

Intel® Ethernet Controller X550

Integrated Single Chip 10GBASE-T Controller Simplifies 10 Gigabit Ethernet (GbE) Server LOM, Converged Network Adapter and Network Add-in Card Designs



Key Features

- Reliable and proven 10 GbE technology from Intel® Corporation.
- Intel's second generation, single and dual-port 10GBASE-T controller with integrated MAC and PHY.
- 10GBASE-T, 1000BASE-T and 100BASE-TX modes.
- Supports NBASE-T* technology (2.5 and 5.0 GbE over CAT 5e).
- IEEE 802.3az Energy Efficient Ethernet (EEE)¹.
- PCI Express* (PCIe*) v 3.0 with up to 8.0 GT/s.
- Unified networking delivering LAN, iSCSI and FCoE over 10GBASE-T.

- Network virtualization stateless offload: VXLAN and NVGRE.
- Intel® Ethernet Flow Director for hardware-based application traffic steering.
- Data Center Bridging (DCB) with FCoE stateless offloads.
- Flexible I/O virtualization for port partitioning and Quality of Service (QoS) of up to 128 virtual ports.
- Integrated IPsec security engines.
- Secure NVM update support.
- Jumbo frames of up to 15.5 KB.
- 11 W maximum power.
- SMBus and NC-SI interfaces with OS2BMC, Management Component Transport Protocol (MCTP) and Wake on LAN (WoL) support.

Product Overview

Intel continues its legacy of Ethernet leadership by introducing the next generation 10 Gb/s family of silicon to support the growing 10GBASE-T ecosystem.

The Intel® Ethernet Controller X550 family is the low cost single-chip 10GBASE-T controller of choice to broadly deploy 10 GbE in platform designs. Now with support for LAN,

FCoE, and iSCSI traffic with added support for VXLAN and NVGRE - 10GBASE-T is the most flexible connectivity for all of your networking requirements. Innovation continues with Intel's second generation integrated 10GBASE-T MAC+PHY, which drives down both cost and power enabling the most cost effective deployment of 10 GbE in the data center. With 10GBASE-T, migration to 10 GbE is dramatically simplified with backward compatibility with your existing GbE network infrastructure.

10GBASE-T Simplifies the Transition from 1 GbE to 10 GbE

The X550 controller provides up to two integrated 10GBASE-T PHYs providing 10 Gb/s of throughput that are also backwards compatible with legacy GbE switches and Cat6A cabling. The ability to auto-negotiate between 100 Mb/s, 1 GbE, and 10 GbE speeds provide the backwards compatibility for a smooth transition and easy migration to 10 GbE.

In addition, the X550 controller also supports NBASE-T* technology. NBASE-T technology can increase network speeds of 2.5 Gb/s or 5 Gb/s using existing CAT 5e/CAT 6 cables at lengths up to 100 m.

Intel's NBASE-T technology can be firmware upgraded to support the emerging IEEE 802.3bz specification¹.

Exciting New Data Center Models

More than just a 10 times per-port increase in performance by using the X550 controller versus a standard 1 GbE controller opens doors for exciting new usage models, including unified networking, I/O virtualization, and flexible port partitioning.

A Complete Unified Network Solution

Converging data and storage onto one fabric eliminates the need for multiple adapters and cables per server. Furthermore, 10 GbE provides the bandwidth to converge these multiple fabrics into a single wire. A key capability that makes all this possible is traffic class separation provided by DCB. DCB provides a collection of standards for additional QoS functionality such as lossless delivery, congestion notification, priority-based flow control, and priority groups. This enables the Intel X550 to provide a one-wire solution with virtual pipes for the different classes of traffic:

- Data: best effort delivery of standard LAN traffic.
- Storage: NAS or SAN including lossless FCoE and iSCSI.
- Management: Guaranteed connectivity of data center IP management.

Unified Networking Principles

Intel's unified networking solutions are built on the principles that have made us successful in Ethernet:

- Open architecture integrates networking with the server, enabling IT managers to reduce complexity and overhead while enabling a flexible and scalable data center network.
- Intelligent offloads lower cost and power while delivering the application performance that customers expect.
- Proven unified networking is built on trusted Intel Ethernet technology, enabling customers to deploy FCoE

or iSCSI with the same quality used in their traditional Ethernet network.

Intel's unified networking solutions are enabled through a combination of Intel Ethernet products along with network and storage protocols integrated in the operating systems. This combination provides proven reliability with the performance that data center administrators around the world have come to expect from Intel.

Best Choice for Server Virtualization

Virtualization changes server resource deployment and management by running multiple applications and operating systems on a single physical server.

With Intel® Virtualization Technology for connectivity (VT-c), the X550 delivers outstanding I/O performance and Quality of Service (QoS) in virtualized data centers and cloud environments. I/O virtualization advances network connectivity used in today's servers to more efficient models by providing FPP, multiple Tx/Rx queues, Tx queue rate-limiting, and on-controller QoS functionality that is useful for both virtual and non-virtual server deployments.

The X550 reduces I/O bottlenecks by providing intelligent offload of networking traffic per VM, enabling near-native performance and VM scalability. The host-based virtualization technologies include:

- VMDq for emulated path: NIC-based VM queue sorting enabling efficient hypervisor-based switching.
- SR-IOV for direct assignment: NIC-based isolation and switching for various virtual station instances enabling optimal CPU usage in virtualized environment.

Additionally, the X550 provides virtual bridging support that delivers both host-side and switch-side control and management of virtualized I/O as well as the following modes of virtualized operation:

- VEPA: IEEE 802.1Qbg support for Virtual Ethernet Port Aggregator.

- VEB: Virtual Ethernet Bridge support with Intel VT.

Networking Virtualization

Network virtualization is the next big trend in creating an agile data center. The Intel X550 family of controllers are ready to help take you to the next level.

- VXLAN and NVGRE offloads: These stateless offloads preserve application performance for overlay networks. With these offloads it is possible to distribute network traffic across CPU core.
- Preserves application performance in network virtualized environment.

Flexible Port Partitioning (FPP)

By taking advantage of the PCI-SIG* SR-IOV specification, FPP enables virtual Ethernet controllers that can be used by a Linux* host directly and/or assigned directly to virtual machines for hypervisor virtual switch bypass. FPP enables the assignment of up to 64 Linux host processes or virtual machines per port to virtual functions. An administrator can use FPP to control the partitioning of the bandwidth across multiple virtual functions. FPP can also provide balanced QoS by giving each assigned virtual function equal access to 10 Gb/s of bandwidth.

10 GbE Performance at Low Cost and Low Power

The Intel® Ethernet Controller X550 brings 10GBASE-T as a cost effective means to bring 10 GbE to embedded, workstation, and server platforms as LAN on Motherboard (LOM) or network add-in card. To reduce cost and power, the X550 controller is manufactured using a 28 nm process with an integrated MAC controller and up to two 10GBASE-T PHYs in a single-chip solution. Integration translates to lower power with reduced per-port power consumption, which can eliminate the need for active fan heatsinks.

Intel® Ethernet Controller X550
Network Connectivity

With lower power, passive heatsinks, and backwards compatibility, the X550 controller and 10GBASE-T are ready for broad deployment. The X550 controller provides bandwidth-intensive applications and virtualized data centers 10 GbE network performance with cost-effective network connectivity.

Network Manageability Interfaces

The X550 controller provides OS2BMC, SMBus and DMTF-defined Network Controller Sideband Interface (NC-SI) for BMC manageability. In addition, it introduces support for Management Component Transport Protocol (MCTP), a new DMTF standard, enabling a BMC to gather information about Intel Ethernet Converged Network Adapters that can include the data rate, link speed, and error counts.

With low power consumption, a small footprint and integrated serial PHYs, the controller is ideally suited for server blades, LOM, NIC, and mezzanine card implementations. The X550 also incorporates the manageability required by IT personnel for remote control and alerting. Communication to the Baseboard Management Controller (BMC) is available either through an on-board SMBus port or the DMTF-defined NC-SI, providing a variety of management protocols, including IPMI, BMC pass-through, OS2BMC, and MCTP.

EXTERNAL INTERFACES

INTERFACE	DESCRIPTION
Intel® Ethernet Controller X550AT	• PCIe interface v3.0 - 8.0 GT/s, 5.0 GT/s and 2.5 GT/s; support for x1 and x4, links widths (lanes)
Intel® Ethernet Controller X550AT2	• PCIe interface v3.0 - 8.0 GT/s, 5.0 GT/s and 2.5 GT/s; support for x1 and x4, links widths (lanes)
Intel® Ethernet Controller X550BT2	• PCIe interface v2.1 - 5.0 GT/s and 2.5 GT/s; support for x1, x4 and x8 links widths (lanes)
Network Interfaces	<ul style="list-style-type: none"> • IEEE 802.3 Ethernet interface for 10GBASE-T, 1000BASE-T, 100BASE-TX (IEEE 802.3an, 802.3ab, 802.3u) • NBASE-T* technology support for 2.5 gigabit and 5.0 gigabit (pre-IEEE 802.3bz)

BOM COST REDUCTION

FEATURE	BENEFIT
Single Chip Design	<ul style="list-style-type: none"> • Designed for passive heatsink thermal solutions • X550AT & X550AT2: Small packaging for easier board layout and design • X550BT2: Footprint compatible with the Intel® Ethernet Controller X540 family
Integrated Copper 10GBASE-T PHYs	• Single chip with integrated PHYs for lower power and simplified component placement

EXTERNAL FEATURES

FEATURE	BENEFIT
NBASE-T Technology Support (IEEE 802.3bz-ready)	• Provides 2.5 GbE and 5.0 GbE link speeds over CAT5e/CAT6 cabling deployments. Firmware upgradable to IEEE 802.3bz ¹
Data Center Bridging (DCB) Support: 802.1Qaz, 802.1Qbb, and 802.1p	<ul style="list-style-type: none"> • Generates congestion notifications messages that slow down offending applications that are causing congestion in the network • Delivers standard-based end-to-end congestion notification that extends lossless Ethernet beyond just a single hop
IEEE 1588 Protocol and IEEE 802.1AS Implementation	<ul style="list-style-type: none"> • Per-packet time-stamping and synchronization of time sensitive applications • Distribute common time to media devices
VLAN Support: • IEEE 801.1Q (VLAN) • IEEE 802.1ad (Double VLAN)	<ul style="list-style-type: none"> • Provide data separation and security between network traffic • Double-tagging can be useful for Internet service providers, allowing the use of VLANs internally while mixing traffic from clients that are already VLAN-tagged
Automatic Cross-over Detection Function (MDI/MDI-X)	• The PHY automatically detects which Media-Dependent Interface (MDI) is required and configures itself accordingly

I/O FEATURES FOR MULTI-CORE PROCESSOR SYSTEMS

FEATURE	BENEFIT
Intel® Flow Director	<ul style="list-style-type: none"> • An advanced traffic steering capability that 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 overhead of interrupts • Load-balancing of interrupt handling between multiple cores/CPU's • Dynamic allocation of up to 64 vectors per port
Multiple Queues: 128 Tx and Rx Per Port	<ul style="list-style-type: none"> • Queues provide QoS for virtualization, DCB, RSS, L2 Ethertype, FCoE redirections, L3/4/5-tuple filters, flow director, and TCP SYN filters • Network packet handling without waiting for buffer overflow providing efficient packet prioritizations
TCP/UDP, IPv4 Checksum Offloads (Rx/Tx/Large-send); Extended Tx Descriptors for More Offload Capabilities	<ul style="list-style-type: none"> • Improve CPU usage • Checksum and segmentation capability to new standard packet type
RSS for Windows Environment Scalable I/O for Linux Environments (IPv4, IPv6, TCP/UDP)	<ul style="list-style-type: none"> • Up to 32 flows per port • Improves the system performance related to handling of network data on multi-processor systems
IPv6 Support for IP/TCP and IP/UDP Receive Checksum Offload	<ul style="list-style-type: none"> • Improved CPU usage
Tx TCP Segmentation Offload (TSO-IPv4, IPv6)	<ul style="list-style-type: none"> • Large TCP I/O is segmented into small packets to increase throughput and reduce CPU overhead • Compatible with large-send offload
FCoE Tx/Rx CRC Offload	<ul style="list-style-type: none"> • Offloads receive FC CRC integrity check while tracking the CRC bytes and FC padding bytes
Large FC Receive	<ul style="list-style-type: none"> • Large FC receive includes two types of offloads that can save a data copy by posting the received FC payload directly to the kernel storage cache or the user application space
FCoE Transmit Segmentation Offloads	<ul style="list-style-type: none"> • Enables the FCoE software to initiate a transmission of multiple FCoE packets up to a complete FC sequence with a single header in host memory (single instruction)
FCoE Coalescing and Direct Data Placement	<ul style="list-style-type: none"> • Hardware can provide DDP offload for up to 2048 concurrent outstanding FC read or write exchanges
Support for Packets Up To 15.5 KB (Jumbo Frames)	<ul style="list-style-type: none"> • Enables higher and better throughput of data
DMA Coalescing	<ul style="list-style-type: none"> • Reduces platform power consumption by coalescing, aligning, and synchronizing DMA • Enables synchronizing power activity and power management of memory, CPU and RC internal circuitry
Flow Director Filters: <ul style="list-style-type: none"> • Up to 32 KB - Two Signature Filters • Up to 8 Kb - Two Perfect-match Filters 	<ul style="list-style-type: none"> • The flow director filters identify specific flows or sets of flows and routes them to specific queues. These filters are an expansion of the L3/L4 5-tuple filters that provide up to additional 32 KB filters
Receive Side Coalescing (RSC)	<ul style="list-style-type: none"> • Merge multiple received frames from the same TCP/IP connection into a single structure

VIRTUALIZATION FEATURES

FEATURE	BENEFIT
Multi-mode I/O Virtualization Operations	<ul style="list-style-type: none"> • Supports two modes of operation of virtualized environments: <ul style="list-style-type: none"> • Direct assignment of part of the port resources to different guest operating systems using the PCI SIG SR-IOV standard (also known as native mode or pass-through mode) • Central management of the networking resources by hypervisor (also known as software switch acceleration mode) • A hybrid model, where some of the VMs are assigned a dedicated share of the port and the rest are serviced by an hypervisor is also supported
VXLAN	<ul style="list-style-type: none"> • A framework for overlaying virtualized layer 2 networks over layer 3 networks. VXLAN enables users to create a logical network for your virtual machines across different networks
NVGRE	<ul style="list-style-type: none"> • Network Virtualization using Generic Routing Encapsulation. The encapsulation of an Ethernet Layer 2 Frame in IP that enables the creation of virtualized Layer 2 subnets that can span physical Layer 3 IP networks
Virtual Machine Device Queues (VMDq)	<ul style="list-style-type: none"> • Offloads data sorting from the hypervisor to silicon, improving data throughput and CPU usage • QoS feature for Tx data by providing round-robin servicing and preventing head-of-line blocking • Sorting based on MAC addresses and VLAN tags
Next Generation VMDq	<ul style="list-style-type: none"> • Enhanced QoS feature by providing weighted round-robin servicing for the Tx data • Provides loopback functionality; data transfers between the virtual machines within the same physical server do not go out to the wire and back in, improving throughput and CPU usage • Supports replication of multicast and broadcast data
64 Transmit (Tx) and Recive (Rx) Queue Pairs Per Port	<ul style="list-style-type: none"> • Supports VMware* NetQueue and Microsoft* VMQ • MAC/VLAN filtering for pool selection and either DCB or RSS for the queue in pool selection
Flexible Port Partitioning: 64 Virutal Functions Per Port	<ul style="list-style-type: none"> • Virtual Functions (VFs) appear as Ethernet Controllers in Linux OSes that can be assigned to VMs, Kernel processes or teamed using the Linux* Bonding Drivers
Support for PCI-SIG SR-IOV Specification	<ul style="list-style-type: none"> • Up to 64 virtual functions per port
IEEE 802.1Q Virtual Local Area Network (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 • Filtering packets belonging to certain VLANs

REMOTE BOOT OPTIONS

FEATURE	BENEFIT
Preboot eXecution Environment (PXE) Flash Interface Support	<ul style="list-style-type: none"> • Enables system boot up via the EFI (32-bit and 64-bit) • Flash interface for PXE 2.1 option ROM
Intel® Ethernet FCoE Boot	<ul style="list-style-type: none"> • Enables system boot up via FCoE
Intel® Ethernet iSCSI Remote Boot	<ul style="list-style-type: none"> • Enables system boot up via iSCSI
Intel Boot Agent software: Linux boot via PXE or BOOTP, Windows* Deployment Services, or UEFI	<ul style="list-style-type: none"> • Allows networked computer to boot using a program code image supplied by a remote server • Complies with the Pre-boot eXecution Environment (PXE) Version 2.1 Specification

SECURITY AND POWER MANAGEMENT

FEATURE	BENEFIT
Receive Packet Filtering	<ul style="list-style-type: none"> Determine which of the incoming packets are allowed to pass to the local machine based on L2, VLAN, or management policies
Integrated IPsec Security Engines for Offloads of up to 1024 Security Associations (SA) for Each Tx and Rx	<ul style="list-style-type: none"> Offloads handle a certain amount of the total number of IPsec flows on the controller in hardware
IEEE 802.3az Energy Efficient Ethernet (EEE) ¹	<ul style="list-style-type: none"> Power consumption by the PHY is reduced by during period of low link utilization to help reduce power consumption
Secure Flexible Firmware Architecture	<ul style="list-style-type: none"> Secure NVM Update that protects the flash from external unauthorized software programming Supports dynamic firmware updating that enables firmware updates without the need for a system reboot
Four Software Definable Pins (SDP) Per Port	<ul style="list-style-type: none"> Software-defined pins (SDP pins) per port that can be used for miscellaneous hardware or software-controlled purposes
LAN Disable Function	<ul style="list-style-type: none"> Option to disable the LAN port and/or PCIe function. Disabling just the PCIe function but keeping the LAN port that resides on it fully active (for manageability purposes and BMC pass-through traffic)
Anti-spoofing for MAC and VLANs	<ul style="list-style-type: none"> Capability insures that a VM always uses a source Ethernet VLAN or MAC address on the transmit path that is part of the set of VLAN tags and Ethernet MAC addresses defined on the Rx path
Full Wake Up Support: <ul style="list-style-type: none"> Advanced Power Management (APM) Support (Formerly Wake on LAN) Advanced Configuration and Power Interface (ACPI) Specification v2.0c Magic Packet* Wake-up Enable with Unique MAC Address 	<ul style="list-style-type: none"> APM - Designed to receive a broadcast or unicast packet with an explicit data pattern (Magic Packet) and assert a signal to wake up the system APCI - PCIe power management based wake up that can generate system wake-up events from a number of sources
Low Power Operation and Power Management	<ul style="list-style-type: none"> Incorporates numerous features to maintain the lowest power possible including PCI Express Link and Network Interface power management
Low Power Link Up - Link Speed Control	<ul style="list-style-type: none"> Enables a link to come up at the lowest possible speed in cases where power is more important than performance

MECHANICAL AND THERMAL

INTERFACE	DESCRIPTION
Intel® Ethernet Controller X550AT	<ul style="list-style-type: none"> 17 mm x 17 mm FC-BGA package
Intel® Ethernet Controller X550AT2	<ul style="list-style-type: none"> 17 mm x 17 mm FC-BGA package
Intel® Ethernet Controller X550BT2	<ul style="list-style-type: none"> 25 mm x 25 mm FC-BGA package
Power:	
<ul style="list-style-type: none"> X550AT X550AT2 X550BT2 	<ul style="list-style-type: none"> 9 W max, 7 W typical (10 Gb/s active) 11 W max, 9 W typical (2 x 10 Gb/s active) 11 W max, 9 W typical (2 x 10 Gb/s active)

MANAGEABILITY FEATURES

FEATURE	BENEFIT
DMTF Network Controller Sideband Interface (NC-SI Pass-through)	<ul style="list-style-type: none"> • Supports pass-through traffic between BMC and controller's LAN functions
Intel® System Management Bus (SMBus) Pass-through	<ul style="list-style-type: none"> • Enables BMC to configure the controller's filters and management related capabilities
Management Component Transport Protocol (MCTP) over SMBus or PCIe	<ul style="list-style-type: none"> • BMC communication between add-in devices within the platform • Allow reporting and controlling of all the information exposed in a LOM device via NC-SI, in NIC devices via MCTP over SMBus or PCIe
<ul style="list-style-type: none"> • OS2MBC Traffic Support: <ul style="list-style-type: none"> • Host-based Application-to-BMC Network Communication Patch • Private OS2BMC Traffic Flow 	<ul style="list-style-type: none"> • Transmission and reception of traffic internally to communicate between the OS and local BMC • Filtering method that enables server management software to communicate with a management controller via standard networking protocols such as TCP/IP instead of a chipset-specific interface • BMC might have its own private connection to the network controller and network flows are blocked
DMTF MCTP Protocol Over SMBus	<ul style="list-style-type: none"> • Enables reporting and controlling information via NC-SI using the MCTP protocol over SMBus
Firmware-based Thermal Management	<ul style="list-style-type: none"> • Can be programmed via the BMC to initiate thermal actions and report thermal occurrences
IEEE 802.3 Management Data Input/Output Interface (MDIO Interface or MII Management Interface)	<ul style="list-style-type: none"> • Enables the MAC and software to monitor and control the state of the PHY
MAC/PHY Control Status	<ul style="list-style-type: none"> • Enhanced control capabilities through PHY reset, link status, duplex indication, and MAC Dx power state
Watchdog Timer	<ul style="list-style-type: none"> • The MAC and each PHY support a watchdog timer to detect a stuck microcontroller
Advanced Error Reporting (AER)	<ul style="list-style-type: none"> • Messaging support to communicate multiple types/severity of errors
Controller Memory Integrity Protection	<ul style="list-style-type: none"> • Main internal memories are protected by Error Correcting Code (ECC) or parity bits
Vital Product Data (VPD) Support	<ul style="list-style-type: none"> • Support for VPD memory area
Flexible MAC Address	<ul style="list-style-type: none"> • MAC address used by a port can be replaced with a temporary MAC address in a way that is transparent to the software layer

INTEL® ETHERNET CONTROLLER X550 PRODUCT CODES

CONFIGURATION	PORTS	PACKAGE	PRODUCT CODE	PRODUCTION MM #	PRODUCT S SPEC	MEDIA TYPE
Intel® Ethernet Controller X550-AT	1	17 x 17 mm	ELX550AT	945964	S LLFT	Tray
Intel® Ethernet Controller X550-AT	1	17 x 17 mm	ELX550AT	945983	S LLFU	Tape & Reel
Intel® Ethernet Controller X550-AT2	2	17 x 17 mm	ELX550AT2	943736	S LL2E	Tray
Intel® Ethernet Controller X550-AT2	2	17 x 17 mm	ELX550AT2	943743	S LL2F	Tape & Reel
Intel® Ethernet Controller X550-BT2	2	25 x 25 mm	ELX550BT2	943742	S LL2G	Tray
Intel® Ethernet Controller X550-BT2	2	25 x 25 mm	ELX550BT2	943744	S LL2H	Tape & Reel

For more information on Intel® Ethernet Controller X550, visit www.intel.com/ethernet



¹Feature to be enabled in a post-launch firmware release.

No license (express or implied, by sstopel 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.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting www.intel.com/design/literature.htm.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

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