



82544EI/82544GC
Gigabit Ethernet Controller
Specification Update

October 10, 2005

The 82544EI/82544GC Gigabit Ethernet Controller may contain design defects or errors known as errata that may cause the product to deviate from published specifications. Current characterized errata are documented in this Specification Update.

82544EI/82544GC SPECIFICATION UPDATE

Information in this document is provided in connection with Intel products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The 82544EI/82544GC Gigabit Ethernet Controller may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an ordering number and are referenced in this document, or other Intel literature may be obtained by calling 1-800-548-4725 or by visiting Intel's web site at <http://www.intel.com>.

Copyright © Intel Corporation, 2001 – 2005

Intel® is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

* Third-party brands and names are the property of their respective owners

CONTENTS

CONTENTS	1
PREFACE	4
NOMENCLATURE	4
COMPONENT IDENTIFICATION VIA PROGRAMMING INTERFACE	5
GENERAL INFORMATION	5
82544EI/82544GC Component Marking Information.....	5
SUMMARY TABLE OF CHANGES	7
Codes Used in Summary Tables.....	7
SPECIFICATION CHANGES	10
1. TCP Segmentation (Large Send Offload).....	10
ERRATA	11
1. Erroneous Response to I/O Cycles	11
2. Bridge I/O Window Aperture May Lead to I/O Access Problems.....	11
3. Device Does Not Always Check for PAR64 Errors	11
4. Undersize Wakeup Packets Can Cause Wakeup.....	11
5. Transmit Delayed Interrupt Canceled but No Immediate Interrupt Generated.....	12
6. Message Signaled Interrupts Require Quad Word Data Alignment	12
7. Receive Descriptor Writeback Problems for Packets Spanning Multiple Buffers.....	12
8. I/O Read Cycles Can Cause Subsequent Incorrect Data Reads	12
9. Certain Registers Cannot Be Written with Particular Alignments in PCI-X Bus Operation.....	13
10. Multiple Buffers Needed for Jumbo Frames Larger than 8 Kilobytes in PCI-X Bus Operation.....	13
11. PCI-X Violation of FRAME# and GNT# Protocol	13
12. EEPROM FLASH Disable Bit Must Not Be Set	14
13. Improper Factory Test Pin Initialization	14
14. PCI-X FLASH Memory Write Problem with Specific Chipset	15
15. Illegal Oversize Packets Overflow Receive FIFO	15
16. Transmit TCP Checksum Incorrectly Modified if Calculated as 0x0000	15
17. Odd Offset Register Writes in PCI-X Bus Operation.....	16
18. Link Failures with Short Cables.....	16
19. System Hang Due to Host Block Requests	16
20. Transmit Descriptor Writeback Problems with Non-Zero WTHRESH.....	17
21. Bus Initialization with Some Chipsets.....	17
22. Packet Reception with APM Enabled before Driver Load.....	17
23. Intermittent Power-on State Due to Decoded High-Impedance Test Modes	18
24. 32-Bit Split-Completion Dependency on subsequent REQ64#	18
25. Message Signaled Interrupt Feature May Corrupt Write Transactions.....	19
26. Link Establishment or Communication Problems in Fiber Mode When Link Partner Does Not Fully Comply with the IEEE 802.3 Specification	19
27. Wakeup Packet Memory (WUPM) cleared upon reset.....	19

82544EI/82544GC SPECIFICATION UPDATE

28.	Unexpected RCMP ACK packets in ASF mode.....	20
29.	Exceeding PCI Power Management Specification Limit of 375mA current during reset and power state transitions.....	20
30.	Memory Access must be enabled in order to Read Device Registers in I/O mode.....	20
31.	Inbound and Outbound reads not fully decoupled in PCI-X mode.....	20
32.	Hang in PCI-X systems due to 2k Buffer Overrun during Transmit Operation	21
33.	CRC Errors due to Rate Adaptation FIFO Overflow in Fiber Mode	21
34.	Transmit Descriptors May Be Written Back to Host, Even Without the RS Bit Set	22
35.	Polarity Detection Error May Cause Inability to Transmit.....	22
SPECIFICATION CLARIFICATIONS		23
1.	Receiver Enabling and Disabling.....	23
2.	LEDs Inactive Until Driver Loads.....	23
3.	PHY Reset Duration.....	23
4.	WUC.APME does not return to value in EEPROM after soft reset.....	23
DOCUMENTATION CHANGES		23
1.	M66EN Signal Must Be Connected	23
2.	Octets Transmitted Counters Adjusted if VLAN Enabled	24
3.	Wake-up Packet Memory Not Writeable for Diagnostics	24
4.	External GMII Mode Operation.....	24
5.	TIPG Register Value Incorrectly Documented	26
6.	Clarification on Settings for PHY Registers 29 and 30.....	26
7.	Remove Transmit Report Status Sent Function.....	27
8.	Remove Transmit DMA Pre-fetching and Preemption Functions	28
9.	Remove Adaptive IFS Throttle Function (AIT)	28

REVISION HISTORY

82544EI/82544GC Gigabit Ethernet Controller Specification Update

Date of Revision	Description
April 1, 2004	Initial public release. Added errata # 25 – 27.
August 3, 2004	Added Errata # 28, 29, 30, and 31. Added spec clarification #4.
January 10, 2005	Added Errata #32 – 33.
October 10, 2005	Added Errata #34 and 35. Added component marking information for the lead-free 82544EI device.

82544EI/82544GC SPECIFICATION UPDATE

PREFACE

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

- 82544EI Gigabit Ethernet Controller Hardware Design Guide Application Note (AP-422), Intel Corporation.
- 82544GC Gigabit Ethernet Controller Hardware Design Guide Application Note (AP-427), Intel Corporation.
- 82544EI/82544GC Gigabit Ethernet Controller Developer's Manual, Intel Corporation.

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

All 82544 family 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.

NOMENCLATURE

Specification Changes are modifications to the current published specifications. These changes will be incorporated in the next release of the specifications.

Errata are design defects or errors. Errata may cause 82544EI/82544GC device behavior to deviate from published specifications. Hardware and software designed to be used with any given stepping must assume that all errata documented for that stepping are present on all devices.

Specification Clarifications describe a specification in greater detail or further highlight a specification's impact to a complex design situation. These clarifications will be incorporated in the next release of the specifications.

Documentation Changes include typos, errors, or omissions from the current published specifications. These changes will be incorporated in the next release of the specifications.

COMPONENT IDENTIFICATION VIA PROGRAMMING INTERFACE

82544EI/82544GC controller stepping can be identified by the following register contents:

82544EI/82544GC Stepping	Vendor ID	Device ID	Revision Number
A0	8086h	1008h	00h
A1	8086h	1008h	01h
A2	8086h	1008h	02h
A3	8086h	1008h	02h*
A4	8086h	1008h	02h*

The device also provides an identification data through the Test Access Port. The version number for A0, A1, A2, A3 and A4 steppings is 0001b.

* The revision number did not change beyond the A2 stepping.

GENERAL INFORMATION

This section covers the 82544EI/82544GC device.

82544EI/82544GC COMPONENT MARKING INFORMATION



Product	Stepping	QDF Number	Top Marking	Notes
82544EI	A0	Q457	FW82544EI	Engineering Samples
82544EI	A1	Q458	FW82544EI	Engineering Samples
82544EI	A2	Q461	FW82544EI	Engineering Samples
82544EI	A3	Q481	FW82544EI	Internal Engineering Only
82544EI	A4	Q483	FW82544EI	Engineering Samples
82544EI	A4	-	FW82544EI	Production Units
82544EI	A4	-	NH82544EI	Lead-Free Production Units
82544GC	A0	Q476	RC82544GC	Engineering Samples
82544GC	A1	Q477	RC82544GC	Engineering Samples
82544GC	A2	Q480	RC82544GC	Engineering Samples
82544GC	A3	Q482	RC82544GC	Internal Engineering Only
82544GC	A4	Q484	RC82544GC	Engineering Samples
82544GC	A4	-	RC82544GC	Production Units

The following page contains topside package drawings.

Note: component marks are subject to change when the device reaches production.

82544EI/82544GC SPECIFICATION UPDATE

Note: Devices that are lead-free are marked with a circled “e1” and have a product code: NH82544EI.



SUMMARY TABLE OF CHANGES

The following table indicates the Specification Changes, Errata, Specification Clarifications or Documentation Changes, which apply to the listed 82544EI/82544GC steppings. Intel intends to fix some of the errata in a future stepping of the component, and to account for the other outstanding issues through documentation or Specification Changes as noted. This table uses the following notations:

CODES USED IN SUMMARY TABLES

X:	Erratum, Specification Change or 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.
NoFix:	There are no plans to fix this erratum.
(No mark) or (Blank Box):	This erratum is fixed in listed stepping or specification change does not apply to listed stepping.
Shaded:	This item is either new or modified from the previous version of the document.

82544EI/82544GC SPECIFICATION UPDATE

No.	A0	A1	A2	A3	A4	Plans	Specification Changes	Page	Notes
1	X	X	X	X	X	Spec Change	TCP Segmentation (Large Send Offload)	10	-
-	-	-	-	-	-	-	-	-	-
No.	A0	A1	A2	A3	A4	Plans	ERRATA	Page	Notes
1	X	-	-	-	-	Fixed	Erroneous Response to I/O Cycles	11	-
2	X	-	-	-	-	Fixed	Bridge I/O Window Aperture May Lead to I/O Access Problems	11	-
3	X	X	X	X	X	NoFix	Device Does Not Always Check for PAR64 Errors	11	-
4	X	X	X	X	X	NoFix	Undersize Wakeup Packets Can Cause Wakeup	11	-
5	X	X	X	X	X	NoFix	Transmit Delayed Interrupt Canceled but No Immediate Interrupt Generated	12	-
6	X	-	-	-	-	Fixed	Message Signaled Interrupts Require Quad Word Data Alignment	12	-
7	X	X	-	-	-	Fixed	Receive Descriptor Writeback Problems for Packets Spanning Multiple Buffers	12	-
8	X	X	-	-	-	Fixed	I/O Read Cycles Can Cause Subsequent Incorrect Data Reads	12	-
9	X	X	-	-	-	Fixed	Certain Registers Cannot Be Written with Particular Alignments in PCI-X Bus Operation	13	-
10	X	X	X	X	X	NoFix	Multiple Buffers Needed for Jumbo Frames Larger than 8Kbytes in PCI-X Bus Operation	13	-
11	X	X	X	X	X	NoFix	PCI-X Violation for FRAME# and GNT# Protocol	13	-
12	X	X	X	X	X	NoFix	EEPROM FLASH Disable Bit Must Not Be Set	14	-
13	X	X	X	X	X	NoFix	Improper Factory Test Pin Initialization	14	-
14	X	X	X	X	X	NoFix	PCI-X FLASH Memory Write Problem with Specific Chipset	15	-
15	X	X	X	X	X	NoFix	Illegal Oversize Packets Overflow Receive FIFO	15	-
16	X	X	X	X	X	NoFix	Transmit TCP Checksum Incorrectly Modified if Calculated as 0x0000	15	-
17	X	X	X	X	X	NoFix	Odd Offset Register Writes in PCI-X Bus Operation	16	-
18	X	X	X	X	X	NoFix	Link Failures with Short Cables	16	-
19	X	X	X	X	X	No Fix	System Hang Due to Host Block Requests	16	-
20	X	X	X	X	X	NoFix	Transmit Descriptor Writeback Problem with Non-Zero WTHRESH	17	-
21	X	X	X	X	X	NoFix	Bus Initialization with Some Chipsets	17	-
22	X	X	X	X	X	NoFix	Packet Reception with APM Enabled before Driver Load	17	-
23	X	X	X	X	X	No Fix	Intermittent Power-on State Due to Decoded High-Impedance Test Modes	18	-
24	X	X	X	X	X	NoFix	32-Bit Split-Completion Dependency on subsequent REQ64#	18	-
25	X	X	X	X	X	NoFix	Message Signaled Interrupt Feature May Corrupt Write Transactions	19	-
26	X	X	X	X	X	NoFix	Link Establishment or Communication Problems in Fiber Mode When Link Partner Does Not Fully Comply with the IEEE 802.3 Specification	19	-
27	X	X	X	X	X	NoFix	Wakeup Packet Memory (WUPM) cleared upon reset	19	-
28	X	X	X	X	X	NoFix	Unexpected RCMP ACK packets in ASF mode	20	-
29	X	X	X	X	X	NoFix	Exceeding PCI Power Management Specification Limit of 375mA current during reset and power state transitions	20	-
30	X	X	X	X	X	NoFix	Memory Access must be enabled in order to Read Device Registers in I/O mode	20	-

82544EI/82544GC SPECIFICATION UPDATE

No.	A0	A1	A2	A3	A4	Plans	ERRATA	Page	Notes
31	X	X	X	X	X	NoFix	Inbound and Outbound reads not fully decoupled in PCI-X mode	20	-
32	X	X	X	X	X	NoFix	Hang in PCI-X systems due to 2k buffer overrun during transmit operation	21	-
33	X	X	X	X	X	NoFix	CRC Errors due to Rate Adaptation FIFO Overflow in Fiber Mode	21	-
34	X	X	X	X	X	NoFix	Transmit Descriptors May Be Written Back to Host, Even Without the RS Bit Set	22	New
35	X	X	X	X	X	NoFix	Polarity Detection Error May Cause Inability to Transmit	22	New
-	-	-	-	-	-	-	-	-	-
No.	A0	A1	A2	A3	A4	Plans	Specification Clarification	Page	Notes
1	X	X	X	X	X	Doc Change	Receiver Enabling and Disabling	23	-
2	X	X	X	X	X	Spec Change	LEDs Inactive Until Driver Loads	23	-
3	X	X	X	X	X	Spec Change	PHY Reset Duration	23	-
4	X	X	X	X	X	Spec Change	WUC.APME does not return to value in EEPROM after soft reset	23	-
-	-	-	-	-	-	-	-	-	-
No.	A0	A1	A2	A3	A4	Plans	DOCUMENTATION CHANGES	Page	Notes
1	X	X	X	X	X	Doc Change	M66EN Signal Must Be Connected	23	-
2	X	X	X	X	X	Doc Change	Octets Transmitted Counters Adjusted if VLAN Enabled	24	-
3	-	-	X	X	X	Doc Change	Wake Up Packet Memory Not Writeable for Diagnostics	24	-
4	X	X	X	X	X	Doc Change	External GMII Mode Operation	24	-
5	-	-	-	-	X	Doc Change	TIPG Register Value Incorrectly Documented	26	-
6	-	-	-	-	X	Doc Change	Clarification on Settings for PHY Registers 29 and 30	26	-
8	X	X	X	X	X	Doc Change	Remove Transmit Report Status Sent Function	27	-
9	X	X	X	X	X	Doc Change	Remove Transmit DMA Pre-fetching and Preemption Functions	28	-
10	X	X	X	X	X	Doc Change	Remove Adaptive IFS Throttle Function	28	-
-	-	-	-	-	-	-	-	-	-

SPECIFICATION CHANGES

1. TCP Segmentation (Large Send Offload)

Problem: Due to multiple issues with regards to TCP segmentation (Large Send Offload) using the 82544EI/GC Ethernet Network Controllers, it was de-featured and thereby not supported.

Affected Docs: 82544EI Gigabit Ethernet Controller Hardware Design Guide Application Note (AP-422), Rev. 0.77, document number A44740-003 and 82544GC Gigabit Ethernet Controller Hardware Design Guide Application Note (AP-427), Rev. 0.77, document number A67149-001.

ERRATA

1. Erroneous Response to I/O Cycles

- Problem:** The 82544EI/82544GC controller will respond to and accept all I/O bus accesses (read or write) with addresses containing AD [2:0] =000b. This behavior will occur regardless of whether the address matches the programmed address range through the I/O Base Address Register or whether IO_EN is set. I/O accesses are associated with legacy PC devices such as serial ports, keyboard controllers and mouse ports.
- Implication:** Systems may fail to boot or will blue screen.
- Workaround:** Place the 82544EI/82544GC controller on a subordinate PCI or PCI-X bus with no other I/O mapped devices. In systems with multiple PCI buses, the subordinate bus is likely to be the bus with the highest speed and bus width.
- Status:** Intel resolved this erratum in the A1 stepping of the 82544EI/82544GC Gigabit Ethernet Controller. However, correcting this erratum revealed erratum #9, "I/O Read Cycles Can Cause Subsequent Incorrect Data Reads".

2. Bridge I/O Window Aperture May Lead to I/O Access Problems

- Problem:** The 82544EI/82544GC device has an I/O window size of eight bytes. Some chipsets or bridges may support a minimum I/O window size of 32 bytes for forwarding I/O cycles. Certain operating systems such as Windows* may not be able to reconcile this difference and may configure the bridge aperture to be zero rather than 32.
- Implication:** In affected systems, the driver will not load correctly because its I/O resources cannot be allocated correctly.
- Workaround:** It may be possible to overcome the problem with BIOS settings. In one case, enabling "hot plug" in the BIOS forced the system to allocate sufficient I/O resources.
- Status:** Intel resolved this erratum in the A1 stepping of the 82544EI/82544GC Gigabit Ethernet Controller.

3. Device Does Not Always Check for PAR64 Errors

- Problem:** During PCI address phases, PCI DAC second address phases, or attribute phases, an 82544EI/82544GC is supposed to check parity on the AD[63:32] bus, the CBE#[7:4] pins and PAR64, reporting errors on SERR# if enabled. Instead, the Ethernet controller ignores PAR64 checking during address and attribute phases.
- Implication:** Errors on 64-bit address cycles will go undetected by the 82544EI/82544GC device, but should be detected by any other target devices on the bus. Errors during data phases on the same signal lines will be captured as usual. Thus, the system still maintains some protection against signal integrity problems on the bus.
- Workaround:** None.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controller.

4. Undersize Wakeup Packets Can Cause Wakeup

- Problem:** Wakeup packets smaller than 64 bytes are supposed to be ignored. However, packets less than 64 bytes that otherwise appear to be legitimate wakeup packets may result in PME# assertion and wakeup.
- Implication:** Spurious wakeups are possible in response to undersize wakeup packets. Such packets are not expected to occur in normal LAN operation.
- Workaround:** None.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controller.

5. Transmit Delayed Interrupt Canceled but No Immediate Interrupt Generated

- Problem:** The 82544EI/82544GC controller has a mechanism to cause delayed (TXDW) interrupts. If IDE is set in the transmit descriptor, a timer value is loaded from the TIDV register. When the timer expires, an interrupt is generated. If a later descriptor is processed with IDE=0, the controller correctly stops the countdown and clears the counter. However, it does not generate an immediate interrupt as expected.
- Implication:** Software drivers that rely on a mixture of delayed and non-delayed transmit interrupts may not work satisfactorily.
- Workaround:** Use the delayed interrupt feature for all transmit descriptor interrupts, or for none of them. Intel drivers typically use IDE=1 for all descriptors, but program diminishing timer values into TIDV when buffer resources are low.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controller.

6. Message Signaled Interrupts Require Quad Word Data Alignment

- Problem:** Message signaled interrupts must be programmed such that the message data is written to addresses aligned on quad word (8 byte) boundaries. The PCI specification only requires DWORD alignment.
- Implication:** If MSI writes are allowed to a DWORD aligned address, data corruption could occur.
- Workaround:** Use care in programming the Message Capability Structure or do not use the MSI capability.
- Status:** Intel resolved this erratum in the A1 stepping of the 82544EI/82544GC Gigabit Ethernet Controller.

7. Receive Descriptor Writeback Problems for Packets Spanning Multiple Buffers

- Problem:** Receive descriptors are typically written back to memory either upon receive interrupts or opportunistically in between writing data buffers. When a received Ethernet packet exceeds the size of a single receive buffer, corrupted receive descriptor writebacks may occur and the controller may hang. The conditions for this erratum are specific:
- The controller is programmed to write back receive descriptors upon a receive interrupt, the interrupt has not yet been triggered, and
 - The controller has a programmed descriptor writeback threshold (RXDCTL.WTHRESH), the number of receive descriptors consumed thus far by the packet is equal to or greater than the threshold.
- Implication:** Corrupted descriptor writebacks may include writing back unconsumed descriptors, descriptor writebacks to incorrect addresses, or writebacks missed altogether. In addition, the device may cease to access the PCI bus or cease packet reception. If the device hangs, a full software or hardware reset is needed.
- Workaround:** If the system uses buffers smaller than the maximum allowed packet size, take the following precautions:
- Configure the receive interrupt to occur immediately on end-of-packet by programming RIDV = 0.
 - Configure the descriptor writeback threshold WTHRESH to a value that will not result in a writeback in the middle of a packet. Packets may be 1514 bytes or up to 16K bytes if long packets are enabled. It is recommended that the RXDCTL.GRAN bit be set to 1 descriptor and WTHRESH set to the maximum number of descriptor buffers the maximum size packet will consume.
- Status:** Intel resolved this erratum in A2 stepping of the 82544EI/82544GC Gigabit Ethernet Controller.

8. I/O Read Cycles Can Cause Subsequent Incorrect Data Reads

- Problem:** For any I/O read cycle with AD [2:0] = 000b, the 82544EI/82544GC controller initiates an internal register read cycle. This behavior occurs even if the I/O cycle is not intended for the 82544EI/82544GC device. Subsequent reads from 82544EI/82544GC controller register space (I/O or memory operations) may return erroneous data. This problem is related to Erratum #1, "Erroneous Response to I/O Cycles", which was corrected.
- Implication:** Systems may boot, but are prone to crash or blue screen later.
- Workaround:** Place the 82544EI/82544GC controller on a subordinate PCI or PCI-X bus with no other I/O mapped devices. In systems with multiple PCI buses, the subordinate bus is likely to be the bus with the highest speed and bus width.
- Status:** Intel resolved this erratum in A2 stepping of the 82544EI/82544GC Gigabit Ethernet Controller.

9. Certain Registers Cannot Be Written with Particular Alignments in PCI-X Bus Operation

- Problem:** In PCI-X operation, if a bridge issues an access to a command/status register with an odd DWORD address (offset ending in 0x4) and REQ64# is asserted, the controller accepts the transaction (ACK64# asserted) but cannot support the access occurring on the high DWORD. Certain register offsets, including FCAH (Flow Control Address High), MTA (Multicast Table Array), and VFSA (VLAN Filter Table Array) cannot be reliably written.
- This problem does not affect bus operation in standard PCI mode.
- Implication:** Software will have limited access to VLAN and multicast filter tables.
- Workaround:** Use the default values for FCAH/FCAL. Intel software drivers changed to prevent accesses to the affected register offsets.
- Status:** Intel resolved this erratum in A2 stepping of the 82544EI/82544GC Gigabit Ethernet Controller.

10. Multiple Buffers Needed for Jumbo Frames Larger than 8 Kilobytes in PCI-X Bus Operation

- Problem:** In PCI-X operation, DMA requests for greater than 8 Kilobytes can result in incorrect DMA operation and 82544EI/82544GC controller can lock up. This problem can occur when very large buffers are used to transmit jumbo frames. The problem occurs only when using Legacy or standard TCP/IP descriptors. Descriptors for TCP Segmentation (Large Send) operations cannot generate DMA requests larger than 8 Kbytes regardless of the size of the data buffers.
- This problem does not affect bus operation in standard PCI mode, nor does it affect packet reception.
- Implication:** Jumbo frame transmissions may be slightly less efficient with buffer size limitations. The performance change will likely be imperceptible.
- Workaround:** Use multiple transmit buffers for jumbo frames exceeding 8 Kbytes in size. Intel software drivers changed to employ the workaround.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controller.

11. PCI-X Violation of FRAME# and GNT# Protocol

- Problem:** As a PCI-X initiator, the 82544EI/82544GC controller will not always assert FRAME# within 6 clocks after GNT# is asserted when the bus is idle. In PCI mode, 82544EI/82544GC controller does meet the 16 clock GNT# to FRAME# protocol.
- Implication:** Implications will vary based on specific chipset / arbiter. Of the PCI-X chipsets tested to date, only one arbiter checked for the 6 clocks delay and stopped servicing requests.
- Workaround:** The chipset has a register setting to allow for greater than 8 clock cycles GNT# to #FRAME timing.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/GC Gigabit Ethernet Controller.

12. EEPROM FLASH Disable Bit Must Not Be Set

Problem: If the *FLASH Disable Bit* is set in the EEPROM (Initialization Control Word 2, offset 0x0F, bit 8), the 82544EI/GC's *Expansion ROM Base Address Register (BAR)* will be set to 00000000 and the 82544EI/GC will incorrectly claim memory cycles from 00000000 to 64kB to 512kB, depending on the value of the *FLASH Size Valid Bits* in PCI configuration space. This will interfere with other devices mapped to that space.

Implication: The *FLASH Disable Bit* must not be set.

- LOM designs typically will not populate a FLASH device on the 82544EI/GC's FLASH memory interface, instead utilizing the main system ROM to store boot code as part of the system BIOS. BIOS facilities (such as CMOS setup programs) are used to enable/disable Boot ROM functionality, and these facilities do not interact with the EEPROM's *Flash Disable Bit*.
- Adapter cards typically will populate a FLASH device on the 82544EI/GC's FLASH memory interface. In systems where multiple adapter cards are used, having more than one device with its FLASH device enabled is problematic due to the fact that the BIOS will shadow boot code from each enabled device into the system's Upper Memory Block (UMB), consuming this finite resource.

Workaround: There are several potential workarounds:

- For LOM designs the workaround is to ensure that the FLASH Disable Bit is set to 0.
- On Intel's Server Adapter designs using the 82544EI/GC LAN controllers, Intel will ship with the FLASH interface enabled and a blank FLASH device on the FLASH interface. The lack of a valid signature and checksum will prevent the boot code from being shadowed into the UMB block.
- If a user wants to enable/disable the adapter's boot code, they will run the BROW utility, which will program/erase the FLASH device. From a user perspective this is 100% consistent with the model used to enable/disable adapter boot code on the LAN Adapters Intel ships today.

Status: Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controller.

13. Improper Factory Test Pin Initialization

Problem: The 82544EI/82544GC Controller has a factory test pin that can be asserted low to place the device in IDDQ test mode. The pin is designated as No Connect in documentation because it has an internal pull-up resistor. However, the pull-up device is gated by logic driven by the test pin input buffer and it is possible for the device to power up with IDDQ test mode enabled.

Implication: The IDDQ factory test mode is designed to test quiescent current leakage. If the device powers up in IDDQ test mode, the controller's I/O pins will be disconnected and it will not be able to communicate.

Workaround: The affected pin is ball N1 on the 82544EI controller and ball L3 on the 82544GC controller. Attach an external pull-up resistor to the test pin to ensure the IDDQ test mode is disabled. Use a common value resistor such as 1K ohms (the value is not critical). Alternatively, the pin may be connected directly to the 3.3V supply. Future revisions of product documentation will indicate these balls as TEST1 instead of No Connect. The factory test pin that is already named "TEST" will be renamed to TEST0 for consistency.

Status: Intel does not plan to resolve this erratum in a future stepping of the 82544EI/GC Gigabit Ethernet Controller.

14. PCI-X FLASH Memory Write Problem with Specific Chipset

- Problem:** In certain PCI-X systems, utility programs cannot properly write FLASH memory connected to the Ethernet controller.
- Corrupted FLASH write cycles have been observed in a particular PCI-X based 8-way server based on the Profusion chipset. The problem occurs when the PB64 PCI-X bridge attempts to write the FLASH using valid 64-bit PCI-X bus cycles. The 82544EI/82544GC controller correctly signals that it does not accept the 64-bit cycle by deasserting ACK64. Then the system re-attempts the write operation with two 32-bit cycles. One of those cycles is a "zero byte write" cycle that does not assert any C/BE# byte enable signals. The zero byte write cycle confuses the FLASH memory controller in the 82544EI/82544GC device, resulting in a spurious bus cycle to the FLASH memory. CS# and WE# are likely to be asserted improperly.
- Several 8-way servers use the Profusion chipset and it is believed that all of them can encounter the problem. Other PCI-X systems are not affected because they do not issue the 64-bit write cycle followed by 32-bit cycles with a zero byte write. Problems have not been observed in standard PCI-based systems. LAN-on-motherboard designs do not generally use FLASH memory connected to the 82544EI/82544GC controller.
- Implication:** Unmodified FLASH programming utilities will not work on certain PCI-X systems.
- Workaround:** Use a FLASH programming utility with I/O FLASH addressing instead of memory-mapped FLASH addressing. Intel has modified the FBOOT utility (v4.16) and the FUTIL (v4.16) utilities to use I/O FLASH addressing. The GT utility already had an I/O addressing mode. Intel's PROSet LAN adapter utility program and the BROW utility (v2.16) are currently being modified.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/GC Gigabit Ethernet Controller.

15. Illegal Oversize Packets Overflow Receive FIFO

- Problem:** The controller should drop invalid Ethernet packets, but frames exceeding the maximum legal size can overflow the receive FIFO, causing a lock up. This problem has only been reported in a test environment with an IXIA packet generator.
- Implication:** The 82544EI/82544GC controller can receive jumbo frames up to 16K without difficulty. If the controller locks up due to an oversize packet, a full software or hardware reset is needed.
- Workaround:** Driver software should ensure that a minimum of 16K is allocated to the receive FIFO. Packets larger than this size should not be present on the LAN.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controller.

16. Transmit TCP Checksum Incorrectly Modified if Calculated as 0x0000

- Problem:** If the controller calculates a transmit TCP checksum as 0x0000, it will automatically change the checksum to 0xFFFF.
- Specifications call for 0xFFFF be substituted for 0x0000 for UDP packets to distinguish UDP packets that carry no checksum. However, the modification does not apply to TCP packets.
- Implication:** If the receiving station is running MS-DOS and calculates a receive checksum of 0x0000, it will flag an error if the checksum contained in the packet is 0xFFFF. Other operating systems treat 0x0000 and 0xFFFF as equivalent in one's complement math. UDP checksums are correct.
- Workaround:** Intel modified the DOS Ethernet driver to check for a received checksum of 0xFFFF on a TCP/IP packet and change it back to 0x0000 before passing the packet to the operating system.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controller.

17. Odd Offset Register Writes in PCI-X Bus Operation

- Problem:** Erratum #10, Certain Registers Cannot Be Written with Particular Alignments in PCI-X Bus Operations, was only partially fixed in the A2 stepping.
- In PCI-X operation, if a bridge issues a 32-bit access to a command/status register with an odd DWORD address (offset ending in 0x4), the controller will alias the write to the even address. For example, a write to offset 5604h will also write 5600h. The affected addresses are in the MTA (Multicast Table Array) and VFSA (VLAN Filter Table Array).
- This problem does not affect bus operation in standard PCI mode.
- Implication:** Software will have limited access to VLAN and multicast filter tables in PCI-X systems.
- Workaround:** Intel software drivers are modified to no longer access the affected register offsets.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controller.

18. Link Failures with Short Cables

- Problem:** When performing external loopback tests on products utilizing the 82544 silicon, if board traces are extremely short, it has been found that high frequency noise on the differential traces does not have enough trace impedance to dissipate. This can cause link failures on loopback tests when cable length is less than 1 meter. This could also be seen in backplane designs utilizing the 82544 silicon.
- Implication:** Link failures could occur during loopback testing, causing erroneous testing results.
- Workaround:** Internal low pass filter values can be adjusted for short cable testing. In order to adjust the filter properly for short cable testing, write a value of 0x0004 to PHY register 29'd and a value of 0x5A40 to PHY register 30'd. Please see documentation change number 7 for clarifications regarding PHY registers 29'd and 30'd.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controller.

19. System Hang Due to Host Block Requests

- Problem:** To improve performance, the host block maintains a queue of four requests that it will execute. If four requests are already in the queue, containing two back-to-back read requests, and another read request comes in (at a very specific time) as the fifth request, the host and the requesting agent become out of sync. The out of sync condition eventually leads to a system hang.
- Investigation of the root cause of the erratum is incomplete. The erratum appears to be located in the DMA unit or its interface to the bus controller.
- Implication:** The problem occurs during high traffic conditions. Using the receive delay timer will increase the likelihood that five requests are queued, increasing the likelihood of a system hang.
- Workaround:** Do not use the receive delay timer ring. Use the default value of '0' in the Receive Delay Timer Ring Register (RDTR) at address offset 0x02820. Other techniques can be used to moderate interrupts: not using descriptor writebacks, or querying the receive descriptor head pointer (approximate location of descriptors used.) Additionally, use the default value for PTHRESH and avoid non-default values of WTHRESH in RXDCTL (0x02828) unless specifically recommended by Intel.
- Intel changed the Linux driver to no longer use RDTR. Other Intel drivers did not previously use the feature.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controllers.

20. Transmit Descriptor Writeback Problems with Non-Zero WTHRESH

- Problem:** Transmit descriptors are typically written back to memory either upon transmit interrupts or opportunistically in between reading data buffers. When the controller has a programmed writeback threshold (TXDCTL.WTHRESH), it will attempt to write back full descriptors instead of just a status byte. The controller may incorrectly calculate the length of the writeback operation, causing corrupted descriptor writebacks. This erratum is closely related to Erratum #8. "Receive Buffer Writeback Problems for Packets Spanning Multiple Buffers."
- Implication:** Corrupted descriptor writebacks may include writing back unconsumed descriptors, descriptor writebacks to incorrect addresses, or writebacks missed altogether. In addition, the device may cease to access the PCI bus or cease packet transmission. If the device hangs, a full software or hardware reset is needed.
- Workaround:** Leave WTHRESH at its default value of 0. Descriptors will be written back immediately.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controller.

21. Bus Initialization with Some Chipsets

- Problem:** Upon initialization, the 82544EI/82544GC controller samples the REQ64# signal on the rising (inactive) edge of RST#. If REQ64# is sampled low (asserted), then the controller starts up with a 64-bit bus width.
- The PCI Local Bus Specification calls for 0 ns minimum input hold time on this signal. However, the 82544EI/82544GC controller requires 1 ns input hold time.
- Implication:** If the signal does not have sufficient hold time, the Ethernet controller could power up with incorrect bus width (64 versus 32 bits).
- Many bridges and chipsets drive the REQ64# signal with a full clock of hold time past the rising edge of RST# and this problem will not be encountered. Other loads on the PCI bus may affect the severity of the problem.
- Workaround:** For embedded designs, verify that the system bridge will deliver a full clock of hold time. If the problem is encountered on an add-in board, try moving the board to a connector on another bus segment.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controllers.

22. Packet Reception with APM Enabled before Driver Load

- Problem:** Upon power-up, the 82544EI/82544GC controller scans EEPROM initialization settings, enabling APM wakeup filtering if the APM Enable Bit is set in Initialization Control Word 2. When APM is enabled and a non-wakeup packet arrives before the software driver has enabled the 82544EI/GC controller's receiver, that packet will become "stuck" and the receiver will lock up. The arriving packet must match the controller's individual address or the broadcast filter. (Multicast packets are not accepted during this time.)
- After the lockup, no more packets (wakeup or non-wakeup) can be received. The condition can be removed by a hardware or software reset of the controller.
- The time window when the problem can occur is between initial power-on and the driver load by the operating system. After system boot and driver load, the controller is reset and the receiver is enabled. Following initial boot and driver load, software can place the controller in the D3 state. ACPI wakeup is not affected by this problem since ACPI can only be enabled through software.
- Implication:** If the EEPROM APM Enable Bit is set as the hardware default, the 82544EI/82544GC controller will be exposed to this problem. Wakeup packets arriving after the lockup condition will fail to properly wake up the device. The likelihood of non-wakeup (directed or broadcast) packets being sent to the 82544EI/82544GC controller will vary depending on the network.
- Workaround:** Program the APM Enable Bit in the EEPROM (Bit 2 of Initialization Control Word 2 at offset 0x0F) to 0. Contact your Intel representative for details on how to change settings in the EEPROM. The driver can still activate APM mode by setting APME in the WUC Register prior to going into a D3 state.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controllers.

23. Intermittent Power-on State Due to Decoded High-Impedance Test Modes

Problem: The 82544EI/82544GC contains internal logic which decodes various test mode(s) for component and/or system-level production test. The decoding of the various test modes is performed based on four component input pins listed in the table below. Three of the decoded pins are designated as NO_CONNECT in documentation because the pins provide internal pull-up resistors. However, two of the test mode operations place I/O pins in tri-state / high-impedance mode, which disable the internal pull-up resistors.

Implication: If on power-up, in configurations where the AUX_PWR indication is asserted to a logic "1", and the following decoded NO_CONNECT pins power on with logic input values shown in the table below, then the high-impedance tri-state mode will be decoded, and the internal pull-up resistors will be disabled resulting in a possible behavior that is intermittent or erratic.

AUX_PWR	NC_L4 (82544EI) / NC_J5 (82544GC)	NC_N2 (82544EI) / NC_L4 (82544GC)	NC_M2 (82544EI) / NC_K3 (82544GC)	Decoded Test Mode
1	0	1	0	Tri-state/ High-Z w/ pull-ups disabled
1	0	0	1	

During this high-impedance tri-state mode, if/when the three NO_CONNECT pins drift to values other than the encodings listed above, the pull-ups will be enabled, and the pins will be pulled up to their nominal operational values. The 82544EI/82544GC may be reset to a mostly operational state by subsequent PCI device resets. However, PCI I/O impedance settings may remain improperly configured until a subsequent LAN_PWR_GOOD reset.

Workaround: The affected NO_CONNECT pin is L4 on the 82544EI controller and pin J5 on the 82544GC controller. Attaching an external pull-up resistor to this pin ensures that the two potential high-impedance modes are not improperly decoded. Future revisions of product documentation will indicate this ball as TEST2 instead of NO_CONNECT.

Status: Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controllers.

24. 32-Bit Split-Completion Dependency on subsequent REQ64#

Problem: The 82544EI/82544GC contains an elasticity FIFO to convert requested data from the PCI/PCI-X bus width into 128-bit internal data widths. When operating as a PCI-X bus segment, an error may occur, if the last few bytes of a requested transmit-DMA operation are returned in a 32-bit split-completion. If the last 32-bit transfer cycle of the requested DMA transfer is followed immediately on the bus by a single-cycle bus turnaround and 64-bit external transaction request, the REQ64# signal from the subsequent transaction improperly affects the internal data path pipeline from the just-completed 32-bit split-completion.

The error may occur when the ending address of a transmit-data descriptor buffer terminates within even-dword alignment (the final data byte of the burst is located at a memory address ending in 0x0..0x3 or 0x8..0xB). When the error occurs, an internal data path pointer fails to be updated properly. The resulting pointer error may result in transmit data being lost, and/or incorrect residual data being delivered to subsequent data fetch operations.

Due to PCI-X bus protocols, transfers of 8 bytes or less which are performed as 32-bit split-completions are not susceptible to this problem, due to the existence of an extra clock cycle between the final data phase and the bus turn-around cycle.

Implication: If the last few bytes of a transmit DMA fetch are lost due to this erratum, the internal DMA client will stall; waiting for its remaining data, in which a transmit hang may be observed. If the subsequent data-fetch operation represents a descriptor-fetch operation, the erroneous residual data path content may corrupt the fetched descriptor. Resulting in potential data addressing errors, including unexpected Dual-Address Cycles (DAC), system memory protection errors, and/or system hangs.

Workaround: To workaround this erratum, each data buffer should be checked to determine whether the final byte of the buffer resides at a host memory offset of 0x0-0x3 or 0x8-0xB within a Quadword memory alignment. If so, terminate the descriptor buffer reference at an odd-word address alignment such as 0x7 or 0xF, and utilize a second data descriptor to reference the final 1-8 bytes of data.

Status: Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controllers.

25. Message Signaled Interrupt Feature May Corrupt Write Transactions

- Problem:** The problem is with the implementation of the Message Signaled Interrupt (MSI) feature in the Ethernet controller. During MSI writes, the controller should use the MSI message data value in PCI configuration space.
- At the same time, for normal write transactions (received packet data and/or descriptor writebacks), the controller temporarily stores the data for write transactions in a small memory until it is granted ownership of the PCI/PCI-X bus. The error condition occurs when during the MSI operation the controller incorrectly pulls data from the memory storing the data waiting to be written. If there are any write transactions waiting when this occurs, these transactions may become corrupted. This, in turn, may cause the network controller to lock up and become unresponsive.
- Implication:** If the affected products are used with an OS that utilizes Message Signal Interrupts and no accommodations are made to mitigate the use of these interrupts, data integrity issues may occur.
- Workaround:** For PCI systems, advertisement of the MSI capability can be turned off by setting the MSI Disable bit in the EEPROM (Init Control Word 2, bit 7).
- For PCI-X systems where MSI support is enumerated as part of the PCI-X specification, Intel is working with OS vendors to ensure that any future implementations of their operating systems can detect these products and avoid using the MSI mechanism. Further details will be communicated as they become available.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controllers.

26. Link Establishment or Communication Problems in Fiber Mode When Link Partner Does Not Fully Comply with the IEEE 802.3 Specification

- Problem:** The following minor compliance issues have been discovered between the TBI/SERDES mode symbol synchronization logic and the IEEE specification:
- When presented with short sequences of malformed code groups, the receive synchronization logic within the Ethernet controller may acquire & indicate link/synchronization prematurely or incorrectly
 - When presented with certain short sequences of malformed code groups, the logic may retain link/synchronization indication through the error sequence instead of immediately detecting and dropping link/synchronization
 - With some specific erroneous sequences of code groups, the auto-negotiation logic may establish link in certain very specific situations where the specification says it should not
 - Finally, the receive error detection logic may not detect and count some symbol errors when malformed idle patterns are received.
- Implication:** If a link partner is not compliant with the IEEE 802.3 Specification in certain very specific ways, the 8254x controller may not be able to establish link or communicate properly with it. If the controller is tested for strict compliance with the IEEE 802.3 Specification, it may fail some of the Clause 36 and Clause 37 test cases.
- However, Intel has performed extensive compatibility testing as an integral part of controller HW validation, and continues to do so with the latest Ethernet devices. To date, these issues have not been shown to cause interoperability problems with any Ethernet devices currently in production.
- Workaround:** None.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controllers.

27. Wakeup Packet Memory (WUPM) cleared upon reset

- Problem:** The 82544EI/GC specifications state that the Wakeup Packet Memory (WUPM) is not cleared on any reset. This is incorrect. Any reset or power-state transition will clear the contents of these registers.
- Implication:** Because a power-state transition takes place on wakeup, the Wakeup Packet Memory will always be cleared before it can be read by software. This makes the memory effectively unable to provide the capability for inspecting the wakeup packet content.
- Workaround:** There is no workaround. WUPM will be considered to be defeatured for the affected controllers.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controllers.

28. Unexpected RMCP ACK packets in ASF mode

- Problem:** According to the RMCP protocol, the response to all RMCP commands (except ACK) should be an RMCP ACK packet. In ASF mode, the Ethernet Controller responds to RMCP ACK packets with a second ACK.
- Implication:** Any management software should be aware of this behavior and not respond to the additional RMCP ACK packets.
- Workaround:** None.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controllers.

29. Exceeding PCI Power Management Specification Limit of 375mA current during reset and power state transitions

- Problem:** During resets and power state transitions the controller may briefly draw more than 375 mA of current as the digital signal processors in the PHY attempt to converge. The excessive current draw persists for approximately 100 milliseconds. Refer to the "Power Specifications -- MAC/PHY" section of this document for specific values.
- Implication:** If an application has current limiting circuitry in place, the Ethernet Controller may trigger these safeguards in power-up or during transitions between D0 and D3 power states.
- Workaround:** None.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controllers.

30. Memory Access must be enabled in order to Read Device Registers in I/O mode

- Problem:** The Ethernet controller will not respond to I/O transaction after a reset until its Memory Access Enable (MAE) bit has been set.
- Implication:** Attempts to access the Ethernet controller via I/O transactions without the MAE bit set will result in a master abort on the PCI bus.
- Workaround:** In order to access registers on the Ethernet Controller using I/O mode, both the I/O Access Enable and the Memory Access Enable bits in PCI configuration space must be set.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controllers.

31. Inbound and Outbound reads not fully decoupled in PCI-X mode

- Problem:** If the Ethernet controller receives a read as a target and signals a split response it will not deliver a completion to this read until its entire outstanding read requests have been satisfied. The device should not make the completion of a sequence for which it is the completer contingent upon another device completing a sequence for which it is a requester.
- Implication:** There is a slight system performance impact due to this erratum. Processors may be stalled while the read transaction is outstanding, so the extra delay may adversely affect CPU utilization.
- If and only if a host bridge also has a similar dependency, the possibility of a deadlock exists. A situation may arise where the bridge is waiting for the controller to respond to a DWord read while the controller is waiting for the bridge to complete a block read.
- Workaround:** None.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controllers.

32. Hang in PCI-X systems due to 2k Buffer Overrun during Transmit Operation

- Problem:** This Ethernet device has an error in the way that it stores data from PCI-X read transactions. If the controller is operating in PCI-X mode and its read data FIFO fills completely then the device can miscalculate the amount of free space in the FIFO and lose all of this data.
- This erratum does not apply to devices running in PCI mode only.
- Implication:** If this device enters this erratum state, the chip loses 2 kilobytes of data. The transmit and receive units of the chip will hang waiting for this data which will never arrive. No data will be corrupted. Once this has occurred, a reset is required to restore the device to normal operation.
- If using larger MTUs (jumbo frames), the chance of reaching this erratum state also increases.
- Workaround:** The issue can occur only when one packet is being completed and the next being started. Therefore, if the first fragment of every packet is limited in size the overflow can be prevented entirely. Drivers can work around this issue by ensuring that the size in the first descriptor of every packet less than 2016 bytes.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controller.

33. CRC Errors due to Rate Adaptation FIFO Overflow in Fiber Mode

- Problem:** In TBI mode and internal-SERDES mode this Ethernet device uses a small FIFO in its receive path to compensate for minute differences between the speed of the link partner's clock and the device's local clock. If the link partner has a faster clock, this FIFO will fill slowly during a packet, and then drain during inter-frame gaps.
- The device has an error in the way that this FIFO empties, causing it to wait several cycles into the inter-frame gap before it begins recovering clock drift.
- This only occurs during operation in fiber mode. Internal PHY mode used for copper applications is unaffected by this erratum.
- Implication:** If the Ethernet device is linked to a partner with a substantially faster clock and multiple frames arrive in sequence with minimal inter-frame spacing, then the device may not have time to recover all of the accumulated drift between frames. The synchronization FIFO will overflow and drop 4 bytes of the packet, which will be visible as a CRC error.
- The larger the difference between the link partner's clock and the Ethernet controller's clock, the fewer back-to-back frames need to be received to see CRC errors. In practice, this will be a very rare occurrence for two reasons. First, most Ethernet devices use clock frequencies near the center of the allowed range, so the difference between clocks will be small. Second, long strings of packets with minimal inter-frame spacing are rare on most networks.
- Workaround:** This erratum may be worked around by setting a larger inter-frame spacing. Specifically, switches must be configured to an inter-frame gap of at least 144 ns (18 symbols) for MTUs less than 10,000 bytes or at least 160 ns (20 symbols) for MTUs between 10,001 and 16,000 bytes.
- Alternatively, a new board design could use a reference clock source with a frequency near the high end of the 802.3 standard's allowed range. This would create a situation where the only way to trigger the erratum was for the link partner to have a faster clock which would violate the 802.3 standard.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controller.

34. Transmit Descriptors May Be Written Back to Host, Even Without the RS Bit Set

- Problem:** If the RS bit is set on at least some transmit descriptors submitted to the device, it is possible that some other transmit descriptors without the RS bit set will be incorrectly written back to host memory.
- Implication:** The unnecessary descriptor write-backs will not cause a functional issue, but they may result in a small amount of unnecessary host bus bandwidth to be consumed.
- Workaround:** None.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controller.

35. Polarity Detection Error May Cause Inability to Transmit

- Problem:** The PHY component inside the 82544EI/GC contains logic to detect the polarity of the cable being used. In this case, "polarity" refers to which wire of each pair is '+' and which is '-'. When the speed is forced to 10Mbps, and the 82544EI/GC attempts to establish link under heavy traffic, the PHY may incorrectly interpret end-of-packet symbols as inverted Normal Link Pulses. When this occurs, the PHY will establish link with the incorrect polarity.
- Implication:** When the PHY detects the incorrect polarity, the MAC will see numerous errors of all kinds. It will not be able to transmit nor receive properly.
- Workaround:** This situation can be avoided in two ways:
1. Disable the Polarity Reversal feature by setting bit 1 of the PHY Specific Control Register (16d) to 1b.
 2. Avoid receiving packets until polarity is detected.
- Drivers can implement option 2 by first, disabling the PHY's transmitter for 150ms, which will cause the link partner to drop link. Then force the PHY to 10Mbps and re-enable the PHY's transmitter.
- In terms of registers, the PHY's transmitter can be powered down by writing:
- ```
Write Register 0x1D = 0x0019
Write Register 0x1E = 0xFFFF
```
- It can be re-enabled by writing:
- ```
Write Register 0x1D = 0x0019
Write Register 0x1E = 0xFFFF0
Write Register 0x1E = 0xFF00
```
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82544EI/82544GC Gigabit Ethernet Controller.

SPECIFICATION CLARIFICATIONS

1. Receiver Enabling and Disabling

Problem: The 82544EI/82544GC controller does not support “throttled” reception by repeatedly disabling/enabling the receiver by programming the Enable (EN) Bit in the Receive Control Register (RCTL). The reason is that the disabling/enabling operation does not re-initialize packet filter logic that demarcates packet start and end locations in the FIFO.

A note will be added to the RCTL register description reminding users to reset at least the receiver before re-enabling it.

Affected Specs: RCTL register description in 82544EI/82544GC Gigabit Ethernet Controller Developer’s Manual, Rev. 1.0.

2. LEDs Inactive Until Driver Loads

Problem: LED indications (link and activity) are not active until the software driver loads even though the PHY has auto-negotiated and established link with a partner on the Ethernet. The LED signals work this way because they are derived from MAC logic and are qualified by the Set Link Up (SLU) bit in the Device Control Register (CTRL). Driver software sets this bit when it initializes. Alternatively, the SLU bit sets automatically when either Advanced Power Management or the SMBus are enabled through EEPROM settings.

Affected Specs: CTRL register description in 82546EB Gigabit Ethernet Controller Networking Silicon Developer’s Manual + Appendices for 82545EM, 82540EM, 82544EI/GC.

3. PHY Reset Duration

Problem: When resetting the PHY through the MAC control register (0x0000 bit 31), the PHY should be held in reset for a minimum of 10 ms before releasing the reset signal.

Affected Specs: 82544GC/EI Gigabit Ethernet Controller Preliminary Datasheet and Hardware Design Guide Rev. 0.5.

4. WUC.APME does not return to value in EEPROM after soft reset

Problem: When software initiates a reset by writing the CTRL.RST bit, the APME bit returns to its hardware default (0) rather than to any value that may have been loaded from the EEPROM upon initial power-up. If the driver does not take this behavior into account, it may inadvertently render the adapter unable to perform wakeup operations when it performs a reset.

There are two options to set up WUC.APME to the desired state. It either needs to re-program this bit to its desired post-reset value, or explicitly initiate a new EEPROM reload and then set only those bits whose values are desired to be different from the EEPROM-based HW defaults.

DOCUMENTATION CHANGES

1. M66EN Signal Must Be Connected

Problem: The pin description tables indicate the M66EN signal is ignored by the 82544EI/82544GC controller, but “should be connected properly for future compatibility.” Actually, the M66EN signal is used by the controllers to sense an operational 66 MHz PCI bus and **must** be connected properly in the system if 66 MHz PCI operation is desired. The text will be changed to show the corrected description.

Reference schematics do show the M66EN signal connected to the system.

Affected Docs: 82544EI Gigabit Ethernet Controller Hardware Design Guide Application Note (AP-422), Rev. 0.77, document number A44740-003 and 82544GC Gigabit Ethernet Controller Hardware Design Guide Application Note (AP-427), Rev. 0.77, document number A67149-001.

2. Octets Transmitted Counters Adjusted if VLAN Enabled

Problem: When the controller is configured to enable transmission of VLAN packets (CTRL.VME = 1), the device automatically adds 4 to the bytes counted in the Good Octets Transmitted (GOTCL and GOTCH) registers and Total Octets Transmitted (TOTL and TOTH) registers.

The register descriptions will change to describe this adjustment and point out that the count is adjusted even for non-VLAN packets transmitted (VLE = 0 in transmit descriptors).

Affected Docs: 82544EI/82544GC Gigabit Ethernet Controller Developer's Manual, Rev. 0.5.

3. Wake-up Packet Memory Not Writeable for Diagnostics

Problem: Changes made to the controller for the A2 stepping and beyond eliminated PCI-X write access to the WUPM register. The register is normally read-only.

Documentation will change to remove diagnostic write capability.

Affected Docs: 82544EI/82544GC Gigabit Ethernet Controller Developer's Manual, Rev. 0.5.

4. External GMII Mode Operation

Problem: The 82544EI/82544GC Gigabit Ethernet Controller contains a full IEEE-compliant Gigabit Media Independent Interface (GMII) connecting the MAC and PHY units. In normal operation, this interface is not exposed on the external pins. However, it is possible to configure the controller to communicate externally with other GMII/MII compliant devices for special applications. This interface has not been previously documented in the design guides.

Documentation will change to reflect external GMII capability. Most of the changes will occur in the Pin Description Table and the Ball to Signal Name Mapping Table. Notes will be added to the AC Specifications to cover the Note that in external GMII operation, the MDIO interface is **not** available.

82544EI Ball Number	82544GC Ball Number	Former Signal Name	Revised Signal Name	Added Function
A14	A10	LOS	LOS/LIND	For external GMII mode and external MII mode, LIND recognizes link indication from the PHY.
L2	K4	GMII_TEST0	GMII_MODE0	For external GMII mode and external MII mode, GMII_MODE0 connects to 3.3V through a pull-up resistor and GMII_MODE1 connects to ground through a pulldown resistor (1K or similar).
L3	K5	GMII_TEST1	GMII_MODE1	

82544EI/82544GC SPECIFICATION UPDATE

82544EI Ball Number	82544GC Ball Number	Former Signal Name	Revised Signal Name	Added Function
A4	D4	TX_DATA0	TX_DATA0	For external GMII mode, TX_DATA [7:0] is the parallel transmit data bus. TX_EN indicates transmission of data on the interface. TX_ER forces propagation of transmit errors and indicates carrier extension.
D5	D5	TX_DATA1	TX_DATA1	
C5	C4	TX_DATA2	TX_DATA2	
A5	E4	TX_DATA3	TX_DATA3	
D6	C5	TX_DATA4	TX_DATA4	
B6	E5	TX_DATA5	TX_DATA5	
A6	B5	TX_DATA6	TX_DATA6	
D7	E6	TX_DATA7	TX_DATA7	
D8	D7	TX_DATA8	TX_DATA8/TX_EN	For external MII mode, TX_DATA [3:0] is the transmit data bus. TX_ER is not used. TX_EN is Transmit Enable.
A7	C7	TX_DATA9	TX_DATA9/TX_ER	
C7	C6	GTX_CLK	GTX_CLK	For external GMII mode, GTX_CLK operates as a 125 MHz transmit clock. For external MII mode, the pin is undefined.
B8	E7	COL_TEST	COL/COL_TEST	For external GMII mode and external MII mode, COL is the collision detection input for half duplex operation.
C8	A6	CRS_TEST	CRS/CRS_TEST	For external GMII mode and external MII mode, CRS is the carrier sense input for half duplex operation.
A8	B7	RX_DATA0	RX_DATA0	For external GMII mode, RX_DATA [7:0] is the parallel transmit data bus. RX_DV indicates data is valid on the interface. RX_ER denotes a receive error.
D9	A7	RX_DATA1	RX_DATA1	
B9	C8	RX_DATA2	RX_DATA2	
A9	E8	RX_DATA3	RX_DATA3	
D10	E9	RX_DATA4	RX_DATA4	
A10	D9	RX_DATA5	RX_DATA5	
C11	C9	RX_DATA6	RX_DATA6	
B11	B9	RX_DATA7	RX_DATA7	
C12	D10	RX_DATA8	RX_DATA8/ RX_DV	For external MII mode, RX_DATA [3:0] is the receive data bus. RX_DV indicates data is valid on the interface. RX_ER denotes a receive error.
D12	A9	RX_DATA9	RX_DATA9/ RX_ER	

82544EI/82544GC SPECIFICATION UPDATE

82544EI Ball Number	82544GC Ball Number	Former Signal Name	Revised Signal Name	Added Function
B13	C11	RBC0	RBC0/RX_CLK	For external GMII mode, RX_CLK is the 125 MHz receive clock for gigabit operation. For external MII mode, RX_CLK is the receive clock for 10/100 Mb/s operation.
A12	B10	RBC1	RBC1/MTX_CLK	For external GMII mode, the pin is not used. For external MII mode, MTX_CLK is the transmit clock for 10/100 Mb/s operation.

Affected Docs: 82544EI Gigabit Ethernet Controller Hardware Design Guide Application Note (AP-422), Rev. 0.77, document number A44740-003 and 82544GC Gigabit Ethernet Controller Hardware Design Guide Application Note (AP-427), Rev. 0.77, document number A67149-001.

5. TIPG Register Value Incorrectly Documented

Problem: The IEEE standard minimum transmit inter-packet gap is 96 bit times. To achieve that gap requires a programmed setting to the IPGT Field in the Transmit Inter packet Gap (TIPG) Register. The actual inter-packet gap in MAC data clocks is the sum of the programmed value and a variable logic synchronization time within the device. Use a recommended programming value of 10 for TBI applications and 10 for 10/100/1000BASE-T applications to assure that the minimum IPG gap will be met under all synchronization conditions.

The Developer's Manual currently indicates a programming value of 8 for 10/100/1000BASE-T implementations and 6 for fiber implementations such as 1000BASE-SX. This value will change to 10. In addition to the register listing, section 12.5 Transmit Initialization text should change.

Affected Docs: 82544EI/82544GC Gigabit Ethernet Controller Developer's Manual, Rev. 1.0, reference number 11050.

6. Clarification on Settings for PHY Registers 29 and 30

Problem: The IEEE specifies 32 PHY registers, and 16 of them are IEEE Reserved. There is not enough space for normal functions and for test modes. Therefore, special settings in Register 29 are used to extend the register space for test modes.

Register 30 is used for test modes, and the different meaning is dependent on how Register 29 Bits 4:0 are programmed. To use the test modes, Register 29 should be programmed appropriately before programming Register 30.

The following tables contain test mode settings for Registers 29 and 30:

Register	Function	Setting	Mode	HW Rst	SW Rst	Notes
29.4:0	Test mode	11010 = gig receiver control	R/W	00000	Retain	
29.4:0	Test mode	11111 = G clks, Fiber, TBT, SD	R/W	00000	Retain	

Gig Receiver control (Test mode = 11010)

Register	Function	Setting	Mode	HW Rst	SW Rst	Notes
30_26.15	Override	1 = Enable 30_26.14:0 Override	W	0	0	note 30_26.a
30_26.14:12	Unused		W	-	-	note 30_26.a
30_26.11:8	enclk_gig_rcv[3:0]	1 = Enable gig receiver clock	W	-	-	note 30_26.a
30_26.7:4	pwrnd_gig_rcv[3:0]	1 = Power down gig receiver	W	-	-	note 30_26.a
30_26.3:0	reset_gig_rcv[3:0]	1 = Reset gig receiver	W	-	-	note 30_26.a
note 30_26.a	Register 29.4:0 must be set to 11010 to program this register Register 30_26 has the following meaning only if register 29.4:0 = 11010 Register 30_26 retains it value even if register 29.4:0 is != 11010 Cannot read since register embedded in power management state machine					

Gig Receiver control (Test mode = 11111)

Register	Function	Setting	Mode	HW Rst	SW Rst	Notes
30_31.15	Override power down	1 = Enable 30_31.11:4 Override	W	0	0	note 30_31.a
30_31.14	Override G clocks	1 = Enable 30_31.3:2 Override	W	0	0	note 30_31.a
30_31.13	Override Fiber	1 = Enable 30_31.1:0 Override	W	0	0	note 30_31.a
30_31.12	Unused		W	-	-	note 30_31.a
30_31.11:8	pwrnd_sd[3:0]	1 = Power down Signal Detect Circuit	W	-	-	note 30_31.a
30_31.7:4	pwrnd_tbt[3:0]	1 = Power down 10BASE-T	W	-	-	note 30_31.a
30_31.3	enclk_10	1 = Enable 10Mb Clocks	W	-	-	note 30_31.a
30_31.2	enclk_non10	1 = Enable non 10Mb Clocks	W	-	-	note 30_31.a
30_31.1	pwrnd_fib	1 = Power down entire fiber channel	W	-	-	note 30_31.a
30_31.0	pwrnd_ftx	1 = Power down fiber transmitter	W	-	-	note 30_31.a
note 30_26.a	Register 29.4:0 must be set to 11010 to program this register Register 30_31 has the following meaning only if register 29.4:0 = 11111 Register 30_31 retains it value even if register 29.4:0 is != 11111 Cannot read since register embedded in power management state machine					

Gig Receiver control (Test mode = 00100)

Register	Function	Setting	Mode	HW Rst	SW Rst	Notes
30_4.15	Reserved	Factory Defaults = 0	R/W	0	Retain	note 30_4.a
30_4.14	Reserved	Factory Defaults = 1	R/W	1	Retain	note 30_4.a
30_4.13	Reserved	Factory Defaults = 0	R/W	0	Retain	note 30_4.a
30_4.12	Reserved	Factory Defaults = 1	R/W	1	Retain	note 30_4.a
30_4.11	Reserved	Factory Defaults = 1	R/W	1	Retain	note 30_4.a
30_4.10	Reserved	Factory Defaults = 0	R/W	0	Retain	note 30_4.a
30_4.9	Reserved	1= Low, 0=High	R/W	0	Retain	note 30_4.a
30_4.8	Reserved	Factory Defaults = 0	R/W	0	Retain	note 30_4.a
30_4.7	Reserved	Factory Defaults = 0	R/W	0	Retain	note 30_4.a
30_4.6	Reserved	Factory Defaults = 1	R/W	1	Retain	note 30_4.a
30_4.5	Reserved	Factory Defaults = 0	R/W	0	Retain	note 30_4.a
30_4.4	Reserved	Factory Defaults = 0	R/W	0	Retain	note 30_4.a
30_4.3	Reserved	Factory Defaults = 0	R/W	0	Retain	note 30_4.a
30_4.2	Reserved	Factory Defaults = 0	R/W	0	Retain	note 30_4.a
30_4.1	Reserved	Factory Defaults = 0	R/W	0	Retain	note 30_4.a
30_4.0	Reserved	Factory Defaults = 0	R/W	0	Retain	note 30_4.a
note 30_26.a	Register 29.4:0 must be set to 00100 to program this register Register 30_4 has the following meaning only if register 29.4:0 = 00100 Register 30_4 retains it value even if register 29.4:0 is != 00100					

Affected Docs: 82544EI/82544GC Gigabit Ethernet Controller Developer's Manual, Rev. 0.5 (10416).

7. Remove Transmit Report Status Sent Function

Problem: The Transmit Report Status Sent function is not implemented to write back descriptor status when packet data goes out on the wire. The Report Packet Sent (RPS) Bit in the transmit descriptor (Bit 4 in TDESC.CMD) is Reserved and should be programmed to 0. The related Report Status function (Bit 3 in TDESC.CMD) may be

82544EI/82544GC SPECIFICATION UPDATE

used to force transmit descriptor status bytes to be written back to memory as the packet data reaches the transmit queue.

Affected text includes 4.3.2 Transmit Descriptor Writeback and references in numerous other sections, including the interrupt description text.

Affected Docs: 82544EI/82544GC Gigabit Ethernet Controller Developer's Manual, Rev. 1.0, reference number 11050.

8. Remove Transmit DMA Pre-fetching and Preemption Functions

Problem: The controller does not implement the ability to start transmit descriptor data fetches before finishing the previous descriptor. In addition, it does not have the ability to disable DMA preemptions during TCP segmentation. The Transfer DMA Control Register (TXDMAC) will be removed from the developer's manual and text in 12.7 Reset Operation should change.

Affected Docs: 82544EI/82544GC Gigabit Ethernet Controller Developer's Manual, Rev. 1.0, reference number 11050.

9. Remove Adaptive IFS Throttle Function (AIT)

Problem: The controller does not have the ability to increase inter-packet gap (beyond TIPG control) during back-to-back transmit operation and Gigabit half-duplex operation. The Adaptive IFS Throttle Register (AIT) will be removed from the developer's manual. Affected sections of the developer's manual include 9.3.2.2 Packet Bursting and 9.5.1 Adaptive IFS.

Note: The controller does not support Gigabit half-duplex mode operation.

Affected Docs: 82544EI/82544GC Gigabit Ethernet Controller Developer's Manual, Rev. 1.0, reference number 11050.