



# Intel<sup>®</sup> Ethernet Controller X550

Specification Update

---

Ethernet Products Group (EPG)

*April 2024*

Revision 3.0  
333717-010

## Revision History

Revision	Date	Comments
3.0	April 22, 2024	<b>Errata added or updated:</b> 32. IEEE Fast Retrain (FR) Mode Failure (Added)
2.9	September 25, 2020	<b>Errata added or updated:</b> <ul style="list-style-type: none"> <li>31. Spurious Uncorrectable ECC Error Detection (Added)</li> </ul>
2.8	July 27, 2020	<b>Errata added or updated:</b> <ul style="list-style-type: none"> <li>29. MCTP over PCIe Communication Stops Working if One LAN Port is Disabled and Swapped PFs Configured (Updated)</li> </ul>
2.7	June 18, 2020	<b>Specification Clarifications added or updated</b> <ul style="list-style-type: none"> <li>6. Sageville Supports Only Auto-Negotiated Link at 100Mb (Added)</li> </ul>
2.6	December 17, 2019	<b>Errata added or updated:</b> <ul style="list-style-type: none"> <li>30. Device ID and MAC Address Cannot be Changed (Added)</li> </ul>
2.5	July 9, 2019	<b>Errata added or updated:</b> <ul style="list-style-type: none"> <li>29. MCTP over PCIe Communication Stops Working if One LAN Port is Disabled and Swapped PFs Configured (Added)</li> </ul>
2.4	November 16, 2018	<b>Errata added or updated:</b> <ul style="list-style-type: none"> <li>19. 10GBASE-T Conformance IEEE 802.3 - Power Spectral Density Test Failure (Updated)</li> <li>24. 10GBASE-T Conformance 55.5.3.2 – Transmitter Linearity (Updated)</li> <li>28. PCIe Phase 2 Fails to Timeout Under Certain Channel Conditions (Added)</li> </ul>
2.3	July 21, 2017	<b>Specification Clarifications added or updated</b> <ul style="list-style-type: none"> <li>3. SAN MAC Address (Added)</li> <li>4. Energy Efficient Ethernet (EEE) Support (Added)</li> <li>5. MAC CRC Errors Resulting from Negotiated Fast Retrain Events (Added)</li> </ul> <b>Specification Changes added or updated</b> <ul style="list-style-type: none"> <li>2. LANx_DIS_N Pins are Input Only During LPG and PERST (Added)</li> <li>3. Enable All PHYs in D3 (Added)</li> </ul> <b>Errata added or updated:</b> <ul style="list-style-type: none"> <li>20. Energy Efficient Ethernet 10/1 GbE (Updated)</li> <li>25. PHY Registers Misconfiguration (Added)</li> <li>26. VFLR on the Fly (Added)</li> <li>27. SDP[1] Pin is Input During PERST (Added)</li> </ul>
2.2	August 2, 2016	<b>Errata added or updated:</b> <ul style="list-style-type: none"> <li>18. 1GBASE-T Conformance IEEE 802.3 - Differential Output Templates Test Failure (Updated)</li> <li>24. 10GBASE-T Conformance 55.5.3.2 – Transmitter Linearity (Added)</li> </ul> <b>Miscellaneous Updates</b> <ul style="list-style-type: none"> <li>Minor formatting changes to Table 1-1, "Markings" and Table 1-3, "MM Numbers".</li> </ul>
2.1	May 10, 2016	<b>Specification Clarifications added or updated</b> <ul style="list-style-type: none"> <li>2. NBASE-T Speed Advertisement (Added)</li> </ul> <b>Specification Changes added or updated</b> <ul style="list-style-type: none"> <li>1. Thermal Sensor Indication (Added)</li> </ul> <b>Errata added or updated:</b> <ul style="list-style-type: none"> <li>21. Internal Thermal Sense Circuit Does not Generate a Thermal Sensor Event Interrupt (Updated)</li> <li>22. PCIe Interrupt Status Bit (Added)</li> <li>23. SMBus Interface Pins Pulled Low when No Power (Added)</li> </ul>
2.0	January 8, 2016	Initial release (Intel Public).

**NOTE:**        ***This page intentionally left blank.***



# 1. Introduction

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

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).

For more information on supported features, see the *Intel® Ethernet Controller X550 Feature Support Matrix*. This document is updated periodically. Please ensure that you have the latest version.

## 1.1 Product Code and Device Identification

**Product Codes:** 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	Description
X550-AT	B0	ELX550AT	S LLFT <sup>1</sup>	17x17 mm package - Single port
			S LLFU <sup>2</sup>	
X550-AT2	B0	ELX550AT2	S LL2E <sup>1</sup>	17x17 mm package - Dual port
			S LL2F <sup>2</sup>	
X550-BT2	B0	ELX550BT2	S LL2G <sup>1</sup>	25x25 mm package - Dual port
			S LL2H <sup>2</sup>	

1. Tray.
2. Tape and Reel.

**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**

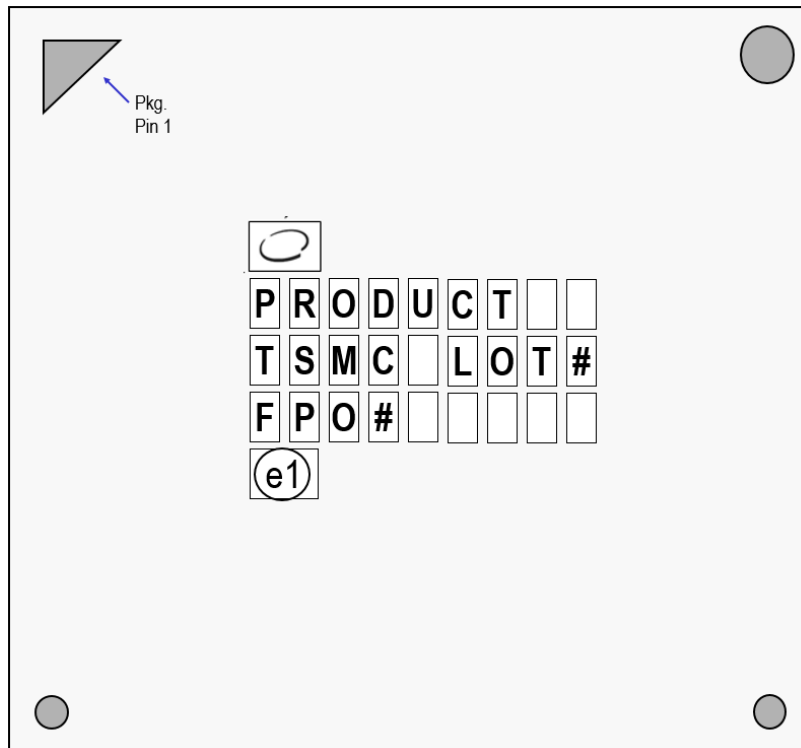
Product	S-Specification	Tray MM#	Tape and Reel MM#
ELX550AT	S LLFT <sup>1</sup>	945964	---
	S LLFU <sup>2</sup>	---	945983
ELX550AT2	S LL2E <sup>1</sup>	943736	---
	S LL2F <sup>2</sup>	---	943743

**Table 1-3. MM Numbers**

Product	S-Specification	Tray MM#	Tape and Reel MM#
ELX550BT2	S LL2G <sup>1</sup>	943742	---
	S LL2H <sup>2</sup>	---	943744

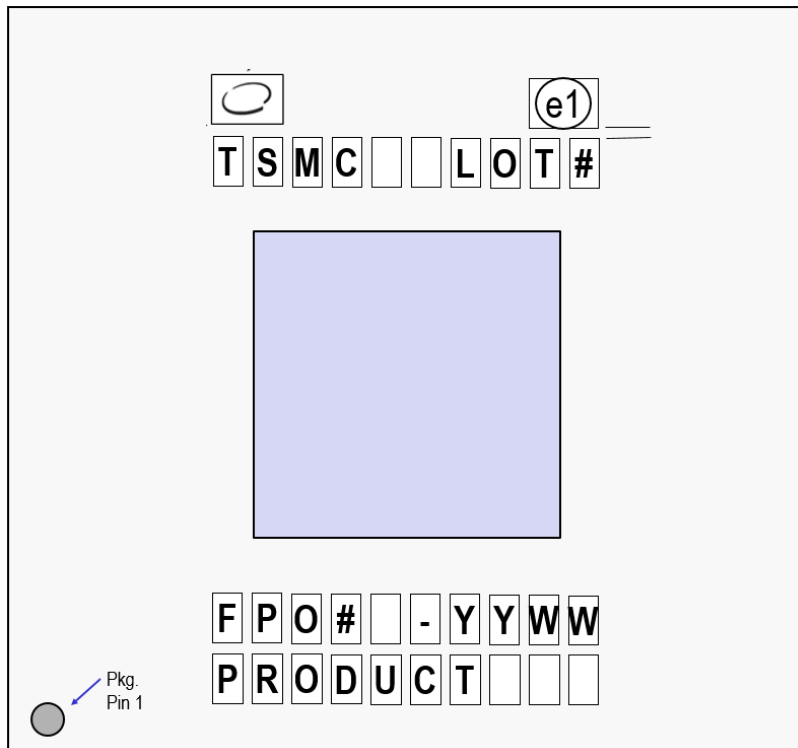
1. Tray.
2. Tape and Reel.

## 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, 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.

**Table 1-4. Nomenclature**

Name	Description
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
2. NBASE-T Speed Advertisement	N/A
3. SAN MAC Address	N/A
4. Energy Efficient Ethernet (EEE) Support	N/A
5. MAC CRC Errors Resulting from Negotiated Fast Retrain Events	N/A
6. Sageville Supports Only Auto-Negotiated Link at 100Mb	N/A

**Table 2-2. Summary of Specification Changes**

Specification Change	Status
1. Thermal Sensor Indication	N/A
2. LANx_DIS_N Pins are Input Only During LPG and PERST	N/A
3. Enable All PHYs in D3	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

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

Erratum	Status
11. SMBus - BMC to LAN Traffic Might be Sent to Wrong Port After Sending First Fragment After First Fragment Condition	B0=Yes; NoFix
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. Sageville Does Not Meet the Timing Requirements for PAUSE Operation in 1 GbE Speed	B0=Yes; NoFix
16. Sageville Does Not Meet the Timing Requirements for PAUSE Operation in 100 Mb/s Speed	B0=Yes; NoFix
17. 10GBASE-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; Fixed in NVM 1.00
19. 10GBASE-T Conformance IEEE 802.3 - Power Spectral Density Test Failure	B0=Yes; NoFix
20. Energy Efficient Ethernet 10/1 GbE	N/A
21. Internal Thermal Sense Circuit Does not Generate a Thermal Sensor Event Interrupt	B0=Yes; Fixed in NVM 1.00
22. PCIe Interrupt Status Bit	B0=Yes; NoFix
23. SMBus Interface Pins Pulled Low when No Power	B0=Yes; NoFix
24. 10GBASE-T Conformance 55.5.3.2 - Transmitter Linearity	B0=Yes; NoFix
25. PHY Registers Misconfiguration	B0=Yes; Fixed in NVM 1.55
26. VFLR on the Fly	B0=Yes; NoFix
27. SDP[1] Pin is Input During PERST	B0=Yes; NoFix
28. PCIe Phase 2 Fails to Timeout Under Certain Channel Conditions	B0=Yes; Fixed in NVM 1.93
29. MCTP over PCIe Communication Stops Working if One LAN Port is Disabled and Swapped PFs Configured	B0=Yes; Fixed in NVM 3.00
30. Device ID and MAC Address Cannot be Changed	B0=Yes; NoFix
31. Spurious Uncorrectable ECC Error Detection	B0=Yes; Fixed in NVM 3.10
32. IEEE Fast Retrain (FR) Mode Failure	B0=Yes; NoFix

## 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 Sageville Replay\_Timer to expire,

causing replay timer timeout correctable errors. Sageville 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. NBASE-T Speed Advertisement

NBASE-T speed advertisement (2.5 GbE/5 GbE) is disabled by default in the Sageville NVM. An NBASE-T enabled driver should be used to link at NBASE-T speeds.

Linux driver version 4.3.9 is currently the only supported software solution for NBASE-T.

For future planned enhancements, contact your Intel representative.

## 3. SAN MAC Address

If SAN MAC Address is enabled in the NVM, it cannot be left with a default value of 0xFFFF or 0x0.

## 4. Energy Efficient Ethernet (EEE) Support

Energy Efficient Ethernet (EEE) is not enabled in Sageville. Sageville does not negotiate to or transition to EEE mode with a EEE-capable link partner.

## 5. MAC CRC Errors Resulting from Negotiated Fast Retrain Events

Electromagnetic Interference (EMI) or Radio Frequency Interference (RFI) events that trigger the Negotiated Fast Retrain protocol within the device result in MAC-layer CRC errors at a rate of approximately 1-2 errors per event. The Negotiated Fast Retrain event results in a minor interruption of the data stream (up to 280 milliseconds) and some packet loss at the physical layer is to be expected. Higher layer protocols are expected to request replays of data that is missing.

Therefore, Bit Error Ratio (BER) tests that are executed in more challenging environments where EMI/RFI is present or transient may experience errors at the MAC layer. Environmental chambers where the heating and cooling elements are in the same cabinet as the DUT are an example of this.

In this case, it is recommended to utilize F/STP CAT6a cables and ensure that chassis design sufficiently shields the device from these events.

## 6. Sageville Supports Only Auto-Negotiated Link at 100Mb

Sageville does not implement the necessary function to detect and link at 100Mb without Auto-Negotiation, as only Auto-Negotiation is supported at any speed.

## 2.2 Specification Changes

### 1. Thermal Sensor Indication

Bit 20 of STATUS registers (offset 0x8) is Reserved.

SDP1 configuration in Table 3-25, "SDP Settings" should be:

SDP	Usage	NVM Setting	ESDP			
			SDPx_NATIVE	SDPx_IODIR	SDP1_Function	SDP23_function
1	SDP	N/A	0	Input/Output	0	N/A
	PCI disable	NVM Control Word 2, SDP_FUNC_OFF_EN bit	N/A	N/A	N/A	
	1588 functionality as defined by the TSSDP register	N/A	1	Input/Output	0	
	Thermal Sensor	TS NVM-based Mode Enable bit in the Common Firmware Parameters word	0	Output	1	
	Reserved	N/A	1	N/A	1	

## 2. LANx\_DIS\_N Pins are Input Only During LPG and PERST

LANx\_DIS\_N pins are inputs only during LAN Power Good reset and PERST and are outputs after that.

If controlled by chipset pins, they should be defined as Open Drain.

Relevant only to the 17x17 mm package.

## 3. Enable All PHYs in D3

If *Enable All PHYs in D3 N* NVM bit (Common Firmware Parameters, bit 3) is set, the internal PHY stays active unconditionally.

The description in the *Intel® Ethernet Controller X550 Datasheet*, Section 6.5.2.2 should be updated to the following:

Enable All PHYs in D3 N

0b = A PHY port will be active during D3 only if APM/ACPI is enabled or if activated by an enable channel MC command.

1b = All PHY ports stay active in D3 unconditionally. This configuration should be used if Thermal Sensor is required during D3.

## 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:**

Sageville 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:**

Sageville 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. Sageville Does Not Meet the Timing Requirements for PAUSE Operation in 1 GbE Speed

### Problem:

In 1 GbE speed, Sageville 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. Sageville Does Not Meet the Timing Requirements for PAUSE Operation in 100 Mb/s Speed

### Problem:

In 100 Mb/s speed, Sageville 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; Fixed in NVM 1.00

## 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; NoFix

The power level failures have been resolved from NVM 1.93 (which includes PHY FW 2.B.B), but power spectral density failures may still be observed on some platforms.

## 20. Energy Efficient Ethernet 10/1 GbE

**Note:** Content for this erratum moved to Specification Clarification #4.

## 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 Sageville 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 Sageville 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; Fixed in NVM 1.00

## 22. PCIe Interrupt Status Bit

**Problem:**

The *Interrupt Status* bit in the Status register of the PCIe configuration space is not implemented and is not set as described in the PCIe specification.

**Implication:**

When using shared legacy PCI interrupts, software might use this bit to determine if Sageville has a pending interrupt. Since the bit is not implemented, the software might not handle the interrupt, resulting in a continuous interrupt assertion.

There is no implication when using MSI or MSI-X.

**Workaround:**

The *Interrupt Status* bit should not be used. Avoid using shared legacy PCI interrupts.

Status: B0=Yes; NoFix

## 23. SMBus Interface Pins Pulled Low when No Power

**Problem:**

The SMBus interface pins, which are supposed to present high impedance to the bus when the device is un-powered, are instead pulling the bus low.

**Implication:**

This prevents the sharing of the bus with other devices that are present on the same bus when Sageville is without power.

**Workaround:**

An external buffer circuit can be implemented on the SMBus pins of the device, as described in Technical Advisory TA-221 (Intel® Ethernet Controller X550 *SMBus Isolation Circuit*).

Status: B0=Yes; NoFix

## 24. 10GBASE-T Conformance 55.5.3.2 – Transmitter Linearity

### Problem:

10GBASE-T transmitter linearity test may fail during conformance testing. The failure may be observed on Dual Tone Test Mode One and Two.

### Implication:

There is no expected interoperability impact.

### Workaround:

None.

Status: B0=Yes; NoFix

In NVM release 1.93 (which include PHY FW 2.B.B), the linearity conformance is improved but failures may still be observed in a small number of corner units.

## 25. PHY Registers Misconfiguration

### Problem:

PHY Registers might be mis-configured in some corner conditions after software Driver initialization.

### Implication:

PHY registers misconfiguration and link down.

### Workaround:

If link is down and present errata scenario is suspected, the PHY configuration can be reloaded by setting PHY register bit 1E.C442.0 to 1b.

Status: B0=Yes; Fixed in NVM 1.55

## 26. VFLR on the Fly

### Problem:

If a VFLR is applied during Tx traffic, some of the Tx queues might hang.

### Implication:

Tx Hang.

### Workaround:

After VFLR and before setting the VFTE (per each VF), the software driver should reset Tx queues by setting and clearing the TXDCTL.ENABLE bit for each queue in this VF.

Status: B0=Yes; NoFix

## 27. SDP[1] Pin is Input During PERST

### Problem:

The SDP[1] pin direction for both ports (a.k.a., SDP0\_1 and SDP1\_1) is configured by the SDP Control Word in the NVM. When PCIe Reset is asserted (PERST), the pin direction is set to input, regardless of NVM configuration.

Relevant only to the 17x17 mm package.

### Implication:

SDP[1] output is not valid during Dr state. When used for thermal indication, no thermal event can be reported in the Dr state.

### Workaround:

Use a different SDP for an output indication that must be valid during Dr.

Status: B0=Yes; NoFix

## 28. PCIe Phase 2 Fails to Timeout Under Certain Channel Conditions

### Problem:

With certain channel conditions it was found that the PCIe Phase 2 coefficient evaluation timer is too short, causing the device to enter a persistent failure state.

### Implication:

If Phase 2 EQ fails to timeout, it will stay in this state for an unlimited amount of time, and the device will either fail to show up on the bus or cause a fatal error.

If the issue occurs, Sageville either fails to establish a PCIe link, or causes a PCIe fatal error.

### Workaround:

None.

Status: B0=Yes; Fixed in NVM 1.93

## 29. MCTP over PCIe Communication Stops Working if One LAN Port is Disabled and Swapped PFs Configured

### Problem:

Sageville Tx drop packets request modeled per Port and not PF.

### Implication:

MCTP over PCIe communication stops working if one LAN port is disabled and swapped PFs configured.

### Workaround:

None.

Status: B0=Yes; Fixed in NVM 3.00

## 30. Device ID and MAC Address Cannot be Changed

### Problem:

As part of the implementation of *Recovery Mode in for Intel® Ethernet Products* (Doc ID: 606286), starting with NVM version 2.00, the device will not change Device ID fields (VID, DID, SVID and SSID) and MAC Address fields during the original factory programming.

### Implication:

After the original factory NVM programming, Device IDs, Factory MAC Address, and SAN MAC Address cannot be changed starting with NVM version 2.00.

### Workaround:

None.

Status: B0=Yes; NoFix

## 31. Spurious Uncorrectable ECC Error Detection

### Problem:

After initial power up, Sageville might report a spurious uncorrectable ECC error for the PCIe Retry Buffer upon the first replay on the PCIe link.

### Implication:

DMA Tx and Rx traffic is halted until a PCIe reset. Additionally, the ECC bit in EICR is set, so this event can be logged by the driver.

### Workaround:

Disable reporting of uncorrectable ECC errors on the PCIe Retry Buffer.

Status: B0=Yes; Fixed in NVM 3.10

## 32. IEEE Fast Retrain (FR) Mode Failure

### Problem:

When the X550 device and the link partner both have IEEE Fast Retrain enabled on a 10GBASE-T Ethernet link and the link partner initiates a fast retrain, the link goes down and is then re-established.

### Implication:

There are link flaps in situations where fast retrain should be sufficient to maintain the link.

### Workaround:

1. Do not enable IEEE Fast Retrain. Use non-standard fast retrain modes supported by the X550 device.
2. Link Partner configured to advertise itself as "multi-port" (i.e., drive device to be secondary clock in link training).



Status: B0=Yes; No Fix



## 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

Sageville 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

When bit banging on a serial interface (such as SPI, I<sup>2</sup>C, 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

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

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.

This document (and any related software) is Intel copyrighted material, and your use is governed by the express license under which it is provided to you. Unless the license provides otherwise, you may not use, modify, copy, publish, distribute, disclose or transmit this document (and related materials) without Intel's prior written permission. This document (and related materials) is provided as is, with no express or implied warranties, other than those that are expressly stated in the license.

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 described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or visiting the [Intel Resource and Documentation Center](#).

© 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.