82540EP Gigabit Ethernet Controller Specification Update

October 10, 2005

The 82540EP 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.

CONTENTS

REVISION HISTORY

82540EP Gigabit Ethernet Controller Specification Update

Date of Revision	Description
April 1, 2004	Initial pubic release. Added Errata #6 and 7.
August 3, 2004	Added Errata # 8 and 9.
January 10, 2005	Added Spec Clarification #3. Removed Spec Changes #1 and 2. Removed Spec Clarification # 1.
October 10, 2005	Added Errata #10.

PREFACE

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

- 82540EP Gigabit Ethernet Controller Datasheet, Intel Corporation.
- 82551QM/82540EM Interchangeable LOM Design Application Note (AP-432), Intel Corporation.
- 82546EB Gigabit Ethernet Controller Networking Silicon Developer's Manual + Appendices for 82545EM, 82540EM, 82544EI/GC, Intel Corporation.

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

All 82540EP 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

82540EP controller steppings will be identified by the following register contents:

Stepping	Vendor ID	Device ID	Revision Number
82540EP A0	8086h	100Eh	03h

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

GENERAL INFORMATION

This section covers the 82540EP device.

82540EP COMPONENT MARKING INFORMATION

Product	Stepping	QDF Number	Top Marking	Notes
82540EP	A0	Q 513	RC82540EP	Engineering Samples
82540EP	A0	-	RC82540EP	Production

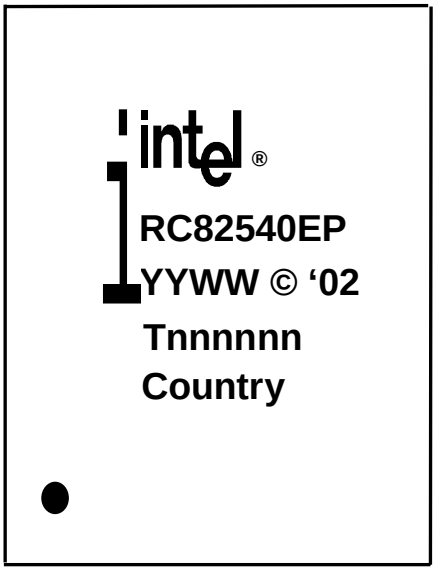

SUMMARY TABLE OF CHANGES

The following table indicates the Specification Changes, Errata, Specification Clarifications or Documentation Changes, which apply to the listed 82540EP 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.

No.	A0	Plans	ERRATA	Page	Notes
1	X	NoFix	Master-Aborts with Some Chipsets During Driver-Initiated Controller Reset	7	-
2	X	NoFix	Intermittent Issues with TCO Receive Packets in IPMI Mode	7	-
3	X	NoFix	LSO Premature Descriptor Write Back	8	-
4	X	NoFix	XOFF from Link Partner can Pause Flow-Control (XON/XOFF) Transmission	8	-
5	X	NoFix	Transmit Descriptor use of RS for non-data (Context & Null) Descriptors	8	-
6	X	NoFix	Message Signaled Interrupt Feature May Corrupt Write Transactions	9	-
7	X	NoFix	Wakeup Packet Memory (WUPM) cleared upon reset	9	-
8	X	NoFix	Unexpected RCMP ACK packets in ASF mode	9	-
9	X	NoFix	Exceeding PCI Power Management Specification Limit of 375mA current during reset and power state transitions	9	-
10	X	NoFix	Legacy Transmit Descriptor Write-Back May Occur Before the Packet Data Associated with the Descriptor is Fetched	10	New
-	-	-	-	-	-
No.	**A0**	**Plans**	**SPECIFICATION CLARIFICATIONS**	**Page**	**Notes**
1	X	Doc Change	LEDs Inactive Until Driver Loads	11	-
2	X	Doc Change	TTL inputs not 5V tolerant	11	-
-	-	-	-	-	-
No.	**A0**	**Plans**	**DOCUMENTATION CHANGES**	**Page**	**Notes**
1	X	Doc Change	1.5V Regulator Control Circuit Start Up	12	-
2	X	Doc Change	Impedance Compensation Resistor Connections	12	-
-	-	-	-	-	-

SPECIFICATION CHANGES

No specification changes reported at this time.

ERRATA

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

Problem: The 82540EP Gigabit Ethernet Controller implements a software-Initiated device reset function through its control register space (CTRL Register Bit 31). 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 internal 64-bit architectures may initiate an additional zero byte write immediately following a 32-bit write to the device 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 out bound 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 82540EP controller can only be accessed using DWord (32-bit) software operations; since the potential zero 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 82540EP Gigabit Ethernet Controller.

2. 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 82540EP 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 82540EP Gigabit Ethernet Controller.

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

Problem: When the 82540EP 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 82540EP Gigabit Ethernet Controller.

5. 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 82540EP Gigabit Ethernet Controller.

6. 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 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).

Status: Intel does not plan to resolve this erratum in a future stepping of the 82540EP Gigabit Ethernet Controllers.

7. Wakeup Packet Memory (WUPM) cleared upon reset

Problem: The 82540EP 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 82540EP Gigabit Ethernet Controllers.

8. Unexpected RCMP 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 82540EP Gigabit Ethernet Controllers.

9. 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 82540EP Gigabit Ethernet Controllers.

10. 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 82540EP Gigabit Ethernet Controller.

SPECIFICATION CLARIFICATIONS

1. LEDs Inactive Until Driver Loads

Problem: LED indications (link, activity, speed) are not active until the software driver loads even though the PHY has auto-negotiated and established link with a partner on the Ethernet.

The LED signals work this way because they are derived from MAC logic and are qualified by the Set Link Up (SLU) bit in the Device Control Register (CTRL). Driver software sets this bit when it initializes. Alternatively, the SLU bit sets automatically when either Advanced Power Management or the SMBus are enabled through EEPROM settings.

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

2. TTL inputs not 5V tolerant

Problem: The TTL inputs on the Ethernet controller are not 5V tolerant. If these inputs are connected to 5V, then damage to the controller is likely to occur. TTL inputs include the JTAG interface pins, the FLASH interface pins, the EEPROM interface pins, the LED pins, the software definable pins, and the LAN_PWR_GOOD pin.

Affected Specs: 82540EP Gigabit Ethernet Controller Datasheet.

DOCUMENTATION CHANGES

1. 1.5V Regulator Control Circuit

Problem: The 82540EP product contains on-chip LDO voltage regulation controllers for both 1.5V and 2.5V. The 1.5V regulator control output is Ball P11 on the 82540EP controller, and may be optionally used to drive an external PNP pass transistor as shown in the figure below.

The earlier 82540EM Gigabit Ethernet Controller contains an erratum which causes the 1.5V regulator control circuit to possibly start up incorrectly. Workarounds involved the use of a resistor attached to the transistor base, replacing the transistor with a similar model carrying tighter gain specifications, enabling the SMBus, and/or adding a second transistor. The 82540EP controller will work properly in a system containing a recommended workaround for the 82540EM device **as long as the workaround is implemented in its entirety.**

The 82540EP controller does not require any additional workaround, however, and may be applied using the circuit below:

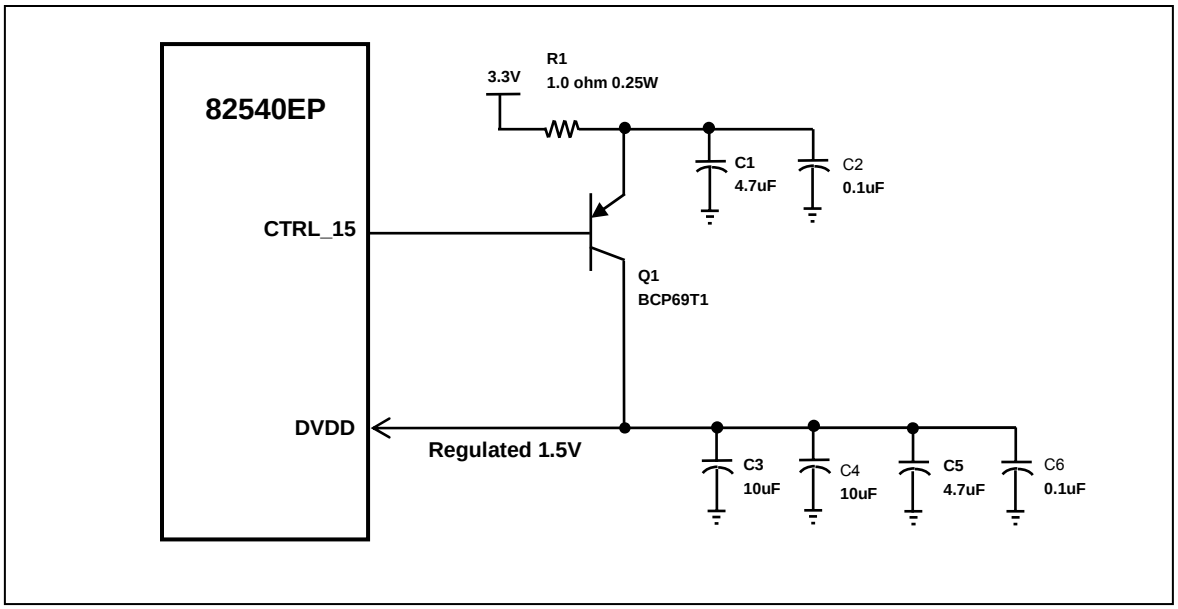

Note that the 82540EP also contains a “Maximum Power-Savings” mode to permit applications to reduce D3 no-wakeup power consumption. However, if “Maximum Power Savings” is utilized, then the on-chip LDO regulation circuits cannot be used, and power must be supplied via external voltage regulation.

Affected Specs: Component selection and values illustrated in the 82551QM/82540EM Interchangeable LOM Design Application Note (AP-432) Rev. 0.5.

2. Impedance Compensation Resistor Connections

Problem: The stuffing option table in the 82551QM/82540EM Interchangeable LOM Design Ap Note suggests incorrect (reversed) connections to the impedance compensation pins for a system built with the 82540EM device.

The correct circuit recommendation for ball G4, ZP_COMP, is:

Pull-down
Stuff precision 53-ohm resistor **pull-down** for 82540EM.

The correct circuit recommendation for ball H4, ZN_COMP, is:

Pull-up
Stuff precision 35-ohm resistor **pull-up** for 82540EM.

This document may be used for the 82540EP Gigabit Ethernet Controller as long as the two documented ballout changes from the 82540EM controller to the 82540EP controller are applied: Ball C8 changes from NC to CLKRUN and Ball C5 changes from APM_WAKEUP to NC.