



Intel[®] Ethernet Controller I350

Specification Update

Networking Division (ND)

March 2016

Revision 2.7
333066-002



Revision History

Revision	Date	Comments
2.7	March 23, 2016	Errata added or updated: <ul style="list-style-type: none"> 38. PCIe Advanced Error Reporting: First Error Pointer (Added)
2.6	September 2, 2015	Errata added or updated: <ul style="list-style-type: none"> 35. NC-SI Output Signals Have Indeterminate Value After Power Up (Added) 36. PCIe SR-IOV Reserved Bits are Writable (Added) 37. Certain Malformed IPv6 Extension Headers are not Processed Correctly by the Device (Added)
2.5	June 23, 2015	Errata added or updated: <ul style="list-style-type: none"> Removed Erratum #26 (duplicate of Erratum #22).
2.4	January 22, 2014	Miscellaneous Updates <ul style="list-style-type: none"> Document converted to Intel standard template. Spec Clarifications added or updated: <ul style="list-style-type: none"> 10. WUFC/PROXYFC NS Bits (Added) 11. "No Match" Firmware Proxying Configuration (Added) 12. SR-IOV Prefetchable Address Space (Added) Spec Changes added or updated: <ul style="list-style-type: none"> 9. SMBus: ARP Auto-Reply Does Not Support SNAP Headers (Added) Errata added or updated: <ul style="list-style-type: none"> 31. VFMPRC is Not Accessible from the VF Memory Space (Added) 32. NC-SI: Count of Dropped Control Packets Could Be Incorrect (Added) 33. Transmit Halt After a D3-to-D0 Power State Transition (Added) 34. Common MDIO Failure when Port 0 is Disabled or Powered Down (Added)
2.3	May 20, 2013	Miscellaneous Updates: <ul style="list-style-type: none"> Updated Table 1-1, "Product and Device Identification"
2.2	March 13, 2013	Spec Clarifications added or updated: <ul style="list-style-type: none"> 9. TimeSync: Ensure That the Programmed Target Time is in the Future (Added) Spec Changes added or updated: <ul style="list-style-type: none"> 6. Overriding MAIN_PWR_OK to Maintain Manageability Session Through a Power Cycle (Updated) 7. Update to SerDes Loopback Configuration (Added) 8. EEPROM Image Release (Added) Errata added or updated: <ul style="list-style-type: none"> 30. NC-SI: Get NC-SI Pass-through Statistics Response Might Contain Incorrect Packet Counts (Added)



Revision	Date	Comments
2.1	November 30, 2012	<p>Spec Clarifications added or updated:</p> <ul style="list-style-type: none"> 3. PCIe Completion Time-out Mechanism Compliance (Updated) <p>Spec Changes added or updated:</p> <ul style="list-style-type: none"> 1. Updates to PXE/iSCSI EEPROM Words (Updated) 6. Overriding MAIN_PWR_OK to Maintain Manageability Session Through a Power Cycle (Added) <p>Errata added or updated:</p> <ul style="list-style-type: none"> 6. NC-SI: Hardware Arbitration Issues (Updated) 27. Enabling a Port Causes a Reset of the Firmware (Added) 28. MCTP: Interface is Not Functional in D3/Dr (Added) 29. MCTP: Command with Duplicate NC-SI IID is Ignored (Added)
2.06	September 12, 2012	<p>Spec Clarifications added or updated:</p> <ul style="list-style-type: none"> 7. ECC Checking of Management RAM is Disabled During PCIe Reset (Added) 8. Dynamic LED Modes Can Only Be Used in an Active Low Configuration (Added) <p>Spec Changes added or updated:</p> <ul style="list-style-type: none"> 5. I²C Timing Parameter Correction (Added) <p>Errata added or updated:</p> <ul style="list-style-type: none"> 22. Firmware Clears Flow Control Setting (Added) 23. PF's MSI TLP Might Contain the Wrong Requester ID When a VF Uses MSI-X (Added) 24. NC-SI: Set Link Command Fails in 1000Base-KX Mode When Not in D0a State (Added) 25. Failures When Functions are Swapped and Ports are Disabled (Added) 26. Removed (Duplicate of Erratum #22) (Added) <p>Software Clarifications added or updated:</p> <ul style="list-style-type: none"> 3. PF/VF Drivers Should Configure Registers that are not Reset By VFLR (Added)
2.05	April 10, 2012	<p>Errata added or updated:</p> <ul style="list-style-type: none"> 21. Firmware Reset with Loss of Manageability Configuration when Using SerDes (Added)
2.04	January 30, 2012	<p>Spec Clarifications added or updated:</p> <ul style="list-style-type: none"> 5. Padding on Transmitted SCTP Packets (Added) 6. MCTP/DMTF Standard Compliance (Added) <p>Errata added or updated:</p> <ul style="list-style-type: none"> 20. VFTA Register Access Might Not Be Performed (Added)
2.03	December 14, 2011	<p>Errata added or updated:</p> <ul style="list-style-type: none"> 18. Alternate MAC Address Port Order (Added) 19. Neighbor Discovery Offload: Override Bit Not Set in Neighbor Advertisement Packet (Added)
2.02	September 14, 2011	<p>Spec Clarifications added or updated:</p> <ul style="list-style-type: none"> 4. Malicious Driver Detection (Added)
2.01	August 18, 2011	<p>Spec Changes added or updated:</p> <ul style="list-style-type: none"> 4. Minimum Value for Flow Control Receive Threshold Low (Added) <p>Software Clarifications added or updated:</p> <ul style="list-style-type: none"> 2. Serial Interfaces Programmed by Bit Banging (Added)



Revision	Date	Comments
2.00	June 21, 2011	<p>Miscellaneous Updates:</p> <ul style="list-style-type: none"> Device Information <ul style="list-style-type: none"> Table 1-1 updated. Combines data in several of the older tables. <p>Errata added or updated:</p> <ul style="list-style-type: none"> 17. IEEE Std 802.3™-2008 Tx Distortion Marginality (Updated)
1.92	May 27, 2011	<p>Spec Clarifications added or updated:</p> <ul style="list-style-type: none"> 3. PCIe Completion Time-out Mechanism Compliance (Added)
1.91	May 6, 2011	<p>Spec Changes added or updated:</p> <ul style="list-style-type: none"> 3. Updated Definition of SW EEPROM Port Identification LED Blinking (Word 0x4) (Added)
1.9	April 6, 2011	<p>Miscellaneous Updates:</p> <ul style="list-style-type: none"> PRQ version (by convention revision 1.9). <p>Spec Clarifications added or updated:</p> <ul style="list-style-type: none"> 2. SR-IOV and DMA Coalescing Cannot Be Set Together (Added) <p>Errata added or updated:</p> <ul style="list-style-type: none"> 5. HOST<-->BMC Only Packets not Counted as Host Sent/Received Packets (Updated — Previously combined under another errata; broken out.) 10. Rx Packet May Be Dropped After Receiving 2 Bit Errors While in an Idle State (updated) 11. SDPx_1 Remains an Input in D3 (Added) 12. NC-SI Enable Global Multicast Filter Command—No Differentiation of Neighbor Discovery Modes (Added) 13. NC-SI: Select Package Command Increments NC-SI Commands Received Counter for Ports 0 and 1 and Not for Ports 2 and 3 (Added) 14. VLAN Tags Byte-Swapped in Loopback Packets (Added) 15. DEV_OFF_N May Not Properly Disable the Device (Added) 16. DMA Coalescing Reacts to MSI-X Vectors Instead of Interrupt Causes (Added) 17. IEEE Std 802.3™-2008 Tx Distortion Marginality (Added)
1.1	February 27, 2011	<p>Miscellaneous Updates:</p> <ul style="list-style-type: none"> Brand string in title updated. <p>Spec Clarifications added or updated:</p> <ul style="list-style-type: none"> 1. SerDes: AN_TIMEOUT Only Works When Link Partner Idle (Updated) <p>Spec Changes added or updated:</p> <ul style="list-style-type: none"> 1. Updates to PXE/iSCSI EEPROM Words (Added) 2. Updates to Software Compatibility EEPROM Word 0x3 (Added) <p>Errata added or updated:</p> <ul style="list-style-type: none"> 4. SGMII: Counters Incorrectly Increment on Collision (Updated) 6. NC-SI: Hardware Arbitration Issues (Updated — Combined six hardware arbitration issues into one master entry. The issues combined are related and minor.) 9. TSYNC: Auxiliary Timestamp from SDP is Unreliable (Updated) <p>Software Clarifications added or updated:</p> <ul style="list-style-type: none"> 1. While in TCP Segmentation Offload, Each Buffer is Limited to 64 KB (Added — Moved to Software Clarifications.)
1.0	January 18, 2011	Initial release



1. Introduction

This document applies to all the I350 variants. The information supplements data provided in the product datasheet.

1.1 Product Code and Device Identification

Product Code: I350

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

Table 1-1 Product and Device Identification

MM#	Top Marking	Spec #	Status & Media	Description	Step
915808	NHI350AM4	S LJ2Z	Tape & Reel	1 Gbs, 4-port, 17x17	A1
915799	NHI350AM4	S LJ3Z	Tray		A1
915809	NHI350AM2	S LJ3R	Tape & Reel	1 Gbs 2-port, 17x17	A1
915801	NHI350AM2	S LJ3S	Tray		A1
915810	NHI350AS4	S LJ3T	Tape & Reel	1 Gbs 4-port, 17x17 SERDES	A1
915802	NHI350AS4	S LJ3U	Tray		A1
915811	NHI350BT2	S LJ3X	Tape & Reel	1 Gbs, 2-port, 25x25	A1
915804	NHI350BT2	S LJ3Y	Tray		A1
913157	NHI350AM4	Q P3L	Engineering, Tray	1 Gbs, 4 port 17x17	A1
913159	NHI350AM2	Q P3R	Engineering, Tray	1 Gbs, 2-port 17x17	A1
913193	NHI350AS4	Q P5A	Engineering, Tray	1 Gbs, 4-port 17x17 SERDES	A1
913194	NHI350BT2	Q P5B	Engineering, Tray	1 Gbs, 2-port 25x25	A1

Table 1-2 Device ID

Device	Vendor ID	Device ID
I350 (EEPROM-Less)	0x8086	0x151F
I350 (SerDes (KX/BX))	0x8086	0x1523
I350 (SGMII)	0x8086	0x1524
I350 (Copper)	0x8086	0x1521
I350 (Fiber)	0x8086	0x1522

1.2 Marking Diagrams



Figure 1-1 Example Component with Identifying Marks

Table 1-3 Marking Diagram Legend

Line Number	Description
1	"intel"
2	Marketing Name
3	Fab Lot Number "XXXXXXXX" (Wafer Lot no. concatenated with Assembler vendor code)
4	Assembly Date Code YYWW; Engineering samples have additional Intel data.
5	Copyright line; includes two number date code and the Pb-free mark (e1)
6	Country of Origin



1.3 Nomenclature Used in This Document

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

Table 1-4 Nomenclature

Name	Description
Specification Clarifications	Greater detail or further highlights concerning a specification's impact to a complex design situation. These clarifications will be incorporated in the next release of the specifications.
Specification Changes	Modifications to the current published specifications. These changes will be incorporated in the next release of the specifications.
Errata	Design defects or errors. Errata 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.
Software Clarifications	Applies to Intel drivers, EEPROM loads.
Documentation Changes	Errors, or omissions in current published specifications. These changes are incorporated in the next release of the applicable document and then dropped from the spec update. You may also check for changes in the revision history of specific documents.
A1, B1, etc.	Stepping to which the status applies.
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.
EEPROM/NVM Fix	This indicates the Errata was in the EEPROM/NVM and is fixed in an updated version.
Eval	Plans to fix this erratum are under evaluation.
Red Change Bar/or Bold	This Item is either new or modified from the previous version of the document.



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. SerDes: AN_TIMEOUT Only Works When Link Partner Idle	N/A
2. SR-IOV and DMA Coalescing Cannot Be Set Together	N/A
3. PCIe Completion Time-out Mechanism Compliance	N/A
4. Malicious Driver Detection	N/A
5. Padding on Transmitted SCTP Packets	N/A
6. MCTP/DMTF Standard Compliance	N/A
7. ECC Checking of Management RAM is Disabled During PCIe Reset	N/A
8. Dynamic LED Modes Can Only Be Used in an Active Low Configuration	N/A
9. TimeSync: Ensure That the Programmed Target Time is in the Future	N/A
10. WUFC/PROXYFC NS Bits	N/A
11. "No Match" Firmware Proxying Configuration	N/A
12. SR-IOV Prefetchable Address Space	N/A

Table 2-2 Summary of Specification Changes

Specification Change	Status
1. Updates to PXE/iSCSI EEPROM Words	N/A
2. Updates to Software Compatibility EEPROM Word 0x3	N/A
3. Updated Definition of SW EEPROM Port Identification LED Blinking (Word 0x4)	N/A
4. Minimum Value for Flow Control Receive Threshold Low	N/A
5. I ² C Timing Parameter Correction	N/A
6. Overriding MAIN_PWR_OK to Maintain Manageability Session Through a Power Cycle	N/A
7. Update to SerDes Loopback Configuration	N/A
8. EEPROM Image Release	N/A
9. SMBus: ARP Auto-Reply Does Not Support SNAP Headers	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. I ² C Data Out Hold Time Violation	A1=Yes; NoFix
2. ECRC Capability Not Set in VF Configuration Space	A1=Yes; NoFix
3. RQDPC Register is RC (Read-Clear)	A1=Yes; NoFix
4. SGMII: Counters Incorrectly Increment on Collision	A1=Yes; NoFix
5. HOST<-->BMC Only Packets not Counted as Host Sent/Received Packets	A1=Yes; NoFix
6. NC-SI: Hardware Arbitration Issues	A1=Yes; NoFix
7. Entering D3 State May Be Delayed	A1=Yes; NoFix
8. PCIe Gen2 Transmitter V _{TX-AC-CM-PP} Max Value Violation	A1=Yes; NoFix
9. TSYNC: Auxiliary Timestamp from SDP is Unreliable	A1=No; Fixed
10. Rx Packet May Be Dropped After Receiving 2 Bit Errors While in an Idle State	A1=Yes; NoFix
11. SDPx_1 Remains an Input in D3	A1=Yes; NoFix
12. NC-SI Enable Global Multicast Filter Command—No Differentiation of Neighbor Discovery Modes	A1=Yes; NoFix
13. NC-SI: Select Package Command Increments NC-SI Commands Received Counter for Ports 0 and 1 and Not for Ports 2 and 3	A1=Yes; NoFix
14. VLAN Tags Byte-Swapped in Loopback Packets	A1=Yes; NoFix
15. DEV_OFF_N May Not Properly Disable the Device	A1=Yes; NoFix
16. DMA Coalescing Reacts to MSI-X Vectors Instead of Interrupt Causes	A1=Yes; NoFix
17. IEEE Std 802.3™-2008 Tx Distortion Marginality	A1=Yes; NoFix
18. Alternate MAC Address Port Order	A1=No; EEPROM/NVM Fix
19. Neighbor Discovery Offload: Override Bit Not Set in Neighbor Advertisement Packet	A1=No; EEPROM/NVM Fix
20. VFTA Register Access Might Not Be Performed	A1=Yes; NoFix
21. Firmware Reset with Loss of Manageability Configuration when Using SerDes	A1=No; EEPROM/NVM Fix
22. Firmware Clears Flow Control Setting	A1=No; EEPROM/NVM Fix
23. PF's MSI TLP Might Contain the Wrong Requester ID When a VF Uses MSI-X	A1=Yes; NoFix
24. NC-SI: Set Link Command Fails in 1000Base-KX Mode When Not in D0a State	A1=Yes; NoFix
25. Failures When Functions are Swapped and Ports are Disabled	A1=No; EEPROM/NVM Fix
26. Removed (Duplicate of Erratum #22)	N/A
27. Enabling a Port Causes a Reset of the Firmware	A1=No; EEPROM/NVM Fix
28. MCTP: Interface is Not Functional in D3/Dr	A1=No; EEPROM/NVM Fix



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

Erratum	Status
29. MCTP: Command with Duplicate NC-SI IID is Ignored	A1=No; EEPROM/NVM Fix
30. NC-SI: Get NC-SI Pass-through Statistics Response Might Contain Incorrect Packet Counts	A1=Yes; NoFix
31. VFMPRC is Not Accessible from the VF Memory Space	A1=Yes; NoFix
32. NC-SI: Count of Dropped Control Packets Could Be Incorrect	A1=Yes; NoFix
33. Transmit Halt After a D3-to-D0 Power State Transition	A1=Yes; NoFix
34. Common MDIO Failure when Port 0 is Disabled or Powered Down	A1=Yes; NoFix
35. NC-SI Output Signals Have Indeterminate Value After Power Up	A1=Yes; NoFix
36. PCIe SR-IOV Reserved Bits are Writable	A1=Yes; NoFix
37. Certain Malformed IPv6 Extension Headers are not Processed Correctly by the Device	A1=Yes; NoFix
38. PCIe Advanced Error Reporting: First Error Pointer	A1=Yes; NoFix



2.1 Specification Clarifications

1. SerDes: AN_TIMEOUT Only Works When Link Partner Idle

Clarification:

The auto-negotiation time-out mechanism (PCS_LCTL.AN_TIMEOUT_EN) only works if the SerDes partner is sending idle code groups continuously for the duration of the time-out period (the usual case).

If the partner is transmitting packets in auto-negotiation, time-out will not occur since auto-negotiation is restarted at the beginning of each packet. If the partner has an application that indefinitely transmits data in spite of the lack of response, it is possible that link will not be established.

Workaround:

If this is a concern, the auto-negotiation time-out mechanism may be considered unreliable and an additional software mechanism can be used to disable auto-negotiation (if sync is maintained without a link being established (PCS_LSTS.SYNC_OK=1 and PCS_LSTS.LINK_OK=0) for an extended period of time).

2. SR-IOV and DMA Coalescing Cannot Be Set Together

Clarification:

DMA coalescing assumes a specific mapping between interrupt causes and MSI-X vectors. This mapping doesn't match the distribution of MSI-x vectors to Virtual functions.

Workaround:

DMA coalescing and SR-IOV cannot be activated together.

3. PCIe Completion Time-out Mechanism Compliance

Clarification:

The I350 Completion Timeout Value[3:0] must be properly set by the system BIOS in the I350 PCIe Configuration Space Device Control 2 register (0xC8; W). Failure to do so can cause unexpected completion time-outs.

The I350 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 μ s to 64 s with a default time range of 50 μ s to 50 ms. The I350 defaults to a range of 16 ms to 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 I350 is installed. This ensures that the I350 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.



4. Malicious Driver Detection

Clarification:

Section 7.8.3.8.3 of the Datasheet [search for “Interrupt on Misbehavior of VM (Malicious Driver Detection)”] describes the actions of the device when an error that might be caused by a malicious driver is detected. These actions are usually sufficient to keep the PF and the other VFs functioning, but there are situations where the transmit DMA function may hang as a result of malicious action by a VF driver.

Workaround:

We recommend that the following actions be performed by the PF driver upon receipt of an MDDET interrupt to ensure that the PF and the VFs are not blocked:

- Clear VFTE and VFRE bits for the offending VF.
- Set VTCTRL.RST of the offending VF.

The above actions were implemented in Intel igb driver v3.1.16

5. Padding on Transmitted SCTP Packets

Clarification:

When using the I350 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 I350 pads the packets after performing the CRC calculation.

6. MCTP/DMTF Standard Compliance

Clarification:

The I350 MCTP protocol implementation is based on the DMTF DSP0236, DSP0237 and DSP0239 standards.

The I350 NC-SI over MCTP implementation is described in Chapter 10 of the Datasheet.

7. ECC Checking of Management RAM is Disabled During PCIe Reset

Clarification:

The PCIEERRCTL.GPAR_EN (Global Parity Enable) bit is cleared during PCIe reset. As a result, the internal memory modules used by the manageability function are not covered by ECC checking during PCIe reset.

8. 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.



9. TimeSync: Ensure That the Programmed Target Time is in the Future

Clarification:

When using the target time feature of the TimeSync logic, software should ensure that the target time being programmed is always in the future. For example, the target time should be larger than the system time except for when the intention is for the target time to be after the system time has wrapped. Unintentionally programming a target time that is smaller than the system time disables the target time functionality until the system time has wrapped.

10. 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.

11. “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.

12. SR-IOV Prefetchable Address Space

Clarification:

In SR-IOV mode, memory space should be allocated to the multiple VFs enabled. To accommodate the full extent of possible memory allocation, 64-bit addressing should be used. The PCI bridge specification requires that a 64-bit BAR be prefetchable.

The “prefetchable” bit at IOV Control Word has been set in the NVM Dev Starter releases since Revision 1.59.



2.2 Specification Changes

1. Updates to PXE/iSCSI EEPROM Words

Change:

Word 0x30, 34, 38, 3A are now defined as follows:

Bit(s)	Value	Port Status	CLP (Combo) Executes	iSCSI Boot Option ROM CTRL-D Menu	FCoE Boot Option ROM CTRL-D Menu
2:0	0	PXE	PXE	Displays port as PXE. Allows changing to Boot Disabled, iSCSI Primary or Secondary.	Displays port as PXE. Allows changing to Boot Disabled, FCoE Enabled.
	1	Boot Disabled	NONE	Displays port as Disabled. Allows changing to iSCSI Primary/Secondary.	Displays port as Disabled. Allows changing to FCoE enabled.
	2	iSCSI Primary	iSCSI	Displays port as iSCSI Primary. Allows changing to Boot Disabled, iSCSI Secondary.	Displays port as iSCSI. Allows changing to Boot Disabled, FCoE Enabled.
	3	iSCSI Secondary	iSCSI	Displays port as iSCSI Secondary. Allows changing to Boot Disabled, iSCSI Primary.	Displays port as iSCSI. Allows changing to Boot Disabled, FCoE Enabled.
	4	FCoE	FCOE	Displays port as FCoE. Allows changing port to Boot Disabled, iSCSI Primary or Secondary.	Displays port as FCoE. Allows changing to Boot Disabled.
	5-7	Reserved	Same as Disabled	Same as Disabled.	Same as Disabled.
4:3	Same as before.				
5	Bit 5, formerly used to indicate iSCSI enable/disable, is no longer valid and is not checked by software.				
15:7	Same as before.				

2. Updates to Software Compatibility EEPROM Word 0x3

Change:

Bit	Name	Default	Description
1:0	Reserved	0	Reserved
2	Reserved	0	Reserved
3	Reserved	0	Reserved
4	Reserved	0	Reserved



Bit	Name	Default	Description
7:5	Reserved	0	Reserved
8	Reserved	0	Reserved.
9	Client	0	Client/Not a Client NIC 0 = Server. 1 = Client. This bit is used by DMIX to verify the NIC is server or client. Server NIC or LOM required.
10	Reserved	0	Reserved
11	LOM	0	Indicates whether dedicated flash for the option ROM is attached to LAN silicon. Used by option ROM update applications, QV and DMIX. 0 = NIC (A dedicated flash is attached) 1 = LOM (No dedicated flash is attached)
12	Reserved	0	Reserved
14:13	Reserved	0	Reserved
15	Reserved	0	Reserved

3. Updated Definition of SW EEPROM Port Identification LED Blinking (Word 0x4)

Change:

Driver software provides a method to identify an external port on a system through a command that causes the LEDs to blink. Based on the setting in word 0x4, the LED drivers should blink between STATE1 and STATE2 when a port identification command is issued.

When word 0x4 is equal to 0xFFFF or 0x0000, the blinking behavior reverts to a default.

Bit	Description
15:12	Control for LED 3 0000b = Default LED Blinking operation is used. 0001b = Default in STATE1 + Default in STATE2. 0010b = Default in STATE1 + LED is ON in STATE2. 0011b = Default in STATE1 + LED is OFF in STATE2. 0100b = LED is ON in STATE1 + Default in STATE2. 0101b = LED is ON in STATE1 + LED is ON in STATE2. 0110b = LED is ON in STATE1 + LED is OFF in STATE2. 0111b = LED is OFF in STATE1 + Default in STATE2. 1000b = LED is OFF in STATE1 + LED is ON in STATE2. 1001b = LED is OFF in STATE1 + LED is OFF in STATE2. 1111b = Default LED Blinking operation is used. All other values are Reserved.
11:8	Control for LED 2—same encoding as for LED 3.
7:4	Control for LED 1—same encoding as for LED 3.
3:0	Control for LED 0—same encoding as for LED 3.



4. Minimum Value for Flow Control Receive Threshold Low

Change:

If FCRTL0.XONE is 1, the minimum value allowed in FCRTL0.RTL is 3 (48 bytes).

5. I²C Timing Parameter Correction

Change:

In the I2C Timing Parameters table, the values of TSU:STA and TSU:STO should be 0.6 μs (Min).

6. Overriding MAIN_PWR_OK to Maintain Manageability Session Through a Power Cycle

Change:

New definition of bit 8 of the Common Firmware Parameters 2 EEPROM word:

Bit	Name	Default	Description
8	Restore KEEP_PHY_LINK_UP Disable	1b	0b = Restore the KEEP_PHY_LINK_UP bit in the MANC_CSR according to the value maintained in firmware. 1b = Legacy behavior (KEEP_PHY_LINK_UP is cleared by MAIN_PWR_OK and is not restored by firmware.)

When the MC wants to maintain a critical manageability session through a reset cycle, it uses a command to set the MANC.KEEP_PHY_LINK_UP bit.

The default behavior of the device is to clear MANC.KEEP_PHY_LINK_UP at the falling edge of the MAIN_PWR_OK input pin. This is intended to prevent the device from drawing too much power from the AUX power rail by reducing the link speed when main power is removed.

In a situation where the AUX power rail provides enough power to maintain one or more 1 GbE links, it might be desirable to maintain a link through a main power cycle in order to keep a manageability session active. This can be done by pulling the MAIN_PWR_OK pin to the AUX power supply and setting the MANC.KEEP_PHY_LINK_UP bit through the SMBus or NC-SI interface.

If the board design has MAIN_PWR_OK pulled to the main power supply, another option is now (starting with EEPROM version 1.63) available to prevent the Ethernet link from going down during the main power cycle.

The necessary settings are:

- Clear the LPLU EEPROM bit for the port (word offset 0x13, bit 6).
- Program the Disable 1000 in non-D0a and Giga Disable EEPROM bits for the port to the same value (word offset 0x20, bits 4:3).
- Clear the Restore KEEP_PHY_LINK_UP Disable EEPROM bit previously defined.
- Set the MANC.KEEP_PHY_LINK_UP bit through the SMBUS or NC-SI interface on the port that should stay linked during reboot.



7. Update to SerDes Loopback Configuration

Change:

1. Add the following definition to the PCS_CFG register (0x4200):

Field	Bit	Initial Value	Description
IGN_SD	0	0b	Ignore Signal Detection 0 = Attempt to establish a link only when a signal is present. 1 = Ignore the signal detection logic and always attempt to establish a link. This bit should be 0b for normal operation. It should be 1b for SerDes loopback operation.

2. Add the following to the procedure for setting SerDes Loopback Mode:
—PCS_CFG.IGN_SD (CSR 0x4200 bit 0) = 1b

Intel LAN drivers and tools implemented Step 2 in Release 18.0 (igb v4.1.2).

8. EEPROM Image Release

Change:

The I350 should provide a data out hold time of 50 ns on the SFPx_I2C_DATA pins. The actual hold time is about 16 ns.

9. SMBus: ARP Auto-Reply Does Not Support SNAP Headers

Change:

The ARP auto-reply functionality does not recognize ARP Request packets that contain an LLC/SNAP header.

2.3 Documentation Updates

None.



2.4 Errata

1. I²C Data Out Hold Time Violation

Problem:

The I350 should provide a data out hold time of 50 ns on the SFPx_I2C_DATA pins. 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: A1=Yes; NoFix

2. ECRC Capability Not Set in VF Configuration Space

Problem:

According to SR-IOV v1.1 specification, the ECRC capability reporting bits (Advanced Error Capabilities and Control Register; bits 5 & 7) should be set in VF if the device supports ECRC generation and checking.

Although the I350 supports ECRC generation and checking, the capabilities are reported in the PF and not in the VF configuration space. In the VF configuration space, these bits are set to zero.

Implication:

ECRC capabilities should be read from PF. The device's behavior (e.g. ECRC generate and check) is correct.

Workaround:

None.

Status: A1=Yes; NoFix

3. RQDPC Register is RC (Read-Clear)

Problem:

The RQDPC register should not be cleared by a read. It is RC for VF and PF.

Implication:

The RQDPC register may be cleared by a VF driver or a PF driver and thus can not be used by both the PF and the VF to gather receive-queue-drop statistics. Only one of the drivers may use it.



Workaround:

The VF driver may request the receive-queue-drop statistics from the PF.

Status: A1=Yes; NoFix

4. SGMII: Counters Incorrectly Increment on Collision

Problem:

In SGMII mode/half duplex, the statistics counters listed below 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 may not be accurate.

Workaround:

None.

Status: A1=Yes; NoFix

5. HOST<-->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: A1=Yes; NoFix



6. NC-SI: Hardware Arbitration Issues

1. Receiving a FLUSH Smaller than the I350 ID Causes Issue.
2. Wait IDLE State Violation.
3. Wait-Idle-to-Normal Operation Time Violation.
4. Timeout Mechanism Stops Upon Receiving Pause Packets from MC.
5. XON Not Sent When Expected.

Problem:

1. In the middle of normal operation, the device may get FLUSH commands with a smaller ID than the Device ID. The device should pass on the received FLUSH; but it sends its own ID for $\sim 2 \mu\text{s}$ and then passes on the lower ID FLUSHs.
2. While the device is in Wait Idle state, it sends IDLE on ARB_OUT even though it did not get anything from the ARB_IN pin.
3. The time from Received-Idle (while in Wait Idle state) until the device sends IDLE on ARB_OUT is $1.7 \mu\text{s}$; the max time allowed (by spec) is $T9=640 \text{ ns}$.
4. Hardware arbitration timeout mechanism stops upon receiving pause packets from the MC. The timer stops counting until the pause indication drops.
5. If a token timeout occurs while the device waits to send an XON packet, the internal state machine is reset and the XON is never sent.

Implication:

1. No implication in actual operation. Eventually, the lower IDs pass and arbitration succeeds.
2. No effect on the functionality of HW-ARB. If the device doesn't receive valid op-codes from ARB-IN, this means that the HW-ARB loop is broken and the system should be moved to SW arbitration; so IDLE op-codes generated in such a situation are irrelevant.
3. The issue is not expected to cause problems because the timeout period is longer. Minor NC-SI compliance violation related to hardware arbitration.
4. Longer than expected timeout (no spec violation).
5. The MC is released by the XOFF timer expiration. Minor NC-SI compliance violation related to hardware arbitration.

Workaround:

None.

Status: A1=Yes; NoFix

7. Entering D3 State May Be Delayed

Problem:

If one of the functions driver accesses the flash and this access is delayed, a request to enter another function (or this one) to D3 will be delayed until the flash transaction is done. In some extreme cases (internal FW is erasing a sector and a function tries to access the flash), this may delay the D3 entrance by hundreds of ms.



Implication:

Entering D3 state may be delayed.

Workaround:

None.

Status: A1=Yes; NoFix

8. PCIe Gen2 Transmitter $V_{TX-AC-CM-PP}$ Max Value Violation

Problem:

According to PCIe 2.1 spec the maximum value for VTX-AC-CM-PP (Tx AC common mode voltage (5.0 GT/s)) is 100 mV. In some cases, the I350 may violate this maximum value. The worst case expected value is 114 mV.

Implication:

Specification violation.

Workaround:

None.

Status: A1=Yes; NoFix

9. TSYNC: Auxiliary Timestamp from SDP is Unreliable

Problem:

The SDP inputs to the timestamp logic are not properly synchronized when using AUX1. As a result, the Auxiliary Timestamp Register values and the Auxiliary Timestamp Taken bits in TSAUXC are sometimes loaded incorrectly when using AUX1

Implication:

The auxiliary timestamp feature should be considered unreliable when using AUX1. Use AUX0 whenever possible.

Workaround:

If AUX1 is used (2 aux timestamps on the same port) in applications that use the auxiliary timestamp feature to synchronize to an external clock, it might be acceptable to drop some of the samples. For such applications, software can filter out many of the incorrect timestamp values by comparing them to an approximate expected timestamp and discarding unreasonable values.

Status: A1=No; Fixed



10. Rx Packet May Be Dropped After Receiving 2 Bit Errors While in an Idle State

Problem:

On SGMII Rx lines, when receiving false carrier indication following an ESD (end of stream delimiter) with no idle cycles (i.e. back to back), the valid packet may be dropped by MAC. According to IEEE802.3 (section 24.3.1, Figure 24-11 “Receive state diagram”) such packets should not be dropped.

In order to experience such failures, the device will need to observe Rx_bit[0] and Rx_bits[9:2]!=11111111 while in IDLE state. This means 2 bit errors in 2 symbols.

The probability for this is much lower than, for example, receiving a single bit error in middle of a valid stream (which will have the same impact—frame drop at the MAC).

Implication:

When performing the End of Stream Delimiter Test (UNH: 24.1.1)—the I350 was not observed to reply to test frames containing code groups between valid ESD and the start of Idle.

Workaround:

None.

Status: A1=Yes; NoFix

11. SDPx_1 Remains an Input in D3

Problem:

The software defined pins SDP0_1, SDP1_1, SDP2_1 and SDP3_1 normally function as an input and output signal. However when the device is in D3, the listed signals behave only as inputs until the device is reset.

Implication:

If using this signal as a Tx-disable in an SFP design the SFP module will not disable its transmitting function.

Workaround:

Use an external pull-up or pull-down resistor to keep the SDP line at the intended voltage level during D3.

Status: A1=Yes; NoFix

12. NC-SI Enable Global Multicast Filter Command—No Differentiation of Neighbor Discovery Modes

Problem:

The NC-SI Enable Global Multicast Filter Command defines two separate enable bits for IPv6 Neighbor Advertisement and IPv6 Router Advertisement packets filtering. The two types of packet are differentiated by their ICMPv6 header message type.



There is a single control that enables forwarding of both types of packets; the setting of one bit allows forwarding of both types of packets.

Implication:

When enabling either IPv6 Neighbor Advertisement or IPv6 Router Advertisement packet filtering, both type of packets are sent through the NC-SI interface to the BMC.

Workaround:

None.

Status: A1=Yes; NoFix

13. NC-SI: Select Package Command Increments NC-SI Commands Received Counter for Ports 0 and 1 and Not for Ports 2 and 3

Problem:

For package-based commands, the NC-SI Commands Received Counter should increase in all channels. When a select package command is received, the counter is increased only for ports 0 and 1 and not for ports 2 and 3.

Implication:

Inaccurate statistics.

Workaround:

None.

Status: A1=Yes; NoFix

14. VLAN Tags Byte-Swapped in Loopback Packets

Problem:

When operating in a Virtualization environment (SR-IOV or VMDq) and sending loopback packets, the VLAN tags in the Rx descriptors may be byte-swapped.

Implication:

Sending loopback packets, VF to VF or PF to PF on the same VLAN may result in a failure to correctly identify the packets.

Workaround:

The software driver can check the Rx descriptors for the loopback bit in the extended status field and byte swap the VLAN tags before processing the packet.

Status: A1=Yes; NoFix



15. DEV_OFF_N May Not Properly Disable the Device

Problem:

When asserting DEV_OFF_N with Power Down Enabled (0x01E:15=1) and a PCIe reset is taken, the device may not load the EEPROM image.

Implication:

The device may not be correctly configured after asserting the DEV_OFF_N signal and all PHYs are disabled.

Workaround:

Using the LAN_DIS_N pins is equally effective. However, if these are not available, enabling Pass-Through/Manageability and/or keeping the PHY enabled in D3 can avoid this condition. In addition, software can clear the CTRL_EXT (Extended Device Control) register, Power Down enable bit (20=0) or set 0x0F:6=0 in the EEPROM image.

Status: A1=Yes; NoFix

16. DMA Coalescing Reacts to MSI-X Vectors Instead of Interrupt Causes

Problem:

DMA coalescing should coalesce interrupt events only if interrupts reflect recurring events like packet receive or transmit events. Interrupts reflecting “onetime events” should not be coalesced.

The DMA coalescing mechanism assumes that the first 16 MSI-X vectors reflect recurring events and the others reflect onetime events. A different mapping of events to MSI-X vector will not achieve the expected behavior.

Implication:

When DMA coalescing is activated, only specific mappings of events to MSI-X vectors are allowed.

Workaround:

Map recurring events to the first 16 vectors and other events to higher vectors.

Status: A1=Yes; NoFix

17. IEEE Std 802.3™-2008 Tx Distortion Marginality

Problem:

The I350 might not meet the IEEE Std 802.3™-2008 specification (40.6.1.2.4) that states that the Tx Distortion must meet the following criteria. “A PHY is considered to pass this test if the peak distortion is below 10 mV for at least 60% of the UI within the eye opening.” The I350 might marginally fail this requirement when operating the device at a high temperature and high voltage corner.

Implication:

IEEE conformance is marginal.

The Tx distortion is less than 10 mV during the critical time when the signal is actually sampled therefore no impact on system performance is observed with the I350 due to this marginality.



Workaround:

None.

Status: A1=Yes; NoFix

18. Alternate MAC Address Port Order

Problem:

The datasheet specifies that the Alternate MAC Address Block in the EEPROM is ordered by PCIe function number. The firmware treats this block as being ordered by physical LAN port. These values are not identical if the LAN Function Sel bit (EEPROM word 0x21, bit 12) is 1b. In this case, the Alternate MAC Address values might be used in the wrong LAN port.

Implication:

Incorrect MAC addresses when the Alternate MAC Address Block is programmed in the EEPROM by software.

Workaround:

Program the Alternate MAC Address Block in LAN port order, rather than PCIe function order.

Status: A1=No; EEPROM/NVM Fix

Fixed in EEPROM version 1.52.

Contact your Intel representative to obtain updated EEPROM images.

19. Neighbor Discovery Offload: Override Bit Not Set in Neighbor Advertisement Packet

Problem:

The Neighbor Advertisement packets transmitted by the I350 when providing the Neighbor Discovery protocol offload have the Override bit cleared. Since the protocol offload is functioning as a host system and not a router, the Override bit should be 1b.

Implication:

The requestor's Neighbor Cache table might not be updated correctly.

Microsoft certification for Windows 8 fails as a result of this issue.

Workaround:

None.

Status: A1=No; EEPROM/NVM Fix

Fixed in EEPROM v1.52.

Contact your Intel representative to obtain updated EEPROM images.



20. VFTA Register Access Might Not Be Performed

Problem:

There is a small probability that a read or write access to a VFTA register might not be performed. This can occur while traffic is being processed if any of the following settings are configured:

- OS-to-BMC and/or BMC-to-OS traffic is enabled by the BMC (NC-SI mode).
- VMDq loopback mode is enabled (MRQC.Multiple Receive Queues Enable is 011b and TXSWC.Loopback_en is 1b).
- Anti-spoofing filters are enabled (MRQC.Multiple Receive Queues Enable is 011b and any of the TXSWC.MACAS or TXSWC.VLANAS bits is set).

Implication:

VFTA reads are not reliable and the wrong data value might be returned.

VFTA writes are not reliable and the new value might not be written to the register.

Workaround:

For VFTA reads: Do not read the VFTA registers. Maintain a shadow copy of the VFTA registers in software.

For VFTA writes: Perform the write access multiple times to ensure that the value has been written with a sufficiently high probability. The probability of failure is less than 1-in-35, so 10 consecutive write accesses should bring the cumulative probability of failure to a negligible level (<10⁻¹⁵).

Intel LAN drivers implemented the above workaround in release 16.6 (igb v3.2.9)

An alternate workaround is to use the following procedure to write to the VFTA while packets are being transmitted.

1. Clear TCTL.EN.
2. Set DTXMXPKTSZ to 0xFFFF.
3. Poll register 0x35B4 until bit (8 + LAN_ID) is 0b.
4. Write to VFTA.
5. Restore DTXMXPKTSZ value.
6. Set TCTL.EN.

Note: This workaround is not intended for use when NC-SI pass-through traffic is enabled.

Status: A1=Yes; NoFix

21. Firmware Reset with Loss of Manageability Configuration when Using SerDes

Problem:

When using SerDes mode, a Link Up event might trigger a Firmware Reset.

This erratum does not apply if the AN_ENABLE and AN_TIMEOUT_EN bits in the PCS_LCTL register are both 0b.



Implication:

Loss of manageability and firmware proxying configuration.

Workaround:

The BMC should periodically confirm the manageability configuration and re-configure when necessary. For NC-SI, it is sufficient to send a command to check if the I350 has returned to the Initial State.

No workaround for loss of the proxying configuration.

Status: A1=No; EEPROM/NVM Fix

Fix Fixed in EEPROM version 1.61.

22. Firmware Clears Flow Control Setting

Problem:

When using Ethernet flow control, it is the responsibility of the driver to program the flow control bits in the CTRL register based on the results of the link establishment process. The firmware might override the driver setting; specifically, the CTRL.TFCE bit could be cleared by the firmware following a link up event.

Implication:

Flow control packets are not transmitted when they should be. The receive packet buffer could overflow.

Workaround:

The driver could re-program the flow control bits 100 ms or more after the link is established to ensure the correct values are used.

Status: A1=No; EEPROM/NVM Fix

Fixed in EEPROM version 1.61.

23. PF's MSI TLP Might Contain the Wrong Requester ID When a VF Uses MSI-X

Problem:

When using IOV, if a PF uses MSI interrupts and one or more VFs use MSI-X interrupts, some of the MSI TLPs for the PF might contain the wrong Requester ID.

Implication:

There could be missing interrupts on the PF since the incorrect Requester ID could result in the virtualization mechanism mis-routing or dropping TLPs.

Workaround:

If any VFs use MSI-X, all PF's should also use MSI-X.

Status: A1=Yes; NoFix



24. NC-SI: Set Link Command Fails in 1000Base-KX Mode When Not in D0a State

Problem:

When using 1000Base-KX mode, Set Link commands fail with the Set Link Power Mode Conflict error code.

Implication:

Unexpected and repeated failure of the Set Link command might prevent the BMC from properly completing its initialization sequence.

Workaround:

Clear the Disable 1000 in non-D0a EEPROM bit (word offset 0x20, bit 3) for each port using 1000Base-KX mode.

Status: A1=Yes; NoFix

25. Failures When Functions are Swapped and Ports are Disabled

Problem:

If the LAN Function Select EEPROM bit (word 0x21, bit 12) is set to 1b and one or more ports are disabled, there might be failures when using NC-SI manageability and/or the Restore MAC Address feature. The exception is the I350-AM2 SKU, which functions correctly with LAN Function Select even though two ports are disabled.

For NC-SI, the channel-to-port mapping is incorrect.

The Restore MAC Address feature uses the wrong Alternate MAC Address field.

Implication:

NC-SI: Commands are executed on the wrong port, so the pass-through does not function.

Restore MAC Address: Incorrect MAC address might result in duplicate MAC addresses on the subnet.

Workaround:

When designing a board using the I350, do not use LAN Function Select to swap the function-to-port mapping if any ports need to be disabled.

Status: A1=No; EEPROM/NVM Fix

Fixed in EEPROM version 1.60.



26. Removed (Duplicate of Erratum #22)

27. Enabling a Port Causes a Reset of the Firmware

Problem:

If a port is disabled (either by strapping or EEPROM setting) when the device is powered up and the port is later enabled, a PCIe reset (without an AC power cycle) can trigger a reset of the firmware.

Implication:

A firmware reset causes a loss of all manageability filter settings that have been configured by the BMC. This likely results in a loss of connectivity for any active manageability sessions unless the BMC can reconfigure the settings quickly enough. It can also cause improper initialization of the device after the PCIe reset.

Workaround:

All ports that might be used should be enabled when the device is powered up. It is OK to disable ports later as required.

Status: A1=No; EEPROM/NVM Fix

Fixed in EEPROM version 1.63.

28. MCTP: Interface is Not Functional in D3/Dr

Problem:

In MCTP mode, the SMBus interface is disabled when the device is in the Dr or D3 state. As a result, the device is unable to process MCTP commands during that time.

Implication:

Sending an MCTP command when the device is in the Dr or D3 state can result in a hang of the MCTP interface.

Workaround:

Do not send MCTP commands when the device is in the Dr or D3 state.

Status: A1=No; EEPROM/NVM Fix

29. MCTP: Command with Duplicate NC-SI IID is Ignored

Problem:

According to the NC-SI specification, when a Network Controller (NC) receives a request with the same Instance ID (IID) as the previous request it must return the previous response. In MCTP mode, the I350 does not respond to a request with the same IID as the previous request.

Implication:

Reduced reliability of the NC-SI Interface.



Workaround:

If a response does not arrive, the MC should repeat the command with a different IID. This might not be effective in cases where repeating the command has side effects.

Status: A1=No; EEPROM/NVM Fix

Fixed in EEPROM version 1.59.

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

Problem:

The I350 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: A1=Yes; NoFix

31. VFMPRC is Not Accessible from the VF Memory Space

Problem:

The VF Multicast Packets Receive Count (VFMPRC) register is not accessible from the VF memory space. It is accessible from the PF memory space.

Implication:

Counter cannot be read by the VF driver.

Workaround:

The VF driver can either count the multicast received packets or ask the PF driver to read the VFMPRC and provide the count value.

The Intel Virtual Function (VF) driver works around this erratum by counting the number of multicast packets in the driver.

Status: A1=Yes; NoFix



32. 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: A1=Yes; NoFix

33. 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 transition.

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.

Status: A1=Yes; NoFix

34. Common MDIO Failure when Port 0 is Disabled or Powered Down

Problem:

The Common MDIO feature (using one pair of MDC/MDIO signals for multiple ports to access an external PHY) uses the Port 0 internal clock. If Port 0 is disabled (by strapping) or in a low-power state (non-D0a and no WoL and NoMng) the clock is gated off and the common MDIO logic does not operate.

Implication:

Ports 1-3 cannot access the external PHY if Port 0 is not operational.

Workaround:

For applications where Port 0 might be disabled by strapping, clear the PHY_in_LAN_disable EEPROM bit for Port 0 so that the MAC is always operational.



Also, one of the following:

- Enable Wake on LAN for Port 0.
- Use an EEPROM image that has manageability enabled.

Status: A1=Yes; NoFix

35. NC-SI Output Signals Have Indeterminate Value After Power Up

Problem:

The NC-SI output signals have an indeterminate value after power up until the first rising edge of the NC-SI input clock. The signals could be tri-stated or driven high or low.

Implication:

Current leakage through the NC-SI I/O buffers.

Workaround:

If the NC-SI input clock is not driven after power up, connect the NC-SI clock input pin so that there is a rising edge after power has stabilized. For example, it could be connected via a resistor to a power-good indication on the board.

Status: A1=Yes; NoFix

36. PCIe SR-IOV Reserved Bits are Writable

Problem:

According to PCIe Specification, RsvdP register fields must be read only and must return 0 (all 0's for multi-bit fields) when read.

In this device the following reserved bits are writable:

- SR-IOV Capability Structure offset 0x08 - SR-IOV Control/Status Register (0x168), bits 15:5.
- SR-IOV Capability Structure offset 0x13 (0x173), bits 7:0.

Implication:

No functional implication. Software should not write reserved bits.

Workaround:

None.

Status: A1=Yes; NoFix



37. Certain Malformed IPv6 Extension Headers are not Processed Correctly by the Device

Problem:

Certain malformed IPv6 extension headers are not processed correctly by the device.

Implication:

If a packet containing such malformed IPv6 extension headers is received, the device may behave unpredictably.

Workaround:

Set bit 16 (IPv6_ExDis) in the RFCTL register to disable the processing of received IPv6 extension headers. Note that with this bit set, checksum calculation and RSS are disabled for IPv6 packets containing extension headers.

This workaround has been implemented in Intel drivers starting from Release 20.2.

Status: A1=Yes; NoFix

38. PCIe Advanced Error Reporting: First Error Pointer

Problem:

The First Error Pointer in the Advanced Error Capabilities and Control Register (PCIe register 0x118 bits 4:0) is a field that identifies the bit position of the first error reported in the Uncorrectable Error Status register. In the I350 implementation, the following bits of the Uncorrectable Status Register are not covered by this field:

- Bit 4 — Data Link Protocol Error Status.
- Bit 13 — Flow Control Protocol Error Status.
- Bit 14 — Completion Timeout Status.

Implication:

PCIe specification compliance issue.

Workaround:

None.

Status: A1=Yes; NoFix



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. 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

Clarification:

The I350 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."



3. PF/VF Drivers Should Configure Registers that are not Reset By VFLR

Clarification:

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):

RDH/T, TDH/T, PSRTYPE, SRRCTL, RXDCTL, TXDCTL, TDWBAL/H, RXCTRL, TXCTRL



LEGAL

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

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

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

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

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

© 2011-2016 Intel Corporation.