

Intel[®] Ethernet Controller X550

Specification Update

Networking Division (ND)

January 2016

Revision 2.0
333717-001



Revision History

Revision	Date	Comments
2.0	January 8, 2016	Initial release (Intel Public).



1. Introduction

This document applies to the Intel® Ethernet Controller X550 (X550).

This document is an update to a published specification, the *Intel® Ethernet Controller X550 Datasheet*. It is intended for use by system manufacturers and software developers. All product documents are subject to frequent revision and new order numbers might apply. New documents might be added. Be sure you have the latest information before finalizing your design.

References to PCIe Express* (PCIe*) in this document refer to PCIe v3.0 (2.5GT/s, 5GT/s, and 8GT/s).

1.1 Product Code and Device Identification

Product Code: ELX550

The following tables and drawings describe the various identifying markings on each device package:

Table 1-1 Markings

Device	Stepping	Top Marking	S-Specification ¹	Description
X550-AT	B0	ELX550AT	S LLFT	17x17 mm package - Single port
			S LLFU	
X550-AT2	B0	ELX550AT2	S LL2E	17x17 mm package - Dual port
			S LL2F	
X550-BT2	B0	ELX550BT2	S LL2G	25x25 mm package - Dual port
			S LL2H	

1. For Tray, Tape, Reel data, see [Table 1-3](#).

Table 1-2 Device ID

Device ID Code	Device ID	Vendor ID	Revision ID
Intel® Ethernet Controller X550-AT	0x15D1	0x8086	0x1
Intel® Ethernet Controller X550-AT2	0x1563	0x8086	0x1
Intel® Ethernet Controller X550-BT2	0x1563	0x8086	0x1

Table 1-3 MM Numbers

Product	Tray MM#	Tape and Reel MM#	Reserved
ELX550AT	945964	945983	
ELX550AT2	943736	943743	
ELX550BT2	943742	943744	

1.2 Marking Diagrams

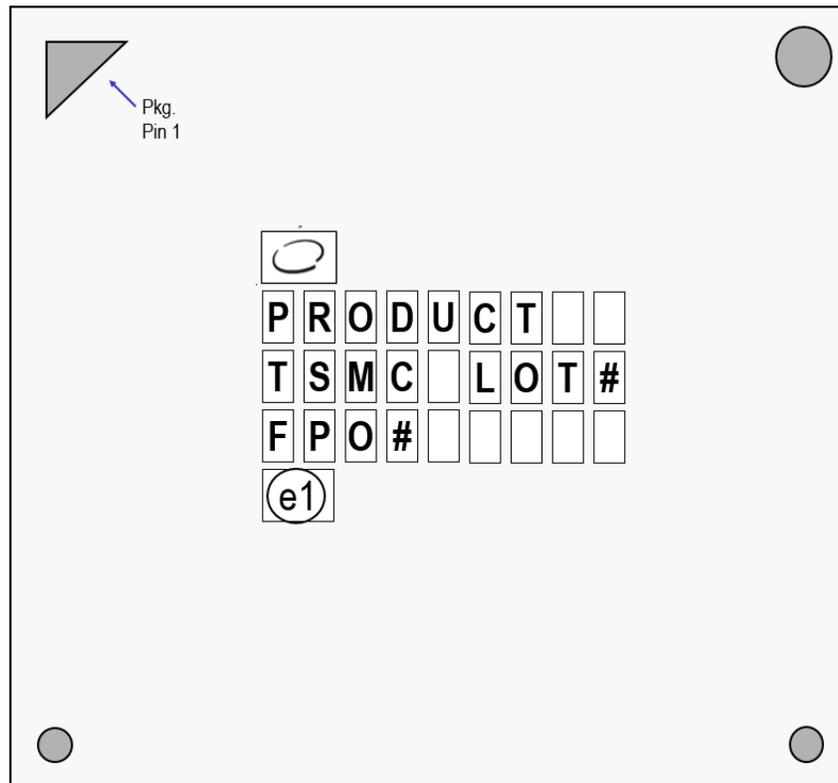


Figure 1-1 X550-AT2

- LINE1: Swirl Logo
- LINE2: Product code
- LINE3: TSMC LOT#
- LINE4: FPO# Trace code
- LINE5: Pb-free mark

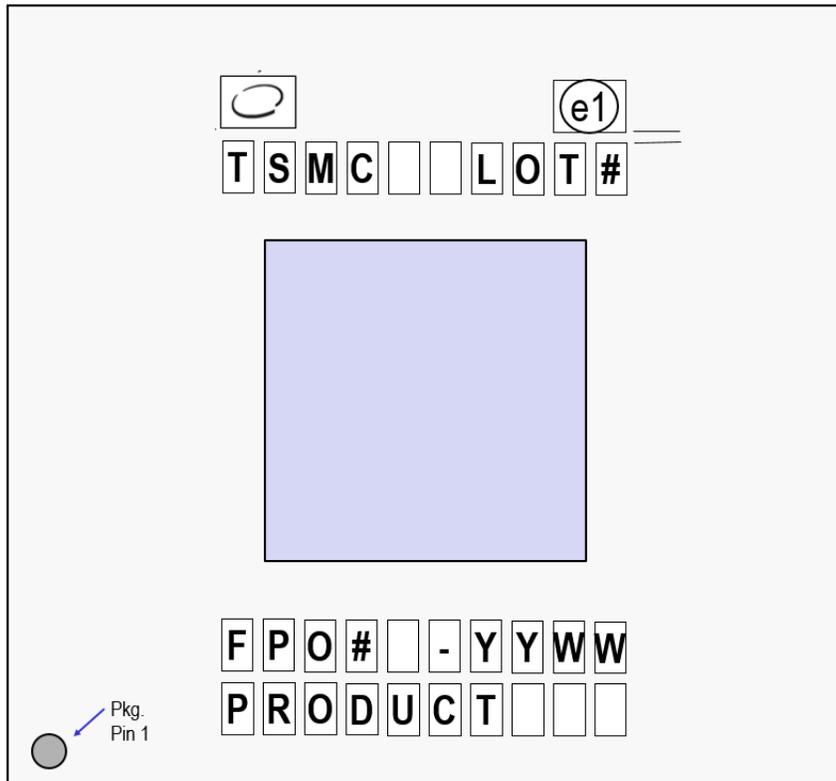


Figure 1-2 X550-BT2

- GRP1LINE1: Swirl Logo, Pb-Free
- GRP1 LINE2: TSMC Lot#
- GRP2LINE1: FPO# YYWW
- GRP2LINE2: Product Code



1.3 Nomenclature Used in This Document

This document uses specific terms, codes, and abbreviations to describe changes, errata, sightings and/or clarifications that apply to silicon/steppings. See [Table 1-4](#) for a description.

Table 1-4 Nomenclature

Name	Description
Specification Clarifications	Greater detail or further highlights concerning a specification’s impact to a complex design situation. These clarifications will be incorporated in the next release of the specifications.
Specification Changes	Modifications to the current published specifications. These changes will be incorporated in the next release of the specifications.
Errata	Design defects or errors. Errata might cause device behavior to deviate from published specifications. Hardware and software designed to be used with any given stepping must assume that all errata documented for that stepping are present on all devices.
Software Clarifications	Applies to Intel drivers, EEPROM loads.
Documentation Changes	Typos, errors, or omissions from the current published specifications. These changes will be incorporated in the next release of the specifications.
A0, B0, etc.	Stepping to which the status applies.
Doc	Document change or update that will be implemented.
Fixed	This erratum has been fixed.
Fix Planned	This erratum is intended to be fixed in a future stepping of the component.
NoFix	There are no plans to fix this erratum.
Fixed in NVM	This erratum has been fixed in NVM X.XX.
Fix Planned in NVM	This erratum is intended to be fixed in a future NVM version.
Eval	Plans to fix this erratum are under evaluation.



2. Hardware Clarifications, Changes, Updates and Errata

See Section 1.3 for an explanation of terms, codes, and abbreviations.

Table 2-1 Summary of Specification Clarifications

Specification Clarification	Status
1. PCIe Re-timers Might Cause Replay Timer Timeout Correctable Errors	N/A

Table 2-2 Summary of Specification Changes

Specification Change	Status
None.	N/A

Table 2-3 Summary of Documentation Updates

Documentation Update	Status
None	N/A

Table 2-4 Summary of Errata; Errata Include Steppings

Erratum	Status
1. AUX Power Detected Register Bit is Always Zero	B0=Yes; NoFix
2. Common Clock Configuration Register Bit is Shared Between Functions	B0=Yes; NoFix
3. PCIe: Read Data from Header Logs 3 & 4 of VDM TLP Are Swapped	B0=Yes; NoFix
4. GbE: Software Reset Causes Link Status Change Interrupt	B0=Yes; NoFix
5. Header Logging in VF Might Be Wrong	B0=Yes; NoFix
6. SMBus Async Notification Timeout Does Not Work	B0=Yes; NoFix
7. ASPM Optional Compliance Bit in Config Space is 0b Instead of 1b	B0=Yes; NoFix
8. SMBus Alert TO Does Not Work if BMC Reads the Status without ARA	B0=Yes; NoFix
9. Transaction Pending Bit Not Functional	B0=Yes; NoFix
10. Wrong Revision ID Value Reported to MNG by PCIe	B0=Yes; NoFix
11. SMBus - BMC to LAN Traffic Might Be Sent to Wrong Port After Sending First Fragment After First Fragment Condition	B0=Yes; NoFix



Table 2-4 Summary of Errata; Errata Include Steppings (Continued)

Erratum	Status
12. No Length Error on VLAN Packets with Bad Type/Length Field	B0=Yes; NoFix
13. GPRC and GORCL/H Also Count Missed Packets	B0=Yes; NoFix
14. Cause of an Interrupt Might Never be Cleared	B0=Yes; NoFix
15. The X550 Does Not Meet the Timing Requirements for PAUSE Operation in 1 GbE Speed	B0=Yes; NoFix
16. The X550 Does Not Meet the Timing Requirements for PAUSE Operation in 100 Mb/s Speed	B0=Yes; NoFix
17. 100BASE-TX Transmitter Conformance ANSI X3.263 - Amplitude Test Failure	B0=Yes; NoFix
18. 1GBASE-T Conformance IEEE 802.3 - Differential Output Templates Test Failure	B0=Yes; Eval
19. 10GBASE-T Conformance IEEE 802.3 - Power Spectral Density Test Failure	B0=Yes; Eval
20. Energy Efficient Ethernet 10/1 GbE	B0=Yes; Fix Planned in NVM
21. Internal Thermal Sense Circuit Does not Generate a Thermal Sensor Event Interrupt	B0=Yes; Fix Planned in NVM

2.1 Specification Clarifications

1. PCIe Re-timers Might Cause Replay Timer Timeout Correctable Errors

The addition of PCIe re-timers add to the total channel latency. According to PCI-SIG ECN extension devices, latency is defined as “the time from when the last bit of a Symbol is received at the input pins of one Pseudo Port to when the equivalent bit is transmitted on the output pins of the other Pseudo Port”. The ECN allows for a maximum of 64 symbol x latency per PCIe re-timer for 8 GT/s speed.

The PCIe ACK/NACK round trip delay is incremented according to the number of re-timers used in Tx/Rx lanes. The extra delay added by a re-timer might cause the X550 Replay_Timer to expire, causing replay timer timeout correctable errors. The X550 design does not take into consideration the extension devices ECN.

If a design must include re-timers, and if Replay_Timer timeout correctable errors are seen, please contact your Intel representative for support.

2.2 Specification Changes

None.

2.3 Documentation Updates

None.



2.4 Errata

1. AUX Power Detected Register Bit is Always Zero

Problem:

Device status register (0xAA) bit 4 should reflect the presence of AUX_PWR. This bit is not functioning and stuck at 0.

Implication:

Compliance failure. If there is a need to read the AUX power state it can be read from register Power Management Capabilities - PMC Register (0x42), field *PME_Support*.

Workaround:

None.

Status: B0=Yes; NoFix

2. Common Clock Configuration Register Bit is Shared Between Functions

Problem:

Common Clock Configuration Register bit — Link Control Register (0xB0) bit 6 should be RW according to PCI SIG and independent between the two functions. The current implementation of this bit is shared between the two functions since they use the same clock source, and hence the bit does not behave as a simple RW.

Implication:

Compliance issue only. No functional implication.

Workaround:

None.

Status: B0=Yes; NoFix

3. PCIe: Read Data from Header Logs 3 & 4 of VDM TLP Are Swapped

Problem:

After command parity occurs on a message of type VDM, all the advanced error reporting registers are read, and reading the header log 3 register returns the data that was supposed to return from header log 4 and vice versa. This is a PCIe spec violation.

Implication:

Compliance issue. Header log is not used by standard operating systems, so no real issue.



Workaround:

Read words in reverse order.

Status: B0=Yes; NoFix

4. GbE: Software Reset Causes Link Status Change Interrupt

Problem:

Setting CTRL.RST bit (software reset) directly or Force TCO BMC command causes interrupt of Link Status Change. EICR.LSC is set and Link Status change Asynch notification is sent to BMC, but no actual link reset occurred.

Implication:

Spurious Link change interrupt registered by software.

Workaround:

Simple software workaround since software creates the reset.

Status: B0=Yes; NoFix

5. Header Logging in VF Might Be Wrong

Problem:

There are places for two error header logs per PF for VFs. These places are supposed to be released when the corresponding error is cleared by the host (by writing 1 in the uncorrectable status reg of the VF), but they are not.

Implication:

Header logging might be wrong.

Workaround:

None.

Status: B0=Yes; NoFix

6. SMBus Async Notification Timeout Does Not Work

Problem:

The timeout for the SMBus asynchronous notification message does not work.

Implication:

If the BMC does not ACK the Async notification, or does not read the status, the DUT continues to send endless SMBus transactions regardless of the NVM Notification TO value.



Workaround:

None.

Status: B0=Yes; NoFix

7. ASPM Optional Compliance Bit in Config Space is 0b Instead of 1b

Problem:

ASPM Optional Compliance bit in config space is 0b. It should be 1b for PCIe 3.0 compliance.

Implication:

Compliance issue. No functional implication.

Workaround:

None.

Status: B0=Yes; NoFix

8. SMBus Alert TO Does Not Work if BMC Reads the Status without ARA

Problem:

Alert TO does not work if two ports have notification for BMC simultaneously, and the BMC reads it (using receive packet command (0xC0)) only from one port without ARA cycle.

Example for scenario is following firmware reset - All the ports want to send status to BMC simultaneously. If the BMC reads the status from only one port, the second/other port(s) continues to pull the alert forever, regardless the notification TO, until the BMC reads the status from all the ports.

Implication:

All BMCs are sending ARA. so no issue.

Workaround:

None.

Status: B0=Yes; NoFix

9. Transaction Pending Bit Not Functional

Problem:

Transaction pending logic does not reflect the right status and might stay set longer than needed.

Implication:

Transaction pending bit in PCIe config space might not reflect the actual state of transactions pending.



Workaround:

None.

Status: B0=Yes; NoFix

10. Wrong Revision ID Value Reported to MNG by PCIe

Problem:

A0 and B0 will have the same revid on SMBus.

Implication:

BMC cannot differentiate between A0 and B0 silicon via the Get UDID command.

Workaround:

The actual revID can be read via the Get Controller Information data command.

Status: B0=Yes; NoFix

11. SMBus - BMC to LAN Traffic Might Be Sent to Wrong Port After Sending First Fragment After First Fragment Condition

Problem:

When BMC is transmitting to both ports, sending two consecutive first fragments causes the internal FIFOs to get out of sync and packets might be sent to the wrong port. Note that this is an error condition, as the first packet is aborted.

Implication:

Traffic might be routed to wrong port.

Workaround:

If such a condition is detected, reset the NIC firmware using the TCO Reset command.

Status: B0=Yes; NoFix

12. No Length Error on VLAN Packets with Bad Type/Length Field

Problem:

The X550 does not assert length error for VLAN packets that have a bad Type/Length field in the MAC header.

Implication:

There is no impact on system-level performance. The packets are posted to the host as with any other packets.



Workaround:

None.

Status: B0=Yes; NoFix

13. GPRC and GORCL/H Also Count Missed Packets

Problem:

GPRC (Good Packets Received Count) and GORCL/H (Good Octets Received Count) count missed packets and missed packets bytes.

Implication:

None.

Workaround:

Statistics are available indirectly for these registers. This workaround is included in Intel drivers.

- For GPRC — Subtract MPC (Missed Packet Count) from GPRC. Alternatively, use QPRC.
- For GORCL/H — Use QBRCL/H (Quad Bytes Received).

Status: B0=Yes; NoFix

14. Cause of an Interrupt Might Never be Cleared

Problem:

If the cause of an interrupt is set by the Extended Interrupt Cause Set (EICS) register writing just before the interrupt line is set, it might not be cleared. This means that there might be a deadlock that prevents the interrupt line from rising.

This erratum only occurs when all three modes referenced are used at the same time: non-PBA mode, Auto Clear (of the cause), No Auto Mask.

PBA is Pending Bit Array mode. During this mode the device is able to capture additional interrupts during the interval between initial interrupt and driver access to the device.

Implication:

The X550 stops issuing interrupts.

Workaround:

When operating using the above configurations, software should manually clear the cause by writing a 1b to the specific bit in the relevant EICR/VTEICR0-63 register (after the interrupt occurs and the EICS was written). This workaround is included in Intel drivers.

Status: B0=Yes; NoFix



15. The X550 Does Not Meet the Timing Requirements for PAUSE Operation in 1 GbE Speed

Problem:

In 1 GbE speed, the X550 responds to a received pause frame after a longer time than defined in the IEEE 802.3 specification.

Implication:

Specification conformance. The response gap is small.

Workaround:

None.

Status: B0=Yes; NoFix

16. The X550 Does Not Meet the Timing Requirements for PAUSE Operation in 100 Mb/s Speed

Problem:

In 100 Mb/s speed, the X550 responds to a received pause frame after a longer time than defined in the IEEE 802.3 specification.

Implication:

Specification conformance. No system impact with low traffic.

Workaround:

None.

Status: B0=Yes; NoFix

17. 100BASE-TX Transmitter Conformance ANSI X3.263 - Amplitude Test Failure

Problem:

100BASE-TX amplitude test might fail during conformance testing.

Implication:

There is no expected performance impact. Conformance test impact only.

Workaround:

None.

Status: B0=Yes; NoFix



18. 1GBASE-T Conformance IEEE 802.3 - Differential Output Templates Test Failure

Problem:

1GBASE-T differential output template test might fail during conformance testing.

Implication:

There is no expected performance impact. Conformance test impact only.

Workaround:

None.

Status: B0=Yes; Eval

19. 10GBASE-T Conformance IEEE 802.3 - Power Spectral Density Test Failure

Problem:

10GBASE-T power spectral density test might fail during conformance testing.

Implication:

There is no expected performance impact. Conformance test impact only.

Workaround:

None.

Status: B0=Yes; Eval

20. Energy Efficient Ethernet 10/1 GbE

Problem:

Energy Efficient Ethernet is not enabled for 10 GbE or 1 GbE modes.

Implication:

The X550 does not negotiate to or transition to EEE mode with a EEE-capable link partner.

Workaround:

None.

Status: B0=Yes; Fix Planned in NVM



21. Internal Thermal Sense Circuit Does not Generate a Thermal Sensor Event Interrupt

Problem:

Internal Thermal Sensor does not generate a high temp threshold/shutdown interrupt required by the drivers to log an over-temperature shutdown event.

Implication:

If the X550 internal temperature sensor detects a temperature condition greater than the programmed threshold temperature, the PHY is shut down as required, and the link is dropped. This reduces device power consumption and is correct PHY behavior. However, because no interrupt was generated, software does not log the cause of the event. To recover, a power-on reset is required to reinitialize the X550 and re-establish link. Note that a thermal event is not expected in a system that follows the thermal design rules.

Workaround:

None.

Status: B0=Yes; Fix Planned in NVM



3. Software Clarifications

Table 3-1 Summary of Software Clarifications

Software Clarification	Status
1. While in TCP Segmentation Offload, Each Buffer is Limited to 64 KB	N/A
2. Serial Interfaces Programmed by Bit Banging	N/A
3. Identity Network Adapter Port by Blinking LED	N/A
4. PF/VF Drivers Should Configure Registers That Are Not Reset by VFLR	N/A

1. While in TCP Segmentation Offload, Each Buffer is Limited to 64 KB

Problem Description:

The X550 supports 256 KB TCP packets. However, each buffer is limited to 64 KB since the data length field in the descriptor is only 16 bits. This restriction can complicate things for the driver if the operating system passes down a scatter/gather element greater than 64 KB in length. This issue can be avoided by limiting the offload size to 64 KB.

Investigation has concluded that the increase in data transfer size does not provide any noticeable improvements in LAN performance. As a result, Intel network software drivers limit the data transfer in all drivers to 64 KB.

Note: Linux operating systems only supports 64 KB data transfers.

2. Serial Interfaces Programmed by Bit Banging

Problem Description:

When bit banging on a serial interface (such as SPI, I²C, or MDIO), it is often necessary to perform consecutive register writes with a minimum delay between them. However, simply inserting a software delay between the writes can be unreliable due to hardware delays on the CPU and PCIe interfaces. The delay at the final hardware interface might be less than intended if the first write is delayed by hardware more than the section write. To prevent such problems, a register read should be inserted between the first register write and the software delay. For example: write, read, software delay, write.



3. Identity Network Adapter Port by Blinking LED

Problem Description:

Intel device drivers and supported tools include a feature that provides network adapter port identification by blinking LED2. This feature assumes that LED2 is connected as the Link/Activity LED as recommended in the reference schematics.

4. PF/VF Drivers Should Configure Registers That Are Not Reset by VFLR

Problem Description:

The following registers are not reset by VFLR and need to be configured by PF or VF in case of a change to a new configuration (such as VF OS transition):

VFRDH/T, VFTDH/T, VFPSRTYPE, VFSRRCTL, VFRXDCTL, VFTXDCTL, VFTDWBAL/H, VFDCA_RXCTRL, VFDCA_TXCTRL, and VFMBMEM.



NOTE: *This page intentionally left blank.*



LEGAL

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 might contain defects or errors which might cause deviations from published specifications.

Copies of documents which have an order number and are referenced in this document might 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 might be claimed as the property of others.

© 2016 Intel Corporation.