Intel® 82573 Family Gigabit Ethernet Controllers Specification Update

Networking Silicon

Revision History

Revision	Revision Date	Description
3.4	June 2012	Removed Table of Contents. Added Erratum #43.
3.3	January 2012	Added Erratum #42. Added Software Clarification #2.
3.2	December 2010	Added Specification Change #3. Added Erratum #41. Added Software Clarification #1.
3.1	November 2010	Added Erratum #40.
3.0	September 2010	Added Erratum #39. Added Specification Clarification #3.
2.9	March 2010	Updated Erratum #8.
2.8	November 2009	Added Erratum #38.
2.7	August 2009	Added Specification Clarification #2. Added Errata #27 through 37.
2.6	June 2009	Added Specification Clarification #1.
2.5	Sept 2008	Updated Erratum #21.
2.4	May 2007	Added Tables 3 and 4 “Production Marks”.
2.3	January 2007	Updated errata 21. Added S5 to problem description. Updated device identification table. Removed Specification Change #1. Removed Sections 7 and 10.
2.2	June 2006	Moved and updated Sighting 1-3 to Errata section. Updated Intel logo.
2.1	March 2006	i) Added errata for Thermal diode issue root cause (errata #19) ii) Added errata for incorrect PET retransmission interval configuration (errata #20) iii) Moved LED link status issue from sightings to errata (errata #21) iv) Removed sighting #12 – “Consuming 40% of CPU cycles during network disconnection” as this was determined not related to Intel® silicon. v) Few wording changes in errata #7 & #17. Fixed some typographical errors.
2.0	Nov 2005	Combined the 82573E/V Specification Update with the 82573L Specification Update. Added errata.
1.0	June 2005	Initial release.

Preface

This document is applicable to the Intel 82573 Gigabit Ethernet controller.

This document is an update to published specifications. Specification documents for this product include:

- 82573 Family of Gigabit Ethernet Controllers Datasheet, Intel Corporation.
- 82573/82562 Dual Footprint Design Guide, Intel Corporation.
- 82573 NVM Map and Programming Information Guide, Intel Corporation.
- PCIe* Family of Gigabit Ethernet Controllers Software Developer's Manual, Intel Corporation.

This document is intended for hardware system manufacturers and software developers of applications, operating systems or tools. It may contain Specification Changes, Errata, and Specification Clarifications.

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.

1.0 Nomenclature

This document uses various definitions, codes, and abbreviations to describe the Specification Changes, Errata, Sightings and/or Specification Clarifications that apply to the listed silicon/steppings:

Table 1 Definitions

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.
Sightings	Observed issues that are believed to be errata, but have not been completely confirmed or root caused. The intention of documenting sightings is to proactively inform users of behaviors or issues that have been observed. Sightings may evolve to errata or may be removed as non-issues after an investigation has been completed.
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.

Table 2 Codes and Abbreviations

Name	Description
X	Specification Change, Erratum, or Specification Clarification that applies to this stepping.
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.
No Fix	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 the listed stepping.
Shaded	This Item is either new or modified from the previous version of the document
DS	Data Sheet
SDM	Software Developer's Manual
EDS	External Data Specification
AP	Application Note

2.0 82573 Device Identification

The following table describes the various identifying markings on each device package.

Device	Stepping	Top Marking	QDF Number	Leaded/Unleaded	Notes
82573E	A-3	RC82573E	n/a	Leaded	Production Units
82573E	A-3	PC82573E	n/a	Lead-Free	Production Units
82573V	A-3	RC82573V	n/a	Leaded	Production Units
82573V	A-3	PC82573V	n/a	Lead-Free	Production Units
82573L	A-0	RC82573L	n/a	Leaded	Production Units
82573L	A-0	PC82573L	n/a	Lead-Free	Production Units

Production Unit Mark Diagrams

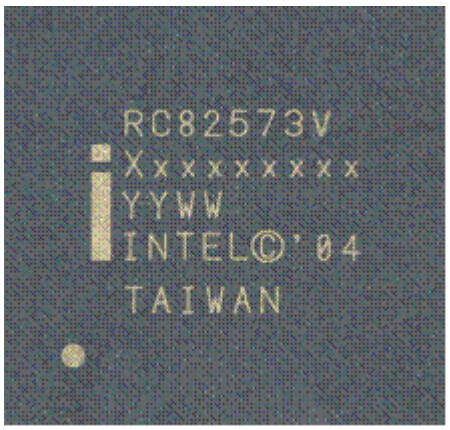

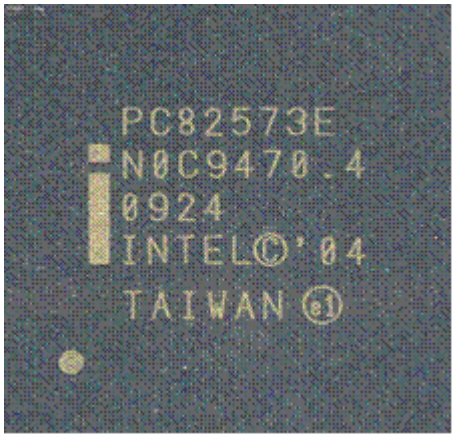

Leaded Version **Lead-Free Version**

Table 3 Production Marks ("RC" Prefix Leaded Parts)

Line 1:	Marketing Name RC82573L, RC82573V, or RC82573E
Line 2:	TSMC Fab Lot Number "Xxxxxxxxx" or "Xxxxxxxx.x"
Line 3:	Assembly Date Code "YYWW"
Line 4:	Copyright line including two number date code
Line 5:	Country of Origin

Table 4 Production Marks ("PC" Prefix Lead-Free Parts)

Line 1:	Marketing Name PC82573L, PC82573V or PC82573E
Line 2:	TSMC Fab Lot Number "Xxxxxxxxx" or "Xxxxxxxx.x"
Line 3:	Assembly Date Code "YYWW"
Line 4:	Copyright line including two number date code
Line 5:	Country of Origin and circled "e1" lead-free mark

3.0 Summary of Table Changes

#	A-0L	A-1E	A-2E	A-3E/V	Plans	Specification Changes	Page
1	X	X	X	X	No Fix	82573E Performs Partial Flash Update with ST Micro Devices	9
2	X	X	X	X	No Fix	Serial Number Capability is Constant Zeros	9
3	X				No Fix	Functionality of Bit 20 of the Interrupt Registers	10

#	A-0L	A-1E	A-2E	A-3E/V	Plans	Errata	Page
1	X	X	X	X	No Fix	Out of Usage-Model PCIe* Specification Compliance Minor Issues	10
2	X	X	X	X	No Fix	Upstream Attempt to Reconfigure the Link by Moving LTSSM from Recovery to Configuration Will Cause a Link Down	11
3	X	X	X	X	No Fix	Request Will Not Be Treated as Completion Abort (CA) the Programming Model Bytes Enable is Violated	11
4	X	X	X	X	No Fix	PCIe* Minimum Transmit and Receive Differential Return Loss is Lower than Specification	12
5	X	X	X	X	No Fix	Crystal Overdrive at High Voltage, High Load Capacitance Corner	12
6	X	X	X	X	No Fix	When the Ethernet Controller Receives a Memory Read32 Locked Request, It uses the Last Completion's Requester ID	13
7	X	X	X	X	No Fix for E/V Fix for L	82573E/82573V IEEE Marginality	13
8	X	X	X	X	No Fix	82573 Disappears in PCI Configuration Space when L0s and L1 PCIe* Link States Are Enabled	14
9	X				No Fix	ASF Agent Answers to Old IP Address After IP Address Was Changed	14
10	X				No Fix	Watchdog Enable Bit Does Not Start Watchdog Timer	15
11	X				No Fix	82573L Link Control Register Requires Word Writes for Proper CLKREQ# Setting	15
12	X				Fix	Incorrect MAC Address is Reported to the OS with the 82573L	15
13		X	X		Fix	Ethernet Controller Doesn't Complete "CFG Write Type 1" Request and Reports "Malformed TLP" Instead of "Unsupported Request"	16
14		X	X	X	No Fix	Time Stamp in PET Alerts is Not Updated with the 82573E/V	16
15		X	X	X	No Fix	Ethernet Controller Halts Requiring the Driver to Issue a Reset if Dynamic DMA Clock Gating is Enabled on the 82573E.	17
16		X	X	X	No Fix	Heartbeat Timer is not Accurate	17

#	A-0L	A-1E	A-2E	A-3E/V	Plans	Errata	Page
17		X	X	X	No Fix for E/V Fix for L[1]	ASPM/Jumbo Frames Disabled Due to Early Receive Threshold Overrun Buffer	18
18	X				Fix	L1 Aggressive with Windows 2000 Spin Lock Issue	18
19	X	X	X	X	No Fix	Thermal Diode Issue	19
20				X	No Fix	ASF Retransmission interval issue	19
21	X	X	X	X	Fix	LED link status issue (fixed through driver changes & extra delay)	20
22		X	X	X	No Fix	Wakeup Failure Occurs after Power is Removed and Re-applied in EEPROM Mode	20
23		X	X	X	No Fix	When the SMBD is Held Low for Longer than Ttimeout MAX, the SMB Will Hang	20
24		X	X	X	No Fix	Packet Data Corruption Upon Recovery (82573E and 82573V only)	21
25		X	X	X	No Fix	TLS Sessions using Intel® AMT SOL and IDER Stress Concurrently Might Disconnect	21
26		X	X	X	No Fix	Intel® AMT Strong Stress on Gigabit Traffic in APT Mode Can Occasionally Cause Firmware Reset	21
27	X	X	X		NoFix	Corruption in LAN-to-BMC Pass Through Data	22
28	X	X	X	X	NoFix	PCIe: Completion with Completer Abort (CA) or Unsupported Request (UR) Status is Considered Malformed	22
29	X	X	X	X	NoFix	PCIe: PCIe Bus May Halt if There are Less Than 16 Posted Data Flow Control Credits (= 256-Byte Memory Writes)	22
30	X	X	X	X	NoFix	PCIe: Surprise Down Following Externally-Initiated Recovery	23
31	X	X	X	X	NoFix	Missing Interrupt Following ICR Read	23
32	X	X	X	X	NoFix	PCIe: Missing Replay Due to Recovery During TLP Transmission	23
33	X	X	X	X	NoFix	PCIe: Reception of Completion That Should Be Dropped May Occasionally Result In Device Hang or Data Corruption	24
34	X	X	X	X	NoFix	LED Turns on at Software Reset of MAC	25
35	X	X	X	X	NoFix	Receive Packet Delayed When Using RDTR or RADV Register	25
36	X	X	X	X	NoFix	The 82573 Overwrites Transmit Descriptors in Internal Buffer	26
37	X	X	X	X	NoFix	Enabling or Disabling RSS with Data in the Rx FIFO May Cause an Rx Hang	26
38	X	X	X	X	NoFix	PCIe: SKP Ordered Set Resets Training Sequence Counter	27
39	X	X	X	X	NoFix	Tx Data Corruption When Using TCP Segmentation Offload	27
40	X	X	X	X	NoFix	PCIe: Spurious SDP/STP Causes Packets to be Dropped	28
41	X				NoFix	Delay of Received Ethernet Packet During ASPM L1	28
42	X	X	X	X	NoFix	Packets Received with an L2+L3 Header Length Greater than 256 Bytes Can Incorrectly Report a Checksum Error	29
43	X	X	X	X	NoFix	Certain Malformed IPv6 Extension Headers are not Processed Correctly by the Device	

[1] Fix available with Restrictions. Refer to Errata 17 for details.

#	A-0L	A-1E	A-2E	A-3E/V	Plans	Specification Clarifications	Page
1	X	X	X	X	No Fix	Spec Clarification for PCIe1 Devices	34
2	X	X	X		NoFix	Manageability: Critical Session (Keep PHY Link Up) Mode Does Not Block All PHY Resets Caused By PCIe Resets	34
3	X	X	X		NoFix	Use of Wake on LAN Together With Manageability (82573E Only)	35
#	A-0L	A-1E	A-2E	A-3E/V	Plans	Software Clarifications	Page
1	X	X	X	X	NoFix	While in TCP Segmentation Offload, Each Buffer is Limited to 64 KB	36
2	X	X	X	X	NoFix	Serial Interfaces Programmed By Bit Banging	36

4.0 Specification Changes

1 82573E Performs Partial Flash Update with ST Micro Devices

Problem: When updating sector 1 from the shadow RAM, only a partial copy is performed.

Implication: When sector 1 is the valid, manageability images could display stability issues.

Workaround: It should be ensured that the valid sector is sector 0. When a Flash update is required, it should be performed twice. This ensures that sector 0 remains as the valid sector.

Status: Intel does not expect to resolve this sighting for the 82573E Gigabit Ethernet Controller.

2 Serial Number Capability is Constant Zeros

Problem:	The 82573E will return zeros when software reads the PCIe* serial capability number.
Implication:	The PCI Express serial number read will return zeros.
Workaround:	None.
Status:	Intel does not expect to resolve this sighting for the 82573E Gigabit Ethernet Controller.

3 Functionality of Bit 20 of the Interrupt Registers

For the 82573L only, changes to the functionality of bit 20 in the interrupt registers are summarized in the table below. Further information on the interrupt registers can be obtained in the *PCIe GbE Controllers Open Source Software Developers Manual, Revision 2.3*.

Register	Bit 20	Description
ICR Register	EDC (82573L only)	Energy Detect Change. This bit is set each time the energy detect status reported by the internal PHY changes.
ICS Register	EDC (82573L only)	Energy Detect Change. Sets the Energy Detect Change Interrupt.
IMS Register	EDC (82573L only)	Energy Detect Change. Sets the mask for the Energy Detect Change.
IMC Register	EDC (82573L only)	Energy Detect Change. Clears the mask for the Energy Detect Change.

If you are using the 82573L, set bit 20 appropriately.

5.0 Errata

1 Out of Usage Model PCI-E Specification Compliance Minor Issues

Problem:	The 82573 has some specification compliance issues that do not have impact on the intended usage of the unit. MAX_PAYLOAD_SIZE is programmed per function. If two PCI functions have different MAX_PAYLOAD_SIZE, the 82573 might use the larger value for all functions. Usage model is to have all functions using same the MAX_PAYLOAD_SIZE. MEM_RD_LK using 64 bit address will be classified as malformed packet. The bridge will filter 64b address space packets so it will not reach the unit.
Implication:	This is not part of the usage model and has no impact on functional flow.
Workaround:	None.
Status:	Intel does not plan to resolve this erratum in the 82573 Gigabit Ethernet Controller.

2 Upstream Attempt to Reconfigure the Link by Moving LTSSM from Recovery to Configuration will Cause a Link Down

Problem:	The 82573 will not move from Recovery to Configuration when it receives TSs with only lane number set to PAD.
Implication:	If the upstream component tries to reconfigure the link by moving the LTSSM from Recovery.Idle to Configuration (sending TS1s with only lane number set to PAD), the link will fail and the units will drop down to Detect states, causing a link down event.
Workaround:	The upstream component should not apply to this option.
Status:	Intel does not plan to resolve this erratum in the 82573 Gigabit Ethernet Controller.

3 Request Will Not Be Treated as Completion Abort (CA) When the Programming Model Bytes Enable is Violated

Problem:	The PCI Express specification allows a device not to accept certain requests. This is under the "programming model" cases. The device needs to issue a Completer Abort error if the specific request violates the programming model. As part of its programming model, the device does not support writes with byte enables to specific memory addresses. Such writes will be fully executed and will not be treated as completion abort.
Implication:	CSR writes with partial Bytes Enables will be executed (in specific address ranges). This scenario will not happen when using the device driver. This functionality is also not needed for the normal operation of the design.
Workaround:	No partial byte enable writing to the device.
Status:	Intel does not plan to resolve this erratum in the 82573 Gigabit Ethernet Controller.

4 PCIe* Minimum Transmit and Receive Differential Return Loss is Lower Than Specification

Problem:	The 82573 PCIe* specification, for minimum transmitter and receiver differential return loss, is lower than required by PCIe* specifications. This will not impact the performance seen by the end user. PCIe* Transmitter Differential Return Loss: According to PCIe* specification, the Transmitter input impedance shall result in a differential return loss greater than or equal to 10 dB over a frequency range of 50 MHz to 1.25 GHz. The measured transmit differential return loss for the 82573 is between 5.30dB and 6.54dB. PCIe* Receiver differential return loss: According to the PCIe* specification, the Receiver input impedance shall result in a differential return loss greater than or equal to 10 dB over a frequency range of 50 MHz to 1.25 GHz. The measured receive differential return loss for the 82573 is between 4.95dB and 6.96dB.
Implication:	This has no impact on the end user system. Increasing the PCIe* return loss will increase the device's sensitivity to ESD.
Workaround:	None.
Status:	Intel does not plan to resolve this erratum in the 82573 Gigabit Ethernet Controller.

5 Crystal Overdrive at High Voltage, High Load Capacitance Corner

Problem:	At the high voltage and high load capacitance corner, the 82573 controller has been found to overdrive crystals with drive levels of 500uW. The measured drive level for a crystal with CL = 26pF and ESR = 50Ω at the high voltage corner reached up to700uW
Implication:	Overdriving the crystal past its specified drive level can decrease its lifetime and lead to the breakdown of the device.
Workaround:	Intel recommends using crystals with maximum load capacitance of 20pF and maximum ESR of 40ohms. Alternatively, a crystal with 1.0mW drive level may be used. Intel will publish a list of recommended devices in the 82573 design guide.
Status:	Intel does not plan to resolve this erratum in the 82573 Gigabit Ethernet Controller.

6 When the 82573 Ethernet Controller Receives a Memory Read32 Locked Request, It Uses the Last Completion's Requester ID

Problem:	When the 82573 Ethernet Controller receives a Memory Read32 Locked request, it uses the last completion's Requester ID.
Implication:	The complete cannot be associated with the request. Additionally the complete is CMPL instead of CMPL_LK.
Workaround:	A workaround is not required since the Ethernet Controller is a PCI Express endpoint and should not receive MEM_RDCLK.
Status:	Intel does not plan to resolve this erratum in the 82573 Gigabit Ethernet Controller.

7 82573E/V IEEE Marginality

Problem:	**10Mb Differential Output Voltage** The IEEE specification requires that the peak differential voltage on the TD (transmit data) circuit when terminated with a 100ohm resistive load be between 2.2V and 2.8V for all data sequences. This specification must be met from 5MHz to 15MHz.
Implication:	On 82573E/V A-2 silicon, the 10Mb output amplitude at 10MHz has been measured as high as 2.8V. From a board-level perspective, output amplitude can be affected by long MDI traces, resistors on the center tap of the magnetic module, incorrect termination on MDI traces, and characteristics of the magnetic module itself.
Workaround:	On 82573E/V A-3, in order to increase the margin allotted for board design variations, the 10Mb amplitude defaults have been modified. Due to the modified settings, the 10Mb output amplitude are lower and in the 2.3-2.5V range.
Status:	This has been resolved in the 82573L Gigabit Ethernet Controller.

Problem:	**100Mb Rise and Fall Times** IEEE specification requires that 100Mb Rise and Fall times be less than 5ns.
Implication:	In 82573E/V A-2, 100Mb rise/fall times can be measured as high as 5.2ns. When linked to an IEEE-compliant device, this errata will not affect a normal, end-user environment.
Workaround:	The 82573E/V A-3 and 82573L rise and fall times will be lowered by 0.2ns.
Status:	This has been resolved in the 82573L Gigabit Ethernet Controller. However, there may be a few failures in systems with trace lengths longer than 4 inches.

8 82573 Disappears in PCI Configuration Space When L0s and L1 PCIe* Link States Are Enabled

Problem: When both L0s and L1 PCIe* link states are enabled, it causes the 82573 to disappear in the PCI Configuration Space. When both the ICH and 82573 enter L0s power state, generating some traffic, the 82573 disappears from the Configuration Space. The time it takes for the device varies from 10 minutes to 24 hours.

Implication: When PCIe* enters L1 state after L0s, noise on the receive line can cause the PCIe* state machine to miss the L1 acknowledge from ICH, resulting in a PCIe* hang.

Workaround: To maximize power savings, L1 is enabled, while L0s is disabled. The workaround includes the following steps:

1. Use an NVM image produced after software driver 10.2 (labeled ASPM).
2. Use software release 10.1 or later.
3. Set the correct bits in the Link Control Register on the LAN and the ICH device. On the LAN, this is achieved by the BIOS setting offset F0h, bits 1:0 to 10b (L1 only setting). On the ICH7, this is achieved by setting register 50h, bits 1:0 to 10b.

This workaround allows these devices to have nearly the same power savings that they would if they fully supported L0s and L1 (about 10 mW to 20 mW difference).

Status: Intel does not plan to resolve this erratum in the 82573 Gigabit Ethernet Controller.

9 ASF Agent Answers to Old IP Address After IP Address Was Changed

Problem: After the IP address of an ASF-enabled platform is changed, the ASF machine continues to answer RMCP requests and commands for both the old and new IP addresses for approximately 20 minutes even though the EEPROM is updated with the new IP address.

Implication: Firmware responds to RMCP commands for the old IP address.

Workaround: After about 20 minutes, the firmware responds only to the new IP address RMCP commands.

Status: Intel does not plan to resolve this erratum in the 82573 Gigabit Ethernet Controller.

10 Watchdog Enable Bit Does Not Start Watchdog Timer

Problem:	In ASF mode only, setting the Watchdog Enable bit in the EEPROM does not activate the Watchdog timer.
Implication:	The firmware will not be able to report any OS or BIOS hangs.
Workaround:	There is no workaround for this issue.
Status:	Intel does not plan to resolve this erratum in the 82573 Gigabit Ethernet Controller.

11 82573L Link Control Register Requires Word Writes for Proper CLKREQ# Setting

Problem:	The CLKREQ# signal will not properly set if the user accesses the Link Control Register (Fh) in the PCIe* Configuration Space through byte access rather than word access. *Note:* This issue relates to the 82573L only.
Implication:	The system reference clock for the LAN PCIe* link will not turn off in L1.
Workaround:	The workaround involves writing to the Link Control Register (Fh) as a word.
Status:	Intel does not plan to resolve this erratum in the 82573 Gigabit Ethernet Controller.

12 Incorrect MAC Address is Reported to the OS with the 82573L

Problem:	In 82573L LOM designs, using an EEPROM and operating with a driver released prior to revision 10.2, the incorrect MAC address is reported to the OS preventing network connection.
Implication:	Platforms that use the 82573L in this configuration may not have network connectivity.
Workaround:	The workaround involves upgrading the device drivers to revision 10.2 driver release or later.
Status:	This erratum was resolved with software release 10.2.

13 Ethernet Controller Doesn't Complete 'CFG Write Type 1' Request and Reports 'Malformed TLP' Instead of "Unsupported Request"

Problem:	The 82573E Gigabit Ethernet Controller is considered an endpoint and replies correctly to 'CFG type 0' access. However, for a 'CFG type 1' access (normally sent to the bridge chip), the 82573E controller has no completion and the uncorrectable error register reports "Malformed TLP."
Implication:	The requester will not receive a completion. Additionally, since the malformed TLP does not update the credits the next request to the Ethernet Controller will not be sent due to insufficient credits. The 82573E controller is an endpoint and should not expect to get 'CFG type 1' access.
Workaround:	If the BIOS does request a 'CFG type 1' access, a branch that jumps over the 'CFG Type 1' access must be implemented if the port(s) has been identified as an endpoint port (replies correctly to 'CFG type 0').
Status:	Intel resolved this erratum in the A-2 stepping of the 82573E Gigabit Ethernet Controller.

14 Time Stamp in PET Alerts is Not Updated with the 82573E/V

Problem:	In ASF mode using the 82573E/V device, all PETs are sent with the same time stamp.
Implication:	All the PETs received in the management console will have the same time stamp.
Workaround:	To avoid this issue, the management console software should use the PET arrival time instead of the PET time stamp.
Status:	Intel does not plan to resolve this erratum in the 82573 Gigabit Ethernet Controller.

15 Ethernet Controller Halts Requiring the Driver to Issue a Reset if Dynamic DMA Clock Gating is Enabled

Problem:	The 82573 Family of Gigabit Ethernet Controllers may hang if the DMA clock is stopped by the D0a Dynamic DMA Clock Gating feature. The controller logic will cause the DMA clock to be gated if, during power state D0, the chip is idle, Dynamic DMA Clock Gating (NVM word 0xF.3) is enabled and NVM word 0xF.0 (disable D0a Dynamic Clock Gating) is not set.
Implication:	DMA Dynamic Clock Gating should be enabled to allow the controller to enter the L1 state, so NVM word 0xF.3 should be set to a 1 in most configurations. If “D0a Dynamic DMA Clock Gating” is enabled also, then there is the potential for the controller to halt unexpectedly.
Workaround:	NVM word 0xF.0 should be set to a 1 to disable “D0a Dynamic DMA Clock Gating” in all configurations.
Status:	The “D0a Dynamic DMA Clock Gating” feature is not supported. Intel has no plans to support this feature in the future.

16 Heartbeat Timer is not Accurate

Problem:	The heartbeats are being sent at intervals of 11.2-11.4 seconds instead of 10.7 seconds. Retransmissions are being sent at intervals of 3.1-3.3 seconds instead of 2.7 seconds.
Implication:	No implication on the user. However, during automated system testing a timer resolution differing from the expected result may be received.
Workaround:	No workaround is available for this issue
Status:	Intel does not plan to resolve this erratum in the 82573 Gigabit Ethernet Controller.

17 ASPM/Jumbo Frames Disabled Due to Early Receive Threshold Overrun Buffer

Problem:	The Jumbo Frames feature uses the Early Receive Threshold (ERT) in the 82573 when ASPM is enabled. When ERT is enabled, Jumbo Frames may possibly corrupt system memory.
Implication:	This only occurs when ASPM is enabled and is most evident when jumbo frames are used. When the 82573 is performing an ERT (depending on the value set in the ERT register) and a corrupted packet is received, there is the possibility that the 82573 can overrun the buffer that it is accessing the receive data into by up to 256 bytes (size of a PCIe* transaction). This overrun occurs because the 82573 is accessing data into the receive buffer but then aborts the transaction due to the error. The buffer will be re-used for the next receive packet.
Workaround:	When ERT is disabled, the decreased FIFO size limits the size of available space for receive traffic. Jumbo Frames (up to packet size of 9KB) with software release 10.3 is supported on the 82573L Gigabit Ethernet Controller as long as ASPM is disabled.[2,3]
Status:	Intel does not plan to resolve this erratum in the 82573 Gigabit Ethernet Controller. Jumbo frames is not supported in 82573E/V & is supported with the workaround above in 82573L.

18 L1 Aggressive with Windows 2000 Spin Lock Issue

Problem:	When L1 Aggressive is enabled on The 82573L Gigabit Ethernet Controller the driver has an error in the code that interacts badly with the spin lock mechanism of Windows 2000 which, under heavy stress traffic, forces a release, resulting in a blue screen.
Implication:	Since L1 Aggressive is disabled by default, the user has to manually enable L1 Aggressive to see the error. The error will only occur with Software Driver 10.2, on Windows 2000 with the 82573L Gigabit Ethernet Controller under heavy stress traffic.
Workaround:	Software Driver 10.2.01 will remove the option and Software Driver 10.3 will fix the issue.
Status:	Intel does not plan to resolve this erratum in the 82573 Gigabit Ethernet Controller.

[2] ASPM can be disabled by setting PCIe* Initial Configuration Word 3 (Word 1A, bits [3:2]) to 00. Refer to 82573 NVM Map & Programming Information Guide (AP475).
[3] Jumbo Frames & Early Receive Threshold (ERT) cannot be enabled at the same time.

19 Thermal Diode Issue

Problem:	When 82573 3.3V supply transitions from high to low, a voltage drop on the thermal pads (ThermN/ThermP) may be observed. This is because the ESD protection diode between the ThermP pad and 3.3V is forward biased when 3.3V is removed and some models of thermal sensor do not provide enough current to drive both ESD and Thermal Diodes. This results in the pulse amplitude from the thermal sensor to be dragged down.
Implication:	Incorrect voltage on the thermal pads can result in inaccurate die temperature measurements.
Workaround:	Intel recommends masking the Thermal sensor reading measurement when 82573 is powered off. Intel also recommends leaving the 3.3V supply greater than or equal to 0.7V to avoid the ThermP pad voltage drops.
Status:	Intel does not plan to resolve this erratum in the 82573 Gigabit Ethernet Controller.

20 ASF Retransmission Interval Issue

Problem: The interval resolution between two consecutive retransmission of Platform Event Trap (PET) alert is 2.7 seconds. If the value of the interval as set in the NVM is larger than 2.7 seconds, the 82573 E/V ASF Firmware will delay more time than it was configured (the original time * ~1.4).

For example: The table below shows the expected interval based on NVM setting and the actual value.

Expected	Actual
2.7 sec	2.7 sec
30 sec	42 sec
57 sec	78 sec
66 sec	92 sec
120 sec	167 sec

Implication:	The PET alerts retransmission will be sent in longer delay than configured
Workaround:	When configuring the PET retransmission time interval, consider the 1.4 factor.
Status:	Intel does not plan to resolve this erratum in the 82573 Gigabit Ethernet Controller.

21 LED Link Status Issue

Problem:	When the system goes to standby (S3 and S5), the link status LEDs might stay on when there is no link.
Implication:	End users could see the LEDs light up when the system is in standby state (S3 and S5) with no link, the LEDs could draw extra power in standby.
Workaround:	This workaround is contained in software release 10.3 or a later version and is disabled, by default, in the driver. Enabling the workaround requires modifying a hidden registry and then enabling WoL in the operating system. Enabling the workaround results in the system taking ~100 ms more time to go into standby. If customers choose to implement this workaround and add the delay, the registry setting to modify is ForceLedOffAtD3. It is a REG_SZ setting and should be set to 100.
Status:	Intel does not plan to resolve this erratum in the 82573 Gigabit Ethernet Controller

22 Wakeup Failure Occurs after Power is Removed and Re-applied in EEPROM Mode

Problem:	Since the internal PCI Express PLL clock gating after PE_RST# is de-asserted, the value of the internal register that holds EEPROM Word 0Ah is reset after the EEPROM loads. A software reset will re-load Word 0Ah register correctly and enable the 82573 Family to work in the regular S0 state.
Implication:	Wakeup in the S5 state may not work when PE_RST# is de-asserted and then asserted while PCI Express clock is running.
Workaround:	BIOS should issue a PCI Express in band reset at least 500 µs after PE_RST# has been asserted. The reset will ensure that the register is properly loaded from EEPROM and correct functionality is restored.
Status:	Intel does not expect to resolve this erratum for the 82573 Family of Gigabit Ethernet Controllers.

23 When the SMBD is Held Low for Longer than Timeout MAX, the SMB Will Hang

Problem:	According to the SMB specification, when the SMB clock is held low for longer than Ttimeout MAX a hang should occur. However, in the 82573E/V this can occur if the SMB clock signal or SMB data signal is held low for a period of Ttimeout MAX.
Implication:	The SMB will hang.
Workaround:	None.
Status:	Intel does not expect to resolve this erratum for the 82573E/V Gigabit Ethernet Controller.

24 Packet Data Corruption Upon Recovery (82573E and 82573V only)

Problem:	If the 82573 (E or V) and the upstream device are configured to enable LOs, a particular sequence of completions can result in a CRC error.
Implication:	PCIe completion data becomes corrupted, which in turn corrupts the Ethernet frame data.
Workaround:	Los has been disabled in the 82573 silicon and no additional workaround is necessary.
Status:	Intel does not expect to resolve this erratum for the 82573E and 82573V Gigabit Ethernet Controller.

25 TLS Sessions using Intel® AMT SOL and IDER Stress Concurrently Might Disconnect

Problem:	When exercising SOL and IDER over a secure connection, the sessions might disconnect.
Implication:	SOL and IDER may stop functioning.
Workaround:	The SOL/IDER sessions should be re-established.
Status:	Intel does not expect to resolve this erratum for the 82573E Gigabit Ethernet Controller.

26 Intel® AMT Strong Stress on Gigabit Traffic in APT Mode Can Occasionally Cause Firmware Reset

Problem:	When sending and receiving heavy traffic in Advanced Pass-Through mode, the 82573E might experience a firmware reset.
Implication:	During reset processing, APT packets may be lost.
Workaround:	This will not occur if the Watchdog feature is disabled. The reset and possible packet loss can be avoided by not enabling this feature.
Status:	Intel does not expect to resolve this erratum for the 82573E Gigabit Ethernet Controller.

27 Corruption in LAN-to-BMC Pass Through Data

Problem:	When transmitting data from the LAN at 10/100 Mb/s to the SMBus, the data might be corrupted. The corruption does not occur if the traffic is limited to single-packet communications, such as a single packet is transmitted to the BMC and it is not followed by another packet until the first one has been transmitted over the SMBus.
Implication:	Loss of communication with the BMC.
Workaround:	Clear the CLK_CNT_1_4 bit in the EEPROM (word 0x0F, bit 7). This has been implemented in EEPROM images starting from version 3.2.2.
Status:	Intel does not expect to resolve this erratum for the 82573E Gigabit Ethernet Controller.

28 PCIe: Completion with Completer Abort (CA) or Unsupported Request (UR) Status is Considered Malformed

Problem:	If the 82573 receives a completion with CA or UR status, and all-zero length field, it recognizes the completion as a malformed completion. According to the PCIe specification, this completion is not malformed.
Implication:	If enabled, an error message is sent upstream (fatal/non-fatal, as implied by the severity of a malformed TLP error). The default is fatal.
Workaround:	None.
Status:	Intel does not expect to resolve this erratum for the 82573 Gigabit Ethernet Controller.

29 PCIe: PCIe Bus May Halt if There are Less Than 16 Posted Data Flow Control Credits (= 256-Byte Memory Writes)

Problem:	The 82573 PCIe transmit bus halts when entering D3/L1 power states if the upstream device issues less than 16 posted data flow control credits.
Implication:	PCIe bus stops with no communication to the upstream device.
Workaround:	Upstream PCIe device must issue at least 16 posted data type credits.
Status:	Intel does not expect to resolve this erratum for the 82573 Gigabit Ethernet Controller.

30 PCIe: Surprise Down Following Externally-Initiated Recovery

Problem:	When the upstream switch initiates link recovery by transmitting TS1 ordered sets and the timing is such that the 82573 is just finishing transmitting a TLP from the alignment layer and the SKP interval timer expires and there is another TLP waiting, the 82573 repeatedly transmits the first 8 bytes of the next TLP, then goes into recovery, then continues to transmit the same 8 bytes until the switch attempts recovery again and times out, resulting in a Surprise Down Error.

Implication: A Surprise Down error might be a fatal error in operating systems that fully support Advanced Error Reporting (AER). However, this is a very rare event in normal operation since recovery during transmission should not occur except under error conditions and only a very specific timing of the recovery relative to a TLP transmission and the SOS interval causes the failure.

Workaround: None.

Status: Intel does not expect to resolve this erratum for the 82573 Gigabit Ethernet Controller.

31 Missing Interrupt Following ICR Read

Problem: If the Interrupt Cause Register (ICR) is read when at least one bit is set in the interrupt mask register and INT_ASSERTED is set to 0b, a new interrupt event occurring on the same clock cycle as the ICR read is ignored.

Implication: Missed interrupts leading to delays in responding to interrupt events. Specifically, this can cause a delay in processing a received packet.

Typically, the ICR is only read in response to an interrupt so this problem does not occur. However, when using legacy interrupts and sharing interrupts between devices, software might poll all the devices to find the source of the interrupt, including those devices that did not assert an interrupt. There might also be other situations in non-Intel drivers where ICR is polled even when no interrupt has been asserted.

Workaround: If reading ICR when there is no active interrupt cannot be avoided, clear the mask register (by writing 0xFFFFFFFF to IMC) before reading ICR. Note that in this case the ICR is cleared when read even if INT_ASSERTED is set 0b.

Status: Intel does not expect to resolve this erratum for the 82573 Gigabit Ethernet Controller.

32 PCIe: Missing Replay Due to Recovery During TLP Transmission

Problem: If the replay timer expires during the transmission of a TLP and the LTSSM moves from L0 to recovery during the transmission of the same TLP, the expected replay does not occur. Additionally, the replay timer is disabled, so no further replays occur unless a NAK is received.

Implication: This situation should not occur during normal operation. If it does occur while the upstream switch is waiting for a replay, the result would be a Surprise Down error which might halt the system.

Workaround: None.

Status: Intel does not expect to resolve this erratum for the 82573 Gigabit Ethernet Controller.

33 PCIe: Reception of Completion That Should Be Dropped May Occasionally Result In Device Hang or Data Corruption

Problem: This erratum can occur when the 82573 Family PCIe receives a completion that should be dropped, while the 82573 is starting a new request with the same TAG as the completion. On an error-free PCIe link, this situation should never occur since the 82573 Family does not assert a second request with the same tag as an outstanding request. Errors that could cause this failure:

- The TAG of a completion is corrupted due to noise on the line. This completion packet will be dropped due to LCRC error, but it could cause a failure if by chance a new request is asserted with the corrupted TAG value at the same time.
- An edge case of ACK timers results in a replay of a completion. This could cause the same case as previously mentioned.

Implication: When the failure occurs, the actual completion data from the new request is corrupted. The implications of this corruption of the read data depend on the type of request the 82573 Family was starting to send and are described as follows:

- TX descriptor with TSO – 82573 Family offload machine might hang.
- TX data or Tx descriptor without offload – 82573 Family transmits a packet on the network with invalid data but a valid CRC.
- RX descriptor – 82573 Family DMA's a receive packet to the wrong memory address.

Workaround: Keeping bit 13 "ACK/NACK scheme", word 0x1A "PCIe Initialization Configuration 3" set to 0b in the EEPROM image will minimize the chances of an ACK timeout.

Status: Intel does not expect to resolve this erratum for the 82573 Family of Gigabit Ethernet Controllers.

34 LED Turns on at Software Reset of MAC

Problem: If the most recent link was at 10/100 Mb/s, performing a software reset (setting CTRL.RST) without a PHY reset causes the LEDs to indicate that a link is up even if there is currently no link.

Implication: The Intel driver performs a software reset as part of the shutdown (S5) sequence. Under the following conditions, a shutdown results in the LEDs indicating link up.

- The most recent established link was at 10/100 Mb/s.
- The link is currently down.
- Manageability and/or Wake on LAN (WoL) is enabled such that the PHY will not be powered down in D3. This LED(s) remains lit until either a link is established or the PHY is reset on power up.

Workaround: There are two workaround options:

1. Enable Wake on Magic Packet from power off in the driver properties of the existing driver. This causes the driver to perform a PHY reset and initialization after the software reset in the shutdown sequence.
2. Modify the driver to perform a PHY reset and initialization as part of every software reset. Note: When using either of these workarounds, the LED(s) might still blink on as a result of the software reset, but will not remain on.

Status: Intel will provide an updated driver that implements workaround #2.

35 Receive Packet Delayed When Using RDTR or RADV Register

Problem: When using the RDTR and/or RADV timer mechanisms, there could be a situation where the write-back timer is incorrectly disabled, which prevents the write-back of a receive descriptor until another packet arrives.

Implication: No packet loss occurs. There might however, be a large delay between the time an Rx packet is received in the device and the time the descriptor is written back to memory, and finally an interrupt generated.

Workaround: It is recommended that the RDTR and RADV registers not be used for moderating Rx interrupts. The preferred solution is to use the Interrupt Throttling Register (ITR).

Status: Intel does not expect to resolve this erratum for the 82573 Gigabit Ethernet Controller.

36 The 82573 Overwrites Transmit Descriptors in Internal Buffer

Problem: This erratum occurs when the internal transmit descriptor buffer is nearly full of descriptors. If the free space in this buffer is smaller than the system cacheline, the calculation of the size of the descriptor fetch might be incorrect.

Implication: Corruption of the transmit descriptor ring; can cause a system crash. In most applications, the descriptors are written back as soon as the data has been read and they are not accumulating in the internal buffer, therefore this issue is not seen. However, in an application where system events such as PCIe flow control prevent the immediate write-back of descriptors, the descriptor buffer could fill up and this issue could be seen.

Workaround: The driver should keep track of the difference between the transmit head and tail and make sure the difference between tail and head is never more than the following values.

Cacheline	Maximum Value (TDT-TDH)
32 bytes	62
64 bytes	60
128 bytes	56
256 bytes	48

Status: Intel does not expect to resolve this erratum for the 82573 Gigabit Ethernet Controller.

37 Enabling or Disabling RSS with Data in the Rx FIFO May Cause an Rx Hang

Problem: Enabling RSS consists of setting both the Multiple Receive Queues Enable bit in MRQC and the Packet Checksum Disable bit in RXCSUM. Changing these settings while there is data in the receive data FIFO could cause the receive DMA to hang. There may be data present in the receive data FIFO even before the driver initialization is executed if the manageability firmware routes some packets to the host using MANC2H.

Implication: No data is received.

Workaround: Do not use RSS.

Status: Intel does not expect to resolve this erratum for the 82573 Gigabit Ethernet Controller. Intel drivers do not enable RSS.

38 PCIe: SKP Ordered Set Resets Training Sequence Counter

Problem: If a SKP ordered set is received during a TS1 or TS2 sequence, the TS counter is cleared. This will generally not be a problem since the upstream device should transmit at least 16 TS2 ordered sets, and the 82573 only needs to detect 8 consecutive TS2 ordered sets to complete the Recovery process, so a single reset of the counter will not cause a failure. A failure can occur if the upstream device is non-compliant and transmits fewer than 16 TS2 ordered sets. In this case, 82573 could fail to complete the Recovery process and then the PCIe link would go down.

Implication: There should be no failure when the upstream device functions according to the PCIe spec. If the upstream device is non-compliant, this issue could result in a Surprise Down error.

Workaround: None.

Status: Intel does not expect to resolve this erratum for the 82573 Gigabit Ethernet Controller.

39 Tx Data Corruption When Using TCP Segmentation Offload

Problem: When using TSO, a situation can occur where a PCIe MRd request is repeated with the same address, resulting in data corruption. At the end of the TCP packet, the Tx DMA hangs because the length doesn't match. This can only occur when the following are true:

- The first buffer of the packet is larger than [3 * (max_read_request - 4)].
- There is a 4 KB boundary within 64 bytes following the end of the header bytes in the buffer

Implication: Possible data corruption since a TCP packet is transmitted containing the wrong data but with the correct checksum.

Data transmission halts as the Tx DMA module enters a hang state.

Workaround: The failure can be avoided by ensuring at least one of the following:

- The buffer containing the headers should not be larger than [3 * (max_read_request - 4)]. To meet this requirement even for the minimum value of 128 bytes for max_read_request, the buffer should not be larger than 372 bytes.
- The alignment of the buffer containing the headers should be such that there is no 4 KB boundary within 64 bytes following the end of the header bytes. Assuming standard Ethernet/IP/TCP headers of 54 bytes, this means that the buffer should not start 54-118 bytes before a 4 KB boundary. For example, 128-byte alignment for this buffer could be used to fulfill this condition.

Status: Intel does not expect to resolve this erratum for the 82573 Gigabit Ethernet Controller.

40 PCIe: Spurious SDP/STP Causes Packets to be Dropped

Problem: When a spurious SDP or STP symbol is received without a corresponding END symbol, the alignment of the received data presented to the link layer might be incorrect. In this case, any following DLLPs or TLPs are dropped. This situation continues until there is an END symbol received that is not immediately followed by an SDP or STP symbol.

During normal operation, the SKP Ordered Sets that are inserted between packets guarantee that the proper alignment is restored within a short time.

Implication: Usually, this issue causes nothing more than a replay of a few TLPs. The 82573 recovers from this situation autonomously.

If the 82573 is connected to an ICH7, a spurious SDP or STP symbol that occurs just before entering L1 could cause a hang of the PCIe link since the ICH7 does not insert SOS when transmitting PM_Request_ACK DLLPs, so the 82573 does not receive them and never enters L1.

Workaround: If the 82573 is connected to an ICH7, ASPM L1 should be disabled. For the 82573L, clearing EEPROM word 0x1E, bit 13 prevents it from repeatedly returning to L1 in the D3 state.

Otherwise, no workaround is required.

Status: Intel does not expect to resolve this erratum for the 82573 Gigabit Ethernet Controller.

41 Delay of Received Ethernet Packet During ASPM L1

Problem: When ASPM L1 is enabled and a single small Ethernet packet is received during L1, the device may stay in L1 and delay the transfer of the packet to memory until some other event returns the link state to L0.

Implication: Large latency on single small received packets. This situation can be seen in large latencies reported when pinging the device.

Workaround: If the latency of single packets is important, disable ASPM L1.

The Intel Linux driver disables ASPM L1 starting with Release 15.6 (e1000e v1.2.17.

Status: Intel does not expect to resolve this erratum for the 82573L Gigabit Ethernet Controller.

42 Packets Received with an L2+L3 Header Length Greater than 256 Bytes Can Incorrectly Report a Checksum Error

Problem:	L2/L3 packets with long/multiple next header extensions incorrectly report a Receive checksum error when the length from Destination Address (DA) to the beginning of the TCP/UDP header is greater than 256 bytes.
Implication:	A receive checksum error can incorrectly be reported by the device, even if there is no checksum error.
Workaround:	When the driver receives a packet with a checksum error reported by the hardware, software should check the L2/L3 header length. If the L2/L3 header length is 256 bytes or greater, software should verify the checksum. The Intel Windows and Linux drivers address this issue by passing packets with bad checksums to the network stack for further examination.
Status:	Intel does not expect to resolve this erratum for the 82573 Gigabit Ethernet Controller.

43 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:	Possible device receive hang if these malformed IPv6 headers are received.
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, hardware no longer offloads the receive checksums correctly for incoming frames with IPv6 extension headers. Also note that software needs to account for this. This issue is addressed in current Intel software device drivers.
Status:	Intel does not expect to resolve this erratum for the 82573 Gigabit Ethernet Controller.

7.0 Specification Clarifications

1. PCIe: Completion Timeout Mechanism Compliance

If the latency for PCIe completions in a system is above 21 ms and PCIe completion timeout mechanism is enabled, there may be unpredictable system behavior.

The 82573 Family complies with the PCIe 1.0a specification for completion timeout mechanism. The PCIe 1.0a specification provides a timeout range between 50 µs to 50 ms with a strong recommendation that it be at least 10 ms. The 82573 uses a range of 21-42 ms.

The completion timeout value in a system must be above the expected maximum latency for completions in the system in which the 82573 Family is installed. This will ensure that the 82573 Family receives the completions for the requests it sends out, avoiding a completion timeout scenario. If the latency for completions is above 21 ms this can result in the device timing out prior to a completion returning. In the event of a completion timeout, per direction in the PCIe specification the device assumes the original completion is lost, and resends the original request. In this condition, if the completion for the original request arrives at the 82573 Family devices, this will result in two completions arriving for the same request, which may cause unpredictable system behavior.

Therefore, if the PCIe completion latency for a system cannot be guaranteed to be lower than 21 ms, the PCIe completion timeout mechanism should be disabled by setting GCR.Disable_timeout_mechanism.

For more details on Completion Timeout operation in the 82573 Family refer to the *Intel® 82573 Family of GbE Controllers Datasheet* and the *PCIe* GbE Controllers Open Source Software Developer's Manual*.

2. Manageability: Critical Session (Keep PHY Link Up) Mode Does Not Block All PHY Resets Caused By PCIe Resets (82573E/V)

When Critical Session Mode (Keep PHY Link Up) is enabled (via the SMBus Management Control command), a transition from D3 to D0 without a general PCIe reset still causes a PHY reset.

3. Use of Wake on LAN Together With Manageability (82573E Only)

The Wakeup Filter Control Register (WUFC) contains the NoTCO bit, which affects the behavior of the wakeup functionality when manageability is in use. Note that if manageability is not enabled, the value of NoTCO has no effect.

When NoTCO contains the hardware default value of 0b, any received packet that matches the wakeup filters will wake the system. This could cause unintended wakeups in certain situations. For example, if Directed Exact Wakeup is used and the manageability shares the host's MAC address, IPMI packets that are intended for the BMC wakes the system, which might not be the intended behavior.

When NoTCO is set to 1b, any packet that passes the manageability filter, even if it also is copied to the host, is excluded from the wakeup logic. This solves the previous problem since IPMI packets do not wake the system. However, with NoTCO=1b, broadcast packets, including broadcast magic packets, do not wake the system since they pass the manageability filters and are therefore excluded.

Effects of NoTCO Settings

WoL	NoTCO	Share MAC Address	Unicast packet	Broadcast Packet
Magic Packet	0b	N/A	OK	OK
Magic Packet	1b	Y	No wake	No wake.
Magic Packet	1b	N	OK	No wake.
Directed Exact	0b	Y	Wake even if MNG packet. No way to talk to the BMC withoug waking host.	N/A
Directed Exact	0b	N	OK	N/A
Directed Exact	1b	N/A	OK	N/A

Intel Windows* drivers set NoTCO by default.

8.0 Software Clarifications

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

The 82573 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

When bit-banging on a serial interface (such as SPI, I^2C, 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”.

9.0 Document Changes

There are no document changes to report at this time.