

82546EB Gigabit Ethernet Controller Specification Update

September 26, 2006

The *82546EB* 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.

82546EB 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 © 2002-2006, 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	1
PREFACE	4
NOMENCLATURE	4
COMPONENT IDENTIFICATION VIA PROGRAMMING INTERFACE	4
GENERAL INFORMATION	5
82546EB Component Marking Information	5
SUMMARY TABLE OF CHANGES	6
Codes Used in Summary Tables	6
SPECIFICATION CHANGES	9
ERRATA	10
1. Collision/ASF/Manageability Data Re-Transmit Problem.....	10
2. DMA and ASF/Manageability Concurrency Problems	10
3. MDI/MDI-X Crossover Auto-Detection Not Working	11
4. ASF / Manageability Transmit and Receive Not Functional in D3 State without APM.....	11
5. Wake-Up Disable Control for TCO / Manageability Packets.....	11
6. ASF State Machines Out of Sync on Power Up	12
7. Short Packets under Heavy Transmit Load.....	12
8. 1.5V Regulator Control Circuit Start Up.....	12
9. PCI-X Bus Collisions with Some Chipsets.....	13
10. Retransmit Requests for 10Mb Half-Duplex Collisions	13
11. Excessive PCI Bus Hold Time	14
12. Bus Initialization with Some Chipsets.....	15
13. SMBALRT# Output Driven in ASF Mode.....	15
14. ASF Lockup upon Resetting MAC.....	15
15. REQ# Pin Requires Pull-Up Resistor in PCI-X Mode	16
16. INTA# / INTB# Configuration	16
17. MWI Transactions May Terminate on Non-Cacheline Boundary	16
18. Some LEDs Asserted in D3 state with Wakeup and Manageability Disabled.....	17
19. PCI-X Maximum Read Burst Size When Programmed Beyond Capability	17
20. Master-Aborts with Some Chipsets during Driver-Initiated Controller Reset	17
21. LSO Premature Descriptor Write Back.....	18
22. XOFF from Link Partner can Pause Flow-Control (XON/XOFF) Transmission	18
23. Transmit Descriptor use of RS for non-data (Context & Null) Descriptors.....	18
24. Address Error Crossing 64KB Boundary during PCI-X Packet Receives or Descriptor Writes	18
25. Intermittent Issues with TCO Receive Packets in IPMI Mode	19
26. Message Signaled Interrupt Feature May Corrupt Write Transactions.....	19
27. Link Establishment or Communication Problems in Fiber Mode When Link Partner Does Not Fully Comply with the IEEE 802.3 Specification	20
28. Wakeup Packet Memory (WUPM) cleared upon reset	20
29. Unexpected RMCP ACK packets in ASF mode	20

82546EB GIGABIT ETHERNET CONTROLLER SPECIFICATION UPDATE

30.	Exceeding PCI Power Management Specification Limit of 375mA current during reset and power state transitions	20
31.	Memory Access must be enabled in order to Read Device Registers in I/O mode	21
32.	Inbound and Outbound reads not fully decoupled in PCI-X mode	21
33.	Hang in PCI-X systems due to 2k Buffer Overrun during Transmit Operation.....	21
34.	CRC Errors due to Rate Adaptation FIFO Overflow in Fiber Mode.....	22
35.	Transmit Descriptors May Be Written Back to Host, Even Without the RS Bit Set.....	22
36.	Legacy Transmit Descriptor Write-Back May Occur Before the Packet Data Associated with the Descriptor is Fetched	22
37.	PCI-X Burst Write Transactions to Memory Mapped Registers at Non-Quadword-Aligned Offsets Fail	23
SPECIFICATION CLARIFICATIONS		24
1.	IPMI Packet Reception Requires SMBus-Compliant I2C Read Transactions	24
2.	WUC.APME does not return to value in EEPROM after soft reset.....	24
3.	82546EB Ports can be Disabled Individually	24
DOCUMENTATION CHANGES		25
1.	On-Board SERDES Feature	25
2.	Thermal Management.....	27

REVISION HISTORY

82546EB Gigabit Ethernet Controllers Specification Update

Date of Revision	Description
September 26, 2006	Removed references to the 82540EM GbE controller.
June 6, 2006	Added Erratum #37.
October 6, 2005	Added Errata #35 and #36. Added Specification Clarifications #3.
August 9, 2005	Added a reference figure for Errata #8. Updated workaround explanation in Errata #8.
June 25, 2005	Removed Spec Change #1 and #2. Removed Spec Clarifications #2 - #4. Removed Documentation Changes #1 - #6, #8, and #10 - #13.
January 10, 2005	Added Errata #33 – 34 and spec clarification #7. Removed Spec Change #1. Removed Spec Clarifications # 1 and 6.
August 3, 2004	Added Errata # 29 – 32.
April 1, 2004	Initial public release. Added errata # 25 - 28 and spec clarification # 5 - 6.

PREFACE

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

- 82546EB Gigabit Ethernet Controller Datasheet, Intel Corporation.
- 82546EB Gigabit Ethernet Controller Design Guide, Intel Corporation.
- 8254x Family of Gigabit Ethernet Controllers, 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 82546EB 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

82546EB controller steppings will be identified by the following register contents:

Stepping	Vendor ID	Device ID	Revision Number
82546EB A0	8086h	1010h	00h
82546EB A1	8086h	1010h	01h

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

GENERAL INFORMATION

This section covers the 82546EB devices.

82546EB COMPONENT MARKING INFORMATION

Product	Stepping	QDF Number	Top Marking	Notes
82546EB	A0	Q466	FW82546EB	Engineering Samples
82546EB	A1	Q467	FW82546EB	Engineering Samples
82546EB	A1	-	FW82546EB	Production



SUMMARY TABLE OF CHANGES

The following table indicates the Specification Changes, Errata, Specification Clarifications or Documentation Changes, which apply to the listed 82546EB 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.

82546EB GIGABIT ETHERNET CONTROLLER SPECIFICATION UPDATE

No.	A0	A1	Plans	ERRATA	Page	Notes
1	X	-	Fixed	Collision/ASF/Manageability Data Re-Transmit Problem	10	-
2	X	-	Fixed	DMA and ASF/Manageability Concurrency Problems	10	-
3	X	-	Fixed	MDI/MDIX Crossover Auto-Detection Not Working	11	-
4	X	-	Fixed	ASF/Manageability Transmit and Receive Not Functional in D3 State Without APM	11	-
5	X	-	Fixed	Wakeup Disable Control for TCO/Manageability Packets	11	-
6	X	-	Fixed	ASF State Machines Out of Sync on Power Up	12	-
7	X	-	Fixed	Short Packets Under Heavy Transmit Load	12	-
8	X	X	NoFix	1.5V Regulator Control Circuit Start Up	12	-
9	X	-	Fixed	PCI-X Bus Collisions with Some Chipsets	13	-
10	X	-	Fixed	Retransmit Requests for 10Mb Half-Duplex Collisions	13	-
11	X	-	Fixed	Excessive PCI Bus Hold Time	14	-
12	X	X	NoFix	Bus Initialization With Some Chipsets	15	-
13	X	-	Fixed	SMBALRT# Output Driven in ASF Mode	15	-
14	X	X	NoFix	ASF Lockup Upon Resetting MAC	15	-
15	X	X	NoFix	REQ# Pin Requires Pull-Up Resistor in PCI-X Mode	16	-
16	X	X	NoFix	INTA# / INTB# Configuration	16	-
17	X	X	NoFix	MWI Transactions May Terminate on Non-Cacheline Boundary	16	-
18	X	-	Fixed	Some LEDs Asserted in D3 State with Wakeup and Manageability Disabled	17	-
19	X	X	NoFix	PCI-X Maximum Read Burst When Programmed Beyond Capability	17	-
20	X	X	NoFix	Master-Aborts with Some Chipsets During Driver-Initiated Controller Reset	17	-
21	X	X	NoFix	LSO Premature Descriptor Write Back	18	-
22	X	X	NoFix	XOFF from Link Partner can Pause Flow-Control (XON/XOFF) Transmission	18	-
23	X	X	NoFix	Transmit Descriptor use of RS for non-data (Context & Null) Descriptors	18	-
24	X	X	NoFix	Address Error Crossing 64KB Boundary during PCI-X Packet Receives or Descriptor Writes	18	-
25	X	X	NoFix	Intermittent Issues with TCO Receive Packets in IPMI Mode	19	-
26	X	X	NoFix	Message Signaled Interrupt Feature May Corrupt Write Transactions	19	-
27	X	X	NoFix	Link Establishment or Communication Problems in Fiber Mode When Link Partner Does Not Fully Comply with the IEEE 802.3 Specification	20	-
28	X	X	NoFix	Wakeup Packet Memory (WUPM) cleared upon reset	20	-
29	X	X	NoFix	Unexpected RMCP ACK packets in ASF mode	20	-
30	X	X	NoFix	Exceeding PCI Power Management Specification Limit of 375mA current during reset and power state transitions	20	-
31	X	X	NoFix	Memory Access must be enabled in order to Read Device Registers in I/O mode	21	-
32	X	X	NoFix	Inbound and Outbound reads not fully decoupled in PCI-X mode	21	-

82546EB GIGABIT ETHERNET CONTROLLER SPECIFICATION UPDATE

No.	A0	A1	Plans	ERRATA	Page	Notes
33	X	X	NoFix	Hang in PCI-X systems due to 2k Buffer Overrun during Transmit Operation	21	-
34	X	X	NoFix	CRC Errors due to Rate Adaptation FIFO Overflow in Fiber Mode	22	-
35	X	X	NoFix	Transmit Descriptors May Be Written Back to Host, Even Without the RS Bit Set	22	-
36	X	X	NoFix	Legacy Transmit Descriptor Write-Back May Occur Before the Packet Data Associated with the Descriptor is Fetched	22	-
37	X	X	NoFix	PCI-X Burst Write Transactions to Memory Mapped Registers at Non-Qword-Aligned Offsets Fail	23	-

No.	A0	A1	Plans	SPECIFICATION CLARIFICATIONS	Page	Notes
1	X	X	Doc Change	IPMI Packet Reception from LAN Controller Requires SMBus-Compliant I2C Read Transactions	24	-
2	X	X	Doc Change	WUC.APME does not return to value in EEPROM after soft reset	24	-
3	X	X	NoFix	82546EB Ports Can be Disabled Individually	24	-
-	-	-	-	-	-	-
No.	A0	A1	Plans	DOCUMENTATION CHANGES	Page	Notes
1	X	X	Doc Change	On-Board SERDES Feature	25	-
2	X	X	Doc Change	Thermal Management	27	-

SPECIFICATION CHANGES

No specification changes reported at this time.

ERRATA

1. Collision/ASF/Manageability Data Re-Transmit Problem

Problem: Re-transmission attempts due to link up/down state transitions or due to collisions in half-duplex mode will have problems due to a logic error in the 82546EB manageability subsystem.

When the problem occurs, an arbiter selecting transmit packets from ASF/management data and regular transmit DMA sources will "lock on" to the data source for an extended sequence of transmit data. There are three problem scenarios:

- The controller is operating with ASF/management transmission and regular data transmission enabled concurrently. In this case, the behavior can appear as a transmit hang.
- The controller is operating with only ASF/management transmission enabled. In this case, the controller will successfully transmit all ASF/management packets. However, an ASF/management packet may occasionally be partially transmitted as a short fragment before being successfully re-transmitted in its entirety.
- The controller is operating in half duplex mode for regular data transmission and collisions occur. In this case, the behavior can appear as a delayed transmit hang after 64Kbytes of successful re-transmission.

Implication: If a lockup occurs, either a hardware or software reset will be required.

Workaround: Avoid operating in half duplex mode. In particular, run ASF/manageability functions with a stable, full-duplex link to avoid re-transmit events.

Implement a software timer to check for a stopped DMA transmit path. If a long interval elapses without expected packet transmission, reset the device.

Intel incorporated a reset timer and an option to disable ASF/management capability in early drivers to allow evaluation of the A0 stepping.

Status: Intel resolved this erratum in the A1 stepping of the 82546EB Gigabit Ethernet Controller.

2. DMA and ASF/Manageability Concurrency Problems

Problem: The 82546EB exhibits some problems when DMA and ASF/management operations occur simultaneously. The following situations are expected to be problematic:

- ASF/management packet becomes available for transmit while an ordinary DMA packet transmission is in progress.
- A hardware-generated XON or XOFF flow-control request occurs while an ASF/management packet transmission is in progress

The problem will result in transmitter lockup, manifested either in the regular DMA transmit path or in the ASF/TCO controller.

Implication: If lockup occurs, either a hardware or software reset will be required.

The rate of transmitter lockups will depend on the rate of event alignment and will increase under stress tests. Lockups due to XON/XOFF generation are expected to be less likely in routine traffic due to the infrequent nature of both XOFF and ASF/management transmits.

Workaround: Implement a software timer to check for a stopped transmitter on both the regular DMA and management paths. If a long interval elapses and transmit progress has not occurred, perform a software reset of the device. If regularly occurring manageability functions do not increment a statistics register, also perform a software reset.

Consider disabling ASF manageability operations while the operating system is running, restricting ASF functions to pre-boot and suspended states only. XON/XOFF flow control generation may be disabled through register programming.

To assist evaluation of early silicon, Intel incorporated a reset timer in the network driver code. Intel also provided an option to disable ASF manageability capability during driver load and re-enable management upon driver unload.

Status: Intel resolved this erratum in the A1 stepping of the 82546EB Gigabit Ethernet Controller.

3. MDI/MDI-X Crossover Auto-Detection Not Working

- Problem:** The 82546EB Controller feature for automatically detecting crossover cables fails to clearly detect the cable type. If a crossover cable is used when it is not required, the controller will not establish link.
- Implication:** Crossover cables to hubs and switches cannot be used unless the transmit/receive pairs are known to be crossed at the link partner. Using MII management registers to manually set the specific MDI or MDI-X configuration corresponding to the cable will allow link to be successfully established.
- Workaround:** Use straight-through cables to hubs and switches, unless the hub or switch port is crossed. Intel's driver code will disable auto MDI-X detection/correction and configure the 82546EB Controller for fixed MDI behavior. Alternatively, MII Management registers can be used to manually configure a particular forced MDI/MDI-X configuration. To do so, bits 6:5 of MII Register 16 should be set to 00 (for uncrossed MDI) or 01 (for crossed MDI-X). Follow this operation with a write to MII Register 0 with bit 15 set to 1 in order to reset the PHY.
- For the A0 stepping, Intel provided network drivers with registry-settable or command-line options to specify the correct manual MDI or MDI-X configuration.
- Status:** Intel resolved this erratum in the A1 stepping of the 82546EB Gigabit Ethernet Controller.

4. ASF / Manageability Transmit and Receive Not Functional in D3 State without APM

- Problem:** If the controller is in the Dr, D0u, or D3 power states, manageability functions do not work unless APM is enabled. Without power management wake-up enabled, part of an internal clock tree is becomes disabled.
- Implication:** Power management wake-up must be enabled if low power and pre-boot manageability operation is desired.
- Workaround:** Enable APM in the EEPROM image by setting bit 10 of the Management Control Word (word 24h) to 1. If APM wake-ups from D3 state are not desired, ignore or mask APM_WAKEUP, and disable APM wakeups from being promoted to PME# assertions by programming bit 15 of the EEPROM Initialization Control Word 2 (word 0Fh) to zero.
- Status:** Intel resolved this erratum in the A1 stepping of the 82546EB Gigabit Ethernet Controller.

5. Wake-Up Disable Control for TCO / Manageability Packets

- Problem:** The 82546EB Controller A0 stepping did not implement the RMCP and Ignore TCO wake-up disable bits planned for the Wakeup Filter Control Register (WUFC). These bits were intended to provide control over whether specific packet types that are received and detected as TCO/management packets per the Manageability Control Register (MANC) settings can initiate ACPI wakeup events.
- Implication:** If ACPI wakeups that overlap with TCO/Management packets are enabled, the TCO/management packets (directed RMCP and/or ARPs) will cause wakeup events from D3 state, even if these packets are expected to be received and serviced quietly by the TCO controller. It may not be possible to configure ACPI wakeups to occur on a specific desired subset of non-TCO packets.
- Workaround:** If ACPI wakeups due to TCO/management packets are unwanted during D3 state ASF/Manageability operation, disable all relevant bits associated with Individual Address (IA) match conditions (EX, BC, MC) as well as ARP and Directed IP match conditions (ARP, IPV4) in the Wakeup Filter Control Register (WUFC).
- Status:** Intel resolved this erratum in the A1 stepping of the 82546EB Gigabit Ethernet Controller. The Ignore TCO bit is implemented to allow all packets classified as TCO packets to be specifically excluded from wake-up generation. Intel has determined that the RMCP bit is an unnecessary control and did not implement it.

6. ASF State Machines Out of Sync on Power Up

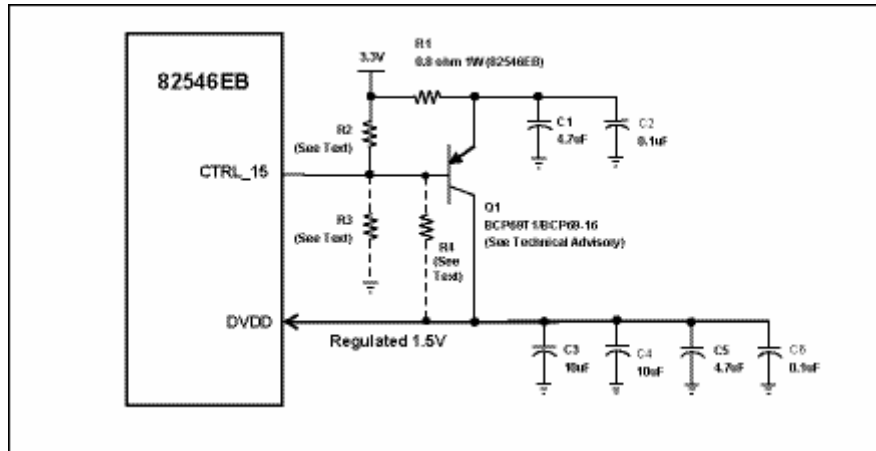
- Problem:** Upon power up, two state machines in the ASF logic may intermittently be out of sync, causing the ASF core to be locked up. The problem is due to an internal signal that does not remain valid long enough to assure sampling.
- Implication:** If the ASF state logic is deadlocked, no ASF/manageability operations can occur.
- Workaround:** Reset the ASF core prior to first use. For OS-present ASF operation, the reset code can be in the ASF driver. For pre-OS ASF operation, the reset code can be placed in the system BIOS.
- To perform the reset operation, issue an SMBus Byte Write with the following parameters:
- SMBus Address = Configured ASF address (typically C8h or 64h)
 - Command Code = F3h (selects the register)
 - Data value = 80h (selects the software reset bit)
- Status:** Intel resolved this erratum in the A1 stepping of the 82546EB Gigabit Ethernet Controller.

7. Short Packets under Heavy Transmit Load

- Problem:** Packets containing 48 or fewer bytes can lock up the transmitter, especially at 10 Mbps and 100 Mbps. The problem occurs because a very small packet can fit completely in a shallow buffer between the DMA unit and the MAC while the previous packet is being transmitted.
- Implication:** The transmitter lock up has only been observed with ARP request packets transmitted by the regular software driver. Note that later in the pipeline the controller pads all packets to the minimum IEEE length of 64 bytes. Transmission of short packets from the SMBus or ASF logic is not affected by the problem.
- Workaround:** Driver software will be modified to pad packets to at least 49 bytes.
- Status:** Intel resolved this erratum in the A1 stepping of the 82546EB Gigabit Ethernet Controller.

8. 1.5V Regulator Control Circuit Start Up

- Problem:** The 82546EB products contain on-chip LDO voltage regulation controllers for both 1.5V and 2.5V. The 1.5V regulator control output, CTRL_15, sometimes powers up in a high impedance state due to a digital portion of the feedback-control circuit which requires 1.5V to be present to operate correctly. The rate of occurrence of this problem may appear to vary from device to device or with different PNP transistors.
- Implication:** When drive is absent on the CTRL_15 output, only leakage current from the CTRL_15 output will be present. Depending on the amount of leakage current and the gain (Beta) of the external PNP transistor, there may be insufficient collector current ($I_C = I_B * \text{Beta}$) to adequately supply the 1.5V rail, and the 1.5V regulator circuit may not turn on. The 1.5V rail requires a minimum collector current of 5 mA to be generated and a 1.5V rail voltage of at least 600 mV to ensure that the digital portion of the circuit turns on correctly.
- When the 1.5V regulator control output fails, the 2.5V regulator control output(s) may also fail, driving the 2.5V supply toward either 0V or the 3.3V rail. The absence of one or both of the 1.5V and/or 2.5V supplies prevents the Ethernet controller from operating correctly.
- Workaround:** The regulator control output is Ball A18 on the 82546EB controller. Intel's original reference designs showed two external resistors, R1 and R2, in this circuit (refer to the following figure)
- The workaround is to completely remove R2 (also shown in the reference schematics as R6 @ 5.1K Ω) and place another resistor, R4, (also shown in the reference schematics as R7 @ 30.1K Ω 1%) to connect between the base of the PNP transistor and the 1.5V VDD output. Existing designs may alternatively use resistor R3 (not shown in the reference schematics) to connect between the CTRL_15 output and ground (refer to the following figure).
- Due to variations in PNP minimum/maximum transistor gain, this circuit must adhere to specific requirements for the PNP transistor, R3 resistor value, and configuration of the 82546EB controller. Follow **workaround instructions** in Technical Advisory TA-149, "Possibility of 1.5 Voltage Regulator Failures on 82546EB LOM Designs." **Note:** Workarounds differ among the 82540EM and 82546EB controllers and each controller has a separate technical advisory.
- Systems that generate 1.5V without using the Ethernet controller's onboard regulator control (at a central power supply, for example) do not need a workaround. No changes are required for the 2.5V regulator circuit.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82546EB Gigabit Ethernet Controller.



9. PCI-X Bus Collisions with Some Chipsets

Problem: On the PCI-X bus, when GNT is removed from a master device prior to the start of a transaction (indicated by FRAME assertion), the specifications allow the bus master one cycle to start the transaction. If FRAME is not asserted within the one clock window, the master should not assert FRAME without another REQ/GNT cycle.

In some cases, the 82546EB controller will assert FRAME and start a bus cycle two clocks after GNT is removed.

Implication: If the Ethernet controller starts a late PCI-X bus cycle, it may conflict with another PCI-X bus cycle mastered by the PCI-X bridge. If there are conflicts, the system will communicate erratically with PCI-X devices on that bus. System lockups are likely, requiring rebooting.

The likelihood of PCI-X bus collisions varies by chipset, since some chipsets do not have the ability to drive new PCI-X bus cycles as soon as two clocks from GNT deassertion. At this time, collisions have been observed only with the Intel P64H2 device.

Workaround: If the system uses an affected chipset, operate the bus in PCI mode instead of PCI-X.

Status: Intel resolved this erratum in the A1 stepping of the 82546EB Gigabit Ethernet Controller.

10. Retransmit Requests for 10Mb Half-Duplex Collisions

Problem: In 10 Mb operation, data transfer handshake synchronization from the DMA clock domain into the transmit MAC clock domain can take more time than it takes for the DMA clock domain to observe a collision event and re-try the transmit request. If a collision occurs at the exact moment that a data transfer from the DMA to MAC occurs, the discrepancy in synchronization delays can cause the logic to become confused about which data has been transferred.

Implication: When the device confuses the order of the data events, the transmitter will lock up. The long synchronization delays inherent in 10 Mb operation can readily cause this situation.

Workaround: Operate the 82546EB controller in full duplex mode. If half duplex operation is required, run the device at the 100 Mbps speed.

Status: Intel resolved this erratum in the A1 stepping of the 82546EB Gigabit Ethernet Controller.

11. Excessive PCI Bus Hold Time

Problem: The PCI Local Bus Specification calls for 0 ns input hold time on the address/data bus (AD [63:0]) and the byte enables (CBE [7:0] #). However, the 82546EB requires approximately 4 ns input hold time for 33MHz operation and approximately 1.5 ns input hold time for 66 MHz operation. This problem is due to incorrect hook-up of the clocks and the affected I/O cells.

Implication: The controller may read data incorrectly, causing parity errors, erratic Ethernet port operation, and possible system lockups. This problem does not affect operation on a PCI-X bus.

The severity of the problem depends on several factors, varying from system to system. In some designs the bridge connected to the 82546EB controller will provide adequate hold time and problems will not be seen. If the 82546EB device is on an add-in board, the problem may vary from slot to slot.

Workaround: Operate the 82546EB controller with 66 MHz or faster timings, since most bridge devices provide enough margin to the PCI hold time to meet the 1.5 ns requirement. There are three ways to force the controller to use 66 MHz or faster timings:

- Place the controller on a 66 MHz PCI bus, with the M66EN (66 MHz Enable) signal connected to the system in the usual way. If the device is on an add-in card, place the card on a 66 MHz slot.
- Leave the controller on a 33 MHz PCI bus, but break the connection between the controller's M66EN signal (Ball C2) and the system. Use a pull-up resistor to connect the controller's M66EN signal to the 3.3V supply. The value of the pull-up resistor is not critical and can be as high as 100K ohms.
- Operate the controller on a PCI-X bus running at 66 MHz or faster. This erratum does not affect PCI-X operation.

Status: Intel resolved this erratum in the A1 stepping of the 82546EB Gigabit Ethernet Controller.

12. Bus Initialization with Some Chipsets

- Problem:** Upon initialization, the 82546EB controller samples the REQ64#, DEVSEL#, STOP# and TRDY# signals on the rising (inactive) edge of RST#. If REQ64# is sampled low (asserted), then controller starts up with a 64-bit bus width. The values sampled on the other pins denote the PCI-X initialization pattern.
- The PCI-X Addendum to the PCI Local Bus Specification calls for 0 ns minimum input hold time on these signals. However, the 82546EB controller requires 1 ns input hold time.
- Implication:** If the signals do not have sufficient hold time, the 82546EB controller could power up with incorrect bus width (64 versus 32 bits) or mode (PCI/PCI-X 66/PCI-X 100/PCI-X 133). However, Intel has not observed this problem in an actual system based on the 82546EB controller.
- The problem cannot occur in systems based on the P64H2 PCI-X bridge because it drives the signals with a full clock of hold time past the rising edge of RST#. Intel expects other contemporary bridges and chipsets to perform the same way. The problem may be encountered in some systems based on legacy bridges and chipsets.
- Other loads on the PCI/PCI-X bus may affect the severity of the problem.
- Workaround:** For LAN on motherboard 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 82546EB Gigabit Ethernet Controller.

13. SMBALRT# Output Driven in ASF Mode

- Problem:** The SMBALRT# output signal has an additional function as a PCI_POWER_GOOD input in ASF mode operation. However, the controller does not disable the SMBALRT# open drain output and contention between the external PCI_POWER_GOOD signal and the internal SMBALRT# signal can result.
- Implication:** The Ethernet controller drives SMBALRT# low during communication between the ASF controller and the main LAN controller logic. Depending on the drive strength of the external PCI_POWER_GOOD signal (or the value of a pull-up resistor attached to the ball), the ASF controller may not observe PCI_POWER_GOOD asserted.
- If PCI_POWER_GOOD cannot be observed following "OS Hang" event detection, the ASF controller may fail to observe reset on the PCI bus, continuing to perform remote control operations such as EEPROM reloads.
- Workaround:** Drive the PCI_POWER_GOOD input from a low impedance source (25 ohms or less).
- Status:** Intel resolved this erratum in the A1 stepping of the 82546EB Gigabit Ethernet Controller.

14. ASF Lockup upon Resetting MAC

- Problem:** The ASF logic has a handshake problem at the expiration of its watchdog timer. When the ASF block attempts to reset the MAC, an ASF lockup can occur.
- Implication:** The watchdog reset manageability function will not be available.
- Workaround:** Intel has workaround ASF software and modified EEPROM settings to accompany the workaround. The workaround disables MAC resets by the ASF logic (EEPROM Word 23h, Bit 2) and enables the ASF logic to request ARP packets following the watchdog event.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82546EB Gigabit Ethernet Controller.

15. REQ# Pin Requires Pull-Up Resistor in PCI-X Mode

- Problem:** In PCI-X operation, the REQ# (Request) output is configured as an open drain output due to a hookup error. Therefore, the output will float when it is not driven to the active (low) state. With REQ# floating to an indeterminate state, the system bridge may become confused and behave erratically on the PCI-X bus.
- Implication:** Since REQ# is a dedicated point-to-point signal, the system may expect it to be driven at all times. If a pull-up resistor is present on REQ#, the resistor may not be strong enough to avoid slow rise times. Intel observed PCI-X master-abort cycles and system hangs due to the erratum.
- In conventional PCI bus operation, the REQ# signal is configured to be correctly driven during its deasserted state. This erratum does not apply.
- Workaround:** Verify that a relatively strong pull-up resistor is present on the REQ# line. Change or supplement the pull-up resistor as needed, aiming for an optimum resistance of approximately 2K ohms. Resistors as large as 10K ohms may be strong enough, however. If the design has an existing pull-up resistor up to 10K ohms that is impractical to change, test carefully to make certain that there is no evidence of failure. The reference schematic attached to this update will change to show a 2K pull-up resistor.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82546EB Gigabit Ethernet Controller.

16. INTA# / INTB# Configuration

- Problem:** The assignment of interrupt signals INTA# and INTB# is not programmable. INTA# is assigned to LAN port A and INTB# is assigned to LAN port B. If either one of the ports is disabled, the remaining port retains its original interrupt signal. However, the PCI configuration space always advertises INTA# in use for a single port configuration.
- Implication:** The operating system may become confused in situations where port A is disabled because the Ethernet controller will continue to use INTB# for port B.
- Workaround:** When the LAN A device is disabled, ensure that BIOS maps LAN B to the interrupt line driven by the INTB# pin, regardless of what pin the LAN B PCI header reports that it is using. Alternatively, use the 82546EB dual port Ethernet controller in such a way that if it becomes necessary to disable one port, the disabled port is always port B.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82546EB Gigabit Ethernet Controller.

17. MWI Transactions May Terminate on Non-Cacheline Boundary

- Problem:** When PCI MWI (Memory Write Invalidate) transactions are enabled, the Ethernet controller will use MWI opcodes whenever possible for highest bus efficiency. For receive packets that do not neatly align to cacheline boundaries, the 82546EB device may continue an MWI burst to end-of-packet, then disconnect at the end-of-packet, resulting in the transfer of a partially filled cacheline in spite of the MWI operation.
- Implication:** Some chipsets may have a problem dealing with prematurely terminated MWI transactions. One legacy PCI chipset was observed to clear the entire cacheline in memory, leading to packet corruption. Contemporary chipsets were observed to handle the transaction as follows:
- Treat it like an ordinary Memory Write (MW) transaction, without data errors, or
 - Pad the remainder of the cacheline with random data. Data errors do not occur because the entire packet is written correctly to the receive buffer memory and the MWI burst remains within the allocated receive buffer region.
- Workaround:** For systems with chipsets that do not tolerate partially filled cachelines, disable the MWI opcode. The 82546EB will then use the MW opcode instead.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82546EB Gigabit Ethernet Controller.

18. Some LEDs Asserted in D3 state with Wakeup and Manageability Disabled

- Problem:** When the controller is in a D3 state (system suspended) with manageability and wake-ups disabled, the clock circuit used for some of the LED sources is stopped. As a result, some of the LED's can potentially exhibit unusual behavior in this configuration. The LED's may remain asserted in this D3 state, despite a change in the state/status of the LED source. The following LED indications are affected: flow-control, full-duplex, collision, activity, link, and activity/link.
- Implication:** Configurations other than Wake on LAN configurations will encounter this behavior.
- Workaround:** Enabling Wake on LAN or SMBus manageability ensures that the clock circuits are not stopped in D3 state and the LED indications remain valid.
- Status:** Intel resolved this erratum in the A1 stepping of the 82546EB Gigabit Ethernet Controller.

19. PCI-X Maximum Read Burst Size When Programmed Beyond Capability

- Problem:** The 82546EB Gigabit Ethernet Controller contains a 2Kb FIFO to receive split-completion read data from PCI-X mode busses. The device advertises this maximum memory read size capability via the Designed Maximum Memory Read Burst Count (DMMRBC) Register, advertising a maximum single read capability of 2Kbytes. The system-configured maximum memory read size is set via the Maximum Memory Read Burst Count (MMRBC) Register. According to PCI-X specifications, the device should use the lower of the two values in determining the size of PCI-X memory read commands issued. However, the 82546EM uses only the value of the programmed MMRBC command field to limit the size of PCI-X memory read commands generated.
- Implication:** When the 82546EB controller is configured for a maximum read size of 4kbytes (MMRBC set to 4kb by BIOS), the controller may issue reads of up to 4Kbytes, which exceed its capability to store the split-completion data (limited to 2Kbytes). As a result, transmit data corruption may occur.
- Workaround:** During device initialization, determine if the configured maximum read size (MMRBC) exceeds the controller's advertised capability (DMMRBC). If so, re-configure the PCI-X command register for the specific controller to not exceed its maximum capability. Alternatively, all transmit descriptors may be restricted to point to data buffers which do not exceed the controller's advertised maximum read size.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82546EB Gigabit Ethernet Controller.

20. Master-Aborts with Some Chipsets during Driver-Initiated Controller Reset

- Problem:** The 82546EB Gigabit Ethernet Controller implements a software-initiated device reset function through the control register space (CTRL Register Bit 26). This software-initiated reset re-initializes all functional state of the controller except for PCI/PCI-X configuration. When the reset is written, the controller requires a few internal clock cycles to complete the reset operation. During this brief time, it will not respond to additional register accesses.
- Some PCI/PCI-X bridge components implemented with 64-bit architectures may initiate an additional zero byte write immediately following the 32-bit write to the CTRL register. This zero byte write is essentially padding for a 64-bit transaction. In cases where the CTRL register access performs a software reset of the controller, the zero byte operation may encounter a master-abort due to the controller reset in progress.
- Implication:** In most system configurations, a master-abort on an outbound 0-byte write operation will not result in any adverse system behavior, though the event may be logged at either the chipset bridge or operating system level.
- If the bridge/chipset is configured to promote the master-abort to a Non-Maskable Interrupt (NMI) **and** the operating system cannot discern and handle NMI events, then a fatal operating system error may occur.
- Workaround:** The 82546EB can only be accessed using DWord (32-bit) software operations. Since a potential 0-byte write immediately following a device CTRL register write is a hardware event "created" by the chipset bridge components, no software mechanism can be used to eliminate the 0-byte write operation.
- When an operating system is incapable of handling the NMI event, the chipset bridge should be configured to avoid promoting this master-abort to a fatal NMI event. Alternatively, the software-initiated device reset may be performed using an I/O access in place of a memory-mapped device access. The latter solution is effective if master-aborts on I/O writes do not result in NMI.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82546EB Gigabit Ethernet Controller.

21. 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 82546EB Gigabit Ethernet Controller.

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

- Problem:** When the 82546EB 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 82546EB Gigabit Ethernet Controller.

23. 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 82546EB Gigabit Ethernet Controller.

24. Address Error Crossing 64KB Boundary during PCI-X Packet Receives or Descriptor Writes

- Problem:** Due to an error in the logic controlling the address counters in the PCIX core, if the intended end of a PCIX block-write transaction (inbound Packet Receive or Descriptor Writeback) crosses a 64KB boundary *and* there is an ADB disconnect that happens to occur at the 64KB boundary, then the address resumed after the ADB disconnect points to the incorrect, previous 64KB page.
- This error applies only to PCIX operation; disconnects at/near 64Kb address boundaries in PCI operation are unaffected.
- Implication:** After resuming from an ADB disconnect at the 64Kbyte boundary, the transaction will resume at an incorrect address. As a result, packet or descriptor data may be written to incorrect locations, resulting in erroneous packet data being received, packet data being seen in descriptor ring contents, or device hangs.
- Workaround:** Ensure that descriptor ring addresses and Receive buffer addresses avoid crossing 64Kbyte alignment boundaries.
- Status:** Intel does not plan to resolve this erratum in a future stepping of the 82546EB Gigabit Ethernet Controller.

25. Intermittent Issues with TCO Receive Packets in IPMI Mode

Problem: The TCO Receive FIFO is implemented as a pair of buffers utilized in an alternating fashion. Under a specific timing condition where a new packet arrives from the Ethernet for the TCO receive FIFO coincident with a previous TCO packet being read from the FIFO to the SMBus, logic tracking the occupied/empty state of the buffers can enter an inconsistent state. This problem does not occur if only a single TCO packet is passed through the TCO Receive FIFO at a time.

Implication: When the LAN controller is in this erratum state, symptoms may include (a) corrupted packets delivered to the SMBus, (b) packets received to the SMBus twice, or (c) a received TCO packet appearing to be "stuck" in the TCO Receive FIFO until a new TCO packet arrives. Most network operations are only mildly affected, as most network protocols allow for lost or late packets and support header/payload integrity checksums.

Workaround(s): To address this problem, a series of steps should be taken in BMC firmware:

- First, the BMC firmware should check integrity of all TCO packets received by checking the IP/UDP checksums and discarding any corrupted packets.
- Second, the BMC firmware should implement a "check for erratum state" function using the following conditions:
 - a) Check to see if two sequential packets received are exact duplicates.
 - b) Attempt to check for "stuck" packets to determine whether a packet received has simply been delayed on the network versus "stuck" in the TCO Receive FIFO. The IPMI specification defines an 8 sequence-number window; received TCO packets exceeding that window may be good indications of being "stuck". Depending upon application, BMC firmware may be able to implement additional mechanisms to detect when a TCO packet received from the LAN controller appears to be one that had been given up as "lost" on the network.
- Finally, upon detecting a likely "TCO erratum state", BMC firmware should implement a specific "TCO Abort" operation on the SMBus to return the LAN controller from an erroneous state back to a normal, operational state without requiring a LAN controller reset. A "TCO Abort" operation is an intentionally abnormally terminated SMBus transaction. The specific TCO Abort transaction recommended by Intel consists of an N-byte SMBus write (N>1) to the LAN controller where the BMC only provides 1 byte of data before initiating a STOP.

Note that issuing this "TCO Abort" operation while the LAN controller is in a normal, healthy operational state can in fact induce the erratum condition. Therefore, checks for the erratum condition should be considered carefully so as to avoid excessive TCO Aborts while the LAN device is in a healthy state. However, if the TCO Abort were to be errantly issued while the LAN controller is in a normal healthy state, and the erratum state were to be induced, it is expected that the erratum-check criteria would again detect & correct the state back to normal, healthy state.

Status: Intel does not plan to resolve this erratum in a future stepping of the 82546EB Gigabit Ethernet Controller.

26. 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 82546EB Gigabit Ethernet Controller.

27. 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 82546EB 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 82546EB Gigabit Ethernet Controller.

28. Wakeup Packet Memory (WUPM) cleared upon reset

- Problem:** The 82546EB 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 82546EB Gigabit Ethernet Controller.

29. 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 82546EB Gigabit Ethernet Controller.

30. 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 82546EB Gigabit Ethernet Controller.

31. 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 resolved this erratum in the 82546GB Gigabit Ethernet Controller.

32. 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 82546EB Gigabit Ethernet Controller.

33. 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 82546EB Gigabit Ethernet Controller.

34. 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 82546EB Gigabit Ethernet Controller.

35. 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 82546EB Gigabit Ethernet Controller.

36. 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 82546EB Gigabit Ethernet Controller.

37. PCI-X Burst Write Transactions to Memory Mapped Registers at Non-Word-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, VFSA, 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:
- 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 (VFSA, 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 82546EB Gigabit Ethernet Controller.

SPECIFICATION CLARIFICATIONS

1. IPMI Packet Reception Requires SMBus-Compliant I2C Read Transactions

Clarification: The 82546EB controller SMBus interface provides the ability to receive packets from the LAN Controller to an external IPMI Controller. The 8254x specifies that SMBus read transactions be used to receive packets, packet status, and (if enabled) LAN Controller status-change alerts. Though the SMBus specification is based upon the I2C specification, SMBus read transactions further require that an IPMI device, when reading from target, read all bytes indicated as available by the target. I2C-based IPMI controllers which do not dynamically interpret the byte-count returned from the target device may violate this specification, resulting in improper behavior of the LAN Controller after reading packet-status fragments. Application Note 430 describes a mechanism for a non-SMBus-compliant I2C controller to properly receive packets from the LAN controller.

Affected Specs: 82546EB Gigabit Ethernet Controller Preliminary Datasheet and Hardware Design Guide Rev. 0.5.

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

Clarification: The 82546EB controller SMBus interface provides the ability to receive packets from the LAN Controller to an external IPMI Controller. The 8254x specifies that SMBus read transactions be used to receive packets, packet status, and (if enabled) LAN Controller status-change alerts. Though the SMBus specification is based upon the I2C specification, SMBus read transactions further require that an IPMI device, when reading from target, read all bytes indicated as available by the target. I2C-based IPMI controllers which do not dynamically interpret the byte-count returned from the target device may violate this specification, resulting in improper behavior of the LAN Controller after reading packet-status fragments. Application Note 430 describes a mechanism for a non-SMBus-compliant I2C controller to properly receive packets from the LAN controller.

Affected Specs: 82546EB Gigabit Ethernet Controller Preliminary Datasheet and Hardware Design Guide Rev. 0.5.

3. 82546EB Ports can be Disabled Individually

Clarification: The 82546EB allows ports to be disabled individually. However, if only one port is disabled (using the LAN disable feature) while the other port is left operational, care must be taken to ensure that the SMBus feature is NOT enabled on the disabled port. If the SMBus is left enabled on a disabled port the 82546EB may prevent the system from booting.

DOCUMENTATION CHANGES

1. On-Board SERDES Feature

Problem: Intel added new SERDES functions to the 82546EB. SERDES stands for Serializer-Deserializer and is a circuit that converts Ethernet packet data from a parallel TBI format to a 1.25 GHz bidirectional serial stream. For 1000BASE-SX or 1000BASE-LX fiber applications, SERDES integration typically saves the designer the expense and board space of a 64-lead external device. The 82546EB Controller contains two integrated SERDES units and the 82545EM Controller contains one SERDES. A '10' needs to be written in ICW3 bits 1 and 0, respectively, in order to enable this feature.

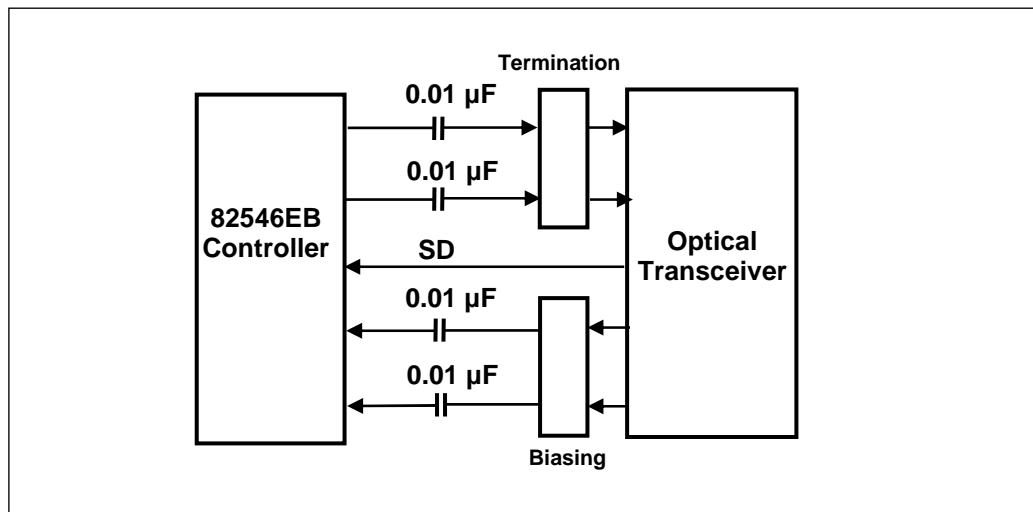
The datasheet for the 82546EB Controller will change to define pins as follows:

Signal Name	Type	Name and Function
RX_A +/- RX_B +/-	I	SERDES Receive Pairs A and B. Differential receive signal pairs for 1.25 GHz serial interface. For serializer-deserializer operation, couple the inputs to ECL voltage levels. If the SERDES interface is not used, leave RX +/- pairs unconnected.
TX_A +/- TX_B +/-	O	SERDES Transmit Pairs A and B. Differential transmit signal pairs for 1.25 GHz serial interface. For serializer-deserializer operation, the outputs drive LVPECL voltage levels. If the SERDES interface is not used, leave TX +/- pairs unconnected.
SIG_DETECT_A SIG_DETECT_B	I	Signal Detects A and B. Indicates that signals have been detected by optical transceivers connected to the 1.25 GHz serial interface. If the SERDES interface is not used, connect unused SIG_DETECT inputs to ground through pull-down resistors.

82546EB GIGABIT ETHERNET CONTROLLER SPECIFICATION UPDATE

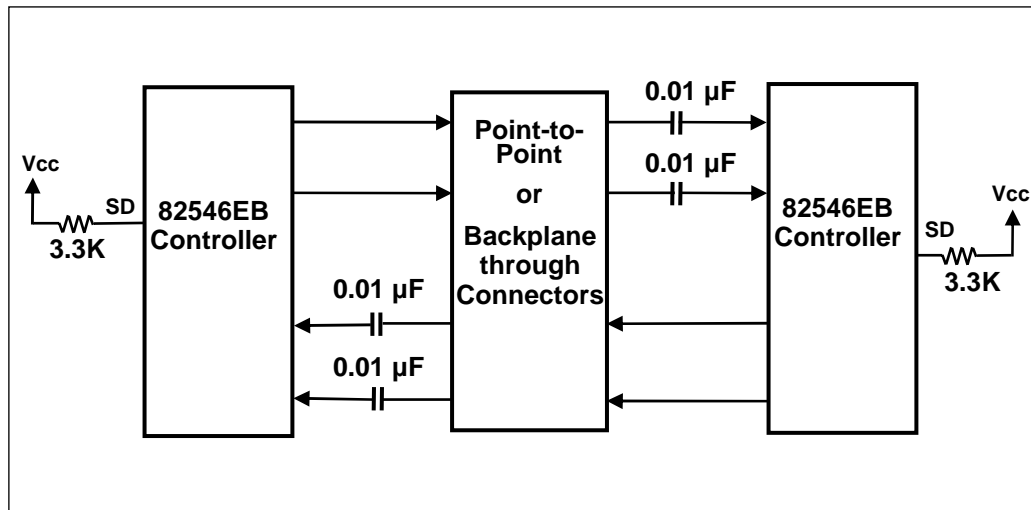
The SERDES ball locations for the 82546EB device are as follows:

RXA +	G19
RXA -	G20
RXB +	J19
RXB -	J20
TXA +	F19
TXA -	F20
TXB +	K19
TXB -	K20
SIG_DETECT A	E20
SIG_DETECT B	L20



82546EB fiber-based applications will use the SERDES interface to drive a 1000BASE-SX or 1000BASE-LX optical transceiver as illustrated in the following figure. The 82546EB controller is internally terminated, so no external termination is required.

Alternatively, the 82546EB controller may be used in short haul gigabit backplane data applications without transceivers and optical fibers. No special termination is required because the 82546EB device is internally terminated. The figure below illustrates a gigabit backplane or point-to-point circuit based on the 82546EB device:



Affected Docs: 82546EB Gigabit Ethernet Controller Preliminary Datasheet and Hardware Design Guide Rev. 0.5.

2. Thermal Management

Problem: The design guide will change to provide additional information on thermal management for the 82546EB controller.

Thermal management is based upon controlling the junction temperature, where

$$T_j = T_{\text{ambient}} + \theta_{JA} * P, \text{ and } \theta_{JA} = \theta_{JC} + \theta_{CA}$$

The goal is to keep the die temperature under a critical temperature of 120C for the given ambient operating environment.

How to Reduce the Case to Ambient Thermal Resistance, θ_{CA} , for LOM Designs

One element that designers can influence in their designs is the thermal resistance θ_{CA} . If increased airflow is not an option, the thermal resistance can be reduced by adding a heat sink to your 82546EB Controller PCB design. The manufacturers of add-on heat sinks should provide detailed literature and technical reports on improvements in θ_{CA} with their products.


An alternative is build a heat sink into the printed circuit board. This is can be accomplished by adding copper to the board and drilling numerous thermal vias for heat dissipation under the BGA package.

Since the GND balls on the 82546EB controller are clustered together, it is fairly easy to use the following layout technique:

- 1) On the first layer of the board, create a copper plane shape under the GND balls of the device. This will create a strong GND connection between balls and provide a large target for the thermal vias.
- 2) Drilling at least 4 thermal vias, which are about 14-18 mils in finished hole size, connect them to the plane shape on the first board layer and tie them to the internal ground plane. The vias should go through all board layers. The vias will act as a conduit for heat dissipation. Bury the thermal via tie legs on all plane layers to maximize heat transfer.
- 3) Minimize any interrupted ground plane under the BGA and within 0.5 inch of the edge of the 82546EB controller package. Solid ground planes help to transfer the heat easily to other regions of the PCB. Large cuts in the plane inhibit the transfer of heat which increases the case temperature of the silicon.
- 4) On the bottom layer, add another copper plane for heat dissipation and good grounding.

Typical Results

Following the above recommendations can dramatically reduce the case temperature of the 82546EB Gigabit Ethernet Controller, keeping the junction temperature within safe limits. In one experiment, using 2 oz. copper instead of 1 oz. copper on six layer circuit boards reduced the measured case temperature from approximately 87°C to 72 °C.



Affected Docs: 82546EB Gigabit Ethernet Controller Preliminary Datasheet and Hardware Design Guide Rev. 0.5.