



82545GM Gigabit Ethernet Controller Specification Update

June 6, 2006

The 82545GM 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.

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.

82545GM 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 © 2003 – 2006, Intel Corporation.

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

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

CONTENTS

CONTENTS	3
PREFACE	5
NOMENCLATURE	5
COMPONENT IDENTIFICATION VIA PROGRAMMING INTERFACE	5
GENERAL INFORMATION	6
82545GM Component Marking Information.....	6
SUMMARY TABLE OF CHANGES	7
Codes Used in Summary Tables	7
SPECIFICATION CHANGES	9
ERRATA	10
1. XOFF from Link Partner can Pause Flow-Control (XON/XOFF) Transmission.....	10
2. Transmit Descriptor use of RS for non-data (Context & Null) Descriptors	10
3. LSO Premature Descriptor Write Back.....	10
4. Message Signaled Interrupt Feature May Corrupt Write Transactions	11
5. Link Establishment or Communication Problems in Fiber Mode When Link Partner Does Not Fully Comply with the IEEE 802.3 Specification.....	11
6. Wakeup Packet Memory (WUPM) cleared upon reset.....	12
7. Unexpected RMCP ACK packets in ASF mode	12
8. Exceeding PCI Power Management Specification Limit of 375mA current during reset and power state transitions	12
9. Inbound and Outbound reads not fully decoupled in PCI-X mode	12
10. Hang in PCI-X systems due to 2k Buffer Overrun during Transmit Operation	13
11. CRC Errors due to Rate Adaptation FIFO Overflow in Fiber Mode.....	13
12. Transmit Descriptors May Be Written Back to Host, Even Without the RS Bit Set	14
13. Legacy Transmit Descriptor Write-Back May Occur Before the Packet Data Associated with the Descriptor is Fetched.....	14
14. PCI-X Burst Write Transactions to Memory Mapped Registers at Non-Qword-Aligned Offsets Fail.....	14
SPECIFICATION CLARIFICATIONS	15
DOCUMENTATION CHANGES	16

REVISION HISTORY

82545GM Gigabit Ethernet Controller Specification Update

Date of Revision	Description
April 1, 2004	Initial public release. Added errata # 4 – 6, spec clarification # 5 – 6, and spec change # 3.
August 3, 2004	Added Errata # 7 – 9.
January 10, 2005	Added Errata #10 – 11 and spec clarification #6. Removed Spec Changes #1 and 2. Removed Spec Clarifications # 1 and 6.
June 30, 2005	Removed Specification Change #1. Removed Specification Clarifications #1 through #4.
October 6, 2005	Added lead-free product information. Added Errata #12 and #13.
June 6, 2006	Added Erratum #14.

PREFACE

This document is an update to published specifications. Specification documents for these products include:

- 82545GM Gigabit Ethernet Controller Datasheet, Intel Corporation
- 82545/82546 Gigabit Ethernet Controller Design Guide, Intel Corporation
- PCI/PCI-X Family of Gigabit Ethernet Controller Software Developer's Manual

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 82545GM 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.

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 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

82545GM controller steppings will be identified by the following register contents:

Stepping	Vendor ID	Device ID	Revision Number
82545GM A0	8086h	1026h	00h
82545GM A1	8086h	1026h	01h
-	-	-	-

These devices also provide identification data through the Test Access Port.

GENERAL INFORMATION

This section covers the 82545GM devices.

82545GM COMPONENT MARKING INFORMATION

Product	Stepping	QDF Number	Top Marking	Notes
82545GM	A0	Q522	RC82545GM	Engineering Samples
82545GM	A1	Q536	RC82545GM	Engineering Samples
82545GM	A1	-	RC82545GM	Production
82545GM	A1	-	PC82545GM	Lead-Free Production

Lead-Free Device



Leaded Device

Note: The black dot in the lower left-hand corner denotes pin 1.

SUMMARY TABLE OF CHANGES

The following table indicates the Specification Changes, Errata, Specification Clarifications or Documentation Changes, which apply to the listed 82545GM 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.

82545GM GIGABIT ETHERNET CONTROLLER SPECIFICATION UPDATE

No.	A0	A1	Plans	ERRATA	Page	Notes
1	X	X	NoFix	XOFF from Link Partner can Pause Flow-Control (XON/XOFF) Transmission	9	-
2	X	X	NoFix	Transmit Descriptor use of RS for non-data (Context & Null) Descriptors	9	-
3	X	X	NoFix	LSO Premature Descriptor Write Back	9	-
4	X	X	NoFix	Message Signaled Interrupt Feature May Corrupt Write Transactions	10	-
5	X	X	NoFix	Link Establishment or Communication Problems in Fiber Mode When Link Partner Does Not Fully Comply with the IEEE 802.3 Specification	10	-
6	X	X	NoFix	Wakeup Packet Memory (WUPM) cleared upon reset	11	-
7	X	X	NoFix	Unexpected RCMP ACK packets in ASF mode	11	-
8	X	X	NoFix	Exceeding PCI Power Management Specification Limit of 375mA current during reset and power state transitions	11	-
9	X	X	NoFix	Inbound and Outbound reads not fully decoupled in PCI-X mode	11	-
10	X	X	NoFix	Hang in PCI-X systems due to 2k buffer overrun during transmit operation	11	-
11	X	X	NoFix	CRC Errors due to Rate Adaptation FIFO Overflow in Fiber Mode	12	-
12	X	X	NoFix	Transmit Descriptors May Be Written Back to Host, Even Without the RS Bit Set	14	-
13	X	X	NoFix	Legacy Transmit Descriptor Write-Back May Occur Before the Packet Data Associated with the Descriptor is Fetched	14	-
14	X	X	NoFix	PCI-X Burst Write Transactions to Memory Mapped Registers at Non-Quadword-Aligned Offsets Fail	14	New

SPECIFICATION CHANGES

No specification changes reported at this time.

ERRATA

1. XOFF from Link Partner can Pause Flow-Control (XON/XOFF) Transmission

- Problem:** When the 82545GM transmitter is paused (by having received an XOFF from link partner), not only is the transmit of normal packets paused, but also of outbound XON/XOFF frames resulting from Receive Packet Buffer levels and Flow-Control Thresholds. Normally, partner's XOFF packets only pause the LAN controller for a finite time interval, after which outbound XON/XOFF's due to Receive Packet-Buffer fullness are again permitted to be sent.
- Implication:** If the transmitter is paused when a Receive FIFO XOFF threshold is reached, the transmission of XOFF frames does not occur and Receive FIFO overrun may potentially occur, resulting in lost packets. This is only expected to be seen with an abnormally high pause time from link partner's XOFF packet(s).
- Workaround:** Receive Flow-Control Thresholds may be tuned/lowered based on the expected maximum pause interval expected from link partner's XOFF packet in order to minimize the likelihood of Receive FIFO overruns.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82545GM Gigabit Ethernet Controller.

2. Transmit Descriptor use of RS for non-data (Context & Null) Descriptors

- Problem:** Due to an internal logic error in the descriptor internal queue, if the internal descriptor queue becomes completely full of pending descriptor status writebacks, the descriptor logic may issue a writeback request with an incorrect writeback amount. The internal descriptor queue may accumulate pending writebacks if transmit descriptors that do not directly refer to transmit data buffers (e.g. context or Null descriptors) are submitted with a status-writeback request (RS asserted) and legacy writeback (status byte writeback only) is utilized.
- Implication:** Due to the invalid internal writeback request size, the PCI logic may hang.
- Workaround:** Ensure that status-writeback reporting (RS) is not set on context or Null descriptors. Alternatively, utilize full-descriptor writebacks (TXDCTL.WTHRESH >= 1). The former workaround is the recommended alternative.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82545GM Gigabit Ethernet Controller.

3. LSO Premature Descriptor Write Back

- Problem:** For large send fetches ONLY (not normal or jumbo frames) the internal DMA engine will decompose the large-send data fetch into a series of individual requests that are completed sequentially. When all read data associated with the first internal DMA request has been fetched, the descriptor is flagged as ready for writeback. Though all data associated with the entire LSO descriptor will eventually be fetched, the descriptor writeback may occur prematurely. The device should wait until all bytes associated with the data descriptor have been completely fetched before writing back the transmit descriptor.
- Implication:** Due to premature write back, an operating system may release and reallocate the buffer, potentially causing buffer re-use and transmission of incorrect data.
- Workaround:** Utilize a second descriptor to point to the last four bytes of the large-send transmit data, and ensure that the buffer is not freed to the operating system/application until the second descriptor has been marked as complete via a status writeback operation.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82545GM Gigabit Ethernet Controller.

4. 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 82545GM Gigabit Ethernet Controller.

5. 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 82545GM 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 82545GM Gigabit Ethernet Controller.

6. Wakeup Packet Memory (WUPM) cleared upon reset

- Problem:** The 82545GM 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 82545GM Gigabit Ethernet Controller.

7. 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 82545GM Gigabit Ethernet Controller.

8. 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 82545GM Gigabit Ethernet Controller.

9. 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 82545GM Gigabit Ethernet Controller.

10. 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 82545GM Gigabit Ethernet Controller.

11. 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 82545GM Gigabit Ethernet Controller.

12. 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 82545GM Gigabit Ethernet Controller.

13. Legacy Transmit Descriptor Write-Back May Occur Before the Packet Data Associated with the Descriptor is Fetched

Problem:	If a legacy transmit operation directly follows a TCP Segmentation Offload transmit operation, the logic may incorrectly associate the successful completion of the TSO transmit with the next descriptor. If the next descriptor is a legacy descriptor, under certain timing scenarios it is possible for the legacy descriptor to be incorrectly written back to host memory with the DD bit set. This might occur even though the packet data for the legacy descriptor has not yet been fetched.
Implication:	Due to the premature write back, an operating system may release and reallocate the transmit buffer, potentially causing buffer re-use or transmission of incorrect data.
Workaround:	Utilize at least two descriptors for any legacy transmit operation. Do not reallocate any buffers associated with the transmit operation until the last descriptor has been written back.
Status:	Intel does not plan to resolve this erratum in a future stepping of the 82545GM Gigabit Ethernet Controller.

14. PCI-X Burst Write Transactions to Memory Mapped Registers at Non-Qword-Aligned Offsets Fail

Problem:	The device does not properly handle burst (greater than 4 bytes in length) write transactions to its memory mapped register space. Registers with addresses ending in 0x0 or 0x8 work, but registers with addresses ending in 0x4 or 0xC cannot be written. For example, a PCI-X Memory Write Block transaction writing 16 bytes to offset 0x2800 properly writes to locations 0x2800 and 0x2808; however, locations 0x2804 and 0x2808 are not updated.
Implication:	The specification for the device states that memory-mapped registers should only be written 32 bits at a time. Software should always be written to follow this rule. It is particularly important during initialization when most of the memory-mapped register accesses take place. Unfortunately, some platforms might perform write combining, which turn consecutive 32-bit writes into a single burst transaction. In this case, the registers with addresses ending in 0x4 or 0xC are not written. Common areas for consecutive adjacent register writes include setup of the large register arrays (MTA, VFTA, and RAR).
Workaround:	If there is platform-level control for a write-combining feature, turn it off. Alternately, software can be written with the possibility of write combining in mind: <ul style="list-style-type: none"> Writes to consecutive registers can be followed by a read transaction, which should flush the posted write from any bridge that might perform combining. Initialization of large register arrays (VFTA, MTA) can be performed in reverse, writing the highest location first and working backward to the lowest. This prevents write combining from occurring since the writes no longer meet the rules that allow combining.
Status:	Intel does not plan to resolve this erratum in a future stepping of the 82545GM Gigabit Ethernet Controller.

SPECIFICATION CLARIFICATIONS

No specification clarifications reported at this time.

DOCUMENTATION CHANGES

No documentation changes reported at this time.