


Intel® Ethernet Network Adapter XXV710 for OCP



10/25GbE adapters provide ultimate flexibility and scalability
for Cloud and Data Center deployments

Key Features

- OCP Spec. v2.0, Type 1 and Type 2
- PCI Express (PCIe) 3.0, x8
- Supports dual port 1/10/25GbE configurations
- IEEE 802.3by and Ethernet Technology Consortium specification compliance
- Network Virtualization offloads including VxLAN, GENEVE, NVGRE, MPLS, and VxLAN-GPE with Network Service Headers (NSH)
- Intel® Ethernet Flow Director for hardware-based application traffic steering
- Dynamic Device Personalization (DDP) enables increased packet processing efficiency for NFV and Cloud deployments
- Data Plane Development Kit (DPDK) optimized for efficient packet processing
- Excellent small packet performance for network appliances and Network Functions Virtualization (NFV)
- Intelligent offloads to enable high performance on servers with Intel® Xeon® processors
- I/O virtualization innovations for maximum performance in a virtualized server
- Adaptive link establishment enables increased interoperability with other 25GbE capable switches and host controllers

Overview

As a founding member of Open Compute Project (OCP), Intel strives to increase the number of open solutions based on OCP specifications. The Intel® Ethernet Network Adapter XXV710 for OCP is part of the Intel® Ethernet 700 Series, and offers 1/10/25GbE port speeds.

700 Series Network Adapters are the foundation for server connectivity, providing broad interoperability, critical performance optimizations, and increased agility for Communications, Cloud, and Enterprise IT network solutions.

- **Interoperability** - Multiple speeds and media types for broad compatibility backed by extensive testing and validation.
- **Optimization** - Intelligent offloads and accelerators to unlock network performance in servers with Intel® Xeon® processors.
- **Agility** - Both Kernel and Data Plane Development Kit (DPDK) drivers for scalable packet processing.

Intel® Ethernet 700 Series delivers networking performance across a wide range of network port speeds through intelligent offloads, sophisticated packet processing, and quality open source drivers.

All Intel® Ethernet 700 Series Network Adapters include these feature-rich technologies:

Flexible and Scalable I/O for Virtualized Infrastructures

Intel® Virtualization Technology (Intel® VT), delivers outstanding I/O performance in virtualized server environments.

I/O bottlenecks are reduced through intelligent offloads, enabling near-native performance and VM scalability. These offloads include Virtual Machine Device Queues (VMDq) and Flexible Port Partitioning using SR-IOV with a common Virtual Function driver for networking traffic per Virtual Machine (VM). Host-based features supported include:

VMDQ for Emulated Path: VMDQ, enables a hypervisor to represent a single network port as multiple network ports that can be assigned to the individual VMs. Traffic handling is offloaded to the network controller, delivering the benefits of port partitioning with little to no administrative overhead by the IT staff.

SR-IOV for Direct Assignment: Adapter-based isolation and switching for various virtual station instances enables optimal CPU usage in virtualized environments.

- Up to 128 virtual functions (VFs), each VF can support a unique and separate data path for I/O related functions within the PCI Express hierarchy.
- Use of SR-IOV with a networking device, for example, allows the bandwidth of a single port (function) to be partitioned into smaller slices that can be allocated to specific VMs or guests, via a standard interface.

Intel® Ethernet Adaptive Virtual Function (Intel® Ethernet AVF): Customers deploying mass-scale VMs or containers for their network infrastructure now have a common VF driver. This driver eases SR-IOV hardware upgrades or changes, preserves base-mode functionality in hardware and software, and supports an advanced set of features in the Intel® Ethernet 700 Series.

Enhanced Network Virtualization Overlays (NVO)

Network virtualization has changed the way networking is done in the data center, delivering accelerations across a wide range of tunneling methods.

VxLAN, GENEVE, NVGRE, MPLS, and VxLAN-GPE with NSH Offloads: These stateless offloads preserve application performance for overlay networks, and the network traffic can be distributed across CPU cores, increasing network throughput.

Flexible Port Partitioning (FPP)

FPP leverages the PCI-SIG SR-IOV specification. Virtual controllers can be used by the Linux host directly and/or assigned to virtual machines.

- Assign up to 63 Linux host processes or virtual machines per port to virtual functions.
- Control the partitioning of per-port bandwidth across multiple dedicated network resources, ensuring balanced QoS by giving each assigned virtual controller equal access to the port's bandwidth.

Network administrators can also rate limit each of these services to control how much of the pipe is available to each process.

Greater Intelligence and Performance for NFV and Cloud deployments

Dynamic Device Personalization (DDP) customizable packet filtering, along with enhanced Data Plane Development Kit (DPDK), support advanced packet forwarding and highly-efficient packet processing for both Cloud and Network Functions Virtualization (NFV) workloads.

- DDP enables workload-specific optimizations, using the programmable packet-processing pipeline. Additional protocols can be added to the default set to improve packet processing efficiency that results in higher throughput and reduced latency. New protocols can be added or modified on-demand and applied at runtime using Software Defined Firmware or APIs, eliminating the need to reset or reboot the server. This not only keeps the server and VMs up, running, and computing, it also increases performance for Virtual Network Functions (VNFs) that process network traffic that is not included in the default firmware. [Download DDP Profiles](#)
- DPDK provides a programming framework for Intel® processors and enables faster development of high-speed data packet networking applications.

Advanced Traffic Steering

Intel® Ethernet Flow Director (Intel® Ethernet FD) is an advanced traffic steering capability. Large numbers of flow affinity filters direct receive packets by their flows to queues for classification, load balancing, and matching between flows and CPU cores.

Steering traffic into specific queues can eliminate context switching required within the CPU. As a result, Intel® Ethernet FD significantly increases the number of transactions per second and reduces latency for cloud applications like memcached.

Features	Description
----------	-------------

General

SFP28 Connectivity	<ul style="list-style-type: none"> Supports SFP28 Direct Attach Copper (DAC), 25GBASE-SR, and 25GBASE-LR physical media as well as SFP+ DAC, 10GBASE-SR and 10GBASE-LR physical media.
Interoperability with other 10/25GbE switches and network adapters	<ul style="list-style-type: none"> Delivers wide interoperability with SFP+, SFP28, and QSFP28 switches and network adapters in the market, enabling smooth adoption and upgrade to 25GbE within a customers' network infrastructure.
Load balancing on multiple CPUs	<ul style="list-style-type: none"> Increases performance on multi-processor systems by efficiently balancing network loads across CPU core when used with Receive-Side Scaling (RSS) from Microsoft or scalable I/O on Linux.
Protect, Detect and Recover	<ul style="list-style-type: none"> The Intel Ethernet 700 Series implements a design philosophy of platform resiliency with 3 attributes supporting the NIST Cybersecurity Framework: Protect, Detect and Recover. These attributes verify the firmware and critical device settings with built-in corruption detection and automated device recovery to return the device to its originally programmed state.
Support for most network operating systems	<ul style="list-style-type: none"> Enables broad deployment for different applications.
RoHS-compliant	<ul style="list-style-type: none"> Complies with the European Union directive 2011/65/EU to reduce the use of hazardous materials.
Time Sync (IEEE 1588, 802.1as)	<ul style="list-style-type: none"> Enables networked Ethernet equipment to synchronize internal clocks according to a network master clock; end point can then acquire an accurate estimate of the master time by compensating for link latency.

I/O Features for Multi-Core Processor Servers

Intel® Ethernet Flow Director (Intel® Ethernet FD)	<ul style="list-style-type: none"> An advanced traffic steering capability increases the number of transactions per second and reduces latency for cloud applications like Memcached.
MSI-X support	<ul style="list-style-type: none"> Minimizes the overhead of interrupts. Load-balancing of interrupt handling between multiple cores/CPU's.
Multiple Queues: 1,536 Tx and Rx queues per device	<ul style="list-style-type: none"> Network packet handling without waiting for buffer overflow providing efficient packet prioritization. Actual number of queues will vary depending upon software implementation.
Tx/Rx IP, SCTP, TCP, and UDP checksum offloading (IPv4, IPv6) capabilities	<ul style="list-style-type: none"> Lower processor usage. Checksum and segmentation capability extended to new standard packet type.

Virtualization Features

Next-generation VMDq	<ul style="list-style-type: none"> Up to 256 maximum VMDq VMs supported. Offloads the data-sorting based on MAC addresses and VLAN tags, functionality from the Hypervisor to the network silicon, improving data throughput and CPU usage.
PCI-SIG SR-IOV Implementation (128 per device)	<ul style="list-style-type: none"> Provides an implementation of the PCI-SIG standard for I/O Virtualization. The physical configuration of each port is divided into multiple virtual ports. Each virtual port is assigned to an individual VM directly by bypassing the virtual switch in the Hypervisor, resulting in near-native performance. Integrated with Intel® VT for Directed I/O (Intel® VT-d) to provide data protection between VMs by assigning separate physical addresses in the memory to each VM. 64/port for dual port.
Virtual Machine Load Balancing (VLMB)	<ul style="list-style-type: none"> VMLB provides traffic load balancing (Tx and Rx) across VMs bound to the team interface, as well as fault tolerance in the event of switch, port, cable, or adapter failure.
Advanced Packet Filtering	<ul style="list-style-type: none"> 1536 exact matched packets (unicast or multicast). 512 hash entries each for unicast and multicast. Lower processor usage. Promiscuous (unicast and multicast) transfer mode support. Optional filtering of invalid frames.
VLAN support with VLAN tag insertion, stripping and packet filtering for up to 4096 VLAN tags	<ul style="list-style-type: none"> Ability to create multiple VLAN segments.
VxLAN, NVGRE, GENEVE, VxLAN-GPE+NSH, MPLS	<ul style="list-style-type: none"> Preserves application performance in network virtualized environments.

Manageability Features

Preboot Execution Environment (PXE) Support	<ul style="list-style-type: none"> Enables system boot via the LAN (32-bit and 64-bit). Flash interface for PXE image.
Unified Extensible Firmware Interface (UEFI)	<ul style="list-style-type: none"> Enables new technologies during the pre-OS boot process and addresses legacy BIOS limitations on hardware.
Simple Network Management Protocol (SNMP) and Remote Network Monitoring (RMON) Statistic Counters	<ul style="list-style-type: none"> Easy system monitoring with industry-standard consoles.
Watchdog Timer	<ul style="list-style-type: none"> Gives an indication to the manageability firmware or external devices that the controller or the software device driver is not functioning.

Specifications

General

Connections ¹	Dual SFP28 cages supporting SFP28 DAC, 25GBASE-SR, and 25GBASE-LR physical media, SFP28 to QSFP28 Direct Attach Breakout Cable, and SFP+ DAC, 10GBASE-SR and 10GBASE-LR physical media.
Network Standard Physical Layer Interfaces	25GBASE-SR/LR (optical transceivers) 25GBASE-CR (DAC). Automatically enables no-FEC, BASE-R FEC and RS-FEC to support CA-N, CA-S and CA-L cables 10GBASE-SR/LR (optical transceivers) 10GbE SFP+ DAC

Technical Features²

Operating Temperature	0 °C to 55 °C (32 °F to 131 °F)
Airflow for Type 1	325 LFM with 55 °C required for CR (DAC) 400 LFM with 55 °C required for SR ² optics (airflow direction: from heatsink to optics)
Airflow for Type 2	300 LFM with 55 °C required for CR (DAC) 350 LFM with 55 °C required for SR ² optics (airflow direction: from heatsink to optics)
Storage Temperature	-40 °C to 70 °C (-40 °F to 158 °F)
Storage Humidity	Maximum: 90% non-condensing relative humidity at 35 °C
LED Indicators	LINK (solid) and ACTIVITY (blinking) LINK SPEED (green = 25Gbps; yellow = 10Gbps)

Adapter Features

Data Rate Supported Per Port	▪ Optical: 1/10/25GbE ▪ Direct Attach: 10/25GbE
Bus Type	PCIe3.0 (8 GT/s)
Bus Width	PCIe x8
Interrupt Levels	INTA, MSI, MSI-X
Hardware Certifications	FCC A, UL, CE, VCCI, BSMI, CTICK, KCC
Controller	Dual port: Intel® Ethernet Controller XXV710

Power Consumption

SKU	Typical Power	Maximum Power
Dual-port 25GBASE-CR	8.8 W	9.7 W
Dual-port 25GBASE-SR	10.7 W	11.8 W

Physical Dimensions

Standard OCP v2.0 form factor, Type 1 and Type 2

Product Order Code

Configuration	Product Code	OCP Type
Dual port	XXV710DA2OCP1	1
Dual port	XXV710DA2OCP2	2

Supported Operating Systems

For a complete list of supported network operating systems for Intel® Ethernet 700 Series Adapters visit: intel.com/support/EthernetOS

Intel® Ethernet Optics

Combine proven, reliable Intel® Ethernet Optics with Intel Ethernet 700 Series and 500 Series Network Adapters for dependable interoperability and consistent performance across the network. Learn more at intel.com/ethernet

Warranty

Intel limited lifetime warranty for retail Ethernet Products, 90-day money-back guarantee (US and Canada).

Customer Support

For customer support options in North America visit: intel.com/content/www/us/en/support/contact-support.html

Product Information

For information about Intel® Ethernet Products and technologies, visit: intel.com/ethernetproducts

¹If you are using optics, please note that extended temperature optics are required for the dual port adapter.

²The Intel Ethernet Network Adapter XXV710 for OCP with SFP28 Open Optics support is designed to support Power Level III modules as defined in the SFF-8419 specification. When Intel® Ethernet SFP28 SR Optics modules are used, adapter use conditions for ambient temperature and airflow have been verified to meet the Standard Temperature Class of Operation as defined in the SFF-8679 specification. For use of other optics modules, it is the system integrator's responsibility to determine the necessary ambient temperature and airflow necessary for the third party optical modules to operate within the Temperature Class of Operation at all times. Operating optical modules outside the supplier specified Temperature Class of Operation range permanently reduces the performance of the optical module over time.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document. Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

The products and services described may contain defects or errors which may cause deviations from published specifications.

© Intel Corporation. Intel, the Intel logo, Xeon, and other Intel marks are trademarks of Intel Corporation or its subsidiaries.

Other names and brands may be claimed as the property of others.