigabits per second(Gbps) max data rate, PCIe 5.0x8 host interconnect support, precise timing capabilities and comprehensive security and manageability features. The network adapter can support higher bandwidth workload requirements. The team is also considering Intel® Infrastructure Processing Unit (IPU) Adapters. Intel IPU Adapters are capable of a wide range of infrastructure-related tasks, including isolation of tenant and provider network and storage (NVME) offloads, security, storage and virtualization in addition to networking.

Learn more

- Explore Intel® Ethernet 800 Series Network Adapters
- Get started with the Intel® Tiber™ Developer Cloud
- Learn more about Intel® Infrastructure Processing Units (IPUs)
- Discover the latest generation Intel® Gaudi® 3 AI Accelerators

Intel IPU Adapters will provide additional features and capabilities that can be leveraged in the developer cloud.

Building a higher-performance data storage system also enables Intel Tiber Developer Cloud to get the most out of the latest generation Intel® Gaudi® 3 AI accelerators as they deploy them. Whether training large models or handling demanding inference tasks, rapid access to large volumes of data will allow users to tackle complex AI workloads faster.

Looking ahead, the team will continue to prioritize customer experience and efficiency. By investing in advanced infrastructure and key technologies, Intel aims to further accelerate developer access to its cloud platform. Through streamlined processes and optimized resource allocation, the Intel Tiber Developer Cloud is committed to providing a seamless and responsive experience that empowers customers to innovate and achieve their goals more rapidly.

> The latest generation Intel® Gaudi® 3 AI accelerators will help the team train large models and handle demanding inference tasks to tackle complex AI workloads faster.

intel®

[1] https://fio.readthedocs.io/en/latest/fio_doc.html

[2] https://io500.org/about

[3] https://ark.intel.com/content/www/us/en/ark/products/192558/intel-ethernet-network-adapter-e810-cqda2.html

[4] https://www.hyperscalers.com/HHHL-100Gb-QSFP28-2-por-MELLANOX-CX516A-ConnectX-5-EN-PCI-E-X16-Gen-3-ADPC0516000-MCX516A-CCAT

All websites last accessed November 2024.

Performance varies by use, configuration, and other factors. Learn more on the Performance Index site.

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.

Intel technologies may require enabled hardware, software or service activation.

Your costs and results may vary.

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

All product plans and roadmaps are subject to change without notice.

Statements in this document that refer to future plans or expectations are forward-looking statements. These statements are based on current expectations and involve many risks and uncertainties that could cause actual results to differ materially from those expressed or implied in such statements. For more information on the factors that could cause actual results to differ materially, see our most recent earnings release and SEC filings at www.intc.com.

1224/PK/CAT/PDF ♻Please Recycle 362350-001EN