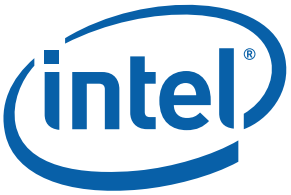

Intel® Ethernet Controller I210 Specification Update

December 2013
Revision 1.3

Revision History

Date	Revision	Description
December 2013	1.3	Added Flash-less device IDs to Table 2. Added Specification Change #3, #4, and #5. Added Documentation Update #2. Added Errata #18 through #28.
July 2013	1.2	Added Specification Changes #1 and #2. Added Specification Clarification #4. Added Documentation Update #1. Added Errata #12 through #17.
January 2013	1.1	Added Errata #8, #9, #10 and #11.
October 2012	1.0	Initial release (Intel public). Added Specification Clarification #1 through #3. Added Errata #1 through #7. Added Software Clarification #1 and #2.

Note: This page intentionally left blank.

1.1 Introduction and Scope

This document applies to the Intel® Ethernet Controller I210.

This document is an update to a published specification, the *Intel® Ethernet Controller I210 Datasheet*. It is intended for use by system manufacturers and software developers. All product documents are subject to frequent revision, and new order numbers will apply. New documents may be added. Be sure you have the latest information before finalizing your design.

1.2 Product Code and Device Identification

Product Codes: WGI210AT (Commercial Temperature Range); WGI210IT and WGI210IS (Industrial Temperature Range).

The following tables and component drawings describe the various identifying markings oneach device package:

Table 1. Markings

Device	Stepping	Top Marking	Description
I210	A2	WGI210AT	Production (Commercial Copper)
I210	A2	WGI210IT	Production (Industrial Temperature Copper)
I210	A2	WGI210IS	Production (Industrial Temperature Fiber)

Table 2. Device IDs

I210 Device ID Code	Vendor ID	Device ID	Revision ID
Not programmed/factory default	0x8086	0x1531	0x3
WGI210AT/WGI210IT (copper only)	0x8086	0x1533	0x3
WGI210IS (fiber, industrial temperature)	0x8086	0x1536	0x3
WGI210IS (1000BASE-KX/BX backplane)	0x8086	0x1537	0x3
WGI210IS (external SGMII, industrial temperature)	0x8086	0x1538	0x3
WGI210IS (Flash-less 1000BASE-KX/BX backplane)	0x8086	0x157C	0x3
WGI210AT/WGI210IT (Flash-less copper only)	0x8086	0x157B	0x3

Table 3. MM Numbers

Product	MM Number	Spec	Media
WGI210AT - Production (Commercial Copper)	925131	SLJXQ	Tape and Reel
WGI210AT - Production (Commercial Copper)	925132	SLJXR	Tray
WGI210IT - Production (Copper and Industrial Temperature Range)	925133	SLJXS	Tape and Reel
WGI210IT - Production (Copper and Industrial Temperature Range)	925138	SLJXT	Tray
WGI210IS - Production (Fiber and Industrial Temperature Range)	925142	SLJXW	Tape and Reel
WGI210IS - Production (Fiber and Industrial Temperature Range)	925143	SLJXX	Tray

1.3 I210 Production Marking Diagram

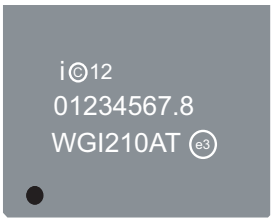

Figure 1. I210 Production Top Marking Example (Commercial Temperature Copper)

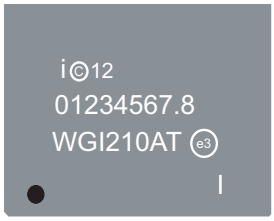

Figure 2. I210 Production Top Marking Example (Industrial Temperature Copper)

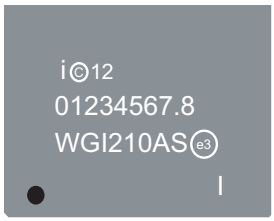

Figure 3. I210 Production Top Marking Example (Industrial Temperature Fiber)

Notes:

Line 1: With no spaces, "i"©YY
Line 2: Fab Lot Trace Code 0123456.78 (10-char max)
Line 3: S-Spec Code and Pb-free mark (e3 or e1)
Line 4: "I" in lower-right corner for industrial temperature rated devices

1.4 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 4 for a description.

Table 4. Terms, Codes, Abbreviations

Name	Description
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 may 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.
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.
Documentation Changes	Typos, errors, or omissions from the current published specifications. These changes will be incorporated in the next release of the specifications.
Yes or No	If the errata applies to a stepping, "Yes" is indicated for the stepping (for example: "A0=Yes" indicates errata applies to stepping A0). If the errata does not apply to stepping, "No" is indicated (for example: "A0=No" indicates the errata does not apply to stepping A0).
Doc	Document change or update that will be implemented.
Fix	This erratum is intended to be fixed in a future stepping of the component.
Fixed	This erratum has been previously fixed.
NoFix	There are no plans to fix this erratum.
Eval	Plans to fix this erratum are under evaluation.
(No mark) or (Blank box)	This erratum is fixed in listed stepping or specification change does not apply to listed stepping.
Red Change Bar / or Bold	This Item is either new or modified from the previous version of the document.
DS	Data Sheet
DG	Design Guide
SDM	Software Developer's Manual
EDS	External Data Specification
AP	Application Note

1.5 Hardware Clarifications, Changes, Updates and Errata

See Section 1.4 for an explanation of terms, codes, and abbreviations.

Table 5. Summary of Hardware Clarifications, Changes and Errata; Errata Include Steppings

Specification Changes	**Status**
1. Proxy: Wake Up on Link Down/Up with MDNS Offload	N/A
2. No Firmware Reset via HICR in Secure Mode	N/A
3. PCIe Timing Parameter Update	N/A
4. Static Device Off Using PCIe Hot Reset	N/A
5. Multicast Listener Discovery (MLD) Protocol Offload is Not Supported	N/A
Specification Clarifications	**Status**
1. PCIe Completion Timeout Mechanism Compliance	N/A
2. Padding on Transmitted SCTP Packets	N/A
3. Dynamic LED Modes Can Only Be Used in an Active Low Configuration	N/A
4. No Match Firmware Proxying Configuration	N/A
5. WUFC/PROXYFC NS Bits	N/A
4. Flash Update Integrity Firmware Enhancements	N/A
6. WUFC/PROXYFC NS Bits	N/A
Documentation Updates	**Status**
1. Port Identification LED Blinking (Word 0x04)	N/A
Errata	**Status**
1. I^2C Data Out Hold Time Violation	A2 NoFix
2. NC-SI Hardware Arbitration Issues	A2 NoFix
3. SGMII: Counters Incorrectly Increment on Collision	A2 NoFix
4. BMC-only Packets Not Counted as Host Sent/Received Packets	A2 NoFix
5. Device Off Deadlock	A2 NoFix
6. Marginal Low 10 Mb Amplitude	A2 NoFix
7. Non-monotonic Integrated SVR Ramp	A2 NoFix
8. Protocol Offload: Incorrect Response to MLDv2 Queries	A2 NoFix
9. Writes to the VPD RW Area are Not Reliable	A2 NoFix
10. NC-SI: Get NC-SI Pass-through Statistics Response Might Contain Incorrect Packet Counts	A2 NoFix
11. MCTP Commands From SMBus are Dropped	A2 Fixed
12. VPD Access During Shadow RAM Load to Flash Causes Firmware Reset and VPD Hang	A2 Fixed
13. NC-SI: Repeated Pause Time After Receiving XOFF	A2 NoFix
14. NC-SI: Set Link Command Failure in Low Power State in SerDes Modes	A2 Fixed
15. NC-SI: Maximum XOFF Renewal Interval Might Be Exceeded	A2 Fixed
16. NC-SI: Set Link and Get Link Status Commands Not Supported in 1000BASE-KX Link Mode	A2 Fixed
17. Proxy: Invalid Neighbor Advertisement Packet with VLAN Tag and SNAP Header	A2 Fixed[1]
18. Failure to Establish PCIe Link After Power Up	A2 NoFix
19. Proxy: Neighbor Solicitation with Multicast Target Address is Not Dropped	A2 Fixed[1]
20. Proxy: Missing Target Link-Layer Address in Neighbor Advertisement	A2 Fixed[1]

Table 5. Summary of Hardware Clarifications, Changes and Errata; Errata Include Steppings

21. NC-SI: Hardware Arbitration Disable is Not Preserved Across Firmware Reset	A2 Fixed
22. NC-SI: Count of Dropped Control Packets Could Be Incorrect	A2 Fixed
23. Transmit Halt After a D3-to-D0 Power State Transition	A2 NoFix
24. Failure of Flash Update from Shadow RAM	A2 Fixed
25. Slow System Clock	A2 NoFix
26. NC-SI: Serdes Link Bit is Clear in Link Status Structure in 1000BASE-KX Link Mode	A2 Fixed
27. Dynamic Device Off is Not Functional	A2 NoFix
28. SMBus: Set Common Filters Command Does Not Set MNGONLY Bit in Shared MAC Address Mode	A2 NoFix

1. I210 Flash-less - A2 NoFix.

1.5.1 Specification Changes

1. Proxy: Wake Up on Link Down/Up with MDNS Offload

As described in the datasheet, when mDNS proxy offload is active, the I210 wakes the system if the LAN link is lost and then re-established. Starting from NVM image release 3.20, the wake-up does not occur unless the link was down for at least 120 seconds. This prevents a spurious wake up triggered by the power state change from D0a to D3 or Dr.

2. No Firmware Reset via HICR in Secure Mode

When the I210 is operating in Secure Mode, the value of the Enable Firmware Reset NVM bit is ignored and firmware reset via the Host Interface Control Register (HICR) is disabled.

This change is implemented starting from NVM image release 3.20.

3. PCIe Timing Parameter Update

In section 5.5.6 of the datasheet, the maximum value of timing parameter $t_{ppg\text{-}clkint}$ (PCIe* PE_RST de-assertion to internal PLL lock) has been updated to 5 ms.

4. Static Device Off Using PCIe Hot Reset

Starting with NVM image release 3.25, the sequence for entering the static device off state can use a PCIe Hot Reset instead of an assertion of the PE_RST_N pin. This change removes the implicit requirement to have a dedicated signal connected to the PE_RST_N pin in order to toggle the pin without resetting the system.

In Section 4.4.4.1 of the datasheet, Step 3 should read, "BIOS issues a PCIe reset, either by asserting and de-asserting the PE_RST_N pin or by generating a PCIe Hot Reset."

Notes:

The sequence for returning from the static device off state cannot use a PCIe Hot Reset since the PCIe link is down in the static device off state.

This change does not apply to flash-less applications.

5. Multicast Listener Discovery (MLD) Protocol Offload is Not Supported

IPv6 Multicast Listener Discovery (MLD) protocol offload is not supported.

The Set Firmware Proxying Configuration host interface command must contain 0 in the Enable MLD field.

1.5.2 Specification Clarifications

1. PCIe Completion Timeout Mechanism Compliance

Clarification:

The I210 Completion Timeout Value[3:0] must be properly set by the system BIOS in the I210 PCIe Configuration Space Device Control 2 register (0xC8; W). Failure to do so can cause unexpected completion timeouts.

The I210 complies with the PCIe 2.0 specification for the completion timeout mechanism and programmable timeout values. The PCIe 2.0 specification provides programmable timeout ranges between 50 µs to 64 s with a default time range of 50 µs - 50 ms. The I210 defaults to a range of 16 ms - 32 ms.

Workaround:

The completion timeout value must be programmed correctly in PCIe configuration space (in Device Control 2 register); the value must be set above the expected maximum latency for completions in the system in which the I210 is installed. This ensures that the I210 receives the completions for the requests it sends out, avoiding a completion timeout scenario. Failure to properly set the completion timeout value can result in the device timing out prior to a completion returning.

The I210 can be programmed to resend a completion request after a completion timeout (the original completion request is assumed to be lost). But if the original completion arrives after a resend request, two completions may arrive for the same request; this can cause unpredictable behavior. Intel NVM images set the resend feature to off. Intel recommends that you do not change this setting.

2. Padding on Transmitted SCTP Packets

Clarification:

When using the I210 to offload the CRC calculation for transmitted SCTP packets, software should not add Ethernet padding bytes to short packets (less than 64 bytes). Instead, the TCTL.PSP bit should be set so that the I210 pads the packets after performing the CRC calculation.

3. Dynamic LED Modes Can Only Be Used in an Active Low Configuration

Clarification:

In any of the dynamic LED modes (FILTER_ACTIVITY, LINK/ACTIVITY, COLLISION, ACTIVITY, PAUSED), LED blinking should only be enabled if the LED signal is configured as an active low output.

4. Flash Update Integrity Firmware Enhancements

Clarification:

The I210 Flash Update integrity feature (Section 3.3.10 of the I210 Datasheet) ensures only Intel digitally signed updates can be applied to I210 products post manufacturing. This is achieved by a combination of hardware and firmware capabilities. NVM image release 3.20 includes firmware enhancements to improve the resilience of this feature.

5. No Match Firmware Proxying Configuration

Clarification:

When the Set Firmware Proxying Configuration command is used and the *No Match Data* field is 0x01, any packet that passes the hardware proxy filters and cannot be processed by the firmware causes a wake up event. Care should be taken when using this setting to prevent the possibility of unintended wake ups.

6. WUFC/PROXYFC NS Bits

Clarification:

The NS and NS Directed bits in both the WUFC and PROXYFC registers enable filters that pass Neighbor Solicitation packets. These filters do not check the ICMPv6 Type field, so they actually pass any ICMPv6 packet that meets all the other requirements. For example, ICMP Echo Request packets can pass these filters. Care should be exercised when setting these bits in WUFC to avoid unintentional system wake-ups.

1.5.3 Documentation Changes

1. Port Identification LED Blinking (Word 0x04)

The default Flash setting for this word can be left at 0xFFFF, which enables software to determine the proper mode of operation(Section 6.8.2 of the I210 Datasheet).

2. Ethernet Controller I210 Supported Flash Parts

Section 11.8.1 of the I210 Datasheet lists a Winbond* validated Flash part as W25Q16DWSSIG. This is incorrect. The correct Winbond validated Flash part is W25Q16CVSSIG.

1.5.4 Errata

1. I^2C Data Out Hold Time Violation

Problem:

The I210 should provide a data out hold time of 50 ns on the SFP_I2C_Data pin. The actual hold time is about 16 ns.

Implication:

Timing specification violation. There have been no reports of failures resulting from this timing. Note that the data input hold time required is zero, so the provided output hold time should be more than enough as long as the I2C CLK and DATA signals are reasonably matched on the board.

Workaround:

None.

Status:

A2 NoFix

2. NC-SI Hardware Arbitration Issues

Problem:

1. During normal operation, the I210 might get FLUSH commands with a smaller ID then the device ID. The I210 should pass on the received FLUSH; but it sends it's own ID for ~2 µs and then passes on the lower ID FLUSHes.
2. The time from received-idle (while in a wait idle state) until the I210 sends IDLE on ARB_OUT is 1.7 µs; the maximum time allowed (by the specification) is T9 = 640 ns.
3. If a token timeout occurs while the I210 waits to send an XON packet, the internal state machine is reset and the XON is never sent.
4. Hardware arbitration timeout mechanism stops upon receiving pause packets from the MC. The timer stops counting until the pause indication drops.
5. The I210 sends XON opcode after the end of the Master Assignment process, even if the XOFF time (~300 ms) has expired.

   If the I210 exits the congested mode during the Master Assignment process it sends XON opcode at the end of the Master Assignment process even if the XOFF has expired.

   The I210 doesn't consider the Master Assignment duration in the XOFF expiration time.
6. When the I210 enters congestion mode it sends XOFF opcode and also makes a request for TOKEN in order to send a XOFF frame.

   When the I210 enters congestion mode it should send a XOFF frame if it has the TOKEN or XOFF opcode even if it hasn’t received the TOKEN.

   The I210 shouldn’t send both of them (opcode and frame) in any case.

Implication:

1. No implication in actual operation. Eventually, the lower IDs pass and arbitration succeeds.
2. The issue is not expected to cause problems because the timeout period is longer. Minor NC-SI compliance violation related to hardware arbitration.
3. The MC is released by the XOFF timer expiration. Minor NC-SI compliance violation related to hardware arbitration.
4. Longer than expected timeout (no specification violation).
5. No implication.
6. Slight delay in traffic coming from the MC but no platform implication.

Workaround:

None.

Status:

A2 NoFix

3. SGMII: Counters Incorrectly Increment on Collision

Problem:

In SGMII mode/half duplex, the following statistics counters incorrectly increment when a collision occurs:

Name	Definition	Location
RLEC	Length error counter.	0x4040
CRCERRS	CRC error counter.	0x4000
RFC	Receive frame counter.	0x40A8

Implication:

Error counters might not be accurate.

Workaround:

None.

Status:

A2 NoFix

4. BMC-only Packets Not Counted as Host Sent/Received Packets

Problem:

When OS2BMC is enabled, packets that do not reach the LAN are not counted as packets sent by the host (HGPTC register). Similarly, packets received from the MC are not counted as packets received by the host (RPTHC register)

Implication:

HGPTC and RPTHC counts are not accurate.

Workaround:

Add the O2BGPTC count to the HGPTC count to get the accurate number of packets sent by the host. Add the B2OGPRC count to the RPTHC count to get the accurate number of packets received by the host.

Status:

A2 NoFix

5. Device Off Deadlock

Problem:

If firmware resets (such as due to a parity error) after entering device off, the I210 does not detect the error and should enter device off but not shut the device down.

This happens only after a firmware reset.

Implication:

The chances of such an event happening while moving to device off are minimal.

Workaround:

None.

Status:

A2 NoFix

6. Marginal Low 10 Mb Amplitude

Problem:

1. 10BASE-T amplitude
   On some designs, the I210 might not meet the IEEE specification that states that the 10 Mb peak differential amplitude be between 2.2V and 2.8V for all data sequences.
2. 10BASE-T TP_IDLE mask failures.
   Some designs might have mask failures on the 10BASE-T TP_IDLE with TPM load.
3. 10 BASE-Te (802.3az section 14.10) amplitude.
   On some designs, the I210 might not meet the IEEE specification that states that when 10BASE-Te is enabled the 10 Mb peak differential amplitude be between 1.54V and 1.96V for all data sequences.
4. 10BASE-Te TP_IDLE and link test pulse waveform mask failures.
   Some designs might have mask failures on the 10BASE-Te TP_IDLE and link test pulse with and without TPM load.

Implication:

No implication on system level performance or interoperability, conformance test only impact.

Workaround:

None.

Status:

A2 NoFix

7. Non-monotonic Integrated SVR Ramp

Problem:

On some designs, both the 0.9V and 1.5V SVR show a non-monotonic start up.

Implication:

No functional impact for systems using an internal SVR, because the system is not vulnerable at the specific time that this non-monotonicity occurs.

Workaround:

None.

Status:

A2 NoFix

8. Protocol Offload: Incorrect Response to MLDv2 Queries

Superseded by Specification Change #5.

9. Writes to the VPD RW Area are Not Reliable

Problem:

VPD write accesses via the PCIe VPD Capability Structure are not always stored in the Flash.

Implication:

VPD writes are not reliable.

Workaround:

Do not use a RW area in the VPD structure. RO areas function correctly.

Status:

Fixed in NVM image release 3.16.

A2 Fixed

10. NC-SI: Get NC-SI Pass-through Statistics Response Might Contain Incorrect Packet Counts

Problem:

The I210 maintains packet counters that are used in the Get NC-SI Pass-through Statistics Response. These counters are cleared by any reset of the port, including the port reset generated by a PCIe reset.

Implication:

If a PCIe reset or port reset has occurred since the previous Get NC-SI Pass-through Statistics Response, the packet count values could be lower than the actual packet counts because the counters were cleared.

Workaround:

The packet counts in the Get NC-SI Pass-through Statistics Response can be used for debug purposes, but they should not be used for maintaining reliable statistics.

Status:

A2 NoFix

11. MCTP Commands From SMBus are Dropped

Problem:

The DMTF MCTP SMBus/I^2C Transport Binding Specification requires that the LSB of the 4th byte of an MCTP over SMBus packet be 1b. Such a packet is dropped by the I210.

Implication:

MCTP over SMBus is not functional since all commands are dropped.

Workaround:

None.

Status:

Fixed in NVM image release 3.16.

A2 Fixed

12. VPD Access During Shadow RAM Load to Flash Causes Firmware Reset and VPD Hang

Problem:

If a VPD read or write access is performed while the firmware is in the process of loading the shadow RAM to the Flash, the firmware hangs. After the firmware watchdog timer expires, the firmware is reset and the VPD access is never completed.

Implication:

- Any manageability configuration from the Manageability Controller (MC) is lost due to the firmware reset.
- No more VPD read or write accesses can be performed until a PCIe reset occurs.

Workaround:

To prevent this scenario, check that EEC.Shadow_modified is 0b before performing a VPD read or write access. If less than 10 ms have passed since the previous VPD write access, it is OK to ignore this bit.

Status:

Fixed in NVM image release 3.20.

A2 Fixed

13. NC-SI: Repeated Pause Time After Receiving XOFF

Problem:

If the I210 receives an XOFF packet from the Manageability Controller (MC) and the next packet is an NC-SI command, the pause timer is restarted when the command is received.

Implication:

The response to the command is delayed until the pause timer expires, which could cause the MC to detect a timeout of the command.

Workaround:

The MC should send an XON packet to explicitly re-enable transmission from the I210 at the end of each congestion event and should not rely on expiration of the pause time in the XOFF packet.

Status:

A2 NoFix

14. NC-SI: Set Link Command Failure in Low Power State in SerDes Modes

Problem:

The I210 checks the Disable 1000 in non-D0a and Disable 100 in non-D0a bits of the PHPM register when determining if the speed(s) requested in a Set Link command are valid in the non-D0a states. If there is a conflict, the command fails with a Set Link Power Mode Conflict status.

This behavior is correct when using the internal PHY, but it is incorrect when using the other link modes.

Implication:

Set Link command is improperly rejected in SerDes modes in low-power states.

Workaround:

Clear the Disable 1000/100 in non-D0a NVM bits in non-copper modes.

Status:

Fixed in NVM image release 3.20.

A2 Fixed

15. NC-SI: Maximum XOFF Renewal Interval Might Be Exceeded

Problem:

When NC-SI flow control is enabled and the MC-to-LAN buffer is congested, the I210 sends XOFF packets to the MC. The NC-SI Specification defines a Max XOFF Renewal Interval after which the XOFF condition must be removed. The I210 violates this specification by continuing to send XOFF packets as long as the congestion condition remains.

Implication:

Unusual congestion on the LAN interface could prevent the MC from communicating with the I210 for extended periods of time.

Workaround:

None.

Status:

Fixed in NVM image release 3.20.

A2 Fixed

16. NC-SI: Set Link and Get Link Status Commands Not Supported in 1000BASE-KX Link Mode

Problem:

When the CTRL_EXT.LINK_MODE is set to 01b (1000BASE-KX), the NC-SI Set Link and Get Link Status commands do not function correctly.

Implication:

The MC cannot properly control the link in 1000BASE-KX link mode.

Workaround:

Use the default link settings defined in the NVM.

Status:

Fixed in NVM image release 3.20.

A2 Fixed

17. Proxy: Invalid Neighbor Advertisement Packet with VLAN Tag and SNAP Header

Problem:

If a Neighbor Solicitation packet is received with a VLAN tag and a SNAP header, the Neighbor Advertisement (NA) response also contains a VLAN tag and a SNAP header. However, the length field in the SNAP header of the NA packet returned by the I210 contains an incorrect value.

Implication:

Neighbor Solicitation proxy offload failure.

Workaround:

Do not use both VLAN tag and SNAP header in a Neighbor Solicitation packet.

Status:

I210 with Flash, fixed in NVM image release 3.20.

I210 Flash-less NoFixTo be fixed in a future SW release.

A2 NoFix

18. Failure to Establish PCIe Link After Power Up

Problem:

If the first de-assertion of PE_RST_N following power-up lasts less than 5 ms, the PCIe PLL might be calibrated incorrectly. When this occurs, the PLL does not lock and the PCIe logic remains in reset until the next power cycle.

Implication:

Failure to establish PCIe link.

Workaround:

Ensure that the duration of the first de-assertion of PE_RST_N after power-up is at least 5 ms.

A firmware workaround for this issue is included in NVM image release 3.25.

Status:

A2 NoFix

19. Proxy: Neighbor Solicitation with Multicast Target Address is Not Dropped

Problem:

According to Section 7.1.1 of RFC 4861, a Neighbor Solicitation packet with a multicast Target Address field should be silently dropped. The I210 accepts and responds to such a packet if the Target Address corresponds to the Solicited Node address provided by the host.

Implication:

No implication when the network is functioning correctly since this is not a valid packet. Reduced immunity to invalid inputs.

Workaround:

None.

Status:

I210 with Flash, fixed in NVM image release 3.25.

I210 Flash-less NoFixA2 NoFix

20. Proxy: Missing Target Link-Layer Address in Neighbor Advertisement

Problem:

If a Neighbor Solicitation packet does not include a source link-layer address option, the Neighbor Advertisement packet sent by I210 in response does not include a target link-layer address option.

Implication:

If the link partner is performing Duplicate Address Detection, the Neighbor Advertisement packet generated by I210 is dropped by the receiver since the target link-layer address is missing. As a result, there could be undetected duplicate addresses on the network.

Workaround:

None.

Status:

I210 with Flash, fixed in NVM image release 3.25.

I210 Flash-less

A2

21. NC-SI: Hardware Arbitration Disable is Not Preserved Across Firmware Reset

Problem:

If NC-SI hardware arbitration is enabled from the NVM and it is disabled by the Select Package command, the hardware arbitration is enabled after a firmware reset.

Implication:

NC-SI interface hang in this situation.

Workaround:

Hardware arbitration should be disabled in the NVM if it is not required. The Select Package command should not be used to disable hardware arbitration.

Status:

Fixed in NVM image release 3.25.

A2 Fixed

22. NC-SI: Count of Dropped Control Packets Could Be Incorrect

Problem:

The NC-SI Control Packets Dropped counter in the Get NC-SI Statistics Response packet does not include control packets that were dropped due to a checksum error.

Implication:

Misleading statistics when debugging.

Workaround:

Add the value of the NC-SI Command Checksum Errors counter to the value of the NC-SI Control Packets Dropped counter when processing a Get NC-SI Statistics Response packet.

Status:

Fixed in NVM image release 3.25.

A2 Fixed

23. Transmit Halt After a D3-to-D0 Power State Transition

Problem:

If EEE is active on a port's transmit path and MANC.KEEP_PHY_LINK_UP is 1b, data transmission from the MAC might halt following a D3-to-D0 power state transmission.

Implication:

Loss of communication over the LAN.

Workaround:

Clear EEER.TX_LPI_EN to disable EEE in the transmit path when going to the D3 state if KEEP_PHY_LINK_UP is 1b.

This workaround is implemented in firmware in NVM image release 3.25.

Status:

A2 NoFix

24. Failure of Flash Update from Shadow RAM

Problem:

If a Flash update from a shadow RAM procedure is performed while there is a management command or proxy packet pending, the Flash update fails and no further updates are performed. This failure is indicated by the value 0x05 in FWSM.Ext_Err_Ind.

During typical operation, this is a low-probability scenario since Flash updates from the shadow RAM are rarely performed and management commands and proxy packets also do not arrive at a high rate.

However, when the Restore MAC Address feature is enabled in the NVM, a Flash update from the shadow RAM is triggered after power-up. If the MC is also polling for the presence of the device at this time, this failure can occur with high probability.

Implication:

Flash updates (by writing EEC.FLUPD) cannot be performed. In Non-Secure Mode the flash can instead be updated directly by software, using the FLSWCTL and FLSWDATA registers.

Workaround:

If the Restore MAC Address feature is enabled in the NVM, the MC should wait 500 ms after power is applied to the I210 before sending commands to the I210.

If this failure is observed, as indicated by FWSM.Ext_Err_Ind, contact your Intel representative.

Status:

Fixed in NVM image release 3.25.

A2 Fixed

25. Slow System Clock

Problem:

On some devices, the internal PLL circuit occasionally provides the wrong clock frequency after power up. The probability of failure is less than one failure per 1000 power cycles. When the failure occurs, the internal clock frequency is around 1/20 of the correct frequency.

The failure can be observed on the NVM_SK output, which will be running at a frequency below 5 MHz.

The failure can be detected from software by either of the following:

- Read internal PHY register 14 from page 252. The failing state is indicated by a value of 0xFF in bits 7:0.
- Read the FRTIMER register twice at a known time interval and see if the difference in the two values matches the interval.

Implication:

No link can be established until the next power cycle.

Workaround:

NVM image release 3.25 contains a workaround in firmware.

For flash-less applications, the following workaround can be performed in software. Note that if a failure occurs, there is no link before the software runs, so APM WoL is not totally reliable in these applications.

1. Set MDICNFG.Destination to 0b.
2. Read PHY register 14 from page 252. If bits 7:0 != 0xFF, go to Step 14. After 5 loops, exit with a fatal error.
3. Set CTRL.PHY_RST to 1b.
4. Set both CTRL_EXT.PHY_Power_Down_Enable and CTRL_EXT.SerDes_Low_Power_Enable to 1b.
5. Clear WUC.
6. Determine the value of auto-load word 0x0A. If this word exists in the iNVM, use that value. Otherwise, use the hardware default value, 0x202F.
7. Perform a bitwise OR of 0x0010 with the value from the Step 6 and write it as an auto-load of 0x0A using EEARBC.
8. Set PCIe configuration space register PMCSR bits 1:0 to 11b. (D3 state).
9. Wait 1 ms.
10. Set PCIe configuration space register PMCSR bits 1:0 to 00b. (D0 state).
11. Write the value from Step 6 as an auto-load of 0x0A using EEARBC.
12. Restore WUC to its original value, if necessary.
13. Go to Step 2.
14. Restore MDICNFG.Destination to its original value, if necessary.

Status:

A2 NoFix

26. NC-SI: Serdes Link Bit is Clear in Link Status Structure in 1000BASE-KX Link Mode

Problem:

The response to a Get Link Status command in 1000BASE-KX link mode (CTRL_EXT.Link_Mode = 01b) has the SerDes Link bit set to 0b when it should be set to 1b. The same applies to a Link Status Change AEN.

Implication:

Incorrect indication of link mode. Other fields in the response might be interpreted incorrectly as a result.

Workaround:

The MC should use another method to determine the link mode.

Status:

Fixed in NVM image release 3.25.

A2 Fixed

27. Dynamic Device Off is Not Functional

Problem:

The I210 does not actually enter the Dynamic Device Off state even if all the necessary conditions are satisfied.

This does not apply to flash-less applications.

Implication:

Power consumption is higher than expected.

Workaround:

None.

Status:

A2 NoFix

28. SMBus: Set Common Filters Command Does Not Set MNGONLY Bit in Shared MAC Address Mode

Problem:

When executing a Set Common Filters command with the CBDM bit set to 0b (Shared MAC Address), the I210 uses the MDEF7 register to enable the IP address and/or port number filters specified in the command, but it does not set bit 7 of the MNGONLY register.

Implication:

Management traffic is duplicated and forwarded to the host in addition to the MC.

Workaround:

After sending the Set Common Filters command, the MC should send an Update Management Receive Filter Parameters command (Parameter Number 0xF) to set bit 7 of the MNGONLY register.

Status:

A2 NoFix

2. Software Clarifications

Table 6. Summary of Software Clarifications

Software Clarifications	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

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

Clarification:

The I210 supports 256 KB TCP packets; however, each buffer is limited to 64 KB since the data length field in the transmit descriptor is only 16 bits. This restriction increases driver implementation complexity if the operating system passes down a scatter/gather element greater than 64 KB in length. This 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 size in all drivers to 64 KB.

Please note that Linux operating systems only support 64 KB data transfers.

For further details about how Intel network software drivers address this issue, refer to Technical Advisory TA-191.

2. Serial Interfaces Programmed By Bit Banging

Clarification:

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 second write. To prevent such problems, a register read should be inserted between the first register write and the software delay, i.e. “write”, “read”, “software delay”, “write”.