

## Optimize Innovation for AI and HPC with Ethernet Products

Maximize scalability, flexibility, and ROI with the performance needed for high performance computing (HPC) and AI workloads, using the Intel® AI connectivity portfolio

### Intel® Ethernet Fabric Suite at a glance:

Proven software that powers AI and HPC clusters via Ethernet

Compatible with compute HW from multiple vendors and with existing application binaries and source code

Third-generation, mature, proven solution that leverages the OpenFabrics Alliance libfabric APIs

Integrated with Intel® oneAPI Toolkit

Includes FastFabric Deployment Tools

### Support advanced fabric demands with the technology you already know

Advanced AI and HPC workloads are introducing new demands for computational horsepower and network throughput. To support these needs, organizations are deploying server clusters to handle the massive amounts of data processing required. Server clusters are then connected via a high-performance network—commonly referred to as a fabric—that supports the demanding requirements of AI connectivity.

Building server clusters introduces a number of budget considerations, including the cost of new server nodes and their supporting infrastructure. Proprietary fabrics supporting AI and HPC application workloads can be costly and limit the number of nodes an organization can afford. Additionally, proprietary fabrics require specialized knowledge to implement and maintain.

To help solve these challenges, Intel offers an Ethernet-based approach to connectivity and high-performance networking across AI, HPC, and front-end data center needs. Our portfolio includes network adapters, innovative infrastructure processing units, and a robust software suite designed to help make your adoption of Ethernet-based connectivity simple and easy.

### Enabling a feature-rich Ethernet fabric with ecosystem partners

Powered by the Intel® Ethernet Fabric Suite (Intel® EFS), our high-performance networking approach extends the capabilities of Remote Direct Memory Access over Converged Ethernet (RoCEv2) to enable an Ethernet-based fabric for HPC and AI. As an evolution of the Intel® Omni-Path Architecture, Intel EFS is tuned to optimal performance with OpenFabrics Interfaces (OFI) to scale and help provide the required function.

In collaboration with our industry and ecosystem partners, Intel offers proven and powerful AI connectivity solutions to support diverse, high-performance Ethernet networking needs. Our portfolio enables you to use Ethernet for scale-out networks that connect server nodes in a cluster as well as front-end data center networks—removing the need to manage separate fabrics for your AI and HPC initiatives.

## Driving AI and HPC fabric innovation

Intel believes the Ethernet approach to fabric is the future. Because Ethernet is an open standard supported by a broad ecosystem, organizations that use it for AI and HPC can benefit from expanded vendor choice, a diverse supplier base, more flexibility, and accelerated innovation.

Our standards-based vision for Ethernet in HPC and AI is aligned with other industry leaders such as the Ultra Ethernet Consortium (UEC) and the Ultra Accelerator Link (UAL) standards group. We're working with these groups and others across the industry to deliver an Ethernet-based, open, interoperable, high-performance, full-communications stack architecture that can meet growing network demands at scale.

### Innovating with Google's Falcon

We're also working to facilitate the adoption of the new Falcon hardware-assisted transport layer for Ethernet-based fabric deployments. Developed by Google and released publicly through the Open Compute Project (OCP), Falcon leverages many of the latest networking technologies to modernize RDMA, delivering better bandwidth, higher utilization, and increased robustness without the complexities of operating Ethernet as a lossless fabric.

## The Intel® AI connectivity portfolio

Our AI connectivity product portfolio is readily available through Intel OEM and channel partners to help support your needs across scale-out and front-end networks.

### Innovative infrastructure acceleration technology

**Take advantage of new ways to help maximize performance for demanding workloads.** The Intel® Infrastructure Processing Unit Adapter E2100 (Intel® IPU E2100) allows you to offload services and infrastructure tasks from the CPU to optimize performance and better support AI and HPC applications. It's ideal for small and large clusters and provides high-performance, efficient RDMA with support for RoCEv2 and Falcon. The Intel IPU E2100 represents the world's first Falcon-compatible device.

### High-performance Ethernet adapters

**Use familiar technology to help support emerging needs and avoid unnecessary complexity.** The Intel® Ethernet 800 series offers standard Ethernet network adapters and controllers designed for the rigors of HPC and AI clusters, as well as many other potential workloads.

### The Intel® Ethernet Approach

#### Open ecosystem

Enable an open ecosystem for AI networking solutions that eliminates single-vendor lock-in.

#### Standard protocols

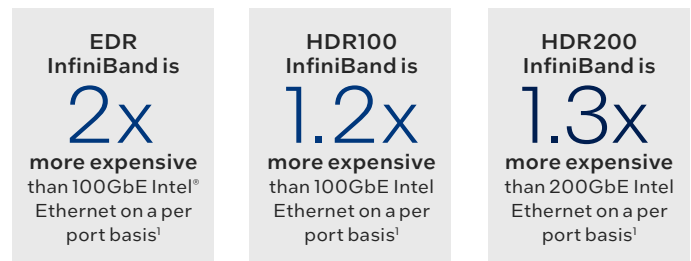
Replace proprietary fabrics with Ethernet and embrace a standards-based approach.

#### Partnerships

Build partnerships to enable complete, open solutions. Differentiate offerings via cost, power, system integration, and open software.

### Why pay more for the same application performance?

Ethernet-based AI and HPC cluster architectures deliver comparable performance to proprietary fabrics—at a fraction of the cost.



Note: Tested on clusters of 100 nodes or fewer. Your costs may vary.

Using our high-performance Ethernet network adapters helps you power AI clusters at scale with the capability to support job sizes of approximately 100 nodes and fabric sizes up to hundreds of nodes. These solutions also help support powerful, AI-optimized front-end data center hardware such as Intel® Gaudi® AI Accelerators, which are specialized compute options designed to enable high performance and enhanced efficiency for large-scale AI efforts.

### Simpler development with Intel EFS

**Seamlessly adopt Ethernet and easily transition your applications.** This suite plays a key role in using Ethernet to support AI and HPC applications. Applications and middleware that support OFI will run via Intel EFS without change. Additionally, applications that work with Intel® Message Passing Interface (Intel® MPI) will see binary application compatibility with various Intel MPI versions and fabrics, even with existing or older MPI applications. Intel EFS is available in most major OS distributions.

## Ethernet for AI and HPC: Benefits



Comparable performance to proprietary fabrics



Familiarity and ease of use to help simplify adoption and operations



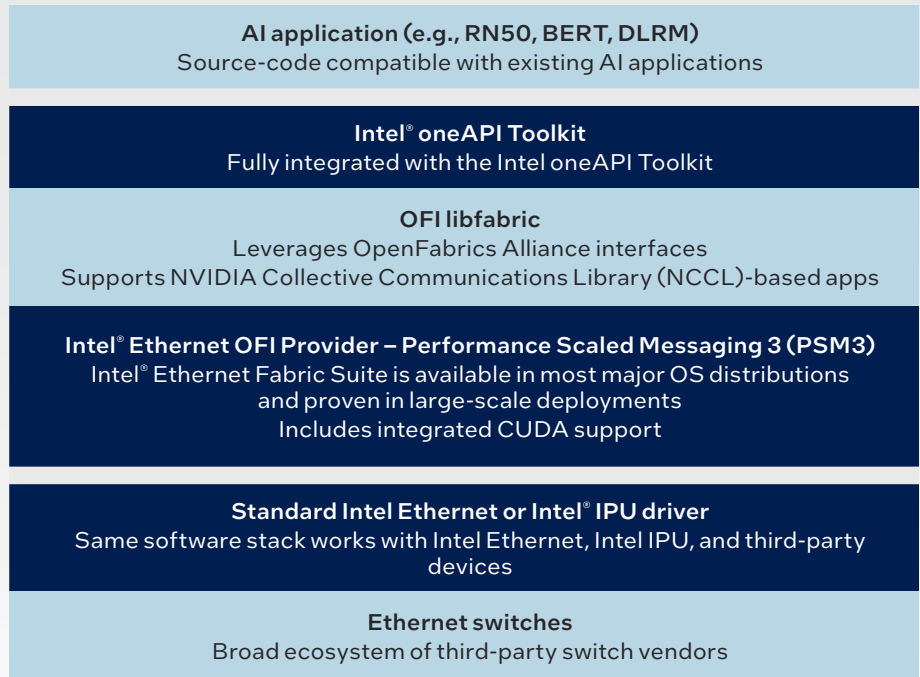
Lower total cost of ownership compared to proprietary solutions



Standards-based and open for flexibility and scalability

### The end-to-end AI and HPC Ethernet stack

By combining Intel® products with the hardware and software components offered by our ecosystem partners, you can build comprehensive, Ethernet-based fabric solutions.



Delivered by Intel  Delivered by third-party vendors

## Start simplifying your AI and HPC connectivity today

Alongside our industry partners, Intel is enabling Ethernet-based fabric solutions. Our hardware and software innovations are ready to help you deliver the fast and scalable connectivity required for AI and HPC workloads.

Learn more about Intel® AI Ethernet connectivity solutions at [intel.com/ethernet](https://intel.com/ethernet)



#### Notices and disclaimers

1. Source: PC Nation.com, April/May 2023.

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 configuration disclosure for details.

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

No product or component can be absolutely secure.

Your costs and results may vary.

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

© 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.

0924/LD/CMD/PDF