



Intel[®] Ethernet Controller E810

Specification Update

February 2024

Revision 3.4
616943-015

Revision History

Revision	Date	Comments
3.4	February 7, 2024	Errata added or updated: <ul style="list-style-type: none"> 35. Imprecise Reporting of Malicious Driver Detection Events (Added)
3.3	August 10, 2023	Errata added or updated: <ul style="list-style-type: none"> 30. Device is Not Enumerated on the PCIe Bus Following Certain Warm Reset Flows (Updated) 32. Inverted FCS Does Not Cause Increment of CRC Error Count (Added) 33. Incorrect Calculation of Inner UDP Checksum on Rx Packet (Added) 34. Incorrect Receive Length Errors (Added)
3.2	March 9, 2023	Errata added or updated: <ul style="list-style-type: none"> 31. PE_WAKE_N is Not Open Drain (Updated)
3.1	March 30, 2022	Errata added or updated: <ul style="list-style-type: none"> 30. Device is Not Enumerated on the PCIe Bus Following Certain Warm Reset Flows (Updated) 31. PE_WAKE_N is Not Open Drain (Added)
3.0	February 24, 2022	Errata added or updated: <ul style="list-style-type: none"> 30. Device is Not Enumerated on the PCIe Bus Following Certain Warm Reset Flows (Added)
2.9	July 27, 2021	Specification Clarifications added or updated: <ul style="list-style-type: none"> 9. Link Activity Indication (Updated) 14. Lack of Link Flap Reporting (Added) Errata added or updated: <ul style="list-style-type: none"> 29. PCIe: Fatal Uncorrectable Errors for 4DW Header TLPs Are Incorrectly Logged (Added)
2.8	May 18, 2021	Specification Changes added or updated: <ul style="list-style-type: none"> 4. Logging of PCIe Correctable Receiver Errors (Added) Miscellaneous updates: <ul style="list-style-type: none"> Added S-Specification information to Table 1-1, "Markings". Added S-Specification information to Table 1-3, "MM Numbers".
2.7	March 22, 2021	Specification Clarifications added or updated: <ul style="list-style-type: none"> 12. PLDM FW Update - Transition to Older NVM Versions (Added) 13. E810 1000BASE-T Operation (Added) Errata added or updated: <ul style="list-style-type: none"> 28. Device Fails to Initialize PCIe Link After Reboot (Added)
2.6	January 22, 2021	Specification Changes added or updated: <ul style="list-style-type: none"> 2. Updates to Table 16-8, "E810-CAM2/CAM1 MAX-Power - PCIe Gen 4 and PCIe Gen 3" (Added) 3. Updates to Table 18-2, "E810 Thermal Specifications" (Added)
2.5	December 21, 2020	Errata added or updated: <ul style="list-style-type: none"> 18. Packet Drops Might Occur Under Receive Stress (Updated) 25. The E810 Fails 10G-KR Rx BER Test Under MTC1 Condition (Added) 26. The E810 Fails 10G-SFI Rx BER Test Under ITT Condition (Added) 27. The E810 Fails 25G-CR Rx BER Test Under Long-Channel Condition (Added)
2.4	November 20, 2020	Specification Clarifications added or updated: <ul style="list-style-type: none"> 10. 50G LAUI-2 Specification, FEC mode, and Alignment Marker Clarification (Added) 11. Link Establishment Attempt on Unsupported Media Through "Lenient Mode" (Added) Specification Changes added or updated: <ul style="list-style-type: none"> 1. Power Up/Down Sequence (Added) Errata added or updated: <ul style="list-style-type: none"> 22. PCIe PM_Active_State_NAK Reception (Added) 23. Some PCIe Messages Are Always Sent with reqid=0 (Added) 24. Thermal Alarm Points for Pluggable Modules Are Using Fixed Values Instead of Reading from Module EEPROM (Added)

Revision	Date	Comments
2.3	September 30, 2020	Errata added or updated: <ul style="list-style-type: none"> 21. The Value from Temperature Sensor Readings of QSFP Module is Incorrect for All Methods That Rely on Device Firmware (Added)
2.2	September 17, 2020	Specification Clarifications added or updated: <ul style="list-style-type: none"> 9. Link Activity Indication (Added) Errata added or updated: <ul style="list-style-type: none"> 16. PLDM Type 2 PDRs for SFP+/SFP28/QSFP28 Ports/Modules Thermal Sensors Are Only Initialized During Firmware Initialization. (Updated) 18. Packet Drops Might Occur Under Receive Stress (Updated) 19. In-Band PCIe Reset Stress Might Cause Device to Enter Recovery Mode (Updated) 20. E810-XXVAM2 Fails 10G-SFI & 10G-KR Tx Trise/Tfall Electrical Test (Added)
2.1	August 10, 2020	Errata added or updated: <ul style="list-style-type: none"> 19. In-Band PCIe Reset Stress Might Cause Device to Enter Recovery Mode (Added)
2.0 ¹	July 23, 2020	Initial public release.

1. There are no previous publicly-available versions of this document.



NOTE: *This page intentionally left blank.*

1. Introduction

This document applies to the Intel® Ethernet Controller E810 (E810).

This document is an update to a published specification, the *Intel® Ethernet Controller E810 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 v4.0 (2.5GT/s, 5GT/s, 8GT/s, and 16GT/s).

For more information on supported features, see the *Intel® Ethernet Controller E810 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: EZE810CAM1, EZE810CAM2, EYE810XXVAM2

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
E810-CAM1	C0	EZE810CAM1	S LNFV ²	Ethernet controller (1x100/2x50/4x25), PCIe v4.0 x16
			S LNFV ³	
E810-CAM2	C0	EZE810CAM2	S LNFV ²	Ethernet controller (2x100/2x50/4x25/8x10), PCIe v4.0 x16
			S LNFY ³	
E810-XXVAM2	C0	EYE810XXVAM2	S LNFZ ²	Ethernet controller (1x50/2x25), PCIe v4.0 x8
			S LNG ³	

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

Table 1-2. Device IDs

Branding String	Interface Type	Compatible Silicon Device	Device ID	Vendor ID	Revision ID
Intel® Ethernet Controller E810-C	N/A	EZE810CAM1 (1x100/2x50/4x25) EZE810CAM2 (2x100/2x50/4x25/8x10)	1590	8086	0x02
Intel® Ethernet Controller E810-C for backplane	Backplane	EZE810CAM1 (1x100/2x50/4x25) EZE810CAM2 (2x100/2x50/4x25/8x10)	1591	8086	0x02
Intel® Ethernet Controller E810-C for QSFP	QSFP28	EZE810CAM1 (1x100/2x50/4x25) EZE810CAM2 (2x100/2x50/4x25/8x10)	1592	8086	0x02
Intel® Ethernet Controller E810-C for SFP	SFP28	EZE810CAM1 (1x100/2x50/4x25) EZE810CAM2 (2x100/2x50/4x25/8x10)	1593	8086	0x02
Intel® Ethernet Controller E810-XXV	N/A	EYE810XXVAM2 (1x50/2x25)	1598	8086	0x02
Intel® Ethernet Controller E810-XXV for backplane	Backplane	EYE810XXVAM2 (1x50/2x25)	1599	8086	0x02

Table 1-2. Device IDs [continued]

Branding String	Interface Type	Compatible Silicon Device	Device ID	Vendor ID	Revision ID
Intel® Ethernet Controller E810-XXV for QSFP	QSFP28	EYE810XXVAM2 (1x50/2x25)	159A	8086	0x02
Intel® Ethernet Controller E810-XXV for SFP	SFP28	EYE810XXVAM2 (1x50/2x25)	159B	8086	0x02

Table 1-3. MM Numbers

Product	S-Specification	Tray MM#	Tape and Reel MM#
E810-CAM1	S LNFV	999W9C	---
	S LNFW	---	999W9D
E810-CAM2	S LNFX	999W9F	---
	S LNFY	---	999W9G
E810-XXVAM2	S LNFZ	999W9J	---
	S LNG2	---	999W9K

1.2 Marking Diagrams



Figure 1-1. Example Component with Identifying Marks

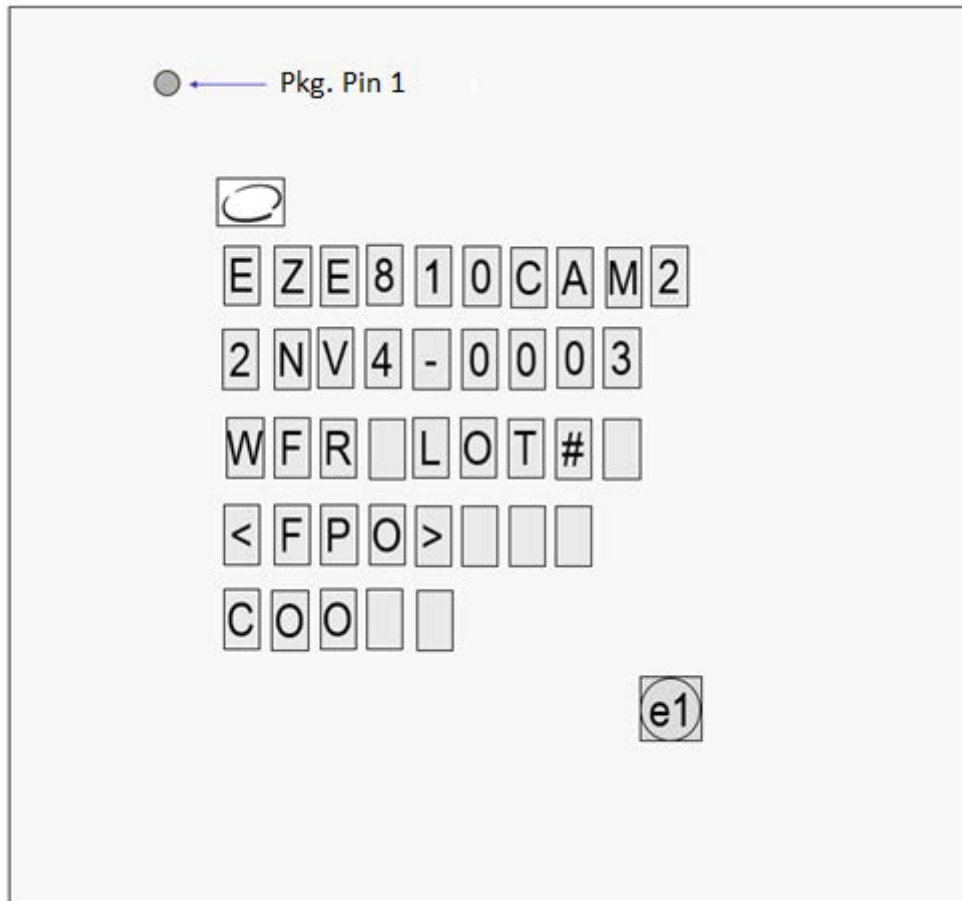


Figure 1-2. Marking Diagram (25x25 mm Package)

- Line 1: Swirl Logo
- Line 2: Product Number
- Line 3: Product ID or designator
- Line 4: Wafer Lot and Wafer ID#
- Line 5: FPO# (Unique Datecode/Lot# ID
- Line 6: Country of Origin (manufacturing assembly site)
- Line 7: RoHS Pb-free symbol

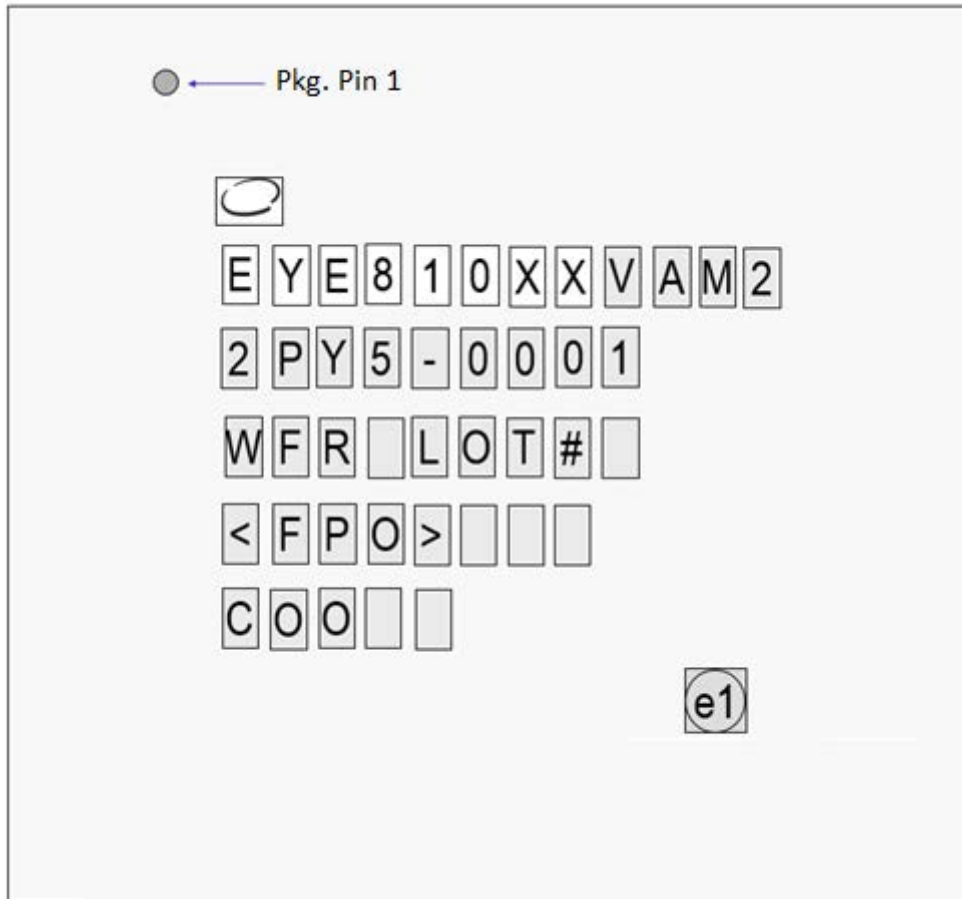


Figure 1-3. Marking Diagram (21x21 mm Package)

- Line 1: Swirl Logo
- Line 2: Product Number
- Line 3: Product ID or designator
- Line 4: Wafer Lot and Wafer ID#
- Line 5: FPO# (Unique Datecode/Lot# ID)
- Line 6: Country of Origin (manufacturing assembly site)
- Line 7: RoHS Pb-free symbol

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
A0, B0, etc.	Stepping to which the status applies.
Doc	Document change or update that will be implemented.
Documentation Changes	Typos, errors, or omissions from 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.
Eval	Plans to fix this erratum are under evaluation.
Fix Planned	This erratum is intended to be fixed in a future stepping of the component.
Fix Planned in NVM	This erratum is intended to be fixed in a future NVM version.
Fixed	This erratum has been fixed.
Fixed in NVM	This erratum has been fixed in NVM X.XX.
NoFix	There are no plans to fix this erratum.
Software Clarifications	Applies to Intel drivers, EEPROM loads.
Specification Changes	Modifications to the current published specifications. These changes will be incorporated in the next release of the specifications.
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.

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. Device Not Seen in PCIe Configuration Space	N/A
2. DCB and RDMA Recommended Configurations	N/A
3. AN73 Parallel Detect Support	N/A
4. Userspace Direct Access (UDA) Support	N/A
5. 1G BASE-T SFP Module Support	N/A
6. Receive Length Error Statistic Counter Does Not Count on Double VLAN(QnQ) Tagged Packets	N/A
7. Handling of FEC Capabilities Across Different Compliance Codes	N/A
8. Firmware Ignores MCTP Instance ID for Response from BMC	N/A
9. Link Activity Indication	N/A
10. 50G LAUI-2 Specification, FEC mode, and Alignment Marker Clarification	N/A
11. Link Establishment Attempt on Unsupported Media Through "Lenient Mode"	N/A
12. PLDM FW Update - Transition to Older NVM Versions	N/A
13. E810 1000BASE-T Operation	N/A
14. Lack of Link Flap Reporting	N/A

Table 2-2. Summary of Specification Changes

Specification Change	Status
1. Power Up/Down Sequence	N/A
2. Updates to Table 16-8, "E810-CAM2/CAM1 MAX-Power - PCIe Gen 4 and PCIe Gen 3"	N/A
3. Updates to Table 18-2, "E810 Thermal Specifications"	N/A
4. Logging of PCIe Correctable Receiver Errors	N/A

Table 2-3. Summary of Documentation Changes

Documentation Update	Status
None.	N/A

Table 2-4. Summary of Errata; Errata Include Steppings

Erratum	Status
1. [TimeSync] Inc/Dec Timer by One Failed	C0=Yes; NoFix
2. AN73 Parallel Detect Support	C0=Yes; NoFix
3. RDMA: Sending Invalid Work Queue Elements (WQE) Might Cause the Device to Halt All Tx/Rx	C0=Yes; NoFix
4. DUT Does Not Receive Packets of Variable Preamble Pattern	C0=Yes; NoFix
5. Device Might Advertise PCIe x16 Support when Configured for x8	C0=Yes; NoFix
6. E810-CAM1 and E810-CAM2 Fail RX-HIGH-IMP-DC-POS/NEG PCIe DC Test	C0=Yes; NoFix
7. E810-CAM1 and E810-CAM2 Fail CR-PAM4 Rx Conformance Spec	C0=Yes; NoFix
8. E810-CAM1 and E810-CAM2 Fail 1G-KX Tx Mask Test	C0=Yes; NoFix
9. E810-CAM1 and E810-CAM2 Fail 10G-SFI Tx Tr/Tf, ucJ, and SNR Electrical Tests	C0=Yes; NoFix
10. E810-CAM1 and E810-CAM2 Fail 10G-KR Tx Tr/Tf, Vpp and a Few Tap Tests	C0=Yes; NoFix
11. E810-CAM1 and E810-CAM2 Fail 25G-KR Tx AC Common Mode Test	C0=Yes; NoFix
12. E810-CAM1 and E810-CAM2 Fail 50G-CR Tx for Some Tap Tests	C0=Yes; NoFix
13. E810-CAM1 and E810-CAM2 Fail 50G-KR Tx in Some Electrical Parametric Testing	C0=Yes; NoFix
14. E810-XXVAM2 Fails 25G-CR Long Channel Rx Configuration Spec	C0=Yes; NoFix
15. NC-SI Might Exceed the Normal Execution Interval (T1=50 ms)	C0=Yes; Fix Planned in NVM
16. PLDM Type 2 PDRs for SFP+/SFP28/QSFP28 Ports/Modules Thermal Sensors Are Only Initialized During Firmware Initialization.	C0=Yes; Fixed in NVM 2.10
17. With Link Level PAUSE Feature it is Required to Set Both Rx and Tx Direction PAUSE Symmetrically	C0=Yes; NoFix
18. Packet Drops Might Occur Under Receive Stress	C0=Yes; Fixed in NVM 2.30
19. In-Band PCIe Reset Stress Might Cause Device to Enter Recovery Mode	C0=Yes; Fixed in NVM 2.10
20. E810-XXVAM2 Fails 10G-SFI & 10G-KR Tx Trise/Tfall Electrical Test	C0=Yes; NoFix
21. The Value from Temperature Sensor Readings of QSFP Module is Incorrect for All Methods That Rely on Device Firmware	C0=Yes; Fixed in NVM 2.14
22. PCIe PM_Active_State_NAK Reception	C0=Yes; NoFix
23. Some PCIe Messages Are Always Sent with reqid=0	C0=Yes; NoFix
24. Thermal Alarm Points for Pluggable Modules Are Using Fixed Values Instead of Reading from Module EEPROM	C0=Yes; Fix Planned in NVM
25. The E810 Fails 10G-KR Rx BER Test Under MTC1 Condition	C0=Yes; NoFix
26. The E810 Fails 10G-SFI Rx BER Test Under ITT Condition	C0=Yes; NoFix
27. The E810 Fails 25G-CR Rx BER Test Under Long-Channel Condition	C0=Yes; NoFix
28. Device Fails to Initialize PCIe Link After Reboot	C0=Yes; NoFix
29. PCIe: Fatal Uncorrectable Errors for 4DW Header TLPs Are Incorrectly Logged	C0=Yes; NoFix
30. Device is Not Enumerated on the PCIe Bus Following Certain Warm Reset Flows	C0=Yes; Fixed in NVM 4.20
31. PE_WAKE_N is Not Open Drain	C0=Yes; NoFix

Table 2-4. Summary of Errata; Errata Include Steppings [continued]

Erratum	Status
32. Inverted FCS Does Not Cause Increment of CRC Error Count	C0=Yes; NoFix
33. Incorrect Calculation of Inner UDP Checksum on Rx Packet	C0=Yes; NoFix
34. Incorrect Receive Length Errors	C0=Yes; NoFix
35. Imprecise Reporting of Malicious Driver Detection Events	C0=Yes; NoFix

2.1 Specification Clarifications

1. Device Not Seen in PCIe Configuration Space

In some systems, the MMIO space requested by the E810 might not be allocated within the 32-bit space (low 4G). In such systems, it is recommended to enable allocation of MMIO in the memory space above 4G.

This option can be found at:

BIOS > Advanced > PCIe/PCI/PnP Configuration > Above 4G Decoding

2. DCB and RDMA Recommended Configurations

When using the device in a configuration with greater than four PFs, DCB and RDMA are not supported. When the number of PFs is four or less, both RDMA and DCB are supported. In this configuration DCB is supported with up to four Traffic classes, with one of them enabled for Priority Flow Control (PFC).

Number of NIC Ports	Traffic Class Recommendation with DCB	RDMA
1, 2, or 4	Up to 4 Traffic classes, with one of them enabled for PFC.	Supported
Greater than 4 ports	No DCB support.	Not supported

3. AN73 Parallel Detect Support

AN73 Parallel detect for 1000BASE-KX support is not supported in E810 devices.

Enablement of LESM allows the E810 device to link with devices that do not support auto-negotiation.

4. Userspace Direct Access (UDA) Support

Userspace Direct Access (UDA) is not supported in Intel® Ethernet 800 Series products. There are no plans to productize UDA as an end-user capability in Intel® Ethernet 800 Series devices. There is an internal queueing mechanism referred to as UDA that is used exclusively by the kernel space driver for iWARP connection setup and error handling mechanisms. The *Intel® Ethernet Controller E810 Datasheet* and related documents will be updated in a future release.

5. 1G BASE-T SFP Module Support

The device currently only works with 1G BASE-T modules that support SGMII connection to the host. Modules using other PHY type connections are not supported.

6. Receive Length Error Statistic Counter Does Not Count on Double VLAN(QnQ) Tagged Packets

Receive Length Error count GLPRT_RLEC[n] ($0x00380140 + 0x8*n$, $n=0..7$) does not count on double VLAN(QnQ) tagged packets. It only increments for packets with a single or no VLAN tag.

7. Handling of FEC Capabilities Across Different Compliance Codes

Depending on the compliance codes advertised by a connected module, the PHY will attempt to link at various combinations of link speed and FEC type. The device uses a single FEC configuration input for all speeds, so if multiple speeds are enabled, certain speed/FEC combinations might be attempted that are not supported by the connected media. If the link partner configuration is set to an unsupported speed and FEC combination as well, link might not establish or a link that does not have a healthy bit error rate might be established.

8. Firmware Ignores MCTP Instance ID for Response from BMC

When receiving an MCTP response, the E810 firmware ignores the instance ID field and accepts the response even if it does not match.

This was done to ensure interoperability with certain BMCs.

9. Link Activity Indication

The device implements a polling mechanism of the Ethernet link activity to determine LED blink behavior. The LEDs are updated every 200 ms based on link activity. Therefore, link activity that is initiated and completed in less than 200 ms might not always result in LED illumination.

10. 50G LAUI-2 Specification, FEC mode, and Alignment Marker Clarification

The Ethernet ecosystem contains a variety of two-lane NRZ-based 50G interface types, defined in multiple specifications and standards. Two common types are LAUI-2 defined by IEEE 802.3cd-2018 Annex 135B/135C, and 50GBASE-R2 defined by the Ethernet Technology Consortium (ETC) Schedule 3. The LAUI-2 uses the IEEE Clause 133 50GBASE-R PCS, whereas the ETC 50GBASE-R2 uses a derivative of IEEE Clause 82 40GBASE-R PCS. These interface types support different FEC encoding capabilities and use different alignment marker spacing, making them incompatible with each other:

- **IEEE 802.3 LAUI-2** — Only supports No FEC data transmission and reception. The alignment marker interval is every 20,479 66-bit blocks.
- **ETC 50GBASE-R2** — Supports No FEC, BASE-R FEC, and RS(528) FEC data transmission and reception. Alignment marker spacing is 16,383 66-bit blocks for No FEC and BASE-R FEC configurations, or 20,479 66-bit blocks for the RS(528) FEC configuration.

For link to be established, the two end points of the link must be configured with both the exact same FEC encoding and the exact same lane alignment marker spacing. The ETC specification assumes the use of Auto-Negotiation Next Page exchanges to exchange FEC ability and requests between devices. However, LAUI-2 does not. Many devices in the Ethernet ecosystem bypass Auto-Negotiation and use a forced configuration similar to LAUI-2, but the ETC 50GBASE-R2 alignment marker spacing is thought to be the more prevalent configuration used in the market. Therefore, to increase interoperability with common devices in the marketplace, Software Release 25.5 implements LAUI-2 with No FEC using the 50GBASE-R2 alignment marker spacing of 16,383 66-bit blocks for two-lane NRZ-based 50G interface forced configurations.

11. Link Establishment Attempt on Unsupported Media Through “Lenient Mode”

The E810 includes a feature referred to as “Lenient Mode”, which allows link establishment to be attempted on unsupported media. Media might be considered unsupported due to lack of inclusion in the Feature Support Matrix or due to an inability to determine the media type (for example, if there are inaccuracies or missing information in the module EEPROM). When an unsupported module is inserted and the E810 is in Lenient Mode, it will attempt a generic link establishment process. This includes cycling through all supported PHY configurations, which is a combination of speed and FEC, and monitoring for a compatible link partner configuration.

If link is able to be established in this manner, it could take a long time to for the E810 to be in a mode compatible with the link partner. For example, a 100G capable E810 port using Lenient Mode and connected to a link partner configured for 50G must cycle through auto-negotiation as well as multiple FEC configurations of 100G before any 50G link modes are attempted. Due to this, long time to link is expected in configurations relying on Lenient Mode. Depending on the total loop time and the position within the loop at the time that the link partner is connected/enabled, TTL can range from 1 second to as high as 60 seconds. Time to link can be much more consistent and reduced by setting the desired speed and/or FEC configuration of the E810, which limits the number of configurations which are being attempted.

Note that Lenient Mode cannot be used to enable modules that exceed the power class support capabilities of the device. Also note that some media is incompatible with certain PHY configurations and will not be able to support link. For example, modules that have an internal FEC encoding layer cannot work if the PHY is also configured to enable FEC. Finally, it is especially critical that the link partner is in a static configuration. If both devices are attempting to support multiple configurations (for example, two E810 devices connected together with both doing Lenient Mode loops), link might never be establish since it is unlikely that both devices will be configured for the media-appropriate configuration at the same time. By default the device has “Lenient Mode” enabled.

12. PLDM FW Update - Transition to Older NVM Versions

When using NVM 2.3x or earlier, transitions to earlier NVM versions are blocked by device firmware when attempted by PLDM FW update. Host NVM update is not affected and allows a downgrade to previous NVM versions. This is fixed in NVM 2.4x and later.

13. E810 1000BASE-T Operation

The E810 currently supports a subset of 1000BASE-T SFP module types, which use SGMII to connect back to the E810. For the E810 to properly know the link status of the module's BASE-T external connection, the module must indicate the BASE-T side link status to the E810. An SGMII link between the E810 and the 1000BASE-T SFP module allows the module to indicate its link status to the E810 using SGMII Auto Negotiation. However, 1000BASE-T SFP modules implement this in a wide variety of ways, and other methods that do not use SGMII are currently unsupported in the E810. Depending on the implementation, link might never be achieved. In other cases, if the module sends IDLEs to the E810 when there is no BASE-T link, the E810 might interpret this as a link partner sending valid data and might show link as being up even though it is only connected to the module and there is no link on the module's BASE-T external connection.

14. Lack of Link Flap Reporting

The device implements a polling mechanism to determine the Ethernet link status, which includes checking the link state every 100 ms. If the Ethernet link is lost and subsequently re-established in less than 100 ms, the link flap might not be detected by the link management engine, and a Link Status Event reporting the link drop is not generated.

2.2 Specification Changes

1. Power Up/Down Sequence

To minimize voltage leakage during power up, Intel has updated the *Intel® Ethernet Controller E810 Datasheet* Chapter 16.3.1.1 Power On/Off Sequence as follows:

“Application of power should begin with the 3.3 V (VDDIO33) and 1.8 V (VDDH) supplies **at the same time**. After 3.3 V reaches 80% of its final value, the digital core 0.8 V supply (VDD) should be enabled and reach its final value within 10 ms.”

2. Updates to Table 16-8, “E810-CAM2/CAM1 MAX-Power - PCIe Gen 4 and PCIe Gen 3”

The following table was updated in a Revision 2.1 of the *Intel® Ethernet Controller E810 Datasheet* (red indicates changes):

Table 16-8. E810-CAM2/CAM1 MAX Power - PCIe Gen 4 and PCIe Gen 3

	Link Speed					
	2x100G PCIe Gen4x16	1x100G/ 2x50G/4x25G PCIe Gen4x16	2x25G PCIe Gen4x8	2x100G PCIe Gen3x16	1x100G/ 2x50G/4x25G PCIe Gen3x16	2x25G PCIe Gen3x8
VDDIO33 (A)	0.02	0.02	0.02	0.02	0.02	0.02
VDDH (A)	0.02	0.02	0.02	0.02	0.02	0.02
AVDDH (A)	0.80	0.50	0.20	0.80	0.50	0.20
AVDD_ETH (A)	3.00	1.80	0.90	3.00	1.80	0.90
AVDD_PCIE (A)	3.904	3.90	2.00	2.60	2.60	1.40
VDD (A)	10.60	10.00	8.80	10.30	9.70	8.50
Power (W)	15.67	13.78	9.97	14.26	12.37	9.19

Notes:

- VCCIO33 current includes on-chip power dissipation only for TDP calculation. System power supply should account for external 3.3 V powered devices, such as LEDs, flash, oscillator, and pluggable modules.
- AVDD_ETH current includes 0.02A for the AVDD_PLL supply.

3. Updates to Table 18-2, “E810 Thermal Specifications”

The following table was updated in a Revision 2.1 of the Intel® Ethernet Controller E810 Datasheet (red indicates changes):

Table 18-2. E810 Thermal Specifications

Parameter	E810-CAM2	E810-CAM1	E810-XXVAM2	Notes
T _{J-MAX}	105 °C	105 °C	105 °C	
TDP	15.7 W	13.8 W	10.8 W	Max TDP; characterized at T _J = 105 °C. Refer to Section 16.4, “Power Dissipation” for a power down by each power rail and operational conditions.
θ _{JC}	0.69 °C /W	0.69 °C /W	0.15 °C /W	Thermal resistance junction-to-case.
θ _{JB}	4.3 °C /W	4.3 °C /W	7.4 °C /W	Thermal resistance junction-to-board.

4. Logging of PCIe Correctable Receiver Errors

The optional error logging of correctable Receiver Errors is disabled in the 800 Series device, as allowed in *PCI Express Base Specification, Revision 4.0, Version 1.0, Section 7.8.4.5, Correctable Error Status Register Footnote 137*.

PHY layer Receiver Error detection and recovery mechanisms are operational such that there is no functional implication to the device or system operation.

Note: Both Correctable Error Status Register[0] and Correctable Error Mask Register[0] are implemented such that the 800 Series device is a PCI-SIG compliant device.

This change is implemented in NVM 2.4x and later.

2.3 Documentation Changes

None.

2.4 Errata

1. [TimeSync] Inc/Dec Timer by One Failed

Problem:

After submitting the command to increment/decrement the master timer by '1', TIME_0 increment/decrement by '1' instead of the TIME_L.

Implication:

Need two register accesses instead of one.

Workaround:

Use Adjust the Timer by 'N', where N=1.

Status: C0=Yes; NoFix

2. AN73 Parallel Detect Support

Problem:

AN73 Parallel Detect support is not supported in E810 devices.

Implication:

Any link requiring Parallel Detect might not link if LESM is not enabled.

Workaround:

Enablement of LESM allows the E810 device to link with devices that do not support auto-negotiation.

Status: C0=Yes; NoFix

3. RDMA: Sending Invalid Work Queue Elements (WQE) Might Cause the Device to Halt All Tx/Rx

Problem:

When using RDMA, sending invalid Work Queue Elements (WQE) in a loop might cause the device to halt all Tx/Rx.

Implication:

All Tx/Rx is stopped until Global reset is issued.

Workaround:

If this issue occurs, apply Global reset.

Status: C0=Yes; NoFix

4. DUT Does Not Receive Packets of Variable Preamble Pattern

Problem:

IEEE compliance UNH- IOL Clause 4 spec fail. Test Suite: Test #4.1.9 - Receive Variable Preamble Pattern, DUT does not accept packets with preamble having patterns other than 55.

Implication:

Received packets with preamble having patterns other than 55 are dropped.

Workaround:

None.

Status: C0=Yes; NoFix

5. Device Might Advertise PCIe x16 Support when Configured for x8

Problem:

If the upper PCIe 8:15 are used for a x8 design, the device advertises x16 link width support instead of x8.

Implication:

Compliance fail. However, the link will be at x8.

Workaround:

None.

Status: C0=Yes; NoFix

6. E810-CAM1 and E810-CAM2 Fail RX-HIGH-IMP-DC-POS/NEG PCIe DC Test

Problem:

The E810 failed RX-HIGH-IMP-DC-POS/NEG PCIe DC Test, which is part of the PCIe base specification.

Implication:

None.

Workaround:

None.

Status: C0=Yes; NoFix

7. E810-CAM1 and E810-CAM2 Fail CR-PAM4 Rx Conformance Spec

Problem:

The E810 failed CR-PAM4 Rx Conformance Spec; ITT - Long & Short channels; JTT - long channel.

Implication:

Compliance test failure. No functional impact expected.

Workaround:

None.

Status: C0=Yes; NoFix

8. E810-CAM1 and E810-CAM2 Fail 1G-KX Tx Mask Test

Problem:

E810-CAM1 and E810-CAM2 fail 1G-KX Tx Eye Mask electrical test. A few K mask hits, where actual eye "touches" the defined mask - in LVHT corner.

Implication:

Compliance test failure. No functional impact expected.

Workaround:

None.

Status: C0=Yes; NoFix

9. E810-CAM1 and E810-CAM2 Fail 10G-SFI Tx Tr/Tf, ucJ, and SNR Electrical Tests

Problem:

E810-CAM1 and E810-CAM2 fail Trise/Tfall, Uncorrelated Jitter, and QSQ-SNR electrical tests.

- Trise/Tfall: 25-30 ps vs. 34 ps spec.
- Uncorrelated Jitter: Up to 3.3 ps vs. 2.3 ps spec.
- QSQ-SNR: 61 db vs. 63 db spec.

Implication:

Compliance test failure. No functional impact expected.

Workaround:

None.

Status: C0=Yes; NoFix

10. E810-CAM1 and E810-CAM2 Fail 10G-KR Tx Tr/Tf, Vpp and a Few Tap Tests

Problem:

E810-CAM1 and E810-CAM2 fail 10G-KR Trise/fall, Vpp, and some Tx Tap tests.

- Trise/fall: 16-22 ps vs. 24 ps spec.
- Vpp: Down to 747 mV vs. 800 mV spec at LVHT corner.
- Tap: Some C(-1) increment/decrement tests, Rpst - all at specific corners.

Implication:

Compliance test failure. No functional impact expected.

Workaround:

None.

Status: C0=Yes; NoFix

11. E810-CAM1 and E810-CAM2 Fail 25G-KR Tx AC Common Mode Test

Problem:

E810-CAM1 and E810-CAM2 fail 25G-KR AC Common Mode Test. Common mode reaches 15 mV RMS in HV corners vs 12 mV spec.

Implication:

Compliance test failure. No functional impact expected.

Workaround:

None.

Status: C0=Yes; NoFix

12. E810-CAM1 and E810-CAM2 Fail 50G-CR Tx for Some Tap Tests

Problem:

E810-CAM1 and E810-CAM2 fail for some Tap tests in 50G-CR-PAM4 Tx conformance. Failures recorded in some C(0)-Hold and C(1)-Hold tests, when other coefficients are incremented/decremented (e.g 0.02 voltage change in C(0), when up to 0.005 change is allowed).

Implication:

Compliance test failure. No functional impact expected.

Workaround:

None.

Status: C0=Yes; NoFix

13. E810-CAM1 and E810-CAM2 Fail 50G-KR Tx in Some Electrical Parametric Testing

Problem:

E810-CAM1 and E810-CAM2 fail 50G KR Tx electrical tests for some SNDR tests, C(-1) setting, and a few Tap tests.

- C(-1) initial conditions: -0.19 vs. -0.2 spec.
- Tap11 SNDR: 28.5-32 db vs. 32.5 db spec.
- Tap10 C(1) Hold: -0.04 average vs. -0.005 spec.
- Tap13 C(1) Hold: 0.01 vs. 0.005 spec.

Implication:

Compliance test failure. No functional impact expected.

Workaround:

None.

Status: C0=Yes; NoFix

14. E810-XXVAM2 Fails 25G-CR Long Channel Rx Configuration Spec

Problem:

E810-XXVAM2 25G-CR fails to establish link when the Rx channel insertion loss reaches the informative spec of 35 db.

Implication:

No link.

Workaround:

E810-XXVAM2 25G-CR succeeds to establish link and meets BER spec when the overall Rx channel insertion loss reaches the mandatory spec of 29.9 db, and with cable insertion loss reaching the max cable spec of 22.5 db insertion loss.

Status: C0=Yes; NoFix

15. NC-SI Might Exceed the Normal Execution Interval (T1=50 ms)

Problem:

When the system transitions between power states, the response time over NC-SI might exceed the Normal Execution Interval (T1=50 ms). In any case, the device responds within Asynchronous Reset Interval (T6=2 s).

Implication:

NC-SI Responses might be longer than expected.

Workaround:

None.

Status: C0=Yes; Fix Planned in NVM

16. PLDM Type 2 PDRs for SFP+/SFP28/QSFP28 Ports/Modules Thermal Sensors Are Only Initialized During Firmware Initialization.

Problem:

PLDM Type 2 PDRs for SFP+/SFP28/QSFP28 ports/modules thermal sensors are only initialized during Firmware initialization.

Implication:

Modules plugged in after firmware initialization are not detected.

Workaround:

An A/C power cycle with modules plugged in causes the firmware to reset and detect new modules.

Status: C0=Yes; Fixed in NVM 2.10

17. With Link Level PAUSE Feature it is Required to Set Both Rx and Tx Direction PAUSE Symmetrically

Problem:

When utilizing Link Level PAUSE feature, it is required to set both Rx and Tx direction PAUSE symmetrically.

Implication:

If PAUSE is not set symmetrically, the device might not generate or respond to PAUSE frames appropriately.

Workaround:

None.

Status: C0=Yes; NoFix

18. Packet Drops Might Occur Under Receive Stress

Problem:

Packet drops might occur under receive stress.

Implication:

Might impact receive stress and performance tests. Devices based on the Intel® Ethernet 800 Series Controllers are designed to tolerate a limited amount of system latency during PCIe and DMA transactions. If these transactions take longer than the tolerated latency, it can impact the length of time the packets are buffered in the device and associated memory, which might result in dropped packets. These packets drops typically do not have a noticeable impact on throughput and performance under standard workloads.

Workaround:

Physical memory should be in a high-performance configuration as recommended by platform vendor. A common recommendation is all channels populated with a single DIMM module. Additionally, if these packet drops appear to affect your workload, changes to the system power management setting might improve this. See the Software Release notes for E810 devices for more details.

Status: C0=Yes; Fixed in NVM 2.30

Note: The NVM change in NVM 2.3x optimized packet handling for typical workloads. This will reduce the packet drop rate in high system latency configurations. However, if the system has other congestion points beyond the E810 (for example, changes in system memory population), then packet drop might occur. The rate of packet drop is system configuration dependent — CPU generation, BIOS settings, memory configuration, PCIe speed and end-user application in use.

19. In-Band PCIe Reset Stress Might Cause Device to Enter Recovery Mode

Problem:

The device firmware might enter recovery mode during stress testing of PCIe in-band resets where many resets are done back-to-back. This is not observed in normal operation but can be seen by running PCI-SIG Configuration Verification test under stress.

Implication:

When this is observed, the device enters Recovery mode and a NVM update and reboot is required.

Workaround:

None.

Status: C0=Yes; Fixed in NVM 2.10

20. E810-XXVAM2 Fails 10G-SFI & 10G-KR Tx Trise/Tfall Electrical Test

Problem:

E810-XXVAM2 fails 10G-SFI & 10G-KR Tx Trise/Tfall electrical test with a Trise/Tfall of 25-30 ps vs. 34 ps specification.

Implication:

Compliance test failure. No functional impact expected.

Workaround:

None.

Status: C0=Yes; NoFix

21. The Value from Temperature Sensor Readings of QSFP Module is Incorrect for All Methods That Rely on Device Firmware

Problem:

Reading the module temperature sensor with NVM 2.10/2.12 gives an incorrect value for all methods that rely on firmware, including NCSI and PLDM. User can expect reported temperature ranges from 0 to 255 with no consistency. QSFP modules with a temperature sensor whose module EEPROMs are aligned to specification SFF-8636, Revision 2.8 or newer could result in port shutdown that is only recoverable with a power cycle and a change of module.

Implication:

Temperature sensor reporting methods that rely on device firmware cannot be used. This could result in port shutdown with modules aligned to SFF-8636, Revision 2.8 or newer.

Workaround:

None.

Status: C0=Yes; Fixed in NVM 2.14

22. PCIe PM_Active_State_NAK Reception

Problem:

If PCIe Active State Power Management (ASPM) L1 is enabled, and the downstream port sends out an PM_Active_State_NAK as a response to the E810 initiating transition to L1 (PM_Active_State_Request_L1), PCIe TLP transactions might halt.

Implication:

PCIe transactions might halt if ASPM L1 is enabled.

Workaround:

Disable ASPM L1 (default in the E810), or disable downstream port ability to NAK PM_Active_State_Request_L1 requests.

Status: C0=Yes; NoFix

23. Some PCIe Messages Are Always Sent with reqid=0

Problem:

Legacy Interrupts (Assert_INTX/Deassert_INTX) and PME_TO_Ack PCIe messages are always sent with reqid=0. PCI Express Root Complex expects a non-zero value and can be considered an uncorrectable error.

Implication:

Drivers using legacy Interrupts might fail. For Intel drivers, this impacts Legacy PXE UNDI driver only. UEFI PXE is not affected.

Workaround:

Do the following in the PCI Express Root Complex to change how this issue handled:

- Clear ACS Violation Severity bit to Non Fatal in Uncorrectable Error Severity Register under AER. This still results in an error reported to the PCI Express Root Complex but the severity is lowered to Non_Fatal.

Status: C0=Yes; NoFix

24. Thermal Alarm Points for Pluggable Modules Are Using Fixed Values Instead of Reading from Module EEPROM

Problem:

Thermal alarm points for pluggable modules are using fixed values defined by the topology netlist instead of reading from module EEPROM.

Implication:

Incorrect thermal alarm points that could result false alarm events or failures to raise an alarms when necessary.

Workaround:

None.

Status: C0=Yes; Fix Planned in NVM

25. The E810 Fails 10G-KR Rx BER Test Under MTC1 Condition

Problem:

Some BER failures recorded in specific PVT conditions when testing 10G-KR under MTC1 requirements.

Implication:

Compliance test failure. No functional impact anticipated under expected use cases.

Workaround:

None.

Status: C0=Yes; NoFix

26. The E810 Fails 10G-SFI Rx BER Test Under ITT Condition

Problem:

Some BER failures recorded in specific PVT conditions when testing 10G-SFI under ITT requirements.

Implication:

Compliance test failure. No functional impact anticipated under expected use cases.

Workaround:

None.

Status: C0=Yes; NoFix

27. The E810 Fails 25G-CR Rx BER Test Under Long-Channel Condition

Problem:

Some BER failures recorded in specific PVT conditions when testing 25G-CR under long-channel requirements.

Implication:

Compliance test failure. No functional impact anticipated under expected use cases.

Workaround:

None.

Status: C0=Yes; NoFix

28. Device Fails to Initialize PCIe Link After Reboot

Problem:

A PCIe Receiver Error that is detected by the E810 just before PERST assertion can cause the internal bus control logic to become unresponsive. This leads to a failure to initialize PCIe link after the PERST, which results in the device not appearing in PCIe configuration space.

This issue has only been observed on platforms that shut down the PCIe link prematurely right before PERST is asserted.

Implication:

The E810 becomes unresponsive during PCIe reset. A cold reset (power cycle of the device or toggle LAN_PWR_GOOD) is required to recover the device and reestablish PCIe link.

Workaround:

Internally masking the reporting of PCIe Receiver Errors prevents this type of failure since no messages are sent on the internal bus at the time of PCIe reset assertion. This workaround is implemented NVM version 2.4x and later.

Status: C0=Yes; NoFix

29. PCIe: Fatal Uncorrectable Errors for 4DW Header TLPs Are Incorrectly Logged

Problem:

Fatal Uncorrectable errors detected on 4DW Header TLPs are incorrectly logged with the 4th Header DW set to zeros. The first 3 DW of the Header are logged with correct values.

This is relevant to some uncorrectable errors and only in the cases the Severity is set to Fatal. This not relevant to those errors that cannot be correlated to a TLP - e.g. Data Link Protocol error, Flow Control Protocol error, Completion Timeout, Receiver Overflow, etc.

The wrong value is logged in the "Header Log Register" address 0x128 of E810 Configuration space.

Implication:

The host or any event collector gets the wrong TLP Header of failing packet.

Workaround:

None.

Status: C0=Yes; NoFix

30. Device is Not Enumerated on the PCIe Bus Following Certain Warm Reset Flows

Problem:

PERST assertion within 20 ms of a Hot Reset (SBR) cycle might cause the device to halt and stay in reset.

Implication:

PCIe link stays down and consequently the device is not enumerated on the PCIe bus. A power cycle of the NIC is required to recover from this issue.

Workaround:

Either of these two conditions should be met to avoid any risk of observing this errata:

- Ensure platform asserts PERST before E810 Hot Reset exit. The E810 exits Hot Reset 2 ms after Downstream port exits Hot Reset according to PCIe specifications.
- Ensure that PERST is not asserted within 20 ms of the E810 exiting Hot Reset.

Status: C0=Yes; Fixed in NVM 4.20

31. PE_WAKE_N is Not Open Drain

Problem:

Output pin PE_WAKE_N is supposed to be implemented as an open-drain signal to allow for wired-OR connections among multiple devices in the system. However, PE_WAKE_N is implemented as a two-state output buffer that drives the output high when the internal wake signal is negated.

Implication:

When multiple devices in the system share the same PCIe WAKE# signal in a wired-OR configuration, the E810 causes electrical contention with any other device that is asserting the WAKE# signal, thereby preventing the wake-up logic from functioning correctly.

Workaround:

Add an open-drain buffer connected to the PE_WAKE_N pin.

Starting from NVM version 3.20, the PE_WAKE_N output buffer is disabled following the NVM auto-load for designs that do not include an external open-drain buffer. The output buffer is then enabled by firmware when the internal wake signal is asserted, and it is disabled following the negation of PCIe reset.

Starting from NVM version 4.10, when using a design that includes an external open-drain buffer as a hardware workaround, the firmware workaround can be disabled from the NVM by setting the *WAKE_N is OD* bit to 1b in the WA Feature Enable TLV of the PFA.

Status: C0=Yes; NoFix

32. Inverted FCS Does Not Cause Increment of CRC Error Count

Problem:

An incorrect value of the Ethernet FCS on a received Ethernet packet is detected by the E810. Normally, the packet is dropped and the Port CRC Error Count (GLPRT_CRCERRS) counter is incremented. In the specific case where the FCS is the bitwise inverse of the correct FCS value, the packet is dropped but the counter is not incremented.

Implication:

Receiving the bitwise inverse of the FCS is almost always due to an intentional modification of the FCS; there is an extremely low probability of random bit errors causing an inverted FCS field. Therefore, the lack of counting these FCS errors affects testing environments, but should not affect detection of bit errors during operation.

Workaround:

When testing the CRC error reporting, do not use the inverted FCS value.

Status: C0=Yes; NoFix

33. Incorrect Calculation of Inner UDP Checksum on Rx Packet

Problem:

If a tunneled UDP packet is received in which there is L2 padding in the inner packet, the offloaded inner UDP checksum calculation is incorrect. A checksum error is reported on a packet that has the correct checksum.

Implication:

When receiving an indication of a checksum error, the Intel OS drivers report to the OS that the checksum was not offloaded, and the OS stack performs the checksum validation. Therefore, there is no effect on the packet reception other than a slight increase in CPU usage.

The driver reports the error count to the OS, possibly resulting in a high error rate that is reported by the OS.

Workaround:

The driver should not report the checksum error count to the OS. This is implemented in Intel drivers starting from Release 28.2.

Status: C0=Yes; NoFix

34. Incorrect Receive Length Errors

Problem:

If an Ethernet packet is received with a packet length larger than 64 bytes and the *Length/Type* field contains a value that is smaller than the number of bytes remaining in the packet, an Rx length error is detected. This type of packet could be generated if a short packet is padded to 64 bytes and then a VLAN tag is inserted without removing any padding bytes.

When an Rx length error is detected, the packet is considered a bad packet and is discarded unless storing bad packets is enabled. Also, the GLPRT_RLEC counter is incremented.

Implication:

Packets with the structure defined above are dropped. These packets are often local diagnostic packets that might be discarded by the software anyway.

The driver reports the error count to the OS, possibly resulting in a high error rate that is reported by the OS.

Workaround:

The driver should not report the Rx length error count to the OS. This is implemented in Intel drivers starting from Release 28.2.

Status: C0=Yes; NoFix

35. Imprecise Reporting of Malicious Driver Detection Events

Problem:

Some of the Tx Data checks performed as part of the Malicious Driver Detection (MDD) are reported as anti-spoof failures in addition to the actual failures. `GL_MDET_TX_TDPU.MAL_TYPE` can indicate an anti-spoof failure instead of the actual cause of the failure.

Implication:

Misleading information when an MDD event occurs can complicate the debug of such events.

Workaround:

If `GL_MDET_TX_TDPU.MAL_TYPE` indicates an anti-spoof failure, `PRTDCB_TDPUC` can be checked for the actual cause. If any bits of `PRTDCB_TDPUC` other than `RCU_ANTISPOOF` are set, it can be assumed that the anti-spoof indication is a side effect of the actual malicious event.

Status: C0=Yes; NoFix

3. Software Clarifications

Table 3-1. Summary of Software Clarifications

Software Clarification	Status
None.	N/A

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

Copies of documents that are referenced in this document can be obtained by visiting the [Intel Resource and Documentation Center](#).

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.

© 2020-2024 Intel Corporation.