

82541PI
Gigabit Ethernet Controller
Specification Update

November 19, 2009

82541PI 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.

82541PI Gigabit Ethernet controllers 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>.

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

Copyright © Intel Corporation 2009

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

CONTENTS

CONTENTS	1
REVISION HISTORY	2
PREFACE	2
NOMENCLATURE	2
COMPONENT IDENTIFICATION VIA PROGRAMMING INTERFACE	3
GENERAL INFORMATION	3
82541PI Component Marking Information.....	3
SUMMARY TABLE OF CHANGES	4
Codes Used in Summary Tables.....	4
SPECIFICATION CHANGES	5
GENERAL OPERATING CONDITIONS FOR 1.8V AND 1.2V.....	5
DC POWER SPECIFICATIONS FOR 1.8V, 1.2V AND 3.3V.....	5
ERRATA	7
1. MASTER-ABORTS WITH SOME CHIPSETS DURING DRIVER-INITIATED CONTROLLER RESET.....	7
2. MARGINAL INTERNAL POWER ON RESET FUNCTION.....	8
3. PARTNER ASSUMED TO BE FULL DUPLEX WHEN DOWNSHIFTING FROM 1000BASE-T TO 100BASE-TX.....	8
4. DHCP NOT SUPPORTED.....	8
5. PET TRANSIENT EVENT PROCESSING DURING LAN LINK DOWN.....	8
6. I2C SHORT TRANSACTIONS NOT SUPPORTED FOR ADDRESS C8H.....	8
7. MESSAGE SIGNED INTERRUPT FEATURE MAY CORRUPT WRITE TRANSACTIONS.....	9
8. FIRST SMB TRANSACTION AFTER SYSTEM POWER-UP IS NACKED.....	9
9. 1000 MBPS SLAVE LINK FAILURES WITH HIGH JITTER LINK PARTNERS AND SHORT CABLES.....	9
10. POSSIBLE RIPPLE ON 1.8V AND 1.2V REGULATOR OUTPUTS.....	10
11. INTERMITTENT ISSUES WITH TCO RECEIVE PACKETS IN IPMI MODE (82541PI).....	10
12. RECEIVE PACKET DELAYED WHEN USING RDTR OR RADV REGISTER.....	11
13. MNG: PASS-THROUGH FROM BMC TO LAN HANGS DURING PCI RESET.....	11
14. MNG: SMBUS HANG WHEN DELAY BETWEEN TRANSACTIONS IS LESS THAN 10 μ S.....	11
15. OVERWRITES TRANSMIT DESCRIPTORS IN INTERNAL BUFFER.....	12
SPECIFICATION CLARIFICATIONS	12
1. RESISTOR VALUE CHANGES FOR 82541PI STEPPING.....	12
2. PCI-X CAPABILITY REGISTER IN PCI CONFIGURATION SPACE.....	12
3. 82541PI SOFTWARE DEVICE DRIVERS IN RELEASE 14.7 ARE THE FINAL WHQL DRIVERS.....	13
DOCUMENTATION CHANGES	13

REVISION HISTORY

82541PI Gigabit Ethernet Controller Specification Update

Date of Revision	Description
November 2009	Initial Public Release.

PREFACE

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

- 82541 Family of Gigabit Ethernet Controllers Datasheet.
- 82541ER Design Guide.
- 82547G(EI)/82541(PI/GI/EI)/82541ER EEPROM Map and Programming Information Guide.

For software driver programming information, please contact your Intel representative.

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 82541PI 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

The 82541PI Gigabit Ethernet controller requires the appropriate EEPROM images described in the *82547G/(E)/82541/(PI/G/EI)/82541ER EEPROM Map and Programming Information Guide* and can be identified by the following register contents:

Stepping	Vendor ID	Device ID	Revision Number
82541PI C0	8086h	1076h	05h
82541PI C0	8086h	107Ch ¹	05h

¹ Lead-free version.


GENERAL INFORMATION

This section lists the 82541PII device ID.

82541PI COMPONENT MARKING INFORMATION

Product	Stepping	QDF Number	Top Marking	Notes
82541PI	C0		541PI	Production Units

Note: 82541PI devices that are lead-free are marked with a circled “e1” and have a product code prefix: LUxxxxxx.



SUMMARY TABLE OF CHANGES

The following table lists the Specification Changes, Errata, Specification Clarifications or Documentation Changes, which apply to the 82541PI Gigabit Ethernet controller.

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 the listed stepping.
Shaded:	This item is either new or modified from the previous version of the document.

No.	C0	Plans	SPECIFICATION CHANGES	Page	Notes
1	X		General Operating Conditions for 1.8V and 1.2V	5	-
2	X		DC Power Specifications for 1.8V, 1.2V, and 3.3V	5	-
No.	C0	Plans	ERRATA	Page	Notes
1	X	NoFix	Master-Aborts with Some Chipsets During Driver-Initiated Controller Reset	7	-
2	X	NoFix	Marginal Internal Power on Reset Function	8	-
3	X	NoFix	Full Duplex Partner Assumed when Shifting from 1000BASE-T to 100BASE-TX	8	-
4	X	NoFix	DHCP Not Supported	8	-
5	X	NoFix	PET Transient Event Processing During LAN Link Down	8	-
6	X	NoFix	I2C Short Transactions Not Supported for Address C8h	8	-
7	X	NoFix	Message Signaled Interrupt Feature May Corrupt Write Transactions	9	-
8	X	NoFix	First SMB Transaction after System Power Up is NACKed	9	-
9	X	NoFix	1000 Mbps Slave Link Failures with High Jitter Link Partners and Short Cables	9	-
10	X	NoFix	Possible Ripple on 1.8V and 1.2V Regulator Outputs	10	-
11	X	NoFix	Intermittent Issues with TCO Receive Packets in IPMI Mode	10	-
12	X	NoFix	Receive Packet Delayed When Using RDTR or RADV Register	11	-
13	X	NoFix	MNG: Pass-Through From BMC to LAN Hangs During PCI Reset	11	-
14	X	NoFix	MNG:SMBus Hang When Delay Between Transactions is Less Than 10 μ s	11	-
16	X	NoFix	Overwrites Transmit Descriptors in Internal Buffer	12	-
No.	C0	Plans	SPECIFICATION CLARIFICATIONS	Page	Notes
1	X	NoFix	Resistor Value Changed for 82541PI Stepping	12	-
2	X	NoFix	PCI-X Capability Register in PCI Configuration Space	12	-
3	X	NoFix	82541PI Software Device Drivers in Release 14.7 Are the Final WHQL Drivers	13	-

SPECIFICATION CHANGES

General Operating Conditions for 1.8V and 1.2V

Problem: This change is related to erratum 48, "Possible Ripple on 1.8V and 1.2V Regulator Outputs," listed in this document. Due to this anomaly, the following specifications were modified to the following:

Symbol	Parameter	Min	Max	Units
VDD(1.8)	DC supply Voltage on 1.8V pins	1.71	1.89	V
VDD(1.2)	DC supply Voltage on 1.2V pins	1.14	1.26	V
Voltage Ramps				
1.8V Ripple	Maximum voltage ripple at frequency below 1 MHz		280	mV _{pk-to-pk}
1.8V Ripple	Maximum voltage ripple at frequency below 1 MHz Note: This is the lowest absolute voltage for the frequency range below 1 MHz.	1.55		V
1.2V Ripple	Maximum voltage ripple at frequency below 1 MHz		180	mV _{pk-to-pk}
1.2V Ripple	Maximum voltage ripple at frequency below 1 MHz Note: This is the lowest absolute voltage for the frequency range below 1 MHz.	1		V

DC Power Specifications for 1.8V, 1.2V and 3.3V

Problem: The subsystem power consumption in D3 cold wake-enabled @ 100Mbps and D0 active @ 10Mbps is not specified. Also power specifications may be rounded off.

The subsystem power consumption has been updated in the subsystem power specification table. The modified Power specifications are listed below. These specifications are updated in the 82541 Datasheet.

Power Specifications - Doa

	D0a							
	Unplugged no link		@ 10 Mbps		@ 100 Mbps		@ 1000 Mbps	
	Typ lcc (mA) ^a	Typ lcc (mA) ^b	Typ lcc (mA) ^a	Typ lcc (mA) ^b	Typ lcc (mA) ^a	Typ lcc (mA) ^b	Typ lcc (mA) ^a	Typ lcc (mA) ^b
3.3V	3	5	5	10	13	15	30	40
1.8V	14	15	85	85	110	115	315	320
1.2V	30	35	85	90	90	100	380	400
Total Device Power	75 mW	85 mW	270 mW	295 mW	350 mW	380 mW	1.1 W	1.2 W

- Typical conditions: operating temperature (TA) = 25 C, nominal voltages, moderate network traffic at full duplex, and PCI 33 MHz system interface.
- Maximum conditions: minimum operating temperature (TA) values, maximum voltage values, continuous network traffic at full duplex, and PCI 33 MHz system interface.

Power Specifications – D3cold

	D3cold – wake-up enabled ^a						D3cold – wake disabled	
	Unplugged link		@ 10 Mbps		@ 100 Mbps			
	Typ lcc (mA) ^b	Typ lcc (mA) ^c	Typ lcc (mA) ^b	Typ lcc (mA) ^c	Typ lcc (mA) ^b	Typ lcc (mA) ^c	Typ lcc (mA) ^b	Typ lcc (mA) ^c
3.3V	2	3	2	3	2	3	4	5
1.8V	14	15	20	25	110	115	1	2
1.2V	21	25	30	35	80	85	7	10
Total Device Power	60 mW	70 mW	80 mW	100 mW	300 mW	320 mW	25 mW	35 mW

- At 1000 Mbps, power consumption is not shown since the controller switches to the 10/100 Mbps state before entering D3 to conserve power.
- Typical conditions: operating temperature (TA) = 25 C, nominal voltages, moderate network traffic at full duplex, and PCI 33 MHz system interface.
- Maximum conditions: minimum operating temperature (TA) values, maximum voltage values, continuous network traffic at full duplex, and PCI 33 MHz system interface.

Power Specifications – Complete Subsystem

Complete Subsystem (Reference Design) including Magnetics, LED, Regulator Circuits												
	D3cold – wake disabled		D3cold wake-enabled @ 10 Mbps		D3cold wake-enabled @ 100 Mbps		D0 @ 10 Mbps active		D0 @ 100 Mbps active		D0 @ 1000 Mbps active	
	Typ lcc (mA) ^a	Typ lcc (mA) ^b	Typ lcc (mA) ^a	Typ lcc (mA) ^b	Typ lcc (mA) ^a	Typ lcc (mA) ^b	Typ lcc (mA) ^a	Typ lcc (mA) ^b	Typ lcc (mA) ^a	Typ lcc (mA) ^b	Typ lcc (mA) ^a	Typ lcc (mA) ^b
3.3 V	4	5	2	3	6	7	7	12	19	21	36	46
1.8 V	1	2	20	25	110	115	85	85	110	115	315	320
1.2 V	7	10	30	35	80	85	85	90	90	100	380	400
Subsystem Power	40 mW	60 mW	170 mW	210 mW	650 mW	685 mW	585 mW	620 mW	725 mW	780 mW	2.4 W	2.5 W

- a. Typical conditions: operating temperature (TA) = 25 C, nominal voltages, moderate network traffic at full duplex, and PCI 33 MHz system interface.
- b. Maximum conditions: minimum operating temperature (TA) values, maximum voltage values, continuous network traffic at full duplex, and PCI 33 MHz system interface.

ERRATA

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

Problem: 82541PI Gigabit Ethernet controllers implement a software-initiated device reset function through the control register space (Control Register, bit 31). This software-initiated reset re-initializes all functional state of the controllers except for PCI/PCI-X configuration. When the reset is written, the controllers require a few internal clock cycles to complete the reset operation. During this brief time, they 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 outbound 0-byte write operation will not result in any adverse system behavior, though the event may be logged at 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 82541PI device 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 or 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 the 82541 Gigabit Ethernet Controller.

2. Marginal Internal Power on Reset Function

- Problem:** Power on (internal) reset is not dependable in the event of a slow power ramp. If the 3.3V supply ramps slowly, the 1.2V supply (from the collector of the external PNP pass transistor) may not be stable when reset occurs. Voltage thresholds are not within design targets.
- Implication:** Internal power on reset is an alternative to using the LAN_PWR_GOOD input. To enable power on reset, LAN_PWR_GOOD must be connected to 3.3V through a pull-up resistor. For this configuration, a power ramp specification **must** be imposed and that the 3.3V supply ramps from its 10% point to its 90% point in less than 15 ms. This requirement is more restrictive than the general recommendation that all power supplies should ramp to within their regulation bands in less than 20 ms.
- Workaround:** The LAN_PWR_GOOD signal should be connected to the system's voltage supervisor function. Using LAN_PWR_GOOD as the external reset input is the preferred method.
- Status:** Intel does not intend to resolve this erratum in the 82541PI Gigabit Ethernet controller.

3. Partner Assumed to Be Full Duplex When Downshifting from 1000BASE-T to 100BASE-TX

- Problem:** For a link partner that advertises 1000BASE-T full duplex and 100BASE-TX half duplex capabilities, the 82541PI controller will attempt to link at 1000 Mbps, followed by 100 Mbps. However, it will assume the link partner is capable of full duplex operation at both speeds.
- Implication:** If the advertisement is correct, a problem could occur obtaining link. Practical Gigabit Ethernet devices are by default capable of full duplex operation. Thus, the possibility of this advertisement is very low.
- Workaround:** There is no workaround for this erratum.
- Status:** Intel does not intend to resolve this erratum in the 82541 Gigabit Ethernet controller.

4. DHCP Not Supported

- Problem:** The management block DHCP function has problems renewing IP addresses at the end of the lease. Therefore, typically an IP address will not be renewed upon link restoration.
- Implication:** The manageability DHCP function is not supported. This issue primarily affects mobile ASF applications.
- Workaround:** An upgrade firmware patch will be available through Intel in the future.
- Status:** Intel does not plan to resolve this erratum in the 82541 Gigabit Ethernet controller.

5. PET Transient Event Processing During LAN Link Down

- Problem:** The manageability block will send transient SOS alert events when the LAN link is down.
- Implication:** A transient event is defined as a change of state followed by a return to the original state. This behavior does not meet ASF specifications.
- Workaround:** There is no workaround for this issue.
- Status:** Intel does not plan to resolve this erratum in the 82541Gigabit Ethernet Controller.

6. I2C Short Transactions Not Supported for Address C8h

- Problem:** In TCO mode, the manageability block will not support I2C short read block transactions for SMB address C8h.
- Implication:** This behavior is a manageability usage limitation. Other addresses are available.
- Workaround:** A firmware workaround patch may be possible.
- Status:** Intel does not plan to resolve this erratum in the 82541Gigabit Ethernet controller.

7. 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 controllers. During MSI writes, the controller incorrectly accesses the write data FIFO. If there are pending write transactions when this occurs, these transactions may become corrupted, which may cause the network controller to lock up and become unresponsive.

For a normal PCI write transaction, the controller's PCI logic receives data to be written from an internal FIFO. Once the controller is given bus ownership, the PCI logic pulls the data out of this FIFO and performs the write transaction.

For systems using MSI writes, the data, which is constant, should be pulled from the controller's PCI Configuration Space rather than the internal FIFO. The affected devices are not pulling this data from PCI Configuration Space. Instead, they are pulling data from the internal FIFO.

Implication: If the affected products are used with a future OS that uses 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 the 82541 Gigabit Ethernet controller.

8. First SMB Transaction after System Power-Up is NACKed

Problem: This issue is relevant only for the first power-up, signified by a change in the auxiliary power state, caused by a physical disconnection or connection of the AC power cord. The 82541 controller will Not Acknowledge (NACK) the address byte of the first SMB access transaction received after the first power up event. All subsequent SMB transactions will be Acknowledged (ACKed), responded and processed correctly.

Implication: Configuration changes may be lost by the 82541 controller, causing synchronization loss between the SMB master and the 82541 controller.

Workaround: The SMB master should implement an SMB retransmission mechanism after all failed SMB transactions. This mechanism will solve this issue as well as SMB failures resulting from a noisy electrical environment. A retransmission count of 2 transactions will suffice to solve this issue completely.

Status: Intel does not plan to resolve this erratum in the 82541 Gigabit Ethernet controller.

9. 1000 Mbps Slave Link failures With High Jitter Link Partners and Short Cables

Problem: When configured to 1000 Mbps slave mode, the 82541 device PHY has an unstable link with high jitter link partners. This issue is relevant only for shorter cable connections, where the cable length is approximately less than 20 meters. The instability is characterized with frequent link drops or high bit error rate if a valid link is present. On master mode the connection is stable and link is established and maintained.

Implication: An unstable and unreliable connection may be obtained when connected to high jitter link partners with short cable. For B0 steppings or earlier, the PHY default configuration is slave mode. Therefore, the problem may be observed if no action is taken.

For the B1 stepping and later versions, the PHY default is automatic, meaning it becomes either master or slave. Thus, the problem is only observed if the PHY was placed into slave mode as result of the Auto-Negotiation process.

Workaround: For proper link with the high jitter link partners, it is required to set the PHY to master mode via a software driver version that can force the PHY into master mode.

Status: Intel does not plan to resolve this erratum in the 82541 Gigabit Ethernet controller.

10. Possible Ripple on 1.8V and 1.2V Regulator Outputs

- Problem:** The internal ground bounce that occurs when the 82541 drives the PCI bus can introduce ripple on the 1.2V and 1.8V linear voltage regulator (LVR) outputs. This ripple violates the published DC specification (nominal voltage of -7% and +5%).
- Implication:** The ripple of concern occurs at a frequency below 1 MHz. When these voltage levels are measured, higher frequency noise components will be seen super-imposed on the lower frequency (less than 1 MHz) ripple. These higher frequency glitches are unrelated to the LVR functionality and should be disregarded. When measuring this ripple, it is recommended that the scope bandwidth is limited to the greatest extent to filter out noise greater than 1 MHz.
- Under normal network data conditions (random data patterns), this ripple does not violate the current DC specification values. However, under heavy simultaneous switching loads presented over prolonged periods of time (such as sequential large packets with worst case data transition patterns received over the network), the ripple can violate the specified DC values. Worst case data patterns are those that cause the majority of 32 PCI data outputs to simultaneously switch from one to zero in a repeated pattern, such as FFFF FFFFh, 0000 0000h, FFFF FFFFh, 0000 0000h, etc.
- Note:** There have been no customer reports related to this ripple component. The specification violation has only been observed during testing using a SmartBits network exerciser to generate sustained worst case loads.
- The 82541 can tolerate a periodic voltage variation beyond the existing DC specification to a degree specified in a revised DC and new AC ripple specification without impact to device functionality or reliability. If the ripple component exceeds the value of the new AC specification, decreased BER performance may result. The Ethernet link may also be reset, causing link to be lost and re-established.
- Root Cause:** Extreme switching (for example, sequences of sequential packets with worst case pattern type data) of the PCI data output buffers causes internal switching currents that disrupt the LVR reference voltage.
- Workaround:** There is no workaround for this switching noise when an internal LVR is used. If external voltage regulators are used, switching noise will not occur since the external regulators are not affected by the internal ground bounce on the LVR.
- Status:** The updated specifications are listed in the Specification Changes section of this document.

11. Intermittent Issues with TCO Receive Packets in IPMI Mode (82541PI)

- 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 82541PI Gigabit Ethernet Controller.

12. Receive Packet Delayed When Using RDTR or RADV Register

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

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

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

Status: Intel does not plan to resolve this erratum in the 82541PI Gigabit Ethernet Controller.

13. MNG: Pass-Through From BMC to LAN Hangs During PCI Reset

Problem: When a PCI reset occurs while a packet is being passed through from the BMC to the LAN, the FIFO control logic may hang.

Implication: BMC can no longer transmit to the LAN.

Workaround: Contact your Intel representative for the latest NVM release which includes an SMBus command that allows the BMC to reset the manageability logic. The BMC should send this command when a hang is detected.

Status: Intel does not plan to resolve this erratum in the 82541PI Gigabit Ethernet Controller.

14. MNG: SMBus Hang When Delay Between Transactions is Less Than 10 μ s

Problem: When the time between the STOP and START is less than 10 μ s, the SMBus logic may hang. This appears externally as a clock stretch which hangs the SMBus. The standard requires a minimum of 4 μ s between STOP and START.

Implication: No SMBus communication.

Workaround: Contact your Intel representative for the latest NVM release.

Status: Intel does not plan to resolve this erratum in the 82541PI Gigabit Ethernet Controller.

15. Overwrites Transmit Descriptors in Internal Buffer

- Problem:** This erratum occurs when the internal transmit descriptor buffer is nearly full of descriptors. If the free space in this buffer is smaller than the system cacheline, the calculation of the size of the descriptor fetch may be incorrect.
- Implication:** Corruption of the transmit descriptor ring; can cause a system crash. In most applications, the descriptors will be written back as soon as the data has been read and they will not be accumulating in the internal buffer, therefore this issue will not be seen. However, in an application where system events prevent the immediate write-back of descriptors, the descriptor buffer could fill up and this issue could be seen.
- Workaround:** The driver should keep track of the difference between the Transmit head and tail and make sure the difference between tail and head is never more than the value shown below:

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

- Status:** Intel does not plan to resolve this erratum in the 82541PI Gigabit Ethernet Controller.

SPECIFICATION CLARIFICATIONS

1. Resistor Value Changes for 82541PI Stepping

- Problem:** The 82541PI (C0 stepping) has new I/O drivers have a lower internal pull-up resistance value than previous versions of the 82541. Thus, some designs will require a change to the value of the pull-down resistors,
- On ball J4, the EE_MODE pin, use a recommended value of 100 Ω , instead of the 1 K Ω used in previous 82541 steppings, if a MicroWire EEPROM is used in the current solution. There is no impact to designs using an SPI EEPROM.
- On ball L13, the JTAG_TRST pin, use a recommended value of 100 Ω , instead of the 1 K Ω used in previous 82541 steppings.
- For 82541PI, verify vendor specifications to ensure that Vol is less than 0.7V on ball N10, EEDO, and ball P9, FLSH_SO/LAN_DISABLE#.
- Affected Specs:** For more information refer to the design schematics (design guide) and datasheet for the 82541 Ethernet Controller.

2. PCI-X Capability Register in PCI Configuration Space

- Problem:** The 82541 controller includes the PCI-X Capability register (capability ID = 0x07) in its PCI configuration space. If the system software were to look only at configuration space it might mistakenly determine that the controller was capable of operation in PCI-X mode, when in fact the 82541 does not support PCI-X mode.
- Affected Specs:** None.

3. 82541PI Software Device Drivers in Release 14.7 Are the Final WHQL Drivers

Clarification: Any changes to the current 82541PI Windows software device driver will no longer be WHQLd. Windows 7 and Windows 2008 Server R2 will only provide in-box base driver support. There will be no support for any future Windows operating systems.

Note that Linux kernel 2.6.31 is the last kernel that is validated for the 82541PI.

DOCUMENTATION CHANGES

There are currently no documentation changes for the 82541PI.